INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 6 Multivariable and Plantwide Control

CHAPTER 6

Orbit Feedback for a Storage Ring Whose Weak Modes Amplify Noise

Pages 138–160

Edited by John D. Hedengren

How to cite this chapter. "Orbit Feedback for a Storage Ring Whose Weak Modes Amplify Noise," in *Multivariable and Plantwide Control*, John D. Hedengren, Ed., vol. 6 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 6, pp. 138–160.

The page numbers in this file are those of the printed volume.

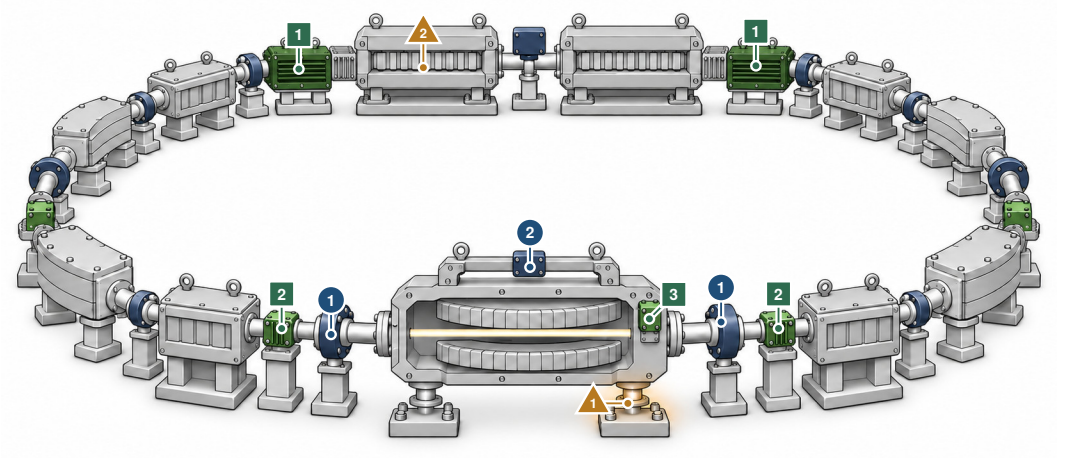
The app for this chapter
https://apmonitor.com/apps/synchrotron-orbit/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v06-multivariable-plantwide/#synchrotron-orbit

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved. Produced by Advanced Process Solutions, LLC.

CHAPTER 6

Orbit Feedback for a Storage Ring Whose Weak Modes Amplify Noise



The controller moves 1 slow correctors 2 fast correctors 3 device end coils (gap feedforward)

The controller reads 1 beam position monitors 2 gap readback

What disturbs it 1 girder drift and vibration 2 insertion device kick as its gap closes

Ninety-six monitors and two corrector sets hold the beam round the whole ring, while the light is made inside the insertion devices, where no monitor sits. Illustration, not to scale.

IN PLAIN TERMS Ninety-Six Sensors Steering One Beam

The main idea. A storage ring keeps a beam of electrons on one closed path so that it can give off bright light for experiments. Small steering magnets around the ring push the beam back wherever ninety-six position sensors see it drift.

The challenge. Some steering patterns move the beam a lot, but others barely move it, so correcting those takes large magnet swings for tiny gains. Sensor noise in those weak patterns turns into magnets that never sit still.

Why it matters. Choosing which patterns to correct decides how steady the light is and how quietly the magnets work.

THREE WAYS TO READ THIS CHAPTER

Explore	13 min	the first page, the two histories in Section 6.2, the fact in the Introduction, the try-its and the fact of Section 6.6 and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 6.6), Technology in Practice (Section 6.8), Limitations (Section 6.9) and the two open challenges

The control question A storage ring's 96 by 96 orbit response is known exactly, yet its singular values span a factor of 405, so a full inverse turns monitor noise into magnet motion. How many directions should the feedback correct, and what must a fast corrector set, a feedforward and a monitor check add?

Learning objectives

- Compute the optics of a thin-lens FODO lattice and the closed-orbit response of one corrector from its β and the tune.
- Explain why the full inverse amplifies monitor noise, and compute the filter factors of a truncated and a Tikhonov-regularized inverse.
- Place the truncation on the measured trade between residual orbit and corrector jitter, and say why it falls at the spectrum's cliff.
- Distinguish what feedback, offloading, a gap feedforward and a residual test each contribute to the graded challenge.
- Show from the loop-gain eigenvalues why a corrector wired backward defeats the full inverse at every bandwidth.

Prerequisites Sections 1.4.7 (singular values and directionality), 1.4.6 (a static inverse of the gain matrix), 1.6.1 (engine and seed) and 1.6.4 (scenarios and data).

The app https://apmonitor.com/apps/synchrotron-orbit/

See also Volume 4, Chapter 10 and Volume 4, Chapter 9 for mid-ranging and valve-position control, a fast handle kept centered by a slow one. Volume 5, Chapter 11 for a controlled temperature that no instrument reaches, as no monitor sits at a source point, and Volume 5, Chapter 12 for a sensor whose reading cannot be taken at face value.

6.1 Introduction

One corrector magnet in a storage ring, asked for a kick of 5 µrad, moves the electron orbit at every monitor in the ring. In the teaching ring of this chapter, corrector 37 sits at a focusing quadrupole, and its kick moves the closed orbit by up to 74.8 µm at monitor 85, far round the ring, and by 40.0 µm rms over all 96 monitors. The same kick at corrector 38, beside a defocusing quadrupole, moves it at most 24.8 µm. Each column of the 96 by 96 response matrix is a pattern of this kind, a cosine whose sign changes 27 times round the 96 monitors, and orbit feedback corrects the beam by inverting the whole matrix [1].

The simulator is built against a natural belief, that a plant model known to machine precision should be inverted exactly. The response of this ring is exact by construction, and its singular values, the gains along the strong and weak directions of Section 1.4.7, still fall from 379 to 0.94 µm/µrad. A controller that divides by all of them turns monitor noise into magnet motion. On the noise of case 2, 1 µm per sample, the full inverse makes the slow correctors jitter 0.198 µrad about their own mean, twelve times the 0.0164 µrad of an inverse truncated to the 48 strongest modes. Truncation and Tikhonov regularization, the classical answers to an ill-posed inverse [2], trade that noise against the part of the orbit the weak modes would have corrected, and the ring shows the trade one mode at a time.

> **DID YOU KNOW? Corrections by the Thousand Each Second**
>
> One light source's orbit feedback, in routine operation for more than 20 years, computes corrections 1.6 thousand times a second from 160 monitors to 38 correctors in each plane. The system planned for the ring's upgrade computes them 22.6 thousand times a second, from 560 monitors to 160 correctors [3].

The same shape appears wherever many actuators act on one field through a matrix whose directions differ greatly in gain. A high-purity distillation column is the classic ill-conditioned plant of process control, where the physics of its poor conditioning decides how much performance a robust controller can keep under actuator uncertainty [4], and the cross-direction profile of a paper

machine is a spatially distributed plant whose controllers are designed in spatial as well as temporal frequency [5]. Two more features of orbit feedback are general. Two actuator sets of different speed and strength share one controlled field, which is the mid-ranging problem of process plants [6], and a failed sensor must be found among many by the residual its readings leave.

The orbit feedback literature is written for accelerator physicists working on real machines, from the decomposition of a measured response [7] to controllers designed per singular mode [8, 9], and a control student rarely meets a hundred-square plant whose modes can be switched on and off one at a time. This chapter traces what an inverse, its truncation or regularization, two corrector sets, a gap feedforward and a residual test each do to the true orbit at the photon source points of a declared storage ring simulated in a browser. The closed-orbit response, the singular value decomposition, Tikhonov regularization, internal model control and mid-ranging are established methods. Every result is a measurement made on this one ring, and two of them overturn the plan the simulator was built from. No lattice error the simulator allows makes the full inverse diverge, so its infeasible case is a corrector wired backward, and the regularization in the graded challenge is not what lets it pass.

The chapter is organized as follows. Section 6.2 follows orbit correction from strong focusing to feedback in mode space. Section 6.3 states the ring: its optics, a closed-form response that agrees with independent tracking at all 11,520 entries, the correctors, the disturbances and what the monitors report. Section 6.4 splits the response into strong and weak directions and shows where the disturbances lie among them. Section 6.5 describes the controllers: a one-sided Jacobi decomposition, a per-mode controller whose closed loop is derived and then measured, the fast set and offloading, the gap feedforward and the residual test, behind a line between measurement and truth that no controller crosses. Section 6.6 measures the ten cases, each credited in the app by a completion rule, and Section 6.8 describes the real monitors, correctors and controllers. Sections 6.7, 6.9 and 6.10 place the ring among the volume's plants, state its limits and summarize. The exercises follow in Section 6.11.

6.2 From Strong Focusing to Feedback in Mode Space

Storage rings focus their beams by alternating gradients. The strong-focusing principle of 1952 showed that a train of lenses that alternately converge and diverge, with equal strengths, focuses on balance, and that the oscillation amplitude and the magnet aperture of a synchrotron shrink with it [10]. The theory that followed wrote the transverse motion with an amplitude function and a betatron phase [11], the β and φ in which every closed-orbit calculation is still written. A dipole error anywhere in such a lattice displaces the closed orbit round the whole ring, by an amount set by β at the error and at the observer and by the tune, which is the response this simulator computes [1].

The first history segment follows the 1993 decomposition of the response matrix into channels. The numerical tools were older. A stable algorithm for the singular values and the pseudo-inverse was published in 1965 [12], and the one-sided Jacobi method, which never forms the product of the matrix with its transpose, was later shown to compute the singular values with high relative accuracy [13]. Because a model's response is only as good as the model, orbit responses measured on the machine were then used to determine a storage ring's optics [14].

> ⧖ HISTORY **1993: The Orbit Split Into Channels**
>
> In 1993 the correction of a storage ring's whole orbit was written as a singular value decomposition of its response matrix, and its use was simulated for a third-generation light source that was to carry 320 corrector magnets and 360 beam position monitors. The decomposition turns the monitors and the correctors into the same number of transformed monitors and transformed correctors. Each transformed monitor is coupled to at most one transformed corrector, with a strength that sets how efficiently that channel corrects, and the channels do not interact. Decoupling the weakly coupled pairs makes the correction more efficient at the expense of accuracy, and it lets the corrector strengths be kept within what their power supplies can deliver [7].

