



Bunch-by-bunch feedback studies at LCLS

APRIL 3–7, 2012

Author:
Dmitry TEYTELMAN

Revision:
1.0

August 31, 2012

Copyright © Dimtel, Inc., 2012. All rights reserved.

Dimtel, Inc.
2059 Camden Avenue, Suite 136
San Jose, CA 95124
Phone: +1 650 862 8147
Fax: +1 603 907 0210
www.dimtel.com

Contents

1	Introduction	2
1.1	Feedback Hardware Configuration	2
2	Results	3
2.1	Front-end calibration	4
2.2	Back-end timing	5
2.3	Measurements near the instability threshold	6
2.4	High current measurements	7
2.4.1	Unstable modes and instability sources	9
2.4.2	Tuner position	9
2.4.3	Cavity temperatures	9
2.4.4	Residual motion	10
2.4.5	Feedback gain estimates	12
2.5	Transverse measurements	14
3	Simulation	14
3.1	Linear model	15
3.2	Time-domain model	16
3.3	Model	17
3.4	Simulated grow/damp	18
3.5	Extrapolation	19
4	Summary	22

Table 1: UVX parameters

Parameter description	Symbol	Value
Nominal RF frequency	f_{RF}	476 MHz
Harmonic number	h	148
Momentum compaction	α	8.3×10^{-3}
Beam energy at injection	E_0	500 MeV
Beam energy in operation	E_0	1.37 GeV
Radiation damping time, longitudinal	τ_{rad}	3.7 ms

1 Introduction

During the week of April 3–7, 2012, bunch-by-bunch feedback system from Dimtel, Inc. [1] was demonstrated in the longitudinal plane of the LNLS UVX storage ring. This brief note summarizes the results of the coupled-bunch instability studies performed in the course of this demonstration.

The main goal of these experiments has been to characterize longitudinal coupled-bunch instabilities in LNLS UVX and to collect sufficient data for defining the parameters of the future permanent feedback installation.

Main parameters of LNLS UVX ring during our studies are listed in Table 1.

1.1 Feedback Hardware Configuration

Bunch-by-bunch feedback signal processing and diagnostics were provided by iGp12 baseband processor [1]. BPM sum signal was processed by a custom low-noise prototype unit, operating at the third harmonic of the ring RF frequency (1428 MHz).

LNLS UVX storage ring is not equipped with a longitudinal feedback kicker at this point. Our plan was to use transverse feedback striplines driven common mode as a weak longitudinal kicker. Our initial setup used a single stripline driven by a baseband amplifier (Amplifier Research 75A250AM2). Initial attempts were quite discouraging — applications of various excitation signals generated no discernible response from the beam.

After some discussion it has been determined that the striplines in question are 15 cm long. Longitudinal shunt impedance of striplines has frequency dependence of the form $\sin(\omega l/c)^2$ with periodic peaks at $(2n + 1)c/4l$. For

$l = 0.15$ m, the first peak is at 500 MHz. Baseband kick signal has the majority of power in the DC–238 MHz band, where the shunt impedance is very low.

In order to increase the kick voltage, an improvised back-end was set up as shown in Fig. 1. Baseband output of the iGp12 was upconverted to 476 MHz using a double-balanced mixer. A remotely adjustable delay line provided carrier phase adjustment. Mixer output was amplified to provides sufficient signal level for driving a power splitter and two power amplifiers: E&I 525LA¹ and Kalmus². To control differential phasing of the striplines two adjustable delay lines were used. Given that the carrier phase is adjustable, only one line was necessary, but a symmetric setup helped provide matched envelope timing in the two paths.

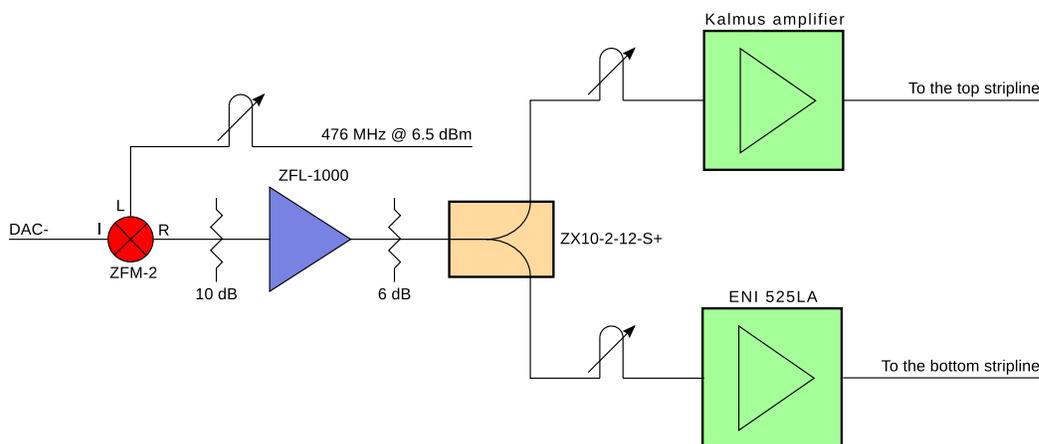


Figure 1: Back-end setup.

2 Results

The main focus of our measurements was on quantifying the longitudinal coupled-bunch instabilities in LNLS UVX. In addition, we have explored bunch cleaning functionality in the vertical plane. In this section the measurements are described in detail and the results are summarized.

¹25 W in 1–500 MHz band

²Model unknown, similar to 550FC, 50 W in 230–520 MHz band

2.1 Front-end calibration

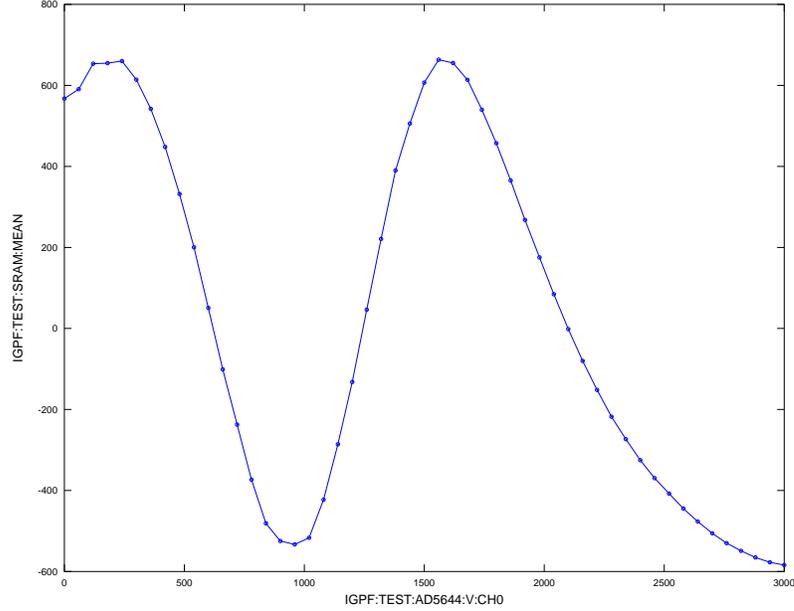


Figure 2: Single-bunch front-end phase shifter sweep at 1.3 mA.

2.1 Front-end calibration

In order to convert the detected beam signal from ADC counts to physical units we performed a front-end calibration. After timing the front end in a single bunch fill configuration, we swept the front-end carrier phase and recorded bunch average signal. With the mixer in the linear regime, the ADC input is given by

$$S_{\text{ADC}} = GI_b \sin(3\Delta\phi) \quad (1)$$

where I_b is the bunch current, $\Delta\phi$ is the phase difference between the local oscillator and the beam signal, and G is the sensitivity in ADC counts per milliamperere. By sweeping $\Delta\phi$ through 360 degrees we extract the full-scale amplitude GI_b . Since bunch current is known, we can calculate the small-signal gain in the phase detector mode — basically the slope at the zero crossing point. Figure 2 shows the phase shifter sweep plot. Calibration

constant is calculated to be 25.2 counts/mA/deg³.

