RCS POLARIZATION CORRECTION USING INFORMATION THEORETIC AND MACHINE LEARNING TOOLS

V. Ranjbar*, W. Bergan, V. Schoefer,
M. Signorelli, E. Gianfelice-Wendt, Brookhaven National Laboratory, Upton, United States

Abstract

One of the main challenges for the EIC Rapid Cycling Synchrotron is to maintain high polarization transmission during acceleration to 18 GeV. While SVD-based orbit smoothing is sufficient below 10 GeV, the stronger residual imperfection spin resonances at higher energy make performance increasingly sensitive to vertical quadrupole misalignments and dipole rolls. A promising approach is therefore to infer the underlying magnet errors from BPM-based orbit-response measurements and use those fitted errors in the resonance-correction workflow. Here we present a direct lattice-fitting method based on simultaneous K_1 dithering and backpropagation through the full accelerator model. In idealized synthetic studies without BPM noise, the method recovers quadrupole offsets and dipole rolls at essentially the numerical precision limit of the model. These results establish a best-case baseline for future studies that will incorporate BPM noise, finite resolution, and Fisher-information-guided optics perturbations to improve conditioning in more realistic scenarios.

INTRODUCTION

High polarization transmission in the EIC Rapid Cycling Synchrotron (RCS) requires control of the imperfection spin resonances driven by vertical closed-orbit errors. Previous tracking studies showed that standard SVD orbit correction is sufficient to maintain more than 99% transmission up to 10 GeV, but that acceleration to 18 GeV becomes sensitive to vertical quadrupole offsets at the 100 µm scale and to dipole rolls at the 100 µrad scale. In that regime, estimating the underlying magnet errors and correcting the corresponding resonance harmonics is more effective than orbit smoothing alone.

The natural starting point is beam-based alignment (BBA), in which small quadrupole-strength variations are used to infer magnetic center offsets from BPM data [1, 2]. For the RCS, the measured vertical orbit can be written as

$$\mathbf{y}_{\mathrm{orb}} = \mathbf{M}_{rQm}\,\mathbf{q}_{\mathrm{mis}}, \tag{1}$$

where $\mathbf{M}_{rQm}$ is the orbit response matrix and $\mathbf{q}_{\mathrm{mis}}$ is the vector of vertical quadrupole misalignments. Once these offsets are estimated, the corresponding imperfection resonance strengths can be computed with DEPOL [3] and reduced with vertical correctors through a second response matrix,

$$\mathbf{w}_{\mathrm{imp}} = \mathbf{M}_{rCW}\,\mathbf{c}. \tag{2}$$

* vranjbar@bnl.gov

Figure 1 shows that this approach already gives a useful reconstruction of the quadrupole offsets and supports high-transmission correction in synthetic RCS error ensembles.

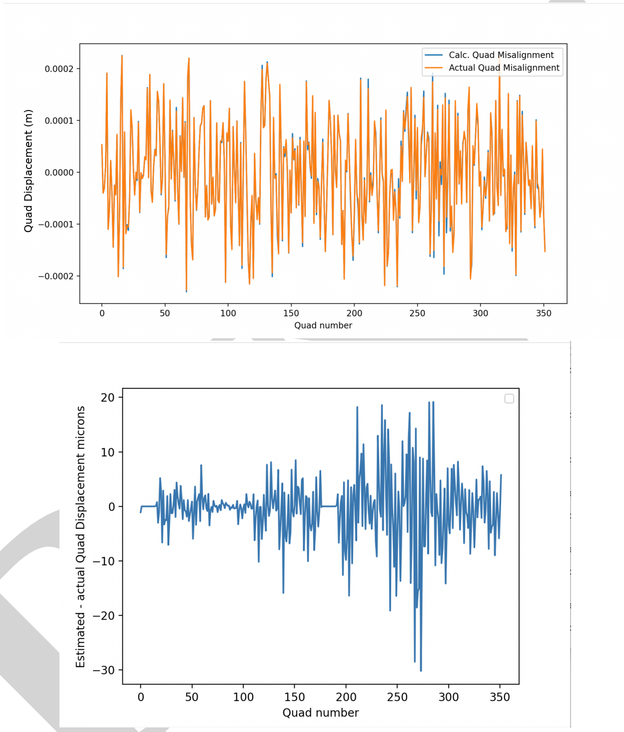

Figure 1: Quadrupole-misalignment reconstruction from a vertical orbit response fit in a misaligned and rolled RCS lattice. Top: fitted and true offsets for a 100 µm RMS quadrupole misalignment ensemble. Bottom: residual difference between the true and reconstructed misalignments.

However, the inversion is not unique: dipole rolls create orbit signatures that can closely resemble those produced by quadrupole vertical offsets. This is closely related issues in BBA [1,2] which is widely used to find the bpm measurement offset with respect to the quadrupole's magnetic center. It is also similar to problems facing LOCO-style optics fits [2], where an orbit response matrix is used to help fit optics and errors.

