ARCHITECTURES AND ALGORITHMS FOR CONTROL AND DIAGNOSTICS OF COUPLED-BUNCH INSTABILITIES IN CIRCULAR ACCELERATORS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Dmitry Teytelman
June 2003

Abstract

Modern light sources and circular colliders employ large numbers of high-intensity particle bunches in order to achieve high luminosity. The electromagnetic coupling of bunches via resonant structures causes coherent instabilities at high beam currents. Achieving high luminosity requires the control of such unstable motion. Feedback control is challenging due to wideband nature of the problem with up to 250 MHz bandwidths required. This thesis presents digital signal processing architectures and diagnostic techniques for control of longitudinal and transverse coupled-bunch instabilities.

Diagnostic capabilities integrated into the feedback system allow one to perform fast transient measurements of unstable dynamics by perturbing the beam from the controlled state via feedback and recording the time-domain response. Such measurements enable one to thoroughly characterize plant (beam) dynamics as well as performance of the feedback system.

Beam dynamics can change significantly over the operating range of accelerator currents and energies . Here we present several methods for design of robust stabilizing feedback controllers. Experimental results from several accelerators are presented.

A new baseband architecture for transverse feedback is described that compactly implements the digital processing functions using field-programmable gate array devices. The architecture is designed to be software configurable so that the same hardware can be used for instability control in different accelerators.

Acknowledgments

First I would like to thank my adviser, Prof. John D. Fox for introducing me to the field of instability control and for skillfully guiding me to the solutions of many exciting and difficult problems. I have enjoyed immensely working with John, be it in the office, lab, or control room. I am also thankful to John for infecting me with the Italian car hobby.

Next I'd like to thank two former graduate students of John, Haitham Hindi and Shyam Prabhakar. For many years, both before and long after his graduation, Haitham has been an invaluable resource for all sorts of control theory and numerical optimization questions. His ability to quickly grasp the salient features of the problem and to come up with attack directions, as well as his intuition helped me numerous times to better understand the issues in question and visualize possible solutions. Shyam Prabhakar designed the core of the data analysis algorithms used in this work. Working in the same office with him was akin to having a perfect reference for all beam dynamics questions. With Shyam, even the late-night machine study shifts were enjoyable.

My graduate studies were supported by Stanford Linear Accelerator Center (SLAC). Participation in B Factory collaboration provided me with many interesting and challenging problems as well as priceless practical experience. I am grateful to Accelerator Research Department of SLAC for the wonderful freedom to pursue diverse research topics. Working at SLAC I enjoyed interacting and collaborating with many people. One of the most rewarding aspects of work at SLAC has been the wide diversity of projects and research directions.

I would also like to thank Ivan Linscott of the STAR lab for insightful comments and for valuable discussions of both hardware design and signal processing topics.

I feel that meeting Mario Serio of Laboratori Nazionali di Frascati early on in my graduate career was crucial for my understanding of the coupled-bunch instabilities as well as of a bigger picture of accelerator physics. I am thankful for his advice to pay attention to the physics of the process, not just the engineering side.

Alessandro Drago has written the first versions of many of the DSP routines used in the longitudinal feedback system. Having his code to work with has simplified implementation of many real-time processing tasks. Most of the DAΦNE measurements presented in this thesis would not be possible without him. I would also like to thank Alessandro for his friendship and help in arranging visits to Frascati, as well as for serving as a guide/translator during these visits.

I have been lucky to have the opportunities to travel to accelerators around the world. These visits allowed me to perform advanced beam dynamics measurements as well as test many new signal processing, data analysis, and control ideas. Visits to Laboratori Nazionali di Frascati encouraged by Mario Serio and Alessandro Drago have been both educational and productive, resulting in many measurements of longitudinal unstable motion in DAΦNE collider and several control filter designs.

A visit to BESSY-II synchrotron light source provided an excellent opportunity for studying both longitudinal and transverse instabilities in an intensive and productive sequence of all-night machine shifts. I would like to thank Shaukat Khan for helping to arrange the visit and for his active participation the the measurements.

I would like to thank John Byrd and Greg Stover of the Advanced Light Source for their support in the long process of control filter development. I am grateful for all the beam time for instability studies and feedback control testing. Without these opportunities many of the techniques described in this thesis would have never reached maturity.

I'd like to thank the operations groups of ALS, BESSY-II, DAΦNE, and PEP-II for their help and patience during numerous beam studies.

None of the experimental measurements presented in this work would be possible without the longitudinal feedback system designed by SLAC/LBL/LNF-INFN collaboration. And this system would not exist without engineering wizardry of Jack Hoeflich, Jeff Olsen, Gerard Oxoby, Bill Ross, Leonid Sapozhnikov, and Andrew Young. I have been lucky to help Ric Claus with system software development, learning a lot in the process. Success of our hardware projects was guaranteed by our technicians Dave Anderson and Boni F. Cordova-Grimaldi.

I have learned a great deal from discussions of technology and physics of accelerator feedback systems with Paul Corredoura, Alessandro Gallo, Sam Heifets, Boris Podobedov, Rich Tighe, Makoto Tobiyama, and Mikhail Zobov.

I would like to thank my mentors during my undergraduate years - Professors Alfred M. Levine, Erlan H. Feria, and Ercument Ozizmir of the City University of New York, College of Staten Island for laying the foundation for my graduate studies and research work at Stanford. I am especially grateful to Professor Levine for taking his time to talk to me late in the summer of 1990 which resulted in me starting the studies at the CSI in the fall.

I am thankful to my parents, Boris and Raisa, for their support of my educational career, their lifelong love and encouragement. I thank my brother Leonid for always being there for me, for his support, attention, and love. My son Lev has also helped during the last year by allowing me sufficient sleep which, in turn, enabled me to finish the dissertation. I am deeply grateful to my dear wife Eugenia for her companionship, love, and empathy. This thesis is dedicated to her.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

In this work we will consider electron and positron circular accelerators and storage rings. These accelerators have many applications ranging from high-energy physics to materials science to protein crystallography. In all cases particle beam intensity is critically important. Coupled-bunch instabilities are an important limitation to increasing beam currents. Such instabilities exist in each of three possible beam motion coordinates. When the beam current crosses the instability threshold the noise-driven position or energy oscillations grow to large amplitudes. Large motion amplitudes often cause beam loss due to the particles encountering some aperture within the machine. Sometimes the growth of instability may saturate and the beam may stay in the ring. However, large oscillations degrade the performance of the machine whether it is the luminosity of a collider or the brightness and spectral purity of a synchrotron light source.

Initially the problem of coupled-bunch instabilities was solved by keeping beam current below the instability threshold. The next step in control of coupled-bunch instabilities was to apply active stabilization (feedback) to the most prominent instability eigenmodes thus enabling one to operate the machine above the threshold current. This approach was sufficient when the number of unstable eigenmodes was relatively small and the accelerator operated not far above instability threshold. With time the number of individual bunches of charged particles increased and with it the number of unstable eigenmodes. Modern storage rings are designed to operate far above the threshold - the Advanced Light Source at LBNL has design current of 400 mA while the instability threshold is 40 mA. In such accelerators feedback is commonly used to extend the thresholds by 10-20 times. At some point per-mode feedback became unpractical. Due to many unstable eigenmodes modern machines require all-mode feedback stabilization to operate at design currents.

This dissertation presents new methods for control of longitudinal coupled-bunch instabilities. The work is organized to logically present the feedback control development cycle. This cycle starts from developing an analytical model of the plant and the feedback hardware as presented in Ch. 2. In the next chapter (Ch. 3) methods for measuring open and closed-loop behavior of the system and techniques for extracting model parameters from such measurements are described. These techniques enable one to characterize in detail the unstable coupled-bunch eigenmodes as shown in Ch. 4. In order to accurately model the closed-loop behavior one needs to know both the plant (the beam) and the feedback system. Chapter 5 introduces techniques for thorough measurement of the feedback channel properties. Taken together, chapters 4 and 5 provide the necessary framework for creating an accurate off-line system model. Such a model provides a way to test and parameterize feedback controller designs paving the way for the optimization-based controller design methods described in Ch. 6.

Finally, Chapter 7 describes a new signal processing hardware architecture which would provide processing and diagnostic capabilities needed to extend the feedback control methods shown in this thesis to transverse coupled-bunch instabilities.

In the following sections the work will be introduced chapter by chapter and the new contributions will be identified.

1.1 Coupled-bunch instabilities and feedback control

This chapter introduces the basic concepts of storage rings and longitudinal beam dynamics. The qualitative and quantitative descriptions of longitudinal focusing and particle motion to large extent follow those developed by M. Sands in "The Physics of Electron Storage Rings: an Introduction" [1]. Small motion of the bunch centroid near the equilibrium is shown to be that of a harmonic oscillator. Next, the effect of bunch-to-bunch coupling on this motion is described and the commonly encountered driving impedances are presented.

Changes in the accelerator parameters as well as interactions with the RF systems affect both the beam motion parameters and its dynamics. The most important dynamics modifiers are discussed in Sec. 2.5.

Next the coupled-bunch instabilities and bunch-by-bunch feedback control are analyzed as a multi-input multi-output (MIMO) dynamic system. From this analysis we develop important model reductions which enable one to predict the behavior of an $N \times N$ MIMO system where N ranges from 120 to 1746 using a single-input single-output (SISO) model.

Finally, this chapter presents a detailed description of a programmable bunch-by-bunch longitudinal feedback system developed by a multi-laboratory collaboration and currently in use at 5 accelerators. Such systems were used to obtain all of the measurements presented in this thesis as well as to implement the novel feedback control algorithms developed in this work.

1.2 Transient diagnostics

In this chapter beam diagnostic techniques based on transient motion recording are described ranging from steady-state (closed-loop) data recording to externally excited grow/damp experiments.

The main new contribution presented here is the estimation of the complex eigenvalue from the exponentially growing or decaying oscillatory motion. Before this work there existed methods developed by S. Prabhakar for estimating the exponential growth or damping rates by considering the envelope of the motion [2] and for estimating the oscillation frequency from phase-space trajectories [3]. Since these extract the real (growth or damping rate) and the imaginary (oscillation frequency) parts of the eigenvalue separately information is lost in the process. Using only the amplitude of the transient to estimate the growth rate discards the phase information. Similarly, extracting oscillation frequency from the phase angle of the transient disregards the amplitude. The new method utilizes joint estimation by fitting growing or damping oscillatory waveform to the data. This technique uses the bandpass nature of the harmonic oscillator response to filter out noise. Consequently, reliable eigenvalue estimation is made possible when the noise levels are high.

Growth and damping measurements made at high beam currents - and showing the fastest instability growth rates as well as the fastest feedback damping - are most important in the feedback system characterization and controller design. In a coupled-bunch feedback system actuator saturation restricts the controllable range of beam motion amplitudes. High feedback gains commonly used result in the limit being relatively close to the system noise floor (2 to 3 times). Combined with fast exponential growth and damping rates this means that most important measurements are made at the high noise levels and benefit greatly from the new eigenvalue estimation method.

1.3 Beam dynamics measurements

Here we develop techniques for beam dynamics parameterization using the eigenvalue measurement technique described in the previous chapter. The most basic characterization approach is that of eigenvalue locus which maps out the positions of modal eigenvalues as a function of beam current. Such mapping is a valuable guide for the feedback designer since it defines the performance requirements for the feedback system. Furthermore, combining the eigenvalue locus with the feedback loop model developed in the next chapter allows one to predict the closed-loop behavior of the system.

In this chapter, in addition to measuring the dependence of the eigenvalues on the beam current, other parametric dependencies are explored. Using the eigenvalues extracted from the growth measurements as well as from the other transient diagnostics we illustrate methods for estimating parameters of the higher-order mode impedances. Such estimation is very important from the longitudinal feedback point of view for several reasons. Knowing the dependence of the driving impedance on environmental and operating factors such as cavity temperature gives user a way to minimize the instability growth rates. Such knowledge also helps extrapolate the measured growth rates to the new operating regimes.

1.4 Feedback loop characterization

In chapter 4 we concentrated on the open-loop part of the transient diagnostic measurements. In this chapter we use both the open and closed-loop measurements to verify and refine a linear discrete-time model of the longitudinal feedback system.

Several important new contributions are presented here. A thorough loop model for the longitudinal feedback system in the eigenmode coordinate system introduced here is critical for development of control algorithms. However any model of a physical system must be verified to insure the agreement between the system and its model. In this chapter a quantitative method is developed for comparison of the feedback system model and the real system. The method provides a measure of error between the model and the system which, in turn, can be used to adjust the model for better matching.

In any physical system one unavoidably encounters parasitic effects. This is especially true for a bunch-by-bunch feedback system with its broadband processing. In this chapter mathematical analysis of the most important parasitic effects is presented and substantiated by the beam and feedback measurements.

1.5 Longitudinal feedback controllers

This chapter presents the most important contribution of the whole thesis - methods for designing longitudinal feedback controllers. We should not forget that the feedback system is there to reliably stabilize the beam, not just collect beam diagnostic data. From that point of view the work presented in the earlier chapters was done to enable one to design practical feedback controllers.

Two methods for controller design are presented in this chapter: frequency-domain design and model-based. Both of these approaches use numeric optimization to generate a feedback controller filter.

Applications of these methods are presented as several case studies for three different accelerators: the Advanced Light Source (ALS), DAΦNE collider in Italy, and BESSY-II light source in Germany. Each of these machines presents a unique feedback control problem. In each case the performance requirements established in beam and feedback loop characterization are posed as a controller design problem and feedback filters are generated. All of these designs have been tested in their respective accelerators and test measurements are presented to confirm the expected operation.

1.6 Transverse baseband processing architecture

Baseband feedback signal processing architecture presented in Ch. 2 is limited to control of longitudinal coupled-bunch instabilities. However modern storage rings require active wideband feedback in all three dimensions: two transverse and one longitudinal. A universal programmable architecture with diagnostics has been proven to be an excellent coupled-bunch instabilities solution. Here a new design is presented which significantly extends the capabilities of the present system. In this chapter algorithms for data demultiplexing are presented as well as possible control filter and diagnostic architectures.

Chapter 2

Coupled-bunch instabilities and feedback control

Basic concepts related to the longitudinal beam dynamics and coupled-bunch instabilities are introduced in this chapter. We start from describing the general layout of a storage ring. Next, concepts of longitudinal focusing and bunching are introduced. Equations governing the motion of the bunches of particles are shown as well as the effects of bunch-to-bunch coupling via electromagnetic fields. Next the driving terms of the longitudinal coupled-bunch instabilities are discussed in Sec. 2.3. Interaction of the accelerating RF cavity and the beam plays an important role in determining longitudinal beam dynamics and is presented in Sec. 2.4. Next, the most important factors that influence the longitudinal beam dynamics and motion parameters are presented in Sec. 2.5. In Secs. 2.6 and 2.9 the linear system models are developed for the longitudinal coupled-bunch instabilities and practical feedback topologies are defined. Finally, in Sec. 2.10 the design of the bunch-by-bunch digital programmable feedback system is presented.

2.1 Circular accelerators

A block diagram of a typical storage ring is shown in Fig. 2.1. Bunches of charged particles injected into a storage ring are steered into "circular" orbit by the magnetic guide field. The guide field has focusing properties in the transverse plane defined by radial or horizontal axis and the vertical axis. The focusing leads to transverse (betatron) oscillations about the design orbit. Charged particles radiate energy under radial acceleration - so called

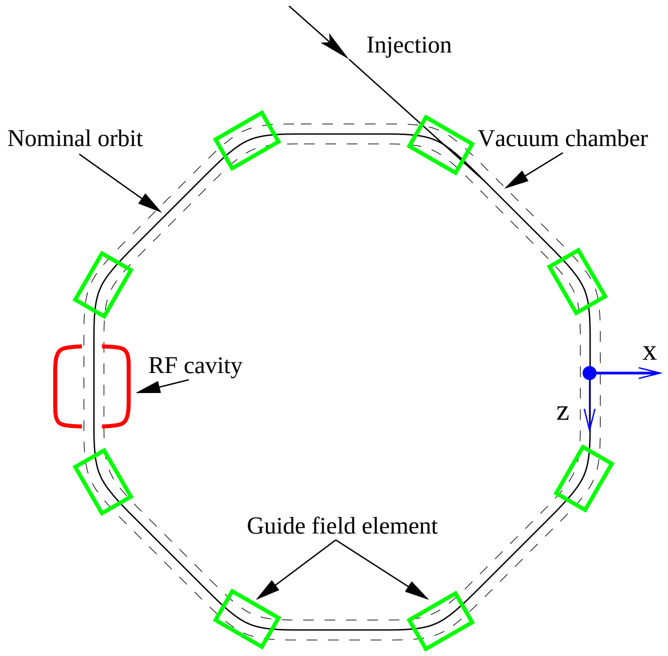

Figure 2.1: Block diagram of a storage ring

synchrotron radiation. In each revolution particle loses some of its energy by this process. This energy loss is compensated in one or more radio frequency (RF) cavities. In addition to restoring the energy lost to the synchrotron radiation periodic accelerating field in the cavities creates longitudinal focusing.

Let us first qualitatively examine the longitudinal focusing effect. In Fig. 2.2 the electric field in the accelerating cavity is shown as a function of time. Times of arrival of bunches of particles are also marked. Dashed line shows the energy U_0 lost by a particle at the nominal energy E_0 in a single turn. Particles passing RF cavity at time τ gain energy $E(\tau) = eV(\tau)$ where e is the charge of the electron. There are two time points within the RF period where particles will gain exactly the energy lost to synchrotron radiation. Particles at the design energy complete a revolution in time T_0. The RF frequency is chosen in such a way that there is an integer number (called harmonic number h) of RF periods $T_{\rm rf}$ in one revolution, i.e.

$$T_0 = hT_{\rm rf}$$

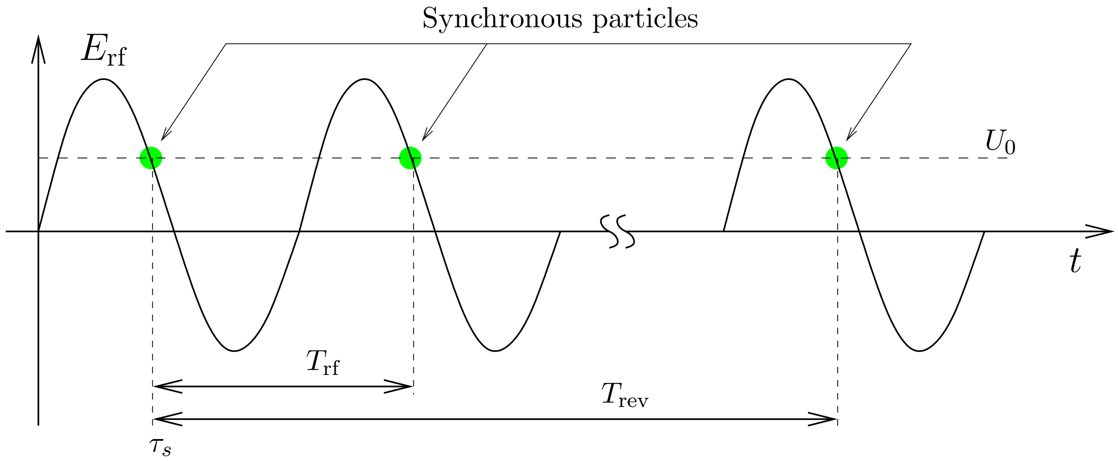

Figure 2.2: Electric field in the RF cavity as a function of time. Times of arrival of synchronous particles are marked by green circles.

As a result, a particle at the nominal energy will gain exactly U_0 on every turn. Without external perturbations the particle would continue sampling RF waveform at T_0 intervals each time gaining exactly the energy lost during that revolution. Perturbations of longitudinal position or energy lead to particle gaining in the cavity more or less energy than U_0. Out of the two time points where energy gain is U_0, only the one with the negative slope of the accelerating field is the stable point. That is the shape of the RF potential near the stable point creates a restoring force for the particle with time of arrival (longitudinal position) error and causes the particle to arrive closer to the stable point in the consecutive turn. That stable point is marked on Fig. 2.2 as τ_s, so called synchronous time. Particles that pass through the cavity at times $\tau_s + kT_{\rm rf}$ are called synchronous particles. The maximum number of the bunches of synchronous particles that can be stored in the ring is given by the harmonic number h.

Let us examine what happens if a particle arrives earlier than the synchronous time. It will gain more energy than it lost in one turn. Highly relativistic particles (above the transition energy) at higher than nominal energy take a longer path around the ring [1, page 73]. Thus, with each revolution, its time of arrival will shift closer to the synchronous time. Similarly, a particle arriving late will take a shorter path and will get closer to the synchronous time as well. Slope of the accelerating field at τ_s creates a potential well causing the particles to execute synchrotron oscillations around the synchronous time. We will call this potential well the RF bucket in which particles can be stored.

Longitudinal focusing groups injected particles near synchronous time points creating bunches of particles. The harmonic number of the ring defines how many synchronous points exist within a single revolution and, therefore, how many bunches can be stored in the ring.

Now we will derive the equations of motion for synchrotron oscillations. Let ϵ be the energy deviation of a particle from the nominal energy E_0. Energy lost by a particle to radiation in one turn is a function of its energy $U_{\text{rad}}(E_0 + \epsilon)$. For small energy deviations we can write

$$U_{\text{rad}}(E_0 + \epsilon) = U_{\text{rad}}(E_0) + \epsilon \dot{U}_{\text{rad}}(E_0) = U_0 + \epsilon \dot{U}_{\text{rad}}(E_0) \tag{2.1}$$

Energy deviation of a particle is related to the change in the path length l around the ring by the momentum compaction factor α as follows [1, page 75]

$$\frac{l}{L} = \alpha \frac{\epsilon}{E_0}$$

where L is the path length for a particle of nominal energy.

Let us observe the motion of the particles at a fixed point in the ring. We will start with a particle that has arrives time τ_1 later than the synchronous particle. On the next turn its time of arrival error will change depending on its energy deviation. For energy deviation ϵ path is lengthened by l. Thus

$$\tau_2 - \tau_1 = \frac{l}{c} = \alpha \frac{\epsilon}{E_0} T_0$$

where c is the speed of light. This change in the time of arrival has been observed in one revolution therefore the rate of change in τ is

$$\frac{d\tau}{dt} = \alpha \frac{\epsilon}{E_0} \tag{2.2}$$

Now consider the energy change of the above particle over one revolution:

$$\delta\epsilon = eV_{\mathrm{rf}}(\tau_s + \tau) - U_{\mathrm{rad}}(E_0 + \epsilon)$$

Again, this change is observed over one turn, so that

$$\frac{d\epsilon}{dt} = \frac{eV_{\mathrm{rf}}(\tau_s + \tau) - U_{\mathrm{rad}}(E_0 + \epsilon)}{T_0}$$

Linearizing RF voltage near $V_{\mathrm{rf}}(\tau_s)$ and substituting Eq. 2.1 we get

$$\frac{d\epsilon}{dt} = \frac{e\tau\dot{V}_{\mathrm{rf}}(\tau_s) - \epsilon\dot{U}_{\mathrm{rad}}(E_0)}{T_0} \tag{2.3}$$

Combining Eqs. 2.2,2.3 we get an equation of motion for a damped harmonic oscillator

$$\ddot{\tau} + 2d_r\dot{\tau} + \omega_s^2\tau = 0 \tag{2.4}$$

$$d_r = \frac{\dot{U}_{\mathrm{rad}}(E_0)}{2T_0}$$

$$\omega_s = \sqrt{-\frac{\alpha e\dot{V}_{\mathrm{rf}}(\tau_s)}{E_0 T_0}} \tag{2.5}$$

where d_r is the radiation damping rate and ω_s is the synchrotron frequency.

The equation of motion derived above describes a stable system. When perturbed either in energy or longitudinal position the particle executes an exponentially decaying oscillation around the nominal point. So far we have considered the motion of a single particle. However in the physical machine many particles are stored in the same RF bucket. The center-of-mass behavior of the stored bunch follows the same dynamics that were derived for a single particle. In the next section we will describe coupling mechanisms through which motion of one bunch can affect other bunches in the ring. This coupling is the root cause of the longitudinal coupled-bunch instabilities.

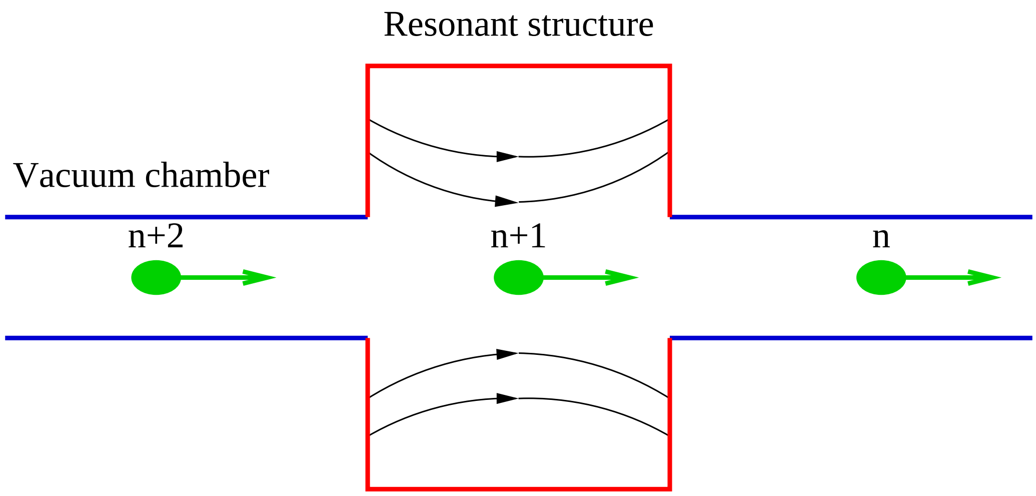

Figure 2.3: Electromagnetic field is excited in the resonant structure by bunch n. The field persists in the structure consistent with the quality factor of the excited resonance and can act on the bunches in the following RF buckets.

2.2 Beam Dynamics

Bunches of charged particles passing through the vacuum chamber of a storage ring leave behind electromagnetic fields. These fields (wake fields) affect the energy of the following bunches providing a bunch-to-bunch coupling mechanism. This is graphically illustrated in Fig. 2.3.

If we monitor electric fields and bunch arrivals within the structure we will observe the sequence of events illustrated in Fig. 2.4.

When bunch passes through the structure is excites an electromagnetic field which persists for some time after the bunch passage. That time is defined by the damping times (related to the resonance quality factors that will be introduced in Sec. 2.3) of the excited resonant modes. In the accelerating RF cavities parasitic resonances can have damping times in hundreds of revolutions. The oscillating field is sampled by the following bunches, e.g. $n+1$ and $n+2$ in Fig. 2.4. If we now apply a modulation to the arrival time of bunch n that will produce a phase modulation of the voltages sampled by the following bunches. Therefore bunch n is now coupled to the following bunches. Since the other bunches also create the wake fields they couple to bunch n closing the feedback loop. For certain combinations of the resonant frequencies and bunch currents the overall system becomes unstable and bunch oscillations will grow to large amplitudes limited either by non-linearities or beam loss.

The theory of coupled-bunch instabilities is well developed [4, 5]. Here we will present the most important analysis results which link the characteristics of the unstable longitudinal motion to the Fourier transform of the wake function known as coupling impedance.

The longitudinal wake function $W^{\|}(t)$ is defined as the integrated longitudinal component of the electric field experienced by a test charge passing through the vacuum chamber a time t later than the particle that excites the field. The field is integrated over an entire turn. Longitudinal impedance $Z^{\|}(\omega)$ is defined as

$$Z^{\|}(\omega) = \int_{-\infty}^{\infty} W^{\|}(t)e^{-i\omega t}dt = \int_{0}^{\infty} W^{\|}(t)e^{-i\omega t}dt \tag{2.6}$$

The impedance is hermitian, that is $Z(-\omega) = Z^{*}(\omega)$ where * denotes complex conjugate as it is a Fourier transform of a real function. This property will be quite important in the following section when we will consider the effect of the longitudinal impedance on beam stability and system eigenvalues.

So far we have considered as a longitudinal coordinate of the bunch its time of arrival

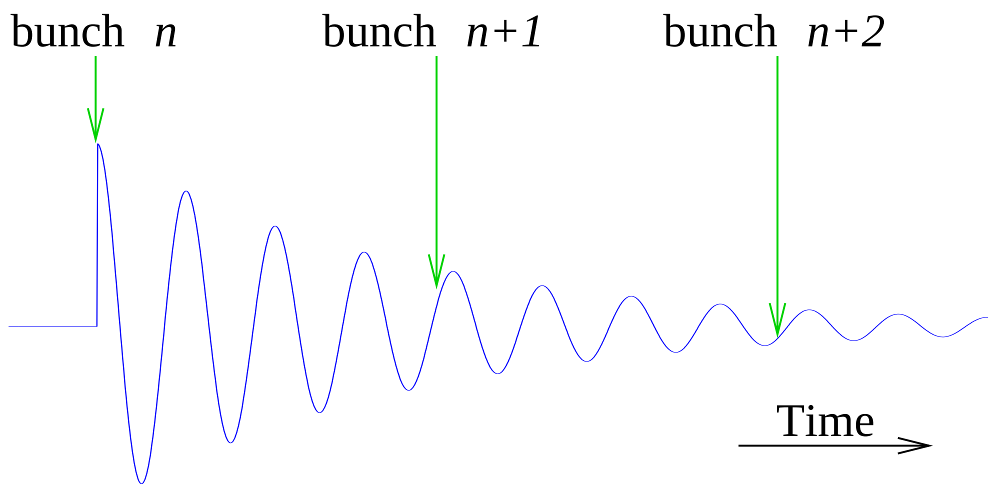

Figure 2.4: Oscillatory electric field excited in the structure by bunch n is sampled by bunches arriving in the structure afterwards.

error τ measured relative to the synchronous time. In practice it is more convenient to use a linearly related coordinate - phase deviation at the RF frequency defined as $\phi = \omega_{\mathrm{rf}}\tau$. Using phase error as a coordinate normalizes the motion to the RF frequency of the ring. The new coordinate carries information on the absolute amplitude of the motion. That is an oscillation amplitude of 20 ps can be large or small depending on the RF frequency. However when expressed as 1 degree at RF it is clearly a small oscillation. Consequently motion amplitudes can be compared between accelerators with different RF frequencies.

The bunch motion in a storage ring can be projected onto an orthonormal basis which will define the eigenmodes. If the ring filling pattern is even, i.e. it has rotational symmetry (for N filled bunches rotation by h/N RF buckets leaves the same buckets filled), the coordinate transformation is the discrete-time finite duration Fourier transform (DFT) given by

$$\varphi_l(t) = \frac{1}{N}\sum_{n=0}^{N-1}\phi_n(t)e^{-i2\pi ln/N} \tag{2.7}$$

The above transforms the motion in bunch domain $\phi_n(t)$ to the motion in the eigenmodal domain $\varphi_l(t)$. We will call the eigenmodal basis for an even filling pattern the even-fill eigenmode (EFEM) basis. In this work the variables in the EFEM basis will be denoted as $\hat{x}$ (x in the bunch basis) except for the phase for which we use φ in the EFEM basis and ϕ in the bunch basis.

The motion of bunch k oscillating in mode l is given by

$$\begin{aligned} \phi_k(t) &= e^{j2\pi kl/N}\varphi_l(t) \\ \varphi_l(t) &= Ae^{\Lambda_l t} \end{aligned}$$

where Λ_l is the modal eigenvalue [6]. If there is no longitudinal impedance affecting mode l its unperturbed eigenvalue is given by

$$\Lambda^0 = -d_r + j\omega_s \tag{2.8}$$

When longitudinal impedances are introduced they result in eigenvalue shift λ_l [4, 2]

$$\lambda_l = \frac{\pi\alpha e f_{\mathrm{rf}}^2 I_0}{E_0 h \omega_s} Z^{\|\mathrm{eff}}(l\omega_0 + \omega_s) \tag{2.9}$$

$$Z^{\|\mathrm{eff}}(\omega) = \frac{1}{2\pi f_{\mathrm{rf}}} \sum_{p=-\infty}^{\infty} (pN\omega_0 + \omega) Z^{\|}(pN\omega_0 + \omega) \tag{2.10}$$

where $f_{\mathrm{rf}} = 1/T_{\mathrm{rf}}$ is the frequency in the accelerating cavities, I_0 is the average (DC) beam current, $\omega_0 = 2\pi/T_0$ is the revolution frequency, and $Z^{\|}(\omega)$ is the total longitudinal impedance. Here we assume that the bunch length is much smaller than the wavelengths of the wake fields driving coupled-bunch instabilities.

For small oscillations the beam samples the longitudinal impedance at the RF frequency. Such sampling aliases the impedance into the frequency band from 0 to $f_{\mathrm{rf}}/2$. Summation over p in Eq. 2.10 describes aliasing of the longitudinal impedance into the frequency range from DC to f_{rf}. The aliased impedance $Z^{\|\mathrm{eff}}$ sampled at the upper synchrotron sideband of the lth revolution harmonic defines the modal eigenvalue shift λ_l. If one wanted to restrict Eqs. 2.9-2.10 to the frequencies below the Nyquist frequency $f_{\mathrm{rf}}/2$ one could replace $Z^{\|\mathrm{eff}}(l\omega_0 + \omega_s)$ by $Z^{\|\mathrm{eff}}((N-l)\omega_0 - \omega_s)$ for $l \geq N/2$. However the original representation is more common in the analysis of longitudinal coupled-bunch instabilities [6, 7, 8, 9].

2.3 Driving terms

Longitudinal impedance $Z^{\|}(\omega)$ defined in Eq. 2.6 is computed from the wake field integrated over the whole ring. This includes contributions from many resonators and broadband impedances. In a physical machine there are usually several dominant resonances which define the unstable eigenmodal structure. Most often these resonances occur in the accelerating RF cavities. The longitudinal impedance of such a cavity consists of the fundamental mode and the higher order modes (HOMs). The fundamental mode is designed to be resonant at the RF frequency of the ring while the frequencies of the higher order modes are not controlled during cavity design and depend on the geometry of the cavity.

Let's consider the effect of the impedance near a given revolution harmonic on the longitudinal stability. The impedance near $\omega_{\mathrm{imp}} = p\omega_{\mathrm{rf}} + l\omega_0$ revolution harmonic (p is integer) affects two eigenmodes: l and $N-l$. According to Eq. 2.9, the eigenvalue of mode l is affected by the longitudinal impedance at the upper synchrotron sideband of ω_{imp}. Passive (lossy) wake fields result in the impedance with positive real part, that is $\Re(Z^{\|}(\omega)) > 0$ for

$\omega > 0$. From Eqs. 2.9-2.10 we see that real part of Λ_l has a shift in value proportional to the real part of the impedance at $\omega_{\text{imp}} + \omega_s$. That shift moves EFEM l toward instability. Since longitudinal impedance is hermitian there is a contribution of the above impedance to $Z^{\|\text{eff}}$ at $-\omega_{\text{imp}} + \omega_s$. Due to frequency scaling effective impedance for that mode will have negative real part since $\Re(Z^{\|}(-\omega)) = \Re(Z^{\|}(\omega)) > 0$. Thus longitudinal impedance near $p\omega_{\text{rf}} + l\omega_0$ drives EFEM l unstable, but acts to stabilize eigenmode $N - l$.

2.3.1 Higher order mode impedances

Often it is useful to parameterize resonant modes in cavity-like structures. Longitudinal impedance of such modes can be expressed as

$$Z^{\|}(w) = \frac{R_s}{1 + jQ[\omega/\omega_r - \omega_r/\omega]} \tag{2.11}$$

where ω_r is the resonance center frequency R_s is the shunt impedance, and Q is the quality factor. Another important parameter of the resonator is the 3 dB bandwidth defined as the difference between upper and lower frequencies at which the absolute value of the impedance is reduced by 3 dB relative to the peak value R_s. For resonators with $Q \gg 1$ the bandwidth is given by

$$\omega_{\text{bw}} = \frac{\omega_r}{Q}$$

To help analyze the effect of higher-order mode impedances on longitudinal stability we will introduce the classes of narrowband and wideband resonances. The quality factor Q in Eq. 2.11 determines the width of the resonance. Traditionally resonators are considered narrowband if $Q \gg 1$ and wideband if $Q \approx 1$. In this context we will define narrowband/wideband resonator grouping by comparing resonator bandwidth to the ring revolution frequency ω_0. Resonances with the bandwidth smaller than the revolution frequency are narrowband and all others - wideband. The two classes call for slightly different control techniques.

Effect of a narrowband resonance on longitudinal stability depends on the tuning of its center frequency. Placing the peak of the resonance midway between two revolution

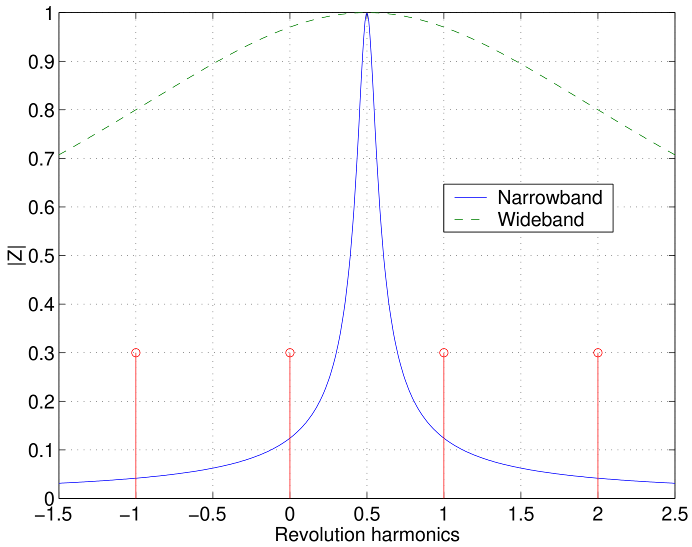

Figure 2.5: Normalized magnitudes of two impedances: narrowband and wideband. The narrowband resonator has the bandwidth of $\omega_0/8$ and the wideband one - $4\omega_0$. Both resonances are tuned midway between two revolution harmonics.

harmonics reduces the resultant instability growth rates. The effect is strongest for the very narrow resonances and becomes less pronounced as the bandwidth increases. For the borderline case of $\omega_{\text{bw}} = \omega_0$ optimal tuning halves the shift of the real part of the eigenvalue relative to the worst-case shift which occurs when peak of the resonance is placed at the upper synchrotron sideband of a revolution harmonic. Narrowband HOMs of the accelerating cavities are commonly tuned by adjusting the temperature of the structure [7, 9]. In some cases cavities are equipped with a separate tuner antenna coupled to the strongest HOMs [10].

Wideband resonances provide less tuning flexibility - when HOM bandwidth is $4\omega_0$ the difference between eigenvalue shift for best and worst tuning is only 6%. Such HOMs cannot be "hidden" between revolution harmonics and always present an effective impedance on the order of $\omega_r R_s/\omega_{\text{rf}}$. For this reason the wideband HOMs require careful attention during accelerator design. The effective impedance presented by these modes must be within the feedback damping limits. The difference between narrowband and wideband resonators is illustrated in Fig. 2.5.

To summarize, narrowband impedances can be controlled via combination of resonant frequency tuning and active feedback while with wideband resonators one has to rely solely on active feedback.

2.3.2 Fundamental mode impedances

The fundamental mode of the accelerating cavity rarely causes longitudinal instabilities even though it presents a large longitudinal impedance. The secret here is in the tuning of the cavity resonant center frequency below the RF frequency. Such tuning is used for two reasons. First, it is used adjust the matching of the cavity with beam to the RF generator. Usually the center frequency of the cavity is tuned so that current reflected to the generator is in phase with generator current (resistive loading). Second, tuning the cavity below the RF frequency ensures that the fundamental impedance does not cause instability of the lowest-frequency mode - also called Robinson instability. Impedances at both upper and lower synchrotron sidebands of the RF frequency contribute to the effective impedance for EFEM 0. The upper sideband figures in the expression with positive sign (unstable) while the lower one has negative sign (stable). When the resonance is tuned below the RF frequency the impedance at the lower sideband is larger than that at the upper one. Thus such tuning guarantees that EFEM 0 is stable. However there are some cases when the impedance of the fundamental mode of the RF cavity can drive longitudinal coupled-bunch instabilities - we will consider two such possibilities here.

The first case is that of a storage ring with large circumference. Such a ring has low revolution frequency. The impedance of the fundamental mode sampled at $\omega_{\rm rf} - \omega_0 + \omega_s$ excites eigenmode $N-1$. The effect is amplified if the RF system is heavily beam loaded, that is the average beam current is large in comparison to the generator current in the RF cavity. Beam loading will be discussed in more detail in Sec. 2.4. Heavy beam loading requires large detuning of the accelerating cavity and, therefore, increases the impedance driving eigenmode $N-1$. In some cases the detuning can be equal to or larger than the revolution frequency [11]. Then the EFEM $N-1$ is driven by the full shunt impedance of the accelerating cavity. To operate in this regime the RF system must be equipped with feedback loops aimed at reducing the impedance presented to the beam [12, 13].

The second case is that of parked cavities. Often there are more accelerating cavities installed in the ring than can be or need to be powered by RF generators. The unused cavities must be detuned away from the RF frequency or its harmonic (for the harmonic RF system). Such detuning of the cavities is called parking. Optimal parking of the cavity is midway between two revolution harmonics to minimize the impedance that affects the beam. When many cavities are parked at the same frequency the cumulative impedance at the nearby revolution harmonics can be large enough to drive coupled-bunch instabilities. The effect of the harmonic RF cavities is amplified by the high resonant frequency - see the scaling factor in Eq. 2.10. One such case will be considered in Ch. 4. There is a method to minimize the effect of parked cavities. This technique is applicable when the number of parked cavities is even. Then parking half of the cavities at $\omega_{\mathrm{rf}} - \omega_{\mathrm{park}}$ and the other half at $\omega_{\mathrm{rf}} + \omega_{\mathrm{park}}$ reduces the effective longitudinal impedance. For illustration suppose that $\omega_{\mathrm{park}} = 2.5\omega_0$ that is the cavities are parked between two and three revolution harmonics away from the RF. Then from Eq. 2.10 we get for the effective impedance for EFEM 2

$$\begin{aligned} Z_2^{\|\mathrm{eff}} &= \frac{1}{\omega_{\mathrm{rf}}}[(\omega_{\mathrm{rf}} + 2\omega_0 + \omega_s) Z_1^{\|}(\omega_{\mathrm{rf}} + 2\omega_0 + \omega_s) - (\omega_{\mathrm{rf}} - 2\omega_0 - \omega_s) Z_2^{\|*}(\omega_{\mathrm{rf}} - 2\omega_0 - \omega_s)] \\ &\approx \frac{4\omega_0 + 2\omega_s}{\omega_{\mathrm{rf}}} Z^{\|}(\omega_{\mathrm{rf}} + 2\omega_0 + \omega_s) \approx \frac{4}{h} Z^{\|}(\omega_{\mathrm{rf}} + 2\omega_0 + \omega_s) \end{aligned} \tag{2.12}$$

Fortunately, for the large rings, where the effect of the parked cavities is most noticeable due to narrow revolution harmonic spacing, the harmonic number is high as well making the impedance reduction in Eq. 2.12 very effective.

2.4 Interaction with the RF system

So far we have developed the longitudinal dynamics using the notion that the RF voltage is provided by a perfect voltage source. In practice the RF accelerating voltage is created by injecting drive currents into the RF resonant cavity. However drive current is not the only input into the cavity. Beam current is the second current source driving the resonator. Consequently the voltage in the accelerating cavity is a result of the total drive current equal to a sum of generator and beam currents. As shown in Sec. 2.1 the voltage in the cavity determines longitudinal dynamics. That is, the beam current source is affected by the voltage in the cavity while that voltage depends on the beam current. This creates a feedback mechanism which can modify the parameters of the longitudinal beam dynamics as well as the dynamics proper. When beam current is small relative to the generator current - called light beam loading - the effect of the interaction between the beam and the RF system is weak. In this case the voltage in the cavity is in most part determined by the constant generator current. However heavy beam loading with the beam current

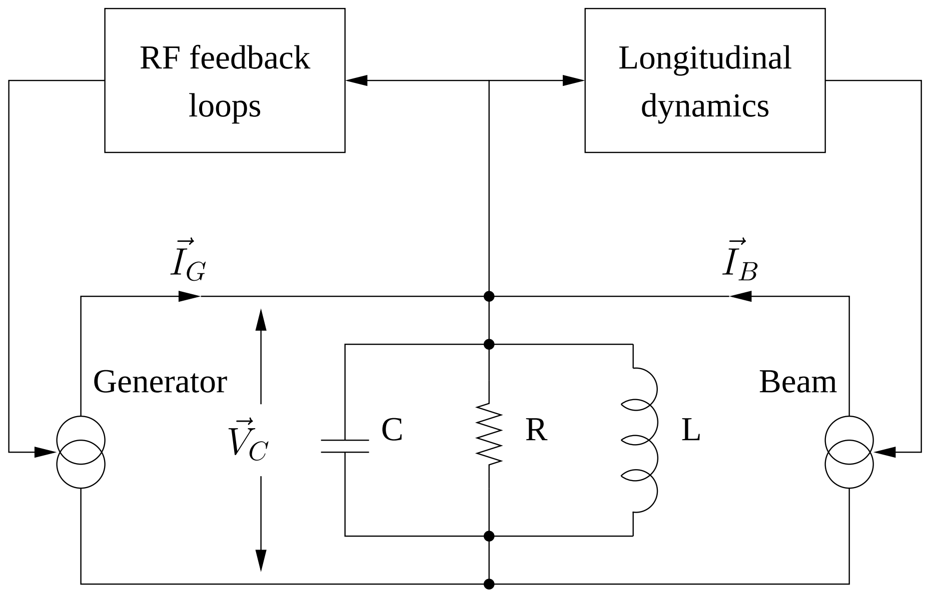

Figure 2.6: Schematic of the RF cavity model with two input currents and feedback loops

being comparable to or larger than the generator current creates strong interaction which significantly modifies longitudinal motion parameters and dynamics for the low-frequency eigenmodes (with mode numbers near 0 or N-1).

Here we will use the Pedersen model [14, 15] to characterize the small-signal behavior of the beam and RF system. In this model the cavity is represented by an equivalent parallel RLC circuit driven by two currents: generator current $\vec{I}_G$ and beam current $\vec{I}_B$. This model is schematically illustrated in Fig. 2.6. The RF system in addition to the RF power generator usually includes additional feedback loops. These can be narrowband amplitude and phase loops that maintain the cavity voltage amplitude and phase at the RF frequency or broader bandwidth feedback to improve system stability margins [15].

The driving current phasors are evaluated at the RF frequency. For short bunches in the lepton storage rings $|\vec{I}_B| = 2I_0$. From here on we will use I_B to represent $|\vec{I}_B|$. In the Laplace domain the impedance of the cavity is given by

$$Z(s) = \frac{2\sigma R s}{s^2 + 2\sigma s + \omega_r^2} \tag{2.13}$$

where $\sigma = \omega_r/2Q$ is the damping time of the cavity. Note that Eq. 2.13 transforms into

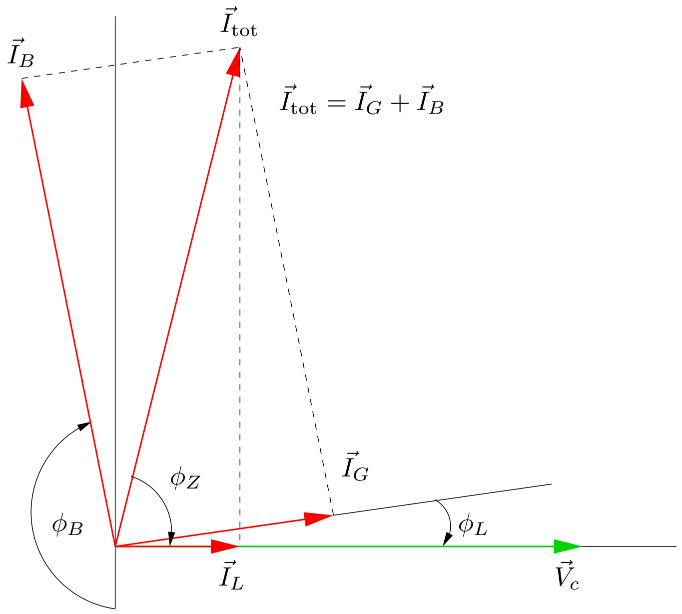

Figure 2.7: Steady-state vector diagram of accelerating cavity currents and voltages

Eq 2.11 with the substitution $s = j\omega$.

The beam loading is characterized by the dimensionless parameter $Y = I_B/I_L$ where I_L is the generator current required to produce the same cavity voltage without beam load and with the cavity resonance tuned to ω_{rf}. From the steady-state vector diagram shown in Fig. 2.7 relating $\vec{I}_B$, $\vec{I}_G$, and cavity voltage $\vec{V}_C$ one gets the following relationship

$$\tan\phi_L = \frac{\tan\phi_Z - Y\cos\phi_B}{1 + Y\sin\phi_B} \tag{2.14}$$

where ϕ_L is the loading angle between the cavity voltage and the generator current, ϕ_B is the synchronous phase angle, and ϕ_Z is the cavity impedance angle. For efficient utilization of the power source loading angle is usually maintained constant and close to 0. To achieve that according to Eq. 2.14 the magnitude of the cavity impedance angle has to increase when Y increases - more obvious if we rewrite that equation as

$$\tan\phi_Z = \tan\phi_L + Y(\tan\phi_L\sin\phi_B + \cos\phi_B)$$

This implies that the tangent of the cavity impedance angle has to change linearly with beam current to keep the loading angle constant. Angle ϕ_Z is adjusted by changing the center frequency of the cavity. The detuning of the cavity resonance is given by $\omega_D = \omega_r - \omega_{\mathrm{rf}} = \sigma \tan \phi_Z$.

The effect of RF cavity impedance on longitudinal dynamics falls outside the framework developed in Section 2.2 due to several factors limiting the applicability of eigenvalue perturbation analysis. First, the cavity fundamental impedance produces large eigenvalue shifts comparable to the eigenvalue itself. Second, the high-Q fundamental impedance depends strongly on the frequency. Shift of the eigenvalue modifies the frequencies at which the beam samples the impedance. Due to this the straightforward application of Eqs. 2.9-2.10 results in large errors.