2.2 Back-end timing

Back-end timing was extremely challenging in this configuration due to several factors. Firstly, the hardware setup generates fairly low kick voltages due to the low shunt impedance and limited amplifier power. Secondly, single bunch beam in LNLS UVX exhibits significant oscillation at the synchrotron frequency driven by the RF systems and exacerbated by the relatively long radiation damping times. As a result, the standard approach of exciting the beam at the synchrotron frequency was unusable — we could not observe the small feedback system excitation on top of the large natural oscillation.

A different approach was then utilized. Rough timing and phasing was performed by observing excitation kick and beam signals at the upstream ports of the striplines using a high-speed oscilloscope. Next, the feedback system was configured with two filters, with phase shift of ± 90 degrees at the synchrotron frequency⁴. EPICS filter set selection control was toggled back and forth once every 20 seconds. We then observed the magnitude of a spectral marker placed at the synchrotron frequency on a stripchart. Periodic switching between negative and positive feedback produced a modulation of the marker magnitude. Next, we adjusted timing and phasing parameters in order to amplify the modulation amplitude. Initially we observed an effect on the order of 1–2 dB. At the end of the procedure, toggling between negative and positive feedback resulted in a 16 dB change in synchrotron oscillation amplitude. Figure 3 shows the feedback action as observed on the stripchart. Initially the feedback loop is open. At around 15:52, negative feedback is turned on, producing a 6–8 dB reduction in oscillation amplitude. Next, at 15:53:30, feedback sign is changed from negative to positive. As a result, we see 10–12 dB oscillation increase. The plot clearly demonstrates why the stripchart modulation technique is necessary. Over time, observed amplitude of externally driven synchrotron motion varies significantly due to the changes in overall excitation amplitude and the beating between the short data acquisition snapshot and the time-varying envelope of oscillation. For example, at the end of the stripchart plot the amplitude drops by 4 dB

³Later during experimental measurements we have replaced a 20 dB attenuator before the front end by a 10 dB one, thus increasing the front-end gain by $\sqrt{10}$.

⁴In the longitudinal plane, due to small phase advance per turn the required controller phase shift is very close to the theoretical 90 degrees

2.3 Measurements near the instability threshold

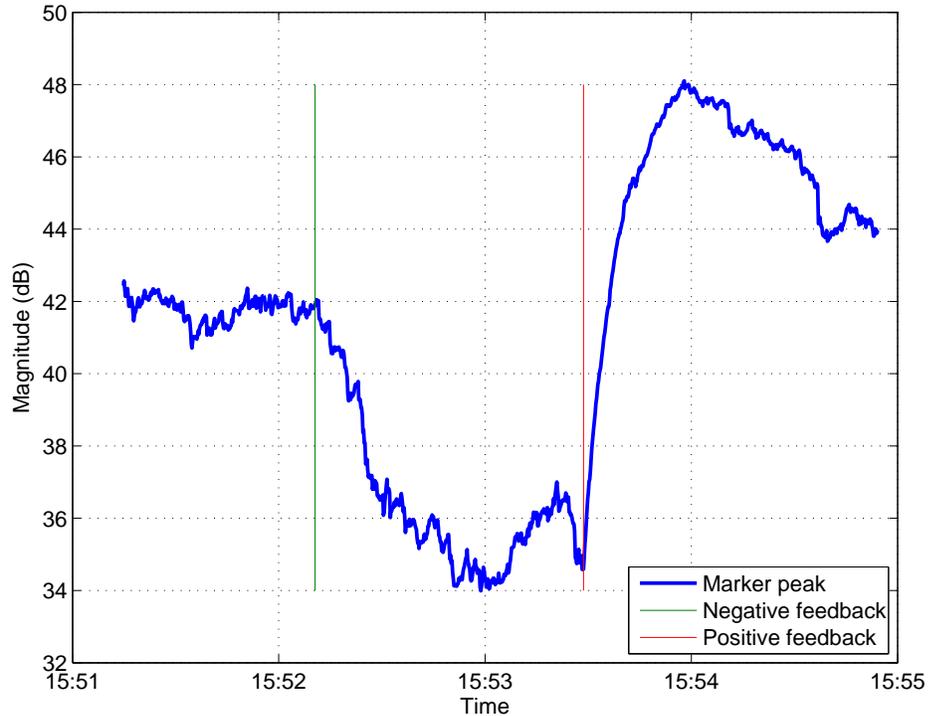


Figure 3: A stripchart plot showing the effect of negative and positive feedback on synchrotron oscillation amplitude.

without any changes in system configuration. Observing small changes in the feedback effectiveness due to timing and phasing adjustments is only feasible with synchronous detection.

2.3 Measurements near the instability threshold

Initial studies were performed at 1.37 GeV at beam currents very close to the instability threshold of 12.5 mA. Below the threshold with the beam stable we performed what is known as drive/damp measurements. The feedback is positive in the first part of the transient and then the loop is open in the second part. From that second part we can extract the open-loop damping rate. Above the threshold, traditional grow/damp measurements were performed [2].

Figure 4 presents the measured open-loop growth and damping rates.

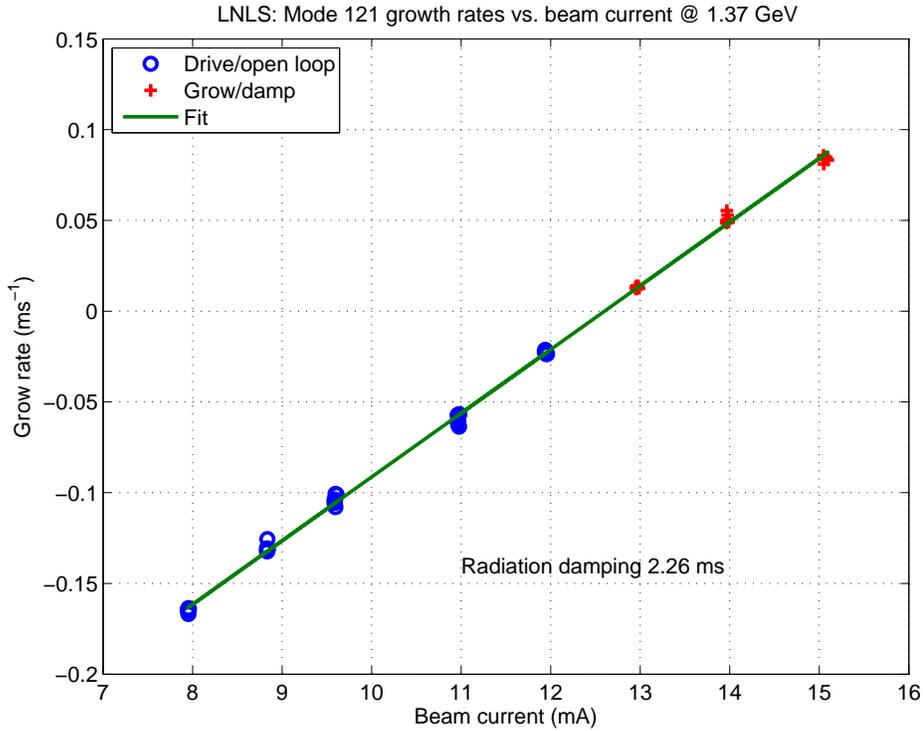


Figure 4: Measurements of eigenmode 121 open-loop growth and damping rates in the vicinity of the instability threshold.

