An Algorithm for Online Optimization of Accelerators

Xiaobiao Huang, Jeff Corbett, James Safranek, Juhao Wu

SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025

Abstract

A general algorithm is developed for online optimization of accelerator performance, *i.e.*, online tuning, using the performance measure as the objective function. This method, named *robust conjugate direction search (RCDS)*, combines the conjugate direction set approach of Powell's method with a robust line optimizer which considers the random noise in bracketing the minimum and uses parabolic fit of data points that uniformly sample the bracketed zone. It is much more robust against noise than traditional algorithms and is therefore suitable for online application. Simulation and experimental studies have been carried out to demonstrate the strength of the new algorithm.

Keywords: optimization, algorithm, SPEAR3

1. Introduction

Accelerators are complex systems that often consist of hundreds of active components such as magnets and rf cavities. Each component may have one or more parameters, *i.e.*, knobs that can be adjusted. Tuning an accelerator is the process of finding a set of parameter values to achieve optimal performance. The machine usually has various diagnostic instrumentation to monitor the performance of its sub-systems. For example, beam position monitors (BPMs) are used to measure the beam orbit and a spectrum analyzer or turn-by-turn BPMs are used to measure the betatron tunes for a circular machine. If the targets of the diagnostics, such as the ideal orbit or the ideal tunes, are known and there is a deterministic approach to turn the knobs to reach the target, then the task of tuning the sub-systems is

Email address: xiahuang@slac.stanford.edu (Xiaobiao Huang)

Preprint submitted to Nuclear Inst. and Methods in Physics Research, A April 29, 2013

straightforward and easy. Response matrices that relate knobs to readings of the diagnostics are often used to calculate the required correction steps. Beam orbit correction and tune correction are examples that fall in this category.

There are other cases when the target of a subsystem is not directly monitored but can be measured indirectly. For example, the optics of a synchrotron or a storage ring is not measured directly but can be inferred from the orbit response matrix (*i.e.*, the differential orbit shifts with respect to individual orbit correctors) or readings of turn-by-turn BPMs for a beam under betatron motion. Fitting the orbit response matrix to calibrate the optics model has been an effective way for optics correction [1, 2].

However, in many cases the diagnostics cannot provide sufficient information to guide the move of the knobs toward the target. One example is the steering of the beam in a transport line prior to injection into a storage ring. The BPMs in the transport line do not measure the effect of the last few steering magnets very well because the distances between these magnets and the downstream BPMs are small. Therefore the launching angle and position of the injected beam are not very well determined with the transport line BPMs. A similar example is the optics matching of the injected beam. There can also be cases where the necessary diagnostics do not exist at all. This is especially possible for new projects in the commissioning phase.

In other cases the ideal target of a subsystem itself may be unknown or not very well known, or there does not exist an effective way to relate the knobs to the target. An important example is the optimization of the nonlinear dynamics of storage rings. Harmonic sextupoles are often used in third generation light source storage rings as free knobs to improve the dynamic aperture and momentum aperture. Ultimately the goal of nonlinear dynamics optimization is to achieve high injection efficiency and high beam lifetime. But there is no direct way to determine the desired parameter adjustment for the harmonic sextupoles. Normally it is desired to set the harmonic sextupoles so that the lattice works as it does in the ideal model. The calibration errors of the sextupole magnets and the nonlinear components in the other magnets can cause discrepancies between the model and the real machine.

Although there have been a few proposed approaches for beam-based nonlinear dynamics calibration [4, 5], the fact that the nonlinear motion detected by BPMs is usually weak and is plagued by the nonlinear response of the BPMs themselves makes it very challenging to achieve the goal. Furthermore, because the ideal model does not contain all the nonlinear effects that

real machine may have (such as the higher order fringe field effects), the real optimal setting may be different from that predicted by the model.

The nonlinear dynamics optimization challenge for the proposed ultimate storage rings (USRs) [7, 8] would be especially severe since they have small dynamic apertures and many strong linear and nonlinear magnets. The calibration errors of the sextupole magnets and the nonlinear errors from the dipole and quadrupole magnets could make a significant difference between the model and the real machine. The ability of achieving the design dynamic aperture and momentum aperture may very well determine the feasibility of the USRs.

