Nuclear Instruments and Methods in Physics Research A 756 (2014) 30–38

Contents lists available at ScienceDirect

Nuclear Instruments and Methods in Physics Research A

journal homepage: www.elsevier.com/locate/nima

In-hardware demonstration of model-independent adaptive tuning of noisy systems with arbitrary phase drift

Alexander Scheinker*, Scott Baily, Daniel Young, Jeffrey S. Kolski, Mark Prokop

Los Alamos National Laboratory, Los Alamos, NM 87545, USA

ARTICLE INFO

Article history:
Received 7 March 2014
Received in revised form
1 April 2014
Accepted 11 April 2014
Available online 22 April 2014

Keywords:
Linac
Adaptive control
Feedback control
Model independent control
RF cavity
Phase drift

ABSTRACT

In this work, an implementation of a recently developed model-independent adaptive control scheme, for tuning uncertain and time varying systems, is demonstrated on the Los Alamos linear particle accelerator. The main benefits of the algorithm are its simplicity, ability to handle an arbitrary number of components without increased complexity, and the approach is extremely robust to measurement noise, a property which is both analytically proven and demonstrated in the experiments performed. We report on the application of this algorithm for simultaneous tuning of two buncher radio frequency (RF) cavities, in order to maximize beam acceptance into the accelerating electromagnetic field cavities of the machine, with the tuning based only on a noisy measurement of the surviving beam current downstream from the two bunching cavities. The algorithm automatically responds to arbitrary phase shift of the cavity phases, automatically re-tuning the cavity settings and maximizing beam acceptance. Because it is model independent it can be utilized for continuous adaptation to time-variation of a large system, such as due to thermal drift, or damage to components, in which the remaining, functional components would be automatically re-tuned to compensate for the failing ones. We start by discussing the general model-independent adaptive scheme and how it may be digitally applied to a large class of multi-parameter uncertain systems, and then present our experimental results.

Published by Elsevier B.V.

1. Introduction

Real systems cannot be perfectly modeled, there are always uncertainties such as time-varying component characteristics (linear/nonlinear operating regimes/damage), environmental conditions (temperature variation and resulting cable length changes), and un-modeled coupling between various sub-systems. For control purposes, the impact of a model's inaccuracy varies from one system to the next, and depends on the variable of interest, the component being controlled, and the sensitivity and level of diagnostics available.

Particle accelerators, in which beams must have very specific characteristics, such as energy and bunch length, and whose age and technology level of hardware greatly varies from one facility to the next, present a number of control problems, both the types that can be accurately modeled, and others, which require adaptive techniques. Traditionally, controllers have been designed as proportional-integral-derivative (PID) [1] feedback loops, whose design depends on a combination of an accurate system model and experience in order to optimally tune the control gains. Depending on the control problem at hand, extensive modeling may be the correct solution, and has been widely successful in the design of control systems, such as recently for final focus stabilization and multi-bunch beam instability compensation [2,3].

However, there are some systems for which modeling is not as successful, and the control of which has moved toward model-independent, learning, or adaptive techniques. The active disturbance rejection technique is one in which a system's nonlinearity, coupling, and noise are dealt with as external disturbances, which the controller attempts to estimate in real time and reject, and has been utilized for compensation of Lorentz force detuning and microphonics in radio-frequency (RF) resonant cavities [4]. Neural networks combined with traditional proportional-integral (PI) techniques have been utilized for stabilization of electron beam energy in free electron laser (FEL) Linacs [5,6]. Classical PI, as well as iterative learning control, and model predictive control, has been utilized in low level RF systems for amplitude and phase control of resonant RF cavities [7].

Due to uncertainties such as hysteresis of magnets, although a machine's operating settings can be written down, and magnet current sources may be returned to exactly their previous values, the actual magnetic fields corresponding to these currents may change. The scheme presented here could be utilized to automatically adjust such magnet settings, until the actual, desired

* Corresponding author. Tel.: +1 914 582 0932.
E-mail address: alexscheinker@gmail.com (A. Scheinker).

http://dx.doi.org/10.1016/j.nima.2014.04.026
0168-9002/Published by Elsevier B.V.

magnetic fields were achieved, as measured by emittance scans, loss monitors, beam position monitors, or any other available diagnostic of beam characteristics.

One particular problem faced by many particle accelerators is arbitrary phase drift, due to factors such as cable length-changes caused by temperature variation. In practice, the arbitrary phase shift of accelerator RF systems results in lengthy phase scan procedures which must be carried out periodically due to constant, slowly varying drift, and following major hardware changes.

In this paper, we demonstrate a model-independent, adaptive control scheme [8,9], which can automatically tune RF systems despite unknown, arbitrary phase shift, as the feedback automatically re-tunes parameters even in time-varying conditions. We demonstrate this algorithm in hardware, on the Los Alamos linear accelerator at the Los Alamos Neutron Science Center (LANSCE), by simultaneously tuning the phases of two RF buncher cavities, in order to maximize the acceptance of the beam into the drift tube Linac (DTL) for acceleration, based only on a noisy cost which is the square of the measured surviving beam current half way down the length of the accelerator, which was implemented entirely with EPICS records, reading a current monitor, and writing motor movement directions.

Remark 1. The benefit of the model independence of this scheme was especially evident in this application because we were not able to directly assign phase values to the buncher cavities, rather we were limited to sending discrete step numbers and directions ($+$ or $-$) to phase changing packages, in which stepper motors would then mechanically adjust the phase, and this step number-to-phase adjustment had different characteristics depending on the position of the stepper motor itself, for example, depending on the location of the motor, a control output of ± 5 steps could result in either $\pm 1^\circ$ or $\pm 0.5^\circ$ actual change in cavity phase, and the phase package of each cavity had its own such characteristics.

The adaptive feedback controller used here is a form of extremum seeking (ES) first conceptualized in the 1920s [10,11], and is closely related to vibrational control [12,13], in which, for example, the upper equilibrium point of a pendulum may be stabilized by rapidly vertically oscillating the pendulum's pivot point. Recently, ES has been extended to perform stabilization and optimization of unknown, possibly unstable systems [14] and further improved and modified into the form presented here [8,9], in which the control effort and convergence rates have analytically known bounds.

We start by giving a brief overview of the controller and how it is implemented. The main benefits of the control algorithm demonstrated here are the following:

- The scheme is completely model independent, it does not rely on an estimate of the system being controlled, and is able to respond to time-variation of the unknown system.
- Despite operating on an unknown system, the update rates and the step sizes prescribed by the controller are analytically known a priori as all the unknowns enter the control scheme as the arguments of $\sin(\cdot)$ or $\cos(\cdot)$ terms.
- As demonstrated in hardware, in this work, the scheme is robust to noise, because on average, the noise does not contribute to parameter dynamics, as long as it does not have an extremely large component at exactly the controller's perturbing frequencies.
- The scheme is very simple and an arbitrary number of coupled parameters can be tuned simultaneously without an increase in complexity for an increasing number of parameters.