From the linear fit to the measurements we estimate zero-current damping of 2.26 ms. The value is smaller than the expected radiation damping time of 3.7 ms [3]. The discrepancy is most likely due to the fact that we have measured coherent damping times. During the damping transients some energy is transferred into intrabunch motion, not observed by the dipole oscillation detector. As a result, dipole motion seems to damp faster.

2.4 High current measurements

At the start of the next experimental period on April 5th, 2012 the machine was already filled with 40 mA. We attempted to stabilize the beam — and succeeded. With the low kick voltage available from the improvised back-end setup we could not expect to capture large amplitude oscillations directly. Instead, we used phase modulation of the RF voltage at the quadrupole fre-

2.4 High current measurements

quency (slightly below the second harmonic of the synchrotron frequency) to excite quadrupole intrabunch motion and damp the dipole oscillations [4, 5]. Once dipole oscillation amplitude comes into the linear range of the bunch-by-bunch feedback channel, the feedback damps it to the noise floor. Next, we turn off the phase modulation, leaving the accelerator in a longitudinally stabilized state. At this point, grow/damp measurements can be performed to quantify the instability growth rates and feedback-induced damping.

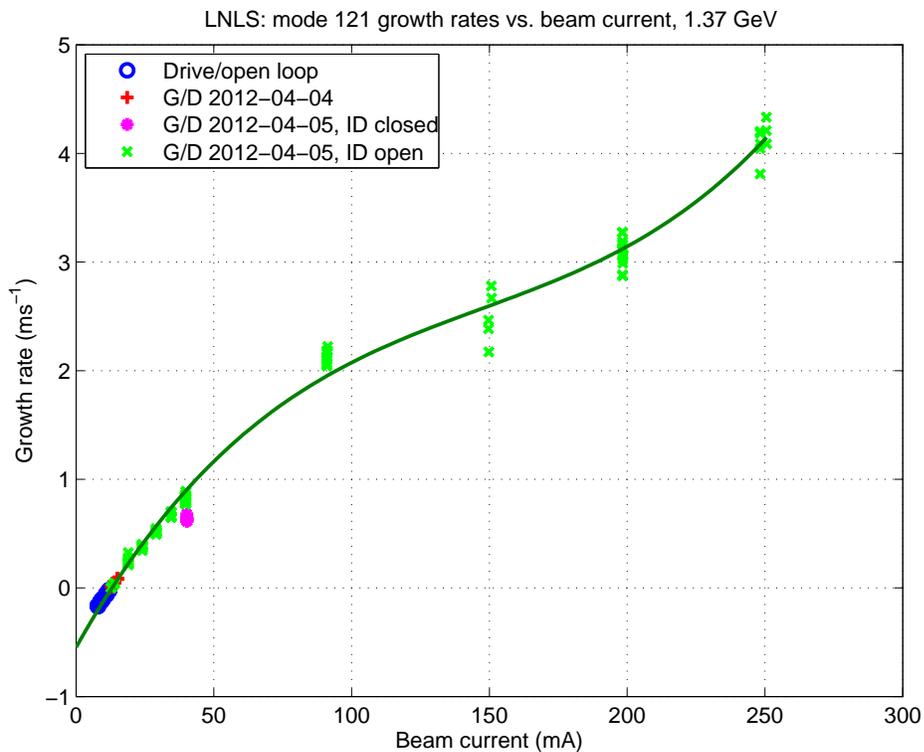


Figure 5: Measurements of eigenmode 121 open-loop growth rates.

Using phase modulation capture mechanism we were able to make measurements at currents up to 250 mA. Measured growth rates are shown in Fig. 5. It is immediately obvious that growth rate dependence on beam current is non-linear. That can indicate that the driving impedance is changing with beam current or that additional damping mechanisms are in play. Two sets of measurements at 40 mA show the growth rates measured with insertion device gaps open and closed. With ID gaps closed the growth rates are

slower, consistent with the expected increase in radiation damping.

2.4.1 Unstable modes and instability sources

Longitudinal instabilities in LNLS UVX are dominated by eigenmode 121. At the very highest current of 250 mA mode 43 was also observed. Figure 6 shows a grow/damp measurement with both modes active. The two observed modes are in rough agreement with the HOMs presented in [4]. Eigenmode 43 is likely to be driven by L9 mode at 2040.125 MHz. Mode L3 is a good candidate for the eigenmode 121 impedance. However there is a significant frequency disagreement: 1356.89 MHz in the reference and 1341.19 MHz in our measurements.

There is a significant difference in the damping rates generated by the feedback system for the two modes. Mode 43 eigenvalue shift is $5316 - 3388i \text{ s}^{-1}$ and mode 121 — $9294 + 4671i \text{ s}^{-1}$. Absolute values of these shifts differ by a factor of 0.61. Stripline kicker impedance ratio at $\omega_{\text{rf}} - 43\omega_{\text{rev}} = 337.7 \text{ MHz}$ and $121\omega_{\text{rev}} = \omega_{\text{rf}} - 27\omega_{\text{rev}} = 389.2 \text{ MHz}$ is 0.86. That explains only part of the difference. The rest is probably due to the gain and phase non-linearity of the two power amplifiers.

Next we have decided to investigate the sensitivity of the longitudinal instability growth rates on cavity operating points, specifically: tuner plunger positions and cavity temperatures.

2.4.2 Tuner position

Figure 7 shows the results of tuner position scans on cavities A and B at 48.6 mA. Each data point presents mean and standard deviation for a set of 7–11 growth rate measurements. The data is consistent with no or small dependence of growth rates on tuner positions.

2.4.3 Cavity temperatures

Measurements of eigenmode 121 growth rates as a function of cavity temperatures are shown in Fig. 8. Over the range of temperatures explored in the scan, the growth rates are constant or weakly dependent on temperature. Considering temperature dependence plots shown in [4], that is to be expected, since temperature sensitivity of longitudinal HOMs is relatively low. All 111 data sets were also checked for appearance of eigenmodes other than 121 — none found.