When there is no direct method to predict the desired changes to the knobs, we usually tune the machine manually by turning the knobs and observing their impact on the machine performance, as measured by output power, injection efficiency, beam lifetime, or luminosity, *etc*, depending on the application. A usual approach is to tune one knob at a time for a few iterations until a good parameter set is obtained. This is basically to solve a multi-variable, usually nonlinear, optimization problem with iterative 1-dimensional parameter scans. The manual tuning approach can be effective in some cases, but it is generally not ideal. First, an automated routine could be more efficient in setting the knobs and evaluating the performance (which is basically an online function evaluation). Second, it is not applicable when the number of knobs or the number of required iterations are high. Third, the 1-dimensional scan algorithm may not be very efficient.

Automatic optimization code has been developed for accelerator control before [6], employing the Nelder-Mead simplex method and iterative 1-dimensional scans. However, these traditional optimization algorithms may not work well for online applications because here the function evaluation is noisy. It is therefore desirable to develop an efficient, robust algorithm for the purpose of online optimization. In this study we propose an efficient algorithm that can effectively overcome the noise issue. We demonstrate its strength by comparing performance with several other algorithms in simulation and apply this new algorithm experimentally to realistic problems.

A discussion of considerations for online optimization is given in section 2. The new algorithm is introduced in section 3 and is compared to other algorithms in simulation in section 4. Experimental results at the SPEAR3 storage ring are showed in section 5.

2. Considerations of online optimization algorithms

Optimization is usually to maximize or minimize the value of a function. Since any maximization problem can be turned into an equivalent minimization problem by putting a minus sign to the definition of the objective function, in the following we consider only the minimization problem for convenience of discussion.

Two basic requirements for an online optimization algorithm are high efficiency and robustness. High efficiency means that it is able to find the optimum with as few function evaluations as possible. Since each function evaluation takes time and the total machine time available for online optimization is usually limited, high efficiency is necessary for relatively large scale problems.

Robustness in this context means the algorithm can find the optimum despite noise in measured function values (and occasional function value outliers) and that the algorithm behaves properly under machine failures. Most traditional algorithms assume the objective function is smooth and use the comparison of function values of different solutions to decide the next step. When noise is introduced, the comparison result may be misleading and the algorithm may fail. Therefore an online algorithm needs to be aware of the noise and take action more cautiously. For some machines or measurements, outliers in the function values are possible as they may be caused by temporary malfunction of a component or an occasional mismatch between sub-systems. A robust algorithm should filter out outliers or the course of the optimization process could be altered. Machine failures are situations when the machine is not working properly as required by the experiment and hence produce meaningless function values, for example, when the injector is down and fails to send beam to the transport line during an injection efficiency measurement. During machine failures a robust algorithm should not proceed. The algorithm should be able to dump the present status and resume the optimization process later. Another aspect of robustness of the algorithm is that it must know the valid boundary of the parameters and avoid suggesting unrealistic solutions.

Traditionally optimization algorithms can be divided into two groups, gradient-based algorithms and non-gradient-based algorithms, depending on if they calculate and take advantage of the gradient of the function in parameter space. In this study we consider only non-gradient-based methods. This is because the gradient calculations built upon numerical differences

(including those algorithms that do not explicitly calculate the gradient) is sensitive to noise and may not be useful.

One candidate algorithm is an iterative parameter scan that basically automates the manual approach as described in the previous section. At each iteration parameters are scanned within a specific range while the other parameters are fixed at the previous best values. After the scan this parameter is set to the value that corresponds to the new minimum. Although this method can sometimes be effective, it is potentially inefficient if there exists a long, narrow valley in the parameter space that is not oriented along or near any parameter direction, in which case each parameter scan will make only a small gain along the valley [9].

Another candidate is Powell’s conjugate direction method [9, 10], which is an iterative line search algorithm that updates the search direction set to make an efficient search. It also employs bracketing and an advanced line optimizer (golden section or quadratic interpolation).

A third candidate we considered is the well known Nelder-Mead simplex method [9, 11]. The simplex method manipulates a non-degenerate geometric body with $N+1$ vertices in the N -variable parameter space, called a simplex, through a series of operations according to function values on the vertices. The operations include reflection, expansion and contraction. This method is known to be easy to use and robust for smooth functions.