The problem can be understood as a tomographic reconstruction of the error sources. Each optics configuration produces a different set of response-matrix columns, or "views," of the same underlying misalignments. The goal is to choose optics perturbations that make those views as independent as possible. For a stack of response matrices $\hat{\mathbf{J}}$, this can be quantified by the weighted Fisher matrix

$$\mathbf{F} = \hat{\mathbf{J}}^{T-1}\hat{\mathbf{J}}, \tag{3}$$

whose smallest eigenvalues govern the resolvability of nearly degenerate error patterns. Here $\hat{\mathbf{J}}$ denotes the stacked response matrix assembled from one or more optics configurations, and is the covariance matrix of the orbit-response measurements. In the idealized studies presented here, may be taken as proportional to the identity.

In the present RCS concept, each quadrupole is assumed to have individual trim-power-supply control, which provides additional response information correlated with quadrupole misalignments. This should help break the degeneracy with dipole rolls once realistic measurement effects are included. We also take the next step of fitting the full accelerator model directly, using SciBmad rather than a reduced linear or higher-order map. In that sense the method closely resembles neural-network training: instead of updating weights through backpropagation, the fit updates misalignments and roll errors. With an appropriate gauge choice, BPM offsets do not need to be determined explicitly in the present idealized study, so the ultimate fit accuracy is set by BPM noise and resolution rather than by static BPM offsets.

BACKPROPAGATION THROUGH ACCELERATOR MODEL

We adopt the simultaneous K_1-dithering framework developed for the RCS BBA studies and cast it as a differentiable inverse problem. A nominal closed orbit $\mathbf{x}_0$ is measured, each quadrupole is perturbed by a small $\delta K_{1,i}$, and the perturbed BPM orbits are assembled into a measurement vector

$$\mathbf{m} = \left[\mathbf{x}_0; \mathbf{R}_1; \mathbf{R}_2; \ldots; \mathbf{R}_{N_q}\right], \qquad \mathbf{R}_i = \frac{\mathbf{x}_i - \mathbf{x}_0}{\delta K_{1,i}}. \tag{4}$$

The fit parameters include vertical quadrupole offsets and dipole rolls, and can optionally include BPM offsets when explicit gauge tracking is required.

For a parameter vector $\mathbf{p}$, the accelerator model produces predicted measurements $\mathbf{m}_{\mathrm{pred}}(\mathbf{p})$ by closed-orbit finding and response evaluation. The residual

$$\mathbf{r} = \mathbf{m} - \mathbf{m}_{\mathrm{pred}}(\mathbf{p}) \tag{5}$$

is minimized with a damped least-squares update,

$$\left(\mathbf{J}^T\mathbf{J} + \lambda\mathbf{I}\right)\Delta\mathbf{p} = \mathbf{J}^T\mathbf{r}, \tag{6}$$

where $\mathbf{J} = \partial\mathbf{m}_{\mathrm{pred}}/\partial\mathbf{p}$ is the local Jacobian of the predicted measurements with respect to the fit parameters, and λ is the usual Levenberg-Marquardt damping parameter [4, 5]. This $\mathbf{J}$ is distinct from the stacked response matrix $\hat{\mathbf{J}}$ used in Eq. (3). In the original implementation these Jacobian columns were computed by parallel finite differences across the full RCS lattice model. Here the same structure is exploited for backpropagation: sensitivities of the orbit, the fitted parameters, and finally the Fisher metric of Eq. (3) are propagated through the accelerator model to candidate optics trims.

This reframes optics design as an optimization problem. Let $\mathbf{k}$ denote a set of trim-quadrupole changes used to generate alternative optics. For each optics variant n, a new response matrix $\mathbf{J}^{(n)}(\mathbf{k})$ is built, normalized, and stacked with the baseline matrix. The optimization objective is to enlarge weak directions of $\mathbf{F}$, or equivalently to increase the minimum singular value of the stacked response matrix, while constraining the tunes away from dangerous resonances and keeping magnet trims within practical limits. In this way, machine-learning tooling is used not as a black box estimator, but as a gradient-based design layer on top of the accelerator physics model.

RESULTS

The emphasis here is the direct lattice fit itself rather than the downstream spin-transport correction. For full-ring synthetic studies with 352 quadrupoles and 320 dipoles, implemented in the SciBmad accelerator-physics framework [6, 7], the optimizer converged after 35 iterations and recovered a self-consistent set of magnet errors that reproduces the measured orbit and K_1-dither responses. The resulting fit statistics are summarized in Table 1.