2.4 High current measurements

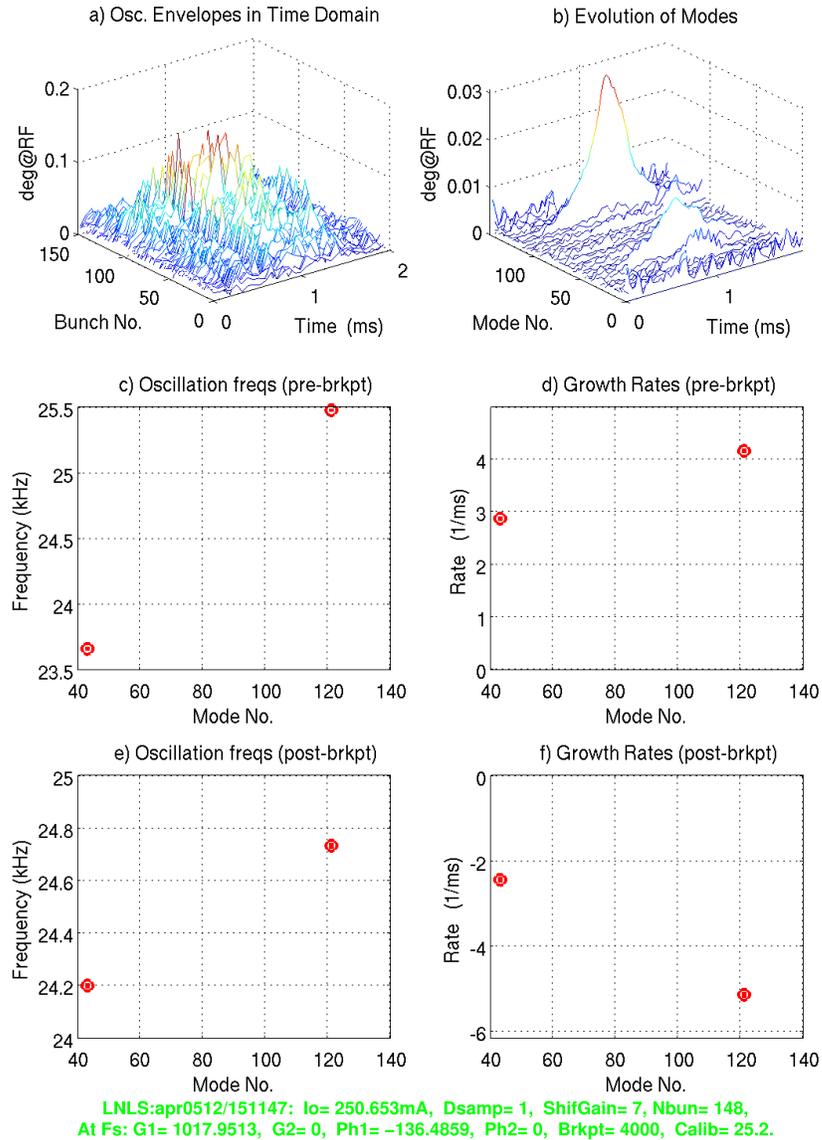


Figure 6: A grow/damp measurement at 250 mA showing growth and damping of modes 43 and 121.

2.4.4 Residual motion

With a bunch-by-bunch feedback system in operation, the dipole beam motion settles at a finite level. In the longitudinal plane that level is typically

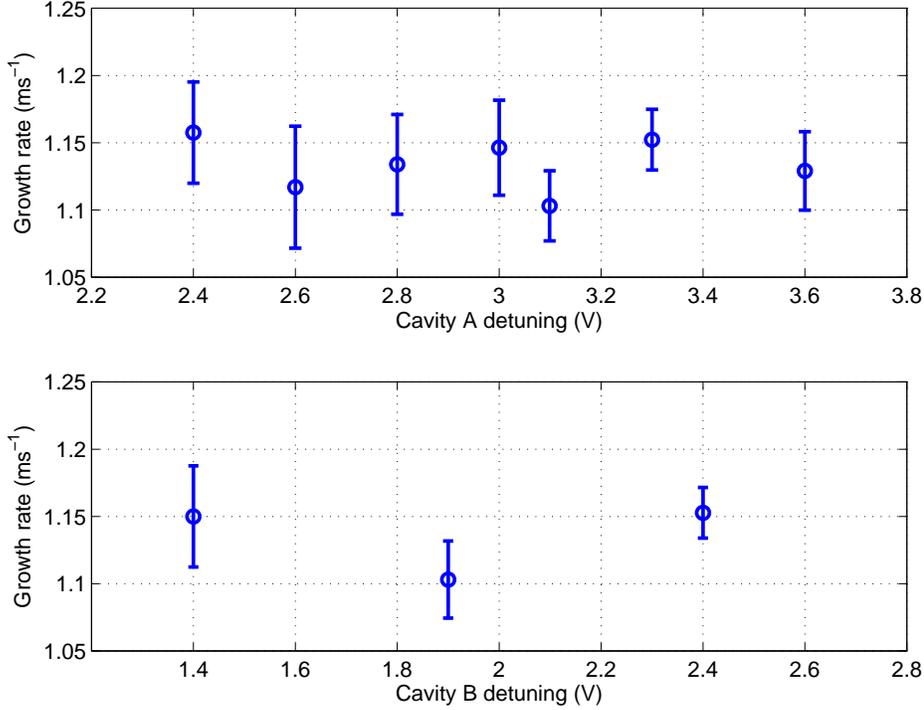


Figure 7: Mode 121 growth rates vs. tuner positions.

defined by the external excitation of the beam in the vicinity of the RF frequency. Such excitation is due to the amplitude and phase noise in the cavity drive signal, both from the low-level electronics and the RF amplifier (klystron, IOT, solid-state) power supply. These spectral components excite beam motion near the lowest synchrotron frequency resonance — eigenmode 0.

While the feedback can have a dramatic effect on the amplitudes of the unstable modes, externally driven motion is only reduced by a factor linearly related to the loop gain. Steady-state residual motion is one of three factors driving the power amplifier and kicker requirements⁵.

Figure 9 shows the measurements of the residual motion as a function of beam current. Residual oscillation decreases significantly at high currents. An inverse linear fit closely follows the data, but the dependence cannot

⁵The other two factors are the open-loop growth rates and the transient excitation amplitudes, e.g. injection

2.4 High current measurements

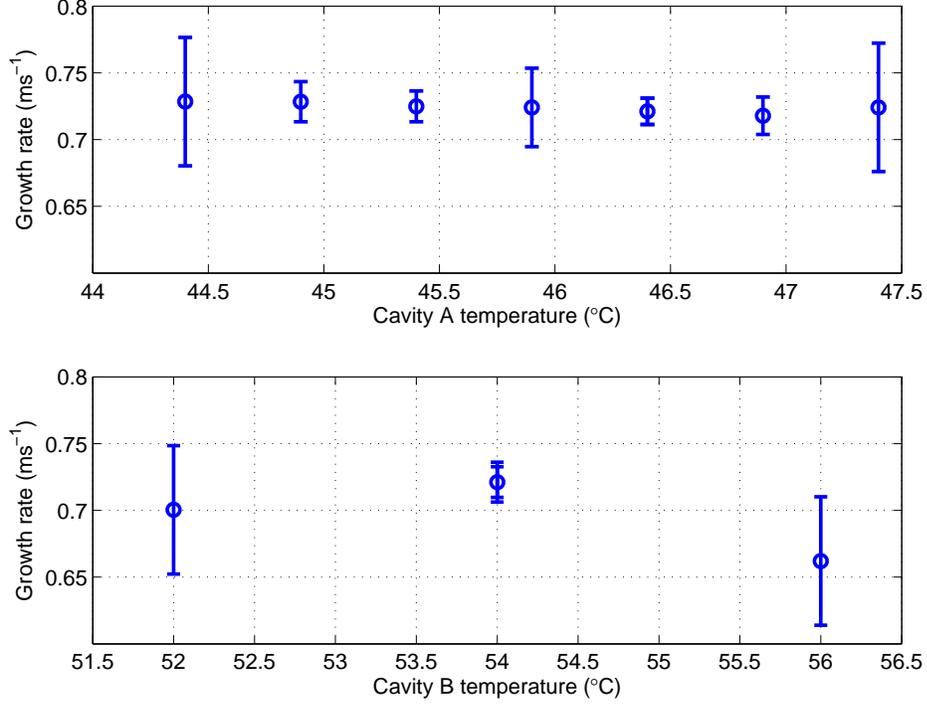


Figure 8: Mode 121 growth rates vs. cavity temperatures.

be explained by the increasing bunch-by-bunch feedback gain. If we are to assume the external excitation is constant, feedback must reach loop gains larger than 30 at 250 mA to achieve this dependence. As we find in Section 3 using grow/damp measurements, peak loop gains at 250 mA are around 1.9. Therefore we must conclude that the reduction in the residual motion with beam current is due to the decrease in the external excitation.

2.4.5 Feedback gain estimates

Grow/damp measurements allow us to estimate the feedback gain and the peak kick voltage. To estimate feedback gain one can use the standard feedback relationship [6]:

$$\lambda_{\text{fb}} = \frac{f_{\text{RF}}^2 \alpha e}{2E_0 h f_s} G_{\text{fb}} \quad (2)$$

where λ_{fb} is the magnitude of the eigenvalue shift due to the feedback action.