The first step in applying this method is to choose a cost, C, whose maximum or minimum corresponds to desired machine performance or beam properties, such as $C=I^2$, where I is a current monitor reading, which we would like to maximize at the end of a particle accelerator, corresponding to the maximum amount of current surviving through the accelerator. Next, we choose the parameters, $p_1,\ldots,p_m$, which we will tune, in order to maximize or minimize the cost C, such as RF cavity phases or amplitude settings, or quadrupole magnet current set points. We start with initial settings $p_1(1),\ldots,p_m(1)$ as are usually chosen, based either on experience or a simplified physics model. Note that typically $C=C(p_1,\ldots,p_m, \hat{p}_{m+1},\ldots,\hat{p}_{m_2},t)$, where the $\hat{p}_j$ terms may be parameters which we are unaware of or unable to control and the time dependence may be due to time variation of the environment or system components. We then record the cost $C(1)$ corresponding to initial parameter settings, and continue by iteratively updating parameter settings according to

$$p_i(n+1)=p_i(n)+\delta_t\alpha_i\sqrt{\omega_i}\,\cos(\omega_i n\delta_t-k_iC(n)). \tag{1}$$

In the above update law, the k_i can be thought of as the controller gains, increasing k_i corresponds to faster convergence and larger control effort, switching $-k_i<0$ to $k_i>0$ causes the system to converge towards a minimum rather than a maximum of C. The α_i are proportional to the size of the perturbation of a parameter's set point, increasing α_i increases both the convergence speed and the size of the steady state oscillation about the extremum value, once optimization has been achieved. The required value of δ_t becomes evident when considering that the parameter update law, Eq. (1), is chosen based on the adaptive feedback scheme described in Refs. [8,9]. Rewriting the update (1) as

$$\frac{p_i(n+1)-p_i(n)}{\delta_t}=\alpha\sqrt{\omega_i}\,\cos(\omega_i n\delta_t-kC(n)), \tag{2}$$

we see that Eq. (2) is a finite difference approximation of

$$\frac{dp_i}{dt}=\alpha\sqrt{\omega_i}\,\cos(\omega_i t-kC(t)), \tag{3}$$

when δ_t is small enough, such as $\delta_t<2\pi/40\max\{\omega_i\}$, in which case, in the absence of the cost function, C, the parameters would perform smooth motion, with at least 40 steps per full oscillation.

Remark 2. Instead of Eq. (3) we may use the adaptive scheme

$$\frac{dp_i}{dt}=\alpha\sqrt{\omega_i}\,\sin(\omega_i t-kC), \tag{4}$$

or

$$\frac{dp_1}{dt}=\alpha\sqrt{\omega_1}\,\cos(\omega_1 t-kC),$$
$$\frac{dp_2}{dt}=\alpha\sqrt{\omega_1}\,\sin(\omega_1 t-kC),$$
$$\frac{dp_3}{dt}=\alpha\sqrt{\omega_2}\,\cos(\omega_2 t-kC),$$
$$\frac{dp_4}{dt}=\alpha\sqrt{\omega_2}\,\sin(\omega_2 t-kC),$$
$$\vdots$$

or any other combination of $\sin(\cdot)$ and $\cos(\cdot)$ terms, as long as the perturbation of each parameter is orthogonal to that of every other one, in the Fourier space sense, such as $\cos(\omega_1 t)$ and $\cos(\omega_2 t)$, where $\omega_1\neq\omega_2$, or $\cos(\omega_1 t)$ and $\sin(\omega_1 t)$. Even sawtooth, triangle, or square waves of varying phase/frequency would work. To ensure independence of parameter oscillations, one should choose $\omega_i=\omega r_i$ where $r_i\neq r_j$ for all $i\neq j$ and ω is a scaling term which may be increased in order to improve stability. See, for example, Ref. [9], in which 24 parameters were tuned simultaneously with perturbations of the form $\cos(\omega_i t+kC)$, $\{\omega_1,\ldots,\omega_{24}\}$.

Typically, for a set of parameters, to whose settings the cost is similarly sensitive and whose operating ranges are similar, the various k_i and α_i are replaced by a common value of k and α, as

was done in the experiment presented here. Also, even for a wide range of parameter sensitivities and ranges, a common value of both k and α can be used following a normalization of all parameters within a prescribed range, as described in more detail in Ref. [9]. Therefore, for notational simplicity, in the remained of the paper, we write k and α in place of k_i and α_i, keeping in mind that unique values may be chosen for each parameter. The actual time between implementation of parameter update $p_i(n+1)$ following the parameter being set to $p_i(n)$ depends on the user, and is typically chosen to be long enough so that the components have had a chance to settle to their new settings (in the results presented here, we waited 5 s between parameter updates). Finally, setting bounds on parameters is straightforward and achieved by choosing maximum and minimum allowed parameter settings $p_{i,\max}$ and $p_{i,\min}$, respectively, and enforcing

$$\text{If } p_i(n+1) > p_{i,\max} \quad \text{then } p_i(n+1) = p_{i,\max},$$
$$\text{If } p_i(n+1) < p_{i,\min} \quad \text{then } p_i(n+1) = p_{i,\min}.$$

Analysis of the convergence of the adaptive scheme is based on the following theorem [15–17].

Theorem 1. *Consider the system*

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x},t) + \sum_{i=1}^{n} \mathbf{e_i} g_i(\mathbf{x},t)\sqrt{\omega_i}\,\cos(\omega_i t) - \sum_{i=1}^{n} \mathbf{e_i} h_i(\mathbf{x},t)\sqrt{\omega_i}\,\sin(\omega_i t), \tag{5}$$

where $\omega_i = \omega_0 r_i$ *and* $r_i \neq r_j$, $\forall i \neq j$. *The terms* $\mathbf{e_i}$ *are the ith basis vectors of* $\mathbb{R}^n$, *and the functions* $\mathbf{f}(\mathbf{x},t) : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}^n$, $g_i(\mathbf{x},t) : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$, *and* $h_i(\mathbf{x},t) : \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ *are continuous and Lipschitz, and continuously differentiable. For* $T \in [0,\infty)$, *and any compact set* $K \subset \mathbb{R}^n$ *for any* $t_0, \delta > 0$, *there exists* $\omega^\star$ *such that for all* $\omega_0 > \omega^\star$, *the distance between the trajectory* $\overline{\mathbf{x}}(t)$ *of the system*

$$\dot{\overline{\mathbf{x}}} = \mathbf{f}(\overline{\mathbf{x}},t) - \frac{1}{2}\sum_{i=1}^{n} \mathbf{e_i}\left(\frac{\partial h_i}{\partial \overline{\mathbf{x}}} g_i - \frac{\partial g_i}{\partial \overline{\mathbf{x}}} h_i\right), \tag{6}$$

and $\mathbf{x}(t)$ *of system (5) has bound*

$$\max_{t \in [t_0, t_0+T]} \|\mathbf{x}(t) - \overline{\mathbf{x}}(t)\| < \delta, \quad \overline{\mathbf{x}}(t_0) = \mathbf{x}(t_0). \tag{7}$$

With a bootstrapping argument, for systems whose trajectories remain within compact sets, the closeness of trajectories property, Eq. (7), is valid for all time for a single choice of $\omega^\star$.

