SLAC-PUB-13686
June 2009

LOCO with constraints and improved fitting technique

Xiaobiao Huang, James Safranek
Mail to: xiahuang@slac.stanford.edu
Stanford Linear Accelerator Center, Menlo Park, CA 94025, USA
Greg Portmann
Lawrence Berkeley National Lab, Berkeley, CA 94720, USA

1.1.1 Introduction

LOCO has been a powerful beam-based diagnostics and optics control method for storage rings and synchrotrons worldwide ever since it was established at NSLS by J. Safranek [1]. This method measures the orbit response matrix and optionally the dispersion function of the machine. The data are then fitted to a lattice model by adjusting parameters such as quadrupole and skew quadrupole strengths in the model, BPM gains and rolls, corrector gains and rolls of the measurement system. Any abnormality of the machine that affects the machine optics can then be identified. The resulting lattice model is equivalent to the real machine lattice as seen by the BPMs. Since there are usually two or more BPMs per betatron period in modern circular accelerators, the model is often a very accurate representation of the real machine. According to the fitting result, one can correct the machine lattice to the design lattice by changing the quadrupole and skew quadrupole strengths. LOCO is so important that it is routinely performed at many electron storage rings to guarantee machine performance, especially after the Matlab-based LOCO code [2] became available.

However, for some machines, LOCO is not easy to carry out. In some cases, LOCO fitting converges to an unrealistic solution with large changes to the quadrupole strengths ΔK. The quadrupole gradient changes can be so large that the resulting lattice model fails to find a closed orbit and subsequent iterations become impossible. In cases when LOCO converges, the solution can have ΔK that is larger than realistic and often along with a spurious zigzag pattern between adjacent quadrupoles. This degeneracy behavior of LOCO is due to the correlation between the fitting parameters – usually between neighboring quadrupoles. The fitting scheme is therefore less restrictive over certain patterns of changes to these quadrupoles with which the correlated quadrupoles fight each other and the net effect is very inefficient χ^2 reduction, i.e., small χ^2 reduction with large changes of ΔK. Under effects of random noise, the fitting solution tends to crawl toward these patterns and ends up with unrealistically large ΔK. Such a solution is not very useful in optics correction because after the solution is dialed in, the quadrupoles will not respond as predicted by the lattice model due to magnet hysteresis. We will show that adding constraints to the fitting parameters is an effective way to combat this problem of LOCO [3-4]. In fact, it improves optics calibration precision even for machines that don't show severe degeneracy behavior.

LOCO fitting is essentially to solve a nonlinear least square problem with an iterative approach. The linear least square technique is applied in each iteration to move the solution toward the minimum. This approach is commonly referred to as the Gauss-Newton method. By using singular value decomposition (SVD) to invert the Jacobian matrix, this method has generally been very successful for LOCO. However, this method is based on a linear expansion of the residual vector over the fitting parameters

Published in the ICFA Beam Dyn. Newslett. 44:49-60, 2007

Work supported in part by US Department of Energy contract DE-AC02-76SF00515

which is valid only when the starting solution is sufficiently close to the real minimum. The fitting algorithm can have difficulties to converge when the initial guess is too far off. For example, it's possible for the χ^2 merit function to increase after an iteration instead of decrease. This situation can be improved by using more robust nonlinear least square fitting algorithms, such as the Levenberg-Marquardt method [7].

We will discuss the degeneracy problem in section 2 and then show how the constrained fitting can help in section 3. The application of Levenberg-Marquadt method to LOCO is shown in section 4. A summary is given in section 5.

1.1.2 The degeneracy problem

A general nonlinear least-square problem is to minimize the merit function

$$f(\mathbf{p}) = \chi^2 = \sum \left[y_i - y(x_i; \mathbf{p}) \right]^2 , \tag{1}$$

where $\mathbf{p}$ is a vector of the fitting parameters, (x_i, y_i) are measured data and $y(x; \mathbf{p})$ is a nonlinear model function. The residual vector is a column vector $\mathbf{r}$ whose components are $r_i = y_i - y(x_i; \mathbf{p})$, with $i = 1,2,\cdots,N$ and N is the number of data points. The Jacobian matrix $\mathbf{J}$ is defined as,

$$J_{ij} = \frac{\partial r_i}{\partial p_j} . \tag{2}$$

Each column of the Jacobian matrix is the derivative of the residual vector over one fitting parameter. In the Gauss-Newton method, the solution is advanced toward the minimum at each iteration by $\Delta\mathbf{p}$, which is determined by