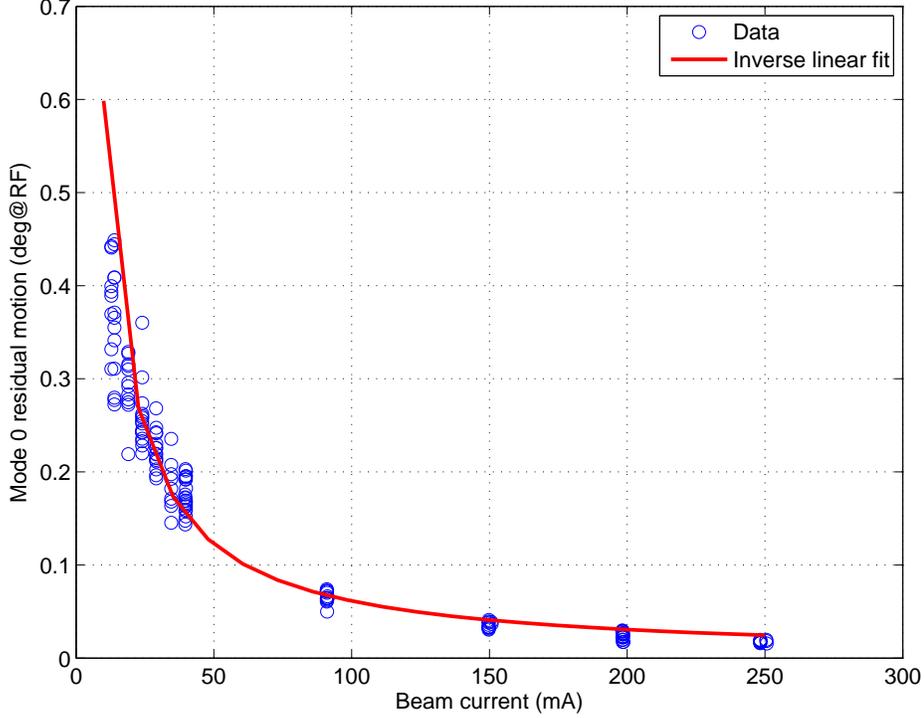


Figure 9: Residual mode 0 longitudinal motion versus beam current.

Feedback-induced shift of mode 121 eigenvalue for the measurement presented in Fig. 6 is $9294 + 4671i \text{ s}^{-1}$, corresponding to a gain of 57 kV/rad. We can use the estimated gain to compute the peak kick voltage. Overall feedback gain is given by

$$G_{\text{fb}} = \frac{V_{\text{max}}}{2047} |H_{\text{fir}}(\omega_s)| i_b G_{\text{fe}} \quad (3)$$

where H_{fir} is the FIR filter transfer function, i_b is the bunch current and G_{fe} is the front-end calibration in counts/rad/A. Using these parameters we estimate $V_{\text{max}} = 23.6 \text{ V}$.

Feedback loop gain and peak kick voltage are further studied in Section 3, where we use grow/damp measurements to calibrate several mathematical models of the beam and the feedback system.

2.5 Transverse measurements

In the transverse plane we have explored the technique of bunch cleaning. In iGp12, bunch cleaning procedure is implemented as follows. For each RF bucket one can enable or disable feedback and excitation signals. Feedback is enabled for the bunches to be preserved and turned off for the ones to clean. Drive signal is enabled only for the bunches we want to eliminate. Drive generator is a swept sinewave NCO, programmed to sweep through the betatron tune range and excite the beam to large transverse oscillation amplitudes, leading to scraping and current loss.

To document the cleaning process we have used a streak camera, swept at 119 MHz — a quarter of the RF frequency. The measurement process presents integrated bunch profiles from groups of bunches spaced by 4 RF buckets. In Figure 10(a), streak camera image is shown for bunch groups 1 and 3⁶, while groups 2 and 4 are shown in Fig. 10(b). Both of these were taken after full ring injection to 145 mA. Next, we applied bunch cleaning to group 1, getting rid of bunches 1, 5, 9, 13, Beam current was reduced to 104 mA and the resulting streak camera image for groups 1 and 3 is shown in Fig. 10(c). In the last step we cleaned group 2 — streak camera result for groups 2/4 at 62 mA is presented in Fig. 10(d).

Further analysis of the streak camera images is presented in Fig. 11. Pixel values are integrated horizontally to generate a profile of the bunch group. In Figure 11(a), bunch groups 1 and 3 are shown before and after the cleaning, while Fig. 11(b) shows groups 2 and 4.

Based on the streak camera data we can conclude that there was no detectable light output from the cleaned bunch groups. To fully characterize the bunch cleaning process, a diagnostic instrument with higher dynamic range is required.

3 Simulation

In order to define the power amplifier and kicker requirements for the permanent feedback system, we need a good mathematical model of the temporary configuration. Once that model is calibrated to the grow/damp and closed-loop measurements, it can be used to estimate the performance with

⁶Bunch group 1 includes bunches 1, 5, 9, 13, ..., similarly, group 2 is 2, 6, 10, 14, ..., group 3 — 3, 7, 11, 15, ..., and group 4 — 4, 8, 12, 16, ...

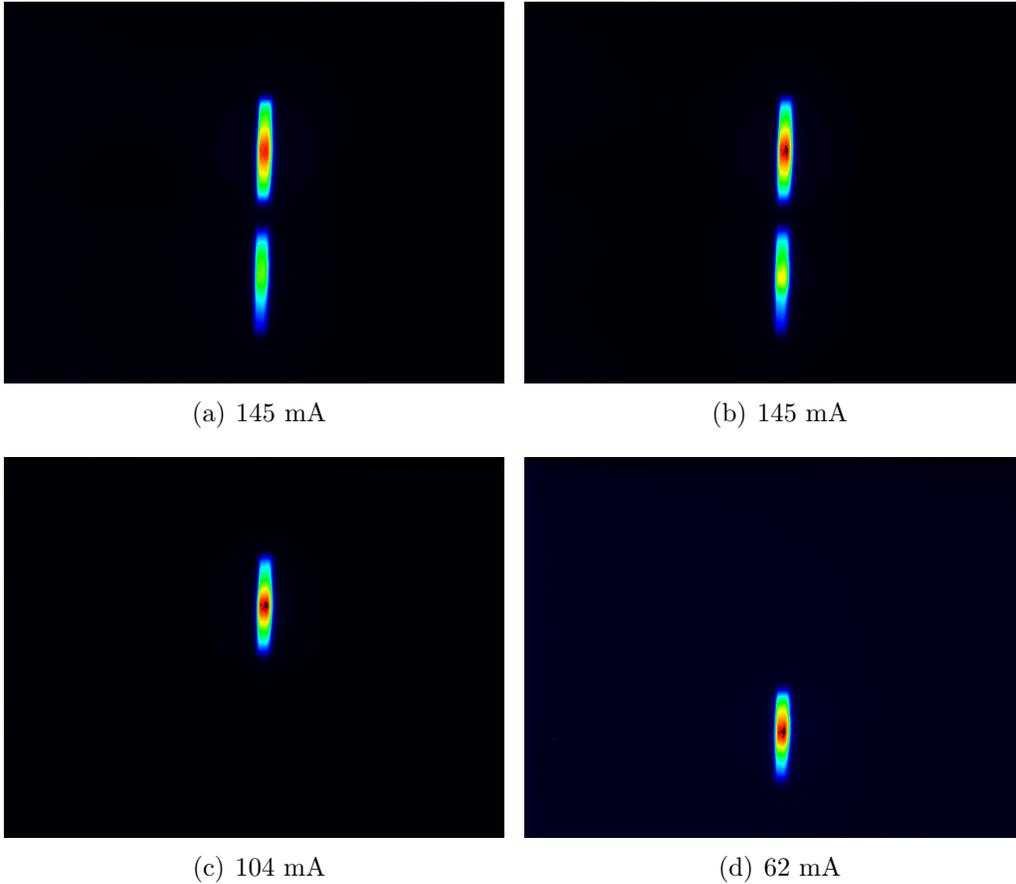


Figure 10: Streak camera record of bunch cleaning. a) Integrated light from bunch groups 1/3; b) Integrated light from bunch groups 2/4; c) Bunch groups 1/3 after cleaning the pattern 1:4:148; d) Bunch groups 2/4 after cleaning the pattern 2:4:148.

a dedicated longitudinal kicker and power amplifier.

