AUTOMATIC ERROR CORRECTION IN OPTICS MODELS USING A NEW METHOD FOR LOCALIZED COMPONENT CALIBRATION

Carl Stern, SandiaView Software, Albuquerque, New Mexico, U.S.A.
Martin Lee, Patrick Lui, SLAC, Stanford, California, U.S.A.

Abstract

This talk describes research on the development of a new prototype system for automatic component calibration and error correction of beamline optics models. The system analyzes orbit response data to decompose a beamline model into "good" and "bad" optics regions. Each "bad" region is then analyzed independently to correct quadrupole calibration errors within that region. We describe a new error correction procedure for regional component calibration that is used to drive this analysis. We also describe important challenges in automating component calibration and the methods we have developed for addressing these challenges. Finally, initial test results on simulated and real data from the SPEAR ring at SLAC are presented.

1 INTRODUCTION

Over the past two decades, a great deal of effort was invested by laboratories in the development of model-based programs for applications in accelerator operation and beam-line commissioning. At SLAC today, model-based programs are used routinely to correct orbit and optics errors in every accelerator beam line. The effectiveness of using these programs depends on the accuracy of the optics model. For example, a commissioning task is to find an accurate optics model for a given beam line in an accelerator or storage ring. We call this process model calibration.

In general, a model calibration procedure involves three tasks: data acquisition, data analysis, and data interpretation. The first two tasks have been partly or wholly automated but not the third. During the past year, we have been working on the development of a fully automated model calibration algorithm. A prototype algorithm has been developed and our experience in this work will be presented.

2 TECHNICAL APPROACH

2.1 Global and Regional Methods

Current methods commonly apply a global orbit fitting procedure, treating the strength of every quadrupole magnet in the ring as a variable. Quadrupole errors are solved by minimizing the discrepancy between the simulated orbit data and the measured data. The measured orbit data is a collection of measured orbit shifts observed at beam position monitors (BPMs) in the ring. In practice, an orbit shift is produced by kicking the beam with one orbit corrector. The simulated orbit shift as a function of the quadrupole strength variable is calculated by using a lattice code. A least-square fitting on the orbit discrepancy is then used to find the quadrupole error.

Since the number of quadrupole magnets and BPMs increases with the size of the ring, the time it takes to solve for quadrupole strength errors also increases with storage ring size in the global orbit data fitting approach. This problem is often exacerbated by the presence of noise in the measurement.

In order to avoid this difficulty with the global approach, we have developed a regional approach. Our current approach builds on a method initially developed at SLAC [1]. In this regional approach, the main objective is to first find all the good regions in the ring. Good regions are those where behavior predicted by the model matches observed behavior of the machine. By definition, bad regions only exist between two neighboring good regions. In our method, the error finding procedure needs only to be applied to one bad region at a time, making model calibration of a beam line less dependent on the ring size and much more tractable.

2.2 The Search for Good Regions

By definition, measured orbit data matches simulated orbit data at every BPM within a good region. In orbit simulation, we consider an orbit as a trajectory within the good region. To describe the trajectory, we use a set of launch parameters (x, x') where x denotes the trajectory at the beginning of the region and x' denotes the slope of the trajectory. To find launch parameters for a good region, we calculate (x, x') by minimizing the discrepancy between the measured and simulated BPM data.

Since there is noise in the measured data, this discrepancy cannot be reduced to zero. In general, the minimum discrepancy is proportional to the noise level. The fitting procedure for finding good regions is thus affected by the noise level. A fitting tolerance is generally used to take the noise effect into account. Therefore, in the search for a good region, we first specify the fitting tolerance to be used and then look for a region that is characterized by having a minimum trajectory discrepancy that is less than the value of this fitting tolerance. After all the good regions have been identified, the bad regions are automatically identified as those which lie between a pair of good regions.

2.3 *Transport Matrix Fitting for Bad Regions*

The objective of model calibration is to find the strength error in the quadrupole magnets within a given bad region. In practice it has been found that the existing orbit-fitting algorithm is very sensitive to the noise in the measured orbit data, particularly for a bad region with a small number of BPMs.

The transport matrix-fitting algorithm that is reported here is less sensitive to BPM noise. The difference between the existing orbit-fitting algorithm and the new algorithm is that it minimizes the transport matrix discrepancy instead of the trajectory discrepancy. We define the transport matrix discrepancy as the mismatch between the transport matrix represented in the model and the measured transport matrix. In practice, the transport matrix can be calculated from observed orbit trajectory.