The application of Theorem 1 to a system of the form (3) is immediate by expanding

$$\begin{aligned}\frac{dp_i}{dt} &= \alpha\sqrt{\omega_i}\,\cos(\omega_i t - kC)\\ &= \alpha\sqrt{\omega_i}\,\cos(\omega_i t)\,\cos(kC)\\ &\quad + \alpha\sqrt{\omega_i}\,\sin(\omega_i t)\,\sin(kC).\end{aligned} \tag{8}$$

System (8) is now in the same form as Eq. (5) with

$$g_i(p,t) = \alpha\,\cos(kC(p,t)),$$

$$h_i(p,t) = -\alpha\,\sin(kC(p,t)).$$

Application of Theorem 1 results in the averaged system

$$\frac{d\overline{p}_i}{dt} = -\frac{k\alpha^2}{2}\frac{\partial C}{\partial \overline{p}_i}(\sin^2(kC) + \cos^2(kC)) \tag{9}$$

$$= -\frac{k\alpha^2}{2}\frac{\partial C}{\partial \overline{p}_i}, \tag{10}$$

and the dynamics of all parameters $\mathbf{p} = (p_1, \ldots, p_m)$ is

$$\frac{d\mathbf{p}}{dt} = -\frac{k\alpha^2}{2}[\nabla C]^T. \tag{11}$$

Remark 3. The above analysis has immediate implications regarding the system's extreme robustness to noise. Consider a noise-corrupted version of the above system of the form

$$\frac{dp_i}{dt} = n_{i,p}(t) + \alpha\sqrt{\omega_i}\,\cos(\omega_i t - k[C + n_{i,c}(t)]), \tag{12}$$

where $n_{i,p}(t)$ is the additive noise contributing to parameter tuning and $n_{i,c}(t)$ is the cost function measurement noise. By the same exact arguments as above, as long as the noise is random and not correlated to the parameter settings p_i, so that $\partial n_{i,c}/\partial \overline{p}_i = 0$, this noisy system will follow average dynamics

$$\frac{d\overline{p}_i}{dt} = n_{i,p}(t) - \frac{k\alpha^2}{2}\frac{\partial C}{\partial \overline{p}_i}(\sin^2(\cdot) + \cos^2(\cdot)) \tag{13}$$

$$= n_{i,p}(t) - \frac{k\alpha^2}{2}\frac{\partial C}{\partial \overline{p}_i}. \tag{14}$$

The remarkable fact that on average the cost measurement noise has no influence on the convergence of Eq. (12) was confirmed and demonstrated in hardware, in the experimental results presented below. Furthermore, if $n_{i,p}(t)$ has non-zero average, convergence towards a local minimum of C is accomplished by choosing sufficiently large values of k, α.

2. Model-independent phase tuning

In our experimental setup we applied the adaptive (1) by prescribing changes, ΔS_1 and ΔS_2 ($\Delta S_i = S_i(n+1) - S_i(n)$: number of steps), for the motors in the phase packages of both the pre-buncher phase, ϕ_1, and the main buncher phase, ϕ_2, as shown in Figs. 2 and 3. The changes, $\Delta S_1(t)$ and $\Delta S_2(t)$, were made based on the value of the current monitor reading, 02CM01 (as shown in Fig. 1), squared: $C = C(\phi_1, \phi_2, t) = I^2(\phi_1, \phi_2, t)$, according to

$$\Delta S_1(n+1) = \delta_t \alpha\sqrt{\omega}\,\cos(\omega t - kC), \tag{15}$$

$$\Delta S_2(n+1) = \delta_t \alpha\sqrt{\omega}\,\sin(\omega t - kC), \tag{16}$$

where the actual phase setting changes, $\Delta\phi_1$ and $\Delta\phi_2$, occurred based on

$$\Delta\phi_1(n+1) = f_1(\Delta S_1(n+1)), \tag{17}$$

$$\Delta\phi_2(n+1) = f_2(\Delta S_2(n+1)), \tag{18}$$

where the functions f_1 and f_2 depend on the particular phase package and motor characteristics, which were close to linear near the mid points, near 180°.

We write $I(\phi_1, \phi_2, t)$ to emphasize that we were only changing the phases of the RF bunchers, but the actual current reading depends on all machine parameters, and varies with time even if none of the machine parameters are intentionally changed, due to component warm up and cool down and other un-modeled disturbances.

Remark 4. In the RF accelerating system, the DTL can be thought of as a filter, through which unmatched beam is improperly accelerated, no longer matched to the quadrupole lattice, and therefore lost as spill along the accelerator before reaching 02CM01. Therefore, the surviving current reading, I, as detected by 02CM01 in Fig. 1 was chosen as the location at which to collect beam measurement in order to ensure that surviving current is correlated with properly phased/matched beam.

The update time interval, between parameter set changes, was chosen as 5 s, so that the phase changes, implemented through mechanical systems, had time to respond and the network had time to re-calculate the new current reading. The value of ω was chosen to be 2000, and δ_t was chosen as $\delta_t = 2\pi/40\omega$ which provides a smooth variation of $S_1(t)$, $S_2(t)$, and therefore of ϕ_1 and ϕ_2, which in the absence of the cost term, would mean that they are oscillating at a smooth rate of 40 phase steps per oscillation,