3.1 Linear model

One method of feedback system modeling is described in [7, chapter 5]. In this method closed-loop eigenvalues are calculated using a discrete-time state-space linear model. Since the model is parametrized and the computation is relatively fast, numeric optimization is used to fit the model to a set of

3.2 Time-domain model

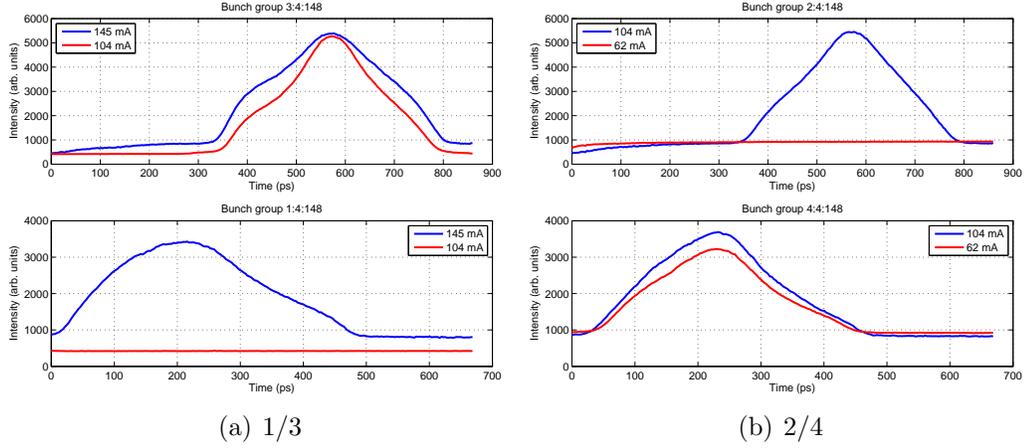


Figure 11: Bunch profiles before and after cleaning

grow/damp measurements. Results of such fitting to grow/damp measurements at currents from 90 to 250 mA are shown in the root locus plot, Fig. 12. The real part of the eigenvalue (growth or damping rate) is plotted on the horizontal axis and the imaginary part (oscillation frequency) — on the vertical axis. There is reasonable agreement between measured and modeled eigenvalues. Estimated kicker peak voltage is 17.7 V — similar to a single data set estimate presented earlier. However, all of these grow/damp measurements were made at a shift gain of 7 (post-FIR gain of 128). At such a high gain the output of the feedback is strongly saturated, making accurate model extraction difficult. A search of all grow/damp data sets found three measurements with a shift gain of 4. These measurements are not in saturation and provide a better modeling base. State-space model fitting comes up with a peak kicker voltage of 39 V — more than twice the value estimated in saturation.

3.2 Time-domain model

A time-domain Simulink model shown in Fig. 13 can be used to more carefully match the physical system, replicating not only loop gains but also the noise sources and steady-state residual levels.

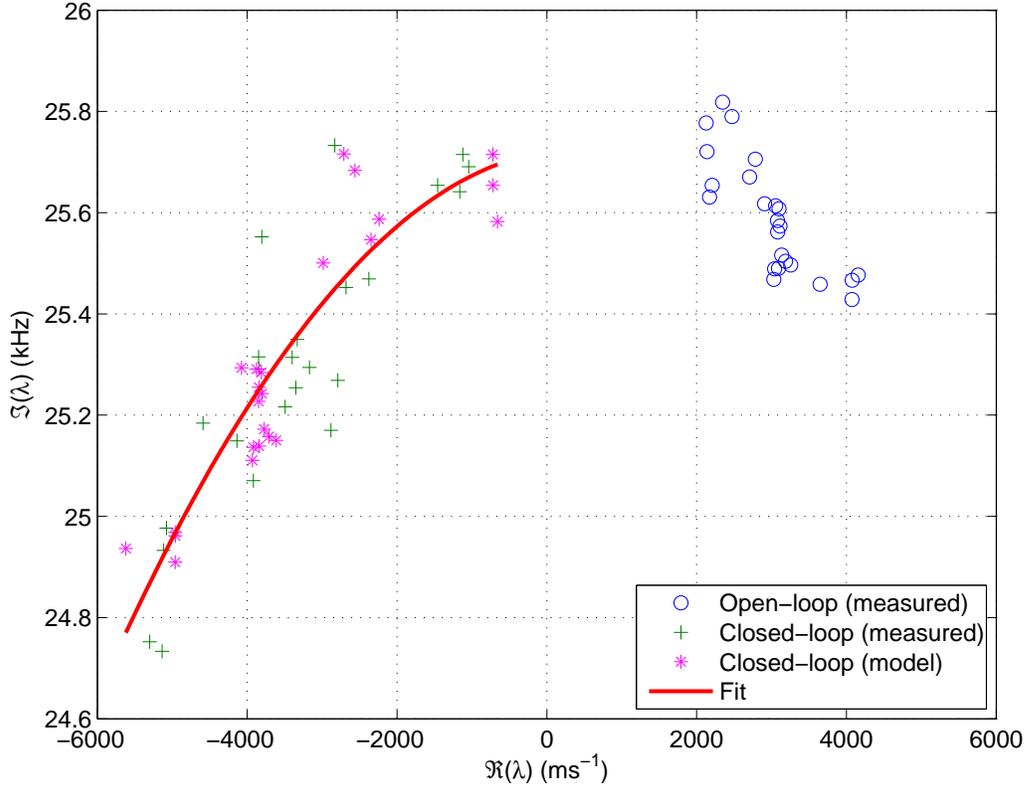


Figure 12: Root locus of open and closed loop eigenvalues (mode 121)

3.3 Model

The model is configured with the accelerator parameters from Table 1. Unstable beam mode is set to the growth rate and the oscillation frequency extracted from a particular grow/damp measurement. Feedback system parameters (front-end calibration) are based on the actual beam calibration measurements.

By adjusting peak kick voltage V_{\max} in the model we match the closed-loop damping to that measured in the real accelerator. In reality many gain factors affect the damping rate - front-end calibration, bunch current, kicker shunt impedance, power amplifier gain, etc. All of the factors that can be directly measured, are quantified and fixed in the model.

3.4 Simulated grow/damp

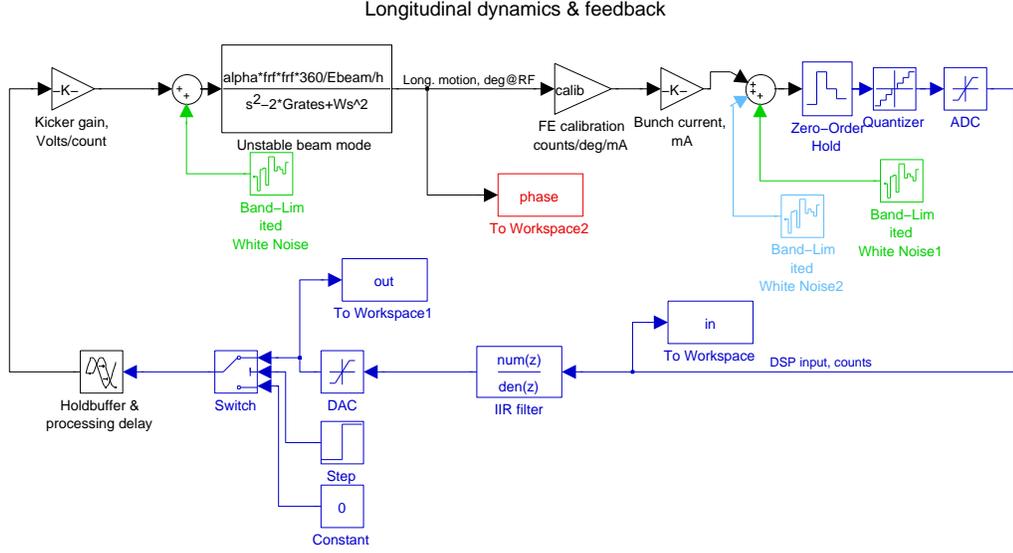


Figure 13: Simulink model of the beam and bunch-by-bunch feedback.

