Determination and correction of the linear lattice of the APS storage ring

V. Sajaev, L. Emery

Motivation

The APS storage ring is a very complex machine containing 400 quadrupoles and 280 sextupoles. Each quadrupole has a separate power supply, which could have calibration errors, and each sextupole could have a non-zero beam orbit in it, which results in additional quadrupole errors. The quadrupole calibration errors and non-zero orbit through the sextupoles are the two main sources of linear optics distortion, and this distortion could seriously affect the performance of the storage ring.

From the beginning of the APS storage ring operation there was a substantial difference between the linear model and the real storage ring. We even have to use some empirical correction factors when transforming the model into the real machine and back to get the betatron frequencies correct. This results in difficulties when tuning the machine to new lattice conditions such as the low-emittance lattice or converging beta function lattice. That is why we decided to develop a fast method for linear lattice calibration using orbit response matrices.

There are several other problems that can be solved using the response matrix fit method:

- Beta function measurements around the ring. Right now the only way to measure the beta functions at the APS is by varying quadrupole gradients. It takes about seven minutes for one quadrupole, so to accurately measure the beta functions at just one point per sector takes almost five hours.
- BPM gain calibration. There are more than 400 beam position monitors (BPMs) around the ring, and many of them have substantial gain errors. Right now there is no reasonably fast way to calibrate them all.
- Local linear coupling characterization and correction.
- Knowing the precise linear model will probably help to better understand the nonlinear character of the ring.

Method description

The orbit response matrix is the change in the orbit at the BPMs as a function of changes in steering magnets. The response matrix is defined by the linear lattice of the machine; therefore it can be used to calibrate the linear optics in a storage ring. Modern storage rings have a large number of steering magnets and precise BPMs, so measurement of the response matrix generates a very large array of precisely measured data. Analysis of this data is a powerful and accurate way to recover the linear optics of a storage ring.

The main idea of the analysis is to adjust the quadrupole gradients of a computer model of the storage ring until the model response matrix best fits the measured response matrix. The method was first suggested (to the author's knowledge) by Corbett, Lee, and Ziemann at SLAC [1] and refined by Safranek at BNL [2]. A very careful analysis of the response matrix was done at the NSLS X-ray ring [3] and at the ALS [4]. A similar method was used at the ESRF for characterization and correction of the linear coupling [5] and to calibrate quadrupole families [6]. There are a number of papers in the Particle Accelerator Conference Proceedings that describe similar model calibration techniques.

The problem of fitting the response matrix is solved in the following way. Let the response matrix M be a function of the vector of variables x. Then we need to solve the equation

$$M_{measured} - M_{model}(x) = 0 \,,$$

which can be solved by Newton's method:

$$M_{measured} - M_{model}(x_0) - \frac{\partial M_{model}}{\partial x} \cdot \Delta x = 0 \,,$$

$$\Delta x = \left(\frac{\partial M_{model}}{\partial x} \right)^{-1} \cdot \left(M_{measured} - M_{model}(x_0) \right) , \qquad (1)$$

where x_0 corresponds to the initial model. Finally, to fit the response matrix, we have to determine all variables that the response matrix depends on, calculate the derivative of the response matrix with respect to the vector of variables, and then invert it. After that, the solution can be found by iteration.

The most obvious and important variables are focusing errors (quadrupole calibration errors or orbit errors in sextupoles), corrector calibration errors, and BPM gain errors. Another obvious but less important set of variables is the energy shift associated with the changing of each steering magnet. These are the variables that are used for the response matrix fit described in this report. An example of other variables can be found in ref. [4], where four parameters were used to fit the signal from each of four BPM buttons, and the longitudinal positions of the steering magnets were also varied. The decision on what variables to use depends on details of the particular storage ring and how accurately the response matrix can be measured. For the APS storage ring, another possible set of variables could come from nonlinearity of the P0s (narrow bandwidth) BPM signals.