Using the small-signal model developed in [13, 14, 16] we can compute the Laplace domain closed-loop poles of the system that includes both the beam and the RF. In Fig. 2.8 a block diagram of the model is shown. The diagram shows propagation of amplitude and phase modulations of beam current - variables a_B and p_B - to the amplitude and phase modulations of the cavity voltage through four transfer functions: G^B_{aa}, G^B_{ap}, G^B_{pa}, and G^B_{pp}. The effective phase of the cavity voltage $p_{V_{\mathrm{eff}}}$ excites the dipole mode beam dynamics represented by the transfer function $B_1(s)$. The RF side of the interaction is omitted here. The resulting model is valid if RF feedback loops have bandwidth much smaller than synchrotron frequency. When wideband feedback is present the complete model has to be examined since the wideband RF feedback loops will affect the longitudinal dynamics.

The transfer functions in the model are obtained by considering propagation of small-signal cavity input current modulations to the cavity voltage. For the resonant cavity with impedance given by Eq. 2.13 these functions are:

$$\begin{aligned} G_{pp}(s) &= G_{aa}(s) = \frac{\sigma^2 + \omega_D^2 + \sigma s}{(s+\sigma)^2 + \omega_D^2} \\ G_{pa}(s) &= -G_{ap}(s) = \frac{\omega_D s}{(s+\sigma)^2 + \omega_D^2} \end{aligned}$$

The transfer functions from $\vec{I}_B$ and $\vec{I}_G$ are obtained by geometrically projecting modulations of these vectors to $\vec{I}_T$. The beam dynamics transfer function $B_1(s)$ includes the dynamics of all coupled-bunch eigenmodes. When fundamental impedance mostly interacts with the lowest-frequency beam mode we use the following function

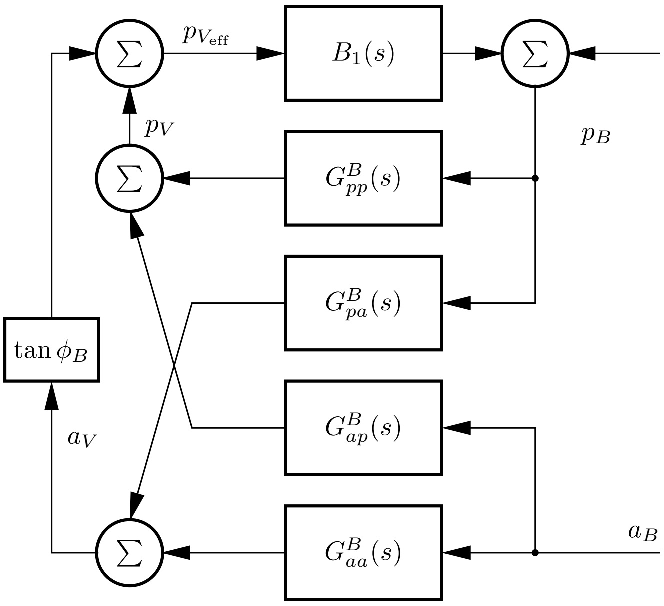

Figure 2.8: Block diagram of the beam cavity interaction

$$B_1(s) = \frac{\omega_s^2}{s^2 + 2d_r s + \omega_s^2}$$

By computing the relevant transfer function at the given beam current and RF configuration we obtain a closed-loop model of beam-cavity interaction and can analyze various dynamic properties of the system such as response to RF generator noise or interaction with longitudinal feedback.

2.5 Dynamics modifiers

Here we will discuss effects such as beam loading, energy ramping, etc. that modify the longitudinal behavior of the beam relative to the model developed in Section 2.2. These important phenomena strongly affect the longitudinal motion parameters and the performance of the stabilizing feedback. We will start the discussion from the effect of the beam loading on eigenmode 0.

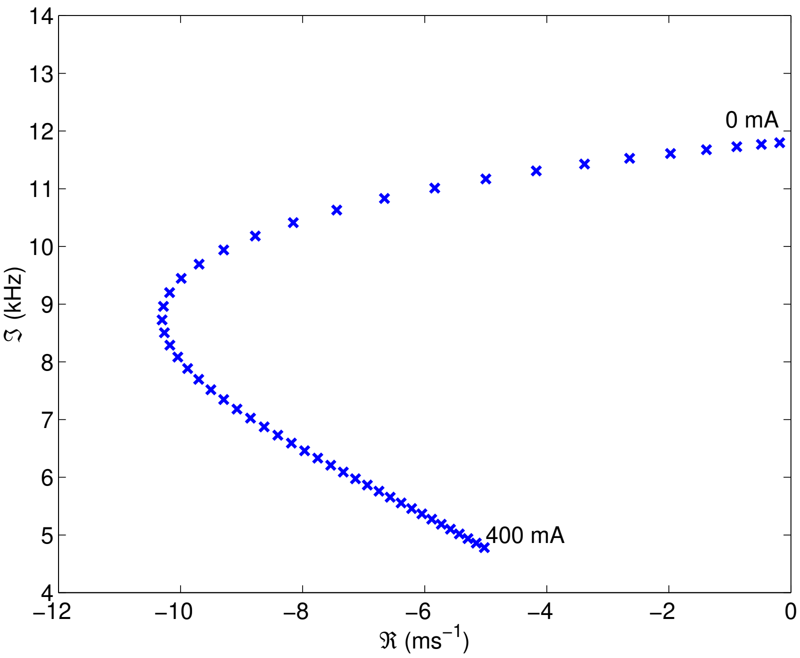

Figure 2.9: Root locus for the dominant closed-loop pole of the cavity-beam mode 0 system at the ALS. The pole is plotted as a function of beam current in the 0 mA to 400 mA range in 8.1 mA steps.

2.5.1 Beam loading of the RF cavity and mode 0 frequency shift

In the previous section we described the interaction of the beam and the RF system. This interaction most strongly affects the lowest-frequency eigenmode (closest to the RF frequency) - mode 0. As the RF cavity is detuned with increasing beam current the impedance driving mode 0 changes. From the eigenvalue perturbation model described in Sec. 2.2 we expect the eigenvalue shift of EFEM 0 to be proportional to $Z^{\|}(\omega_{\mathrm{rf}} + \omega_s) - Z^{\|*}(\omega_{\mathrm{rf}} - \omega_s)$. The detuned RF fundamental impedance has a large negative imaginary part. Thus we expect a downward shift in the mode 0 oscillation frequency. Precise changes in the dynamics are available from a small-signal model described in Sec. 2.4. Applying the model to the case of the Advanced Light Source we obtain the closed-loop pole locations shown in Fig. 2.9.

The root locus shows that mode 0 oscillation frequency shifts from 11.8 kHz nominal synchrotron frequency to 4.8 kHz at 400 mA beam current. The large frequency shift has been also observed at other accelerators with heavy beam loading - for example, in DAΦNE positron ring the frequency shift is from 33 kHz to 11 kHz at 1 A.

The large frequency shift of mode 0 changes the interaction of the beam dynamics and active longitudinal feedback system and places additional constraints on feedback controller design.

2.5.2 Energy ramping

Not all storage rings are equipped with the full energy injection system. That is the energy provided by the injection accelerator is lower than the desired operating energy of the ring. Charge can be stored in the ring only when the nominal and injection energies are closely matched. These contradicting requirements give rise to energy ramping technique. In this method the ring optics is initially configured for the injection energy and the ring is filled to the full operating current. Next the magnetic lattice is modified to slowly raise (ramp) the ring energy to the desired value. The ramping process is reversed when injection is needed again.

Energy ramping affects longitudinal dynamics in several ways. Both synchrotron frequency ω_s and modal eigenvalue shift λ_l depend on the beam energy. From Eq. 2.5 we see that synchrotron frequency changes as $1/\sqrt{E}$. Modal oscillation frequency is given by the imaginary part of the modal eigenvalue $\Lambda_l = \Lambda^0 + \lambda_l$ and the modal growth rate - by the real part. From Eqs. 2.8,2.9 we have

$$\begin{aligned} \varpi_l &= \Im(\Lambda_l) = \omega_s + \frac{\pi\alpha e f_{\mathrm{rf}}^2 I_0}{E_0 h \omega_s} \Im(Z^{\|\mathrm{eff}}(l\omega_0 + \omega_s)) \\ \sigma_l &= \Re(\Lambda_l) = -d_r + \frac{\pi\alpha e f_{\mathrm{rf}}^2 I_0}{E_0 h \omega_s} \Re(Z^{\|\mathrm{eff}}(l\omega_0 + \omega_s)) \end{aligned}$$

where ϖ_l and σ_l are the oscillation frequency and the exponential damping rate of EFEM l. Substituting ω_s from Eq. 2.5 and assuming that effective impedance changes little with the changes in synchrotron frequency we get

$$\varpi_l = \omega_s(1 + \frac{\pi f_{\mathrm{rf}} I_0}{-\dot{V}_{\mathrm{rf}}} \Im(Z^{\|\mathrm{eff}}(l\omega_0 + \omega_s^0)) \tag{2.15}$$

$$\sigma_l = -d_r + \omega_s \frac{\pi f_{\mathrm{rf}} I_0}{-\dot{V}_{\mathrm{rf}}} \Re(Z^{\|\mathrm{eff}}(l\omega_0 + \omega_s^0)) \tag{2.16}$$

where ω_s^0 is the zero-current synchrotron frequency at the injection energy. From these equations we see that when energy is ramped up the modal oscillation frequency has the same scaling as the synchrotron frequency. Reduction in the oscillation frequency expands the feedback controller bandwidth requirement with wideband controllers presenting bigger design challenge. A positive side effect of energy ramping is the reduction of the induced growth rates according to Eq. 2.16.

2.5.3 High harmonic RF systems

A high harmonic RF system is used to apply a voltage to the beam at a harmonic of the ring RF frequency. Such system can be active - powered by an external generator, or passive - harmonic voltage is generated by the beam current. A harmonic RF system can be used to modify the longitudinal potential well and, consequently, to lengthen or shorten the bunch size.

High harmonic RF systems are often used in the storage rings for improving beam lifetime and for introducing Landau damping for controlling longitudinal coupled-bunch instabilities [17, 18, 19]. In storage rings with dense electron bunches and moderate nominal energy the beam lifetime is often determined by large-angle intrabeam (Touschek) scattering [10]. Higher harmonic RF system provides a way to stretch the bunches and lower the charge density thereby improving the lifetime.

The higher harmonic RF system generates an accelerating voltage in addition to the main RF system. The additional voltage modifies the longitudinal potential well depending on the relative phase between the two RF systems. By adjusting the relative phase one can narrow or widen the potential well and change the longitudinal bunch size [20].

Let us consider the RF voltage seen by the beam when both main and harmonic RF systems are operational. It is given by [10, 21]

$$V(\tau) = V_{\mathrm{rf}} \sin(\omega_{\mathrm{rf}}(\tau + \tau_s)) + V_h \sin(n\omega_{\mathrm{rf}}(\tau + \tau_h)) \tag{2.17}$$

where V_h is the voltage in the harmonic cavities, τ_h is the phase of that voltage seen by a synchronous particle, and n is the integer RF harmonic.

Introduction of the harmonic voltage not only changes the bunch length but also the synchrotron frequency. In bunch lengthening mode the slope of the total RF voltage at the synchronous phase is reduced and so is the synchrotron frequency. When V_h is held constant lower ω_s causes increased eigenvalue shifts according to Eq. 2.9. The same longitudinal impedances result in faster growth rates making the feedback stabilization more difficult.

In many cases due to budget constraints the harmonic RF system is passive. In this case there is no external generator and the harmonic voltage is generated by the beam itself. Desired relative phase of the harmonic voltage is adjusted by harmonic cavity tuning. In this configuration the harmonic voltage is dependent on beam current. For the fixed cavity tuning $V_h \propto I_0$. This causes the synchrotron frequency to change with beam current. A stabilizing feedback controller for such a system must be adaptive or be able to handle significant changes in plant dynamics. In fact, most of the advanced controller design methods described in Chapter 6 were brought to life by the installation of the third harmonic passive RF system at the ALS [19, 22, 23].

2.6 Beam model as a multi-input multi-output system

In order to design a stabilizing feedback controller for the longitudinally unstable storage ring we need to build a model of our unstable system. In Section 2.2 we considered dynamics of an unstable eigenmode as a second-order harmonic oscillator, characterized by the modal eigenvalue Λ. The real part of the eigenvalue is the growth or damping rate of the mode while the imaginary part is the oscillation frequency.

The undamped natural frequency of a second-order system is defined as [24, page 72]

$$\omega_l = \sqrt{\varpi_l^2 + \sigma_l^2}$$

The following differential equation derived using Eqs. 2.4 and 2.7 governs the motion of an EFEM

$$\ddot{\varphi}_l + 2\sigma_l \dot{\varphi}_l + \omega_l^2 \varphi_l = 0 \tag{2.18}$$

To stabilize the above system we need a way to affect the internal dynamics, i.e. an actuator. Since longitudinal coupled-bunch dynamics of the ring are defined by the electric field in the accelerating RF cavities and modified by the parasitic wake fields it is natural to stabilize the system using longitudinal electric field. Let's consider an actuator (kicker) that creates a longitudinal electric field such that a bunch of charge q gains energy $qv(t)$. Let $v_n(t)$ be the kick of bunch n and $\hat{v}_l(t)$ be the same kick transformed to the EFEM basis.

Then Eq. 2.18 becomes

$$\ddot{\varphi}_l + 2\sigma_l \dot{\varphi}_l + \omega_l^2 \varphi_l = \frac{\alpha e \omega_{\mathrm{rf}}}{E_0 T_0} \hat{v}_l(t) \tag{2.19}$$

By taking a Laplace transform we get a transfer function from the kicker voltage to beam phase

$$\hat{G}_l(s) = \frac{\hat{\Phi}_l(s)}{\hat{V}_l(s)} = \frac{\alpha e \omega_{\mathrm{rf}}}{E_0 T_0} \frac{1}{s^2 + 2\sigma_l s + \omega_l^2} \tag{2.20}$$

The overall beam system has N eigenmodes so that it can be represented as a MIMO linear system with N inputs and N outputs. The input vector $\vec{\hat{V}}$ and the output vector $\vec{\hat{\Phi}}$ are related by the transfer matrix $\mathbf{G}(s)$:

$$\vec{\hat{\Phi}} = \hat{\mathbf{G}}(s)\vec{\hat{V}}$$

By definition the eigenmodes are linearly independent, therefore $\hat{\mathbf{G}}(s)$ is a diagonal matrix of the following form:

$$\hat{\mathbf{G}}(s) = \begin{bmatrix} \hat{G}_0(s) & 0 & \dots & 0 \\ 0 & \hat{G}_1(s) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \hat{G}_{N-1}(s) \end{bmatrix} \tag{2.21}$$

2.7 Earlier work in coupled-bunch instability control

Feedback control of coupled-bunch instabilities has been extensively studied and described in literature. Control architectures ranging from mode-by-mode [25] to bunch-by-bunch [26, 27, 28] have been explored and their performance documented. Within author's research group H. Hindi performed extensive analysis and modeling of coupled-bunch instability control addressing beam model measurements [29], optimal LQG controller design [30] and performance of optimal and simplified controllers [31].

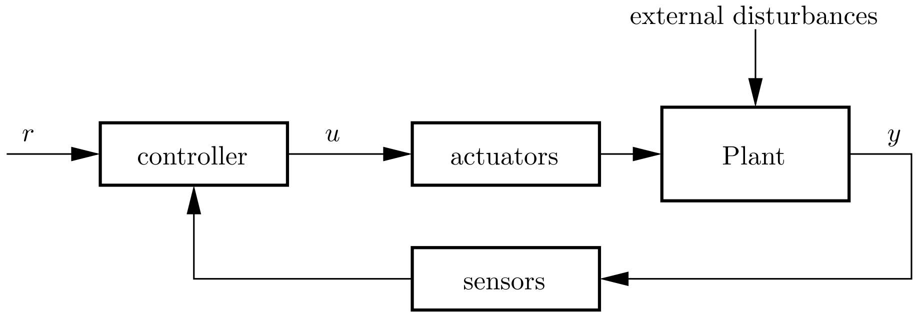

Figure 2.10: Generalized block diagram of a feedback control system

2.8 Feedback control: an introduction

In general, the objective in a feedback control system is to make some output of a dynamic system behave in a desired way by manipulating the input of that system. A general block-diagram of such a system is shown in Fig. 2.10. The system consists of the physical system (plant) the output of which we want to control. The output signal y is measured by the sensors and sent to the controller. The control objective might be to keep y small (or close to some constant value) - this is defined as a *regulator problem*. A different objective is to make plant output y follow some reference signal r - a *servomechanism problem*. Controller in Fig. 2.10 can be a regulator - then input r is omitted - or a servo. In any case controller determines the error between plant output and desired value and, based on the knowledge of plant dynamics, computes the control output u. The control signal is then applied to the plant via actuators.

The plant is subject to external disturbances which affect the output y. As one of the performance criteria of the control system one can consider the reduction of the transfer gain from external disturbance input to plant output.

2.9 Feedback control of coupled-bunch instabilities

The longitudinal coupled-bunch instability is a regulator problem aimed at keeping bunch phase excursions from the synchronous position small. In this case the control problem is that of stabilization since the plant is open-loop unstable. In order to design a feedback controller we need to chose our sensors and actuators. A good choice driven by sensitivity and implementation feasibility is to measure bunch phase error relative to synchronous phase.

A good actuator choice is an electromagnetic structure that applies a longitudinal electric field to the beam. In the previous section coupled-bunch instabilities were formulated as a MIMO linear system. In practice it is more convenient to observe bunch signals at a single point in the ring. Consequently the sensor and controller outputs are multiplexed at T_{rf}. Such multiplexing defines the bandwidth requirements for the actuator and the sensor. According to the Nyquist criterion in order to independently measure and correct individual bunches both the sensor and the actuator must have bandwidth of at least $f_{\mathrm{rf}}/2$. Note, that in modeling we will still use the multi-input multi-output model since sequential sampling instead of simultaneous does not significantly change the model. The difference between the two sampling schemes is $T_0 - T_{\mathrm{rf}}$ between the first and last bunches. During that time longitudinal position changes very little due to the fact that $\omega_s \ll \omega_0$.

In order to stabilize the system described by the transfer matrix in Eq. 2.21 we need to apply feedback to the unstable eigenmodes. Feedback systems which act only on the unstable eigenmodes - so called mode-by-mode systems - do exist and utilize the properties of the bunch spectrum which separate the different eigenmodes in the frequency domain [32, 33, 25]. The drawback of such systems is that they are designed for a particular unstable eigenmode spectrum. Consequently these feedback systems are not portable between different accelerators. Another operational difficulty is that changes in ring components, such as installation of new RF cavities, require feedback hardware redesign. In a general case one would like to design a system that allows independent control of each eigenmode. In the EFEM basis such system has a diagonal transfer matrix. Unfortunately in the physical world longitudinal motion is observed in the bunch basis and the real-time conversion to the EFEM basis is impractical. A feedback system that has independent eigenmode control results in a fully populated feedback matrix, that is correction signal for a given bunch depends on the motion of all bunches. Such a system is very computationally intensive. An obvious simpler topology is a bunch-by-bunch feedback in which the correction signal for a given bunch depends only on the motion of that bunch. Then the feedback matrix is diagonal in the bunch basis. This layout is illustrated in Fig. 2.11. The feedback system shown here is not only diagonal, but also acts equally on all bunches. Such system can be described by the following transfer matrix

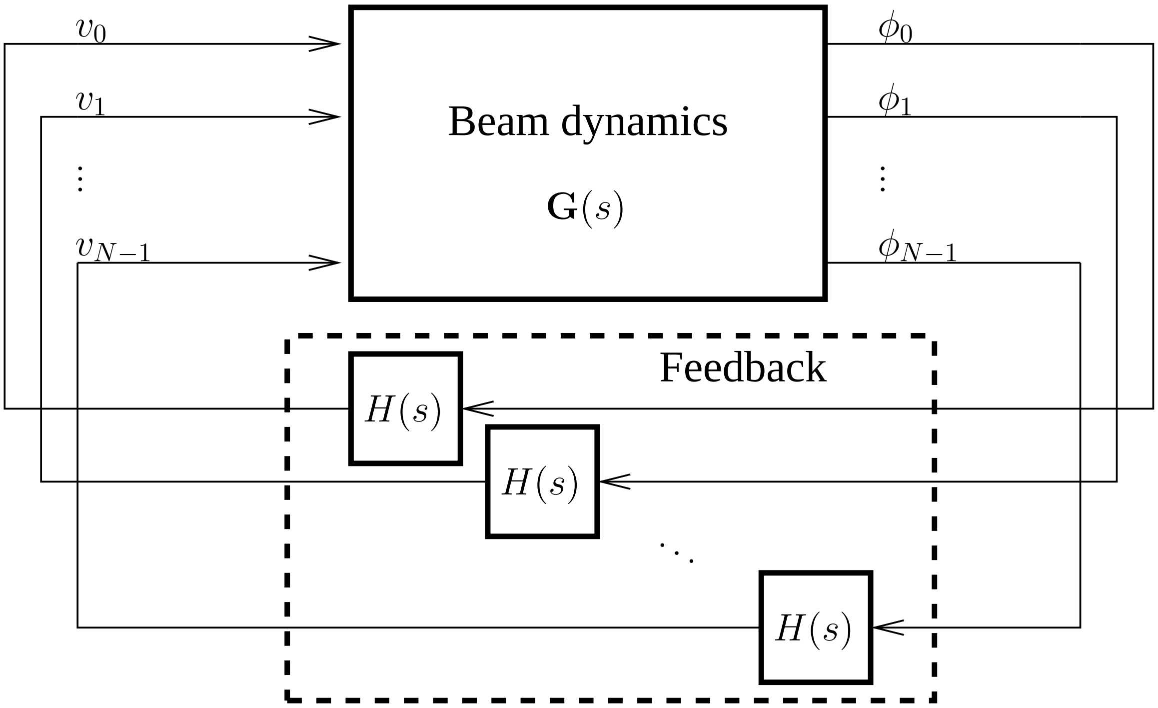

Figure 2.11: Block diagram of the beam and the bunch-by-bunch feedback system

$$\mathbf{H}(s) = \begin{bmatrix} H(s) & 0 & \dots & 0 \\ 0 & H(s) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & H(s) \end{bmatrix} = H(s)\mathbf{I}_N \tag{2.22}$$

where $\mathbf{I}_N$ is $N \times N$ identity matrix. An important property of such a feedback is invariance under coordinate transformations. Let coordinate transformation $\mathbf{T}$ be the transformation from bunch to EFEM basis. Elements of matrix $\mathbf{T}$ are given by the definition of the DFT in Eq. 2.7. Note that in the following discussion exact nature of the transformation matrix is unimportant as long as it is nonsingular. Then we have

$$\vec{\varphi} \;\; = \;\; \mathbf{T}\vec{\phi}$$

$$\begin{aligned}\hat{\vec{v}} &= \mathbf{T}\vec{v} \\ \hat{\mathbf{G}} &= \mathbf{T}\mathbf{G}\mathbf{T}^{-1} \\ \hat{\mathbf{H}} &= \mathbf{T}\mathbf{H}\mathbf{T}^{-1}\end{aligned} \tag{2.23}$$

If we substitute the expression for the feedback matrix from Eq. 2.22 into Eq. 2.23 we get

$$\hat{\mathbf{H}} = \mathbf{T}H(s)\mathbf{I}_N\mathbf{T}^{-1} = H(s)\mathbf{T}\mathbf{T}^{-1} = H(s)\mathbf{I}_N = \mathbf{H} \tag{2.24}$$

This shows that uniform bunch-by-bunch feedback is invariant under coordinate transformations. A system that applies identical feedback to each bunch applies the same feedback to each eigenmode. If the feedback is resistive, i.e. it only modifies the real part of the plant pole, its action will shift all eigenvalues along the real axis by the amount dependent on the loop gain. For the negative resistive feedback eigenvalues shift left and at some gain setting all of them are shifted to the left half plane (LHP) resulting in a stable closed-loop system.

An important property of Eq. 2.24 is that it holds true even for asymmetric ring fill patterns. As long as there is a linear coordinate transformation from the bunch coordinates to the eigenmodes the uniform bunch-by-bunch feedback results in the same feedback being applied to the eigenmodes of the uneven fill.

Let us consider a single-input single-output system consisting of an eigenmode plant and a feedback system. The open-loop transfer function for eigenmode l is given by $K_l\hat{G}_l(s)H(s)$ where K_l is the loop gain parameter. In a physical feedback system with realizable transfer function and transport delays there will be some maximum loop gain $K_l^{\max}$ above which the closed-loop system becomes unstable. There is also a minimum gain $K_l^{\min}$ below which the system is unstable. For unstable eigenmodes this minimum gain is positive while for the stable ones it is negative corresponding to positive feedback. To achieve longitudinal stability one has to configure each eigenmode so that $K_l^{\min} < K_l < K_l^{\max}$. With the uniform bunch-by-bunch feedback there is a single adjustable loop gain since the same feedback is applied to each eigenmode. Then for the overall system stability we need

$$\max_l K_l^{\min} < \min_l K_l^{\max}$$

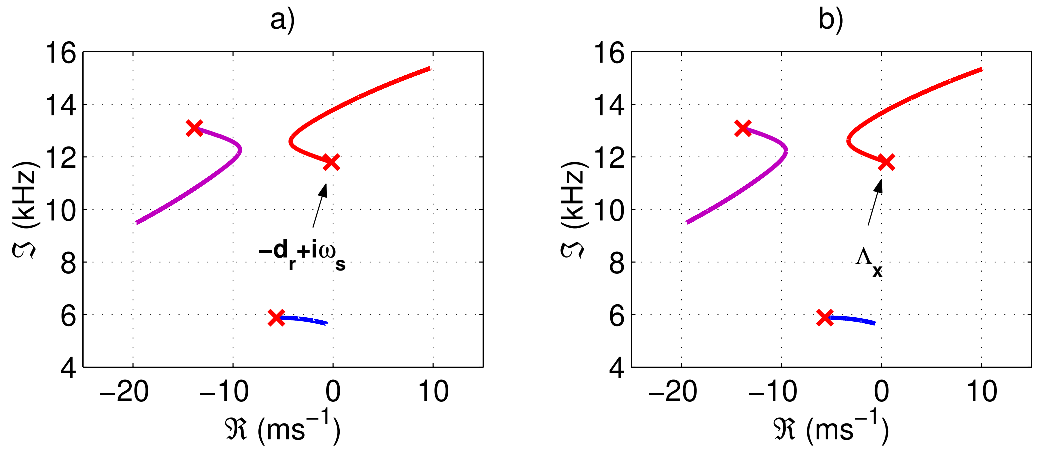

Figure 2.12: Fragments of a root loci for a) an unperturbed (stable) eigenmode with pole at $-d_r + i\omega_s$. The open-loop poles of the system are marked by $\times$. b) an unstable eigenmode.

The above inequality requires that there is a range of loop gains where all unstable plant modes are stabilized and no eigenmodes are driven to instability by too high a gain. Usually this is not a limitation since difference between the eigenvalues of the stable and unstable modes is much smaller than the unperturbed eigenvalue. Let's consider the root locus for an unstable eigenmode with feedback. The locus starts from the unstable pole and moves to the LHP. If we substitute a stable eigenmode the starting point is shifted by a small amount without strongly changing the overall locus. This effect is illustrated in Fig. 2.12 for two eigenmodes - one with eigenvalue of $-d_r + i\omega_s$ and another one with unstable eigenvalue Λ_x estimated on the basis of experimental measurements. Both root loci are plotted for the same range of loop gains. In this case the unstable eigenmode has narrower stability range than the unperturbed mode.

In general, arbitrary eigenmode feedback can achieve better damping since feedback can be optimized on the per-mode basis. In practice, for the systems of relatively weakly coupled oscillators (eigenvalue shifts are small) the restrictions of the bunch-by-bunch feedback do not limit system performance. Only when the mode-to-mode differences between eigenvalues are large relative to the unperturbed eigenvalue the bunch-to-bunch feedback becomes a limitation.

2.10 Digital bunch-by-bunch feedback system

A programmable system was developed to control longitudinal instabilities in the ALS storage ring and PEP-II and DAΦNE colliders [34] - it has since been installed and operated at SPEAR, BESSY-II and the PLS [35, 36].

The system to be described here uses bunch-by-bunch digital signal processing to generate feedback correction signal. Since this is a diagonal bunch-by-bunch system it is sufficient and complete to consider its operation as a feedback loop around any one bunch. Fig. 2.13 shows a conceptual block diagram of such a feedback loop. Longitudinal position of a bunch is measured in the front-end detector and digitized by the analog-to-digital converter (ADC). Signal of a single bunch is digitized at the revolution frequency. In this feedback system we take advantage of the fact that longitudinal oscillations are slow relative to the revolution frequency, so a bunch oscillating at the synchrotron frequency takes many revolutions to complete one period. Longitudinal bunch motion observed on every turn is oversampled and one can reduce the sampling rate (downsample) without significantly affecting the performance of the feedback system [37]. Downsampling is done by processing one out of every N_{ds} samples - N_{ds} is called downsampling factor. Downsampling factor is chosen so that there are 4 to 6 samples per synchrotron oscillation period. Table 2.1 shows the revolution and synchrotron frequencies for several accelerators as well as the downsampling factors being used. After downsampling the bunch signal is processed by a digital signal processor (DSP) using discrete filtering process such as finite impulse response (FIR) or infinite impulse response (IIR). The computed actuator signal (kick) is sent to the holdbuffer. The function of the holdbuffer is to maintain the kick value for a given bunch between updates.

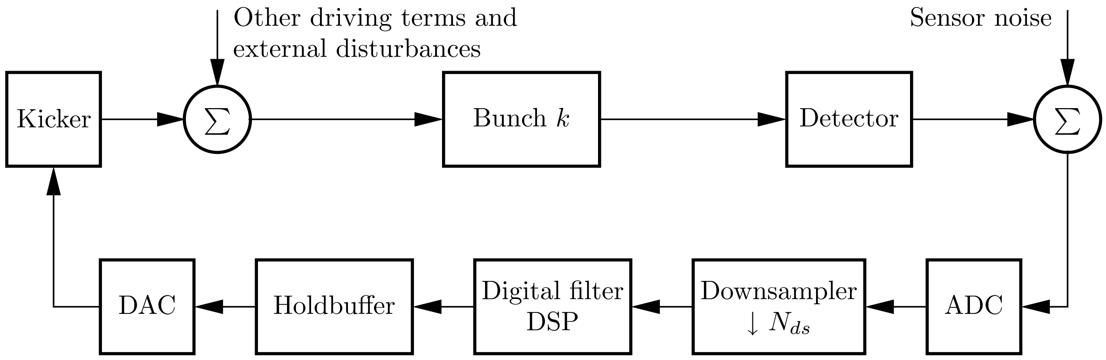

Figure 2.13: One bunch slice of a digital bunch-by-bunch feedback system

Table 2.1: Longitudinal feedback configuration parameters for several installations

Machine	PEP-II	BESSY-II	ALS	DAΦNE
Bunch crossing rate, MHz	238	500	500	368
Number of bunches	1746	400	328	120
Revolution frequency, kHz	136	1250	1524	3067
Synchrotron frequency, kHz	6	8	12	33
Downsampling factor	6	29	31	11
Bunch sampling rate, kHz	22	43	49	272

If one were to produce actuator signal only once per downsampling period the feedback gain would be reduced by approximately $1/N_{\mathrm{ds}}$. Holdbuffer output is converted to an analog signal by the digital-to-analog converter (DAC) and applied to the bunch via back-end modulator and the kicker.

2.10.1 System architecture

The feedback processing channel is implemented in a mix of VXI and VME modules. The following modules are implemented in the VXI format: timing, front-end, downsampler, holdbuffer, and back-end. VME module list consists of the interface board and the DSP board. The overall architecture is illustrated in Fig. 2.14. Next we will consider functionality of each system module.

Timing module

This module receives the RF master oscillator signal and the revolution clock (fiducial). The fiducial signal is first synchronized to the RF clock. Every revolution the timing module counts the number of RF clock edges between the fiducials. If the counted value differs from the programmed number of bunches one of the two error bits is set - missing or extra fiducial. Additional functionality in this module includes programmable bucket trigger signals and temperature monitoring. Due to high density of the ECL logic on the VXI modules thermal management and monitoring is critical for stable system operation. The modules are equipped with multiple temperature sensors which are monitored via the timing module. The temperature signals are compared to the adjustable trip limits and if these are exceeded VXI power supply is turned off to prevent overheating. The control system polls the temperature readings at 1 Hz rate for operator information.

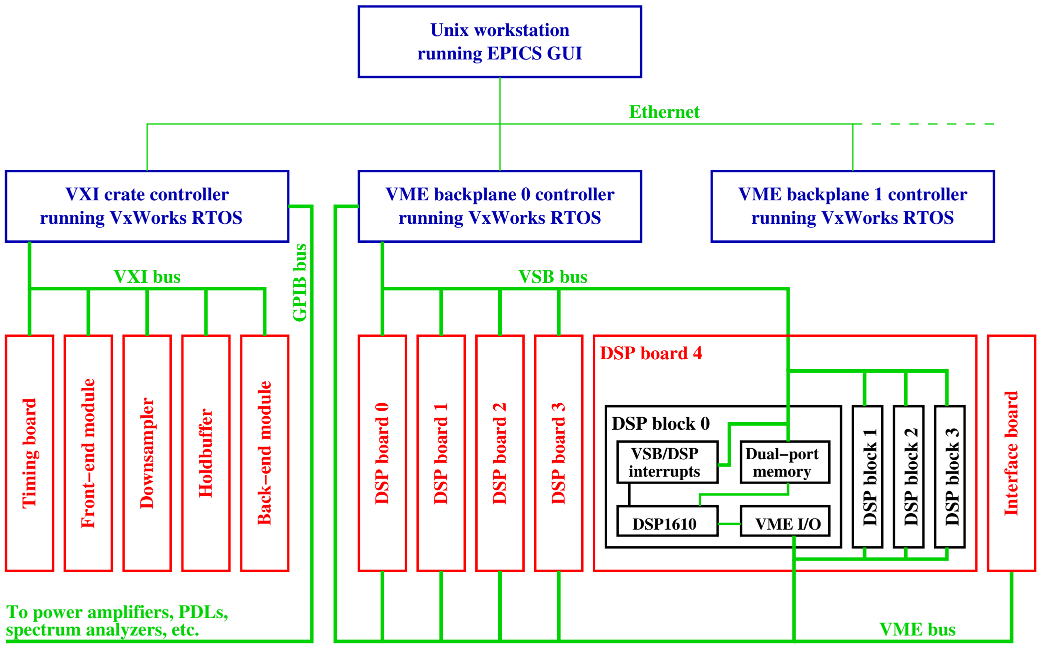

Figure 2.14: Overall feedback system block diagram showing different modules and the interconnect buses. Bunch data transmission from the downsampler to the interface boards and from the interface boards to the holdbuffer via 1.3 Gbps links is not shown.

Front-end module

Front-end module houses most of the analog signal processing involved in measuring longitudinal bunch positions. Some relatively bulky functions, notably the comb generator, are housed in the system oscillator chassis described later. In spite of that we will present here the complete front-end analog signal processing chain.

Bunches passing through the button beam position monitor (BPM) structure generate capacitively coupled (differentiated) fast pulses at the output electrodes [38]. The signals from all four electrodes are summed to reduce sensitivity to transverse beam position and passed through a comb generator filter. This filter is a passive stripline structure which produces a series of uniformly spaced pulses (burst) from a single input pulse [39]. The repetition rate of the burst is chosen to be the $T_{\rm rf}/6$ so that the output signal is bandpass filtered around $6f_{\rm rf}$. The number of pulses is set to 4 producing a 2/3 of the RF bucket long burst. This burst is mixed with the sixth harmonic of the RF frequency locked to the master oscillator. Phasing of the RF-derived carrier is adjusted to produce phase detection - the output signal for bunch m with current i_m is proportional to $i_m sin\phi_m$. For small oscillations around synchronous phase the phase detector output is close to $i_m\phi_m$. The baseband pulse produced in this manner has a flat top lasting $2/3T_{\mathrm{rf}}$ which minimizes sensitivity of the system to sampling clock jitter and synchronous phase shifts which change the time of arrival of the bunch and the timing of the phase detector output. The output of the double-balanced mixer is low-pass filtered to eliminate $12f_{\mathrm{rf}}$ as well as carrier leakage.

Front-end module includes several service functions in addition to the broadband phase detection. These are phase servo loop, RMS detector, and system oscillator digital interface. The phase servo loop integrates the phase-detector output signal and adjusts the phase shifter in the carrier path to maintain zero DC at the output. The loop helps to eliminate DC offsets due to synchronous beam phase drifts. The front-end broadband RMS detector is monitored by the control system and provides information about system stability.

Downsampler

The downsampler module is a programmable stream processor operating on a basic unit of four samples - a group. The downsampling sequence is defined in the random-access memory (RAM) by a series of 32-bit words - a program. The address of the RAM is driven by two counters: the group counter and the turn counter. The group counter is incremented at $f_{\mathrm{rf}}/4$ and is synchronized to the ring fiducial signal. The turn counter is controlled by the downsampler program. At the end of each turn a bit is set in the program to increment the counter. After N_{ds} revolutions the counter is reset to 0. The address into the RAM defines the group of four bunches and the turn in the downsampling sequence. At each address the program word defines the operation to be performed on the particular group on that turn. One can ignore the group or direct it to one of the four possible 1.3 Gbps serial links. Each link can be connected to a VME bus mastering interface board within a separate VME backplane. The interface board accesses up to five four-processor DSP modules which perform the feedback computation. The downsampler program word includes bits to select the physical link for the data as well as the logical number of the DSP module to send the data to. Additional bits are used to define the address in the holdbuffer memory (discussed below), generate holdbuffer control signals and to control program flow. Definitions of the bits making up the downsampler program word are shown in Table 2.2.

The program for the downsampler is generated based on several inputs. The system configuration - map of available DSP boards and links is used to schedule the transactions.

Table 2.2: Bit definitions in the downsampler program word

Bits	Definition
0:4	DSP buffer pointer
5:7	DSP base address
8	Link 0 enable
9	Link 1 enable
10	Link 2 enable
11	Link 3 enable
12:13	Reserved
14	Holdbuffer stop bit
15	Holdbuffer start bit
16:26	Holdbuffer address
27:28	Reserved
29	Increment turn count
30	Clear turn count
31	Halt

The goal is defined by a sampling pattern consisting of all groups to be processed. Normally the pattern will include all groups except for a small ion-clearing gap. Program generation constraints include the minimum time between accesses to one link and the DSP processing time.

Interface board

Data from the downsampler module arrives to the interface board within a VME backplane. The interface board is equipped with 1.3 Gbps link receiver and transmitter. The transmitter is used to send the DSP output to the holdbuffer module. The interface board exchanges information with the DSP boards in a VME read-modify-write cycle. During this transaction the feedback output for the previous group is read out and the new data for the current group of 4 is sent to the DSP board. Bunch data is transferred as four 8-bit wide samples for a total of 32 bits. Address field in the read-modify-write transaction is made up from the base address of the DSP board and the additional information - DSP buffer address, filter and exception bits.

Together with group data the downsampler sends some service information on the serial links. The holdbuffer address is sent to the interface board which forwards it to the holdbuffer with the correction data obtained from the DSP board. DSP buffer address provides a unique group number for a given DSP board. Two filter bits are sent to all DSPs and are used to select one of 4 sets of feedback coefficients. This capability can be used to adapt the system to changes in beam dynamics. It is also utilized to trigger diagnostic measurements. A single exception bit is sent to only one DSP. Within the downsampler it is defined by a register value holding the bunch number of interest. This provides a way to modify the feedback processing for a single bunch.

DSP board

A DSP board is populated with four AT&T 1610 DSPs running at 25 ns cycle time. Each byte of the 32 bit VME bus is connected to the upper eight bits of the parallel input/output (PIO) port of one DSP via a simple interface and the lower eight bits of the PIO receive DSP buffer address from the VME address bus. DSP1610 processors are paired with 32 Kbyte dual-port memories - a total of 4 per board - used to download the filter algorithm to the DSPs and to exchange data between DSPs and the backplane controller. During system operation the VME bus is in continuous use and cannot be utilized for control or diagnostics of the DSP boards. Each DSP board is equipped with the VME subsystem bus (VSB) interface used to control the DSP1610 processors as well as access the diagnostic memory. The dual-port memory is routinely used to make beam diagnostic measurements without interrupting the feedback process.

DSP processors execute a tight, hand-optimized processing loop which computes FIR or IIR filter output based on programmable coefficients. The code executes 12-tap FIR computation including data I/O and synchronization in 1675 ns. The host can use the dual-port memory to reload the filter coefficients to the DSP while the feedback process is running. Coefficient table is first loaded into the memory and then an interrupt is issued to the DSP which then copies the new coefficients to the on-chip RAM. This is an invasive process which interrupts feedback action for a short period of time on the order of 200-300 μs. Only one DSP board is reloaded at a time, so that feedback control is removed from only a small part of all bunches. This technique allows one to modify the feedback controller on-line without disturbing beam stabilization.

Holdbuffer

Functionality of the holdbuffer centers around a block of static RAM (SRAM) large enough to hold data for all of the groups in the ring. The SRAM is being continuously read out at sequential addresses incremented at the group rate of $f_{\mathrm{rf}}/4$. From each location correction values for a group of four bunches are sent to the DAC and converted to analog signal at the RF frequency after parallel to serial conversion. Holdbuffer SRAM has 4 ns cycle time which allows two access transactions per 8 ns group cycle at 500 MHz RF frequency. When kick data is received from the serial link it is written to the holdbuffer SRAM at the address indicated by the holdbuffer address sent from the downsampler via interface board. Each group gets updated once per downsampling period. The kick value is read out on every turn and applied to the bunch between updates - it is being "held".

Back-end

Output signal of the holdbuffer DAC has baseband spectral structure with largest components in the DC to $f_{\mathrm{rf}}/2$ band. Longitudinal kickers typically have bandpass response centered near 1 GHz. The kick signal has to be translated in frequency to efficiently drive the kicker. This is accomplished in the back-end module. Baseband kick signal amplitude modulates a carrier signal to produce the kicker drive. Carrier signal used in this system is a quadrature phase shift keying (QPSK) modulated waveform at an odd multiple of $f_{\mathrm{rf}}/4$. Different installations use QPSK carriers at $9/4f_{\mathrm{rf}}$ (PEP-II, ALS), $11/4f_{\mathrm{rf}}$ (BESSY-II), and $13/4f_{\mathrm{rf}}$ (DAΦNE, PLS). At the first glance this seems to be an odd frequency choice, but there are several strong reasons to select such a frequency. Suppose carrier frequency is an RF harmonic. Optimal timing of a bunch relative to the kick voltage is at the peak. The following bunch arrives at the peak of the carrier as well due to the T_{rf} periodicity. Due to the finite bandwidth of the back-end processing channel and the kicker there is residual kick from the preceding bunch when the new one passes through the kicker. The residual kick is at the peak and produces unwanted coupling between the bunches. Now consider the $9/4f_{\mathrm{rf}}$ carrier. The kick from the preceding bunch is sampled at a zero crossing by the following one therefore the coupling is minimal. However one cannot use such continuous wave (CW) carrier directly since timing bucket 0 on a positive peak we will get zero voltage for buckets 1 and 3, and negative voltage for bucket 2 due to the quarter-period phase shift over T_{rf}. To solve this problem the carrier is phase-modulated with -90 degrees phase shift (QPSK) every RF period.

Another reason for using QPSK carrier modulation is that it equalizes the feedback gain for different eigenmodes as shown in [40]. Yet another motivation is that this modulation scheme allows to center the kicker response away from the RF harmonic where the beam has significant power, reducing the power deposited in the kicker structure from the beam.

Back-end module generates QPSK carrier from $9/4 f_{\mathrm{rf}}$ and modulates it with the hold-buffer output signal. Additional functionality on the module includes optical "woofer" link, back-end RMS power monitoring, and programmable attenuation of the output signal. The "woofer" link is a low-pass filtered version of the holdbuffer output used to modulate the ring RF system reference phase. This "woofer" path uses the RF system as a high-gain low-frequency kicker - analogous to an audio woofer channel. The "woofer" functionality is extensively used to help stabilize PEP-II rings which are equipped with wideband RF system capable of affecting approximately 10 lowest frequency eigenmodes.

Additional components

In addition to the modules described above, a longitudinal feedback system includes system oscillator chassis, VXI and VME/VSB slot-0 processors, programmable delay lines, power amplifiers, and diagnostic equipment.

System oscillator chassis generates carrier signals used in both the front-end and the back-end. Sixth harmonic of the RF frequency is generated in a 24× frequency multiplier driven by the $f_{\mathrm{rf}}/4$ signal. The multiplier combines PLL stage with a cavity-filtered step recovery diode multiplier in order to minimize the output phase noise. The QPSK carrier at 9/4, 11/4, or 13/4 is produced by exciting a step recovery diode with $f_{\mathrm{rf}}/4$ and filtering around the necessary harmonic. Front-end comb generator is also located in the system oscillator. Another important function in this chassis is a fake beam generator. Using a step recovery diode a simulated bunch signal is generated with repetition rate every four buckets. The $f_{\mathrm{rf}}/4$ carrier at the input of the diode can be phase modulated at a several kHz frequency to simulate longitudinal oscillations. A set of coaxial RF switches is used to select between real and fake beam signals. These switches are controlled from the front-end module using a dedicated digital interface. The interface also includes several status signals such as carrier multiplier PLL lock status and output RF signal detectors.

Each VME and VXI backplane in the longitudinal feedback system is equipped with a slot-0 processor. These single board computers (SBCs) use VME and VSB buses to configure and monitor the feedback modules. The VXI crate is controlled by a National Instruments VXIcpu-030 SBC based on Motorola 68030 processor. VME/VSB backplanes are managed by Force CPU-40 or Motorola MVME166. All slot-0 controllers run VxWorks real-time operating system and use Experimental Physics and Industrial Control System

(EPICS) software for operator interface. Each SBC is linked to the outside world by the Ethernet connection.

Programmable delay lines are used in both the front and the back ends to adjust signal timing. In the front end a delay line is placed directly before the ADC input and allows one to adjust the timing of the bunch signal relative to the sampling clock. In the back-end a delay line is located between the modulator output and the power amplifier input. This device provides adjustment of the kick signal relative to the bunch which samples the voltage in the kicker gap. The delay lines are controlled via GPIB interface from the VXI slot-0 controller. GPIB is also used to configure and monitor the MilMega solid-state power amplifiers. The interface allows one to turn amplifiers on and off and to monitor forward and reflected power levels and the fault status.

Several diagnostic instruments such as FFT and spectrum analyzers are also controlled via GPIB. These devices are used in automated timing procedures which determine the optimal settings for the front and back end delay lines.

Chapter 3

Transient diagnostics

Measurements of parameters of the unstable longitudinal beam dynamics present special challenges and are vital parts of feedback system configuration and maintenance. In the open-loop configuration there are several possible situations - the beam may be oscillating longitudinally with amplitudes sufficient to cause nonlinear saturation of the exponential growth, or the beam may be lost when oscillation amplitudes exceed some aperture. In either case the open-loop measurement provides little information on the small-amplitude beam dynamics which determine the performance of the feedback. In the closed-loop configuration the motion is normally damped to the noise floor and little information can be gained from the beam signals. Certainly one can measure the closed-loop transfer function by exciting the system at a summing input and observing the response at some point in the loop. However to extract beam dynamics exact knowledge of the feedback loop transfer function is needed. In addition determining all the unstable modes of interest requires numerous network analyzer sweeps at each revolution harmonic from 0 to $f_{\mathrm{rf}}/2$.

In order to avoid these problems we developed a family of transient diagnostic techniques. In these techniques the system is perturbed and the response is recorded by the feedback system. The longitudinal coordinates of all bunches are sampled and can be recorded in the digital memory for tens of thousands of turns. The individual measurement of the instantaneous phase of each bunch, in conjunction with the long time record of motion is a very powerful source of information about the beam dynamics and machine impedances. The frequency resolution available in these long recordings allows measurement of modal oscillation frequencies with resolution of a few Hz, while sampling all revolution harmonics over the full RF bandwidth. A single transient measurement without external excitation

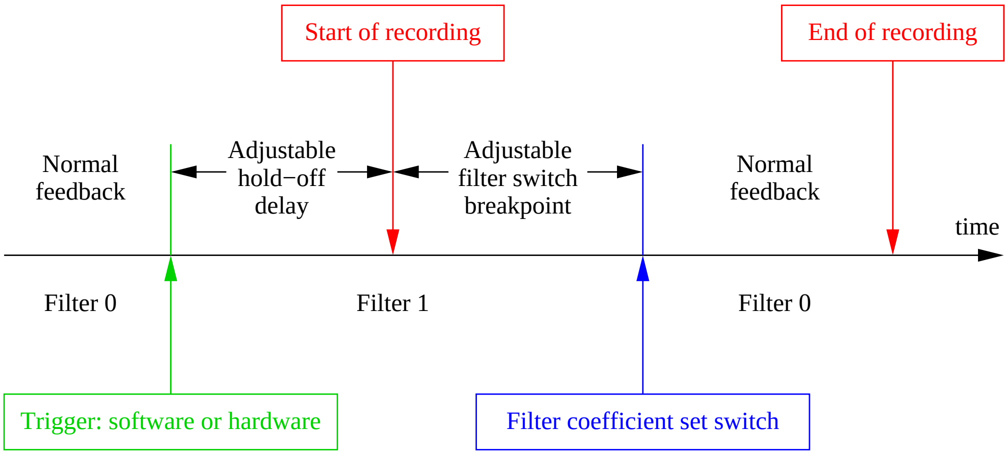

Figure 3.1: Time sequence of a transient diagnostic measurement.

can provide information about the fastest unstable eigenmodes. To study slower unstable or even stable modes narrowband or wideband external excitation is applied.

The general time sequence of events for such a measurement is presented in Fig. 3.1. The diagnostic measurement starts upon a software or a hardware trigger. The software trigger is normally used to acquire beam data under operator control. It is also possible to configure the external software to periodically trigger acquisition of the longitudinal beam motion. The hardware trigger is used when the recording needs to be synchronized to events external to the feedback system such as the injection process or manipulation of other feedback loops, e.g. transverse bunch-by-bunch feedback.

The trigger event causes the feedback system to switch the active coefficient set of the control filter and enter the hold-off period. The hold-off delay provides a way to position the recording window with respect to the longitudinal transient. For the open-loop grow/damp measurement (filter 1 has zero gain) growth of the unstable motion starts after the trigger event. The time before the motion rises sufficiently above the noise floor of the digitized signal depends on the feedback loop gain and the growth rates. The adjustable hold-off period is used to delay the recoding until the motion reaches detectable amplitudes. In cases when the growth rates are large no hold-off delay is used and data acquisition starts immediately after the trigger.