Table 1: BBA Fit Statistics for 352 Quadrupoles and 320 Dipoles

Parameter	RMS Error	Max Error	R^2
Quad X (μm)	0.000000	0.000001	1.000000
Quad Y (μm)	0.000001	0.000003	1.000000
Dipole Roll (μrad)	0.000021	0.000106	1.000000

In this idealized study the recovered quadrupole offsets and dipole rolls agree with the applied values at essentially the numerical precision limit of the model. The RMS reconstruction errors are below the nanometer level for quadrupole offsets and at the 2.5×10^{-5} μm level for dipole rolls, with $R^2 = 1$ for all reported parameter classes. This demonstrates that, when the forward model is exact and measurement noise is absent, the direct lattice fit contains sufficient information to recover the underlying machine error state.

Table 2 gives a representative sample of the fitted and true quadrupole offsets and shows the same near-exact agreement on an element-by-element basis.

Table 2: Representative Comparison for the First 10 Quadrupoles

Index	True X (μm)	Fit X (μm)	True Y (μm)	Fit Y (μm)	Error (μm)
1	-113.72	-113.72	69.44	69.44	0.000000
2	-258.71	-258.71	114.80	114.80	0.000000
3	50.31	50.31	192.30	192.30	0.000000
4	-148.40	-148.40	-60.20	-60.20	0.000000
5	-29.11	-29.11	-128.43	-128.43	0.000000
6	297.47	297.47	204.60	204.60	0.000000
7	121.97	121.97	-92.91	-92.91	0.000000
8	48.96	48.96	-19.57	-19.57	0.000000
9	-78.44	-78.44	297.24	297.24	0.000002
10	32.30	32.30	-193.50	-193.50	0.000000

The applied error ensemble used for this study is summarized in Table 3. The synthetic data correspond to nominal 200 μm quadrupole offsets and 200 μrad dipole rolls, with the realized RMS values set by the random seed.

These results should therefore be interpreted as a best-case baseline for the method. They establish that the remaining

Table 3: Applied Misalignment Statistics

Parameter	Applied σ	Realized RMS	Max |Applied|
Quad X (μm)	200	180.0	399.3
Quad Y (μm)	200	168.9	399.1
Dipole Roll (μrad)	200	175.0	399.0

limitations are not in the optimizer or the direct-lattice formulation itself, but in the realism of the measurement model and in the conditioning of the inverse problem once experimental effects are included.

The present study also uses an idealized measurement model, so the fit accuracy is not yet limited by BPM resolution. A natural next step is to include realistic BPM noise and finite resolution in the synthetic data and then quantify the resulting Fisher-information floor. That will set the practical lower bound on recoverable quadrupole offsets and dipole rolls, and determine how many optics variants are needed before additional measurements cease to improve the fit.

CONCLUSION

For the RCS, polarization preservation above 10 GeV is limited less by the availability of correctors than by the quality of the inferred magnet-error model used to target imperfection resonances. Simple orbit-response fitting already enables substantial improvement, but the similarity between quadrupole-offset and dipole-roll signatures is expected to limit performance once realistic measurement noise and resolution are included.

The differentiable framework described here provides a practical path forward. It combines simultaneous BBA, response-matrix-based spin correction, and information-theoretic optics design in a single workflow. The present synthetic-data results show that, in the absence of BPM noise and with an exact forward model, direct lattice fitting can recover the underlying quadrupole offsets and dipole rolls to essentially machine precision. The next step is therefore to introduce realistic BPM noise and resolution limits, and to optimize a small set of lattice perturbations that maximize Fisher information once those experimental constraints are present.

REFERENCES

[1] M. G. Minty and F. Zimmermann, *Measurement and Control of Charged Particle Beams*, Springer, 2003. doi:10.1007/978-3-662-08581-3

[2] J. Safranek, "Experimental determination of storage ring optics using orbit response measurements,"*Nucl. Instrum. Methods Phys. Res., Sect. A*, vol. 388, no. 1–2, pp. 27–36, Mar. 1997. doi:10.1016/s0168-9002(97)00309-4

[3] E. D. Courant and R. D. Ruth, "The acceleration of polarized protons in circular accelerators", BNL, Upton, NY, United States, Rep. BNL 51270, Sep. 1980. https://www.osti.gov/servlets/purl/7034691

[4] K. Levenberg, "A method for the solution of certain non-linear problems in least squares," *Quart. Appl. Math.*, vol. 2, pp. 164–168, 1944; doi:10.1090/qam/10666

[5] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters," *SIAM J. Soc. Indust. Appl. Math.*, vol. 11, pp. 431–441, Jun. 1963. doi:10.1137/0111030

[6] The Bmad Consortium, "SciBmad.jl: Modular, polymorphic, differentiable, and portable (CPU/GPU compatible) accelerator physics code," https://github.com/bmad-sim/SciBmad.jl,

[7] M. G. Signorelli, J. Devlin, G. H. Hoffstaetter, D. Sagan, "SciBmad: A differentiable and GPU-parallelized software library for particle accelerator design, nonlinear analysis, and machine learning", presented at IPAC'26, Deauville, France, May 2026, paper THP5325, this conference.