$$\mathbf{J}^T \mathbf{J} \Delta\mathbf{p} = -\mathbf{J}^T \mathbf{r}_0 , \tag{3}$$

where $\mathbf{r}_0$ is the residual vector of the previous iteration. This is essentially the method adopted by the original LOCO, although where the equation was $\mathbf{J}\Delta\mathbf{p} = -\mathbf{r}_0$. However, it is much faster to do SVD on matrix $\mathbf{J}^T\mathbf{J}$ than on $\mathbf{J}$ since the latter has tens of times more rows [3].

In a fitting problem, two parameters can be deeply coupled such that their contributions to the merit function are very difficult to separate. In an extreme case, for example, it is impossible to determine the two fitting parameters p_1 and p_2 in the problem defined by $\chi^2 = \sum \left[y_i - y(x_i, p_1 - p_2) \right]^2$ because the merit function has no dependence on $p_1 + p_2$. The corresponding columns of the Jacobian matrix for the two parameters differ by only a scaling constant. Therefore, the Jacobian matrix is rank deficient. In a less severe case, the merit function may have weak dependence over $p_1 + p_2$ so that in principle it can be determined. But it is susceptible to noise in the experimental data and tends to have large error bar. Consequently the two parameters p_1 and p_2 both have

larger error bars. The coupling of adjacent quadrupole gradient parameters in LOCO is very similar to the above case. If two quadrupoles are placed next to each other without

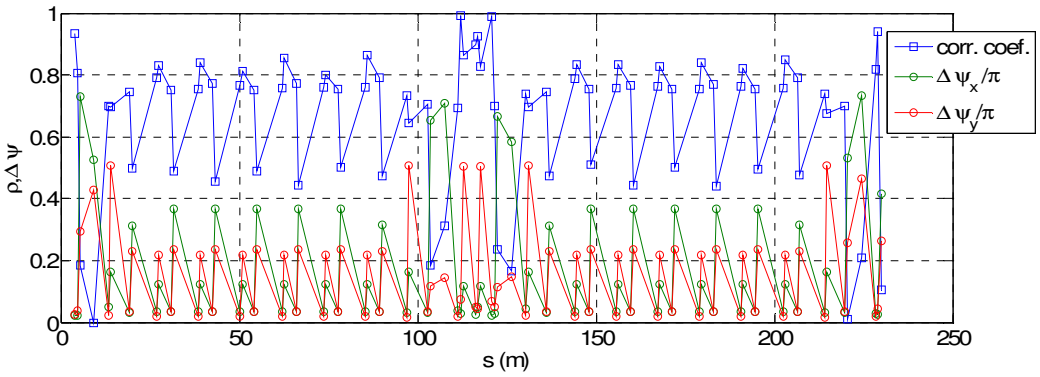

Figure 1: the correlation coefficients and betatron phase advances between neighboring quadrupoles (excluding QFC) in SPEAR3.

much space between, they perturb the linear optics of the machine in essentially the same manner. LOCO may be able to accurately fit the combined integrated gradient of the two magnets but it would not be able to distinguish the individual contributions. When quadrupoles are separated with drift spaces or other components, the coupling between their gradients gets weaker but remains existent. Detailed analysis shows that it is the betatron phase advances between two quadrupoles that determine their coupling strength. Adding BPMs in the vicinity can alleviate the problem to some extent, but cannot eliminate it.

The coupling between fitting parameters is reflected in the similarity of their corresponding columns in the Jacobian matrix, which may be characterized by their correlation coefficient,

$$\rho_{12} = \frac{\mathbf{J}_1^T \mathbf{J}_2}{\|\mathbf{J}_1\| \|\mathbf{J}_2\|}, \tag{4}$$

where $\mathbf{J}_{1,2}$ are corresponding columns of the parameters p_1 and p_2, and $\|\bullet\|$ stands for the 2-norm of its argument. Fig. 1 shows the correlation coefficients between neighboring quadrupoles in SPEAR3. The horizontal and vertical betatron phase advances ($\mathrm{mod}\, 2\pi$) between these quadrupoles are also plotted. The correlation coefficients between neighboring QF, QD magnets in a DBA cell are around 0.8. Quadrupole magnets in the double waist straight section (in the center of the plot), [6], have even stronger correlation between their neighbors. Clearly, strong correlation is the result of small betatron phase advances between them. Stronger correlation can occur between two quadrupoles that are physically set apart but both betatron phase advances between them are close to a multiple of π. One may calculate a correlation matrix and examine the coupling relations between the fitted quadrupole gradients.

