

Modeling and simulation of longitudinal dynamics for Low Energy Ring–High Energy Ring at the Positron-Electron Project

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

Stanford Linear Accelerator Center, Stanford, California 94309, USA

(Received 22 September 2006; published 21 February 2007)

A time domain dynamic modeling and simulation tool for beam-cavity interactions in the Low Energy Ring (LER) and High Energy Ring (HER) at the Positron-Electron Project (PEP-II) is presented. Dynamic simulation results for PEP-II are compared to measurements of the actual machine. The motivation for this tool is to explore the stability margins and performance limits of PEP-II radio-frequency (RF) systems at future higher currents and upgraded RF configurations. It also serves as a test bed for new control algorithms and can define the ultimate limits of the low-level RF (LLRF) architecture. The time domain program captures the dynamic behavior of the beam-cavity-LLRF interaction based on a reduced model. The ring current is represented by macrobunches. Multiple RF stations in the ring are represented via one or two macrocavities. Each macrocavity captures the overall behavior of all the 2 or 4 cavity RF stations. Station models include nonlinear elements in the klystron and signal processing. This enables modeling the principal longitudinal impedance control loops interacting via the longitudinal beam model. The dynamics of the simulation model are validated by comparing the measured growth rates for the LER with simulation results. The simulated behavior of the LER at increased operation currents is presented via low-mode instability growth rates. Different control strategies are compared and the effects of both the imperfections in the LLRF signal processing and the nonlinear drivers and klystrons are explored.

DOI: [10.1103/PhysRevSTAB.10.022801](https://doi.org/10.1103/PhysRevSTAB.10.022801)

PACS numbers: 29.20.-c, 29.27.Bd

I. INTRODUCTION

High-current accelerators exhibit dynamics between the beam-loaded radio-frequency (RF) systems and the particle beams. To counteract coupled-bunch instabilities due to the RF cavity fundamental impedance, some RF systems employ feedback techniques which act to reduce the impedance interacting with the particle beam and hence increase the stability of the beam.

The RF feedback techniques mentioned above (“direct feedback” and “comb feedback”) have implications for the dynamics and stability of the closed-loop RF systems as well as for the particle beam. The analysis of the RF system dynamics and the interaction with the dynamics of the particle beam are nonlinear problems. Understanding the effectiveness of these techniques and the requirements for the operation of the RF systems are fundamental aspects of the design and operation of these heavily beam-loaded accelerator RF systems [1–3].

This paper presents results from a nonlinear simulation study of the RF systems in the Positron-Electron Project (PEP-II) B-Factory collider, and highlights the design and topology of the low-level RF (LLRF) feedback systems. The simulation model is verified against measured accelerator dynamics, and the likely operational limits for the existing LLRF system implementation are predicted. Several methods for improving the performance of the LLRF systems are explored, as part of a study of operation of the PEP-II facility at higher luminosity and currents. Results specific to PEP-II are presented, though the general

form of the simulation model, and the simulation technique itself, are generally applicable to high-current accelerators. Indeed, our results have been applied at the Stanford Positron Electron Accelerating Ring (SPEAR III). The analysis in this paper focuses on the Low Energy Ring (LER) since it is closer to instability limits than the High Energy Ring (HER).

The paper is organized as follows. Section II describes the physical system to be modeled. Section III describes the simulation and relates it with the system through equations, assumptions, and simplifications. Section IV defines the RF impedance and its control via the LLRF feedback paths. Section V presents the means to compare the physical system and simulation, and shows that a close relationship of the two corresponds to a close relationship of their RF impedances. Section VI defines the instability growth rates and the way they are measured. Then, it compares the growth rates of simulation and physical system for the same operating points. Section VII describes quantitatively the sensitivity of growth rates on different parameters, thus presenting possibilities for improvements and modifications. Section VIII describes some implementation details of the LLRF system, and highlights how imperfections in the systems and differences that exist between the multiple RF stations influence the dynamics. Section IX describes the limitations on the physical system and presents the predicted growth rates for higher currents as well as the estimated system margins. Finally, Sec. X mentions some of the potential projects involving the simulation and future measurements related to this work.

II. SYSTEM DESCRIPTION

The PEP-II RF system block diagram is shown in Fig. 1. The RF stations are comprised of 1.2 MW, 476 MHz klystrons with either 2 or 4 normal-conducting RF cavities with high-order mode dampers and an R/Q ratio of 116. In heavy loaded rings, there is a strong interaction between the multiple-bunched beam and the RF station. This beam loading is mainly produced by the effective cavity impedance seen by the beam current. Feedback systems around the stations are needed to reduce that impedance and consequently minimize the fast unstable growth of the low-order modes in the beam.

The LLRF systems include direct and comb loop feedback paths to reduce impedances seen by the beam. The stations also incorporate numerous low bandwidth regulating loops which control the cavity tuners, the high-voltage power supply voltage, and compensate for gap transient effects [4,5]. The tuner loop adjusts the cavity for minimum reflected power, whereas the klystron saturation loop maintains constant saturation headroom by controlling the high-voltage power supply to the klystron. The gap feedback loop removes revolution harmonics from the feedback error signal to avoid saturating the klystron.

The direct loop causes the station to follow the RF reference adding regulation to the cavity voltage, thus extending the beam-loading Robinson stability limit and lowering the effective fundamental impedance seen by the beam. The comb loop consists of a second order digital infinite impulse response (IIR) filter that adds narrow gain peaks at synchrotron sidebands around revolution harmonics to further reduce the residual impedance. Despite the LLRF feedback, the beam exhibits low-mode coupled-bunch instabilities at operating currents due to the fundamental impedance, and a special “woofers” feedback channel is required to control low-mode instabilities [6], seen as the “longitudinal low group-delay woofers” in the block diagram.




FIG. 1. (Color) System block diagram. Fast dynamics (modeled) appear in blue, slow dynamics (fixed parameters in simulation) in green, and not modeled components in red.

III. MODEL DESCRIPTION

The simulation is focused on understanding the interaction among the low-order dynamics of the beam, the cavities, and the fast LLRF feedback loops. This tool is developed as a block system in SIMULINK, which uses the system parameters calculated in MATLAB [7] to set the initial conditions of the slow loops and to provide measurement/estimation tools. The simulation is an update of a previous work developed by Tighe [8].