Application of the method to the APS storage ring

Difficulties: model size and degeneracy

Up to this time, the most comprehensive analysis of the response matrix has been done at the NSLS X-ray ring and the ALS. These two storage rings are much smaller than the APS. In the case of the X-ray ring, 626 variables were varied to fit 8,640 elements in the response matrix. At the ALS a total of 500 variables were varied to fit 15,744 elements in the matrix. In case of the APS, if one would try to use all correctors and BPMs, there would be 2,240 variables to vary and about 560,000 elements to fit. The size of the response matrix derivative would be 9 Gb and is much larger than the available size of memory in average computers. In addition, the computation time would be many days.

There are two sources of model degeneracy that the APS storage ring has that the other smaller rings lack. First, the smaller rings are able to store the beam without the sextupoles powered. This allows them to separate two kinds of gradient errors: quadrupole imperfections and orbit errors in sextupoles. Second, the average betatron phase advance between quadrupoles at APS is 0.088, while for the NSLS X-ray ring it is 0.17 and for the ALS it is 0.28.

Choice of variables

As mentioned above, use of the entire APS storage ring response matrix for the fit would require inversion of a huge 2,240 × 560,000 matrix. Right now this is impossible without using a "supercomputer." That is why we have to limit the size of the measured response matrix.

In general, the size of the uncoupled response matrix is

$$N_{elements} = N_{corr\,x} \times N_{BPM\,x} + N_{corr\,y} \times N_{BPM\,y} ,$$

where N_{corr} and N_{BPM} are the number of correctors and BPMs in the x and y planes The number of variables is

$$N_{variables} = N_{quads} + 2 \cdot N_{corr\,x} + N_{BPM\,x} + N_{corr\,y} + N_{BPM\,y} ,$$

where the factor of 2 in front of N_{corrx} comes from the energy variation.

One way to decrease the size of the response matrix is to fit the response matrix of only a part of the ring and to represent the rest of the ring as a linear transfer matrix. We were unable to get a high precision fit and do not understand why. We probably have to use nonlinear transfer matrices.

Another way to decrease the response matrix size is to reduce the number of steering magnets and BPMs. However, in order to achieve a high precision fit, we have to use as many BPMs as possible. This leaves only one available option – reducing the number of steering magnets.

The most obvious minimal set of steering magnets is to use one steering magnet per sector. Then in case of one corrector, nine BPMs (narrow bandwidth BPMs excluded), and nine quadrupoles (B:Q5 excluded) per sector per plane, the size of the response matrix derivative is

$$N_{variables} \times N_{elements} = 1200 \times 28800 \ .$$

For double-precision calculations the size of the response matrix derivative is about 260 Mb. The size of the computer memory required to invert a matrix and then manipulate it is about four times the size of the matrix, i.e., in this case about 1.0 Gb. This example is presented as one set of variables that fits 1 Gb of memory. This set of variables is usually used for our calculations.

Measurements and fitting

Two applications are used for the response matrix measurement and fit: the Operations Analysis Group (OAG) application *SROrbitResponse* is used to measure the response matrices, and the *SRLOCOFitting* application performs the model fitting. The output of the *SROrbitResponse* application is a file of the measured response matrix of the storage ring; this file is used as input for the fitting application. The output of the *SRLOCOFitting* application is a file of fitted variables in the format of the "parameter" file of the `elegant` code [7]. This file is used to update the ideal `elegant` model of the storage ring. When this is done, the calculated `elegant` response matrix fits (within the accuracy of the fit) to the measured one. This updated model can then be used for all kinds of calculations in `elegant`, including calculation of the beta function around the ring.

The measurement takes about 35 seconds per steering magnet, for a total of about 50 minutes for the measurements of one steering magnet per sector per transverse plane. This time does not depend on the number of BPMs used, so all BPMs are always used in the measurements. The accuracy of the response matrix measurement (or repeatability) can be estimated by measuring the "zero" response, i.e., response measured at the BPMs without any changes of the steering magnets. According to such measurements, the accuracy of the response matrix measurement is about 1 μm. It is interesting to note here that the accuracy of the narrow bandwidth BPMs is, as expected, better; however, their nonlinearity prevents them from being used in the fit.