Once the hold-off delay elapses the feedback system starts recording input beam motion in the dual-port memories attached to each DSP. A programmable coefficient set switch breakpoint divides the acquired data into two parts measured under different feedback conditions. For the open-loop grow/damp the first part is the open-loop (growth) transient and the second part is the closed-loop (damping) transient. Upon filling all of the available dual-port memory with acquired data the recording stops and the data can be read out via control bus. After the recording has stopped the DSPs continue providing feedback corrections using the original control filter and the transient measurement can be triggered again.

This transient diagnostic interrupts the normal feedback process for a precisely controlled period of time. There is some risk of losing feedback control if the beam motion grows to large amplitudes in the open-loop portion of the transient. When the feedback loop closes large beam motion saturates the feedback correction signal and the effective loop gain drops. If the reduced gain is insufficient to overcome the instability growth rates, the oscillations continue to grow and the feedback control is lost. If this occurs the beam current must be subsequently reduced to the point where the feedback damping becomes larger than the instability growth rates to recover control. In most cases the risk of losing control can be minimized by conservatively selecting the breakpoint timing as well as the hold-off delay and then iteratively adjusting these parameters.

3.1 Types of transient diagnostics

Here we will concisely discuss each of the distinct transient diagnostics. Every technique is tailored to the measurement of a particular set of beam parameters.

3.1.1 Closed-loop recording

This is a widely used measurement technique and for a very good reason - this diagnostic is completely non-invasive. In this case the coefficient set 1 is the same as set 0. Thus the feedback control conditions are not affected by this measurement and there is no chance of adversely influencing the beam. Several important parameters are measured in this diagnostic. First, this technique is used to verify the feedback control and measure the steady-state motion within the loop. Second, the recording can be used to compute the steady-state actuator effort and check for saturation. Third, bunch-by-bunch currents and synchronous phases can be extracted from the same data set. Asymmetric fill patterns excite longitudinal impedances at the multiples of the revolution frequency. These asymmetric fills produce

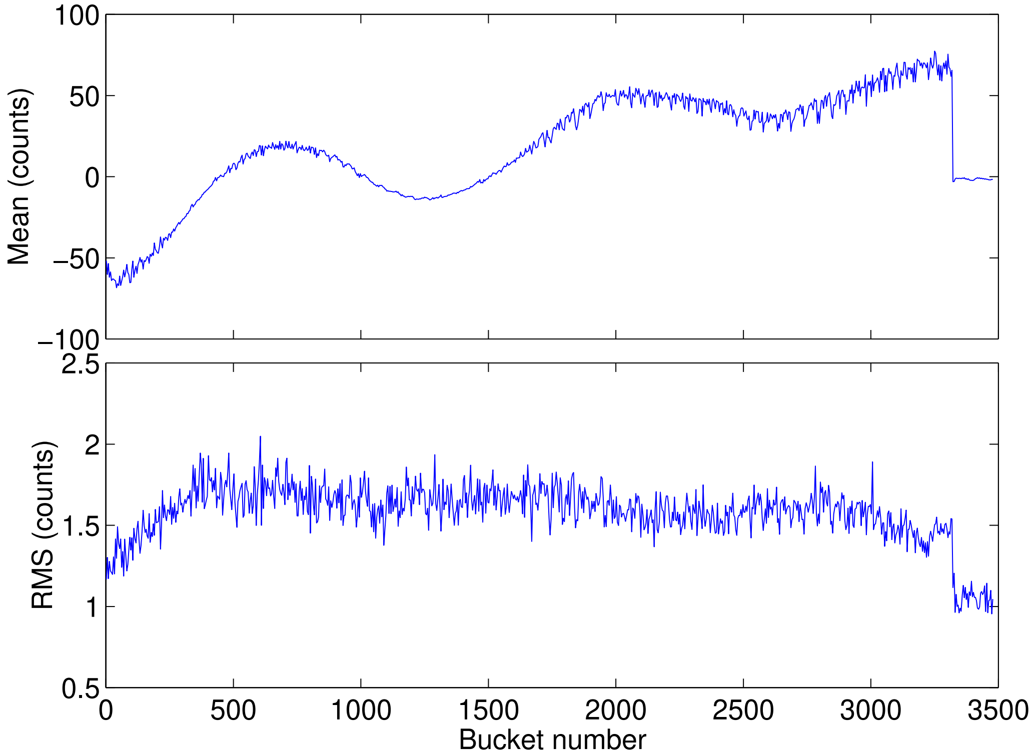

Figure 3.2: Average (top) and RMS (bottom) of bunch data recorded in PEP-II LER at 1732 mA.

a resulting voltage transient which shifts individual bunch synchronous phases, generating the so called gap transient. From bunch currents and synchronous phases impedances at the multiples of the revolution frequency can be extracted as we will demonstrate in Sec. 4.2.2.

A closed-loop recording is illustrated in Fig. 3.2 showing average and RMS values for the bunch data in PEP-II LER. The mean positions of bunches are non-zero due to the synchronous phase transient. Using the record one can check for proper centering of the transient within the ADC input range. Low RMS amplitude of motion reflects good control of unstable modes and is determined by the feedback gain and external noise. From the same data we extract bunch-by-bunch currents and synchronous phases as illustrated in Fig. 3.3. Bunch phases are determined relative to the phase of the master oscillator so there is an overall DC offset from the physical synchronous phases.

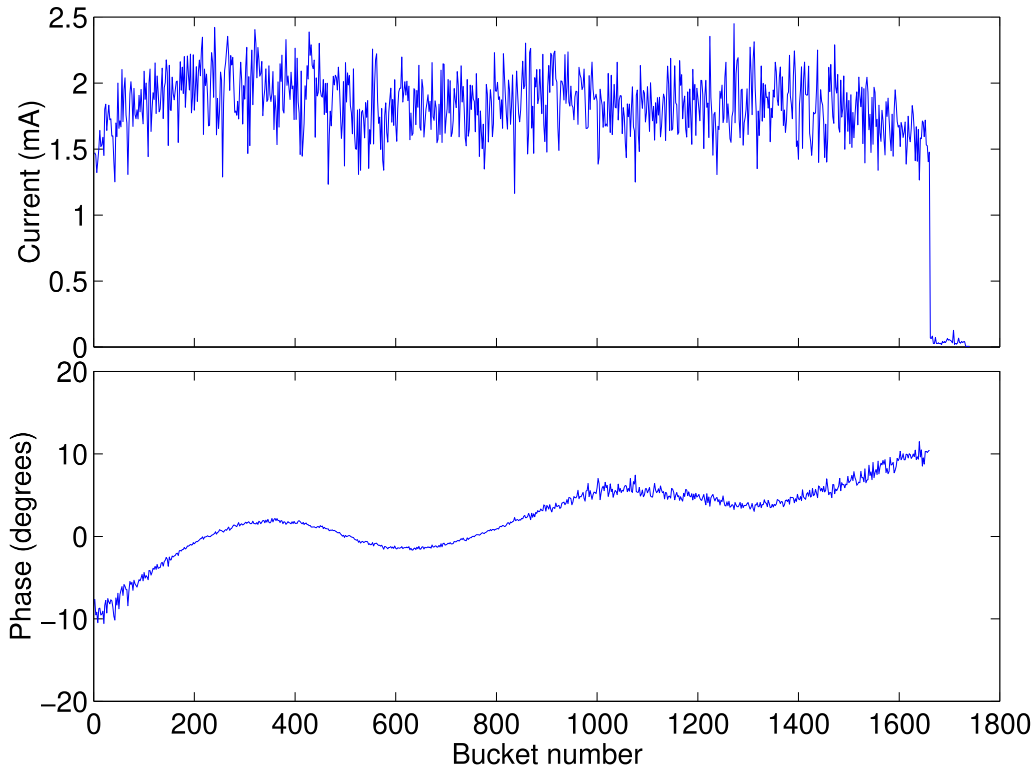

Figure 3.3: Bunch-by-bunch currents (top) and phases (bottom)

3.1.2 Grow/damp measurement

Another fundamental measurement is the grow/damp. In this measurement the feedback conditions are modified during the transient. A traditional choice of filter 1 is a zero gain filter which results in an open-loop measurement before the coefficient switch breakpoint. In the open-loop conditions unstable modes grow exponentially due to noise and the feedback system records the motion of the bunches during the transient.

In some cases the expected growth rates are very high and opening the feedback loop even for a short time is considered risky. Then a modified grow/damp measurement is used with filter 1 configured to provide feedback control at a reduced gain. This experiment does not directly quantify the instability growth rates. These can be estimated using the knowledge of filters 1 and 0 and by comparing the growth rates measured with filter 1 to the damping rates determined directly from the second portion of the transient.

Conversely, when the growth rates are very slow one might use positive feedback for filter 1 as a way to excite the motion to a measurable level.

3.1.3 Driven grow/damp measurement

Stable eigenmodes can be characterized by driving the mode of interest to a measurable oscillation level before the transient. In this case a function generator is configured to output a sine wave at the frequency corresponding to the upper synchrotron sideband of the l-th revolution harmonic in order to excite EFEM l. The excitation signal is given by $A\sin((l\omega_0 + \omega_s)$. The signal is applied to the inverting input of the feedback ADC (in this case used as a summing junction). The amplitude of the excitation is adjusted to achieve steady-state motion significantly above the noise floor of the ADC. Too large an amplitude can lead to saturation of the output kick and must be avoided. The excitation signal is amplitude modulated by the general-purpose TTL output of one of the DSP processors. The effect of the modulation is to turn off the excitation during data acquisition. The resulting data set provides information about the open and closed-loop damping rates of the selected eigenmode.

3.1.4 Injection transient

An injection transient is an example of a measurement synchronized to the external trigger signal, in this case the injection system clock. The goal of such a measurement is to record the bunch motion after charge is injected into the bucket. Such a measurement can be used to quantify the phase and energy errors in the injection system. Complete analysis of the data is complicated since the motion is not limited to centroid oscillations. Injection into an unfilled or empty bucket produces mostly centroid energy and position oscillations. If the bucket is partially filled the motion after injection is a superposition of motions of existing and newly injected charges. The feedback system measures the centroid longitudinal position and, thus, provides incomplete information about the distribution of particles in the longitudinal phase space. More detailed studies of the injection transients are made possible by the use of a streak camera [41].

3.2 Projection onto the even-fill eigenmode basis

The feedback system records the motion of all bunches in a transient measurement. The raw data provides information about the instabilities in the bunch basis. Since analysis of the motion is much simpler in the eigenmode basis a coordinate transformation is needed. As shown in Ch. 2 the transformation from the bunch coordinates to the EFEM basis is a DFT. This transformation is efficiently computed using the fast Fourier transform (FFT). However the transformation has to be applied to a vector of bunch positions sampled simultaneously. The longitudinal feedback uses a single pickup and the bunches are observed sequentially over the period of one revolution. In addition data is recorded with downsampling so that a given bunch is sampled once in a downsampling period consisting of multiple turns. Thus the bunch data must be time aligned before applying the FFT. In addition we bandpass filter the bunch signals around the synchrotron frequency and convert the measured real trajectories to complex signals. The conversion is accomplished using the Hilbert transform and provides estimates of the phase-space motion trajectories [3].

3.3 Estimation of modal eigenvalues

As shown in the previous section from a grow/damp measurement we obtain the complex trajectories of the even-fill eigenmodes. Next we would like to estimate underlying system dynamics that determine the motion. In each transient there are two distinct segments: growth (open-loop) and damping (closed-loop). In the open-loop case, as shown in Section 2.6, motion is defined by a single complex pole. Position of that pole is the modal eigenvalue Λ. For the closed-loop we will assume that system dynamics are dominated by a single complex pole, i.e. the motion is still described by a complex exponential. Thus, for both parts of the transient measured trajectories follow the analytical function

$$y(t) = ae^{\Lambda t} + \epsilon_r + i\epsilon_i$$

where a is the complex factor describing modal magnitude and phase at $t = 0$, Λ is the complex eigenvalue, ϵ_r and ϵ_i are Gaussian-distributed measurement noises. Measurement consists of complex samples $y_k, k = [0, N-1]$ taken at times $t_k = kT_s$. Such a measurement is illustrated in Fig. 3.4 showing real and imaginary parts of the open-loop trajectory of mode 233 at the ALS. Our goal is to estimate values of a and Λ such that some distance function between analytical exponential and y_k is minimized. Let's write noise-free trajectory as a function of real parameters

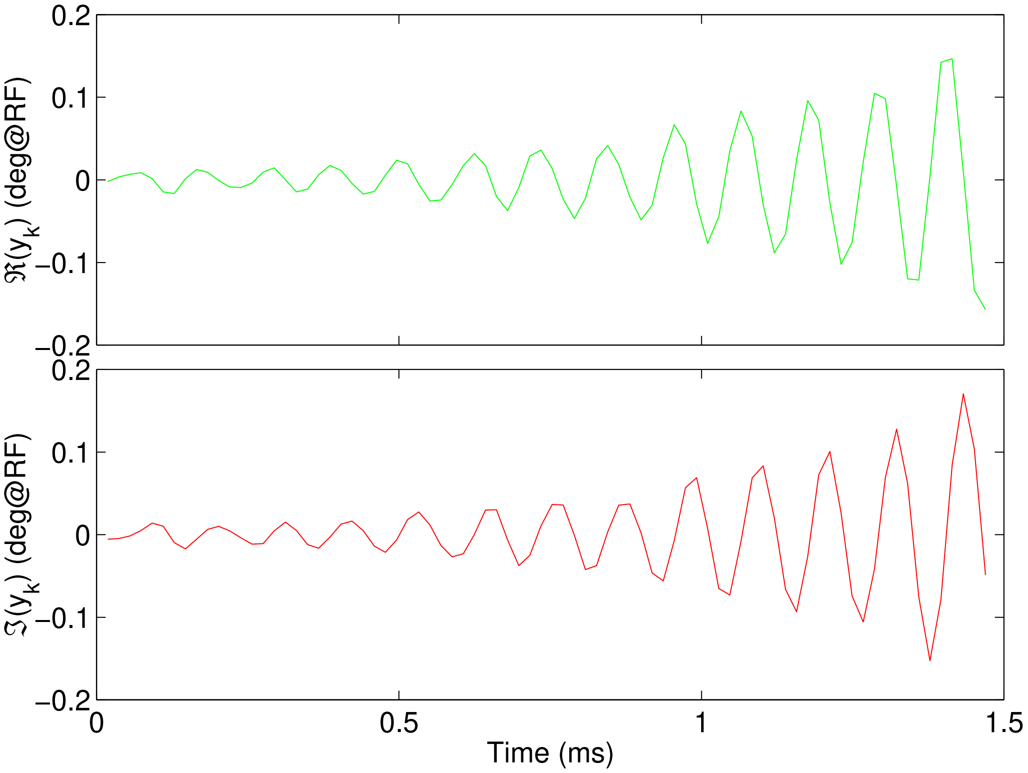

Figure 3.4: Trajectory of mode 233 at the ALS during a growth transient.

$$\begin{aligned} f_k(x) &= x_3 e^{ix_4} e^{(x_1+ix_2)t_k} \\ a &= x_3 e^{ix_4} \\ \Lambda &= x_1 + ix_2 \end{aligned} \tag{3.1}$$

A convenient choice of a distance function is sum of squares:

$$L(x) = \sum_{k=0}^{N-1} |f_k(x) - y_k|^2$$

Minimization of $L(x)$ was implemented using quasi-Newton line search method [42]. In order to improve the optimization efficiency the gradient of the distance function was

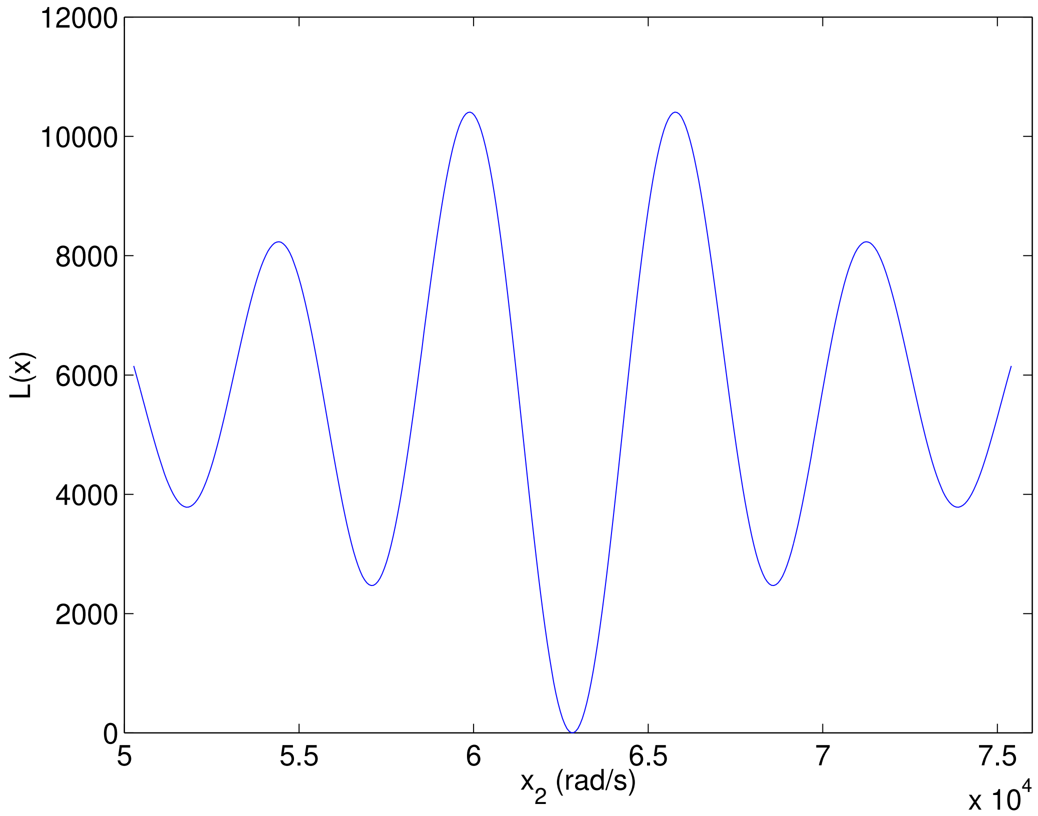

Figure 3.5: Fit distance $L(x)$ versus oscillation frequency for simulated transient data without noise. At the global minimum $L(x) = 0$.

computed:

$$\begin{aligned}
\nabla L(x) &= 2\sum_{k=0}^{N-1}\{\Re(f_k(x)-y_k)\Re(\nabla f_k(x)) + \Im(f_k(x)-y_k)\Im(\nabla f_k(x))\} \\
\nabla f_k(x) &= \begin{bmatrix} t_k f_k(x) \\ i t_k f_k(x) \\ e^{(x_1+ix_2)t_k+ix_4} \\ i f_k(x) \end{bmatrix}
\end{aligned}$$

Unfortunately function $L(x)$ has many local minima. For example its behavior versus x_2 is shown in Fig. 3.5. In order for the numerical optimization to converge on the global minimum, starting point x^0 must be carefully chosen. The following method has been developed to estimate the starting point. The growth rate (x_1^0) is estimated by doing a linear least-squares fit to $\ln|y_k|$. Oscillation frequency is obtained by examining the FFT of the transient and choosing the location of the peak of the magnitude as x_2^0. An obvious choice for an estimate of amplitude and phase is y_0. However for slow growth transients that measurement is dominated by noise and has little relationship to the phase of the exponential. A better estimate is again obtained from the FFT. The sequence given in equation Eq. 3.1 has the following transform

$$\begin{aligned} F(\omega_m) &= x_3 e^{ix_4} \frac{1 - e^{NT_s(x_1 + i(x_2 - \omega_m))}}{1 - e^{T_s(x_1 + i(x_2 - \omega_m))}} \\ \omega_m &= \frac{2\pi m}{NT_s}, m = 0, 1, \ldots, N-1 \end{aligned}$$

Let ω_l be the location of the peak of the magnitude of the FFT. Then $x_2 - \omega_l \approx 0$. Consequently amplitude and phase are estimated as the polar coordinates of

$$x_3 e^{ix_4} = F(\omega_l) \frac{1 - e^{T_s x_1^0}}{1 - e^{NT_s x_1^0}}$$

Numeric optimization provides us with $\hat{x}$ that minimizes the distance function. It is important to know the accuracy of that estimate. Let us consider Taylor series expansion up to quadratic term of the distance function $L(x)$ at $\hat{x}$. At a minimum $\nabla L(\hat{x}) = 0$, so that

$$L(x) = L(\hat{x}) + \frac{1}{2} \sum_{k,l} \frac{\partial^2 L(\hat{x})}{\partial x_k \partial x_l} (x_k - \hat{x}_k)(x_l - \hat{x}_l) + \cdots$$

Disregarding higher-order terms we can express the covariance matrix for the least-squares estimate as [43]

$$V = \frac{2L(\hat{x})}{N-4} H(\hat{x})^{-1}$$

where $H(x)$ is the Hessian matrix of the distance function defined as

$$H_{kl}(x) = \frac{\partial^2 L(x)}{\partial x_k \partial x_l}$$

Standard deviation of $\hat{x}_k$ is given by $\sqrt{V_{kk}}$. Elements of the Hessian are given by

$$\begin{aligned}
\frac{\partial^2 L(x)}{\partial x_m \partial x_n} &= 2\sum_k \left[\Re\left(\frac{\partial f_k(x)}{\partial x_m}\right) \Re\left(\frac{\partial f_k(x)}{\partial x_n}\right) + \Re(f_k(x) - y_k) \Re\left(\frac{\partial^2 f_k(x)}{\partial x_m \partial x_n}\right) \right. \\
&+ \left. \Im\left(\frac{\partial f_k(x)}{\partial x_m}\right) \Im\left(\frac{\partial f_k(x)}{\partial x_n}\right) + \Im(f_k(x) - y_k) \Im\left(\frac{\partial^2 f_k(x)}{\partial x_m \partial x_n}\right) \right] \\
\frac{\partial^2 f_k(x)}{\partial x_m \partial x_n} &= \begin{bmatrix} x_3 t_k^2 & i x_3 t_k^2 & t_k & i x_3 t_k \\ i x_3 t_k^2 & -x_3 t_k^2 & i t_k & -x_3 t_k \\ t_k & i t_k & 0 & i \\ i x_3 t_k & -x_3 t_k & i & -x_3 \end{bmatrix} e^{(x_1 + i x_2) t_k + i x_4}
\end{aligned}$$

In Fig. 3.6 estimation of the eigenvalues is shown for a grow/damp transient taken at the ALS. The feedback is off during the first 15 ms of the transient and unstable mode 233 grows due to noise. Next feedback turns on and starts damping the unstable motion. Absolute value of the error is plotted to demonstrate that $f_k(\hat{x})$ and y_k agree in both magnitude and phase. In this section we've made an assumption that measurement noise is Gaussian-distributed. In order to verify this assumption we will examine the residual errors for the above fit. Figure 3.7 presents a histogram of the real parts of the residuals for the damping part of the transient in Fig. 3.6. It shows excellent agreement with the scaled Gaussian probability density function with variance computed from the residuals.

Joint estimation of growth rate and oscillation frequency has several advantages over techniques presented in [2, 3] that separately estimate each quantity. Exponential fitting to the envelope of motion has large errors when noise levels are high. Joint estimation is frequency-selective and, as a result, is less sensitive to noise. Conversely, when phase angle of the transient is used to determine the oscillation frequency, useful signal-to-noise ratio information contained in the amplitude is discarded.

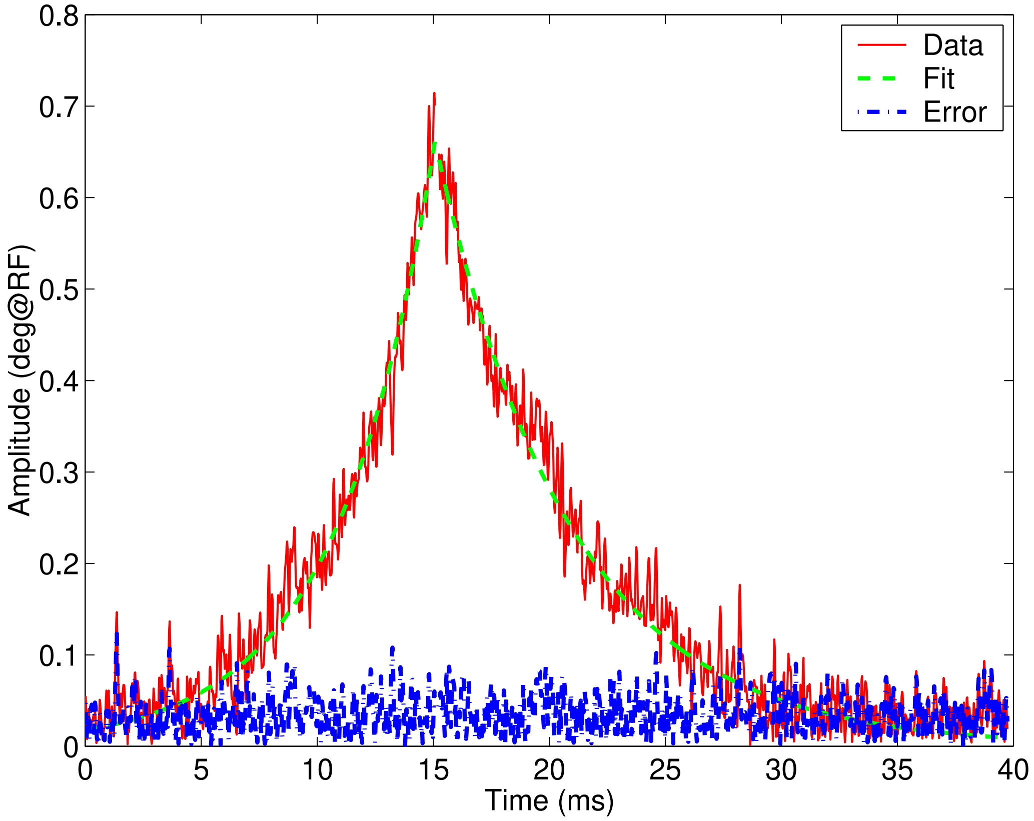

Figure 3.6: A transient at the ALS with mode 233 growing and damping. Magnitudes of the data, fit, and error ($|f_k(\hat{x}) - y_k|$) are plotted. Open-loop eigenvalue is estimated as $(240 \pm 3) + (72014 \pm 3)i\ s^{-1}$ while for the closed-loop we get $-(173 \pm 2) + (72286 \pm 2)i\ s^{-1}$

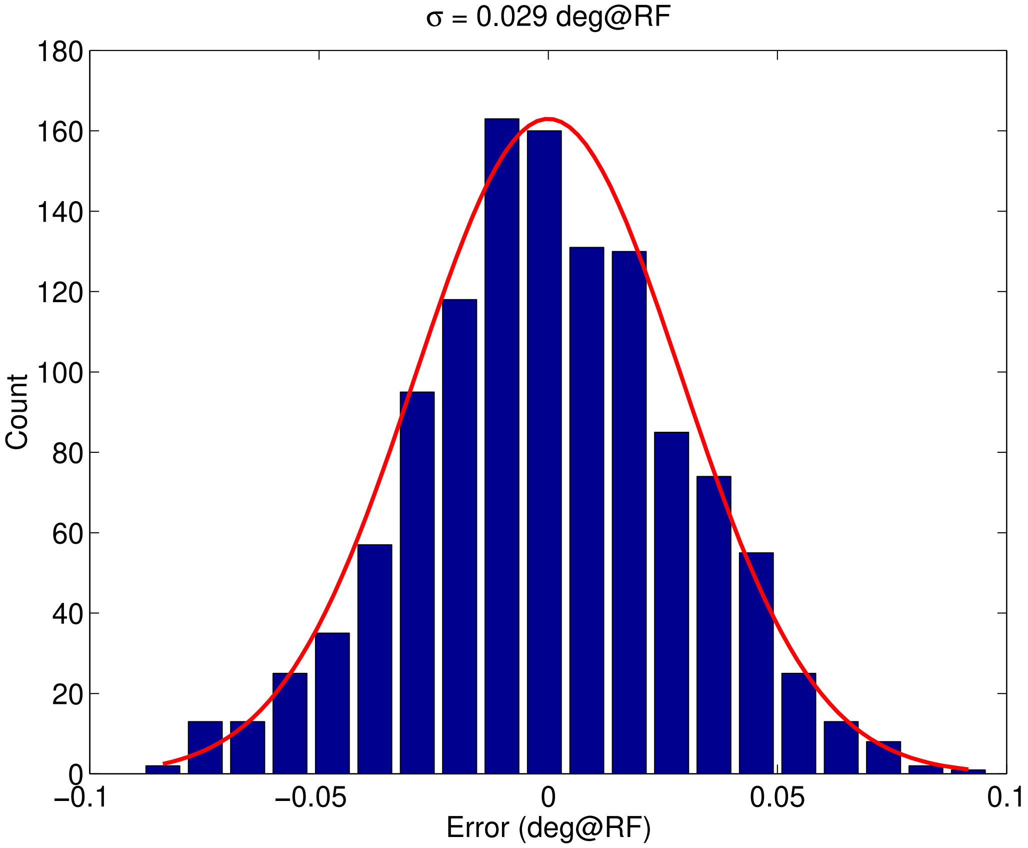

Figure 3.7: Histogram of the real parts of the residuals $f_k(\hat{x}) - y_k$. Expected distribution is shown as well.

Chapter 4

Beam dynamics measurements

Coupled-bunch instabilities must be carefully characterized to be successfully controlled. As we showed in Ch. 3 a digital bunch-by-bunch feedback system provides a powerful tool in the analysis of these instabilities. The ability to quickly measure the complex eigenvalues of the unstable modes allows one to map the dependence of these eigenvalues on the operating conditions. One of the most important dependencies is the relationship of the growth rates and oscillation frequencies of the EFEMs to the beam current. This relationship defines the required performance of the longitudinal feedback system. The range spanned by the imaginary parts of the eigenvalues must be contained within the bandwidth of the feedback controller. Similarly, the growth rates define the minimum gain of the feedback.

Positions of the eigenvalue Λ_l as a function of beam current I_0 trace a curve in the complex plane. We will call this curve an eigenvalue locus based on its similarity to the root locus plot widely used in the analysis of the feedback systems. The shape of the eigenvalue locus depends on many factors. In the simplest case it is a straight line defined by Eq. 2.9. This is true if both the synchrotron frequency and the effective impedance are independent of beam current. However both of these parameters can change with beam current. As we will show in Ch. 6 for accelerators with high-harmonic passive RF systems the synchrotron frequency changes significantly with beam current. The effective impedance can also change due to shifts in the higher-order mode (HOM) frequencies. These shifts are driven by the changes in the resonator temperatures, RF cavity tuner positions, and other effects. In accelerators with broadband RF feedback loops designed to control fundamental impedances of the RF cavities [11], the performance of these loops varies with beam current causing changes in the effective impedance $Z^{\| eff}$.

In general, the functional form of the eigenvalue locus is not known. For the purposes of the feedback controller design as well as feedback simulation having the exact dependence is not critical. If the range of the operating currents of the accelerator is sampled with reasonable density one can interpolate the eigenvalue locus using a low-order polynomial. This technique is presented in Sec. 4.1. However in some cases the functional form is known and can be used to characterize the effective impedance. This brings us to the second important task in the beam dynamics measurements - identifying and quantifying impedances that excite coupled-bunch instabilities.

Stable and reliable operation of coupled-bunch feedback is usually achieved via a combination of several methods. These include impedance control in the design stage of the accelerator components, selection of the feedback hardware (sensors, power amplifiers, kickers), optimization of the control algorithms, and impedance management during machine operation. Impedance management is defined as optimally tuning the resonances to minimize the induced growth rates of the coupled-bunch instabilities. Resonances include the fundamental impedances of the RF cavities influenced via cavity tuners, parasitic HOMs controlled by the resonator temperature or dedicated tuner, and other impedances. In order to optimize the tuning of a resonator we need a way to measure the dependence of its impedance on the control parameter. In Sec. 4.2 we will present two methods for characterizing the complex impedances and optimizing their tuning with respect to the longitudinal growth rates.

4.1 Measurement of the eigenvalue locus

Measurements of the eigenvalue locus have several applications. The locus serves as a foundation of the controller design process by defining the required specifications. Once the controller is designed, the eigenvalue locus is used to test the design and predict the performance in the physical system. Information about the effective impedance can be extracted from the eigenvalue locus.

Sec. 3.3 described how the complex eigenvalue is estimated based on the data from a grow/damp transient. Here we will use multiple transients taken at different beam currents to determine the eigenvalue locus. Figure 4.1 shows the growth rates and oscillation frequencies for EFEM 214 in the ALS versus beam current. Linear fits to the data are shown in both cases. Combining the real and imaginary parts into a complex eigenvalue we obtain

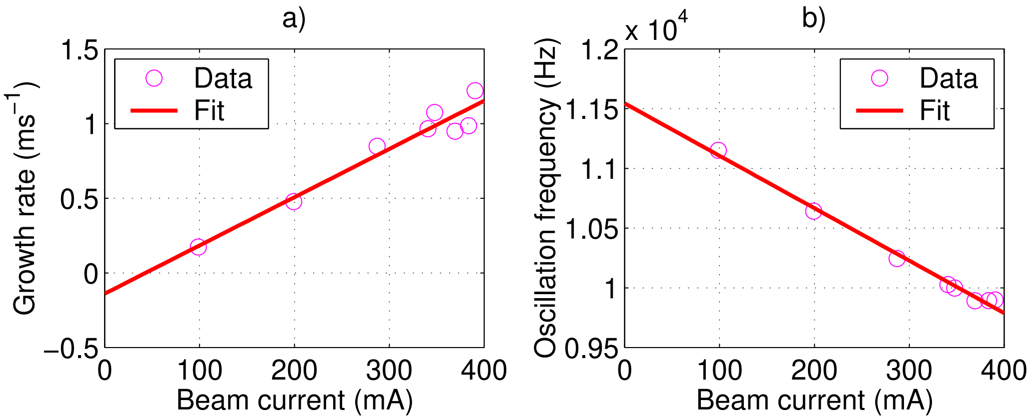

Figure 4.1: a) Plot of measured growth rates of EFEM 214 in the ALS versus beam current. First-order polynomial fit is shown as well; b) EFEM 214 oscillation frequency as a function of beam current.

the locus plot shown in Fig. 4.2. At any given beam current I_0 one obtains the full beam model from the locus value $\Lambda_l(I_0)$ while the current provides the necessary information to compute the loop gain. As a result modeling feedback system performance becomes possible over the full range of beam currents.

By considering the derivatives of the real and imaginary parts of $\Lambda_l(I_0)$ with respect to current we can estimate the effective impedance. In fact, the impedance can be estimated from a single eigenvalue measurement. However using the slope information we avoid errors due to the uncertain knowledge of the synchrotron frequency and the radiation damping. In addition, the derivative is a natural way to use multiple data points in order to reduce the measurement error. This technique is not directly applicable to the locus presented in Fig. 4.2 due to the effect of the high harmonic cavities in the ALS on the synchrotron frequency. Results obtained by applying this technique to the BESSY-II data will be presented in Sec. 4.3.2.

As mentioned before, the locus is not necessarily a straight line. As an example we will examine EFEM 240 in the BESSY-II. In this set of measurements two passive third harmonic RF cavities were tuned in. As a result the synchrotron frequency changes with beam current and the overall dependence of the oscillation frequency on the beam current is nonlinear. Figure 4.3 shows the eigenvalue locus for this mode. The dependence of the oscillation frequency on beam current is fit with the second-order polynomial, while the

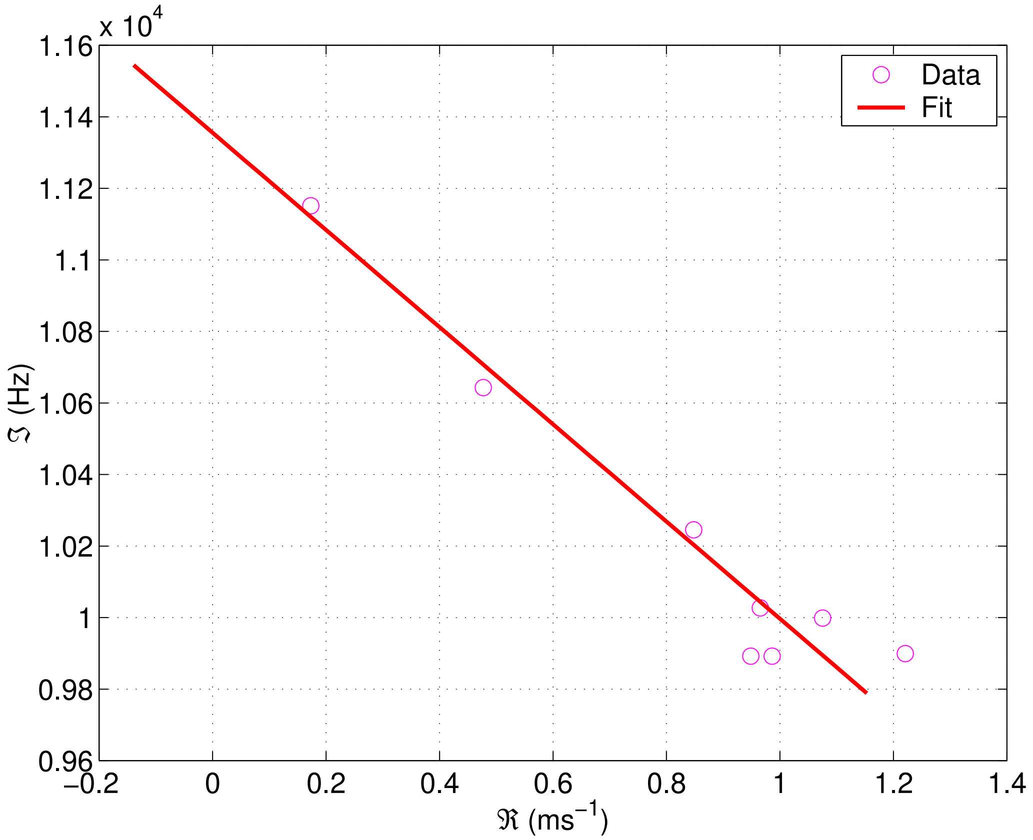

Figure 4.2: Eigenvalue locus for EFEM 214 in the ALS. Polynomial fit is used to plot the locus line from 0 to 400 mA. Measured complex eigenvalues are shown as well.

growth rate is fit with a linear model.

4.2 Measurement of driving impedances

The interaction of charged particles in a storage ring or circular accelerator with the ring impedance determines many important accelerator dynamics parameters. Single and multi-bunch instabilities are the result of interactions of the bunches with the impedance of the machine, and achieving high stored currents requires knowledge and control of the ring components which produce the dominant narrow-band impedances. Traditionally longitudinal impedances in an accelerator have been studied using both bench measurements of the accelerator components as well as beam-based measurements. There are several laboratory

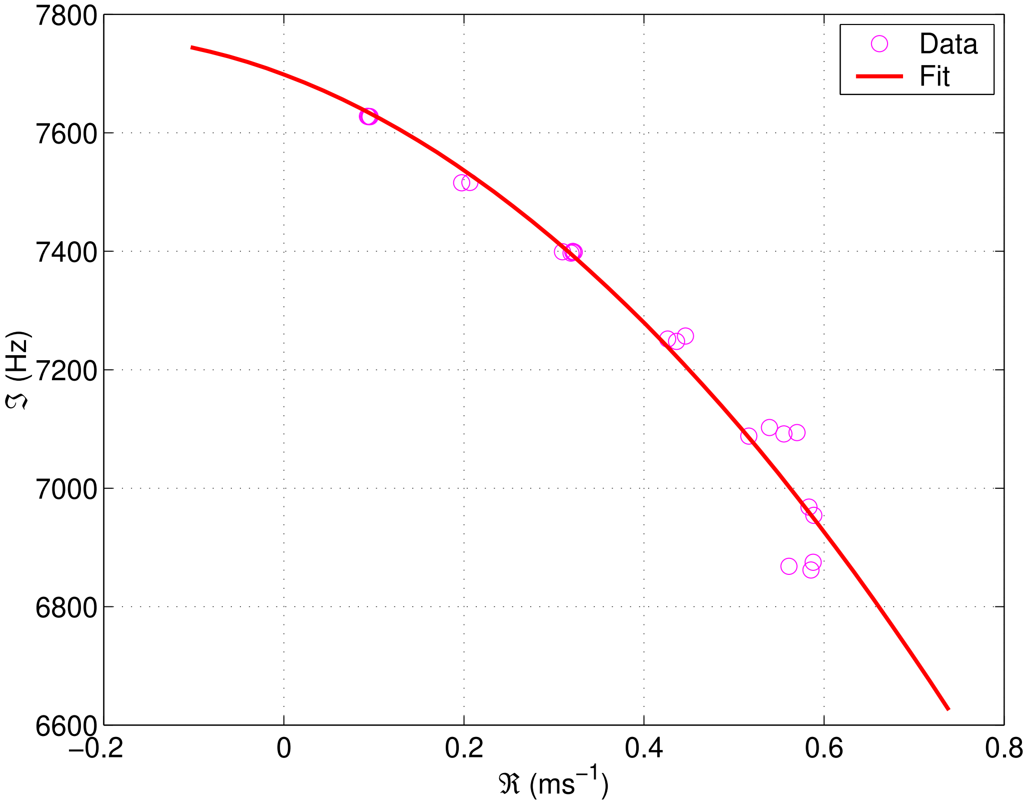

Figure 4.3: Eigenvalue locus for EFEM 240 in the BESSY-II. Polynomial fits are used to plot the locus line from 0 to 400 mA.

techniques to measure impedances of physical components based on propagating pulses on wires through structures, and on measurements of transfer functions of structures as ferrite beads or conducting needles are moved within a cavity volume [44, 45]. Beam-based impedance measurement techniques exist as well. The integrated longitudinal impedance of a circular machine can be measured via the shift in synchronous phase vs. current of a single test bunch [46, 47]. By varying bunch length in such measurements one can estimate the parameters of the broad-band equivalent impedance [48, 49]. However such techniques do not resolve individual higher-order modes. Information about these modes is very important for predicting the thresholds of collective instabilities and for configuring accelerator and feedback systems to stabilize the beam. Beam-based techniques are an important adjunct to numerical calculations and lab bench measurements of RF cavities and components, in that they can measure the actual installed impedance, which is strongly influenced by coupling ports, parasitic components, and environmental factors which can be difficult to include in simulations or lab tests.

Frequency-resolved information about the impedance can be extracted from a measurement of the beam transfer function (BTF) [50, 51]. However such a measurement can only be performed below the instability threshold. In addition network analyzer sweeps have to be repeated for each unstable mode making the BTF approach slow and cumbersome for machines with large numbers of coupled-bunch modes. Yet another method for characterizing the impedance is through observation of cavity coupling probe signals excited by the beam [52]. This approach allows determination of the center frequencies and quality factors of resonant modes, but the coupling of these resonances to the beam is measured only qualitatively. That is, resonances that do not couple to the beam will not be excited, however the effect of the excited resonances on the beam is not clear from such a measurement. In addition, certain resonances within the RF cavity may be weakly coupled to the probe.

In this chapter we present several beam-based longitudinal impedance measurement techniques. These multi-bunch techniques measure the effective longitudinal impedance as a function of frequency in a sampling bandwidth up to the RF frequency. This effective impedance represents a sum of physical impedance components in the frequency bands $Nf_{\rm rf}$ to $(N+1)f_{\rm rf}$ shifted to DC to $f_{\rm rf}$ band and summed (the aliasing process). By comparing the effective (aliased) impedance to the bench measurements of the accelerator components various higher-order mode resonators can be identified and their complex impedance (and parameters such as center frequency and Q) measured.

In Sec. 4.2.1 we describe how impedances can be estimated based on coupled-bunch instability measurements. The relation between impedances and synchronous phase transients is explored in Sec. 4.2.2. The experimental results are presented in Sec. 4.3.

4.2.1 Estimation of longitudinal impedances from coupled-bunch modal eigenvalues

In Chapter 2 we presented the relationship between the longitudinal impedance and the coupled-bunch eigenvalue. Recall that the eigenvalue shift of mode l is given by

$$\lambda_l = \frac{\pi\alpha e f_{\rm rf}^2 I_0}{E_0 h\omega_s} Z^{\|eff}(l\omega_0+\omega_s) \tag{4.1}$$

Once the eigenvalues of the unstable modes Λ are measured via a grow/damp transient it is possible to extract the effective (aliased) impedance vs frequency according to Eq. 4.1. The aliased beam-derived impedance, combined with knowledge about the impedances from bench measurements of ring components may then be properly assigned as an unaliased impedance vs. frequency.

The periodic filling pattern of the beam samples the resonances at the upper synchrotron sidebands of revolution harmonics creating a periodic sampling grid. Wideband impedances are thus sampled at multiple points since they excite multiple eigenmodes. Consequently, for the wideband resonators a single measurement can be sufficient to estimate the three defining parameters: resonant frequency, quality factor, and shunt impedance. Narrowband resonances are sampled by these eigenvalue measurements at a low number of points spaced by ω_0. Thus it is necessary to move the resonant center frequency with finer resolution to fully resolve and measure the impedance parameters. This adjustment of the center frequency can be achieved via adjustment of cavity temperature or via motion of a dedicated tuning probe. Some of the techniques described here are best suited to the measurement of the narrowband impedances while other methods are better for the wideband cases.

4.2.2 Longitudinal wake potential and impedance from measured synchronous phase transients

For the cases when ring fill pattern is uneven additional information about the impedance can be obtained from analyzing the dependence of synchronous phases on bunch currents. Previous work by Prabhakar [53] presents the relationship between the bunch currents, impedances, and synchronous phases. This work is applicable to fill patterns where all buckets are populated, however unevenly. For empty buckets, synchronous phase is not measurable. Theoretically one can define the steady-state synchronous phase for empty buckets, for example as that of a test particle of infinitesimally small charge. However physical measurement of that sort is infeasible. Empty buckets (gaps) are present in most uneven fill patterns and here we will extend the analysis of synchronous phase transients to such fills.

Synchronous phase variations are caused by the beam excitation of the longitudinal impedances at the revolution harmonics. The effect is illustrated in Fig. 4.4 showing 4 bunches with bunch 0 having much larger charge than others. Bunch 0 excites an oscillatory wake field. The synchronous phases of the following bunches shift by the amount

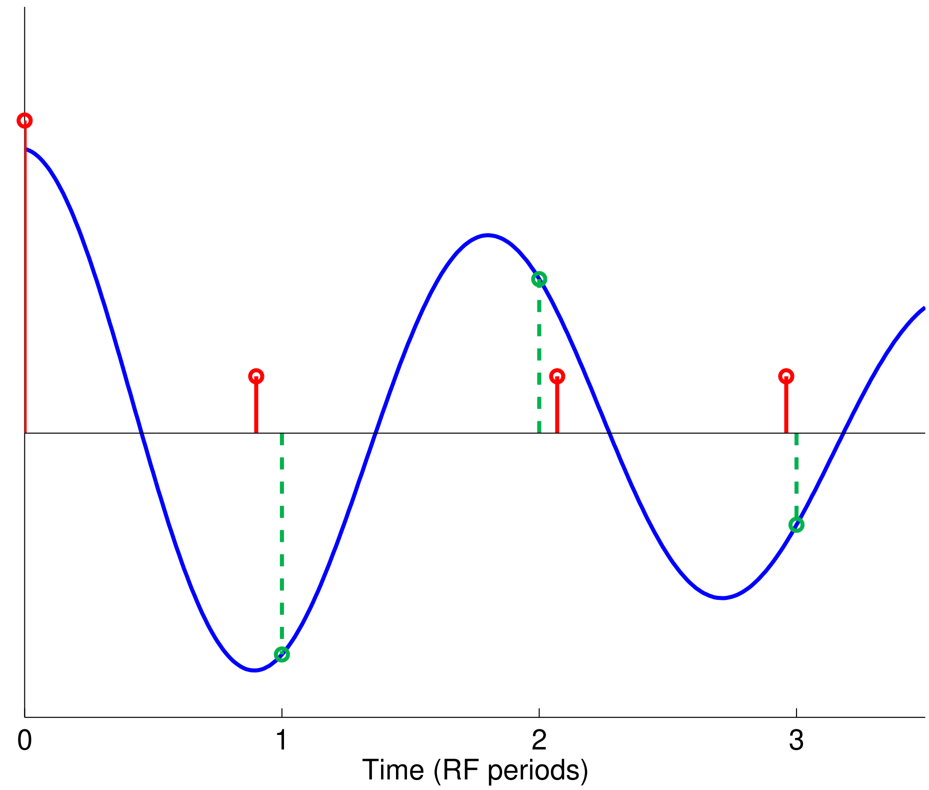

Figure 4.4: Effect of the wake field on the synchronous positions of the bunches. The sketch shows the synchronous positions (spaced by T_{rf}) by green dashed lines, while the bunch arrivals are indicated by the red stems. The height of these stems is proportional to bunch charge. A large bunch (at time 0) produces a decaying sinusoidal wake field. The equilibrium positions of the following small bunches (near times 1, 2, and 3) are altered according to the wake field sampled by the bunches.

proportional to the wake potential sampled at times $t = kT_{\mathrm{rf}}$ (for small shifts). Let us define i_k and ϕ_k, $k = 0, 1, \ldots, N-1$ as vectors of bunch currents and steady-state synchronous phases respectively. Then according to [53] we have

$$\Phi_n = \frac{-N}{|V_c cos(\phi_s^0)|} I_n Z_n^\dagger \tag{4.2}$$

$$Z_n^\dagger = \sum_{m=-\infty}^{\infty} Z^{\|}((mN+n)\omega_0) \tag{4.3}$$

where Φ_n and I_n are DFTs (Discrete Fourier Transforms) of ϕ_k and i_k respectively, V_c is the peak RF cavity voltage and ϕ_s^0 is the synchronous phase in absence of wake fields. Summation in Eq. 4.3 is analogous to that in Eq. 2.10 and describes aliasing of the longitudinal impedance into the 0 to ω_{rf} band. However in this case unlike Eq. 2.10 longitudinal impedance is not scaled by the frequency since in this case we are considering the equilibrium mode.

