

Beam-Based BPM Alignment

R. Talman

September 2002

Collider Accelerator Department
Brookhaven National Laboratory

U.S. Department of Energy

USDOE Office of Science (SC)

Notice: This technical note has been authored by employees of Brookhaven Science Associates, LLC under Contract No.DE-AC02-98CH10886 with the U.S. Department of Energy. The publisher by accepting the technical note for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this technical note, or allow others to do so, for United States Government purposes.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.



Beam-Based BPM Alignment

BNL/SNS TECHNICAL NOTE

NO. 116

Richard Talman and Nikolay Malitsky

September 16, 2002

COLLIDER-ACCELERATOR DEPARTMENT
BROOKHAVEN NATIONAL LABORATORY
UPTON, NEW YORK 11973

Beam-Based BPM Alignment

Richard Talman and Nikolay Malitsky

Accelerator-Collider Department
Brookhaven National Laboratory

ABSTRACT

An operational, beam-based, null-measurement, control room procedure designed to steer the closed orbit through the effective (no steer) center of every quadrupole is described. Performance of the procedure is simulated using UAL (Unified Accelerator Libraries). Matching SNS hardware availability, quadrupole strengths are assumed to be trimmable, but only in families, not individually. The accuracy of the procedure is unaffected by geometric and/or electrical misalignment of BPM's (beam position monitors) but calibration of their misalignments is a byproduct of the procedure. Some of the many possible failure mechanisms have been modeled, and have been found not to invalidate the procedure. A mini-introduction to the use of UAL is also given.

Table of Contents

1. Definition of the Task	2
2. Orbit Smoothing Algorithms	4
3. BPM Alignment at SNS with Quadrupoles Ganged in Families	5
4. Fundamental Limitations	7
5. UAL Simulation Tools	9
6. Numerical Output From the Simulation	13
7. Relaxation of Hardware Requirements	15
8. Conclusions and Comments	16

1. Definition of the Task

We are concerned with steering the closed orbit through the centers of all quadrupoles. Presumably the geometric centers of the quadrupoles themselves and the geometric centers of their BPM's have been aligned fairly accurately during the manufacture and installation phases. But magnet imperfection may cause significant displacement of the effective quadrupole center (position of no particle deflection) from its geometric center and electronic imperfection may cause significant displacement of the effective BPM center (position the data acquisition system reports to be zero) relative to *its* center.

In this report, to avoid ambiguity, when “center” is used, be it quad center or BPM center, the meaning will always be *effective* center that is meant. From the control room the absence of steering is the determinant of beam passage through effective quadrupole center and output of zero[†] from the BPM is the determinant of beam passage through the effective center of the BPM.

If all geometric alignments, and all effective-relative-to-geometric alignments have been performed with satisfactory accuracy during construction and installation the BPM's can be used to achieve the desired closed orbit steering. The present paper describes a control room, beam-based, procedure that is independent of the BPM alignment and can therefore corroborate, or even supercede, the installation accuracy.

In the control room one has *passive* access to the displacement of the closed orbit position at every BPM location and one can *actively* alter the closed orbit using horizontal and vertical steering elements, also referred to as kickers.[‡] Also, for those quadrupoles having trim windings, it is possible to influence the lattice actively.

An ideal arrangement would deliver every quadrupole with a full “detector/adjuster package” consisting of trim winding, horizontal and vertical kickers, and horizontal and vertical BPM.[¶] There is a natural “null measurement” operational procedure that can be

[†] To be boringly explicit, “zero” output from the BPM means that the actual digital number produced by the data acquisition system, from the analog BPM signal(s), has the value that corresponds to the beam's having passed through the geometric center of the BPM to the best knowledge of the control system. But, in what follows, it will be the ability of BPM's to accurately measure *changes in position* that will be important. Their absolute readings will be irrelevant.

[‡] In this paper the terms BPM and detector are used interchangeably. Also the terms kicker, steerer, adjuster, and corrector are synonyms.