Table 6.1 Variables by Role, With Their Values When a Case Starts

Variable	Role	Start	Range
x_d, x'_d	CV, source point d (truth)	0	±3 µm, ±1 µrad
y (rms)	CV, true orbit at the 96 BPM places	0	≤ 2 µm (challenge)
y_m	measurement, 96 BPM readings at 1 kHz	0	noise σ_n
g_d	measurement, gap readbacks	40 mm	10–40 mm
$\theta_{s,cmd}$	MV, 96 slow commands	0	±50 µrad (5 A)
$\theta_{f,cmd}$	MV, 24 fast commands	0	±5 µrad (5 A)
θ_{ff}	MV, end coils of the devices	0	from the table
θ_s, θ_f	states, kicks after the lags	0	—
δ_g	DV, girder drift (24)	0	0–0.41 µm/$\sqrt{\text{s}}$
v_g	DV, girder vibration	0	0–2 µm, 10–50 Hz
r	DV, dipole ripple at 50 Hz	0	0–0.2 µrad
κ_d	DV, device kicks	0	to 1.6, −2.2, 1.9 µrad
fault	DV, BPM 37	none	+60 µm stuck, or noisy
ν, ε_k	lattice tune, quadrupole error	14.2, 0	14.05–14.35, 0–2 %

The users set the target. The stability an experiment needs depends on how it samples, integrates and scans, over time scales from microseconds to hours, and at third-generation rings it can come to micron and submicroradian stability of the electron orbit [15]. The second history segment follows fast orbit feedback to a controller per singular mode. Internal model control came from the process industries, where its structure was reviewed and unified in 1982 [16]. Two actuator arrays acting on one beam were decomposed by their controllable subspaces, with mid-ranging where the arrays' control directions align [9], an internal-model form of the mid-ranging that process plants use when a quick, small handle and a slow, large one share a single controlled variable [6, 17]. Rate limits on the corrector supplies brought an anti-windup problem of their own [18].

> ⧖ **HISTORY** **2013: A Controller for Every Mode**
>
> Fast orbit feedback splits its work into space and time. The spatial part inverts a response matrix that is normally ill-conditioned, so the less significant singular values are removed or filtered, and the usual dynamics are proportional-integral. Most systems apply the same controller to every mode with a gain per mode, because separate dynamics per mode cost two more matrix products, too heavy for most storage-ring systems with more than a hundred actuators and update rates above a few kilohertz. A study of 2013 gave each mode its own internal-model controller, chosen from the disturbance spectrum in that mode, and ran it on a booster synchrotron of 22 sensors and 22 actuators operated with stored beam, the test bed for a storage ring with 172 of each [8].

6.3 A Declared Ring of 96 Monitors and 120 Correctors

The simulated ring is a declared teaching lattice, not a model of any facility, as shown in Figure 6.1. It has 24 cells of 20 m, 480 m round, each carrying two periods of a thin-lens FODO lattice: a focusing quadrupole (QF), a drift of $L = 5$ m, a defocusing quadrupole (QD) and another drift. A beam-position monitor (BPM) sits 0.4 m downstream of each of the 96 quadrupoles and a slow corrector at each quadrupole. One fast corrector sits at 12.5 m in each cell, two dipoles at 2.5 and 7.5 m, and in cells 3, 11 and 19 an insertion device occupies the middle of the straight from 15 to 20 m, where its photon source point lies. Only the horizontal plane is modeled. The role of each variable is given in Table 6.1, and the Ring view of the simulator, which draws the same ring live, is shown in Figure 6.2.

6.3.1 Optics of the Teaching Lattice

For thin quadrupoles of focal length $\pm f$ separated by drifts L, the phase advance μ of one FODO period and the extreme values of the amplitude function are given by Equation (6.1).

$$\sin\frac{\mu}{2} = \frac{L}{2f}, \qquad \mu = \frac{2\pi\nu}{48}, \qquad \beta_{\pm} = \frac{2L\,(1 \pm \sin(\mu/2))}{\sin\mu} \tag{6.1}$$

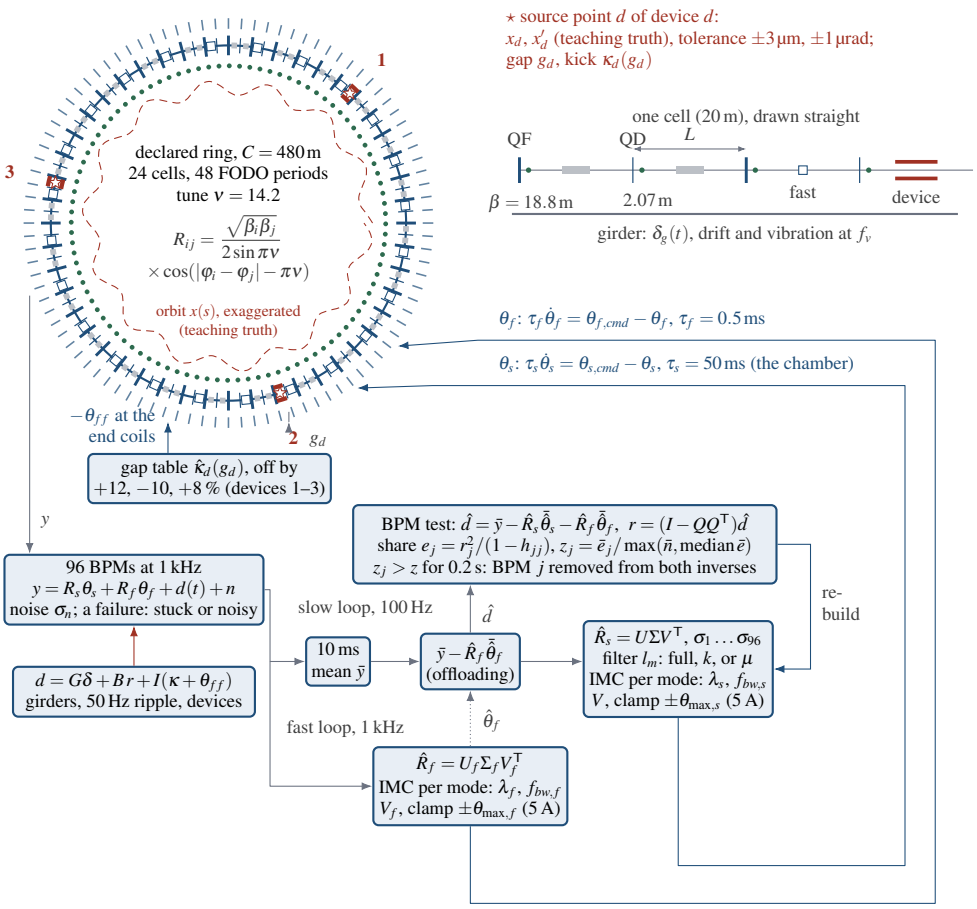


Figure 6.1 The declared ring as modeled, with what its controllers read and drive. The ring carries 24 cells (one is drawn straight at the right) with focusing quadrupoles thick and defocusing ones thin, dipoles shaded, the 96 BPMs as dots inside, the 96 slow correctors as bars outside and the 24 fast correctors as open boxes. The insertion devices 1 to 3 kick the beam by $\kappa_d(g_d)$ at their source points (stars). The dashed orbit and the source points' x_d and x'_d are teaching truth that no controller reads. The controllers see the BPM readings y and the gap readbacks g_d. The slow loop at 100 Hz and the fast loop at 1 kHz each invert their own response in mode space, offloading removes the fast set's share from the slow loop's input, the slow loop's disturbance estimate $\hat{d}$ feeds the BPM test, and the gap table feeds $-\theta_{ff}$ forward. Not to scale.

where ν is the horizontal tune, the number of betatron oscillations in one turn, and 48 is the number of periods. At $\nu = 14.2$ the period advances 106.5°, $f = 3.12$ m, and β is 18.79 m at a QF and 2.07 m at a QD. One period per cell would need 213°, past the 180° at which a FODO period stops being stable, which is why each cell carries two. The simulator does not use Equation (6.1) directly. It

builds the ring from transfer matrices and carries the optics round it, as shown in Equation (6.2).

$$\begin{aligned}
&\cos\mu_t = \tfrac{1}{2}(M_{11}+M_{22}), \quad \beta_0 = \frac{M_{12}}{\sin\mu_t} \\
&\alpha_0 = \frac{M_{11}-M_{22}}{2\sin\mu_t} \\
&\text{drift } \ell: \ \beta \leftarrow \beta - 2\alpha\ell + \gamma\ell^2, \quad \alpha \leftarrow \alpha - \gamma\ell \\
&\qquad\qquad \varphi \leftarrow \varphi + \arctan\frac{\ell}{\beta - \alpha\ell} \\
&\text{quadrupole: } \alpha \leftarrow \alpha \pm \beta/f, \quad \gamma = \frac{1+\alpha^2}{\beta}
\end{aligned} \tag{6.2}$$

where M is the one-turn product of the drift matrices $\left[\begin{smallmatrix} 1 & \ell \\ 0 & 1 \end{smallmatrix}\right]$ and the thin-lens matrices $\left[\begin{smallmatrix} 1 & 0 \\ \mp 1/f & 1 \end{smallmatrix}\right]$, taken from just before the first QF. The integrated phase gives $\nu = 14.200000$ and $\cos\mu_t = 0.309017$. A BPM 0.4 m after a QF sees $\beta = 16.46\,\text{m}$, one after a QD 2.42 m, and a fast corrector and a source point both 7.08 m. With the lattice error on, each quadrupole strength $1/f$ is multiplied by $1+\varepsilon_k g_q$, with g_q a fixed seeded unit normal, and the optics are carried round again without assuming periodicity.

6.3.2 The Closed-Orbit Response

A thin kick θ_j at a place with optics (β_j, φ_j) moves the closed orbit at a place i and changes its angle, as shown in Equation (6.3).

$$\begin{aligned}
R_{ij} &= \frac{\sqrt{\beta_i\beta_j}}{2\sin\pi\nu}\cos\Delta_{ij}, \qquad \Delta_{ij} = |\varphi_i - \varphi_j| - \pi\nu \\
R'_{ij} &= \frac{\sqrt{\beta_j/\beta_i}}{2\sin\pi\nu}\big[-\alpha_i\cos\Delta_{ij} - \operatorname{sgn}(\varphi_i-\varphi_j)\sin\Delta_{ij}\big]
\end{aligned} \tag{6.3}$$

where R is in m/rad, which is µm per µrad, and R' in µrad per µrad. The first line is the closed-orbit response of Equation 15.47 in Wiedemann's text [1]. From it the simulator builds R_s (96 BPMs by 96 slow correctors), R_f (96 by 24), the rows of the three source points' positions and angles, and three disturbance matrices. A girder displaced by δ_g moves its four quadrupoles, which kick the beam by $\pm\delta_g/f$, so the column c of G sums R_{iq}/f over the cell's QFs and $-R_{iq}/f$ over its QDs. The vector B sums the response to a unit kick at every one of the 48 dipoles, and I holds the response to a kick at each device's middle.

