SLAC-R-554

New Diagnostics and Cures for Coupled-Bunch Instabilities

Shyam Prabhakar

Stanford Linear Accelerator Center
Stanford University
Stanford, CA 94309

SLAC-Report-554

Prepared for the Department of Energy
under contract number DE-AC03-76SF00515

Printed in the United States of America. Available from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Road, Springfield, VA 22161.

NEW DIAGNOSTICS AND CURES FOR COUPLED-BUNCH INSTABILITIES

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

By
Shyam Prabhakar
August 2001

Abstract

Electromagnetic interaction between a charged particle beam and its surroundings causes collective instabilities, which must be controlled if the new light sources and colliders are to meet their design goals. Control requires a combination of passive damping and fast active feedback on an unprecedented technological scale. Efficient instability diagnosis techniques are also needed for machines with large numbers of bunches. This thesis describes new methods of measuring and analyzing coupled-bunch instabilities in circular accelerators, and demonstrates the existence of a new cure.

A new technique is demonstrated for simultaneous measurement of growth rates, damping rates and coherent tune shifts of all unstable coupled-bunch eigenmodes from a single 10–25-ms transient snapshot of beam motion. The technique has been used to locate and quantify beam impedance resonances at PEP-II, ALS and SPEAR. This method is faster than existing spectral scan methods by at least an order of magnitude, and has the added advantage of revealing coupled-bunch dynamics in the linear small-signal regime. A method is also presented for estimating beam impedance from multi-bunch fill shape and synchronous phase measurements.

Phase space tracking of multi-bunch instabilities is introduced as a "complete instability diagnostic." Digit ised multi-bunch data is analyzed offline, to estimate the phase space trajectories of bunches and modes. Availability of phase space trajectories is shown to open up a variety of possibilities, including measurement of reactive impedance, and diagnosis of the fast beam-ion instability.

Knowledge gained from longitudinal measurements (all made using a digital longitudinal feedback system) has been used to optimise cavity temperatures, tuner positions and feedback parameters, and also to identify sources of beam noise at the three machines.

A matrix-based method is presented for analyzing the beneficial effect of bunch-to-bunch tune variation on instability growth rates. The method is applicable to the calculation

of instability eigenvalues in machines with more than one unstable coupled-bunch mode. This technique is useful in studying machines like PEP-II and KEK-B, which do not lend themselves to tune spread analysis by conventional methods.

A similar mathematical formalism is used to understand the dynamics of azimuthally asymmetric beams. Simple formulae are derived for asymmetry-induced growth rate reduction. "Optimal" fill shapes based on these ideas have been experimentally verified at the ALS and SPEAR, where the longitudinal instability threshold has been raised by factors of six and two, respectively. Thus we have a new, zero-cost, easily implementable cure for coupled-bunch instabilities.

Acknowledgements

As a student who entered the field of beam physics with only the slightest awareness of particle accelerators, I have many people to thank for the completion of my thesis work. I will do my best to acknowledge all of them here, though I must apologise in advance for the omissions that such efforts often entail.

The largest amount of guidance and steering has come from my advisor John Fox. It was he who initially suggested that I analyse experimental data on longitudinal beam dynamics. The opportunity to analyse such data and make measurements based on the longitudinal feedback system was a unique one, and I learnt a lot of what I know of beam dynamics by trying to make sense of the experimental results. John had a large hand in my academic socialisation: he impressed upon me the importance of publishing my work promptly, giving talks, and writing code that other physicists and engineers could easily use. It is not possible to list all of his contributions to this thesis. Suffice it to say that all of the measurements, concepts, and techniques discussed here have benefited greatly from his creative ideas, constructive criticism, and skillful engineering.

Starting with Leonid Sapozhnikov, who shared his knowledge of DSPs and electronic circuits with me, I have benefited a lot from my fellow graduate students. Haitham Hindi, who shared an office with me for quite a while, gave me my nodding familiarity with optimal control theory, and worked with me to solve numerous feedback and instability problems. We enjoyed long conversations on beam dynamics and on life, in equal measure. The matrix-based formalism in the final chapters is a direct consequence of his excellent informal lectures on paper we rescued from the recycling bin. Dmitry Teytelman has also been an invaluable resource in designing new experiments. I have scarcely made a measurement that he did not contribute to; either directly, or in his capacity as our harware and software guru. Without his wizardly skills, many of the diagnostics presented here would have remained pipe dreams. He also contributed important ideas to the analysis of feedback system data.

I have been lucky to have an office in the same building as Alex Chao and Sam Heifets, who bore the brunt of my frequent questions on beam dynamics. A few minutes in conversation with them has always been worth many hours, or even days, of reading papers and textbooks. Albert Hofmann of CERN has also been generous in answering questions and discussing ideas, on his visits to SLAC. In addition, he has helped by teaching his first-class course on beam experiments, and by collaborating on some of the early SPEAR experiments. Gennady Stupakov strengthened my understanding of the fast beam-ion instability with his precise and well-considered elaborations on the subject.

Ivan Linscott of Stanford University devised clever data analysis methods that were precursors to the current crop of algorithms. I am also indebted to him for showing me a paper on Josephson junctions, which inspired the development of the uneven-fill theory.

Others who contributed by vetting ideas and supplying insightful comments and advice include: Martin Lee, Mike Zisman, Uli Wienands, Gennady Stupakov, David Whittum, Helmut Wiedemann, Ron Ruth, Karl Bane, and Kathy Thompson of SLAC, Boris Podobedov and Chun-xi Wang, who were fellow students at SLAC, Glen Lambertson, John Byrd, and R. Rimmer of LBNL, Bruno Zotter and Fleming Pederson of CERN, Mario Serio and Mikhail Zobov of INFN-LNF, Bob Hettel, Jeff Corbett, and Andrei Terebilo of SSRL, and Stefan Bilbao of Stanford University.

Paul Corredoura and Rich Tighe, who are responsible for the PEP-II RF feedback system, collaborated on many of the SLAC measurements. They also helped by teaching me about beam loading and RF feedback. Andrew Young of SLAC took part in many of the measurements, and oversaw the maintenance of the PEP-II feedback power amplifiers.

The experiments performed at the various machines were team efforts. Participants in the measurements include: Mario Serio and Alessandro Drago of INFN-LNF, Uli Wienands, Mike Zisman, Michiko Minty, and Bill Ross of SLAC, Greg Stover, John Byrd, Walter Barry, and John Corlett of LBNL, and Jim Sebek, Bob Hettel, and Jim Safranek of SSRL.

I am particularly grateful to Mario and Alessandro for inviting me for a stint in Frascati. I gained valuable experimental experience by working with them, and also had the time of my life as a tourist.

I would like to thank the operations groups at PEP-II, SPEAR and the ALS, for their cooperation and support in delivering the beam conditions required for the experiments.

Finally, I would like to thank my parents for making me what I am today, and my wife Anchal for her wonderful support.

Contents

Abstract iv

Acknowledgements vi

1 Introduction 1

1.1 Coupled-Bunch Instabilities 2

1.2 Overview of Contributions 2

1.2.1 Bunch Currents and Synchronous Phases 3

1.2.2 Measuring Instability Growth and Damping Rates 4

1.2.3 Phase Space Tracking: A Complete Diagnostic 5

1.2.4 A Matrix Formalism for Landau Damping 5

1.2.5 Optimal Uneven Fills: A New Cure 5

2 Basic Concepts, Feedback 7

2.1 Storage Rings, Single-Particle Dynamics 7

2.2 Wake Fields and Impedance 11

2.3 Coupled-Bunch Dipole Instabilities 15

2.4 Existing Cures 20

2.4.1 Impedance Minimization 20

2.4.2 Landau Damping 21

2.4.3 Active Feedback 22

2.5 Digital Bunch-by-Bunch Longitudinal Feedback System 25

2.5.1 General Approach, Downsampling 25

2.5.2 Front End 26

2.5.3 Filter Algorithm 27

2.5.4 Back End 28

3 Bunch Currents and Synchronous Phases **29**

3.1 Single-Bunch Synchronous Phase 30

3.2 Multi-bunch Synchronous Phase 31

3.3 Experimental results 33

3.3.1 Time Domain 34

3.3.2 Transfer Function 36

3.3.3 Coupled-Bunch Instability 38

3.3.4 Matching of PEP-II Gap transients 39

3.4 Summary 40

4 Measuring Instability Growth and Damping Rates **41**

4.1 ALS Measurements 42

4.1.1 Signal Processing 43

4.1.2 Longitudinal 44

4.1.3 Transverse 52

4.2 PEP-II Results 55

4.3 SPEAR Results 61

5 Phase Space Tracking: A Complete Diagnostic **63**

5.1 Introduction 63

5.2 Signal Processing 64

5.3 Phase space tracking of bunch trains 67

5.4 Phase space tracking of coupled-bunch modes 71

5.5 Measuring Reactive Feedback 75

5.6 Summary 77

6 A Matrix Formalism for Landau Damping **79**

6.1 Equivalent State Matrix 80

6.1.1 Single Unstable Mode 81

6.1.2 Multiple Unstable Modes 82

6.2 Application to PEP-II 83

6.3 Summary 84

7 Optimal Uneven Fills: A New Cure **86**
7.1 Derivation of Coupling Matrix 87
7.2 Modulation Coupling 90
7.2.1 Simplest Case 90
7.2.2 Algorithm for Optimising Fill Shape 91
7.2.3 Sample Calculation: PEP-II 92
7.2.4 Intuitive Explanation for Modulation Coupling 93
7.3 Landau Damping 93
7.3.1 Interbunch Tune Spread Formula 94
7.3.2 Narrowband Impedance Spectrum 94
7.3.3 Broadband Impedance Spectrum 96
7.4 Some Special Cases 96
7.5 Brief Summary of the Theory 97
7.6 Measurements 97
7.6.1 ALS Modulation Coupling 98
7.6.2 ALS Landau Damping 100
7.6.3 SPEAR 102
7.6.4 PEP-II Landau Damping 103
7.7 Conclusion 105

8 Conclusion, Ideas for Further Study **106**

A Computer Programs for Data Analysis **109**
A.1 synchp.m 110
A.2 Modes.m 111
A.2.1 checkfit.m 112
A.2.2 PhasSpac.m 113
A.3 sideband.m 114
A.3.1 Memory Limitations 115
A.4 Program Listings 116

Bibliography **164**

List of Tables

4.1 ALS Parameters 43
4.2 PEP-II Parameters 55
4.3 SPEAR Parameters 61

List of Figures

2.1 Schematic diagram of a storage ring. 8

2.2 Example of energy gain from an RF system as a function of the starting time t of a revolution. 9

2.3 Example longitudinal impedance and wake function of a resonant cavity mode (RLC band pass filter) with $Q = 10$ and $R_s = 1$. 14

2.4 The locus of possible values of λ, due to a fixed high-Q resonator. 20

2.5 Block diagram of longitudinal bunch-by-bunch feedback system. 26

3.1 Averaged bunch signal spectrum up to 2 kHz, for a 291-bunch fill with a total current of 122 mA. Two circles mark the 720-Hz line, onto which bunch signals are projected for current monitoring. 34

3.2 (a) Bunch current measurement, 291-bunch fill, 122 mA total current. Note the step discontinuity at the 175th filled bucket. (b) Raw mean bunch phase, ADC counts. (c) Synchronous phase variation around the ring. The discontinuity in the fill causes ϕ_s to oscillate with a peak-to-peak amplitude of 3° at the RF frequency. 35

3.3 (a) Bunch current measurement, 291-bunch fill, 96 mA total current. The discontinuity in the fill looks like an impulse at low frequencies. (b) Raw mean bunch phase, ADC units (c) Synchronous phase variation around the ring. The "impulse response" contains about three oscillations in one revolution period. 36

3.4 Estimate of Z_n obtained by averaging transfer functions from four consecutive data sets. The impedance is in the $M\Omega$ range due to the fundamental resonance of the parked cavities. The cavity resonant frequencies seem to be closer to $3f_o$ than $2f_o$. 37

3.5 Beam pseudospectrum for a 291-bunch 84-mA fill, taken a few days before the data displayed in previous figures. The pseudospectrum shows that mode 3 is unstable, with a steady state amplitude of 2° at the RF frequency. . . . 38
3.6 Synchronous phase transients in the PEP-II HER and LER, measured during collision (10% gap). 39

4.1 ALS grow-damp measurement. a) Bunch oscillation envelopes. b) Modal amplitudes. c,d) Exponential fits before break point yield growth rates without feedback. e,f) Exponential fits after break point yield feedback-induced damping rates. 45
4.2 Example of external transient technique for measuring (negative) growth rates of naturally stable modes with and without feedback. Mode 61 is externally excited and then allowed to decay naturally until $t = 7$ ms, at which point feedback is turned on and the mode is rapidly damped. 48
4.3 Log plot of experimentally measured and predicted growth rates as a function of modal frequency, with error bars around measured points. The predicted rates are based on the radiation damping rate and measurements of the impedance of a model RF cavity. 49
4.4 Longitudinal growth rates vs. cavity temperature at the ALS, $I_o = 100$ mA. 50
4.5 (a) Magnitude of high-resolution DFT of growing transient in Fig. 4.1, showing all the spectral components from 0 to 250 MHz. (b) Real part of Z^{eff}. . 51
4.6 90-kHz sections of the high-resolution spectrum in Fig. 4.5(a), showing an unstable upper sideband and a damped lower sideband at $204f_o$ and also at $233f_o$. 52
4.7 Block diagram of experimental setup for "grow-damp" measurements of transverse motion. Longitudinal feedback system used to gate the transverse feedback system, and record transverse motion. 53
4.8 ALS horizontal grow-damp at $I_o = 94$ mA. Resistive wall instability is the strongest. Feedback gain margin is comfortable. 54
4.9 PEP-II longitudinal growth rates from RF cavity HOMs at nominal bunch spacing and tune, design current. (a) HER, 1 A, 20 cavities. (b) LER, 2.25 A, 4 cavities. 56

4.10 PEP-II HER grow-damp measurement at I_o = 361 mA. Feedback is positive during the "grow" portion of this grow-damp, and negative during the "damp" portion. Modes from 760 to 790 show significant growth and damping. 58
4.11 PEP-II HER beam pseudospectrum. Positive feedback excites a band of roughly 60 sidebands, in the vicinity of the 770th revolution harmonic (105 MHz). The aliased impedance of the largest cavity HOM is superimposed. . 59
4.12 PEP-II LER open-loop longitudinal growth rates versus I_o, with 90% of the ring filled. 60
4.13 Grow-damp measurement of longitudinal instabilities at SPEAR, I_o = 29 mA. Positive feedback is used during the "grow" portion, to uncover the modes closest to instability. 62

5.1 Growing vertical instability transient in a 150-bunch train in the PEP-II HER at I_o = 52 mA. Feedback is switched off at $t = 0$. The oscillation amplitude a^k of each bunch k grows exponentially over the 20 ms interval, with trailing bunches growing to larger amplitudes than leading bunches. 68
5.2 Color-coded representation of magnitude spectra of the 150 bunches in previous figure. Peaks of Fourier transform of bunch transients lie at the same tune, indicating that tune spread across the train is $\leq$ the frequency resolution (50 Hz). 69
5.3 Phase space angle differentials ($\phi_n^k - \phi_n^{150}$) for all 150 bunches. Differentials for first few bunches are noisy due to smaller oscillation amplitudes. Differential angles are almost constant over 20 ms, indicating that the bunches oscillate coherently. 70
5.4 (a) Relative tunes of bunches 46 to 150, calculated using linear fits to the phase space angle differentials in the previous figure. (b) Growth rates of the same bunches, calculated using exponential fits to the magnitude transients a^k. 70
5.5 Linear evolution of modal phase space angles Φ^{204} and Φ^{233} at the ALS (longitudinal instabilities, I_o = 157 mA). *dashdot*: Φ^{204}, even fill. *dotted*: Φ^{233}, even fill. *solid*: Φ^{204}, square-wave fill. *dashed*: Φ^{233}, square-wave fill. Square-wave fill couples the two frequencies and creates a mixed eigenmode, so that U^{204} and U^{233} are phase-locked. 72

5.6 Growth in magnitude of two sets of projections of a longitudinal instability transient in PEP-II LER. Uneven fill, $I_o = 703$ mA. (a) A^{787} to A^{794}, beating is evidence of at least two uneven-fill eigenmodes in this frequency range. (b) A^{807} to A^{815}, quasi-exponential growth indicates that this set of sidebands oscillates coherently as a single eigenmode. 73

5.7 (a) Average of modal phase space magnitudes A^m; m = 775, 776, ... 815 (same data as previous figure). (b) Average coherent tunes, calculated using linear fits to Φ^m; m = 775, 776, ... 815. Fitted tunes show negligible variation above m = 798, implying that the band of projections on the right side of the dotted line contains only one eigenmode. 74

5.8 Modal phase space trajectories of growing PEP-II LER longitudinal instability (same data as previous figure). (a) Representative selection of "modes" above m = 798. Expanding spirals about the origin indicate a single uneven-fill eigenmode. (b) Trajectories of a few "modes" below m = 798, whose magnitude transients show beating. 76

5.9 Evolution of modal phase space angles Φ^{204} and Φ^{233}, during a 238 mA ALS grow-damp. Change in slope at break point indicates that feedback has a reactive component. 77

6.1 Eigenvalues of longitudinal coupled-bunch modes in PEP-II with and without bunch-to-bunch tune spread, 1746 bunches: (a) HER, 1 A, 600 rad/s tune spread (b) LER, 2.25 A, 600 rad/s tune spread. 84

7.1 Illustration of fill optimization. N = 1000, I_o = 500 mA. Solid lines: 50% fill and 25% fill maximise C_4 for i_{max} = 1 mA, 2 mA. Dashdot: Reference sinusoid at $4\omega_o$. 91

7.2 PEP-II HER cavity-induced modal growth rates vs. mode frequency ($l\omega_o + \omega_s$) at I_o = 1 A for: a) Even fill at nominal 4.2 ns spacing (feedback required). b) Even fill at 6 × 4.2 ns spacing (stabilised by modulation coupling). . . . 92

7.3 Graphical look-up table for fill-induced damping of eigenvalue of unstable longitudinal EFEM n as C_n is increased from 0 (100% of ring filled) to 0.5 (61% filled). Dashed lines: Evolution of λ_n from a few even-fill starting points. 95

7.4 Uneven fill coupling EFEMs 204 and 233 by maximising C_{29} for $x = 0.5$. Solid line: ideal; 'x's: measured. a) Typical section of bunch current profile. b) Fill spectrum. 98
7.5 Projection onto EFEMs of measured ALS signal soon after feedback is stopped. a) Even fill, 158 mA. EFEMs 204 and 233 grow independently. b) Fill shown in Fig. 1, 156 mA. The composite eigenmode grows faster due to modulation coupling of unstable EFEMs 204 and 233. 99
7.6 Uneven "Landau fill" that maximises C_{95}, (*i.e.* C_{233}) for $I_o = 175$ mA, $x = 0.47$. 100
7.7 Measured growth of ALS longitudinal instabilities in EFEM basis. a) Even fill, 172 mA. b) "Landau fill" (see Fig. 3), 175 mA. Modes barely grow. . . 101
7.8 a) Measured fill spectrum of regular SPEAR fill. b) Calculated spectrum of SPEAR "Landau fill". Peak at C_{50} generates tune spread, doubles instability threshold. 103
7.9 a) Solid line: Meaured PEP-II HER longitudinal tune spread due to 7% gap, $I_o = 605$ mA. Dashed lines: 3 representative eigenvectors (calculation). b) Calculated eigenvalues of 605 mA beam. '+'s: No tune spread. 'o's: Tune spread included. Dashdot line: Radiation damping. 104

Chapter 1

Introduction

Modern synchrotron light sources and circular colliders require the storage of high-current charged particle beams for the attainment of their design goals. Colliders require intense beams for frequent production of collision events, and radiation sources generate more photons when they accelerate larger numbers of charged particles.

The storage of high-current charged particle beams presents a variety of challenges, such as increased power requirements and increased heating of the vacuum chamber surrounding the beam. The most serious consequence of high beam intensities is the possibility of collective instabilities, which result from electromagnetic interaction between the beam and its surroundings.

Collective instabilities can be understood as follows: The beam interacts with its environment to produce electromagnetic fields known as wake fields, which change the beam trajectory. The perturbed trajectory modifies the wake field, which in turn changes the beam trajectory yet again, and so on. If the beam intensity is below a certain threshold, this cycle of interaction between the beam and its wake field reaches a stable steady state. However, if successive disturbances enhance each other enough to steadily increase in magnitude, the trajectory becomes unstable. If this happens, the beam is pushed into nonlinear large-amplitude oscillations, or is lost by collision with the vacuum chamber. The instability mechanism is a collective one, since it involves cooperative particle motion and the collective electromagnetic field.

Collective instabilities are one of the main factors that limit the performance of high energy accelerators. Even if nonlinear mechanisms avert beam loss, collision rates and radiation quality are adversely affected by large-amplitude beam oscillations relative to the

ideal trajectory. For this reason, various collective instability mechanisms have been studied theoretically and experimentally over the last four decades, and a number of techniques have been developed to allow stable high-current operation.

1.1 Coupled-Bunch Instabilities

Modern circular accelerators achieve high total currents by distributing the current among large numbers of circulating bunches, so as to avoid single-bunch instabilities, improve the beam lifetime, and reduce two-beam effects (in the case of colliders). In such cases, unstable coupling between bunches through long-range wake fields is often the main current-limiting factor.

The instability of longitudinal coupled-bunch oscillations in circular accelerators was first studied by K. Robinson in 1964 [1]. The fundamental resonance of radio frequency accelerating cavities was found to cause beam bunches to oscillate coherently in energy under certain conditions. Transverse coupled-bunch instabilities driven by the resistance of vacuum chamber walls were identified by E. Courant and A. Sessler in 1966 [2]. The concept of a beam impedance has since been developed, and more general approaches have been used to study longitudinal and transverse coupled-bunch instabilities excited by long-range wake fields of almost any origin [3, 4, 5]. The approach typically involves working out the eigenmodes and eigenvalues of the system in the neighbourhood of equilibrium, where the dynamics are linear. The objective is to calculate instability growth rates and threshold currents in the presence of natural damping mechanisms and active feedback.

Recently, new kinds of transverse instabilities, excited by transient charged particles in the vicinity of the beam, have been discovered [6, 7, 8].

1.2 Overview of Contributions

New generations of accelerators pose a formidable challenge to physicists and engineers who study and combat coupled-bunch instabilities. The existence of unstable eigenmodes at hundreds or thousands of possible frequencies makes instability diagnosis by conventional techniques difficult, if not completely impractical. There is a real need for improved measurement and analysis techniques, which must be followed by imaginative ways of damping coherent motion.

This thesis presents new experimental and theoretical methods for measuring and predicting coupled-bunch instabilities, and demonstrates the efficacy of a new cure. Although the focus is on longitudinal rigid-bunch motion, some of the theoretical results are directly applicable to longitudinal and transverse instabilities of any order. Measurements of transverse dipole instabilities are also presented.

A programmable longitudinal feedback (LFB) system has been the primary beam manipulation and data collection tool at the ALS[1] [9], PEP-II[2] [10], and SPEAR[3] [11, 12], because of its ability to record the oscillations of all stored bunches while controlling beam motion.

The following subsections constitute an outline of the contributions of this thesis, in the order of the chapters. Brief statements are included about historical background and comparitive merits of older approaches. Those who are unfamiliar with the material are encouraged to read Ch. 2, which introduces the basic ideas and equations that underlie subsequent treatments of beam diagnostics, coupled-bunch instabilities and active feedback.

1.2.1 Bunch Currents and Synchronous Phases

The shape of a storage ring fill (bunch current vs. bunch number) has a significant effect on coupled-bunch instabilities. Chapter 3 describes a technique for extracting bunch currents from LFB system data. The method makes use of the fact that recorded signals are proportional to the product of bunch current and longitudinal phase error. Results are shown from PEP-II.

Knowledge of bunch currents allows the calculation of bunch synchronous phases from LFB data. A simple formula is derived, for calculating the beam impedance from multi-bunch currents and synchronous phases. The single-bunch variant of this method is already well known. The multi-bunch approach yields more spectral information, since the wake function is sampled at a higher frequency. This method has been used to identify inaccurately parked idle cavities as the cause of a mode 3 longitudinal instability in the PEP-II HER (High Energy Ring). Synchronous phase measurements are also important at colliders like PEP-II, because of the need to match beam loading transients in the two rings.

[1] Lawrence Berkeley National Laboratory Advanced Light Source
[2] Stanford Linear Accelerator Center (SLAC) B Factory
[3] Stanford Synchrotron Radiation Laboratory Light Source

1.2.2 Measuring Instability Growth and Damping Rates

A new "grow-damp" technique is shown to efficiently solve the problem of diagnosing coupled-bunch instabilities in large accelerators. Growth rates are directly measured in the linear small-signal regime by turning off feedback and recording the oscillations of all bunches in the ring as the unstable modes grow out of the noise floor. Simultaneous characterisation of all unstable modes from a single 10–25-ms transient is a unique feature of such a measurement.

Grow-damp measurements based on LFB system data have been used to identify coupled-bunch instabilities and measure their amplitudes and growth rates under a variety of beam and cavity conditions at the ALS, PEP-II and SPEAR. The observations are compared with estimates of the cavity impedance. This method has been used find the ALS cavity temperatures most conducive to longitudinal stability, and to calculate the effective beam impedance at SPEAR. The grow-damp technique has also been used to measure transverse coupled-bunch instabilities. The impedance estimates from such measurements have been used to create "optimal" fill shapes at the ALS and SPEAR [13, 14].

Damping rates of stable modes are measured by externally exciting them and observing their natural decay rates.

LFB system data are also used to reconstruct beam "pseudospectra," *i.e.* beam spectra without revolution harmonics, from 10–25-ms pieces of data. These pseudospectra cover the coupled-bunch eigenmode frequency range with a resolution of 100–40 Hz. A conventional heterodyned spectrum analyzer would take more than half an hour to measure a beam spectrum of comparable width and resolution.

Data analysis is performed using MATLAB programs to filter raw data, remove sampling time offsets, calculate bunch and mode amplitudes as a function of time and fit exponentials to modal transients. Other MATLAB tools look for correlations between even-fill eigenmodes, check the quality of exponential fits, and calculate beam spectra from the data.

The "damp" portion of the grow-damp transients is used to measure feedback-induced damping rates of all unstable modes. Damping rate measurements have been made at various currents at ALS, SPEAR and PEP-II. Grow-damp measurements constitute a quick and comprehensive test of feedback system performance. They have been used to optimise feedback gain and phase and detect amplifier irregularities.

1.2.3 Phase Space Tracking: A Complete Diagnostic

Chapter 5 describes an instability diagnostic that exploits the information contained in the angular evolution of coupled-bunch oscillations in phase space. In addition to enabling measurement of coherent tunes and bunch tunes with accuracy of a few Hz, phase space tracking allows new kinds of comparisons between instability theory and experiment.

Phase space evolution of bunches participating in a low-threshold vertical instability in the PEP-II High Energy Ring is used to distinguish between the fast beam-ion instability and conventional instabilities. Tracking of longitudinal motion at the ALS and PEP-II is used to measure coherent tunes and gain new insights into uneven-fill instabilities.

This technique has also been used to measure the reactive component of the beam impedance, and to measure and minimise the reactive component of longitudinal feedback.

1.2.4 A Matrix Formalism for Landau Damping

Existing methods of calculating the effect of bunch-to-bunch tune variation on longitudinal coupled-bunch instability growth rates only address the problem of a single unstable eigenmode. In addition, they do not directly yield the growth rate reduction afforded by a given tune distribution.

A more general approach is presented in Ch. 6, that involves computing the eigenvalues of a state matrix of reduced size. The method is applied to the analysis of PEP-II longitudinal coupled-bunch modes, a large number of which are unstable in the absence of feedback.

This technique can be used just as easily to analyse the effects of transverse bunch tune variation.

1.2.5 Optimal Uneven Fills: A New Cure

As mentioned earlier, coupled-bunch instabilities in unevenly filled rings have not been very well understood so far. R. Kohaupt has calculated an upper bound for the instability growth rate in rings with a square-wave fill (even fill with one continuous group of bunches removed) [15]. K. Thompson has derived expressions for the elements of the bunch coupling matrix whose eigenvalues determine beam stability [16].

S. Bogacz calculated the longitudinal tune spread introduced by square-wave fills in 1994, and suggested that parasitic impedance resonances be tuned to increase the tune spread effect [17]. The problem with such an approach is that it would simultaneously

increase the bunch coupling that drives instability.

In many cases, these analyses fall short of a general, practically useful treatment of coupled-bunch instabilities in unevenly filled (azimuthally asymmetric) rings. In practice, however, empirically determined azimuthal asymmetries have been seen to have a significant effect on instability growth rates.

Chapter 7 describes the physics of two mechanisms for the influence of uneven fills on coupled-bunch instability growth rates. The resultant damping is calculated approximately using back-of-the envelope formulae and a graphical look-up table. The first mechanism is the one studied by Bogacz, generalised to the case of an arbitrary fill shape. The second mechanism involves coupling of stable and unstable even-fill eigenmodes to each other by means of spatial harmonics in the fill shape. A heuristic algorithm is presented for designing easily implementable "optimal" fill shapes, that exploit tune spreads and mode coupling to minimise instability growth rates.

Results are shown from the ALS and SPEAR, where specially tailored uneven fills have raised the longitudinal instability threshold by factors of six and two, respectively. A theoretical calculation of Landau damping from the PEP-II HER gap transient is also shown to be in good agreement with experimental data.

Chapter 2

Basic Concepts, Feedback

This chapter introduces the basic ideas and equations that underlie subsequent treatments of beam diagnostics, coupled-bunch instabilities and active feedback. The relevant elements of a circular accelerator or storage ring are listed, and the concept of phase focusing is explained. Individual charged particles are shown to behave like damped harmonic oscillators, when in the vicinity of equilibrium. The concept of wake fields and impedance is discussed, and a brief derivation of the equations governing longitudinal coupled-bunch dynamics is given.

Existing cures for coupled-bunch instabilities are mentioned, along with formulae describing the amount of damping that they provide. This is followed by some background material on the digital bunch-by-bunch feedback system that was used to perform most of the diagnostics.

2.1 Storage Rings, Single-Particle Dynamics

A storage ring contains lumps of charged particles injected into a vacuum chamber, which encloses the subsequent particle trajectories. Dipole magnets bend the trajectories, and guide the particles at the nominal beam energy E_o along a closed orbit known as the reference orbit, or design orbit. Various magnets along the circumference of the ring provide transverse focusing, so that off-orbit particles are pushed back towards the design orbit. This gives rise to betatron (transverse) oscillations about the reference, in the horizontal and vertical planes.

During each revolution, the magnetic fields cause the particles to lose some energy

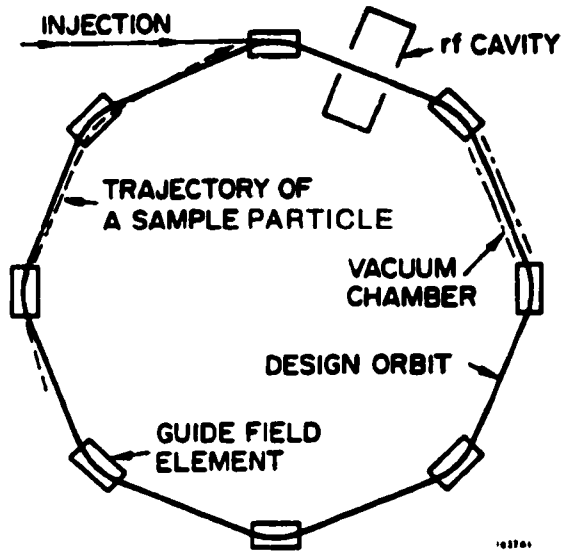

Figure 2.1: Schematic diagram of a storage ring.

through radiation, known as synchrotron radiation. This is balanced by an energy kick from accelerating cavities. The accelerating field has to be time-varying, since a DC field would not give circulating charged particles any net energy gain. The cavities are usually driven by high power microwave sources. Figure 2.1 shows the main components of a storage ring, viewed from above[1].

The RF cavity field is periodic, so that successive bunches of charged particles traverse the accelerating gap at intervals of roughly one period, *i.e.* T_{rf}. Thus, the stored beam can consist of at most h bunches, where the harmonic number h is the ratio of the RF frequency $f_{rf} = \omega_{rf}/2\pi$ to the revolution frequency $f_o = \omega_o/2\pi$. In other words,

$$h = \frac{\omega_{rf}}{\omega_o} = \frac{f_{rf}}{f_o} = \frac{T_o}{T_{rf}} \tag{2.1}$$

A synchronous particle is defined as a particle with energy E_o, that samples the accelerating voltages so that it gains exactly as much energy (U_o) as it loses through synchrotron radiation, on each turn. The behaviour of the synchronous particle, as well as the longitudinal focusing properties of the accelerating field in the RF cavity, can be understood

[1]This figure is from M. Sands' note on electron storage rings [18].

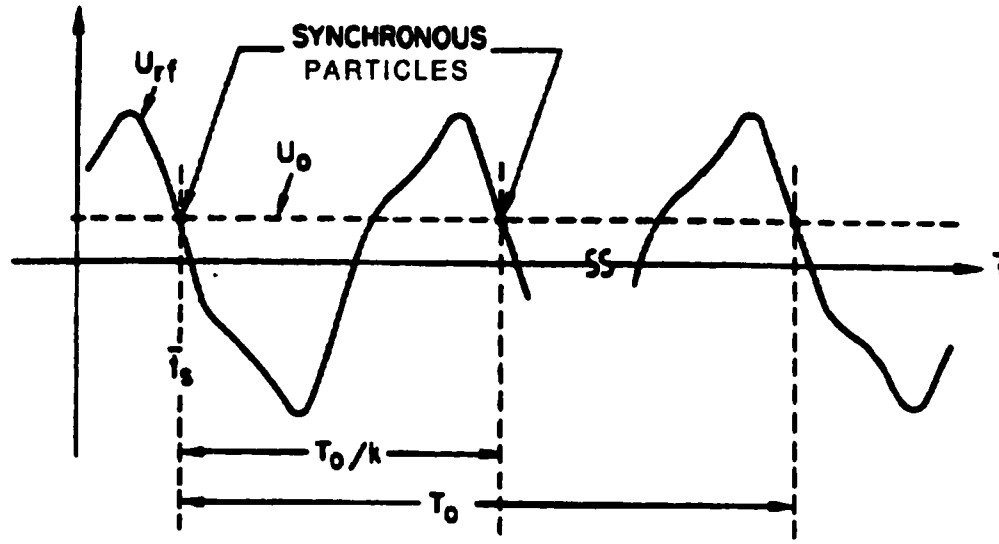

Figure 2.2: Example of energy gain from an RF system as a function of the starting time t of a revolution.

from Fig. 2.2, which shows an example of the energy gain from an RF system as a function of the starting time t of a revolution[2]. A particle arriving at the cavity a bit later than the synchronous particle will receive less energy than U_o from it. If the particle is highly relativistic, the net energy loss will cause it to follow a shorter path around the ring, without any significant change in speed (more bending in the bend magnet fields). Successive revolutions of decreasing circumference will push the delayed particle back in the direction of the synchronous particle. Conversely, particles sampling the RF waveform a bit early receive more energy than U_o, thus delaying their next arrival at the cavity. In this way, the accelerating field creates a longitudinal potential well, giving rise to synchrotron (longitudinal) oscillations.

All particles within a certain neighbourhood of the synchronous particle in longitudinal phase space will remain confined to that neighbourhood, due to phase focusing. This neighbourhood is known as an RF bucket, since any particle placed inside it will remain confined to it. Particles outside the bucket are usually lost, unless they are kicked into a bucket by forces other than the accelerating field in the cavities. The energy loss per turn by synchrotron radiation increases with the energy of the particle. In highly relativistic machines, this gives rise to radiation damping of all oscillation amplitudes. The trajectory

[2]This figure is from M. Sands' note on electron storage rings [18].

of each particle spirals inwards in phase space. The inward spiral is balanced in the steady state by quantum excitation from the stochasticity of synchrotron radiation [18].

For convenience, we shall refer specifically to high energy electron storage rings from here on, with the understanding that most results apply equally well to any other storage rings (with the exception of radiation damping being negligible, in the case of many hadron machines).

To the first order, the revolution frequency of stored electrons varies linearly with the energy error ϵ. Thus, if τ is the arrival time (time delay) of an electron relative to the synchronous particle, then

$$\dot{\tau} = \alpha \frac{\epsilon}{E_o}, \tag{2.2}$$

where the constant of proportionality α, known as the momentum compaction factor, is a function of E_o and the magnetic guide field. In the case of machines that are not highly relativistic, α is replaced by $\eta = \alpha - 1/\gamma^2$, where γ is the ratio of E_o to the rest energy.

If we assume that the electron completes many revolutions around the ring and receives many kicks from the accelerating cavities in one synchrotron (longitudinal) oscillation cycle, then energy kicks from the cavities can be smoothed over an entire turn. Thus,

$$\dot{\epsilon} = \frac{e\,V_{rf}(\tau) - U_{rad}\,(\epsilon)}{T_o}, \tag{2.3}$$

where V_{rf} is the total cavity voltage waveform shown in Fig. 2.2, $U_{rad}(\epsilon)$ is the energy lost to synchrotron radiation in one turn, and T_o is the revolution period.

If we ignore external perturbations and linearise these equations for small oscillations about the synchronous electron, we get the following equation for a damped harmonic oscillator.

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

where the radiation damping rate d_τ equals $\dot{U}(0)/2T_o$, and

$$\omega_s^2 = -\frac{\alpha e \dot{V}_{rf}(0)}{E_o T_o} \tag{2.5}$$

The resonance angular frequency ω_s of these oscillations is known as the synchrotron frequency. The synchrotron tune ν_s is defined as ω_s/ω_o, where the revolution frequency ω_o equals $2\pi/T_o$.

It is clear from the above equation that $\dot{V}_{rf}(0)$ is negative if the beam is highly relativistic

($\eta \approx \alpha > 0$). In other words, high energy particles experience phase focusing only on the negative slope of the cavity voltage. In most cases, the accelerating voltage $V_{rf}(\tau)$ in the accelerating cavities is sinusoidal, and can be written as:

$$V_{rf}(\tau) = V_c \sin[\omega_{rf}(\tau_s + \tau)], \tag{2.6}$$

where τ_s is the arrival time of the synchronous particle. Thus, Eq. 2.5 is equivalent to:

$$\omega_s^2 = -\frac{\alpha e \omega_{rf}}{E_o T_o} V_c \cos(\omega_{rf}\tau_s) \tag{2.7}$$

If the average total energy loss of the particle over one turn is U, then

$$eV_c \sin(\omega_{rf}\tau_s) = U \tag{2.8}$$

$$\Rightarrow \tau_s = \frac{1}{\omega_{rf}} \sin^{-1}\left(\frac{U}{eV_c}\right) \tag{2.9}$$

Typically, τ_s is slightly below $T_{rf}/2$. Clearly, eV_c has to be greater than U, for the beam to be stable. The ratio eV_c/U is known as the overvoltage factor. As the overvoltage is increased, ω_s increases, τ_s increases towards $T_{rf}/2$, and the size of the RF bucket in phase space increases. The overvoltage is limited by the peak power output of the klystrons (microwave amplifiers) that drive the RF cavities. The klystrons need to be operated well below saturation to allow RF feedback loops to function.

2.2 Wake Fields and Impedance

Charged particles leave behind electromagnetic fields as they fly through the vacuum chamber [19]. These fields arise from discontinuities in the vacuum chamber, and the finite resistivity of the vacuum pipe. If the particles move at about the speed of light, causality dictates that the fields should be excited after the particle has passed by. The momentum of trailing particles is affected by these "wake fields," that introduce coupling between the individual particle oscillations. The total wake force acting on a trailing charge is calculated by integrating over the contributions of all the particles that preceded it.

If a highly relativistic test charge e follows a similar particle of unit charge at a distance z, then it loses energy equal to $eW^{\parallel}(z)$ on every turn, due to the wake fields. The longitudinal wake function $W^{\parallel}(z)$ is calculated by integrating the longitudinal component of the electric

field experienced by the test charge over an entire turn. If the lead particle has a transverse displacement x in the horizontal direction, then the integrated horizontal wake force equals $-exW^x(z)$, where $W^x(z)$ is the horizontal wake function.

The wake function is the aggregate impulse response of the vacuum chamber. It does not depend upon properties of the charged particle beam. Due to causality, $W^{\parallel}(z) = 0$ and $W^{\perp}(z) = 0$ when $z < 0$.

If the particles are highly relativistic, we can use the approximation $z = ct$ to write the wake functions as functions of the time lag t rather than the distance z, where c is the speed of light. A more complete discussion of wake functions and impedances is given in [3].

Theoretical analyses of wake field perturbations of longitudinal and transverse dynamics are simplified by the assumption that the wake fields from each element in the vacuum chamber are smoothed over the entire ring, instead of being localised to the vicinity of the source. On the other hand, simulations are conveniently performed if we assume that the one-turn integrated wake is applied as a lumped kick at a single point in the ring. Thus, it is often desirable to assume that the dynamics are not significantly affected by smoothing of the wake forces, or their translation by less than a turn. This assumption is implicit in the calculation of a wake function that does not depend on the longitudinal position of the test particle in the ring. It is shown in the appendix to [20] that this is a reasonable approximation, as long as the tune perturbation induced by wake fields is small.

The longitudinal beam impedance $Z^{\parallel}(\omega)$ is defined as the Fourier transform of the longitudinal wake function $W^{\parallel}(t)$:

$$Z^{\parallel}(\omega) = \int_{-\infty}^{\infty} W^{\parallel}(t)\, e^{-j\omega t}\, dt \tag{2.10}$$

Similarly,

$$Z^{\perp}(\omega) = j \int_{-\infty}^{\infty} W^{\perp}(t)\, e^{-j\omega t}\, dt \tag{2.11}$$

In general, the beam impedance is complex. Since the wake function is the impulse response of the beam environment, the beam impedance is the transfer function from the beam current to the wake voltage.

RF cavities are among the main sources of impedance in storage rings. In addition to the fundamental accelerating mode, they contain additional trapped modes, known as parasitic modes or Higher-Order Modes (HOMs). The impedance of these resonances is similar to that of an RLC (resistor-inductor-capacitor) band pass filter. The impedance of

longitudinal monopole modes in the cavities is typically expressed as:

$$Z^{\parallel}(\omega) = \frac{R_s}{1 + jQ[\omega/\omega_r - \omega_r/\omega]}, \tag{2.12}$$

where ω_r is the resonance angular frequency of the mode and R_s is the shunt impedance [21]. $Z^{\parallel}(\omega) = R_s$ at resonance, *i.e.* at $\omega = \omega_r$. The quality factor Q determines the width of the impedance. If $Q \gg 1$ (narrowband mode), then the 3-dB bandwidth of the resonant mode is approximately ω_r/Q. A mode is considered to be broadband if $Q \approx 1$. The velocity of a damped harmonic oscillator responds to the externally applied force through a similar transfer function.

An inverse Fourier transform gives the corresponding longitudinal wake function

$$\begin{aligned} W^{\parallel}(t) &= 0; && t < 0 \\ &= \Gamma R_s; && t = 0 \\ &= 2\Gamma R_s\, e^{-\Gamma t}\left(\cos(\omega_p t) - \frac{\Gamma}{\omega_p}\sin(\omega_p t)\right); && t > 0, \end{aligned} \tag{2.13}$$

where the damping rate Γ is given by:

$$\Gamma = \omega_r/2Q, \tag{2.14}$$

and

$$\omega_p = \sqrt{\omega_r^2 - \Gamma^2} \tag{2.15}$$

Thus, narrowband modes persist for a long time when excited, whereas broadband modes are more rapidly damped. The ratio R_s/Q for a given cavity mode is determined solely by the cavity geometry, while the Q of the mode is determined by the resistance of the cavity walls.

The magnitude and phase of a sample band pass impedance $Z^{\parallel}(\omega)$ are shown in Figs. 2.3(a) and (b) respectively. In this example, $Q = 10$ and $R_s = 1$. The approximate 3-dB frequencies $\omega_r(1 \pm 1/2Q)$ are marked by vertical dashed lines. The angle of the impedance goes from 90^o at low frequencies to -90^o at high frequencies, crossing zero exactly at $\omega = \omega_r$. As expected from Eq. 2.12, $\angle\, Z^{\parallel} \approx \pm 45^o$ when $\omega = \omega_r(1 \mp 1/2Q)$.

Figs. 2.3(c) and (d) show the real and imaginary parts of $Z^{\parallel}(\omega)$. The real part of the impedance reaches its maximum at $\omega = \omega_r$, and drops to $R_s/2$ at around $\omega_r(1 \pm 1/2Q)$. It

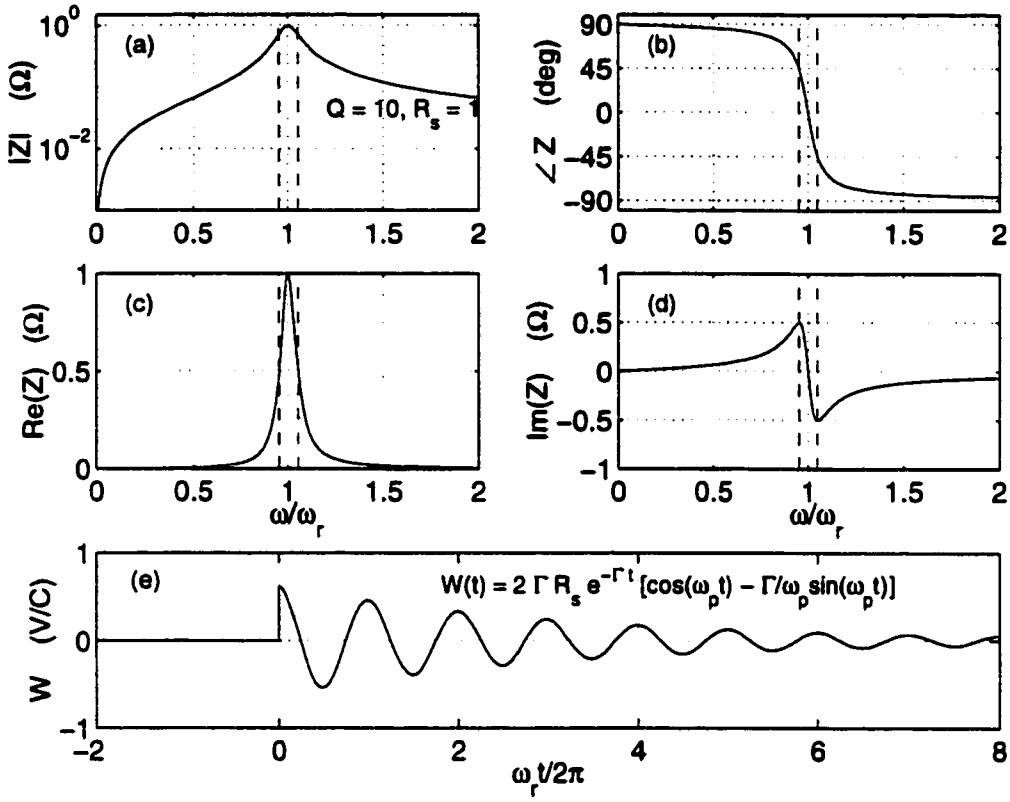

Figure 2.3: Example longitudinal impedance and wake function of a resonant cavity mode (RLC band pass filter) with $Q = 10$ and $R_s = 1$.

is always positive, since the cavity is a passive device. $Im[Z^{\parallel}(\omega)]$ is positive below ω_r and negative above it. It reaches its extreme values of $\pm R_s/2$ at approximately $\omega_r(1 \mp 1/2Q)$, with a steep slope in between.