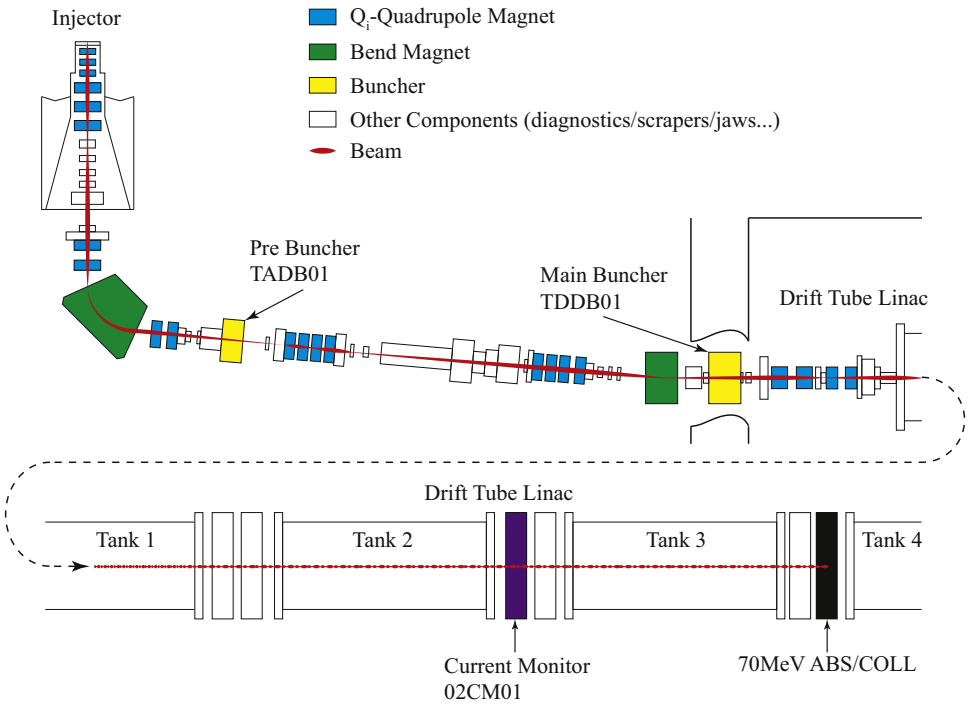

Fig. 1. Detailed overview of the entire system. The experiment consisted of maximizing the surviving beam current, entering at the injector, as measured by the current monitor (02CM01), following the first two accelerating structures of the accelerator, tanks one and two of the drift tube linac, by tuning the phase set points of the pre-buncher and main buncher cavities.

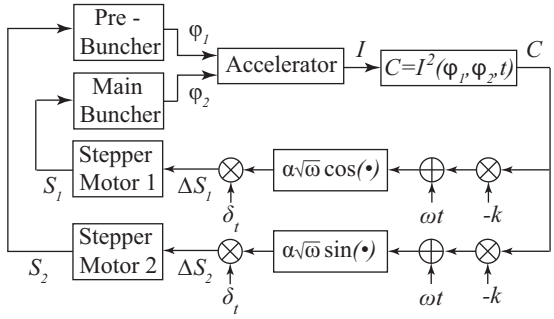

Fig. 2. The phases of the pre-buncher and the main buncher were tuned by sending step numbers and directions to the motors of the buncher phase packages.

which, as described above, is the finite difference approximation of the dynamics

$$\frac{dS_1}{dt}=\alpha\sqrt{\omega}\,\cos(\theta(t)), \tag{19}$$

$$\frac{dS_2}{dt}=\alpha\sqrt{\omega}\,\sin(\theta(t)), \tag{20}$$

where $\theta(t)=\omega t-kC$. Considering the trajectory $(S_1(t),S_2(t))$, of Eqs. (19) and (20), in a 2-dimensional space, the heading angle of the trajectory, $\theta(t)$, is a function of time, with rate of change

$$\frac{d\theta}{dt}=\omega-k\frac{dC}{dt}, \tag{21}$$

while the speed of the trajectory is constant:

$$\left|\left(\frac{dS_1}{dt},\frac{dS_2}{dt}\right)\right|=\sqrt{\left(\frac{dS_1}{dt}\right)^2+\left(\frac{dS_2}{dt}\right)^2}=\alpha\sqrt{\omega}. \tag{22}$$

Therefore, when the trajectory is heading in the correct direction such that the cost is increasing, that is $dC/dt>0$ and therefore $-kdC/dt<0$, the rate of rotation $d\theta/dt$ is decreased. On the other hand, when the trajectory is heading in the wrong direction such that the cost is decreasing, that is $dC/dt<0$ and therefore $-kdC/dt>0$, the rate of rotation $d\theta/dt$ is increased. Therefore, the trajectory spends more time moving towards the maximum, and on average performs a gradient ascent towards the maximum (by the same reasoning, for minimum instead of maximum seeking we replace $-k$ with k).

The mathematical analysis in Refs. [8,9] and the discussion above, including Theorem 1, show that on average the parameter dynamics are given by

$$\frac{dS_1}{dt}=\frac{k\alpha^2}{2}\frac{\partial C}{\partial S_1},\quad \frac{dS_2}{dt}=\frac{k\alpha^2}{2}\frac{\partial C}{\partial S_2}, \tag{23}$$

and therefore the system performs a gradient ascent towards a maximum value of C.

As an illustrative example, consider Fig. 4, which shows the simulation of the 2-dimensional system:

$$\dot{x}=\alpha\sqrt{\omega}\,\cos(\omega t+kC(x,y)) \tag{24}$$

$$\dot{y}=\alpha\sqrt{\omega}\,\sin(\omega t+kC(x,y)) \tag{25}$$

$$\dot{\bar{x}}=-\frac{k\alpha^2}{2}\frac{\partial C(\bar{x},\bar{y})}{\partial\bar{x}}=-k\alpha^2\bar{x} \tag{26}$$

$$\dot{\bar{y}}=-\frac{k\alpha^2}{2}\frac{\partial C(\bar{x},\bar{y})}{\partial\bar{y}}=-k\alpha^2\bar{y}, \tag{27}$$

where $C(x,y)=x^2+y^2$, $\alpha=\sqrt{0.5}$, $\omega=50$, and $k=5$.

Remark 5. In order for the averaging result, Eq. (23), to hold, the value of the $\omega_i=\omega r_i$ must be "big enough," relative to kdC/dt. As C is analytically unknown, the choice of ω is usually made in an iterative fashion, in which k is first chosen and fixed, and ω is then increased until the scheme is stable. Admittedly this is not ideal, but once a large enough ω is chosen, it typically can be left fixed for the particular system at hand. From experience, even limited

