Chapter 9
Automatic Tuning and Control for Advanced Light Sources

Alexander Scheinker

Abstract The next generation of X-ray Free Electron Laser (FEL) advanced light sources allow users to drastically change beam properties for various experiments. The main advantage of FELs over synchrotron light sources is their ability to provide more coherent, brighter flashes of light by tens of orders of magnitude with custom bunch lengths down to tens of femtoseconds. The wavelength of the brighter, more coherent light produced by an FEL is extremely dependent on both the electron beam energy, which must be adjusted between different experiments, and maintaining minimal electron bunch emittance. A large change in beam energy and bunch length usually requires a lengthy manual re-tuning of almost the entire accelerator. Therefore, unlike traditional machines which can operate for months or years at fixed energies, RF, and magnet settings FELs must have the ability to be completely re-tuned very quickly. For example, the Linac Coherent Light Source (LCLS) FEL can provide electrons at an energy range of 4–14 GeV and 1 nC pulses with 300 fs pulse width down to 20 pC pulses with 2 fs pulse width. The next generation of X-ray FELs will provide even bright, shorter wave-length (0.05 nm at EuXFEL, 0.01 nm at MaRIE), more coherent light, and at higher repetition rates (2 MHz at LCLS-II and 30,000 lasing bunches/second at EuXFEL, 2.3 ns bunch separation at MaRIE) than currently possible, requiring smaller electron bunch emittances than achievable today. Therefore, the next generation of light sources face two problems in terms of tuning and control. In parallel with the difficulties of improving performance to match tighter constraints on energy spreads and beam quality, existing and especially future accelerators face challenges in maintaining beam quality and quickly tuning between various experiments. We begin this chapter with a brief overview of some accelerator beam dynamics and a list of control problems important to particle accelerators. In the second half of this chapter we introduce some recently developed model-independent techniques for the control and tuning of accelerators with a focus on a feedback based extremum seeking method for automatic tuning and optimization which can tune multiple coupled parameters simultaneously and is incredibly robust to time-variation of system components and noise.

A. Scheinker (✉)
Los Alamos National Laboratory, Los Alamos, NM, USA
e-mail: ascheink@lanl.gov

9.1 Introduction

Particle accelerators are large complex systems composed of many thousands of coupled components which include radio frequency (RF) electromagnetic accelerating cavities, magnets, cooling systems, and detectors. For many decades accelerators have been designed with specific, static, operating conditions in mind, such as specific beam energies, currents, repetition rates, and bunch separations. For example, the Los Alamos Neutron Science Center accelerator is a ∼1 km long linear accelerator that has two fixed design energies of 100 and 800 MeV, and 8 fixed beam types which vary in terms of bunch length, charge/bunch, and repetition rate. Once the accelerator is tuned up following a maintenance outage, it is mostly continuously run with the various beam types accommodated by a fixed magnet/RF system setup, with intermediate tuning by operators to make up for small disturbances and fluctuations. The advanced photon source (APS) is a ∼1 km circumference synchrotron with magnet and RF systems tuned for a fixed 7 GeV electron beam which can be sent to various user stations with unique magnet and optic systems including monochromators for the production of specific light energy ranges from 3.5 to 100 keV. The Large Hadron Collider at CERN is the world's most powerful accelerator with a circumference of 27 km and beam energy of 6.5 TeV per beam for two counter circulating proton beams. The machine is run at a fixed energy for years at a time while massive detectors at four collision points collect data for fundamental particle physics research. The three machines described above encompass a majority of existing accelerators, which are designed for and operated at fixed settings, providing very specific beam types and energies.

Unlike the static machines described above, the next generation of X-ray Free Electron Laser (FEL) advanced light sources are being designed and operated with the fundamentally different approach of allowing users to drastically change beam properties for various experiments. The main advantage of FELs over synchrotron light sources such as the APS is their ability to provide more coherent, brighter flashes of light by tens of orders of magnitude with custom bunch lengths down to tens of femtoseconds. The wavelength of the brighter, more coherent light produced by an FEL is extremely dependent on the electron beam energy, which must be adjusted between different experiments. A large change in beam energy and bunch length requires the re-tuning of almost the entire accelerator. For example, the shortest, few femtosecond electron bunches require an adjustment of the source in lowering the total electron bunch charge so that the space charge forces of such short pulses are manageable. The bunch compressor system and RF energy settings and offsets must then also be adjusted to provide the new, shorter bunch length. Finally, depending on the required light and therefore electron beam energy of the given experiment, the magnet focusing systems throughout the accelerator and the undulator must be retuned. Therefore, unlike traditional machines which can operate for months or years at fixed energies, RF, and magnet settings FELs must have the flexibility to be completely re-tuned. For example, the Linac Coherent Light Source (LCLS) FEL

can provide electrons at an energy range of 4–14 GeV and 1 nC pulses with 300 fs pulse width down to 20 pC pulses with 2 fs pulse width.

The next generation of X-ray FELs will provide even bright, shorter wave-length (0.05 nm at EuXFEL, 0.01 nm at MaRIE), more coherent light, and at higher repetition rates (2 MHz at LCLS-II and 30000 lasing bunches/second at EuXFEL, 2.3 ns bunch separation at MaRIE) than currently possible, requiring smaller electron bunch emittances than achievable today. Existing light sources are also exploring new and exotic schemes such as two-color operation (LCLS, FLASH, SwissFEL). To achieve their performance goals, the machines face extreme constraints on their electron beams. The LCLS-II requires $<0.01\%$ rms energy stability, a factor of $>10\times$ more than the existing LCLS [1], while the EuXFEL requires <0.001 deg rms RF amplitude and phase errors, respectively (current state of the art is ~0.01) [2].

Therefore, the next generation of light sources face two problems in terms of tuning and control. In parallel with the difficulties of improving performance to match tighter constraints on energy spreads and beam quality, existing and especially future accelerators face challenges in maintaining beam quality and quickly tuning between various experiments. It can take up to 10 h to retune the low energy beam sections (<500 MeV) and they still achieve sub-optimal results, wasting valuable beam time. Future accelerators require an ability to quickly tune between experiments and to compensate for extremely closely spaced electron bunches, such as might be required for MaRIE, requiring advanced controls and approaches such as droop correctors [3, 4].

While existing and planned FELs have automatic digital control systems, they are not controlled precisely enough to quickly switch between different operating conditions [5]. Existing controls maintain components at fixed set points, which are set based on desired beam and light properties, such as, for example, the current settings in a bunch compressor's magnets. Analytic studies and simulations initially provide these set points. However, models are not perfect and component characteristics drift in noisy and time-varying environments; setting a magnet power supply to a certain current today does not necessarily result in the same magnetic field as it would have 3 weeks ago. Also, the sensors are themselves noisy, limited in resolution, and introduce delays. Therefore, even when local controllers maintain desired set points exactly, performance drifts. The result is that operators continuously tweak parameters to maintain steady state operation and spend hours tuning when large changes are required, such as switching between experiments with significantly different current, beam profile (2 color, double bunch setups), or wavelength requirements. Similarly, traditional feed-forward RF beam loading compensation control systems are limited by model-based beam-RF interactions, which work extremely well for perfectly known RF and beam properties, but in practice are limited by effects which include un-modeled drifts and fluctuations and higher order modes excited by extremely short pulses. These limitations have created an interest in model-independent beam-based feedback techniques that can handle time-varying uncertain nonlinear systems [6–13], as well as machine learning, and other optimization techniques [14–18].

We begin this chapter with a list of control problems important to particle accelerators and a brief overview of simple beam dynamics, including longitudinal and

transverse effects and the coupling between them and an overview of RF systems. The second half of this chapter introduces some recently developed techniques for the control and tuning of accelerators with a focus on a feedback based extremum seeking method for automatic tuning and optimization.

9.1.1 Beam Dynamics

The typical coordinate system for discussing particle accelerator beam dynamics is shown in Fig. 9.1. The Lorentz force equation:

$$\frac{d\mathbf{P}}{dt} = e\left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B}\right), \tag{9.1}$$

describes charged particle dynamics. In (9.1) e is electron charge, $\mathbf{v}$ is velocity, $v = |\mathbf{v}|$, $\mathbf{P} = \gamma m \mathbf{v}$ the relativistic momentum, $\gamma = 1/\sqrt{1 - \frac{v^2}{c^2}}$ the Lorentz factor, c the speed of light, $\mathbf{E}$ the electric field and $\mathbf{B}$ the magnetic field. In a particle accelerator $\mathbf{E}$ and $\mathbf{B}$ sources include electromagnetic accelerating fields, other charged particles, and magnets used for steering and focusing of the beams. While electric fields are used to accelerate particles, magnetic fields guide the particles along a design trajectory and keep them from diverging transversely. We start by reviewing betatron oscillations, a form of oscillatory motion which is common to all particle accelerators [19–24].

Betatron oscillations are a general phenomenon occurring in all particle accelerators and are of particular importance in circular machines. For a particle traveling at the designed beam energy, $p = p_0$, the transverse equations are given by Hill's equation

$$x'' = K_x(s)x, \qquad y'' = K_y(s)y, \tag{9.2}$$

with (x, y) being the transverse particle locations relative to the accelerator axis (see Fig. 9.1), s (or z) is a parametrization of particle location along the axis of the accelerator, and $x'(s) = dx(s)/ds$. In a ring, the function $K_{x,y}(s + L) = K_{x,y}(s)$ is L-periodic, where L is the circumference of the accelerator, and depends on magnetic field strengths. Equation (9.2) resembles a simple harmonic oscillator with position-
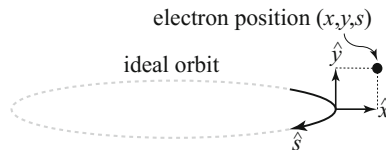

Fig. 9.1 A coordinate system centered on the ideal particle orbit. Distance along the orbit is parametrized by s. Transverse offset from the axis of the orbit is given by x and y

dependent $K_{x,y}(s)$. The solution of (9.2) is of the form

$$p_{x,y}(s) = A\sqrt{\beta_{x,y}(s)} \cos\left(\psi_{x,y}(s) + \delta\right), \quad \psi_{x,y}(s) = \int_0^s \frac{d\sigma}{\beta_{x,y}(\sigma)}, \tag{9.3}$$

where $\beta_{x,y}(s)$ are the periodic solutions of the system of equations

$$\beta_{x,y}'''(s) + 4K_{x,y}(s)\beta_{x,y}'(s) + 2K_{x,y}'(s)\beta(s)_{x,y} = 0, \tag{9.4}$$

$$\frac{1}{2}\beta(s)_{x,y}\beta_{x,y}''(s) - \frac{1}{4}\left(\beta_{x,y}'(s)\right)^2 + K_{x,y}(s)\beta_{x,y}^2(s) = 1. \tag{9.5}$$

The solutions of (9.3) are known as betatron oscillations and are periodic functions of s with varying amplitude and frequency [20].

In general, betatron motion is governed by equations of the form:

$$x''(s) = -K_x(x, y, s, \mathbf{P}, t)x(s) + F_x(x, x', y, y', s, \mathbf{P}, t), \tag{9.6}$$

$$y''(s) = -K_y(x, y, s, \mathbf{P}, t)y(s) + F_y(x, x', y, y', s, \mathbf{P}, t). \tag{9.7}$$

The nonlinear coupling between x and y depends not only on particle position, trajectory, energy deviation, and time.

Typically, quadrupole magnets focus the beam transversally, maintaining a tight bunch along the accelerator axis, and dipole magnets having only a non-zero y component of magnetic field direct the particles in a circular orbit in the (x, s) plane. The linear quadrupole and dipole magnetic field components give (9.6), (9.7) of the form

$$x'' = -\frac{p_0}{p}\left(\frac{1}{\rho^2} - K_1(s)\right)x + \frac{p - p_0}{p}\frac{1}{\rho}, \tag{9.8}$$

$$y'' = -\frac{p_0 K_1(s)}{p}y. \tag{9.9}$$

$K_1(s)$ is periodic and proportional to quadrupole field strength. The value $p = \sqrt{E^2/c^2 - m^2c^2}$ is the total kinetic momentum. p_0 is the designed kinetic momentum. The value ρ is the local radius of curvature [20].