Table 2: Eigenvalues for the measured and simulated grow/damp transients

Parameter description	Measured	Simulated
Open-loop growth rate, ms^{-1}	0.75	0.74
Open-loop oscillation frequency, Hz	25768	25765
Closed-loop damping rate, ms^{-1}	0.17	0.17
Closed-loop oscillation frequency, Hz	25848	25838

3.4 Simulated grow/damp

Figure 14 compares measured and simulated grow/damp transients. Kicker voltage was adjusted to match measured and simulated closed-loop damping rates. Growth transient amplitude is matched to the measurement by adjusting the wideband RF noise excitation level in the model (Band-limited white noise block). Table 2 summarizes the estimated eigenvalues for the two transients.

Matching the damping rate in the model to that measured in the physical system requires $V_{\text{max}} = 44.1$ V. That is reasonably close to 39 V estimated

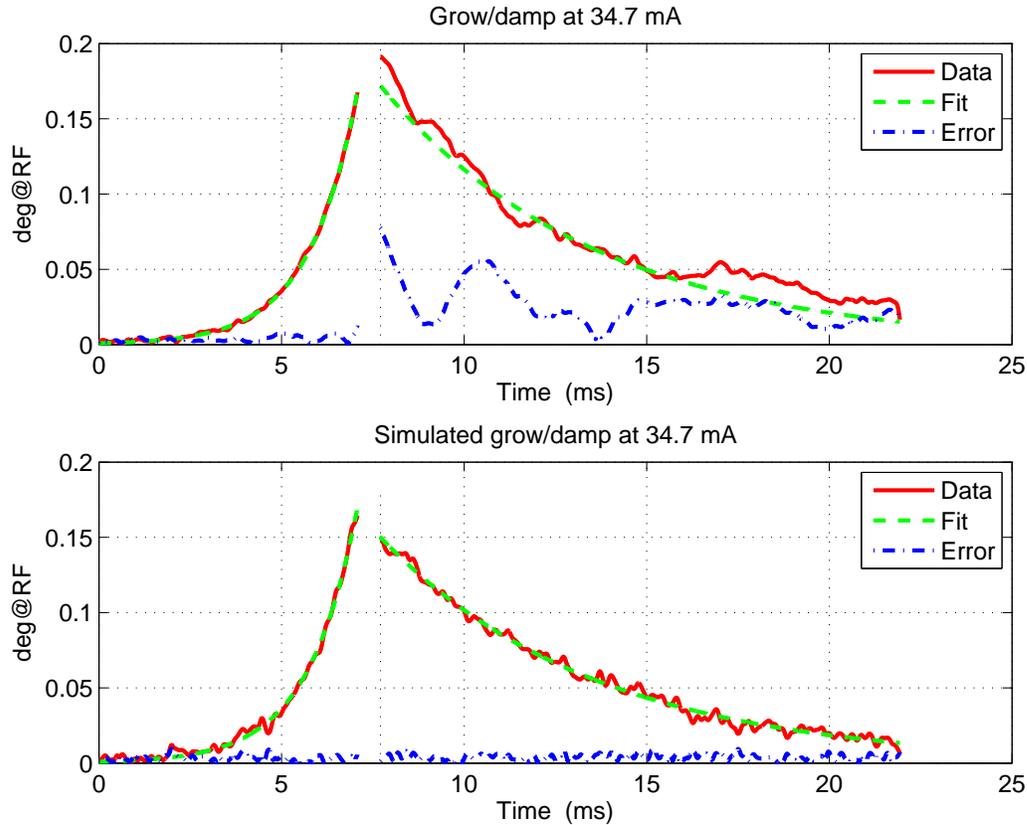


Figure 14: Measured and simulated grow/damp transients at 34.7 mA.

earlier. If we assume that our output stage was producing 75 W of power at the maximum output, we can estimate the effective shunt impedance of the striplines as 13Ω .

3.5 Extrapolation

Having calibrated a time-domain model with measurements, we can simulate the expected performance of a longitudinal feedback system using proposed kicker and power amplifier parameters. Assuming 700Ω kicker shunt impedance driven by a 30 W power amplifier we get $V_{\max} = 205$ V. Figure 15 shows a simulated grow/damp measurement at 250 mA. Instability growth

3.5 Extrapolation

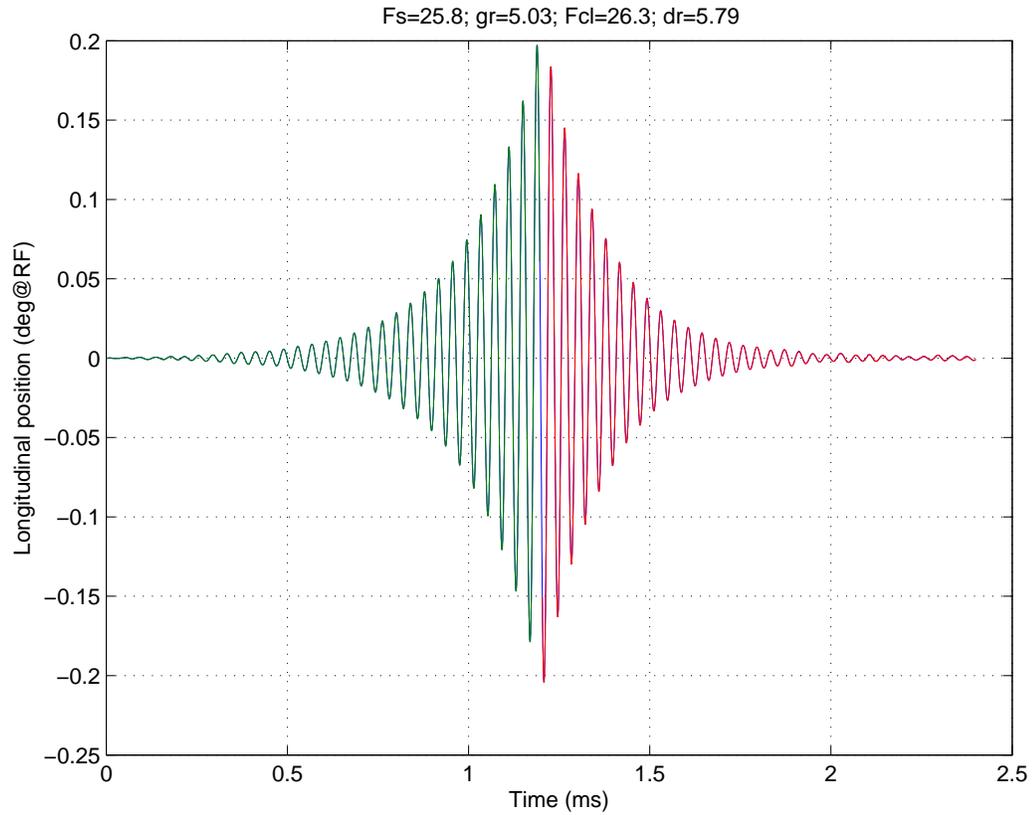


Figure 15: Simulated grow/damp transient at 250 mA, kick within the linear range.

rate was set at 5 ms^{-1} — 25% higher than the measured rates shown in Fig. 5. In this measurement the feedback gain was set to achieve roughly the same closed-loop damping rate as the open-loop growth rate. This is a normally recommended setting for robust system operation. Feedback closure time was set so that kicker output was just at the edge of saturation. Thus, the feedback system in this configuration can handle transients of up to 0.2 degrees without saturation.