[¶] It will turn out that missing BPM's will not invalidate the proposed procedure but the kickers must be present, for the procedure to work.

performed using a quadrupole (call it quad i , and let its inverse focal length be q_i) endowed with such a package. Taking the quad center as origin, let the closed orbit position be (x_i, y_i) . The effect of the quadrupole is to cause angular orbit deflections

$$\Delta x'_i = -q_i x_i, \quad \Delta y'_i = q_i y_i. \quad (1.1)$$

The effect of making fractional change f (absolute change $f q_i$) in the quadrupole's strength is to introduce a further kink ($\delta \Delta x'_i = -f q_i x_i$, $\delta \Delta y'_i = f q_i y_i$) that changes the closed orbit. This changes not only (x_{di}, y_{di}) but also the complete set of closed orbit measurements at all N_d BPM locations, (x_{dj}, y_{dj}) , $j = 1, 2, \dots, N_d$. A simple operational procedure is then to adjust the local kicker values to “null out” this closed orbit change; the kicker strengths will be $\delta x'_i = -\delta \Delta x'_i$, $\delta y'_i = -\delta \Delta y'_i$. (In principle the nulling could employ a single BPM (which need not be the i 'th) but a more robust (less subject to noise) procedure would be to average numerous BPM's.) From the available data one obtains

$$x_i = \frac{\delta x'_i}{f q_i}, \quad y_i = -\frac{\delta y'_i}{f q_i}. \quad (1.2)$$

This information can be used to center the beam on the quad, or more practically, if there is a local BPM, to calibrate the BPM so that its electrical center coincides with the quad center, both horizontally and vertically. The BPM will then serve as a secondary, or transfer, standard. After all BPM's have been calibrated in this way they can be used for a grand smoothing that puts the beam through the centers of all quads. Even if a quad lacks a BPM the closed orbit can be adjusted to pass through the quad center, provided there is a local steering elements that can be used for the null measurement.

Unfortunately, in practice, all quadrupoles are not necessarily delivered with the full detector/adjuster package. In the case of SNS, though there are trim windings on all quadrupoles, the quadrupole trims are not individually powered. Rather the quadrupoles are grouped in families of 8 having trim windings powered from a single power supply. The purpose of this report is to generalize the null calibration procedure in this circumstance and to simulate its performance using UAL.

2. Orbit Smoothing Algorithms

It has to be assumed that there is a control program which uses N_d detectors (BPM's) and N_a adjusters to smooth the orbit horizontally, where “smooth” may mean that all measured offsets are zero. More commonly there is a redundantly generous distribution of BPM's so that $N_d > N_a$, so the “badness”

$$\mathcal{B}(\delta x'_1, \delta x'_2, \dots, \delta x'_{N_a}) = \sum_{i_d=1}^{N_d} \frac{x_{i_d}^2}{\beta_{i_d}^{(x)}} \quad (2.1)$$

can be minimized but not be made to vanish. Here \mathcal{B} is expressed as a function of the adjuster deflections $\delta x'_1, \delta x'_2, \dots, \delta x'_{N_a}$ since they are the quantities to be varied in order to minimize \mathcal{B} . Mathematically this leads to the equations

$$\frac{\partial \mathcal{B}}{\partial (\delta x'_{i_a})} = 0, \quad i_a = 1, 2, \dots, N_a. \quad (2.2)$$

In this report the UAL algorithm (TEAPOT module) that simulates orbit smoothing is called `hsteer` and the corresponding vertical algorithm is `vsteer`. Programs like this rely on the optical model of the lattice to calculate an “influence function” $T_{i_a}(i_d)$ which is the closed orbit displacement at detector i_d caused by unit deflection at adjuster location i_a . Starting from closed orbit displacements $x_{i_d}^{(0)}$, the effect of applying kicks $\delta x'_{i_a}$ is to produce displacements $x_{i_d}^{(0)}$ given by

$$\frac{x_{i_d}}{\sqrt{\beta_{i_d}^{(x)}}} = \frac{x_{i_d}^{(0)}}{\sqrt{\beta_{i_d}^{(x)}}} + \sum_{i_a=1}^{N_a} \delta x'_{i_a} T_{i_a}(i_d). \quad (2.3)$$