Sources of nonlinearity and coupling in the functions F_x and F_y in (9.6), (9.7) are nonlinear magnetic field components, misaligned magnets, solenoid fields, magnetic field errors, and skew components of magnetic field gradients. Furthermore, all manufactured magnets are non-ideal and introduce nonlinear field components, higher order coupling terms given by [23]:

$$\Delta B_y + j\Delta B_x = B_0 \sum_{n=0}^{\infty} (b_n + ja_n)(x + jy)^n. \tag{9.10}$$
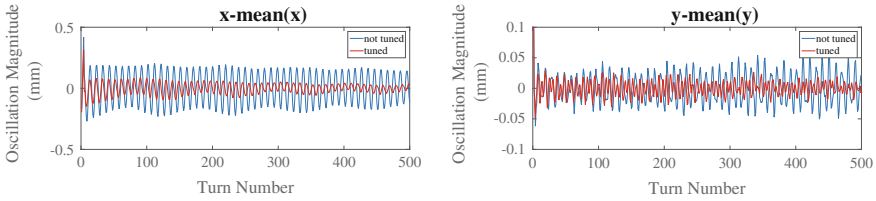
Fig. 9.2 BPM readings of x and y beam displacement over 500 turns, before and during tuning

Sometimes nonlinear magnets are purposely introduced into the accelerator lattice. For example, sextuple magnets are placed in regions of high dispersion to mitigate the fact that particles with various momentums experience non-equal forces from the same magnetic fields and their trajectories diverge (chromatic effects). Such magnets result in nonlinear coupling terms such as $(x^2 - y^2)$ and $(1 - \Delta)xy$, where $\Delta = (p - p_0)/p$ [20].

Betatron motion occurs in all accelerators, magnetic lattices are designed to minimize betatron oscillations. However, some regions of accelerators require large amplitude transverse particle motion. If this motion is not carefully, precisely controlled, excessive betatron oscillations are generated. One such section is a group of pulsed kicker magnets used to horizontally kick the beam out and then inject back into a machine. During injection kicks an imperfect match of parameters of the magnets results in the extremely large betatron oscillations, as shown in Fig. 9.2.

9.1.2 RF Acceleration

Particle acceleration in an RF field. For a particle passing through an RF cavity gap of length L, the energy gain due to an electromagnetic standing wave along the axis is given by

$$\Delta W = q \int_{-L/2}^{L/2} E(z) \cos(\omega t(z) + \phi)dz, \quad t(z) = \int_0^z \frac{dz}{v(z)}, \tag{9.11}$$

where $t(z)$ has been chosen such that the particle is at the center of the accelerating gap at $t = 0$, $\phi = 0$ if the particle arrives at the origin when the field is at a crest, and $v(z)$ is the velocity of the particle. This energy gain can be expanded as

$$\Delta W = q \int_{-L/2}^{L/2} E(z) \left[\cos(\omega t(z)) \cos(\phi) - \sin(\omega t(z)) \sin(\phi)\right] dz \tag{9.12}$$

and rewritten in the form

$$\Delta W = q V_0 T \cos(\phi) = q E_0 T L \cos(\phi), \quad E_0 = \frac{V_0}{L}, \tag{9.13}$$

where

$$V_0 = \int\limits_{-L/2}^{L/2} E(Z) dz, \quad T = \frac{\int_{-L/2}^{L/2} E(z) \cos(\omega t(z)) dz}{V_0} - \tan(\phi) \frac{\int_{-L/2}^{L/2} E(z) \sin(\omega t(z)) dz}{V_0}, \tag{9.14}$$

and T known as the transit-time factor. For typical RF accelerating cavities, the electric field is symmetric relative to the center of the gap and the velocity change within an accelerating gap for a relativistic particle is negligible so $\omega t(z) \approx \omega z / v = 2\pi z / \beta \lambda$, where $\beta = v/c$ and $\beta \lambda$ is the distance a particle travels in one RF period. We can then rewrite the transit-time factor as

$$T = \frac{\int_{-L/2}^{L/2} E(z) \cos(2\pi z / \beta \lambda) \, dz}{V_0}. \tag{9.15}$$

Assuming that the electric field is constant $E(z) \equiv E_0$ within the gap, we get

$$T = \frac{\sin(\pi L / \beta \lambda)}{\pi L / \beta \lambda}, \tag{9.16}$$

and plugging back into (9.13) we get

$$\Delta W = \frac{q E_0 \beta \lambda}{\pi} \cos(\phi) \sin\left(\frac{\pi L}{\beta \lambda}\right), \tag{9.17}$$

which is, as expected, maximized for $\phi = 0$ and $L = \beta \lambda / 2$, that is for a particle that spends the maximal half of an RF period being accelerated through the cavity. This however would not be an efficient form of acceleration as most of the time the particle would see a much smaller than maximal RF field. For a given voltage gain V_0, we get a maximum $T = 1$ with $L = 0$, which is not realizable. Actual design values of T depend on individual cavity geometries and desired efficiency.

9.1.3 Bunch Compression

For maximal acceleration, we typically choose $\phi = 0$, especially for highly relativistic electrons. However, sometimes a nonzero ϕ is chosen either for longitudinal bunching or to purposely introduce an energy gradient along the electron bunch which can then be utilized for bunch compression. We define ϕ as the relative phase

between a particle and the zero crossing of the RF field, such that earlier particles, with $\phi < 0$ will receive a higher energy gain than later particles with $\phi > 0$. The energy offset of a particle at phase ϕ at the exit of the RF compressor cavity, relative to the reference particle, is given by

$$\Delta E_1 = \Delta E_0 - \frac{q V_{\text{rf}}}{E} \sin(\phi), \tag{9.18}$$

where V_{rf} is the compressor voltage, E is beam energy, ΔE_0 is the initial energy offset. Next the beam is transported through a dispersive section with non-zero R_{56}, where

$$R_{56}(s) = \int_{s_0}^{s} \frac{R_{16}(s')}{\rho(s')} ds', \tag{9.19}$$

where R_{16} is the transverse displacement resulting from an energy error in a dispersive region of the accelerator. The energy offset is then translated to a longitudinal position offset according to

$$\Delta z_1 = \Delta z_0 + R_{56} \Delta E_1 = \Delta z_0 + R_{56} \left(\Delta E_0 - \frac{q V_{\text{rf}}}{E} \sin(\phi) \right). \tag{9.20}$$

For an RF field of frequency ω_{rf}, the phase ϕ relative to the RF at position offset Δz_0 is given by $\phi = -\omega_{\text{rf}} \Delta z_0 / c$. If this phase is small, we can expand sine and rewrite both the energy and position change as

$$\Delta z_1 \approx \left(1 + R_{56} \frac{\omega_{\text{rf}} V_{\text{rf}}}{cE} \right) \Delta z_0 + R_{56} \Delta E_0, \tag{9.21}$$

$$\Delta E_1 = \Delta E_0 - \frac{e V_{\text{rf}} \omega_{\text{rf}}}{cE} \Delta z_0. \tag{9.22}$$

Therefore the final bunch length can be approximated as

$$\sigma_{zf} = \sqrt{\left(1 + R_{56} \frac{e V_{\text{rf}} \omega_{\text{rf}}}{Ec} \right)^2 \sigma_{z0}^2 + R_{56}^2 \sigma_{\Delta E_0}^2}, \tag{9.23}$$

where σ_{z0} is the initial bunch length and $\sigma_{\Delta E_0}$ is the initial beam energy spread [26], with maximal compression for an RF system adjusted such that $R_{56} \frac{e V_{\text{rf}} \omega_{\text{rf}}}{Ec} \approx -1$.

9.1.4 RF Systems

For a right-cylindrical conducting cavity of radius R_c, as shown in Fig. 9.3, the 010 transverse-magnetic resonant mode, referred to as **TM$_{010}$**, is used for acceleration
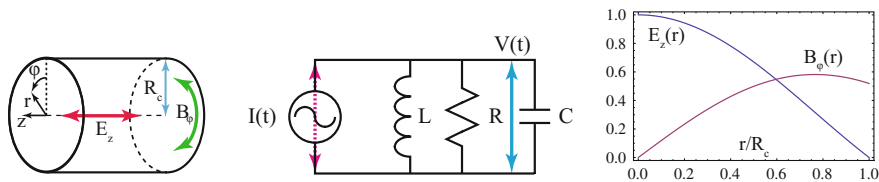
Fig. 9.3 Left: Electromagnetic field orientations for **TM$_{010}$** accelerating mode of a right cylindrical RF cavity. **Center**: RLC circuit approximation of the dynamics of a single RF mode. **Right**: The axial electric field is maximal on axis and zero at the walls of the cavity and the opposite is true of the azimuthal magnetic field

because along the axis this mode has a large oscillating electric field and no magnetic field, as shown in Fig. 9.3. The electromagnetic fields of the **TM$_{010}$** mode are:

$$\mathbf{E}(r, t) = E_0 J_0 \left(\frac{2.405r}{R_c} \right) e^{i\omega_0 t} \hat{\mathbf{z}} = \mathbf{E}_z(r) e^{i\omega_0 t} \hat{\mathbf{z}}, \tag{9.24}$$

$$\mathbf{B}(r, t) = -i E_0 \sqrt{\frac{\epsilon}{\mu}} J_1 \left(\frac{2.405r}{R_c} \right) e^{i\omega_0 t} \hat{\varphi} = \mathbf{B}_\varphi(r) e^{i\omega_0 t} \hat{\varphi}, \tag{9.25}$$

where J_0 and J_1 are Bessel functions of the first kind with zero and first order, respectively, and the resonant frequency is given by

$$\omega_0 = \frac{2.405c}{R_c}, \qquad c = \text{speed of light.} \tag{9.26}$$

The dynamics of such a single mode of an RF cavity with resonant frequency f_0 can be approximated as

$$\ddot{V}_{\text{cav}} + \frac{\omega_0}{Q_L} \dot{V}_{\text{cav}} + \omega_0^2 V_{\text{cav}} = \frac{1}{C} \dot{I}, \tag{9.27}$$

where $\dot{V} = \frac{dV}{dt}$, $\ddot{V} = \frac{d^2V}{dt^2}$, $\omega_0 = 2\pi f_0$, Q_L is the loaded quality factor of the resonant cavity, L and C are the inductance and capacitance of the cavity structure, respectively, such that $\sqrt{LC} = \frac{1}{\omega_0}$, and $I = I_c + I_b$ is the input current driving the RF fields, the sources of which are both the RF generator, I_c, and the beam itself I_b [19, 27, 28].

For a driving current of the form

$$I_u(t) = I_0 \cos(\omega_0 t), \tag{9.28}$$

after the fast decay of some transient terms, the cavity response is of the form

$$V_{\text{cav}}(t) = R I_0 \left(1 - e^{-t/\tau} \right) \cos(\omega_0 t), \quad \tau = \frac{2Q_L}{\omega_0}. \tag{9.29}$$
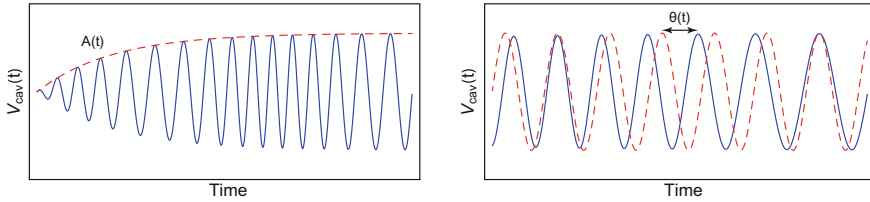
Fig. 9.4 Amplitude of the cavity field and its phase relative to a reference signal

9.1.5 Need for Feedback Control

Although (9.29) implies that for a desired accelerating gradient one must simply choose the correct input power level and drive the cavity, as shown in Fig. 9.4. However, in the real world simply choosing set points for an RF drive signal does not work because of un-modeled time varying disturbances which perturb cavity fields from their desired set points. These disturbances include:

1. Temperature variation-induced resonance frequency drifts on the time scales of minutes to hours.
2. Mechanical vibrations which alter the cavity resonance frequency on the times scale of milliseconds.
3. RF source voltage and current fluctuations on the time scale of microseconds.
4. RF source voltage droop on the time scale of microseconds.

Furthermore, even if a desired accelerating voltage could be reached within a desired rise time, when the beam that is to be accelerated shows up, it itself perturbs the fields both by interacting with the oscillating electrons in the cavity walls and by drawing energy out of the cavity via the electric field which accelerates it, causing both amplitude and phase changes on the time scales of nanoseconds which must be compensated for in order to maintain proper acceleration of subsequent beam bunches.

Therefore real time active feedback control is always necessary, both to bring cavity voltage amplitudes and phases to their required set points before beam can be properly accelerated and during beam acceleration in order to maintain tight bounds on beam-induced cavity field errors, known as beam loading.

From the above discussions it is clear that all of the disturbances experienced by the RF systems immediately couple into the transverse and longitudinal beam dynamics. Similarly, many of the beam dynamics, including the effects of space charge forces, magnet misalignments, and energy deviations alter a particle's position within a bunch and therefore the phase of the RF system relative to the particle's arrival time and therefore the entire accelerator is a completely coupled system in terms of the final beam phase space distribution relative to the RF systems, magnet systems, and the forces due to the particles in the beam itself.