The longitudinal wake function (impulse response) corresponding to the resonator impedance (second-order band pass transfer function) is shown in Fig. 2.3(e). It is an exponentially decaying sinusoid whose frequency ω_p is very close to ω_r when $Q^2 \gg 1$. It is necessarily positive at $t = 0+$, because the energy loss $q^2 W^{\parallel}(0)$ of the charge that produces the wake is necessarily positive.

2.3 Coupled-Bunch Dipole Instabilities

Since this dissertation mainly addresses rigid (dipole) oscillations of bunches of charged particles, it is sufficient to assume that each bunch behaves like a macroparticle with the appropriate charge. More general approaches can be found in [3, 4, 5, 22].

Consider a stored beam that consists of N identical macroparticle bunches of charge q, filled in N RF buckets, with a constant interbunch spacing $T_b = T_o/N$. Such a fill is known as an "even fill." N is necessarily a factor of h, *i.e.*, N divides h evenly. The total beam current I_o is given by

$$I_o = \frac{1}{T_o} \sum_{k=0}^{k=N-1} q_k = \frac{q}{T_b} \tag{2.16}$$

Using a derivation similar to that of Eq. 2.4, we can write the following equation to describe the effect of wake fields on the longitudinal oscillations of the n^{th} bunch in the ring:

$$\ddot{\tau}_n + 2d_\tau \dot{\tau}_n + \omega_s^2 \tau_n = -\frac{\alpha e}{E_o T_o} V_n^{wk}(t); \qquad n = 0, 1, ..., N-1, \tag{2.17}$$

where $eV_n^{wk}(t)/T_o$ is the rate of energy loss of macroparticle n due to the superposition of the wake forces of all preceding macroparticles. It is assumed that all sources of wake field coupling are smoothed over the entire ring. The following equations also contain the assumption that the wavelengths of the relevant wake fields are much larger than the bunch length, so that the wake of a bunch is approximated by that of a point charge q, located at the bunch centroid. Thus,

$$V_n^{wk}(t) = q \sum_{p=-\infty}^{\infty} \sum_{k=0}^{N-1} W[t_{n,k}^p + \tau_n(t) - \tau_k(t - t_{n,k}^p)], \tag{2.18}$$

where the time interval between bunch n, on the present turn, and bunch k, p turns ago, is $t_{n,k}^p = (pN + n - k)T_b$. If τ is small compared to the period of the highest relevant frequency in the beam impedance, we can use the Taylor expansion

$$W[t_{n,k}^p + \tau_n(t) - \tau_k(t - t_{n,k}^p)] \approx W(t_{n,k}^p) + [\tau_n(t) - \tau_k(t - t_{n,k}^p)]\dot{W}(t_{n,k}^p) \tag{2.19}$$

The first term in the above expansion is a constant; it only shifts the equilibrium position of the bunch, *i.e.*, the arrival time of the synchronous particle. Since τ is the bunch position relative to equilibrium, and since the azimuthal (rotational) symmetry of the even fill forces

the equilibrium positions of all bunches to shift by the same amount, we can ignore this static term altogether. Later on, we will see that uneven fills (lacking azimuthal symmetry) produce equilibrium shifts that vary from bunch to bunch. This variation will be used to estimate the longitudinal impedance in Ch. 2. In Ch. 7, it will be assumed that this variation has a negligible effect on beam stability, since the equilibrium shifts induced by wake fields are typically small compared to the wavelengths of the significant beam impedances. Thus,

$$V_n^{wk}(t) \approx \sum_{p=-\infty}^{\infty} \sum_{k=0}^{N-1} q_k[\tau_n - \tau_k(t - t_{n,k}^p)]\dot{W}(t_{n,k}^p) \tag{2.20}$$

By combining Eqs. 2.17 and 2.20, we get

$$\ddot{\tau}_n + 2d_r\,\dot{\tau}_n + \omega_s^2\,\tau_n \approx -\frac{\alpha eq}{E_o T_o} \sum_{p=-\infty}^{\infty} \sum_{k=0}^{N-1} [\tau_n - \tau_k(t - t_{n,k}^p)]\dot{W}(t_{n,k}^p); \quad n = 0, 1, ..., N-1 \tag{2.21}$$

This is a system of coupled linear equations, with solutions of the form $e^{\lambda t}$, or $e^{j\Omega t}$. The eigenvalue λ is related to the coherent oscillation frequency Ω by: $\lambda = j\Omega$.

Although it is not immediately obvious from the last equation, we know from the rotational symmetry of the fill shape that the eigenvectors must remain unchanged, except for multiplication by a scalar, when rotated by one bucket. Let this scalar be r. Rotation by N buckets brings us back to the original eigenvector, so

$$r^N = 1 \tag{2.22}$$

Thus, r must be one of the N^{th} roots of unity:

$$r_l = e^{j2\pi l/N}; \qquad l = 0, 1, ..., N-1 \tag{2.23}$$

These N roots of unity define the N eigenvectors of an N-bunch even-fill, which are merely the Fourier vectors

$$\begin{aligned} v_l &= [1\ e^{jl\theta}\ e^{2jl\theta}..\ e^{(N-1)jl\theta}]^T/\sqrt{N}; \qquad l = 0, 1, ..., N-1 \\ \theta &= 2\pi/N \end{aligned} \tag{2.24}$$

It is clear that a discrete Fourier transform is all that is required to project even-fill coupled-bunch oscillations onto the even-fill eigenmodes (EFEMs). As we shall see in Ch. 7, the

symmetry argument used here for even fill shapes can be extended to yield useful results for periodic uneven fills. It will also be shown in the same chapter that the N EFEMs form a convenient set of basis vectors for studying uneven-fill dynamics.

Since we now know the even-fill eigenvectors, we only need to plug them into Eqs. 2.21 to find the corresponding eigenvalues. If the bunches are oscillating in mode l, then[3]

$$\begin{aligned} \tau_k(t) &= Ae^{j2\pi kl/N}e^{j\Omega t} \\ \Rightarrow \tau_k(t-t^p_{n,k}) &= \tau_n(t)e^{-j2\pi(n-k)l/N}e^{-j\Omega t^p_{n,k}} \\ &= \tau_n(t)e^{-j(l\omega_o+\Omega)t^p_{n,k}}, \end{aligned} \tag{2.25}$$

since $\omega_o T_b = 2\pi/N$. If we use the substitution $u = pN + n - k$, then

$$\tau_k(t-t^p_{n,k}) = \tau_k(t-uT_b) = \tau_n(t)e^{-j(l\omega_o+\Omega)uT_b} \tag{2.26}$$

The analysis is further simplified by assuming that the radiation damping rate is small ($d_r \ll \omega_s$), and that the effect of the wake fields can be treated as a small first-order perturbation to the uncoupled bunch dynamics [$|\Omega - \omega_s| \ll \omega_s$]. These are realistic approximations, which can be checked after the first-order calculation has been performed. They allow us to simplify the LHS (left-hand side) of Eq. 2.21 as follows:

$$\ddot{\tau}_n + 2d_r\dot{\tau}_n + \omega_s^2\tau_n \approx 2j\omega_s[\dot{\tau}_n + (d_r - j\omega_s)\tau_n] \tag{2.27}$$

From Eqs. 2.21, 2.26 and 2.27, we get

$$\dot{\tau}_n + (d_r - j\omega_s)\tau_n \approx \tau_n\frac{j\alpha eq}{E_o 4\pi\nu_s}\sum_{u=-\infty}^{\infty}[1 - e^{-j(l\omega_o+\Omega)uT_b}]\dot{W}(uT_b) \tag{2.28}$$

Using Fourier identities and Eq. 2.16, the equation above can be transformed to:

$$\begin{aligned} \dot{\tau}_n+(d_r-j\omega_s)\tau_n \approx -\tau_n\frac{\alpha eI_o}{E_o 4\pi\nu_s}\sum_{p=-\infty}^{\infty}&[pN\omega_o Z^{\parallel}(pN\omega_o) \quad \ldots \\ &-[(pN+l)\omega_o+\Omega]Z^{\parallel}[(pN+l)\omega_o+\Omega] \end{aligned} \tag{2.29}$$

$$\Rightarrow \dot{\tau}_n+(d_r-j\omega_s)\tau_n = \lambda_l\tau_n \tag{2.30}$$

[3]The terms "mode l" and "EFEM l" will henceforth be used interchangeably, where the meaning is clear from the context. A more clear distinction will be made while discussing the eigenmodes of unevenly filled rings.

If we treat the RHS of this equation as a first order perturbation by assuming that $\Omega \approx \omega_s$, then the coherent eigenvalue shift λ_l of mode l is

$$\lambda_l = \frac{\alpha e f_{rf} I_o}{2E_o \nu_s}[Z^{\parallel eff}(l\omega_o + \omega_s) - Z^{\parallel eff}(0)]; \tag{2.31}$$

$$Z^{\parallel eff}(\omega) = \frac{1}{\omega_{rf}} \sum_{p=-\infty}^{\infty} (pN\omega_o + \omega) Z^{\parallel}(pN\omega_o + \omega) \tag{2.32}$$

We see that λ_l is proportional to the product of I_o and the effective impedance $Z^{\parallel eff}$. This makes sense, since the coherent eigenvalue shift produced by wake fields must be proportional to the wake voltage induced by the beam. Although $Z^{\parallel eff}$ has the dimensions of an impedance, it is actually an aliased and scaled version of $\omega Z^{\parallel}(\omega)$. Proportionality to ω reflects the fact that the bunch dynamics are coupled through the derivative of the longitudinal wake voltage, rather than its actual value (this is not true for transverse coupled-bunch instabilities).

Although the actual eigenvalue of mode l is given by

$$\Lambda_l = (-d_\tau + j\omega_s) + \lambda_l, \tag{2.33}$$

the coherent eigenvalue shift λ_l will often be referred to as the "eigenvalue" of mode l, for convenience, and because λ contains all of the variation between eigenvalues of different modes.

Since the eigenvectors evolve as $e^{\Lambda_l t}$, the growth rate of mode l is $Re(\Lambda_l) = Re(\lambda_l) - d_\tau$. If the product of I_o and the real part of the effective impedance at one of the modal frequencies $(l\omega_o + \omega_s)$ is large enough, the growth rate is positive, and bunch oscillations will grow exponentially until they are limited by nonlinearities or beam loss. The threshold value of the beam current I_o is defined as the value at which the largest net growth rate $Re(\Lambda_l)$ equals 0. Since instabilities degrade accelerator performance quite significantly, most machines operate below their instability thresholds. The coherent frequency shift of mode l is $\Omega_l - \omega_s = Im(\lambda_l)$.

We see from Eq. 2.32 that $Re(Z^{\parallel eff})$ is an odd function of ω and $Im(Z^{\parallel eff})$ is an even function. Thus, if ω_s is much smaller than the widths of the peaks in the impedance spectrum, then the EFEM eigenvalues come in pairs $\{\lambda_l, \lambda_{N-l}\}$, where

$$\lambda_{N-l} \approx -\lambda_l^* \tag{2.34}$$

The existence of a highly stabilised counterpart to every unstable mode will be exploited in Ch. 7.

The exact contribution of an individual impedance resonance of the form of Eq. 2.12 depends on the location of ω_r with respect to the modal frequencies. However, it is possible to draw some general conclusions from Eqs. 2.12 and 2.31. If an impedance resonance has $Q \gg 1$, then we can use the approximation

$$\frac{\omega R_s}{1 + jQ[\omega/\omega_r - \omega_r/\omega]} \approx \pm \frac{\omega_r R_s}{1 + jQ[\omega/\omega_r - \omega_r/\omega]}, \tag{2.35}$$

since the impedance is negligible far from $\omega = \pm\omega_r$. Hence,

$$\lambda \propto \pm \frac{\omega_r I_o R_s}{1 + jQ[\omega/\omega_r - \omega_r/\omega]} \tag{2.36}$$

This equation indicates that the real and imaginary parts of λ due to a high-Q resonator behave like the real and imaginary parts of a second order band pass transfer function (see Fig. 2.3). It is well known that the real and imaginary parts of a second order band pass impedance trace out a circle in the complex plane. This implies that the possible values of λ due to a fixed high-Q resonator must also lie on circles, as illustrated in Fig. 2.4. This is a very useful property of resonator eigenvalues, since it allows us to estimate the peak resonator impedance R_s from measurements of λ. We can see from the figure that the largest growth that an impedance resonance can produce is twice the largest tune shift.

Transverse coupled-bunch instabilities behave in a very similar way. The transverse analog of Eq. 2.31 is [3]

$$\lambda_l = \frac{\beta f_o I_o}{2(E_o/e)} Z^{\perp eff}(l\omega_o + \omega_\beta); \tag{2.37}$$

$$Z^{\perp eff}(\omega) = \sum_{p=-\infty}^{\infty} Z^{\perp}(pN\omega_o + \omega), \tag{2.38}$$

where ω_β is the betatron (horizontal or vertical) oscillation frequency, and β is the beta function. Note: ω_β is typically greater than 1. As opposed to the longitudinal effective impedance, the transverse effective impedance is simply an aliased version of $Z^{\perp}(\omega)$. The resistance of vacuum chamber walls is one of the important sources of transverse coupled-bunch instability [2].

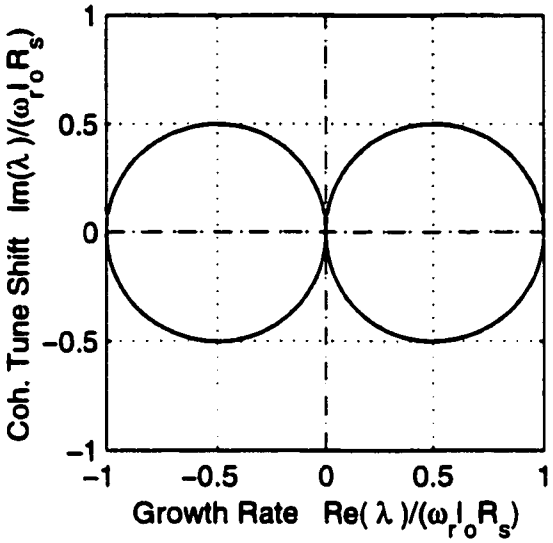

Figure 2.4: The locus of possible values of λ, due to a fixed high-Q resonator.

2.4 Existing Cures

Coupled-bunch instabilities are conventionally cured using: a) Minimization of impedance seen by the beam [23, 24]. b) Landau damping [25, 26]. c) Active feedback [27, 28, 29, 30]. For overviews of cures for instabilities, see [31, 32, 33, 34].

Miscellaneous cures that do not fall in any of the above categories include raising the beam energy (to increase radiation damping), and changing the scaling from Z to λ by changing the oscillation frequency.

2.4.1 Impedance Minimization

Accelerators are generally designed within an impedance budget, which sets an upper limit on the total impedance of the vacuum chamber. This sets an upper limit on the growth rates of the various beam instabilities. Landau damping and/or feedback must take care of the instabilities that remain after impedance minimization.

Techniques for reducing the beam impedance are too varied to discuss in detail here. The impedance of the metallic pipe surrounding the beam can be reduced by using a material such as copper, which has low resistivity, and by optimising the geometry of tube [35]. Discontinuities in the beam pipe, such as beam position monitors (BPMs), bellows, ion pumps, etc., are minimised, and designed to reduce unwanted contributions to the beam

impedance. RF cavities are among the main sources of coupled-bunch instabilities. Cavity HOMs can be damped by means of damping ports [36, 37], which reduce the Q of the modes by coupling high-frequency power out of the cavities and into resistive loads. Since the beam only sees the impedance at sidebands of revolution harmonics, one commonly used approach is to tune the worst cavity resonances so that they land in the gap between consecutive sidebands [38, 39]. Tuning is accomplished by adjusting the depth of plungers that intrude into the cavities, or by changing the temperatures of the cavities (thermal expansion changes the mode frequencies). It is also sometimes useful to build cavities with a systematic variation in geometry, so that the resonances of different cavities land at slightly different frequencies [40].

2.4.2 Landau Damping

Landau damping is the damping of coherent oscillations by means of a spread in the resonant frequencies of individual oscillators. The eigenvalue equations in the previous section are derived under the assumption that the individual oscillators all have the same nominal tune. If there is a spread in their tunes, it is clear that the instability growth rate will be smaller, since they will not couple to each other as strongly.

The term "Landau damping" is most often used to describe damping of bunch motion by means of a spread in the tunes of charged particles within a bunch (intrabunch tune spread). Although damping from bunch-to-bunch (interbunch) tune spreads is often referred to by the same name, opinion is divided on whether or not it qualifies as "Landau damping." For the purposes of this dissertation, however, the distinction is merely a semantic one, since the two cases are described by the same mathematical equation.

Let us assume that N electrons/macrobunches are coupled to each other by a single dominant impedance resonance, which gives their coherent motion an eigenvalue shift of λ, when the electron/bunch tunes are all identical. If the natural oscillation frequency ω_s (or ω_β) of each electron/bunch k is now shifted by an amount δ_k, then the new Landau-damped eigenvalue (shift) λ^L is given by the dispersion relation [26]

$$\frac{1}{\lambda} = \langle \frac{1}{\lambda^L - j\delta_k} \rangle \approx \int_{\delta_{min}}^{\delta_{max}} \frac{\rho(\delta)}{\lambda^L - j\delta}\, d\delta, \tag{2.39}$$

where $\langle . \rangle$ denotes the mean over all values of k. The integral version of this equation is based on the approximation that the δ_ks are closer to their neighbours than they are to λ^L,

in which case we can replace the discrete averaging over k by an average over a fictitious continuous distribution $\rho(\delta)$. There are many ways of deriving Eq. 2.39. One of them will be shown in Ch. 6.

Intrabunch tune spreads are generated by deliberately distorting the potential wells occupied by the bunches [41, 42]. Interbunch tune spreads can be created by RF quadrupoles in the transverse case [43] and RF drive modulation in the longitudinal case [44]. For example, if the cavity voltage $V_{rf}(t)$ were amplitude modulated at f_o, different bunches would see different slopes $\dot{V}_{rf}(0)$, and δ_k would complete one sinusoidal cycle from $k = 0$ to $k = N-1$.

2.4.3 Active Feedback

In most modern accelerators, active feedback is necessary for controlling coupled-bunch instabilities induced by resistive beam pipe walls and parasitic trapped modes. The general strategy of feedback is to sense the parameter of the system that must be controlled, perform some operation on the sensed signal to produce a correction signal, and then apply the correction signal back to the system through an actuator. For example, a longitudinal feedback system might sense the longitudinal arrival time error τ of a bunch by mixing (multiplying) the beam position monitor (BPM) signal with a reference sinusoid that is locked to the RF frequency. The low-pass-filtered output would go through an analog or digital processing block that estimates $\dot{\tau}$ (*i.e.* ϵ). This feedback signal would then be amplified and fed to a longitudinal kicker, which would deliver the prescribed energy kick to each bunch. See [45] for an overview of pickup electrodes and kickers.

Feedback systems for curing coupled-bunch instabilities can be divided into frequency-domain (mode-by-mode) and time-domain (bunch-by-bunch) systems.

Mode-by-mode systems [27, 46] are sometimes used when the offending impedance resonances are narrowband, and small in number. The feedback front end has a narrowband filter at the frequency of each unstable mode. Similarly, the processing block treats each mode independently, while applying the required phase shift and gain. Thus, there are as many parallel channels as there are modes to be damped. Clearly, the mode-by-mode approach is cumbersome in machines with large broadband impedance resonances, or many narrowband resonances.

Bunch-by-bunch systems resolve the individual bunch signals in the front end, and act on each bunch individually. In other words, the feedback signal for a bunch depends only

on it's own error signal. A full-bandwidth time-domain system would have to cover at least $N/2$ revolution harmonics in its passband. It can be shown that a full-bandwidth time-domain system with identical damping on all bunches is equivalent to a frequency-domain system with identical damping on all modes [28].

The action of longitudinal bunch-by-bunch feedback can be explained with the help of the following extension of Eq. 2.17:

$$\ddot{\tau}_n + 2d_\tau \dot{\tau}_n + \omega_s^2 \tau_n = \frac{\alpha e}{E_o T_o}[V_n^{fb}(t) - V_n^{wk}(t)], \tag{2.40}$$

where $V_n^{fb}(t)$ is the feedback kick given to bunch n.

The problem of coupled-bunch instabilities would be completely solved if we could only make $V_n^{fb}(t)$ equal $V_n^{wk}(t)$. Of course, things are never that simple. Firstly, there is the problem of detecting V^{wk}. The wake voltage is composed of contributions from all parts of the storage ring, and it is clearly not feasible to put detectors everywhere, even if we take for granted the ability to reconstruct wake fields from electrode signals. It is often the case, though, that most of the troublesome wakes originate from a few well-known locations in the ring, such as the RF cavities. A few strategically placed detectors could then be used to sense the wake fields in the cavities. We are then faced with the problem of reconstructing $V_n^{wk}(t)$ fast enough to cancel it out, over a broad range of frequencies. Cable delays and the speed of the electronics are limiting factors in this case. There is also the practical problem of adding special couplers for delivering feedback signals to all the troublesome elements of the vacuum chamber. In addition, we need feedback amplifiers with sufficient power, that span the range of frequencies over which the impedance is significantly large.

Clearly, the approach of canceling $V_n^{wk}(t)$ out at the source is fraught with problems. There is, however, one important band-limited implementation of this approach that is often used, namely, RF feedback [47, 48, 21]. With this exception, instabilities are usually controlled by detecting the relevant beam oscillation signal, and applying a suitable correction through a kicker, as described at the beginning of this subsection. Such beam-based feedback systems are limited by their group delay, bandwidth, detection noise, and amplifier power. Broadband amplifiers are among the most expensive components of modern feedback systems.

If the magnitude of the feedback gain at the frequency ω_s is G^{fb}, then

$$V_n^{fb} = -jG^{fb}\tau_n \approx -\frac{G^{fb}}{\omega_s}\dot{\tau}_n \tag{2.41}$$

From Eqs. 2.40 and 2.41, the feedback-induced damping rate is given by

$$d^{fb} = \frac{\alpha e}{2E_oT_o}\frac{G^{fb}}{\omega_s} \tag{2.42}$$

Note: The growth and damping rates from impedances, feedback, and radiation damping are all assumed to be first order perturbations that add up linearly. Thus, a beam with no Landau damping is stable only if the sum of radiation damping and feedback damping is larger than the impedance-induced growth rate, *i.e.*, if the net growth rate

$$Re(\Lambda) = Re(\lambda) - d_r - d^{fb} < 0 \tag{2.43}$$

We know that λ is proportional to I_o. If there is no normalisation in the front end, and time-domain feedback is used, d^{fb} must be proportional to the bunch current $i_b = I_o/N$. If frequency-domain feedback is used, $d_{fb} \propto I_o$. Neglecting radiation damping for the moment, we can therefore say that the beam must be stable at all currents, if it is stable at any one current. In other words, damping coupled-bunch instabilities is simply a matter of increasing the feedback gain G^{fb} until the net growth rate becomes negative. This seems simple enough to do. However, the maximum power output of the back end amplifiers places a limit on G^{fb}.

Rule of thumb: The feedback system is at it's maximum achievable gain when sensor noise (detection noise) in the front end causes saturation in the back end.

Once feedback is saturated, G^{fb} is approximately inversely proportional to the beam oscillation amplitude. Thus, for a given no-feedback growth rate $Re(\lambda)$, the maximum achievable kicker voltage determines the maximum oscillation amplitude at which Eq. 2.43 is still satisfied. The power amplifier is therefore selected so that the feedback actuator is strong enough to damp the largest expected growth rate, even when the beam oscillates at the worst-case amplitude expected from noise excitation and injection errors [49].

If the feedback system has a gain that scales with the bunch current, it is equivalent to an additional beam impedance [50]. From Eqs. 2.31 and 2.42, the effective damping

impedance of the feedback system is

$$Z^{fb} = \frac{G^{fb}}{I_o \, \omega_{rf}} \tag{2.44}$$

This dissertation is mostly concerned with the longitudinal and transverse feedback systems at the ALS and PEP-II. These are time-domain systems. The ALS transverse system [51] calculates a weighted sum of the signals from two different BPMs and delays the correction signal by the appropriate amount by means of a long cable. The PEP-II transverse system [52] is slightly different, in that it has a digital delay.

2.5 Digital Bunch-by-Bunch Longitudinal Feedback System

The PEP-II/ALS/DAΦNE longitudinal feedback (LFB) system [53, 54, 55] is the main source of data presented in the following chapters. This digital, programmable, time domain system has been installed at the ALS, PEP-II, DAΦNE [56, 57], BESSY-II [58, 59], and the PLS [60, 61]. A block diagram of the main components of the system is shown in Fig. 2.5.

Programmability is a necessary feature of the DSP-based system, since beam parameters can and do change with time. The LFB system has great utility as a diagnostic tool, since it has the capacity to record the digitised oscillation coordinate of each bunch in the ring, while simultaneously manipulating feedback parameters. The LFB user interface and control software are described in [62].

2.5.1 General Approach, Downsampling

The LFB system has been designed for machines with up to a few thousand bunches, circulating at revolution frequencies in the MHz or hundreds of kHz range. Conceptually, a bunch-by-bunch system is equivalent to a system with N parallel channels for the N bunches. For this reason, the correction signal is calculated in parallel by an array of digital signal processors (DSPs). This makes the system scalable. The sequence of bunch oscillation coordinates emanating from the front end is demultiplexed, and sent to the designated DSPs in small packets. The correction signals calculated by the DSPs are written to addresses in a hold buffer, which contains the latest correction values, in the order of the bunch numbers.

Scalability is limited by bottlenecks in the motion of data to and from the DSPs. The data-handling problem is solved by downsampling the bunch signals. Since a synchrotron

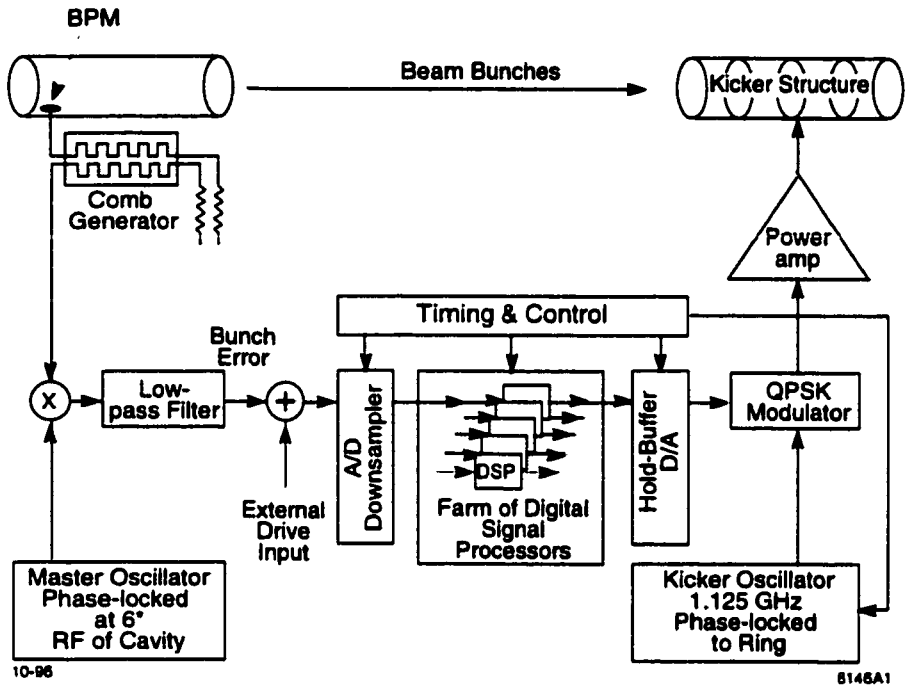

Figure 2.5: Block diagram of longitudinal bunch-by-bunch feedback system.

oscillation cycle usually takes tens or bundreds of revolutions, it is not necessary to compute a correction signal on every turn. Instead, the system looks at the oscillation coordinate of a bunch and updates its correction signal in the hold buffer only once every D turns, where D is a programmable downsampling factor. Downsampling reduces data flow in and out of the DSP farm, and also reduces the number of DSPs required to implement the feedback algorithm. Note: The hold buffer is read continuously, so that bunches receive correction kicks on every turn.

2.5.2 Front End

The front end detects the sequence of arrival time errors τ_n, and digitises it at the bunch crossing rate. It must be broadband, to accept the entire range of modal frequencies. Equivalently, there must be minimal cross talk between the signals of adjacent bunches [63].

It is convenient to view the longitudinal oscillation coordinate as a phase error, where

the phase of a bunch is defined by

$$\phi = \omega_{rf}\,\tau \tag{2.45}$$

The front end is thus a phase detector.

As shown in Fig. 2.5, the sum signal from four button-type BPM electrodes is passed through a comb generator. The comb generator converts the impulsive input from each bunch traversal into a tone burst of a few cycles at Hf_{rf}, where H is the harmonic of the RF frequency at which the phase is detected. The comb generator functions as a band pass filter centered at Hf_{rf}. Interbunch cross talk is minimised by keeping the tone burst shorter than the bunch spacing. The comb filter output is mixed with a signal at Hf_{rf} from a master oscillator that is locked to the ring RF clock. The mixer output contains the baseband detected phase signal, together with a mixing product at $2Hf_{rf}$, which is removed by a low pass filter. The ouput of the low pass filter is digitised by an 8-bit analog-to-digital convertor (ADC).

Let us assume that the master oscillator signal is proportional to $\cos[H(\omega_{rf}t + \phi_{m.o.})]$, and the comb generator output is proportional to $i_b \sin[H(\omega_{rf}t + \phi_b)]$, where i_b is the bunch current, and ϕ_b is the bunch phase error. The low-pass-filtered mixer output can then be described by

$$s(t) \;\propto\; i_b \sin[H(\phi_b - \phi_{m.o.})] \tag{2.46}$$

$\phi_{m.o.}$ is adjusted by means of a phase shifter, so that it approximately equals the average equilibrium bunch phase. Thus, for small oscillations about equilibrium,

$$s(t) \;\propto\; i_b\,(\phi_b - \phi_{m.o.}) \tag{2.47}$$

This proportionality holds as long as

$$H(\phi_b - \phi_{m.o.}) \ll \pi/2 \tag{2.48}$$

These equations form the basis of techniques described in Ch. 3 for measuring bunch currents and synchronous phase (equilibrium phase) transients.

2.5.3 Filter Algorithm

The DSPs implement a bunch-by-bunch discrete-time feedback algorithm. For the results discussed in this dissertation, they were programmed to act as finite impulse response (FIR)

filters [64, 65]:

$$u_i(n) = \sum_{k=0}^{N-1} h(k)\phi_i(n-k), \tag{2.49}$$

This equation represents a discrete-time convolution in which the correction signal u_i of the i-th bunch at turn n is computed using a weighted average of the last N measurements[4], $\phi_i(n\text{–}N\text{+}1), \cdots, \phi_i(n)$, of the phase of bunch i. The bunch-by-bunch nature of this algorithm is evident from the fact that u_i is calculated from ϕ_i alone, without considering the phases of other bunches. The filter coefficients, $h(0), \cdots, h(N-1)$, are selected to provide zero DC response[5] and maximum gain G_{fb} at f_s. Their phase shift is adjusted for net negative feedback at f_s.

More sophisticated infinite impulse response feedback filters are being tested, that optimise feedback performance under constraints that cannot be met by FIR filters [66, 67].

2.5.4 Back End

The back end consists of the portion of the feedback system downstream of the digital-to-analog converter (DAC). It takes the baseband correction signal from the hold buffer, modulates it, and passes it to the kicker through a broadband power amplifier (see Fig. 2.5).

The DAC runs at the bunch crossing frequency f_b, with a resolution of 8 bits. The DAC output must be impressed on the beam with a bandwidth of at least $f_b/2$, if each bunch is to receive an independent kick. High-power baseband (0 to $f_b/2$) amplifiers with the required bandwidth (hundreds of MHz) are hard to come by. It is also not easy to design a baseband kicker with high shunt impedance and a flat response over the required frequency range. For these reasons, the DSP output is made to amplitude modulate a quadrature phase shift keyed (QPSK) carrier, before it is amplified and fed to the kicker. The carrier signal is a sinusoid at $n \pm 1/4$ times the RF frequency. It is locked to the beam timing signal, so as to maintain synchronisation between the peaks of the correction signal and the arrival of bunches in the kicker. See [68] for a discussion of the merits of QPSK modulation in this context. Kicker structures used with this feedback system are discussed in [69, 70].

[4]Equation 2.49 seems to indicate that the calculation of the correction signal in the DSPs is instantaneous, since $u_i(n)$ depends on $\phi_i(n)$. In reality, the calculation of u_i takes roughly $1/n_b$ samples, where n_b is the number of bunches handled by each DSP.

[5]The DC response of the filters contributes nothing to beam stability. It should be minimised, to avoid saturation of the correction signal.

Chapter 3

Bunch Currents and Synchronous Phases

Electron bunches in circular accelerators lose energy to synchrotron radiation and exchange energy with wake fields on each turn. The average energy change of each bunch over a turn is exactly compensated by the average kick it receives from the RF cavity holding voltage. The synchronous phase of a bunch is the phase of the RF voltage at which the energy kick equals this average. This is the equilibrium phase about which bunches oscillate longitudinally.

If we increase the charge in a single bunch, its synchronous phase will ride up the RF voltage waveform to keep up with the increasing loss of energy to wake fields. If we know the slope of the RF voltage, we can easily calculate the energy lost to wake fields per unit current from the synchronous phase increase. This gives us a measure of the integrated beam impedance, as we shall see in the next section. In itself, the integral reveals nothing about the shape of the impedance. However, some information about the frequency spectrum of the impedance can be gleaned from repeating this measurement at various bunch lengths.

If there are many bunches in the ring, with varying charge or spacing, then each bunch could see a different steady state wake voltage. Thus, the synchronous phase varies from bunch to bunch, if the fill is not azimuthally symmetric.

This chapter describes a new method for measuring the longitudinal impedance spectrum $Z(j\omega)$ using synchronous phase data from multi-bunch fills. Derivation of the necessary transfer functions is followed by experimental results from PEP-II HER (High Energy Ring) commissioning. The measurements suggest an explanation for the observation of

coupled-bunch instabilities at beam currents of the order of 100 mA (Sep., Oct. 1997). Multi-bunch synchronous phases are extracted from data taken using the HER longitudinal feedback system, which can digitally sample and record the phase of all bunches in the ring simultaneously [54, 55].

Synchronous phase transients need to be matched in the PEP-II HER and LER (Low Energy Ring), to achieve high luminosity. Matching of gap transients is complicated by incomplete knowledge of the impedance seen by the beam, or by distortions in the fill shape (beam current profile). Measurement and analysis of multi-bunch synchronous phases is therefore useful as a diagnostic tool during commisioning and normal operation of colliders. Such measurements are also very useful as a feedback diagnostic, since bunches with large synchronous phase offsets receive less than the nominal feedback gain.

3.1 Single-Bunch Synchronous Phase

As mentioned earlier, an electron bunch adjusts its synchronous phase ϕ_s so that the average kick it receives from the RF voltage cancels the average energy loss over a turn due to wake fields (V^{wk}) and synchrotron radiation (Uo). This gives the following equation:

$$V_c \sin(\phi_s) = Uo + V^{wk},$$

where V_c is the peak RF cavity voltage. The synchronous phase ϕ_s^o in the absence of wake fields is given by:

$$\phi_s^o = \sin^{-1}(Uo/V_c)$$

If ϕ_s^o is not very different from ϕ_s, i.e., if V^{wk} is small compared to V_c, we can write:

$$\phi = \phi_s - \phi_s^o \approx \frac{V^{wk}}{V_c \cos(\phi_s^o)}$$

$$\Rightarrow \phi \approx \frac{-V^{wk}}{|V_c \cos(\phi_s^o)|}, \tag{3.1}$$

since the PEP-II beams are above transition. For a single bunch of charge q, we have the following relations:

$$V^{wk} = q \sum_{k=0}^{\infty} W(kT_o)$$

$$
\begin{aligned}
&= q\int_{-\infty}^{\infty} W(t)\sum_{k=-\infty}^{\infty}\delta(t-kT_o)\ dt \\
&= q/T_o\int_{-\infty}^{\infty} W(t)\sum_{n=-\infty}^{\infty} e^{-jn\omega_o t}\ dt \\
\Rightarrow V^{wk} &= i_o\sum_{n=-\infty}^{\infty} Z(jn\omega_o), \qquad (3.2)
\end{aligned}
$$

where the longitudinal impedance $Z(j\omega)$ is the Fourier transform of the longitudinal wake-function $W(t)$, T_o is the revolution period, and i_o is the bunch current. On combining Eqs. 3.1 and 3.2, we get:

$$\phi/i_o \approx \frac{-1}{|V_c\cos(\phi_s^o)|}\sum_{n=-\infty}^{\infty} Z(jn\omega_o), \qquad (3.3)$$

The above equations contain the assumption that the bunch is vanishingly small in length. To take the finiteness of the bunch length into account we must replace $Z(j\omega)$ in Eq. 3.3 with $|F(j\omega)|^2 Z(j\omega)$, where $F(j\omega)$ is the Fourier transform of the normalised bunch line density.

It is thus evident that measurement of single-bunch synchronous phase versus bunch current yields an integral over the longitudinal impedance. Such a measuement yields no information on the shape of $Z(j\omega)$. Although we can improve the situation a little by repeating the experiment with varying bunch lengths (and therefore varying $F(j\omega)$), this method still falls short of directly quantifying resonances in the impedance.

The next section discusses the relationship between multi-bunch synchronous phases and the shape of the fill, which is used in this note to directly measure the longitudinal impedance spectrum $Z(j\omega)$ in the HER.

3.2 Multi-bunch Synchronous Phase

If a beam is filled with N bunches at an even spacing T_b, we can rewrite Eq. 3.1 as:

$$\phi_k \approx \frac{-V_k^{wk}}{|V_c\cos(\phi_s^o)|}; \qquad k = 0, 1, ...N\text{–}1 \qquad (3.4)$$

The steady state wake voltage V_k^{wk} seen by bunch k is the superposition of the wakes left by all preceding bunches:

$$V_k^{wk} \approx \sum_{m=0}^{\infty} q_{k-m} W(mT_b), \tag{3.5}$$

where q_{k-m} is the charge of the m-th bunch preceding bunch k. It should be noted here that the variation in ϕ_k from bunch to bunch results in a small amount of unevenness in the bunch spacing. This unevenness is neglected in computing the RHS of the above equation, so as to keep it linear. If f_{max} is the largest frequency of interest in $|F(j\omega)|^2 Z(j\omega)$, then the approximation is valid as long as $f_{max}\phi_k/f_{rf} \ll \pi/2$ for all k, where f_{rf} is the RF frequency (476 MHz at PEP-II).

Since the above equation is a discrete convolution, it is useful to rewrite it in terms of Fourier transforms. The DFT (Discrete Fourier Transform) of V_k^{wk} is defined as:

$$\mathbf{V}_n = \sum_{k=0}^{N-1} V_k^{wk} e^{-j2\pi kn/N}; \qquad n = 0, 1, \ldots N\text{-}1 \tag{3.6}$$

Using Eq. 3.5, this can be rewritten as:

$$\begin{aligned} \mathbf{V}_n &= \sum_{m=0}^{\infty} W(mT_b) \sum_{k=0}^{N-1} q_{k-m} e^{-j2\pi kn/N} \\ &= \sum_{m=-\infty}^{\infty} W(mT_b) e^{-j2\pi mn/N} \sum_{k=0}^{N-1} q_{k-m} e^{-j2\pi(k-m)n/N} \\ \Rightarrow \mathbf{V}_n &= \mathbf{W}_n \mathbf{Q}_n = T_o \mathbf{W}_n \mathbf{I}_n, \end{aligned} \tag{3.7}$$

where $\mathbf{Q}_n$ and $\mathbf{I}_n$ are the DFTs of q_k and i_k respectively, and $\mathbf{W}_n$ is the DTFT (Discrete-Time Fourier Transform) of $W(kT_b)$. This can be rewritten as a transfer function from i to ϕ using Eq. 3.4:

$$\frac{\mathbf{\Phi}_n}{\mathbf{I}_n} = \frac{-T_o}{|V_c \cos(\phi_s^o)|} \mathbf{W}_n, \tag{3.8}$$

where $\mathbf{\Phi}_n$ is the DFT of ϕ_k. Now all that remains is to express $\mathbf{W}_n$ in terms of the longitudinal impedance $Z(j\omega)$. The derivation shown in the previous section can be repeated to produce the following result:

$$\begin{aligned} \mathbf{W}_n &= 1/T_b \sum_{m=-\infty}^{\infty} Z[j(mN+n)\omega_o] \\ \Rightarrow \mathbf{W}_n &= Z_n/T_b, \end{aligned}$$

where Z_n is the longitudinal impedance aliased at a sampling rate of $1/T_b$. From this equation and Eq. 3.8, we get:

$$\frac{\Phi_n}{\mathbf{I}_n} = \frac{-N}{|V_c \cos(\phi_s^o)|} Z_n = \frac{-N}{|V_c \cos(\phi_s^o)|} \sum_{m=-\infty}^{\infty} Z[j(mN+n)\omega_o] \tag{3.9}$$

This is the multi-bunch analog of Eq. 3.3. The two equations are equivalent in the single-bunch case ($N = 1$). Again, we need to replace $Z(j\omega)$ with $|F(j\omega)|^2 Z(j\omega)$, if we want to take the finiteness of the bunch length into account. Of course, this need not be done if the bunch length is much smaller than the wavelength of the relevant resonances in $Z(j\omega)$.

If we know q_k and ϕ_k for all k, we can calculate the aliased longitudinal impedance Z_n using the above formula. As we increase N, the loss of information due to aliasing decreases. The best we can do is to fill every bucket, in which case N equals the harmonic number h. In practice, measurement noise prevents us from accurately calculating the entire aliased longitudinal impedance spectrum from just one multi-bunch measurement with $N = h$. However, as we shall see in the next section, it is possible to measure really large impedance resonances at revolution harmonics that are excited by the shape of the multi-bunch fill.

3.3 Experimental results

In this section we present a retrospective analysis of multi-bunch synchronous phases and bunch currents (fill shapes), using data from measurements made during the PEP-II HER commissioning run from Sep. to Oct. 1998. The measurements were originally aimed at identifying unstable coupled-bunch modes and quantifying noise-driven beam motion, and are therefore not ideally suited to our present purpose of estimating Z_n. In particular, our knowledge of fill shapes during the run comes only from the serendipitous presence of 60-Hz harmonics in the klystron output, which makes crude current monitoring possible. We will see, however, that it is still possible to obtain useful information about the longitudinal impedance from the available data. Bunch current monitoring using the feedback system was first demonstrated at the ALS [71]. Since then, variants of the method described in this section have been used at all of the machines that use the PEP-II-ALS-DAΦNE feedback system.

A sample measurement of the PEP-II HER and LER gap transients during collision (Feb. 1999) is also shown.

A typical PEP-II data set consists of around 660 samples (one every 6 turns) of the phase of each bunch in the ring. The fill pattern being examined here consists of 291 bunches, at an even spacing of 12 RF buckets between bunches. The front end phase measurement contains a gain that is proportional to bunch current. Multi-bunch synchronous phases are therefore calculated by averaging the digitised signals for each bunch and dividing the averages by the corresponding bunch currents.

Line harmonics from the klystron impose the same low-frequency motion on all the bunches. During the commissioning run from Sep. to Oct. 1998, 360-Hz and 720-Hz lines from the klystron were large enough to be detected in the bunch data. These spectral lines afford a crude current monitor, since the bunch signals are proportional to charge times longitudinal phase[1]. Bunch currents are estimated by projecting individual bunch signals onto a line harmonic spectrum calculated by averaging over all the bunch signals.

3.3.1 Time Domain

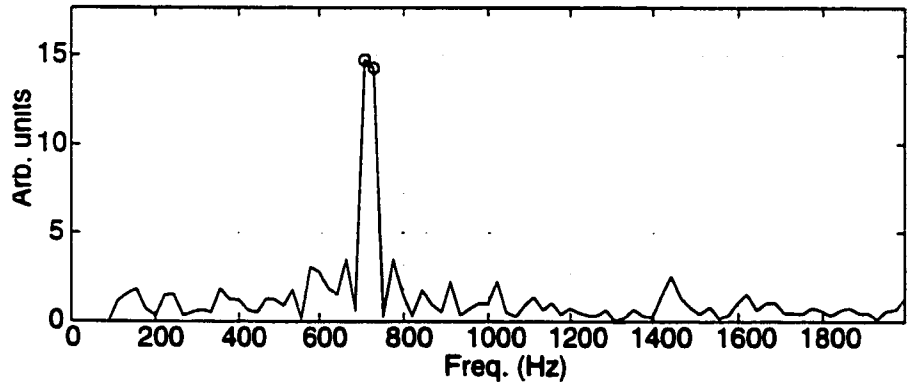

Figure 3.1: Averaged bunch signal spectrum up to 2 kHz, for a 291-bunch fill with a total current of 122 mA. Two circles mark the 720-Hz line, onto which bunch signals are projected for current monitoring.

Figure 3.1 shows the averaged low-frequency bunch signal spectrum for a 291-bunch 122-mA fill. In this case we calculate bunch currents by projecting individual bunch signals onto the 720-Hz line in the averaged spectrum, and then scaling the result so that the calculated

[1] Actually, the signal of bunch k is proportional to $i_k \sin[H(\phi_k + \phi_s^o - \phi_{m.o.})]$, where H is the harmonic of the RF frequency at which phase detection is performed, and $\phi_{m.o.}$ is the phase of the master oscillator phase reference (see Eq. 2.46). However, the assumption is reasonable, since $i_k \sin[H(\phi_k + \phi_s^o - \phi_{m.o.})] \approx i_k \phi_k \cos[H(\phi_s^o - \phi_{m.o.})]$, when $H\phi_k \ll \pi/2$.

total beam current agrees with that measured by the DCCT (DC Current Transformer). The bunch currents i_k so obtained are shown in Figure 3.2(a). There is a step discontinuity in the fill at the 175th filled bucket ($k = 175$). Figure 3.2(b) shows the averaged bunch signals in units of ADC (Analog to Digital Converter) counts. The synchronous phase graph, calculated by dividing the average signal by the bunch currents (and by a calibration factor), is shown in Figure 3.2(c). The step in the fill can be seen to cause the synchronous phase to ring a few times and then drift back to its initial value. This is evidence of a resonance in the impedance $Z(j\omega)$.