Letting $\mathcal{Q} = (\delta x'_1, \delta x'_2, \dots, \delta x'_{N_a})^T$ be the (transpose of the) vector of unknowns and substituting Eqs. (2.3) into Eqs. (2.2) yields equations (in matrix form)

$$\mathbf{M} \mathbf{Q} = \mathbf{V}, \quad (2.4)$$

where

$$\mathbf{M}_{ab} = \sum_{i_d=1}^{N_d} T_a(i_d) T_b(i_d), \quad V_a = - \sum_{i_d=1}^{N_d} \frac{x_{i_d}^{(0)}}{\sqrt{\beta_{i_d}^{(x)}}} T_{i_a}(i_d). \quad (2.5)$$

Solving Eq. (2.4) yields kicker values which minimize the badness.

In an uncoupled lattice containing no nonlinear elements this operation would need to be done just once, and the same for vertical. But with skew elements and/or sextupoles

or other nonlinear elements present in the lattice there are horizontal deflections due to vertical offsets (and vice versa). This typically makes it necessary to proceed by successive iterations of horizontal steering, vertical steering, and decoupling. But, for simplicity in this report, the lattice is assumed to be uncoupled.

3. BPM Alignment at SNS with Quadrupoles Ganged in Families

Consider, for example, the family consisting of the $N_q(=8)$ quadrupoles labeled QFH in the SNS lattice whose MAD lattice description file is `BmBasedBPMAAlign.mad`. This file differs only from file `ff_sext_latnat.mad` by name changes made for the present simulation. Both files are available at <http://www.ual.bnl.gov>. The task is to measure all N_d horizontal misalignments and all N_d vertical misalignments. Of course there are also many other quadrupoles, grouped in other families. The procedure described here is immediately applicable to all such families. It is not even required that all nominal quadrupole strengths in the same family be equal or that the fractional trim strengths be equal. But, in this report, these simplifications will be made.

The strategy to be followed is much the same as with a single quadrupole trim. An intentional systematic change of the strengths of the quadrupoles in a single family causes the closed orbit to shift because of the (random and unknown) displacements of the quadrupoles in the family. Using an orbit smoothing algorithm the associated kicker magnets can be adjusted to undo this change. Then the individual quad misalignments can be inferred from the kicker strengths and the nominal quadrupole strengths using Eqs. (1.2). At that time all BPM offsets would be recorded to enable subsequent use of the BPM's as "secondary standards".

Concentrating first on the horizontal measurement, it is essential now to restrict the adjusters being used to precisely those associated with the quadrupoles in the QFH family. Therefore $N_a = N_q$. To avoid distraction it also seems appropriate to concentrate attention only on the associated detectors, so $N_d = N_q$.[†] So from here on the term "perfectly smooth closed orbit" is equivalent to $\mathcal{B} = 0$ where \mathcal{B} is given by Eq. (2.1) with $N_d = N_q$. Because

[†] In the single quadrupole algorithm it was said to be appropriate, for noise averaging etc., to utilize detectors other than the local one. For the same reasons, even in the multiple quad problem, detectors outside the family could be put to beneficial use. To simplify setting up the lattice this was not done in our simulation and it was not needed because the simulation did not include random BPM noise.

$N_a = N_d$, the number of equations (2.2) is equal to the number of unknowns. Therefore the equations have a unique solution. If the lattice were ideal (except for the misalignments being investigated) this would be the end of the story. But, for reasons described in other sections, when the calculated kicker values are installed the value of \mathcal{B} is still found to differ from zero. This may necessitate proceeding by successive approximation.

In any case one eventually achieves the result $\mathcal{B} = 0$, be it in simulation or in the control room. Repeating, for emphasis, what has already been implied, this only means that the orbit is perfect *as far as the QFH detectors are concerned*. Let us refer to this restricted closed orbit as the “QFH closed orbit”. The orbit shown by all BPM’s in the ring will not necessarily improve in the successive approximations described in the previous paragraph. In fact, our simulation shows that the closed orbit at points outside the QFH family frequently is made worse by a next approximation. Though disconcerting this is what is to be expected.