9.1.6 *Standart Proportional Integral (PI) Control for RF Cavity*

The vast majority of accelerator systems, such as RF feedback and power converters are typically controlled at fixed set points with simple, classical, proportional integral (PI) controllers. Therefore we start with a detailed overview of RF cavity phase and amplitude PI control. To develop feedback controllers we must consider the coupled beam-cavity-RF source system. We consider only the ω_0 frequency component of the beam, $A_b(t)\cos(\omega_0 t + \theta_b(t))$, an RF driving current of the form $I_c(t) = A_c(t)\cos(\omega t + \theta_c(t))$, and a cavity field of the form $V_{\mathrm{cav}}(t) = A_{\mathrm{cav}}(t)\cos(\omega t + \theta_{\mathrm{cav}}(t))$. The single second order differential equations describing the cavity dynamics, (9.27), can then be simplified to two coupled, linear, first order differential equations:

$$\dot{I} = -\omega_{\frac{1}{2}} I - \Delta\omega Q + \beta_{I,c} I_c + \beta_{I,b} I_b, \tag{9.30}$$

$$\dot{Q} = \Delta\omega I - \omega_{\frac{1}{2}} Q + \beta_{Q,c} Q_c + \beta_{Q,b} Q_b, \tag{9.31}$$

where $\Delta\omega = \omega - \omega_0$ is the difference between RF generator and cavity resonance frequencies, $\omega_{\frac{1}{2}} = \omega_0 / 2Q_L$, and the I and Q quantities represent

$$I(t) = A_{\mathrm{cav}}(t)\cos(\theta_{\mathrm{cav}}(t)), \quad I_c(t) = A_c(t)\cos(\theta_c(t)), \quad I_b(t) = A_b(t)\cos(\theta_b(t)), \tag{9.32}$$

$$Q(t) = A_{\mathrm{cav}}(t)\sin(\theta_{\mathrm{cav}}(t)), \quad Q_c(t) = A_c(t)\sin(\theta_c(t)), \quad Q_b(t) = A_b(t)\sin(\theta_b(t)), \tag{9.33}$$

from which amplitudes and phases can be calculated according to

$$A_\bullet(t) = \sqrt{I_\bullet^2(t) + Q_\bullet^2(t)}, \qquad \theta_\bullet(t) = \arctan\left(\frac{Q_\bullet(t)}{I_\bullet(t)}\right). \tag{9.34}$$

Equations (9.30), (9.31) can be written in the compact linear form

$$\dot{\mathbf{x}} = A\mathbf{x} + B_c\mathbf{u} + B_b\mathbf{d}, \quad \mathbf{x} = \begin{bmatrix} I \\ Q \end{bmatrix}, \quad A = \begin{bmatrix} -\omega_{\frac{1}{2}} & -\Delta\omega \\ \Delta\omega & -\omega_{\frac{1}{2}} \end{bmatrix}, \quad \mathbf{u} = \begin{bmatrix} I_c \\ Q_c \end{bmatrix}, \quad \mathbf{d} = \begin{bmatrix} I_b \\ Q_b \end{bmatrix}, \tag{9.35}$$

where $\mathbf{u}$ refers to the control vector, and the beam itself, $\mathbf{d}$, is thought of as a disturbance. The goal of RF feedback control is typically to maintain the cavity field as given by $\mathbf{x}$ at a desired set point thereby ensuring proper acceleration of the beam. In addition to providing a simple, linear approximation of the dynamics of the beam, cavity, and RF generator system, (9.35) is very useful because a typical digital RF system does not have access to the raw cavity voltage signal $V_{\mathrm{cav}}(t)$, but rather to $I_{\mathrm{cav}}(t)$ and $Q_{\mathrm{cav}}(t)$, which are provided by down sampling the cavity field signal. For example, at the Los Alamos Neutron Science Center (LANSCE) linear accelerator, $f_{RF} = 201.25\,\mathrm{MHz}$ RF signals of the form $V_{\mathrm{cav}}(t) = A_{\mathrm{cav}}(t)\cos(2\pi f_{RF} t + \theta_{\mathrm{cav}}(t))$ are first mixed down via local oscillators to signals at an intermediate frequency

$f_{IF} = 25\,\text{MHz}$, of the form $A_{\text{cav}}(t)\cos(2\pi f_{IF}t + \theta_{\text{cav}}(t))$, which can be expanded in the I, Q formalism as:

$$
\begin{aligned}
A_{\text{cav}}&(t)\cos(2\pi f_{IF}t + \theta_{\text{cav}}(t)) \\
&= \underbrace{A_{\text{cav}}(t)\cos(\theta_{\text{cav}}(t))}_{I_{\text{cav}}(t)}\cos(2\pi f_{IF}t) - \underbrace{A_{\text{cav}}(t)\sin(\theta_{\text{cav}}(t))}_{Q_{\text{cav}}(t)}\cos(2\pi f_{IF}t) \\
&= I_{\text{cav}}(t)\cos(2\pi f_{IF}t) - Q_{\text{cav}}(t)\cos(2\pi f_{IF}t)\sin(2\pi f_{IF}t).
\end{aligned}
\tag{9.36}
$$

Then, by oversampling the signal (9.36) at the rate $f_s = 4 \times f_{IF}$, the analog to digital converter (ADC) collects samples at time steps $nt = \frac{n}{f_s}$:

$$
V_{\text{cav}}\left(\frac{n}{4f_{IF}}\right) = I_{\text{cav}}\left(\frac{n}{4f_{IF}}\right)\cos\left(\frac{n\pi}{2}\right) - Q_{\text{cav}}\left(\frac{n}{4f_{IF}}\right)\sin\left(\frac{n\pi}{2}\right), \tag{9.37}
$$

directly receiving the samples:

$$
\{I_{\text{cav}}(0),\, -Q_{\text{cav}}(t_s),\, -I_{\text{cav}}(2t_s),\, Q_{\text{cav}}(3t_s),\, \dots\}. \tag{9.38}
$$

The job of the RF control system is to maintain the cavity fields at amplitude and phase set points, $A_s(t)$ and $\theta_s(t)$, respectively, which translate into I and Q set points: $I_s(t) = A_s(t)\cos(\theta_s(t))$, $Q_s(t) = A_s(t)\sin(\theta_s(t))$. The most simple typical RF feedback control system first compares the cavity I and Q signals to their set points and calculates error signals $I_e(t) = I_{\text{cav}}(t) - I_s(t)$, $Q_e(t) = Q_{\text{cav}}(t) - Q_s(t)$, and then performs proportional-integral feedback control of the form

$$
I_c(t) = -k_p I_e(t) - k_i \int_0^t I_e(\tau)d\tau, \qquad Q_c(t) = -k_p Q_e(t) - k_i \int_0^t Q_e(\tau)d\tau.
\tag{9.39}
$$

Typically particle accelerators are pulsed at rates of tens to hundreds of Hz. For example, in the LANSCE accelerator, the RF drive power is turned on for 1 ms at a rate of 120 Hz. Once RF is turned on, cavity fields build up and reach steady state within a few hundred microseconds, after which the cavities are ready to accelerate the beam, whose sudden arrival perturbs the cavity fields, as shown in Fig. 9.5.

Although the initial I and Q set points are in the forms of smooth ramps, as seen from the shape of the cavity field amplitude in Fig. 9.5, once the field has reached steady state and before the beam has arrived, the set points are fixed in order to maintain a precise field amplitude and phase offset of the bunches relative to the RF zero crossing. Therefore, in what follows we consider the cavity set points only after steady state has been reached and they are therefore constants of the form:

$$
I_s(t \geq T_{\text{rise}}) \equiv I_s(T_{\text{rise}}) = I_r, \quad Q_s(t \geq T_{\text{rise}}) \equiv Q_s(T_{\text{rise}}) = Q_r. \tag{9.40}
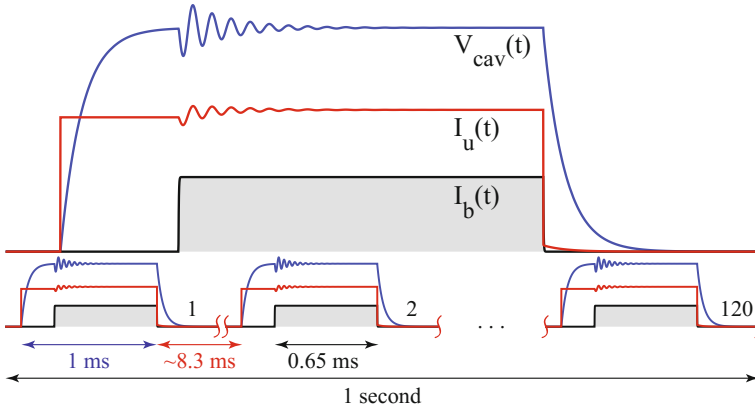$$

Fig. 9.5 The RF source, $I_c(t)$, is turned on at a rate of 120 Hz, for $\sim$1 ms per pulse. The beam, $I_b(t)$, arrives around $\sim$350 μs into the pulse after the cavity field, $V_{\text{cav}}(t)$, has had time to settle. The beam's arrival disrupts the cavity field's steady state

Plugging the feedback (9.39) into the cavity dynamics (9.35) and rewriting the dynamics in terms of the error variables, we are then left with the closed loop system

$$\dot{\mathbf{x}}_e = A\mathbf{x}_e + A\mathbf{x}_r - k_p B_c \mathbf{x}_e - k_i B_c \int_0^t \mathbf{x}_e(\tau)d\tau + B_b \mathbf{d}, \quad \mathbf{x}_e = \begin{bmatrix} I_e \\ Q_e \end{bmatrix}, \quad \mathbf{x}_r = \begin{bmatrix} I_r \\ Q_r \end{bmatrix}. (9.41)$$

Taking the Laplace transform of both sides of (9.41), assuming that we are at steady state so that $\mathbf{x}_e(T_{\text{rise}}) = 0$, we get

$$s\mathbf{X}_e(s) = A\mathbf{X}_e(s) + \frac{1}{s}A\mathbf{x}_r - k_p B_c \mathbf{X}_e(s) - \frac{1}{s}k_i B_c \mathbf{X}_e(s) + B_d \mathbf{D}(s)$$
$$\Longrightarrow$$
$$\mathbf{X}_e(s) = \left(s^2 I - s\left(k_p B_c - A\right) + k_i B_c\right)^{-1} \left(A\mathbf{x}_r + s B_d \mathbf{D}(s)\right). \qquad (9.42)$$

The gains, k_i and k_p of the simple PI feedback control loop are then tuned in order to maintain minimal error despite the disturbances $A\mathbf{x}_r$ and $s B_d \mathbf{D}(s)$. The constant term $A\mathbf{x}_r$ is due to the natural damping of the RF cavity and is easily compensated for. The more important and more difficult to deal with term is $s B_d \mathbf{D}(s)$, which, in the time domain is proportional to the derivative of the beam current $B_d \dot{\mathbf{d}}(t)$. Because the beam is typically ramped up to an intense current very quickly (tens of microseconds) or consists of an extremely short pulse, the derivative term is extremely disruptive to the cavity field phase and amplitude. Some typical beam current and bunch timing profiles are shown in Fig. 9.6. Currently LCLS is able to accelerate 1 nC during extremely powerful $\sim$3 μs RF pulses, with a separation of 8.3 ms between bunches. The European XFEL is pushing orders of magnitude beyond the LCLS bunch timing with 1 nC pulses separated by only 220 ns. This is extremely challenging for an
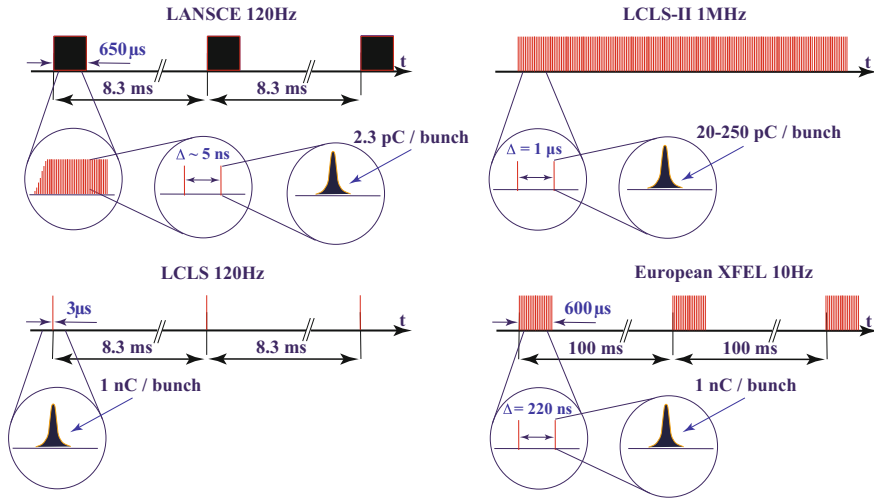
Fig. 9.6 Beam current time profiles of several accelerators are shown

RF system which must maintain field amplitude and phase set points and recover between bunches. The proposed MaRIE accelerator will push this problem another order of magnitude in attempting to accelerate high charge pulses with only ~2.5 ns of separation.