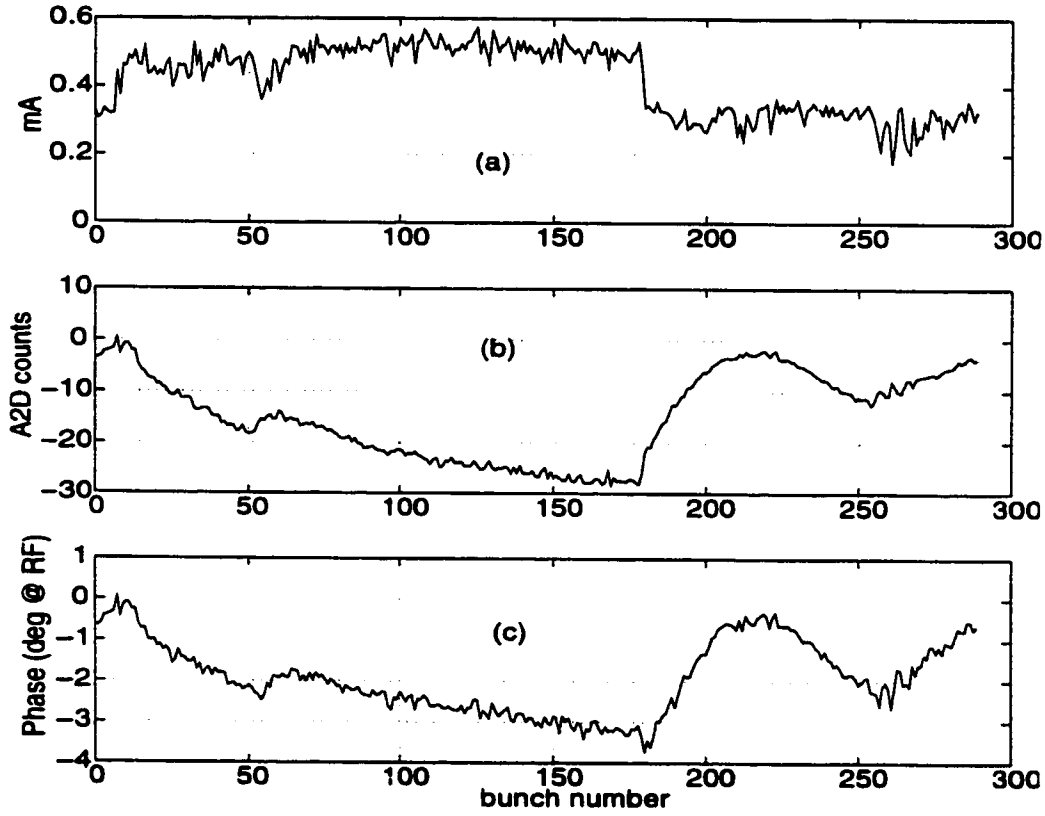

Figure 3.2: (a) Bunch current measurement, 291-bunch fill, 122 mA total current. Note the step discontinuity at the 175th filled bucket. (b) Raw mean bunch phase, ADC counts. (c) Synchronous phase variation around the ring. The discontinuity in the fill causes ϕ_s to oscillate with a peak-to-peak amplitude of 3° at the RF frequency.

In the time domain we would like to observe the response of the synchronous phase to impulses in the fill shape. Figure 3.3(a) shows a 96-mA 291-bucket fill with a discontinuity

that is impulsive at low frequencies. The resultant synchronous phase ringing is shown in Figure 3.3(c). We can see from the figure that the "impulse response" goes through about three oscillations and dies out in one revolution period. This indicates that $Z(j\omega)$ has a strong resonance three revolution harmonics away from some multiple of the bunch frequency, which is a twelfth of the RF frequency in this case.

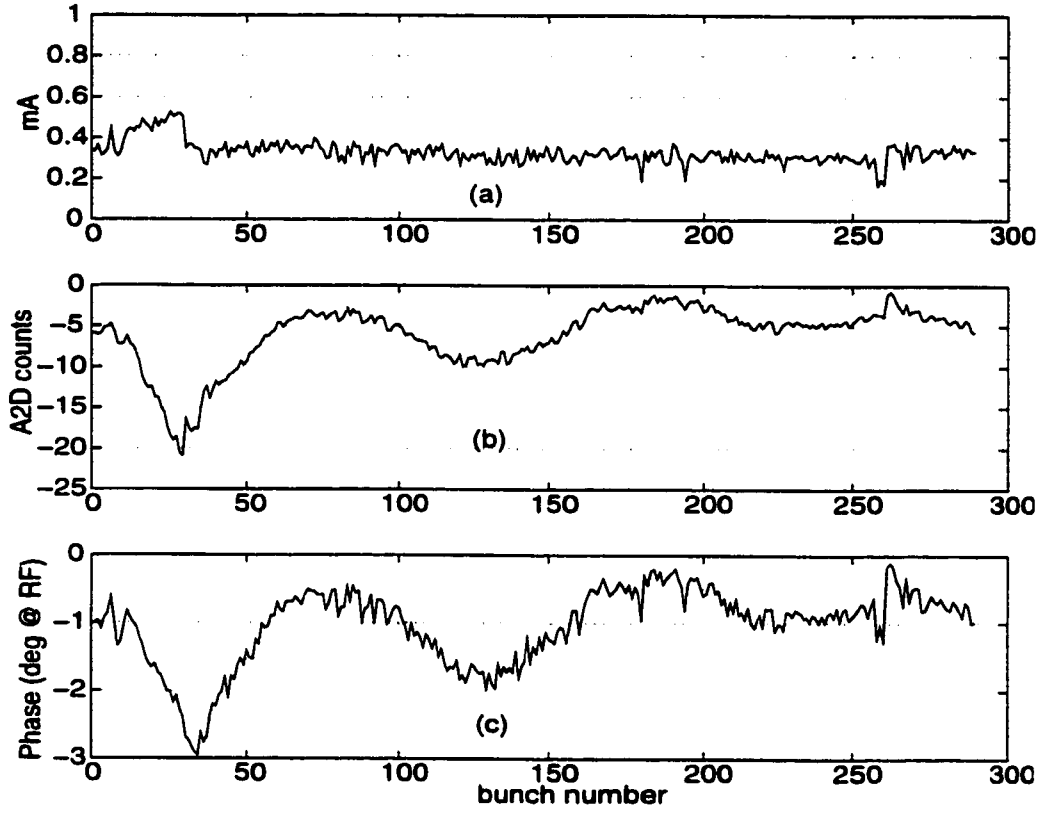

Figure 3.3: (a) Bunch current measurement, 291-bunch fill, 96 mA total current. The discontinuity in the fill looks like an impulse at low frequencies. (b) Raw mean bunch phase, ADC units (c) Synchronous phase variation around the ring. The "impulse response" contains about three oscillations in one revolution period.

3.3.2 Transfer Function

We are now ready to take the DFTs of i_k and ϕ_k and calculate the transfer function from i_k to ϕ_k. As shown in Eq. 3.9, the scaled transfer function is the aliased longitudinal impedance Z_n. Of course, we should only calculate Z_n at revolution harmonics n that have a reasonably

good signal to noise ratio (SNR). By looking at the shape of the fill in Figure 3.2(a) and Figure 3.3(a), we can tell that $\mathbf{I}_n$ is relatively large at the first few revolution harmonics and small elsewhere. If we calculate Z_n wherever $\mathbf{I}_n$ is above the noise floor, we should therefore expect to have a reasonable estimate for small values of n.

It must be pointed out here that the DC synchronous phase is not known, since it is canceled in the feedback front end by a DC offset designed to prevent the phase signal from saturating the digitiser. This precludes the calculation of Z_o from the data presently available. Z_o can be measured by keeping the offset fixed at a nominal value and varying the total beam current.

Transfer functions have been calculated from 15 different sets of data. The resulting impedance estimates are consistent to within 20% of each other. Repeatability of the impedance measurements will be improved by more accurate bunch current monitoring. Figure 3.4 shows the estimate of Z_n obtained by averaging transfer functions from four consecutive data sets with similar fill shapes. The aliased impedance is calculated only for the first four values of n, since the excitation is close to the noise floor everywhere else. As expected, there is a strong resonance in Z_n at $n = 3$, with $Z_3 = 8.6M\Omega$.

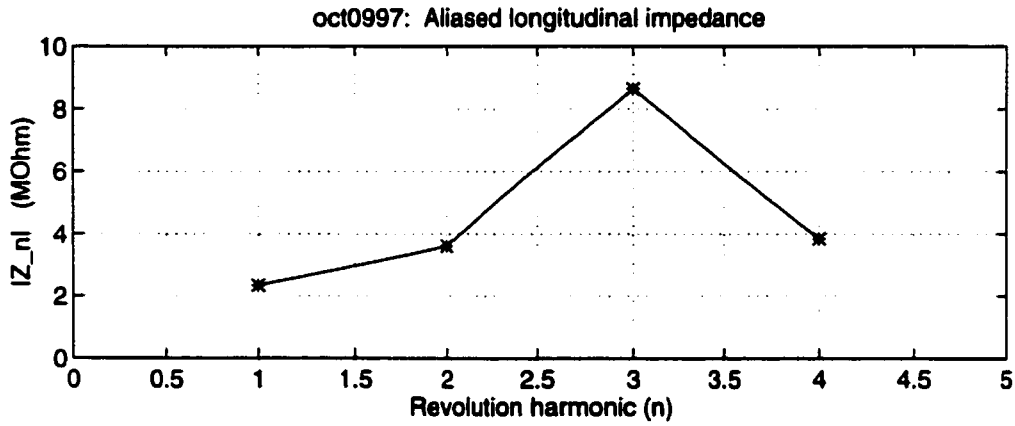

Figure 3.4: Estimate of Z_n obtained by averaging transfer functions from four consecutive data sets. The impedance is in the $M\Omega$ range due to the fundamental resonance of the parked cavities. The cavity resonant frequencies seem to be closer to $3f_o$ than $2f_o$.

The HER has 20 installed RF cavities, each with a loaded shunt impedance of $761k\Omega$. The experiment was performed with eight active cavities, tuned about 5 kHz away from the RF frequency. This detuning is small compared to the revolution frequency, which is 136.3 kHz. Six idle cavities were nominally parked exactly halfway between the second and third

revolution harmonics above f_{rf}, while the other six occupied a symmetric location below f_{rf}. If, however, they were all parked exactly three revolution harmonics away from f_{rf}, their impedances would add up to $9.2M\Omega$ at $n = 3$. The asymmetry between Z_2 and Z_3 in Figure 3.4, together with the fact that $Z_3 = 8.6M\Omega$, indicates that the 12 idle cavities were indeed parked closer to the third revolution harmonic than the second.

3.3.3 Coupled-Bunch Instability

Ideally, idle cavities should be parked symmetrically around f_{rf} so that they do not drive coupled-bunch instabilities. The impedance estimates shown in Figure 3.4 suggest that they might not have been parked accurately. This conclusion is also borne out by the fact that coupled-bunch modes 2 and 3 were sometimes seen to be unstable. Figure 3.5 shows the beam pseudospectrum [72] (beam spectrum without revolution harmonics, calculated from digitised data) for a 291-bunch 84-mA fill, taken a few days before the data displayed in the previous figures. The pseudospectrum shows that mode 3 is unstable, with a steady state amplitude above 2° at the RF frequency.

The conclusion that incorrect parking of the idle cavities was responsible for the mode 3 instability has been borne out by subsequent measurements [73].

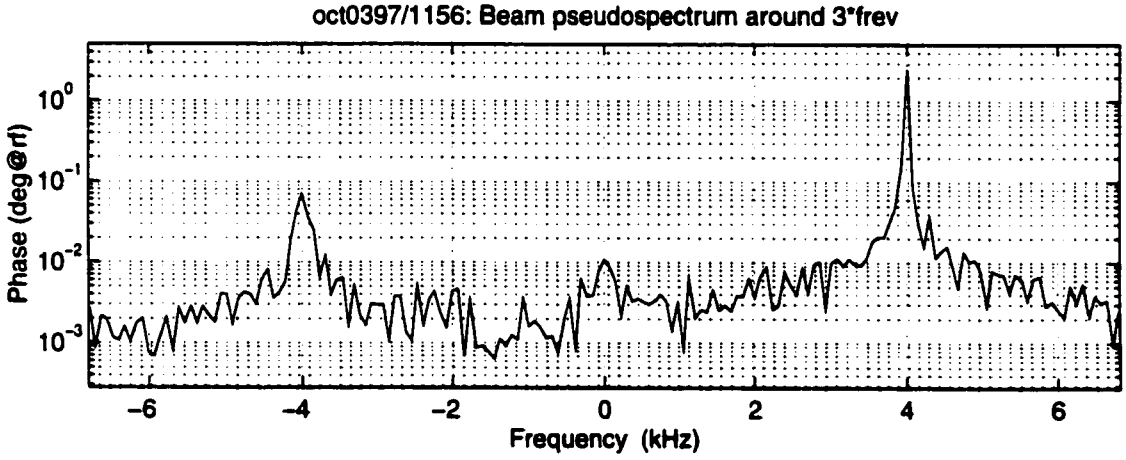

Figure 3.5: Beam pseudospectrum for a 291-bunch 84-mA fill, taken a few days before the data displayed in previous figures. The pseudospectrum shows that mode 3 is unstable, with a steady state amplitude of 2° at the RF frequency.

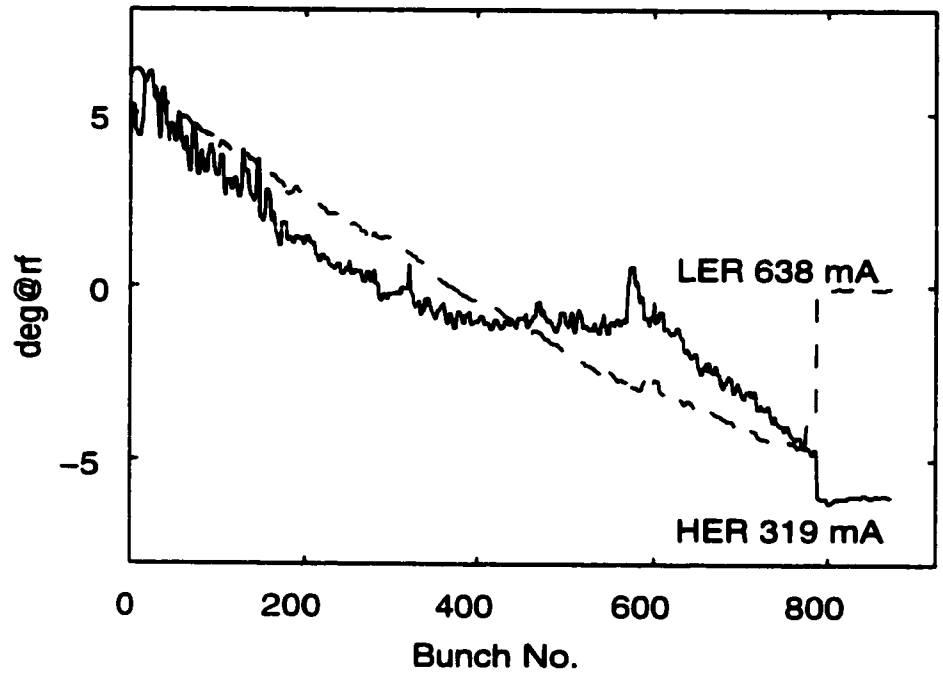

Figure 3.6: Synchronous phase transients in the PEP-II HER and LER, measured during collision (10% gap).

3.3.4 Matching of PEP-II Gap transients

At the design current, the PEP-II gap transients are of the order of 10-20 RF degrees. Although this makes the problem of feedback harder to solve, it does not impact the luminosity, as long as the HER and LER transients are the same. Transient matching between the two rings is a non-trivial task, requiring diagnostic measurements of the actual bunch currents and synchronous phases.

Figure 3.6 shows a measurement of HER and LER synchronous phase transients during commissioning, with the HER at 319 mA and the LER at 638 mA. Every fourth RF bucket is populated in this case. The flat portions at the ends of the two traces coincide with the 10% gap at the end of each bunch train. Since the feedback system receives no beam signal in the gap interval, these flat segments merely reflect the offsets in the processing.

We see that the gap transients in the two rings are mismatched by up to 2 deg@RF, possibly due to irregularities in the fill shape[2]. Although mismatches of this magnitude

[2]The DC component of the synchronous phase transient is canceled out in the LFB front end. Hence, DC phase mismatches between the two rings are not detected in this measurement.

have no noticeable effect on the PEP-II luminosity [74], it must be kept in mind that this measurement was made at a third of the design current, and that larger mismatches are possible at higher currents.

3.4 Summary

A novel beam-based technique for measuring the longitudinal impedance spectrum $Z(j\omega)$ has been demonstrated at PEP-II. The technique involves calculation of the transfer function from fill shape to multi-bunch synchronous phase. Bunch currents and synchronous phases have been extracted from a retrospective analysis of data taken using the longitudinal feedback system during PEP-II HER commissioning.

The presence of line harmonics in the klystron output has been used to extract bunch currents from feedback system data. Impulsive discontinuities in the fill were seen to cause the synchronous phase to ring at thrice the revolution frequency. The corresponding transfer function agrees well with the impedance of the parked cavities, if we assume that they were tuned closer to $f_{rf} + 3f_o$ than to $f_{rf} + 2f_o$.

Our ability to measure the longitudinal impedance in these retrospective data sets has been limited mainly by noise in the bunch current measurements. Cleaner current monitoring by injecting low-frequency signals into the feedback front end has been implemented, and would improve the situation significantly. Of course, this method is not suitable for measuring small impedances, which produce synchronous phase transients that are much smaller than a 1 deg@RF.

Specific regions of the impedance spectrum can be explored by adjusting the fill shape to excite the targeted revolution harmonics. For example, we could investigate the impedance around $100f_o$ with a 582-bucket fill by creating a periodicity of approximately 5.82 buckets in the fill. We could get a good measurement by injecting a little extra charge into every sixth bucket (i.e. every 18th bucket at 238 MHz).

In colliders, measurements of multi-bunch synchronous phases are also useful in estimating the loss in luminosity due to shifts in the collision point.

Chapter 4

Measuring Instability Growth and Damping Rates

When a beam exhibits unstable coupled-bunch motion, one needs to perform diagnostic measurements to identify the source of the problem, before deciding upon a cure. The goal of instability diagnostics is to identify the modal pattern of the instability, the growth rate, the coherent tune shift, the extent of nonlinearities, and the correlations of all of these observables with conditions such as cavity temperatures, cavity tuner positions, vacuum pressure, fill shape (bunch current profile), synchronous phase variation, etc. If a feedback system is supposed to damp the instabilities, it is also necessary to perform feedback diagnostics, to track down potential non-idealities and sources of gain reduction.

Chapter 3 described a method for using LFB data to measure the fill shape, which has a direct effect on beam stability, and the multi-bunch synchronous phase transient, which has an effect on stability because it affects feedback gain.

The modal structure of an unstable beam can be qualitatively studied by examining the beam spectrum once nonlinearities limit mode growth [75, 76]. It is more useful, though, to quantify instabilities in the linear small-oscillation region, since this directly yields the impedance of external resonant structures (see Eqs. 2.31 and 2.37). Such measurements have been made by recording the output of narrowband filters tuned to the frequencies of the relevant modes while oscillations grow [27]. Unfortunately, a machine with hundreds or thousands of bunches, and external resonances driving a large number of unstable modes, would pose practical problems for such an approach. Even in machines with a small number of unstable modes, this technique has the disadvantage of requiring prior knowledge of the

location of the most unstable mode.

In this chapter, an offline modal decomposition of digitised data is used to study the coupled longitudinal motion of electron bunches in storage rings. The technique allows simultaneous measurement of linear-region growth rates and (feedback induced) damping rates of all unstable beam modes via a time-domain transient technique. Such measurements have been useful in identifying the ALS cavity temperatures most conducive to longitudinal stability. The damping rates of naturally stable modes are measured using externally excited transients. Longitudinal growth and damping rates at the ALS and PEP-II, measured under a variety of beam conditions, are compared to projections based on the estimated cavity impedance [77, 78]. A measurement of unstable longitudinal modes at SPEAR is shown, as well as a sample grow-damp measurement of horizontal instabilities at the ALS.

All of the data displayed in this chapter were digitised using the programmable LFB system described in Ch. 2. The transverse measurement was made by gating the transverse feedback signal, and simultaneously recording the transverse front end signal in the LFB DSPs. See [51] for a description of the ALS transverse feedback system.

Brief recapitulation: The eigenmodes of longitudinal motion of the bunch centers in the case of an azimuthally symmetric beam with N evenly spaced bunches of equal charge consist of all bunches oscillating at the same amplitude. The N eigenmodes are characterised by the angular separation in phase space between successive bunches, which is a constant in any given mode. At a fixed azimuth in the ring, mode l produces a beam signal at $(l+\nu_s)f_o$ (upper sideband) and at $(N-l-\nu_s)f_o$ (lower sideband). This pattern is repeated with a periodicity of Nf_o, if the bunches are assumed to be point-like.

4.1 ALS Measurements

The ALS is a synchrotron light source, which stores electrons at 1.0–1.9 GeV. It has demonstrated longitudinal coupled-bunch instabilities since commissioning in 1993 [79].

Table 4.1 lists the relevant parameters of the ALS machine and the ALS LFB system. The high bandwidth required for simultaneous tracking of all beam modes is achieved by a bunch sampling rate of 500 MHz, at the ALS. After downsampling, the feedback signal is calculated in a digital processing array composed of 40 processors. The processors implement a bunch-by-bunch discrete-time FIR algorithm at an aggregate multiply-accumulate rate of 1.6×10^9 operations/s.

Table 4.1: ALS Parameters

Parameter	Description	Value
E_o	Beam energy	1.5 GeV
f_{rf}	RF frequency	499.65 MHz
h	Harmonic number	328
f_o	Revolution frequency	1.5233 MHz
α	Momentum compaction factor	1.594e-3
I_o	Design current	400 mA
f_s	Nominal synchrotron frequency	11.3–12 kHz
d_r	Radiation damping rate	74 s^{-1}
-	Bunch sampling rate	499.65 MHz
D	Downsampling factor	21–31
P	Feedback output power	200–500 W
-	Output amplifier bandwidth	1–2 GHz

Growth and damping rates are extracted from "grow-damp" measurements made by switching feedback off, allowing unstable modes to grow spontaneously from the noise floor, and then turning feedback on again to damp them back to the noise floor. This sequence of events spans only a few tens of milliseconds. Naturally stable modes are studied by exciting them through the external drive input of the feedback system and observing the resulting decay transients. The bunch oscillations stay linear throughout.

The results presented in this section are for a 320-bunch beam followed by an eight-bucket gap. Charge variation within the populated RF buckets is of the order of 15%. However, we will project the recorded beam motion onto the symmetric-beam eigenmodes, since they are well known and simple. It will become apparent later in this chapter that this projection is quite useful, in spite of the unevenness in the fill[1].

4.1.1 Signal Processing

A typical ALS data set consists of around 1000 samples (one every D turns, where D is the downsampling factor) of the phase of each of the 320 bunches. The first step in modal analysis is to filter the data with a band pass filter centered at the synchrotron frequency, to improve the signal-to-noise ratio.

Due to the requirements of downsampling, the bunches are not all sampled on the same

[1]In Ch. 7, it is shown that the set of even-fill eigenmodes forms a natural basis for observing coupled-bunch instabilities, even when fill unevenness affects the eigenvalues and eigenvectors.

turn. In addition, even those bunches that are sampled on the same turn are sampled at different instants, since they don't all fly past the BPM at the same time. Hence, the second step is to delay or advance the bunch signals using appropriate phase shifts in the frequency domain, so that we approximate simultaneous sampling on every D-th turn. Transverse feedback signals that are downsampled and recorded by the LFB system need special care at this stage of the processing, since they are usually aliased to a lower frequency. A further complication arises when aliasing converts an upper betatron sideband into a lower sideband. The data analysis programs have been designed to compensate for these effects, when given a rough estimate of the unaliased betatron tune (for details, see the appendix).

Although real bunch oscillations vary as $\cos(\Omega t)$, it is convenient to represent bunch oscillations as phasors of the form $e^{j\Omega t}$. This is accomplished by masking negative frequencies in the discrete Fourier transform (DFT) of the individual bunch signals. Instabilities are usually studied in this form, with beam spectra that consist only of upper sidebands.

It is evident from Eq. 2.24 that the projection onto symmetric-beam modes is achieved by taking the DFT of all the bunch phases at a single instant, on a single turn. The empty buckets are assigned a phase of zero. The strength of the symmetric-beam modes is tracked by observing the magnitude of the DFTs over time.

4.1.2 Longitudinal

Fig. 4.1 illustrates the measurement of growth rates of unstable modes with and without feedback, using the grow-damp technique. The 238-mA beam is initially stable under the action of negative feedback. At $t = 0$ ms, the feedback system is turned off under software control, and the exponential growth of unstable modes begins. At $t = 7$ ms the feedback is turned on again, and the oscillation amplitudes damp back to their initial steady state level. The bunch motion is recorded in a dual-port memory which is read by an external processor. The data is then processed offline. After the data is read, the DSP processors can be triggered again to record another transient.

Fig. 4.1(a) shows the envelopes of the longitudinal oscillations of the 320 bunches. We see growth of unstable oscillations up to $t = 7$ ms, followed by damping of the motion. The envelopes do not contain information about the phase relationship between individual bunch oscillations. They do not grow or damp exponentially, indicating that the bunches are oscillating in a superposition of two or more coupled-bunch eigenmodes.

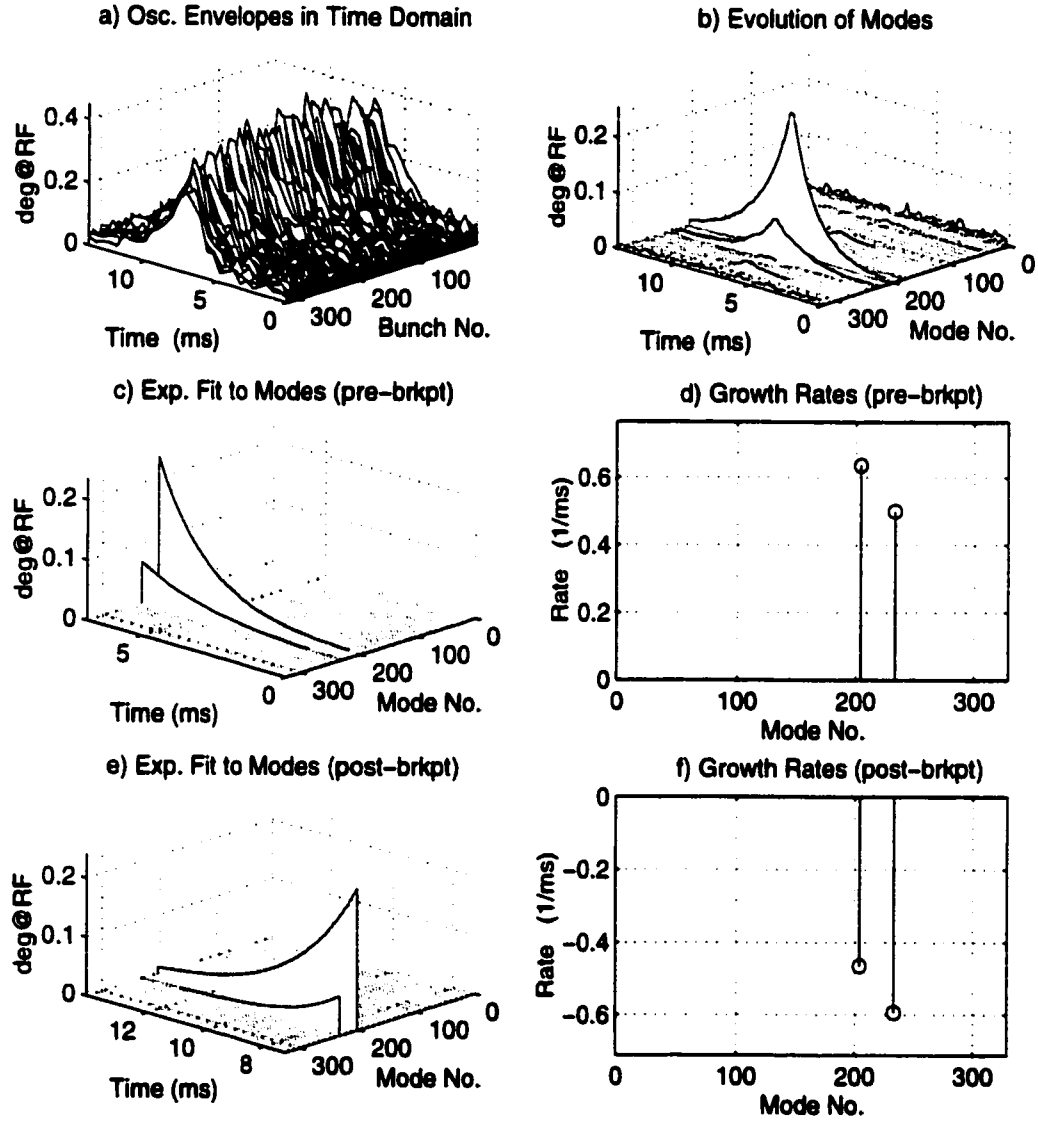

Figure 4.1: ALS grow-damp measurement. a) Bunch oscillation envelopes. b) Modal amplitudes. c,d) Exponential fits before break point yield growth rates without feedback. e,f) Exponential fits after break point yield feedback-induced damping rates.

The unstable modes of oscillation are revealed by taking turn-by-turn DFTs of the bunch phases, as described earlier. Fig. 4.1(b) traces the evolution over time of the 328 "modes" during the grow-damp transient. Although the fill is slightly uneven, we see from this figure that the projection onto symmetric-beam eigenmodes is successful in dramatically simplifying the picture. It is apparent that the beam motion is the result of exponential growth and damping of unstable modes 204 and 233.

The growth rates of the modes were found by curve fitting to be 0.64 ms^{-1} and 0.5 ms^{-1} respectively [Fig. 4.1(c),(d)]. Exponential fits to the tails of the transients in Fig. 4.1(b) revealed damping rates of 0.46 ms^{-1} and 0.59 ms^{-1} respectively [Fig. 4.1(e),(f)]. We see here the action of the feedback system in turning the net growth rate from positive to negative. The growth rates correspond to effective cavity impedances of 67 kΩ and 84 kΩ respectively (see Eq. 2.31). We can see that the action of feedback shifts the two open-loop growth rates down by the same amount (approximately 1.1 ms^{-1}). This is consistent with the expectation that bunch-by-bunch feedback should damp all symmetric-beam modes equally. From Eqs. 2.42 and 2.44, we can now estimate that the feedback system has a gain of 97 V/ps, or 31 V/mrad, and an effective damping impedance of 130 kΩ.

For the measurement of naturally stable modes, a narrowband excitation at the desired mode frequency is injected into the feedback system at the external drive input (see Fig. 2.5). This excitation is impressed on the beam through the power amplifier and kicker, and bunch motion at the desired frequency is excited (a single longitudinal mode). When the excitation and feedback are turned off, the excited mode decays at its natural rate. The growth rate of the mode becomes more negative when feedback is turned on again. The transient is recorded and processed as before. Fig. 4.2 shows a measurement of mode 161, which is naturally stable. At $I_o = 145$ mA, the natural growth rate of the mode is -0.07 ms^{-1}, which is close to what is expected from radiation damping, and the feedback-induced growth rate is -1.14 ms^{-1}.

In this experiment, the 200-W ALS feedback power amplifier has been used. The previously described experiment was performed with a borrowed 500-W amplifier. Thus, for identical DSP gains, the feedback impedance in this case should be lower by a factor of $\sqrt{5/2}$. However, the gain in the DSPs is greater by a factor of 2 in this case. The two factors combine to give an expected feedback impedance of $130 \times 2 \times \sqrt{2/5} = 164$ kΩ. However, the growth rate shift in this case implies a damping impedance of 207 kΩ (see Eqs. 2.42 and 2.44). The slight discrepancy between these two numbers is explained by the

feedback-induced tune shifts in the two cases, as discussed in Ch. 5.

If the above measurement is repeated for several modes (or an excitation is applied to several modes simultaneously) the gain of the feedback system can be measured as a function of frequency. This is a useful system check, since it can be used to examine the combined frequency response of the power amplifier and kicker.

Growth and damping rates of unstable modes were measured at currents ranging from 60 to 250 mA at a variety of cavity temperatures. The growth rates varied by up to a factor of 3 due to cavity temperature changes alone. Small changes in the beam current over a few minutes occasionally caused significant changes in growth rates, possibly due to movement of the cavity tuners. Open-loop (no feedback) damping rates of a number of stable modes were measured by exciting them to a measurable amplitude and tracking their decay.

It is instructive to compare measured growth rates to predictions based on the measured impedance $Z(\omega)$ of a model RF cavity [77] and Eq. 2.31. We expect to see unstable modes where the measured cavity resonances land on upper sidebands of revolution harmonics. It must be kept in mind, however, that the cavities installed in the ring could differ from the model cavity in geometry, temperature and tuner position. Another potential source of deviations from the expected modal structure is the difference between the exact resonance frequencies of the two installed RF cavities, which could blur the effective impedance seen by the beam.

The comparison between measured and predicted growth rates (at 100 mA, without feedback) is shown in Fig. 4.3. The cavity-induced growth rate is antisymmetric about 250 MHz, but the addition of radiation damping breaks this symmetry. Of the cavity resonances in the figure that could potentially drive instabilities at currents up to 400 mA, only two have been seen to do so at the nominal cavity temperature and filling pattern[2]. These resonances drive modes 204 and 233, which appear as 'x's at 311 MHz and 355 MHz respectively, in Fig. 4.3. Their growth rates have been measured at 15 different beam currents, and normalised to a current of 100 mA. Error bars for the measured positive growth rates are one standard deviation wide on each side. The sharpest resonance is the TM-011 mode at 808 MHz, which is aliased down to 308 MHz. This resonance correlates fairly well with the instability at mode 204, although an exact calculation places it between modes 202 and 203. The correspondence is greater between mode 233 and the resonances

[2]The situation has changed quite a bit with the introduction of passive Landau cavities, which have larger HOM resonances than the active cavities. The study of instability and feedback issues in the presence of the Landau cavities is ongoing [80]

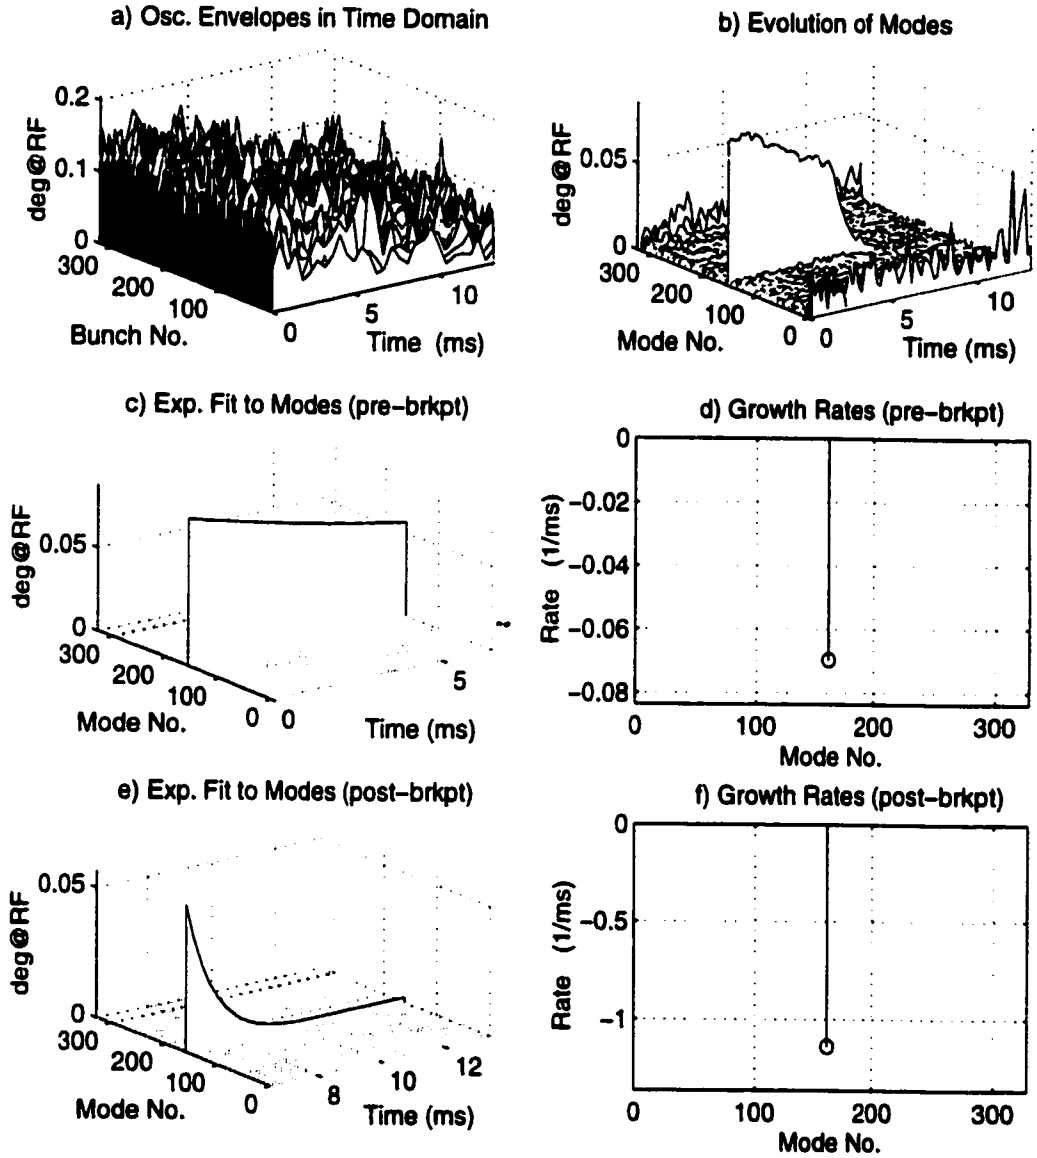

Figure 4.2: Example of external transient technique for measuring (negative) growth rates of naturally stable modes with and without feedback. Mode 61 is externally excited and then allowed to decay naturally until $t = 7$ ms, at which point feedback is turned on and the mode is rapidly damped.

at 355 MHz (aliased from 2.3 and 2.8 GHz). Neither of the unstable modes shows the worst-case growth rate, which implies that the cavity resonances do not land exactly on a mode frequency.

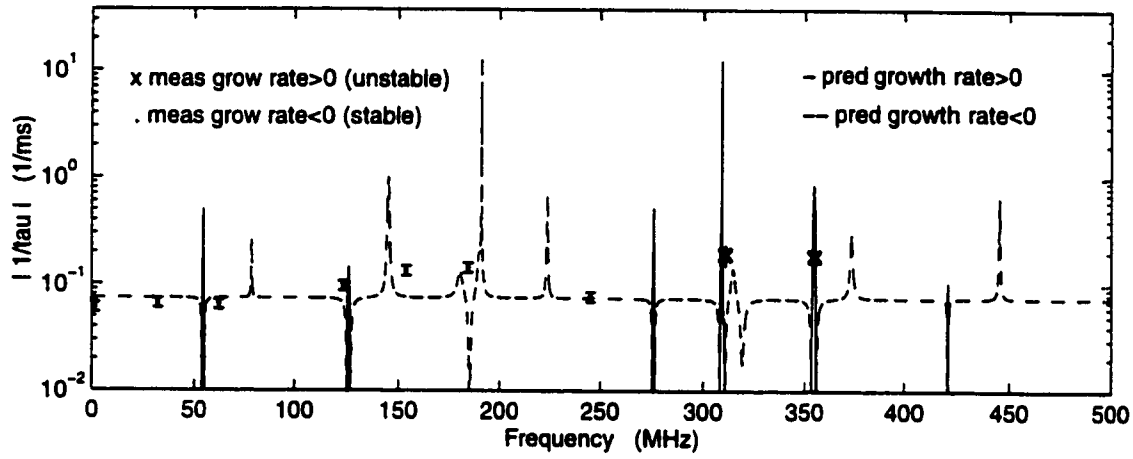

Figure 4.3: Log plot of experimentally measured and predicted growth rates as a function of modal frequency, with error bars around measured points. The predicted rates are based on the radiation damping rate and measurements of the impedance of a model RF cavity.

The measured (negative) growth rates of naturally stable modes are mostly close to the radiation damping rate (Fig. 4.3). These rates are less sensitive to variations in cavity temperature and tuner position, since the modal frequencies do not correspond to large cavity resonances. The corresponding error bars are conservatively estimated at 10%.

When the temperature of an RF cavity is changed, so is its size, and consequently, the resonant frequency of each parasitic mode. The HOM frequencies are also affected by the cavity tuners, which are adjusted to keep the beam loading small at the new temperature setting. Thus, the temperatures (T1 and T2) of the two ALS cavities constitute knobs with which to reduce the overlap between the HOM impedance and the modal frequencies.

The temperature of the cooling water exiting the ALS cavities can be controlled with a precision of around 0.5°C. Grow-damp measurements were used to study the behaviour of longitudinal coupled-bunch growth rates, as the cavity water temperatures were swept from 40°C to 48°C, in steps of 1°C. T2 was kept at 48°C while T1 was swept, and T1 was kept at 45°C while T2 was swept. Figure 4.4 shows the measured growth rates, scaled to a beam current of 100 mA. We see that the maximum growth rate decreases, as the

two temperatures are increased. Of course, transverse HOM frequencies also move with temperature. As a compromise between longitudinal and transverse requirements, the ALS cavities are kept at the highest transversely stable temperature.

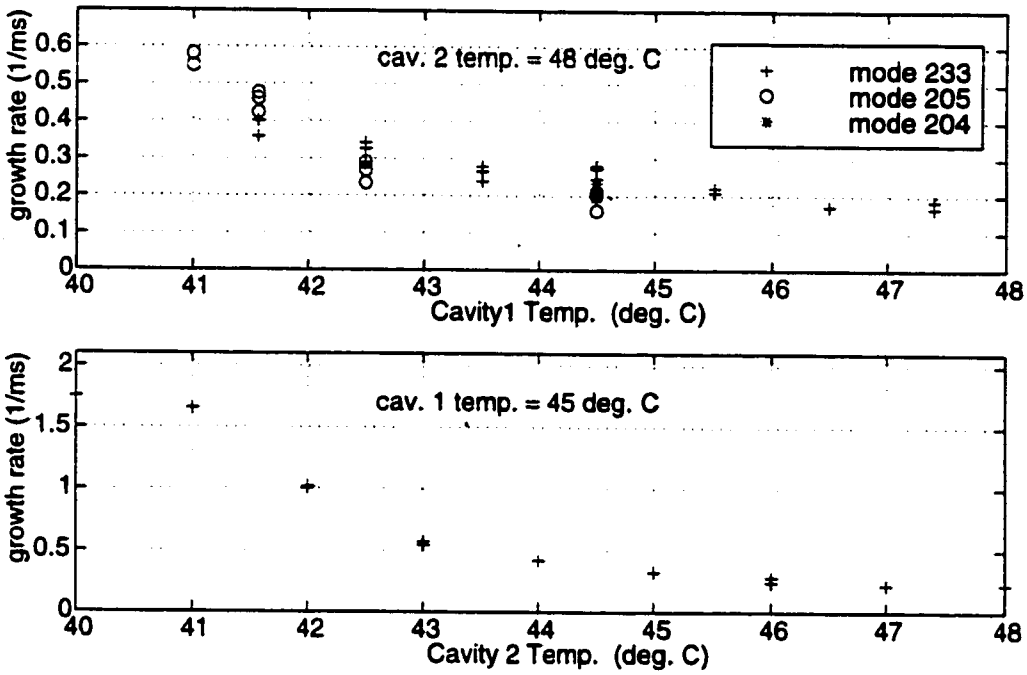

Figure 4.4: Longitudinal growth rates vs. cavity temperature at the ALS, $I_o = 100$ mA.

Another way of analysing transient data is to concatenate the sampled bunch phases over a few ms, and take the DFT of the resulting vector. The resulting spectrum, which was originally called a "pseudospectrum" [72], is similar to a conventional BPM signal spectrum. The main difference is that synchrotron sidebands are easier to identify in the pseudospectrum, since revolution harmonics are suppressed. This suppression comes from removal of the DC component of the bunch signals in the LFB front end, as well as in offline data processing. The distance of the sidebands from the revolution harmonics gives us the frequency shifts of the modes, and the width of the sidebands gives us an equivalent way of calculating their growth rates. See the appendix, for details of the algorithm used for data analysis.

The beam pseudospectrum resulting from a single 4–8-ms transient covers the entire 500-MHz range of the modes (0 to f_{rf}) with a resolution of 250–125 Hz. A heterodyned spectrum analyser would take at least a few minutes to perform the 328 narrowband sweeps

required to produce an equivalent spectrum, by which time the oscillations would have reached a damped or saturated steady state.

Fig. 4.5(a) shows the magnitude of the DFT of the bunch-phase signal over the last 4 ms of mode growth (Δf = 250 Hz). Only the section from 0 to $f_{rf}/2$ is shown, since the magnitude of the pseudospectrum is symmetric about $f_{rf}/2$. This is compared to the real part of the effective impedance of the RF cavity [Fig. 4.5(b)]. Radiation damping is converted to an effective impedance using Eq. 2.31. The two plots agree in several respects. The two largest impedance peaks drive longitudinal motion at (or close to) their aliased resonant frequencies. They correspond to the two prominent modes in Fig. 4.1. The impedance plot suggests that a few more cavity resonances should poke up above radiation damping, but this has never been observed, possibly for the reasons mentioned above.

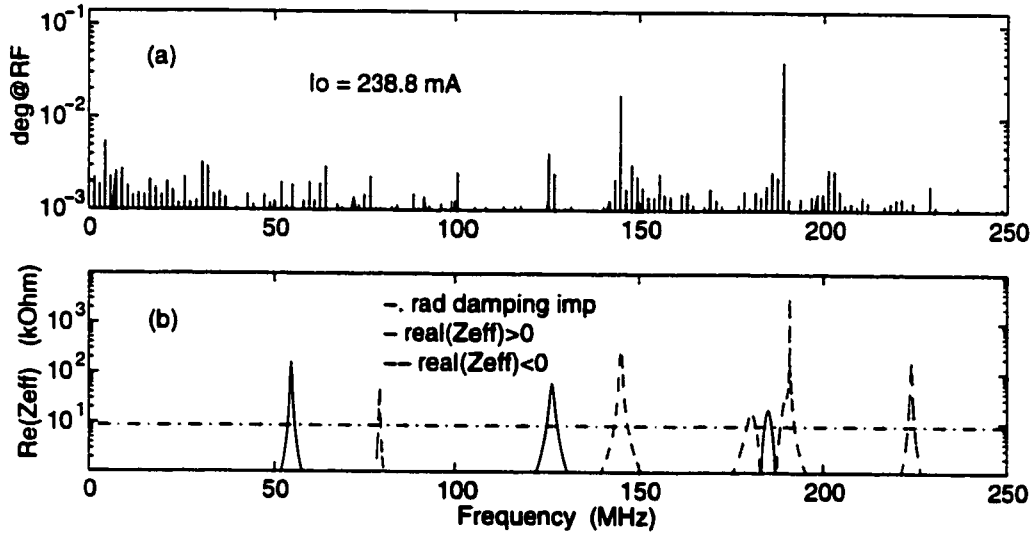

Figure 4.5: (a) Magnitude of high-resolution DFT of growing transient in Fig. 4.1, showing all the spectral components from 0 to 250 MHz. (b) Real part of Z^{eff}.

Fig. 4.6 zooms in on two 90-kHz sections of the 500-MHz pseudospectrum, the first 250 MHz of which are shown in Fig. 4.5(a). These sections contain the two largest upper sidebands in the pseudospectrum. The lower sidebands of the same modes are visible in Fig. 4.5(a), at 189 MHz and 145 MHz. Fig. 4.6 shows that the most unstable mode is an upper sideband at $204f_o$ (and therefore a lower sideband at $328 - 204 = 124f_o$), with a

linewidth corresponding to the previously measured growth rate. The lower sideband at $204f_o$ is damped to the noise floor by the TM-011 cavity mode. The next most unstable mode is an upper sideband at $233f_o$ (and therefore a lower sideband at $95f_o$). The smaller linewidth of this mode is consistent with its smaller growth rate.