Let us define set U as follows: $m \in U$ if and only if $i_m \neq 0$ (set of RF bucket numbers with non-zero stored charge). In the following discussion we will assume that a measurement of ϕ_m is available for all $m \in U$. By definition of inverse DFT

$$\phi_k = \frac{1}{N}\sum_{n=0}^{h-1}\Phi_n e^{j2\pi\frac{n}{N}k} \tag{4.4}$$

Let us define $N \times N$ matrix $\mathbf{T}^{-1}$ as follows

$$T^{-1}_{kn} = \frac{1}{N}e^{j2\pi\frac{n-1}{N}(k-1)}, n = 1, \ldots, N, k = 1, \ldots, N$$

Then Eq. 4.4 can be written as $\vec{\phi} = \mathbf{T}^{-1}\vec{\Phi}$. Combining this with Eq. 4.2 we get

$$\vec{\phi} = \frac{-N}{|V_c cos(\phi_s^0)|}\mathbf{A}\vec{Z}^{\dagger} \tag{4.5}$$

$$\mathbf{A} = \mathbf{T}^{-1}diag(\vec{I}) \tag{4.6}$$

Equation 4.5 describes an $N \times N$ linear system of equations with complex coefficients and unknowns. Since ϕ is measured only for $m \in U$, equations described by rows of matrix $\mathbf{A}$ with indexes outside set U are not defined. In addition it is evident from Eq. 4.6 that elements of $\vec{Z}^{\dagger}$ corresponding to weakly excited components of $\vec{I}$ are not well defined. We will choose threshold level $I_{\mathrm{threshold}}$ to separate large components of $\vec{I}$. Selection of the threshold is dependent of the signal-to-noise ratio in the measurement of bunch currents and synchronous phases as well as the values of N, V_c, and ϕ_s^0. Let us define set V such that $k \in V$ if and only if $|I_k| > I_{\mathrm{threshold}}$ (set of revolution harmonic numbers with sufficient excitation). Defining matrix $\mathbf{A}^{UV}$ as consisting of rows U and columns V of matrix $\mathbf{A}$ we can write

$$\vec{\phi}_U = \frac{-N}{|V_c cos(\phi_s^0)|} \mathbf{A}^{UV} \vec{Z}_V^\dagger \tag{4.7}$$

In the DFT of bunch currents DC term (I_0) is the largest, so it is always included in V. Also, as mentioned in section 4.2.1 impedance is hermitian, so that $Z_k^\dagger = Z_{N-k}^{\dagger *}$. Equation 4.7 can be rewritten as a real linear system of M_U equations with M_V unknowns, where M_U and M_V are sizes of of sets U and V respectively. Since only a few frequencies are strongly excited by the fill pattern we normally have $M_V \ll M_U$. Thus we have an overdetermined system of equations. Solving that system in the least-squares sense we obtain $\vec{Z}_V^\dagger$.

4.3 Experimental results

We will illustrate impedance characterization techniques described above with measurements performed at the Advanced Light Source and BESSY-II.

4.3.1 ALS measurements

The goal of the first measurement is to quantify the HOM impedances of the two 500 MHz main RF cavities installed at the ALS. Past longitudinal instability measurements have determined that there are two dominant EFEMs, modes 205 and 233, excited by the impedances in the main RF cavities [2]. Using the lab bench measurements made on the spare cavity identical to the ones installed in the ring mode 205 had been identified as driven by the TM_{011} longitudinal mode at 812 MHz. Mode 233 has two potential driving HOMs, at 2.353 GHz and 2.853 GHz [54]. The aliased effective impedance for ALS, as determined from the bench measurements, is illustrated in Fig. 4.5[1]. The three abovementioned resonances in the main RF cavities are marked as well as a strong HOM at 2324 MHz in the bunch-lengthening third harmonic cavities.

Due to technical limitations it is only possible to fill 320 RF buckets at the ALS. All of the transient measurements described here were taken with 320 equally filled buckets followed by a gap of 8 RF buckets. Since the gap is small the eigenmodes of the fill are

[1]Courtesy of John Byrd

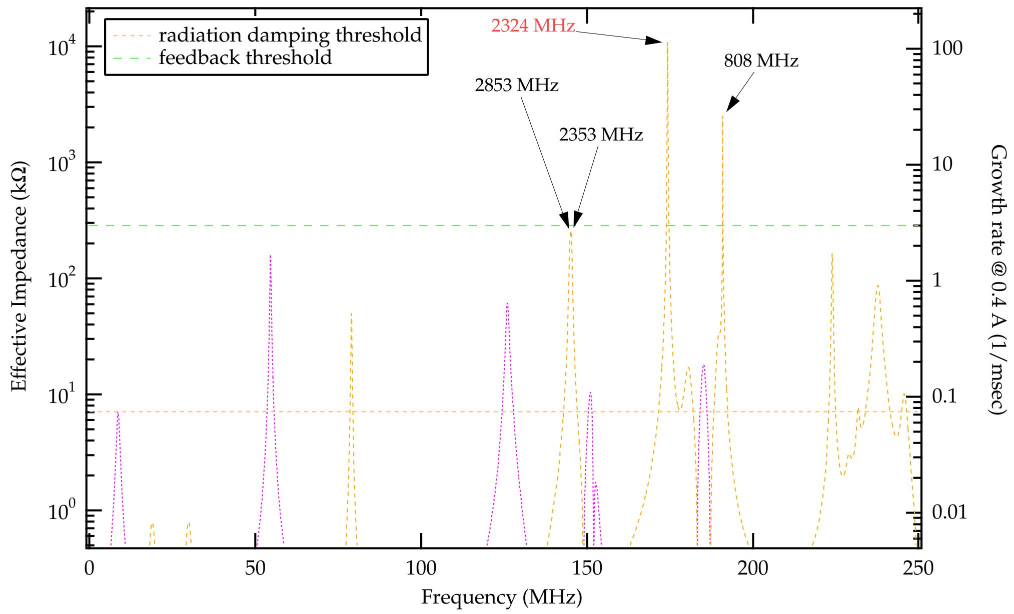

Figure 4.5: Real part of the aliased longitudinal impedance obtained from the laboratory bench measurements. $Z^{\|eff}$ between 250 MHz and 500 MHz is folded into 0 to 250 MHz range and shown in yellow.

close to those of an even-fill[2]. A group of instability measurements was made in order to quantify the measurement noise. From 17 independent growth transients at 80 mA total beam current we find $\lambda_{233} = (0.43 \pm 0.02) - (3.36 \pm 0.02)j$ ms^{-1}. Both the real and the imaginary parts of eigenvalue shift have very small standard deviations. The errors are due to the measurement noise and the variations in operating conditions of the accelerator that cause changes in the growth rates.

In order to characterize the frequency dependence of the impedance we shifted the center frequencies of the cavity HOM resonances by changing the temperature of the cavity. In these measurements we swept the cavity temperature set-point in 11 steps over a range of $15°C$. At each point the temperature was allowed to stabilize and the open-loop eigenvalues of the unstable modes were measured using the transient grow-damp technique. During the sweep of each cavity the temperature of the other cavity was held constant. In Fig. 4.6

[2]According to Eqs. 7.10-7.13 in [6] diagonal coupling matrix of an even fill is modified in this case with off-diagonal entries that are $8/320 = 1/40$ of the magnitude of the on-diagonal entries. Resulting difference between the eigenvalues of the uneven and even fills is negligible.

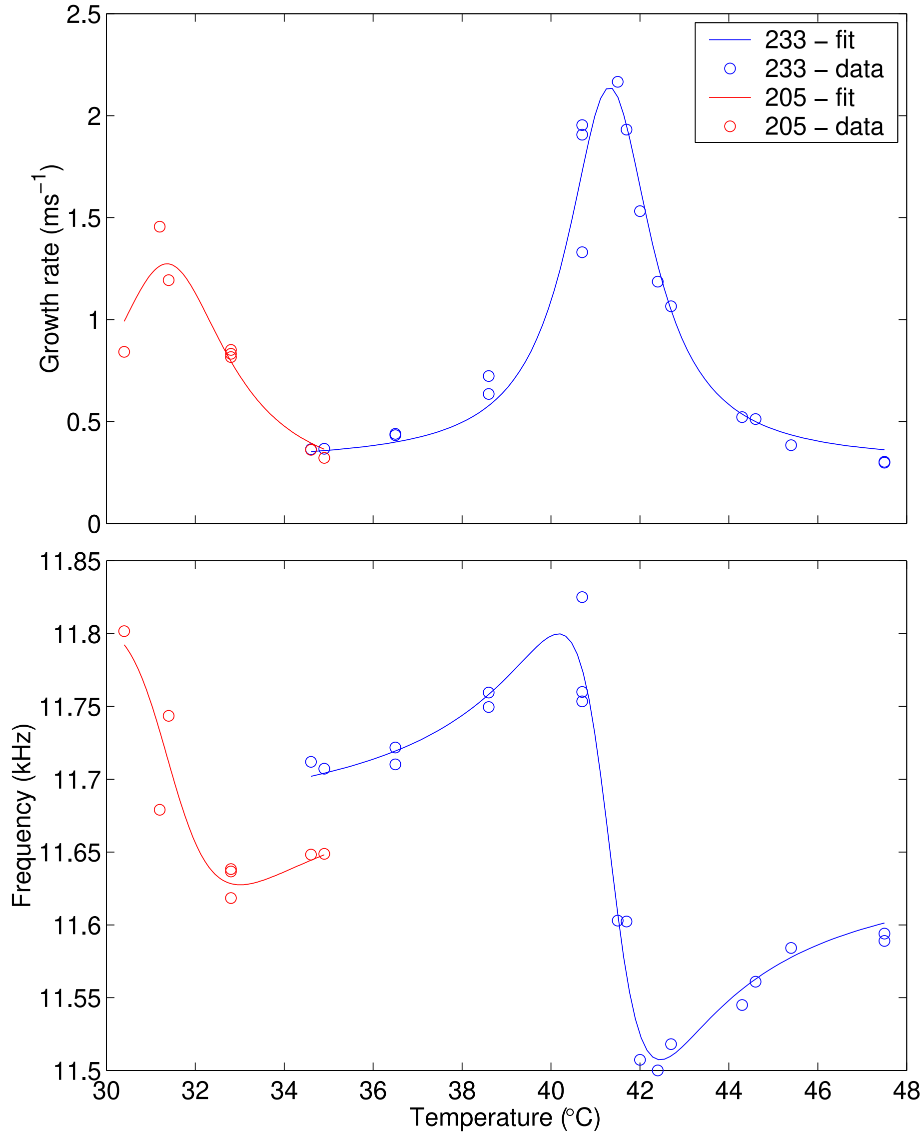

Figure 4.6: Growth rates (top) and oscillation frequencies (bottom) of modes 205 and 233 in main RF cavity 2 normalized to $I_{\mathrm{ref}} = 100\ mA$.

the growth rates and oscillation frequencies of modes 205 and 233 are plotted versus temperature of cavity 2. As the total beam current changes during the experiment, all of the measurements were scaled to a reference current I_{ref} (100 mA) as follows:

$$\begin{aligned}\lambda_m^* &= (\lambda_m + d_r)\frac{I_{\mathrm{ref}}}{I_m} - d_r \\ \omega_m^* &= (\omega_m - \omega_s)\frac{I_{\mathrm{ref}}}{I_m} + \omega_s\end{aligned}$$

where λ_m is the measured growth rate, ω_m is the measured oscillation frequency, and I_m is the beam current during the transient measurement.

The expected effect of changes in cavity temperature is to change the center frequency ω_r of the narrowband HOMs [7, 9]. Since the relative change in frequency is small the relationship between cavity temperature and center frequency is linear of the form

$$\omega_r(T) = \omega_r(T_0) + K(T - T_0) \tag{4.8}$$

where T_0 is the arbitrary reference temperature and K is the temperature to frequency conversion gain. The measurements presented above agree well with this model. However these measurements do not provide a means to distinguish between the two possible HOMs at 2.353 and 2.853 GHz as the source of the aliased impedance. To resolve this ambiguity the ring was filled with a single bunch while a cavity probe signal was monitored on a spectrum analyzer. We observed that change of cavity temperature had very small effect on the magnitude of the revolution harmonics excited within the 2.353 GHz resonance while the amplitude of the probe signal at 2.853 GHz revolution harmonic scaled with temperature in agreement with the growth rate measurements. This leads to two conclusions. First, the resonance measured in the temperature scan is at 2.853 GHz. Second, we can consider the impedance presented by the 2.353 GHz HOM as constant when temperature is changed.

Once the HOM has been identified we can determine the shunt impedance R_s from the λ_m^*. However, in order to quantify the quality factor Q we need to determine the dependence of the resonator center frequency ω_r on the cavity temperature. We will use the relationship presented in Eq. 4.8 and select the reference temperature T_0 to correspond to the center frequency of the HOM at the upper synchrotron sideband of the 233rd revolution harmonic above $5\omega_{\mathrm{rf}}$ (critical temperature as defined in [7]). In this case the induced growth rate is largest at T_0. Equation 4.8 becomes

$$
\begin{aligned}
\omega_r &= \omega_{233} + \omega_s + K(T - T_0) \\
\omega_{233} &= 5\omega_{\mathrm{rf}} + 233\omega_0
\end{aligned}
$$

In order to measure K a single bunch in the ring was filled to a nominal per bunch charge and the cavity probe signal at ω_{233} was measured at two temperatures, T_1 and T_2. At the cavity temperature T_1 the ring RF frequency was increased by $\Delta\omega$ so that the probe signal magnitude equaled that measured at T_2. The temperature to frequency conversion gain K can then be computed as

$$
K = \frac{\omega_{233}\Delta\omega}{\omega_{\mathrm{rf}}(T_2 - T_1)}
$$

Let us consider the eigenvalue of mode 233, Λ_{233}. From Eqs. 2.8,2.9, and 2.10:

$$
\begin{aligned}
\Lambda_{233} &= \Lambda^0 + \lambda_{233} = -d_r + j\omega_s + \lambda_{\mathrm{ext}} + \\
&+ \frac{\pi\alpha e f_{\mathrm{rf}}^2 I_0}{E_0 h\omega_s} Z^{\| eff}_{2.853\,\mathrm{GHz}}(233\omega_0 + \omega_s) \\
Z^{\| eff}_{2.853\,\mathrm{GHz}}(\omega) &= \frac{5\omega_{\mathrm{rf}} + \omega}{\omega_{\mathrm{rf}}} Z^{\|}_{2.853\,\mathrm{GHz}}(5\omega_{\mathrm{rf}} + \omega) \\
Z^{\|}_{2.853\,\mathrm{GHz}}(\omega) &= \frac{R_s}{1 + iQ[\omega/\omega_r - \omega_r/\omega])}
\end{aligned}
$$

where λ_{ext} is mode 233 eigenvalue shift due to impedances other than 2.853 GHz resonance in the cavity of interest. We can parameterize Λ_{233} as a function of 5 real variables: R_s, Q, T_0, $Re(\lambda_{\mathrm{ext}})$, $Im(\lambda_{\mathrm{ext}})$. Using nonlinear least-squares estimation we extract parameter values. Figure 4.7 shows the measured and fitted growth rates and oscillation frequencies for the main RF cavity 1.

Additionally we compared the impedance parameters with the cavity probe measurements taken with a single bunch fill. The signal level at ω_{233} and neighboring revolution harmonics excited by the beam is measured with the spectrum analyzer. These measurements spanning 3 different temperatures were normalized for current and scaled to T_0 using the measured temperature to frequency coefficient K. In Fig. 4.8 the cavity probe signal

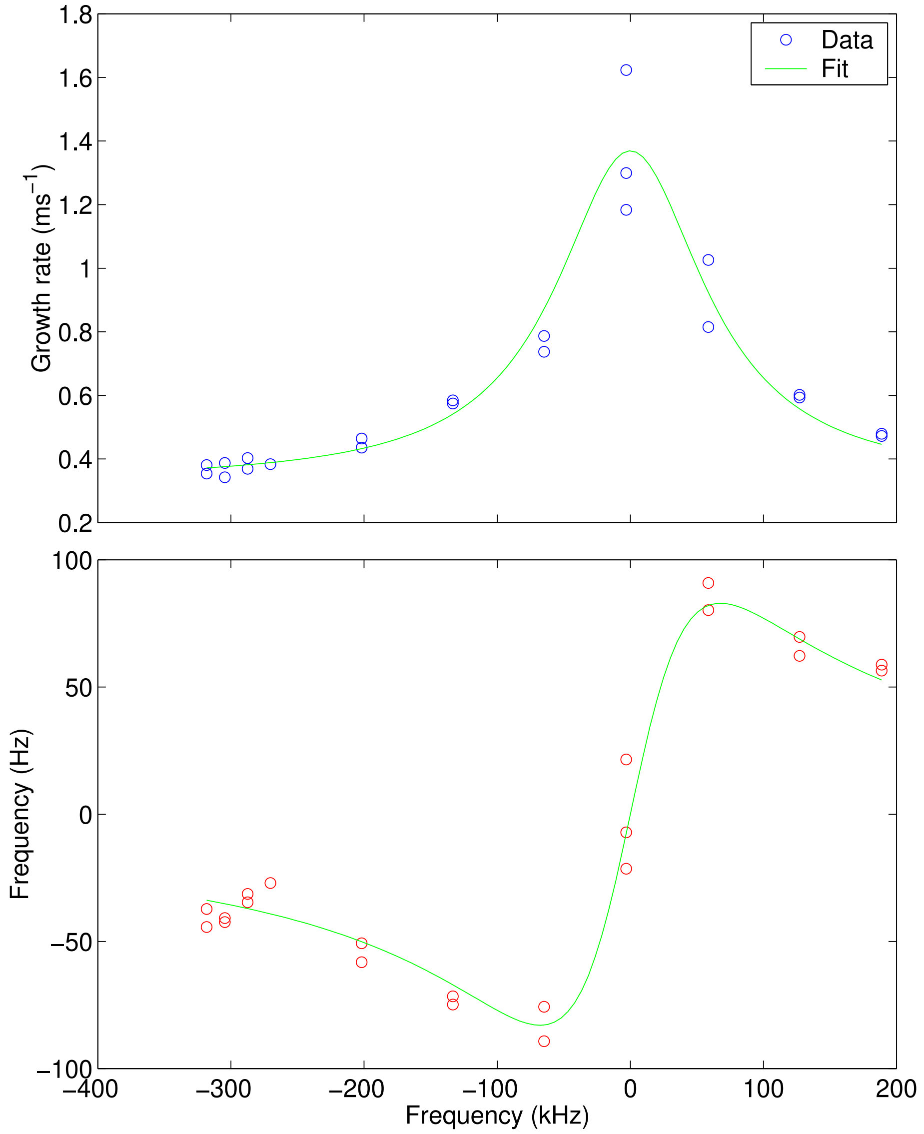

Figure 4.7: Least-squares fit to the measured growth rates (top) and the measured coherent frequency shifts (bottom) versus resonator center frequency for cavity 1.

is shown to have good agreement with the magnitude of the impedance extracted from the transient measurements.

In Table 4.1 results for both cavities are summarized. Data from numerical computations

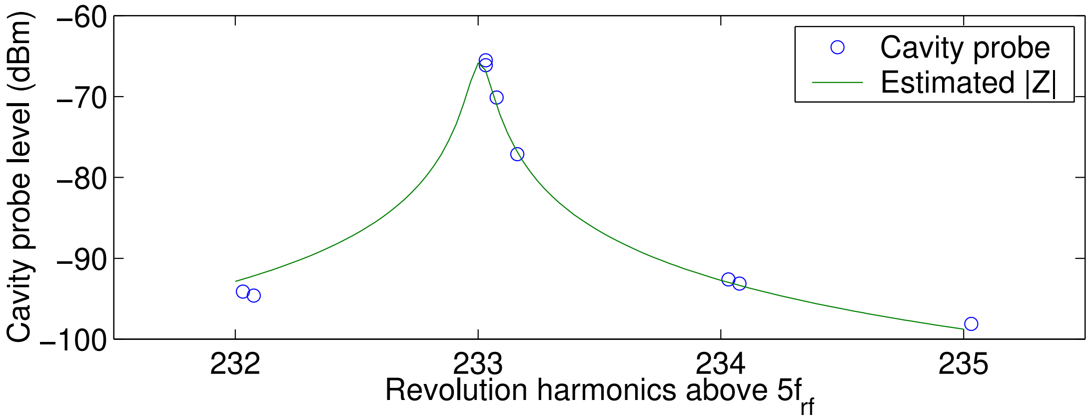

Figure 4.8: Cavity 1 probe signal measured on the spectrum analyzer and the magnitude of the 2.853 GHz impedance.

and bench measurements of a spare RF cavity are included [45, 55]. Both numerical modeling and bench measurements have large errors in their estimates of the shunt impedances [3]. As for the center frequency, error of bench-based measurement is quite small, only 600 kHz. However this degree of accuracy is insufficient for prediction of the coupled-bunch instability growth rates as well as determination of the optimal cavity operating temperatures. Note that characteristics of the 2.853 GHz resonances in two cavities differ significantly. The cavities have RF windows of different designs which can cause variations in the R/Q values. Additionally, the mode in question is close to the beam pipe cut-off frequency and is strongly affected by the evanescent coupling of the fields in the cavity out the beam pipe aperture. The beam pipe attached to each cavity is different and can also lead to changes in the R/Q. Both of the above effects, while not definitive, can cause the observed differences in R/Q.

Using growth rates vs. RF cavity temperature results it is possible to optimize operating temperatures of the main RF cavities. Since temperatures affect the transverse impedances as well as longitudinal impedances, mapping growth rates in horizontal and vertical planes is necessary for a full understanding of the tradeoff. Results in all 3 planes have to be compared with the damping rates generated by the feedback systems in order to select the operating temperatures.

[3] In the bench measurement the shunt impedance is computed by multiplying the measured Q value by the R/Q calculated by URMEL modeling

Since transverse growth rates have not been characterized, we have defined for each cavity a range of temperatures where longitudinal growth rates do not change significantly. As seen in Fig. 4.6 there are two such ranges for cavity 2, 32 $°C$ to 36 $°C$ and 45 $°C$ to 48 $°C$. The lower range of temperatures is not a good choice for machine operation due to closeness to mode 205, so the higher temperature range was selected. Within this range the temperatures have additionally been optimized empirically based on the performance of the transverse feedback systems. Based on these optimizations cavity 1 was set to 46.5$°C$ and cavity 2 to 45.5$°C$ for nominal ALS operation.

At the selected temperatures we can estimate the relative contribution of the four different resonances driving mode 233. The growth rate can be expressed as follows:

$$\begin{aligned}\frac{1}{\tau_{233}} = Re(\Lambda_{233}) &= \frac{\pi\alpha e f_{\text{rf}}^2 I_0}{E_0 h\omega_s} Re\left(Z_1^{\|eff\ 2.3} + Z_1^{\|eff\ 2.8} + \right.\\ &\left. Z_2^{\|eff\ 2.3} + Z_2^{\|eff\ 2.8}\right) - d_r \\ &= \frac{1}{\tau_1^{2.3}} + \frac{1}{\tau_1^{2.8}} + \frac{1}{\tau_2^{2.3}} + \frac{1}{\tau_2^{2.8}} - d_r\end{aligned}$$

where τ_{233} is the growth time of EFEM 233, $\tau_1^{2.3}$ and $\tau_2^{2.3}$ are the contributions to the growth time from the 2.353 GHz HOM in cavities 1 and 2 respectively. Similarly $\tau_1^{2.8}$ and $\tau_2^{2.8}$ are the contributions from the 2.853 GHz resonances in the main RF cavities.

Using extracted fit data we find:

$$\frac{1}{\tau_1^{2.8}} = 0.06\ \text{ms}^{-1}$$

Table 4.1: Extracted impedance parameters for ALS main RF cavities compared to results of numerical computation using URMEL-T and bench measurements of a spare cavity

Cavity	1	2	URMEL	Spare
F_r, GHz	2.8532	2.8532	2.8081	2.85375
$T_0,\ °C$	38.61 ± 0.07	41.32 ± 0.04	-	-
$R_s,\ k\Omega$	55 ± 2	97 ± 3	379	-
$Q,\ \times 10^3$	21 ± 2	24 ± 2	80.9	4
$R/Q,\ \Omega$	2.6 ± 0.2	4.0 ± 0.3	4.7	-

$$\begin{aligned} \frac{1}{\tau_2^{2.8}} &= 0.13 \text{ ms}^{-1} \\ \frac{1}{\tau_1^{2.3}} + \frac{1}{\tau_2^{2.3}} &= 0.37 \text{ ms}^{-1} \end{aligned}$$

As evident from these numbers at the optimized cavity temperatures the contribution to the growth rate of the 2.853 GHz impedance is 1/2 of the effect of the 2.353 GHz impedance.

4.3.2 BESSY-II measurements

These measurements were aimed at quantifying longitudinal impedances at BESSY-II [56]. The machine was filled with 350 uniformly filled consecutive bunches out of 400 to a current of 165 mA. A series of 15 transient grow/damp experiments was conducted over a period of 10 minutes. During the measurements the machine configuration remained unchanged except for the current decaying to 154 mA. There are three unstable EFEMs seen in the data: 281, 396, and 397. According to Eq. 2.9 complex longitudinal impedances can be extracted from the measured growth rates and oscillation frequencies. However these measured values are offset by Λ_0. Since the eigenvalue shift is small relative to the nominal synchrotron frequency ω_s, errors in measurement of ω_s cause large systematic errors in the imaginary part of the extracted impedance. This difficulty can be avoided if we consider the derivative

$$\begin{aligned} \frac{\partial Im(\Lambda_l)}{\partial I_0} &= \frac{\partial Im(\lambda_l)}{\partial I_0} = \\ & \frac{\pi \alpha e f_{\rm rf}^2}{E_0 h \omega_s} Im(Z^{\| eff}(l\omega_0 + \omega_s)) \end{aligned} \tag{4.9}$$

A least-squares fit to the oscillation frequency component returns two coefficients: slope and offset. The slope will be used to extract the imaginary part of the impedance according to Eq. 4.9. The offset part corresponds to the nominal synchrotron frequency. In Fig. 4.9 oscillation frequencies and least-squares fits are plotted for the three unstable modes. Data for mode 281 indicates a very small imaginary part of the impedance (zero slope). At the same time modes 396 and 397 have significant imaginary impedances of opposite signs. Table 4.2 shows coefficients of the linear fits. All three fits agree on the estimate of ω_s. Using the slope information to compute the imaginary part of the impedance and the growth rates for the real parts we get:

Table 4.2: Coefficients of the linear fits to the oscillation frequencies vs. beam current

Mode	Slope, Hz/mA	Offset, Hz
281	0±0.3	7411 ± 46
396	3.43±0.16	7434 ± 26
397	-2.24±0.11	7380 ± 18

$$\begin{aligned} Z_{281}^{\|eff} &= (63.2 \pm 8.1) + (0 \pm 94)j\ k\Omega \\ Z_{396}^{\|eff} &= (59.0 \pm 3.3) + (1115 \pm 53)j\ k\Omega \\ Z_{397}^{\|eff} &= (59.6 \pm 3.7) - (726 \pm 36)j\ k\Omega \end{aligned}$$

In BESSY-II there are four third harmonic RF cavities that are designed to improve beam lifetime. During the above measurements center frequencies of all four cavities were tuned between $3\omega_{\rm rf} - 4\omega_0$ and $3\omega_{\rm rf} - 3\omega_0$ (parked). Our impedance measurement for modes 396 and 397 correlates well with the impedance of the parked cavities. Since for both modes the impedance is sampled far from the resonance, the imaginary component is much larger than the real one. The larger imaginary part seen at mode 396 indicates that some of the cavities are parked closer to the fourth revolution harmonic below $3f_{\rm rf}$. Knowing the original frequencies of these impedances allows us to compute the unaliased longitudinal impedance:

$$\begin{aligned} Z_{396}^{\|} &= (19.7 \pm 1.1) + (373 \pm 18)j\ k\Omega \\ Z_{397}^{\|} &= (19.9 \pm 1.2) - (243 \pm 12)j\ k\Omega \end{aligned}$$

As described in Sec. 4.2.2 we can estimate the impedance by analyzing the synchronous phase transient. In this case we estimate the impedance sampled at the revolution harmonics rather than at their upper synchrotron sidebands. However for the high-Q resonance parked between two revolution harmonics there is little change over the ω_s interval. In Fig. 4.10 the synchronous phase transient in BESSY-II is presented. Top graph shows bunch-by-bunch currents with 350 consecutive buckets filled nearly equally. The periodic excitation of the impedance from the fill with gap generates oscillatory behavior of the synchronous phases

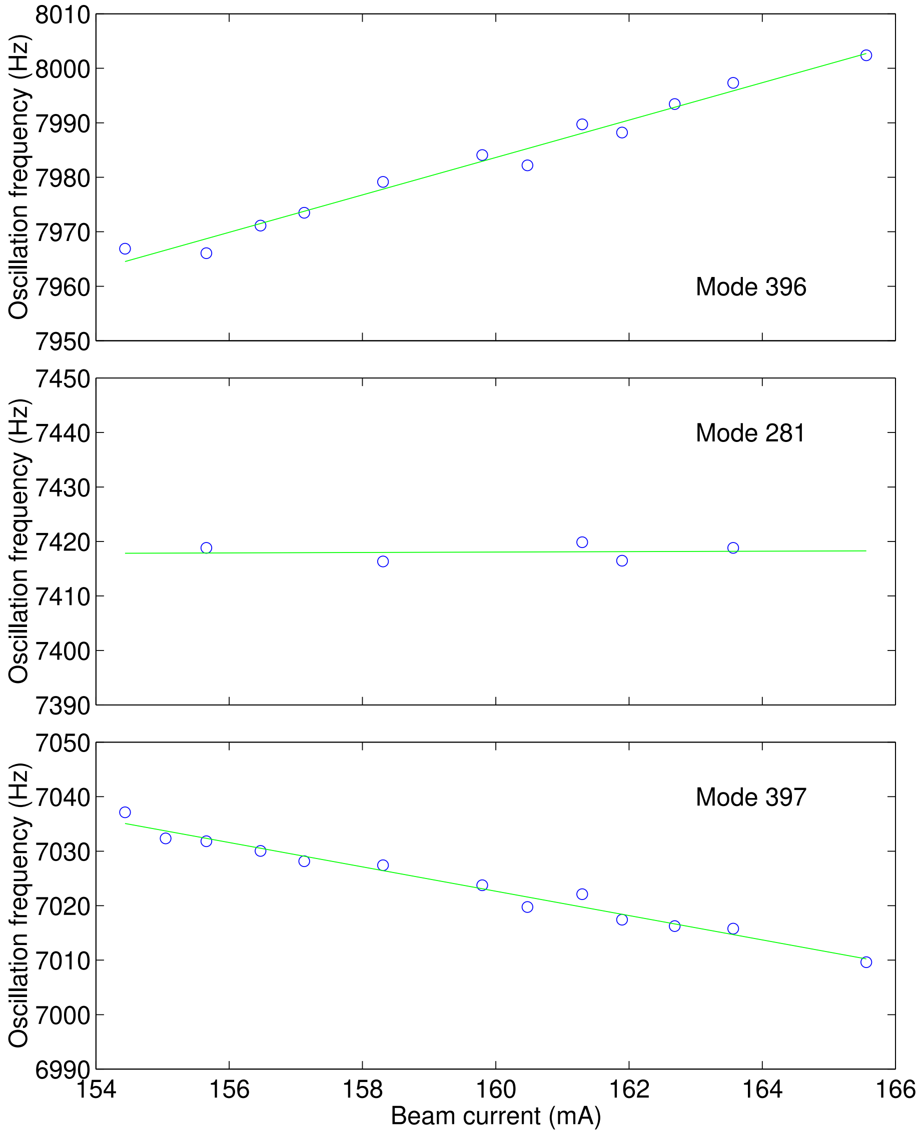

Figure 4.9: Measurements of open-loop oscillation frequencies of three unstable modes and least-squares linear fits to the data. Data for mode 396 shows positive slope of frequency change with beam current indicating positive imaginary part of the impedance according to Eq. 4.9, for mode 281 - nearly zero imaginary part, while mode 397 shows negative imaginary impedance (negative slope).

shown in the bottom plot. Solving Eq. 4.7 in the least-squares sense we obtain the aliased impedances. Least-squares estimate of the synchronous phases is also shown in Fig. 4.10 for

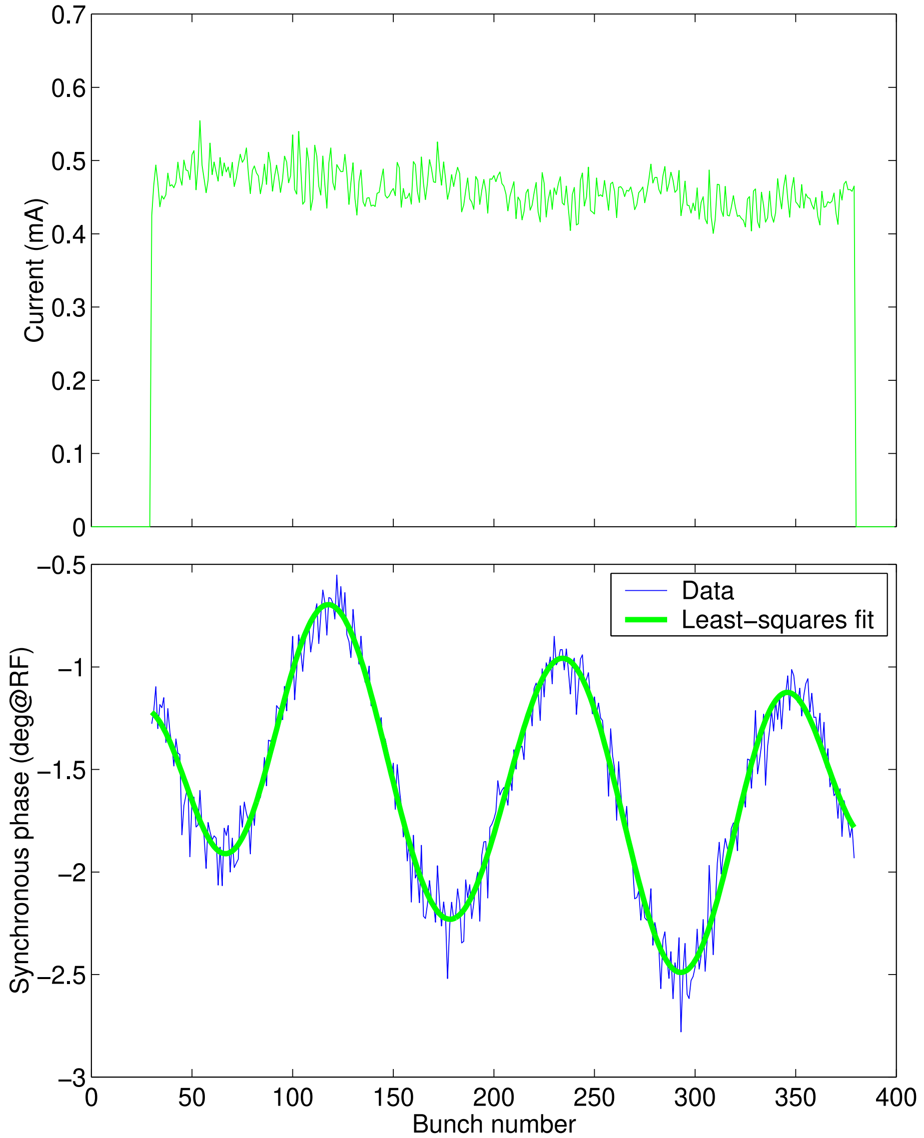

Figure 4.10: Bunch-by-bunch currents (top) and synchronous phases (bottom) extracted from BESSY-II dataset.

comparison with experimental data. Using 15 BESSY transient measurements described above we get the following values:

$$Z_{396}^{\dagger} = (35 \pm 22) + (344 \pm 14)j \ k\Omega$$

$$Z^{\dagger}_{397} = (22 \pm 6) - (233 \pm 15)j \; k\Omega$$

These values show reasonable agreement with the results obtained from the growth rates and tune shifts. The two methods of measuring the impedance can be used together in order to determine unaliased frequencies. This is possible due to the fact that during aliasing into $Z^{\|eff}$ the impedance is scaled by resonant frequency, while in $Z^{\dagger}$ it is unscaled. Thus, from the ratio of $Z^{\|eff}$ to $Z^{\dagger}$ we can estimate the frequency of the physical impedance. From Eq. 2.10 we have (assuming that one physical impedance dominates the aliased function)

$$\begin{aligned} |Z^{\|eff}_{396}| &= \frac{(pN + 396)\omega_0}{N\omega_0} |Z^{\dagger}_{396}| \qquad (4.10) \\ p_{\mathrm{exp}} &= \frac{|Z^{\|eff}_{396}|}{|Z^{\dagger}_{396}|} - \frac{396\omega_0}{400\omega_0} = 2.2 \end{aligned}$$

Since p in Eq. 4.10 is an integer by definition, comparison above indicates that the physical impedance is at $2\omega_{\mathrm{rf}} + 396\omega_0 = 3\omega_{\mathrm{rf}} - 4\omega_0$. This conclusion agrees perfectly with the expected position of the parked third-harmonic cavities.

4.4 Summary

In this chapter the eigenvalue locus has been defined. We presented a locus measurement method based on the growth transient data acquired by the longitudinal feedback systems.

We have also demonstrated two new methods for measuring the impedance of accelerator components using transient diagnostic capabilities of the DSP-based longitudinal feedback systems. The essential feature of these methods is the use of the data sets of bunch motion coordinates sampled over long time intervals. In one approach impedance is computed via measurements of complex eigenvalues of the unstable coupled-bunch modes. The second method extracts the impedance from the relationship between bunch-by-bunch currents and steady-state synchronous phases. The methods extend the capabilities of laboratory bench measurements by quantifying the physical impedances as installed in the accelerator. Dependence of the impedances on operating conditions such as temperature or tuner position can be extracted and used to select optimal working points. By comparing information obtained from growth transients with the analysis of the synchronous phase transients for uneven fills it is possible to determine the spectral position of the driving impedance.

There are certain limitations on the applicability of these methods. Transient measurement depends on the beam being unstable for it relies on growing modes of oscillation to quantify the impedances. Measurement of stable modes is possible via external excitation, but is relatively noisy due to the slow transients involved. Synchronous phase transients can be reliably observed only for large impedances. In addition, only a few revolution harmonics are usually excited by the Fourier components of the bunch current distribution. Both techniques sample the impedances at the spacing of the rotation frequency. Consequently for the measurement of the narrowband resonances a controlled way to shift the resonator center frequency is needed. For the cases when only one of the two techniques is applicable additional information is needed to determine unaliased spectral locations.

Chapter 5

Feedback loop characterization

In order to operate the feedback system efficiently it is important to know the feedback loop dynamics and be able to model the dynamics off-line. In Chapter 2 we developed a model for the longitudinal coupled-bunch instabilities and the digital feedback system. In this chapter the model will be extended to include the detailed models of the sampling process, feedback processing, loop gains and transport delays. Characterization of the beam and the feedback system will be presented using techniques described in Ch. 3. The feedback loop system identification will be demonstrated with the measurement results from several storage rings.

5.1 Feedback system model

The model of a digital bunch-by-bunch feedback system was presented in Sec. 2.10. In that model we considered a one bunch slice of the full system. Such model is conceptually simple since it replaces a MIMO system by a SISO one. However the motion of a given bunch is a linear combination of multiple eigenmodes and is not easily modeled. To get around this problem we will use the property of the bunch-by-bunch feedback system derived in Sec. 2.9. That derivation showed that a bunch-by-bunch feedback system is invariant under coordinate transformations. Thus we can replace the bunch dynamics in our model with the single eigenmode dynamics. The block diagram for the overall system model is shown in Fig. 5.1. The sampling of the bunch signal here happens at the downsampled rate. After feedback computation a zero-order hold (ZOH) is created by the holdbuffer action. All cable and processing delays are lumped together in a single block following ZOH. If the pickup

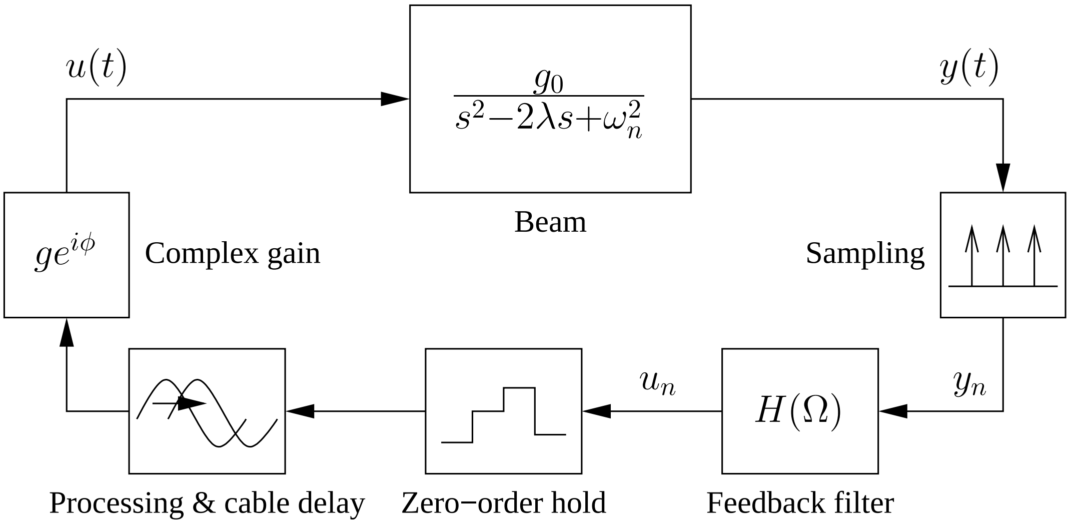

Figure 5.1: Single mode feedback model

and the kicker are right next to each other the overall delay must be close to an integer multiple of the revolution time so that the pickup signal from a given bunch generates the kick signal acting on that same bunch. In practice the locations of the BPM and the kicker can be quite arbitrary and so can be the loop delay.

All of the loop gains have been coalesced into g_0. These gains include the front-end gain in counts/deg, the back-end gain in V/count, and the beam transfer function gain in degrees/V. The overall equation for g_0 is as follows

$$g_0 = \frac{V_{\rm max}}{128} \times \frac{2\pi\alpha e f_{\rm rf}^2}{E_0 h} \times \frac{360 G_{\rm fe} I_{\rm bunch}}{2\pi} = \frac{360\alpha e V_{\rm max} G_{\rm fe} I_{\rm bunch} f_{\rm rf}^2}{128 E_0 h}$$

The first term in the product is the gain from DAC counts to the voltage on the beam, the second term is the transfer function gain in rad/V and the third term is the front-end gain based on the calibration factor $G_{\rm fe}$ in counts/degree/mA. The back-end gain assumes optimal setup where full-scale signal at the DAC produces fully saturated power amplifier output. The amplifier power together with the shunt impedance of the kicker structure gives us the maximum kick voltage [57]. Transfer function parameters σ_l and ω_l are determined by the eigenvalue of the unstable mode. There is an arbitrary complex loop gain parameter $ge^{i\phi}$

used to match the model to the physical system. This gain parameter is used to account for variations in gain and phase nonlinearities of various components over the operating band of $f_{\mathrm{rf}}/2$. Here we assume that the gain and phase for a given eigenmode are constant, that is the changes over the modal bandwidth around $l\omega_0 + \omega_l$ are negligible. This assumption is justified since these parameters vary smoothly over the full $f_{\mathrm{rf}}/2$ bandwidth. Changes in a much narrower modal bandwidth are necessarily small.

5.2 Feedback loop characterization

Once the simulation model of the feedback loop is defined two important questions need to be answered. First, is the model behavior consistent with that of a physical system? Second, if there are differences between the physical system and the model, how does one match the model to the system? In this section both of these questions will be addressed in detail.

Comparison of the feedback system to the off-line model is based on the grow/damp measurements. A properly configured grow/damp measurement provides information about one or more fastest growing eigenmodes in both open-loop and closed-loop configurations. Using the analysis techniques described in Ch. 7 one can parameterize the motion in either case using a complex exponential function. This parameterization produces two eigenvalues: open and close-loop ones with complex conjugate pairs implied. Thus the dynamics in both cases are reduced to those of a second-order linear system. For the open loop case this description is perfectly adequate since the small-signal behavior of an eigenmode is properly described by a complex conjugate pole pair. However in the closed-loop configuration the system is more complicated and the complete model will contain many poles. If there is one strongly dominant pole - that is the pole with damping time much slower than the damping times of other poles - second-order description can still be used. In this case the dynamic behavior of other poles settles very quickly relative to the damping transient of the dominant pole. In some cases there will be several poles with comparable damping times. The resultant damping transient can be strongly non-exponential and the second-order parameterization will break down.

Fortunately pole positions of the closed-loop system can be adjusted by changing the loop gain or, more invasively, by changing the feedback controller dynamics. For the purpose of establishing the model one can reconfigure the feedback so that there is one dominant closed-loop pole. Measurements of the closed-loop dynamics are then consistently described by the second-order parameterization. Once the model is created, however, one would like to verify the feedback performance in the worst-case conditions, that is at the highest possible stored current (which generates the fastest growth rates) and with the optimal feedback system setup. The optimal configuration of a stabilizing feedback system is the one that produces maximum damping of the closed-loop poles. In many cases that leads to multiple poles with identical or very similar damping making impossible model extraction described in Ch. 3. Future work on this subject could extend the model parameterization by using system identification methods.

Above we described a method to quantitatively compare the physical system to our model provided that there is one dominant closed-loop pole. If the two poles being compared are different we need to modify the model to agree with the measurements. For that purpose we use the complex loop gain. A single grow/damp measurement is not sufficiently reliable to verify the model since one can almost always match two complex numbers using two parameters controlling the model. To avoid these issues we verify system model based on a large set of grow/damp measurements using different feedback controller responses. Two sets of closed-loop poles are matched using numerical optimization in a two dimensional space with the two parameters being the absolute value and the phase of the complex loop gain. The goal function is defined as follows:

$$f(g,\phi) = \sum_{k=1}^{M} \left| 2\frac{p_{\mathrm{ex}} - p_{\mathrm{model}}(g,\phi)}{p_{\mathrm{ex}} + p_{\mathrm{model}}(g,\phi) - 2\lambda_l} \right|^2$$

For each of M measurements we consider the distance between measured pole p_{ex} and modeled pole p_{model} normalized by the average distance of these two poles from the open-loop eigenvalue. The normalization compensates for the necessarily larger measurement errors in the high-gain configurations which significantly shift the open-loop pole. This can be represented geometrically as shown in Fig. 5.2. When taken on its own the error magnitude $|\vec{D}|$ does not tell us whether the error is acceptably small or unacceptably large. On the other hand, if the error is normalized to the magnitude of the feedback effect on the system - the very effect we are trying to quantify - it becomes a relative measure. For a single measurement error of 1 or 100% ($|\vec{D}| = |\vec{C}|$) is obviously large. Normally after model complex gain optimization we expect the errors under 0.2.

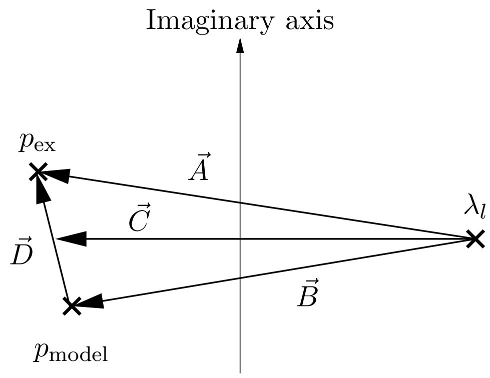

Figure 5.2: Geometric representation of the error function. Average distance of the two closed-loop poles (measured and modeled) from the open-loop pole is given by vector $\vec{C}$. Error function contribution for one measurement is given by $|\vec{D}|/|\vec{C}|$.

Next we will illustrate the feedback loop characterization methods with measurements from several accelerators starting from the Advanced Light Source.

5.2.1 ALS results

The loop characterization experiment at the ALS performed at relatively low currents in the range of 69-76 mA consists of 16 individual grow/damp measurements. To make sure that the agreement is independent of a particular controller multiple feedback controllers were used. To generate these controllers standard FIR filters based on a sampled sine wave were used. Variation of controller characteristics was achieved by changing the phase shift of the sampled sinusoid resulting in changes of the controller phase response at the synchrotron frequency. In all measurements eigenmode 233 driven by the impedances in the main RF cavities grew significantly above the noise floor. The following results are obtained by analyzing and modeling the motion for that eigenmode.

In Fig. 5.3 a histogram of error values is shown for the above 16 measurements. Optimized loop gain is 1.54 ± 0.02 and loop phase is -1.0 ± 0.9 degrees. A gain larger than unity indicates that the power amplifier is being over-driven and will reach full output power at the DAC setting below full scale. From these measurements we also extract the parameters for other unstable eigenmodes. Model optimization for the EFEM 212 driven by the impedances in the third harmonic cavities produces the gain of 1.20 ± 0.04 and the phase of 13.1 ± 1.9 degrees. The complex loop gain difference between the two eigenmodes

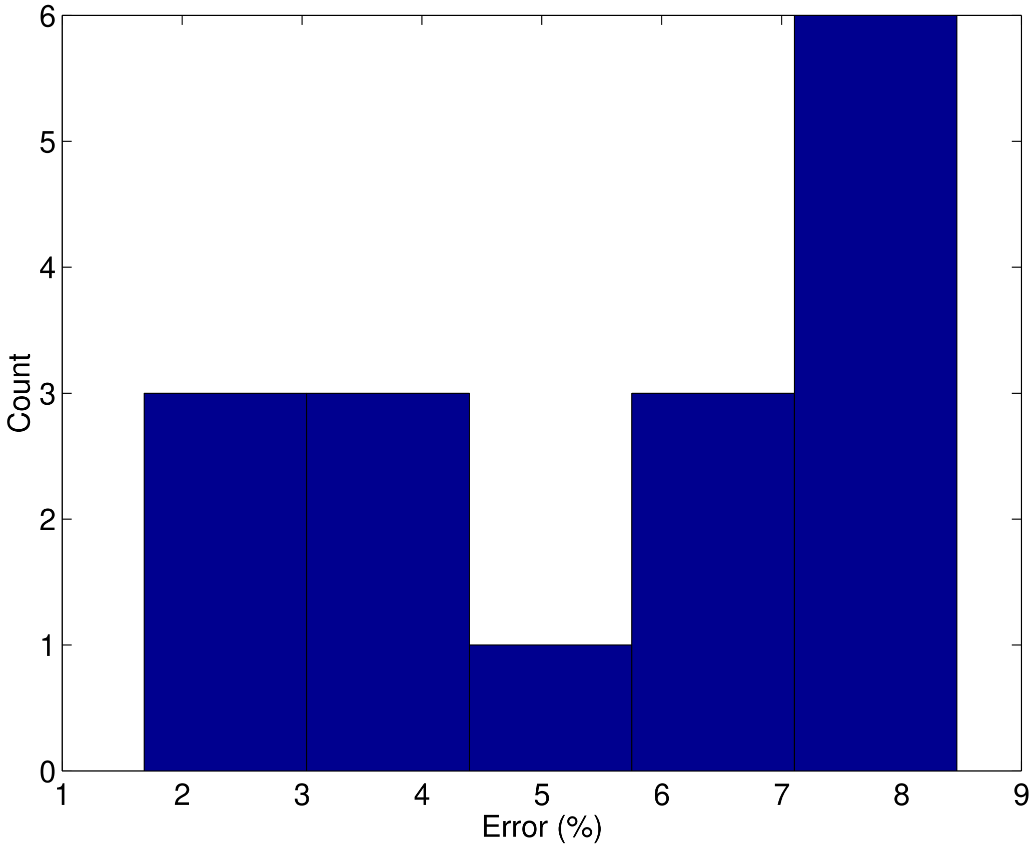

Figure 5.3: Histogram of relative errors for ALS feedback loop characterization experiment

points to non-ideal behavior of some feedback loop components and was later traced to an unterminated monitor port sending back a significant reflection.