6.3.3 Correctors, Disturbances and the True Orbit

The orbit is algebraic in the kicks, because the beam settles on its closed orbit within a few hundred turns, a few microseconds. The kicks follow their commands through first-order lags, the slow ones through the eddy currents of the vacuum chamber they sit over, as shown in Equation (6.4).

$$\begin{aligned}
y &= R_s\theta_s + R_f\theta_f + d \\
d &= G\delta + B\,r + I(\kappa + \theta_{ff}) \\
\theta(t+\Delta t) &= \theta_{app} + \big(\theta(t) - \theta_{app}\big)\,\mathrm{e}^{-\Delta t/\tau} \\
\theta_{app} &= \operatorname{clip}(\theta_{cmd}, \pm\theta_{\max})
\end{aligned} \tag{6.4}$$

where τ is $\tau_s = 50\,\text{ms}$ for the slow set and $\tau_f = 0.5\,\text{ms}$ for the fast one. The commands are held between control executions, so the update is exact for any step. The disturbances run on their own

Table 6.2 Parameters and Defaults. A Is Adjustable in the Parameter Panel, H Hidden From It, S a Control or Disturbance Slider, F Fixed. All Are Teaching Values

Symbol	Meaning	Default (range)	
ν	horizontal tune	14.2 (14.05–14.35)	A
ε_k	rms quadrupole error	0 (0–2 %)	A
L, f	drift, focal length (derived)	5 m, 3.12 m	F
τ_s, τ_f	slow and fast coil lags	50 (10–200), 0.5 (0.1–2) ms	A
$\theta_{\max,s}$	slow limit, 5 A	50 (10–200) µrad	A
$\theta_{\max,f}$	fast limit, 5 A	5 (1–20) µrad	A
σ_n	BPM noise per 1 ms sample	0.5 (0–5) µm	A
κ_s	scale of the device kicks	1 (0–3)	H
q	ground drift rate	0 (0–0.41) µm/$\sqrt{\text{s}}$	S
A_v, f_v	girder vibration	0 µm (0–2), 10 Hz (10–50)	S
r_0	dipole ripple	0 (0–0.2) µrad	S
k, μ	modes kept, Tikhonov weight	48 (1–96), 10 (0.1–1000)	S
$f_{bw,s}$, $f_{bw,f}$	loop bandwidths	1 (0.1–10), 100 (10–300) Hz	S
z	BPM test threshold	100 (10–1000)	S
δ_m, N_m	response measurement	10 µrad, 200	F
	gap table error, devices 1–3	+12, −10, +8 %	F

clocks so that a run does not depend on the physics step, as shown in Equation (6.5).

$$\begin{aligned} w_g &\leftarrow w_g + q\sqrt{\Delta t_d}\,\varepsilon_g, \qquad \delta = Kw + \delta_{step} \\ v_g(t) &= A_v\,\mathrm{Re}\{p_g\,z(t)\,\mathrm{e}^{\mathrm{i}2\pi f_v t}\}, \qquad r(t) = r_0\sin(2\pi 50 t) \\ \kappa_d(g) &= \kappa_{\max,d}\left(\frac{g_{open}-g}{g_{open}-g_{\min}}\right)^2 \end{aligned} \tag{6.5}$$

where each of the 24 girders walks every $\Delta t_d = 10$ ms with rate q in µm/$\sqrt{\text{s}}$, and K smooths the walk along the ring with a row-normalized Gaussian of 3 cells, so that each girder's rms after a time T is $q\sqrt{T}$. The vibration has a fixed complex girder pattern p_g, strongest round cells 7 to 9, and an envelope $z = 1 + \zeta_1 + \mathrm{i}\zeta_2$ whose two parts are first-order colored noise of time constant 1 s and standard deviation 0.35, updated every 1 ms, which makes a narrow line at f_v. A burst multiplies A_v by 4 for 10 s. A ground step moves girders 6 to 9 by 3 µm. Each gap moves toward its target at 5 mm/s between $g_{open} = 40$ mm and $g_{\min} = 10$ mm, and $\kappa_{\max,d}$ is 1.6, −2.2 and 1.9 µrad for devices 1 to 3. The position and angle at each source point are the source-point rows applied to the same kicks and disturbances.

6.3.4 What the Monitors Report

Each BPM is read every 1 ms as the true orbit plus white noise of $\sigma_n = 0.5$ µm per sample (1 µm in case 2), drawn from the case's seed. A failed BPM 37 either jumps 60 µm from its last reading and freezes, or reads the true orbit plus 30 µm rms of noise. The gap readbacks are exact. The correctors' currents are the controllers' own commands, and the source points carry no monitor. The alternative to the model's response is a measured one, obtained as Equation (6.6).

$$\tilde{R}_{ij} = R_{ij} + \frac{\sigma_n}{\delta_m\sqrt{2N_m}}\,\varepsilon_{ij} \tag{6.6}$$

where each corrector would be kicked by $\pm\delta_m = 10$ µrad and the BPMs averaged over $N_m = 200$ samples each way, an element error of 0.0025 µm/µrad. The simulator computes it at once from the ring's true response and that noise, as a measurement on a quiet beam would read, and does not spend the 48 s of beam time a real one would take.

6.3.5 Clocks, Starting State and Numerics

The BPMs are read and the fast loop runs every 1 ms, the physics step, and the slow loop runs every 10 ms on the mean of the ten readings since its last update. A log row is written every 10 ms with

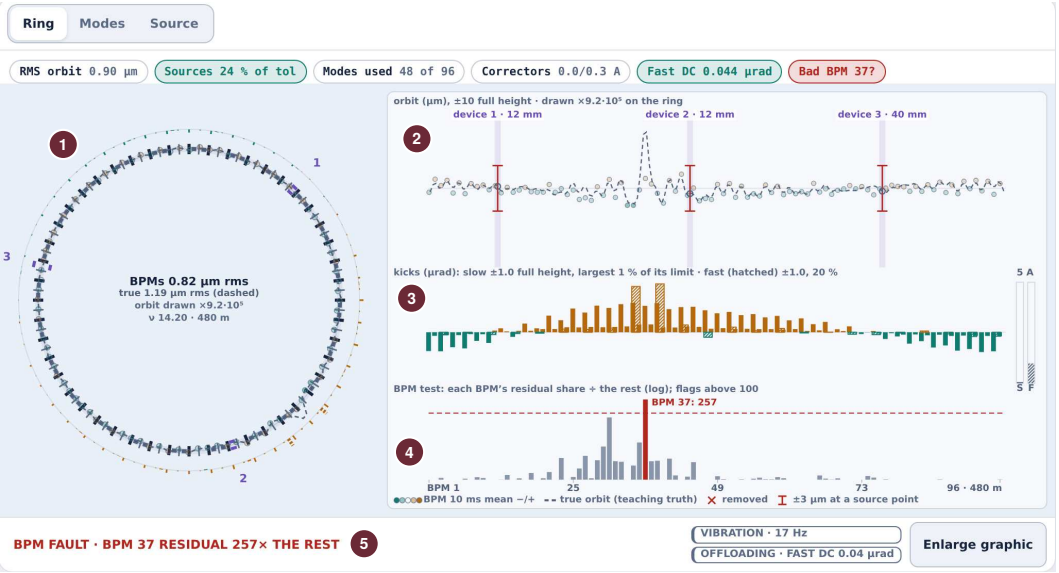


Figure 6.2 The challenge at $t = 300.19$ s on the Ring view with the suggested settings, 0.19 s after BPM 37 turned to noise and 0.02 s before the test removes it, cropped to the plant view. (1) The ring with its BPMs and correctors, the true orbit at 1.19 µm rms in its middle. (2) The unrolled strip with devices 1 and 2 closed to 12 mm and device 3 open at 40 mm, the BPM readings and the true orbit dashed. (3) The kick strip, with the fast set at 20 % of its limit and the slow set at 1 %. (4) The BPM test's strip, which points at BPM 37 with 257 times the median share against a threshold of 100. (5) The status line, which reads the fault. The challenge's results are in Table 6.5.

the minimum and maximum of the source-point positions over its ten steps, so that a 10 to 50 Hz vibration survives the log. Every case starts on the golden orbit, zero at every BPM, with every kick at zero except where the case sets one, and the case's seed drives every noise stream (Section 1.6.1). The parameters are listed in Table 6.2. The scripted events (a gap change, a burst, the BPM failure, the ground step, the lattice error) keep the engine's order of Figure 1.10. Each fires once, at the start of the first step at or after its time, and the controllers then run before the exact update.

6.3.6 Simplifications

The quadrupoles are thin lenses, the dipoles' weak focusing is left out of the optics, and only the horizontal plane is modeled, without coupling, dispersion from the correctors or energy change, so the model says nothing about vertical stability or beam size. Each corrector is a linear kick behind one first-order lag, whereas a real chamber's eddy currents have several time constants and real supplies have slew and rate limits [18]. BPM noise is white and independent, with no gain error or drift of its own. The tolerances at the source points (± 3 µm, ± 1 µrad) and of the orbit (2 µm rms) are teaching values, while a real light source derives its own from its experiments' needs [15]. The lattice error and the tune are bounded together. The seeded error raises the tune by 0.067 at 2 %, and beyond the slider ranges the true tune would reach the half-integer, where these optics have no stable solution, so that resonance stays out of reach.

6.4 Strong and Weak Directions of the Response

The singular value decomposition of Section 1.4.7, $R_s = U\Sigma V^{\mathsf{T}}$, writes the response as 96 independent modes. Column m of V is a pattern of corrector kicks, column m of U the orbit pattern it makes, and σ_m its gain, in µm per µrad. For this ring $\sigma_1 = 379.17$ and $\sigma_{96} = 0.9357$, a condition number of

405.2, and the 96 values fall in two groups of 48 with a cliff between $\sigma_{48} = 8.740$ and $\sigma_{49} = 1.282$, as shown in Figure 6.3a. The cliff comes from the β pattern. Modes 1 to 48 put a mean 89 % of their corrector weight on the 48 correctors at the QFs ($\beta = 18.8$ m), and modes 49 to 96 the same share on those at the QDs ($\beta = 2.07$ m). With every $\sqrt{\beta}$ replaced by the ring's mean and the phases kept, the ratio at modes 48 and 49 falls from 6.82 to 1.29 and the condition number to 192.8. Within each group the gain falls with the distance of the mode's betatron harmonic from the tune. Modes 1 to 6 are harmonics 13 to 15, modes 45 to 48 harmonics 0 to 2, and modes 91 to 96 harmonics 45 to 48. The 24 fast correctors' response R_f runs from 157 to 5.84 µm/µrad.

The disturbances do not spread evenly over these modes, and that decides what an inverse can buy. Projected on the orbit modes U, the girder drift as the plant makes it, smoothed along the ring, puts 86.0 % of its power in modes 45 to 48, the lowest harmonics of the strong group, 6.4 % in modes 1 to 44, 0.34 % in modes 49 to 92, and 7.3 % in the four weakest, modes 93 to 96. An inverse truncated anywhere from 48 to 92 modes therefore removes the same 92 to 93 percent of the drift's power, and the rest lies in the directions where each micrometer of orbit costs about a microradian of kick.

6.5 Feedback in Mode Space as Implemented

Every controller reads only a measurement packet: the 96 BPM readings with the BPM system's declared noise rating, the three gap readbacks, and its own corrector commands. A test of the application fails if a controller file reads the plant, the truth, the lattice error, the failed BPM's identity, the device kick scale or the event schedule, and a second test replays 320,000 recorded packets into a controller with no plant behind it and obtains every command again.