Finally, multi-objective genetic algorithms (MOGA) recently have been widely used in the accelerator community, mostly for optimizing accelerator designs [12]. Genetic algorithms are powerful for simulation because the time-consuming evaluation of new solutions can be done in parallel. However, they may not be promising candidates for online optimization since usually it takes many function evaluations for genetic algorithms to work. Nonetheless we include a genetic algorithm, NSGA-II [15], in this study for comparison.

After exploring the candidate algorithms for noisy problems in simulation, we came to realize that an effective and efficient method can be reached by combining the conjugate direction search of Powell’s method with a line optimizer that is robust against random noise and outliers. In the following section we describe this new method and compare its performance to the other candidate methods.

3. The Robust Conjugate Direction Search (RCDS) algorithm

The goal of our optimization algorithm is to minimize a multi-variable function $f(\mathbf{x})$ where $\mathbf{x}$ is an N -dimensional vector that represents a solution. Each element of $\mathbf{x}$ is a parameter which for online applications must have a valid range. Since the scales of the parameters can be very different and a general algorithm should not be affected by the parameter scales for a particular problem, it is ideal to normalize the parameter space. We choose to normalize each parameter to the range $[0, 1]$ so that the parameter space can be represented as an N -dimensional unit cube.

The iterative search along conjugate directions is generally more efficient than searching along individual parameter directions [10]. By definition two vectors $\mathbf{u}$ and $\mathbf{v}$ are conjugate if they satisfy the condition $\mathbf{u} \cdot \mathbf{H} \cdot \mathbf{v} = 0$ with the Hessian matrix $\mathbf{H}$ defined as the matrix of second order derivatives of the objective function $f(\mathbf{x})$, i.e., $H_{ij} = \partial^2 f / \partial x_i \partial x_j$. Successive search on a conjugate direction set is efficient since the previous result would not be spoiled by a new line search. Each search direction is orthogonal to the others in terms of their impact to the objective function.

In our algorithm we adopt the same routine as used in Powell's method to build up and update the conjugate direction set as detailed in Ref. [9]. Basically after each iteration, the direction that has the largest decrease of the objective function is replaced with the direction between the best solutions before and after the iteration, provided that the direction is still effective (which is to be tested with an extrapolation). Powell's method usually suggests the initial direction set to be unit vectors of the parameter space. Consequently it takes the algorithm some time to build up the conjugate direction set. Because efficiency is important for online optimization, we recommend providing an initial estimate for the conjugate direction set whenever possible, which may be obtained through a model. For example, one may calculate the Hessian matrix with the model and use the eigenvectors of the Hessian matrix as the initial conjugate direction set. It is also possible to use a fixed conjugate direction set without updating.

The main modification we made to Powell's method was in the line optimizer, which is now made aware of the rms noise level σ_f of the measurement. The line optimizer is a one-dimensional optimization algorithm aimed at minimizing $g(\alpha) = f(\mathbf{x}_0 + \alpha \mathbf{u})$, where $\mathbf{x}_0$ is the present solution and $\mathbf{u}$ represents the search direction. Similar to Powell's method we first bracket the minimum, *i.e.*, find a zone $[\alpha_l, \alpha_h]$ in which a local minimum of $g(\alpha)$ exists.

We determine the lower and upper bounds of the zone by starting from the present solution (i.e., $\alpha = 0$), going along one direction until the minimum, which is updated as the search goes, is less than the function value of the last solution by more than $3\sigma_f$ and setting the bound. The initial step size is given externally. The step size then grows by a factor of 1.618 (the golden section) for each additional step (following Ref. [9]). If a solution larger than the minimum by $3\sigma_f$ is not found before the search enters the invalid range (in which case the function value should return NaN, *i.e.*, not a number), the last valid solution is set as the bound. The bound in the other direction is similarly found.

The next task of the robust line optimizer is to determine the minimum within the bracket $[\alpha_l, \alpha_h]$. To efficiently recover information from noisy data, we adopt an approach to uniformly sample within the bounded zone and fit the data to a parabolic curve. Typically we take six data points, including the two end points of the zone. The solutions found during bracketing are saved and reused for fitting. If any of the previous solutions is close to a proposed sample point (say, within 10% of the zone size), the evaluation of the latter can be skipped. Often times only 2 or 3 more extra evaluations are required. After the parabolic fit, we look for outliers by examining the errors, *i.e.*, the differences between the actual function values (*i.e.*, measured data) and fitted values. An outlier is defined as a data point whose error is significantly larger than the average error of the centroid of the data sets. If an outlier is identified, it is removed and the remaining data points are re-fitted. The minimum of the line search is then calculated from the parabolic curve.