In the experiment the controller phase was swept over 100 degrees and we can present the results as a function of that parameter. To do so requires making one assumption. Changing controller phase produces some variation of controller gain at the synchrotron frequency. In addition the beam current decays during the experiment. To meaningfully present the results we need to normalize for these variations in the loop gain. We will assume that the pole shift from open to closed loop scales linearly with loop gain. This assumption is valid unless the gain is very high or the gain changes over a wide range. In our case the gain changes by only 7 dB. Normalized pole shifts are presented in Fig. 5.4. The real parts of vectors $\vec{A}$ and $\vec{B}$ are plotted as the damping rate shift while the imaginary parts give us the frequency shift. The plot shows that both the real and the imaginary parts

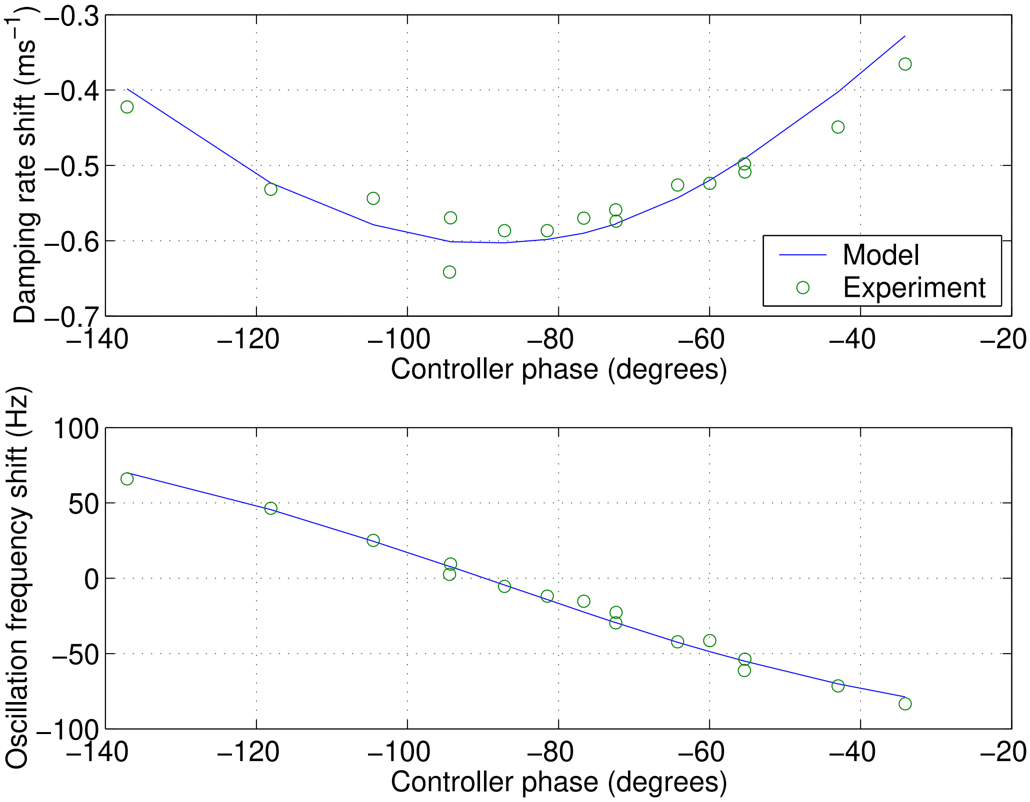

Figure 5.4: Shifts in damping rate and oscillation frequency between open and closed loop states (real and imaginary parts of the pole shifts respectively).

of the closed-loop poles are well estimated by the model. The experiment as well as the model confirm that optimal damping is achieved when the feedback path phase shift (from $y(t)$ to $u(t)$ in Fig. 5.1) is 90 degrees. This results in resistive damping with zero frequency shift between open and closed loop configurations.

5.2.2 DAΦNE results

A similar set of measurements has been performed at the DAΦNE collider in order to verify the system model to be used in controller design. At beam currents from 53 to 90 mA 11 grow/damp transients were recorded. The accelerator was filled with 20 bunches with 4 RF bucket spacing leaving a gap of 1/3 of the ring. In Fig. 5.5 results of matching the model to the measurements for EFEM 21 are shown. The numerical optimization returned loop gain of 1.04 ± 0.03 and phase of -1 ± 2 degrees.

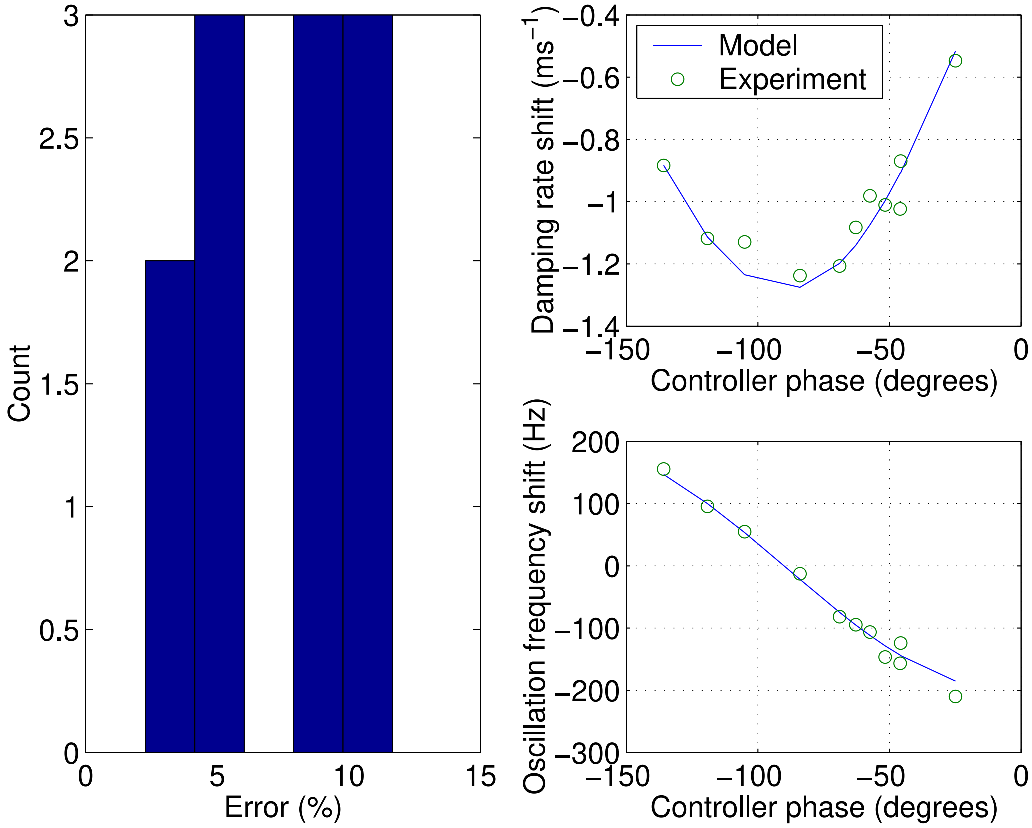

Figure 5.5: Loop characterization results for DAΦNE positron ring.

In the electron ring it is possible to uniformly fill all 120 RF buckets. Loop characterization was done for the three fastest-growing eigenmodes: 21, 89, and 101. Table 5.1 shows the model reconstruction results. This installation shows relatively little gain variation. There is some phase nonlinearity, however its better characterization requires more grow/damp measurements to reduce the statistical uncertainty.

Table 5.1: Complex gain parameters for three EFEMs in DAΦNE e^- ring

Mode	Gain	Phase
21	0.91 ± 0.04	-9.4 ± 2.7
89	0.84 ± 0.04	3.6 ± 2.6
101	1.0 ± 0.1	2.7 ± 6.3

Table 5.2: Complex gain parameters for three EFEMs in BESSY-II

Mode	Gain	Phase
281	1.04 ± 0.04	-5.4 ± 2.0
396	1.31 ± 0.02	-1.6 ± 0.8
397	1.23 ± 0.03	-2.1 ± 1.3

5.2.3 BESSY-II results

As described in Sec. 4.3.2 BESSY-II has three unstable eigenmodes: 281, 396, and 397. Feedback plant characterization results for these EFEMs are based on the same set of measurements presented in Sec. 4.3.2. Extracted loop gains and phases are presented in Table 5.2. Three eigenmodes show much less variation of the loop phase than the two modes measured in the ALS. Gain differences have not been investigated further since the measured 2 dB difference between modes 281 and 396 is much less than stability margin of the BESSY-II feedback system.

5.3 Parasitic effects

It is evident from the previous section that a physical feedback system deviates from an ideal model, sometimes in a significant manner. Here important mechanisms that can cause these deviations will be explored. Some of these effects are easier to analyze in the time domain while others are best visualized in the frequency domain. One mechanism that can cause non-ideal loop response is bunch-to-bunch coupling in the analog front and back-end channels. We will start from analyzing this parasitic effect in the time domain.

5.3.1 Bunch-to-bunch coupling

To analyze the coupling in the front-end processing consider the impulse response of that channel - closely recreated by its response to a single-bunch fill - as sampled by the ADC. The ideal front-end would produce a signal in only one ADC sample per turn leaving neighboring samples at their noise floor level. Finite rise and fall times, filter ringing, and reflections all lead to non-zero signals appearing in additional samples. Let's define signal levels in the filled sample and the following parasitic ones as $x_0, x_1, \ldots, x_K$. If the ideal front-end transfer function has only the first element non-zero, the coupled case can be considered as the ideal channel followed by the coupling transfer function. Coupling transfer function is that of an FIR filter with coefficients $1, x_1/x_0, \ldots, x_K/x_0$. Since the ADC is timed for maximum response in x_0 we have $x_m < x_0, m = 1, \ldots, K$. In fact, in a properly working channel $x_m \ll x_0$. Once the distortion transfer function is determined one can compute the response at $l\omega_0 + \omega_s$ to find the effect of the coupling on the EFEM l. Since significant coupling lasts only for a few RF periods and the FIR response is dominated by the first (unity) coefficient, the frequency response of this parasitic effect does not have sharp, narrowband features comparable to the width of the synchrotron response. Consequently, a single complex gain characterizes well the effect of front-end coupling on the feedback loop dynamics for a given eigenmode.

Analysis of parasitic coupling in the back-end channel is identical to the above description of the front-end coupling with one difference. In the back-end case the beam samples the kicker voltage rather than the ADC sampling the beam-derived signal. This sampling establishes the distortion FIR filter.

5.3.2 Shifts in synchronous phase

In a sampled system - especially a system with dual samplers - it is important to maintain timing alignment of the sampling clocks and the analog signals being sampled. For longitudinal feedback systems described here sophisticated timing procedures have been developed for both sampling points, that is the ADC in the front-end and the beam sampling the kicker voltage in the back-end [58]. Unfortunately there are uncontrollable effects that shift the timing during system operation. Let us consider the mechanisms that can cause these shifts and the effect of the latter on the feedback loop transfer function.

There are many causes for the synchronous timing shifts, some of which move all bunches identically (common-mode shifts) and others which produce differential bunch-to-bunch shifts. Poor RF cavity phase regulation, setpoint changes due to operator action, misphasing of multiple RF stations - all of these cause common-mode synchronous phase (and time) shifts. Differential shifts are caused by beam loading transients described in detail in Sec. 4.2.2. Common mode and differential shifts cause distinctly different problems in longitudinal feedback system operation. While common-mode shifts can be avoided by adhering to proper operating procedures and maintaining RF cavity phase regulation loops, the differential shifts can be reduced only via fill unevenness reduction, higher RF voltage, or driving impedance reduction. The above actions are usually infeasible due to other ring operating requirements as well as high costs.

Bunch timing shifts considered here are of the ±15 degrees@$f_{\rm rf}$ order. The reason for such a range lies in the acceptance range of the feedback front-end. Since the detection frequency is $6 \times f_{\rm rf}$ the range is ±15 degrees at the RF frequency. Larger timing shifts make the feedback system completely inoperable and are not considered here. For a storage ring with 500 MHz RF frequency these shifts correspond to ±83 ps timing shifts.

Front-end sensitivity

In case of the ADC we are concerned with the relative timing of the ADC sampling clock and the baseband phase detector output. ADC clock is derived from the ring master oscillator and is stably phase locked to that reference. The phase detector output is driven by the beam, thus changes in the synchronous timing appear directly as timing shifts of this signal. Two effects are caused by the beam timing shifts in the front-end. First, the phase of the $6 \times f_{\rm rf}$ burst changes relative to the carrier resulting in a DC offset and lower detector gain. Common mode shifts are eliminated in the channel via phase servo loop as described in Sec. 2.10.1. However differential shifts do modify the loop gain from bunch to bunch. Second, the baseband pulse is shifted relative to the ADC clock. Since the main pulse has flat top, timing shifts within $\approx T_{\rm rf}/2$ have little effect. However shifts in timing modify the way the reflections and filter ringing is sampled by the ADC (in the following buckets). This changes the front-end distortion FIR transfer function and, consequently, the phases and gains across the processing band. Normally in the front-end bunch-to-bunch coupling is very small, with the distortion transfer function close to unity. Timing shifts in this case can be neglected.

Back-end sensitivity

As mentioned earlier, in the back-end the beam samples the voltage in the kicker. The sampled waveform is at the QPSK carrier frequency of 9, 11, or 13 times $f_{\rm rf}/4$. Nominally beam is timed to sampled a peak of the kicker voltage. Shift in bunch arrival time slides the sampling point along the kick waveform. On one hand timing shifts have less effect in the back-end than in the front-end simply due to the ratio of the carriers used. For example, if the QPSK carrier is at $9f_{\rm rf}/4$, gain reduction due to the timing shift of ϕ degrees at RF frequency is $cos(2.25\phi)/cos(6\phi)$ times less in the back-end. On the other hand the effect is more significant due to several reasons. First, the back-end does not have phase servo loop to reject common-mode timing shifts. These shifts simultaneously affect all bunches and can produce significant loop transfer function changes. Second, in the back-end timing shifts cause not only gain reduction but also a loop phase change. Thus synchronous phase variations act to erode both gain and phase stability margins. Systems running at the highest loop gains suffer most from these variations. Note that feedback systems often can handle much smaller simultaneous changes in loop gain and phase than the changes of each parameter separately [59, page 1-23].

We will illustrate the effects of back-end timing shifts with measurements made at the Advanced Light Source. Longitudinal feedback system there uses a drift-tube kicker design [57]. Analytical response of such kicker is given by the following expression

$$\begin{aligned} V_{\mathrm{kick}} &= R_s I_{\mathrm{QPSK}} = R_s \frac{V_{\mathrm{QPSK}}}{50} \\ R_s &= 4Z_l \left\{ \frac{n \sin(2n\frac{\omega}{9f_{\mathrm{rf}}})}{2n\cos(\frac{\omega}{9f_{\mathrm{rf}}})} e^{i(\frac{\pi}{2}-(2n-1)\frac{\omega}{9f_{\mathrm{rf}}})} \right\}^2 \end{aligned}$$

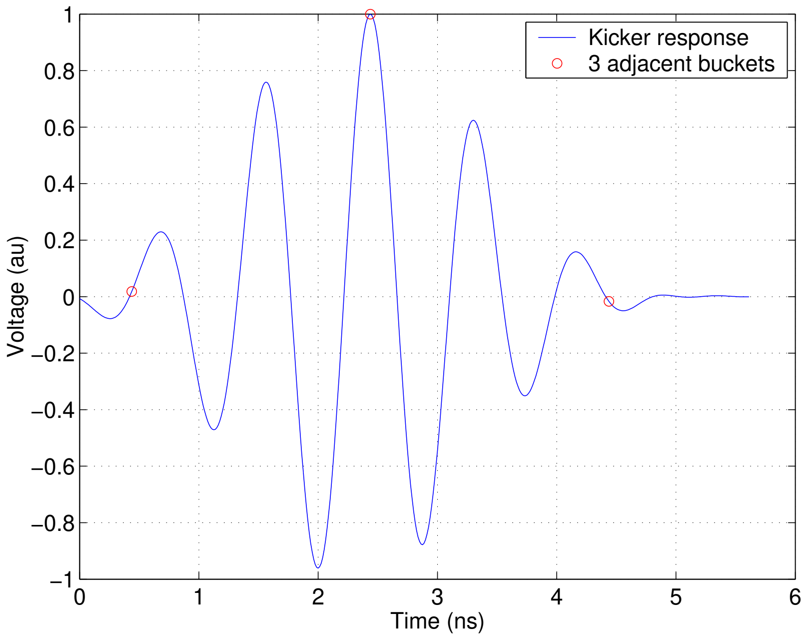

Figure 5.6: Simulated ALS kicker response for single-bunch excitation. The kick is timed to the bunch with ≈ 2.5 ns arrival time. Two adjacent buckets at ± 2 ns are shown as well.

where n is the number of drift-tubes per kicker and Z_l is the impedance of a drift-tube. Fig. 5.6 shows the analytically computed kicker voltage for a single-bunch (2 ns long) QPSK burst excitation. The bunch in the center of the waveform is timed perfectly - at the largest peak. QPSK modulation places adjacent bunches at zero crossings minimizing the coupling between bunches.

To check this model we measured response of the physical kicker at the ALS. The measurement was done using a single bunch in the ring to sample kicker voltage. Single bucket kick was generated by the feedback system. The kick was modulated at the synchrotron frequency to maximize coupling to the beam. Back-end delay line was swept over a 10 ns range with 25 ps steps. At each setting the amplitude of bunch oscillation was measured in the front-end using an FFT spectrum analyzer. This measurement is identical to the back-end timing procedure [58] with the exception of the large sweep range. As a result we get rectified kicker response to a QPSK burst. Using the knowledge of the oscillatory nature of the original waveform it is possible to "unwrap" the measurement. The results are shown in Fig. 5.7. Note the additional ringing in the experimental measurement, most likely introduced by a mismatch in the power distribution network between the power amplifier

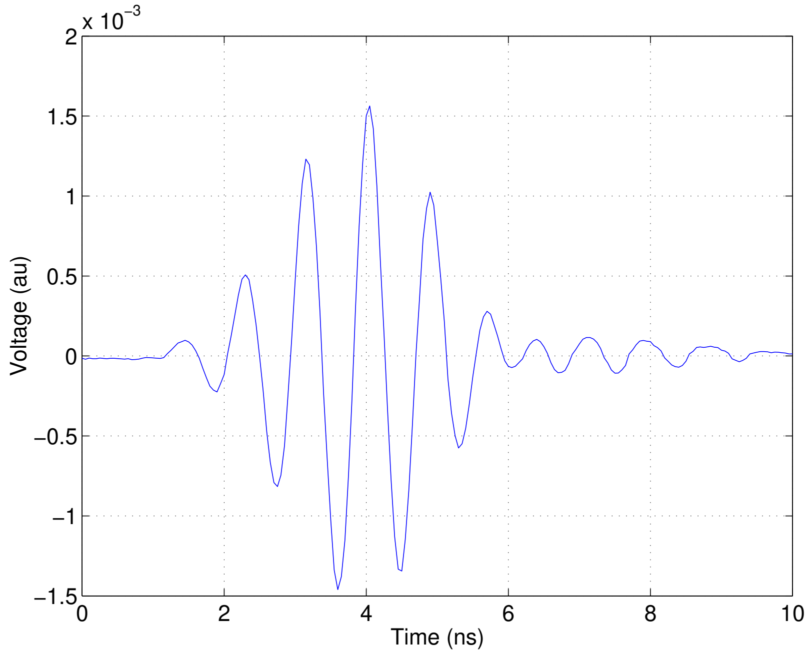

Figure 5.7: Measured ALS kicker response.

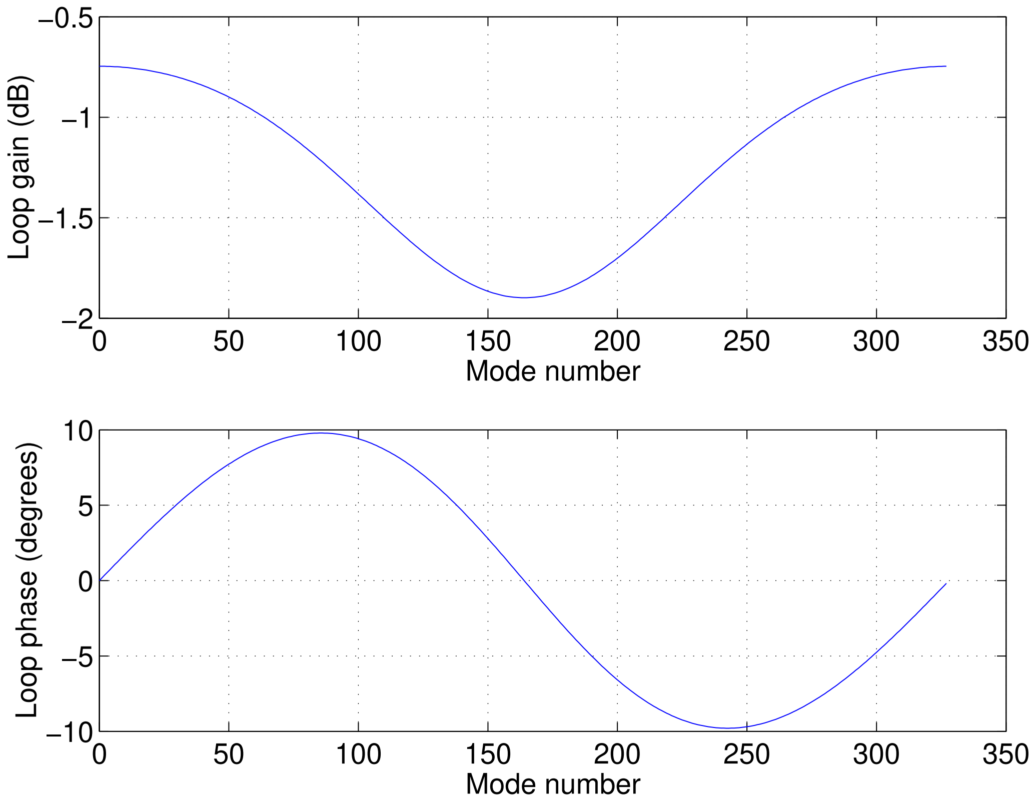

Figure 5.8: FIR distortion filter response when bunch timing advanced 75 ps. Zero dB gain corresponds to nominal back-end response.

and the kicker.

Now let us qualitatively examine what happens when bunch timing shifts from optimal point. The center bunch "slides" off the peak and samples a lower kick voltage. At the same time adjacent bunches move off zero crossings creating bunch-to-bunch coupling. By sampling the single-bucket kicker response at the appropriate points we get the coefficients for the 3-tap FIR distortion filter. In Fig. 5.8 the frequency response of such a filter is shown for the case when bunch timing is advanced 75 ps. Phase shift in the modal range of 200-250 is most significant with main RF cavity impedances driving modes 205 and 233 and with harmonic cavities exciting EFEMs 212-218. Changes in loop gain and phase are unacceptably large for the system running very close to stability margins. These issues will be further discussed in Ch. 6.

An experimental measurement of back-end timing shift effects was made at the ALS.

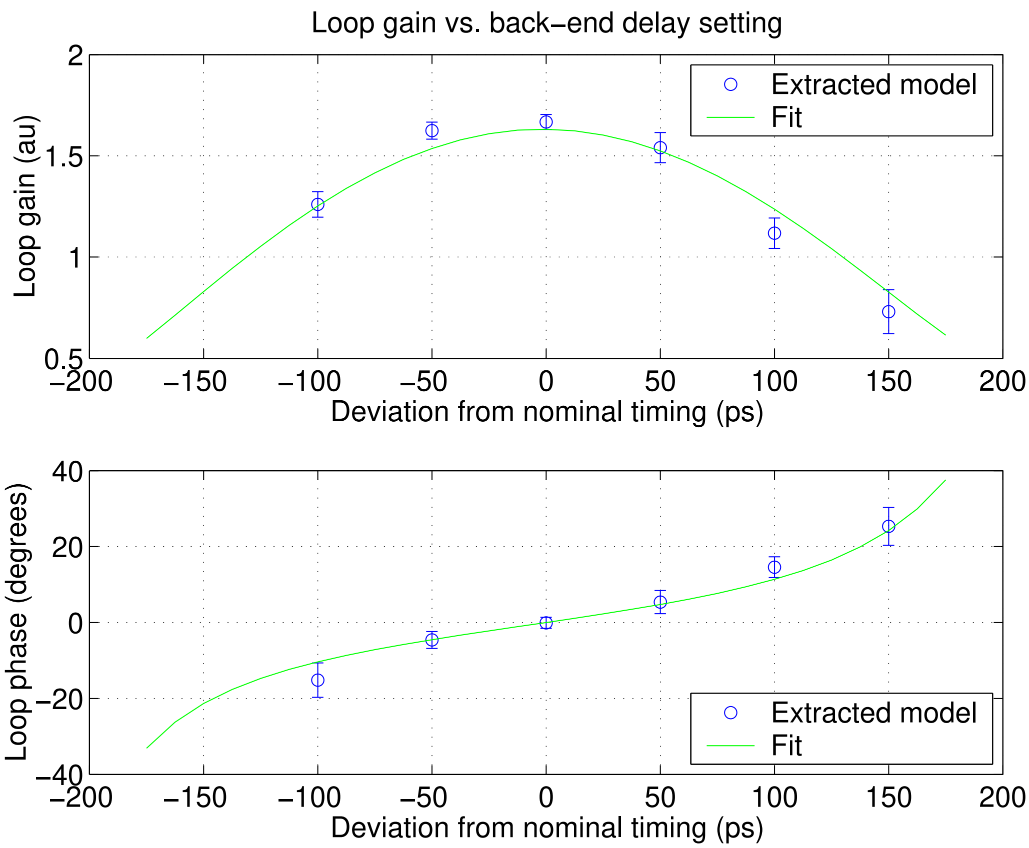

Figure 5.9: Feedback loop gain and phase dependence on back-end timing for EFEM 233

To simulate the timing shifts back-end delay line was offset from the nominal position. At each setting multiple grow/damp measurements were made at beam currents around 80 mA. From these measurements we extracted EFEM 233 loop gain and phase for that particular timing. The results are compared in Fig. 5.9 with gains and phases computed using 3-tap FIR distortion model. The FIR gains were scaled by 1.63 to match extracted values, similarly, FIR phases were shifted by -1.9 degrees. The measurements agree very well with the simulation - possibly agreement could be refined by using measured kicker QPSK response.

5.3.3 Effect of power stage transfer function

In building longitudinal feedback system a lot of care is taken to ensure proper broadband response of the processing channel. Output power amplifier is a challenging part of that

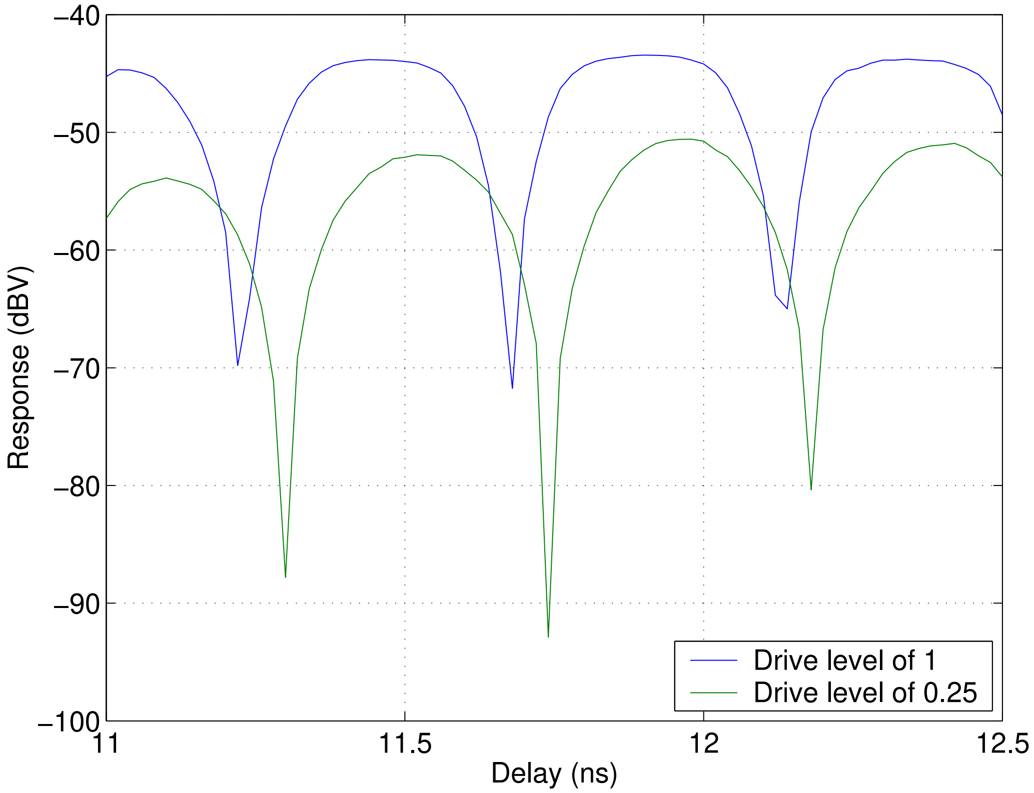

Figure 5.10: Measurement of the AM-to-PM conversion in the ALS TWT amplifier. Response for the 0.25 drive level is shifted relative to the full-scale response by 60 ps which corresponds to 24 degrees at the excitation frequency of 1125 MHz.

response simply because of high power required (200 to 500 W) and broad 250 MHz bandwidth. Standard transfer function measurements can be used to verify the bandwidth and phase linearity. Another important characteristic is the transient response of the device which is not necessarily measured by the classic swept sine network analyzer. For the transient response once can use AM modulated signals using network analyzer to verify correct carrier and sideband reproduction.

At the Advanced Light Source a 200 W traveling wave tube (TWT) amplifier had been used at the output of the feedback system (the amplifier has been replaced since by a solid-state GaAs device for various reasons). Such amplifiers have AM-to-PM conversion characteristics, that is changing input signal amplitude causes changes in the phase shift through the device. To characterize the extent of AM-to-PM conversion we performed back-end timing sweeps at different drive levels. In Fig. 5.10 two such sweeps are shown, at full- and quarter-scale drive levels. There is an obvious shift between the two curves quantified by comparing the two unwrapped sinusoids as 24 degrees at 1125 MHz. This is equivalent to a 60 ps timing shift and will affect the response of the feedback. Fortunately, since the system timing is performed at the full-scale drive level the back-end is correctly phased at maximum output. Phase shift with decreasing amplitude leads to higher steady-state kick level than would be achievable with a better amplifier.

Chapter 6

Longitudinal feedback controllers

6.1 Recursive digital filters

In this chapter we will discuss the design of feedback controllers intended to stabilize coupled-bunch instabilities. These controllers will be implemented as discrete-time systems. An important class of discrete-time processing structures consists of the structures described by linear constant-coefficient difference equations, e.g.

$$y[n] = a_0 x[n] + a_1 x[n-1] + a_2 x[n-2] \tag{6.1}$$

$$y[n] + b y[n-1] = a_0 x[n] + a_1 x[n-1] \tag{6.2}$$

where $x[n]$ is the input sequence and $y[n]$ is the output one. These difference equations can be represented graphically by block diagrams. Such diagrams for Eqs. 6.1-6.2 are shown in Fig. 6.1.

The block diagrams point to an important distinction between two classes of systems defined by difference equations. Equation 6.1 defines a filter without feedback, i.e. output signal depends only on a finite number of the input sequence terms. For the system defined by Eq. 6.2 $y[n]$ depends also on the past values of the output. It can be shown that in such a structure the output is a function of an infinite number of past inputs for $b \neq 0$.

In general all discrete filters can be separated into these two distinct classes, finite and infinite impulse response. Finite impulse response (FIR) filters are implemented by networks without feedback. As a result, the length of the impulse response of such a filter

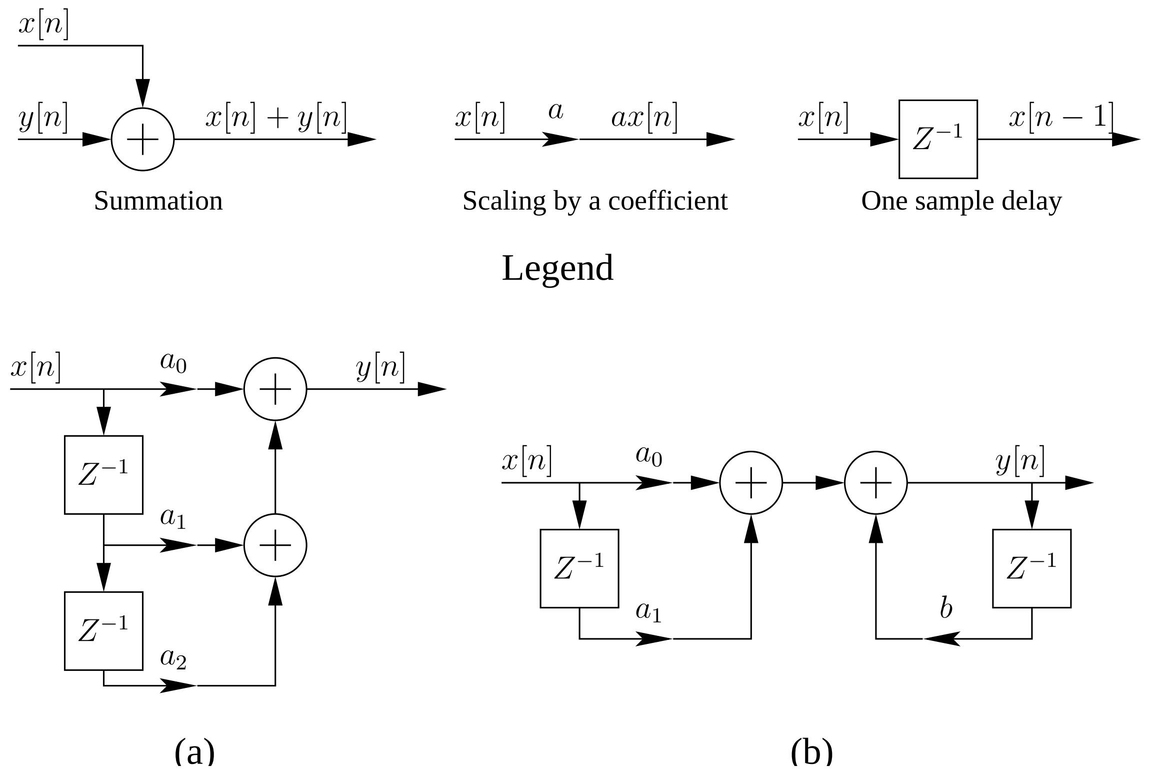

Figure 6.1: (a) Block diagram corresponding to Eq. 6.1; (b) Block diagram corresponding to Eq. 6.2

is determined by the number of the delay elements. When feedback terms are added to the filter network it is possible to obtain infinitely long impulse responses (IIR). This means that for the same number of memory elements a filter with feedback allows the implementation of more complex responses. However such flexibility comes at a price. As for any system with feedback, the designer of the IIR filters needs to worry about stability.

A convenient tool in the analysis of the FIR or IIR filters is the z-transform. It is defined for sequence $x[n]$ as

$$X(z) = \sum_{n=-\infty}^{\infty} x[n]z^{-n} \tag{6.3}$$

Taking z-transform of Eq. 6.2 we get

$$Y + bz^{-1}Y = a_0X + a_1z^{-1}X \tag{6.4}$$

where X and Y are the z-transforms of the corresponding sequences. Rearranging the terms in Eq. 6.4 we get the relationship between X and Y or the system function

$$H(z) = \frac{Y}{X} = \frac{a_0 + a_1z^{-1}}{1 + bz^{-1}}$$

In general, for an IIR or FIR system $H(z)$ is given by

$$H(z) = \frac{\sum_{k=0}^{N} a_k z^{-k}}{1 + \sum_{k=1}^{M} b_k z^{-1}} \tag{6.5}$$

For an FIR system M is equal to 0. Expressing the system function as a ratio of polynomials in z provides a simple way to check for the stability of an IIR system. Such a system is stable if all of the roots of the denominator polynomial (poles) are contained within the unit circle. Rewriting Eq. 6.5 in terms of its poles and zeros (roots of the numerator polynomial) we get

$$H(z) = g_0 \prod_{n=1}^{N} \frac{z - z_n}{z - p_n} \tag{6.6}$$

Now the stability requirement can be expressed as $|p_n| < 1$. There are no restrictions on the zeros of the system function.

If the z-transform is evaluated on the unit circle, i.e. for $z = e^{i\Omega}$ Eq. 6.3 becomes equivalent to the discrete-time infinite duration Fourier transform. Thus, to obtain the transfer function of a system one has to evaluate the system function on the unit circle:

$$H(\Omega) = H(z)|_{z=e^{i\Omega}}$$

6.2 Implementation-specific performance optimization

Ultimately all of the controller design methods result in a specification of the feedback filter system function. Performance of the feedback depends on the way that these system

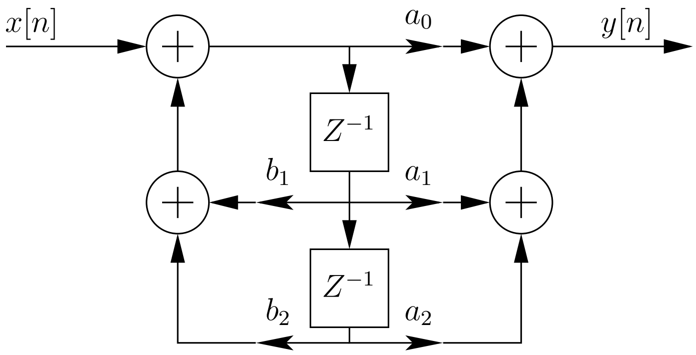

Figure 6.2: Direct form II realization of a second-order section

functions get implemented. For a given difference equation there are many possible implementation topologies differing in efficiency, quantization sensitivity, and noise performance. For the DSP-based system considered here IIR filtering was implemented using a cascade of second-order sections. A cascade of second-order sections has been selected for its insensitivity to coefficient quantization and computational efficiency [60, pages 300–344]. Due to memory and execution time constraints the maximum number of sections was set to 6. Thus, the search space for the controller design methods includes at most 12 poles and 12 zeros. Each second-order section (SOS) was implemented in the direct form II. Topology of such implementation is illustrated in Fig. 6.2. All computation is done using 16 bits for both data and coefficients. Coefficients are represented in the Q14 format that is with two bits to the left of the binary point. In such representation coefficients can span from -2 to $2-2^{-14}$. The reason for such a choice becomes clear if we examine the relationship between the coefficients and roots of a second-order polynomial.

For a denominator polynomial defined by $z^2 + b_1 z + b_2$ and its two roots p_1 and p_2 we have:

$$(z-p_1)(z-p_2) = z^2 - (p_1+p_2)z + p_1 p_2 \Rightarrow \begin{cases} b_1 = -(p_1+p_2) \\ b_2 = p_1 p_2 \end{cases} \tag{6.7}$$

Since for a stable IIR filter poles must be within the unit circle we get

$$
\begin{aligned}
|b_1| &= |p_1 + p_2| \leq |p_1| + |p_2| < 2 \\
|b_2| &= |p_1||p_2| < 1
\end{aligned}
$$

Therefore poles of any stable filter can be implemented in direct form II with coefficients smaller than 2.

Input and output samples are 8 bits wide while the IIR computation is done with 16 bits of resolution. Alignment of the input data within the 16 bit word is driven by a trade-off of two parameters: dynamic range and roundoff noise. If we leave empty bits in the upper part of the 16 bit word, a filter can support very large gains before saturating. However roundoff noise is introduced at the level of the LSB of the input. Leaving free bits in the lower part of the word cuts down the roundoff noise, but also reduces the dynamic range. In our implementation we chose place the 8 bit input word in bits 4 through 11 leaving 4 bits below and above. In this case roundoff noise is kept at $\frac{1}{16}$ of the input LSB and the node values within the SOS cascade can be 16 times the full-scale input level.

Starting from specification of a controller by 12 poles and 12 zeros we first need to group poles into 6 pairs. There is no ambiguity in pairing complex poles since they occur in conjugate pairs. No constraints are placed on the pairings of the real poles. Similarly complex zeros are paired unambiguously. Since real zeros can occur outside the unit circle, certain combinations of zeros can lead to large coefficients a_1 and a_2. These are represented in the Q14 format so the large values will have to be scaled down to the $[-2, 2-2^{-14}]$ range reducing the gain of the section. Since the overall gain has to be maintained at the design value, reduced gain of one section requires increased gains elsewhere, possibly amplifying the quantization noise. From Eq. 6.7 we see that we need to concentrate on reducing the product of the paired zeros. Thus we sort all real zeros by absolute values and sequentially pair largest ones with the smallest. Such pairing minimizes the range of values of the coefficients a_1 and a_2 for all sections with real zeros.

In implementing a filter we will need to define a distance in z-plane between pole and zero pairs. When both pairs are complex we will define the distance between a pole and a zero in the upper half-plane. When poles, zeros or both are real, we will use minimum distance between the two sets:

$$D = \min_{mn} |p_m - z_n|, \ m = 1, 2; n = 1, 2 \tag{6.8}$$

Our distance definition for complex pairs is equivalent to that given in Eq. 6.8 for the real pairs. Thus we will use Eq. 6.8 for all distance computations.

The next step in defining the implementation is to combine pole and zero pairs into second order sections. We follow a variant of the rule presented in [61] that suggests sequentially pairing the largest magnitude pole with the zero closest to it in the z-plane. Complex poles have equal magnitudes and for real poles we use the larger absolute value.

We modified the rule by selecting first a zero pair closest to zero frequency, i.e. to the point $z = 1$. That pair is combined with the minimum magnitude pole pair. The need for this modification will become obvious if we consider typical signal levels for the feedback system in operation, for example at the ALS. Feedback controllers there are normally configured with the in-band gains of 25-50 so that an input oscillation amplitude of 2-4 counts produces a full scale 8 bit output correction signal. At the same time the amplitude of the synchronous phase transient which appears as DC position of each bunch is around 100 counts. Since we desire to provide high gain for small signals riding on the large DC offsets it is important to remove the DC before any amplification takes place. Notice that in the direct form II pole response is computed first, followed by the zeros. If a pair of poles close to DC is chosen for the first section and these poles provide DC gain of more than 16 we risk saturating the processing before DC rejection takes place. Choosing a small magnitude pole pair minimizes this risk. This is still an imperfect solution since there are pole placements that can lead to saturation. One could escape DC saturation by selecting the pole pair most distant from $z = 1$. However such pole pair could be close to unit circle and provide high gain at some other frequency. Since in the first section DC is not yet removed such a strategy is likely to lead to saturation. Another way to avoid these risks is to sacrifice one SOS to DC rejection placing both poles at the origin. Then one section later on would need to have both zeros at the origin. This approach is wasteful since by eliminating a section one can reduce processing time with a possible improvement in control performance. Yet another way to alleviate DC saturation problems is to use transposed direct form II structure for the second-order section. Block diagram of such an implementation is shown in Fig. 6.3. Notice that computation of zeros is done first, solving the abovementioned DC

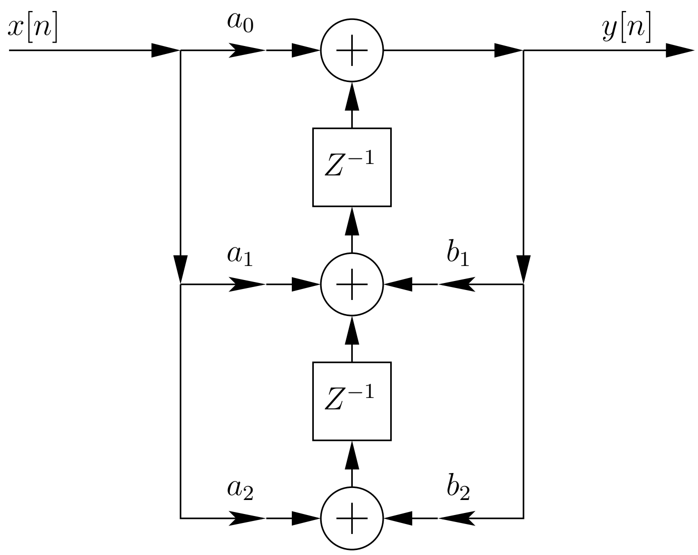

Figure 6.3: Transposed direct form II realization of a second-order section

saturation issues.

Another DC rejection issue stems from numerical truncation. In our IIR implementation result of the final accumulation is truncated by retaining upper 16 bits of the 32 bit word. Thus, a small negative DC term gets truncated to the LSB of the 16 bit word. If DC poles are placed after DC zeros that error can be amplified. Let's consider an example with real poles and zeros:

$$H(z) = \frac{(z-0.999)(z-0.999)}{(z-0.98)(z-0.98)}$$

At DC we expect to have $H(1) = 0.0025$. Let's consider the layout where the zeros precede the poles and the input signal is constant at -100 counts. After quantizing the coefficients we expect the output of the zero section to be 0.00006 before the truncation. After the 32-to-16 bit conversion we get -0.0625. Passing that through the pole section we get output of -287 counts instead of expected -0.25. Limiting the filter output to 8 bits results in DC saturation. If the poles and zeros are swapped saturation happens during pole computation. Even though there is no universal solution to this problem, controllers encountered in practice rarely need to have poles so close to DC. In addition one can confine the poles to a circle of radius smaller than unity limiting the gain of a single pole.

Once the first SOS is configured, the rest is paired according to the rule quoted above . Next we define the order in which these sections occur. Two possible choices are presented in [61]: sorting by increasing or decreasing pole magnitudes. Sorting by increasing magnitudes provides better overall saturation behavior while the other choice leads to lower roundoff noise. Since in our implementation roundoff noise is much smaller than the LSB of the input word we chose to sort the poles by increasing magnitudes.

6.3 Frequency-domain controller design

In this approach to feedback controller design we start from specifying a desired frequency-domain complex transfer function. Then a recursive digital filter is designed to approximate such a transfer function. Filter design problem has been extensively studied and there many approaches to solving it [62, 63, 64, 65, 66]. Traditionally these methods concentrate on approximating the magnitude of the desired transfer function. In cases where the phase response is considered the goal is obtaining a linear phase characteristic or, equivalently, a constant delay through the filter. The magnitude of the resulting delay is usually not important as long as it is constant. In our case the task is more demanding since in general for the wideband control filters we desire constant phase shift within the control bandwidth. Since there are delays in the loop associated with sampling as well as transport delays, we need to provide a positive slope of phase versus frequency, i.e. negative delay. Obviously negative delay cannot be a characteristic of a causal system. However it is possible to approximate phase behavior of a negative delay within a limited bandwidth.

6.3.1 Problem formulation

Let $H_d(\Omega)$ be the desired complex transfer function where Ω is the digital frequency of a sampled system. For the sampling interval T_s digital frequency is related to physical one as $\Omega = \omega T_s$. Transfer functions are defined on the set $(\mathcal{U} : 0 \leq \Omega \leq \pi)$. Let vector A consist of n parameters defining an infinite impulse response (IIR) filter and $H(A, \Omega)$ be the transfer function of that filter.

In order to determine how well $H(A, \Omega)$ approximates $H_d(\Omega)$ we need to define a distance function. A common choice is a weighted L_p norm defined for two real-valued functions $f_1(x)$ and $f_2(x)$ on a set $\mathcal{X}$ as:

$$\|L\|_p = \left\{ \int_{\mathcal{X}} W(x)|f_1(x) - f_2(x)|^p dx \right\}^{1/p} \tag{6.9}$$

where $W(x)$ is a positive weighting function. Since in our case functions are complex-valued, Eq. 6.9 has to modified. An obvious choice is to consider $|H(A,\Omega) - H_d(\Omega)|$. However such a distance function has a single weighting function. In controller design we need to be able to weight independently magnitude and phase. For example, within control bandwidth one needs both magnitude and phase of $H(A,\Omega)$ to approximate the desired response. Outside the control band, only magnitude response needs to be optimized. These considerations lead to a choice of two distance functions: weighted norms of magnitude and phase errors.

$$\|L^M(A)\|_p = \left\{ \int_0^{\pi} W_M(\Omega)||H(A,\Omega)| - |H_d(\Omega)||^p d\Omega \right\}^{1/p} \tag{6.10}$$

$$\|L^\phi(A)\|_q = \left\{ \int_0^{\pi} W_\phi(\Omega)|\angle H(A,\Omega) - \angle H_d(\Omega)|^q d\Omega \right\}^{1/q} \tag{6.11}$$

If we define the overall distance function as a linear combination of $\|L^M(A)\|_p$ and $\|L^\phi(A)\|_q$ our approximation problem is reduced to minimization of a real-valued function of n real parameters. In order to avoid numerical integration we will approximate the norms in Eqs. 6.10-6.11 by summations:

$$\begin{aligned} L^M_{2p}(A) &= \sum_{k=1}^{K} W_M(\Omega_k)(|H(A,\Omega_k)| - |H_d(\Omega_k)|)^{2p} \\ L^\phi_{2q}(A) &= \sum_{k=1}^{K} W_\phi(\Omega_k)(\angle H(A,\Omega_k) - \angle H_d(\Omega_k))^{2q} \end{aligned}$$

Here we take advantage of a fact that a global minimum of $\|L(A)\|_{2p}$ is also a global minimum of $\|L(A)\|^{2p}_{2p}$. Note that the above formulation does not require the points Ω_k to be evenly spaced. One important requirement on the maximum spacing of the frequency points is that significant features of both transfer functions do not fall between the samples. For example, if there is a sharp resonance or notch between two frequency samples, it will not affect the distance function indicating that frequency sampling is too coarse.