Because of the coupling between quadrupoles, some patterns of changes of the quadrupole gradients are much less restricted by LOCO fitting. If these patterns form a null-space in the parameter space, meaning that they correspond to singular values considerably smaller than others, then they can be simply removed by proper selection

of singular values. However, the less restrictive patterns have various severity

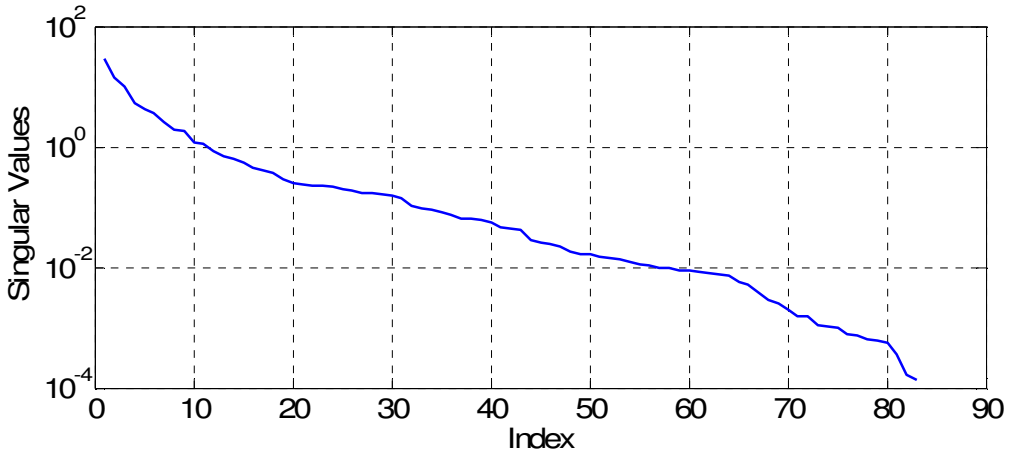

Figure 2: singular values of the correlation matrix of SPEAR3 quadrupoles.

and are rarely orthogonal. The patterns through SVD have to be orthogonal to each other. Consequently the singular value spectrum is usually a smooth curve without a clear cut. For example, the singular values for the correlation matrix in the SPEAR3 example are shown in Fig. 2. The less restrictive patterns tend to be clustered toward the lower end of the spectrum. When degeneracy becomes a problem, removing some of the very low singular value modes is justifiable. However, completely removing any mode is a loss of information and thus would reduce the fitting accuracy to some level. It often takes much work to find an optimal threshold. And our experience shows that some less restrictive patterns still leak into the solution even if a seemingly optimal threshold is applied. It seems not possible to retain good accuracy of fitting and keep the quadrupole strength changes reasonably low at the same time by merely selecting the singular values. This has also been the observation at Soleil, [10].

The response matrix measurement always has errors because of random BPM noises and random machine fluctuation. There are also systematic errors due to, for example, nonlinearity of the machine. The solution may drift along a less restrictive direction by a large step to gain a small reduction of χ^2 that is under the error level. Since the fitting algorithm does its best to find a minimum, the final solution tends to acquire large excursions toward the less restrictive directions. Such excursions are not necessarily reflected in the standard error bar calculation of LOCO because there only random BPM noises are considered.

The degeneracy caused by coupling between the fitting parameters is intrinsic to the problem. Strictly speaking, there is no single “best solution”. Instead, any solution whose χ^2 differs from the global minimum by less than a certain amount determined by the noise level is a valid, equivalent solution. These solutions should in principle give the same machine lattice. We should pick the reasonable ones from this set of equivalent solutions. In the next section, we will show how this can be done by the constrained fitting method.