An illustration of the bracketing and fitting procedures is shown in Figure 1, using a line search from the vertical emittance minimization problem to be discussed in section 4.1 as an example. Starting from the initial point $\alpha = 0$ toward the positive direction, the upper bound is found with one step. It then searches in the opposite direction and finds the lower bound with three steps, each step with a larger step size. Two new points are added to uniformly sample the bracketed zone and are used for fitting. Finally the new minimum is found from the fitted parabolic curve.

Our line optimization algorithm is different from that of the original Powell's method in several aspects. First, the bracketing approach guarantees the bounds are robust against noise. We use golden section extrapolation to look for the bounds instead of a three-point parabolic extrapolation since the latter may be unreliable for noisy data. We also make sure the bounds are

Figure 1: Illustration of bracketing and quadratic fitting of the line optimizer.

above the noise level so that they truly bracket the minimum. Second, by uniformly sampling the bracket and making a parabolic fit, we not only improve the accuracy of the solution by gaining a statistical advantage from multiple data points, but also obtain a reliable minimum within the bracket by global sampling. This is unlike the iterative golden section search or three-point parabolic interpolation methods, for which a single noisy data point may lead the search to the wrong section of the bracket and hence a false minimum. Third, our algorithm also handles occasional outliers properly. These differences give our optimization method the power to steer through noisy data and find the optimum with high precision, as will be demonstrated in the next sections.

The termination condition for the iterative algorithm may be either exceeding a pre-specified total number of evaluations, exceeding a total time limit or when the solution stays within the vicinity of the previous iteration for a specified number of times. Or the algorithm can be terminated manually. It is worth noting that monitoring the progress during online optimization is important. Printing out and logging the trial solutions and algorithm decisions provide useful guidance for fine tuning the optimization routine parameters.

Our new algorithm may be referred to as the robust conjugate direction search (RCDS) method.

4. Simulation studies

To test the robustness against noise and the convergence rate of the RCDS method, we applied it to two realistic accelerator optimization problems in simulation and compared its performance to other methods, including the Nelder-Mead simplex method, the original Powell's method and NSGA-II. We also include an algorithm that uses our robust line optimizer for iterative parameter scans (referred to as *IMAT*), which is basically the same as RCDS except that the direction set is fixed as unit vectors in the parameter space.

The two problems are minimization of the vertical emittance of the SPEAR3 ring with skew quadrupole magnets and optimization of the SPEAR3 Booster-to-SPEAR (BTS) transport line optics.

4.1. Minimization of vertical emittance for SPEAR3

The vertical emittance of a storage ring comes from the vertical dispersion inside dipole magnets and linear betatron coupling between horizontal and vertical planes. Both contributing factors can be corrected with skew quadrupoles. In SPEAR3 there are 13 free skew quadrupoles for vertical emittance control. Usually the skew quadrupole setting for vertical emittance correction is obtained by fitting the orbit response matrix data [1], in particular, the off-diagonal elements. Since beam loss is dominated by Touschek scattering for most third generation light sources, the beam loss rate provides an indirect way to minimizing vertical emittance. We have tested this approach in experiments (see section 5) and also simulated the optimization process.

In the simulation, errors are seeded on 42 skew quadrupoles (including the 13 that are used for correction) in the lattice. The vertical emittance is computed with the code Accelerator Toolbox [13] which implements Ohmi's procedure [14]. The loss rate is then calculated, assuming the gas scattering lifetime is 40 hours and the Touschek lifetime for a 500 mA beam is 10 hours for a coupling ratio (i.e., the vertical to horizontal emittance ratio) of 0.2%. The coupling ratio for the error seed when the 13 correcting skew quadrupoles are turned off is 0.9%.