Overall distance function can be expressed as

$$L(A) = \beta L_{2p}^{M}(A) + (1-\beta)L_{2q}^{\phi}(A) \tag{6.12}$$

where parameter $\beta \in [0, 1]$ controls relative weights of the phase and magnitude errors.

Now let us consider the parameterization of the controller system function. We desire a stable controller, thus it is important to be able to easily check for stability. Since a system is stable if all of its poles are within the unit circle, parameterization of Eq. 6.6 seems to be the obvious choice. However in physical implementation filter coefficients must be real, so complex poles or zeros must occur in conjugate pairs. This requirement naturally leads to selecting a second-order section as a unit of cascade implementation. Then Eq. 6.5 can be written as

$$H(z) = g_0 \prod_{n=1}^{N_s} \frac{z^2 + a_{1n}z + a_{2n}}{z^2 + b_{1n}z + b_{2n}} \tag{6.13}$$

In this case parameter vector A is given by

$$A = (a_{11}, a_{12}, b_{11}, b_{12}, \ldots, a_{1k}, a_{2k}, b_{1k}, b_{2k}, \ldots, g_0) \tag{6.14}$$

As stated above we need a way to check for stability. Due to coefficient quantization poles just inside unit circle can become unstable in the physical system. In addition, a pole near the unit circle causes a gain peak within a section. To avoid these problems we will require all filter poles to lie within the disk of radius ρ, $\rho < 1$. For each second-order section this requirement translates into three linear inequality constraints:

$$\begin{aligned} b_2 &\le \rho^2 \\ \rho b_1 - b_2 &\le \rho^2 \\ -\rho b_1 - b_2 &\le \rho^2 \end{aligned}$$

For efficient numerical minimization it is important to be able to compute analytically the gradient of the distance function. Partial derivatives of the distance function in Eq. 6.12 with respect to parameters $(A_1, A_2, \ldots, A_{2N_s+1})$ are given by

$$\begin{aligned}
\frac{\partial L(A)}{\partial A_m} &= \beta \frac{\partial L_{2p}^M(A)}{\partial A_m} + (1-\beta)\frac{\partial L_{2q}^{\phi}(A)}{\partial A_m} \\
\frac{\partial L_{2p}^M(A)}{\partial A_m} &= 2p \sum_{k=1}^{K} W_M(\Omega_k)(|H(A,\Omega_k)| - |H_d(\Omega_k)|)^{2p-1} \frac{\partial |H(A,\Omega_k)|}{\partial A_m} \\
\frac{\partial L_{2q}^{\phi}(A)}{\partial A_m} &= 2q \sum_{k=1}^{K} W_M(\Omega_k)(\angle H(A,\Omega_k) - \angle H_d(\Omega_k))^{2q-1} \frac{\partial \angle H(A,\Omega_k)}{\partial A_m}
\end{aligned}$$

For the parameterization of the filter system function given in Eq. 6.13 we have

$$\begin{aligned}
|H(A,\Omega)| &= |H(e^{i\Omega})| = |g_0| \prod_{n=1}^{N_s} \left| \frac{e^{i\Omega} + a_{1n} + a_{2n}e^{-i\Omega}}{e^{i\Omega} + b_{1n} + b_{2n}e^{-i\Omega}} \right| = |g_0| \prod_{n=1}^{N_s} \left\{ \frac{P(a_{1n}, a_{2n}, \Omega)}{P(b_{1n}, b_{2n}, \Omega)} \right\}^{\frac{1}{2}} \\
\angle H(A,\Omega) &= \angle H(e^{i\Omega}) = \sum_{n=1}^{N_s} \arctan \frac{(1-a_{2n})\sin\Omega}{(1+a_{2n})\cos\Omega + a_{1n}} - \sum_{n=1}^{N_s} \arctan \frac{(1-b_{2n})\sin\Omega}{(1+b_{2n})\cos\Omega + b_{1n}}
\end{aligned}$$

where

$$P(a_1, a_2, \Omega) = 1 + a_1^2 + a_2^2 + 2a_2 \cos 2\Omega + 2a_1(1+a_2)\cos\Omega$$

Then partial derivatives of the magnitude and the phase of the transfer function at a frequency Ω are given by

$$\begin{aligned}
\frac{\partial |H(e^{i\Omega})|}{\partial a_{1m}} &= |H(e^{i\Omega})| \frac{2a_{1m} + 2(1+a_{2m})\cos\Omega}{2P(a_{1m}, a_{2m}, \Omega)} \\
\frac{\partial |H(e^{i\Omega})|}{\partial a_{2m}} &= |H(e^{i\Omega})| \frac{2a_{2k} + 2\cos 2\Omega + 2a_{1k}\cos\Omega}{2P(a_{1m}, a_{2m}, \Omega)} \\
\frac{\partial |H(e^{i\Omega})|}{\partial b_{1m}} &= -|H(e^{i\Omega})| \frac{2b_{1m} + 2(1+b_{2m})\cos\Omega}{2P(b_{1m}, b_{2m}, \Omega)} \\
\frac{\partial |H(e^{i\Omega})|}{\partial b_{2m}} &= -|H(e^{i\Omega})| \frac{2b_{2k} + 2\cos 2\Omega + 2b_{1k}\cos\Omega}{2P(b_{1m}, b_{2m}, \Omega)} \\
\frac{\partial \angle H(e^{i\Omega})}{\partial a_{1m}} &= \frac{(a_{2m}-1)\sin\Omega}{P(a_{1m}, a_{2m}, \Omega)}
\end{aligned}$$

$$\frac{\partial \angle H(e^{i\Omega})}{\partial a_{2m}} = -\frac{\sin 2\Omega + a_{1m}\sin\Omega}{P(a_{1m}, a_{2m}, \Omega)}$$
$$\frac{\partial \angle H(e^{i\Omega})}{\partial b_{1m}} = -\frac{(b_{2m}-1)\sin\Omega}{P(b_{1m}, b_{2m}, \Omega)}$$
$$\frac{\partial \angle H(e^{i\Omega})}{\partial b_{2m}} = \frac{\sin 2\Omega + b_{1m}\sin\Omega}{P(b_{1m}, b_{2m}, \Omega)}$$

The approximation problem can then be set up as a nonlinear minimization with linear inequality constraints and solved using Sequential Quadratic Programming [42].

6.3.2 Specifying desired response

One of the most significant difficulties of the controller design lies in specifying desired transfer function. There are several objectives that have to be met by $H_d(\Omega)$. First of all, it needs to stabilize longitudinal dynamics for all possible accelerator operating parameters. Second, as described in Section 6.2, $H_d(\Omega)$ must provide DC rejection. Third, broadband noise as well as power line harmonics present in the input signal have to be suppressed.

Model developed in Chapter 4 describes beam dynamics as an unstable harmonic oscillator. A controller of such a plant can be implemented as a differentiator [24, 67]. Advantages of such a controller include insensitivity to resonant frequency and DC rejection. Consider the beam model as follows:

$$G(s) = \frac{g_0}{s^2 - 2\lambda s + \omega_n^2}$$

With controller $H(s) = Ks$ closed-loop characteristic equation becomes

$$C(s) = s^2 + (g_0 K - 2\lambda)s + \omega_n^2$$

Thus differentiator feedback acts to stabilize the system independent from ω_n^2. However gain of such a controller rises linearly with frequency leading to actuator saturation. Transfer function of differentiator has constant 90 degree phase shift at all frequencies, i.e. zero delay. Discrete-time controller cannot replicate such response due to signal processing delays as well as delay due to the zero-order hold.

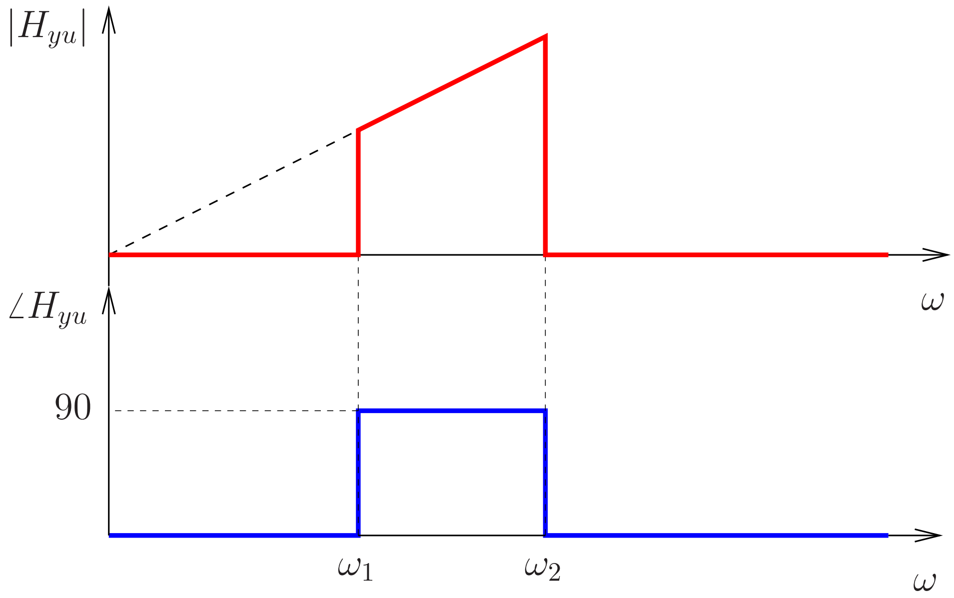

Figure 6.4: Magnitude and phase of a transfer function of a band-limited differentiator

In specifying desired response we chose to approximate the differentiator transfer function in a band-limited fashion. Let us consider feedback loop model developed in Chapter 4 as illustrated in Fig. 5.1. Our desired transfer function $H_d(\Omega)$ should result in band-limited differentiator dynamics between $y(t)$ and $u(t)$. If longitudinal oscillation frequency is expected to be in the range ω_{d1} to ω_{d2} we define control band to extend by ω_{guard} below and above. Desired transfer function from $y(t)$ to $u(t)$ is shown in Fig. 6.4. Here $\omega_1 = \omega_{d1} - \omega_{\text{guard}}$ and $\omega_2 = \omega_{d2} + \omega_{\text{guard}}$. Magnitude of the transfer function is linear with frequency within the control band and zero elsewhere. Phase response is 90 degrees in band. Phase response outside the control band is undefined since transfer function magnitude is zero.

Once desired continuous-time transfer function $H_{yu}(i\omega)$ has been defined we can compute the discrete-time transfer function $H_d(\Omega)$. In order to obtain that function we need to divide the continuous time function by the transfer functions of the zero-order hold and the delay. These are given by

$$
\begin{aligned}
H_{\text{zoh}}(\Omega) &= e^{-i\Omega/2}\text{sinc}(\frac{\Omega}{2\pi}) \\
H_{\text{delay}}(\Omega) &= e^{-i\Omega T_d/T_s}
\end{aligned}
$$

Magnitude of the resulting transfer function is amplified at the high frequencies to compensate for the zero-order hold. In-band phase response is then a negative group delay of $-(T_s/2 + T_d)$.

Useful insight into the optimization process can be gained by factoring controller system function into minimum-phase and all-pass components. As described in Section 6.1 transfer function is stable and causal if all of the poles lie within the unit circle. Transfer function $H(z)$ is minimum-phase if $1/H(z)$ is stable and causal. In other words, both poles and zeros of minimum-phase transfer function lie within the unit circle. System is called all-pass if the magnitude of its transfer function is unity at all frequencies. All-pass systems have equal numbers of poles and zeros related as $p_k = 1/z_k^*$. Thus a stable and causal all-pass system has all of its zeros outside the unit circle. Consequently one can represent any stable and causal system function $H(z)$ as a product of minimum-phase and all-pass functions:

$$H(z) = H_{\min}(z) * H_{\mathrm{ap}}(z)$$

here $H_{\min}(z)$ contains the poles and zeros of $H(z)$ that lie inside the unit circle and zeros that are conjugate reciprocals of the zeros of $H(z)$ located outside the unit circle. Then $H_{\mathrm{ap}}(z)$ contains the zeros of the $H(z)$ that lie outside the unit circle. In the product poles of $H_{\mathrm{ap}}(z)$ cancel conjugate reciprocal zeros of $H_{\min}(z)$ resulting in the original transfer function. Group delay of $H(\Omega)$ can be expressed as

$$grd[H(e^{i\Omega})] = grd[H_{\min}(e^{i\Omega})] + grd[H_{\mathrm{ap}}(e^{i\Omega})]$$

A general property of an all-pass system is that $grd[H_{\mathrm{ap}}(e^{i\Omega})] > 0$ at all frequencies [60, page 238]. Thus if the goal of the optimization was in minimizing group delay, resulting transfer function would be minimum phase. However since we require certain phase shift within the control band optimum filters can be non-minimum-phase. One can separate phase response error into two parts: DC (mean error) and AC. Larger group delay will increase the AC part of the error. However if the resulting filter has zero mean phase deviation from $H_d(\Omega)$, overall error can be smaller.

In the next three subsections three controller designs will be presented. These designs were created for three different accelerators: ALS, DAΦNE, and BESSY-II. In each of these machines longitudinal coupled-bunch dynamics varies as a function of several parameters such as beam current, beam energy, main and third-harmonic RF system configurations. There variations are different in each case requiring unique controller designs.

6.3.3 ALS

SLAC designed DSP-based longitudinal feedback systems have been commissioned at ALS, BESSY-II, DAΦNE, PEP-II, PLS using finite impulse response (FIR) filters. The filters were sufficient to control longitudinal instabilities in all machines listed above. The ALS system was used successfully with an FIR filter for a period of 5 years.

In June 1999 a third harmonic RF system consisting of 5 passive cavities was installed at the ALS. Following the installation the original FIR feedback filter had difficulties maintaining longitudinal control over the operating range of currents and energies. We made a series of instability studies to identify the effects that caused the loss of control. From these studies we determined several problem areas. Let us consider them one by one starting from the most significant.

Synchrotron frequency shifts

The most important effect we observed is the change of synchrotron frequency with beam current and energy. The third harmonic RF system at the ALS is passive, that is, it is powered by the beam. Since the harmonic voltage V_h in Eq. 2.17 scales linearly with beam current, the slope of the RF voltage at the synchronous phase changes as a function of current. The slope decreases monotonically as the beam current I_0 increases. According to Eq. 2.5 synchrotron frequency is proportional to the square root of that slope. Over a range of values of the beam current from zero to full design current of 400 mA synchrotron frequency changes significantly.

Measurements of the mode 233 oscillation frequencies as a function of current are shown in Fig. 6.5 for two conditions: without and with third harmonic cavities. Data without third harmonic cavities was taken on the 28th of May of 1998. Using a linear fit we estimate the frequency shift over the 400 mA range as 413 Hz. In the second case, with third harmonic cavities the shift in frequency is much larger over the same range of currents reaching 3322 Hz.

In the ALS one has to take into account the shift in the longitudinal oscillation frequency with beam energy in addition to the synchrotron frequency shift due to the third harmonic

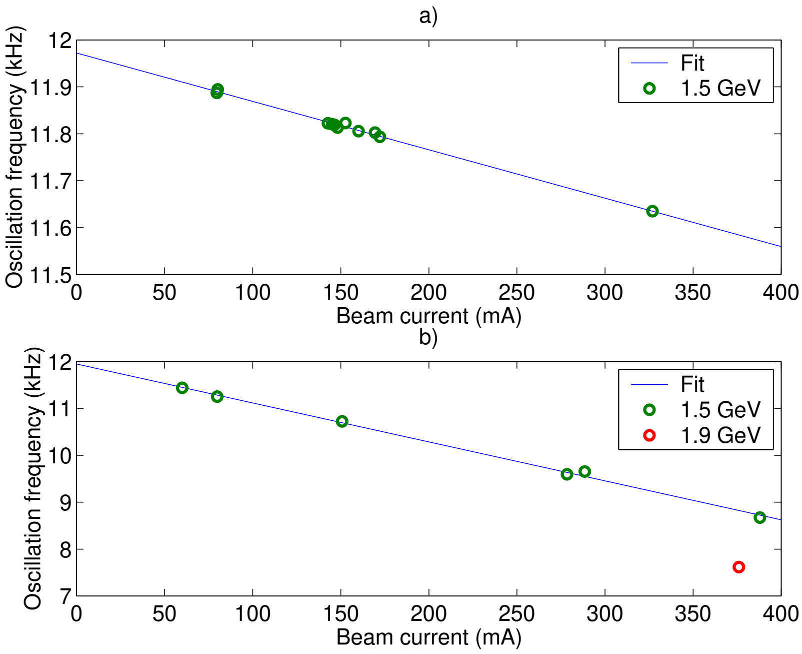

Figure 6.5: Measured oscillation frequencies for the EFEM 233 in the ALS. a) Measurements taken on 5/28/1998 (before third harmonic cavities were installed) at the 1.5 GeV nominal energy. Frequency shift over the 400 mA range is 413 Hz. b) Measurements taken on 3/27/2000 with 2 third harmonic cavities tuned in and 3 cavities tuned away from $3f_{\mathrm{rf}}$ (parked). At 1.5 GeV frequency shift is 3322 Hz between 0 and 400 mA. When ring energy is ramped to 1.9 GeV oscillation frequency drops further to 7.6 kHz.

cavities. The Advanced Light Source is often operated with energy ramping from the injection energy of 1.5 GeV to 1.9 GeV. As shown in Sec. 2.5.2 changes in ring energy lead to significant changes in longitudinal dynamics. From Eq. 2.15 we see that energy ramping from E_0 to E_1 lowers the modal oscillation frequency by a factor of $\sqrt{E_1/E_0}$. As shown in Fig. 6.5 oscillation frequency drops from 9 kHz to 7.6 kHz after ramping.

The effect of these frequency shifts on the longitudinal feedback is best understood if we consider the transfer function of the ALS FIR feedback controller illustrated in Fig. 6.6. Let us examine the change in the transfer function between 12 kHz and 7.6 kHz. There is little difference in the gain with the lower frequency being 0.4 dB higher. However phase changes from -90 degrees at 12 kHz to -23 degrees at 7.6 kHz. For resistive damping controller should have phase shift of -90 degrees - at 7.6 kHz FIR controller is mostly reactive. Ignoring reactive tune shifts and resulting changes in dynamics we estimate that at 7.6 kHz

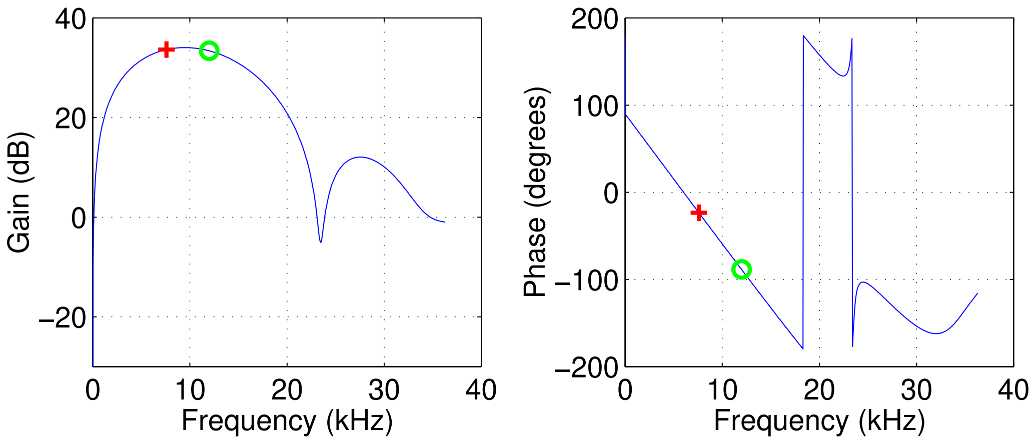

Figure 6.6: Transfer function of the ALS 6 tap FIR filter. The left plot shows the magnitude of the transfer function and the right - the phase. The nominal synchrotron frequency of 12 kHz is marked by the green circles while the shifted oscillation frequency of 7.6 kHz is marked by the red crosses.

feedback gain drops by $\cos(90-23) \approx 0.4 \approx -8$ dB. Even this optimistic estimate of the gain change is sufficient to make the system unstable since feedback is normally configured for a damping margin of 2 when closed-loop damping rate is equal to the open-loop growth rate. In that configuration gain drop of 6 dB will result in the closed-loop damping rate of 0 with further gain reduction leading to instability. For a more rigorous test we can compute the closed-loop damping rate as a function of modal eigenvalue. Using the measurements of the eigenvalues versus beam current and energy we construct the open-loop eigenvalue curves on the complex plane with the beam current as an independent variable. Then, closed-loop dominant pole for each open-loop eigenvalue is computed using the system model described in Ch. 5. In Fig. 6.7 both open-loop and closed-loop poles are shown for the 1.5 GeV and 1.9 GeV beam energies. Note that the real part of the closed-loop pole approaches zero at 400 mA and 1.9 GeV bringing the system close to instability. At that point small variations in the growth rates or feedback gain can lead to loss of beam control.

Mode 0 tune shifts

An additional complication in operating bunch-by-bunch longitudinal feedback at the ALS comes from the large tune shift experienced by the even-fill eigenmode 0 - the lowest-frequency mode of motion in which all bunches move in phase. This eigenmode is strongly

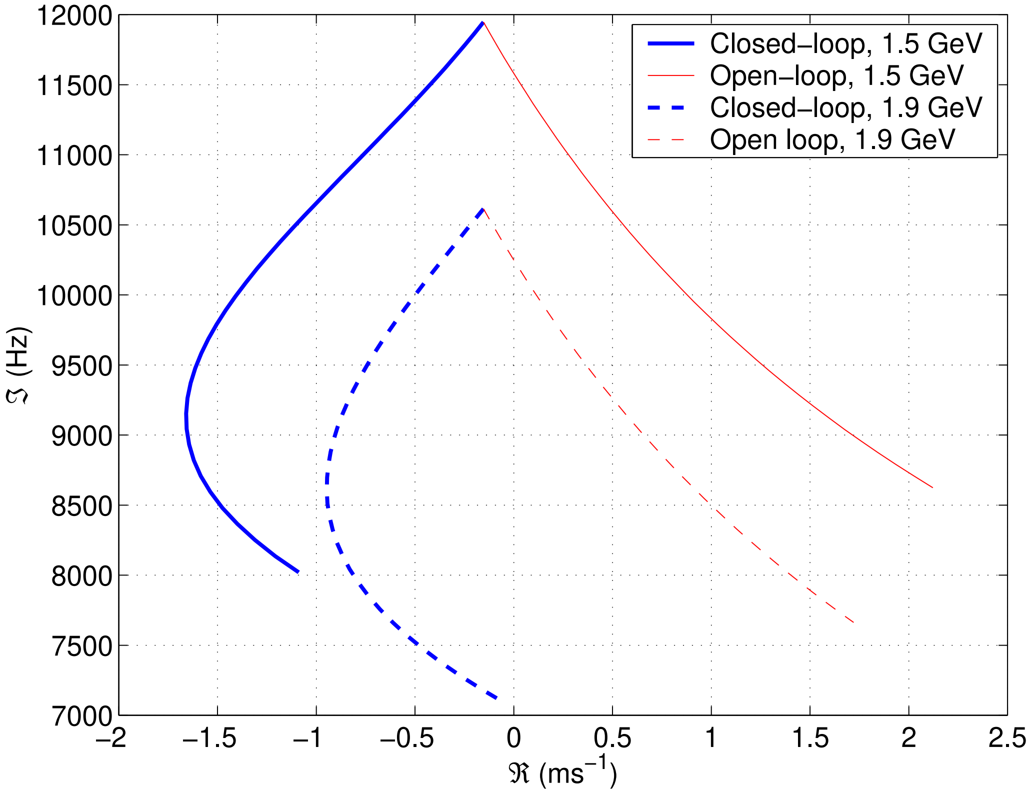

Figure 6.7: Open and closed loop dominant poles for beam currents ranging from 0 to 400 mA. The points where open-loop and closed-loop curves meet are at zero current (resulting in zero feedback gain). From those points beam current increases monotonically along each curve. The closed-loop poles were computed using 6 tap FIR filter.

affected by the fundamental impedances of the main RF cavities. Normally center frequencies of the RF cavities are tuned below $\omega_{\rm rf}$. Such tuning guarantees stability of mode 0 - so called Robinson stability [68]. Tuning of the RF cavities is a function of beam current as it is used to compensate for the beam loading [68]. Let's consider the effect of cavity detuning on the eigenvalue of mode 0. According to Eq. 2.9 the eigenvalue shift is given by

$$\begin{aligned}\lambda_0 &= \frac{\alpha e f_{\rm rf} I_0}{2E_0 h\omega_s}[(\omega_{\rm rf}+\omega_s)Z^{\|}(\omega_{\rm rf}+\omega_s)-(\omega_{\rm rf}-\omega_s)Z^{\|*}(\omega_{\rm rf}-\omega_s)] \\ &\approx \frac{\pi\alpha e f_{\rm rf}^2 I_0}{E_0 h\omega_s}[Z^{\|}(\omega_{\rm rf}+\omega_s)-Z^{\|*}(\omega_{\rm rf}-\omega_s)] \qquad (6.15)\end{aligned}$$

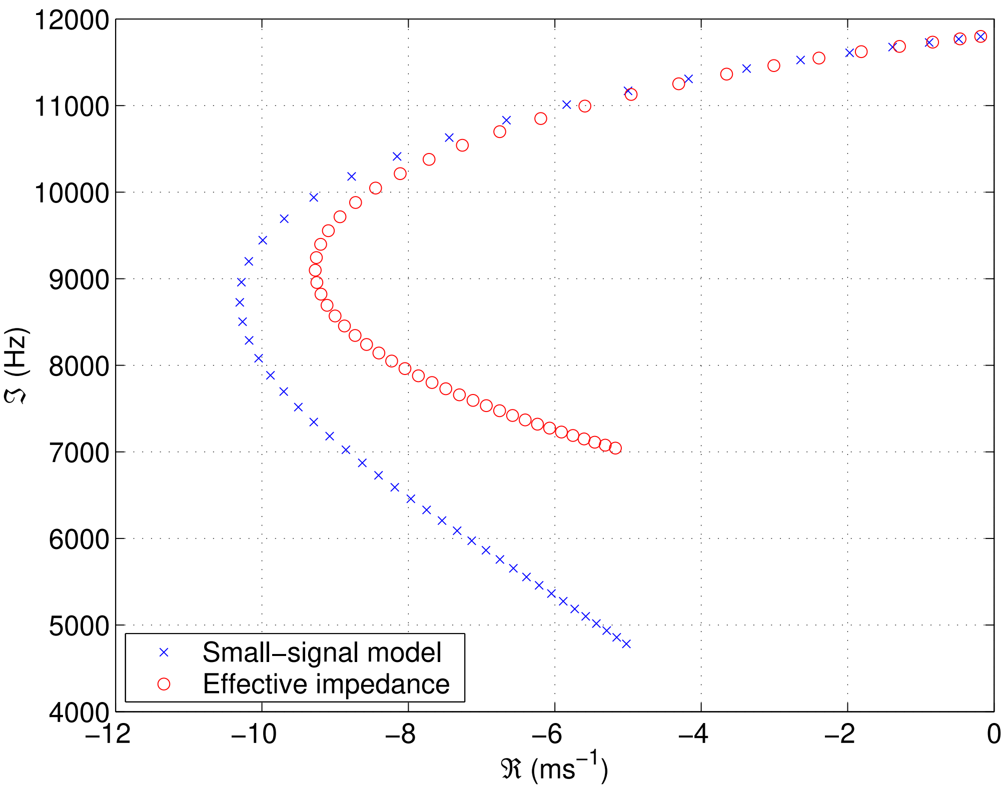

Figure 6.8: Mode 0 open-loop eigenvalues in ALS vs. beam current from 0 to 400 mA in 8.14 mA steps.

In practice the eigenvalue shift is larger than predicted by the above equation due to the large magnitude of the shift. In Fig. 6.8 mode 0 eigenvalues computed using Eq. 6.15 and the small-signal Pedersen model introduced in 2.4 are plotted vs. beam current. Note that large RF cavity detuning has two effects: reduction of the damping rate $\frac{1}{\tau_0} = \Re(\lambda_0)$ and increase of the tune shift $\Im(\lambda_0)$. As the ring is filled with charge mode 0 tune is shifted to lower frequencies. It is known from the operational experience at the ALS as well as from simulations that oscillation frequency of mode 0 shifts from 12 kHz (nominal synchrotron frequency) to 4 kHz at 400 mA [69]. If we examine the transfer function of the FIR controller shown in Fig. 6.6 we see that at 4 kHz the phase shift is 30 degrees or 120 degrees away from the optimal resistive damping phase. Therefore the FIR controller provides positive rather than negative feedback for mode 0. The effect has been noted before the installation of the third harmonic cavities as seen in the larger than expected level of phase noise at high currents [69].

When the third harmonic cavities were added several effects combined to make the stability of mode 0 a significant limitation to the operation of the machine. First, as described in the previous section, third harmonic cavities lower the frequencies of all modes. As mode 0 shifts below 4 kHz positive feedback effect of the FIR controller is enhanced. For example, from 4 kHz to 3 kHz phase response moves from 30 degrees to 45 while the gain drops by 2 dB. At the phase of 30 degrees anti-damping gain is scaled by $\cos(90-30) = 0.5$ which at 45 degrees it is $\cos(90-45) = 0.707$ or 3 dB higher. Thus overall gain is increased by 1 dB by shifting the frequency from 4 kHz to 3 kHz. Second, large frequency shifts of the other EFEMs necessitated increase in the overall controller gain - thereby increasing positive feedback applied to mode 0. Third, third harmonic cavities are tuned above $3\omega_{\mathrm{rf}}$ to obtain bunch lengthening. Such tuning has two consequences: reduced damping and slight positive tune shift of mode 0. This effect has not been quantitatively analyzed to determine whether mode 0 stability is improved or impaired by these counteracting influences. However the combination of the three effects described above has been sufficient to drive mode 0 to instability.

Synchronous phase transients and feedback gain limitations

Due to user requirements ALS operates with a large gap in the fill pattern. As described in Ch. 4 uneven fill pattern excites the impedances at the revolution harmonics. Resulting wake fields shift the equilibrium positions of the bunches producing a synchronous phase transient. Such a transient strongly affects the operation of the longitudinal feedback system via several mechanisms. As described in Ch. 2 front-end phase detection is performed using a double-balanced mixer with the carrier frequency of $6f_{\mathrm{rf}}$. The detected signal (at the input of the ADC) for bunch k is $G_{\mathrm{fe}} i_k \sin(6\phi_k)$. This expression assumes that the phase of the carrier is adjusted in such a way that the signal of a bunch with its centroid at synchronous position is perfectly in phase with the carrier. For small oscillations of the bunch centroid detected signal is

$$x_k \approx 6 G_{\mathrm{fe}} i_k \phi_k \tag{6.16}$$

The situation changes if there are differences in the synchronous position from bunch to bunch. Then the carrier cannot be adjusted to be in phase with the signals of different bunches, at best we can aim for zero average phase error. Consequently some bunches will be in phase which the carrier and others will be offset. If we denote equilibrium phase of bunch k relative to the average of the synchronous phase around the ring as $\bar{\phi}_k$, then we can write for the ADC input

$$x_k = G_{\mathrm{fe}} i_k \sin(6\phi_k) = G_{\mathrm{fe}} i_k \sin(6(\hat{\phi}_k + \bar{\phi}_k)) \tag{6.17}$$

where $\hat{\phi}_k$ is the motion of bunch k about the synchronous phase. For small oscillation amplitudes we can rewrite Eq. 6.17 as

$$x_k \approx G_{\mathrm{fe}} i_k \sin(6\bar{\phi}_k) + 6G_{\mathrm{fe}} i_k \cos(6\bar{\phi}_k)\hat{\phi}_k \tag{6.18}$$

Consider the above equation for the bunch signal in phase with the carrier. Then $\hat{\phi}_k = 0$ and Eq. 6.18 reduces to Eq. 6.16. Non-zero synchronous position $\bar{\phi}_k$ has two effects. First of all, the detection gain is reduced by $\cos(6\bar{\phi}_k)$. Second, ADC input signal contains a DC offset of $G_{\mathrm{fe}} i_k \sin(6\bar{\phi}_k)$.

The gain reduction due to the synchronous phase transient is a function of beam current since the wake voltage is excited by the beam current. At low beam currents both the transient and the gain reduction are small. However as the current increases the magnitude of the synchronous phase transient goes up and the feedback gain is reduced for some of the bunches. Figure 6.9 shows the synchronous phase transient in the ALS measured using the LFB facilities as well as the front-end gain multiplier due to the transient. Gain for the bunches at the extremes of the transient is reduced by 30 to 40%. Remember that the feedback system uses gain scaling linearly with beam current to combat the linear increase in the growth rates. Synchronous phase transient results in the feedback gain scaling less than linearly. Thus at some current feedback gain will become insufficient to control the coupled-bunch growth rates. Fortunately, for the ALS this deviation from linearity is only 5% at 400 mA for the phase transient similar to the one shown in Fig. 6.9. However tuning in more third harmonic cavities (only two out of five are actively tuned in in the above data set) will increase the magnitude of the synchronous phase transient as well as accelerate the nonlinear gain loss. In addition to the gain reduction in the front-end there is gain lowering in the back end since part of the bunches arrives in the kicker out of phase with the kick

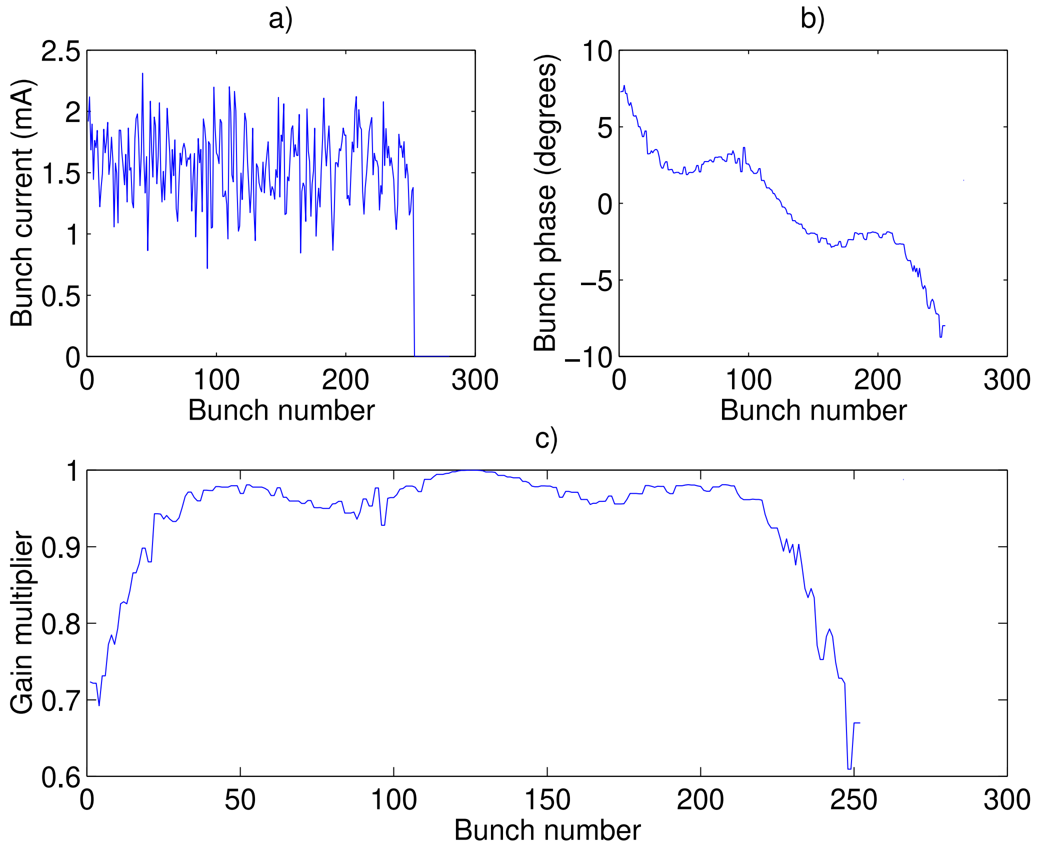

Figure 6.9: Effect of the synchronous phase transient on the front-end gain. This measurement was taken at the ALS at 1.5 GeV, 397 mA. a) Bunch-by-bunch currents showing 252 filled RF buckets. b) Equilibrium phases of the bunches relative to the phase detector carrier in degrees at RF frequency. c) Front-end gain reduction resulting from the non-zero equilibrium phases. Bunches with offsets near zero in plot b) have gain multipliers near 1 (little reduction) while the bunches at the extremes of the transient lose 30 to 40% of the front-end gain.

waveform. However back-end operates at the much lower frequency ($9/4f_{\rm rf}$ vs. $6f_{\rm rf}$), so the gain reduction is a smaller effect. At the same time, as shown in Ch. 5, mis-timed back-end kick on the bunches with non-zero equilibrium phase produces some feedback loop phase shift for those bunches further reducing feedback effectiveness.

The second effect introduced by the synchronous phase transients is the DC offset of the bunch signal at the ADC input. This offset depends on the equilibrium phase and current of the bunch. Offsets over the whole ring behave similarly to the synchronous phase transient as illustrated in Fig. 6.9. These offsets limit the allowable feedback front-end gain G_{fe}.

If the amplitude of the synchronous phase transient is small, the gain is limited by the noise level in the front-end circuitry. Empirical rule for setting G_{fe} has been to make RMS noise of the digitized signal equal one ADC count. Such a gain for the advanced light source was $G_{\mathrm{fe}} = 202$ counts/mA. At that gain setting with nominal per bunch current of 1.5 mA 4 degrees of motion would saturate the ADC input. This saturation is not a concern in the operation since digital signal processing section of the LFB normally operates with gain much larger than unity, so its output saturates long before the input. Now consider a case when the synchronous phase transient is large, on the order on 20 degrees peak-to-peak. Then some bunches have $+10$ degree offsets and some - -10 degrees. At the high front-end gain of 202 counts/mA ADC input signal for these bunches and for all other bunches with offsets more than 4 degrees will be saturated all the time. To operate the feedback system with such a transient the gain has to be lowered. If we define i_{nom} as the nominal per bunch current, $\bar{\phi}_{\mathrm{max}}$ as the amplitude of the transient and $\hat{\phi}_{\mathrm{min}}$ as the minimum oscillation amplitude we would like to detect without ADC saturation, we get for the front-end gain

$$G_{\mathrm{fe}} = \frac{128}{i_{\mathrm{nom}} \sin(6(\bar{\phi}_{\mathrm{max}} + \hat{\phi}_{\mathrm{min}}))} \tag{6.19}$$

For the ALS with $i_{\mathrm{nom}} = 1.5$ mA, $\bar{\phi}_{\mathrm{max}} = 10$ degrees, $\hat{\phi}_{\mathrm{min}} = 2$ degrees we get $G_{\mathrm{fe}} =$ 90 counts/mA - a reduction in gain of 7 dB.

This significant reduction in the available front-end gain has to be compensated by the gain increase elsewhere in the processing chain if the damping is to be maintained at the same level. There are two places where one can realistically achieve this gain increase: digital signal processing section and the output power stage. Theoretically, one could also modify the feedback kickers for higher shunt impedance or increase their number. However that path is less desirable due to the expense and the prolonged machine down time required to install these vacuum components. As for the output power stage, wideband power amplifiers in the 1-1.125 GHz range are quite expensive and only increase the loop gain by the square root of total power increase. Thus we are faced with the desire to increase the digital gain. Avoiding saturation of the output signal by the input and ADC quantization noise is quite challenging due to the combination of large gain and wide bandwidth required of the controller.

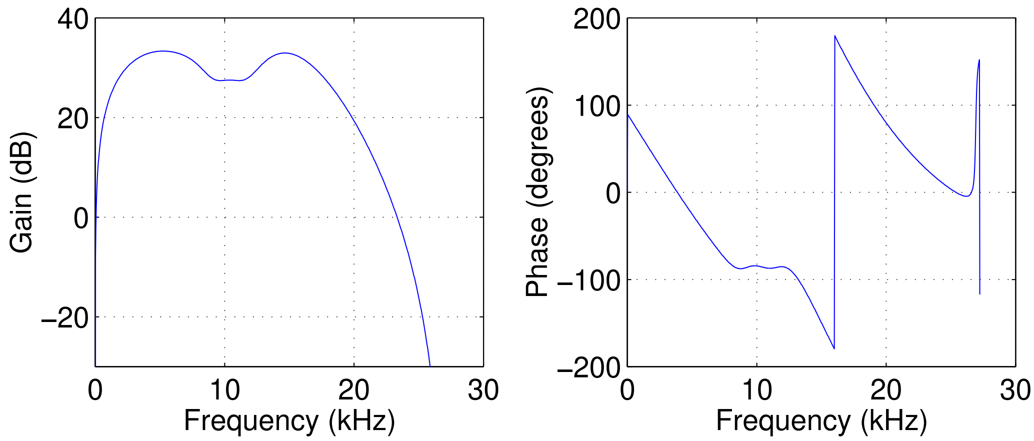

Figure 6.10: Transfer function of the ALS IIR controller designed using frequency-domain method

Increases in growth rates with reduction in synchrotron frequency

Another side effect of the reduction in the synchrotron frequency is the increase of the growth rates. If we examine Eq. 2.9 it is clear that the growth rates are inversely proportional to ω_s. Since synchrotron frequency decreases linearly with beam current (Fig. 6.5) growth rates increase faster than linearly. That nonlinearity accounts for the slight curvature in the open-loop poles plot in Fig. 6.7.

Controller design

We will describe here the feedback controller designed to longitudinally stabilize the ALS over the full range of beam currents and energies. Design was carried out using the frequency-domain optimization procedure. The transfer function of the controller is shown in Fig. 6.10. In selecting the weighting vectors for magnitude and phase emphasis was placed on achieving the phase response. Notice that the controller is wideband relative to the span of modal oscillation frequencies of 7-12 kHz. Maintaining the transfer phase near -90 degrees over a wide range of frequencies provides necessary damping for the combinations of modal growth rates and oscillation frequencies encountered at the ALS. In addition this design, named `landau_41`, has phase shift of -3 degrees at 4 kHz and, therefore, has a much smaller positive feedback effect on mode 0 than the FIR controller. The filter is implemented in a cascade of 6 second-order sections for a total of 12 poles and 12 zeros.

Its pole-zero map is illustrated in Fig. 6.11. Absence of poles near the unit circle results in high saturation margins throughout the cascade.

This controller design has been shown to provide ample damping of the unstable HOMs. Plot of the dominant closed-loop poles of the beam-feedback system obtained from the model is shown in Fig. 6.12. Note that this design provides large stability margins (fast damping) both at 1.5 and 1.9 GeV.

The `landau_41` controller has been used in the normal operation of the longitudinal feedback at the ALS from January 2000. A grow/damp measurement illustrating the performance of this controller is presented in Fig. 6.13. In this measurement three eigenmodes grow to the measurable amplitudes: 213, 217, and 233. Modes 213 and 217 are driven by the HOMs in the third harmonic cavities. One of these two modes is driven by the HOMs in the two beam powered harmonic cavities, while the other mode is due to the three parked cavities. Note the open-loop modal oscillation frequencies are down to 7.5 kHz for these modes - result of synchrotron frequency decrease with both beam loading and energy ramping. Damping rate provided by the feedback is slightly lower than that expected from Fig. 6.12. The discrepancy is caused by the gain reduction effects due to the synchronous phase transients which are not modeled in computing the closed-loop poles.

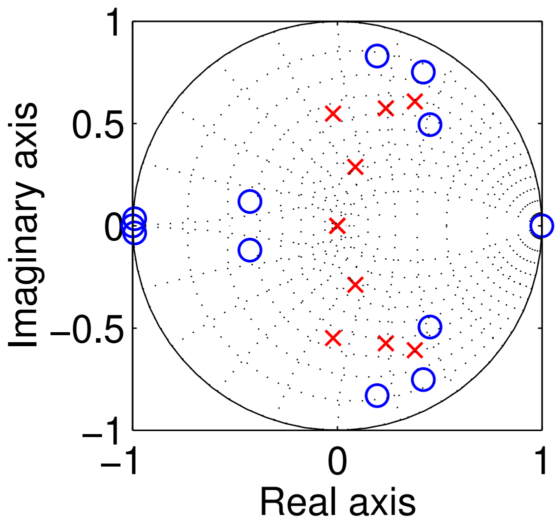

Figure 6.11: Poles and zeros of the ALS controller

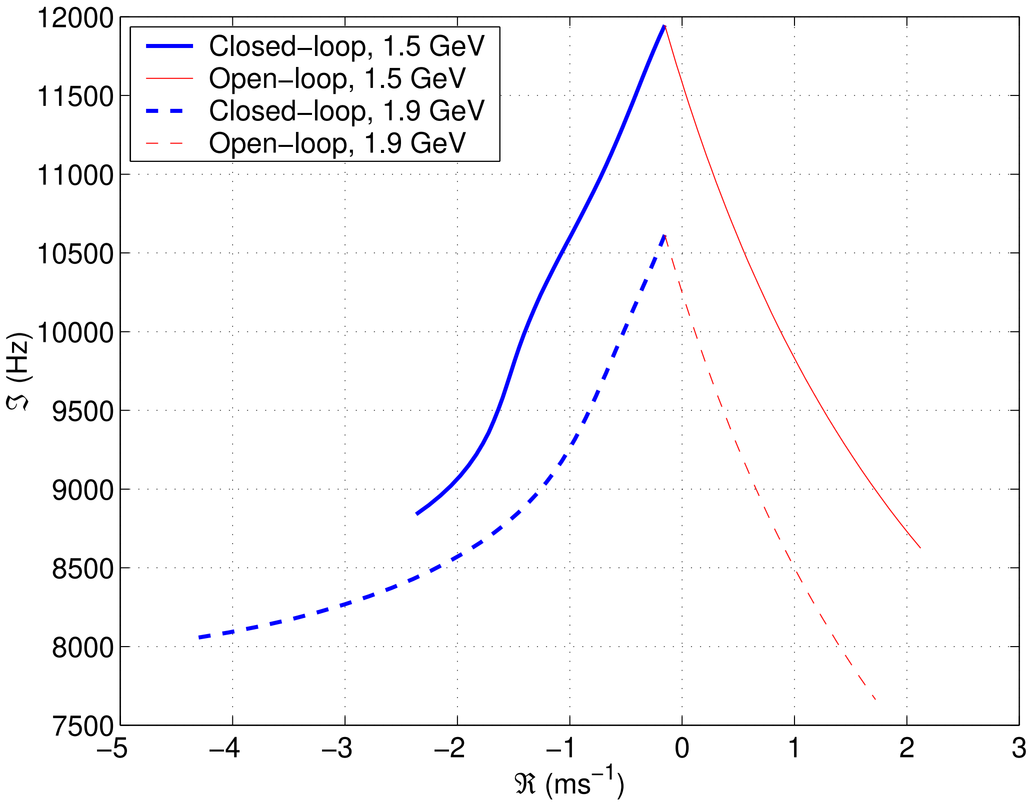

Figure 6.12: Open-loop and dominant closed-loop poles for the ALS IIR controller

The unfortunate feature of the `landau_41` controller is large gain at low frequencies (1-4 kHz). Since controller phase more than 90 degrees away from the resistive feedback setting this results in certain amount of positive feedback for mode 0. As loop gain increases one approaches instability of that eigenmode while lowering the gain leads to loss of stabilization of other eigenmodes. Thus, the system has a limited region of loop gains resulting in stability. Experiments at the ALS show that the range of acceptable gains varies between 6 and 8 dB.

In general, any practical feedback controller will have a limited range of allowable gains. The lower limit is dictated by the stability of the closed-loop system while the upper is either a stability or a noise saturation limit. That is, as the gain is increased either the system becomes unstable or the feedback output saturates with sensor or quantization noise. The size of the gain window is an important factor in the operation of the accelerator. Larger gain window makes it possible to retain feedback control independent from the variations in

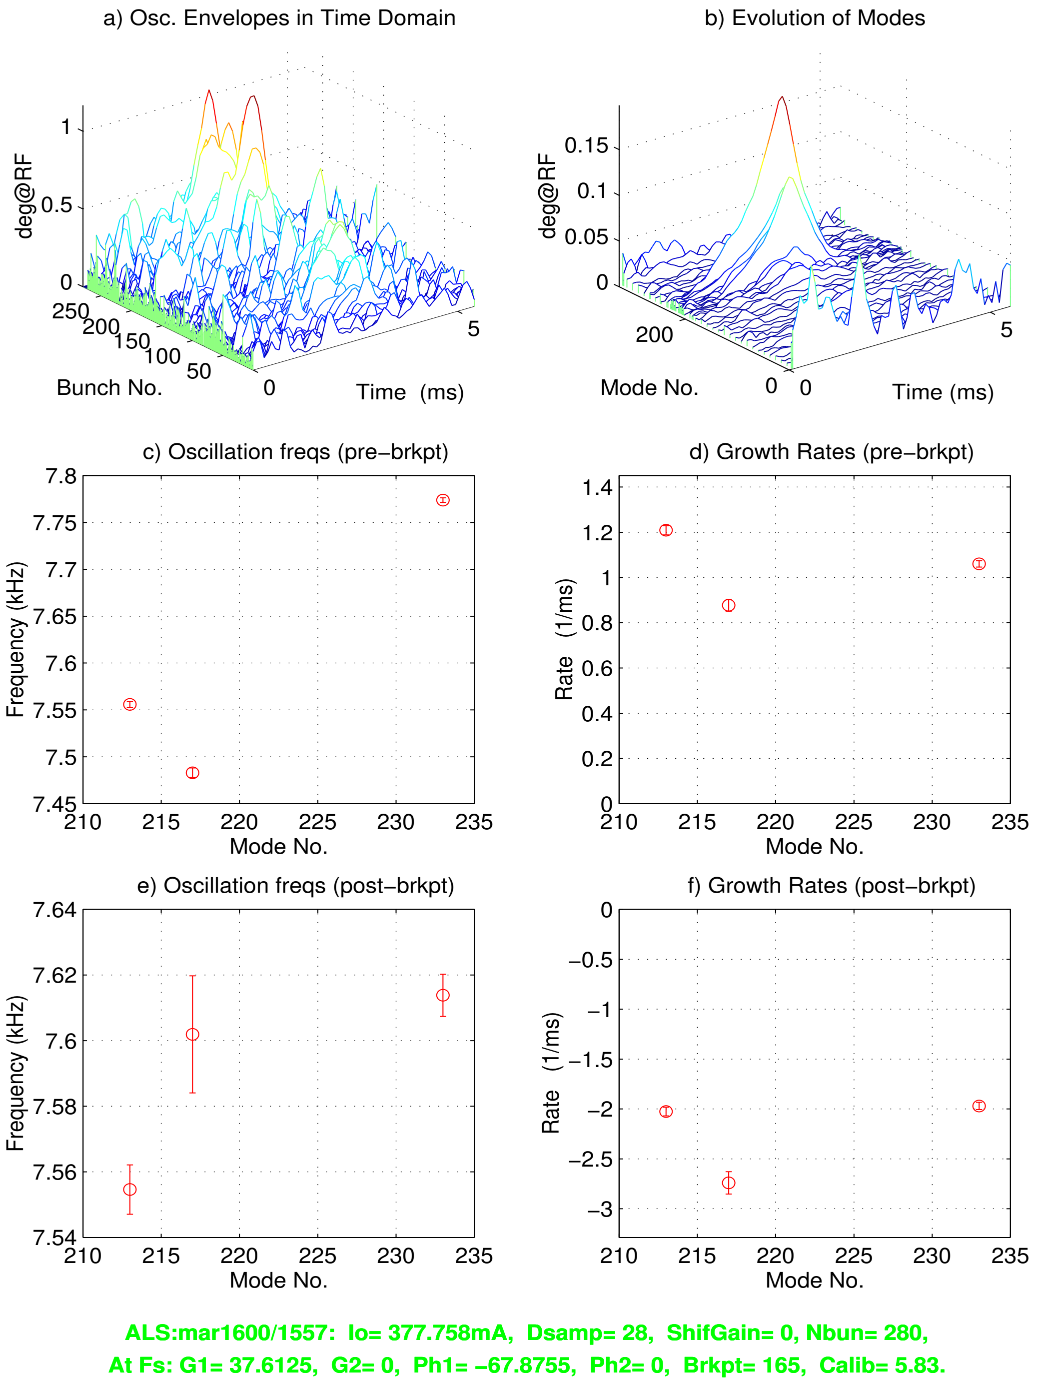

Figure 6.13: A grow/damp transient measurement at the ALS, 377 mA, 1.9 GeV, `landau_41` controller.