We define the transport matrix as the (2x2) matrix which transforms the trajectory coordinates (x, x′) across a bad region. To find the measured transport matrix, we first solve for the measured values for the coordinates at the beginning and end of a given bad region. The values of (x, x') at the beginning of a bad region are computed from the measured orbit trajectory in the good region upstream of the bad region. Similarly, the values of (x, x') at the end of a bad region are computed from the measured data in the good region downstream of the bad region.

With these values known, we compute the transport matrix which transforms the measured (x, x') from the end of the upstream good region to the beginning of the downstream good region (Figure 1). Finally, we try to match this *measured transport matrix* in the bad region by adjusting quadrupole field strengths of selected quadrupoles in the model.

We define an *error hypothesis* as a subset of quadrupoles in the bad region whose calibration error is conjectured to explain the discrepancy between the measured transport matrix and the transport matrix represented in the model. We evaluate *error hypotheses* through a simulated fitting procedure that varies the field strengths of the quadrupoles in the error hypothesis to attempt to match the measured transport matrix. We currently use a lattice code, COMFORT, to implement this fitting procedure, although any lattice code with good fitting services would serve the same purpose.

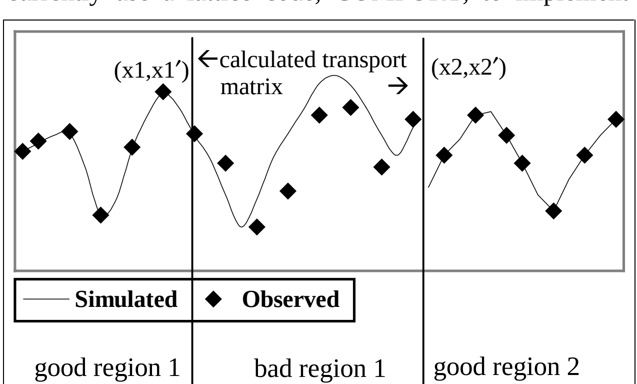

Figure 1. Calculated Transport Matrix

This fitting procedure allows us to search over *error hypotheses*, evaluating the adequacy of each hypothesis in terms of the quality of fitting to the measured transport matrix, or equivalently the degree of error reduction, that it supports. This is a comparative evaluation procedure, where a superior quality of fitting to that of competing error hypotheses, particularly hypotheses involving a similar number of quadrupoles, provides a method for hypothesis selection. We have applied heuristics to the process of searching over error hypotheses that will be discussed in the next section.

We have also experimented with the application of further constraints to this fitting procedure, with good results. We compute intermediate transport matrices, derived from measured orbit trajectories at intermediate BPMs within the bad region, and add elements from these matrices to the fitting requirements. This implicitly coerces the fitting not only to beginning and endpoints of the bad region but also to intermediate locations.

3 AUTOMATING DATA ANALYSIS AND INTERPRETATION

Automating model calibration involves a number of challenges. While fitting and error minimization algorithms provide the mathematical backbone of model calibration, sophisticated data analysis, data interpretation, and search are also required. In human model calibration, these latter are heavily dependent on the knowledge and skill of a modeling expert, leading to a strong correlation between the accuracy of calibration and the level of expertise of the human who performs it.

We describe very briefly a few of the heuristic methods we have adopted in the effort to automate expert human performance in data analysis and interpretation.

3.1 *Data Analysis and Filtering*

Our system employs a well-known multi-track method of analysis [2]. A *track* is a set of orbit shifts produced by the kick from a single corrector. Using correctors from around the ring generates a set of tracks, each probing quadrupole fields around the ring at different locations.

Despite the differences in orbit trajectory, every track is determined by the same set of simulated and measured transport matrices. This provides useful data redundancy and a robust environment for data analysis and filtering. Averaging over tracks reduces the effects of noise in any one track. Inconsistency between tracks allows elimination of bad data when, for example, one track is "outvoted" by the remaining tracks. Data anomalies across tracks, for example, consistent near-zero readings at one BPM, allows data pruning at that BPM.

We have observed, in the course of our research, that data analysis is a dynamic and interactive process. Data interpretation feeds back into data analysis. For example, a method for finding *good regions* implicitly validates blocks of good data and provides a new consistency-based check for identifying anomalous data. Because it is desirable to add new methods for data analysis as they are discovered, we have decided to implement an extensible rule system for encoding data analysis expertise.