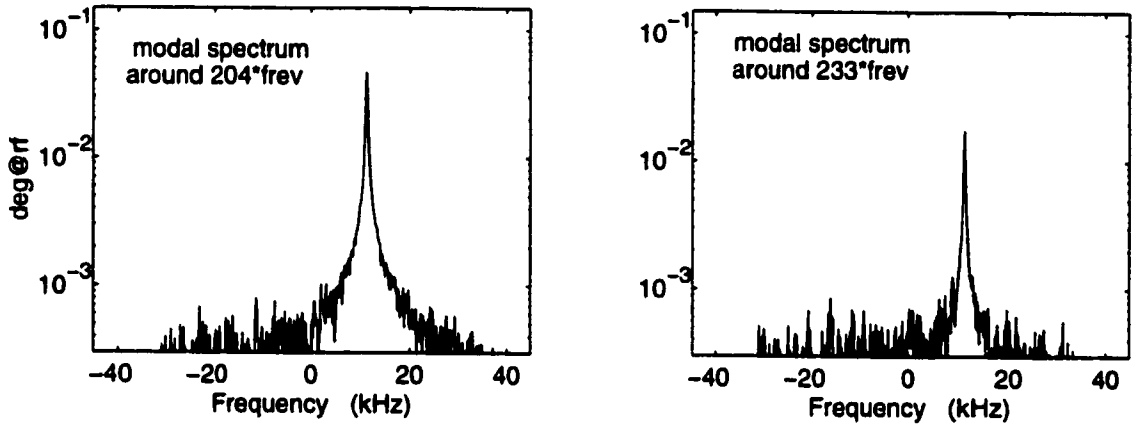

Figure 4.6: 90-kHz sections of the high-resolution spectrum in Fig. 4.5(a), showing an unstable upper sideband and a damped lower sideband at $204f_o$ and also at $233f_o$.

4.1.3 Transverse

"Grow-damp" measurements of transverse instabilities have been made by using the LFB system in conjunction with the transverse feedback system. As shown in Fig. 4.7, the baseband beam motion monitor signal from the transverse front end is fed to the LFB analog-to-digital convertor[3]. The transverse feedback path is gated by a fast TTL signal from the LFB system. Upon receipt of a software trigger, the gate is opened for a few milliseconds, and then closed again. The LFB DSPs simultaneously record the bunch oscillation signals for offline analysis.

In the absence of a spare LFB system, this measurement requires the interruption of longitudinal feedback. More details on such measurements, as well as the design of a new transverse feedback system with built-in diagnostic capabilities, can be found in [66].

Figure 4.8 shows an example of a horizontal grow-damp measurement at the ALS, at $I_o = 94$ mA. The fill pattern is the same as before; an even fill with an 8-bucket gap. We

[3]Figure 4.7 has been prepared by D. Teytelman.

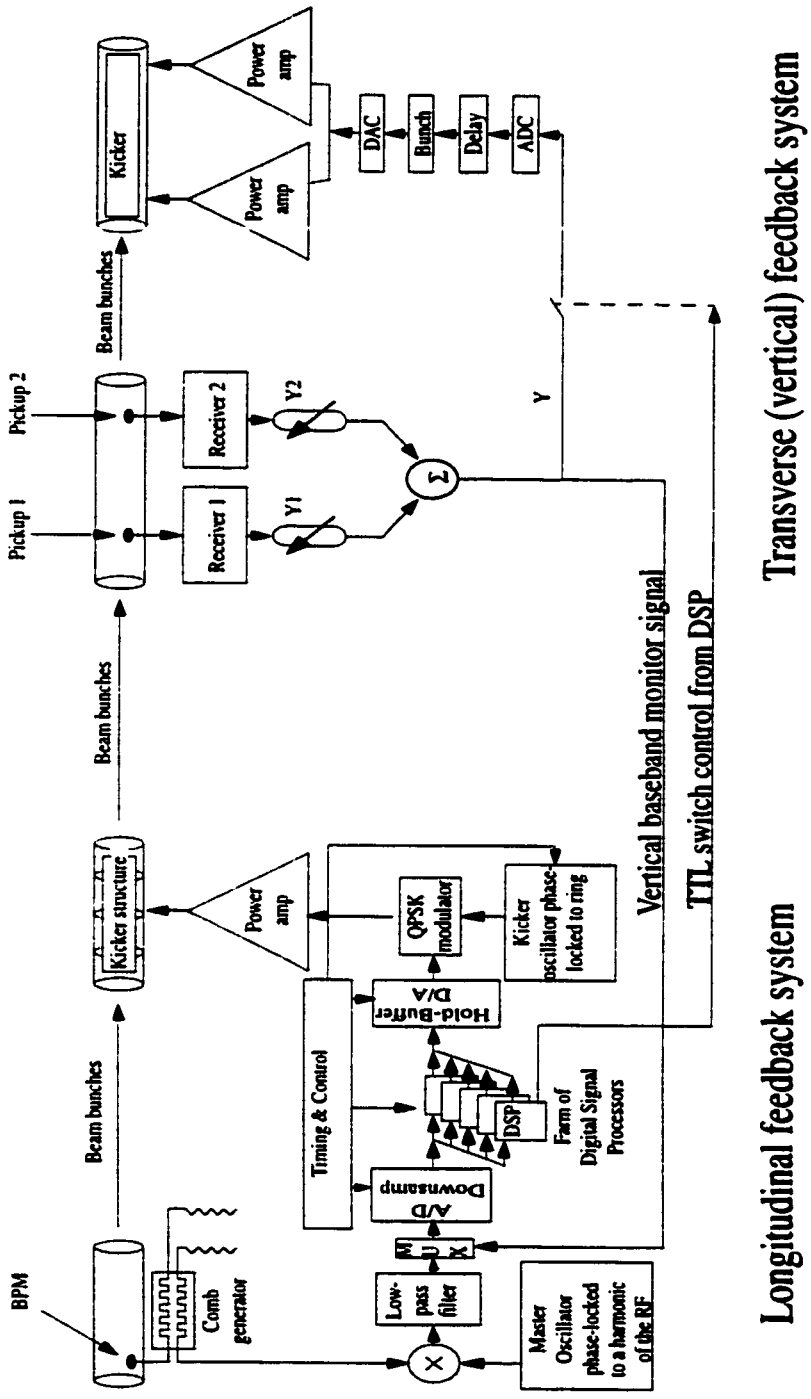

Figure 4.7: Block diagram of experimental setup for "grow-damp" measurements of transverse motion. Longitudinal feedback system used to gate the transverse feedback system, and record transverse motion.

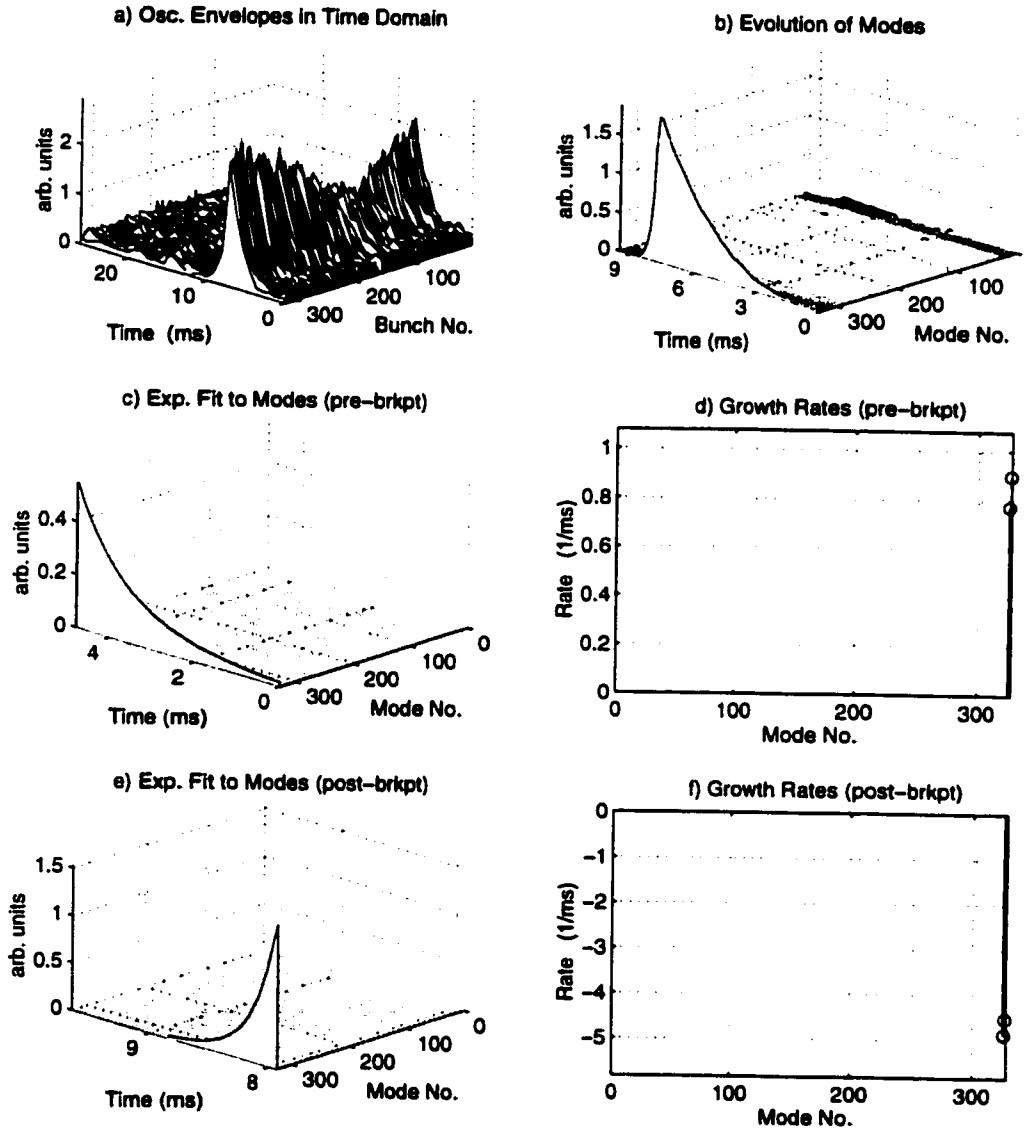

Figure 4.8: ALS horizontal grow-damp at $I_o = 94$ mA. Resistive wall instability is the strongest. Feedback gain margin is comfortable.

see that the growth and damping rates are faster than in the longitudinal case. Only modes 326 and 327 (-2 and -1) are unstable. Note: These modes are numbered according to the conventions of longitudinal modes, so they are upper sidebands of revolution harmonics 326 and 327, and lower sidebands of revolution harmonics 2 and 1. They are probably driven by the impedance of the vacuum chamber walls, which is dominant at low frequencies. Their open-loop growth rates are 0.77 ms^{-1} and 0.9 ms^{-1}, respectively. The growth rates have been calculated from fits to the initial parts of the growing transients, since nonlinear effects make the transients non-exponential at the highest amplitudes seen in this measurement. The damping rates are much faster, indicating that the transverse feedback system has a comfortable gain margin.

4.2 PEP-II Results

The PEP-II B Factory is an asymmetric electron-positron collider. The High Energy Ring (HER) stores electrons at 9 GeV, while the Low Energy Ring (LER) stores positrons at 3.1 GeV. Table 4.2 summarises the relevant beam and feedback parameters. So far, beam currents up to 750 mA and 1750 mA have been achieved in the HER and LER respectively. During commissioning, a variety of longitudinal and transverse beam dynamics experiments have been performed with the help of the PEP-II LFB system [54, 55, 81, 82, 83]. This section focuses on longitudinal dynamics in the two rings.

Table 4.2: PEP-II Parameters

Parameter	Description	HER	LER	Unit
E_o	Beam energy	9	3.1	GeV
f_{rf}	RF frequency	476	476	MHz
h	Harmonic number	3492	3492	-
f_o	Revolution frequency	136.3	136.3	kHz
α	Momentum compaction factor	2.44e-3	1.23e-3	-
I_o	Design current	1	2.25	A
f_s	Nominal synchrotron frequency	5.1–6	3.3–4.5	kHz
d_r	Radiation damping rate	54	34	s^{-1}
-	No. of RF cavities	20	4	-
-	Bunch sampling rate	238	238	MHz
D	Downsampling factor	6	6	-
P	Feedback output power	1500	1500	W
-	Output amplifier bandwidth	1–2	1–2	GHz

The two main sources of longitudinal motion identified in PEP-II are coupled-bunch instabilities induced by the cavity impedance and noise from the klystron. Coupled-bunch instabilities are usually caused by unwanted Higher-Order Modes (HOMs) in the RF cavities, or by impedance sources elsewhere in the beam surroundings. At PEP-II however, the large beam current and small revolution frequency combine to produce "low-mode" instabilities, *i.e.*, instabilities within the bandwidth of the detuned RF cavity fundamental mode [21, 82]. Low-mode motion is damped by a combination of RF feedback loops acting through the klystron [84]. The HOM-induced instabilities have been successfully damped by the above-mentioned longitudinal feedback system.

The longitudinal growth rates expected from the RF cavity HOMs [78], in the absence of feedback, are shown in Fig. 4.9. The rates are calculated at the design current and tune, with 1746 bunches evenly filled at a spacing of 4.2 ns. Since the two rings use similar cavities, their even-fill eigenvalue spectra (impedance-induced growth rates and coherent tune shifts) differ only by a scale factor. The spectra extend from DC to 119 MHz (half the bunch frequency). The damped cavity HOMs cover broad ranges of frequency.

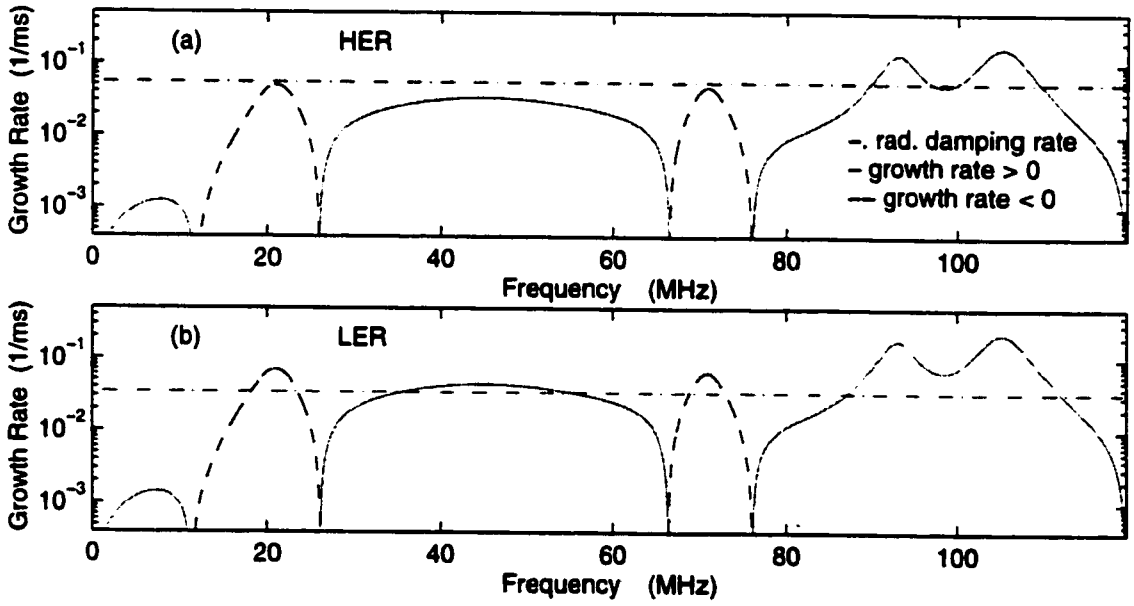

Figure 4.9: PEP-II longitudinal growth rates from RF cavity HOMs at nominal bunch spacing and tune, design current. (a) HER, 1 A, 20 cavities. (b) LER, 2.25 A, 4 cavities.

Only two HOMs are strong enough to cause instabilities in the HER. The strongest

is aliased to 105 MHz (mode 770), and drives roughly 65 modes in a 9-MHz band. The second potentially troublesome mode is aliased to 93 MHz (mode 683), and drives roughly 51 modes in a 7-MHz band. The two unstable bands are expected to broaden and coalesce in the LER, [Fig. 4.9(b)]. A new unstable band, centered at 45 MHz, is expected in the LER at 2.25 A, albeit with a very small growth rate. This band of modes is expected to be naturally stable at the beam currents achieved so far. If radiation damping is taken into account, the largest expected growth rates in the HER and LER are 0.12 ms^{-1} and 0.19 ms^{-1}, respectively.

Evidence of bands of impedance-driven modes was first uncovered in the HER. Since most of the initial commissioning was done at currents below the HER instability threshold, grow-damps were performed with low-gain positive feedback during the "grow" interval, so as to add to the destabilising effect of the HOM resonances. Figure 4.10 shows such a grow-damp, with 1650 bunches filled at the nominal spacing, and $I_o = 361$ mA. The time-domain bunch envelopes plot does not show clear growing and damping sections, since the unstable bunch motion is superimposed on noise-excited oscillations at modes within the bandwidth of the RF cavity fundamental resonance[4]. The modal plot in Fig. 4.10 zooms in on the section from mode 720 to mode 830, which contains the band of modes that grew under positive feedback, and waned under negative feedback[5]. Most of the unstable motion is localised to the band from mode 760 to mode 790. This agrees fairly well with the aliased impedance of the largest cavity HOM, which was predicted to excite modes around 770.

Interestingly, the impedance resonance that aliases to mode 683 (see Fig. 4.9) has not been seen to excite coupled-bunch motion so far, in either of the two rings. This might indicate that this mode is better damped than expected. Of course, this could also be due to the fact that grow-damp measurements have not been performed at the highest beam currents achieved so far.

The growth and damping rates of the modes in Fig. 4.10 do not reflect the full impedance of the feedback system, since the noise-excited low-mode motion in this measurement is sufficient to cause significant feedback saturation.

The technique of investigating the impedance spectrum at below-threshold currents, by means of positive feedback, is also illustrated by Fig. 4.11, which shows pseudospectra of the HER longitudinal motion below the open-loop (no feedback) threshold. As described

[4]Noise-excited oscillations were reduced significantly later on, once the RF feedback loops were fully commissioned.

[5]Due to erratic hidden-line removal, some of the modal lines in this figure are missing.

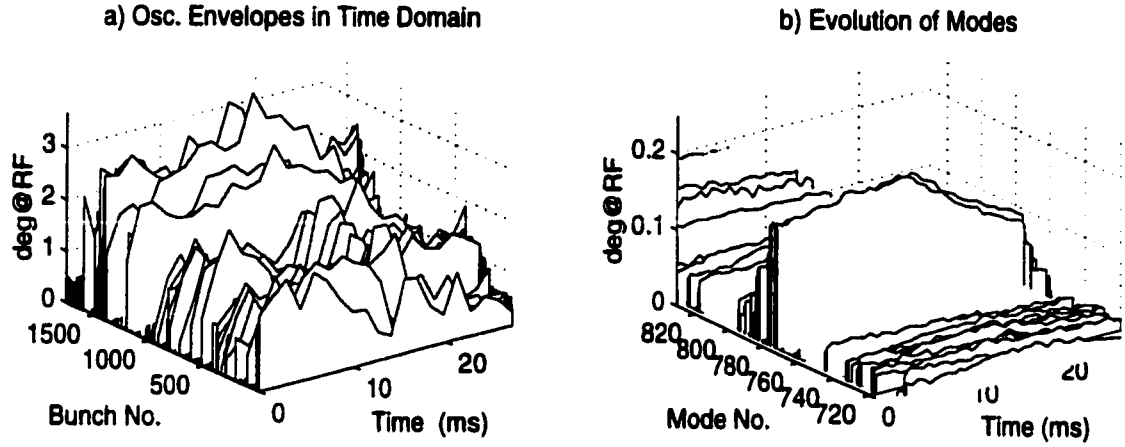

Figure 4.10: PEP-II HER grow-damp measurement at $I_o = 361$ mA. Feedback is positive during the "grow" portion of this grow-damp, and negative during the "damp" portion. Modes from 760 to 790 show significant growth and damping.

earlier, revolution harmonics are suppressed in this figure. The pseudospectra extend from 0 to 119 MHz, which is half the bunch crossing frequency in this case (every other RF bucket is filled). With positive feedback, we see excitation of a broad band of synchrotron sidebands from 100 MHz to 110 MHz. These are clearly the modes closest to instability, when there is no feedback. The impedance of the largest expected cavity HOM resonance in Fig. 4.9 is overlaid on the driven pseudospectrum. It coincides quite well with the band of driven modes.

Measurement of open-loop growth rates at machines like PEP-II presents a unique challenge. As will be explained in Ch. 7, the combination of broad impedance resonances and uneven filling patterns creates eigenmodes that are significantly different from those of an evenly filled ring. Thus, the projection onto even-fill eigenmodes by means of a Fourier transform results in non-exponential beating between superimposed uneven-fill eigenmodes. The problem of eigenmode and eigenvalue resolution is made even harder by the fact that superimposed PEP-II eigenmodes tend to have very similar eigenvalues.

D. Teytelman has devised an analysis technique that gets around this problem by calculating the growth rate of the rms oscillation amplitude, averaged over a band of "modes." If we assume that the band of projections onto even-fill modes contains uneven-fill modes

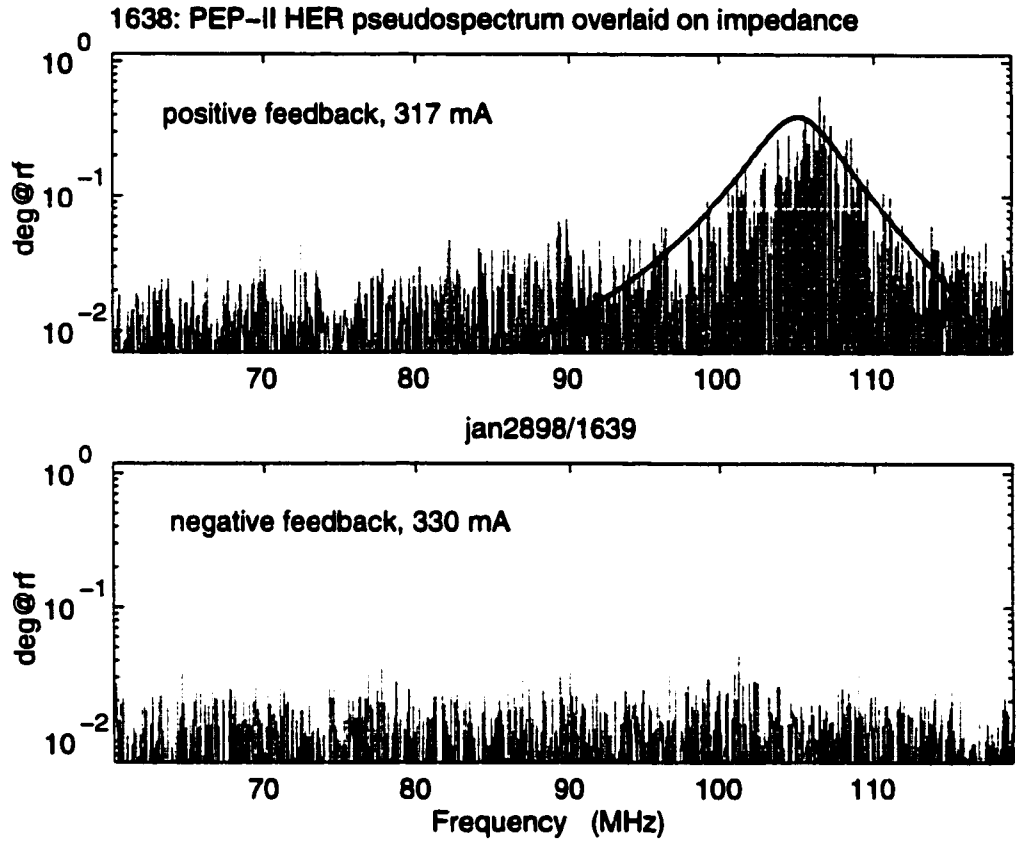

Figure 4.11: PEP-II HER beam pseudospectrum. Positive feedback excites a band of roughly 60 sidebands, in the vicinity of the 770th revolution harmonic (105 MHz). The aliased impedance of the largest cavity HOM is superimposed.

of roughly the same eigenvalue, then this gives us a good estimate of the average of the true growth rates. In practice, the rms oscillation amplitude within the unstable band has been seen to grow exponentially, even when the individual projections onto even-fill basis vectors show significant beating. This method has been applied to an HER grow-damp measurement at $I_o = 605$ mA, which shows open-loop growth in the band from mode 740 to mode 790. The growth rate is very slow; it is a mere 6 s^{-1}. In this case, 93% of the ring is filled, at a bunch spacing of two RF buckets. The instability threshold is approximately 550 mA, which is significantly higher than the value of 310 mA expected from impedance measurements, at $f_s = 5800$ Hz. It will be shown in Ch. 7 that this discrepancy is explained by Landau damping from the interbunch tune spread generated by the HER gap transient.

Figure 4.12 shows the measured open-loop growth rates of the band of LER modes from 780 to 800, over a range of beam currents, with 90% of the ring filled at a spacing of 4 buckets. The rates scale fairly linearly with I_o, with an instability threshold of 310 mA. The expected even-fill threshold is 385 mA with $f_s = 3300$ Hz, and 4 RF cavities installed. In this case, the beam is less stable than expected, indicating that Landau damping is probably not significant at these low currents. Another possibility is that the actual radiation damping in the LER is slightly smaller than the design value. The linear fit to the growth rates in Fig. 4.12 intersects the vertical axis at -0.02 ms^{-1}. The design value of the radiation damping rate is 0.03 ms^{-1}.

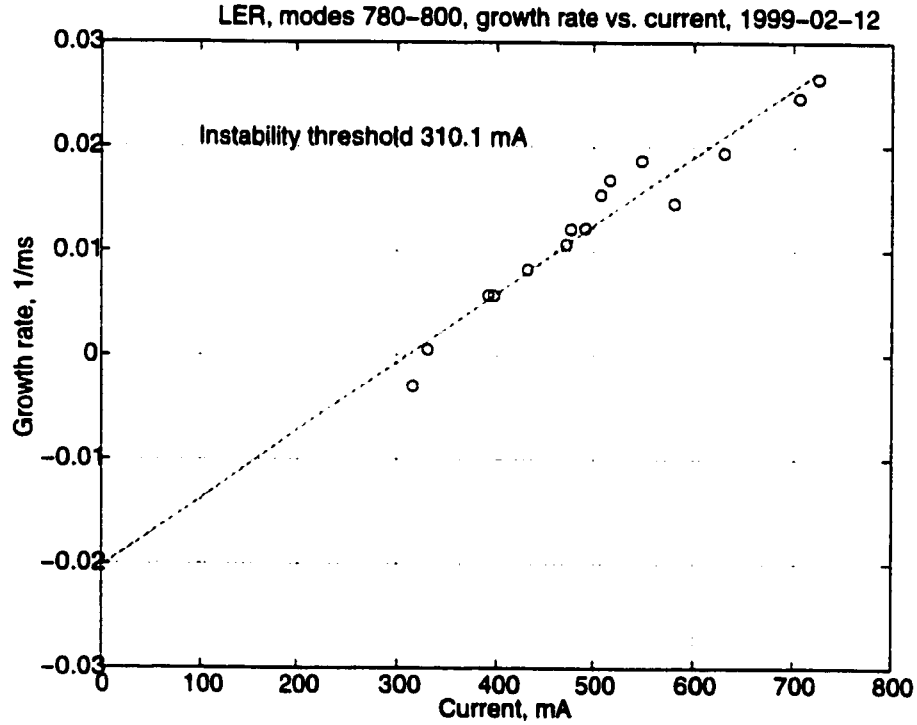

Figure 4.12: PEP-II LER open-loop longitudinal growth rates versus I_o, with 90% of the ring filled.

In general, the beam impedance is not a constant, since cavity tuners are continuously adjusted as I_o is changed. As a result, cavity-induced coupled-bunch instability growth rates often fail to scale linearly with I_o. The beam impedance is also affected by changes in the temperature of resonant structures in the ring. If the impedance is caused by a resonance that is narrow compared to f_o, then small changes in the resonant frequency will result in

noticeable growth rate variation. In the case of PEP-II, motion of HOM resonances has less of an effect on growth rates, since the dominant cavity resonances are much broader than f_o. However, the amount of Landau damping afforded by beam loading of the fundamental does vary with the tuner position.

4.3 SPEAR Results

The SPEAR storage ring functions as a 3-GeV light source, with a stored current of 100 mA. Electrons are injected at 2.3 GeV, and then ramped up to the design energy. The energy lost on each turn is replenished in two 5-cell RF cavities, which have a rich spectrum of high-Q parasitic resonances [85]. Longitudinal instability measurements at SPEAR were performed using temporarily installed LFB electronics, in conjunction with a 50-W amplifier and a stripline kicker. Most of the measurements were made at the injection energy. The pertinent beam and feedback parameters are shown in Table 4.3.

Table 4.3: SPEAR Parameters

Parameter	Description	Value	Unit
E_o	Beam energy	2.3	GeV
f_{rf}	RF frequency	358.5	MHz
h	Harmonic number	280	-
f_o	Revolution frequency	1.28	MHz
α	Momentum compaction factor	1.5e-2	-
I_o	Design current	100	mA
f_s	Nominal synchrotron frequency	28.5	kHz
d_r	Radiation damping rate	110	s^{-1}
-	No. of active RF cavities	1	-
-	No. of passive RF cavities	1	-
-	Bunch sampling rate	358.5	MHz
D	Downsampling factor	14	-
P	Feedback output power	50	W

Since the temporary LFB system had low gain, it could not control instabilities at currents much higher than the open-loop threshold, which was roughly 30 mA[6]. Grow-damp measurements above the open-loop threshold were also complicated by the fact that the feedback system was too weak to damp the beam, once the oscillations grew large enough to saturate the back end.

[6]This threshold varied significantly with temperature and cavity tuner settings.

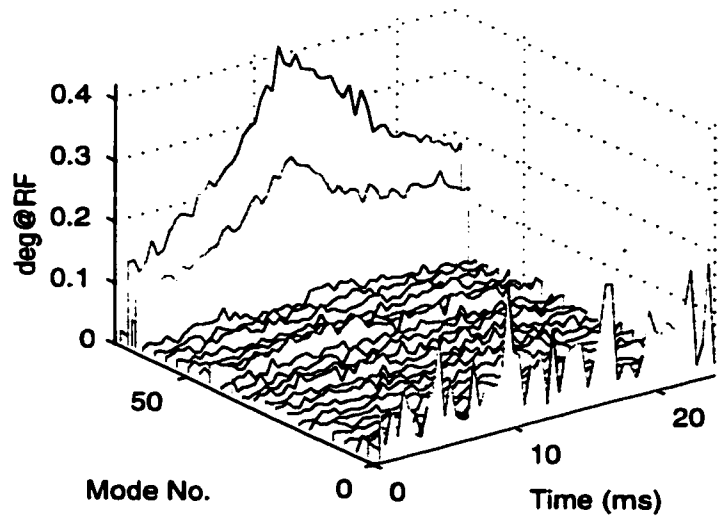

Figure 4.13: Grow-damp measurement of longitudinal instabilities at SPEAR, I_o = 29 mA. Positive feedback is used during the "grow" portion, to uncover the modes closest to instability.

Figure 4.13 shows a grow-damp measurement at the below-threshold current of 29 mA. The ring is evenly filled with 70 bunches, at an interbunch spacing of 4 RF buckets. In this case, beam oscillations grow under the action of positive feedback, until negative feedback is restored at t = 10 ms. Only modes 65 and 67 are excited by positive feedback, indicating that these modes are closest to instability at I_o = 29 mA. Many other unstable modes have been seen at SPEAR, under various conditions of temperature and tuner position.

Surprisingly, the beam is longitudinally stable in the production configuration, with I_o = 100 mA and E_o = 3 GeV. The elevated threshold is partly explained by the increased radiation damping at 3 GeV. The positions of the active cavity tuners might also be playing a part in elevating the threshold from 30 mA at E_o = 2.3 GeV to above 100 mA at E_o = 3 GeV. It is suggested in Ch. 7 that the unevenness of the production fill is yet another contributing factor.

Chapter 5

Phase Space Tracking: A Complete Diagnostic

This chapter describes an instability diagnostic that exploits the information contained in the angular evolution of coupled-bunch oscillations in phase space. In addition to enabling measurement of coherent tunes and bunch tunes with accuracy of a few Hz, phase space tracking allows new kinds of comparisons between instability theory and experiment.

Phase space tracking is used to compare the signature of a low-threshold vertical instability in the PEP-II HER [83] to those of the fast beam-ion instability (FBII) [6, 86] and conventional instabilities. It is shown that this method has the potential to distinguish between the two instability mechanisms.

Tracking of longitudinal motion at the ALS and PEP-II is used to measure coherent tunes and gain new insights into uneven-fill instabilities.

An ALS example is used to demonstrate measurement of feedback-induced tune shifts, using phase space tracking. Such measurements are now commonly used to minimise the reactive component of the LFB impedance.

5.1 Introduction

Diagnosis of the nature and cause of unstable bunch motion is the first step towards a cure. Observation of BPM power spectra under various beam conditions is the most common diagnostic. Other recently developed techniques include streak camera imaging of bunch motion, and offline analysis of digitised bunch oscillation data.

Theoretical analyses of coupled-bunch instabilities yield qualitative and quantitative predictions about bunch trajectories in phase space. Ideally, an experimenter who wants to diagnose dipole instabilities would like to be able to track the phase space positions of all bunches under various beam conditions. This would require a measurement system of bandwidth $1/(2T_b)$, where T_b is the bunch spacing.

When viewed in the light of the bandwidth or information rate requirement, streak camera measurements are seen to be unsuitable for multi-bunch phase space tracking, since they suffer from update rate limitations. However, they provide excellent time resolution, and are very useful in studying the bunch size and shape. By the same criterion, the traditional technique of observing BPM signals on a heterodyned spectrum analyser is limited by the resolution bandwidth of the spectrum analyser. To identify the instability mode number, one needs resolution bandwidths comparable to the synchrotron or betatron tune. All information outside the resolution bandwidth is lost, as is information contained in the phase of the Fourier transform of the BPM signal. For these reasons, heterodyned spectrum analyser measurements are used mainly for steady state measurements of instability frequency and amplitude, or for narrowband detection of the phase space magnitude transient of a single coupled-bunch mode.

Phase space tracking of multibunch motion is only possible with fast digitisation and storage of the oscillation coordinate of each bunch in the machine. Diagnostic techniques that satisfy this criterion suffer from loss of bunch centroid information only to the extent that measurements are never noise free. In this chapter, tracking is based on LFB system data. As mentioned earlier, LFB "grow-damp" measurements facilitate observation of the growth of unstable oscillations in the linear small-amplitude region, for which theoretical predictions exist. Another advantage of the LFB system is its ability to store downsampled data. Conventional digital oscilloscopes with data storage capabilities cannot easily use downsampling to take advantage of the fact that the beam signal only contains information near synchrotron and betatron sidebands of revolution harmonics.

5.2 Signal Processing

Coupled-bunch instability data from a single BPM is often used to measure average tunes, oscillation amplitude envelopes, average phase shifts from bunch to bunch, etc. However, it is not too difficult to estimate the approximate phase space angle corresponding to each

sample of bunch data, if the beam motion is approximately sinusoidal [87, 88].

We assume that the sampled longitudinal or transverse bunch oscillation coordinate s_n^k of bunch k is given by

$$s_n^k = a_n^k \cos(2\pi n D\nu + \phi_n^k) = Re\{u_n^k\}; \tag{5.1}$$

$$u_n^k = a_n^k e^{j(2\pi n D\nu + \phi_n^k)}, \tag{5.2}$$

where n is the sample number, D is the downsampling factor, u_n^k is the analytic signal corresponding to s_n^k, and a_n^k and ϕ_n^k vary slowly compared to the nominal tune ν. Although there is no unique solution for the phase space magnitude a and the normalised phase space angle ϕ, we can use the approximation

$$u_n^k \approx s_n^k - j\tilde{s}_n^k, \tag{5.3}$$

where $\tilde{s}_n^k$ is the Hilbert transform of s_n^k. The Hilbert transform is calculated by taking the discrete Fourier transform of s_n^k, rotating all positive-frequency components by $+90^o$ and all negative-frequency components by -90^o, and then taking the inverse Fourier transform. From Eqs. (5.2,5.3), we get

$$a_n^k \approx |s_n^k - j\tilde{s}_n^k|, \tag{5.4}$$

$$\phi_n^k \approx \angle(s_n^k - j\tilde{s}_n^k) - 2\pi n D\nu \tag{5.5}$$

This approach is equivalent to that of calculating the quadrature component of a narrow-band signal from its in-phase component. In the terminology of transverse diagnostics, we are estimating the signal at a fictitious BPM that is 90^o ahead of the original BPM in betatron phase.

When used alone, the discrete Hilbert transform produces significant errors at the edges of discrete data sets, since the corresponding filter has a long impulse response. This effect is minimised in practice by simultaneously subjecting the data to a smooth band pass filter centered at the synchrotron or betatron sideband. The band pass filter rejects noise outside the frequency band of interest. We also introduce appropriate delays in the bunch signals by means of a phase shift that is proportional to frequency. This compensates for the fact that

the bunches are sampled at different instants and on different turns, due to the requirements of downsampling.

With the processing techniques described above, we have approximations to the amplitude a_n^k and normalised angle ϕ_n^k of each bunch k at each sample instant n. If the data are digitised soon after feedback is switched off, we have enough information to test almost any theoretical prediction about the coupled-bunch instability. The most immediate application of this technique is in the diagnosis of fast transverse instabilities in short bunch trains. For example, one could distinguish between conventional instabilities and the FBII by matching the angle variation along a bunch train with the frequencies of various kinds of ions in the vacuum chamber. Tune shifts along the bunch train are also a strong indicator of the FBII. Bunch tunes can be tracked continuously by taking the derivative of smooth fits to the bunch angle.

Instabilities in beams with most of the buckets filled are well described by projections of the beam motion onto even-fill eigenmodes (EFEMs). These projections are calculated by taking the discrete Fourier transform of the sequence of analytic signals at each turn. The modal phase space coordinate U_n^m of the m-th EFEM at turn n is thus given by

$$U_n^m = A_n^m e^{j(2\pi n D\nu + \Phi_n^m)} = \sum_{k=0}^{N-1} u_n^k \, e^{-j2\pi mk/N}, \tag{5.6}$$

where N is the ratio of the harmonic number to the bunch spacing, and A_n^m and Φ_n^m are the magnitude and normalised angle respectively in modal phase space.

There is a subtle but important difference between the analytic signals u_n^k used for studying bunch motion in the time domain and those used for calculating modal projections using the equation shown above. In the former case, data is processed to estimate the phase space coordinate of each bunch as it crosses the BPM (coincidence in space). This is how a localised impedance sees bunch motion. In the latter case, signals are processed to recreate instantaneous snapshots of the analytic signals of all bunches each time the first bunch crosses the BPM (coincidence in time). This is because bunch oscillations must be projected onto the Fourier domain simultaneously, for the discrete Fourier transform to correspond to a modal decomposition.

In the past, measurements of modal magnitude transients A_n^m have been used to determine growth rates of coupled-bunch eigenmodes. Here we shall use normalised modal angle transients Φ_n^m to precisely measure coherent tunes in the linear regime. In addition to

affording direct measurements of the imaginary part of the beam impedance, this approach yields new insights into uneven-fill instability dynamics.

5.3 Phase space tracking of bunch trains

The PEP-II HER has exhibited vertical and horizontal instabilities at surprisingly low beam currents. The vertical instability was seen to grow and then saturate at amplitudes of around 100-300 μm at currents as low as 5 mA [83]. The low threshold and small saturation amplitudes triggered a search for a possible FBII.

The FBII is usually distinguished from conventional instabilities by studying the effect of variations in gas pressure, bunch spacing, train length, and bunch currents on the spectrum of betatron sidebands. Although such measurements have been used at the Advanced Light Source [89] and the Pohang Light Source [90], they are not always conclusive during commissioning, since conditions such as beam orbit, vacuum pressure, coupling, beam size, feedback, etc. are sometimes not well controlled.

To diagnose the HER vertical instability, we make use of the prediction that FBII growth in bunch trains is characterised by variation in the growth rate and bunch tune along the train. Conventional coupled-bunch instabilities caused by wake fields that persist over the length of the gap are not expected to exhibit such behavior. The approach of distinguishing between the two kinds of instability on the basis of a single growing transient has the advantage of being insensitive to artefacts such as parameter drift.

The vertical instability was investigated using digitised records of the oscillations of each bunch, immediately after switching off feedback. Figure 5.1 shows a typical growing transient in a 150-bunch train with a 4.2 ns (nominal) spacing and a total beam current I_o of 52 mA[1]. Feedback is switched off approximately at $t = 0$. The bunch oscillation amplitudes increase exponentially with time, with bunches at the tail of the train reaching higher amplitudes than those at the head. Although growth along the train is sometimes thought to be a symptom of the FBII, it is also a feature of conventional instabilities driven by an impedance resonance whose fill time is comparable to the length of the bunch train.

Bunch tune variation can most easily be examined by locating the peak in the Fourier spectrum of each bunch signal. Figure 5.2 zooms in on the bunch spectra in the region of

[1]The PEP-II HER design goal is 1 A in 1658 buckets. At a spacing of 4.2 ns, the entire ring can be filled with 1746 bunches.

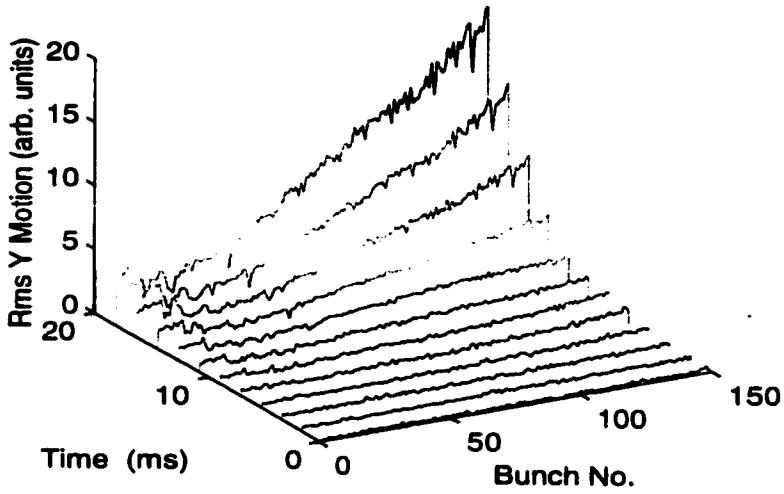

Figure 5.1: Growing vertical instability transient in a 150-bunch train in the PEP-II HER at $I_o = 52$ mA. Feedback is switched off at $t = 0$. The oscillation amplitude a^k of each bunch k grows exponentially over the 20 ms interval, with trailing bunches growing to larger amplitudes than leading bunches.

the vertical tune peak. The trailing bunches show a more pronounced spectral peak than the leading bunches, since they oscillate at larger amplitudes. Since the data record is 20 ms long, these spectra have a frequency resolution of no more than 50 Hz, *i.e.* 0.0004 in tune units. At this resolution, we see no tune shift along the bunch train. In addition to the limited resolution of the discrete Fourier transform, such measurements are also often complicated by power supply ripple, which imposes a 60 Hz modulation on the betatron tune.

We can measure tune variations with greater sensitivity by tracking the normalised bunch phase space angles ϕ_n^k. For example, we could subtract ϕ_n^{150} from all the other angles to get the phase space angle of each bunch relative to that of the last bunch. This automatically masks the tune variation due to power supply ripple. The slope of this angle differential directly yields the tune of the corresponding bunch relative to that of bunch 150. Figure 5.3 shows the phase space angle differentials $(\phi_n^k - \phi_n^{150})$ for all bunches. Only the last 7 ms of data are shown, since the signal to noise ratio is worse during the initial section of the growing transient. We see that the differential angles are almost constant over these 7 ms, with a small positive slope in the section from bunch 60 to bunch 140[2]. This

[2]See [91] for an animation depicting the evolving bunch magnitudes and relative phases in this transient.

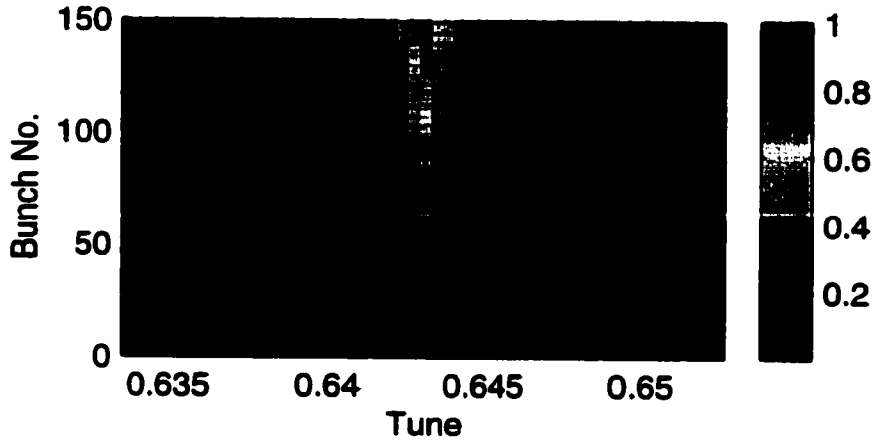

Figure 5.2: Color-coded representation of magnitude spectra of the 150 bunches in previous figure. Peaks of Fourier transform of bunch transients lie at the same tune, indicating that tune spread across the train is $\leq$ the frequency resolution (50 Hz).

implies that the bunches oscillate coherently, with very little tune variation along the train. The exact tune variation can be extracted from linear fits to the relative angles $(\phi_n^k - \phi_n^{150})$. Instantaneous tunes are not calculated, since the relative angles vary linearly with time in this piece of data.

Figure 5.4(a) shows the fitted tunes of bunches 46 to 150, relative to the tune of bunch 150. The peak-to-peak variation is less than 50 rad/s. The first 45 bunches are excluded because they grow to smaller amplitudes and have smaller signal-to-noise ratios. Exponential fits to the magnitude transients a^k yield the instability growth rates shown in Fig. 5.4(b). As can be expected of conventional instabilities, the growth rate variation across the train is small enough to be accounted for by the presence of other eigenmodes at small amplitudes. The original theoretical studies [6, 86] of the FBII predicted that oscillations should grow as $\exp(\sqrt{t/\tau})$, where τ is the growth time. However, the experimental data quite clearly shows exponential growth. A more detailed theoretical analysis, which incorporates β-function variation around the ring, predicts exponential FBII growth and a linear variation in tune shift and growth rate along the train [92]. The linear variation of growth rates and tunes is not borne out by the data, as can be seen from Fig. 5.4. It is of course possible that some of the approximations made in [92] do not apply to the relative time scales of this experiment.