1.1.3 LOCO fitting with constraints

Since coupling between LOCO fit parameters can cause excursions of the solution in unconstrained directions, it is natural to put a penalty on such excursions. This can be done by modifying the merit function with additional penalty terms. Ideally we would like the penalty terms to represent the unconstrained patterns as was done in Ref. [3-4]. However, it is not easy and also not necessary to identify these patterns. A much simpler approach is to put penalty on the change of gradient ΔK of each quadrupole directly. Therefore the least-square problem now reads,

$$\chi^2 = \sum_{i,j} \frac{(M_{\mathrm{mod},ij} - M_{\mathrm{meas},ij})^2}{\sigma_i^2} + \frac{1}{\sigma_{\Delta K}^2} \sum_k w_k^2 \Delta K_k^2 , \qquad (5)$$

where $\sigma_{\Delta K}$ is an overall normalization constant and w_k^2 are individual weighting factors that represent the needs to constrain their corresponding quadrupoles. The weighting factors should be adjusted according to the performance on a trial and error basis. But once a suitable set of weighting factor is found for one lattice, there is usually no need to change it later on. Removing singular values is equivalent to put infinite penalty weight on the corresponding patterns. This extreme measure seems to work less efficiently than the more "gentle" approach here.

It is straightforward to implement this modified LOCO fitting scheme. The input and output are the same data as the original LOCO. Only the minimization algorithm needs a slight modification. The additional terms in χ^2 amount to simultaneous linear equations, $\Delta K_k = 0$, $k = 1,2,\cdots,N_q$, each with a weight $w_k^2 / \sigma_{\Delta K}^2$, where N_q is the number of constrained quadrupoles. Consequently the residual vector and the Jacobian matrix are extended. Suppose the original Jacobian matrix has N rows (i.e., N data points). The Jacobian matrix will have N_q more rows with nonzero elements $J_{N+k,k} = w_k / \sigma_{\Delta K}$ and the corresponding additional elements of the previous residual vector $\mathbf{r}_0$ are zeros. Then Eq. (3) is applied as usual to find the step to the next solution.

The additional constraint terms changes the solution to the linear problem of each iteration. However, since only the gradient changes between successive iterations are constrained with a cost, the global minimum of the original problem remains the same. What has been changed is the convergence path. Fig. 3 serves as an illustration of this picture, where point 0 represents the initial guessed solution; point M is the global minimum within a sea of equivalent solutions under the noise level. The unconstrained path (solid arrows) takes large excursions and reaches the global minimum quickly, while the constrained path (dashed arrows) touches the nearest edge of the noise sea and slows down.

Fig. 4 is an example of real LOCO data at SPEAR3 fitted for 12 iterations with both algorithms. Here 12 iterations are run just to show the behavior of the fitting methods. We usually run two or three iterations and apply the solution for optics correction. The convergence paths are shown on an rms relative gradient change vs. residual χ^2 (normalized by degree of freedom) plot. The unconstrained algorithm converges in

three iterations but ends with a 2.1% rms relative gradient change. The constrained algorithm brings χ^2 down to the same level in three steps with an rms relative gradient change of only 0.6%. The additional gradient changes of the unconstrained algorithm do not result in a very different lattice. In fact, the rms relative differences of horizontal and vertical beta functions between the two lattices at point 3 and point 3' are only 1% and 0.3%, respectively. Most of the additional gradient changes are cancelled by fighting each other. The lattice difference would be far larger if the gradient changes were random.

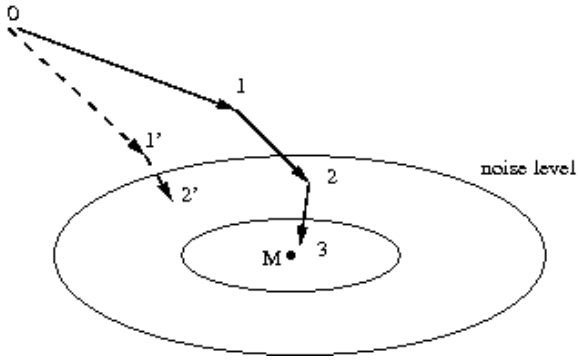

Figure 3: An illustration of the changes to the convergence path with or without constraints. Solid: no constraints; Dashed: with constraints.