The fitting process, which is just solving equation (1), consists of two main parts. First, the inverse derivative of the response matrix is calculated. This step is time consuming and can take up to several hours of computation time (even when we use parallel computing on many workstations). However, we must calculate the inverse derivative of the response matrix only once for a particular response matrix configuration. The second part is the fitting itself. This is done in several iterations using the calculated inverse derivative of the response matrix and takes about half an hour.

Figure 1 (top) is an example of the measured orbit response together with the ideal model response. The bottom plot of the Figure 1 shows the same trajectory after the fitting of the model is done. Before the correction, a typical rms difference between the ideal model and the measured response matrix is 50 μm. After the fit is done, the rms difference is decreased to the noise level of the BPMs, which is about 1 μm.

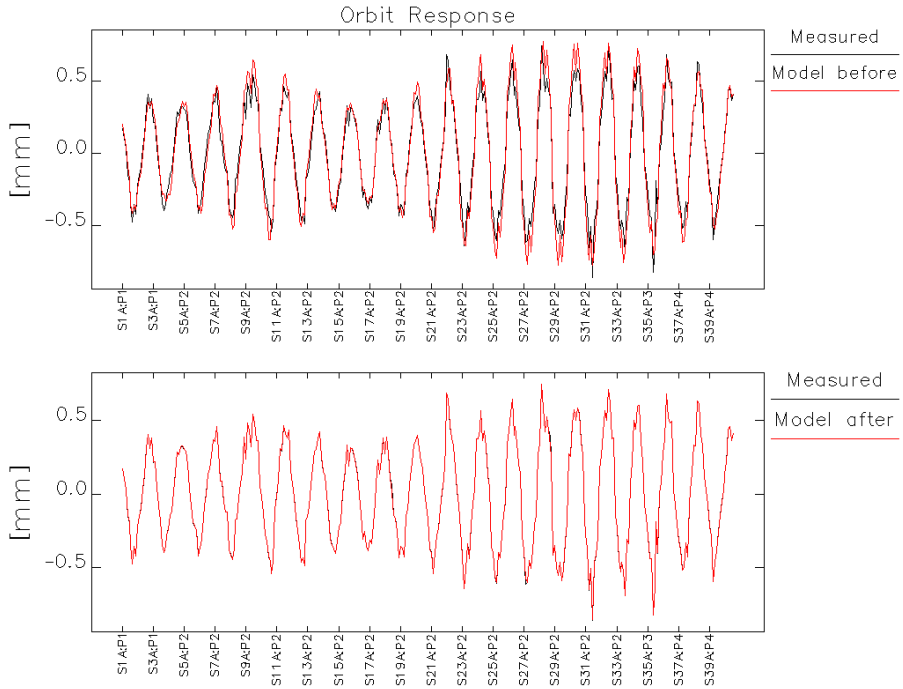

Figure 1. An example of the orbit shift due to a change in a steering magnet (black line is the measurement, red line is the model). Before fitting (top plot), there is significant discrepancy between the model and the measurements. After fitting (bottom plot), the rms difference between the model and the measured orbit is reduced to 1 μm.

After this nearly perfect agreement between the model and measured data is achieved, we have to ask: does this agreement necessarily imply a good agreement between the fitted model and the real elements in the storage ring? Although the number of data points (28800) is much greater than the number of fit parameters (1200), this does not guarantee the solution is unique. In particular, the redundancy of quadrupoles in the storage ring and the lack of ability to store the electron beam without sextupoles make it impossible to determine separate quadrupole errors. In other words, it is very likely that the measured response matrix can be reproduced (within the accuracy of the measurements) using different sets of quadrupole gradients. However, one would expect that the BPM gains and steering magnet calibrations should be unique.

The easiest way to confirm the above statement is to measure several orbit response matrices, analyze each one separately, and see how much variation there is between the fit parameters for the different data sets. Unfortunately, this requires too much study time. In practice, the response matrix for A:H2s, B:H2s, A:V1s, and A:V4s has been measured. It was then was split in two parts and analyzed by pieces using two different

configurations of the steering magnets: A:H2s + A:V4s and B:H2s + A:V1s. Such a procedure guarantees that, on the one hand, the measurements were done exactly at the same time with the same electron beam parameters. On the other hand, not only are the measurements independent, but the inverse derivative of the response matrix must be recalculated for both steering magnet configurations, thus ensuring that there is a significant difference between the two calculations.