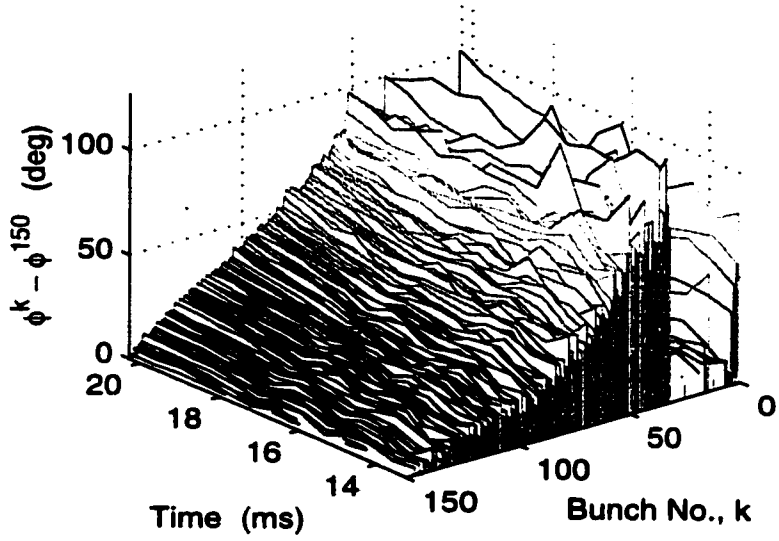

Figure 5.3: Phase space angle differentials ($\phi_n^k - \phi_n^{150}$) for all 150 bunches. Differentials for first few bunches are noisy due to smaller oscillation amplitudes. Differential angles are almost constant over 20 ms, indicating that the bunches oscillate coherently.

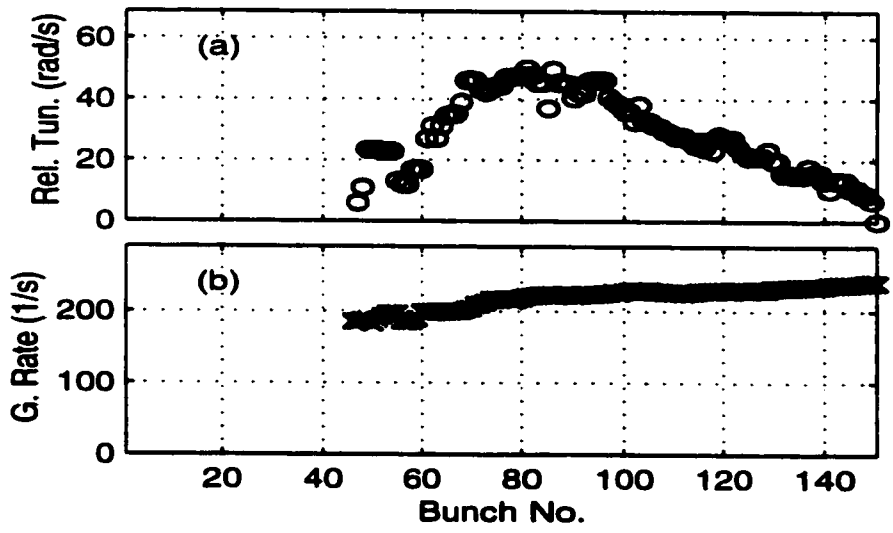

Figure 5.4: (a) Relative tunes of bunches 46 to 150, calculated using linear fits to the phase space angle differentials in the previous figure. (b) Growth rates of the same bunches, calculated using exponential fits to the magnitude transients a^k.

5.4 Phase space tracking of coupled-bunch modes

The previous section focused on applications of time domain phase space tracking, *i.e.* tracking of bunch trajectories u_n^k in phase space. Coherent instabilities in rings with more filled buckets than empty buckets are better described by the projection of these trajectories onto the eigenmodes of an evenly filled ring (see Eq. 5.6). The magnitude A^m of each projection U^m corresponds to the magnitude of the sideband of the m-th revolution harmonic in the bunch spectrum. If the coherent tune is a constant, the angle of U^m evolves linearly with a constant slope $\nu + d\Phi^m/dt = \nu + \Delta\nu^m$. The coherent tune shift $\Delta\nu^m$ can thus be measured accurately by measuring the slope of the modal phase space angle as the instability grows linearly out of the noise floor. This is a direct measurement of the imaginary part of the beam impedance.

Of the 328 longitudinal coupled-bunch modes at the ALS, only modes 204 and 233 are unstable in most cases (Ch. 4). Figure 5.5 shows the measured linear evolution of Φ^{204} (dashdot) and Φ^{233} (dotted line) as the two modes grow out of the noise floor. The ring is evenly filled at $I_o = 157$ mA. The slopes give the coherent frequency shifts, which are -132 Hz and -196 Hz respectively.

From Eq. 2.31, we know that the effective longitudinal impedance Z^{eff}, for a beam with N evenly spaced bunches, is related to the even-fill coherent tune shift by

$$Im\{Z^{eff}(m\omega_o + \omega_s)\} = -\frac{4\pi E\nu_s}{\alpha e h I_o}\Delta\nu^m; \qquad m = 0, 1, ..., N-1 \tag{5.7}$$

By scaling the measured tune shifts according to the above equation, we get $Im\{Z^{eff}(204\omega_o + \omega_s)\} = -157k\Omega$ and $Im\{Z^{eff}(233\omega_o + \omega_s)\} = -232k\Omega$. Together with the measured growth rates, these numbers have been used to estimate the shunt impedance of the cavity resonances that drive the instabilities (see Ch. 7).

In addition to aiding in impedance measurement, graphs of modal phase space angles provide information about the shape and nature of the eigenvectors of an unevenly filled ring. It is shown in Ch. 7 that two even-fill eigenmodes could be coupled to each other by means of uneven fills that contain Fourier components at their spatial beat frequency. For example, the ALS even-fill modes at $204f_o$ and $233f_o$ can be coupled together using a square-wave fill with a periodicity of roughly $1/(29f_o)$. The coupling creates a new pair of eigenmodes which are linear combinations of the two even-fill eigenmodes. If we measure the growth of one of the "mixed" eigenmodes, we should naturally see that projections

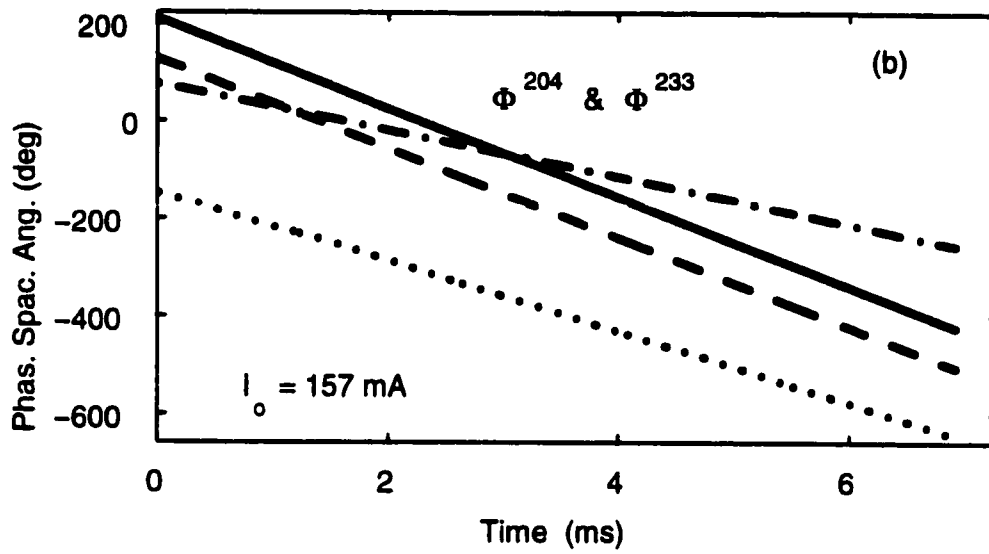

Figure 5.5: Linear evolution of modal phase space angles Φ^{204} and Φ^{233} at the ALS (longitudinal instabilities, I_o = 157 mA). *dashdot*: Φ^{204}, even fill. *dotted*: Φ^{233}, even fill. *solid*: Φ^{204}, square-wave fill. *dashed*: Φ^{233}, square-wave fill. Square-wave fill couples the two frequencies and creates a mixed eigenmode, so that U^{204} and U^{233} are phase-locked.

of the motion onto even-fill modes 204 and 233 show exactly the same growth rate and coherent tune shift. Such a measurement was performed at the ALS on the same day and at the same beam current (157 mA) as the above-mentioned even-fill measurement. The normalised phase space angles Φ^{204} (solid) and Φ^{233} (dashed line) of the projections U^{204} and U^{233} are shown in Fig. 5.5. The existence of a mixed eigenmode is confirmed by the fact that Φ^{204} and Φ^{233} have identical slopes (the slopes were different by 64 Hz in the even-fill case). As will be explained in Ch. 7, mixtures of unstable even-fill eigenmodes are generally to be avoided, since the mixed mode is more unstable than either of the even-fill modes.

Longitudinal coupled-bunch instabilities in PEP-II exhibit more complicated behaviour, since they are driven by damped cavity resonances which span tens of revolution harmonics [78]. Other complications include irregular fill shapes during commissioning and gap-induced interbunch tune spreads [55, 93], which tend to couple neighbouring even-fill eigenmodes to each other. Conventional measurements of instability growth rates are difficult to interpret under such circumstances.

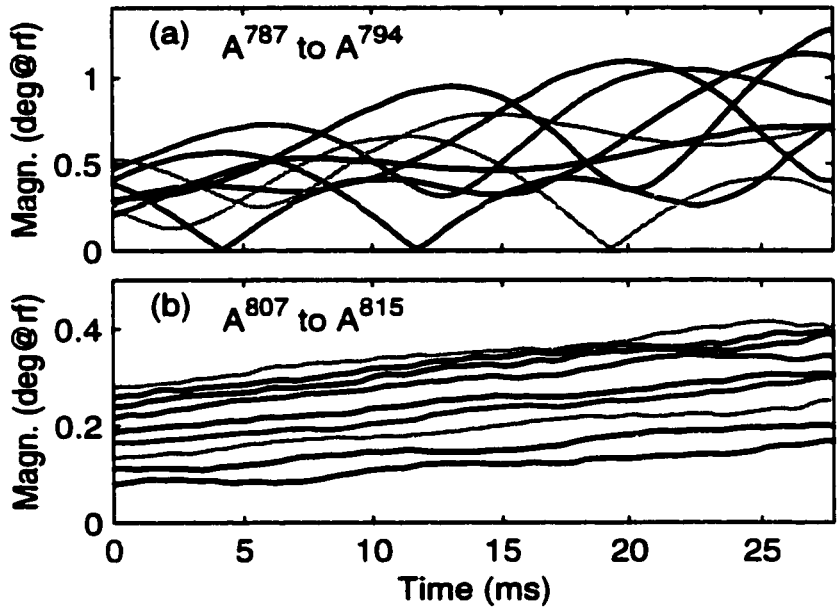

Figure 5.6: Growth in magnitude of two sets of projections of a longitudinal instability transient in PEP-II LER. Uneven fill, $I_o = 703$ mA. (a) A^{787} to A^{794}, beating is evidence of at least two uneven-fill eigenmodes in this frequency range. (b) A^{807} to A^{815}, quasi-exponential growth indicates that this set of sidebands oscillates coherently as a single eigenmode.

The uneven-fill eigenmodes can be thought of as linear combinations of even-fill eigenmodes. Thus an uneven-fill eigenmode could show up at more than one sideband in the beam spectrum, and a single sideband could be a superposition of many eigenmodes. Since different modes in general have different coherent frequencies, we should see beating of sideband amplitudes on a spectrum analyser in zero span mode, and beating of the A^ms in the reconstructed phase space trajectories.

Figure 5.6 shows the magnitude growth of two sets of projections of a single longitudinal instability transient in the PEP-II LER. The data was taken at an above-threshold beam current of 703 mA. There is a clear qualitative difference between the upper traces, which show beating at a frequency of 50 to 80 Hz, and the lower traces, which show slow quasi-exponential growth. The obvious conclusion is that the "modes" in Fig. 5.6(a) are actually superpositions of two or more uneven-fill eigenmodes with slightly different coherent frequencies. The "modes" in Fig. 5.6(b) look like projections of a single uneven-fill eigenmode, since they all show about the same growth rate.

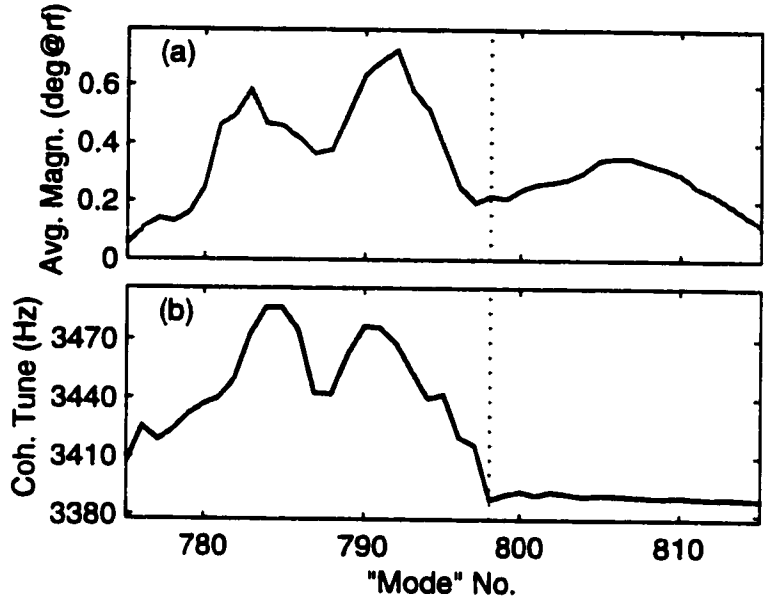

Figure 5.7: (a) Average of modal phase space magnitudes A^m; m = 775, 776, ... 815 (same data as previous figure). (b) Average coherent tunes, calculated using linear fits to Φ^m; m = 775, 776, ... 815. Fitted tunes show negligible variation above m = 798, implying that the band of projections on the right side of the dotted line contains only one eigenmode.

The LER instability transient described above is localised to the frequency range between $775f_o$ and $815f_o$, which agrees with the aliased frequency band of the largest cavity resonance (Ch. 4). Figures 5.7(a) and (b) show the average magnitudes and tunes of the modal phase space trajectories $\{U^{775}, U^{776}, ..., U^{815}\}$ in the same piece of data. We see a clear transition at "mode" 798, which is marked with a dotted line. The coherent tune spectrum to the right of the dotted line shows no tune variation, confirming our earlier conclusion that this band of projections onto even-fill modes contains just a single uneven-fill eigenmode. The other possibility, which is much less likely, is that this band contains multiple eigenmodes whose growth rates are very close and whose coherent tunes are identical to within 4 Hz.

The "modes" to the left of the dotted line in Fig. 5.7(b) seem to have a relatively large coherent tune variation. This frequency band contains two or more eigenmodes that beat against each other over time scales comparable to the length of the data set. The modal phase space angles Φ^m do not evolve linearly in this band, and therefore these calculated tunes have errors of the order of the beat frequency.

The phase space trajectories of some of the "modes" which comprise the single uneven-fill eigenmode above $798f_o$ are shown in Fig. 5.8(a). Here we see the simple exponentially growing single-frequency spirals that we usually expect. Although the actual modal phase space trajectories complete roughly 95 revolutions around the origin in the duration of this piece of data, the figure shows less than a single revolution for each "mode". This is because the phase space angle of a reference mode has been subtracted from the angles of each of the displayed trajectories, to reduce clutter in the graphical representation. In other words, we plot $U^m(t)\exp(-j\omega_{ref}t)$ in the complex plane rather than $U^m(t)$.

The phase space trajectories of a few beating "modes" U^m are shown in Fig. 5.8(b). Most of these trajectories look approximately like circles with a stationary or slowly rotating center. This indicates that the complicated beating in Fig. 5.6 is largely explained by the superposition of just two uneven-fill eigenmodes. The slowly rotating (and diverging) centers of the circles are the tips of phasors that represent an eigenmode whose coherent frequency is very close to ω_{ref}, which is $2\pi\times$ 3416 rad/s in this case[3]. The circular orbits are formed when another eigenmode with a slightly larger coherent frequency is superimposed on the original mode. It should be possible to use such plots as visual aids in precisely measuring the growth rates and coherent tunes of unstable uneven-fill modes that beat against each other.

5.5 Measuring Reactive Feedback

The LFB system should ideally act on instability growth rates alone, without affecting coherent tunes, so that the available feedback power is optimally utilised. In other words, the equivalent impedance of the feedback system should be purely resistive at the synchrotron frequency. The feedback impedance can be adjusted by adjusting filter coefficients in the DSPs. However, one does not always know in advance the exact phase shift required from the filters, to make feedback purely resistive. For this reason, measurements of the reactive component of feedback are a useful diagnostic.

It is not easy to measure the tune shift induced by linear feedback by conventional methods, since a damped beam often has no measurable motion, and an unstable beam is subject to extraneous nonlinear tune shifts. Again, the grow-damp technique, combined with phase space tracking, comes in handy.

[3]The evolution of the two bands of phasors is animated in [91].

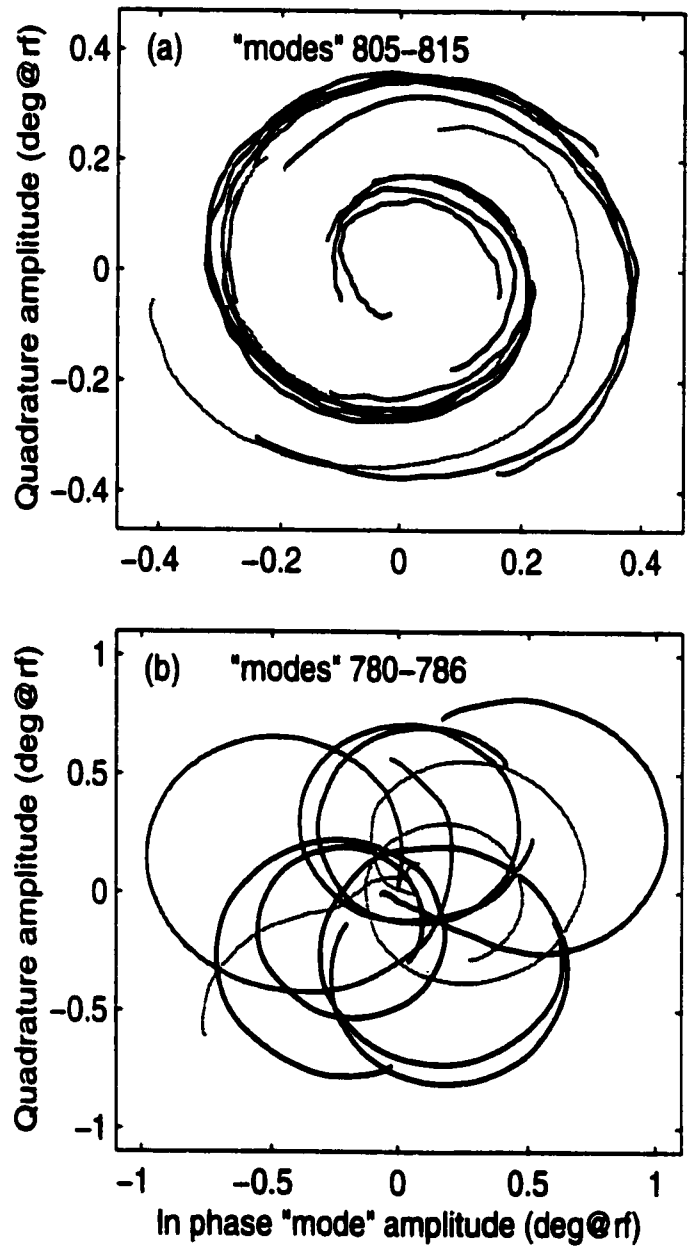

Figure 5.8: Modal phase space trajectories of growing PEP-II LER longitudinal instability (same data as previous figure). (a) Representative selection of "modes" above $m = 798$. Expanding spirals about the origin indicate a single uneven-fill eigenmode. (b) Trajectories of a few "modes" below $m = 798$, whose magnitude transients show beating.

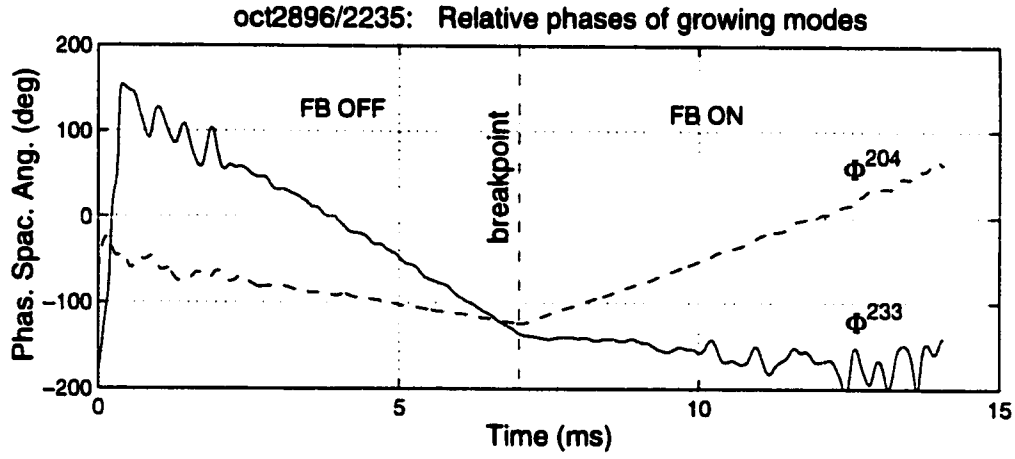

Figure 5.9: Evolution of modal phase space angles Φ^{204} and Φ^{233}, during a 238 mA ALS grow-damp. Change in slope at break point indicates that feedback has a reactive component.

Figure 5.9 shows the evolution of the modal phase space angles Φ^{204} and Φ^{233}, during the 238 mA ALS grow-damp measurement illustrated in Fig. 4.1. The slopes of these lines give the coherent tune shifts of the modes. There is a kink in the phase space angle graphs, at the break point. This implies that feedback changes the slope of the phase space angle graphs by introducing a tune shift. Linear fits to the pre- and post-break point segments reveal feedback-induced tune shifts of 104 Hz and 99 Hz respectively, for the two modes. The tune shifts of the two modes are very similar, since they sample the transfer function of the feedback system at about the same frequency.

The LFB system is tuned by adjusting the feedback filter so that the difference between pre- and post-break point tunes is minimised. Note: This diagnostic is reliable only when feedback dynamics are linear, or mildly saturated. When the feedback system is heavily saturated, it produces no noticeable tune shift, since the gain is very low. This gives the mistaken impression that the feedback system presents no reactive impedance to the beam.

5.6 Summary

Phase space tracking is a powerful new diagnostic for coupled-bunch instabilities. Tracking in the frequency (modal) domain has been shown to be useful in accurately measuring the imaginary part of the effective beam impedance at the ALS. Measurements of modal phase

space trajectories at the ALS and PEP-II confirm qualitative predictions about uneven-fill coupled-bunch eigenmodes. These trajectories can conceivably be used to measure coherent tunes and growth rates in cases where conventional methods are frustrated by beating between multiple uneven-fill eigenmodes.

Phase space tracking of grow-damp data facilitates measurement of the reactive component of feedback, which should be minimised, for optimal utilization of the available feedback power.

Tracking in the time domain has been used to study a low-threshold vertical instability in the PEP-II HER, and compare features of the bunch phase space trajectories to characteristics of conventional instabilities and the FBII. The trajectories fail to match qualitative features described in the existing literature on FBII theory. The method shows promise as a tool for analyzing data from future FBII experiments, and for revealing aspects of instability growth that have hitherto remained unexamined.

Chapter 6

A Matrix Formalism for Landau Damping

Coupled-bunch instabilities can be cured by introducing a tune spread between the bunches (Landau damping), so that they cannot organise a growing coherent oscillation. Some of the techniques for inducing such tune spreads are mentioned in Ch. 2.

Existing methods of analysing the effect of bunch-to-bunch tune shifts on coupled-bunch instabilities [26] are strictly applicable only when the beam impedance overlaps just a single synchrotron or betatron sideband. This is not usually the case. Even if there is just a single, narrow impedance resonance to reckon with, the resonance often overlaps both the upper and lower sidebands of the revolution harmonic closest to its center frequency. However, it is shown in [94] that the conventional dispersion relation is a good approximation, as long as the tune variation is smooth on a time scale of $1/\Delta f$, where Δf is the frequency difference between the mode under examination, and the nearest other mode with a non-negligible impedance.

Another (relatively minor) problem with the conventional approach is that the dispersion relation has no explicit solution.

This chapter presents a more general approach to the Landau damping problem, that involves computation of the eigenvalues of a reduced state matrix. Evenly filled rings are assumed throughout, *i.e.*, the interbunch spacing and the bunch charge are assumed to be constant. The method is applied to the analysis of PEP-II longitudinal coupled-bunch modes. Since a large number of closely-spaced PEP-II modes are naturally unstable, conventional techniques do not apply to this case. Together with decreasing revolution

frequencies in new high energy accelerators, the increasing use of damped RF cavities with broad HOM resonances makes the PEP-II example quite relevant.

The most direct and general way of calculating Landau-damped eigenvalues for a system of N bunches is to solve for the eigenvalues of the $N \times N$ state matrix A. If N is large, this becomes computationally cumbersome. If we assume slow tune variation around the ring, we can make the eigenvalue problem more manageable by creating an equivalent $M \times M$ matrix that models the dynamics of the most unstable modes of A. This is the reduced state matrix referred to in the previous paragraph.

6.1 Equivalent State Matrix

In general we can write the equations of motion of a linear system as:

$$\dot{X} = AX; \quad X(t) = X_o\, e^{\Lambda t}, \tag{6.1}$$

where X_o is any eigenvector of A, and Λ is the corresponding eigenvalue. Consider N identical evenly spaced rigidly oscillating bunches with oscillation coordinates x_k. For mathematical convenience, the variables x_k shall be considered to be complex, so that the state matrix A has size N, and the eigenvectors are merely the N Fourier vectors of Eq. 2.24:

$$v_l = [1 \;\; e^{jl\theta} \;\; e^{2jl\theta} \;\; ... \;\; e^{(N-1)jl\theta}]^T/\sqrt{N}; \qquad \theta = 2\pi/N; \qquad l = 1, 2, ...N-1 \tag{6.2}$$

The corresponding eigenvalues Λ_l ($j\Omega_l$, usually) are given by: $\Lambda_l = (-d_r + j\omega_z) + \lambda_l$, if we assume that the coherent eigenvalue shifts λ_l and the radiation damping rate d_r are small compared to the longitudinal or transverse oscillation frequency ω_z. We can calculate λ_l for each mode l by scaling the effective impedance at the corresponding revolution harmonic [5].

From here on we shall drop the common additive term $(-d_r + j\omega_z)$ from the eigenvalues Λ_l, so that we are left with only the part that contains the coherent tune shift of mode l, *i.e.*, λ_l. This merely shifts the eigenvalue spectrum of A, without changing the eigenvectors. We now have:

$$A = \sum_{l=0}^{N-1} \lambda_l\, v_l\, v_l^H, \tag{6.3}$$

where v_l^H denotes the complex conjugate of v_l^T. If we now add a small tune shift δ_k ($\delta_k \ll \omega_z$)

to the tune of each bunch k, we get the following modified matrix:

$$A = \mathrm{diag}(j\delta_0 \;\; j\delta_1 \;\; \ldots \;\; j\delta_{N-1}) + \sum_{l=0}^{N-1} \lambda_l v_l v_l^H, \tag{6.4}$$

where the "diag" function places a vector of length N onto the diagonal of an otherwise empty $N \times N$ matrix. The eigenvalues of this A matrix reveal the damping effect of a tune spread on the unstable coupled-bunch modes. Unfortunately, if N is very large ($N = 1746$ at PEP-II), the eigenvalue problem becomes difficult, or even insoluble, on most computers. The next two subsections describe the construction of an equivalent A-matrix of reduced size, whose eigenvalues approximate the most unstable eigenvalues of A.

6.1.1 Single Unstable Mode

The physics behind the approximation is illustrated by the simple case of a beam with only one unstable coupled-bunch mode v_o. Equation 6.4 reduces to:

$$A = \mathrm{diag}(j\delta_0 \;\; j\delta_1 \;\; \ldots \;\; j\delta_{N-1}) + \lambda_0 v_0 v_0^H$$

Here it is assumed that only λ_0 is non-negligible. By summing the rows of the eigenvalue equation, we get:

$$1 = \lambda_0 \langle \frac{1}{\lambda^L - j\delta_k} \rangle_k, \tag{6.5}$$

where λ^L is the Landau-damped eigenvalue, and $\langle u_k \rangle_k$ denotes the mean of u over all k. The common approach at this stage is to make the approximation that the δ_ks are closer to their neighbours than they are to λ^L, in which case we can replace the discrete averaging in the above equation by an average over a fictitious continuous distribution $\rho(\delta)$ [26]:

$$1 \approx \lambda_0 \int_{\delta_{min}}^{\delta_{max}} \frac{\rho(\delta)}{\lambda^L - j\delta} \, d\delta \tag{6.6}$$

Physically, this approximation is equivalent to the statement that neighbouring tunes are blurred together by the speed of evolution of the unstable mode. The matrix reduction method inverts this approximation by going from N discrete tunes to M, where $M < N$. We could, for example, choose M to be $N/2$ by averaging pairs of adjacent tunes. We would

then have the following average over $N/2$ fictitious tunes δ_m^1:

$$1 \approx \lambda_0 \langle \frac{1}{\lambda^L - j\delta_m^1} \rangle_m \tag{6.7}$$

The physical interpretation of this approximation is the same as before, with the criterion that the δ_m^1s are closer to their neighbours than they are to λ^L. We now have a smaller matrix A^1 of size $N/2 \times N/2$ whose largest eigenvalue is about the same as that of A. We can progressively reduce the size of A as long as the closeness criterion holds, until the eigenvalues become easy to compute.

6.1.2 Multiple Unstable Modes

In the case where the beam impedance hits more than one coupled-bunch mode, we need to transform the state vector X and the state matrix A to the Fourier basis:

$$Y = V^H X; \quad B = V^H A V; \quad BY = \lambda^L Y, \tag{6.8}$$

where the columns of V are the normalised Fourier eigenvectors of a beam with no tune spread. With a little manipulation, we can arrive at the following dispersion relation from Equations 6.4 and 6.8:

$$Y = CY; \qquad C_{m,n} = \lambda_n \langle \frac{e^{j(n-m)2\pi k/N}}{\lambda^L - j\delta_k} \rangle_k \tag{6.9}$$

This dispersion relation is hard to solve in its present form. If we assume that δ is a smooth function of k, and therefore so is $\lambda^L - j\delta$, then the terms far from the main diagonal drop out of C. If there are only a few unstable modes excited by narrow impedance resonances, their eigenvalues can be calculated independently as in the previous section, provided that the mode numbers are not too close, and the unstable unperturbed (no tune spread) eigenvalues are far from degeneracy.

Unfortunately, in the case of rings with low revolution frequency and/or damped RF cavities such as PEP-II, the unstable modes are clustered together, and their interaction through the tune spread must be considered. Let mode p be the most unstable unperturbed eigenmode. Consider the set of unperturbed modes from $(p-q)$ to $(p+r-1)$, where q and $(r-1)$ are larger than the number of non-negligible diagonals above the main diagonal in C. If the modes outside this set are stable or have eigenvalues far from λ_p, they do not couple

to mode p. We could truncate C so that only the portion that couples modes within the set to each other remains. We could now make use of the smoothness of δ_k to downsample it by a factor $N/(q+r) = N/M$, making sure that the closeness criterion is still satisfied. The obvious next step is to transfer back to the regular basis to get the following equivalent state matrix:

$$A^1 = \mathrm{diag}(j\delta_0^1 \;\; j\delta_1^1 \;\; \ldots \;\; j\delta_{M-1}^1) + \sum_{m=0}^{M-1} \lambda_m^1 \, v_m \, v_m^H, \tag{6.10}$$

where $\{\delta_m^1\}$ is the downsampled version of $\{\delta_k\}$ and $\lambda_m^1 = \lambda_{p-q+m}$. The matrix A^1 models the truncated C-matrix. It is most accurate close to row p, if $q \approx r$, while it introduces an artificial "wrap around" coupling between modes at either end of the truncated C-matrix due to the downsampling of δ_k.

We now have a reduced matrix whose eigenvalues approximate those of a Landau-damped beam in the general case, if bunch tune variation is smooth.

6.2 Application to PEP-II

In this section we apply the equivalent matrix method to the study of longitudinal tune spreads in PEP-II. The impedance of the two strongest HOMs in the damped PEP-II RF cavities produces a broad spectrum of unstable longitudinal coupled-bunch modes, which are expected to stabilise with feedback. Here we examine the effect of bunch tune spreads as the only longitudinal damping mechanism in the HER and LER. The effect of bunch-by-bunch feedback can be added on as an increase in radiation damping. The rings have a harmonic number of 3492, with every other bucket filled. We will assume that all 1746 buckets are equally filled.

The design current of the HER is 1 A. Based on the measured cavity HOMs, we have a band of roughly 60 unstable modes about mode 770, and another band of roughly 40 unstable modes around mode 682 in the absence of tune spreads. The most unstable mode is at $p = 770$, with $Re(\lambda_p) = 115$ 1/s. With a uniform tune distribution between $\delta = -300$ rad/s and $\delta = 300$ rad/s, we cannot easily compute the eigenvalues of the 1746 x 1746 A-matrix directly, so we reduce it by a factor of 6 ($M = 291$). We can choose $q = 145$, $r = 146$. Figure 6.1(a) shows the eigenvalues (*including radiation damping*) of the unperturbed HER beam and the approximate eigenvalues of the beam with a tune spread of 600 rad/s across the bunches. The perturbed eigenvalue spectrum, shown with 'x's, is most accurate at

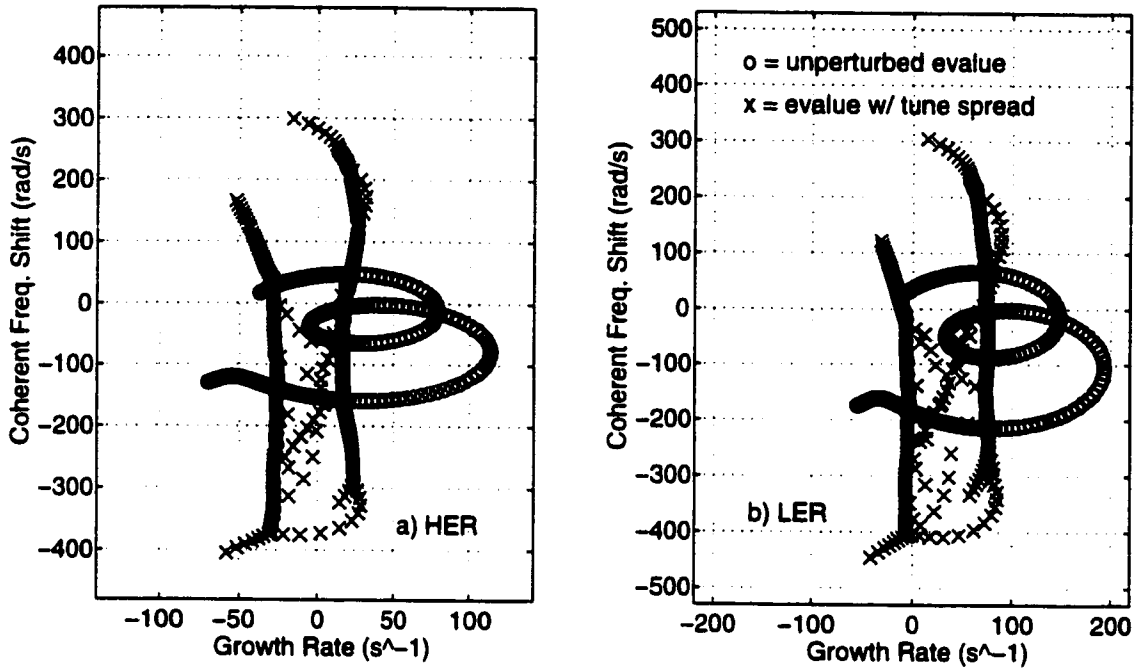

Figure 6.1: Eigenvalues of longitudinal coupled-bunch modes in PEP-II with and without bunch-to-bunch tune spread, 1746 bunches: (a) HER, 1 A, 600 rad/s tune spread (b) LER, 2.25 A, 600 rad/s tune spread.

its center, since $q \approx r$. We can see from the figure that the most unstable modes are Landau-damped down to a growth rate of roughly 25 1/s.

The LER has a design current of 2.25 A. Since the cavities in the two rings are identical, the LER is most unstable at the same value of p, with $Re(\lambda_p) = 195$ 1/s. If we assume the same tune distribution as in the case of the HER, we could use the same values of q and r. Figure 6.1(b) shows the perturbed and unperturbed eigenvalues of the LER longitudinal coupled-bunch modes. In this case, the most unstable mode is damped down to a growth rate of 75 1/s.

6.3 Summary

Existing methods of analysing the effect of bunch-to-bunch tune shifts on coupled-bunch instabilities are applicable to beams with a single unstable mode, or a few non-interacting modes. Unfortunately, in the case of rings with low revolution frequency and/or damped RF cavities such as PEP-II, we are faced with multiple unstable modes.

We have presented a more general approach to the Landau damping problem that involves computing the eigenvalues of a reduced state matrix. The application of the method to the case of longitudinal coupled-bunch modes in PEP-II has shown that a tune spread of 95 Hz across the bunches damps the most unstable HER mode from a growth rate of 115 1/s to a growth rate of 25 1/s. The corresponding numbers for the LER at full current are 195 1/s and 75 1/s respectively, given the same tune spread.

Chapter 7

Optimal Uneven Fills: A New Cure

As explained in previous chapters, coupled-bunch instabilities are conventionally cured using (see [31, 33, 34] and references therein): a) Minimization of impedance seen by the beam. b) Landau damping. c) Active feedback. This chapter describes yet another cure.

Studies of coupled-bunch instabilities [5, 22] have traditionally considered even bunch spacings and equal bunch currents (even fills), since the uneven-fill eigenvalue problem has no general analytic solution. Approaches to the uneven-fill problem have included numerical computation of the eigenvalues of the $N \times N$ bunch coupling matrix in the general N-bunch case [2, 16], and an upper bound on growth rates for an even fill with one gap [15].

Unfortunately, the existing theories do not intuitively explain why some fill shapes are more stable than others, or how one could systematically design an uneven fill to reduce the severity of instabilities. The latter problem can be restated as follows: given an effective beam impedance $Z^{eff}(\omega)$ and a maximum allowable bunch current i_{max}, how does one distribute the desired beam current I_o among the h RF buckets to minimise the largest instability growth rate $\max[Re(\lambda_l)]$?

The space of possible fill shapes has $h-1$ dimensions, since the bunch currents i_k are constrained by: $\sum_{k=0}^{h-1} i_k = I_o$, and $0 \leq i_k \leq i_{max}$; $k = 0, ..., h-1$. In some machines, it is not practical to fill consecutive buckets. Hence, we shall use the generic variable $N = h/S_b$ in place of h, where S_b is the minimum bunch spacing. If N is in the hundreds or thousands, it is clearly impractical to search for the most stable fill shape by evaluating the eigenvalues of an $N \times N$ matrix at a sufficient number of points in the $(N-1)$-dimensional search space.

Analytical complexity notwithstanding, empirically selected uneven fills have successfully raised instability thresholds at the Cornell Electron Storage Ring [95], the SPEAR

storage ring, the Advanced Photon Source [96] and other machines. Recently, the longitudinally stabilising effect of interbunch tune spreads, arising from RF cavity transients induced by gaps in the fill, has been noted [93, 97]. Though important in large rings, this effect is weak in small rings with revolution frequencies beyond the tunable range of the cavities.

This chapter presents a new, general theoretical framework for coupled-bunch instabilities in unevenly filled rings. Uneven-fill longitudinal dynamics are explained in terms of two physical phenomena: modulation coupling of strong even-fill eigenmodes (EFEMs) and Landau damping from fill-induced interbunch tune spreads. The former effect is also present in the transverse plane. These concepts are utilised to devise a simple, easily implementable algorithm for shaping fills to cure instabilities.

The theoretical predictions are verified by uneven-fill instability measurements from the ALS, SPEAR, and PEP-II.

7.1 Derivation of Coupling Matrix

As is mentioned in Ch. 2, the N Fourier vectors $v_l = [1\ e^{jl\theta}\ e^{2jl\theta}..\ e^{(N-1)jl\theta}]^T$; $\theta = 2\pi/N$; $l = 0,..N-1$, make up the eigenmodes of an N-bunch even fill. In the absence of wake fields, all modes have the same eigenvalue $-d_\tau + j\omega_s$. From here on, we shall use the word "eigenvalue" only for the coherent eigenvalue shift produced by wake fields.

As described in Ch. 2, the longitudinal arrival-time error τ_n of the n^{th} bunch centroid is given by

$$\ddot{\tau}_n + 2d_\tau\dot{\tau}_n + \omega_s^2\tau_n = -\frac{\alpha e}{E_o T_o}V_n, \tag{7.1}$$

where $V_n(t)$ is the total wake voltage seen by bunch n. If the bunches are short,

$$V_n(t) = \sum_{p=-\infty}^{\infty}\sum_{k=0}^{N-1} q_k W[t_{n,k}^p + \tau_n(t) - \tau_k(t - t_{n,k}^p)],$$

where q_k is the charge of bunch k, $t_{n,k}^p = (pN + n - k)T_b$, and T_b is the bunch spacing ($T_b = S/f_{rf} = T_o/N$). Note: uneven fills are analyzed as N-bunch evenly spaced fills with varying charge. If τ is small,

$$V_n = \sum_{p=-\infty}^{\infty}\sum_{k=0}^{N-1} q_k[\tau_n - \tau_k(t - t_{n,k}^p)]\dot{W}(t_{n,k}^p) \tag{7.2}$$

The zero-order term in this Taylor expansion has been dropped, since it merely produces a small shift in equilibrium position (synchronous phase), which we ignore. From this point on, one could derive an $N \times N$ matrix describing the coupling of every bunch to every other. This matrix is far from sparse. In addition, it offers no obvious insights into the dynamics of uneven fills. We shall instead use the EFEM basis

$$\tau^m = \sum_{n=0}^{N-1} \tau_n e^{-j2\pi \frac{mn}{N}} \tag{7.3}$$

$$\tau_n = \frac{1}{N} \sum_{m=0}^{N-1} \tau^m e^{j2\pi \frac{mn}{N}}. \tag{7.4}$$

We assume eigenvectors of the form

$$\tau_k = B_k e^{j\Omega t}; \qquad k = 0, ..N-1. \tag{7.5}$$

Similarly,

$$\tau^m = D_m e^{j\Omega t}; \qquad m = 0, ..N-1. \tag{7.6}$$

Thus,

$$V_n = \sum_{p=-\infty}^{\infty} \sum_{m,k=0}^{N-1} q_k \frac{\tau^m}{N} e^{j2\pi \frac{mn}{N}} [1 - e^{-j(m\omega_o + \Omega)t_{n,k}^p}] \dot{W}(t_{n,k}^p)$$

With the substitution $u = pN + n - k$, we get $t_{n,k}^p = uT_b$, and $q_k = q_{n-u}$. If we now project V_n onto the l^{th} EFEM,

$$\ddot{\tau}^l + 2d_\tau \dot{\tau}^l + \omega_s^2 \tau^l = -\frac{\alpha e}{E_o N} \sum_{m=0}^{N-1} \sum_{u=-\infty}^{\infty} I_{l-m} \tau^m e^{-j2\pi \frac{(l-m)u}{N}} \times [1 - e^{-j(m\omega_o + \Omega)uT_b}] \dot{W}(uT_b), \tag{7.7}$$

where the complex amplitude of the p^{th} revolution harmonic in the beam spectrum is

$$I_p = \sum_{k=0}^{N-1} i_k e^{-j2\pi \frac{pk}{N}} \tag{7.8}$$

Since the total ring impedance is

$$Z(\omega) = \int_{-\infty}^{\infty} W(t) e^{-j\omega t} dt, \tag{7.9}$$

we have

$$\ddot{\tau}^l + 2d_r\dot{\tau}^l + \omega_s^2\tau^l = \frac{j\alpha e\omega_{rf}}{E_oT_o}\sum_{m=0}^{N-1} I_{l-m}Z_{lm}(\Omega)\tau^m, \tag{7.10}$$

where the coupling impedance $Z_{lm}(\omega)$ is given by

$$\begin{aligned} Z_{lm}(\omega) &= Z^{eff}(l\omega_o + \omega) - Z^{eff}[(l-m)\omega_o]; \\ Z^{eff}(\omega) &= \frac{1}{\omega_{rf}}\sum_{p=-\infty}^{\infty}(pN\omega_o + \omega)Z(pN\omega_o + \omega) \end{aligned} \tag{7.11}$$

If $d_r \ll \omega_s$ and $|\Omega - \omega_s| \ll \omega_s$, then

$$\ddot{\tau}^l + 2d_r\dot{\tau}^l + \omega_s^2\tau^l \approx 2j\omega_s[\dot{\tau}^l + (d_r - j\omega_s)\tau^l]. \tag{7.12}$$

From Eqs. (7.10) and (7.12),

$$\begin{aligned} \dot{\tau}^l + (d_r - j\omega_s)\tau^l &= \sum_{m=0}^{N-1} A_{lm}\tau^m; \\ A_{lm} &= \frac{\alpha e f_{rf}}{2E_o\nu_s} I_{l-m}Z_{lm}(\omega_s) \end{aligned} \tag{7.13}$$

If the fill is even, $I_k = 0$ for $k \neq 0$, and the coupling matrix A becomes diagonal. The diagonal elements yield the following well-known expressions [5] for the even-fill eigenvalues (see Eq. 2.31):

$$\lambda_l = A_{ll} = \frac{\alpha e f_{rf}}{2E_o\nu_s} I_o[Z^{eff}(l\omega_o + \omega_s) - Z^{eff}(0)]; \qquad l = 0, ..N-1. \tag{7.14}$$

It is apparent from Eq. (7.13) that:

- The sum of the eigenvalues ($\sum A_{ll}$) is independent of fill shape.
- Uneven-fill eigenvalues vary linearly with I_o.
- Radiation damping merely shifts all eigenvalues by d_r, regardless of fill shape.
- If all filled buckets have the same charge q_k, then broadband bunch-by-bunch feedback also damps all uneven-fill modes equally, since it behaves like radiation damping.
- The EFEM basis yields a sparse A-matrix. The ALS, for example, has significant impedance only at 4 revolution harmonics [72]. Thus the ALS A-matrix, which is of

size 328 x 328, has non-negligible entries only along 4 rows and 8+1 diagonals. Even these are thinly populated, since I_k is usually negligible for most k.

In addition to being sparse and computationally more tractable, Eq. (7.13) reveals the two main uneven-fill phenomena: modulation coupling and Landau damping.

7.2 Modulation Coupling

Modulation coupling, which is represented by the terms $I_{l-m}Z^{eff}(l\omega_o + \omega_s)$ [see Eqs. 7.11 and 7.13], arises from the fact that the longitudinal beam signal is proportional to $i_k\tau_k$. If τ_k has a frequency component at $m\omega_o + \omega_s$ (*i.e.* τ^m) and i_k has a component at $(l-m)\omega_o$, then their product must excite the effective impedance at $l\omega_o + \omega_s$, which in turn drives τ^l. In other words, τ^l "sees" τ^m if $I_{l-m}Z^{eff}(l\omega_o + \omega_s)$ is large. The loop is closed if τ^m sees τ^l through $I_{m-l}Z^{eff}(m\omega_o + \omega_s)$. Note: $I_{m-l} = I^*_{l-m}$.