6.5.1 Filter Factors on One Decomposition

The controllers design with the model response $\hat{R}_s = U\Sigma V^{\mathsf{T}}$ of the declared lattice at the present tune (or the measured $\tilde{R}$), restricted to the BPMs in use, decomposed by the one-sided Jacobi method [13]. A pair of columns is rotated when $|w_p \cdot w_q| > 10^{-15}\|w_p\|\|w_q\|$, and the sweeps stop when one rotates nothing, which takes 14 sweeps for R_s and about 12 ms in a server-side JavaScript runtime. The inverse in force is a filter factor on each mode, for the full inverse, a truncation at k modes and Tikhonov regularization in turn, as shown in Equation (6.7).

$$l_m = 1, \qquad l_m = [m \le k], \qquad l_m = \frac{\sigma_m^2}{\sigma_m^2 + \mu} \tag{6.7}$$

where μ is in $(\mu\text{m}/\mu\text{rad})^2$, so a mode with $\sigma = \sqrt{\mu}$ is corrected halfway, which is the form of the mode-space paper [8] and of Tikhonov regularization in the filter-factor form [2]. At $\mu = 10$ the strongest modes are corrected by 1.00 of the way, mode 48 by 0.88, mode 49 by 0.14 and the weakest by 0.08.

6.5.2 Internal Model Control Per Mode

Each mode is a sampled first-order plant. A command held over a period T reaches the kick through the lag, with $a = \mathrm{e}^{-T/\tau}$, and the slow loop reads the mean of $n = 10$ fast samples, so its mode seen at the BPMs is $\sigma_m z^{-1}(b_0 + b_1 z^{-1})/(1 - a z^{-1})$ with $b_0 = 1 - c$, $b_1 = c - a$, and c the mean of a_f^j over $j = 1$ to 10, $a_f = \mathrm{e}^{-1\,\text{ms}/\tau}$. The fast loop reads the latest sample, so $b_0 = 1 - a$ and $b_1 = 0$. The controller is internal model control in each mode [8, 16], as shown in Equation (6.8).

$$\begin{aligned} \hat{d} &= \bar{y} - \hat{R}_s \bar{\hat{\theta}}_s, \qquad v_{m,k} = -u_m^{\mathsf{T}} \hat{d}_k \\ \bar{u}_{m,k} &= \lambda\, \bar{u}_{m,k-1} + K_m\,(v_{m,k} - a\, v_{m,k-1}) \\ \theta_{cmd} &= \textstyle\sum_m v'_m\, \bar{u}_m \\ K_m &= \frac{l_m(1-\lambda)}{\sigma_m(b_0 + b_1)}, \qquad \lambda = \mathrm{e}^{-2\pi f_{bw} T} \end{aligned} \tag{6.8}$$

where u_m and v'_m are the columns of U and V, $\bar{y}$ is the mean BPM reading over the period, and $\bar{\hat{\theta}}_s$ is the controller's own model of the kicks over the same period, the applied (clamped) commands passed through the lag in closed form. Because the model is driven by the clamped commands, a corrector held at its limit winds nothing up. With the model right and $l_m = 1$ the mode's closed loop is $(1-\lambda)z^{-1}\beta(z^{-1})/(1-\lambda z^{-1})$, with $\beta = (b_0 + b_1 z^{-1})/(b_0 + b_1)$, a first-order response at the declared bandwidth whatever τ_s, and with $l_m < 1$ the mode settles on the fraction l_m of the full correction. The step decay of one simulated mode gives 0.5000, 1.0002 and 3.0002 Hz for slow bandwidths of 0.5, 1 and 3 Hz, and 50.000, 100.000 and 200.000 Hz for the fast loop. A mode with $l_m = 0$ keeps its command where it is. When a BPM is removed both decompositions are rebuilt on the remaining rows, the controllers' states are projected onto the new modes, and the part of each command that no remaining BPM can see decays at the loop's own rate λ.

The noise cost follows from the same loop. White BPM noise of σ_n per sample, averaged over n samples, moves the kicks with the variance of Equation (6.9).

$$\begin{aligned} \sum_j \operatorname{var}\theta_j &= \frac{\sigma_n^2}{n}\,\frac{1-\lambda}{1+\lambda}\sum_m \frac{l_m^2}{\sigma_m^2} \\ \frac{\operatorname{rms}\theta_{cmd}}{\operatorname{rms}\theta} &= \frac{\sqrt{1+a^2-2a\lambda}}{1-a} \end{aligned} \tag{6.9}$$

The commands lead the lag, so they jitter 2.01 times as much as the kicks at a slow bandwidth of 1 Hz, 3.09 times at 3 Hz and 4.92 times at 10 Hz. The full inverse's kicks therefore exceed those of a truncation at k modes by $(\sum_{all}\sigma_m^{-2}/\sum_{m\le k}\sigma_m^{-2})^{1/2}$, a root-sum-square of $1/\sigma$, which is 16.19 at 40 modes and 12.77 at 48. A ratio of plain sums of $1/\sigma$ would give 21.1 at 40 modes.

6.5.3 The Fast Set and Offloading

The fast loop applies Equation (6.8) to $\hat{R}_f = U_f \Sigma_f V_f^{\mathrm{T}}$ every 1 ms on the latest reading, with the same kind of inverse on its own 24 modes, a default bandwidth of 100 Hz and limits of ± 5 µrad. Without offloading, each loop integrates the same error, so the two share a pole at $z = 1$ and the fast set keeps whatever steady kick it took first. Offloading, the mid-ranging of process control [6] that Volume 4, Chapter 10 applies to a power cycle, here in the internal-model form proposed for two corrector arrays [9, 17], changes only the slow loop's input, as shown in Equation (6.10).

$$\hat{d}_s = \bar{y} - \hat{R}_f\,\bar{\hat{\theta}}_f - \hat{R}_s\,\bar{\hat{\theta}}_s \tag{6.10}$$

The slow loop then corrects the orbit the beam would have without the fast set, independent of both loops' kicks, and the fast loop corrects what remains, so the two form a cascade and the fast set's steady kick drains to the slow set wherever the slow inverse reaches.

6.5.4 Gap Feedforward

Each device's end coils kick $-\hat{\kappa}_d(g_d)$ from a declared table read at the gap readback. The table is the kick law of Equation (6.5) scaled by $1+\varepsilon_d$, with $\varepsilon_d = +12$, -10 and $+8$ percent for devices 1 to 3, a declared table error. It does not read the plant's hidden kick scale.

6.5.5 A Residual Test for a Failed Monitor

At every slow update the test compares the disturbance estimate $\hat{d} = \bar{y} - \hat{R}_s\bar{\hat{\theta}}_s - \hat{R}_f\bar{\hat{\theta}}_f$ with the 28 directions the declared sources can make, the 24 girder columns of G, the dipole pattern B and the three device columns of I, all from the model lattice. A monitor that fails reads a number at one place,

which no girder, dipole or device can make. The statistics are given in Equation (6.11).

$$\begin{aligned} r &= (I - QQ^{\mathsf{T}})\hat{d}, \qquad e_j = \frac{r_j^2}{1 - h_{jj}} \\ \bar{e}_j &\leftarrow \bar{e}_j + \alpha_e (e_j - \bar{e}_j) \\ z_j &= \frac{\bar{e}_j}{\max(\bar{n},\ \mathrm{median}_i\, \bar{e}_i)} \\ \bar{n} &= \max(\sigma_n, 0.05\,\mu\mathrm{m})^2 / 10 \end{aligned} \tag{6.11}$$

where Q is an orthonormal basis of the 28 directions on the rows in use, h_{jj} the leverage of row j, and $\alpha_e = 1 - \mathrm{e}^{-10\,\mathrm{ms}/0.5\,\mathrm{s}}$. A monitor is removed from both inverses when it has the largest z_j and z_j exceeds the threshold $z = 100$ for 0.2 s. The test then starts a fresh window and cannot act for 1 s. The median guards against a model that is wrong everywhere, and the floor $\bar{n}$ keeps a start-up residual from flagging a healthy monitor when the noise rating is zero. With the test off the statistic still runs and the Ring view still shows where the residual points.

6.5.6 Hand Kicks, a Helper and Three Models

With the loops in manual the slow commands hold, and one corrector, chosen by number, takes the hand kick within $\pm 50\,\mu$rad. A helper applies one correction step from the last 10 ms mean, $\theta \leftarrow \theta - V\,\mathrm{diag}(l_m/\sigma_m)\,U^{\mathsf{T}}\bar{y}$, with 48 modes if the slow feedback is off. In automatic mode the hand station follows the command requested of the chosen corrector, and in manual the loops track the held commands, so a transfer either way is bumpless. The plant uses the true lattice, with any quadrupole error and, in case 10, one corrector's column reversed. The controllers' model is the declared lattice at the present tune, which knows neither, and the identified model is the measured response of Equation (6.6). With the model wrong, $\hat{d}$ contains $(R - \hat{R})\theta$ and each mode's loop closes through the eigenvalues η of $L(M - I)$, with $L = \mathrm{diag}(l_m)$ and $M = \Sigma^{-1}U^{\mathsf{T}}RV$, whose eigenvalues are those of $R\hat{R}^{+}$, and each eigenvalue gives a loop whose poles solve Equation (6.12).

$$z^2 - \lambda z + \eta\,(1 - \lambda)\,\frac{b_0 z + b_1}{b_0 + b_1} = 0 \tag{6.12}$$

For the fast loop ($b_1 = 0$) the root is $z = 1 - (1 - \lambda)e$, with $e = 1 + \eta$, stable for $0 < e < 2/(1 - \lambda)$. A loop gain e that is negative diverges at every bandwidth, and one far above 1 diverges only at a high one.

6.6 Ten Cases, Measured

The ten cases are listed in Table 6.3. Each runs on the case's seed, 20261030, with any setting its row names entered before Run, as Section 1.6.4 describes. Every number in this section comes from the simulator's own engine, run without a browser, and the scorecard judges the true orbit, sampled at the BPMs' 1 kHz. The orbit rms is over the 96 BPM places, the slow jitter is the rms of each slow command about its own 2 s mean, averaged over the 96, and the fast-set DC is the rms over the 24 fast correctors of each one's 1 s mean. The headline numbers are collected in Table 6.4.

6.6.1 One Kick and the Full Inverse

In case 1, 5 µrad at corrector 37 moves the orbit by up to 74.8 µm at BPM 85 and by 48.8 µm at a source point, and only 2 of the 96 BPMs read less than 1 % of the peak. The kick makes a closed-orbit distortion all the way round the ring, not a closed bump. At corrector 38, where β is 2.07 m instead of 18.8 m, the largest excursion is 3.0 times smaller, the ratio $\sqrt{18.8/2.07}$ that Equation (6.3) predicts. In case 2 the full inverse at 3 Hz holds the orbit at 0.09 µm rms while the slow commands jitter 0.198 µrad, and the noise model of Equation (6.9) predicts 0.199 µrad. Truncated to 48 modes for the next 20 s, the jitter falls 12 times, to 0.0164 µrad, against 12.77 predicted, and the orbit grows to

Table 6.3 The Ten Cases

#	Settings and events	Run	Teaching point
1	Manual, all feedback off, kick by hand	20 s	a column of R
2	Full inverse at 3 Hz, $\sigma_n = 1$ µm, drift 0.05	40 s	noise amplification
3	Truncated at 40 modes, drift 0.3	60 s	where to cut
4	Tikhonov, $\mu = 10$, drift 0.3	60 s	a filter, not a cut
5	Slow set only, 0.5 µm at 10 Hz	20 s	the slow bandwidth
6	Fast set added, ground step at 5 s	40 s	offloading
7	Device 2 closes to 12 mm at 3 s	20 s	feedforward
8	Fast set, offloading, BPM 37 stuck at 3 s	30 s	a residual that points
9	Challenge, full inverse, slow set only	600 s	the whole problem
10	Full inverse, corrector 38 reversed, drift 0.15	30 s	infeasible: a sign error

Drift rates q in $\mu\text{m}/\sqrt{\text{s}}$. Unless a row says otherwise the slow loop is truncated at 48 modes at 1 Hz, and the fast set, offloading, the feedforward and the BPM test are off.