Although the PI controller used in (9.41) can theoretically hold the error $\mathbf{x}_e$ arbitrarily close to zero arbitrarily fast by choosing large enough gains k_i and k_p relative to the magnitude of the beam disturbance $\left\| B_d \dot{\mathbf{d}}(t) \right\|$, in practice all control gains are limited by actuator saturation, response time, and most importantly, delay in the feedback loop. A typical RF feedback loop is shown in Fig. 9.7 and may experience as much as 5 μs of round trip delay, which is an large delay relative to beam transient times.

Consider for example the following scalar, delay system, where the goal is to quickly drive $x(t)$ to zero from an arbitrary initial condition, but only being able to do so based on a controller which uses a delayed measurement of $x(t)$, $x(t - D)$. Considering a simply proportional feedback control, $u = -kx$, for the system

$$\dot{x}(t) = u(x(t - D)) \quad \implies \quad \dot{x}(t) = -kx(t - D), \tag{9.43}$$

taking Laplace transforms we get

$$sX(s) - x(0) = -ke^{-Ds}X(s) \quad \implies \quad X(s) = \frac{x(0)}{s + ke^{-Ds}}. \tag{9.44}$$

If we assume the delay is small, $D \ll 1$, we can approximate $e^{-Ds} \approx 1 - Ds$, invert the Laplace transform and get the solution
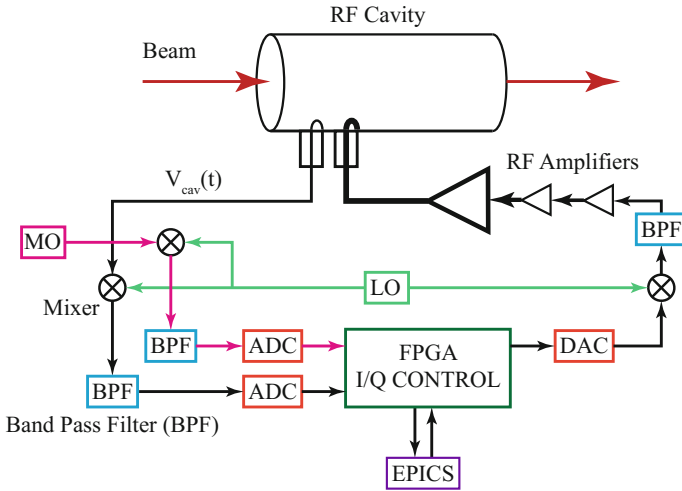
Fig. 9.7 Typical digital RF control setup with signals coming from the cavity into the digital FPGA-based controller and then back out through a chain of amplifiers
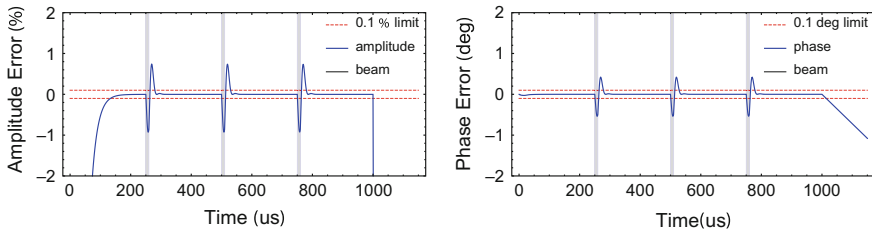

Fig. 9.8 Cavity field errors with frequency shift, RF power droop, beam loading, and simple proportional-integral feedback control

$$x(t) = x(0)e^{\gamma t}, \quad \gamma = \frac{-k}{1 - kD}, \tag{9.45}$$

which exponentially converges to 0 for $\gamma < 0$, requiring that k satisfy $\frac{1}{D} > k > 0$, a limit on possible stabilizing values of the feedback control gain. If our system (9.43) had an external disturbance, $d(t)$ the gain limit would be a major limitation in terms of compensating for large or fast $d(t)$.

Because of such limitations, a feedback only LLRF system's response to beam loading would typically look like the results shown in Fig. 9.8, where each intense beam pulse causes a large deviation of the accelerating field's voltage from the design phase and amplitude, which must be restored before the next bunch can be properly accelerated.

9.2 Advanced Control and Tuning Topics

For problems which can be accurately modeled, such as systems that do not vary with time and for which extensive, detailed diagnostics exist, there are many powerful optimization methods such as genetic algorithms (GA), which can be used during the design of an accelerator by performing extremely large searches over parameter space [29]. Such multi-objective genetic algorithms (MOGA) have been applied for the design of radio frequency cavities [30], photoinjectors [31], damping rings [32], storage ring dynamics [33], lattice design [34], neutrino factory design [35], simultaneous optimization of beam emittance and dynamic aperture [36], free electron laser linac drivers [37] and various other accelerator physics applications [38]. One extension of MOGA is multi-objective particle swarm optimization, has been used for emittance reduction [39]. Brute force approaches such as GA and MOGA search over the entire parameter space of interest and therefore result in global optimization, however, such model-based approaches are only optimal relative to the specific model which they are using, which in practice rarely exactly matches the actual machine when it is built. Differences are due to imperfect models, uncertainty, and finite precision of construction. Therefore, actual machines settings undergo extensive tuning and tweaking in order to reach optimal performance. Recently efforts have been made to implement a GA method on-line for the minimization of beam size at SPEAR3 [40]. Robust conjugate direction search (RCDS) is another optimization method. RCDS is model independent, but at the start of optimization in must learn the conjugate directions of the given system, and therefore is not applicable to quickly time-varying systems [41, 42]. Optimization of nonlinear storage ring dynamics via RCDS and particle swarm has been performed online [43].

Although many modern, well behaved machines can possibly be optimized with any of the methods mentioned above, and once at steady state, the operation may not require the fast re-tuning future light sources will require algorithms with an ability to quickly switch between various operating conditions and to handle quickly time-varying systems, based only on scalar measurements, rather than a detailed knowledge of the system dynamics, when compensating for complex collective effects. If any of the methods above were used, they would have to be repeated every time component settings were significantly changed and it is highly unlikely that they would converge or be well behaved during un-modeled, fast time-variation of components. Therefore, a model-independent feedback-based control and tuning procedure is required which can function on nonlinear and time varying systems with many coupled components.

The type of tuning problems that we are interested in have recently been approached with powerful machine learning methods [15, 44], which are showing very promising results. However, these methods require large training sets in order to learn how to reach specific machine set points, and interpolate in between. For example, if a user requests a combination of beam energy, pulse charge, and bunch length, which was not a member of a neural network-based controller's learning set, the achieved machine performance is not predictable. Furthermore, machine

components slowly drift with time and un-modeled disturbances are present and limit any learning-based algorithm's abilities. Extremum seeking (ES) is a simple, local, model-independent algorithm for accelerator tuning, whose speed of convergence allows for the optimization and real-time tracking of many coupled parameters for time-varying nonlinear systems. Because ES is model independent, robust to noise, and has analytically guaranteed parameter bounds and update rates, it is useful for real time feedback in actual machines. One of the limitations of ES is that it is a local optimizer which can possible be trapped in local minima.

It is our belief that the combination of ES and machine learning methods will be a powerful method for quickly tuning FELs between drastically different user desired beam and light properties. For example, once a deep neural network (NN) has learned a mapping of machine settings to light properties for a given accelerator based on collected machine data, it can be used to quickly bring the machine within a local proximity of the required settings for a given user experiment. However, the performance will be limited by the fact that the machine changes with time, that the desired experiment settings were not in the training data, and un-modeled disturbances. Therefore, once brought within a small neighborhood of the required settings via NN, ES can be used to achieve local optimal tuning, which can also continuously re-tune to compensate for un-modeled disturbances and time variation of components. In the remainder of this chapter we will focus on the ES method, giving a general overview of the procedure and several simulation and in-hardware demonstrations of applications of the method. Further details on machine learning approaches can be found in [15, 44] and the references within.

9.3 Introduction to Extremum Seeking Control

The Extremum seeking method described in this chapter is a recently developed general approach for the stabilization of noisy, uncertain, open-loop unstable, time-varying systems [6, 7]. The main benefits of this approach are:

1. The method can tune many parameters of unknown, nonlinear, open-loop unstable systems, simultaneously.
2. The method is robust to measurement noise and external disturbances and can track quickly time-varying parameters.
3. Although operating on noisy and analytically unknown systems, the parameter updates have analytically guaranteed constraints, which make it safe for in-hardware implementation.

This method has been implemented in simulation to automatically tune large systems of magnets and RF set points to optimize beam parameters [11], it has been utilized in hardware at the proton linear accelerator at the Los Alamos Neutron Science Center to automatically tune two RF buncher cavities to maximize the RF system's beam acceptance, based only on a noisy measurement of beam current [12], it has been utilized at the Facility for Advanced Accelerator Experimental Tests, to

non-destructively predict electron bunch properties via a coupling of simulation and machine data [13], it has been utilized for bunch compressor design [45], and has been used for the automated tuning of magnets in a time-varying lattice to continuously minimize betatron oscillations at SPEAR3 [8]. Furthermore, analytic proofs of convergence for the method are available for constrained systems with general, non-differentiable controllers [9, 10].

9.3.1 Physical Motivation

It has been shown that unexpected stability properties can be achieved in dynamic systems by introducing fast, small oscillations. One example is the stabilization of the vertical equilibrium point of an inverted pendulum by quickly oscillating the pendulum's pivot point. Kapitza first analyzed these dynamics in the 1950s [46]. The ES approach is in some ways related to such vibrational stabilization as high frequency oscillations are used to stabilize desired points of a system's state space and to force trajectories to converge to these points. This is done by creating cost functions whose minima correspond to the points of interest, allowing us to tune a large family of systems without relying on any models or system knowledge. The method even works for unknown functions, where we do not choose which point of the state space to stabilize, but rather are minimizing an analytically unknown function whose noisy measurements we are able to sample.

To give an intuitive 2D overview of this method, we consider finding the minimum of an unknown function $C(x, y)$. We propose the following scheme:

$$\frac{dx}{dt} = \sqrt{\alpha\omega} \cos\left(\omega t + kC(x, y)\right) \tag{9.46}$$

$$\frac{dy}{dt} = \sqrt{\alpha\omega} \sin\left(\omega t + kC(x, y)\right). \tag{9.47}$$

Note that although $C(x, y)$ enters the argument of the adaptive scheme, we do not rely on any knowledge of the analytic form of $C(x, y)$, we simply assume that it's value is available for measurement at different locations (x, y).

The velocity vector,

$$\mathbf{v} = \left(\frac{dx}{dt}, \frac{dy}{dt}\right) = \sqrt{\alpha\omega} \left[\cos\left(\theta(t)\right), \sin\left(\theta(t)\right)\right], \tag{9.48}$$

$$\theta(t) = \omega t + kC(x(t), y(t)), \tag{9.49}$$

has constant magnitude, $\|\mathbf{v}\| = \sqrt{\alpha\omega}$, and therefore the trajectory $(x(t), y(t))$ moves at a constant speed. However, the rate at which the direction of the trajectories' heading changes is a function of ω, k, and $C(x(t), y(t))$ expressed as:
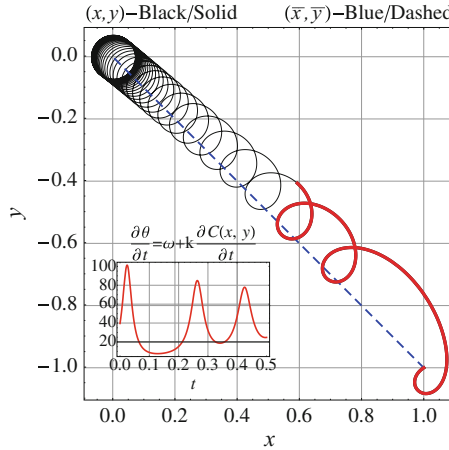
Fig. 9.9 The subfigure in the bottom left shows the rotation rate, $\frac{\partial \theta}{\partial t} = \omega + \frac{\partial C(x,y)}{\partial t}$, for the part of the trajectory that is bold red, which takes place during the first 0.5 s of simulation. The rotation of the parameters' velocity vector $\mathbf{v}(t)$ slows down when heading towards the minimum of $C(x,y) = x^2 + y^2$, at which time $k\frac{\partial C}{\partial t} < 0$, and speeds up when heading away from the minimum, when $k\frac{\partial C}{\partial t} > 0$. The system ends up spending more time heading towards and approaches the minimum of $C(x,y)$

$$\frac{d\theta}{dt} = \omega + k\left(\frac{\partial C}{\partial x}\frac{dx}{dt} + \frac{\partial C}{\partial y}\frac{dy}{dt}\right). \tag{9.50}$$

Therefore, when the trajectory is heading in the correct direction, towards a decreasing value of $C(x(t), y(t))$, the term $k\frac{\partial C}{\partial t}$ is negative so the overall turning rate $\frac{\partial \theta}{\partial t}$ (9.50), is decreased. On the other hand, when the trajectory is heading in the wrong direction, towards an increasing value of $C(x(t), y(t))$, the term $k\frac{\partial C}{\partial t}$ is positive, and the turning rate is increased. On average, the system ends up approaching the minimizing location of $C(x(t), y(t))$ because it spends more time moving towards it than away.