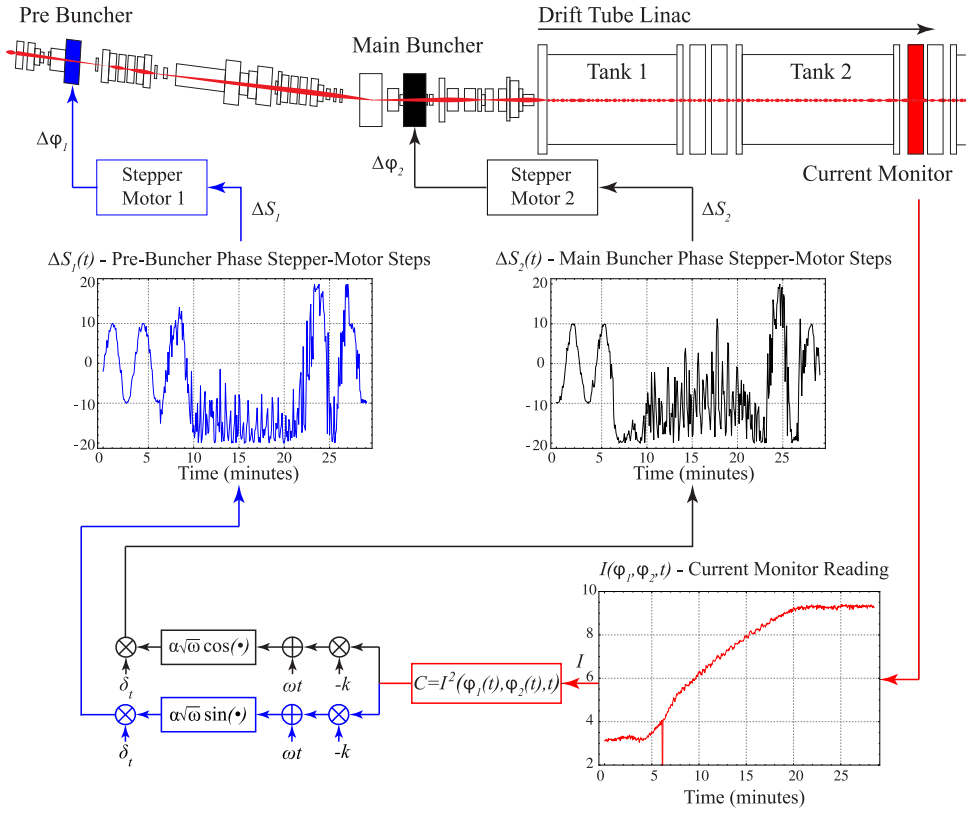

Fig. 3. The phases of the pre-buncher TADB01 and the main buncher TDDB01 were tuned by sending step numbers and directions to the motors of the buncher phase packages.

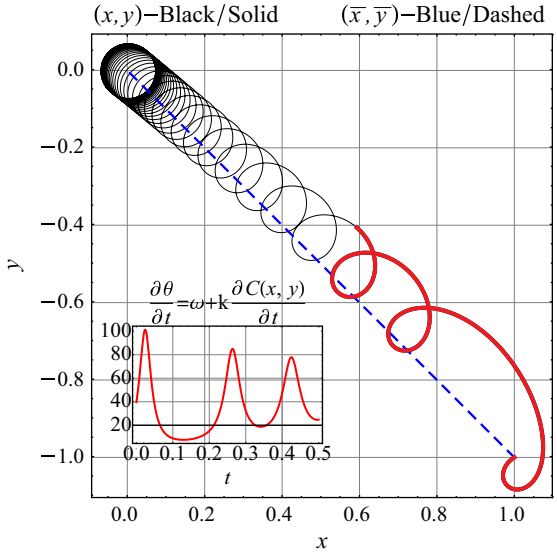

Fig. 4. The rotation rate can be seen to decrease and increase corresponding to moving towards or away from the minimum at $(0,0)$.

simulations can be very helpful in choosing the range of ω for a given system.

In summary, the adaptive control scheme presented here is digitally implemented through a control system such as EPICS in the following steps:

- Choose a cost, C, whose maximum or minimum corresponds to desired machine performance/beam properties.
- Choose parameters $p_1, \ldots, p_m$ which will be tuned in order to influence the value of the cost C and upper and lower bounds on all parameter settings $p_{i,\max}$ and $p_{i,\min}$.
- Choose perturbation sizes α_i for each parameter based on its allowable range.
- Choose control gains k_i, whose magnitude will influence the speed of convergence. Typically a single value k is used for a family of parameters for whom the sensitivity of the cost, C, is comparable.
- Choose perturbing functions, such as $\cos(\cdot)$, $\sin(\cdot)$, square, sawtooth, or triangle waves, of varying frequencies ω_i.
- Choose ω large enough so that the system is smoothly varying with perturbing frequencies $\omega_i = \omega r_i$ such that $r_i \neq r_j$ for all $i \neq j$.

3. Experiment I: one parameter slightly detuned, extremely noisy cost, slow adaptation

The experiment setup was as shown in Figs. 1 and 3, with a 150 μs H+ production, 12.12 mA beam injected from the source and then bunched by the pre-buncher and the main buncher. If the phase of either buncher is incorrect, the beam's bunching does not match up with the phase of the accelerating RF electric field in the tanks of the DTL. Besides possibly slowing rather than accelerating the beam, such a phase mismatch creates a mismatch between the beam energy and the designed beam energy relative to the quadrupole lattice throughout the accelerator, leading to beam breakup and loss. After being accelerated through the second tank of the DTL, after which the surviving current was measured by current monitor 02CM01, beam then passed through tank 3 of the DTL, which was delayed, without acceleration and terminated at

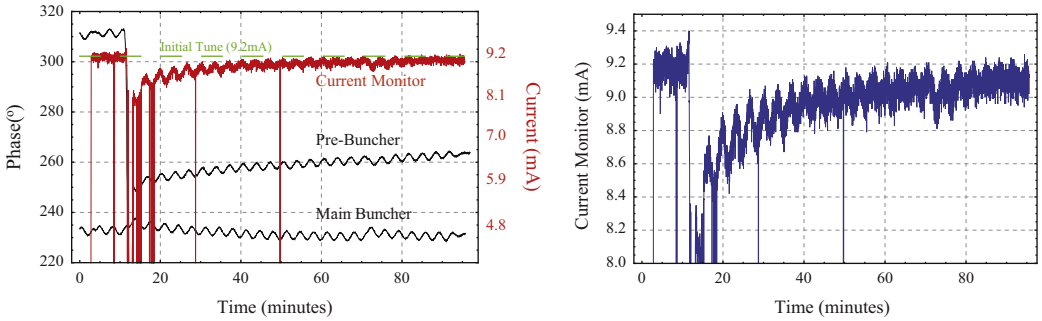

Fig. 5. *Experiment* I results: The top figure shows the pre-buncher and main buncher tuned settings and automatic re-tuning after the pre-buncher was manually de-tuned. The corresponding initial drop in surviving current to ~ 8.5 mA and gradual current growth is also shown by the current monitor readings. While the main buncher is seen to experience a temporary shift due to the large initial system perturbation, before returning to the neighborhood of its original set point, the pre-buncher phase has settled at a new location. The bottom figure shows the current monitor reading in more detail, to emphasize that the "cost" which is this reading squared, was very noisy, with the few instantaneous drops to zero caused by the machine auto-protection system, whose triggering was unrelated to this experiment.