Table 6.4 Measured Outcomes of the Cases, on the True Orbit Over Each Case's Run Unless Stated

#	Configuration	Measured
1	5 µrad at corrector 37 / 38	largest 74.8 / 24.8 µm, rms 40.0 / 13.4 µm
2	full / 48 modes, 20 s each	jitter 0.198 / 0.0164 µrad, orbit 0.09 / 0.21 µm
3	k = 20 / 40 / 48 / 96	2.17 / 2.14 / 0.61 / 0.09 µm, jitter 0.0009 / 0.0043 / 0.0278 / 0.0875 µrad
4	μ = 1 / 10 / 100	0.33 / 0.61 / 1.32 µm, jitter 0.0484 / 0.0256 / 0.0123 µrad
5	slow bandwidth 1 / 10 Hz	10 Hz line kept 105 / 111 %, in tolerance 92.8 / 89.1 %
6	fast set, offloading off / on at 10 s	line 22 %, fast DC at 40 s 0.18 / 0.0022 µrad
7	feedback / with feedforward	peak at source 2: 2.68 / 0.30 µm (×8.8)
8	test off / on at 5 s	41.8 µm at 30 s / within 2 µm 5.9 s after, 0.18 µm
10	full at 0.1–10 Hz / k 48, 60 / k 70	132.6–136.0 µm / 0.38, 0.38 / 16.28 µm
	μ 10 / μ 1 / measured $\tilde{R}$	0.38 / 25.45 / 0.05 µm

Case 2's orbit and case 8's and 10's values are the true orbit's rms through a 1 s filter at the end of the run. Case 9 is in Table 6.5.

0.21 µm. Measured on noise alone, the ratio of the full inverse to a truncation at 40 modes is 16.14, against 16.19 predicted.

Pause. Try Exercise 6.1 before reading on.

6.6.2 Where to Cut

Case 3 adds a minute of girder drift ($q = 0.3\,\mu\text{m}/\sqrt{\text{s}}$, about 7.3 µm per girder in ten minutes) and asks where to cut. Each point of Figure 6.3b is one run of that minute, and the measured trade falls in steps where the spectrum and the drift's content put them. Up to 44 modes the orbit stays at 2.11 to 2.17 µm rms, because the drift lies in modes 45 to 48. At 48 modes it falls to 0.61 µm with a jitter of 0.0278 µrad, and from 48 to 88 modes the orbit does not move while the jitter rises to 0.0430 µrad, because modes 49 to 88 carry almost none of the drift. Only the four weakest modes buy the rest, from 0.59 µm at 92 modes to 0.09 µm at all 96, for a jitter of 0.0875 µrad. The 48-mode cut is therefore the corner of the trade on this ring, and it sits at the cliff because the β pattern puts both the cliff and the drift there.

Tikhonov regularization in case 4 draws a smooth curve through the same plane. At $\mu = 10$ it leaves 0.61 µm for 0.0256 µrad of jitter, the residual of the 48-mode truncation with 8 % less jitter, because its filter falls where the singular values fall. A lower μ buys more of the weak modes, 0.33 µm for 0.0484 µrad at $\mu = 1$ and 0.11 µm for 0.0798 µrad at $\mu = 0.1$, and $\mu = 100$ leaves 1.32 µm for 0.0123 µrad. At every jitter the sweep covers, the Tikhonov points lie below the truncation curve, since a filter corrects a fraction of every weak mode while a cut corrects each one fully or not at all.

6.6.3 A Line Faster Than the Chamber

Case 5 vibrates the girders by 0.5 µm at 10 Hz with only the slow set on. At a slow bandwidth of 1 Hz the 10 Hz line at source point 2 is 105 % of the disturbance's own, and at the 10 Hz maximum it is 111 %, while the share of samples with all three source points in tolerance falls from 92.8 to 89.1 %.

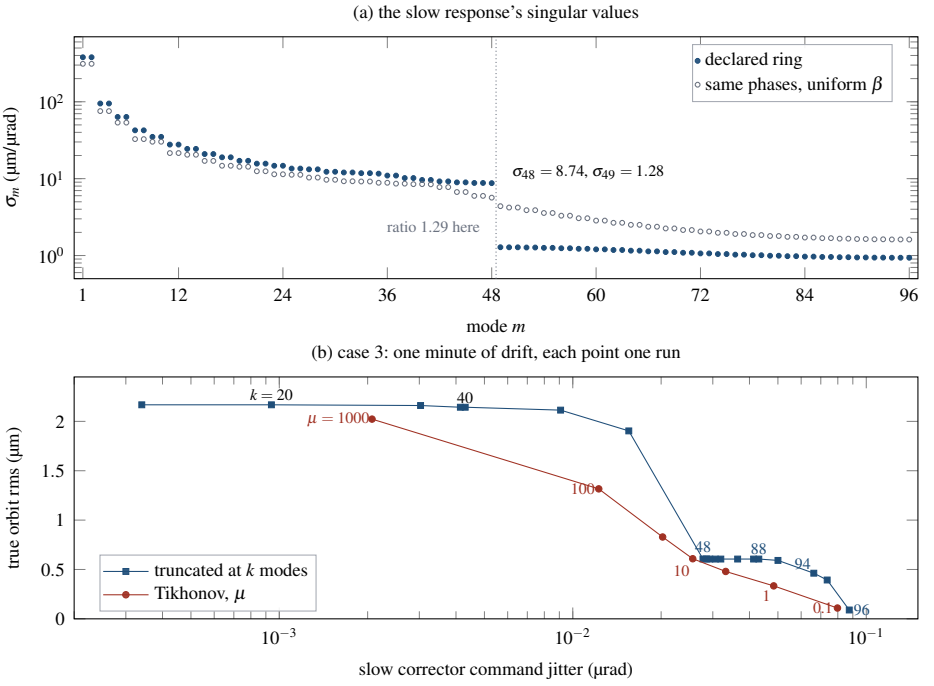


Figure 6.3 The spectrum and the trade it sets. (a) The 96 singular values of the slow response (filled dots), and of the same ring with β made uniform (open circles), where the cliff between modes 48 and 49 disappears. (b) Case 3's minute of drift, each point a run: the true orbit's rms against the slow command jitter for truncation at k modes (squares) and for Tikhonov regularization at every μ the simulator offers (dots).

The cause is the loop and not the chamber. The internal model inverts the 50 ms lag, so what limits the loop is its 10 ms period and 10 ms mean, and its sensitivity at 10 Hz derived from Equation (6.8) is 1.049 at a 1 Hz bandwidth and 1.114 at 10 Hz. An ideal first-order loop would keep 71 % of a line at its own bandwidth, and the same loop run every 5 ms would keep 92 %. The chamber is the reason the slow supplies stop at 10 Hz, and the fast set is the remedy. In case 6 the fast set at 100 Hz keeps 22 % of the line at source point 2. After the ground step at 5 s it holds the beam, and with offloading off its 1 s mean kick is still 0.18 μrad rms at 40 s. Switched on at 10 s, offloading brings that below 0.05 μrad in 1.7 s and to 0.0022 μrad by 40 s, which leaves the fast set in the middle of its ±5 μrad range. Offloading reaches only where the slow inverse reaches. Truncated at 40 or 47 modes the fast set keeps 0.25 or 0.11 μrad, from 48 modes 0.002 μrad, and every μ from 1 to 100 drains it, to 0.002 to 0.036 μrad.

6.6.4 What Feedback Cannot See

In case 7 device 2 closes from 40 to 12 mm, where it kicks 1.92 μrad. With feedback alone source point 2 reaches 2.68 μm and is still 2.63 μm off at 20 s, while every BPM reads close to zero. Ninety-six correctors can null 96 readings, but the device kicks between two BPMs, and the bump it leaves inside its own straight is invisible to all of them. The gap table cancels the kick where it occurs, and the peak falls to 0.30 μm, 8.8 times less. The rest, 0.25 μm at the end, is about the 10 % error of device 2's table applied to its 2.6 μm, which the BPMs cannot see either, so the reduction is measured with an imperfect table and not with a perfect one.

Table 6.5 The Challenge on Seven Paired Records. In Tolerance Is the Share of Samples With All Three Source Points Inside, and Jitter the Slow Commands' at the End

Variant	Pass	In tolerance (%)	Orbit rms (µm)	Jitter (µrad)
shipped: full inverse, slow set only, no test	0/7	67.70–97.48	0.72–1.75	0.126–0.153
suggested: $\mu = 10$, fast set, offloading, feedforward, test	6/7	98.89–100.00	0.51–1.13	0.011–0.015
suggested, full inverse	6/7	99.89–100.00	0.36–0.78	0.047–0.056
suggested, truncated at 48	6/7	96.75–100.00	0.52–1.21	0.012–0.017
suggested, no offloading	7/7	99.51–100.00	0.51–1.14	0.010–0.013
suggested, no feedforward	2/7	60.79–99.26	0.52–1.14	0.011–0.015
suggested, no fast set	0/7	83.63–98.55	0.99–2.00	0.011–0.015
suggested, no BPM test	0/7	86.48–99.56	1.00–1.41	0.025–0.030
shipped with the fast set and test	5/7	97.75–99.89	0.36–0.86	0.047–0.056

TRY IT IN THE APP **7 · The Insertion-Device Gap: Feedback, Then Feedforward**

Load the scenario, press Run and watch the source-point trend. Device 2 closes its gap at 3 s, and with feedback alone source point 2 reaches 2.68 µm while every monitor reads close to zero. Switch Gap feedforward On, press Reset and run again: the peak falls to 0.30 µm.

Pause. Try Exercise 6.2 before reading on.

6.6.5 A Monitor That Lies

When BPM 37 freezes 60 µm off in case 8, both loops believe it. A corrector reaches its limit 0.01 s later and the true orbit runs to 41.8 µm rms by 30 s, while the residual that no girder, dipole or device can make points at BPM 37 with 4,238 times the median share. Switched on at 5 s, the test removes BPM 37 at once, since it has stood above the threshold for 2 s, and the orbit is back within 2 µm rms 5.9 s later and at 0.18 µm by 30 s. With the test on from the start the removal comes 0.20 s after the failure. Before any failure the largest healthy share on the challenge's seven records is 2.1 times the median, against a threshold of 100. The test trusts the model lattice, and a 2 % quadrupole error raises the largest healthy share to 55, without a removal.