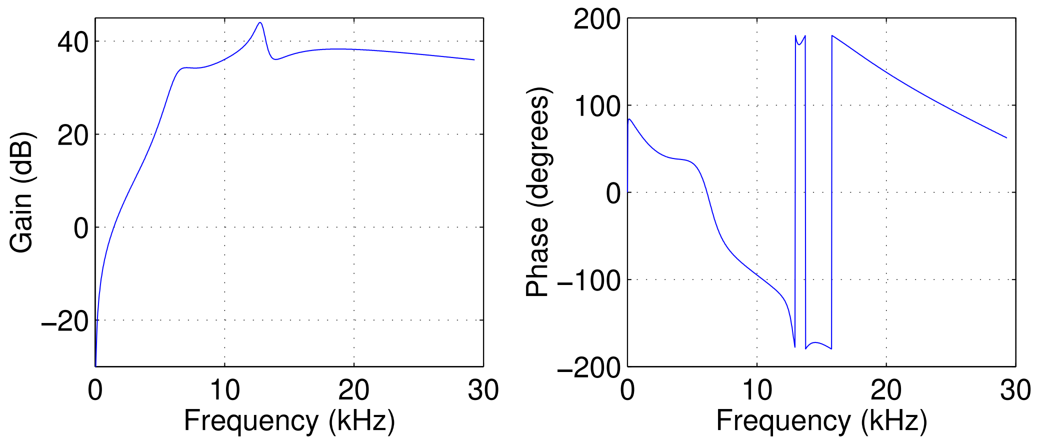

Figure 6.14: Transfer function of the `alsFD3` IIR controller designed to avoid excitation of mode 0

the growth rates of unstable modes as well as changes in the feedback gain. In this respect, controller `landau_41` has a relatively small gain window which adversely affects robustness of the longitudinal feedback at the ALS. The upper limit of the window is defined by stability of mode 0 and we can try to improve the robustness by changing the response of the controller to damp mode 0. Remembering that without feedback mode 0 is stable another approach becomes feasible - rolling off the gain of the controller so that eigenvalue of mode 0 is only weakly affected.

The latter approach was used to design another feedback controller for the ALS, design named `alsFD3` - referring to a frequency-domain (FD) design with 3 second-order sections. The frequency response of the controller is shown in Fig. 6.14. Note that the gain of the controller falls rather quickly below the lower control frequency of 7 kHz. Down at 4 kHz this controller has a gain of 12 dB - much lower than the 32 dB for `landau_41`. The design has been tested at the Advanced Light Source and the measured open and closed-loop poles as well as the model-predicted curves are shown in Fig. 6.15. The tests indicate that the controller provides sufficient damping of the coupled-bunch modes inducing the damping rates on the order of the open-loop growth rates (approximately 6 dB gain margin). The phase response of the `alsFD3` is not as well controlled as that of `landau_41`. However within the frequency band of interest deviation from resistive feedback is minor, with reactive frequency shifts not exceeding 50 Hz. At higher currents some controller saturation is evident when the measured damping rates start falling short of the values expected from

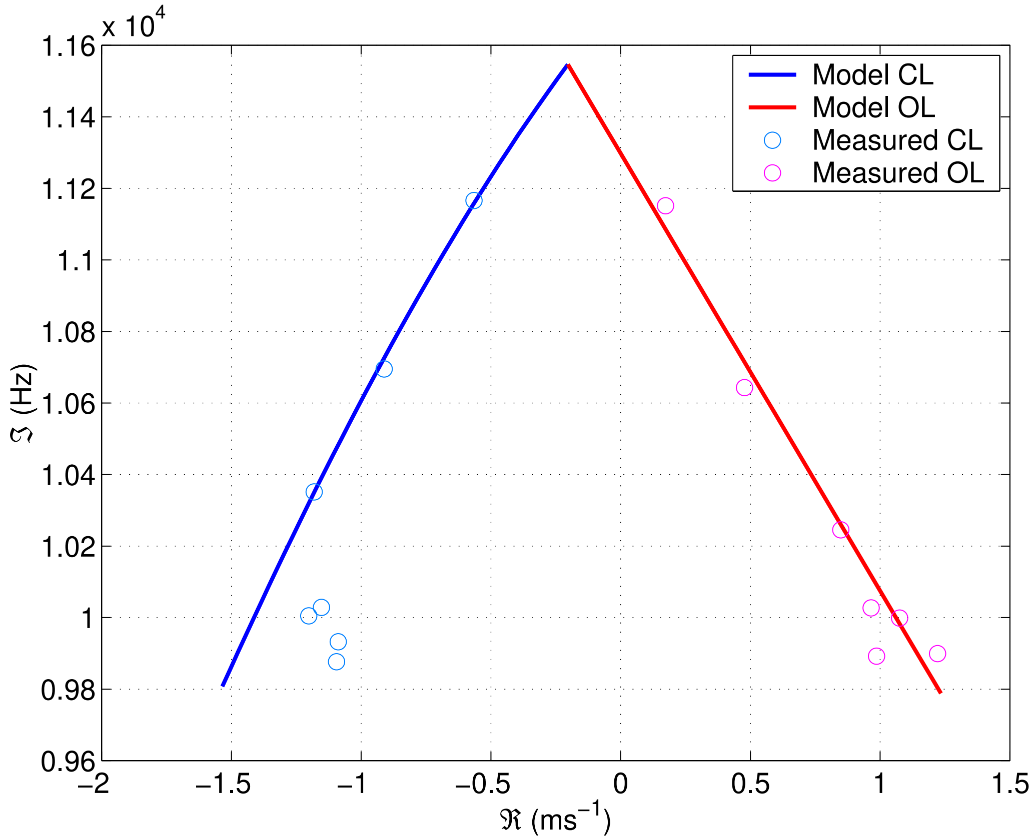

Figure 6.15: Measured open and closed-loop poles for mode 214 in the ALS. Model prediction of the closed-loop values is made using a linear fit to the measured open-loop eigenvalues.

the model.

As described earlier, the main goal of the new design is to avoid excitation of strongly tune-shifted eigenmode 0. In order to verify that the goal was achieved we compared the performance of the new controller with the existing design. Beam motion was recorded with the feedback system in the steady-state closed-loop situation. Measurements were taken at nearly the same beam current (395 and 400 mA) with each of the controllers in a 276 bunch fill pattern. The fill pattern was chosen as one of the standard ALS operating patterns with significant mode 0 problems. In Fig. 6.16 quadratically averaged bunch spectra for the two cases are plotted on the same axis. Data for the `landau_41` shows mode 0 motion excited to the level of 0.6 ADC counts at 4.8 kHz. The new design shows very little motion in that

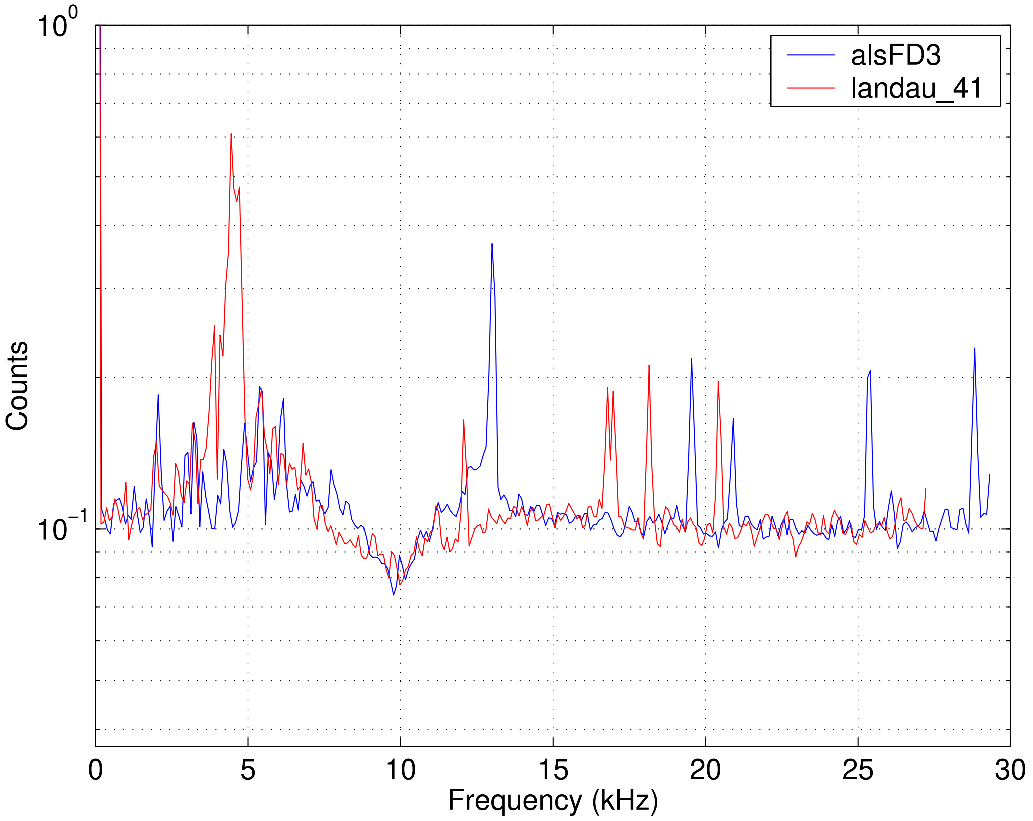

Figure 6.16: Bunch spectra quadratically averaged for the two IIR controller designs. Large mode 0 motion at 4.8 kHz excited by the `landau_41` controller is not present in the spectrum of the new `alsFD3` design. Narrowband peaks above 10 kHz are due to the external noise signals present at the input of the ADC. Since the two controllers operate at different downsampling factors and, consequently, different bunch sampling frequencies, the external signals are aliased differently. In each case the same number of 5 narrow lines appears in the spectrum.

frequency band confirming its better performance with respect to eigenmode 0.

6.3.4 DAΦNE

Longitudinal feedback problem at DAΦNE is very different from that at the ALS. There are no high harmonic cavities in this machine and the change in the synchrotron frequency over the operating range of currents is very small, going from 33 kHz at 1 mA to 32.5 kHz at 800 mA. However the beam loading of the main RF cavities is quite strong and leads to

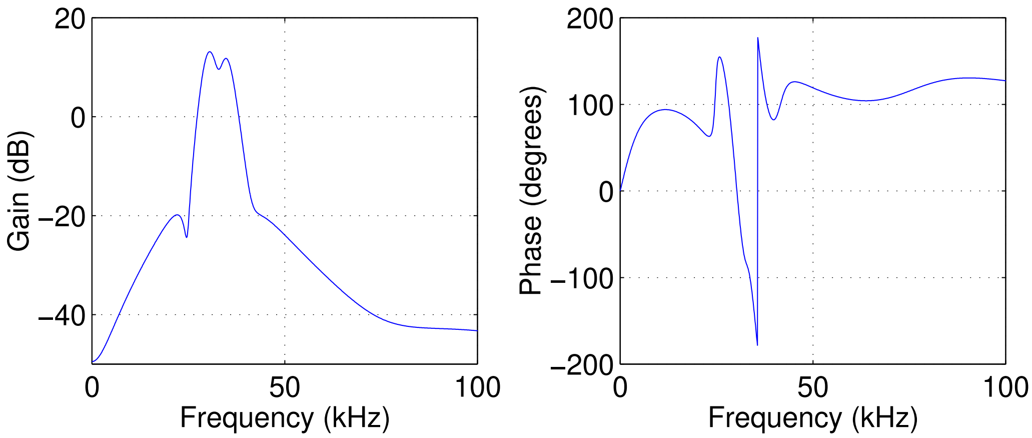

Figure 6.17: Frequency response of the DAΦNE electron ring longitudinal feedback controller.

a significant shift in the frequency of mode 0 going to 11 kHz at 1 A beam current. Due to its wideband response standard 6-tap FIR filter interacts with mode 0 driving it unstable. Several solutions have been tested before the IIR design was introduced.

The first method used higher center frequency of the FIR filter as a way to reduce controller gain at low frequencies. This approach partially alleviated the problem - partially since achievable gain reduction is limited by the shallow slope of the magnitude response vs. frequency. The second approach developed at DAΦNE is based on using a dedicated mode 0 feedback. Using the RF system power amplifier (klystron) as an actuator the system provides additional damping of mode 0 sufficient to combat the effect of the LFB.

Design of the IIR controller used the approach described in the previous subsection - rolling off the gain of the feedback controller so that mode 0 is weakly affected by the feedback. Mode 0 oscillation frequency changes in a continuous fashion from the nominal synchrotron frequency near zero current to the low value at the high current. Thus beam "samples" the response of the feedback controller over that range of frequencies. Roll off of the controller gain is accompanied by the fast change in the controller phase which can lead to instability of mode 0 if there is insufficient Robinson damping at that frequency of mode 0 (determined by the detuning of the RF cavities). Range of frequencies where gain of the controller is rolled off was selected based on the results of simulations of mode 0 frequency and damping [70].

In Fig. 6.17 the frequency response of the DAΦNE electron ring controller is shown.

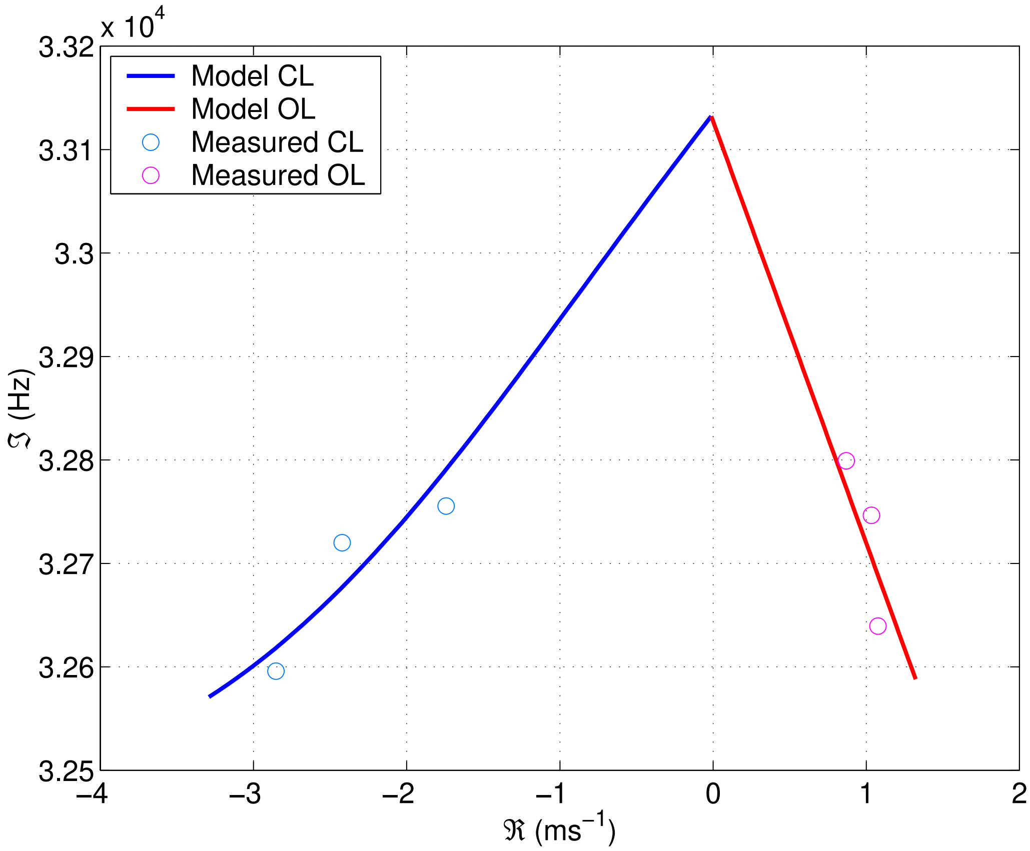

Figure 6.18: Open and closed-loop eigenvalues for mode 21 using DAΦNE electron ring controller. Model based curves are plotted for beam currents from 0 to 600 mA. Three data points were taken at 381, 409, and 546 mA.

Magnitude response of this controller is very narrowband. The lower cut-off starts at 29 kHz (-3 dB) and is below -20 dB at 25 kHz. According to the calculations and the simulation mode 0 damping peaks between 22 and 18 kHz oscillation frequency depending on the tuning angle of the RF system. Thus the controller shown here has little chance of exciting mode 0 instability. Shallower cut-off at high frequencies starts at 36 kHz and is at -20 dB at 43 kHz. The narrowband structure (7 kHz -3 dB bandwidth) of the controller rejects the sensor and quantization noise thus better utilizing actuator power. Phase shift in the control band of 32.5 to 33.5 kHz is near 90 degrees with only 6 degrees of variation. Phase slope does get steeper at the edges of the control band. Damping performance of the controller is illustrated in Fig. 6.18 for the beam currents from 0 to 600 mA. Gain margin is nearly 12 dB

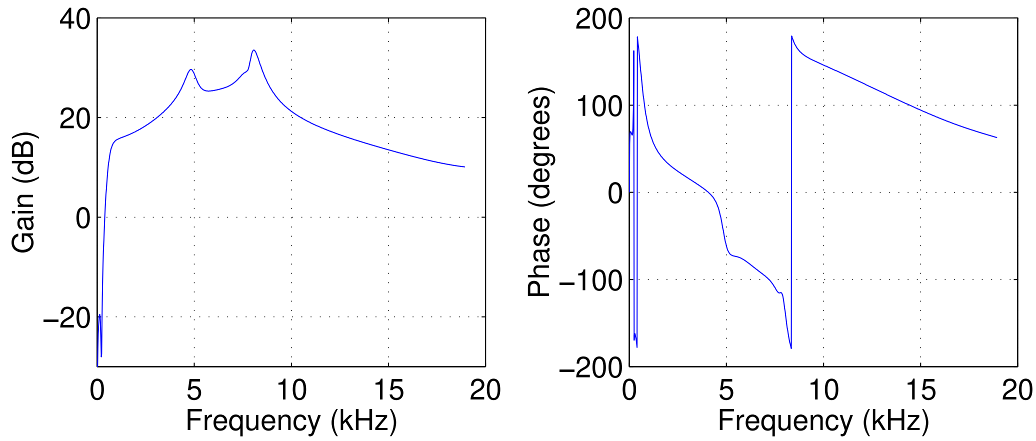

Figure 6.19: Frequency response of the BESSY-II longitudinal feedback controller.

as far as gain reduction is concerned. This performance is achieved at relatively low DSP gain peaking at 12 dB in-band leading to excellent noise performance of the closed-loop system. This controller has been successfully used in the day-to-day operation of DAΦNE.

6.3.5 BESSY-II

Longitudinal feedback control problem in the BESSY-II is similar to that in the ALS. There are 4 third harmonic cavities installed in the BESSY-II ring used for improving the lifetime. The third harmonic RF system is passive and, therefore, powered by the beam. As shown earlier, passive high-harmonic RF system causes changes in the synchrotron frequency as a function of current. The standard FIR controller has a limited control range due to the linear phase slope within the passband. An IIR controller has been designed using the frequency-domain method to provide resistive damping in the frequency band from 8 kHz nominal synchrotron frequency down to 4 kHz expected oscillation frequency of the unstable higher-order modes at the 400 mA beam current. Response of the controller implemented in 6 second-order sections is shown in Fig. 6.19.

The design was tested at beam currents up to 350 mA with two high-harmonic cavities actively powered and the remaining two parked. Only two cavities were used due to the technical (tuner) problems with the other half of the third harmonic RF system. As a result we observed oscillation frequencies shifting only to 6.8 kHz. Throughout the achievable operating range of the machine the IIR controller provided sufficient damping with the

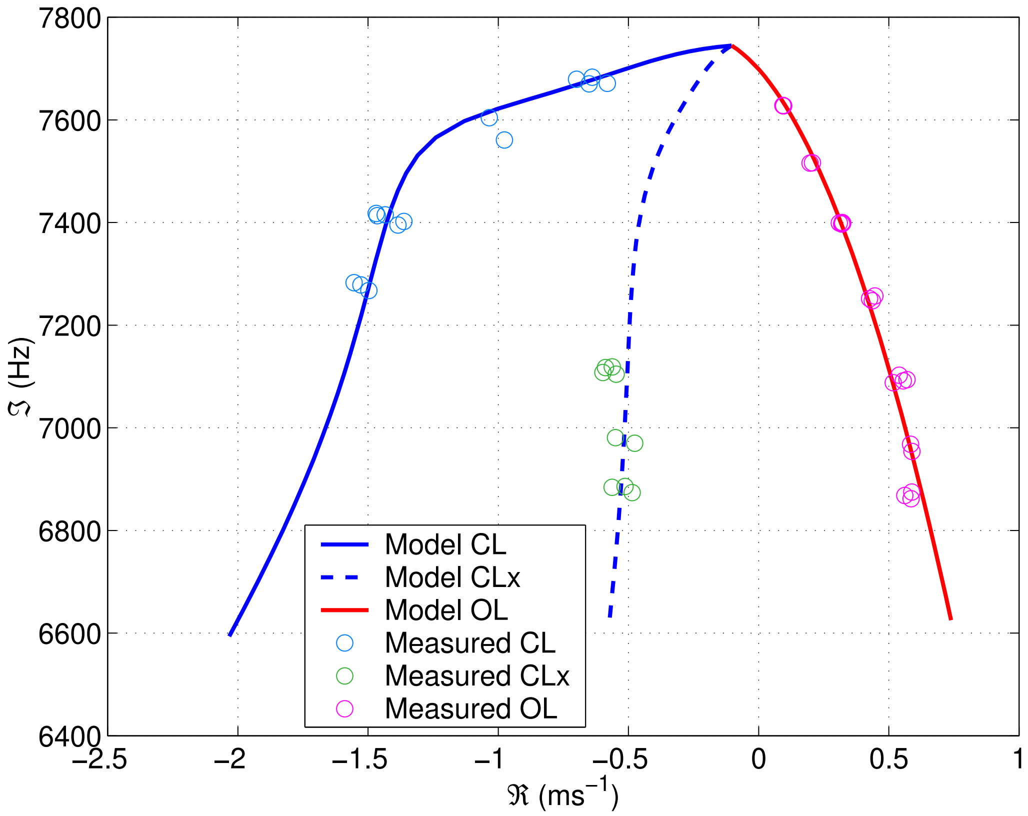

Figure 6.20: Open and closed-loop eigenvalues of mode 240 in BESSY-II using the IIR feedback controller. Model based curves are evaluated for beam currents up to 400 mA. Designation CLx refers to a set of measurements as well as modeling of the closed-loop system with 6 dB lower loop gain.

estimated gain margin of 11 dB at 400 mA. Eigenvalues of the EFEM 240 in both open and closed-loop configurations are shown in Fig. 6.20.

6.4 Model-based design

The main disadvantage of the frequency-domain design method presented in the previous section is that the result is only as good as the goal transfer function specified by the designer. In other words, the optimization distance function is not directly dependent on the closed-loop stability of the beam-feedback system. As the experience with the frequency-domain design grew this shortcoming of the optimization procedure became quite obvious.

In order to alleviate this deficiency a new optimization method described here has been developed.

Positions of the poles of the closed-loop system are the natural indicators of the stability. We will consider here a discrete-time model consisting of the discretized harmonic oscillator (the beam) and the feedback controller - the model that was developed in Ch. 5. The damping rate of a given pole in the z-plane is given by $\ln |p|$. Stable poles are within the unit circle, that is

$$|p| < 1 \Rightarrow \ln |p| < 0$$

In this optimization technique we will evaluate the closed-loop poles for a set of beam models defined by their eigenvalues spanning the desired range of growth rates and oscillation frequencies. By selecting the least damped pole (with the largest $\ln |p|$) we get the worst-case measure of the closed-loop stability with the given controller. The value of the damping rate can be used as a distance function for a numeric minimization procedure resulting in a minimax type of optimization. Using only the damping rate is insufficient, however, since such a distance function can lead to unacceptably large gains outside the frequency band determined by the eigenvalues of the our set of beam models.

We will make a simplifying assumption while selecting the set of beam models. As shown in the previous section, the set of the eigenvalues of interest forms a curve in the s- or z-plane. The independent variable is the beam current causing changes in the real and imaginary parts of the modal eigenvalues. In addition loop gain changes proportionally with beam current. Thus each point on the eigenvalue curve has a well-defined parameter set consisting of the real part of the eigenvalue (growth rate), imaginary part (oscillation frequency), beam current and loop gain. Unfortunately specifying this set requires many grow/damp measurements and the specifications change depending on the tuning conditions of the RF systems. To simplify the problem we will define the eigenvalue set using only 4 parameters: maximum beam current $I_{\max}$, largest expected growth rate at that current $\gamma_{\max}$, and the minimum and maximum oscillation frequencies over all beam currents $\omega_{\min}$ and $\omega_{\max}$. The all the eigenvalues in the set lie on the piece of a straight line parallel to the imaginary axis defined by

$$\begin{aligned} \Re(\Lambda) &= \gamma_{\max} \\ \Im(\Lambda) &\in [\omega_{\min}, \omega_{\max}] \end{aligned}$$

To define the finite set we will sample the above curve in some manner. In this sampling it is important not to leave large inter-point spaces so that the system stability could not be compromised between the two points. As the loop gain we will select the gain at the largest beam current.

This simplification is reasonable if one assumes that feedback damping scales linearly with the loop gain. Since growth rates and loop gain are linear with current, a system stable at the largest current is stable at all beam currents.

The minimax optimization method is illustrated in the sketch shown in Fig. 6.21. The open-loop poles are determined as described above. On the first iteration the closed-loop poles are evaluated for each of the 6 beam models. In reality there are more poles than shown in the sketch - here we only show the least damped pole (dominant pole) for each closed-loop system. Based on the damping rate the circled pole is selected for optimization. At some later point a snapshot of dominant poles shows a different arrangement of the poles and a different pole used for optimization. At the end the minimax optimization leads to equal damping of the dominant poles.

6.4.1 Optimization problem formulation

Let us mathematically define the goal function used in this optimization. We will start from defining N beam plants as described above. For each of these plants we will compute the closed-loop poles in the z-plane for the feedback controller specified by the parameter vector A. Let $\mathcal{P}(A)$ be a set of all computed poles for N closed-loop systems. Then the least damped pole in that set has damping given by $\max\{\ln|p|\}, p \in \mathcal{P}(A)$. The distance function $L(A)$ will be a linear combination of the above damping and the function $g(A)$ reflecting the gain of the controller outside the $[\omega_{\min}, \omega_{\max}]$ band. $g(A)$ is computed by first defining a set of discrete frequencies $\mathcal{W}$ which sample the bands where we would like to maintain low controller gain e.g. DC and low frequencies. Then we have

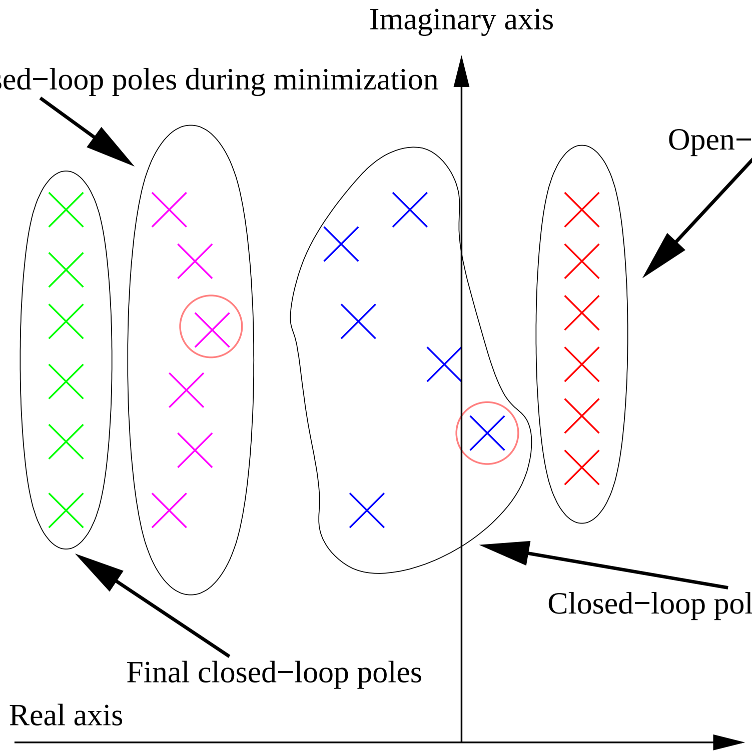

Figure 6.21: A sketch of the open and closed-loop poles in the s-plane for the model-based optimization

$$g(A) = \frac{1}{N_{\mathcal{W}}} \sum_{k \in \mathcal{W}} w_k |H(A, \Omega_k)|$$

where $N_{\mathcal{W}}$ is the number of elements in $\mathcal{W}$, w_k is the weight associated with discrete frequency Ω_k. In other words $g(A)$ is the weighted average of the the magnitude of the controller transfer function. This choice is by no means unique - one can select other functions, such as sum of squares, for gain-limiting. Overall distance function is given then by

$$L(A) = \beta \max_{\mathcal{P}(A)} \{\ln |p|\} + (1 - \beta) g(A) \tag{6.20}$$

where parameter $\beta \in [0, 1]$ controls the relative weights of the in band damping and the out of band gain.

The numerical optimization was implemented using the simplex search algorithm (Matlab function `fminsearch()`) [42], selected since it does not use numerical or analytic gradients. This feature is helpful in ensuring the convergence of the the optimization given that the problem is not continuous. In addition to the distance function minimization we need to guarantee the stability of the controller by confining the poles to a circle of radius ρ. The constraints are implemented as steep "walls" in the distance function by modifying Eq. 6.20 as follows

$$L_c(A) = L(A) + 10^6 \sum_{\mathcal{D}(A)} \frac{|p_k|}{2\rho} (\mathrm{sgn}(|p_k| - \rho) + 1) \tag{6.21}$$

where $\mathcal{D}(A)$ is the set of all controller poles and $\mathrm{sgn}(x)$ is the signum function defined as

$$\mathrm{sgn}(x) = \begin{cases} 1 & x > 0 \\ 0 & x = 0 \\ -1 & x < 0 \end{cases}$$

If all controller poles are within radius ρ the right hand side of Eq. 6.21 reduces to $L(A)$. Otherwise a large penalty is added for each pole violating the boundary. Coefficient 10^6 is chosen as being much larger than the values of $L(A)$ in a practical optimization.

Controller parameterization A used in this optimization is different from that used in the frequency-domain design and described in Eq. 6.14. In this case we implemented A as follows

$$A = (r_1, r_2, \ldots, r_N, \phi_1, \phi_2, \ldots, \phi_N, a_0, a_1, \ldots, a_{2N})$$

where r_k and ϕ_k are the polar coordinates of the k-th pole and a_0 through a_{2N} are the numerator coefficients. The overall controller transfer function is given by

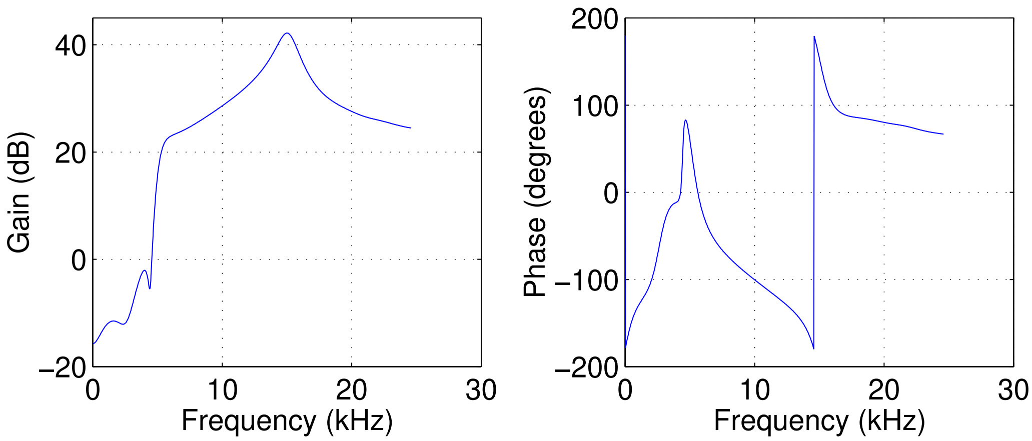

Figure 6.22: Frequency response of the ALS `dccut2` feedback controller.

$$H(A,z) = \frac{\sum_{k=0}^{2N} a_k z^{2N-k}}{\prod_{m=1}^{N}(z - r_m e^{i\phi_m})(z - r_m e^{-i\phi_m})}$$

6.4.2 ALS results

Controller named `dccut2` was designed for the ALS using the model-based optimization technique. The specification called for the control band from 6.5 to 12 kHz. The transfer function of this controller is shown in Fig. 6.22. The sharp gain roll-off below 5 kHz is used to avoid affecting mode 0 stability. Positions of the open-loop eigenvalues used in the design as well as the modeled closed-loop dominant poles are shown in Fig. 6.23. The optimization does a very good job of equalizing the damping of the closed-loop poles. Reactive tune shifts are not constrained in this optimization and can be quite large - note 1.6 kHz shift at 12 kHz oscillation frequency.

The `dccut2` controller has been tested at the ALS. The open and closed-loop eigenvalues are plotted in Fig. 6.24. Achieved damping is much lower than expected based on Fig. 6.23. The reduction is due to the fact that the design model was based on a different feedback configuration. Measurements were taken on a reconfigured hardware with 6 dB lower gain and 10 degrees of phase shift. As a result damping rates have very little margin and the system became unstable at beam currents above 340 mA.

The testing results pinpoint the main problem with this optimization method. Since

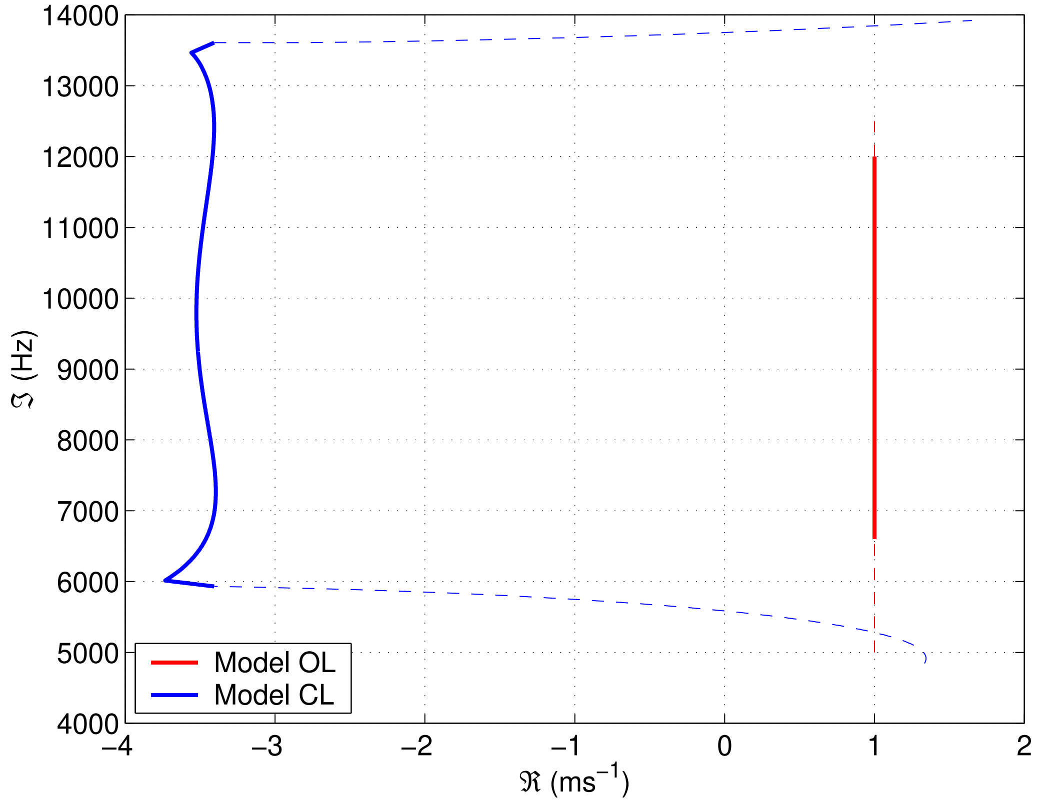

Figure 6.23: Open-loop eigenvalues and the modeled closed-loop dominant poles for the `dccut2` controller. Continuous lines show the behavior within the design control band while the dashed lines illustrate the performance just outside that band. Note that the damping (real part) of the poles within the control band is well equalized by the optimization. Some reactive tune shift is seen in the difference in vertical positions of the edges of the open and closed-loop lines.

the design is done by only considering the damping of the closed-loop system resulting controllers can have significant reactive effects. Controller performance is then strongly dependent on the feedback loop characteristics being the same as those used in optimization. In other words, model-based designs can suffer from small phase and gain margins. Let's consider what happens if the damping is achieved with a reactive angle of 60 degrees. In the optimization process the gain reduction of $\cos 60 = 0.5$ is compensated by making the gain at that frequency factor of two larger. If in the physical feedback system a loop phase shift of 10 degrees occurs feedback gain drops by $\cos 70/\cos 60 = 0.68$ rather than $\cos 10 = 0.98$. This reduction is usually coupled with significant increase in reactive tune shifts. As a

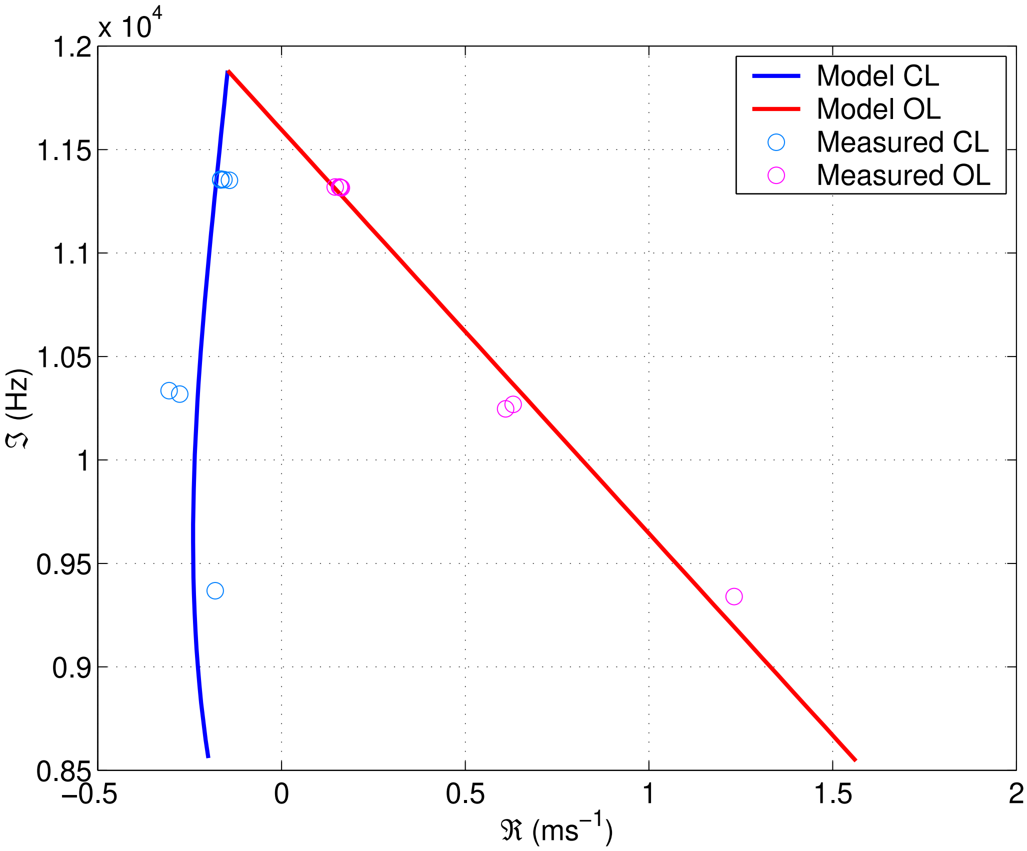

Figure 6.24: Open and closed-loop eigenvalues of mode 233 in the ALS measured with `dccut2` controller. The measurements span the range of currents from 66 to 300 mA while the model curves are shown from 0 to 400 mA.

result, small changes in the feedback system lead to instability.

6.5 Summary

In this chapter we've presented two controller design methods. Both of these techniques are based on the numeric optimization and generate a controller specification as a stable IIR filter of predetermined order.

In practical use the frequency-domain controller design method produces optimal approximations of the goal transfer function. The method is computationally efficient and enables the user to explore multiple goal functions. The main weakness of this approach, as pointed out in Sec. 6.4, is that the resulting controller is only as good as the goal transfer function since the method does not utilize the dynamic model of the beam and the feedback system.

The second design algorithm is based on such a model and uses closed-loop pole locations as a cost function. However this method is much slower than the frequency-domain optimization and, more importantly, does not address the question of robustness of the resulting designs.

Experience with the above methods suggests that good practical controllers could be designed using H-infinity approach which allows one to optimize robustness under parametric perturbations of the plant, the sensor, and the actuator.

Chapter 7

Transverse baseband processing architecture

7.1 Introduction

Modern circular accelerators operate above the coupled-bunch instability threshold. Such operation requires special passive and active measures to damp the unstable motion. Among passive methods are techniques of tuning troublesome impedances away from revolution harmonics [7, 9]. Another passive instability damping approach is based on designing fill patterns that couple unstable and stable modes [6]. Active damping using electronic feedback provides stability when passive methods are not sufficient. When the number of unstable modes is small and well known it is possible to design a mode-by-mode system that provides feedback gain only at the frequencies of the known unstable modes. However this approach is ill suited to cases when the number of unstable modes is large. Additionally, mode-by-mode system has to be modified if the machine impedance changes, e.g. when RF cavities or insertion devices are altered. An bunch-by-bunch feedback system acts independently on all the bunches. Such a system providing identical damping for each bunch is equivalent to a mode-by-mode system with the same damping on all modes.

Thus, a bunch-by-bunch feedback provides control of unstable coupled-bunch motion independently of the modal structure of the exciting impedance. A system proposed here carries an additional advantage of being easily adaptable to different accelerators due to the programmable baseband processing architecture.

In contrast to the all analog processing implemented for the ALS, BESSY-II, and PEP-II transverse systems, this DSP-based proposal uses and extends the beam and system diagnostic functions as well as system commissioning tools that were developed for the PEP-II longitudinal systems.

7.2 Overall system architecture

A block diagram of a bunch-by-bunch feedback system is shown in Fig. 7.1. As shown in the diagram there are three main blocks: analog front-end, baseband digital processing and analog back-end.

The main function of the front-end is to detect the motion of the bunches in the appropriate plane, longitudinal or transverse. It outputs pulse-amplitude modulated (PAM) signal that is sampled by the ADC. Pulse shape defines the sensitivity of the system to the shifts in clock timing or in the beam synchronous phase. Additional functionality can include orbit rejection to avoid saturating the ADC, servo circuitry in heterodyned detectors to correctly phase the carrier to the beam signal, monitoring outputs, etc. Experience with the longitudinal feedback systems at PEP-II, ALS, and BESSY-II indicates that integrated system diagnostics, such as the "fake-beam" generator, are extremely useful in signal path testing and adjustment. Such a generator simulates bunch signals seen at the bunch position monitor (BPM) outputs and can be modulated to include synchrotron or betatron motion. Diagnostic and controller development techniques described in Chapters 3, 4, 5, and 6 rely heavily on ability of the feedback system to record long beam motion sequences

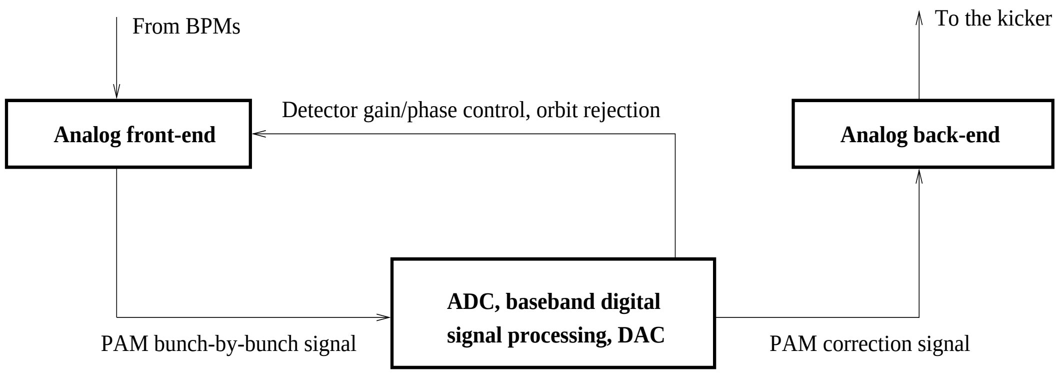

Figure 7.1: Block diagram of complete system

while manipulating feedback parameters.

The next block in the processing chain is the baseband digital signal processing. This module samples the PAM signal supplied by the front-end. Each bunches' position samples are processed using a digital filter. There are several objectives the above filter must achieve. First of all it has to provide proper phase shift at the coupled-bunch oscillation frequency for resistive damping of unstable motion. Second, user must be able to set the gain of the system. Rejection of the DC component of bunch motion is important since this component often exceeds the magnitude of coupled-bunch motion and can saturate the output. Filter output signal is converted by the DAC at the bunch crossing rate. DSP also has to provide some form of adjustable delay in order to apply correction to the right bunch.

Output of the DSP module is fed to the back-end. There baseband signal is modulated to place it within kicker bandwidth. After power amplifiers output signal is applied to the beam through a kicker.

7.3 Baseband DSP

In this chapter I will concentrate on the description of the baseband signal processing architecture. Block diagram of the proposed system is shown in Fig. 7.2. Input PAM signal is sampled at the bunch crossing rate. Resulting data stream is demultiplexed to bring down the processing rate. Output of the demultiplexer is sent to a number of processing modules. Computed correction samples are combined by the multiplexer into a bunch crossing rate stream which is then sent to the DAC. In this design ADC and DAC are assumed to be 8 bit wide.

7.3.1 Sampling and demultiplexing

The industry-standard front-panel data port (FPDP) interface has been selected to carry data between modules. This interface is 32 bit wide so 4 samples can be transferred over one link simultaneously. Maximum clock rate for FPDP is specified as 40 MHz. For 500 MSPS system at least 4 links are needed ($500/4/40 = 3.125$).

Standard serial-to-parallel demultiplexing illustrated in Fig. 7.3 imposes a severe limitation on the ring harmonic number. Since bunch data on every turn has to be sent to the same processing element, ring size has to be an integer number of samples processed in parallel. With 4 links harmonic number will be limited to multiples of 16. This is

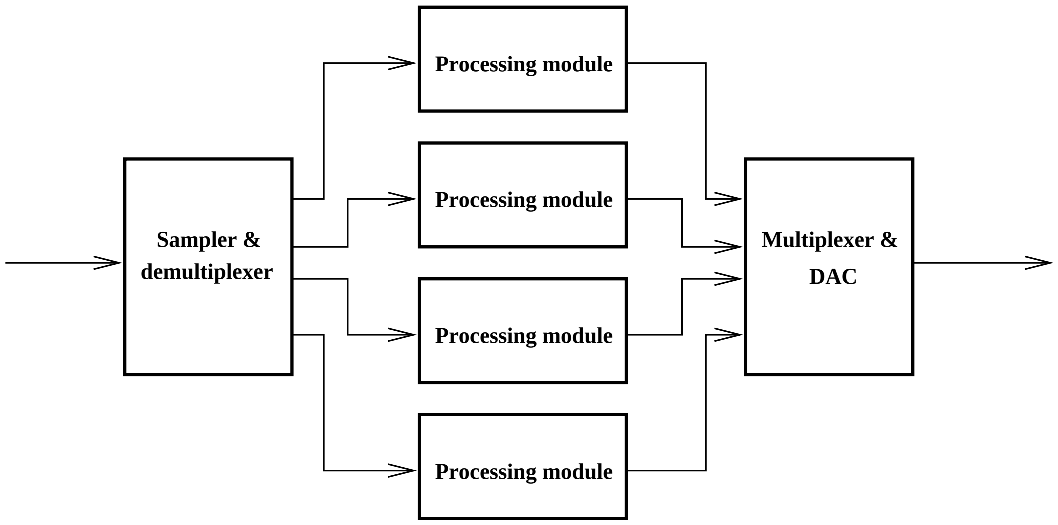

Figure 7.2: Block diagram of baseband processing section

very severe indeed: out of 6 machines supported by SLAC longitudinal feedback (PEP-II, ALS, DAΦNE, BESSY-II, PLS, SPEAR-3) [36] only one has harmonic number divisible by 16. A feedback system using standard demultiplexing into 24 channels is being built for Sincrotrone Trieste and Swiss Light Source since both machines have acceptable harmonic numbers [71]. The KEK B-factory feedback systems use demultiplexing by 16 and are limited to the multiple of 16 harmonic numbers.