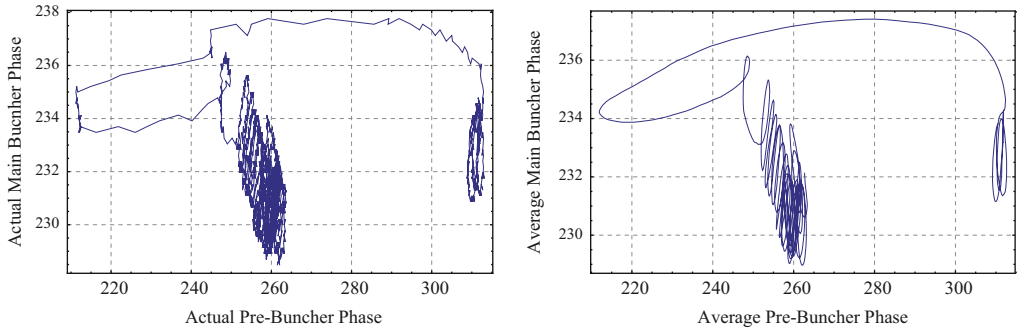

Fig. 6. *Experiment* I results: The top figure shows the actual phase settings of the pre-buncher and the main buncher. Because of the stepper motor, we can see sharp, non-smooth changes in each of the settings. The bottom figure shows a windowed average of each component, and we see that on average, they perform a smooth flow towards cost-maximizing settings, as predicted by the analysis.

the 70 MeV absorber collector. The initial phase settings of the pre-buncher and the main buncher were 307.83° and 234.78°, respectively. At these initial settings, approximately 9.2 mA ($\sim 76\%$) of current was being detected by 02CM01.

The first part of the experiment included machine setup and initial parameter sensitivity calibration. Only 1 component was purposely, slightly detuned from its set point, and the controller was then turned on with a small gain and a perturbation size, $k=0.1, \alpha=1423$, which, for the chosen frequency, $\omega=2000$ and $\delta_t=2\pi/40\omega$ corresponded to a maximum assigned movement of 5 steps for each motor per iteration:

$$|\Delta S_1|, |\Delta S_2| = |\delta_t \alpha \sqrt{\omega}\ \cos(\cdot)|, |\delta_t \alpha \sqrt{\omega}\ \sin(\cdot)|$$
$$\leq |\delta_t \alpha \sqrt{\omega}| = \left|\frac{2\pi\alpha}{40\sqrt{\omega}}\right| \approx 5. \quad (28)$$

The results of the experiment are shown in Figs. 5 and 6. In Fig. 5 we see how the current monitor reading and the phase settings evolve as we start with the tuned up settings, detune the pre-buncher phase set-point only, and activate the feedback loop in order to re-maximize the current monitor reading. It is evident that the current reading is very noisy, with multiple sudden drops to zero, which were caused by a machine auto-protect which was unrelated to the experiment (this issue was fixed in subsequent experiments). This extremely noisy and discontinuous current reading was actually beneficial as it clearly demonstrated the schemes robustness to random noise. This experiment confirmed that on average, as seen in the bottom portion of Fig. 5, as long as the perturbing frequency components are present, the scheme is able to automatically adjust the phase settings, and random, sudden changes are averaged out to have no effect.

Fig. 6 shows the evolution of the pre-buncher and main buncher phase settings, their actual and window averaged versions. As the main buncher was initially unchanged, we see that it undergoes only a slight perturbation, remaining in the neighborhood of 230°, while the pre-buncher, which was manually de-tuned, is tuned to a new set point, one which maximizes the current monitor reading. The elliptical oscillations in the beginning and the end of the adjustment process are what we expect as described in the mechanism responsible for the convergence. During the first few minutes of optimization we see a very large drift in the phase settings, as the gradient of the cost is evidently very large relative to the pre-buncher's initial de-tuned set point. As the system approaches the maximum the parameter oscillation settles to a steady state about the maximizing points.

4. Experiment II: both parameters slightly detuned, noisy cost, slow adaptation

In the second experiment, both components were manually de-tuned from their initial tune up settings, so that the surviving current value dropped to 8.3 mA. Again, ω was set at $\omega=2000$, $k=0.1$, and initially, $\alpha=1400$. Around 40 min into the experiment the value of α was adjusted to $\alpha=700$ to allow the phase settings to zoom in closer to their optimal values, and finally to 0 so that the system could settle.

Figs. 7 and 8 show the results of the second experiment. In Fig. 7 we see the initial surviving current drop from 9.2 mA to 8.3 mA as both the pre-buncher and the main buncher phase were manually detuned. We then see the adaptive feedback quickly re-phase both bunchers in order to re-maximize the surviving current. As the perturbing amplitude, α, is decreased, the phase settings settle at their optimal values. Alternatively, the phase

readings, in steady state oscillation, could have been averaged, to find the optimal settings.

In Fig. 8 we see the evolution of the phase settings. Despite the very jagged stepper motor set point changes in the phase settings, the system behaves as we expect, with the windowed average of the phase settings clearly showing a drift towards and settling around the maximizing phase set points.

5. Experiment III: both parameters extremely detuned, noisy cost, aggressive/fast adaptation

The results of the final experiment are shown in Figs. 9–11. In Fig. 9 we see the surviving beam value drop from 9.2 mA to 3 mA, as both the pre-buncher and main buncher phase settings were drastically changed. Again, ω was set at $\omega = 2000$ and the controller was then activated for approximately 25 min with very gentle adaptation parameters $k = 0.1$, $\alpha = 1400$. At approximately the 30 min mark, as shown in the bottom part of Fig. 9 the parameters were changed to $k = 0.4$ and $\alpha = 5694$, and we see that the scheme quickly found new optimal phase settings within ~ 25 min, at which point the surviving current of 9.3 mA was higher than what the initial machine tune-up had achieved. The perturbation amplitude was then gradually reduced to $\alpha = 2847$, $\alpha = 1428$, $\alpha = 714$ and finally $\alpha = 0$ as we let the system settle to equilibrium.

In Fig. 10 we focus on the actual step numbers sent by the adaptive controller to the phase-setting stepper motors. We see that during the rapid convergence period, during the 25 min interval focused in the bottom portion of Fig. 9, the sinusoidal oscillations have disappeared as the terms $k\partial C/\partial p_1$ and $k\partial C/\partial p_2$ in the parameter vector's heading angle's rate of change

$$\frac{d\theta}{dt} = \omega - k\frac{dC}{dt} = \omega - k\frac{\partial C}{\partial \phi_1}\frac{\partial \phi_1}{\partial t} - k\frac{\partial C}{\partial \phi_2}\frac{\partial \phi_2}{\partial t}, \tag{29}$$

are large enough to keep the system flowing towards the maximum.

In Fig. 11 the initial oscillation and the sudden flow toward the maximizing settings are evident once the controller's gains and perturbation amplitude have been increased. The system is then again settled, oscillating about the equilibrium points.