Since in experiment the loss rate is measured by monitoring the beam current drop in a fixed time interval, the noise in the loss rate data is mostly from the uncertainty in the beam current measurement. This is simulated by assuming a gaussian random noise with $\sigma_I = 0.03$ mA. We also add a 1% random noise to the beam current in the lifetime calculation to simulate

the beam current variation between successive top-off injections. We keep a record of the coupling ratio for each evaluated solution, which does not have the random noise in the loss rate data and is thus a true measure of the quality of the solution. Apparently longer intervals in loss rate measurement leads to less noise. We chose an interval of 6 seconds, which corresponds to a 0.06 mA/min rms uncertainty in the loss rate. The loss rate for a 500 mA beam with all 13 correcting skew quads turned off is 0.6 mA/min.

The skew quadrupole parameter (normalized focusing gradient) range in the simulation is set to $-0.3 \sim 0.3 \text{ m}^{-2}$, while the magnet length is 20 cm. For the simplex method, the initial simplex has a vertex that corresponds to all skew quadrupoles set to zero. The other vertices are each off in one parameter, by 10% of the range of the parameter. The starting point for Powell's method is for all skew quadrupoles set to zero. The NSGA-II method is applied here to optimize only one objective.

Figure 2 shows the performances of the simplex method and Powell's method. We plot the history of the objective function value of the best solution during each optimization run. Because of the noise in function values, the course of each run is different. Three runs with random noise and a run without random noise are shown for each method. It is clear the performance of the simplex method degrades significantly when noise is introduced to the objective functions. This is understandable because the algorithm stops to work as planned as soon as the size of the simplex shrinks to the point when the comparison results of the function values on the vertices are altered by the random noise.

Powell's method turns out to be more sensitive to noise than the simplex method. With the noise level given by the 6-second beam loss time interval, the algorithm fails to make any meaningful reduction of the objective function. The noise-free run also does not converge to a solution as good as the simplex method. This is probably because the objective function is evaluated with numerical noise. For Powell's method we also show data for a run with a 60-second interval which indicate that the algorithm works better with reduced noise level.

We then applied the multi-objective genetic algorithm NSGA-II [15] to the problem. The population of solutions was 100 seeds and the algorithm was run for 60 generations. The results are shown in Figure 3, for noise levels corresponding to a 6-second interval, 10-second interval and no random noise in the data. Although the objective functions (left plot) reach the same level at the end of about 6000 evaluations, the coupling ratio (right plot) indicates

Figure 2: The objective function (loss rate with reversed sign) of the best solution to-date during optimization for the simplex method (left) and Powell's method (right)

the algorithm performs better with lower noise. In fact, the solutions favored by noise become elite seeds and dominate the population. For example, the average error in the objective function for the 6-second case is -0.028 mA/min in the second generation and -0.080 mA/min in the 20th. The optimization results become more distorted as noise level gets higher. It is also seen that NSGA-II converges very slowly, with or without noise in the data.

Figure 3: The objective function (left) and the coupling ratio (right) of the best solution to-date during optimization for the NSGA-II method.

Figure 4 shows the performance for our RCDS method. The robustness of the RCDS method against noise is compelling. All three runs with random noise converge to solutions with the same level of objective function value found by the noise-free run. It is worth noting that the real difference is even smaller when noise is removed from the objective functions, as shown in the coupling ratio plots (Figure 4 right). The simulation runs converged to

nearly the same solution. Differences between the skew quadrupole strengths of the final solutions are small, less than 2% of the strengths themselves. The initial conjugate direction set is obtained from the Jacobian matrix of the orbit response matrix with respect to the skew quadrupole parameters. If the initial directions are the unit vectors, the convergence would take longer, similar to Powell's method in the noise-free case (Figure 2 right).

Figure 4: The objective function (left) and the coupling ratio (right) of the best solution to-date during optimization for the RCDS method.

The best performances for all algorithms are compared in Figure 5, which shows the coupling ratio of the best solution during the runs. The coupling ratio is without noise but the loss rate evaluation during optimization include random noise level for the 6-second interval case. The two algorithms that employ our line optimizer, RCDS and IMAT, are robust against noise, but the latter converges more slowly because the search directions are not pair-wise conjugate.