In the control room the QFH closed orbit *appears to be* perfectly smooth when all BPM outputs from the QFH family are zero. But this only means that the closed orbit has been adjusted to pass through the electrical centers of every QFH detector. Let us assume that this has been done (even though, in this circumstance, because of misalignments, the actual closed orbit may deviate badly from the true design closed orbit, which is defined to pass through quad centers.) The smoothing algorithms used to achieve this were described in the previous section.

Next we apply the systematic fractional strength change f to all quadrupoles in the QFH family. This causes the QFH closed orbit to be no longer smooth. Applying `hsteer` and `vsteer` again yields the kicker strengths needed to re-smooth the orbit. Finally the misalignments being sought are given by Eq. (1.2). This completes the determination of closed orbit displacement relative to quad centers at all quads in the family.

4. Fundamental Limitations

As with all operational procedure, the BPM calibration can be compromised by world realities. Some BPM's may not function at all and electronic noise will cause fluctuation of the measured positions. These effects are perhaps the ones most likely to limit the practicality of the procedure being described here. For this procedure the absolute accuracy of BPM's is irrelevant, but it is important that their least count correspond to a very small distance—or rather that they be capable of stably and reproducibly recording very small beam position *changes*.

Nevertheless, since one has no way of predicting such electronic malfunctioning, except from empirical data or at such low signal levels that shot noise dominates, we have so far assumed in the simulation that all BPM's function perfectly.

Another practical complication is that BPM's, though physically close to their associated quadrupoles, cannot be precisely superimposed. But, provided they are reasonably close and that signals are available from other nearby BPM's, values can be accurately interpolated to the precise quadrupole locations. Even if vertical quads are restricted to vertical focusing quad locations, and horizontal to horizontal (as is common) there are reliable interpolation procedures to produce the signals assumed in this report. Similarly, even though kickers are not precisely in their ideal locations, the kicker strengths can be appropriately adjusted. In the simulation all such complications are ignored; it is assumed that all detectors and adjusters are ideally superimposed on their associated quadrupoles. How to account for non-optimal kicker locations is considered briefly below.

There are other more fundamental effects that potentially limit the practicality of the proposed method. The effect of increasing a quadrupole strength is not just to cause a steering effect proportional to the quad offset. There are also changes in the lattice optics, both tunes and beta functions. At worst the change in quad strength could make the lattice unstable and, at a minimum, the changes in lattice optics will cause the closed orbit fitting programs to be somewhat inaccurate. Such a limitation is already present in the single quadrupole procedure of section 1. In the interest of increased signal to noise ratio one wishes to make the quad strength increment $f q_i$ as large as possible, but the need to limit lattice distortion forces one to compromise. Nonlinear elements present in

the ring also limit the accuracy of the procedure. Nonlinear elements would not affect the one quadrupole, null measurement of section 1, but their presence reduces the accuracy of the closed orbit algorithms. All effects mentioned in this paragraph are subject to investigation using UAL or another simulation code. Our (very limited) investigations started with the guess that a one percent alteration of quad strengths ($f = 0.01$) would be satisfactory. At this level we find the algorithm to be essentially unaffected by changing the chromaticities from their natural (all chromaticity sextupoles off) values to being zero in both planes. Similarly the procedure is little affected by the inclusion or exclusion of magnet imperfections at anticipated levels.

The achievable accuracy can be estimated as follows. Let us concentrate on vertical orbit smoothing. If the lattice is taken to consist of nothing but 90 degree FODO cells and the tune is Q there will be $8Q$ quads altogether, each with its local BPM. But of these only half are close to vertical quads where their accuracy is high and only about half of those are favorably located (approximately an odd number of half-betatron-periods from the quad) relative to a particular vertical steering that is being nulled. If the r.m.s. position error at a quad of strength q_i is σ_y , the r.m.s. deviation of the deflection to be nulled for trim factor f is $f q_i \sigma_y$. The downstream displacement caused by such a deflection is

$$\sigma_d \lesssim \beta_{\text{typ.}} f q_i \sigma_y \quad (4.1)$$

For individually trimmed quads the $2Q$ “useful” detectors would improve the nulling precision by a factor $1/\sqrt{2Q}$. The effect of being forced to trim the N_q quads in a family will exact a loss of accuracy which will erode this factor to $1/\sqrt{2Q/N_q}$ (or worse.) Incorporating this estimate in Eq. (4.1), using the estimate $\beta_{\text{typ.}} q_i \approx 1$ and solving for σ_y yields