The overall dynamic system is of complex structure, including a large number of state variables with different dynamics that makes simulating at this level cumbersome. The beam at PEP-II is composed of 1746 physical bunches. The longitudinal dynamics of individual bunches can be modeled, based on energy considerations, by

$$\ddot{\tau}_n + 2d_r \dot{\tau}_n + \frac{\alpha(ev_{\text{rf}}(\tau_s + \tau_n) - U_{\text{rad}}(E_o))}{E_o T_{o\psi}} = 0 \quad (1)$$

$$\text{for } n \neq 1, \dots, 1746, \psi$$

where $\tau_{n\psi}$ is the time deviation of the n^{th} bunch centroid with respect to the arrival time of the synchronous particle τ_s , $2d_r \dot{\tau}_n \approx \dot{U}_{\text{rad}}(E_o)/T_{o\psi}$ is the radiation damping rate, $\alpha\psi$ is the momentum compaction factor, $T_{o\psi}$ is the harmonic revolution period, and $ev_{\text{rf}}(t)$ is the total energy, including wake fields, transmitted to the beam by all the RF stations per revolution period. The goal of the simulation is to study the low-order mode behavior of the beam induced by the interaction with the RF stations. Thus, the particle beam is modeled via a variable number of macrobunches $N\psi$ comparable to the IIR comb filter samples per turn, rather than the 1746 physical bunches. This approach reduces the number of state variables assigned to model the beam dynamics, but allows keeping the same abort gap in the filling pattern and fully resolves all the low-order beam modes and interactions with the RF fundamental impedance.

The energy $ev_{\text{rf}}(\tau_s + \tau_n)$ applied per turn to the n^{th} bunch is the net contribution of all the RF cavities in the ring. The voltage v_{rf} can be expressed by

$$v_{\text{rf}}(\tau_s + \tau_n) = \sum_{i=1}^{ST\psi} \sum_{j=1}^{K\psi} v_{c_{i,j}}(\tau_{s_{i,j}} + \tau_n), \psi \quad (2)$$

where $ST\psi$ is the number of stations, $K\psi$ is the number of cavities per station ($K\psi = 2$ in the LER and $K\psi = 2$ or $K\psi = 4$ in the HER), and $v_{c_{i,j}}$ is the instantaneous voltage corresponding to the j^{th} cavity in the i^{th} RF station. In nominal operation, the cavities per station are detuned by the same magnitude which allows us to group either the two- or the four-cavity station in a unique dynamic macromodel (a 2 or 4 cavity macromodel). This simplification defines the voltage per station as

$$\begin{aligned}
 v_{s_i}(t) &= \sum_{j=1}^{K\psi} v_{c_{i,j}}(t) \leftarrow \text{with } K\psi \neq 2 \text{ or } 4 \\
 &= 2v_{c_i}(t) \text{ or } = 4v_{c_i}(t) \cdot \psi
 \end{aligned}$$

Further simplification in the simulation is possible by considering that in normal operation the voltages of all the stations present almost the same relative phase with respect to the beam. In that case, (2) can be simplified to

$$v_{\text{rf}}(\leftarrow_{s\psi} \quad n) = \sum_{i=1}^{ST\psi} v_{s_i}(\leftarrow_{s\psi} \quad n) = \sum_{i=1}^{ST\psi} v_{s_i}(\leftarrow_{s\psi} \quad n) \cdot \psi \quad (3)$$

These simplifications represent the cavities for all the ring RF stations in two macromodels (macrocavities). All the 4 cavity RF station interactions are lumped into a single 4 cavity macrostation and all the 2 cavity RF stations are similarly modeled via a single 2 cavity macrostation. The reason behind the development of two separate macrocavities is the differences in operation point of a two-cavity RF station and a four-cavity RF station. An example is the tuned resonance frequency of the cavities in each case.

The simulation models the RF signals in baseband and uses the in-phase/quadrature formalism to represent them. Macro-cavities modeled under this formalism are represented by a reduced model defined by

$$\frac{dV_m(t)}{dt\psi} = A(\omega_r)V_m(t) + BI_{\text{kly}}(t - t_w, V_{\text{DC}}) + BI_{\text{beam}}(t),$$

where $[A, \mathbf{B}]$ is the state representation of the cavity, $V_m(t) = [V_{m_I}, V_{m_Q}]^T$ is the in-phase/quadrature macrocavity voltage vector, ω_r is the resonance frequency of the cavity, t_w is the delay of the waveguide between the klystron and the cavities, V_{DC} is the station high-voltage bias, and $I_{\text{kly}}(t, \psi_{\text{DC}})$, $I_{\text{beam}}(t)$ are the in-phase/quadrature klystron and beam current vectors, respectively. As it is depicted in Fig. 1, the important blocks that affect the dynamic interaction between the beam and the RF station are the direct loop, the comb loop, and the group delays associated with the signal propagation around the station. Those blocks are represented in the simulation as discrete blocks. The frequency-dependent elements in the LLRF processing (such as the klystron driver amplifier, the LLRF processing filters with lead/lag networks, etc.) are implemented in the model, as are features which allow nonlinear responses (such as the klystron saturation effects).

Loops in green in Fig. 1 are slow in nature and set the high-voltage power supply magnitude or the tuner position for the cavities. In the time frame in which the simulation characterizes the dynamic interaction between the beam and the RF stations, the changes in the variables controlled by the slow loops are negligible. These slow variables are set via parameters in the simulation, which can be calculated from the initial conditions to define the operation point of the system. Based on these described model simplifications, the simulation complexity is scaled to the




FIG. 2. (Color) Simplified system block diagram. The system transfer function is measured between I (input) and O (output).

minimum required to reproduce the essential physical dynamics. The reduced model described for the system is depicted in the simplified block diagram in Fig. 2. Only blocks in blue are modeled, the components in red/dashed are not modeled.

Thus, the simulation includes the effective impedance presented by all the stations to the beam, representing the collective effect of all the cavities and their feedback loops through the combination of 2 and/or 4 macrocavity stations. Detailed models of klystrons and driver amplifiers, including nonlinearities and the frequency responses, are utilized to analyze the limits in growth rate reduction due to the feedback system and to understand discrepancies between stations. The simulation can be used to predict stability in future operation points, as well as to study the effectiveness of possible additions and modifications to the RF stations.