6. Conclusions and future work

We have demonstrated the effectiveness of a model-independent technique to simultaneously, adaptively tune multiple parameters in a very noisy and non-linear environment. The adaptive controller automatically, simultaneously tuned the phase of both the pre-buncher and the main buncher in the H+ transport section, so that the amount of current surviving to the

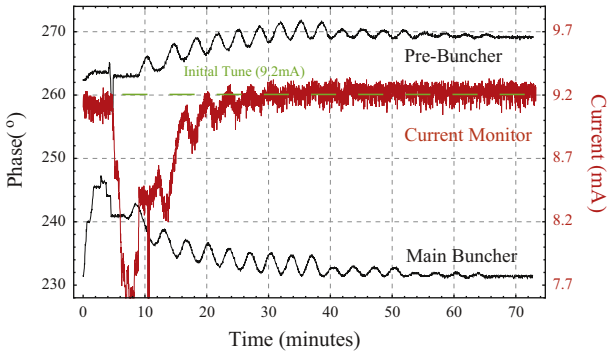

Fig. 7. *Experiment* II results: Both the pre-buncher and main buncher phase settings were detuned until only ~ 8.3 mA of current was surviving at the current monitor. The adaptive scheme is then seen to re-tune both the settings to new values at which the surviving current is again maximized, despite a very noisy current monitor reading. As the value of α was decreased from 1400, to 700, to 0, the phase settings settled to steady states.

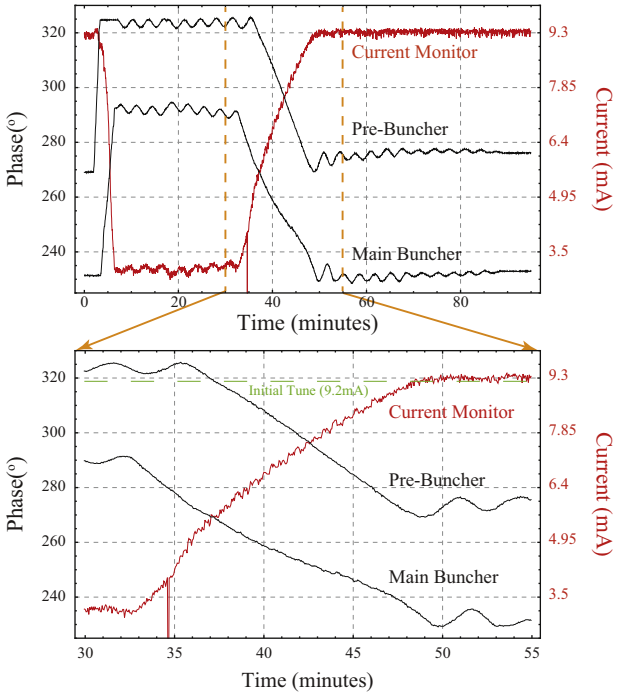

Fig. 9. *Experiment* III results: As the phase of both the pre-buncher and main buncher are drastically de-tuned the surviving current is diminished. Once the adaptive controller is turned on, it starts to slowly re-tune the phase settings. Once the controller's gains are sufficiently increased, the adaptation is very fast. Unlike the first two cases, we see that the buncher phases return to almost their exact initial settings.

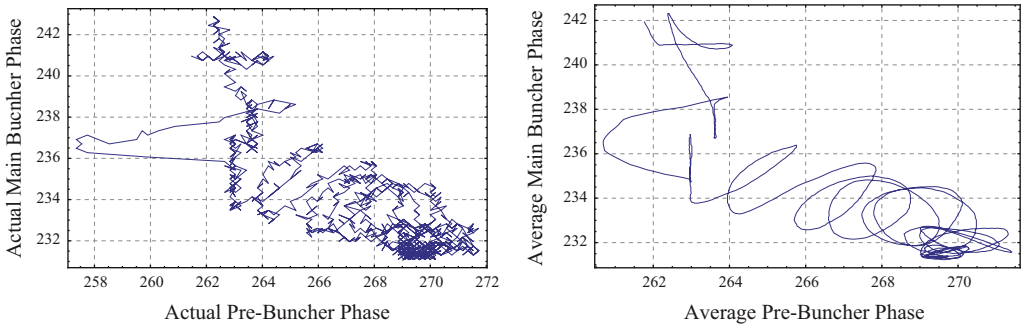

Fig. 8. *Experiment* II results: The top figure shows the actual phase settings of the pre-buncher and the main buncher. Because of the stepper motor, we can see sharp, non-smooth changes in each of the settings. The bottom figure shows a windowed average of each component, and we see that on average, they perform a smooth flow towards cost-maximizing settings, as predicted by the analysis.

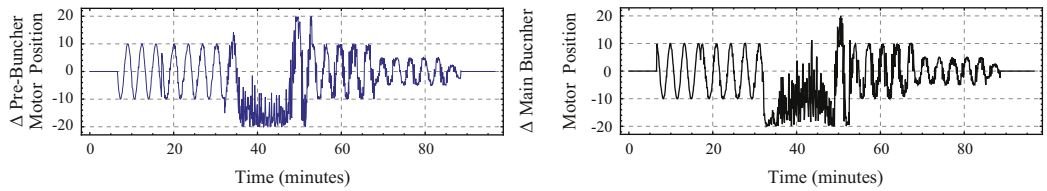

Fig. 10. *Experiment* III results: This figure shows the values of the step numbers sent by the controller to the phase-setting stepper motors. We see that during the rapid convergence, as the gradient of the cost function is high and amplified by the gains, the parameters are not perturbed in a sinusoidal fashion, but are instead quickly heading directly toward the maximizing values.

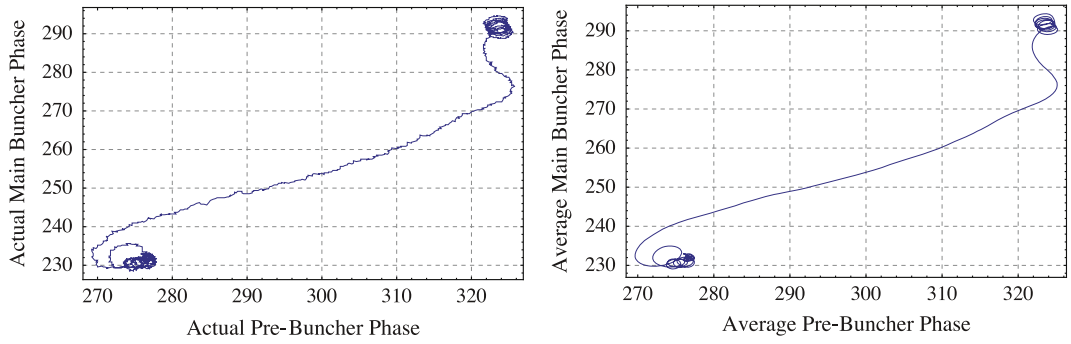

Fig. 11. *Experiment* III results: The tope figure shows the actual phase settings of the pre-buncher and the main buncher. Because of the stepper motor, we can see sharp, non-smooth changes in each of the settings. The bottom figure shows a windowed average of each component, and we see that on average, they perform a smooth flow towards cost-maximizing settings, as predicted by the analysis.