6.6.6 Ten Minutes, Graded

The challenge runs ten minutes with drift at $q = 0.15\,\mu\text{m}/\sqrt{\text{s}}$, a 0.3 µm girder vibration at 17 Hz, 0.05 µrad of dipole ripple, device 1 closing at 40 s, bursts at 150 and 520 s, device 2 closing at 220 s, BPM 37 turning to noise at 300 s, the ground step at 380 s and device 3 closing at 440 s. It passes if all three source points are within ± 3 µm and ± 1 µrad on at least 99 % of the 1 kHz samples, the true orbit stays within 2 µm rms over the run, no corrector command touches its limit at any sample, and BPM 37 and no other is removed within 2 s of its failure. The verdict is judged on teaching truth, since no monitor sits at a source point, and the attempt is void once any model parameter or disturbance slider moves. The variants on seven paired records, the case's seed and seeds 1 to 6, each seeing the same disturbances, are compared in Table 6.5.

The shipped settings fail every record, at 67.70 to 97.48 % in tolerance. The suggested settings pass 6 of 7, and on the case's own record they hold 99.73 % in tolerance and 0.78 µm rms and remove BPM 37 0.21 s after it fails, as shown at the moment before the removal in Figure 6.2. The seventh record, seed 5, reaches 98.89 %, and 6,580 of its 6,682 lost samples are source point 2's angle after the ground step and device 3's gap closing, the drift a regularized slow inverse leaves in its weak modes at a place no monitor sees.

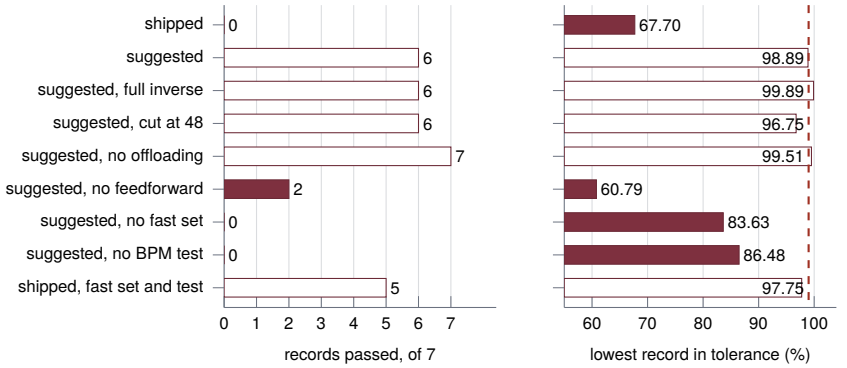


Figure 6.4 The challenge variants of Table 6.5 on its seven paired records: (left) how many records pass, (right) the lowest share of samples in tolerance among them, with the challenge's 99 % as a dashed line. Dark bars pass 2 of 7 or fewer, white bars 5 or more. Notice that the full inverse passes as often as the suggested Tikhonov setting, while removing the fast set, the feedforward or the BPM test drops the passes to 2 or none.

> ✦ **DID YOU KNOW? Quieter Magnets, Not a Pass**
>
> From the simulator: in the challenge, the suggested settings with the full inverse in place of Tikhonov regularization pass the same 6 of 7 records. What the regularization changes is the slow jitter, 0.011 to 0.015 µrad against 0.047 to 0.056 µrad (Table 6.5).

Taking the suggested settings apart one piece at a time shows what the ring needs. Without the fast set no record passes (83.63 to 98.55 %, the bursts), and without the BPM test none does either, because the noisy monitor is chased for five minutes. Without the feedforward only 2 of 7 pass (60.79 to 99.26 %), since feedback alone leaves each source point up to about 2.6 µm off once its gap has closed, as in case 7. The regularization is not what passes the challenge. The full inverse in its place passes the same 6 of 7, failing seed 5 on one 1 ms command sample at its limit at 300.149 s, before the noisy monitor is removed, and dropping offloading passes all 7. What Tikhonov regularization buys is quieter magnets, a slow jitter of 0.011 to 0.015 µrad against 0.047 to 0.056 µrad for the full inverse, and what offloading buys is a centered fast set, a mean fast kick of 0.02 to 0.10 µrad against 0.07 to 0.29 µrad without it. An independent referee written for the review of the simulator, judging the true orbit and the commands at every 1 ms, agrees with the simulator's verdict on all ten records it checks. Figure 6.4 draws the passes and the worst records of Table 6.5.

6.6.7 A Sign No Bandwidth Fixes

The plan the simulator was built from expected its infeasible case from a lattice error without a measured response. On this ring that does not happen. A 1 % rms quadrupole error beats β by 11 % rms and keeps the real part of every loop gain e of the full inverse above 0.82 and its magnitude below 1.17. The slider's 2 % beats β by 23 % rms (57 % at one BPM) and keeps them above 0.69 and below 1.41, so by Equation (6.12) every mode stays stable at every bandwidth, and on 30 s of the challenge's drift the full inverse at 10 Hz holds the orbit at 0.09 µm. A tune error from -0.15 to $+0.15$ with the β pattern right keeps $\mathrm{Re}\, e \geq 0.58$ in the slow loop. Case 10 uses a model error of another kind, corrector 38 wired backward. Its reversed column makes $M = I - 2V^{\mathsf{T}} e_{38} e_{38}^{\mathsf{T}} V$, with one eigenvalue $e = -1$, so the full inverse diverges at every slow bandwidth. It reaches a corrector limit 1.6 s in and an orbit of 132.6 to 136.0 µm at bandwidths from 0.1 to 10 Hz. Truncated at k modes the eigenvalue becomes $1 - 2\|V_k^{\mathsf{T}} e_{38}\|^2$, which holds while the reversed corrector keeps less than half its weight in the kept modes. That weight is 0.107 at 48 modes, 0.347 at 60 and 0.556 at 70,

and the runs hold at 0.38, hold at 0.38 and diverge to 16.28 µm. Under Tikhonov the smallest loop gain is 0.61 at $\mu = 10$, which holds at 0.38 µm, and -0.18 at $\mu = 1$, which diverges to 25.45 µm. The full inverse on a measured response holds at 0.05 µm. With the reversed column it tolerates element errors up to about 0.1 µm/µrad, 40 times the measurement's, where its smallest loop gain is still 0.61, and fails at 0.3 µm/µrad. The BPM test misreads the case. On from the start it removes BPMs 39 and 41, both healthy and beside the reversed corrector, after which the full inverse holds at 0.32 µm only because their rows are gone, and switched on at 5 s it removes BPMs 39, 41 and 40.

> **TRY IT IN THE APP** **10 · A Corrector Wired Backward: The Full Inverse Diverges (Infeasible)**
>
> Load the scenario and press Run: the full inverse drives a corrector to its limit 1.6 s in. Lower Slow bandwidth to 0.1 Hz, press Reset and run again, and the orbit still runs away. Then set Slow feedback to Truncated with Modes kept at 48, press Reset and run once more: the orbit holds at 0.38 µm.

Pause. Try Exercise 6.3 before reading on.

6.6.8 Checks, Agreement and the Control Rate

Code verification asks whether the simulator solves its equations. Closed orbits tracked independently with transfer matrices agree with Equation (6.3) at all 11,520 entries of R_s and R_f to 1.3×10^{-12} µm/µrad. The decomposition reconstructs R_s to 2.6×10^{-14} of its largest entry, and its singular values agree with the eigenvalues of $R_s^{\mathsf{T}} R_s$ from a QR algorithm to 9.9×10^{-15} of σ_1 and with the monitoring studio's own Jacobi solver to 2.2×10^{-15}. One simulated slow mode reproduces the kick standard deviation implied by Equation (6.9) within 1 % at 1, 3 and 10 Hz. Halving the physics step to 0.5 ms moves the rms orbit of cases 5, 6 and 7 by at most 1.4×10^{-11} percent, because every state is updated exactly over a held command and every disturbance keeps its own clock. The refinement that matters is the control rate. In a copy of the simulator whose BPMs and fast loop run at 2 kHz and slow loop at 200 Hz, with the noise per sample raised by $\sqrt{2}$ to keep its density, case 5 keeps 102.3 and 91.3 % of the line, case 6 21.1 %, and the challenge ranks the variants as before, the suggested settings passing 7 of 7 (the lowest at 99.04 %), without the feedforward 3 of 7, and without the fast set or the test none. All 40 of the simulator's checks pass, 22 written for this ring. In a desktop browser at 20 times real time the page draws 58.9 to 60.1 frames a second, and the ten minutes of the challenge take 30.0 s. Agreement with published relations covers the lattice and the inverse. The thin-lens values of Equation (6.1) and the integrated optics agree to four digits, the response is Wiedemann's closed form [1], and the filter factors are the mode-space and filter-factor forms of Tikhonov [2, 8]. No parameter is fitted to a real ring, so the model is verified and not validated. No study of learning has been made with the simulator, so no educational gain is claimed.

6.7 The Ring on the Multivariable Ladder

The ring ends a sequence in this series that runs from a two-by-two temperature laboratory (Volume 1, Chapter 4), through the identification of a many-by-many model (Volume 1, Chapter 12) and a paper machine's profile, to this hundred-by-hundred inverse. It is best taken once multivariable feedback and the singular value decomposition of Section 1.4.7 are familiar. A simulation can draw beside every monitor the truth that no real ring measures [19], here the orbit between the BPMs and at the source points. Three cases carry most of the lesson. Case 3 places the cut on a measured trade, case 7 shows a disturbance that feedback cannot see, and case 10 a model error that no tuning repairs. The residual test of case 8 echoes the gross-error tests of data reconciliation, which the monitoring workbench of Volume 1, Chapter 13 applies to stored records. A corrector wired backward, or a monitor that freezes 60 µm off in one sample, would be rare on a real ring, but each isolates one mechanism. At top speed the ten-minute challenge plays in about 30 s in a browser tab, and the export puts the true orbit and source points beside the readings under column names that begin with `truth`.

Each case is credited by a completion rule that needs the ring's response to the reader's own action (Section 1.6.4). Case 3 asks for a cut between 44 and 52 modes and a full minute run with it, and case 10 asks for the full inverse run to a corrector limit, then a truncation at 60 modes or fewer or a measured response, and an orbit held within 1 µm rms for 5 s after Reset. The challenge is graded by the function the scorecard shows.

> **OPEN CHALLENGE** **Stability at a Tenth of the Beam**
>
> The upgrade of one light source expects beam sizes of the order of 4 µm in both planes, and holding the orbit to a tenth of that pushes the state of the art in fast orbit feedback, in space and in time. Its planned system runs twelve times faster than the present one with a target bandwidth of 1 kHz, and its designers look beyond the classic PID regulator to model-based methods that use the known spectra of orbit motion, which first need a reliable identified model [3]. Progress would be such a controller shown holding a running ring.

6.8 Technology in Practice

What Measures It

A ring reads its orbit only where it has monitors. At one ring two button monitors are planned close to each new magnet so that the beam's position in it can be measured accurately, and their sensitivity and linearity are worked out from the shape of the chamber and the buttons [20]. The monitors are sampled and averaged into fast data at a fixed rate, 10 kHz at one ring, and that average delays the reading a little behind the beam's motion, an effect seen both in simulation and on the machine [21]. The chapter's BPMs are read every 1 ms, and the slow loop works on the mean of the ten readings since its last update.