IV. RF CAVITY IMPEDANCE AND MODAL GROWTH RATES

The effect of the accelerating fundamental RF impedance on the coupled-bunch instability has been studied [9–11]. The dynamic interaction between the beam and the fundamental longitudinal impedance can be quantified by linearizing (1) around the operation point. The longitudinal arrival-time error of the n^{th} bunch centroid becomes

$$\begin{aligned}
 \omega_{s\psi} 2d_{r\psi n\psi} \omega_{s\psi n\psi}^2 &= -\frac{\alpha e\psi}{E_o T_o \psi} \sum_{p=-\infty}^{\infty} \sum_{k=0}^{N-1} q_{k\psi} \\
 &\times W[t_{n,k\psi}^{p\psi} \quad n(t) - \quad k(t - t_{n,k}^{p\psi})], \psi
 \end{aligned} \quad (4)$$

where $\omega_{s\psi}$ is the unperturbed synchrotron frequency, $q_{k\psi}$ is the charge of bunch k , $t_{n,k\psi}^p = (pN\psi - n - k)T_b$, with T_b the bunch spacing (T_o/N), and $W(t)$ is the wake field generated by bunch k and seen by bunch n . It is more useful to express the time deviation of the n^{th} bunch as the phase deviation at the RF frequency, $\phi_{n\psi} = \omega_{\text{rf}} t_n$, and analyze the phase deviations transforming the bunch domain into

the eigenmodal domain. As such, the phase deviation ϕ_n of the n^{th} bunch in the bunch domain can be transformed to the N -filled bunch base by

$$\varphi_l(t) = \frac{1}{N} \sum_{n=0}^{N-1} \phi_n(t) e^{i\psi^{2\pi n/N}}, \quad (5)$$

where φ_l is the phase deviation of the l^{th} mode in the eigenmodal domain. For this analysis, we assume there is no abort gap, but it is possible to include it by suppressing macrobunches. Assuming equal charge for all the bunches and introducing the relation of the wake field to the overall longitudinal impedance Z^{eff} by

$$Z^{\text{eff}} = \int_{-\infty}^{\infty} W(t) e^{i\psi^{\omega t}} dt, \quad \psi$$

following [11], (4) can be simplified to

$$\begin{aligned} \varphi_l = 2d_r \dot{\varphi}_l + \omega_s^2 \varphi_l = i\psi \frac{\alpha e I_0 \omega_{\text{rf}}}{E_o T_o} \\ \times [Z^{\text{eff}}(l\omega_0 - \omega_l) - Z^{\text{eff}}(0)] \varphi_l \end{aligned} \quad (6)$$

where $\omega_{\text{rf}} = 2\pi/T_{\text{rf}}$ is the frequency in the accelerating cavities, I_0 is the average DC beam current, ω_l is the oscillation frequency of mode l , and $Z^{\text{eff}}(\omega)$ is the total effective longitudinal impedance defined as

$$Z^{\text{eff}}(\omega) = \frac{1}{\omega_{\text{rf}}} \sum_{p=-\infty}^{\infty} (pN\omega_{o\psi} - \omega) Z^{\text{eff}}(pN\omega_{o\psi} - \omega), \quad \psi$$

Assuming $d_r \ll \omega_s$ and $\omega_l \approx \omega_s$, the left-hand side of (6) becomes

$$\begin{aligned} \varphi_l = 2d_r \dot{\varphi}_l + \omega_s^2 \varphi_l \approx 2i\omega_s [\dot{\varphi}_l + (d_r - i\omega_s) \varphi_l] \\ = i \frac{\alpha e I_0 \omega_{\text{rf}}}{E_o T_o} [Z^{\text{eff}}(l\omega_0 - \omega_s) \\ - Z^{\text{eff}}(0)] \varphi_l \end{aligned} \quad (7)$$

which can be rewritten as

$$\dot{\varphi}_l - \Lambda_l \varphi_l = 0$$

with

$$\Lambda_l = d_r - i\omega_s - \frac{\alpha e I_0 \omega_{\text{rf}}}{2E_o T_o \omega_s} [Z^{\text{eff}}(l\omega_0 - \omega_s) - Z^{\text{eff}}(0)], \quad \psi$$

$\Lambda_l = \sigma_l - i\omega_l$ is the complex natural frequency. The modal growth rate of the l^{th} characteristic beam mode and the modal oscillation frequency are then given by

$$\begin{aligned} \sigma_l = -d_r + \frac{\alpha e I_0 \omega_{\text{rf}}}{2E_o T_o \omega_s} \mathcal{R}[Z^{\text{eff}}(l\omega_0 - \omega_s) - Z^{\text{eff}}(0)] \\ \omega_l = \omega_s - \frac{\alpha e I_0 \omega_{\text{rf}}}{2E_o T_o \omega_s} \mathcal{I}[Z^{\text{eff}}(l\omega_0 - \omega_s) - Z^{\text{eff}}(0)]. \end{aligned} \quad (8)$$

Equation (8) defines the eigenvalues of the beam dynamics

in the beam modal frame. The effect of the longitudinal effective impedance is evident on the modal damping and the deviation of the synchrotron frequency of the individual modes with respect to ω_s .

The longitudinal effective impedance is determined by the RF cavity impedance and the action of the fast feedback loops at each RF station. The RF cavity impedance per station is modified by the feedback loops as

$$Z_i(\omega) = [I\psi G(\omega)H(\omega)]^{-1} Z_{st_i}(\omega), \quad \psi \quad (9)$$

where $G(\omega)H(\omega)$ corresponds to the return ratio of the station and $Z_{st_i}(\omega)$ is the frequency response of the RF cavities. Adjustable parameters in the control loops and the stations define the frequency response of the system and the stability of the RF feedback loops. The overall station impedance $Z_i(l\omega_0 - \omega_s)$ at frequencies $l\omega_0 - \omega_s$ corresponds to the beam perturbation. This impedance is minimized by optimizing the LLRF station parameters, compatible with stability performance criteria for the RF loops. It is important to recognize that the overall stability of the system is comprised not only of the stability of the LLRF control loops, but also of the beam stability affected by the interaction with the longitudinal impedance of the RF stations.

V. FREQUENCY DOMAIN MODELING

In PEP-II operations, station parameters are configured using a noninvasive method that starts with the identification of the closed-loop transfer function of each station [12]. The motivation behind this method is the inability to measure the open-loop transfer function of the station with beam in the machine, because opening the LLRF control loops causes loss of the impedance control. While it is possible at zero current to measure the LLRF open-loop transfer function, and hence study the closed-loop stability margins, as the machine is filled the RF station dynamics change, since many station parameters vary with the operation point. To best configure the LLRF parameters at operating currents, the closed-loop system transfer function is first measured injecting a complex time domain excitation at the input, as marked in Fig. 2. The time domain response of the station is sampled at the output and recorded. The closed-loop transfer function $H_{\text{meas}}(\omega)$ is estimated using the correlation method based on the measured input/output files.

