APPLICATION OF DIRECT METHODS OF OPTIMIZING STORAGE RING DYNAMIC AND MOMENTUM APERTURES

M. Borland †, L. Emery, V. Sajaev, A. Xiao, ANL, Argonne, IL 60439, USA
W. Guo, BNL, Upton, NY, 11973, USA

Abstract

Optimization of dynamic and momentum apertures is one of the most challenging problems in storage ring design. For storage-ring-based x-ray sources, large dynamic aperture is important in obtaining high injection efficiency, which leads to efficient operation and protects components from radiation damage. X-ray sources require large momentum aperture to obtain sufficiently long Touschek lifetimes with low-emittance beams. We have developed effective methods of optimizing dynamic and momentum apertures that rely directly on tracking using a moderately sized Linux cluster. After reviewing the method, we describe examples of its application to APS operations, upgrades, and next-generation storage rings.

INTRODUCTION

One of the most desirable characteristics of storage-ring-based x-ray light sources is low emittance. To achieve this, lattice designers use strong focusing to obtain large horizontal phase advance per cell, leading to large chromatic aberrations and thus strong chromaticity correcting sextupoles in order to obtain adequate momentum aperture (MA). In addition, low emittance means small dispersion, requiring yet stronger chromatic sextupoles. This leads to small dynamic aperture (DA), making it more difficult to accumulate beam. In extreme cases, the dynamic aperture may be so small that sufficient lifetime is not achieved.

Ring designers commonly add extra families of sextupoles to correct the effect of the chromatic sextupoles [1]. The challenge is to adjust the sextupoles to simultaneously maximize both DA and MA. Perhaps the most common approach is to minimize many resonance and tune variation driving terms [2]. However, one must carefully choose the weights for these terms, based on experience and, ultimately, tracking. Further, we commonly want non-zero linear chromaticity to suppress instabilities, which challenges the assumptions of the perturbative approach, since then one does not want the higher-order chromaticities to be minimized, but rather one needs to use them to reduce the chromatic tune spread.

In this paper, we discuss further a tracking-based optimization method [3] that has proven very successful and is a considerable improvement over previous attempts discussed in [4] and, in part, in [3]. Following an explanation of the method, we discuss application to the Advanced Photon Source (APS) storage ring and the NSLS-II ring.

Although our method could use any tracking code, the ability to create fully scripted simulations is essential, since matching and tracking must run without human intervention. Thus, we use the tracking program `elegant` [5, 6], as well as the SDDS Toolkit [7] and `geneticOptimizer` [8].

OPTIMIZATION METHOD

In this method we use many computers simultaneously to evaluate the DA and MA for various lattice tunings (e.g., tunes and sextupole settings). DA and MA computation includes radiation damping, synchrotron oscillations, and physical apertures. After completion of a sufficient number of evaluations, a genetic algorithm is used to "breed" more candidate configurations based on the best configurations seen so far. The process continues until a sufficiently good solution is obtained or until the results stop improving.

Dynamic Aperture

For the DA, we've found that the area of the stable region is a good parameter to use, with some limitations and conditions. We first determine the DA by performing line scans outward from the origin. (Scanning outward is used instead of scanning inward in order to avoid being fooled by stable islands.) Once the stable boundary is found, we analyze the boundary points to clip off any regions that "stick out" in a manner that indicates a poorly behaved boundary. An example is shown in Figure 1: The region that sticks out on the right side is probably related to a stable island and is not considered a useful contribution to the DA. Finally, having found the clipped DA boundary, we compute its area and its contribution to the penalty function. Because the area computation ignores contributions from useless regions, the optimized results are unlikely to display such regions.

The contribution to the penalty function is computed by comparing the area A to the desired area A_d using a weighting factor ΔA

$$P(A) = \begin{cases} (A - A_d)^2/\Delta A^2 & A < A_d \\ 0 & A \geq A_d \end{cases} . \qquad (1)$$

For APS we typically want an aperture $-13\text{mm} \leq x \leq 7\text{mm}$ and $|y| \leq 1.5\text{mm}$, giving $A_d = 30\mu\text{m}^2$.