> **TECHNOLOGY** **Sensor: The Button Beam Position Monitor**
>
> A button monitor sets four small electrodes into the wall of the vacuum chamber. Each passing bunch induces an image charge on them, the front-end electronics filter out a single harmonic of the bunch train's spectrum, and the beam's coordinates follow from the normalized differences of the four button voltages. The relation is linear near the center, and an electrostatic model of the exact chamber and buttons corrects its nonlinearity farther out [20]. At one storage ring the monitors read to about 100 to 250 nm in fast acquisition at 10 kHz [21]. The chapter's 96 monitors read the true orbit plus white noise of 0.5 µm per 1 ms sample, with no gain error or drift of their own.

What Moves It

Two corrector sets act on the beam. At one ring two versions of fast supplies run side by side, the newer with a bandwidth of 6.3 kHz in the horizontal plane and the original with 3.2 kHz. Their converters take a 16-bit command over a range of 1 A, a step of about 15 µA, and the feedback there tolerated supply noise up to about 0.1 mA without a change in the integrated spectrum of the orbit [21]. The chapter's correctors are limited at 5 A, which is 50 µrad for the slow set and 5 µrad for the fast one.

> **TECHNOLOGY** **Final Element: The Corrector and Its Supply**
>
> A fast corrector is a small dipole magnet fed by a bipolar power supply. One supply built for an upgraded ring is a four-quadrant converter, an H-bridge of four transistor switches under pulse-width modulation, whose output current is read by a Hall-effect sensor and held by a proportional-integral regulator. Its key requirement, a small-signal bandwidth of 10 kHz, is hard to meet because the magnet's inductance is high and the supply's input voltage low [22]. The field must then pass the vacuum chamber, whose eddy currents act at one ring as a delay of about 70 to 80 µs [21]. The chapter lumps supply, magnet and chamber into one lag, 50 ms for the slow set and 0.5 ms for the fast one.

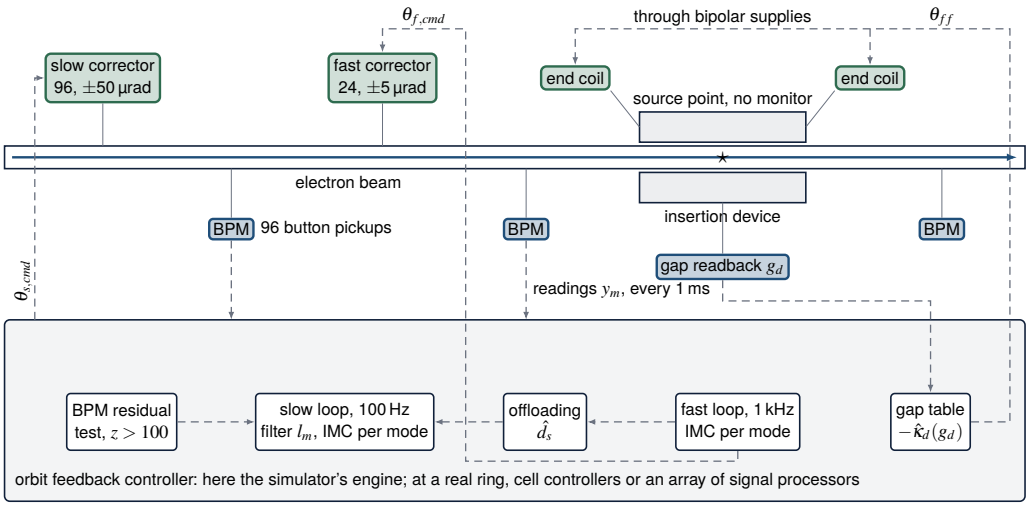


Figure 6.5 The loops of the chapter on their devices, drawn for one stretch of the ring. The 96 button BPMs (blue) feed both loops, the gap readback feeds the table, and the slow correctors, fast correctors and end coils (green) act on the beam through their supplies. Notice that the source point between the end coils carries no monitor, so only the gap table acts there, and that the slow loop reads the fast set's share through offloading.

Where the Controller Runs

At one ring the chain runs from the monitors to cell controllers, which compute the inverse response and a PID compensation, then to an interface that converts each command to an analog signal and through a current limiter to the supply, all at a correction rate of 10 kHz [21]. At another, a distributed array of digital signal processors computes the corrections at 1.6 kHz with an integral-only regulator for a unity-gain bandwidth of 60 to 80 Hz [3]. The chapter's structure is of the same kind at lower rates: the slow loop every 10 ms and the fast loop every 1 ms, each an internal-model controller per singular mode, with offloading between them, the gap table ahead of them and the residual test beside them. Figure 6.5 draws the loops on their devices.

> **IN PRACTICE** **Measure the Response, Then Choose the Cut**
>
> A ring's response matrix is measured by reading the change of beam position while the correctors are varied one by one [7]. Measuring how well the closed loop suppresses motion needs an excitation larger than the beam's own noise, yet a large one disturbs the gain it measures, so its amplitude is a compromise. At one ring with 90 modes, the first few did most of the noise suppression and about ten were enough [21]. The threshold that drops the weak modes is set mainly by the correction needed and the corrector limits, usually as small as keeps every supply out of saturation [7].

What Is Modernizing It

Three directions stand out. Faster and larger systems are one: the feedback planned for one upgraded ring runs twelve times faster than its predecessor, and its model is identified from beam measurements that include the supply, the magnet, the vacuum chamber and the monitors [3]. Supplies that switch an H-bridge at high frequency are the second, built to reach a 10 kHz bandwidth into an inductive magnet [22]. The third brings the experiment into the loop. At one nanoimaging beamline, vibrations at 27 and 120 Hz came from the beamline optics, where the electron monitors cannot see them, and a data link fed the beamline's own position monitor into the electron orbit feedback, whose bandwidth

of about 400 Hz reduced the vibration markedly [23]. As in case 7, feedback corrects only what some monitor sees.

6.9 Limitations

The ring is a declared teaching lattice of thin lenses in the horizontal plane, without coupling, dispersion or energy change, so it says nothing about vertical stability, beam size or a real machine's optics. The BPMs and the fast loop run at 1 kHz and the slow loop at 100 Hz, where the plan asked for 10 kHz and 1 kHz and real fast feedback runs at several kilohertz [8]. The plan's structure cost 103 % of each wall second at 20 times real time in a server-side runtime, and the effect of doubling the rates is the one measured in Section 6.6. Each corrector is one lag without rate limits. The measured response is the true matrix plus element noise of the right size, not a simulated 48 s measurement that would also record the drift. The feedforward table carries a declared error of 8 to 12 %, and case 7's factor of 8.8 belongs to that table. The BPM test trusts the model lattice, removes healthy monitors under a model error, as in case 10, and needed a noise floor to stay quiet at zero noise. The tune and lattice-error sliders stop at 14.35 and 2 %, which keeps the half-integer resonance out of reach. One recorded run per case is graded, seven records are checked for the challenge, and the suggested settings fail one of them. The note on what the reader changed appears 3 s after the change and describes the state at that later moment. A simulated pass is not a specification for any light source.

> **OPEN CHALLENGE** **Using the Weak Modes Without Their Noise**
>
> The small singular values of a measured response are dominated by measurement errors, yet they still carry information that orbit feedback could use. Tikhonov regularization, used at many light sources, lowers their gain instead of dropping them, and at one ring a simulation that matched its measurements pointed to a smoother spatial gain from it, with the dynamics of each mode still to be optimized [21]. Progress would be a way to correct each weak direction as far as it is known. In this chapter the four weakest modes hold 7.3 % of the drift's power.

6.10 Conclusion

A storage ring's orbit feedback inverts a matrix whose weak directions turn noise into magnet motion, and this simulator lets each choice about that inverse be made and measured on a declared ring of 96 monitors and 120 correctors. One 5 µrad kick moves the orbit by up to 74.8 µm round the whole ring, the singular values fall from 379 to 0.94 µm/µrad with a cliff at mode 48 that the β pattern makes, and the full inverse makes the slow correctors jitter 12 times as much as a 48-mode truncation. The girder drift puts 86 percent of its power in modes 45 to 48 and 7.3 percent in modes 93 to 96, so the cut at the cliff leaves 0.61 µm and only the four weakest modes buy the rest. The slow loop amplifies a 10 Hz line by its own period, a gap table corrects what no feedback sees at a point without a monitor, and the residual test finds a frozen monitor at 4,238 times the median share. In the graded ten minutes the fast set, the BPM test and the feedforward decide the verdict, while regularization and offloading buy quieter and centered magnets. No realistic lattice error breaks the full inverse here, and one corrector wired backward breaks it at every bandwidth. The ring is a thin-lens teaching model run at reduced control rates. A fuller model would add the vertical plane and coupling, rate-limited supplies, chambers with several time constants, a response measured through the drifting plant, and the rates real rings run. The simulator shows that how far to trust a model's weak directions is a design decision with a measurable price.

KEY POINTS

- A storage ring's orbit feedback inverts a matrix whose weak directions turn monitor noise into magnet motion.
- The full inverse makes the slow correctors jitter 12 times as much as a 48-mode truncation.
- The girder drift lies mostly in modes 45 to 48, so the cut at the cliff leaves 0.61 µm and only the four weakest modes buy the rest.
- A gap table corrects what no feedback sees at a point without a monitor, and one corrector wired backward breaks the full inverse at every bandwidth.

Where else this applies. Volume 4, Chapter 10 keeps a fast handle centered by a slow one, as offloading does here for the fast correctors.

References

[1] H. Wiedemann, *Particle Accelerator Physics*, ser. Graduate Texts in Physics. Springer International Publishing, 2015.

[2] P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*. Society for Industrial and Applied Mathematics, 1998.

[3] P. Kallakuri, D. Paskvan, J. Carwardine, N. Sereno, and N. Arnold, "Modeling the fast orbit feedback control system for APS upgrade," in *Proceedings of IBIC2017*. Grand Rapids, MI, USA: JACoW, 2018, pp. 196–198.

[4] S. Skogestad, M. Morari, and J. C. Doyle, "Robust control of ill-conditioned plants: high-purity distillation," *IEEE Transactions on Automatic Control*, vol. 33, no. 12, pp. 1092–1105, 1988.

[5] G. E. Stewart, D. M. Gorinevsky, and G. A. Dumont, "Feedback controller design for a spatially distributed system: the paper machine problem," *IEEE Transactions on Control Systems Technology*, vol. 11, no. 5, pp. 612–628, 2003.

[6] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5–6, pp. 469–474, 1998.

[7] Y. Chung, G. Decker, and K. Evans, Jr., "Closed orbit correction using singular value decomposition of the response matrix," in *Proceedings of International Conference on Particle Accelerators*. IEEE, 1993, pp. 2263–2265.

[8] S. Gayadeen, S. R. Duncan, and M. T. Heron, "Fast orbit feedback control in mode space," in *Proceedings of ICALEPCS2013*, San Francisco, CA, USA, 2013, pp. 1082–1085, paper THMIB07.

[9] S. Gayadeen, S. R. Duncan, and W. P. Heath, "Design of multi-array controllers for electron beam stabilisation on synchrotrons," in *2013 American Control Conference*. IEEE, 2013, pp. 1201–1206.

[10] E. D. Courant, M. S. Livingston, and H. S. Snyder, "The strong-focusing synchroton—a new high energy accelerator," *Physical Review*, vol. 88, no. 5, pp. 1190–1196, 1952.