end of DTL Tank 3, as detected by current monitor 02CM01, was automatically maximized from 3 mA to 9.3 mA in $\sim$25 min, as shown in Fig. 9. The achieved 9.3 mA of current, $\sim$77.5% of the injected 12.12 mA of current, was slightly (0.1 mA) higher than the initial machine tune, which was achieving 9.2 mA, and was very close to the $\sim$80% as achieved in the best case using typical tune up methods. We believe that our result was achieved despite limitations such as the fact that some of the magnets had been set for optimal H$-$ transport, and that we were limited to only tuning the pre-buncher and the main buncher phases and not their amplitudes, and we did not have a chance to optimize the phases of DTL tanks 1 and 2.

From the above experiments, and analysis, it is clear that the choices of k, α, and the cost C have a strong influence on the rate of convergence of the adaptive scheme. Furthermore, because the average convergence speed of the ith parameter is proportional to $k\alpha\partial C/\partial p_i$, as we approach a maximum or a minimum of C, if the function levels out and its gradient decreases, so will the convergence rate. Therefore, it would be beneficial to design adaptive control amplitude and gain, $\alpha(t)$ and $k(t)$, to try and maintain a constant rate of convergence. We point out that one has to be careful about, for example, in the case of minimum seeking adaptively increasing the gain, $k(t)$, as the cost decreases. While this has the potential to keep the convergence rate higher or more consistent throughout, it can also destabilize the algorithm if not done carefully, in particular, one should maintain $dk/dt \ll \omega_i$, because the "cost" seen by the algorithm is the product $k(t)C(x,t)$.

Another possibility is of changing both $k(t)$ and $\alpha(t)$ at equal and opposite rates, such that $k(t)$ is increasing while $\alpha(t)$ is decreasing while the product $k(t)\alpha(t) \equiv k(0)\alpha(0)$ is a constant, and the increase in $k(t)$ is on a slower time scale than the perturbations $\cos(\omega_i t)$. In this case, the overall average dynamics would be unchanged, but the convergence would have better characteristics, both converging faster and decreasing in parameter perturbation size as it approaches equilibrium. We have not yet had a chance to try such a scheme outside of simulation, but we do plan on experimenting with this approach in the future.

Our next goal is to expand this to tune more parameters simultaneously, namely the entire RF system of the accelerator, which we believe, has the potential to both save time and may achieve superior results in terms of amount of surviving beam current compared to what has been done by manual hand-tuning, and would include both the H+ and H$-$ beams of the accelerator. Furthermore we point out that due to limited development time (only one six hour shift), we did not have a chance to push the controller to its limit, to test if faster convergence can be achieved with even higher gains. From the experimental results, and the smoothness of the averaged parameter drifts, we believe that we can substantially increase the convergence speed, by a factor of two at the minimum.

Finally, we point out that although this scheme may be useful for initial machine tune-up, as demonstrated by the experiments discussed in this work, it may have even greater potential as an online system for fine-tuning preset or pre-tuned machines, continuously compensating for time-variation in machine changes, such as thermal drift, and also, considering, for example, returning a machine to a tune up, based on previously saved machine settings. Due to uncertainties such as hysteresis of the magnets, although a machine's settings, such as the magnet current sources may be returned to exactly their previous values, the actual magnetic field corresponding to these currents may have changed. The scheme presented in this case would automatically adjust magnet settings, until the actual, desired magnetic fields were achieved, as measured by emittance scans, loss monitors, beam position monitors, or any other available diagnostic of beam characteristics. Furthermore, in the case of unexpected change in operating conditions, such as the failure of one or more components, this adaptive scheme would attempt to automatically retune the remaining, functional system components, in order to compensate for those which are failing.

References

[1] K. Astrom, T. Hagglund, PID Controllers: Theory, Design, and Tuning, Instrument Society of America, Research Triangle Park, NC, 1995.
[2] C. Collette, S. Janssens, D. Tshilumba, Nuclear Instruments and Methods in Physics Research Section A 684 (2012) 7.
[3] W.Z. Wu, Y. Kim, J.Y. Li, D. Teytelman, M. Busch, P. Wang, G. Swift, I.S. Park, I.S. Ko, Y.K. Wu, Nuclear Instruments and Methods in Physics Research Section A 632 (2011) 32.

[4] J. Vincent, D. Morris, N. Usher, Z. Gao, S. Zhao, A. Nicoletti, Q. Zheng, Nuclear Instruments and Methods in Physics Research Section A 643 (2011) 11.
[5] E. Meier, S.G. Biedron, G. LeBlanc, M.J. Morgan, J. Wu, Nuclear Instruments and Methods in Physics Research Section A 610 (2009) 629.
[6] E. Meier, S.G. Biedron, G. LeBlanc, M.J. Morgan, J. Wu, Nuclear Instruments and Methods in Physics Research Section A 609 (2011) 79.
[7] T. Poggi, Nuclear Instruments and Methods in Physics Research Section A 729 (2013) 506.
[8] A. Scheinker, Model independent beam tuning, in: Proceedings of the 2013 International Particle Accelerator Conference, Shanghai, China, 2013.
[9] A. Scheinker, X. Pang, L. Rybarcyk, Physical Review Special Topics—Accelerators and Beams 16 (2013) 102803.
[10] M. Leblanc, Sur l'electri"cation des chemins de fer au moyen de courants alternatifs de frequence elevee, Revue Generale de l'Electricite, 1922.
[11] W.H. Moase, C. Manzie, D. Nesic, I.M.Y. Mareels, Extremum seeking from 1922 to 2010, in: 29th Chinese Control Conference, 2010, p. 14.
[12] P.L. Kapitza, Soviet Physics JETP 21 (1951) 588.
[13] S. Meerkov, IEEE Transactions on Automatic Control 14 (1980) 755.
[14] A. Scheinker, Extremum seeking for stabilization (UCSD Ph.D. thesis), November, 2012.
[15] J. Kurzweil, J. Jarnik, Journal of Applied Mathematics and Physics 38 (1987) 241.
[16] H.J. Sussmann, W. Liu, Limits of highly oscillatory controls and approximation of general paths by admissible trajectories, in: Proceedings of the 30th IEEE CDC, Brighton, UK, 1991.
[17] H.J. Sussmann, New differential geometric methods in nonholonomic path finding, in: Progress Systems and Control Theory, 1992, pp. 365–384.