In some cases, the DA area may be misleading, for example, a solution with large vertical aperture but small hori-

*Work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357.

† borland@aps.anl.gov

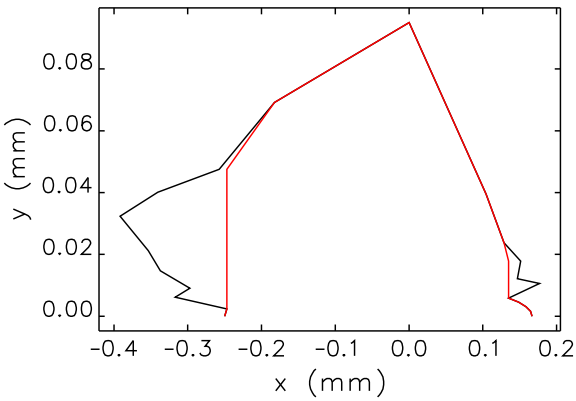

Figure 1: Illustration of algorithm for removing pathological features from DA results. The red boundary is used for DA area computations.

zontal aperture. To prevent this, we limit the vertical search region to the maximum desired vertical DA.

Momentum Aperture

To assess the Touschek lifetime, we need to know the MA for particles scattered at many positions around the ring, i.e., the s-dependent MA [9]. We don't need the MA for every lattice element, but simply for a representative set. A good choice for a double-bend cell is to compute the MA at several locations on either side of the dipoles. For the APS, we compute the MA at the exit of the S1, S3, and S4 sextupoles in the first six (out of 40) sectors.

We compute the MA's contribution to the penalty function in a fashion similar to what was done for the DA:

$$P(\delta_{min}) = \begin{cases} (\delta_{min} - \delta_{des})^2/\Delta\delta^2 & \delta_{min} < \delta_{des} \\ 0 & \delta_{min} \geq \delta_{des} \end{cases} . \quad (2)$$

where δ_{min} is the minimum absolute value of the MA at any location, δ_{des} is the desired value, and $\Delta\delta$ is a weighting factor. For APS, we typically choose $\delta_{des} = 2.35\%$, equal to the bucket half-height.

An alternative is to compute the Touschek lifetime from the s-dependent MA using the `touschekLifetime` program [10], which operates on output from `elegant`. As of now, this additional step has not seemed necessary.

Importance of Errors

It is well known that the DA and MA are affected by magnet strength and alignment errors. For example, without errors the effect even of the half-integer resonance may not be seen in a tracking simulation. Hence, we must include lattice errors in the simulations. It would seem that in order to be realistic, we must not only include errors, but also correct those errors using real-world techniques.

Effective methods exist for correcting linear optics [11] and coupling (e.g., [12, 13, 14]), which is important in light sources because of the small insertion device vertical apertures. In the APS, for example, we correct lattice function errors to the 1% rms level [15] and coupling to the 1% level.

However, correction is not essential in the simulations. Instead, we simply use random errors that give lattice function and coupling errors at post-correction levels. This neatly side-steps a considerable complication.

To prevent the optimizer from being misled by variations resulting from different ensembles, we use a fixed error ensemble for all simulations. We impart strength and roll errors to quadrupoles and sextupoles only, which gives all the essential features of a post-correction lattice. Strength errors are typically 0.02% rms, which typically gives lattice function errors of 1% rms or more. Roll errors are typically 0.5 mrad rms, which gives coupling of roughly a few percent. We prefer to make the errors somewhat on the high side, since this helps ensure a robust solution. Following optimization, we evaluate the lattice with typically 20 to 50 ensembles to verify that the solution is robust. This is a simple precaution that can be easily carried out with, e.g., `elegantRingAnalysis` [16]. To date, we have seen only one case where this step yielded a surprise. This was traced to a poor choice of integer tunes.