Figure 2 shows the results of fitting using the two different configurations of steering magnets. As expected, the solution for the quadrupoles is ambiguous and shows considerable difference between the two sets, while the solution for BPM gains does not depend on the response matrix configuration.

As a confirmation that the two different quadrupole sets indeed represent the same response matrix, Figure 3 shows the vertical beta functions calculated by `elegant` for both quadrupole sets. We can see that these two plots are for all practical purposes identical. This figure demonstrates that, in spite of the quadrupole calibration ambiguity, the fit provides a unique storage ring model in terms of beta functions. The rms difference between the beta functions calculated using different sets are 0.1 m for horizontal and 0.2 m for vertical beta functions. These rms numbers can be used as an estimation of the accuracy of the beta-function determination. Assuming the average beta function around the ring is about 10 m, the relative rms accuracy would be of the order of a few percent.

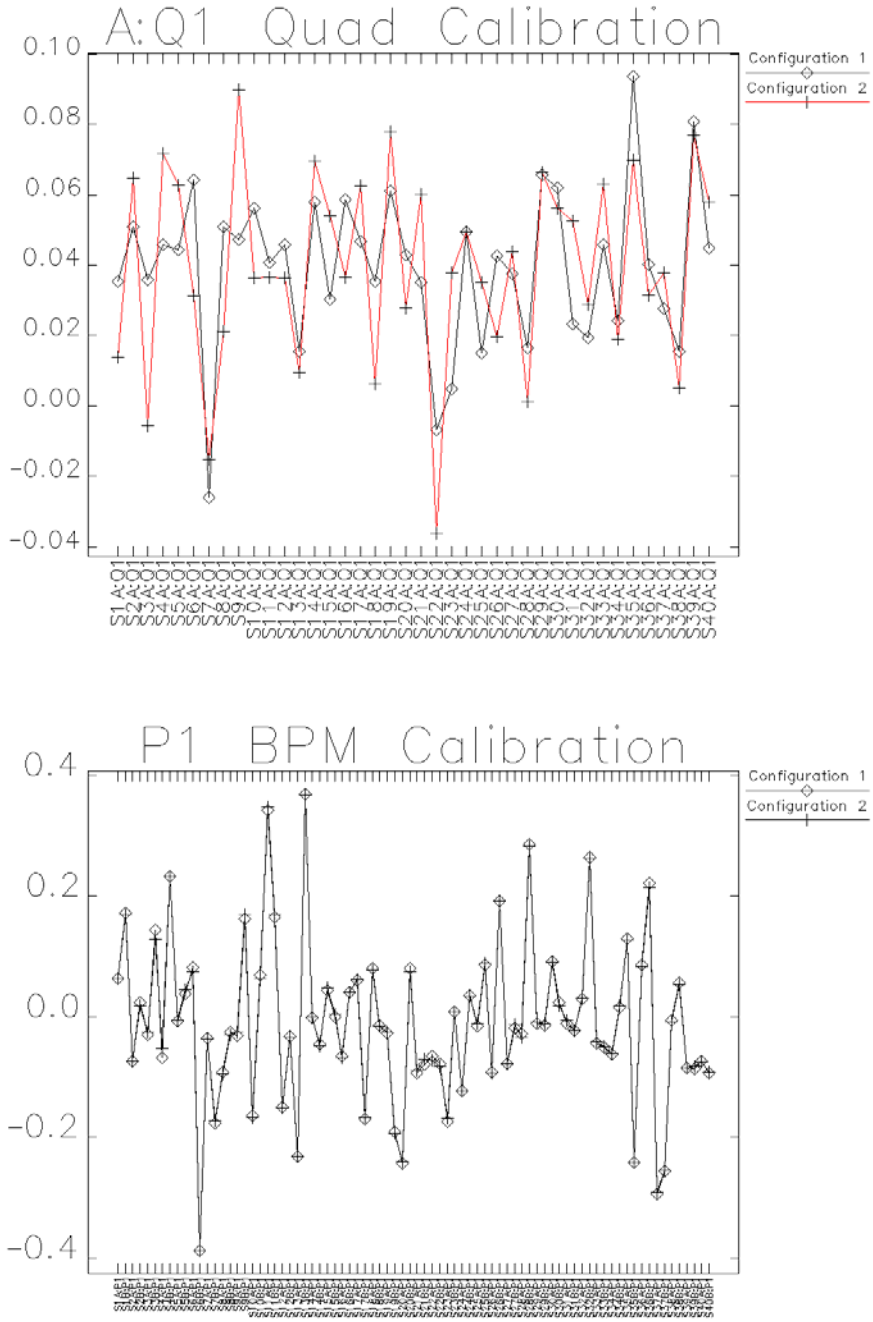

Figure 2. Some results of quadrupole calibrations (top) and BPM gains (bottom) using two different configurations of the response matrix. As expected, different configurations resulted in different quadrupole calibrations but the same BPM gains.

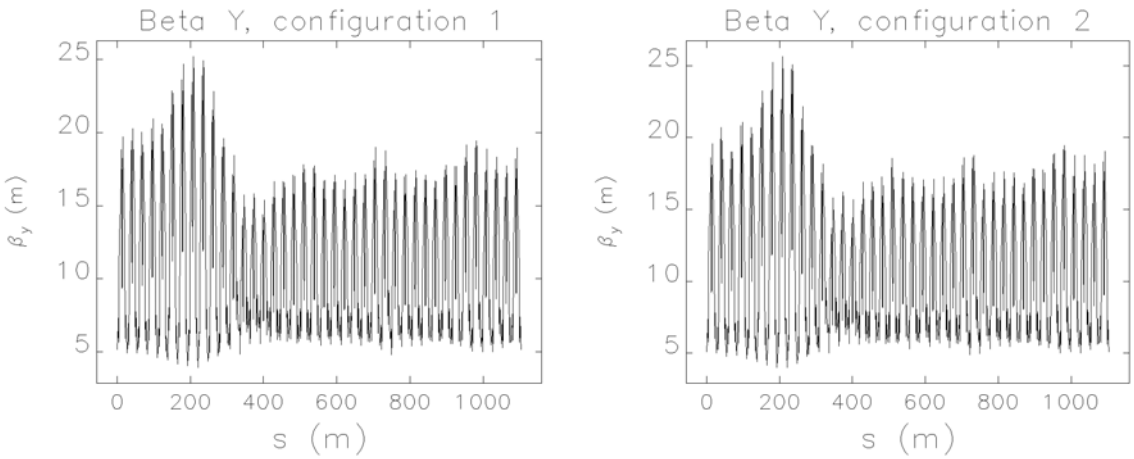

Figure 3. Vertical beta function of the storage ring (8 nm×rad lattice) according to the response matrix fit. The two plots represent two different response matrix configurations used for fitting the model. In spite of having different quadrupole gradients, these two models give the same beta functions.

Beta function beating correction

The model created by the response matrix fit represents the real machine in terms of beta functions and can be used for all kinds of calculations. First, the model can be used to calculate beta function beating corrections. This work has been done with both the low-emittance and the high-emittance lattices. The fitted model was used to calculate the beta functions, and then the *SRbetaCorrection* application was used to compute the quadrupole corrections. This application reads a set of desired beta function values at quadrupoles and a set of measured beta-function values at quadrupoles, and then uses an inverse matrix multiplication to determine a set of gradient corrections. This application was previously used to correct beta function modulation using a sparse set of measured beta functions from quadrupole scans [8]. The corrections were then applied to the storage ring, and the response matrix measurement and fit were performed again. The beta functions resulting from the "after correction" response matrix fit are shown in Figures 4 and 5.

It is important to note that the effect of the quadrupole corrections on the storage ring exactly coincides with the effect predicted by the model. In other words, the change in the beta functions after the correction is exactly the same as that intended by the beta function correction application.

Corrected beta functions improve the symmetry of the machine; this in turn should improve the nonlinear beam dynamics and overall performance of the storage ring. The positive effect of the correction was indeed observed for the low-emittance lattice; the lifetime was increased by 40% and the injection efficiency was improved.