Figure 16 shows simulation result with a longer open-loop period. Here the feedback is pushed deep into saturation — in fact the damping rate is barely negative when the loop is first closed around 1.4 ms. Therefore, this configuration will remain stable with input transients under 0.5 degrees.

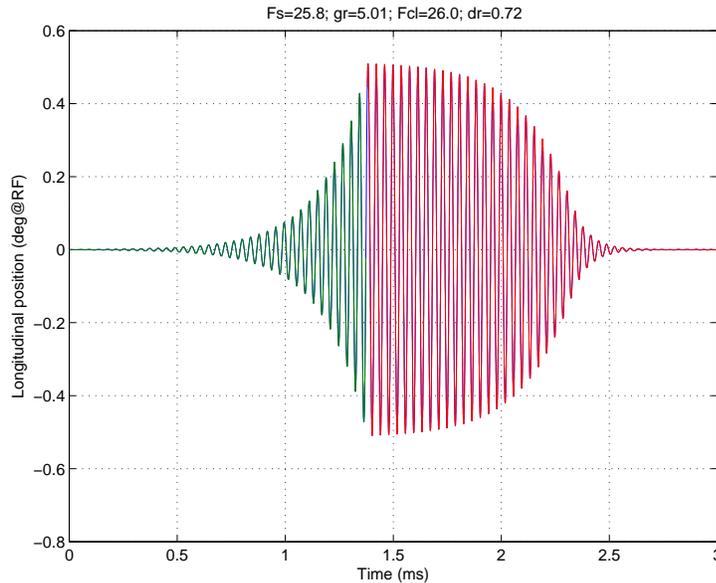


Figure 16: Simulated grow/damp transient at 250 mA, feedback output far in saturation.

Anything above that level will cause the feedback system to lose control.

Next, a slightly different configuration was investigated. The same kicker was driven by two 30 W power amplifiers, producing $V_{\max} = 290$ V. Figure 17 shows a grow/damp transient in high saturation. Here the maximum oscillation amplitude is around 0.7 degrees — consistent with the expected $\sqrt{2}$ increase.

Finally, a configuration with a 500 W power amplifier has been simulated, shown in Figure 18. As expected, more than an 8-fold increase in power produces less than a factor of 3 improvement in maximum acceptable transient.

Observed steady-state residual motion at 250 mA is around 0.02 degrees. Both single and dual power amplifier configurations have good headroom before stability is compromised. Simulation results for various power amplifier configurations are summarized in Table 3.

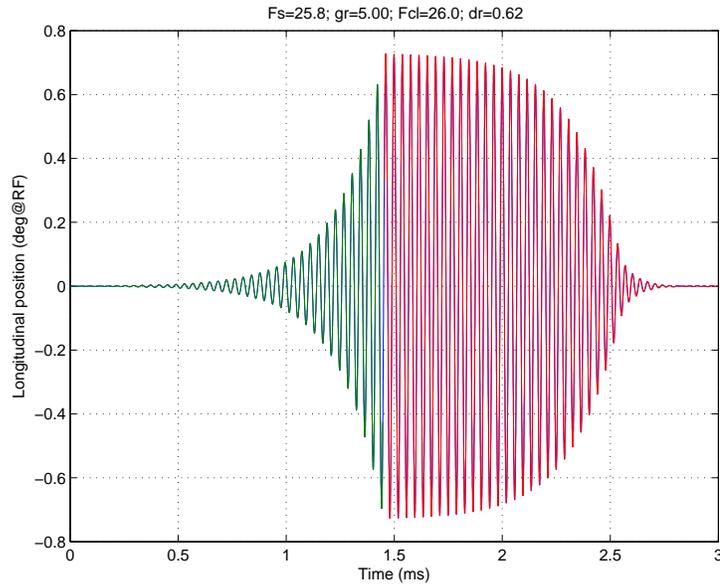


Figure 17: Simulated grow/damp transient at 250 mA with two 30 W power amplifiers.

Table 3: Amplifier power and acceptable transient amplitude

Output power, W	Maximum transient, deg@RF
30	0.51
60	0.73
500	2.06

4 Summary

Operation of the iGp12-148F bunch-by-bunch feedback processor and the FEP-476 front-end has been successfully demonstrated at LCLS UVX. Longitudinal coupled-bunch instabilities have been characterized both below and above the instability threshold. Bunch-by-bunch diagnostic data has also been used to analyze the transverse modal patterns above the respective instability thresholds.

Measured instability growth rates and closed-loop steady-state spectra have been used to configure an off-line model of the beam and the feedback

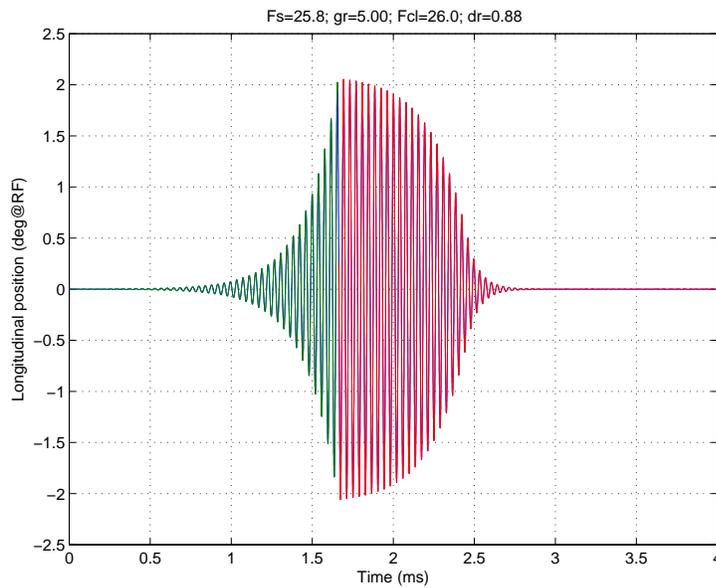


Figure 18: Simulated grow/damp transient at 250 mA with 500 W power amplifier.

system. The model has then been used to estimate power amplifier and kicker requirements for production running of the accelerator under feedback control.

References

- [1] <http://www.dimtel.com/products/igp12>.
- [2] S. Prabhakar, *New diagnostics and cures for coupled bunch instabilities*. PhD thesis, Stanford University, 2000. SLAC-R-554.
- [3] N. Abreu, O. Bagnato, R. Farias, M. Ferreira, C. Pardine, *et al.*, “Conceptual design of a 3rd harmonic cavity system for the LNLS electron storage ring,” in *EPAC’06: Proceedings*, pp. 1316–1318, 2006.
- [4] N. Abreu, R. Farias, L. Jahnel, L. Liu, C. Pardine, *et al.*, “Upgrade and Commissioning of the LNLS RF System,” in *9th European Particle Accelerator Conference (EPAC 2004)*, pp. 950–952, European Physical Society Accelerator Group, 2004.
- [5] N. Abreu, N. Abreu, R. Farias, and P. Tavares, “RF phase modulation at the LNLS electron storage ring,” in *EPAC’06: Proceedings*, pp. 1905–1907, 2006.
- [6] S. Prabhakar *et al.*, “Observation and modal analysis of coupled-bunch longitudinal instabilities via a digital feedback control system,” *Part. Accel.*, vol. 57, p. 175, 1997.
- [7] D. Teytelman, *Architectures and algorithms for control and diagnostics of coupled-bunch instabilities in circular accelerators*. PhD thesis, Stanford University, 2003. SLAC-R-633.