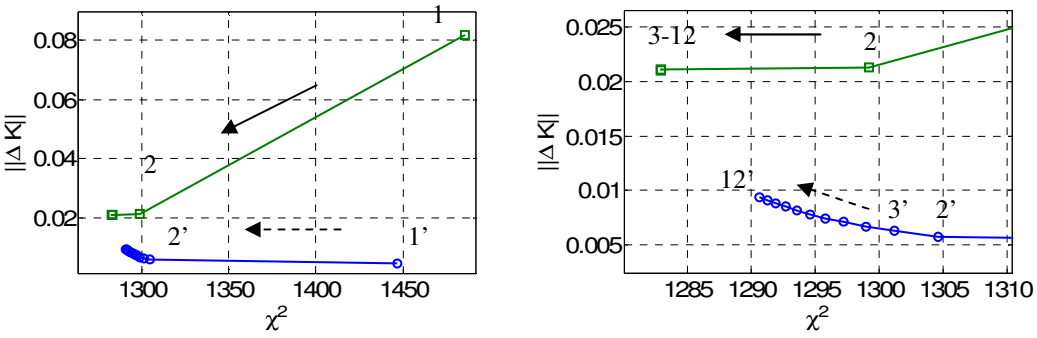

Figure 4: The rms relative change of gradients vs. the residual χ^2 for a SPEAR3 data set. Green: no constraints; Blue: with constraints. Point 0 is located at (2.0×10^6 , 0). The right plot is a blowup show of the left plot.

A LOCO solution with smaller gradient changes is preferred for optics control. When quadrupole setting corrections are dialed in according to the solution, it is assumed that the magnets respond with a linear field-current relation. This assumption holds only when the current changes are sufficiently small because of magnet saturation and hysteresis. Therefore, a solution with large changes of gradients may not produce the expected lattice in reality. For the nominal SPEAR3 lattice the coupling problem is not severe. A beta beating under 1% rms was achieved with proper selection of singular values using the original LOCO method. The constrained fitting, however, still proves to be a valuable tool by improving the precision of optics control. Presently, the machine optics are corrected to a beta beat less than 0.2% beta beating without special

care to singular value selections. For the low alpha lattice, [5], this tool is indispensable because the lattice couldn't be calibrated without it. For one LOCO data set from the low alpha lattice, the rms relative ΔK change is fitted to be 14.5% by keeping 530 out of 542 singular values and larger if more singular values are kept with the unconstrained method. The constrained method finds a solution with only 3.5% rms relative ΔK change and the rms relative beta function difference between the two resulted lattices is only 2.0% horizontal and 0.7% vertical. After correction, the beta beat was reduced to 0.5% horizontal and 0.3% vertical.

It is worth noting that the ability of LOCO to identify large gradient errors is not sacrificed by adding constraints. The constraints endanger primarily the existence of the less restrictive patterns in the solution because they are "cheap" in terms of χ^2. Real gradient errors usually cause large χ^2 contributions and they rarely form a less restrictive pattern. For example, during the 2007 shutdown of SPEAR3, an insertion device was moved to a new location. The perturbation to the lattice corresponds to normalized χ^2 contributions over 8×10^4. The nearby quadrupole magnets stand out in the fitting solution to account for the change, even if we put 20 times more weight on these quadrupoles than we normally do.

The penalty terms, the costs, are essential to the improved performance of LOCO. It is important to set the cost factors properly. If the penalty is too high, the solution would converge too slowly. If the penalty is too low, the benefit of constraints would not be seen. The average cost may be a good indicator of the choice of weighting factors. In Fig. 5 we plot the average cost for each iteration for the example shown in Fig. 4. For SPEAR3, equal weighting factors were initially chosen for all quadrupoles and then adjustments are made as needed. The more sensitive a parameter is to coupling issues, the more weight gets added. The proper overall weighting factor may be found by adjusting it to produce an average cost at the first iteration that is comparable to the normalized residual χ^2. The correlation coefficient plot (Fig. 1) could be useful in identifying the sensitive quadrupoles. However, a more straightforward and easier method is to study the individual χ^2 contribution of the fitting parameters, [3]. The χ^2 contribution of a fitting parameter is defined as the increase to χ^2 if we set this parameter to its initial value at iteration 0 and keep all other parameters unchanged. We can define the contribution of a group of parameters in the same manner. Ideally, if there is no correlation between the fitting parameters so they don't fight each other, then the sum of their individual contributions should be nearly equal to their group contribution, assuming a linear expansion is valid. In contrast, if a parameter drifts along a less restrictive direction by a large amount from iteration 0, then it will cause a sizable χ^2 contribution because the other parameters are not present to cancel its effect in this calculation. This is what was observed in the same SPEAR3 data set. Fig. 6 shows the χ^2 contribution after three iterations with or without constraints. Because quadrupoles that share the same power supply are combined as one fitting parameter, there are 72 quadrupole fit parameters. Also shown in Fig. 6 are contributions of the 14 skew quadrupoles. It is first noted that the group contribution of all fitting parameters for the two cases are about the same ($\Delta\chi^2 \approx 3600$).This is not surprising because the