In order to get around this limitation a new demultiplexing scheme is proposed here. I will call this approach "uneven stepping". Basic idea of uneven stepping is to vary the size of parallelized transaction. If, instead of always sending N samples through parallel links, we allow N and $N-1$ samples to be sent, then any harmonic number h above some minimum $h_{\min}$ can be handled by this system. Below the minimum some harmonic numbers cannot be supported. Let's prove the above statement and determine $h_{\min}$ as a function of N. Define K as the total number of parallel transactions per revolution.

$$K = \left\lfloor \frac{h}{N-1} \right\rfloor$$

If K transactions of $N-1$ samples are used, we will have some number of bunches M left over. M is defined as the remainder of dividing h by $N-1$:

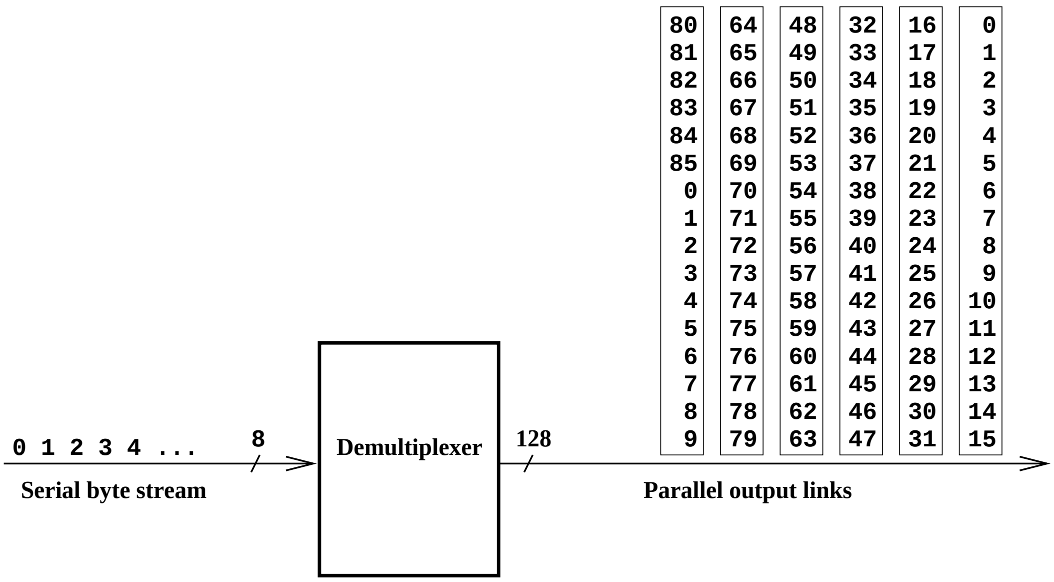

Figure 7.3: Standard demultiplexing for h=86. Notice that bunches do not end up in the same processing channel on the next turn

$$M = \mathrm{rem}(h, N-1) = h - \left\lfloor \frac{h}{N-1} \right\rfloor (N-1) = h - K(N-1)$$

If, however, M transactions are N samples wide, we get

$$(K-M)(N-1) + M * N = K(N-1) + M = h \tag{7.1}$$

Therefore this algorithm allows one to match any harmonic number. Now let us derive $h_{\min}$. For Eq. 7.1 to be valid we need $(K-M) \geq 0$. Thus

$$\begin{aligned} K - M \geq 0 &\Rightarrow K \geq M \\ \min(K) = \max(M) &= N-2 \\ \min\left(\left\lfloor \frac{h}{N-1} \right\rfloor\right) = \frac{h_{\min}}{N-1} &= N-2 \\ h_{\min} &= (N-1)(N-2) \end{aligned}$$

Selecting $N = 16$ gives $h_{\min} = 210$. This is more than 120 which is harmonic number

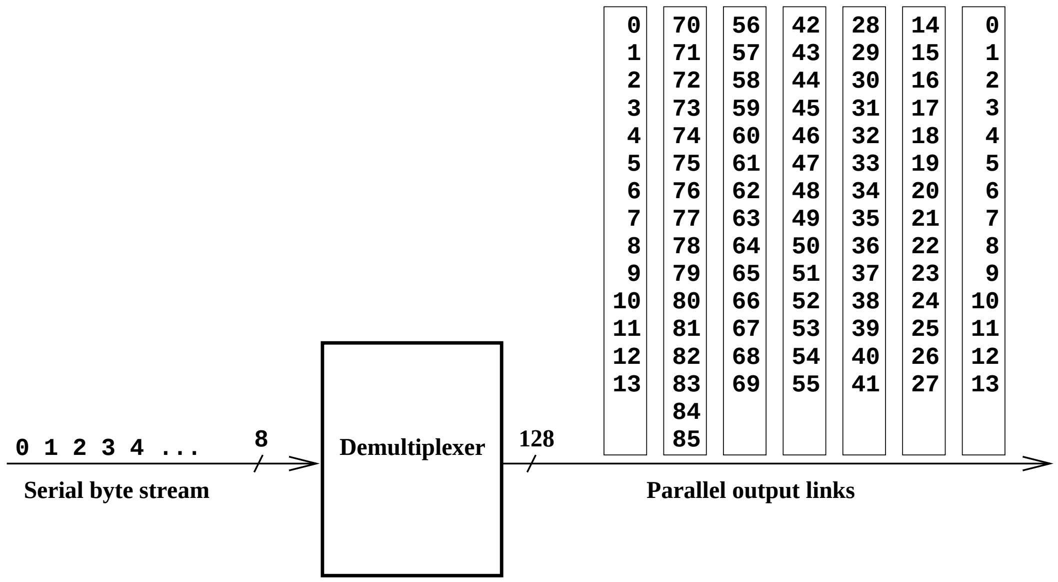

Figure 7.4: Uneven stepping demultiplexing for h=86. Parallel transactions of width 16 and 14 are used. Compare to Fig. 7.3 for turn-by-turn sample alignment.

for one of the target machines, DAΦNE. However 120 is evenly divisible by 15, so it can be matched. Another possible approach is to limit harmonic numbers to be only even. Then using transaction sizes of 16 and 14 samples any even h larger than 84 is achievable. This approach is additionally attractive since it avoids clocking data through the demultiplexer at full bunch crossing rate. With dual clock/dual-output ADC clock rates can be kept to half the full bunch crossing rate.

Figure 7.4 illustrates the uneven stepping demultiplexing. In this example $h = 86$ is matched using transactions of 14 and 16 samples. Computations can be done in the $\frac{h}{2}$ basis. Thus we get

$$\begin{aligned} h_{\frac{1}{2}} &= 43; N_{\frac{1}{2}} = 8 \\ M_{\frac{1}{2}} &= \mathrm{rem}(43, 7) = 1 \\ K_{\frac{1}{2}} - M_{\frac{1}{2}} &= 5 \end{aligned}$$

Five transactions of width 14 and one 16 sample wide are needed for proper matching. As shown in Fig. 7.4 bunches are sent to the same processing channels on each turn.

7.3.2 Signal processing and diagnostics

As described previously sampled data is sent over FPDP links to signal processing modules. For the 14/16 uneven stepping maximum clock rate of the samples on the link is given by 500 MHz/14 = 35.7 MHz. Evaluation of various options lead to the processing architecture based on Xilinx DSP Core. Data from an input FPDP link is directed to 4 Xilinx Virtex field-programmable gate array (FPGA) chips. Each chip handles one out of four samples from the link.

FPDP features

FPDP links use TTL level 32-bit wide data bus [72]. Data is strobed using PECL differential clock with 40 MHz maximum clocking rate. There are several handshaking signals as well as a frame synchronization signal. This signal can be used to send fiducial information to the processing modules. Additionally there are two general-purpose I/O lines per link. One of these lines can be used as a trigger bit for transient measurements.

Signal processing

Figure 7.5 illustrates internal FPGA architecture for a 5 tap FIR filter. FIFO blocks are used to store past samples. For processing multiple bunches length of each FIFO block has to be set equal to the number of bunches handled by the chip. Using example illustrated in Fig. 7.4 we conclude that each chip processes 6 bunches except the last two FPGAs each of which only handles one bunch. Adjustable length FIFOs would allow to reconfigure the system through software for operation with a particular harmonic number, while fixed length delays would require separate FPGA layouts for each machine. Shift gain block shown at the output is a barrel shifter for controlling the loop gain in 6 dB steps.

All addition blocks in the drawing include saturation in order to maintain correct feedback phase in presence of large input signals. Similarly, shift gain includes output saturation.

For a feasibility test an evaluation system has been built using Xilinx XC4085XLA FPGA chip. Signal processing architecture similar to the one shown in Fig. 7.5 has been implemented and tested at 40 MHz processing rate.

Filter layout illustrated is not the only one possible. Flexibility of FPGA-based design allows various FIR and IIR structures to be implemented. Main limitation to the size and complexity of the processing structure comes from cell count of the FPGA and routing

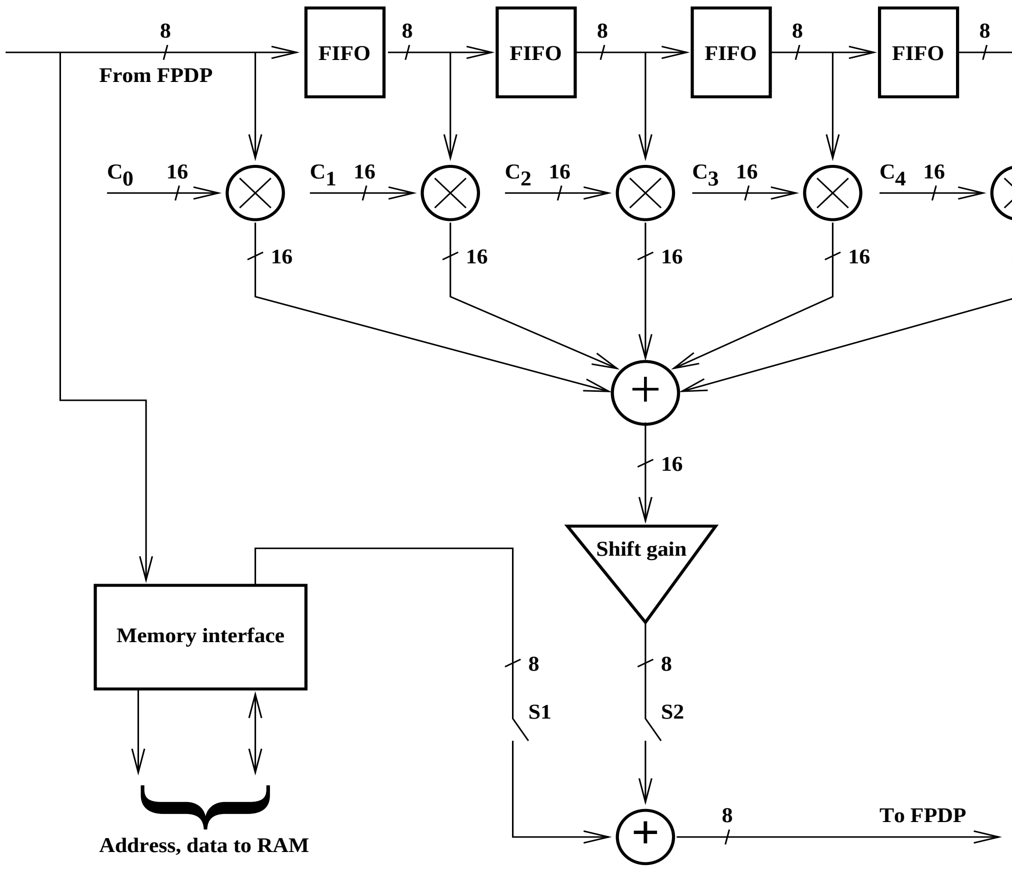

Figure 7.5: Signal processing architecture within a single FPGA

restrictions.

Diagnostics

Figure 7.5 illustrates some of the diagnostic features of the DSP modules. Diagnostic requirements are based on the experience with longitudinal DSP-based feedback [73, 74]. Here is a list of desired diagnostic features in the DSP module:

- Input data acquisition
- Arbitrary output waveform generation with bunch-by-bunch control
- Changing filter coefficients on the fly
- Transient diagnostics:

Table 7.1: Data acquisition and playback

S2	S1	Function	Applications
Off	Off	Feedback off, drive off	Recording growing transient, feedback off on a given bunch for tune measurement
Off	On	Feedback off, drive on	System timing, open-loop transfer function
On	Off	Feedback on, drive off	Normal operation, damping transient recording, closed-loop recording
On	On	Feedback on, drive on	Transfer function measurement, closed-loop transfer function

  - Grow/damp
  - Single-bunch tune measurement

- Closed-loop transfer function
- Open-loop transfer function

Data acquisition and waveform playback functions are controlled by the memory interface block. Partitioning memory in two sections will allow simultaneous recording and playback as needed for transfer function and tune measurements. Two switches S1 and S2 control the output of the FPGA. Table 7.1 lists possible settings of switches S1 and S2 and resulting functionality. In physical implementation S2 can be eliminated since setting filter coefficients to all zeros will achieve the same goal.

Filter coefficients will be kept in two register banks. Each bank can be updated by the software when it is not used by the real-time processing. Bank select can be generated by multiple sources, e.g. filter bit from FPDP link, software, bunch-by-bunch register. Let's consider these one by one. Filter bit carried by the FPDP link allows one to record data for all bunches simultaneously. Software control of this bit within the demultiplexer module provides unsynchronized grow/damp measurement capability. Additionally this bit can be hardware-synchronized to enable simultaneous grow/damp measurements in both transverse planes as well as injection transient measurements. Under software control coefficients can be switched for modification of the online filter during system tuning or in adaptive control mode. Bunch-by-bunch register has one bit per each bunch processed by the FPGA. This feature allows modification of feedback parameters for a given bunch or a set of bunches.

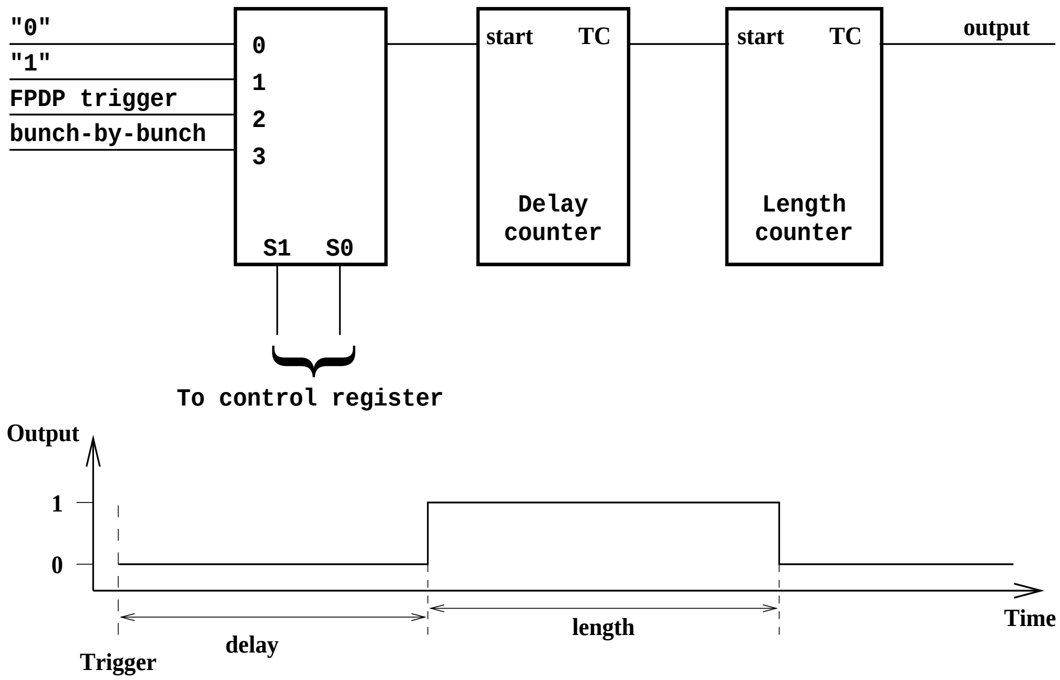

Figure 7.6: Trigger block

Trigger controls of data recording, playback and filter select are very similar. All of them are based on a programmable pulse generator with a multiplexed edge-sensitive trigger. There are several parameters that can be independently controlled within each trigger block. Figure 7.6 shows a possible layout of such a block and indicates controllable parameters. Since controls of recording, playback and filter select are independent, many experimental configurations can be defined.

In diagnostic measurements recording data for each bunch on every turn provides complete information on beam dynamics. However at 500 MSPS acquisition rate system will produce very large data sets. In grow/damp measurements length of the transient is determined by the growth and damping times of unstable modes. For most machines these times are in the 0.1–10 ms range. If the system can record 20 ms worth of data, even the slowest growth and damping times can be measured. This requires 10 MB of memory available for data acquisition. With a little margin 1 MB per FPGA satisfies this requirement. If we assume that all groups are 14 samples wide (worst case), then a total of 14 MB of data can be acquired.

However in many cases recording data on every turn is redundant. Since the oscillation frequency is well known (the tune) and the envelope of motion is relatively slow, details of the motion can be reconstructed from undersampled data. Thus in acquisition data can be downsampled. There are several benefits of downsampling in this context. For the fixed duration of the transient much less data has to be handled. Conversely, long transients can be captured using all the memory.

Note that downsampling only extends to the diagnostic functions. Feedback control is still computed on every turn. However one could use downsampling in the feedback path, for example in a longitudinal feedback application.

Downsampled transverse diagnostics have been done using existing longitudinal feedback systems at PEP-II and ALS. These experiments successfully demonstrated the feasibility of the approach and provided experience on advantages and limitations of such diagnostics.

Let's investigate downsampling process in more detail and determine limitations driving downsampling factor selection. Let us consider motion of a single bunch with one even-fill eigenmode growing exponentially.

$$x(t) = e^{\lambda t} \sin \omega_x t \tag{7.2}$$

Taking Fourier transform of the above signal we get

$$X(j\omega) = \frac{\omega_x}{\omega_x^2 + \lambda^2 - 2\lambda j\omega - \omega^2} \tag{7.3}$$

If the signal is sampled at the downsampling rate M the following sequence is obtained

$$x_s[n] = x(nT_s) = e^{\lambda T_s n} \sin(\omega_x T_s n)$$

where $T_s = MT_0$ and T_0 is ring revolution time. The Fourier transform of this sequence is related to the Fourier transform of the continuous-time signal as

$$X_s(e^{j\omega}) = \frac{1}{T_s} \sum_{k=-\infty}^{\infty} X(j\frac{\omega}{T_s} - j\frac{2\pi k}{T_s}) \quad \text{for } -\pi \leq \omega < \pi \tag{7.4}$$

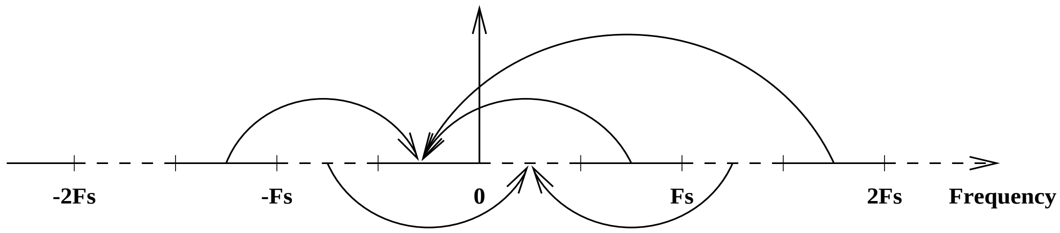

Figure 7.7: Aliasing during sampling. Spectral components in frequency regions drawn in the same line styles are summed together. Arrows show aliasing of high frequency components to baseband.

This equation describes aliasing of the continuous-time signal during sampling. The process is graphically illustrated in Fig. 7.7. Under certain conditions one can reconstruct $X(j\omega)$ from $X_s(e^{j\omega})$. These conditions are well-known for the case when $X(j\omega)$ has low-pass form. Then reconstruction will be perfect for $X(j\omega)$ such that $X(j\omega) = 0$ for $|\omega| \geq \pi/T_s$ Now we will formulate similar conditions for the bandpass $X(j\omega)$. Since $x(t)$ is a real function its Fourier transform is Hermitian, i.e. $X(-j\omega) = X^*(j\omega)$. Thus, to avoid adding non-zero components in Eq. 7.4 we need to have $X(j\omega) = 0$ everywhere, except for one frequency band Ω_l, where

$$\Omega_l = \{\, \omega : \frac{\pi l}{T_s} \leq |\omega| < \frac{\pi(l+1)}{T_s} \,\}$$

In this case we would be able to perfectly reconstruct continuous-time signal from its samples. However from Eq. 7.3 it is obvious that the spectrum is non-zero at all frequencies. Therefore some errors will be introduced in reconstructing the signal. In order to quantify these error let us consider the magnitude of the spectrum $|X(j\omega)|$. This second-order function has one maximum:

$$\max_{\omega} |X(j\omega)| = |X(j\hat{\omega})| = \frac{1}{2\lambda}$$
$$\hat{\omega} = \sqrt{\omega_x^2 - \lambda^2}$$

Let us define integer l such that

$$\hat{\omega} \in \Omega_l$$

If we define spectral attenuation level R we can determine frequencies ω_R where $|X(j\omega)|$ falls below $|X(j\hat{\omega})|/R$.

$$\omega_R = \sqrt{\omega_x^2 - \lambda^2 \pm 4\lambda\omega_x\sqrt{R^2 - 1}} \tag{7.5}$$

For $\lambda \ll \omega_x$ we can simplify the above as

$$\omega_R \approx \omega_x \pm 2\lambda\sqrt{R^2 - 1}$$

Notice that for large R Eq. 7.5 only has one real root. That is due to the fact that $|X(0)| = \frac{\omega_x}{\omega_x^2+\lambda^2}$. Thus

$$R_{\max} = \frac{|X(j\hat{\omega})|}{|X(0)|} = \frac{\omega_x^2 + \lambda^2}{2\lambda\omega_x} \approx \frac{\omega_x}{2\lambda}$$

In order to minimize reconstruction errors downsampling factor should be selected so that peak of the signal spectrum is placed at the center of the frequency band Ω_l. Thus

$$\begin{aligned} \frac{\pi l}{T_s} &= \hat{\omega} - \frac{\pi}{2T_s} \\ \frac{\pi(2l+1)}{2MT_0} &= \hat{\omega} \end{aligned}$$

Given desired downsampling rate M we can compute nearest optimal rate $\hat{M}$ as follows

$$\begin{aligned} \hat{l} &= round\left[\frac{1}{2}\left(\frac{2MT_0\hat{\omega}}{\pi} - 1\right)\right] \\ \hat{M} &= round\left[\frac{\pi(2\hat{l}+1)}{2T_0\hat{\omega}}\right] \end{aligned}$$

For a given attenuation level R one can determine maximum downsampling rate M_R such that at the edges of the aliasing band spectrum magnitude is below $1/2\lambda R$. We have

$$\Delta(\omega_R) \approx 4\lambda\sqrt{R^2-1} < \frac{\pi}{M_R T_0}$$
$$M_R < \frac{\pi}{4\lambda T_0 \sqrt{R^2-1}}$$

These calculations are all done for the ideal case of motion described in Eq. 7.2. The same applies to the damping transient since the only difference between the two spectra is in phase component. However in a real machine there are factors that need to be taken into consideration when selecting downsampling factor. First of all, one has to take into account coherent and incoherent tune shifts which widen the required bandwidth. Additionally, tune shifts with amplitude have to be considered. Another factor is synchrotron modulation sidebands on the transverse signals. Large level of longitudinal motion will add to the aliasing. Finally, effect of the wideband noise needs to be counted since downsampling sums broadband noise from multiple frequency bands.

7.3.3 Multiplexing

Feedback bunch-by-bunch corrections are computed from the bunch samples. These corrections are to be assembled together in a single data stream feeding output DAC. This task is accomplished by the multiplexer module as shown in Fig. 7.2.

Another function multiplexer has to perform is buffering. Data from the demultiplexer is sent out in real-time, so there are two different transaction durations, long and short corresponding to N and $N-1$ wide transactions. Due to pipeline delays within processing it is possible that several wide transactions are sent to the multiplexer at short transaction timing. Conversely, there will be the same number of narrow transactions arriving at long intervals.

Let's consider demultiplexing scheme of 14 and 16 sample groups described in section 7.3.1. Transaction timing is shown in Fig. 7.8. Data cycles from the demultiplexer module to the signal processing boards are shown on the top. Cycles 6, 7, 8, and 9 are 16 sample wide (wide transactions) while other cycles are 14 sample wide (narrow transactions). In the middle are data cycles from signal processing boards to the multiplexer. Since processing is clocked by the input FPDP, timing of the cycles is identical to the ones

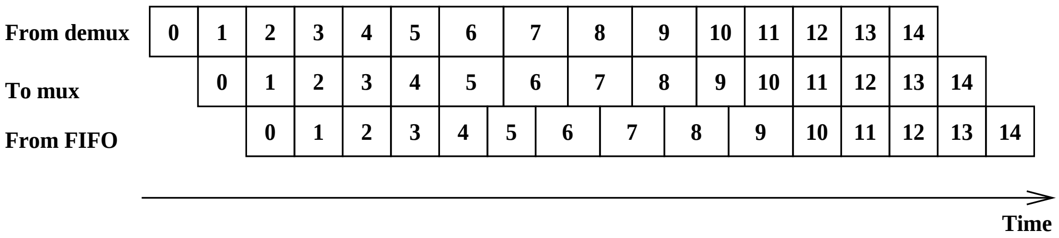

Figure 7.8: Transaction timing with one cycle processing delay

on the top. However due to the processing delay of one cycle there is mismatch in transaction width and timing. For example, transaction 5 is narrow, but has long timing while transaction 9 is wide and has short timing. For this configuration having a FIFO filled with one cycle will solve the problem. The bottom row shows transaction reads from the FIFO. When transaction 5 is written into the FIFO, demultiplexer fetches transaction 4 from the FIFO and starts outputting it. When that is done, transaction 5 is fetched from the FIFO. That will happen 2 RF clocks before transaction 6 is written into the FIFO. Thus FIFO will be empty for a short period (2 clocks). Similarly it will empty on each cycle until transaction 11 is written.

In order to quantify the nominal depth of the FIFO let's consider a worst-case scenario. This is when processing delay is more than the number of wide transactions. Let's define K as the number of wide transactions. Then initially K narrow transactions are sent to the FIFO with long timing. If N is the size of the wide transaction, then $K(N-1)$ samples are written to the FIFO and KN are read in that period. Thus FIFO should have K sample reserve. For N and $N-1$ uneven stepping we can have at most $\max(\mathrm{rem}(h, N-1)) = N-2$ wide transactions. Therefore keeping one transaction in the FIFO (at least $N-1$ samples) is sufficient since the largest reserve required is $N-2$ samples. Similar argument applies to the uneven stepping with N and $N-2$ sample transactions.

An additional function the multiplexer is required to perform is digital delay. This is needed to allow alignment of the feedback output with the beam. This delay should cover a range of one revolution to compensate all possible sensor/kicker placements and interconnect delays. Buffering FIFO described above can support this function. By adjusting the read clock delay relative to the revolution fiducial signal with RF clock cycle granularity one can move the feedback output with the same granularity. Then maximal depth of the FIFO defines the largest delay achievable. For 14/16 sample transaction design a 256 word FIFO

(256×128) provides up to 3584 sample delay. Since the FIFO is clocked at the FPDP rates, it can be easily implemented using CMOS technology.

Chapter 8

Conclusions and future research directions

8.1 Summary

This work mainly focused on the problem of the feedback control of the longitudinal coupled-bunch instabilities. In Chapter 2 the general framework of the problem was presented, a linear dynamic model of the unstable motion and multiple factors that affect the parameters or the dynamics of this model were described. The same chapter gave a detailed description of a longitudinal bunch-by-bunch feedback system used for all of the measurements and feedback control developments presented in this thesis.

Chapters 3, 4, and 5 described several new techniques developed to thoroughly and accurately characterize the dynamic system - the beam and the overall feedback loop. These techniques were illustrated with many measurements from several accelerators showing the wealth of information these methods provide.

In the following chapter, Ch. 6, two methods for designing longitudinal feedback controllers were presented. Applications of these methods to three accelerators, each with unique control requirements, were shown. Measurements verifying correct operation of the resulting designs were also presented.

In chapter 7 a new design of a programmable baseband processing channel was introduced. This design extends the control and diagnostic features of the existing longitudinal feedback system to all three dimensions.

8.2 Main new contributions

A significant part of this thesis was brought to life by a desire to effectively control coupled-bunch instabilities in the ALS after the third harmonic cavities were added to the ring. The search for feedback controller architectures capable of robustly stabilizing the longitudinal motion led to the development of a thorough system model and required careful parametric mapping of instabilities. Inevitable failures of intermediate controller designs led to the realization that the overall system had small stability margins further eroded by the loop parasitics. In order to create a robust controller, analysis of parasitic effects in the feedback loop and experimental verification of such analysis was necessary. Finally, the search for better controllers culminated in the development of two optimization-based design methods.

The new method for extracting the complex eigenvalue from an exponentially growing or damping modal oscillation is a foundation for all the beam and feedback characterization techniques presented in this thesis. Since these techniques most often require multiple measurements to obtain accurate results, a fast and reliable way to analyze grow/damp measurements is essential.

Another important new contribution presented in this work is the development of a thorough mathematical model of the longitudinal feedback system. Such a model is critical for development and testing of feedback controllers. In this thesis the model has been developed and verified against the physical system. In the iterative process of controller design it is important to understand the causes of failures to maintain stability. A detailed analysis of the most significant parasitic effects, such as bunch-to-bunch coupling, bunch synchronous phase shifts, etc., has been presented here. Such analysis also serves as a guide for the future feedback channel designs helping to chose front-end detection and kick signal modulation methods.

As indicated earlier, one of the most important new developments presented in this work is the optimization-based IIR controller design methodology. The two techniques presented here use numerical optimization to meet the specified feedback requirements such as control bandwidth, overall gain, etc. These design methods have been successfully used to produce feedback controllers for three accelerators: ALS, DAΦNE, BESSY-II.

8.3 Promising research directions

The research work that has been presented in this thesis is by no means complete. Here I will outline several promising new research directions starting from the transient diagnostics.

A significant limitation of the existing transient grow/damp diagnostic is the sensitivity to the modal growth rates. During the measurement the feedback loop is open for a predefined period of time. If the bunch oscillations grow very large during that period the feedback may be unable to damp the motion and the feedback control will be lost. Choosing the duration of the open-loop section is difficult since the final motion amplitude depends on both the growth rates of the unstable modes and the initial modal amplitudes (due to the noise excitation) when the feedback is turned off. Since the growth rates depend strongly on environmental variables such as resonator temperatures it is impossible to determine in advance the optimal growth transient duration. In this situation one must err on the side of caution and use conservatively short time intervals. As a result, captured transients are lower in amplitude and closer to the noise floor making eigenvalue estimation more difficult.

A promising method for addressing this problem is automatic growth termination. In this technique one would turn on the feedback automatically when the beam motion grew to significant amplitudes. Unfortunately using the input motion signals is difficult due to DC offsets and out-of-band noise. However, the feedback system is already equipped with the necessary signal - the controller output. If the feedback loop is opened at the output of the controller all one has to do is observe the kick signal requested by the controller. When that signal approaches DAC saturation the loop should be closed to maintain feedback control. Such an approach implicitly scales the transient cut-off point to the existing feedback loop gain.

If automatic growth termination method is implemented it opens the way for addressing another shortcoming of the existing grow/damp technique. In order to get a reliable measurement of both open and closed-loop eigenvalues one needs to record several transients, usually from 5 to 10. This is to eliminate errors due to growth rate variations and initial conditions dependence especially significant for the fast transients. At the same time the fast transients rarely use all of the recording time available, most often occupying from one-tenth to one-quarter of the record. If transient measurements are automated several grow/damps can be performed in series thus cutting down the number of separate experiments that need to be performed. In addition, taking several measurements in quick succession helps avoid the problems with beam current decaying between the transients.

The eigenvalue extraction method described in Ch. 3 assumes that the modal motion is adequately described by a second-order model. This assumption is true for the open-loop case when the motion is small. However in a closed-loop situation there can be multiple system poles with similar damping times. In fact, configuring a longitudinal feedback system for maximum damping usually results in two poles close in frequency and with the same damping. In this case closed-loop transient might not be an exponentially decaying oscillation and the eigenvalue extraction can fail. Sometimes the damping transient is well fit by the second-order response since one pole is excited to a much lower amplitude than the pole nearest the open-loop oscillation frequency. However even in this case the extracted eigenvalue cannot be used for comparisons with the system model with multiple poles. To avoid this problem one can consider using more sophisticated system identification techniques to extract the closed-loop system model. For example, it is possible to measure the modal closed-loop transfer function using DSPs to add an excitation signal to the feedback output signal and to record the response of the beam. The excitation sequence would have to be designed to correctly place the power in the spectrum taking into account the sensitivity of the plant to different frequencies.

In this thesis to analyze bunch motion we always projected it into the EFEM basis. Unfortunately the ring filling patterns most often encountered in practice are uneven. In some cases the unevenness is small and does not significantly perturb the EFEM analysis. However in many situations even-fill eigenmodes are not suitable for analyzing the motion. It is important to develop methods to compute the uneven fill eigenmodes. Theoretically, such computation is possible but it requires precise knowledge of the impedance seen by the beam and of the bunch currents. Such information is rarely available making that approach impractical. Some of my recent research (still unpublished) uses linear combinations of EFEMs to reconstruct the uneven fill eigenmodes. This method has been successfully applied to the analysis of data from BESSY-II and PEP-II but will benefit from further development.

Feedback controller design methods presented in this work have certain limitations and further work in this direction is quite important. Significant advances in controller design could be made by applying the tools of μ-synthesis and optimal H_∞ control.

Diagnostic methods presented in this work are absent from the existing transverse bunch-by-bunch feedback systems. Implementation of the processing architecture described in

Ch. 7 will significantly extend the reach of these diagnostics.

Bibliography

[1] M. Sands, "The physics of electron storage rings: An introduction," Tech. Rep. SLAC-R-121, SLAC, Stanford, CA, USA, 1970.

[2] S. Prabhakar *et al.*, "Observation and modal analysis of coupled-bunch longitudinal instabilities via a digital feedback control system," *Part. Accel.*, vol. 57, p. 175, 1997.

[3] S. Prabhakar, J. D. Fox, D. Teytelman, and A. Young, "Phase space tracking of coupled bunch instabilities," *Phys. Rev. ST Accel. Beams*, vol. 2, p. 084401, 1999.

[4] J. L. Laclare, "Bunched beam instabilities," in *Proceedings, High Energy Accelerators*, (Geneva), pp. 526–539, Birkhauser, 1980.

[5] A. W. Chao, *Physics of Collective Beam Instabilities in High Energy Accelerators.* New York: Wiley, 1993.

[6] S. Prabhakar, *New diagnostics and cures for coupled bunch instabilities.* PhD thesis, Stanford University, 2000. SLAC-R-554.

[7] M. Svandrlik *et al.*, "The cure of multibunch instabilities in ELETTRA," in *Proceedings of the 1995 Particle Accelerator Conference and International Conference on High - Energy Accelerators*, (Piscataway, NJ, USA), pp. 2762–2764, IEEE, 1996.

[8] M. Serio *et al.*, "Multibunch instabilities and cures," in *EPAC 96: proceedings*, (Bristol), pp. 148–152, IOP Publishing, 1996.

[9] M. Kwon *et al.*, "HOM suppression of the cavities with plunger tuner by the temperature tuning," in *Proceedings of the 1999 Particle Accelerator Conference*, (Piscataway, NJ, USA), pp. 899–901, IEEE, 1999.

[10] J. M. Byrd *et al.*, "Design of a higher harmonic RF system for the Advanced Light Source," *Nucl. Instrum. Meth.*, vol. A439, pp. 15–25, 2000.

[11] P. Corredoura, L. Sapozhnikov, and R. Tighe, "RF feedback development for the PEP-II B factory," in *EPAC 94: proceedings*, (River Edge, NJ, USA), pp. 1954–1956, World Scientific, 1994.

[12] D. Boussard and G. Lambert, "Reduction of the apparent impedance of wide band accelerating cavities by RF feedback," *IEEE Trans. Nucl. Sci.*, vol. 30, p. 2239, 1983.

[13] F. Pedersen, "RF cavity feedback," in *B factories: the state of the art in accelerators, detectors and physics: proceedings*, pp. 192–207, SLAC, 1992.

[14] F. Pedersen, "Beam loading effects in the CERN PS booster," *IEEE Trans. Nucl. Sci.*, vol. 22, pp. 1906–1909, 1975.

[15] F. Pedersen, "A novel RF cavity tuning feedback scheme for heavy beam loading," *IEEE Trans. Nucl. Sci.*, vol. 32, pp. 2138–2140, 1985.

[16] D. Boussard, "RF power requirements for a high intensity proton collider, parts 1 (chapters 1,2,3) and 2 (chapters 4,5,6)," in *1991 IEEE particle accelerator conference: accelerator science and technology: proceedings*, pp. 2447–2449, IEEE, 1991.

[17] R. Biscardi, S. L. Kramer, and G. Ramirez, "Bunch length control in the NSLS VUV ring," *Nucl. Instrum. Meth.*, vol. A366, pp. 26–30, 1995.

[18] M. Georgsson, A. Andersson, and M. Eriksson, "Landau cavities at MAX II," *Nucl. Instrum. Meth.*, vol. A416, pp. 465–474, 1998.

[19] J. M. Byrd *et al.*, "Commissioning of a higher harmonic RF system for the Advanced Light Source," *Nucl. Instrum. Meth.*, vol. A455, pp. 271–282, 2000.

[20] J. M. Byrd and M. Georgsson, "Lifetime increase using passive harmonic cavities in synchrotron light sources," *Phys. Rev. ST Accel. Beams*, vol. 4, p. 030701, 2001.

[21] A. Hofmann and S. Myers, "Beam dynamics in a double RF system," in *International Conference on High Energy Accelerators, 11th, Geneva, Switzerland, Jul 7-11, 1980. PROCEEDINGS.*, (Geneva), pp. 610–614, Birkhauser, 1980.

[22] D. Teytelman *et al.*, "Design and implementation of iir algorithms for control of longitudinal coupled-bunch instabilities," in *Beam Instrumentation Workshop : Proceedings*, (Melville), pp. 218–225, AIP, 2000.

[23] J. M. Byrd *et al.*, "Harmonic cavities and longitudinal beam stability in electron storage rings," in *Proceeding of the 2001 Particle Accelerator Conference*, pp. 380–384, IEEE, 2001.

[24] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems*. Reading, Massachusetts: Addison-Wesley, 1991.

[25] J. M. Steimel and D. McGinnis, "Damping in the fermilab booster," in *1993 IEEE Particle Accelerator Conference: proceedings*, (Piscataway, NJ, USA), pp. 2100–2102, IEEE, 1994.

[26] D. Boussard, "Control of coupled-bunch instabilities in high-current storage rings," in *1991 IEEE particle accelerator conference: accelerator science and technology: proceedings*, pp. 2537–2541, IEEE, 1991.

[27] A. Pei *et al.*, "Longitudinal multibunch feedback experiment with switched filter bank," in *Proceedings of the 1995 Particle Accelerator Conference and International Conference on High - Energy Accelerators*, (Piscataway, NJ, USA), pp. 3094–3096, IEEE, 1996.

[28] M. Tobiyama and E. Kikutani, "Development of a high speed digital signal process system for bunch-by-bunch feedback systems," *Phys. Rev. ST Accel. Beams*, vol. 3, p. 012801, 2000.

[29] H. Hindi *et al.*, "Measurement of multibunch transfer functions using time domain data and fourier analysis," in *Beam instrumentation workshop : proceedings*, (N.Y.), pp. 170–179, AIP, 1994.

[30] H. Hindi *et al.*, "A formal approach to the design of multibunch feedback systems: LQG controllers," in *EPAC 94: proceedings*, (River Edge, NJ, USA), pp. 1622–1624, World Scientific, 1994.

[31] H. Hindi, S. Prabhakar, J. Fox, and D. Teytelman, "Design and verification of controllers for longitudinal oscillations using optimal control theory and numerical simulation: Predictions for PEP-II," in *1997 IEEE Particle Accelerator Conference: Proceedings*, (Piscataway, NJ, USA), pp. 2362–2364, IEEE, 1998.

[32] D. A. Herrup, D. McGinnis, J. Steimel, and R. Tomlin, "Analog dampers in the fermilab booster," in *Proceedings of the 1995 Particle Accelerator Conference and International Conference on High - Energy Accelerators*, (Piscataway, NJ, USA), pp. 3010–3012, IEEE, 1996.

[33] K. Junck, J. Marriner, and D. McGinnis, "A study of the longitudinal coupled bunch instability in the fermilab main ring," in *Proceedings of the 1995 Particle Accelerator Conference and International Conference on High - Energy Accelerators*, (Piscataway, NJ, USA), pp. 3013–3015, IEEE, 1996.

[34] G. Oxoby *et al.*, "Bunch-by-bunch longitudinal feedback system for PEP-II," in *EPAC 94: proceedings*, (River Edge, NJ, USA), pp. 1616–1618, World Scientific, 1994.

[35] S. Khan and T. Knuth, "BESSY II feedback systems," in *Proceedings of the 1999 Particle Accelerator Conference*, (Piscataway, NJ, USA), pp. 1144–1146, IEEE, 1999.

[36] J. Fox *et al.*, "Programmable DSP-based multi-bunch feedback: Operational experience from six installations," in *Beam Instrumentation Workshop 2000*, (Melville, New York, USA), pp. 226–233, AIP, 2000.

[37] H. A. Hindi, D. Briggs, J. Fox, W. Hosseini, and A. Hutton, "Downsampled bunch by bunch feedback for PEP-II," in *B factories: the state of the art in accelerators, detectors and physics: proceedings*, pp. 216–219, SLAC, 1992.

[38] S. R. Smith, "Beam position monitor engineering," in *Beam Instrumentation Workshop: Proceedings*, (Woodbury, New York, USA), pp. 50–65, AIP, 1997.

[39] A. Young, J. Fox, and D. Teytelman, "VXI based multibunch detector and QPSK modulator for the PEP-II/ALS/DAPHNE longitudinal feedback system," in *1997 IEEE Particle Accelerator Conference: Proceedings*, (Piscataway, NJ, USA), pp. 2368–2370, IEEE, 1998.

[40] A. Gallo, A. Ghigo, F. Marcellini, M. Migliorati, L. Palumbo, and M. Serio, "Simulations of the bunch-by-bunch feedback operation with a broadband RF cavity as longitudinal kicker," Tech. Rep. G-31, DAΦNE, 1995.

[41] J. M. Byrd and S. De Santis, "Longitudinal injection transients in an electron storage ring," *Phys. Rev. ST Accel. Beams*, vol. 4, p. 024401, 2001.

[42] *Optimization Toolbox User's Guide.* Natick, MA, USA: The MathWorks, Inc., 2000.

[43] A.G.Frodesen, O.Skjeggestad, and H. Tøfte, *Probability and Statistics in Particle Physics.* Bergen, Norway: Universitetsforlaget, 1979.

[44] L. Palumbo and V. G. Vaccaro, "Wake fields measurements," in *Frontiers of Particle Beams: Observation, Diagnosis and Correction*, (Berlin), pp. 312–354, Springer-Verlag, 1989.

[45] J. N. Corlett and J. M. Byrd, "Measurement and computation of the higher order modes of the ALS 500-MHz accelerating cavities," in *1993 IEEE Particle Accelerator Conference: proceedings*, (Piscataway, NJ, USA), pp. 3408–3410, IEEE, 1994.

[46] M. A. Allen, J. M. Paterson, J. R. Rees, and P. B. Wilson, "Beam energy loss to parasitic modes in SPEAR II," *IEEE Trans. Nucl. Sci.*, vol. 22, pp. 1838–1842, 1975.

[47] B. Podobedov and R. Siemann, "New apparatus for precise synchronous phase shift measurements in storage rings," *Phys. Rev. ST Accel. Beams*, vol. 1, p. 072801, 1998.

[48] B. W. Zotter and S. A. Kheifets, *Impedances and Wakes in High-Energy Particle Accelerators.* Singapore: World Scientific, 1998.

[49] W. Anders, P. Kuske, and T. Westphal, "The longitudinal impedance of BESSY determinated by measurements of the synchronous phase shift," in *High Energy Accelerators (HEACC 92): proceedings*, (Singapore), pp. 1121–1123, World Scientific, 1993.

[50] D. Möhl and A. Sessler, "The use of RF-knockout for determination of the characteristics of the transverse coherent instability of an intense beam," in *Proceedings, 8th International Conference on High Energy Accelerators*, (Geneva), pp. 334–337, CERN, 1971.

[51] A. Hofmann and B. Zotter, "Measurement of beam stability and coupling impedance by RF excitation," *IEEE Trans. Nucl. Sci.*, vol. 24, p. 1487, 1977.

[52] A. Temnykh, D. Hartill, S. Belomestnykh, and R. Kaplan, "The HOM study of CESR RF cavities using single circulating bunch," Tech. Rep. CBN 97-5, Cornell U., Ithaca, NY, USA, 1997.

[53] S. Prabhakar, J. Fox, H. Hindi, D. Teytelman, and A. Young, "Calculation of impedance from multibunch synchronous phases: Theory and experimental results," in *Proceedings of the Sixth European Particle Accelerator Conference*, (Bristol), pp. 996–998, IOP Publishing, 1998.

[54] J. M. Byrd, W. Barry, J. N. Corlett, J. Fox, and D. Teytelman, "Observations of collective effects at the Advanced Light Source," in *KEK-Proceedings-96-6*, (Ibaraki, Japan), pp. 527–545, National Lab. High Energy Physics, 1996. Presented at International Workshop on Collective Effects and Impedance for B factories, Tsukuba, Japan, 12-17 Jun, 1995.

[55] J. M. Byrd and J. Corlett, "Spectral characterization of longitudinal coupled-bunch instabilities at the Advanced Light Source," *Part. Accel.*, vol. 51, p. 29, 1995.

[56] R. J. Bakker, "Status and commissioning-results of BESSY-II," in *Proceedings of the 1999 Particle Accelerator Conference*, (Piscataway, NJ, USA), pp. 197–199, IEEE, 1999.

[57] J. N. Corlett, J. Johnson, G. Lambertson, and F. Voelker, "Longitudinal and transverse feedback kickers for the als," in *EPAC 94: proceedings*, (River Edge, NJ, USA), pp. 1625–1627, World Scientific, 1994.

[58] D. Teytelman and J. Fox, "Setup of pep-ii longitudinal feedback systems for even/odd bunch spacings," in *Beam instrumentation workshop 2002 : proceedings*, (Melville), pp. 474–482, AIP, 2002.

[59] Y. R. Chiang and M. G. Safonov, *Robust Control Toolbox User's Guide.* Natick, MA, USA: The MathWorks, Inc., 1998.

[60] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing.* New Jersey: Prentice-Hall, 1989.

[61] L. B. Jackson, "Roundoff-noise analysis for fixed-point digital filters realized in cascade on parallel form," *IEEE Trans. Audio Electroacoust.*, vol. AU-18, no. 2, pp. 107–122, 1970.

[62] A. G. Deczky, "Synthesis of recursive digital filters using the minimum p-error criterion," *IEEE Trans. Audio Electroacoust.*, vol. AU-20, no. 4, pp. 257–263, 1972.

[63] J.-P. Thiran, "Recursive digital filters with maximally flat group delay," *IEEE Trans. Circuit Theory*, vol. CT-18, no. 6, pp. 659–664, 1971.

[64] X. Chen and T.W.Parks, "Design of IIR filters in the complex domain," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 38, no. 6, pp. 910–920, 1990.

[65] G. Cortelazzo, C. De Simone, and G. Mian, "IIR transfer function design with a double-objective differential correction algorithm," *IEEE Trans. Circuits Syst.*, vol. 38, no. 7, pp. 791–796, 1991.

[66] A. Tarczynski and G. Cain, "A new algorithm for designing near-optimal chebyshev IIR and FIR filters," in *38th Midwest Symposium on Circuits and Systems. Proceedings*, (New York, NY, USA), pp. 584–587, IEEE, 1996.

[67] H. Hindi *et al.*, "Analysis of DSP-based longitudinal feedback system: Trials at SPEAR and ALS," in *1993 IEEE Particle Accelerator Conference: proceedings*, (Piscataway, NJ, USA), pp. 2352–2354, IEEE, 1994.

[68] K. Robinson, "Stability of beam in radiofrequency system," Tech. Rep. CEAL-1010, CEA, Cambridge, Massachusetts, 1964.

[69] J. M. Byrd, "Effects of phase noise in heavily beam loaded storage rings," in *Proceedings of the 1999 Particle Accelerator Conference*, (Piscataway, NJ, USA), pp. 1806–1808, IEEE, 1999.

[70] A. Gallo. Private communication.

[71] C. G. M. Lonza, D. Bulfone, "Digital processing electronics for the ELETTRA transverse multi-bunch feedback system," in *Accelerator and Large Experimental Physics Control Systems*, (Trieste, Italy), pp. 255–257, Comitato Conferenze ELETTRA, 1999.

[72] VITA, *Front Panel Data Port Specifications, VITA 17-199x, Rev. 1.7*, November 1998.

[73] D. Teytelman *et al.*, “Feedback control and beam diagnostic algorithms for a multiprocessor DSP system,” in *Beam Instrumentation Workshop: Proceedings*, (Woodbury, New York, USA), pp. 390–397, AIP, 1997.

[74] D. Teytelman *et al.*, “Accelerator diagnostic techniques using time-domain data from a bunch-by-bunch longitudinal feedback system,” in *DIPAC 97: beam diagnostics and instrumentation for particle accelerators: Proceedings*, (Frascati (Rome) Italy), pp. 3–5, Laboratori Nazionali di Frascati: SIS–Pubblicazioni, 1997.