4.2. Optimization of injection optics match in simulation

As a second example, we use the optimization of the SPEAR3 Booster-to-SPEAR (BTS) transport line optics to demonstrate the performance of the algorithms. The phase space distribution of the injected beam at the injection point (*i.e.*, the septum magnet) should match the acceptance of the storage ring for the injection efficiency to be maximized. The horizontal phase space configuration for the SPEAR3 at the injection point is illustrated in Figure 6. The ellipse represents the acceptance of the storage ring and the vertical line at $x = -10$ mm represents the septum wall. Only particles that are injected into the acceptance and are to the left side of the septum

Figure 5: The coupling ratio of the best solution to-date during optimization for several methods, all with loss rate noise for a 6-second interval.

line will survive. The acceptance and the position of the septum wall in this simulation are only representative for SPEAR3. To leave room for injection efficiency improvement by the optimization algorithms, we intentionally reduced the horizontal dynamic aperture to 12.5 mm in the simulation. The vertical acceptance is not shown but is considered in the simulation to have a vertical dynamic aperture of 3 mm. The dots represent injected particles, which are generated randomly according to the optics functions at the end of the transport line and the emittances of the Booster ring. The chromatic contribution to the phase space distribution due the momentum spread and dispersion functions are also accounted for. The injection efficiency is calculated by counting the number of surviving particles among 1000 particles. Because of the finiteness of the total particles, the injection efficiency has a random noise level of 1.6%. The left plot in Figure 6 shows the situation for the initial, present BTS optics with a simulated injection efficiency of 61.7%.

In the simulation we use the last six quadrupole magnets to adjust the optics functions at the end of the transport line. The range for each parameter (quadrupole gradient) is constrained around the initial value by $\pm 0.2 \text{ m}^{-2}$. The injection efficiency with a reversed sign is used as the objective function. All five algorithms are again applied to this problem. The size of the initial simplex is 20% of the range for the simplex method. For the NSGA-II al-

gorithm (MOGA) the population size is 100 and it runs for 60 generations. Only 1000 evaluations are shown since there is very little improvement afterward. The initial conjugate direction set for the RCDS method is derived from the Jacobian matrix of the six σ -matrix elements (*i.e.*, σ_x^2 , $\sigma_{xx'}$, $\sigma_{x'}^2$ and likewise for the y-plane) with respect to the quadrupole parameters, as calculated from the initial optics model.

Figure 7 shows the best-to-date injection efficiency during the optimization runs for various algorithms. Once again the RCDS method converges fast and is robust against noise. The 1-dimensional scan with the robust line optimizer (IMAT) converges slower. Powell's method is more sensitive to random noise and the simplex method becomes sensitive to noise when the simplex is small. The genetic algorithm is inefficient and does not converge to the best solution since the population gradually becomes dominated by solutions favored by noise. The best solution of BTS optics found by the RCDS method has an injection efficiency of 85.0%, as illustrated in Figure 6 (right plot).

Figure 6: The horizontal injection optics match at the septum for the initial transport line optics (left) and the optics found by RCDS (right).

5. Experimental applications of the RCDS method

We have applied the RCDS method experimentally at the SPEAR3 storage ring to demonstrate online accelerator optimization. The applications include the minimization of the vertical emittance with skew quadrupoles, as described in sub-section 4.1. In the experiment the stored beam current is

Figure 7: The history of the best injection efficiency during optimization for the various algorithms.

maintained at a maximum of 500 mA with top-off injection every 5 minutes. The parameter range for the skew quadrupole currents is between -20 A and 20 A with the initial strengths for all 13 magnets set to zero. The initial loss rate is 0.43 mA/min, corresponding to a coupling ratio of about 1%. The loss rate is evaluated by computing the beam current loss in 6 seconds. The corresponding rms noise of the loss rate as measured by the standard deviation of 100 evaluations is 0.04 mA/min. The loss rate evaluation code waits until injection is over if there is not enough time before the next injection.

Figure 8 shows the objective function for the trial solutions during the run. With 200 evaluations, the loss rate reaches 1.75 mA/min. The best solution corresponds to a minimum beam lifetime of 4.60 hours at 500 mA beam current. For comparison, an earlier solution found by fitting orbit response matrix has a loss rate of 1.50 mA/min and a lifetime of 5.25 hours at 500 mA. According to the loss rate, the average vertical beam size is reduced by about 15% by the RCDS method relative to the orbit response matrix method. The two solutions are close in the parameter space, with an average difference of 1.7 A for the skew quadrupoles.