two resulting model lattices are nearly the same (see discussion before Fig. 3). However, for the unconstrained case, the individual contributions are considerably higher, especially for a few quadrupole magnets. It is reasonable to assign more weights to these parameters. These magnets turn out to be located around the double waist straight section, where the coupling between magnets is stronger. The constrained solution has small individual χ^2 contribution for the fitting parameters and their sum is on the same order as their group contribution.

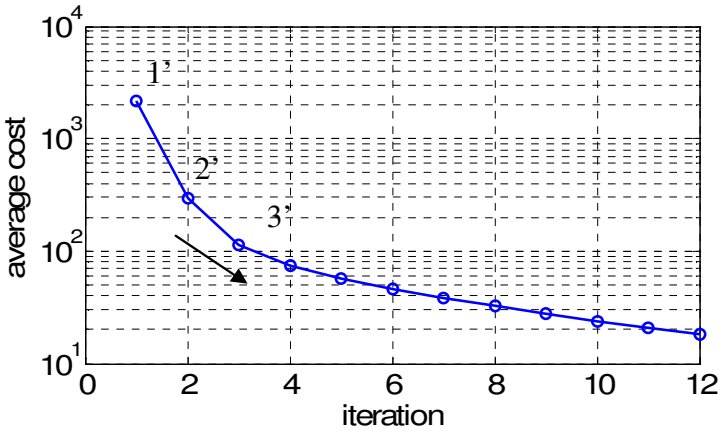

Figure 5: the average cost for the SPEAR3 example shown in Fig. 4

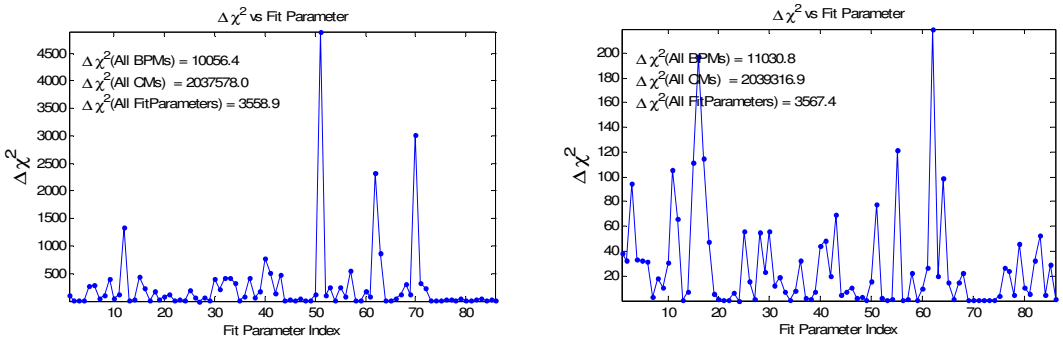

Figure 6: χ^2 contribution for individual fitting parameters at point 3 (left) and point 3' (right) for the SPEAR3 example shown in Fig. 4. The first 72 parameters are quadrupoles. The last 14 are skew quadrupoles.

Constraints can be easily applied to other parameters such as BPM gains and rolls, corrector gains and rolls but this is not necessary for the case of SPEAR3.

1.1.4 The Levenberg-Marquadt fitting algorithm

The Gauss-Newton method that is adopted by the original LOCO converges quickly if the initial solution is close to the minimum. But there are situations when a good initial guess is not available. In those situations, the fitting algorithm may behave unexpectedly, for example, an increased χ^2 or one that oscillates around the minimum. In such cases the steepest descent is more suitable because it can move the solution toward the minimum. The Levenberg-Marquadt algorithm combines the two methods in

an elegant way so that either one is selected in its applicable region [7]. This is achieved by replacing Eq. (3) with

$$(\mathbf{J}^T\mathbf{J} + \lambda\mathbf{D}^T\mathbf{D})\Delta\mathbf{p} = -\mathbf{J}^T\mathbf{r}_0, \tag{6}$$