[11] E. D. Courant and H. S. Snyder, "Theory of the alternating-gradient synchrotron," *Annals of Physics*, vol. 3, no. 1, pp. 1–48, 1958.

[12] G. Golub and W. Kahan, "Calculating the singular values and pseudo-inverse of a matrix," *Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis*, vol. 2, no. 2, pp. 205–224, 1965.

[13] J. Demmel and K. Veselić, "Jacobi's method is more accurate than QR," *SIAM Journal on Matrix Analysis and Applications*, vol. 13, no. 4, pp. 1204–1245, 1992.

[14] J. Safranek, "Experimental determination of storage ring optics using orbit response measurements," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 388, no. 1–2, pp. 27–36, 1997.

[15] R. O. Hettel, "Beam stability at light sources (invited)," *Review of Scientific Instruments*, vol. 73, no. 3, pp. 1396–1401, 2002.

[16] C. E. Garcia and M. Morari, "Internal model control. a unifying review and some new results," *Industrial & Engineering Chemistry Process Design and Development*, vol. 21, no. 2, pp. 308–323, 1982.

[17] S. Gayadeen and W. Heath, "An internal model control approach to mid-ranging control," *IFAC Proceedings Volumes*, vol. 42, no. 11, pp. 542–547, 2009.

[18] S. Gayadeen and S. R. Duncan, "Discrete-time anti-windup compensation for synchrotron electron beam controllers with rate constrained actuators," *Automatica*, vol. 67, pp. 224–232, 2016.

[19] J. Park, R. A. Martin, J. D. Kelly, and J. D. Hedengren, "Benchmark temperature microcontroller for process dynamics and control," *Computers & Chemical Engineering*, vol. 135, p. 106736, 2020.

[20] V. Smaluk, B. Kosciuk, and B. Bacha, "A button-type beam position monitor for the complex bend test in NSLS-II storage ring," Brookhaven National Laboratory, Upton, NY, USA, Tech. Rep. NSLSII-ASD-TN-446 (BNL-230229-2026-TECH), 2026.

[21] S. Kongtawong, Y. Tian, L. Yu, X. Yang, G. Wang, K. Ha, and T. Shaftan, "Numerical simulation of NSLS-II fast orbit feedback system," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 997, p. 165175, 2021.

[22] J. Wang and G. Sprau, "A high bandwidth bipolar power supply for the fast correctors in the APS upgrade," in *Proceedings of NAPAC2016*. Chicago, IL, USA: JACoW, 2017, pp. 96–98.

[23] S. Kongtawong, Y. Hidaka, H. Yan, G. Wang, Y. Tian, K. Ha, J. Mead, D. Padrazo, Y. Chu, and T. Shaftan, "Stabilization of x-ray beam for nanoimaging using source electron fast orbit feedback," in *X-Ray Nanoimaging: Instruments and Methods VII*. San Diego, CA, USA: SPIE, 2025.

6.11 Exercises

6.1 Concept check · 2 min

In scenario 2 the full inverse makes the slow correctors jitter 0.198 µrad about their own mean. Truncated to 48 modes on the same BPM noise and the same 3 Hz bandwidth, they jitter 0.0164 µrad, twelve times less. What makes the difference?

(a) The full inverse divides the noise in the 48 weak modes by singular values near 1 µm/µrad

(b) The full inverse corrects the weak modes, which carry most of the girder drift

(c) The 50 ms chamber amplifies the commands of the full inverse more than those of the truncated one

(d) The full inverse runs the slow loop every millisecond instead of every 10 ms

6.2 Concept check · 1 min

In scenario 7 device 2 closes its gap with feedback alone. After 20 s every BPM reads close to zero, but source point 2 is still 2.63 µm off. Why does the slow loop leave it there?

(a) The 1 Hz slow loop has not yet caught up with a gap that closes over 5.6 s

(b) The device kicks between two BPMs, and the bump inside its straight is invisible to all of them

(c) The 48-mode truncation drops the modes in which the device's kick lies

(d) The gap table is off by 10 % for device 2

6.3 Concept check · 2 min

In scenario 10 corrector 38 is wired backward and the full inverse drives the orbit to more than 130 µm rms. Why does lowering the slow bandwidth to 0.1 Hz not stop it?

(a) The correctors reach their 50 µrad limit, and the integrators wind up behind it

(b) A 0.1 Hz loop is too slow for the ground drift of the case

(c) One eigenvalue of the loop gain is -1, and a negative loop gain diverges at every bandwidth

(d) The BPM test removes two healthy monitors and leaves the inverse without rank

6.4 Theory and calculation · 8 min

The teaching ring has a tune $\nu = 14.2$, 48 thin-lens FODO periods and drifts of $L = 5\,\mathrm{m}$. Use Equations (6.1) and (6.3).

(a) Find the phase advance μ of one period, the focal length f and the amplitude function $\beta_{\pm}$ at a QF and at a QD. Why does each 20 m cell carry two periods? *Checkpoint:* 106.5°, 3.12 m, 18.79 m and 2.07 m; one period per cell would need 213°.

(b) A BPM 0.4 m after a QF sees $\beta = 16.46\,\text{m}$. Taking $|\cos\Delta_{ij}| = 1$, find the largest orbit move at such a BPM for a 5 µrad kick at a corrector at a QF and at one at a QD. *Checkpoint:* 14.96 and 4.97 µm/µrad, so 74.8 and 24.8 µm.

(c) Find the ratio of the two moves and say what sets it. *Checkpoint:* 3.0, $\sqrt{18.79/2.07}$.

6.5 Theory and calculation · 7 min

The slow loop runs every $T = 10\,\text{ms}$ behind a chamber lag of $\tau_s = 50\,\text{ms}$. Its singular values include $\sigma_1 = 379.17$, $\sigma_{48} = 8.740$, $\sigma_{49} = 1.282$ and $\sigma_{96} = 0.9357\,\mu\text{m}/\mu\text{rad}$. Use Equations (6.7) and (6.9).

(a) Find the Tikhonov filter factor l_m at $\mu = 10$ for these four modes. *Checkpoint:* 1.00, 0.88, 0.14 and 0.08.

(b) At what singular value is a mode corrected halfway, and where does it fall in the spectrum? *Checkpoint:* $\sigma = \sqrt{10} = 3.16\,\mu\text{m}/\mu\text{rad}$, inside the cliff between modes 48 and 49.

(c) Find the ratio of command jitter to kick jitter at slow bandwidths of 1 and 10 Hz. *Checkpoint:* 2.01 and 4.92.

6.6 Theory and calculation · 6 min

In scenario 10 corrector 38 is reversed. With the inverse truncated at k modes, the reversed mode's loop gain is $e = 1 - 2\|V_k^{\mathsf{T}} e_{38}\|^2$. The weight $\|V_k^{\mathsf{T}} e_{38}\|^2$ is 0.107 at 48 modes, 0.347 at 60 and 0.556 at 70.

(a) Find e at each truncation and predict which runs hold. *Checkpoint:* 0.786, 0.306 and -0.112: 48 and 60 hold, 70 diverges.

(b) For the fast loop ($b_1 = 0$, $T = 1\,\text{ms}$) at 100 Hz, find λ, the range of e for which the loop is stable, and the pole for $e = -1$. *Checkpoint:* $\lambda = 0.533$, $0 < e < 4.29$, $z = 1.47$.

6.7 App challenge · 8 min

Load scenario 2, The full inverse: the magnets chase the noise, at the default seed and a speed of 1. The scenario sets Slow feedback to Full inverse at a slow bandwidth of 3 Hz and a BPM noise of 1 µm.

(a) Run to 20 s and pause. Read the Slow corrector jitter tile, the noise prediction it shows beside the measurement, and the RMS orbit through the 1 s filter. *Checkpoint:* 0.198 µrad against 0.199 µrad predicted, orbit 0.09 µm.

(b) Set Slow feedback to Truncated (Modes kept stays at 48) and run to the end of the case at 40 s. Read the same tiles. *Checkpoint:* 0.016 µrad against 0.016 µrad, orbit 0.21 µm.

(c) Compare the jitter ratio with the 12.77 that Equation (6.9) predicts, and say what the truncation costs. *Checkpoint:* About 12; the orbit doubles, from 0.09 to 0.21 µm.

Accept jitters within 0.002 µrad and orbits within 0.03 µm. A switch up to half a second from 20 s stays inside these tolerances.

6.8 App challenge · 8 min

Load scenario 3, Truncation at 40 modes: residual against noise, at the default seed and a speed of 5. The case runs one minute of girder drift and pauses at 60 s. For each run read the RMS orbit over the run and the Slow corrector jitter at the end.

(a) Run the case with its 40 modes. Then set Modes kept to 48, press Reset and run it again. *Checkpoint:* 2.14 µm and 0.004 µrad; 0.61 µm and 0.028 µrad.

(b) Repeat with 88 and with 96 modes kept. *Checkpoint:* 0.61 µm and 0.043 µrad; 0.09 µm and 0.088 µrad.

(c) Explain why 40 extra modes from 48 to 88 buy nothing while the last eight buy most of the remaining residual. *Checkpoint:* The drift's power by mode: modes 45 to 48 and 93 to 96.

Accept orbits within 0.02 µm and jitters within 0.002 µrad.

6.9 App challenge · 7 min

Load scenario 7, The insertion-device gap: feedback, then feedforward, at the default seed and a speed of 2. Device 2 closes from 40 to 12 mm starting at 3 s, and the case pauses at 20 s.

(a) Run the case with feedback alone. Read the Worst source point tile, and the position of source point 2 at 20 s from the column `truth_src2_um` of the exported file. *Checkpoint:* 2.69 µm; 2.63 µm.

(b) Set Gap feedforward to On, press Reset and run again. Read the same two values. *Checkpoint:* 0.31 µm; 0.25 µm.

(c) Account for the 0.25 µm that the feedforward leaves. *Checkpoint:* The table's declared −10 % error for device 2.

Accept positions within 0.03 µm.

6.10 Application area · 4 min

Scenario 10 can be held by inverting a measured response instead of the model. A real measurement kicks each of the 96 slow and 24 fast correctors by ±10 µrad and averages the BPMs over 200 samples each way at 1 kHz.

(a) How long does the measurement take, and what element error does Equation (6.6) give at a BPM noise of 0.5 µm? *Checkpoint:* 48 s; 0.0025 µm/µrad.

(b) The full inverse on the reversed ring tolerates element errors up to about 0.1 µm/µrad. Name two things a measurement on a working ring would add that the simulator leaves out. *Checkpoint:* Drift during the measurement, and beam time.

6.11 Application area · 3 min

A beamline scientist reads that the suggested settings keep all three source points within ±3 µm and ±1 µrad on 99.73 % of the samples of the graded challenge. She asks whether the ring would meet her experiment's needs. Give three reasons from the chapter why this pass is not a specification.

6.12 Review the agent · 5 min

An AI assistant reviews the graded challenge and writes: "The shipped full inverse fails because it amplifies BPM noise through the weak modes. The fix is Tikhonov regularization at $\mu = 10$. It is the setting that makes the suggested configuration pass." Find the flaw in the reasoning and say how you would check it on the app.