To obtain a parametric model of the closed-loop transfer function, the transfer function of a linearized model of the station, $H_{\text{model}}(\omega)$, is fitted to the estimated function $H_{\text{meas}}(\omega)$ for the given operation point by adjusting characteristic parameters in $H_{\text{model}}(\omega)$. The linearized model is parametrized by only 8 unknown parameters: cavity resonant frequency ω_r , cavity loaded quality factor Q_l , direct loop gain G_d , direct loop delay T_d , direct loop phase shift ϕ_d , comb loop gain G_c , comb loop delay T_c , and comb loop phase ϕ_c . The model is characterized by only those

parameters because the frequency responses of the lead-lag compensation, comb filter, and the equalizer's finite impulse response (FIR) filter can be accurately modeled from the known hardware implementation. An estimated system response is then derived via least-squares fitting from the model parameters. In this process, some parameters are not adjusted by the fitting routine: $\omega_{r\psi}$ is determined by the average cavity detuning as measured by the tuner position read-back and $Q_{l\psi}$ is set to the nominal value based on the design cavity Q and coupling factor β_0 . The fitting routine is then based on a six-dimensional optimization including a frequency-weighted error function given by

$$f(\vec{x}) = \sum_{\omega, \psi} W(\omega) |H_{\text{meas}}(i\omega) - H_{\text{model}}(\vec{x}, \psi\omega)|^2, \psi \quad (10)$$

where \vec{x} is a vector of 6 optimization parameters and $W(\omega)$ is a weighting function. This is performed for a bandwidth of 1.5 MHz, yielding the 8 parameters that define the linear model. The linearized model for the station $H_{\text{model}}(\vec{x}, \psi\omega)$ in Laplace domain is expressed by

$$H_{\text{model}}(s) = \frac{L_{\text{dir}}(s)}{1 - [L_{\text{comb}}(s)]L_{\text{dir}}(s)}, \psi \quad (11)$$

where $L_{\text{dir}}(s)$ is the open-loop transfer function of the direct loop and $L_{\text{comb}}(s)$ is the open-loop transfer function of the comb loop. These are given by

$$L_{\text{dir}}(s) = \frac{2\sigma\psi}{s^2} \frac{H_{L-L}(s)G_d e^{i\phi_d - (s-i\omega_{rf})T_d}}{2\sigma\psi \omega_{r\psi}^2}$$

$$L_{\text{comb}}(s) = H_{\text{comb}}(s)H_{\text{eq}}(s)G_c e^{i\phi_c - (s-i\omega_{rf})T_c},$$

where $\sigma\psi = \omega_{rf}/2Q_{l\psi}$ is the damping time of the cavity, $H_{L-L}(s)$ is the transfer function of the lead-lag compensation, $H_{\text{comb}}(s)$ is the transfer function of the comb filter, and $H_{\text{eq}}(s)$ is the transfer function of the equalizer FIR filter.

This parametrization of the transfer function allows the calculation of the open-loop transfer function of the station around the actual operation point. With the open-loop estimate, optimal values are calculated for both the direct loop gain and phase rotation parameters (G_d , ϕ_d) and the comb filter gain and phase rotation (G_c , ϕ_c). These optimal values are defined through the specification of open-loop gain and phase margins (consistent with relatively “flat” and controlled frequency response of the closed-loop system). Through this model-based technique, the LLRF systems of the physical machine are configured and studied over the range of high-current operating points, and the RF systems are periodically adjusted and “tuned” in operation.

A similar method is used in the time domain simulation to specify the parameters of the macrostation. To achieve agreement between the simulation and the physical system in the estimation of impedances and growth rates, it is important that the simulation defines an effective impedance interacting with the beam equal to the physical impedance presented by the RF stations to the beam. From (8)

and (9), it is important to observe that this is possible only if there is agreement between the transfer function measured per station and the transfer function and return ratio defined in the simulation. Consequently, there should be close agreement between the linear model parameters fit to the physical station and the linear model parameters fit to the time domain simulation data.

Since the transfer function relationship between model and physical system implies a growth rate consistency, it is reasonable to use the transfer functions for verifying the simulation model. This was done for different operating points, but the analysis below is for the LER at 1400 mA. The transfer function of the simulated station is measured by playing a complex noise file in the input and reading the response at the output as marked in Fig. 2, just as was done for the physical system. The time domain simulation includes estimates of klystron nonlinearities and driver amplifier frequency responses.

In Figs. 3 and 4 the collected data are shown in red for measured (physical station) and simulated (model) transfer functions, respectively. Also shown in green are the fitted 6 parameter linear model responses $H_{\text{model}}(\omega)$ defined above.

From these figures we can clearly see the agreement between the data and the fit, which demonstrates the accuracy of the fitting tools. An important feature of the fitted linear model is its ability to compare the resulting sets of parameters extracted from the physical system and simulation. Their close agreement provides evidence of convergence of the simulation with the physical system. The measured and simulated transfer functions as well as the fitted parameters are very close providing confidence that the growth rates will also be comparable.

These transfer functions, while similar, are not identical. It is apparent from the studies of the multiple RF stations




FIG. 3. (Color) Transfer function and parameters of operating station in LER at 1400 mA.




FIG. 4. (Color) Transfer function and parameters of macrostation from nonlinear simulation in the LER at 1400 mA.

via these transfer functions and models that there are subtle but important variations in the physical RF stations. So the macromodel scheme must represent some sort of weighted average of all the stations to accurately model the beam dynamics. However, this macromodel would not fully capture the limits and dynamics of the LLRF stations themselves, or predict the limits of a particular physical station. The efforts to better understand the small discrepancies between the macromodel and the individual klystron physical transfer functions led to important further measurements of the physical stations which are analyzed later in Sec. VII.

VI. GROWTH RATE MEASUREMENTS

The essential beam dynamic measurements from the simulation are the modal growth rates since these are used to quantify beam stability. The technique for measuring the growth/damping rates in PEP-II operations were first presented by Prabhakar [11] and then refined by Teytelman [13]. Using a related technique, naturally stable modes were studied by injecting a narrowband excitation at the desired mode frequency into the feedback system and observing the resulting decay transients. In PEP-II operations, the damping performance of the operating system is evaluated by opening the longitudinal feedback, letting the unstable beam modes grow for a few milliseconds and then closing that feedback to damp the instability out. In this growth/damp technique, the transient process where the longitudinal loop is open should last a few milliseconds such that the unstable bunch amplitude does not exceed the recapture range of the longitudinal channel. Via transient-domain measurements of the bunches during both steps on the process, it is possible to measure the free growth rate of unstable modes and the overall damping performance of the closed-loop system.