The ability of this direction-dependent turning rate scheme is apparent in the simulation of system (9.46), (9.47), in Fig. 9.9. The system, starting at initial location $x(0) = 1, y(0) = -1$, is simulated for 5 s with update parameters $\omega = 50, k = 5, \alpha = 0.5$, and $C(x,y) = x^2 + y^2$. We compare the actual system's (9.46), (9.47) dynamics with those of a system performing gradient descent:

$$\frac{d\bar{x}}{dt} \approx -\frac{k\alpha}{2}\frac{\partial C(\bar{x}, \bar{y})}{\partial \bar{x}} = -k\alpha\bar{x} \tag{9.51}$$

$$\frac{d\bar{y}}{dt} \approx -\frac{k\alpha}{2}\frac{\partial C(\bar{x}, \bar{y})}{\partial \bar{y}} = -k\alpha\bar{y}, \tag{9.52}$$

whose behavior our system mimics on average, with the difference

$$\max_{t \in [0,T]} \|(x(t), y(t)) - (\bar{x}(t), \bar{y}(t))\| \tag{9.53}$$

made arbitrarily small for any value of T, by choosing arbitrarily large values of ω.

Towards the end of the simulation, when the system's trajectory is near the origin, $C(x, y) \approx 0$, and the dynamics of (9.46), (9.47) are approximately

$$\frac{\partial x}{\partial t} \approx \sqrt{\alpha \omega} \cos(\omega t) \implies x(t) \approx \sqrt{\frac{\alpha}{\omega}} \sin(\omega t) \tag{9.54}$$

$$\frac{\partial y}{\partial t} \approx \sqrt{\alpha \omega} \sin(\omega t) \implies y(t) \approx -\sqrt{\frac{\alpha}{\omega}} \cos(\omega t), \tag{9.55}$$

a circle of radius $\sqrt{\frac{\alpha}{\omega}}$, which is made arbitrarily small by choosing arbitrarily large values of ω. Convergence towards a maximum, rather than a minimum is achieved by replacing k with $-k$.

9.3.2 General ES Scheme

For general tuning, we consider the problem of locating an extremum point of the function $C(\mathbf{p}, t) : \mathbb{R}^n \times \mathbb{R}^+ \to \mathbb{R}$, for $\mathbf{p} = (p_1, \ldots, p_n) \in \mathbb{R}^n$, when only a noise-corrupted measurement $y(t) = C(\mathbf{p}, t) + n(t)$ is available, with the analytic form of C unknown. For notational convenience, in what follows we sometimes write $C(\mathbf{p})$ or just C instead of $C(\mathbf{p}(t), t)$.

The explanation presented in the previous section used $\sin(\cdot)$ and $\cos(\cdot)$ functions for the x and y dynamics to give circular trajectories. The actual requirement for convergence is for an independence, in the frequency domain, of the functions used to perturb different parameters. In what follows, replacing $\cos(\cdot)$ with $\sin(\cdot)$ throughout makes no difference.

Theorem 1 *Consider the setup shown in Fig. 9.10 (for maximum seeking we replace k with $-k$):*

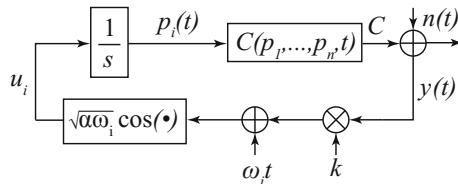$$\dot{p}_i = \sqrt{\alpha \omega_i} \cos(\omega_i t + k y), \qquad y = C(\mathbf{p}, t) + n(t) \tag{9.56}$$

Fig. 9.10 Tuning of the ith component p_i of $\mathbf{p} = (p_1, \ldots, p_n) \in \mathbb{R}^n$. The symbol $\frac{1}{s}$ denotes the Laplace Transform of an integrator, so that in the above diagram $p_i(t) = p_i(0) + \int_0^t u_i(\tau) d\tau$

where $\omega_i = \omega_0 r_i$ such that $r_i \neq r_j \; \forall i \neq j$ and $n(t)$ is additive noise. The trajectory of system (9.56) approaches the minimum of $C(\mathbf{p}, t)$, with its trajectory arbitrarily close to that of

$$\dot{\bar{\mathbf{p}}} = -\frac{k\alpha}{2}\nabla C, \quad \bar{\mathbf{p}}(0) = \mathbf{p}(0) \tag{9.57}$$

with the distance between the two decreasing as a function of increasing ω_0. Namely, for any given $T \in [0, \infty)$, any compact set of allowable parameters $\mathbf{p} \in K \subset \mathbb{R}^m$, and any desired accuracy δ, there exists $\omega_0^\star$ such that for all $\omega_0 > \omega_0^\star$, the distance between the trajectory $\mathbf{p}(t)$ of (9.56) and $\bar{\mathbf{p}}(t)$ of (9.57) satisfies the bound

$$\max_{\mathbf{p}, \bar{\mathbf{p}} \in K, t \in [0, T]} \|\mathbf{p}(t) - \bar{\mathbf{p}}(t)\| < \delta. \tag{9.58}$$

Remark 1 One of the most important features of this scheme is that on average the system performs a gradient descent of the actual, unknown function C despite feedback being based only on its noise corrupted measurement $y = C(\mathbf{p}, t) + n(t)$.

Remark 2 The stability of this scheme is verified by the fact that an addition of an un-modeled, possibly destabilizing perturbation of the form $\mathbf{f}(\mathbf{p}, t)$ to the dynamics of $\dot{\mathbf{p}}$ results in the averaged system:

$$\dot{\bar{\mathbf{p}}} = \mathbf{f}(\bar{\mathbf{p}}, t) - \frac{k\alpha}{2}\nabla C, \tag{9.59}$$

which may be made to approach the minimum of C, by choosing $k\alpha$ large enough relative to the values of $\left\|(\nabla C)^T\right\|$ and $\|\mathbf{f}(\bar{\mathbf{p}}, t)\|$.

Remark 3 In the case of a time-varying max/min location $\mathbf{p}^\star(t)$ of $C(\mathbf{p}, t)$, there will be terms of the form:

$$\frac{1}{\sqrt{\omega}}\left|\frac{\partial C(\mathbf{p}, t)}{\partial t}\right|, \tag{9.60}$$

which are made to approach zero by increasing ω. Furthermore, in the analysis of the convergence of the error $\mathbf{p}_e(t) = \mathbf{p}(t) - \mathbf{p}^\star(t)$ there will be terms of the form:

$$\frac{1}{k\alpha}\left|\frac{\partial C(\mathbf{p}, t)}{\partial t}\right|. \tag{9.61}$$

Together, (9.60) and (9.61) imply the intuitively obvious fact that for systems whose time-variation is fast, in which the minimum towards which we are descending is quickly varying, both the value of ω and of the product $k\alpha$ must be larger than for the time-invariant case.

Remark 4 In the case of different parameters having vastly different response characteristics and sensitivities (such as when tuning both RF and magnet settings in the same scheme), the choices of k and α may be specified differently for each component p_i, as k_i and α_i, without change to the above analysis.

Fig. 9.11 ES for simultaneous stabilization and optimization of an unknown, open-loop unstable system based on a noise corrupted scalar measurement
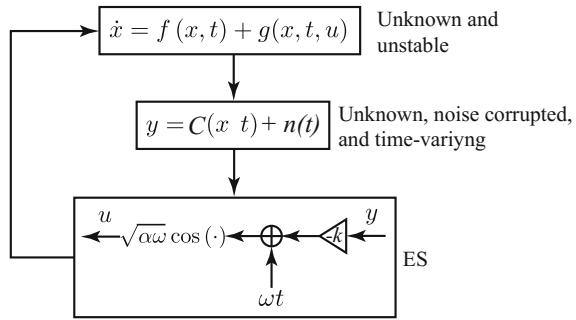
A more general form of the scheme for simultaneous stabilization and optimization of an n-dimensional open-loop unstable system with analytically unknown noise-corrupted output function $C(\mathbf{x}, t)$ is shown in Fig. 9.11, but will not be discussed in detail here.

9.3.3 ES for RF Beam Loading Compensation

The ES method described above has been used both in simulation and optimization studies and has been implemented in hardware in accelerators. We now return to the RF problem described in Sect. 9.1.6, where we discussed the fact that due to delay-limited gains and power limitations, the sudden transient caused by beam loading greatly disturbs the RF fields of accelerating cavities which must be re-settled to within prescribed bounds before the next bunches can be brought in for acceleration. ES has been applied to this beam loading problem in the LANSCE accelerator via high speed field programmable gate array (FPGA).

In order to control the amplitude and phase of the RF cavity accelerating field, the $I(t) = A(t) \cos(\theta(t))$ and $Q(t) = A(t) \sin(\theta(t))$ components of the cavity voltage signal were sampled as described in Sect. 9.1.6, at a rate of 100 MS/s during a 1000 μs RF pulse. The detected RF signal was then broken down into 10 μs long sections and feed forward $I_{\mathrm{ff},j}(n)$ and $Q_{\mathrm{ff},j}(n)$ control outputs were generated for each 10 μs long section, as shown in Fig. 9.12.

Remark 5 In the discussion and figures that follow, we refer to $I_{\mathrm{cav}}(t)$ and $Q_{\mathrm{cav}}(t)$ simply as $I(t)$ and $Q(t)$.

The iterative extremum seeking was performed via finite difference approximation of the ES dynamics:

$$\frac{x(t+dt) - x(t)}{dt} \approx \frac{dx}{dt} = \sqrt{\alpha\omega}\cos(\omega t + kC(x,t)), \qquad (9.62)$$
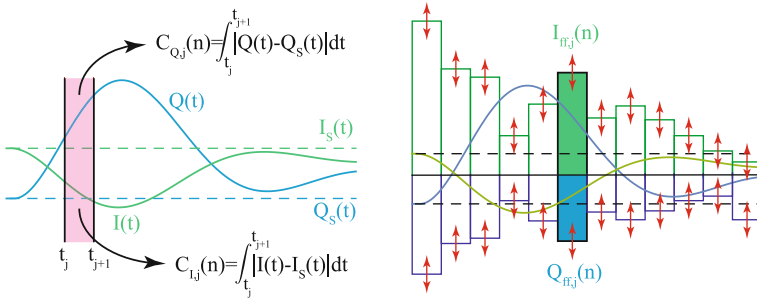
by updating the feedforward signals according to

Fig. 9.12 Top: Iterative scheme for determining I and Q costs during 1–$10\,\mu$s intervals. **Bottom**: ES-based feedforward outputs for beam loading transient compensation

$$I_{\text{ff},j}(n+1) = I_{\text{ff},j}(n) + \Delta\sqrt{\alpha\omega}\cos\left(\omega n\Delta + kC_{I,j}(n)\right), \qquad (9.63)$$

and

$$Q_{\text{ff},j}(n+1) = Q_{\text{ff},j}(n) + \Delta\sqrt{\alpha\omega}\sin\left(\omega n\Delta + kC_{Q,j}(n)\right), \qquad (9.64)$$

where the individual I and Q costs were calculated as

$$C_{I,j}(n) = \int_{t_j}^{t_{j+1}} |I(t) - I_s(t)|\, dt, \qquad (9.65)$$

$$C_{Q,j}(n) = \int_{t_j}^{t_{j+1}} |Q(t) - Q_s(t)|\, dt. \qquad (9.66)$$

Note that although the I_j and Q_j parameters were updated on separate costs, they were still dithered with different functions, $\sin(\cdot)$ and $\cos(\cdot)$, to help maintain orthogonality in the frequency domain. The feed forward signals were then added to the PI and static feed forward controller outputs. Running at a repetition rate of 120 Hz, the feedback converges within several hundred iterations or a few seconds.