Since the transverse beam signal is the product of the transverse oscillation coordinate and i_k, modulation coupling is also present in the transverse plane. Similarly, higher bunch shape oscillations are also affected by modulation coupling.

This insight immediately suggests a new cure for coupled-bunch instabilities, namely, coupling of unstable EFEMs to stable EFEMs through uneven fills tailored to maximise the difference frequency. It also suggests that fill shapes that couple unstable EFEMs to each other should definitely be avoided (experimental evidence follows).

7.2.1 Simplest Case

If $I_kZ^{eff}(k\omega_o)$ is negligible for all $k \neq 0$, the modulation coupling terms are the only manifestation of fill unevenness, and the problem simplifies considerably. In addition, if $Z(\omega)$ is made up of a small number n of sharp resonances, we can approximate the most unstable eigenvalues by those of an equivalent A-matrix consisting only of the n corresponding rows and columns. This is a dramatic simplification in large rings with hundreds or thousands of stored bunches. If we now design a fill so that I_k coincides only with the beat frequency of the most stable EFEM m and the most unstable EFEM l, we get an equivalent A-matrix which is diagonal except for the coupling between τ^m and τ^l. This reduces the eigenvalue problem to a quadratic equation with the solution:

$$\lambda = \frac{1}{2}(\lambda_l + \lambda_m) \pm \frac{1}{2}\sqrt{(\lambda_l - \lambda_m)^2 + 4C^2_{l-m}\lambda_l\lambda_m}, \tag{7.15}$$

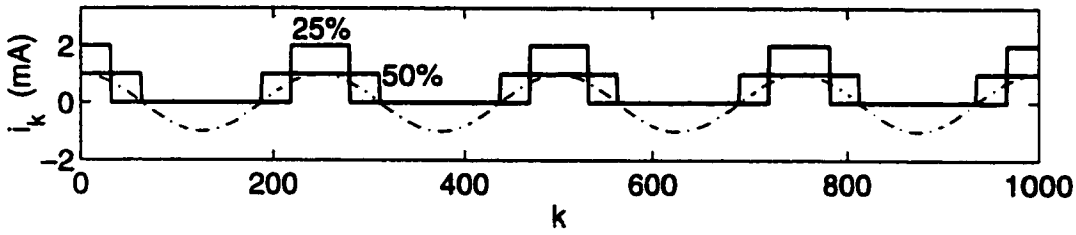

Figure 7.1: Illustration of fill optimization. $N = 1000$, $I_o = 500$ mA. Solid lines: 50% fill and 25% fill maximise C_4 for $i_{max} = 1$ mA, 2 mA. Dashdot: Reference sinusoid at $4\omega_o$.

where C is a modulation parameter defined by $C_p = |I_p|/I_o$. If $C_{l-m} = 0$, the even-fill eigenvalues λ_l and λ_m are unperturbed. As C_{l-m} approaches unity (it can never exceed unity), one eigenvalue approaches zero and the other approaches $\lambda_l + \lambda_m$. This yields the maximum damping. Equation 7.15 is equally valid in the transverse case.

7.2.2 Algorithm for Optimising Fill Shape

To damp instabilities by modulation coupling, we need a way of identifying unstable EFEMs and their eigenvalues, or at least their growth rates. We also need an algorithm for designing fill shapes to maximise C_p at desirable values of p, given I_o and i_{max}.

It can be shown that the following procedure maximises C_p under the above-mentioned constraints:

1) For each bucket n in the N-bucket pattern, calculate a corresponding "weight" $\cos(2\pi\frac{pn}{N})$.

2) Fill each of the "heaviest" I_o/i_{max} buckets to the same current i_{max}.

If N is not divisible by p, this gives

$$C_p \approx \sin(\pi x)/(\pi x), \tag{7.16}$$

where x is the fraction of the ring filled. Figure 7.1 shows two example fills which maximise C_4 when $N = 1000$, $I_o = 500$ mA, and $i_{max} = 1$ mA, 2 mA.

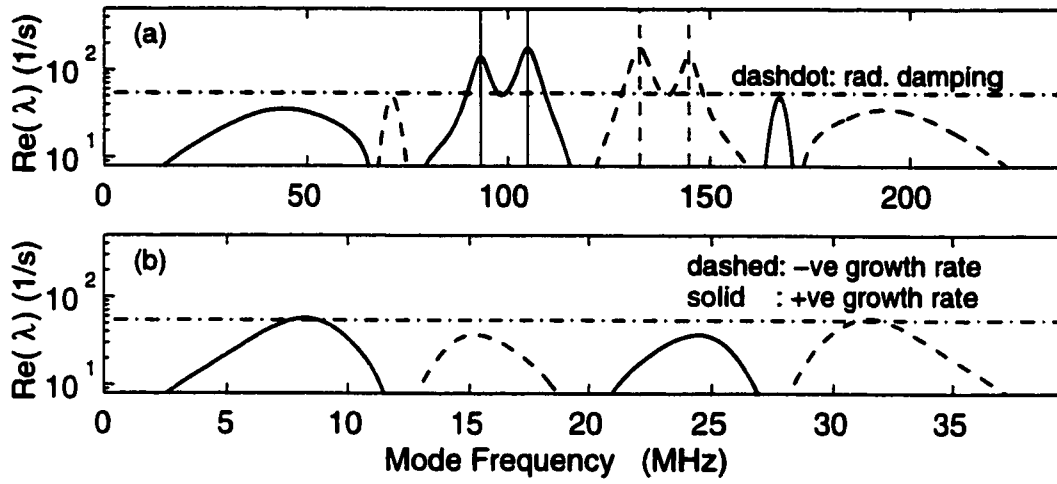

Figure 7.2: PEP-II HER cavity-induced modal growth rates vs. mode frequency ($l\omega_o + \omega_s$) at $I_o = 1$ A for: a) Even fill at nominal 4.2 ns spacing (feedback required). b) Even fill at 6×4.2 ns spacing (stabilised by modulation coupling).

7.2.3 Sample Calculation: PEP-II

The PEP-II RF cavities have two prominent parasitic resonances [78] which drive broad bands of modes unstable. The bands are centered at 93.1 MHz (EFEM 683) and 105 MHz (EFEM 770), if the ring is filled evenly with the nominal bunch spacing ($T_b = 4.2$ ns, $N = 1746$). The cavity resonances also stabilise corresponding bands of modes centered at 144.9 MHz (EFEM 1063) and 133 MHz (EFEM 976). Figure 7.2(a) shows the estimated cavity-induced growth and damping rates for the High Energy Ring (HER) at the design current of 1 A. The most obvious approach is to couple the modes around 105 MHz to those near 133 MHz. This would damp the most unstable band, but the band around 93.1 MHz would remain undamped. A better approach is to couple 105 MHz to 144.9 MHz by maximising C_{293}. This automatically couples 93.1 MHz to 133 MHz. In general, if τ^a couples to τ^{N-b}, then τ^b couples to τ^{N-a}.

In machines like PEP-II with low ω_o and damped cavities, quite a few EFEMs see a non-negligible impedance. The problem of exactly calculating the most unstable eigenvalue for various values of C_{293} is therefore non-trivial. If C_{293} is small, we can assume that τ^{770} only couples to τ^{1063}, τ^{771} only couples to τ^{1064}, etc., so that Eq. (7.15) is sufficient. It is possible to show that as C_{293} becomes large, $C_{2\times293}$ starts growing and then $C_{3\times293}$, and so on, if one uses the fill optimization algorithm described earlier. Thus the smaller, more

distant impedance resonances start to figure in the exact calculation.

Here we will use the fact that optimising C_{291} is almost as useful, since ω_o is small compared to the bandwidths of the two large impedance resonances. C_{291} can be optimised using fills with a periodicity of $6T_b$, since $N/6 = 291$. This simplifies the eigenvalue problem considerably, since we can extend even-fill symmetry arguments. If a fill has a periodicity of MT_b, its eigenvectors must remain unchanged except for a scale factor when rotated by M. The N/M scale factors are given by $e^{j2\pi rM/N}$; $r = 0,..N/M-1$. Each scale factor defines a family of M eigenmodes. Since these families are orthogonal to each other, the eigenvalue problem reduces to N/M problems of size $M \times M$, *i.e.* A separates out into N/M unconnected matrices. Physically, this corresponds to the problem of finding out the relative phases and amplitudes of M adjacent bunches in an eigenstate, given that the complex amplitudes of successive periods are related by the ratio $e^{j2\pi rM/N}$.

In the limit, only every sixth slot is filled, $C_{291} = 1$, and the 1746-bucket uneven fill reduces to a 291-bucket even fill. $Z(w)$ is then aliased into the frequency band from 0 to 39.7 MHz, and the two main impedance resonances almost cancel each other. We can see from Fig. 7.2(b) that such a fill is completely stabilised by modulation coupling and radiation damping, without recourse to active feedback or Landau damping.

7.2.4 Intuitive Explanation for Modulation Coupling

In general, as C_p approaches 1, the fill shape approaches that of an evenly filled ring with $\omega_{rf} = p\omega_o$. Such a fill would alias $Z(w)$ into the band from 0 to $p\omega_o$, so that impedances separated by $p\omega_o$ overlap. This intuitively explains modulation coupling. Equivalently, if we sample τ^l at all filled buckets, it looks increasingly like τ^{l+p} sampled at the same buckets as C_p increases.

7.3 Landau Damping

Landau damping, which is represented by the tune-spread terms $I_{l-m}Z^{eff}[(l-m)\omega_o]$ [see Eqs. 7.11 and 7.13], is another uneven-fill cure for coupled-bunch instabilities.

For illustration, consider a ring with one sharp impedance resonance which coincides only with the n^{th} revolution harmonic, where $n\omega_o$ is not a multiple of the bunch frequency $N\omega_o$. If we design a fill optimised for C_n, we excite a sinusoidal ringing in the wake voltage at the frequency $n\omega_o$, which gives each bunch a unique frequency shift proportional to

$n\omega_o I_n Z(n\omega_o)$. The tune spread, which is proportional to the impedance (and therefore to the even-fill instability growth rate), provides Landau damping.

Of course, narrowband resonances are not required for producing a tune spread, though they simplify the explanation of fill-induced Landau damping. The proportionality of tune spread to impedance is quite general, as we shall see in the following subsection. There is some degeneracy in the frequency shifts, if n and N are not mutually prime.

Unfortunately, there is no analog of this phenomenon in the transverse plane.

This method has the advantage that all unstable dipole modes are damped by the interbunch tune spread. Bunch shape oscillations are also damped, since the distorted potential wells produce interbunch spreads in quadrupole, sextupole and higher tunes as well.

7.3.1 Interbunch Tune Spread Formula

The tune shift of bunch k relative to the mean tune is

$$\delta\omega_s^k = j\frac{\alpha e f_{rf}}{E_o \nu_s} \sum_{l=1}^{N-1} \left[Z^{eff}(l\omega_o) I_l e^{j2\pi kl/N} \right] \tag{7.17}$$

$\delta\omega_s$ is purely real, since the real part of the summand is an odd function of l, with period N.

7.3.2 Narrowband Impedance Spectrum

The analysis of Landau damping in unevenly-filled rings with just a few sharp impedance resonances is relatively simple. If n is the most unstable EFEM, a good strategy would be to design a fill that optimises C_n.

The best value of C_n for damping EFEM n is different from the optimum for other EFEMs.

EFEMs other than n

Landau damping of EFEMs other than n can be calculated in the usual way [26, 94], as long as they are not coupled to other prominent EFEMs by modulation coupling or by Landau

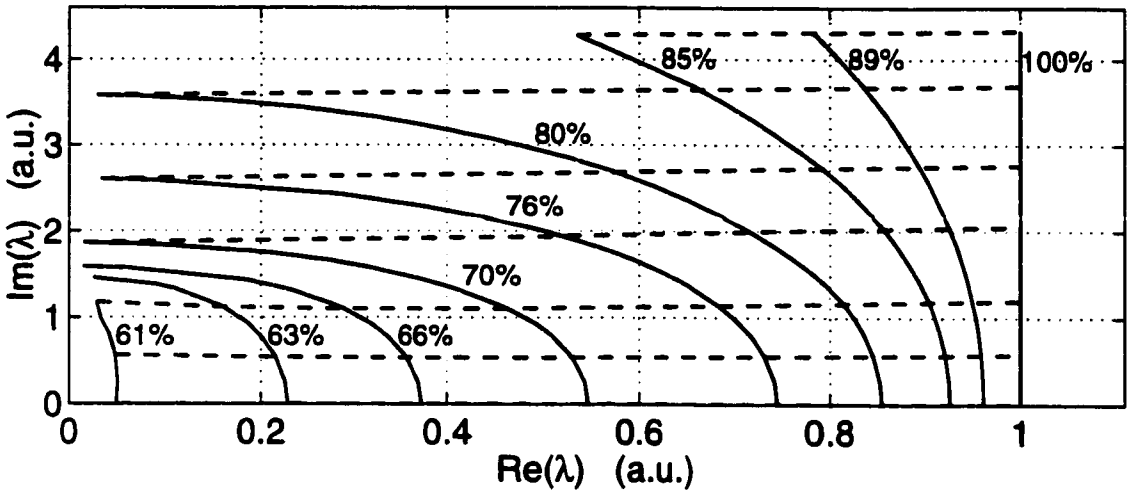

Figure 7.3: Graphical look-up table for fill-induced damping of eigenvalue of unstable longitudinal EFEM n as C_n is increased from 0 (100% of ring filled) to 0.5 (61% filled). Dashed lines: Evolution of λ_n from a few even-fill starting points.

terms on the n^{th} diagonal of A. Alternatively, one could use the formula

$$\lambda(\sigma_\omega) \approx \lambda_l - \sigma_\omega^2/\lambda_l;\ l \neq n, \tag{7.18}$$

where σ_ω is the rms of $\delta\omega_s^k$ over all filled buckets. If C_n is close to 1, EFEMs other than n are barely damped, since all bunches see almost the same tune shift. Since σ_ω is close to its maximum when $x \approx 0.45$, the best value of C_n for EFEMs other than n is approximately $\sin(0.45\pi)/(0.45\pi) = 0.7$.

EFEM n

The damping of EFEM n is larger than that of the other modes, since the combination of tune spread and fill unevenness introduces coupling between τ^n and τ^{N-n}. If Landau damping and coupling to τ^{N-n} are the only significant effects and λ_{-n} is assumed to equal $-\lambda_n^*$ (this is true from Eq. 2.3!, if $2\omega_s$ is small compared to the bandwidth of the HOM), then the variation of λ_n with fill fraction x is shown in Fig. 7.3 (numerical computation). Dashed lines show the evolution of λ_n from a few even-fill starting points. This figure is symmetric about both axes. Interestingly, eigenvalues with large imaginary parts are completely damped even by 80% fills. EFEM n is best damped by maximising C_n.

7.3.3 Broadband Impedance Spectrum

In large high-current machines, the gap transient is often the strongest source of tune spreads [93, 97]. Since this spread arises from the lowest frequencies in i_k, the Landau terms in A tend to couple neighbouring EFEMs to each other. This reduces the efficacy of Landau damping, if the dominant impedance resonances are broad enough to straddle many revolution harmonics (coupling of unstable modes through the tune spread).

Intuitively the relative efficacies of various tune distributions can be understood on the basis of the fact that bunches couple mainly to their immediate neighbours, if cavity resonances are damped. Thus, small, rapidly oscillating tune distributions perform as well as large, smoothly varying tune distributions from gap transients.

Coupling of neighbouring EFEMs makes conventional one-mode-at-a-time analysis of Landau damping invalid in machines like PEP-II, which have broad bands of impedance-driven modes. If the gap is small, one could neglect modulation coupling and calculate Landau-damped eigenvalues using matrix reduction [98], as explained in the previous chapter. If the gap is large, there are currently no alternatives to numerical computation for such machines, unless one extends the matrix-reduction technique.

7.4 Some Special Cases

If an uneven fill excites no revolution harmonics in $Z(w)$ and has no significant components at modulation coupling frequencies, its eigenvalues are no different from those of an even fill.

In cases where $Z(w)$ consists of many narrow resonances that vary significantly with temperature and tuner position, fills that contain a "cocktail" of important frequencies should prove useful. If, for example, one wanted both C_u and C_v to be large, one could give bunch n a weight $\cos(2\pi\frac{un}{N}) + R\cos(2\pi\frac{vn}{N})$ in step (1) of the fill-generating algorithm, where R is a monotonic function of the desired ratio C_v/C_u.

Other ideas for further study include tuning of superconducting cavity HOMs to land on revolution harmonics, and symmetric tuning of Landau cavities around multiples of f_{rf}. If a narrowband resonance of a superconducting cavity is tuned onto a revolution harmonic, it might be possible to achieve a large interbunch tune spread without having any significant impedance at the neighbouring sidebands.

A similar effect could be achieved by tuning the fundamental modes of Landau cavities

onto the stabilising sidebands of (say) EFEMs +1 and -1. A fill with a gap would then create some interbunch tune spread.

7.5 Brief Summary of the Theory

Uneven-fill theory is most easily explained in terms of relative amplitudes of revolution harmonics in the bunch spectrum. If we define $I_p = \sum_{k=0}^{N-1} i_k e^{-j2\pi\frac{pk}{N}}$, where i_k is the current in bunch k, then the amplitude of any revolution harmonic p in the bunch spectrum relative to the lines at multiples of $N\omega_o$ is given by $C_p = |I_p|/I_o$.

The theory identifies two important classes of revolution harmonics p in the beam spectrum:

1) Harmonics of the form $l-m$, where l and m are even-fill eigenmodes (EFEMs) that couple strongly to the beam impedance. If C_{l-m} is comparable to 1 (it never exceeds 1), EFEMs l and m couple to each other, forming two composite eigenmodes which are linear combinations of the two EFEMs (modulation coupling). The corresponding eigenvalues are given by Eq. 7.15.

2) Harmonics that coincide with peaks in Z^{eff}. Fills with significant components at such frequencies drive a steady state ringing in the wake field, which creates potential well distortion that varies from bunch to bunch. The resulting interbunch tune spread damps instabilities (Landau damping). The amount of damping achievable is shown as a graphical look-up table (Fig. 7.3).

The best solution to the problem of maximising C_p for a certain desired beam current I_o and maximum bunch current i_{max} is to fill a suitable subset $M = I_o/i_{max}$ of the N available buckets. If the filled fraction x is defined as $x = M/N$, then $C_p \approx \sin(\pi x)/(\pi x)$.

7.6 Measurements

This section presents uneven-fill instability measurements from the ALS, PEP-II and SPEAR. The ability to measure the eigenvalues (growth rates and coherent tune shifts) of all unstable modes simultaneously has greatly improved the accuracy and completeness of instability diagnostics at these machines.

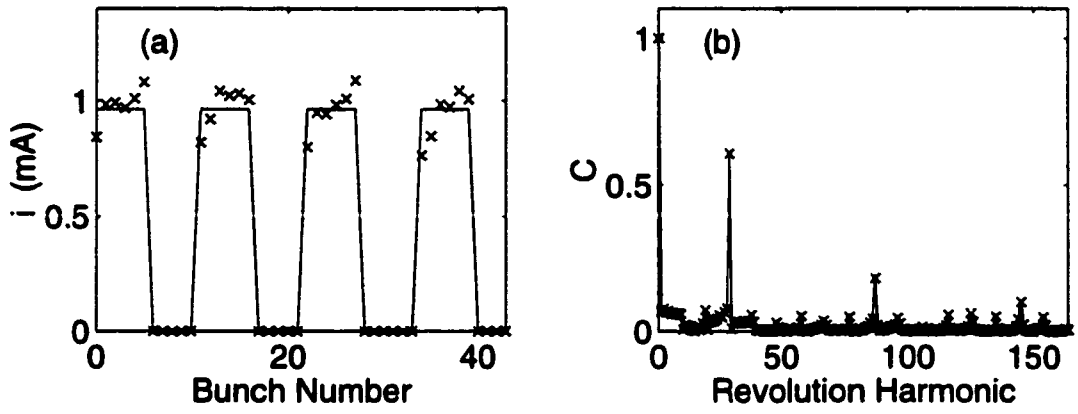

Figure 7.4: Uneven fill couples EFEMs 204 and 233 by maximising C_{29} for $x = 0.5$. Solid line: ideal; 'x's: measured. a) Typical section of bunch current profile. b) Fill spectrum.

7.6.1 ALS Modulation Coupling

The theoretical predictions of modulation coupling were first tested at the ALS. As was shown in Ch. 4, only two of the 328 ALS modes are unstable: modes 204 and 233. The two were coupled using the fill shown in Fig. 7.4(a), which optimises C_{29} for $x = 162/328$. The measured and desired bunch currents differ by less than 20% in most cases. Since the difference has the form of white noise, the error in C_{29} is very small (see Fig. 7.4(b)).

Figure 7.5(a) shows a transient measurement of the growth of the two modes at $I_o = 158.5$ mA in an even 328-bucket fill with no longitudinal feedback. The technique of digitising and analyzing such transients in the linear small-oscillation regime is unique in that it allows accurate measurement of the instability growth rate and coherent tune, *i.e.*, the most unstable eigenvalues. Growth rates are measured using exponential fits to the modal transients in the EFEM basis. Coherent tunes are measured using linear fits to the angular evolution of individual EFEMs in phase space.

Fits to the growth rates and coherent tunes in Fig. 7.5(a) give us the following eigenvalues[1]: $\lambda_{204} = (0.39 \pm 0.02) - (0.83 \pm 0.03)j$ ms^{-1} and $\lambda_{233} = (0.35 \pm 0.02) - (1.23 \pm 0.03)j$ ms^{-1} (assuming that $d_\tau = 0.074$ ms^{-1}). For the uneven fill shown in Fig. 7.4, $C_{29} = 0.62$. If we scale the even-fill eigenvalues to $I_o = 156.3$ mA and apply Eq. 7.15, we get $\lambda = (0.59 \pm 0.03) - (1.68 \pm 0.05)j$ ms^{-1} for the most unstable uneven-fill eigenvalue. The

[1]If the impedance is constant, eigenvalues vary linearly with I_o. Deviations of the measurements from linearity arise from fluctuations in the cavity temperatures and tuner positions, and from errors in fitting curves to measurements. Here, RMS deviations from linearity are used as error bars for eigenvalue measurements.

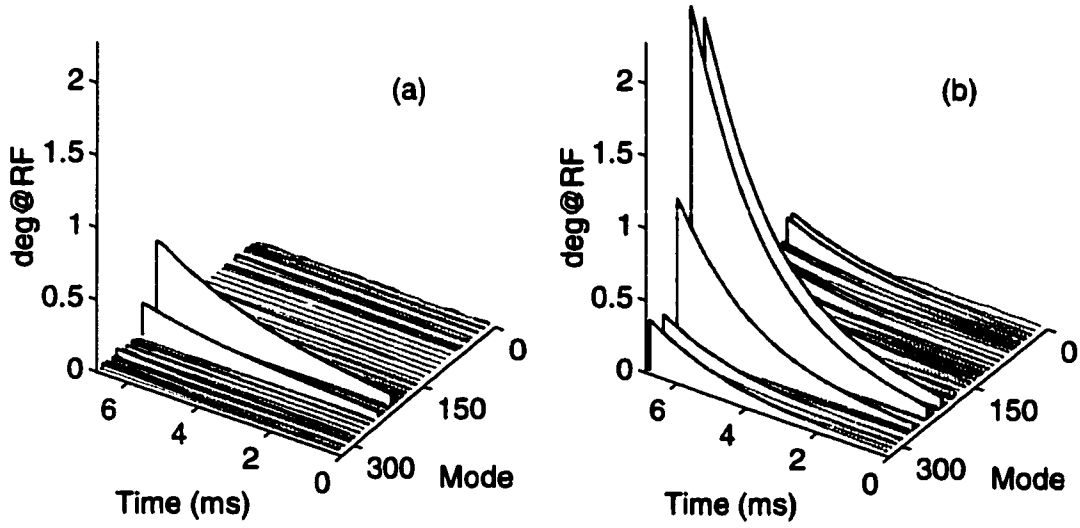

Figure 7.5: Projection onto EFEMs of measured ALS signal soon after feedback is stopped. a) Even fill, 158 mA. EFEMs 204 and 233 grow independently. b) Fill shown in Fig. 1, 156 mA. The composite eigenmode grows faster due to modulation coupling of unstable EFEMs 204 and 233.

uneven-fill mode grows faster than the EFEMs, since two unstable EFEMs are coupled. Figure 7.5(b) shows the projection of the growing uneven-fill signal onto the 328 EFEMs. We expect the dominant eigenmode to be some linear combination of v_{204} and v_{233}. In the figure, we also see projections onto v_{262}, v_{291}, v_{175}, etc., since the recorded DSP signal is proportional to the convolution of I_k and the modal spectrum. The measured eigenvalue is $\lambda = (0.59 \pm 0.02) - (1.6 \pm 0.03)j$ ms^{-1}, which is close to the theoretically predicted value of $(0.59 \pm 0.03) - (1.68 \pm 0.05)j$ ms^{-1}.

This piece of data illustrates the resolution of a problem that experimenters measuring instabilities often face. When the fill is significantly uneven, a single unstable eigenmode created by a single narrowband cavity resonance shows up as a multitude of sidebands in the beam spectrum. It is thus difficult to distiguish between multiple eigenmodes and multiple projections of a single eigenmode if one merely observes the beam spectrum. The problem is solved by digitising the growth of beam motion and calculating an eigenvalue for the projection of the motion onto each significant EFEM. If the calculated growth rates and coherent tunes of the projections are identical, it is fair to conclude that they all correspond to just one uneven-fill eigenmode. On the other hand, if the projections have different

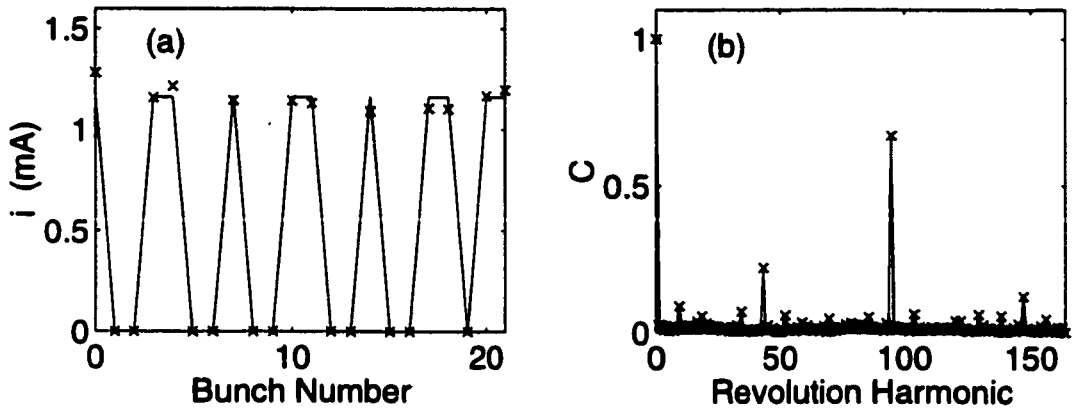

Figure 7.6: Uneven "Landau fill" that maximises C_{95}, (*i.e.* C_{233}) for $I_o = 175$ mA, $x = 0.47$.

eigenvalues, they must correspond to distinct eigenmodes.

7.6.2 ALS Landau Damping

Another cure for the instability of longitudinal modes 204 and 233 at the ALS is to use fill shapes with large values of C_{233} or C_{204}. Although both kinds of fills yield significant damping through the interbunch tune spread, the best choice is not immediately obvious. In general, if one uses the fill-optimising rule described earlier to maximise C_p, EFEM p is stabilised more than other EFEMs, since the combination of tune spread and fill unevenness introduces coupling to EFEM $N - p$. Since the growth rates of EFEMs 204 and 233 are comparable, the deciding factor in picking the best value of p is the amount by which the "non-p" EFEM is Landau-damped. A quick calculation based on Eqs. (7.17) and (7.18) shows that C_{233} is more effective than C_{204} when x is in the neighbourhood of 0.5. Intuitively, C_{233} produces more tune spread than C_{204} because $|\lambda_{233}| > |\lambda_{204}|$, and EFEM 204 is more easily Landau-damped because $|Im(\lambda_{204})| < |Im(\lambda_{233})|$.

Figure 7.6(a) shows the desired and measured bunch currents in a fill which optimises C_{233} for $x = 154/328$ and $I_o = 175.2$ mA. Figure 7.6(b) shows the desired and expected normalised fill spectrum, which is symmetric about $p = 164$. The difference between the expected and measured values of $C_{95} = C_{233}$ is negligible.

A baseline even-fill instability measurement was first made at $I_o = 172$ mA (see Fig. 7.7(a)). Exponential fits to the mode amplitudes and linear fits to the mode phases in this growing transient give the following eigenvalues: $\lambda_{204} = (0.47 \pm 0.02) - (0.05 \pm 0.03)j$ ms^{-1} and

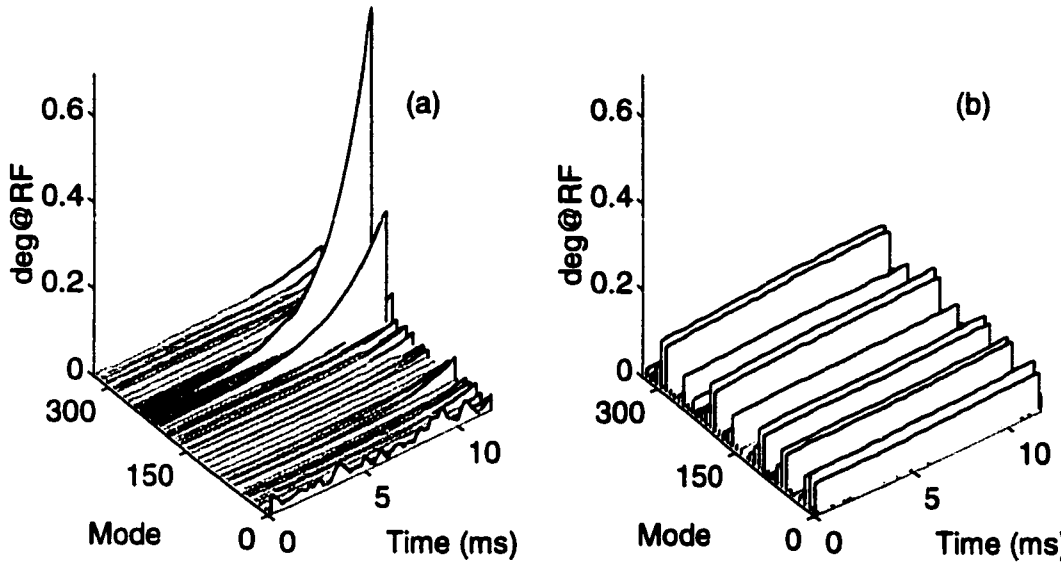

Figure 7.7: Measured growth of ALS longitudinal instabilities in EFEM basis. a) Even fill, 172 mA. b) "Landau fill" (see Fig. 3), 175 mA. Modes barely grow.

$\lambda_{233} = (0.61 \pm 0.02) - (1.16 \pm 0.03)j$ ms^{-1}. These numbers are not merely scaled up from the previously measured even-fill eigenvalues, since the two measurements were made on different days at different cavity tuner positions and temperature settings.

We can see from Fig. 7.3 that "Landau fills" with $x < 0.6$ almost completely damp the primary target mode, which is EFEM 233 in this case. Thus, any residual instability in the fill illustrated in Fig. 7.4 must correspond to the Landau-damped mode 204.

Although many methods exist for calculating the instability growth rate once the bunch tune distribution is calculated [26, 94], we use numerical computation of the eigenstructure of the mode coupling matrix, since it is the most exact. For this we need to know the shunt impedance R_s, the resonant frequency f_r and the quality factor Q of the two cavity modes responsible for the measured values of λ_{204} and λ_{233}.

If the effective impedance corresponding to an even-fill eigenvalue is $R + jX$, then the shunt impedance of the cavity mode is given by $(\frac{f_r}{f_{rf}}R_s - R)^2 + X^2 = [\frac{f_r}{f_{rf}}R_s]^2$ (see Fig. 2.4, and the accompanying discussion). By correlating this result with data on ALS cavity modes [77], we get (nominally) $R_s = 11.36$ kΩ, $f_r = 1809.69$ MHz and $Q = 2900$ for EFEM 204 and $R_s = 43$ kΩ, $f_r = 2852.92$ MHz and $Q = 9149$ for EFEM 233. The numerical calculation then gives us an eigenvalue of $(0.1 \pm 0.04) + (1.62 \pm 0.06)j$ ms^{-1} for the Landau-damped

mode 204. Error bars are calculated by assuming that errors in measured eigenvalues arise from fluctuations in f_r. Note: The real part of the most unstable eigenvalue is 6 times smaller than in the even-fill case.

Figure 7.7(b) shows a measurement of the growth of instabilities in a 175 mA beam whose fill shape is illustrated in Fig. 7.6. Although there is only one unstable eigenmode, the slowly-growing uneven-fill signal has projections onto many EFEMs of the form $[204 \pm n233]$ *modulo* 328. The measured eigenvalue is $(0.09 \pm 0.003) + (1.63 \pm 0.005)j$ ms^{-1}, which compares well with the theoretically predicted value of $(0.1 \pm 0.04) + (1.62 \pm 0.06)j$ ms^{-1}.

7.6.3 SPEAR

Uneven-fill cures have also been tested on longitudinal instabilities at the SPEAR storage ring, which has two 5-cell RF cavities with a rich spectrum of high-Q (see Eq. 2.12) parasitic resonances [85]. The SPEAR beam has an even-fill longitudinal threshold of 20 to 30 mA at the injection energy of 2.3 GeV. Considerable variation is seen in the threshold and spectral location of the instability over time scales of hours, as a result of variations in cavity temperatures and active-cavity tuner positions. The passive-cavity tuners are positioned so as to minimise the impedance seen by the beam [39]. The regular configuration, which has 54 of the 280 buckets filled at $I_o = 100$ mA and $E_o = 3$ GeV, is longitudinally stable due to increased radiation damping at the higher beam energy.

Although the SPEAR production fill was not specifically designed to reduce longitudinal instabilities, it turns out to have a spectral line at a modulation-coupling frequency. The measured spectrum of the production fill is shown in Fig. 7.8(a). The prominent peaks are at $p = 20, 60, 80$ and 100. Even-fill instability measurements [54, 55] performed using temporarily installed longitudinal feedback hardware indicate that EFEM 230 is usually the most unstable mode. Since $C_{100} = 0.75$ in the production fill, EFEM 230 is damped by modulation coupling to its stable counterpart EFEM 50 (see Eq. 7.15). At SPEAR, $h = 280$, so $[230+100]$ *modulo* $h = 50$. In general, EFEM n and EFEM $h-n$ are coupled if C_{2n} is large.

The proposed upgrade of SPEAR [99] aims for higher beam currents. This would necessitate activation of the second RF cavity and reduce tuner flexibility, if the present cavities are retained. In addition, bunch currents might exceed allowable limits if only 54 buckets are filled. Thus, if the same RF cavities are used, there is no guarantee of stability for SPEAR 3, unless active feedback is employed. However, one could experiment with new

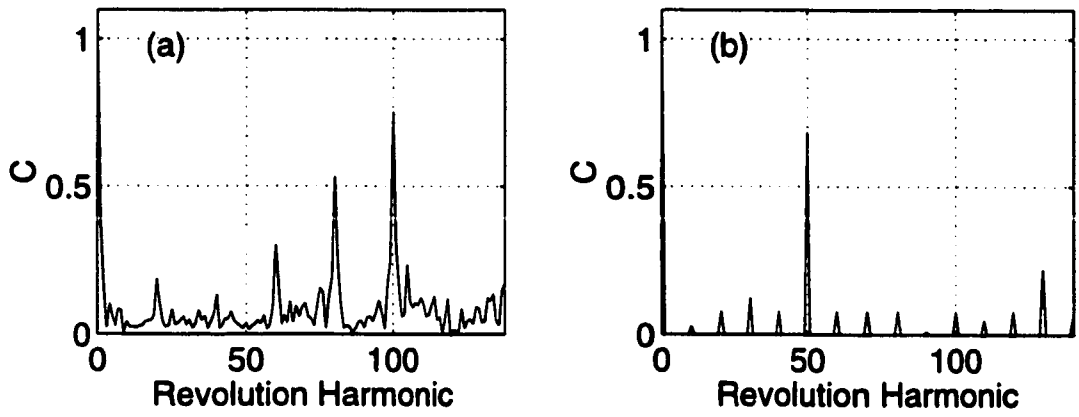

Figure 7.8: a) Measured fill spectrum of regular SPEAR fill. b) Calculated spectrum of SPEAR "Landau fill". Peak at C_{50} generates tune spread, doubles instability threshold.

fill patterns that raise the instability threshold, while simultaneously increasing the number of filled buckets. Figure 7.8(b) shows the calculated spectrum of a pattern that maximises C_{50}, *i.e.* C_{230}, for a 46% fill (130 buckets filled). At first, this fill was unsuccessful in raising the instability threshold at 2.3 GeV. Upon observing the signal from cavity probes, it was discovered that the idle cavity resonance that often drives EFEM 230 was far from its expected position. When this resonance was moved onto the relevant revolution harmonic by adjusting cavity tuners, all modes were completely stabilised at $I_o = 60$ mA (at least twice the even-fill threshold).

It seems unlikely right now that the present RF cavities will be retained in SPEAR 3 [100]. The discussion of optimal fills for SPEAR 3 will have to be revisited when the parameters of the new system are known.

7.6.4 PEP-II Landau Damping

As mentioned earlier, the fill that maximises C_1, *i.e.*, the conventional even fill with a gap at one end, is likely to be the best "Landau fill," in large high-current machines.

Offline RF cavity measurements [78] indicate an even-fill longitudinal threshold of 310 mA for the PEP-II High Energy Ring (HER). The measured threshold is at 550 mA due to the gap transient. Figure 7.9(a) shows a measurement of interbunch tune spread in a 605 mA uneven fill with a 7% gap (solid line). The tune spread has been extracted from feedback system data through Lorentzian fits to individual bunch spectra in a record of

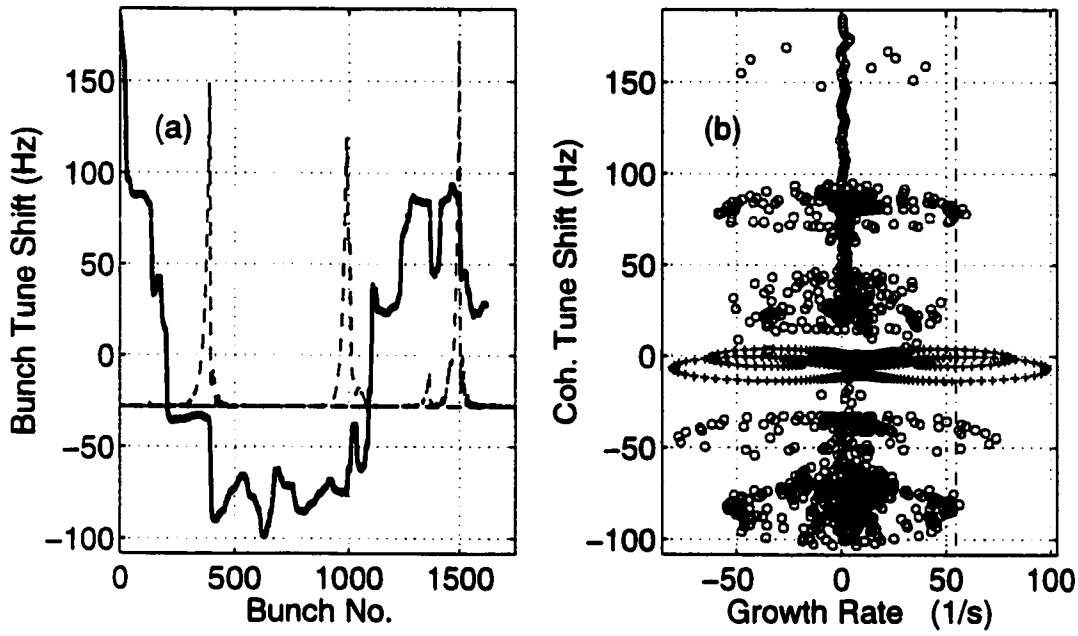

Figure 7.9: a) Solid line: Meaured PEP-II HER longitudinal tune spread due to 7% gap, I_o = 605 mA. Dashed lines: 3 representative eigenvectors (calculation). b) Calculated eigenvalues of 605 mA beam. '+'s: No tune spread. 'o's: Tune spread included. Dashdot line: Radiation damping.

growing beam motion[2]. Ideally, the tune distribution should be smooth. However, this measurement was made with a ragged looking fill before the RF feedback loops were fully commissioned.

The nominal eigenvalues of this HER beam (calculated using matrix reduction [98]) are shown as circles in Fig. 7.9(b). We see that eigenvalue clusters coincide in tune with moderately flat regions in the bunch tune graph shown in Fig. 7.9(a). Only a few eigenvalues are unstable. The largest calculated eigenvalue is (73 ± 15) s^{-1}. The uncertainty in this calculation comes largely from error bars on the estimated cavity impedance [101]. Representative eigenvectors, shown as dashed spikes in Fig. 7.9(a), are localised to short segments within the bunch train that correspond to flat regions in the tune distribution. These segments are decoupled from the rest of the train by the tune spread. The calculated

[2]These are not the true incoherent bunch tunes. However, we assume that the coherent motion of the bunches is made up of localised eigenvectors whose coherent tunes approximately equal the average incoherent tune of the participating bunches. This is valid only if the impedance is broadband and the incoherent tune spread is much larger than the coherent tune shift. Locally flat regions in Fig. 7.9(a) reflect eigenvector localization.

eigenvalue agrees with the measured eigenvalue of (60 ± 5) s^{-1}. The imaginary parts of these eigenvalues are not considered, since the single-bunch tune for this day is not known with sufficient accuracy.

7.7 Conclusion

The theory described in this chapter opens the door to further theoretical and experimental investigations of the stability of longitudinal and transverse dipole modes and bunch shape oscillations. It reveals a new cure for coupled-bunch instabilities, which has been verified experimentally at ALS, SPEAR and PEP-II.

The uneven-fill cure increases in efficacy, as the maximum allowable bunch current i_{max} increases, and the beam current I_o is distributed among fewer and fewer buckets. Factors that limit i_{max} include heating of vacuum chamber elements, intrabunch scattering, and beam-beam effects (in colliders).

Chapter 8

Conclusion, Ideas for Further Study

The goal of this thesis has been to exploit the capabilities of the SLAC/LBNL/INFN-LNF longitudinal feedback system, and to close the gap between what is theoretically predicted, and what can be conveniently measured, in the field of coupled-bunch instabilities. The pursuit of this goal has also led to the development of new theoretical approaches to the instability problem. This chapter briefly summarises the previous chapters, and suggests directions for further study.

A method has been described for extracting bunch currents and multi-bunch synchronous phases from feedback system data. Since the two are related through the beam impedance, such measurements constitute a novel beam-based impedance measurement. This impedance measurement technique has been used to identify the cause of a low-mode instability in PEP-II. Knowledge of multi-bunch currents and phases is also useful in itself, since the two distributions affect instability growth rates, feedback performance, and collider luminosity, among other things.

The above-mentioned technique for impedance measurement is limited by the fact that we do not see the "synchronous phase" of the empty buckets. Also, noise and systematic errors currently get in the way of measuring impedances that induce synchronous phase transients much smaller than 1 deg@RF. A dedicated diagnostic system with higher sensitivity than the regular feedback system should be able to measure smaller impedances.

Modal decompositions of digitised data have been used to study longitudinal and transverse coupled-bunch transients, induced by switching the feedback state. Transient records

of beam motion, lasting just 10-20 ms, have been analyzed offline to reveal the frequencies and growth rates of unstable modes, the damping rates of stable modes, and the effective impedance of the feedback system, under a variety of beam conditions. The measured longitudinal growth and damping rates at the ALS and PEP-II are in agreement with projections based on the estimated cavity impedance at the two machines. Longitudinal grow-damp data have been used to minimise ALS instability growth rates by optimising cavity temperatures.

It has been shown that offline analysis can also be used to reconstruct the full phase space trajectories of the oscillating bunches, or modes, from LFB system data. Phase space tracking provides a complete description of beam motion, and allows new kinds of comparisons between theory and experiment. The technique has been used to accurately measure the imaginary part of the effective beam impedance at the ALS, and confirm qualitative predictions about uneven-fill dynamics at the ALS and PEP-II. It also facilitates measurement of the reactive component of feedback, which should be minimised, for optimal utilization of the available feedback power.

Measurement of nonlinear effects, such as amplitude-dependent growth rates and tune shifts, should be possible with this method.

Time-domain phase space tracking has been used to study a low-threshold vertical instability in the PEP-II HER. The trajectories fail to match qualitative features described in the existing literature on FBII theory. The method shows promise as a tool for analyzing data from future FBII experiments, and for revealing aspects of instability growth that have hitherto remained unexamined. The diagnostics shown in this thesis are probably just the first harvest from the field of possibilities opened up by the LFB system. More extensive transverse measurements, as well as beam-beam experiments with similar data recording, are sure to bear fruit.

The initial ALS grow-damp experiments provided excellent illustrations of the fact that the dynamics of hundreds of bunches can sometimes be represented by just a few independent variables. The use of "modal" projections, in mathematical descriptions of Landau damping and uneven-fill dynamics, is a direct consequence of the attempt to formalise this concept.

A matrix-based method has been developed, for calculating the effect of bunch-to-bunch tune shifts on coupled-bunch instabilities. It has more general applicability than the dispersion relation that is conventionally used to describe Landau damping, and is especially

useful in situations where many even-fill modes are coupled to each other.

A related mathematical formalism, based on projections of beam motion onto even-fill eigenmodes, has been used to simplify the uneven-fill instability problem. Uneven-fill longitudinal dynamics have been explained in terms of two physical phenomena: modulation coupling of strong even-fill eigenmodes (EFEMs) and Landau damping from fill-induced interbunch tune spreads. The former effect is also present in the transverse plane.

Based on the above-mentioned concepts, a simple, easily implementable algorithm has been designed, for shaping fills to cure instabilities. The new uneven-fill cure has been verified experimentally at the ALS, PEP-II and SPEAR.

The uneven-fill theory opens the door to further investigations of the stability of longitudinal and transverse dipole modes and bunch shape oscillations.

Appendix A

Computer Programs for Data Analysis

This appendix lists some of the MATLAB codes that have been used to analyse LFB system data, and produce most of the figures shown in this thesis. The listings are preceded by functional descriptions of the main programs, and some information that new users might need, to use the programs effectively. The material presented here complements the tutorial on data analysis codes prepared by D. Teytelman [102].

The main programs are `synchp.m`, `Modes.m`, `PhasSpac.m`, and `sideband.m`. Except for `PhasSpac.m`, which must be executed after `Modes.m`, they all draw their data directly from a file called `gd.mat`, which is in MATLAB's data format. The file contains the following 11 variables[1]: **bunches**, **beamCurrent**, **damp_brkpt**, **downsamp**, **gains**, **phases**, **rf_freq**, **ring_size**, **shift_gain**, **taps**, and **turn_offsets**.

bunches is a matrix that contains the sampled oscillations of the bunches in the ring, in units of ADC counts. Each column contains the signal of a single bunch, sampled once every **downsamp** turns. The turn number of the initial sample of each bunch is stored in the **turn_offsets** vector. **rf_freq** is the ADC sampling frequency, in MHz. If the LFB system processes bunches with a minimum spacing of b buckets, then **rf_freq** $= f_{rf}/b$, and the harmonic number equals $b \times$ **ring_size**. The total beam current in mA is stored in **beamCurrent**.