The grow/damp analysis is performed in the beam modal domain transforming the measured beam phase into the even-filled modal base defined by (5). The system dynamics when the longitudinal loop is open is characterized by (6), where the growth/damping rate for the l^{th} beam mode is defined by $\sigma_{l\psi}$ in (8). When the longitudinal feedback loop is closed, the growth/damping rate for the l^{th} beam mode can be defined by

$$d_{l\psi} = -d_{r\psi} - \frac{\alpha e I_0 \omega_{\text{rf}}}{2 E_0 T_o \omega_{s\psi}} \mathcal{R}[Z^{\text{eff}}(l\omega_0 - \omega_s) - Z^{\text{eff}}(0)] \Delta_l, \quad (12)$$

where $\Delta_{l\psi}$ is the effective feedback damping rate due to the longitudinal feedback. This technique allows measuring both $\sigma_{l\psi}$ of the unstable beam modes and $d_{l\psi}$ for the same modes. The first reveals the interaction between the beam and the longitudinal impedance, while the second measures the net damping of the system, quantifying the performance of the longitudinal feedback loop.

To characterize the modal growth rates in the time domain simulation, a procedure similar to the first part of the growth/damp measurement technique is used. From an initial position of the beam near the equilibria, we let the beam naturally evolve in time and study the interaction between the RF station and the beam [6]. The advantage in the simulation is that, due to the absence of the instrumental noise floor and the ability to start with appropriate initial beam conditions, the stable and unstable modes can be estimated concurrently.

The growth and damping rates can be extracted from the transformed time domain data collected through the natural complex frequency Λ_l , whose real part corresponds to the growth/damping rate and the imaginary part to the oscillation frequency, as shown in (8). The natural frequency is therefore fitted to the evolving modes (shown on the left in Fig. 5); revealing stable/unstable modes and identifying the beam mode with the highest growth rate. To achieve consistency between the physical system and the simulation, the same growth rate extraction tools are used to analyze the time domain data in both cases. The simulated growth rates for modes -10 to 10 and their oscillation frequencies can be seen on the right in Fig. 5 for the LER running at 2500 mA. Modes -3 and -4 are usually the most unstable modes; the shift in mode number results from the change in cavity detuning with increasing beam current.

To compare the results from growth/damp measurements performed in the LER at different currents and the simulation, RF stations at the LER and macrostations in the simulation were set with similar parameters and the growth rates were studied. In this case, results correspond to the LER operating with 4 RF stations each running at 1.25 MV, and with beam currents from 1400 to 2500 mA. In these measurements the loop parameters of the RF stations were not set to the optimum values due to imperfections in the




FIG. 5. (Color) Simulated low frequency beam modes (LER at 2500 mA).

klystron driver and LLRF controllers (Sec. VIII). The same linear model is fitted to both the real station and the simulation. In the simulation, several cases were analyzed to set the real klystron nonlinear static transfer function per station. In each of these cases, the klystron nonlinearity and the frequency response of the driver amplifier are included in the macrostation model, and the parameters of the closed loop are set so that the linear model fit to the macrostation is equal to that of the corresponding physical station. Under these conditions, the growth rates are measured in the simulation. This process is repeated in order to evaluate the interaction with the beam dynamics of each individual station in the LER.

Results of this validation are depicted in Fig. 6, where multiple growth rates measured from the physical system are compared with the simulation. The drifting of the growth rates in the physical system (which will be explained in Sec. IX) can be seen in this graph. The individual klystron and driver amplifier characteristics of each station are used as a model in the simulation to calculate growth rates. This process is repeated for each station, and




FIG. 6. (Color) Measured and simulated growth rates for the LER (simulated for each station as well as the macroklystron). The most unstable mode is -3 .

the individual results are depicted in Fig. 6 in green. The macroklystron is defined as the average of these results. The simulation not only reproduces the form of the most unstable growth rates for various beam currents, but it also agrees with the physical system in the number of the most unstable mode. The discrepancy at low currents (underestimation of the growth rates) remains to be better understood via additional dedicated measurements on the physical system.

One application of this simulation is that the free growth rates can then be compared to the expected effective feedback damping rate $\Delta_{l\psi}$ from the longitudinal loop, providing a quantitative measure of stability margins for each mode (in contrast to earlier work [8]). The effective feedback damping rate $\Delta_{l\psi}$ is used as a metric because it is in the first order proportional to the feedback gain, and thus to the beam current, until other effects become dominant at higher currents. The growth and damping rates $\sigma_{l\psi}$ and $d_{l\psi}$ on the contrary have some nonlinear dependence on the beam current. In Fig. 7 we can see an extrapolated line for the effective feedback damping rate with current based on three sets of measurements from the physical machine, as well as the estimated maximum achievable rate. The current estimate for the maximum effective feedback damping rate is -6 to -8 ms^{-1} , but work is in progress to find the exact limit. The difficulty in estimating the limit is related to the fact that it cannot be directly measured in the physical system and that it changes with the system architecture. In this paper the more conservative group-delay limit value of -6 ms^{-1} will be used for estimations and predictions.

The effect of operating with different gap voltage or number of stations at a given beam current was also studied. In this case, the station includes a nonlinear klystron, an ideal driver amplifier, and the closed-loop parameters are set to the optimal condition. The results are shown in Fig. 8. As expected, we see that the growth rate drops as we increase the gap voltage per cavity. This is a result of




FIG. 7. (Color) Plot of Δ_1 with current in mA. The -6 ms^{-1} limit is from the low-mode longitudinal control path.




FIG. 8. (Color) Growth rates for a different number of station and gap voltage for the LER.

direct reduction due to the increase in synchrotron frequency and other changes due to the operation point variations. That can happen either by increasing the total gap voltage, as seen from the first three lines, or as depicted by the first and last line by reducing the number of active stations for the same total gap voltage.

VII. GROWTH RATE SENSITIVITY ANALYSIS

An important use of this simulation tool is the study of the effect of different LLRF parameters on the stability and performance of the LLRF and accelerator system, without requiring time from the operating machine. Additionally, simulations allow analysis of different system configurations and parameter combinations that are not directly applicable to the physical machine without major system changes. These studies have assisted in understanding the




FIG. 9. (Color) Modes -10 to 10 for the nominal case.

sensitivity of the growth rates to certain control loop parameters.

To better understand options for improved performance, the sensitivity on the adjustable parameters is considered, initially assuming an ideal linear model for the direct loop controller, driver amplifier, and klystron. Based on this model, the LER operating at 4.5 MV and 1400 mA was simulated. The loop parameters were adjusted to satisfy the original operational criterium, that is to configure the LLRF direct and comb loops by maximizing the stability margin (gain and phase). The growth rates and synchrotron frequency for modes -10 to 10 for the LER operating in those conditions were estimated and are depicted in Fig. 9. The maximum growth rate, corresponding to mode -4 , is indicated by the nominal value in Table I. The optimal case was modified changing individually the loop parameters to understand their impact in the interaction between the RF station and the beam dynamics. The maximum growth rates resulting from adjusting each of these parameters are shown in Table I.