Software Details

The optimization uses the general-purpose script `geneticOptimizer`. The user provides an input file listing the independent variables along with their initial values, allowed ranges, and randomization levels. Typically the indepenent variables are the tunes and the sextupole strengths. (Although not required by our method, allowing the tune to vary has been very effective and is advisable.) The user also provides two scripts, one to run a configuration and another to postprocess it. The first script (the "run script") performs any required matching and assembles the full lattice, then performs tracking to determine the DA and MA. The second script postprocesses the results and returns the value of the penalty function.

As an example, for some future long-straight-section lattices for the APS, the run script invokes `elegant` four times to perform matching of different types of sectors. Rematching permits wide variation of the tune without changing other essential lattice features, such as maximum lattice functions or lattice functions at the insertion devices. It also ensures that there are no spurious lattice function errors. Following matching, the sector solutions are combined into a ring solution that is evaluated for DA/MA in a single run. Often this run also performs the final chromaticity correction using whichever sextupoles are left free.

Executing and combining the results of multiple program runs is greatly facilitated by `elegant`'s thorough use of SDDS files and the SDDS toolkit. It permits complete automation of the process, regardless of the complexity, using only commands in a script.

APS OPTIMIZATION

The APS storage ring has 280 sextupole magnets with individual power supplies. Because of the symmetry of the lattice, in the past these sextupoles have been powered in four families. Because we run in modes with fairly high single-bunch current while lacking a bunch-by-bunch feedback system, we operate with significant non-zero chromaticities $\xi_x = d\nu_x/d\delta$ and $\xi_y = d\nu_y/d\delta$. In 24-bunch, 100-mA mode, we have $\xi_x \approx 7$ and $\xi_y \approx 6$, while for hybrid mode we require $\xi_x \approx \xi_y \approx 11$ to achieve the required 16-mA single-bunch current.

Operational Lattices

We previously reported [3] on successful application of this technique to improvement of the APS operational lattice, where we realized a 25% improvement in lifetime for the 24-bunch mode. One of the most surprising aspects of this work was that in several cases the optimization clearly favored a sextupole configuration that did not have the same symmetry as the lattice itself. One advantage of the new configuration is that the lifetime is now long enough to be used in 324-bunch, non-top-up mode, which reduces the number of lattices we must maintain. Since then, we've also optimized the hybrid mode lattice and obtained a 10% improvement in lifetime, which is less than the predicted improvement of 20%, but still significant.

APS Renewal Lattices

The APS Renewal is a project to update the APS accelerators and beamlines. One of the most interesting aspects of the accelerator improvements is the provision of a number of long straight sections (LSS). These will allow 7.7 m for insertion devices compared to 4.8 m at present. We previously reported [3] on development of several such lattices and successful testing of a mockup lattice with eight symmetrically placed, emulated long straight sections.

Since then, we've explored several additional lattices. Among these are four groups of two long straights separated by a short straight (4x2LSS) and 10 symmetrically placed long straights (10LSS). Both of these yielded solutions that can be expected to work in practise. One advantage of our optimization method is that it works with little human involvement beyond initial setup, which is similar for most configurations. Hence, we can quickly look at many different possibilities. Further, if the optimization converges, the solution is almost certain to be valid as it is based to begin with on tracking.

The 4x2LSS configuration is a complex case that illustrates how the commandline nature of `elegant` and `SDDS` contribute to the implementation. APS has two types of sectors: those with Decker distortion [17] (sectors 1 through 35) and those without (sectors 36-40). Hence, we must match two kinds of Decker-distorted sectors, one with a short straight on both ends and a second with a short straight on one end and long straight on the other. A third solution is required for the non-distorted sector. We constrained all sectors to have very similar tunes, finding a solution for the normal sector first and then matching the other sectors to that solution. Using SDDS tools allowed us to overlay the reflected short-to-long solution on the sectors requiring long-to-short optics, which is necessary because we wish to independently vary all the sextupoles without symmetry constraints and hence have different element names in these sectors. Following this, we match the non-distorted sector to the same initial and final lattice functions and the same phase advance as the Decker-distorted sectors. The linear optics solutions are then loaded over the full lattice and the sextupoles are set according to the values provided by the optimizer. The free sextupoles are adjusted to give the desired chromaticity, subject to sextupole strength limits. Finally, we track to obtain the DA and MA. The DA, MA, and chromaticity are then used in the penalty function computation, as described above.