$$\sigma_y \gtrsim \frac{\sigma_d}{f} \sqrt{\frac{N_q}{2Q}}, \quad (4.2)$$

as the estimated accuracy with which the closed orbit can be steered through the quadrupole center. Taking the square root factor as 1, the precision with which the orbit can be steered through the quad is approximately the BPM precision eroded by factor $1/f$. With f being of order 0.01 the steering accuracy is 100 times worse than the measurement accuracy. To achieve 0.1 mm steering accuracy will require something like $1 \mu\text{m}$ BPM reproducibility.

Note, though, that it is *short term reproducibility* not absolute or even long term relative accuracy that is required. Perhaps the required precision could be attained using very low frequency excitation with lock-in detection. Least count precision of the steering power supplies may also be an issue, as the required deflections are very small.

5. UAL Simulation Tools

This is not the place for full documentation of the UAL simulation environment. But to give the flavor of the user interface a few fragments of code will be given. Two points are especially important:

- UAL has no proprietary input language, neither for lattice description nor to specify the simulation commands to be executed. Rather the input takes the form of a valid PERL program, where PERL is a popular “scripting” language. This program can read in a standard lattice description, such as a MAD description or SXF or ADXF, make ad hoc specialized changes not included in these descriptions, and then proceed to the desired simulation.
- The mechanism for specifying families of elements (such as all quadrupoles of type QFH) is very different from the `type=` mechanism of MAD and other lattice programs. Rather a family is specified by the “explicit” listing of the names of the elements making up the family.[†]

A crude outline of the SNS lattice is shown in Fig. 5.1 Here are a few fragments from the MAD lattice description file:

[†] The word “explicit” is placed in quotation marks because the so-called “regular expression” mechanism is used to specify families and the element selection by regular expression may not look all that explicit to a reader unfamiliar with regular expressions. This will be clarified somewhat by examples to be given shortly. The consistent naming of lattice elements has always been an important consideration in writing lattice description files. It is important for conveying the intended purposes for the various elements and for the exact association of elements in external models with elements in the tunnel. UAL makes naming scheme discipline all the more important. It is very convenient for all elements of a family (for example because they share the same bus) to have the same name. But this is too much to ask in general as it fails to allow for “overlap” of the different sorts of family that need to be defined. The regular expression mechanism permits the efficient selection of elements even in the face of such type overlap, but only if a consistent naming scheme is carefully respected. As a last resort a family can, in fact, be defined within UAL as a *really* explicit list of all of its elements. An example of this will be given below. When first encountered the regular expression approach may seem awkward to the user but it is a “feature”, not a “bug”, as it solves a really hard simulation problem—how to specify families without the need for ad hoc tampering with the lattice description language? Such tampering frequently leads to errors and always erodes portability.




Figure 5.1: SNS lattice showing the particular family of quadrupoles used in the paper to illustrate beam-based quadrupole alignment when quadrupoles cannot be trimmed individually. The QFH elements are half-quads. In real life the BPM's and kickers would be somewhat displaced from the actual full quads.

```

! BmBasedBPMAAlign.mad
...
! Half Quadrupoles
QFH  : QUADRUPOLE, L = lq/2, K1 = KF/Brho ! defocusing arc quad (21Q40)
QFBH : QUADRUPOLE, L = lq/2, K1 = KF/Brho ! large focusing arc quad (26Q38)
...
! Steering elements and BPM's
kickha : HKICKER
bpmha  : HMONITOR
...
QF_a  : LINE = (QFH1,QFH,kickha,kickva,bpmha,bpmva,QFH,QFH2)
QF_b  : LINE = (QFH1,QFH,kickhb,kickvb,bpmhb,bpmvb,QFH,QFH2)
...
ACF_a : line = (OARC,BND,OARC,QF_a)
...
ARC_a : line = (ACF_a,ACFM,ACS1,ACS2,ACS3,ACS4,ACF_e,ACFL)
...
SP_a  :      line = (INSERT,ARC_a)
...
RING  : line = (SP_a,SP_b,SP_c,SP_d)