The grow-damp measurements involve switching to a new set of feedback parameters,

[1] More variables will be added soon, to accommodate the use of infinite impulse response (IIR) filters in the feedback algorithm.

and then returning to the original parameters. Data recording starts after the first switch, and ends after the second. For example, in a typical grow-damp experiment, data recording starts soon after feedback is switched off. The break point is the point in the data at which negative feedback is restored.

`damp_brkpt` is the scaled break point. The gains and phases of the finite impulse response (FIR) feedback filters are stored in **gains** and **phases**, and **taps** is the length of the filter. `shift_gain` is the number of bits by which the DSP output is left-shifted, to augment the filter gain.

A.1 synchp.m

The program `synchp.m` has been used to calculate bunch currents and synchronous phases from the data in `gd.mat` (see Ch. 3). The program assumes that there is some low-frequency motion that is common to all the bunches, in the frequency range from 0 to `fCM`. The variable `fCM` is initialised in `Recogn.m`. This "motion" could come from power supply ripple in the klystron, or from a small-amplitude external modulation of the LFB front end phase shifter.

The subprogram `curr_mon.m` calculates and plots the average of the low-frequency bunch spectra, from 0 to `fCM`. The user enters the number of frequencies in the plot that seem to represent coherent bunch motion. The bunch signal amplitudes at these frequencies are used to make an initial estimate of the bunch currents.

The earlier version of `synchp.m` calculated synchronous phases by dividing the mean bunch signals by the product of the calibration (ADC counts per mA-deg) and the bunch currents returned by `curr_mon.m`. The new version uses a more accurate method, based on the fact that the average bunch signal is proportional to $i_k \sin(H\phi_k)$, and not to $i_k\phi_k$, where i_k and ϕ_k are the bunch current and synchronous phase (relative to the LFB phase reference), respectively. H is the harmonic of f_{rf} at which phase detection is performed. It also takes into account the fact that the low-frequency signal amplitudes are proportional to $i_k \cos(H\phi_k)$, and not to i_k. Phases are not calculated for bunches with detected currents below a threshold that is fixed in `synchp.m`.

The program plots the bunch currents and synchronous phases. It also plots the averages of the bunch signals, which help in identifying saturation of the ADC.

`synchp.m` works well only when:

– There is some power supply noise or external phase modulation on the bunch signal.

The low-frequency motion should be small compared to the detection range of the front end, but large enough to be detected above the noise floor.

- The `FE_freq` variable is set correctly in `Recogn.m`. `FE_freq` = H = (front end phase detection frequency)/f_{rf}.
- The `mix_ofst` (mixer offset) variable is set correctly in `Recogn.m`. `mix_ofst` = average DSP signal of the empty buckets, in ADC counts.
- The `calib` variable is set correctly in `Recogn.m`. This is the front end calibration in ADC counts per mA-deg@RF.

The subprogram `synchTF.m` is optionally called by `synchp.m`. It calculates and plots the transfer function from bunch currents to synchronous phases, at all frequencies that have a good signal-to-noise ratio. Please see Ch. 3 for equations describing the relation between this transfer function and the beam impedance. Ch. 3 also describes conditions under which the estimated impedance is believable.

A.2 Modes.m

`Modes.m` is used to perform modal analysis of recorded multi-bunch motion, and plot the results. Figure 4.2 is an example of a figure generated by `Modes.m`.

The program starts by loading oscillation data, and initialising machine-specific parameters. The oscillation frequency `fr_osc` is then detected, by locating the peak of the averaged bunch spectra. The bunch signals are converted to units of deg@RF, based on the calibration stored in `Recogn.m`. Before changing units, the program looks for bunch currents in a file called `BunCurr.mat`. If this file is not found, all of the sampled bunches are assumed to have equal charge, for the purposes of unit conversion.

The signals are then passed through a band pass filter, to remove some of the noise. The filter width is determined by variables initialised in `Recogn.m`. See the first few lines of code in `params.m`, for the relation between the variables `freq_ratio` and `delta_f`, and the width of the band pass filter. A smooth gaussian filter is employed, so as to minimise transients at the beginning and end of the data record.

The bunch signals are also delayed by appropriate amounts, to compensate for the sampling time offsets. For example, if bunch k is sampled on turn T_k , and if bunch l is

sampled on turn T_l, then signal of bunch k could be delayed by $\Delta t = [(T_l - T_k + (l-k)/h]T_o$, to simulate simultaneous sampling. Alignment of the sampling instants is necessary for projecting the motion onto EFEMs. The delay is implemented by means of a Fourier-domain phase shift that increases linearly with frequency. To compensate for aliasing, transverse signals are given an additional phase shift $\omega_o \Delta\nu \Delta t$, where $\Delta\nu$ is the difference between the true and aliased fractional tunes.

Before projecting the motion onto the modal basis, the real bunch signals, which have the form $a\cos(\omega t)$, are converted to complex signals of the form $a\,e^{j\omega t}$, by means of a Hilbert transform (see Eq. 5.3). The program then calculates and plots the evolution of the bunch oscillation amplitudes, *i.e.*, the oscillation envelopes. Modal projections are calculated by taking the FFT (Fast Fourier Transform) of the vector of bunch oscillation coordinates, on each turn[2]. We thus obtain the complex amplitude of all the EFEMs, on each turn. The evolution of the magnitudes of the EFEMs is plotted next to the graph of bunch envelopes.

The program offers to fit exponentials to the portions of the modal transients before and after the break point. If the offer is accepted, it prompts the user for the fit cutoff. Exponential fits and growth rates are plotted for all modes whose amplitude exceeds the fit cutoff. The user must set the cutoff at a level that is high enough to exclude modes in the noise floor.

Note: The function used to fit the transients is not a simple exponential. Rather, it is an exponential added in quadrature to a constant noise level. Thus, there are three fit parameters - initial amplitude, growth rate, and noise amplitude.

A.2.1 checkfit.m

It is very important to visually check the quality of exponential fits, to make sure that the calculated growth rates are believable. This can be done by executing `checkfit.m` after `Modes.m` terminates. The `checkfit` program produces a two-dimensional graph of the exponential fits, superimposed on the measured modal transients. It also prints the calculated growth rates and rms fit errors to the command line. Clearly, the rates are accurate only if the transients coincide with the fits.

Most commonly, the quality of the fits is affected by nonlinear effects, such as saturation of feedback, or of the instability mechanism. The modal transient becomes non-exponential as a result of such saturation. The solution is to move the break point, so that fits are

[2]The term "FFT" refers to a specific algorithm for implementing the DFT (Discrete Fourier Transform).

only calculated to the small-amplitude portion of the transient, where the dynamics are linear, and the evolution is exponential. After `Modes.m` has been executed (with or without `checkfit.m`), the break point can be changed to sample number n by typing `breakpt = n` at the command prompt. The user can then redo the fits before or after the break point by typing `modefitg` or `modefitd`, respectively.

In short, exponential fits are valid *only* when oscillations are small, and the dynamics are linear.

Transients that rise sharply out of the noise floor just before the break point sometimes cause problems, because the fit routine may not have enough data to do a good job. In this case, as in all other cases, the user must exercise discretion in judging the accuracy of the fit, based on the results of `checkfit.m`.

The most subtle problems arise from the fact that uneven-fill eigenmodes are sometimes very different from EFEMs. Thus, the FFTs performed by `Modes.m` do not in general constitute a modal decomposition. In other words, the Fourier vectors are not necessarily eigenvectors of uneven-fill motion, as explained in Ch. 7. Thus, a single "modal" transient could actually be a superposition of two or more uneven-fill eigenmodes, which beat against each other, and fail to evolve exponentially. In such cases, exponential fits can be misleading.

A.2.2 PhasSpac.m

The `Modes` program only displays the mode magnitudes, *i.e.*, the absolute values of the complex phase space coordinates of the modes. The evolution of the phase space angles of the modes can be studied by executing `PhasSpac.m` after `Modes.m` has terminated.

The program first asks the user to specify the section of the transient that is to be examined. Section 1 is the section before the break point, and section 2 is contains the portion after the break point. The sizes of sections 1 and 2 can be changed by moving the break point, as shown above. The user is then prompted for a vector containing the modes of interest.

It would not be very helpful to plot the evolution of the phase space angles directly, since the modes generally complete a revolution in phase space every few samples. Instead, the program plots the normalised angles, which are calculated by subtracting the angle of a fictitious reference mode (see Ch. 5). The reference mode oscillates exactly at the nominal synchrotron or betatron frequency `fr_osc`, which is the frequency detected by `Modes.m`.

Analysis of longitudinal modes is simpler, since longitudinal tunes are not aliased by

downsampling. Thus, a longitudinal reference mode is always an upper sideband, whose angle evolves as $2\pi \times \texttt{fr\_osc} \times t$. However, in 50% of the cases, aliasing converts transverse upper sidebands into lower sidebands. More precisely, aliasing inverts the sidebands when the fractional tune lies between $k/2D$ and $(k+1)/2D$, where D is the downsampling factor, and k is an odd integer. When this happens, `Modes.m` detects the inversion and responds by working with analytic signals of the form $a\,e^{-j\omega t}$, so as to avert the confusion that would be created if any mode l were to be mistaken for its counterpart mode $-l$. In such cases, the transverse reference mode must be a lower sideband, whose angle evolves as $-2\pi \times \texttt{fr\_osc} \times t$.

`PhasSpac.m` prompts the user for the desired reference mode. An upper sideband can be chosen by typing `u`, and a lower sideband by typing `l`. Alternatively, the angles are normalised by subtracting the angle transient of an actual, measured mode, that is chosen as a reference mode. This can be done by entering the reference mode number, instead of `u` or `l`. The uses of this program are discussed in Ch. 5.

The plotted phase space angles sometimes have ripples at the beginning or end of the data record. This happens because of transients from the band pass filtering performed by `Modes.m`. If the transients are large enough to cause problems, the bandwidth of the filter should be increased by changing the appropriate variable in `Recogn.m`. The variables `freq_ratio` and `delta_f` control the widths of filters for longitudinal and transverse data, respectively.

A.3 sideband.m

The program `sideband.m` is sometimes used as an alternative, or complement to `Modes.m`. It useful when one needs to examine the beam spectrum. For example, Figs. 4.5 and 4.6 have been generated using `sideband.m`.

This program works on the same principle as an FFT analyser, which is a kind of spectrum analyser. An FFT analyser samples and stores the incoming signal, and then takes the FFT of sampled vectors, to calculate the beam spectrum. A typical FFT analyser has a sampling rate in the region of 10 MHz. The processed vectors are roughly 10000 samples long, so the frequency resolution is of the order of 1 kHz, when the full frequency span is used.

Data stored by the LFB system is sampled at 500 MHz. Since the data records are

10–25-ms long, frequency resolutions of 100–40 Hz are achievable. Note: Downsampling increases the frequency resolution that can be achieved, with a fixed amount of memory space.

sideband.m first removes the DC component of each individual bunch signal, so that revolution harmonics are suppressed in the spectrum. This is useful, because it is often hard to distinguish between revolution harmonics and sidebands, in a conventional beam spectrum.

The program then creates an "FFT-able" vector, *i.e.*, a vector of continuous samples, from the raw DSP data. This is done by upsampling the bunch data — to estimate the bunch signals at each turn — and then laying the samples out in a continuous sequence, as seen by a BPM. The interpolated samples are calculated using MATLAB's **interp** function.

The FFT of the above-mentioned vector is a beam spectrum with suppressed revolution harmonics, also known as a beam pseudospectrum.

If steady state motion is recorded, it is useful to analyse the entire data record as a single block. However, when analysing grow-damp data, it is useful to be able to look at the spectrum of specific segments of the growing or damping transients. For this reason, the program only calculates the pseudospectrum of the section of data specified by the user. The user specifies a section by entering the starting sample number and the final sample number. To analyse the entire data record, the user enters **1**, for the starting number, and **M**, for the final sample number. **M** is a variable in the workspace that contains the number of samples per bunch, in the **bunches** matrix.

The pseudospectrum is plotted, together with the difference of upper and lower sideband amplitudes at each revolution harmonic. The zoomed spectrum around the 8 revolution harmonics with the largest sideband asymmetry is plotted on a separate figure. Presumably, these 8 sidebands represent the most unstable modes.

A.3.1 Memory Limitations

Of all the data analysis programs described here, **sideband.m** requires the most memory space. Consider a machine with $h = 400$, and $D = 20$. If the data record in **gd.mat** contains $M = 1000$ samples of each bunch, then the size of the time domain vector of continuous samples is $D \times M \times h = 8 \times 10^6$. Since MATLAB uses 8 bytes for each number, this vector alone would require 64 MB of memory space. In addition, the program needs 64 MB for the calculated spectrum, and a comparable amount of space for storing intermediate

computation results.

In practice, it is advisable to use a computer with as much RAM as possible, and then run `sideband.m` on a suitable subset of the data record. If the RAM capacity of the machine is as low as 32 MB, it is sometimes tedious even to run `Modes.m` on data from large machines, such as PEP-II. A capacity of 128 MB is generally sufficient, although it really helps to have even more, when analysing PEP-II data.

A.4 Program Listings

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% synchp.m

% Loads bunch data, removes turn offsets and calculates bunch
% currents and synchronous phases.
% Plots the currents, the averages of the bunch signals (check
% for ADC saturation), and the synchronous phases.

% Copyright 2000 by Shyam Prabhakar

clear
global bun_curr mb Io
loadgd
bcmin = 0.1;
CALIB = calib*180/pi; % Counts per mA-rad@rf

[bunches,M] = rmTrnOfst(bunches,turn_offsets,downsamp);
mb1 = mean(bunches);
mb = (mb1 - mix_ofst)*FE_freq/CALIB;
bunches = bunches - ones(M,1)*mb1;

curr_mon
scal = fmin('curr_scal',1,10,[0 1e-5]);
bun_curr = bun_curr/scal;
```

```
synch_p = atan2(mb,bun_curr)*180/pi/FE_freq;
bun_curr = (mb.^2 + bun_curr.^2).^0.5;
v = find(bun_curr<max(max(bun_curr)/10,bcmin)|mb1>126|mb1<-127);
if ~isempty(v),
  synch_p(v) = zeros(size(synch_p(v)));
end

figure(2), clf, orient tall
subplot(3,1,1), plot(0:N-1,bun_curr,'g','linewidth',1), grid on
ax = axis; ax(3)=0; axis(ax)
title([Dir ': Bunch Current Monitor, Io=' num2str(Io) ...
       'mA' ', nCM=' num2str(round(nCM))])
ylabel('mA')

subplot(3,1,2), plot(0:N-1,mb1,'g','linewidth',1), grid on
ax = axis; ax(3) = max(ax(3),-128); ax(4) = min(ax(4),127); axis(ax);
title('Averages of bunch signals')
ylabel('ADC counts')

synch_p1 = synch_p;
q_medflt = 'y'; med_win = 1;
while strcmp(q_medflt,'y'),
  u = find(~synch_p1);
  synch_p1(u) = NaN + synch_p1(u);
  subplot(3,1,3), plot(0:N-1,synch_p1,'g','linewidth',1), grid on
  title('Synchronous phase (relative to reference oscillator)')
  xlabel('bunch number'), ylabel(osc_unit)
  q_medflt   = input('Median filter synch_p? (y/n): ','s');
  if strcmp(q_medflt,'y'),
    med_win  = input('Enter median filter window: ');
    synch_p1 = med_flt(synch_p,med_win);
  end
end
```

```
synch_p  = med_flt(synch_p,med_win);

bnum = find(mb1>126|mb1<-127)-1;
if size(bnum),
  for j = 1:3,
    subplot(3,1,j), hold on, ax = axis;
    y = ax(3:4)'*ones(size(bnum));
    plt = plot(ones(2,1)*bnum,y,'-.','linewidth',0.25);
    set(plt,'col',[0.75 0 0])
end, end

q_TF = input('Plot Transfer function? (y/n): ','s');
if strcmp(q_TF,'y'), synchTF, end
```

****

```
% loadgd.m

% This program loads data from the file gd.mat, and does the
% initial sorting and renaming of variables. Calls Recogn.m to
% initialise machine-specific parameters. Calls Reduce.m to
% reduce the "bunches" matrix to it's smallest possible width.
% Automatically called at the start of Modes.m and synchp.m

% Copyright 2000 by Shyam Prabhakar

fprintf(['\n---- ', pwd, ' ----\n'])
fprintf('\nLoading and sorting by bunch no. ...\n')

load gd
[Dir,Year] = getDir;
Recogn

[M,N]    = size(bunches);
```

```
[R,C]    = find(bunches == -1000 | bunches == 12345678);
if ~isempty(R),
  R = min(R); bunches(R:M,:) = []; M = R-1;
end
lto      = length(turn_offsets);
if N~=lto,
  if N<lto,
    fprintf('\nWarning: turn_offsets has more columns than bunches\n')
  else
    fprintf('\nCAUTION!! bunches has more columns than turn_offsets\n')
    bunches(:,lto+1:N)=[];   N = lto;
end, end

Reduce

[Y,I]    = sort(bunches(1,:));
bunches  = bunches(2:M,I);
M        = M-1;
Io       = beamCurrent;
while Io <= 0, Io = input('Enter estimated Io (in mA): '); end

breakpt  = round(samples_per_bunch*damp_brkpt/63);
if gains(2) == gains(1) & phases(2) == phases(1),
  breakpt = M;
end
harm_no    = ring_size/red_fact;
Frf        = rf_freq/red_fact*1e6;
fitrecord  = [0 0]; PDpos = 0;
cnt        = ceil(harm_no/15);
fsamp      = Frf/harm_no/downsamp;
clear beamCurrent damp_brkpt ring_size rf_freq lto Y I R C

*******************************************************************************
```

```
% getDir.m

% Initialises Dir and Year for labeling and other
% purposes. Assumes that it is being called from a time-
% stamped directory of the form mmmddyy/tttt. For example,
% jan3098/1024.

% Copyright 2000 by Shyam Prabhakar

function [Dir,Year] = getDir
global New

str      = pwd;
lstr     = length(str);
if New, fsep = filesep;
else,   fsep = '/';
end
fseppos  = find(str==fsep);
fseppos2 = fseppos(length(fseppos));
strchk   = str2num(str(fseppos2-2:fseppos2-1));
if isempty(strchk),
  fseppos1 = fseppos(length(fseppos)-2);
  fseppos2 = fseppos(length(fseppos)-1);
else,
  fseppos1 = fseppos(length(fseppos)-1);
end
Dir      = str(fseppos1+1:lstr);
if nargout > 1,
  Year     = str(fseppos2-2:fseppos2-1);
  Year     = str2num(Year);
  if isempty(Year),
    Year    = 999;
end, end
```

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% Recogn.m

% Initialises a list of machine-specific parameters like
% calibration, data size, etc. Automatically called by Modes.m
% and synchp.m (via loadgd.m).

% Copyright 2000 by Shyam Prabhakar

if exist('q_trans')~=1,
  q_trans = input('Is this transverse data (y/n)?: ','s');
end
if strcmp(q_trans,'y'),
  if exist('tun_guess')~=1,
    tun_guess = input(['Enter estimate of fractional ' ...
                       'betatron tune: ']);
  end
  tun_alias = fix(2*tun_guess*downsamp);
  osc_unit  = 'mm';
else
  tun_alias = 0;
  osc_unit  = 'deg@RF';
end
fit_ord = 4; % Order of polynomial fits
FE_freq = 6; % Detection at 6*Frf

if ring_size == 328,
  machine  = 'ALS';
  calib    = 7.88; % counts/mA/osc_unit
%  calib    = 21.2;
  Uo       = 134.1e3;
  WIN_dflt = 30;
  n_pts_1  = 30;
```

```
  n_pts_2  = 70;
  samples_per_bunch = 1008;
  if strcmp(q_trans,'y'),
    delta_f    = 3.5e3;
    fCM        = 3e3;
  else
    freq_ratio = 1.2;
    fCM        = 2.5e3;
  end
  mix_ofst     = -16.5;
elseif ring_size == 280,
  machine  = 'SPEAR';
  calib    = 8.5;
  Uo       = 560e3;
  WIN_dflt = 30;
  n_pts_1  = 30;
  n_pts_2  = 60;
  samples_per_bunch = 2267;
  if strcmp(q_trans,'y'),
    delta_f    = 4.5e3;
    fCM        = 300;
  else
    freq_ratio = 1.05;
    fCM        = 10e3;
  end
  mix_ofst     = 0;
elseif ring_size == 1746,
  WIN_dflt = 30;
  n_pts_1  = 20;
  n_pts_2  = 45;
  if exist('q_m')~=1,
    q_m      = input('Is this HER data (y/n)?: ','s');
  end
```

```
  if strcmp(q_m,'n'),
    machine = 'PEP-II LER';
    Uo      = 770e3;
    samples_per_bunch = 610;
    if strcmp(q_trans,'y'),
      delta_f    = 2e3;
      fCM        = 800;
      calib      = 280.6; % cnts/mA-mm
    else
      freq_ratio = 1.3;
      fCM        = 2e3;
      calib      = 8.9;
    end
    mix_ofst     = -4;
  else
    machine = 'PEP-II HER';
    Uo      = 3.58e6;
    samples_per_bunch = 661;
    if strcmp(q_trans,'y'),
      delta_f    = 2e3;
      fCM        = 800;
%      calib      = 280.6; % cnts/mA-mm
      calib      = 280.6/sqrt(10); % 10 dB pad
    else
      freq_ratio = 1.2;
      fCM        = 3e3;
      calib      = 16.06;
    end
    mix_ofst     = -4;
  end
elseif ring_size == 120,
  machine  = 'DAFNE';
  calib    = 1.166;
```

```
  Uo       = 134.1e3; %To be changed later
  WIN_dflt = 20;
  n_pts_1  = 200;
  n_pts_2  = 300;
  samples_per_bunch = 4032;
  if strcmp(q_trans,'y'),
%   Filter bandwidth for TRANSV, default = 6e3
    delta_f    = 6e3;
    fCM        = 300;
  else
%   Filter freq_ratio for LONG = f_high/fsynch = fsynch/f_low
    freq_ratio = 1.1;
    fCM        = 10e3;
  end
  mix_ofst     = 0;
elseif ring_size == 468,
  machine  = 'PLS';
  calib    = 11.02;
  Uo       = 223e3;
  WIN_dflt = 30;
  n_pts_1  = 30;
  n_pts_2  = 60;
  samples_per_bunch = 1763;
  if strcmp(q_trans,'y'),
    delta_f    = 3e3;
    fCM        = 300;
  else
    freq_ratio = 1.3;
    fCM        = 3e3;
  end
  mix_ofst     = 0;
elseif ring_size == 400,
  machine  = 'BESSY II';
```

```
  calib     = 13.4;          %sk% to be changed later
  Uo        = 900e3;        %sk% check the log files
  WIN_dflt = 30;
  n_pts_1   = 30;
  n_pts_2   = 60;
  samples_per_bunch = 1588;
  if strcmp(q_trans,'y'),
    delta_f     = 3e3;
    fCM         = 300;
  else
    freq_ratio = 1.3;
    fCM         = 1e3;
  end
  mix_ofst      = 0;
else
  machine   = '?????';
  calib     = 15;
  Uo        = 150e3;
  WIN_dflt = 30;
  n_pts_1   = 30;
  n_pts_2   = 45;
  samples_per_bunch = 1e3;
  if strcmp(q_trans,'y'),
    delta_f     = 6e3;
    fCM         = 300;
  else
    freq_ratio = 1.3;
    fCM       = 2e3;
  end
  mix_ofst      = 0;
end

fprintf(['\nRecognised ' machine ' ring\n'])
```

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% Reduce.m

% Reduce size of "bunches" matrix. Example: If every 4th bucket is
% filled and the fill starts at bucket 3, the user can select
% bunches 3, 7, 11, 15, ... by entering a downsampling factor of
% 4 and a starting bucket number of 3 at the prompts. Automatically
% called by synchp.m and Modes.m (via loadgd.m).

% Copyright 2000 by Shyam Prabhakar

if exist('r_inp')~=1,
  r_inp  = input('Downsample columns of ''bunches'' ? (y/n): ','s');
end
if strcmp(r_inp,'y'),
  if exist('red_fact')~=1,
    red_fact = ring_size + 1;
    while rem(ring_size,red_fact),
      red_fact   = input('Enter the reduction factor: ');
      red_fact   = round(red_fact);
    end
  else
    if rem(ring_size,red_fact), error('Illegal bunch spacing!!!'), end
  end
  if exist('red_start')~=1,
    red_start    = input('Starting bucket# (1st bucket = #1): ');
  end
  red_start    = rem(red_start-1,red_fact) + 1;
  bnums        = bunches(1,:);
  bnumDS       = red_start:red_fact:max(bnums);
  bnumidx      = find(rem(bnums-red_start,red_fact)==0);
  if max(size(bnumidx)) < max(size(bnumDS)),
    error('Some of the desired buckets are not sampled!!!');
```

```
  end
  bunches      = bunches(:,bnumidx);
  turn_offsets = turn_offsets(bnumDS);
  N            = max(size(bnumidx));
else
  red_fact     = 1;
end

*******************************************************************************

% rmTrnOfst.m

% Removes the sampling time offsets between the bunches, using
% filters that implement the appropriate fractional sample delay.
% The sampling time offsets are in the turn offsets vector "to".
% Does not take aliasing of transverse data into account.

% Copyright 2000 by Shyam Prabhakar

function [bunches,M] = rmTrnOfst(bunches,to,D)

% Realign all the samples so that each bunch is sampled at turn D-1
% There will be an error if max(fix_to) > 2*l

fprintf('\nInterpolating and resampling data ...\n')

[M,N]    = size(bunches);
l        = 4;
alpha    = .5;
select   = zeros(M,1);
extrapol = zeros(M+2*l,1);
resamp   = zeros(M+2*l,1);

s1  = toeplitz(0:l-1) + eps;
```

```
s2  = hankel(2*l-1:-1:1);
s2p = hankel([1:l-1 0]);
s2  = s2 + eps + s2p(l:-1:1,l:-1:1);
s1  = sin(alpha*pi*s1)./(alpha*pi*s1);
s2  = sin(alpha*pi*s2)./(alpha*pi*s2);
ap  = s1 + s2;
am  = s1 - s2;
ap  = inv(ap);
am  = inv(am);
d   = zeros(2*l,1);
d(1:2:2*l-1,:) = ap + am;
d(2:2:2*l,:)   = ap - am;
x   = (0:D-1)/D;
y   = zeros(2*l,1);
y(1:2:2*l-1)   = (l:-1:1);
y(2:2:2*l)     = (l-1:-1:0);
X1  = ones(2*l,1);
X1(1:2:2*l-1)  = -ones(l,1);
XX  = eps + y*ones(1,D) + X1*x;
y   = X1 + y + eps;
h   = .5*d'*(sin(pi*alpha*XX)./(alpha*pi*XX));
b   = zeros(2*l*D+1,1);
b(1:l*D) = h';
b(l*D+1) = .5*d(:,l)'*(sin(pi*alpha*y)./(pi*alpha*y));
b(l*D+2:2*l*D+1) = b(l*D:-1:1);
b   = b(1:2*l*D);   % Since b(2lD+1)=0. Be careful about indices now!

rem_to = rem(to,D);
fix_to = floor(to/D);
for k  = 1:N,
  select       = bunches(:,k);
  uppr_extrap  = 2*select(1) - select(l+1:-1:2);
  lwr_extrap   = 2*select(M) - select(M-1:-1:M-l);
```

```
  extrapol     = [uppr_extrap; select; lwr_extrap];
  interp_filt  = b(D-rem_to(k):D:2*l*D);
  resamp       = filter(interp_filt,1,extrapol);
  bunches(:,k) = resamp((2*l+1:2*l+M)-fix_to(k));
end

bunches(M-1:M,:) = [];
M = M-2;

*****************************************************************************

% curr_mon.m

% Automatically called by synchp.m. Estimates bunch currents
% by looking at the amplitudes of line harmonics (harmonics of
% 60 Hz/50 Hz), or low-frequency phase shifter modulation in
% the data.

% Copyright 2000 by Shyam Prabhakar

len         = 2^(ceil(log(M)/log(2)));
bunff       = fft(bunches,len);
freq        = (0:len-1)/len*fsamp;
u           = find(freq<fCM);
lenu        = length(u);
sbunff      = sum(bunff(u,:)');
sbunff(1:5) = zeros(1,5);
asbunff     = abs(sbunff);

figure(1), clf
plot(freq(u),asbunff,'g'), grid on
xlabel('Freq (Hz)'), ylabel('Arb. units')
axis([min(freq(u)) max(freq(u)) 0 1.2*max(asbunff)+eps])
title([Dir ': Averaged Bunch Signal Spectrum'])
```

```
[srt,v1] = sort(asbunff);
ansCM    = 'y';
while strcmp(ansCM,'y'),
  nCM      = input('Enter no. of freq.s for CM: ');
  if isempty(nCM)|~nCM, nCM = 1; end
  v        = v1(lenu-nCM+1:lenu);
  hold on, plot(freq(v),asbunff(v),'ro')
  drawnow
  ansCM    = input('Change no. of freq.s? (y/n): ','s');
  if strcmp(ansCM,'y'),
    ch = get(gca,'children'); delete(ch(1))
end, end

bun_curr    = real(sbunff(v)*bunff(v,:));
u           = find(bun_curr<0);
bun_curr(u) = zeros(size(u));
bun_curr    = bun_curr*(Io/sum(bun_curr));
```

```
*************************************************************************
```

```
% curr_scal.m - calculates the error, for a given current scaling.

% Copyright 2000 by Shyam Prabhakar

function y = curr_scal(k),
global bun_curr mb Io

y = abs(sum((mb.^2 + (bun_curr/k).^2).^0.5) - Io);
```

```
*************************************************************************
```

```
% med_flt.m - performs median filtering on the input x.

% Copyright 2000 by Shyam Prabhakar
```

```
function y = med_flt(x,win)

[M,N] = size(x);
len   = max(M,N);
halfwin = fix((win-1)/2);
if ~halfwin | len<win,
  y = x;
else
  if min(M,N)~=1,
    error('Check size of input to med_flt!!')
  end
  x        = x(:);
  extend1  = ones(halfwin,1)*x(1);
  extend2  = ones(halfwin,1)*x(len);
  x        = [extend1;x;extend2];
  y        = zeros(M,N);
  for j=1:len,
    y(j) = median(x(j:j+2*halfwin));
  end
end

*************************************************************************

% synchTF.m

% Called by synchp.m. Calculates and plots transfer function
% from bunch currents to synchronous phases (aliased impedance
% at revolution harmonics), at all frequencies that seem to
% have a good signal-to-noise ratio. The "clean" revolution
% harmonics are determined on the line "vTF = find(....);".

% Copyright 2000 by Shyam Prabhakar
```

```
x    = [bun_curr bun_curr(N)*ones(1,harm_no-N)];
x    = x-mean(x);
y    = [synch_p synch_p(N)*ones(1,harm_no-N)];
y    = y-mean(y);
N1   = harm_no;

fftx    = fft(x);
ffty    = fft(y);
N2      = floor(harm_no/2) + 1;
vTF     = find(abs(fftx(1:N2))>N1/50 | abs(ffty(1:N2))>N1/5);
TF      = ffty./fftx;
[xxx,u] = max(abs(fftx));
if real(TF(u)<0), TF = -TF; end
TFab    = abs(TF);

figure(3), clf
if ~isempty(vTF),
  plot(vTF-1,TFab(vTF),'g'), grid on, hold on
  plot(vTF-1,TFab(vTF),'m*'), hold off
  title([Dir ': Transfer Function (synch_p/bun_curr)'])
  xlabel('Revolution harmonic'), ylabel('deg@RF/mA')
  axis([0 5*ceil((max(vTF)-1)/5) 0 5*ceil(max(TFab(vTF))/5)])
  MOhm = input('Display TF in MOhm instead of deg/mA (y/n)? ','s');
  if strcmp(MOhm,'y'),
    Vc  = input('Enter total cavity voltage in MV: ');
    Vc  = 1e6*Vc;
    Vcc = sqrt(Vc^2-Uo^2);
    Zn  = TFab*pi/180*1e3*Vcc/harm_no/1e6;
    plot(vTF-1,Zn(vTF),'g'), grid on, hold on
    plot(vTF-1,Zn(vTF),'m*'), hold off
    title([Dir ': Aliased longitudinal impedance'])
    xlabel('Revolution harmonic (n)'), ylabel('|Z_n|   (MOhm)')
    axis([0 5*ceil((max(vTF)-1)/5) 0 1.3*max(Zn(vTF))])
```

```
end, end

************************************************************************

%    Modes.m

%    This program loads longitudinal oscillation data, filters it to
%    remove some of the noise, removes the offsets in the sampling times
%    of the bunches, plots the oscillation envelopes of the bunches,
%    plots the strengths of modes as a function of time, and fits
%    growing/ damping exponentials to the largest modal transients (so
%    as to estimate growth/damping rates).
%    Last revision: Mar 25 1997.
%    Revision on July 22 1998: Hilbert transform of bunch transients added
%    to remove folding of modal spectrum about f_bunch/2. FltOfstHil
%    created to perform hilbert transform, filter data, remove samp. time
%    offsets and compensate for aliasing of transverse data due to
%    downsampling.

% Copyright 2000 by Shyam Prabhakar
%   ------------------------------------------------------------------

clear
global mode_f N fsamp mmm New Npts osc_unit

Ver
loadgd
findFs, params, FltOfstHil

if exist('q_timdom')~=1,
  q_timdom = input('Plot bunch envelopes in time domain (y/n)? ','s');
end
if strcmp(q_timdom,'y'), bunenv, end
```

```
if N < harm_no, bunches(1,harm_no) = 0; N = harm_no; end

fprintf('\nCalculating modal strengths ...\n')
if N>1,
  mode_ff = 1/nbun*fft(bunches.').';
else
  mode_ff = 1/nbun*bunches;
end
Npts      = N;
mode_f    = abs(mode_ff(:,1:Npts));

fprintf('\nPlotting results ...\n')
waterplot1(M,ceil(M/n_pts_2))
drawnow

modes_to_fit = 0:Npts-1;
if exist('qfitg')~=1,
  qfitg = input('Fit exponentials to modes before breakpoint (y/n)? ','s');
end
if strcmp(qfitg,'y'), modefitg, end
if exist('qfitd')~=1,
  qfitd = input('Fit exponentials to modes after breakpoint (y/n)? ','s');
end
if strcmp(qfitd,'y'), modefitd, end

fprintf('\n')
ParmDisp

*************************************************************************

% Ver.m

% Executes MATLAB commands specific to the version of MATLAB that
% is being used (for compatibility with old and new versions).
```

```
% WARNING: MATLAB 5.2 has better hidden line removal in most
% cases, but it sometimes takes much longer to execute than
% MATLAB 4.2. The fully fulled PEP-II ring is an extreme example
% of execution time. Smaller rings may not have this problem.
% Ver.m is automatically called by Modes.m, etc. where reqd.

% Copyright 2000 by Shyam Prabhakar

VER        = version;
New        = str2num(VER(1)) > 4;
ch         = get(0,'children');
if New&isempty(ch),
  eval('colordef none'),
  figure(1), clf
else
  figure(1), clf
end
set(1, 'DefaultTextFontsize', 11, 'DefaultAxesFontsize', 11,  ...
       'name', '   MODES', 'DefaultPatchErasemode', 'background')

*************************************************************************

% findFs.m - finds the synchrotron/betatron frequency.

% Copyright 2000 by Shyam Prabhakar

len = 2^(ceil(log(M)/log(2)));
bunches = bunches - ones(M,1)*mean(bunches);
bunches = fft(bunches,len);
sum_bunches = abs(bunches)*ones(N,1);
freq = (0:len-1)/len*fsamp;
f_idx = find(freq<fCM | freq>fsamp/2);
sum_bunches(f_idx) = zeros(size(f_idx));
[sbffmax,maxidx]   = max(sum_bunches);
```

```
fr_osc = round(freq(maxidx));
fprintf(['\nThe oscillation frequency is ' num2str(fr_osc) ' Hz.\n'])
if exist('fs_corr')~=1,
  fs_corr = input('Is the estimate of oscill. freq. correct? (y/n): ','s');
  if strcmp(fs_corr,'n'),
    fr_osc = input('Enter your estimate of oscill. freq. in Hz: ');
  end
else
  if strcmp(fs_corr,'n'),
    fr_osc = fr_osc1;
    fprintf('\nThe corrected oscillation frequency is ')
    fprintf([num2str(fr_osc) ' Hz.\n'])
  end
end

*******************************************************************************

% params.m

% Sets band pass filter width, according to parameters created
% by Recogn.m. Uses front end calibration and beam/bunch current
% information to change units of oscillations from counts to
% deg@RF or mm. The variable nbun should be set to the proper
% value in case of a bunch train fill. The nominal value for
% nbun is N, the number of columns in the "bunches" data
% matrix. If a bunch train fill is analysed without entering
% the correct value of nbun, the change of units will introduce
% a scale factor of N/nbun in the calculated oscillation
% amplitudes. Automatically called by Modes.m, etc.

% Copyright 2000 by Shyam Prabhakar

% Set bandpass filter cutoffs
if strcmp(q_trans,'y'),
```

```
  flt_low_freq  = fr_osc - delta_f/2;
  flt_high_freq = fr_osc + delta_f/2;
else
  flt_low_freq  = fr_osc*(2-freq_ratio);
  flt_high_freq = fr_osc*freq_ratio;
end

nbun = N;   % Change this if fewer than N buckets are filled ++++++<<<<<<

% Change units of 'bunches' from counts to osc_units
if exist('BunCurr.mat'),
  load BunCurr
  Ncurr = size(bun_curr,2);
  if Ncurr == N,
    fprintf('\nUsing Current Monitor Result to Scale DSP Signals !!\n')
    bc        = bun_curr;
    ubc       = find(bun_curr<.12);
    bc(ubc)   = .12*ones(size(ubc));
    count2deg = 1./(bc*calib);
    M1        = size(bunches,1);
    bunches   = (ones(M1,1)*count2deg).*bunches;
  else
    fprintf('\nCurrent monitor result has wrong size (ignore)!!\n')
    count2deg = 1/(Io*calib/nbun);
    bunches   = count2deg*bunches;
  end
else
  count2deg = 1/(Io*calib/nbun);
  bunches   = count2deg*bunches;
end

*******************************************************************************

% FltOfstHil.m
```

```
% This program filters bunch signals, and removes sampling time
% offsets by adding the appropriate phase shift in the frequency
% domain. The bunch data are converted from real signals to
% analytic signals by means of a hilbert transform.

% Copyright 2000 by Shyam Prabhakar

fprintf(['\nFiltering signals, performing hilb. transf.,' ...
        ' removing samp. time offsets...\n'])

center_freq1  = (flt_low_freq+flt_high_freq)/2;  % Same as fr_osc
center_freq2  = fsamp - center_freq1;
delta_f_by2   = (flt_high_freq-flt_low_freq)/2;
Freq          = (0:len-1)/len*fsamp;
H             = exp(-(Freq-center_freq1).^2/(2*delta_f_by2^2));
flt_mask      = H + exp(-(Freq-center_freq2).^2/(2*delta_f_by2^2));
[mx,u_cent_f] = max(H);
flt_mask      = flt_mask/flt_mask(u_cent_f);
f_len         = fix((len-1)/2);
hil_mask      = [1 2*ones(1,f_len) ones(1,1-rem(len,2)) zeros(1,f_len)];
n_smooth      = ceil(len/30);
n_smby2       = floor(n_smooth/2);
n_smooth      = 2*n_smby2+1;
if n_smooth > 1 & len > 2*n_smooth,
  idx1            = -n_smby2:n_smby2;
  vec1            = 1 + sin(pi/2*idx1/(n_smby2+1));
  idx1            = rem(len+idx1,len)+1;
  if rem(len,2),
    idx2          = -n_smby2:n_smby2-1;
    vec2          = 1 + fliplr(sin(pi/2*(idx2+.5)/n_smby2));
    idx2          = f_len + 2 + idx2;
  else
```

```
    vec2          = fliplr(vec1);
    idx2          = f_len + 2 + (-n_smby2:n_smby2);
  end
  hil_mask(idx1) = vec1;
  hil_mask(idx2) = vec2;
end
if rem(tun_alias,2), hil_mask = 2 - hil_mask; end

flt_hil_mask  = flt_mask.*hil_mask;

nadv            = max(turn_offsets(1:N)) - turn_offsets(1:N);
nadv            = nadv + (N-1:-1:0)'/N;
Tadv            = nadv'*harm_no/Frf;
omeg            = zeros(1,len+1);
idx_max         = ceil(len/2);
omeg(1:idx_max) = 2*pi*(Freq(1:idx_max)  +                    ...
                   (-1)^tun_alias*ceil(tun_alias/2)*fsamp);
omeg            = omeg - fliplr(omeg);
omeg(len+1)     = [];

bunches       = bunches.*(flt_hil_mask.'*ones(1,N)).*exp(i*omeg'*Tadv);
bunches       = ifft(bunches);

bunches([1:15 M-14:len],:) = [];
M             = M-30;
breakpt       = max(breakpt-15,1);
breakpt       = min(breakpt,M);

*************************************************************************

% bunenv.m - calculate and plot envelopes of bunch
% oscillations.

% Copyright 2000 by Shyam Prabhakar
```

```
fprintf('\nCalculating envelopes of bunch oscillations ...\n')

WIN       = fsamp/fr_osc;
WIN       = round(round(WIN_dflt/WIN)*WIN);
if ~WIN, WIN = WIN_dflt; end
WIN       = ceil(WIN/10);

decim     = ceil(M/n_pts_1);
env       = zeros(fix((M-WIN)/decim)+1,nbun);
mean_vec  = ones(1,WIN)/WIN;

for j = 1:decim:M-WIN+1,
  bun_squared          = abs(bunches(j:j+WIN-1,1:nbun)).^2;
  env((j-1)/decim+1,:) = sqrt(mean_vec*bun_squared);
end
t_env = (0:decim:(M-WIN))*1e3/fsamp;

fprintf('\nPlotting results ...\n')

figure(1), subplot(3,2,1), hold off
waterfall(t_env, 1:nbun, env'),grid on
axis([0 max(t_env) 1 nbun+(nbun==1) 0 max(max(env))+eps])
ylabel('Bunch No.'), xlabel('Time  (ms)'), zlabel(osc_unit)
title('a) Osc. Envelopes in Time Domain')
drawnow

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% waterplot1.m - creates a 3-D plot of the evolution of the
% mode amplitudes.

% Copyright 2000 by Shyam Prabhakar
```

```
function waterplot1(max_samp,decim)
global mode_f N fsamp mmm New Npts osc_unit

mode_f1     = mode_f(decim:decim:max_samp,1:Npts);
[m,n]       = size(mode_f1);
X           = 0:n-1;
Y           = decim*1e3/fsamp*(1:m);

mm          = max(mode_f1);
mmm         = max(mm);
mm1         = mm;
dn          = ceil(n/30);
mm1(1:dn:n) = 3*mmm*ones(size(mm1(1:dn:n)));
medmm2      = 2*median(mm);
v           = find(mm1>medmm2);

figure(1), subplot(3,2,2), reset(gca)
waterfall(Y,X(v),mode_f1(:,v).')
axis([0 max(Y) 0 n-1+(n==1) 0 mmm])
grid on, ylabel('Mode No.'), xlabel('Time (ms)'), zlabel(osc_unit)
title('b) Evolution of Modes')
if ~New,
  drawnow
  ch  = get(gca,'children');
  chidx = fliplr(length(v) - find(mm(v)>medmm2) +1);
  fc = get(gca,'color');
  if strcmp(fc,'none'), fc = get(gcf,'color'); end
  set(ch(chidx),'face',fc)
end

*******************************************************************************

% modefitg.m
```

```
% This program fits exponentials to modal transients that grow
% larger than the fit cutoff typed in by the user. Only the section
% of the transients before the break point is used for calculating
% the fit. Fits and growth rates are plotted by calling plotfit.m.

% Copyright 2000 by Shyam Prabhakar

global mode_f N fsamp strength t
fit_type = 'exp'; q_cutoff   = 'y';

fprintf('\nFitting exponentials to modes before breakpoint ...\n')

[M,N] = size(mode_f);
u       = (1:2:min(M,breakpt*(1-(Year>93&Year<96)/9)))';
lenu    = length(u);
t       = (u-min(u))/fsamp*1000;        % !!(in millisec.)!!
option2 = [0 1e-4 lenu*1e-8 zeros(1,10) 1100];
option1 = [0 1e-2 lenu*1e-5 zeros(1,10) 400];
Scale   = 10/mmm;

while strcmp(q_cutoff,'y') & lenu > 2,
  fit_cutof = Scale*input(['Enter fit cutoff in ' osc_unit ': ']);
  grate     = zeros(1,Npts);
  fit       = zeros(lenu,Npts);
  RMS       = zeros(1,Npts);
  Noyz_flr  = zeros(1,Npts);
  modnumg   = [];
  fprintf('\n')
  for j=modes_to_fit+1,
    strength = Scale*mode_f(u,j);
    if max(strength) > fit_cutof,
      modnumg     = [modnumg j];
      check       = fmins('f1',[.3 .01*Scale],option2);
```

```
      x0          = [check .01*Scale];
      sol         = fmins('f2',x0,option2);
      fit(:,j)    = sqrt(sol(3)^2+sol(2)^2*exp(2*sol(1)*t))/Scale;
      grate(j)    = sol(1);  RMS(j) = sqrt(f2(sol)/lenu)/Scale;
      Noyz_flr(j) = abs(sol(3))/Scale;
    end
    if floor(j/cnt)==j/cnt, fprintf(['j=',num2str(j),'  ']), end
  end
  fprintf('\n')
  if ~isempty(modnumg), fitrecord(1) = 1; plotfit, end

  fprintf(['\nPresent fit cutoff is ' num2str(fit_cutof/Scale) ...
osc_unit '\n'])
  q_cutoff = input('Change the fit cutoff (y/n)? ','s');
end

if lenu <= 2, fprintf('\nNot enough data before breakpoint'), end

if exist('Txhdl') & PDpos ~= 3,
  if PDpos == 1, set(Txhdl,'erasemode','normal'), delete(Txhdl), end
  ParmDisp
end

ug       = u;
rateg    = grate(modnumg);
fitg     = fit(:,modnumg);
RMSg     = RMS(modnumg);
Nyz_flrg = Noyz_flr(modnumg);

*************************************************************************

% modefitd.m

% This program fits exponentials to modal transients that grow
```

```
% larger than the fit cutoff typed in by the user. Only the section
% of the transients after the break point is used for calculating
% the fit. Fits and damping rates are plotted by calling plotfit.m.