NSLS-II OPTIMIZATION

The NSLS-II ring is currently in the construction stage. The lattice of this 3-GeV machine features 30 sectors with alternating long and short straight sections. Strong damping wigglers are employed to bring the emittance below the 1-nm level. Considerable work has been done on design and optimization of this lattice for near-zero chromaticity [18, 19] using minimization of resonant driving terms. Here, we report results of optimization using our direct technique. Unlike the previous work, we allowed all sextupoles to vary independently. (Hence, our results do not necessarily indicate an advantage of our method.) This choice was based on experience with APS optimization, which taught us that symmetry constraints on the sextupoles are not necessarily desirable [3].

We began with a somewhat mistuned lattice having $\xi_x \approx 1$ and $\xi_y \approx 3$, but exhibiting poor DA and MA. The first stage of optimization targeted $\xi_x = \xi_y = 2$, a DA area of 100 μm^2, and an energy aperture of 3% using 40 processors, with tunes free to vary. The optimization converged after about 500 runs (see Figure 2), moving the tunes from $\nu_x = 33.1$ and $\nu_y = 16.20$ to $\nu_x = 33.25$ and $\nu_y = 16.09$. This optimization gave an unnecessarily large vertical aperture, so we modified the parameters to only scan to 1.5 mm in the vertical instead of 3.0 mm, and re-optimized starting from the previous best value, thus increasing the horizontal aperture. We evaluated the result with 20 random ensembles (using the typical error levels given above), giving excellent results as seen in Figures 3 and 4. Note that multipole errors have yet to be included.

Starting from the $\xi_x = \xi_y = 2$, we continued the optimization with a target of $\xi_x = \xi_y = 4$. The resulting dynamic aperture was nearly identical to that shown in Figure 3, while the momentum aperture was slightly reduced. Obtaining this new solution was relatively effortless, requiring changing only a few values in the penalty function.

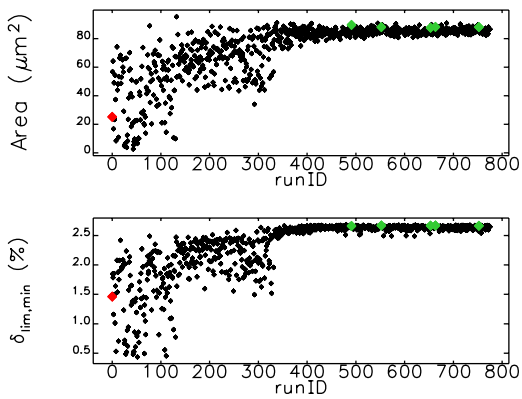

Figure 2: Progress of initial optimization of NSLS-II for $\xi_x = \xi_y = 2$. To aid convergence, the error level was adjusted by hand at around run 150 and again at around run 350.

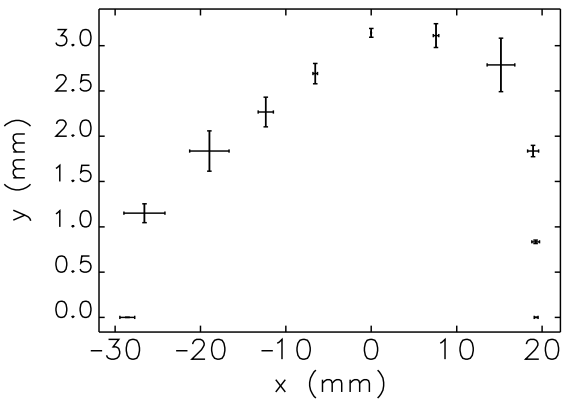

Figure 3: Dynamic aperture for 20 ensembles for final optimization of NSLS-II for $\xi_x = \xi_y = 2$.