These preliminary experimental results are shown in Fig. 9.13 and summarized in Table 9.1. The maximum, rms, and average values are all calculated during a 150 μs window which includes the beam turn on transient to capture the worst case scenario. The ES-based scheme is a $>2\times$ improvement over static feed-forward in terms of maximum errors and a $>3\times$ improvement in terms of rms error. With the currently used FPGA, the ES window lengths can be further reduced from 10 μs to 10 ns and with the latest FPGAs down to 1 ns, which will greatly improve the ES performance.
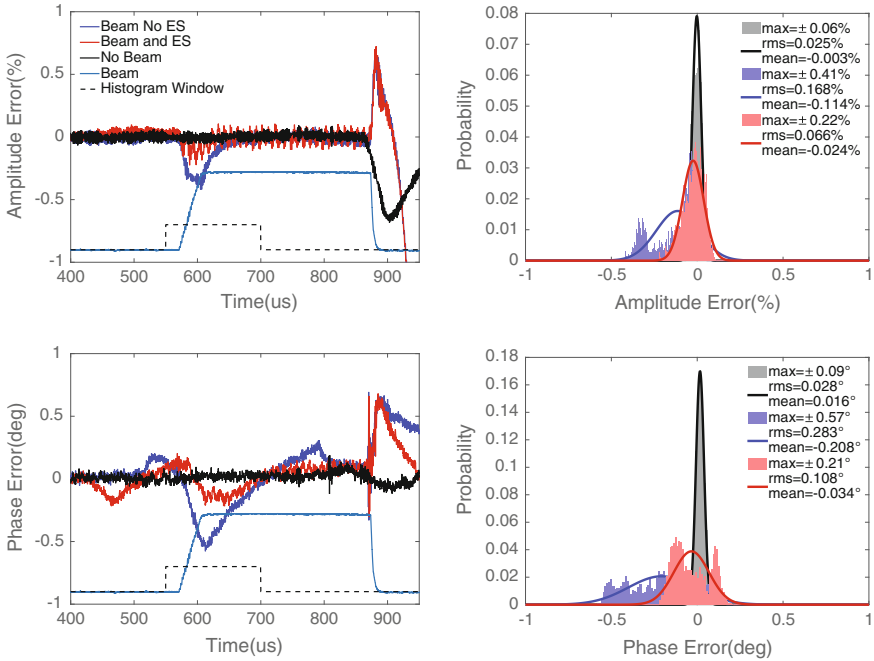
Fig. 9.13 Phase and amplitude errors shown before, during, and after beam turn-on transient. The histogram data shown is collected during the dashed histogram window, and cleaned up via 100 point moving average after raw data was sampled at 100 MS/s. **Black**: Beam OFF. **Blue**: Beam ON, feedback, and static feed-forward only. **Red**: Beam ON, feedback, static feed-forward, and iterative ES feed-forward

Table 9.1 ES performance during beam turn on transient

	No Beam	Beam, No ES	Beam and ES
max A error (%)	±0.06	±0.41	±0.22
rms A error (%)	0.025	0.168	0.066
mean A error (%)	−0.003	−0.114	−0.024
max θ error (%)	±0.09	±0.57	±0.21
rms θ error (%)	0.028	0.283	0.108
mean θ error (%)	0.016	−0.208	−0.034

9.3.4 ES for Magnet Tuning

ES has also been tested in hardware for magnet-based beam dynamics tuning, as described in Sect. 9.1.1. At the SPEAR3 synchrotron at LCLS, ES was used for continuous re-tuning of the eight parameter system shown in Fig. 9.14, in which the delay, pulse width, and voltage of two injection kickers, K_1 and K_2, as well as
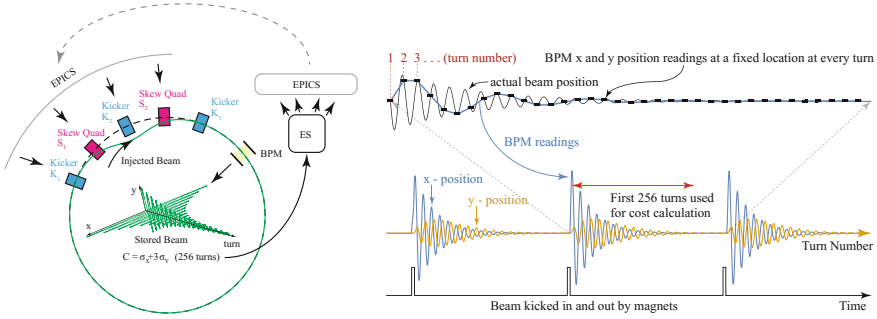
Fig. 9.14 Kicker magnets and skew quadrupole magnets. When the beam is kicked in and out of orbit, because of imperfect magnet matching, betatron oscillations occur, which are sampled at the BPM every time the beam completes a turn around the machine

the current of two skew quadrupoles S_1 and S_2, were tuned in order to optimize the injection kicker bump match, minimizing betatron oscillations. At SPEAR3, we simultaneously tuned 8 parameters: (1). $p_1 = K_1$ delay. (2). $p_2 = K_1$ pulse width. (3). $p_3 = K_1$ voltage. (4). $p_4 = K_2$ delay. (5). $p_5 = K_2$ pulse width. (6). $p_6 = K_2$ voltage. (7). $p_7 = S_1$ current. (8). $p_8 = S_2$ current. The parameters are illustrated in Figs. 9.14, 9.15. While controlling the voltage for the kicker magnets K_1, K_2, and the current for the skew quadrupole magnets S_1, S_2, in each case a change in the setting resulted in a change in magnetic field strength.

The cost function used for tuning was a combination of the horizontal, σ_x, and vertical, σ_y, variance of beam position monitor readings over 256 turns, the
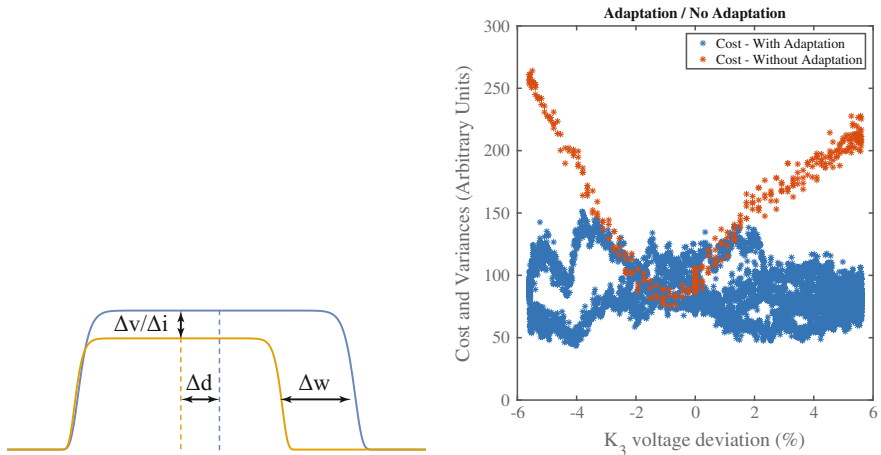

Fig. 9.15 Left: Kicker magnet delay (d), pulse width (w), and voltage (v) were adaptively adjusted, as well as the skew quadrupole magnet currents (i). **Right**: Comparison of beam quality with and without adaptation

minimization of which resulted in decreased betatron oscillations,

$$C = \sqrt{\frac{1}{256} \sum_{i=1}^{256} (x(i) - \bar{x})^2} + \sqrt{\frac{9}{256} \sum_{i=1}^{256} (y(i) - \bar{y})^2}$$
$$= \sigma_x + 3\sigma_y, \tag{9.67}$$

where the factor of 3 was added to increase the weight of the vertical oscillations, which require tighter control since the vertical beam size is much smaller and therefore users are more sensitive to vertical oscillations.

The cost was based on beam position monitor (BPM) measurements in the SPEAR3 ring based on a centroid x and y position of the beam recorded at each revolution, as shown in Fig. 9.14. Variances σ_x and σ_y were calculated based on this data, as in (9.67). Feedback was implemented via the experimental physics and industrial control system (EPICS) [47].

To demonstrate the scheme's ability to compensate for an uncertain, time-varying perturbation of the system, we purposely varied the voltage (and therefore resulting magnetic field strength) of the third kicker magnet, $K_3(t)$. The kicker voltage was varied sinusoidally over a range of $\pm 6\%$ over the course of 1.5 h, which is a very dramatic and fast change relative to actual machine parameter drift rates and magnitudes. The ES scheme was implemented by setting parameter values, kicking an electron beam out and back into the ring, and recording beam position monitor data for a few thousand turns. Based on this data the cost was calculated as in (9.67), based on a measurement of the horizontal and vertical variance of beam position monitor readings. The magnet settings were then adjusted, the beam was kicked again, and a new cost was calculated. This process was repeated and the cost was iteratively, continuously minimized.

Figure 9.14 shows the cost, which is a function of betatron oscillation, versus magnet setting $K_3(t)$, with and without ES feedback. For large magnetic field deviations, the improvement is roughly a factor of 2.5.

9.3.5 ES for Electron Bunch Longitudinal Phase Space Prediction

The Facility for Advanced Accelerator Experimental Tests (FACET) at SLAC National Accelerator Laboratory produces high energy electron beams for plasma wakefield acceleration [48]. For these experiments, precise control of the longitudinal beam profile is very important. FACET uses an x-band transverse deflecting cavity (TCAV) to streak the beam and measure the bunch profile (Fig. 9.16a). Although the TCAV provides an accurate measure of the bunch profile, it is a destructive measurement; the beam cannot be used for plasma wakefield acceleration (PWFA) once it has been streaked. In addition, using the TCAV to measure the bunch profile requires
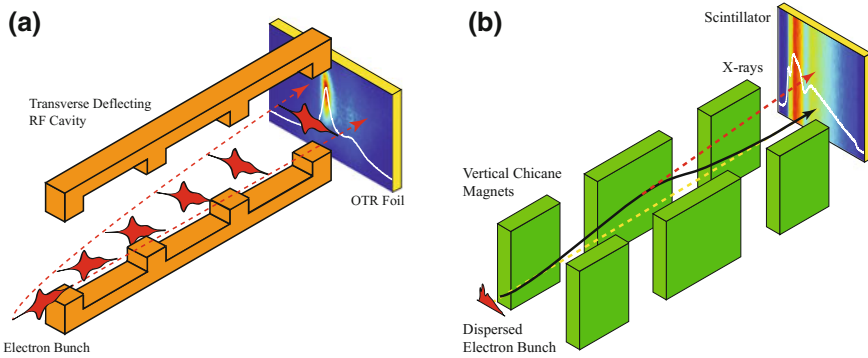
Fig. 9.16 The energy spectrum is recorded as the electron bunch passes through a series of magnets and radiates x-rays. The intensity distribution of the X-rays is correlated to the energy spectrum of the electron beam (**a**). This non-destructive measurement is available at all times, and used as the input to the ES scheme, which is then matched by adaptively tuning machine parameters in the simulation. For the TCAV measurement, the electron bunch is passed through a high frequency (11.4 GHz) RF cavity with a transverse mode, in which it is streaked and passes through a metallic foil (**b**). The intensity of the optical transition radiation (OTR) is proportional to the longitudinal charge density distribution. This high accuracy longitudinal bunch profile measurement is a destructive technique

adjusting the optics of the final focus system to optimize the resolution and accuracy of measurement. This makes it a time consuming process and prevents on-the-fly measurements of the bunch profile during plasma experiments.

There are two diagnostics that are used as an alternative to the TCAV that provide information about the longitudinal phase space in a non-destructive manner. The first is a pyrometer that captures optical diffraction radiation (ODR) produced by the electron beam as it passes through a hole in a metal foil. The spectral content of the ODR changes with bunch length. The pyrometer is sensitive to the spectral content and the signal it collects is proportional to $1/\sigma_z$, where σ_z is the bunch length. The pyrometer is an excellent device for measuring variation in the shot-to-shot bunch profile but provides no information about the shape of the bunch profile or specific changes to shape. The second device is a non-destructive energy spectrometer consisting of a half-period vertical wiggler located in a region of large horizontal dispersion. The wiggler produces a streak of X-rays with an intensity profile that is correlated with the dispersed beam profile. There X-rays are intercepted by a scintillating YAG crystal and imaged by a CCD camera (Fig. 9.16b). The horizontal profile of the x-ray streak is interpreted as the energy spectrum of the beam [49].