From these data, it is obvious that the direct and comb loop gains as well as the comb loop delay do not affect the growth rates significantly. However, the direct and comb loop phases do have a significant effect on the growth rates.

TABLE I. Growth rate sensitivity table.

Parameter	Adjustment	Growth rate	Change
Nominal value	...	0.263	
Direct loop gain	20%	0.233	-11%
Comb loop gain	20%	0.221	-16%
Comb loop delay	50 ns	0.258	-2%
Direct loop phase	10°	0.119	-55%
Direct loop phase	-10°	0.408	55%
Comb loop phase	10°	0.106	-60%
Comb loop phase	-10°	0.415	58%




FIG. 10. (Color) Direct loop phase: 10° rotation in green, -10° rotation in red, and the nominal case for reference in blue.

This insight from the model suggests one method of influencing growth rates via adjustments of the direct and comb loop phases. The loop parameters not only influence the interaction of the beam with the RF station, but also affect the intrinsic stability of the station, as defined by $G(\omega)H(\omega)$ in (9).

The deviation in magnitude and mode number can be seen for the direct loop phase in Fig. 10 and for the comb loop phase in Fig. 11 for changes of $\pm 10^\circ$. We can see how the rotation of the phase affects the impedance. The growth rates are reduced in the positive rotation case and increased with the negative rotation. With even larger phase rotations (not shown here) the number of the most unstable mode changes. The margin of variation of the loop parameters to improve the beam stability is restricted by the stability margin of the closed-loop RF feedback.




FIG. 11. (Color) Comb loop phase: 10° rotation in green, -10° rotation in red, and the nominal case for reference in blue.




FIG. 12. (Color) Growth rates with direct and comb loop phase rotation for the LER at 1400 mA.

These studies led to the first application of insight gained from the simulation to the physical system. The impact of the direct and comb phase rotation on the growth rates was studied in the LER. As predicted from the simulation studies, an improvement of machine growth rates is possible by adjusting the loop parameters. The simulation studies showed that the original optimal criterium to maximize the stability of the station feedback loops comes with a tradeoff to the growth rates. Figure 12 shows the effect on the simulated growth rates due to the direct and comb loop phase rotation. It can be seen that the optimal setting for direct and comb loop phase $[0^\circ 0^\circ]$, based on the RF station stability, does not correspond with the minimum growth rate.

We now understand that it is possible to achieve great improvement in the growth rates with a relatively small reduction of the LLRF loop stability margins. The comb loop phase rotation was studied, since it has a smaller effect on the stability margins. Details of the study of the comb loop rotation are summarized in Fig. 13, where the growth rate of the dominant unstable mode is plotted versus the comb filter phase rotation for the LER operating with 3 RF stations at 4.5 MV and 1400 mA. This plot combines simulation results with the average growth rate measured from the LER operating at the same conditions. As for Fig. 6, consistency was achieved by setting the simulation operation point so that its fit to the linear model matched that of the physical system. After the 0° case was checked, the comb phase was rotated by 5° steps in both physical system and simulation. From the resulting figure, the optimal comb phase rotation (based on minimum growth rate) is determined to be between 15° and 20° .

Based on these results, a comb phase rotation has been applied since April of 2006 in the LER RF systems, allowing an increased beam stability margin. Only part of the optimal phase rotation as presented in Fig. 13 was implemented, due to loss of stability margin observed in the closed-loop transfer function of some stations. This




FIG. 13. (Color) Measured and simulated growth rates vs comb phase rotation for the LER at 1400 mA. Agreement in both the general form and most unstable mode number.

effect, which did not allow implementation of the desired amount of rotation, was caused by an unexpected behavior of the driver amplifier transfer function near the carrier frequency. The source of this behavior, studied as a result of the insight from the deviations between the simulation and the physical system, is described in Sec. VIII. The phase offset of about 5° seen in the plot between the physical system and the simulation has not yet been explained adequately. Further measurements will be necessary to determine the cause, possibly imperfections in the LLRF signal processing hardware—the radio-frequency processor module (RFP).

VIII. LLRF SYSTEM IMPERFECTIONS AND IMPLICATIONS

From this work, useful insight was gathered from the performance of the physical PEP-II stations and the development of the simulation macromodels. For example, it was shown that there are important variations in the frequency responses and saturation curves between stations. This is a result of the actual variations in station klystron responses, as well as variations and imperfections in the LLRF electronic systems. These variations in turn have a significant effect on the estimated growth rates.

TABLE II. Growth rate vs station table.

Station	Growth rate (ms ⁻¹)
LER 4-2	0.385
LER 4-3	0.240
LER 4-4	0.211
LER 4-5	0.260
Macroklystron	0.274

This effect can be seen in Table II where the growth rates from LER simulated at 1400 mA are shown. From these data, one can see the dependence of the growth rates on the klystron and driver amplifier model used. The frequency response of LER 4-2 is the furthest from desired, showing onset of unstable loop behavior (due to its far from ideal driver amplifier as will be shown later). Consequently, it has to be operated with a lower direct loop gain (-5 dB) than the rest, leading to the large deviation between its maximum modal growth rate and the average. The variation among stations can also be seen in Fig. 6, for different operation conditions, where the growth rates from 4 different stations have been plotted around the value of the macroklystron. Because of these variations, to best estimate the behavior of the physical system, with multiple individual RF stations, the growth rates have to be computed either for each station and averaged or through calculating the effective impedance from each station and averaging over the whole ring. This way the effective growth rate of a macroklystron that represents the whole ring is calculated.

These variations described above as well as the small discrepancies between the physical system and the simulation as presented in Sec. V prompted further measurements of the klystron transfer functions. In these tests a full-power klystron with low-level driver circuitry was evaluated on a test stand. The first series of tests focused on the nonlinear amplitude saturation characteristic of the power klystrons. The data was used to understand the impact of this klystron nonlinearity on growth rates and develop compensation techniques to correct for this effect, which resulted in the klystron linearizer [14,15]. While there is an impact on growth rates from klystron nonlinearity, it could not explain the magnitude of the deviation from expected behavior. This conclusion, together with additional measurements of the LLRF and the klystron transfer functions, led to an empirical simulation result. After a small 2–3 dB frequency response variation was added in the model klystron (in the form of a small band-pass of increased gain near the center frequency), much better agreement between the model and the physical transfer functions and growth rates was found.

A second series of full-power tests concentrated more closely on understanding the deviations, and the results showed that the apparent large deviations among klystrons were strongly related to the nonlinear behavior of the LLRF system 120 W solid-state driver amplifier. These amplifier functions were specified and tested for frequency response and gain uniformity in the initial development of the RF stations. However, their large-signal behavior was measured, not a small-signal measurement in the presence of a large-signal carrier. In operation, the RF station must deliver a large RF power at the carrier (ring operation frequency), but still pass faithfully the small modulation signals within the bandwidth of the direct and comb loops.