FUTURE DEVELOPMENT

Although this technique is highly successful, there are still opportunities for improvement. Using parallel DA and MA computation[20], we could use many cores for each job submitted by the genetic optimizer, which promises convergence in hours instead of weeks. We could also perform optimization with several error ensembles rather than a single ensemble, in order to further ensure robustness of the final result. We are interested in trying other optimization algorithms besides the genetic technique, e.g. parallel simplex [4]. A related option is to use `elegant`'s internal optimizer, which will be possible in the near future once DA and MA optimization is supported in the parallel version. As mentioned above, a refinement of the penalty function would be to compute the momentum aperture directly using the `touschekLifetime` program. One could also envision simulating injection efficiency instead of only the dynamic aperture, which might yield interesting new solutions with highly asymmetric dynamic apertures (i.e., large on the side where beam is injected, but smaller on the other side). Incorporating lattice and coupling correction is also desirable because it will ensure that choices of tune are not overly constrained by correctable effects of errors.

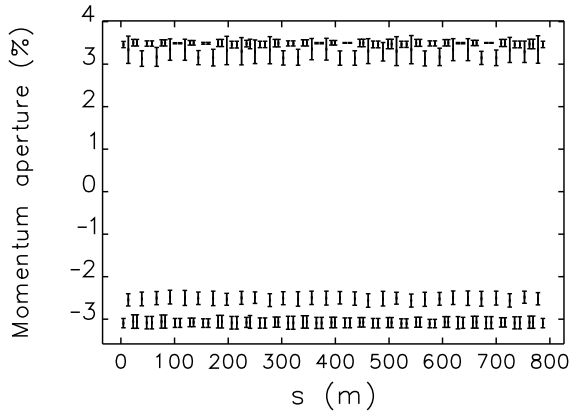

Figure 4: Momentum aperture for 20 ensembles for final optimization of NSLS-II for $\xi_x = \xi_y = 2$.

CONCLUSION

We have developed a practical and robust tracking-based method of optimizing storage ring nonlinear dynamics. Details of the algorithm and implementation were presented. An update on APS-related applications was given, followed by examples of successful application to the NSLS-II ring design.

REFERENCES

[1] E. A. Crosbie, Proc. PAC87, 443-445 (1988).
[2] J. Bengtsson, SLS Note 9/97 (1997).
[3] M. Borland *et al.*, Proc. PAC09, TH6PFP062, to be published.
[4] H. Shang and M. Borland, Proc. PAC05, 4230-4232 (2005).
[5] M. Borland, APS LS-287, September 2000.
[6] M. Borland, these proceedings.
[7] M. Borland, Proc. PAC95, 2184-2186 (1996).
[8] M. Borland and H. Shang, unpublished program.
[9] M. Belgrounne *et al.*, Proc. PAC03, 896-898 (2003).
[10] A. Xiao and M. Borland, Proc. PAC07, 3453-3455 (2007).
[11] J. Safranek, NIM A 388, 27 (1997).
[12] P. Bagley *et al.*, Proc. PAC89, 874-876 (1990).
[13] J. Safranek and S. Krinsky, Proc. PAC93, 1491-1493 (1993).
[14] L. Emery *et al.*, Proc. PAC03, 2330-2332 (2003).
[15] V. Sajaev and L. Emery, Proc. EPAC 2002, 742-744 (2002).
[16] M. Borland, Proc. PAC05, 4200-4202 (2005).
[17] G. Decker and O. Singh, Phys. Rev. ST Accel. Beams 2, 112801 (1999).
[18] S. Kramer and J. Bengtsson, Proc. PAC05, 3378-3380 (2005).
[19] W. Guo *et al.*, Proc. PAC09, TU5RFP008, to be published.
[20] Y. Wang *et al.*, these proceedings.