Another application is the optimization of the injection kicker bump match. SPEAR3 has three injection kickers that make a closed bump for the stored beam during injection. In the top-off injection mode the kicker

Figure 8: History of all evaluated solutions and the best-to-date solution during the loss rate optimization experiment with the RCDS algorithm.

bump should be accurately matched to eliminate perturbation to the stored beam. The parameters to be adjusted are the voltage, pulse delay and pulse width for kicker 1 and 3 while the kicker 2 parameters are fixed. Two skew quadrupoles that are within the bump are also adjusted to minimize the perturbation to the vertical plane. There are a total of 8 parameters.

To test the algorithm, we first deliberately put in errors to these 8 parameters to cause betatron oscillations in both the horizontal and vertical planes. The oscillation amplitudes are measured with a turn-by-turn BPM. The objective function is the sum of the horizontal and vertical rms orbit deviation from the first 100 turns after the kick. When the RCDS algorithm is used, it quickly brings the perturbations down to a minimum, as shown in Figure 9. The horizontal oscillation amplitude is reduced from 0.3 mm to 0.03 mm within 40 evaluations.

We have also applied the RCDS method to optimize the launching angles and positions of the injected beam to SPEAR3 with two horizontal and two vertical steering magnets, using injection efficiency as the objective function. The algorithm demonstrated robust performance, usually bringing the injection efficiency to a maximum value within about 20 evaluations.

Figure 9: History of all evaluated solutions and the best-to-date solution during the kicker bump matching optimization experiment with the RCDS algorithm.

6. Conclusions

In this paper we discussed the need of an automated online optimization algorithm and the general requirements for such an algorithm. Execution efficiency and robustness in the presence of noise are the two most important challenges. Using realistic accelerator optimization problems, we demonstrated with simulation that traditional methods, such as the simplex method, Powell's method and genetic algorithms are sensitive to noise in function values. The genetic algorithms tend to be inefficient for online application since the convergence is slow and it may fail to converge to the optimum because the population will be dominated by solutions favored by noise. The traditional 1-dimensional scan may be inefficient since the successive line scans may partially cancel the gains previously made.

We developed a method called the robust conjugate direction search (RCDS) method. The new method combines the conjugate direction set approach of Powell's method with a new robust line optimizer which considers the random noise in bracketing the minimum and uses parabolic fit of data points that uniformly sample the bracketed zone. In simulation this algorithm demonstrated high efficiency and strong robustness against noise.

The RCDS method was then successfully applied to online optimization experiments at the SPEAR3 storage ring. The applications include min-

imization of the vertical emittance, optimization of injection kicker bump matching and optimization of injection beam steering.

We should point out that this method is applicable to automatic tuning of complex systems other than accelerators. It may also be used for design optimizations which evaluate the objective functions with numerical simulation. In the latter case the line optimizer may be partially parallelized.

Acknowledgment

The study is supported by DOE Contract No. DE-AC02-76SF00515.

References

- [1] J. Safranek, Nucl. Instrum. Methods Phys. Res., Sect. A, **388**, 1-2, 27 (1997)
- [2] X. Huang, *et al*, ICFA Newsletter on Beam Dynamics, **44**, 60 (2007)
- [3] X. Huang, *et al*, Proceedings of PAC'09, Vancouver, Canada (2009)
- [4] R. Bartolini, *et al*, Phys. Rev. ST - Accel. Beams, **11**, 104002 (2008)
- [5] X. Huang, *et al*, Phys. Rev. ST - Accel. Beams, **13**, 114002 (2010)
- [6] L. Emery, *et al*, Proceedings of PAC'03, Portland, Oregon (2003)
- [7] Y. Cai, *et al*, Phys. Rev. ST - Accel. Beams, **15**, 054002 (2012)
- [8] M. Borland, ICFA Newsletter on Beam Dynamics, **57**, 48 (2012)
- [9] W. H. Press, *et al*, Numerical Recipes, 3rd ed, Cambridge University Press, (2007)
- [10] M.J.D. Powell, Computer Journal **7** (2), 155-162 (1964)
- [11] J. Nelder, R. Mead, Computer Journal, **7**, 308 (1965)
- [12] I. Bazarov, *et al*, Proceedings of PAC'05, Knoxville, TN (2005)
- [13] A. Terebilo, Proceedings of PAC'01, Chicago, IL (2001)
- [14] K. Ohmi, *et al*, Phys. Rev. E, **49**, 751 (1994)
- [15] K. Deb, IEEE Trans. on Evolutionary Computation, vol 6, no 2, (2002).