It is these small modulation signals which serve to achieve the impedance control.

As can be seen from the large- and small-signal transfer functions of the driver amplifier in Fig. 14, when the amplifier is driven by just the RF carrier the transfer function is almost flat and the amplifier behavior unremarkable. However, for the two tone case of a large carrier combined with a small test signal (the way the amplifier is driven in normal operations), there is some unusual behavior around the carrier frequency. This small-signal transfer function distortion is very similar to the empirical result, and highlights the value of the simulation model in understanding the behavior of the physical systems.

Additional tests of several driver amplifiers installed in the LER RF stations showed that the level of the distortion

varies from station to station. More significantly, LER 4-2, the station with the most deviation from ideal response, is seen to have the most distorted small-signal gain characteristic [16].

With the simulation model we can predict that improving these amplifier responses will lead to a direct decrease of the growth rates. It will also allow applying the full optimal comb rotation, thus further decreasing the growth rates (as presented in Sec. IX).

These results highlight that the simulation is more than a tool to imitate the physical system. It can help diagnose small imperfections and nonidealities of implementation in the LLRF systems, since it compares the physical system with the expected behavior through several physical measurements. As such, it can identify the necessity for updates or modifications to the LLRF implementation, and help evaluate possible modifications to the systems.

IX. PREDICTIONS OF HIGH-CURRENT OPERATIONAL MARGINS AND SYSTEM LIMITS

In Sec. VI, the effective damping rate $d_{i\psi} = d_{i\psi} - \sigma_{i\psi}$ was defined. A comparison of $d_{i\psi}$ with the maximum modal growth rate $\sigma_{i\psi}$ can be used to determine stability. Mathematically, for stability the damping rate provided by the system $d_{i\psi}$ has to be larger in absolute value than the growth rate $\sigma_{i\psi}$ (for equality the exponential envelope of the beam motion would simply be a fixed amplitude of zero time constant). Running the beam instability feedback has provided us with experience with the practical and operational limits of control. According to this experience, the effective damping provided by the system should be roughly twice the growth rate, so that the absolute values of the growth rate $\sigma_{i\psi}$ and the damping rate $d_{i\psi}$ are almost equal. The rationales for this factor of 2 operational margin include several arguments. First, the growth rate does not have an exact value, but fluctuates around an average value as various system parameters drift around a controlled value (e.g. in the physical system there is power supply ripple, other perturbations which modulate the system effectiveness). There are variations in the woofer feedback system gain due to gap transient effects, and other system factors. Our experience operating these systems suggests that when the empirical limit is crossed there is an increasing probability of losing control of the beam. Therefore, our predictions are not for hard limits; rather they are operating points past which it is increasingly difficult to operate the stations and maintain control.

Another important limitation for the LER is the available klystron power. Currently, the implemented maximum power of an LER klystron is of the order of 1 MW. For an operating gap voltage of 4.5 MV, the klystron power delivered exceeds 1 MW when the beam current approaches 3700 mA. However, even with the higher gap voltage of 5.4 MV (and with the associated implications for reduced bunch length), the beam current cannot exceed 3800 mA




FIG. 14. (Color) Driver transfer functions. (a) Driver amplifier transfer function driven by carrier (large-signal response). (b) Driver amplifier transfer function driven by carrier and modulation (small-signal response). The modulation is swept across the band at a level -30 dB below the power carrier at 476 MHz.

before the klystron power delivered goes over 1 MW. At this power region of it is difficult to tune the stations in the simulation, for the type of klystrons installed in the LER. Since this difficulty arises for a somewhat ideal system as the simulation, we expect worse behavior from the physical system. Therefore, in this paper the klystron power delivered is set to a maximum of 1 MW in our estimates of system limits.

The consistency of the error margin between measured and estimated growth rates at achieved currents gives a certain confidence in predicting the growth rates of the machine at higher currents. As part of this prediction, we exploit the advantage of the simulation to test several possible high-current configurations.

In Fig. 15, growth rates are plotted in red for the LER with the operating RF station configuration of 4.5 MV versus current. This case is the same and consistent with measured growth rates depicted in Fig. 6. It should be noted that, as the beam current is increased in the simulation, the voltage of the high-voltage power supply is also increased to match that of the physical system for similar operating points, up to the limit of 82 kV. The black line presents the same operating conditions, but assuming that the LLRF, the driver amplifier, and the klystron are ideal and the loop parameters are adjusted to the optimal conditions. This case is consistent with the system conditions of the nominal case in Fig. 9. In magenta, the same system configuration as the red line of 4.5 MV are used, but with a higher gap voltage of 5.4 MV. In blue the 4.5 MV system configuration is analyzed by using improved driver amplifiers, with flatter small-signal gain, which allows setting optimal loop parameter conditions. This later case allows us to also apply a 10° comb phase rotation, with further improvements shown in the figure in green. Here a conservative rotation is limited to 10° due to RFP module distortion. From this figure, the big reductions in growth rates related with the two possible improvements (driver amplifier re-




FIG. 15. (Color) Estimated growth rates. Values up to 2500 mA validated for LER with the physical system as shown in Fig. 6.




FIG. 16. (Color) Plot of $i_{\psi} 2\sigma_{i_{\psi}}$ for LER.

sponse and comb loop operating point rotation) are apparent.

A quantitative estimate of net controlled margin can now be made, by comparing these numbers to the expected damping rates from Fig. 7 to determine stability. Since the margin has been set from the empirical (margin of two) limit, in Fig. 16 the sum $i_{\psi} 2\sigma_{i_{\psi}}$ is plotted. When the sum is equal to zero, the absolute values of the damping $d_{i_{\psi}}$ and growth $\sigma_{i_{\psi}}$ rates are equal, thus reaching our empirical limit. From this figure we see that for the nominal RF station configuration (red curve) the empirical margin is at 3150 mA. Therefore, for any increases of beam current beyond 3150 mA we expect that the existing physical system will be hard to maintain in operation. On the same figure, one can see that if the physical system is instead run at the higher gap voltage of 5.4 MV (magenta curve), there is an improvement as expected from the analysis in Sec. VI. However, the empirical margin is still well below 4000 mA, with an approximate value of 3800 mA. An exact value is difficult to calculate since the klystrons at this point are run close to our power limit of 1 MW, thus deteriorating the ability to tune and operate the stations. In the same figure, the curves for a system with improved driver amplifiers and implemented comb rotation are also shown. The growth rates are significantly reduced at this point, showing that with these updates, and the installation of 1.2 MW SLAC-type power klystrons, the physical system could exceed 4000 mA. This power level has been reached on the test stand in absence of beam. Also, the 1.2 MW SLAC-type klystrons installed in some of the HER stations have been operating up to 1 MW without limiting RF station operation.