The measured energy spectrum is observed to correlate with the longitudinal bunch profile in a one-to-one manner if certain machine parameters, such as chicane optics, are fixed. To calculate the beam properties based on an energy spectrum measurement, the detected spectrum is compared to a simulated spectrum created with the 2D longitudinal particle tracking code, LiTrack [50]. The energy spread of short electron bunches desirable for plasma wakefield acceleration can be uniquely
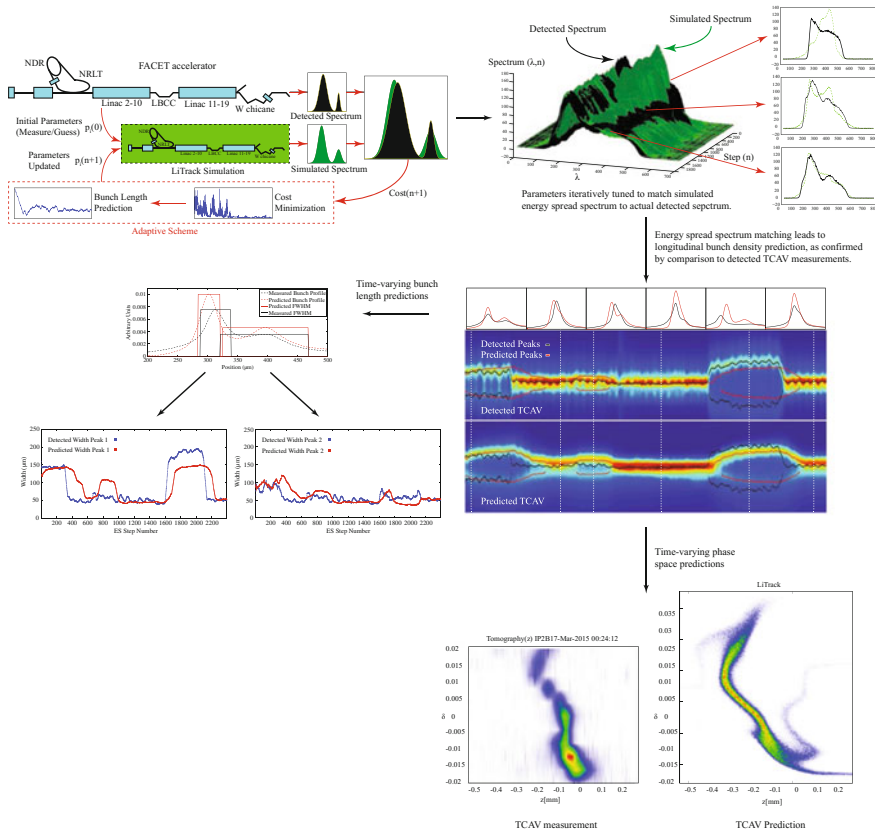
Fig. 9.17 ES scheme at FACET

correlated to the beam profile if all of the various accelerator parameters which influence the bunch profile and energy spread are accounted for accurately. Unfortunately, throughout the 2 km facility, there exist systematic phase drifts of various high frequency devices, mis-calibrations, and time-varying uncertainties due to thermal drifts. Therefore, in order to effectively and accurately relate an energy spectrum to a bunch profile, a very large parameter space must be searched and fit by LiTrack, which effectively limits and prevents the use of the energy spectrum measurement as a real time measurement of bunch profile.

Figures 9.16 and 9.17 show the overall setup of the tuning procedure at FACET. A simulation of the accelerator, LiTrack is run in parallel to the machines operation. The simulation was initialized with guesses and any available measurements of actual machine settings, $\mathbf{p} = (p_1, \ldots, p_n)$. We emphasize that these are only guesses because even measured values are noisy and have arbitrary phase shift errors. The electron beam in the actual machine was accelerated and then passed through a series of deflecting magnets, as shown in Figs. 9.16b and 9.17, which created X-rays, whose

intensity distribution can be correlated to the electron bunch density via LiTrack. This non-destructive measurement is available at all times, and used as the input to the ES scheme, which is then matched by adaptively tuning machine parameters in the simulation. Once the simulated and actual spectrum were matched, certain beam properties could be predicted by the simulation.

Each parameter setting has its own influence on electron beam dynamics, which in turn influenced the separation, charge, length, etc, of the leading and trailing electron bunches.

The cost that our adaptive scheme was attempting to minimize was then the difference between the actual, detected spectrum, and that predicted by LiTrack:

$$C(\mathbf{x}, \hat{\mathbf{x}}, \mathbf{p}, \hat{\mathbf{p}}, t) = \int \left| \tilde{\psi}(\mathbf{x}, \mathbf{p}, t, \nu) - \hat{\psi}(\hat{\mathbf{x}}, \hat{\mathbf{p}}, t, \nu) \right|^2 d\nu, \qquad (9.68)$$

in which $\tilde{\psi}(\mathbf{x}, \mathbf{p}, t, \nu)$ was a noisy measurement of the actual, time-varying (due to phase drift, thermal cycling…) energy spectrum, and $\hat{\psi}(\hat{\mathbf{x}}, \hat{\mathbf{p}}, t, \nu)$ was the LiTrack, simulated spectrum, $\mathbf{x}(t) = (x_1(t), \ldots, x_n(t))$ represents various aspects of the beam, such as bunch length, beam energy, bunch charge, etc. at certain locations throughout the accelerator, $\mathbf{p}(t) = (p_1(t), \ldots, p_n(t))$ represents various time-varying uncertain parameters of the accelerator itself, such as RF system phase drifts and RF field amplitudes throughout the machine, $\mathbf{x}(t)$ are approximated by their simulated estimates $\hat{\mathbf{x}}(t) = (\hat{x}_1(t), \ldots, \hat{x}_n(t))$ and actual system parameters, $\mathbf{p}(t)$, are approximated by virtual parameters $\hat{\mathbf{p}}(t) = (\hat{p}_1(t), \ldots, \hat{p}_n(t))$.

The problem was then to minimize the measurable, but analytically unknown function C, by adaptively tuning the simulation parameters $\hat{\mathbf{p}}$. The hope was that, by finding simulation machine settings which resulted in matched spectrums, we would also match other properties of the real and simulated beams, something we could not simply do by setting the simulation parameters to the exact machine settings, due to unknowns, such as time-varying, arbitrary phase shifts.

LiTrackES simulates large components of FACET as single elements. The critical elements of the simulation are the North Damping Ring (NDR) which sets the initial bunch parameters including the bunch length and energy spread, the North Ring to Linac (NRTL) which is the first of three bunch compressors, Linac Sectors 2–10 where the beam is accelerated and chirped, the second bunch compressor in Sector 10 (LBCC), Linac Sectors 11–19 where the beam is again accelerated and chirped, and finally the FACET W-chicane which is the third and final bunch compressor.

We calibrated the LiTrackES algorithm using simultaneous measurements of the energy spectrum and bunch profile while allowing a set of unknown parameters to converge. After convergence we left a subset of these calibrated parameters fixed, as they are known to vary slowly or not at all and performed our tuning on a much smaller subset of the parameters:

- p_1: NDR bunch length
- p_2: NRTL energy offset
- p_3: NRTL compressor amplitude

- p_4: NRTL chicane T_{566}
- p_5: Phase Ramp

"Phase ramp" refers to a net phase of the NDR and NRTL RF systems with respect to the main linac RF. Changing the phase ramp parameter results in a phase set offset in the linac relative to some desired phase.

LiTrackES, the combination of ES and LiTrack, as demonstrated, is able to provide a quasi real time estimate of many machine and electron beam properties which are either inaccessible or require destructive measurements. We plan to improve the convergence rate of LiTrackES by fine tuning the adaptive scheme's parameters, such as the gains k_i, perturbing amplitudes α_i and dithering frequencies ω_i. Furthermore, we plan on taking advantage of several simultaneously running LiTrackES schemes, which can communicate with each other in an intelligent way, and each of which has slightly different adaptive parameters/initial parameter guesses, which we believe can greatly increase both the rate and accuracy of the convergence. Another major goal is the extension of this algorithm from monitoring to tuning. We hope to one day utilize LiTrackES as an actual feedback to the machine settings in order to tune for desired electron beam properties.

9.3.6   ES for Phase Space Tuning

For the work described here, a measured XTCAV image was utilized and compared to the simulated energy and position spread of an electron bunch at the end of the LCLS as simulated by LiTrack. The electron bunch distribution is given by a function $\rho(\Delta E, \Delta z)$ where $\Delta E = E - E_0$ is energy offset from the mean or design energy of the bunch and $\Delta z = z - z_0$ is position offset from the center of the bunch. We worked with two distributions:

$$\text{XTCAV measured} : \rho_{\text{TCAV}}(\Delta E, \Delta z),$$
$$\text{LiTrack simulated} : \rho_{\text{LiTrack}}(\Delta E, \Delta z).$$

These distributions were then integrated along the E and z projections in order to calculate 1D energy and charge distributions:

$$\rho_{E,\text{TCAV}}(\Delta E), \quad \rho_{z,\text{TCAV}}(\Delta z),$$
$$\rho_{E,\text{LiTrack}}(\Delta E), \quad \rho_{z,\text{LiTrack}}(\Delta z).$$

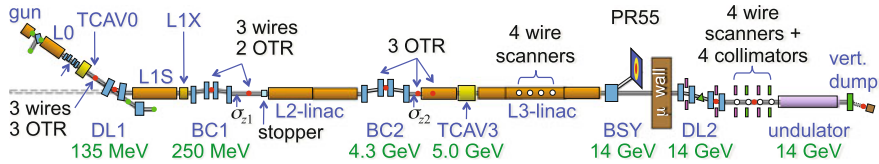Finally, the energy and charge spread distributions were compared to create cost values:

Fig. 9.18 Components of the LCLS beamline

$$C_E = \int \left[\rho_{E,\text{TCAV}}(\Delta E) - \rho_{E,\text{LiTrack}}(\Delta E) \right]^2 d\Delta E, \tag{9.69}$$

$$C_z = \int \left[\rho_{z,\text{TCAV}}(\Delta z) - \rho_{z,\text{LiTrack}}(\Delta z) \right]^2 d\Delta z, \tag{9.70}$$

whose weighted sum was combined into a single final cost:

$$C = w_E C_E + w_z C_z. \tag{9.71}$$

Iterative extremum seeking was then performed via finite difference approximation of the ES dynamics (Fig. 9.18):

$$\frac{\mathbf{p}(t + dt) - \mathbf{p}(t)}{dt} \approx \frac{d\mathbf{p}}{dt} = \sqrt{\alpha\omega} \cos(\omega t + kC(\mathbf{p}, t)), \tag{9.72}$$

by updating LiTrack model parameters, $\mathbf{p} = (p_1, \ldots, p_m)$, according to

$$p_j(n + 1) = p_j(n) + \Delta\sqrt{\alpha\omega_j} \cos\left(\omega_j n \Delta + kC(n)\right), \tag{9.73}$$

where the previous step's cost is based on the previous simulation's parameter settings,

$$C(n) = C(\mathbf{p}(n)). \tag{9.74}$$

The parameters being tuned were:

1. **L1S phase**: typically drifts continuously and is repeatedly corrected via an invasive phase scan. Within some limited range a correct bunch length can be maintained by the existing feedback system. This parameter is used for optimizing machine settings and FEL pulse intensity. When the charge off the cathode is changed, L1S phase must be adjusted manually.
2. **L1X phase**: must be changed if L1S phase is changed significantly. This linearizes the curvature of the beam.
3. **BC1 energy**: control bunch length and provides feedback to L1S amplitude.
4. **L2 phase**: drifts continuously with temperature, is a set of multiple Klystrons, all of which cycle in amplitude and phase. Feedback is required to introduce the correct energy chirp required for BC2 peak current/bunch length set point. Tuned to maximize FEL intensity and minimize jitter.
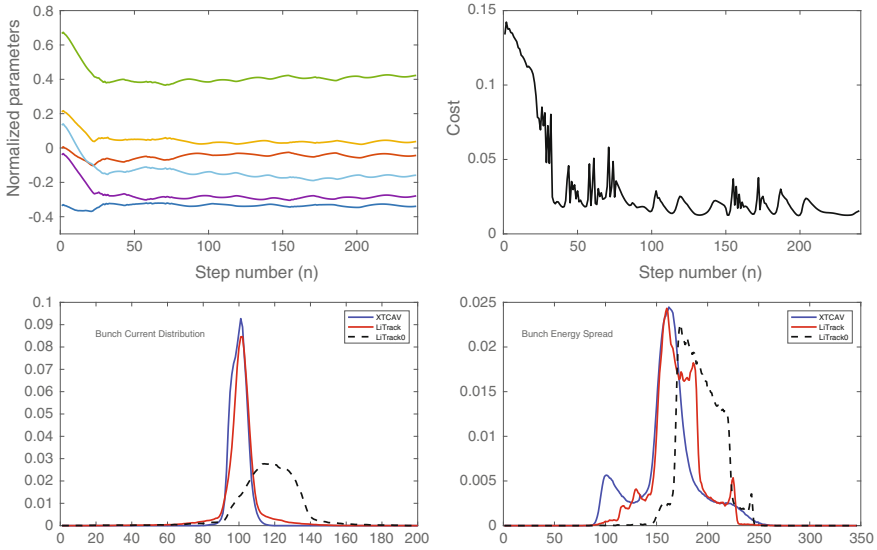
Fig. 9.19 Parameter convergence and cost minimization for matching desired bunch length and energy spread profiles

5. **BC2 energy**: drifts due to Klystron fluctuations, must be changed to optimize FEL pulse intensity for exotic setups.
6. **L3 phase**: drifts continuously with temperature, based on a coupled system of many Klystrons.