% Copyright 2000 by Shyam Prabhakar

global mode_f N fsamp strength t
fit_type = 'exp'; q_cutoff   = 'y';

fprintf('\nFitting exponentials to modes after breakpoint ...\n')

[M,N] = size(mode_f);
u       = (min(M,breakpt):2:M)';
lenu    = length(u);
t       = (u-min(u))/fsamp*1000;        % !!(in millisec.)!!
option2 = [0 1e-4 lenu*1e-8 zeros(1,10) 1100];
option1 = [0 1e-2 lenu*1e-5 zeros(1,10) 400];
Scale   = 10/mmm;

while strcmp(q_cutoff,'y') & lenu > 2,
  fit_cutof = Scale*input(['Enter fit cutoff in ' osc_unit ': ']);
  grate     = zeros(1,Npts);
  fit       = zeros(lenu,Npts);
  RMS       = zeros(1,Npts);
  Noyz_flr  = zeros(1,Npts);
  modnumd   = [];
  fprintf('\n')
  for j=modes_to_fit+1,
    strength = Scale*mode_f(u,j);
    if max(strength) > fit_cutof,
      modnumd     = [modnumd j];
      check       = fmins('f1',[-.3 .01*Scale],option2);
      x0          = [check .01*Scale];
```

```
      sol         = fmins('f2',x0,option2);
      fit(:,j)    = sqrt(sol(3)^2+sol(2)^2*exp(2*sol(1)*t))/Scale;
      grate(j)    = sol(1);  RMS(j) = sqrt(f2(sol)/lenu)/Scale;
      Noyz_flr(j) = abs(sol(3))/Scale;
    end
    if floor(j/cnt)==j/cnt, fprintf(['j=',num2str(j),'  ']), end
  end
  fprintf('\n')
  if ~isempty(modnumd), fitrecord(2) = 1; plotfit, end

  fprintf(['\nPresent fit cutoff is ' num2str(fit_cutof/Scale) ...
osc_unit '\n'])
  q_cutoff = input('Change the fit cutoff (y/n)? ','s');
end

if lenu <= 2, fprintf('\nNot enough data after breakpoint'), end

if exist('Txhdl') & PDpos ~= 1,
  if PDpos == 2, set(Txhdl,'erasemode','normal'), delete(Txhdl), end
  ParmDisp
end

ud       = u;
rated    = grate(modnumd);
fitd     = fit(:,modnumd);
RMSd     = RMS(modnumd);
Nyz_flrd = Noyz_flr(modnumd);
```

****

```
% f1.m

% Calculates the error between the modal transient and the
% exponential fit.
```

```
% Copyright 2000 by Shyam Prabhakar

function y = f1(x);
global strength t

f = x(2)*exp(x(1)*t) - strength;
y = f'*f;

********************************************************************************

% f2.m

% Calculates the error between the modal transient and the
% (exponential + noise floor) fit.

% Copyright 2000 by Shyam Prabhakar

function y = f2(x);
global strength t

f = sqrt(x(3)^2+x(2)^2*exp(2*x(1)*t)) - strength;
y = f'*f;

********************************************************************************

% plotfit.m

% Calls waterplot2.m or waterplot3.m to plot the exponential
% fits to the modal transients, and plots the calculated
% growth rates or damping rates.

% Copyright 2000 by Shyam Prabhakar

fprintf('\nPlotting fit ...\n')
```

```
figure(1)

X = 0:Npts-1;
if max(u) <= breakpt,
   waterplot2
   subplot(3,2,4),hold off
else
   waterplot3
   subplot(3,2,6),hold off
end

lv    = length(v);
gg    = [zeros(1,lv) grate(v) NaN*zeros(1,lv)];
w     = [v v v];
[w,j] = sort(w);
plot(X(w),gg(j),'r'), grid on, hold on
plot(X(v),grate(v),'go'), hold off
set(gca,'fontsize',11)
xlabel('Mode No.')
ylabel('Rate   (1/ms)')
axis([0 length(grate) 1.2*min(grate) 1.2*max(grate)+eps])
if max(u) <= breakpt,
   title('     d) Growth Rates (pre-brkpt)')
else
   title('     f) Growth Rates (post-brkpt)')
end

*******************************************************************************

% waterplot2.m

% Creates 3-D plot of exponential fits to modes before the
% break point.
```

```
% Copyright 2000 by Shyam Prabhakar

decim        = ceil(lenu/30);
Y            = 2e3/fsamp*(1:decim:lenu);          % time in ms
fit1         = fit(1:decim:lenu,:);
mf           = max(fit1);
mmf          = max(mf);
v            = [1 find(mf)];
X1           = X(1:ceil(Npts/40):Npts);
Y1           = 2e3/fsamp*[1 lenu/2 lenu];
Z1           = eps*([1:length(X1)]'*[1 1 1] + ones(length(X1),1)*[1:3]);

figure(1), subplot(3,2,3), reset(gca)
waterfall(Y1,X1,Z1), hold on
waterfall(Y,X(v),fit1(:,v).'), hold off, v(1) = [];
axis([0 max(Y) 0 max(X) 0 mmf])
grid on, ylabel('Mode No.'), xlabel('Time (ms)'), zlabel(osc_unit)
title('c) Exp. Fit to Modes (pre-brkpt)')
if ~New,
  drawnow
  ch  = get(gca,'children');
  fc = get(gca,'color');
  if strcmp(fc,'none'), fc = get(gcf,'color'); end
  set(ch(1:length(v)),'face',fc)
end

*************************************************************************

% waterplot3.m

% Creates 3-D plot of exponential fits to modes after the
% break point.

% Copyright 2000 by Shyam Prabhakar
```

```
decim          = ceil(lenu/30);
Y              = 1e3/fsamp*(breakpt+4) + 2e3/fsamp*(1:decim:lenu);
fit1           = fit(1:decim:lenu,:);
mf             = max(fit1);
mmf            = max(mf);
v              = [1 find(mf)];
X1             = X(1:ceil(Npts/40):Npts);
Y1             = 1e3/fsamp*(breakpt+4) + 2e3/fsamp*[1 lenu/2 lenu];
Z1             = eps*([1:length(X1)]'*[1 1 1] + ones(length(X1),1)*[1:3]);

figure(1), subplot(3,2,5), reset(gca)
waterfall(Y1,X1,Z1), hold on
waterfall(Y,X(v),fit1(:,v).'), hold off, v(1) = [];
axis([min(Y) max(Y) 0 max(X) 0 mmf])
grid on, ylabel('Mode No.'), xlabel('Time (ms)'), zlabel(osc_unit)
title('e) Exp. Fit to Modes (post-brkpt)')
if ~New,
  drawnow
  ch  = get(gca,'children');
  fc = get(gca,'color');
  if strcmp(fc,'none'), fc = get(gcf,'color'); end
  set(ch(1:length(v)),'face',fc)
end

*************************************************************************

% ParmDisp.m - displays relevant parameters at figure bottom

% Copyright 2000 by Shyam Prabhakar

str1     = [machine '/' Dir ':  Io= ' num2str(Io) 'mA'  ',  Dsamp= ' ...
            num2str(downsamp) ',  ShifGain= ' num2str(shift_gain)    ...
          ', Nbun= '  num2str(nbun), ','];
```

```
str2     = ['Gain1= '  num2str(gains(2))     ...
           ',  Gain2= '  num2str(gains(1))     ...
           ',  Phase1= '  num2str(phases(2))   ...
           ',  Phase2= '  num2str(phases(1))   ...
           ',  Brkpt= '  num2str(breakpt)      ...
           ',  Calib= '  num2str(calib) '.'];

figure(1)
pos_fig1 = get(1,'position');
width    = pos_fig1(3);
height   = pos_fig1(4);

if     fitrecord(2), subplot(326), PDpos = 3;
elseif fitrecord(1), subplot(324), PDpos = 2;
else   subplot(322), PDpos = 1;
end

ax_posn  = get(gca,'position');
ax_X     = width*ax_posn(1);
ax_Y     = height*ax_posn(2);

TxtX     = .5*width - ax_X;
TxtY1    = -40;
TxtY2    = TxtY1 - 14;
Txhdl    = 1e3*ones(1,2);
Txhdl(1) = text(TxtX, TxtY1, str1, 'color', 'g', 'fontweight', ...
'demi', 'units', 'pixels', 'horizontalalignment', ...
'center', 'erasemode', 'xor');
Txhdl(2) = text(TxtX, TxtY2, str2, 'color', 'g', 'fontweight', ...
'demi', 'units', 'pixels', 'horizontalalignment', ...
'center', 'erasemode', 'xor');
set(Txhdl,'units','normalized')
```

```
orient tall

*******************************************************************************

% checkfit.m

%  This program should be executed after Modes.m, if the accuracy
%  of  the  exponential  fit to  any  particular  section  of the
% modal transients needs to be checked.

% Copyright 2000 by Shyam Prabhakar

lg = 0; ld = 0;
if exist('modnumg'), lg = length(modnumg); end
if exist('modnumd'), ld = length(modnumd); end

dash = ['-------------']; dash = [dash dash dash];
if ~(lg+ld),
  error('No fits have been calculated!')
elseif ~lg,
  modnumgd = modnumd; ugd = ud; rategd = rated; lgd = ld;
  fitgd = fitd; RMSgd = RMSd; Nyz_flrgd = Nyz_flrd;
  strdisp  = ['Mode#   ' 'D.rate   ' 'RMS error  ' 'Noise floor'];
  ckformat = '%5.3g %8.4g   %9.3g   %9.3g\n';
elseif ~ld,
  modnumgd = modnumg; ugd = ug; rategd = rateg; lgd = lg;
  fitgd = fitg; RMSgd = RMSg; Nyz_flrgd = Nyz_flrg;
  strdisp  = ['Mode#   ' 'G.rate   ' 'RMS error  ' 'Noise floor'];
  ckformat = '%5.3g %8.4g   %9.3g   %9.3g\n';
else
  ugd = [ug; ud];
  tmp = zeros(1,N); tmp(modnumg) = tmp(modnumg) + 1;
  tmp(modnumd) = tmp(modnumd) + 2;
  modnumgd = find(tmp); lgd = length(modnumgd);
```

```
  rate1 = rateg; Nyz_flr1 = Nyz_flrg; RMS1 = RMSg; fit1 = fitg;
  rate2 = rated; Nyz_flr2 = Nyz_flrd; RMS2 = RMSd; fit2 = fitd;
  if lgd-lg,
    [temp,indx] = sort([modnumg find(tmp==2)]);
    rate1(lgd) = 0; Nyz_flr1(lgd) = 0; RMS1(lgd) = 0; fit1(1,lgd) = 0;
    rate1 = rate1(indx); Nyz_flr1 = Nyz_flr1(indx); RMS1 = RMS1(indx);
    fit1 = fit1(:,indx);
  end
  if lgd-ld,
    [temp,indx] = sort([modnumd find(tmp==1)]);
    rate2(lgd) = 0; Nyz_flr2(lgd) = 0; RMS2(lgd) = 0; fit2(1,lgd) = 0;
    rate2 = rate2(indx); Nyz_flr2 = Nyz_flr2(indx); RMS2 = RMS2(indx);
    fit2 = fit2(:,indx);
  end
  rategd = [rate1; rate2]; Nyz_flrgd = [Nyz_flr1; Nyz_flr2];
  RMSgd = [RMS1; RMS2]; fitgd = [fit1; fit2];
  strdisp  = ['Mode#   '  'G.rate   '  'D.rate          ' ...
              'RMS error            '  'Noise floor'];
  ckformat = '%5.3g %8.4g %8.4g   %9.3g   %9.3g   %9.3g   %9.3g\n';
  dash = [dash dash(1:length(dash)-6)];
end

fprintf(['\n' dash '\n' strdisp '\n' dash '\n'])
disp(sprintf(ckformat,[modnumgd-1; rategd; RMSgd; Nyz_flrgd]))
disp(dash)

PlotCheck

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

% PlotCheck.m - plots the superimposed modal transients and
% exponential fits. Called by checkfit.m.

% Copyright 2000 by Shyam Prabhakar
```

```
Dir = getDir;

figure(2), clf
set(2, 'name', '   CHECK FIT', 'DefaultAxesFontsize', 11)
tbreak = min(M,breakpt)*1e3/fsamp; tbreak = [tbreak tbreak];
for j = 1:lgd,
  subplot(lgd,1,j)
  plot(1e3/fsamp*ugd, mode_f(ugd, modnumgd(j)), 'g'), hold on
  plot(1e3/fsamp*ugd, fitgd(:,j), 'm'), grid on
  ylim = axis; ylim = [ylim(3)*(1+eps) ylim(4)*(1-eps)];
  plot(tbreak, ylim, 'w:')
  if Year > 93 & Year < 96, plot(tbreak*8/9,ylim,'w:'), end
  ylabel(osc_unit)
  if j == 1,
    title([Dir '   Data + Fit for Mode #' num2str(modnumgd(1)-1)])
  else
    title(['Mode #'  num2str(modnumgd(j)-1)])
  end
end
xlabel('Time  (ms)')
orient tall

*************************************************************************

% sideband.m

% This program interpolates between the samples and interleaves
% bunch data to produce one long vector which contains the
% sequence of bunch oscillation coordinates seen by the BPM. The
% FFT of this vector is the beam pseudospectrum, which contains
% oscillation sidebands, but no revolution harmonics (ideally). The
% pseudospectrum is plotted, together with the difference of
% upper and lower sideband amplitudes at each revolution
```

```
% harmonic. The zoomed spectrum around the 8 revolution harmonics
% with the largest sideband asymmetry is plotted on a separate
% figure. No filtering is done.

% Copyright 2000 by Shyam Prabhakar

clear
global bunHR bunmax Frf harm_no len Dir NyzFlr osc_unit

VerHR
loadgdHR
paramsHR
bunches  = bunches';

resampl
[srtpk,mnum] = sortpeaks(fr_osc, .5*fr_osc);

figure(2), clf
set(2, 'name', '  ZOOM', 'DefaultTextFontsize', 11, ...
'DefaultAxesFontsize', 11)
for j = 1:8,
  subplot(4,2,j), Zoom(mnum(j),1.7*fr_osc)
  if j < 7, xlabel(''), end
end
orient tall

fprintf('\n')

*************************************************************************

% VerHR.m - equivalent to Ver.m. Called by sideband.m

% Copyright 2000 by Shyam Prabhakar
```

```
VER        = version;
New        = str2num(VER(1)) > 4;
ch         = get(0,'children');
if New&isempty(ch),
  eval('colordef none'), figure(1)
else
  figure(1), clf
end
set(1, 'name', '  SIDEBANDS', 'DefaultTextFontsize', 11, ...
'DefaultAxesFontsize', 11)

*****************************************************************************

% loadgdHR.m - equivalent to loadgd.m. Called by sideband.m.

% Copyright 2000 by Shyam Prabhakar

fprintf(['\n---- ', pwd, ' ----\n'])
fprintf('\nLoading and sorting by bunch no. ...\n')

load gd
[Dir,Year] = getDir;
Recogn

[M,N]     = size(bunches);
[R,C]     = find(bunches == -1000 | bunches == 12345678);
if ~isempty(R)
  R = min(R);  bunches(R:M,:) = []; M = R-1;
end
lto       = length(turn_offsets);
if N~=lto,
  if N<lto,
    fprintf('\nWarning: turn_offsets has more columns than bunches\n')
  else
```

```
    fprintf('\nCAUTION!! bunches has more columns than turn_offsets\n')
    bunches(:,lto+1:N)=[];   N = lto;
end, end

Reduce

[Y,I]     = sort(bunches(1,:));
bunches   = bunches(2:M,I);
M         = M-1;
bunches   = bunches-ones(M,1)*mean(bunches);
Io        = beamCurrent;
while Io <= 0, Io = input('Enter estimated Io (in mA): '); end

breakpt   = round(samples_per_bunch*damp_brkpt/63);
if gains(2) == gains(1) & phases(2) == phases(1),
  breakpt = M;
end
harm_no     = ring_size/red_fact;
Frf         = rf_freq/red_fact*1e6;
frev        = Frf/harm_no;
cnt         = ceil(harm_no/15);
fsamp       = frev/downsamp;
clear beamCurrent damp_brkpt ring_size rf_freq lto Y I R C

*******************************************************************************

% paramsHR.m

% Detects the oscillation frequency, and uses the front end
% calibration and beam/bunch current information to change units
% of oscillations from counts to deg@RF or mm. Automatically called
% by sideband.m.

% Copyright 2000 by Shyam Prabhakar
```

```
% Find the synchrotron frequency
[xx,yy]       = max(abs(bunches));
[xx,yy]       = sort(xx);
yy            = yy(max(1,N-15):N);
len           = 2^(ceil(log(M)/log(2)));
yuk           = abs(fft(bunches(:,yy),len));
yuk           = mean(yuk');
freq          = (0:len-1)/len*fsamp;
yukidx        = find(freq<1e3 | freq>fsamp/2);
yuk(yukidx)   = zeros(size(yukidx));
[duk,muk]     = max(yuk);
fr_osc        = round(freq(muk));
fprintf(['\nThe oscillation frequency is ' num2str(fr_osc) ' Hz.\n'])
fs_err = input('Is the estimate of oscill. freq. correct? (y/n): ','s');
if strcmp(fs_err,'n'),
  fr_osc = input('Enter your estimate of oscill. freq. in Hz: ');
end

% Change units of 'bunches' from counts to osc_units
if exist('BunCurr.mat'),
  load BunCurr
  Ncurr = size(bun_curr,2);
  if Ncurr == N,
    fprintf('\nUsing Current Monitor Result to Scale DSP Signals !!\n')
    bc        = bun_curr;
    ubc       = find(bun_curr<.12);
    bc(ubc)   = .12*ones(size(ubc));
    count2deg = 1./(bc*calib);
    bunches   = (ones(M,1)*count2deg).*bunches;
  else
    fprintf('\nCurrent monitor result has wrong size (ignore)!!\n')
    count2deg = 1/(Io*calib/N);
```

```
    bunches    = count2deg*bunches;
  end
else
  count2deg = 1/(Io*calib/N);
  bunches   = count2deg*bunches;
end

*************************************************************************

% resampl.m

% Upsamples the bunch signals, to get a sample per turn. The
% expanded "bunches" matrix is then strung out into a single
% vector, which contains the sequence of bunch phases seen by the
% BPM. The sequence is then FFT-ed, to get the beam pseudospectrum.
% The pseudospectrum is plotted as a function of frequency.

% Copyright 2000 by Shyam Prabhakar

fprintf('\nInterpolating and resampling data ...  ')
fprintf(['(breakpt = ' num2str(breakpt) ')\n'])

n_pts   = 0;
while n_pts < 9,          % Need at least 9 to interpolate
  minn    = input('Enter starting sample #: ');
  minn    = max(1,minn);
  maxx    = input('Enter final sample #: ');
  maxx    = min(M,maxx);
  n_pts   = maxx - minn + 1;
end
len     = (n_pts-5)*harm_no*downsamp;
bunches = (2/len)*bunches;
bunHR   = zeros(harm_no,(n_pts-5)*downsamp);
NyzFlr  = 2*2*sqrt(downsamp/len)*mean(count2deg);  % Assume 2 cnt rms noise
```

```
fprintf('\n')

for k = 1:N,
  intr       = interp(bunches(k,minn:maxx), downsamp);
  bunHR(k,:) = intr(3*downsamp-turn_offsets(k): ...
               (n_pts-2)*downsamp-1-turn_offsets(k));
  if floor(k/cnt)==k/cnt, fprintf(['k=',num2str(k),'  ']), end
end
fprintf('\n')

bunHR = bunHR(:);

fprintf('\nPerforming FFT ...\n')
bunHR  = fft(bunHR);
bunHR  = abs(bunHR);
%[u,bunmax,bunHR] = db(bunHR,80);
bunmax = max(bunHR);
u      = find(bunHR > NyzFlr/5);
freq   = (Frf/len/1e6)*(u-1);

subplot(2,1,1), hold off
semilogy(freq,bunHR(u),'g')
axis([0 Frf/2/1e6 NyzFlr bunmax*2])
xlabel('Frequency   (MHz)'), ylabel(osc_unit)
title([Dir ':     Beam pseudospectrum    ' '(Sample# ' ...
       num2str(minn) ':' num2str(maxx) ')'])
```

****

```
% sortpeaks.m

% Plots the difference between upper and lower sidebands,
% as a function of frequency (revolution harmonic).
```

```
% Returns the sorted differences, and the corresponding
% vector of mode numbers.

% Copyright 2000 by Shyam Prabhakar

function [srtpk,mnum] = sortpeaks(fs, deltaf)
global bunHR Frf harm_no len osc_unit

uofst   = round(fs/Frf*len);
uspan   = round(deltaf/2/Frf*len);
peak    = zeros(1,harm_no);

uusb    = (uofst+1-uspan):(uofst+1+uspan);
peak(1) = max(bunHR(uusb));
hby2    = floor(harm_no/2);
for j = 1:hby2,
  umidu          = round(j/harm_no*len) + 1 + uofst;
  uusb           = umidu - uspan:umidu + uspan;
  umidl          = round(j/harm_no*len) + 1 - uofst;
  ulsb           = umidl - uspan:umidl + uspan;
  peak(j+1)      = max(bunHR(uusb)) - max(bunHR(ulsb));
  j_conj         = harm_no - j;
  peak(j_conj+1) = -peak(j+1);
end

j = harm_no/2;
if fix(j) == j,
  umidu          = round(j/harm_no*len) + 1 + uofst;
  uusb           = umidu - uspan:umidu + uspan;
  peak(j+1)      = max(bunHR(uusb));
end

subplot(2,1,2), hold off
```

```
semilogy(0:hby2, abs(peak(1:hby2+1)), 'g')
axis([0 hby2 min(abs(peak)) max(peak)*2])
title('Peak Values of Synchrotron Sidebands (upper - lower)')
ylabel(osc_unit), xlabel('Mode')

[srtpk,mnum] = sort(peak);
srtpk        = fliplr(srtpk);
mnum         = fliplr(mnum) - 1;

********************************************************************************

% Zoom.m

% Plots the portion of the pseudospectrum from mnum*frev - flim to
% mnum*frev + flim

% Copyright 2000 by Shyam Prabhakar

function Zoom(mnum, flim, col, Ymax)
global bunHR bunmax Frf harm_no len Dir NyzFlr osc_unit

if nargin < 4, Ymax = 2*bunmax;
  if nargin < 3, col = 'g';
end, end
frev  = Frf/harm_no;
ulow  = round((mnum*frev-flim)/Frf*len) + 1;
ulow  = max(ulow,1);
uhigh = round((mnum*frev+flim)/Frf*len) + 1;
u     = ulow:uhigh;
freq  = (u-1)/len*Frf/1e3 - mnum*frev/1e3;
flim  = flim/1e3;
semilogy(freq,bunHR(u),col), grid on
axis([min(freq) max(freq) NyzFlr/5 Ymax])
xlabel('Frequency   (kHz)'), ylabel(osc_unit)
```

```
title([Dir ': ' num2str(mnum) '*frev'])

************************************************************************

% PhasSpac.m

%  This program is to be run after Modes.m. Displays phase space
% angles of selected modes relative to a fictitious upper (u) or
% lower (l) sideband at the nominal synchrotron/betatron frequency
% (fr_osc). Phases can also be plotted relative to the phase of an
% actual mode.

% Copyright 2000 by Shyam Prabhakar

sect = input('Which section of the transient do you want (1/2)?: ');
if sect==2,
  u  = breakpt+1:M;
else
  u  = 1:breakpt;
end
lu   = size(u,2);
bun_ff    =  mode_ff(u,:);

mlist     = input('Enter vector of mode#s : ');
mlen      = length(mlist);
Xrec      = real(bun_ff(:,mlist+1));
Yrec      = imag(bun_ff(:,mlist+1));
mphas     = atan2(Yrec,Xrec);
tphas     = (u-1)'*1e3/fsamp;

clear u l;
mref      = input('Enter reference mode# : ','s');
if strcmp(mref,'u'),
  phasref = angle(exp(i*2*pi*fr_osc/1e3*tphas));
```

```
elseif strcmp(mref,'1'),
  phasref = angle(exp(-i*2*pi*fr_osc/1e3*tphas));
else
  mref    = str2num(mref);
  Xref    = real(bun_ff(:,mref+1));
  Yref    = imag(bun_ff(:,mref+1));
  phasref = atan2(Yref,Xref);
end
mphas     = unwrap(mphas-phasref*ones(1,mlen),pi);
%mphas     = unwrap(mphas-phasref*ones(1,mlen),3*pi/2);
mmp       = mean(mphas);
mmp1      = angle(exp(i*mmp));
mphas     = 180/pi*(mphas+ones(lu,1)*(mmp1-mmp));

hold off
plot(tphas,mphas),grid on
xlabel('Time (ms)')
ylabel('Relative phase (deg)')
title([Dir  ':   Relative phases of growing modes'])

*******************************************************************************
```

Bibliography

[1] K.W. Robinson, "Stability of Beam in Radiofrequency System," Cambridge Electron Accel. Report No. CEAL-1010 (1964).

[2] E. Courant and A. Sessler, "Transverse Coherent Resistive Instabilities of Azimuthally Bunched Beams in Particle Accelerators," Rev. Sci. Instrum. **37**, 1579 (1966).

[3] A. Chao, *Physics of Collective Instabilities in High Energy Accelerators* (Wiley, 1993).

[4] J.-M. Wang, "Modes of Storage Ring Coherent Instabilities," in *Physics of Particle Accelerators*, proceedings of the SLAC summer school, 1985, and the Fermilab summer school, 1984 (AIP, 1987), p. 697.

[5] J. Laclare, "Bunched Beam Instabilities," in *11th International Conference on High Energy Accelerators*, Geneva, 1980, edited by W.S. Newman (Birkhauser Verlag, 1980), p. 526.

[6] T. Raubenheimer and F. Zimmermann, "A Fast Beam-Ion Instability in Linear Accelerators and Storage Rings," Phys. Rev. E **52**, 5487 (1995).

[7] K. Ohmi, "A Coupled Bunch Instability due to Beam-Photoelectron Interactions in Positron Storage Rings," in *Fifth European Particle Accelerator Conference, Sitges, 1996*, proceedings, edited by S. Myers *et al.* (Institute of Physics Publishing, 1996), p. 1069.

[8] J. Rogers, "Photoelectron Instabilities in Electron Positron Factories," in *Beam Dynamics Issues for e^+e^- Factories*, proceedings of the 14th Advanced ICFA Workshop, Frascati, Italy, 1997, edited by L. Palumbo and G. Vignola (Frascati Physics Series Vol. X, 1998), p. 413.

[9] LBNL Report No. LBL-PUB-643-Rev.2, 1989.

[10] "PEP-II: An Asymmetric B Factory. Conceptual Design Report," SLAC Report No. SLAC-418, 1993.

[11] M.A. Allen, Rudolf A. Ecken, Leonard Genova, B. Humphrey, L.G. Karvonen, Gerhard T. Konrad, J.V. Lebacqz, R.A. McConnell, and C.W. Olson, "Design and Operation of the SPEAR-II RF System," IEEE Trans. Nucl. Sci. **22**, 1269 (1975).

[12] J. Safranek and H. Wiedemann, "Low-Emittance in SPEAR," in *IEEE 1991 Particle Accelerator Conference*, proceedings, San Francisco (IEEE 1991), p. 1104.

[13] S. Prabhakar, "Curing Coupled-Bunch Instabilities with Uneven Fills: Theory" (to be published).

[14] S. Prabhakar, J. Fox, and D. Teytelman, "Curing Coupled-Bunch Instabilities with Uneven Fills: Experiment" (to be published).

[15] R.D. Kohaupt, "On Multi-Bunch Instabilities for Fractionally Filled Rings," DESY Report No. DESY 85-139, 1985.

[16] K. Thompson and R.D. Ruth, "Transverse and Longitudinal Coupled Bunch Instabilities in Trains of Closely Spaced Bunches," in *1989 IEEE Particle Accelerator Conference: Accelerator Science and Technology* (IEEE, 1989), p. 792.

[17] S.A. Bogacz, "Potential Well Distortion Effects for a Partially Filled Ring," Part. Accel. **48**, 19 (1994).

[18] M. Sands, "The Physics of Electron Storage Rings: an Introduction," SLAC-121 (1970).

[19] P. Wilson, "Introduction to Wakefields and Wake Potentials," in *Physics of Particle Accelerators*, proceedings of the Fermilab summer school, 1987 and the Cornell University summer school, 1988, edited by M. Month and M. Dienes (AIP, 1989), p. 524; SLAC-PUB-4547.

[20] K. Thompson and R. Ruth, "Transverse Coupled-Bunch Instabilities in Damping Rings of High-Energy Linear Colliders," Phys. Rev. D **43**, 3049 (1991)

[21] F. Pedersen, "RF Cavity Feedback," SLAC-400, 192 (1992).

[22] F. Sacherer, "A Longitudinal Stability Criterion for Bunched Beams," IEEE Trans. Nucl. Sci. **20**, 825 (1973).

[23] M. Zobov, R. Boni, A. Gallo, A. Ghigo, F. Marcellini, L. Palumbo, M. Serio, B. Spataro, and G. Vignola, "Measures to Reduce the Impedance of Parasitic Resonant Modes in the DAPHNE Vacuum Chamber," in *Beam Dynamics Issues for e^+e^- Factories*, proceedings of the 14th Advanced ICFA Workshop, Frascati, Italy, 1997, edited by L. Palumbo and G. Vignola (Frascati Physics Series Vol. X, 1998), p. 371.

[24] W. Chou and J. Griffin, "Impedance Scaling and Impedance Control," in *Proceedings of the 1997 Particle Accelerator Conference*, Vancouver (IEEE, 1998), p. 1724.

[25] L.D. Landau, "On the Vibration of the Electronic Plasma," J. Phys. USSR **10**, 25 (1946).

[26] see H.G. Hereward, "Landau Damping," in *CERN Accelerator School: Advanced Accelerator Physics*, proceedings, Oxford, 1985, CERN 87-03, p. 255 (1987), and references therein.

[27] F. Pedersen and F. Sacherer, "Theory and Performance of the Longitudinal Active Damping System for the CERN PS Booster," IEEE Trans. Nucl. Sci. **24**, 1396 (1977).

[28] R.D. Kohaupt, "Theory of Multibunch Feedback Systems," DESY-91-071, 1991.

[29] D. Boussard, "RF and Feedback Systems," in *Tau-Charm Factory*, proceedings of the 3rd workshop, Marbella (Editions Frontieres, 1994), p. 579.

[30] K. Balewski, "Review of Feedback Systems," in *Sixth European Particle Accelerator Conference, Stockholm, 1998*, proceedings, edited by S. Myers *et al.* (Institute of Physics Publishing, 1998), p. 169.

[31] A. Mosnier, "Cures of Coupled Bunch Instabilities," in *Proceedings of the 1999 Particle Accelerator Conference* (IEEE, 1999), p. 628.

[32] D. Boussard, "Cures of Instabilities," in *CERN Accelerator School: Advanced Accelerator Physics Course*, proceedings of the 5th, Rhodes, 1993, edited by S. Turner, CERN 95-06, p. 391 (1995).

[33] G. Jackson, "Identification and Correction of Fermilab Instabilities," in *Proceedings of the Workshop on Beam Instabilities in Storage Rings* (Press of University of Science and Technology of China, 1994), p. 194.

[34] W. Chou, "Beam Instability Studies for the SSC," in *Proceedings of the Workshop on Beam Instabilities in Storage Rings* (Press of University of Science and Technology of China, 1994), p. 176.

[35] T. Weiland, "Low Impedance Vacuum Chambers," Part. Accel. **51**, 53 (1995); PEP-II-TECH-NOTE-59, 1994.

[36] R. Rimmer, D. Goldberg, G. Lambertson, F. Voelker, K. Ko, N.M. Kroll, R. Pendleton, F. Adams, and M. Dejong, "Higher Order Mode Damping Studies on the PEP-II B Factory RF Cavity," in *Third European Particle Accelerator Conference, Berlin, 1992*, proceedings (Editions Frontieres, 1992), p. 1289.

[37] R. Boni, A. Gallo, F. Marcellini, and G. Vignola, "Operational Experience with the DAPHNE Radio-Frequency Systems," in *Proceedings of the 1999 Particle Accelerator Conference* (IEEE, 1999), p. 866.

[38] M. Svandrlik, A. Fabris, and C. Pasotti, "Improved Methods of Measuring and Curing Multibunch Instabilities in ELETTRA," in *Fifth European Particle Accelerator Conference, Sitges, 1996*, proceedings, edited by S. Myers *et al.* (Institute of Physics Publishing, 1996), p. 1144.

[39] J. Sebek and C. Limborg, "Measurement of RF cavity HOMs with Beam," in *Beam Dynamics Issues for e^+e^- Factories*, proceedings of the 14th Advanced ICFA Workshop, Frascati, Italy, 1997, edited by L. Palumbo and G. Vignola (Frascati Physics Series, 1998), p. 365.

[40] H. Ego, M. Hara, Y. Kawashima, Y. Ohashi, T. Ohshima, H. Suzuki, I. Takeshita, and H. Yonehara, "Suppression of the Coupled-Bunch Instability in the SPring-8 Storage Ring," Nucl. Instrum. Meth. A **400**, 195 (1997).

[41] A. Hofmann and S. Myers, "Beam Dynamics in a Double RF System," in *11th International Conference on High Energy Accelerators*, Geneva, 1980, edited by W.S. Newman (Birkhauser Verlag, 1980), p. 610.

[42] R.A. Bosch and C.S. Hsue, "Suppression of Longitudinal Coupled-Bunch Instabilities by a Passive Higher Harmonic Cavity," Part. Accel. **42**, 81 (1993).

[43] R. Averill, A. Hofmann, R. Little, H. Mieras, J. Paterson, K. Strauch, G-A. Voss, H. Winick, "Synchrotron and Betatron Instabilities of Stored Beams in the CEA," in *Proceedings of the 8th International Conference on High Energy Accelerators*, Geneva, 1971, edited by M.H. Blewett (CERN, 1971), p. 301.

[44] D. Boussard, J. Gareyte, and D. Mohl, "Study and Compensation of Coherent Longitudinal Instability in CERN PS," IEEE Trans. Nucl. Sci. **18**, 1073 (1971).

[45] D. Goldberg and G. Lambertson, "Dynamic Devices: A Primer on Pickups and Kickers," in *Physics of Particle Accelerators*, edited by M. Month and M. Dienes (AIP, 1992), p. 537; LBL-31664-mc (microfiche).

[46] A. Renieri and F. Tazzoli, "The Longitudinal Feedback System in ADONE," in *Proceedings of the 9th International Conference on High Energy Accelerators*, Stanford, 1974, p. 370 (1974).

[47] D. Boussard and G. Lambertson, "Reduction of the Apparent Impedance of Wide Band Accelerating Cavities by RF Feedback," IEEE Trans. Nucl. Sci. **30**, 2239 (1983).

[48] D. Boussard, "Control of Cavities with High Beam Loading," IEEE Trans. Nucl. Sci. **32**, 1852 (1985).

[49] S. Khan and T. Knuth, "Longitudinal and Transverse Feedback Systems for BESSY-II," in *Beam Instrumentation Workshop*, proceedings of the 8th, Stanford, 1998, edited by R. Hettel, S. Smith, and J. Masek (AIP, 1999), p. 537.

[50] S.V. Ivanov, "Impedance Treatment of Longitudinal Coupled-Bunch Feedbacks in a Protron Synchrotron," IHEP Preprint No. IFVE-96-8, 1996.

[51] W. Barry, J. Byrd, and J. Corlett, "The LBL Advanced Light Source (ALS) Transverse Coupled Bunch Feedback System: Recent Commisioning Results," in *Beam Instrumentation Workshop*, proceedings of the 6th, Vancouver, 1994 (AIP, 1995), p. 501.

[52] W. Barry, J. Byrd, J. Corlett, M. Fahmie, J. Johnson, G. Lambertson, M. Nyman, J. Fox, and D. Teytelman, "Design of the PEP-II Transverse Coupled Bunch Feedback

System," in *Proceedings of the 1995 Particle Accelerator Conference*, Dallas (IEEE, 1996), p. 2681.

[53] G. Oxoby, R. Claus, J. Fox, H. Hindi, J. Hoeflich, I. Linscott, J. Olsen, S. Prabhakar, L. Sapozhnikov, J. Corlett, G. Lambertson, A. Drago, and M. Serio,"Bunch-by-Bunch Longitudinal Feedback System for PEP-II", in *4th European Particle Accelerator Conference, London, 1994*, proceedings, edited by V. Suller and Ch. Petit-Jean-Genaz (World Scientific, 1994), p. 1616.

[54] J. Fox, R. Larsen, S. Prabhakar, D. Teytelman, A. Young, A. Drago, M. Serio, W. Barry, G. Stover, "Multibunch Instability Diagnostics via Digital Feedback Systems at PEP-II, DAPHNE, ALS and SPEAR," in *Proceedings of the 1999 Particle Accelerator Conference* (IEEE, 1999), p. 636.

[55] D. Teytelman, J. Fox, H. Hindi, C. Limborg, I. Linscott, S. Prabhakar, J. Sebek, A. Young, A. Drago, M. Serio, W. Barry, and G. Stover, "Beam Diagnostics Based on Time-Domain Bunch-by-Bunch Data," in *Beam Instrumentation Workshop*, proceedings of the 8th, Stanford, 1998, edited by R. Hettel, S. Smith, and J. Masek (AIP Conference Proceedings Volume #451, 1999), p. 222.

[56] Frascati INFN-LNF-90-031(R) (1990).

[57] M. Serio *et al.*, Multibunch Instabilities and Cures," in *Fifth European Particle Accelerator Conference, Sitges, 1996*, proceedings, edited by S. Myers *et al.* (Institute of Physics Publishing, 1996), p. 148.

[58] E. Jaeschke for the BESSY-II Project Team, "Status of the High Brilliance Synchrotron Light Source BESSY-II," in *Proceedings of the 1997 Particle Accelerator Conference*, Vancouver (IEEE, 1998), p. 713.

[59] S. Khan, T. Knuth, "BESSY II Feedback Systems," in *Proceedings of the 1999 Particle Accelerator Conference* (IEEE, 1999), p. 1144.

[60] J. Choi, J.Y. Huang, M.G. Kim, T.-Y. Lee, E.S. Park, and S.S. Chang, "The Operational Status of PLS," in *Proceedings of the 1999 Particle Accelerator Conference* (IEEE, 1999), p. 2418.

[61] Y.J. Kim, J.Y. Huang, M. Kwon, I.S. Ko, "Status of Longitudinal Feedback System for the PLS Storage Ring," in *Proceedings of the 1999 Particle Accelerator Conference* (IEEE, 1999), p. 1076.

[62] R. Claus, J. Fox, I. Linscott, G. Oxoby, W. Ross, L. Sapozhnikov, and D. Teytelman, "Software Architecture of the Longitudinal Feedback System for PEP-II, ALS and DAPHNE," in *Proceedings of the 1995 Particle Accelerator Conference*, Dallas (IEEE, 1996), p. 2660.

[63] D. Briggs, P.L. Corredoura, J.D. Fox, A. Gioumousis, W. Hosseini, L. Klaisner, J.L. Pellegrin, K. Thompson, and G.G. Lambertson, "Prompt Bunch by Bunch Synchrotron Oscillation Detection via a Fast Phase Measurement," in *IEEE 1991 Particle Accelerator Conference*, proceedings, San Francisco (IEEE 1991), p. 1404.

[64] A. Oppenheim, R. Schafer, "Discrete-Time Signal Processing," Prentice Hall (1989).

[65] Franklin, Powell, Workman, "Digital Control of Dynamic Systems," Addison Wesley (1990).

[66] D. Teytelman, PhD thesis, Stanford University (to be published).

[67] H. Hindi, J. Fox, S. Prabhakar, L. Sapozhnikov, G. Oxoby, I. Linscott, and D. Teytelman, "A Formal Approach to the Design of Multibunch Feedback Systems: LQG Controllers," in *4th European Particle Accelerator Conference, London, 1994*, proceedings, edited by V. Suller and Ch. Petit-Jean-Genaz (World Scientific, 1994), p. 1622.

[68] A. Gallo, A. Ghigo, F. Marcellini, M. Migliorati, L. Palumbo, M. Serio, "Simulations of the Bunch-by-Bunch Feedback Operation with a Broadband RF Cavity as Longitudinal Kicker," DAΦNE Technical Note G-31, 1995.

[69] J. Corlett, J. Johnson, G. Lambertson, and F. Voelker, "Longitudinal and Transverse Feedback Kickers for the ALS," in *4th European Particle Accelerator Conference, London, 1994*, proceedings, edited by V. Suller and Ch. Petit-Jean-Genaz (World Scientific, 1994), p. 1625.

[70] R. Boni, A. Drago, A. Gallo, A. Ghigo, F. Marcellini, M. Migliorati, M. Serio, and M. Zobov, "Kickers and Power Amplifiers for the DAΦNE Bunch by Bunch Longitudinal

Feedback System," in *Fifth European Particle Accelerator Conference, Sitges, 1996*, proceedings, edited by S. Myers *et al.* (Institute of Physics Publishing, 1996), p. 1881.

[71] S. Prabhakar, D. Teytelman, J. Fox, and H. Hindi, "Use of Digital Feedback System as a Bunch Current Monitor: Results from ALS," SLAC-PEP-II-AP-NOTE-96-29 (1996).

[72] S. Prabhakar, R. Claus, J. Fox, H. Hindi, I. Linscott, J. Olsen, W. Ross, and D. Teytelman, "Observation and Modal Analysis of Coupled Bunch Longitudinal Instabilities via a Digital Feedback Control System," Part. Accel. **57**, 175 (1997); SLAC Report No. SLAC-PUB-7717.

[73] U. Wienands *et al.*, "First Beam-Commissioning Results from the PEP-II B-factory High Energy Ring: Intensity Effects," in *Beam Dynamics Issues for e^+e^- Factories*, proceedings of the 14th Advanced ICFA Workshop, Frascati, Italy, 1997, edited by L. Palumbo and G. Vignola (Frascati Physics Series Vol. X, 1998), p. 379.

[74] J. Seeman, private communication.

[75] R. Kohaupt, "Single Beam Instabilities in DORIS," IEEE Trans. Nucl. Sci. **22**, 1456 (1975).

[76] R. Stiening and J. Griffin, "Longitudinal Instabilities in the Fermilab 400-GeV Main Accelerator," IEEE Trans. Nucl. Sci. **22**, 1859 (1975).

[77] J. Corlett and J. Byrd, "Measurement and Computation of the Higher Order Modes of the ALS 500-MHz Accelerating Cavities," in *1993 IEEE Particle Accelerator Conference*, proceedings, Washington, D.C. (IEEE, 1994), p. 3408.

[78] R.A. Rimmer, J. Byrd, M. Irwin, and D.A. Goldberg, "Updated Impedance Estimate of the PEP-II RF Cavity," in *Fifth European Particle Accelerator Conference, Sitges, 1996*, proceedings, edited by S. Myers *et al.* (Institute of Physics Publishing, 1996), p. 2035.

[79] A. Jackson, "Commissioning and Performance of the Advanced Light Source," in *1993 IEEE Particle Accelerator Conference*, proceedings, Washington, D.C. (IEEE, 1994), p. 1432.

[80] J. Byrd (to be published); D. Teytelman (to be published).

[81] S. Prabhakar, J. Fox, H. Hindi, D. Teytelman, and A. Young, "Calculation of Impedance from Multibunch Synchronous Phases: Theory and Experimental Results," in *Sixth European Particle Accelerator Conference, Stockholm, 1998*, proceedings, edited by S. Myers *et al.* (Institute of Physics Publishing, 1998), p. 996.

[82] S. Prabhakar et al, "Low-Mode Longitudinal Motion in the PEP-II HER," SLAC-PEP-II-AP-NOTE-98-06, 1998.

[83] S. Prabhakar, D. Teytelman, J. Fox, M. Minty, U. Wienands, and A. Young, "Measurements of a Fast Vertical Instability in the PEP-II HER," SLAC-PEP-II-AP-NOTE-99-04 (1999).

[84] P. Corredoura, "Architecture and Performance of the PEP-II Low Level RF System," in *Proceedings of the 1999 Particle Accelerator Conference* (IEEE, 1999), p. 435.

[85] P. Wilson, "Mode Impedances for PEP- and Petra-Type Cavities," CERN Report No. LEP/70/69, 1978.

[86] G. Stupakov, T. Raubenheimer, and F. Zimmermann, "Effect of Ion Decoherence on Fast Beam-Ion Instability," Phys. Rev. E **52**, 5499 (1995).

[87] A.N. Dubrovin, A.S. Kalinin, D.N. Shatilov, E.A. Simonov, and V.V. Smaluk, "Applications of Beam Diagnostic System at the VEPP-4," in *Fifth European Particle Accelerator Conference, Sitges, 1996*, proceedings, edited by S. Myers *et al.* (Institute of Physics Publishing, 1996), p. 1585.

[88] G. Morpurgo, "The BOM 1000 Turn Display: a Tool to Visualize the Transverse Phase-Space Topology at LEP," in Proceedings of the 1999 Particle Accelerator Conference, New York, 1999, p. 1571.

[89] J. Byrd, A. Chao, S. Heifets, M. Minty, T.O. Raubenheimer, J. Seeman, G. Stupakov, J. Thomson, and F. Zimmermann, "First Observations of a Fast Beam Ion Instability," Phys. Rev. Lett. **79**, 79 (1997).

[90] J.Y. Huang, M. Kwon, T.-Y. Lee, I.S. Ko, Y.H. Chin, and H. Fukuma, "Direct Observation of the Fast Beam-ion Instability," Phys. Rev. Lett. **81**, 4388 (1998).

[91] S. Prabhakar, J.D. Fox, D. Teytelman, and A. Young, "Phase Space Tracking of Coupled-Bunch Instabilities," Phys. Rev. ST Accel. Beams **2**: 084401, 1999.

[92] G. Stupakov, "Fast Ion Instability in Real Lattice," in *Proceedings of the 1997 Particle Accelerator Conference*, Vancouver (IEEE, 1998), p. 1632.

[93] S. Prabhakar, D. Teytelman, J. Fox, A. Young, P. Corredoura, and R. Tighe, "Commissioning Experience from PEP-II HER Longitudinal Feedback," in *Beam Instrumentation Workshop*, proceedings of the 8th, Stanford, 1998, edited by R. Hettel, S. Smith, and J. Masek (AIP, 1999), p. 529.

[94] Y.H. Chin and K. Yokoya, "Landau Damping of a Multibunch Instability due to Bunch to Bunch Tune Spread," DESY Report No. DESY 86-097, 1986.

[95] M. Billing, "Observation of a Longitudinal Coupled Bunch Instability with Trains of Bunches in CESR," CLNS Report No. CLNS 98/1564, 1998.

[96] K. Harkay, A. Nassiri, J.J. Song, Y.W. Kang, and R.L. Kustom, "Compensation of Longitudinal Coupled-Bunch Instability in the Advanced Photon Source Storage Ring," in *Proceedings of the 1997 Particle Accelerator Conference*, Vancouver (IEEE, 1998), p. 1575.

[97] O. Naumann and J. Jacob, "Fractional Filling Induced Landau Damping of Longitudinal Instabilities at the ESRF," in *Proceedings of the 1997 Particle Accelerator Conference*, Vancouver (IEEE, 1998), p. 1551.

[98] S. Prabhakar, J.D. Fox, and H. Hindi, "A Matrix Formalism for Landau Damping," in *Beam Dynamics Issues for e^+e^- Factories*, proceedings of the 14th Advanced ICFA Workshop, Frascati, Italy, 1997, edited by L. Palumbo and G. Vignola (Frascati Physics Series Vol. X, 1998), p. 385; SLAC Report No. SLAC-PUB-7978.

[99] Spear 3 Conceptual Design Report (in press).

[100] H. Winick, private communication.

[101] R. Rimmer, private communication.

[102] D. Teytelman (to be published).