3.2 Hypothesis search and selection

The problem of determining which quadrupole magnets in a bad region have calibration errors, i.e., the problem of selecting the best error hypothesis, is non-trivial. There are two related challenges to this problem. The first is to construct justifiable criteria for deciding between error hypotheses. The second challenge is to develop a focused search procedure that searches over a potentially large space of error hypotheses and finds optimal or near optimal hypotheses within a reasonable time.

The transport matrix fitting procedure described in the previous section provides one decision criterion. It eliminates large numbers of hypotheses by returning a very poor convergence metric for matrix fittings based on those hypotheses. However we found that this criterion by itself does not always select a unique hypothesis. We discovered in testing that several different sets of error hypotheses often exhibited comparable convergence behavior during transport matrix fitting, with no one error hypothesis standing out uniquely from the rest.

We are currently confronting both aspects of this problem by taking a probabilistic approach. We define the optimal error hypothesis as that with the highest estimated probability. Taken in isolation, the probability of a single quadrupole error is correlated with the size of the calibration error: ($p(e) \approx 1-(\alpha e)^2$). The prior probability of an error hypothesis $\cup e_i$ would be the product of the single error probabilities: ($p(\cup e_i) \approx \Pi_i(p(e_i))$). The posterior probability would be the prior probability scaled by the convergence metric γ: ($p(\cup e_i) \approx (-\ln(\gamma)) * \Pi_i(p(e_i))$).

However, this method of probability estimation is inaccurate because error hypotheses for the same bad region are *competing* explanations for the same discrepancy. It is easily seen, by consideration of Bayes' rule, that the probability of each causal hypothesis is thereby also conditioned on that of the competing hypotheses.

We have thus been motivated to explore the use of Bayesian belief networks to better implement probability estimation. In these networks the flow of causal influence and the interactions between competing causal explanations are made explicit, providing a principled method for calculating posterior probabilities. We will explore a recently developed entropy-based probing strategy in Bayesian models that supports efficient search for maximum likelihood explanations. [3].

4 TEST RESULTS

A prototype system for correcting quadrupole field strength errors has been implemented and tested. Initial testing was performed on simulated data from the SPEAR ring at SLAC. On noise-free data, our system correctly identified up to three simulated quadrupole errors using the convergence metric from transport matrix fitting as the selection criterion. The convergence metric for the correct error hypothesis was often four orders of magnitude better than that of the nearest competing hypothesis.

Next we tested our system on real data from SPEAR. An estimated 5 – 10% noise level in the data reduced the convergence separation between the best hypotheses. The use of probabilistic methods, specifically the assignment of higher probabilities to hypotheses with smaller numbers of quadrupole errors, allowed us to recalibrate a SPEAR model with two bad regions and a total of three quadrupole calibration errors. The recalibrated model exhibits good agreement between predicted and observed orbit displacements. Other performance characteristics of the recalibrated model will soon be tested.

5 SUMMARY AND FUTURE WORK

A first generation system automating quadrupole error correction in beamline models has been developed. Three aspects of this system have been discussed: 1) the invention of a new method for regional error correction using transport matrix fitting, 2) the automation of heuristic methods for data analysis, and 3) the application of probabilistic methods for finding optimal error hypotheses.

Future plans include the implementation of a second generation error correction system with an improved interface, a wider range of heuristic methods for data analysis, and a fully embedded probabilistic inference system for error analysis.

ACKNOWLEDGEMENTS

Some of the work reported here was performed by Vista Control Systems in collaboration with SLAC. It was supported by U.S. Department of Energy SBIR grant DE-FG03-98ER82708 and contract DE-AC03-76SF00515.

REFERENCES

[1] M.J. Lee. GLAD: A Generic Lattice Debugger. SLAC-PUB-5700, Nov 1991. Invited talk given at International Conference on Accelerator and Large Experimental Physics Control Systems, Tsukuba, Japan, Nov. 11-15, 1991.

[2] W.J. Corbett, M. J. Lee, V. Ziemann. A Fast Model Calibration Procedure for Storage Rings. Presented at 1993 Particle Accelerator Conference (PAC 93), Washington, DC, 17-20 May 1993. Published in IEEE PAC 1993:108-110(QCD183:P3;1993).

[3] G. Luger, D. Pless, C. Stern. An entropy-based probing strategy for finding most probable diagnoses in model-based BBNs. Technical Report (in preparation). Department of Computer Science, University of New Mexico, Albuquerque, NM. 1999.