Machine tuning work has begun with general analytic studies as well as simulation-based algorithm development focused on the LCLS beam line, using SLACs LiTrack software, a code which captures most aspects of the electron beams phase space evolution and incorporates noise representative of operating conditions. The initial effort focused on developing ES-based auto tuning of the electron beam's bunch length and energy spread by varying LiTrack parameters in order to match LiTrack's output to an actual TCAV measurement taken from the accelerator by tuning bunch compressor energies and RF phases. The results are shown in Figs. 9.19 and 9.20. Running at a repetition rate of 120 Hz, the simulated feedback would have converged within 2 s on the actual LCLS machine.

Preliminary results have demonstrated that ES is a powerful tool with the potential to automatically tune an FEL between various bunch properties such as energy spread and bunch length requirements by simultaneously tuning multiple coupled parameters, based only on a TCAV measurement at the end of the machine. Although the simulation results are promising, It remains to be seen what the limitations of the method are in the actual machine in terms of getting stuck in local minima and time of convergence. We plan on exploring the extent of parameter and phase space through which we can automatically move.
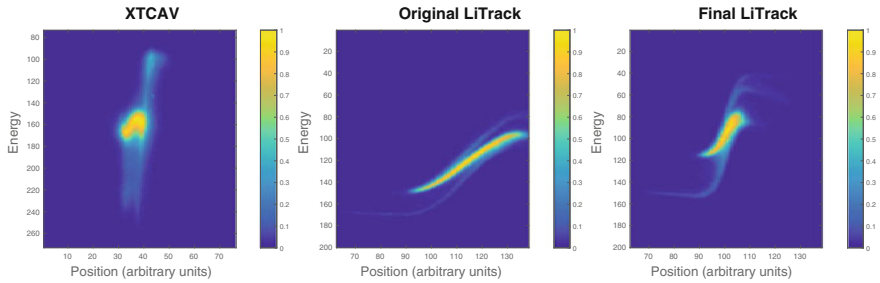
Fig. 9.20 Measured XTCAV, original LiTrack and final, converged LiTrack energy versus position phases space of the electron bunch shown

9.4 Conclusions

The intense bunch charges, extremely short bunch lengths, and extremely high energies of next generation FEL beams result in complex collective effects which couple transverse and longitudinal dynamics and therefore all of the RF and magnet systems and their influence on the quality of the light being produced. These future light sources, especially 4th generation FELs, face major challenges both in achieving extremely tight constraints on beam quality and in quickly tuning between various, exotic experimental setups. We have presented a very brief and simple introduction to some of the beam dynamics important to accelerators and have introduced some methods for achieving better beam quality and faster tuning. Based on preliminary results it is our belief is that a combination of machine learning and advanced feedback methods such as ES have great potential towards meeting the requirements of future light sources. Such a combination of ES and machine learning has recently been demonstrated in a proof of principle experiment at the Linac-Coherent Light Source FEL [51]. During this experiment we quickly trained a simple neural network to obtain an estimate of a complex and time-varying parameter space, mapping longitudinal electron beam phase space (energy vs time) to machine parameter settings. For a target longitudinal phase space, we used the neural network to give us an initial guess of the required parameter settings which brought us to within a neighborhood of the correct parameter settings, but did not give a perfect match. We then used ES-based feedback to zoom in on and track the actual optimal time-varying parameters settings.

References

1. T.O. Raubenheimer, Technical challenges of the LCLS-II CW X-RAY FEL, in *Proceedings of the International Particle Accelerator Conference*, Richmond, VA, USA (2015)
2. C. Schmidt et al., Recent developments of the European XFEL LLRF system, in *Proceedings of the International Particle Accelerator Conference*, Shanghai, China (2013)

3. J. Bradley III, A. Scheinker, D. Rees, R.L. Sheffield, High power RF requirements for driving discontinuous bunch trains in the MaRIE LINAC, in *Proceedings of the Linear Particle Accelerator Conference*, East Lansing, MI, USA (2016)
4. R. Sheffield, Enabling cost-effective high-current burst-mode operation in superconducting accelerators. Nucl. Instrum. Methods Phys. Res. A **758**, 197–200 (2015)
5. R. Akre, A. Brachmann, F.J. Decker, Y.T. Ding, P. Emma, A.S. Fisher, R.H. Iverson, Tuning of the LCLS Linac for user operation, in *Conf. Proc. C110328: 2462-2464*, 2011 (No. SLAC-PUB-16643) (SLAC National Accelerator Laboratory, 2016)
6. A. Scheinker, Ph.D. thesis, University of California, San Diego, Nov 2012
7. A. Scheinker, Model independent beam tuning, in *Proceedings of the 4th International Particle Accelerator Conference*, Beijing, China (2012)
8. A. Scheinker, X. Huang, J. Wu, Minimization of betatron oscillations of electron beam injected into a time-varying lattice via extremum seeking. IEEE Trans. Control Syst. Technol. (2017). https://doi.org/10.1109/TCST.2017.2664728
9. A. Scheinker, D. Scheinker, Bounded extremum seeking with discontinuous dithers. Automatica **69**, 250–257 (2016)
10. A. Scheinker, D. Scheinker, Constrained extremum seeking stabilization of systems not affine in control. Int. J. Robust Nonlinear Control (to appear) (2017). https://doi.org/10.1002/rnc.3886
11. A. Scheinker, X. Pang, L. Rybarcyk, Model-independent particle accelerator tuning. Phys. Rev. Accel. Beams **16**(10), 102803 (2013)
12. A. Scheinker, S. Baily, D. Young, J. Kolski, M. Prokop, In-hardware demonstration of model-independent adaptive tuning of noisy systems with arbitrary phase drift. Nucl. Instrum. Methods Phys. Res. Sect. A **756**, 30–38 (2014)
13. A. Scheinker, S. Gessner, Adaptive method for electron bunch profile prediction. Phys. Rev. Accel. Beams **18**(10), 102801 (2015)
14. S.G. Biedron, A. Edelen, S. Milton, Advanced controls for accelerators, in *Compact EUV & X-ray Light Sources* (Optical Society of America, 2016), p. EM9A-3
15. A.L. Edelen, S.G. Biedron, B.E. Chase, D. Edstrom, S.V. Milton, P. Stabile, Neural networks for modeling and control of particle accelerators. IEEE Trans. Nucl. Sci. **63**(2), 878–897 (2016)
16. Y.B. Kong, M.G. Hur, E.J. Lee, J.H. Park, Y.D. Park, S.D. Yang, Predictive ion source control using artificial neural network for RFT-30 cyclotron. Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrom. Detect. Assoc. Equip. **806**, 55–60 (2016)
17. M. Buchanan, Depths of learning. Nat. Phys. **11**(10), 798–798 (2015)
18. X. Huang, J. Corbett, J. Safranek, J. Wu, An algorithm for online optimization of accelerators. Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrom. Detect. Assoc. Equip. **726**, 77–83 (2013)
19. T.P. Wangler, *RF Linear Accelerators* (Wiley, 2008)
20. R. Ruth, Single particle dynamics in circular accelerators, in *AIP Conference Proceedings*, vol. 153, No. SLAC-PUB-4103 (1986)
21. H. Wiedemann, *Particle Accelerator Physics* (Springer, New York, 1993)
22. D.A. Edwards, M.J. Syphers, *An Introduction to the Physics of High Energy Accelerators* (Wiley-VCH, 2004)
23. S.Y. Lee, *Accelerator Physics* (World Scientific Publishing, 2004)
24. M. Reiser, *Theory and Design of Charged Particle Beams* (Wiley-VCH, 2008)
25. C.X. Wang, A. Chao, Transfer matrices of superimposed magnets and RF cavity, No. SLAC-AP-106 (1996)
26. M.G. Minty, F. Zimmermann, *Measurement and Control of Charged Particle Beams* (Springer, 2003)
27. J.C. Slater, Microwave electronics. Rev. Modern Phys. **18**(4) (1946)
28. J. Jackson, *Classical Electrodynamics* (Wiley, NJ, 1999)
29. M. Borland, Report No. APS LS-287 (2000)
30. R. Hajima, N. Taked, H. Ohashi, M. Akiyama, Optimization of wiggler magnets ordering using a genetic algorithm. Nucl. Instrum. Methods Phys. Res. Sect. A **318**, 822 (1992)

31. I. Bazarov, C. Sinclair, Multivariate optimization of a high brightness dc gun photo injector. Phys. Rev. ST Accel. Beams **8**, 034202 (2005)
32. L. Emery, in *Proceedings of the 21st Particle Accelerator Conference*, Knoxville, 2005 (IEEE, Piscataway, NJ, 2005)
33. M. Borland, V. Sajaev, L. Emery, A. Xiao, in *Proceedings of the 23rd Particle Accelerator Conference*, Vancouver, Canada, 2009 (IEEE, Piscataway, NJ, 2009)
34. L. Yang, D. Robin, F. Sannibale, C. Steier, W. Wan, Global optimization of an accelerator lattice using multiobjective genetic algorithms. Nucl. Instrum. Methods Phys. Res. Sect. A **609**, 50 (2009)
35. A. Poklonskiy, D. Neuffer, Evolutionary algorithm for the neutrino factory front end design. Int. J. Mod. Phys. A **24**, 5 (2009)
36. W. Gao, L. Wang, W. Li, Simultaneous optimization of beam emittance and dynamic aperture for electron storage ring using genetic algorithm. Phys. Rev. ST Accel. Beams **14**, 094001 (2011)
37. R. Bartolini, M. Apollonio, I.P.S. Martin, Multiobjective genetic algorithm optimization of the beam dynamics in linac drivers for free electron lasers. Phys. Rev. ST Accel. Beams **15**, 030701 (2012)
38. A. Hofler, B. Terzic, M. Kramer, A. Zvezdin, V. Morozov, Y. Roblin, F. Lin, C. Jarvis, Innovative applications of genetic algorithms to problems in accelerator physics. Phys. Rev. ST Accel. Beams **16**, 010101 (2013)
39. X. Huang, J. Safranek, Nonlinear dynamics optimization with particle swarm and genetic algorithms for SPEAR3 emittance upgrade. Nucl. Instrum. Methods Phys. Res. Sect. A **757**, 48–53 (2014)
40. K. Tian, J. Safranek, Y. Yan, Machine based optimization using genetic algorithms in a storage ring. Phys. Rev. Accel. Beams **17**, 020703 (2014)
41. X. Huang, J. Corbett, J. Safranek, J. Wu, An algorithm for online optimization of accelerators. Nucl. Instrum. Methods Phys. Res. A **726**, 77–83 (2013)
42. H. Ji, S. Wang, Y. Jiao, D. Ji, C. Yu, Y. Zhang, X. Huang, Discussion on the problems of the online optimization of the luminosity of BEPCII with the robust conjugate direction search method, in *Proceedings of the International Particle Accelerator Conference*, Shanghai, China (2015)
43. X. Huang, J. Safranek, Online optimization of storage ring nonlinear beam dynamics. Phys. Rev. ST Accel. Beams **18**(8), 084001 (2015)
44. A.L. Edelen et al., Neural network model of the PXIE RFQ cooling system and resonant frequency response (2016). arXiv:1612.07237
45. B.E. Carlsten, K.A. Bishofberger, S.J. Russell, N.A. Yampolsky, Using an emittance exchanger as a bunch compressor. Phys. Rev. Spec. Top.-Accel. Beams **14**(8), 084403 (2011)
46. P.L. Kapitza, Dynamic stability of a pendulum when its point of suspension vibrates. Sov. Phys. JETP **21**, 588–592 (1951)
47. R.L. Dalesio, J.O. Hill, M. Kraimer, S. Lewis, D. Murray, S. Hunt, W. Watson, M. Clausen, J. Dalesio, The experimental physics and industrial control system architecture: past, present, and future. Nucl. Instrum. Methods Phys. Res. Sect. A **352**(1), 179–184 (1994)
48. M.J. Hogan, T.O. Raubenheimer, A. Seryi, P. Muggli, T. Katsouleas, C. Huang, W. Lu, W. An, K.A. Marsh, W.B. Mori, C.E. Clayton, C. Joshi, Plasma wakefield acceleration experiments at FACET. New J. Phys. **12**, 055030 (2010)
49. J. Seeman, W. Brunk, R. Early, M. Ross, E. Tillman, D. Walz, SLC energy spectrum monitor using synchrotron radiation. SLAC-PUB-3495 (1986)
50. K. Bane, P. Emma, *LiTrack*: a fast longitudinal phase space tracking code. SLAC-PUB-11035 (2005)
51. A. Scheinker, A. Edelen, D. Bohler, C. Emma, and A. Lutman, Demonstration of model-independent control of the longitudinal phase space of electron beams in the Linac-coherent light source with Femtosecond resolution. Phys. Rev. Lett. **121**(4), 044801 (2018)