where $\mathbf{D} = \mathrm{diag}(\|\mathbf{J}_1\|, \|\mathbf{J}_2\|, \cdots, \|\mathbf{J}_{NP}\|)$, NP is the number of columns of the Jacobian matrix and $\lambda > 0$ is a scaling constant. In other words, the diagonal elements of the matrix $\mathbf{J}^T\mathbf{J}$ are scaled up by $1+\lambda$. The constant λ controls the behavior of the algorithm. If λ is very small (much less unity), it is the same as the Gauss-Newton method. If λ is much larger than unity, it becomes the steepest descent method. Usually λ is set to a small value initially, for example $\lambda = 0.001$. Then after every iteration it is adjusted, depending on the result of the solution found in that iteration. A simple way is to scale it down by a factor of 10 if the solution reduces χ^2 and update the Jacobian matrix; and if χ^2 is increased then scale λ up by a factor of 10 until a solution is found to reduce χ^2. In this way χ^2 is guaranteed to decrease after every iteration. The application of the Levenberg-Marquadt method to LOCO has been suggested in Ref. [3-4].

It is interesting that the fitting with constraints scheme can be cast into the same framework as the Levenberg-Marquardt algorithm. In fact, Eq. (6) applies to the former case with $\lambda = 1$ fixed and $\mathbf{D} = \mathrm{diag}(0, \cdots, 0, w_1, w_2, \cdots, w_{N_q})/\sigma_{\Delta K}$, where the 0's are for non-constrained parameters. This is not surprising because both methods want to limit the step sizes toward the next solution, but for different reasons! For the Levenberg-Marquardt method, the goal is to find the global minimum reliably. But the constrained fitting actually wants to avoid the global minimum by going with slower paces.

A recent report [8] indicates that a more delicate form of the Levenberg-Marquardt algorithm [9] could have better performance for LOCO. This algorithm is based on explicit control of a trust-region: it finds the best solution for the iteration within the region specified by

$$\|\mathbf{D}\Delta\mathbf{p}\| \le \Delta, \tag{7}$$

where Δ on the right hand side represents the size of the trust region and it is adjusted according to the efficiency of χ^2 reduction after every iteration.

1.1.5 Summary

Two ways to improve the LOCO technique have been discussed – the constrained fitting method and the Levenberg-Marquadt algorithm. The constrained fitting is introduced to cure the degeneracy problem caused by the coupling between fitting parameters (mainly the neighboring quadrupoles). This is a common problem that occurs to many machines in different severity. The constraints are implemented by putting penalties for the step sizes between the solutions of successive iterations. It has

shown to be an efficient way to remove the less restrictive patterns from the solution and it results in fitted lattices with small changes from the starting point. This enables precise control over the machine optics, even for machines where the degeneracy problem has made the original form of LOCO not useful.

The Levenberg-Marquadt algorithm is a robust solver for general nonlinear least square problems. It is useful for LOCO in cases when the initial guessed solution is not close enough to the minimum. In such cases the Gauss-Newton solver may fail because the solution it finds could lie outside of the region where linearization of the model is valid. The Levenberg-Marquadt algorithm is based on the trust-region strategy – the solution it finds is confined in a region where the linear model is valid. The size of the trust region is controlled implicitly or explicitly and is adjusted after every iteration.

1.1.6 References

1. J. Safranek, "Experimental determination of storage ring optics using orbit response measurements", Nucl. Instr. and Methods, A 388 (1997) 27-36
2. J. Safranek, et al, "Matlab-based LOCO", Proceedings of EPAC02, Paris, France
3. X. Huang, Indiana University PhD thesis (unpublished), 2005
4. X. Huang, et al, "Fitting the fully coupled ORM for the Fermilab Booster", Proceedings of PAC05, Knoxville, TN
5. X. Huang, et al, "Low alpha mode for SPEAR3", Proceedings of PAC07, Albuquerque, NM
6. J. Corbett, "Implementation of the double-waist chicane optics in SPEAR3", Proceedings of EPAC06, Edinburgh, Scotland
7. W. H. Press, et al, "Numerical recipes in C", 2nd Ed. Cambridge University Press.
8. L. Yang, et al, "A new code for orbit response matrix analysis", Proceedings of PAC07, Albuquerque, NM
9. J. J. More, "The Levenberg-Marquadt algorithm: implementation and theory", Lecture Notes in Mathematics (G. Wastson ed.), vol 630, Springer, 1978, p. 105
10. See contribution by L. Nadolski in this ICFA beam dynamics newsletter.