In terms of the earlier work on the klystron linearizer, we should note that the black line shown in Fig. 15 is for an ideal amplifier/perfectly linear klystron, whereas the system built for evaluation was solely an amplitude linearizer [15]. The distortion caused by the imperfections of the driver amplifier, as described in Sec. VIII, contains signifi-

cant phase distortion which could not be corrected by the prototype linearizers. The effect of this added phase distortion has not been evaluated by itself. A further study with the model may provide further insights into the incremental effect of this phase distortion as well as the potential improvement in the linearizer behavior when applied to a system with improved driver amplifiers (less phase distortion). Even though the prototype linearizers did not provide a decrease in growth rates, they did provide extra gain margin in the closed-loop LLRF system. This extra gain margin may be beneficial as currents are pushed towards 4000 mA in the LER.

X. FUTURE DIRECTIONS

The model validation will be continued for the HER via estimates of current limits and study of expected growth rates. The detailed study of transfer functions of the physical system must be continued to fully determine the small discrepancies and explain the remaining unexpected behaviors. Future possible upgrades of PEP-II (e.g. a new asymmetric comb loop filter) can be implemented in the simulation and their impact on stability studied. The simulation tool offers a path to evaluate the necessary noise performance and implementation possibilities of new system implementations. The pragmatic goal is to first quantify and resolve all nonidealities and problems, including quantifying longitudinal stability issues, and then move in new directions.

One important purpose of this simulation tool is to understand the impact of imperfections in the LLRF processing functions, and understand how possible implementations, with various noise mechanisms, imperfect channel isolation, and other nonidealities, affect the overall system performance. This tool is of great value in understanding future LLRF implementation options: for example, the impact of quantizing noise in an all-digital implementation can be studied, and the necessary arithmetic resolution of digital processing specified. Similarly, the impact of known imperfections can be understood via the model, which can compare possible alternate implementations and can help determine a useful path to system improvements. This effort can be useful to future accelerators as well as to PEP-II.

It is important to note that all the analysis formalism and measurement tools presented up to and including Sec. V are not limited to PEP-II, but can be applied to other high-current accelerators. The different design constants and characteristics are formed into a structure which is called by the simulation, making a transition between machines relatively straightforward.

XI. CONCLUSIONS

The simulation of the PEP-II RF system is a close representation of the actual system. As such, it can predict

the performance limits of the LLRF systems at higher currents and study the effectiveness of upgrades or their optimal configurations. It also provides insight into subtle behaviors of the system and suggestions for optimal tuning (as with the comb rotation). The simulation model was very helpful in obtaining insight into the effect of the variations in the klystron responses and system imperfections. Another aspect of the simulation is the ability to separate the stability of the particle beam from that of the LLRF, and study various possible operating points and upgrades.

One of the most important features of this tool is the adaptability to simulate the interaction between the RF stations and the beam for other systems and accelerators. The interaction between the different parts of the algorithm and SIMULINK is through a parameter structure, which can be easily modified. Thus, the simulation was easily adapted to be used for modeling SPEAR to study Robinson instability and it was also modified to help with the design of the klystron linearizer [14].

ACKNOWLEDGMENTS

The authors want to thank the PEP-II operations and accelerator groups for their consistent help with these measurements, and Sam Heifets, the SLAC Accelerator Research and Technology department, Beam Physics department, and Stanford Synchrotron Radiation Laboratory's (SSRL) accelerator physics department for many helpful collegial discussions which helped focus this work. We also thank the SLAC klystron department and staff who helped with the full-power klystron test stand measurements. We acknowledge our colleagues at CERN and KEK who share their ideas and interests in these RF impedance control architectures. The original efforts by Paul Corredoura and Rich Tighe in these areas provided our group with a substantial base to build on. We would like to thank Flemming Pedersen for providing the initial ideas on the subject and for his continuing work. This work was supported by DOE Contract No. DOE-AC03-76SF00515.

-
- [1] F. Pedersen, in Proceedings of B Factories—the State of the Art in Accelerators, Detectors and Physics, Stanford, CA, 1992, pp. 192–207.
 - [2] B. R. Cheo and S. P. Jachim, *IEEE Trans. Electron Devices* **38**, 2264 (1991).
 - [3] P. B. Wilson, in Proceedings of the Ninth International Conference On High Energy Accelerators, Springfield, 1975, pp. 57–62; SLAC Stanford, SLAC-PUB-1456, 1974.
 - [4] P. Corredoura, in Proceedings of the IEEE Particle Accelerator Conference, 1999; SLAC-PUB-8124, 1999.
 - [5] D. Teytelman, PEP-II Low-level RF Tutorial, <http://www.slac.stanford.edu/~dim/lecture1.pdf>, 2003.

- [6] D. Teytelman, D. Van Winkle, and J. Fox, in Proceedings of the IEEE Particle Accelerator Conference, Knoxville, TN, 2005; SLAC-PUB-11253, 2005.
- [7] SIMULINK, The Mathworks, Natick, MA 01760, USA.
- [8] R. Tighe and P. Corredoura, Proceedings of the IEEE PAC, Dallas, TX, 1995; SLAC-PUB-950-6856, 1995.
- [9] A. Chao, *Physics of Collective Beam Instabilities in High Energy Accelerators* (John Wiley & Sons, Inc., New York, 1993), Chap. 4
- [10] F. Pedersen, IEEE Trans. Nucl. Sci. **32**, 2138 (1985).
- [11] S. Prabhakar, R. Claus, J. Fox, H. Hindi, I. Linscott, J. Olsen, W. Ross, and D. Teytelman, Part. Accel. **57**, 175 (1997).
- [12] D. Teytelman, in Proceedings of the IEEE Particle Accelerator Conference, Knoxville, TN, 2005; SLAC-PUB-11252, 2005.
- [13] D. Teytelman, Ph.D. thesis, Stanford University Electrical Engineering Department, 2003; SLAC-Report-633, 2003.
- [14] J. Fox, T. Mastorides, D. Teytelman, D. Van Winkle, and Y. Zhou, in Proceedings of the IEEE PAC, Knoxville, TN, 2005; SLAC-PUB-11228, 2005.
- [15] D. Van Winkle, M. Browne, J. Fox, T. Mastorides, C. Rivetta, and D. Teytelman, in Proceedings of the European Particle Accelerator Conf, Edinburgh, Scotland, 2006; SLAC-PUB-11945, 2006.
- [16] D. Van Winkle and C. Rivetta, SLAC Internal Note.