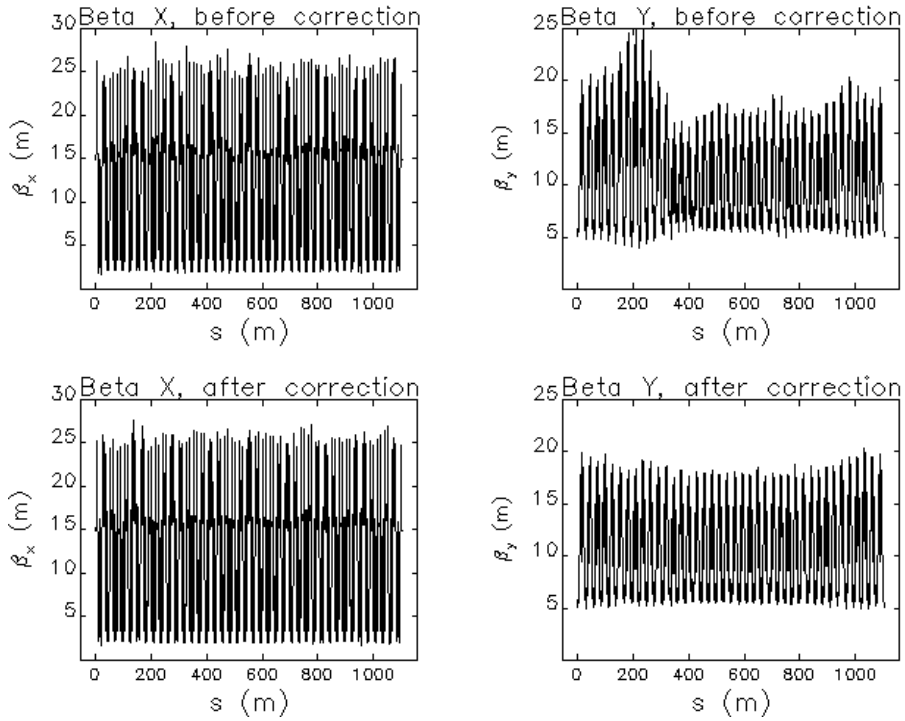

Figure 4. Beta functions of the "high-emittance" lattice before correction (top row) and after correction (bottom row).

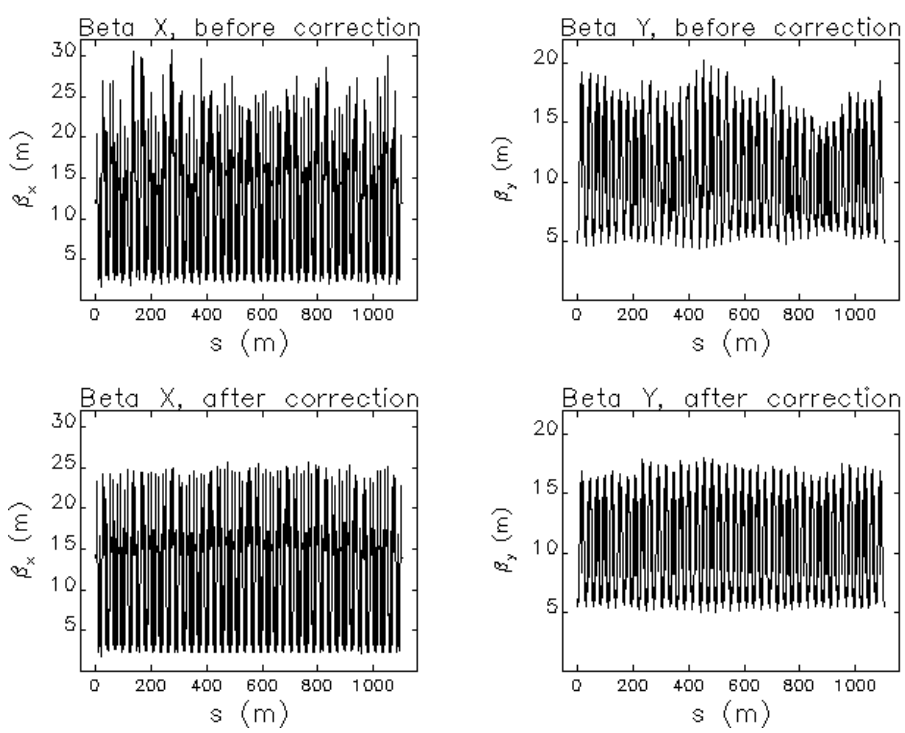

Figure 5. Beta functions of the "low-emittance" lattice before correction (top row) and after correction (bottom row).

At the APS storage ring the lifetime is defined by the nonlinear energy acceptance. To confirm that the energy acceptance was increased after the beta function correction, the lifetime dependence on the rf voltage was measured. Figure 6 shows the lifetime vs. rf voltage taken on three different dates. The important feature of this plot is not the absolute lifetime, but the gap voltage where the lifetime achieves a maximum (the overall lifetime is dependent on bunch pattern and coupling, which we didn't reproduce for all measurements). This voltage is a measure of the energy acceptance. The first curve corresponds to the initial low-emittance lattice without any quadrupole corrections; the best lifetime is achieved at 8.0 MV. The second curve was measured after the first correction was applied; the maximum is achieved at 8.5 MV. This correction was based on the beta functions measured in several quadrupoles by the quadrupole scan method. The third curve was measured after the correction based on the response matrix fit. The lifetime maximum is achieved on or beyond 9.5 MV.

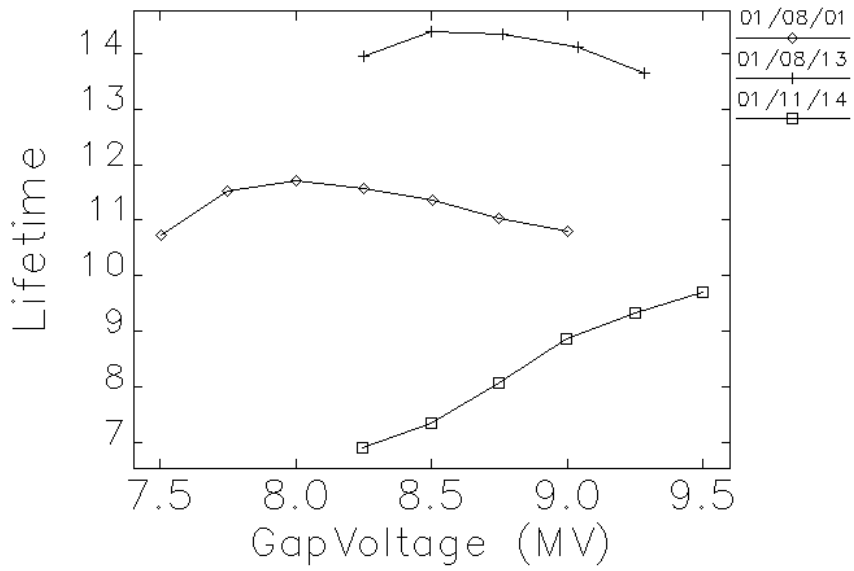

Figure 6. Lifetime dependence on the rf voltage. The first curve corresponds to the initial low-emittance lattice without any quadrupole corrections, the second curve was measured after the first correction was applied, and the third curve was measured after the correction based on the response matrix fit. The important feature of this plot is the gap voltage where the lifetime achieves its maximum.

BPM gain calibration

APS has over 400 BPMs in each horizontal and vertical plane. Some of these BPMs have substantial gain errors. That is why the BPM gain errors are used as variables to fit the measured response matrix. As mentioned above, the fit is reliable and provides BPM gain errors that are independent of the solution configuration. Figure 7 shows the results of the gain calculations for all BPMs (except narrow bandwidth BPMs) summarized in histograms. Here a gain of 1 means 0% gain error, 1.1 means 10% gain error and so on. The example of the gain correction applied to the dispersion measurements is presented in Figure 8. This plot shows the horizontal dispersion of the low-emittance lattice at the

B:P5s locations as measured by the dispersion measurement application (black line) and the same measurements after applying the gain corrections (red line). We can see that after correcting the gain, the dispersion appears less perturbed and smoother (i.e., more physical).

Source not found!

Figure 7. Horizontal and vertical calibrations of the BPM gains.

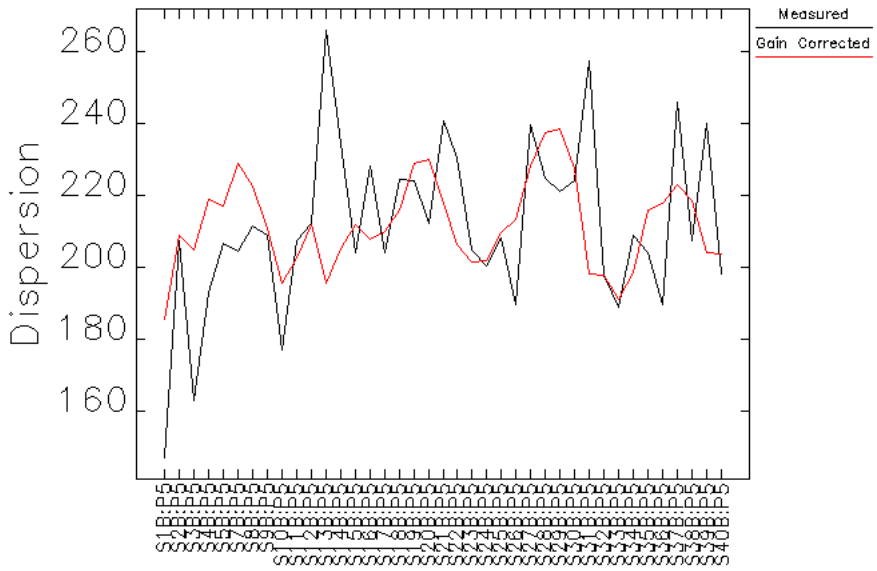

Figure 8. Dispersion function of the low-emittance lattice at the B:P5 locations around the ring. Black line is direct measurements; red line is after the gain correction.

Conclusions and future plans

We have created precise linear models of the storage ring in terms of beta functions for both low-emittance and high-emittance lattices. Using these models, the beta function beating corrections have been successfully applied. The lifetime was increased by 40% for the low-emittance lattice as a result of the corrections.

The models allow the user to apply predictable and precise changes to the existing lattice. For example, after applying the beta function corrections, the beta-function changes exactly coincide with the changes predicted by the model.

The generated machine model contains not only the real beta functions of the ring but it also provides BPM gain calibrations. This could be useful for a number of applications that rely on the BPM readings.

Unfortunately, one important task has not been accomplished. The excessive number of quadrupoles per unit phase advance does not allow us to determine the real gradients for every quadrupole. Work on correcting this shortfall will continue in the future.

Among other future plans is use of the response matrix fit to characterize and correct both global and local linear coupling. The implementation of a coupling calculation is rather straightforward but requires additional programming.

This work would not be possible without the help provided by many APS people. In particular, one of the authors (VS) would like to thank S. Milton for stimulating and supporting the work, and M. Borland for his tremendous support with regard to the storage ring operation and software implementation.

References

[1] W.J. Corbett, M. Lee and V. Ziemann, Proceedings of the PAC'93, Washington, DC, p. 108 (1993).
[2] J. Safranek, Proceedings of the PAC'95, Dallas, TX, p. 2817 (1995).
[3] J. Safranek, NIM A 388, p. 27 (1997).
[4] D. Robin, J. Safranek, W. Decking, Phys. Rev. Special Topics – Accelerators and Beams, v. 2, 044001 (1999).
[5] R. Nagaoka, Proceedings of the EPAC'2000, Vienna, Austria, p. 131 (2000).
[6] L. Farvacque and R. Nagaoka, Proceedings of the EPAC'98, Stockholm, Sweden, p. 909 (1998).
[7] M. Borland, "`elegant`: A Flexible SDDS-Compliant Code for Accelerator Simulation," Advanced Photon Source LS-287, September 2000.
[8] L. Emery, Proceedings of the PAC'99, New York, p. 401 (1999).