```

As cryptic and abbreviated as it is, along with the figure this fragment is intended to show, more or less, how the eight QFH elements are distributed around the ring and how their associated kickers and BPM's are given similar, but distinguishable, names, kickha, kickhb, ..., kickhh, and bpmha, bpmhb, ..., bpmhh, and so on.[†] The ring consists of four similar sectors, with two QFH elements in each. Already in the MAD description the QFH quads are, in fact, half quads. This is convenient in that the kickers and BPM's have been situated at the centers of the actual physical quadrupoles. Of course this is an unphysical idealization of what is possible in practice but it avoids the need for the sort of “housekeeping” corrections described elsewhere in this report.

Next to be listed are some fragments from the UAL command file (i.e. PERL program) `BmBasedBPMAAlign.pl` used for the simulation being described. This program is available at the same site as the files listed previously. The little circled numbers are included for convenience of reference in the following section. First the line ①

```
$shell->readMAD("file" => "./data/BmBasedBPMAAlign.mad");
```

reads in the MAD lattice description file. To make the simulation more realistic it is appropriate to introduce random field errors with statements like ②

```
$shell->addFieldError("elements" => "^(q[fd][lc]h|qfbh)$", "R" => 0.12,
    "b" => [0.0, 0.0, 0.4, 0.1, 0.7, -12.10, 0.0, 0.0, 0.0, 0.0]);
```

where the assigned erect multipole errors (in conventional units) at reference radius $R = 12$ cm are enclosed in square brackets. (For skew errors “b” would be replaced by “a”.) The actual values used come from early magnet measurements; they are probably not up to date. The cryptic expression ‘`^(q[fd][lc]h|qfbh)$`’ is the regular expression specifying the elements to which multipole errors are to be assigned. Only four comments will be made about this, but they should be enough to give the general idea of how families are defined. (i) Element names are case-insensitive. Uppercase names are converted to lowercase before processing. Hence, for example, “QFBF” becomes “qfbh”. (ii) The symbol ‘`^`’ forces name matching to start at the beginning of the name. So the ‘`^(q...`’ part

[†] The critical reader may complain about the inconsistency of naming QF_a with underscore but kickha, bpmha, etc. without. This was, in fact, the first notation adopted, but note that, after lower case conversion QFh becomes qfh, which clashes with the QFH name. This is the sort of care in naming whose necessity was warned of earlier.

of the regular expression allows names beginning with “q” to be candidates for inclusion (provided they also meet other later requirements). (iii) The symbol “...\$” forces name matching to end at the end of the name. So the “^(...)\$” part of the regular expression limits both the beginning and ending match. (iv) The part of the regular expression (... |qfbh) gives qfbh as one of the strings whose presence makes the name a candidate for inclusion. This (along with the requirements already mentioned) shows that the element “QFBH” from the code fragment above will be one of the magnets to have field errors added; its name (after conversion to lower case) includes the string “qfbh” and matches both at beginning and end. This is an example of the “last resort” explicit name inclusion mentioned above in that element name “QFBH” is specified “explicitly”.

The next UAL command begins the actual simulation: ③

```
$shell->addMisalignment("elements" => "^qfh$",
                        "dx" => 0.01, "dy" => 0.01, "engine" => $rgenerator);
```

This applies random, Monte Carlo (version specified by \$rgenerator) generated alignment errors, with r.m.s. deviations of 1 cm in both planes, to all QFH elements. The subsequent task of the simulation is to recover these numbers using the beam-based calibration procedure.

The next instructions are ④

```
$shell->hsteer("adjusters" => "^kickh", "detectors" => "^bpmh");
$shell->vsteer("adjusters" => "^kickv", "detectors" => "^bpmv");
```

These adjust the kicker magnets to smooth the orbit. As explained previously the resulting orbit is “smooth” only in the sense that it passes through the electrical centers of the QFH BPM’s, but this completes the simulation of the operational set-up phase.

The next step is ⑤

```
$f = 0.01;
$dqInUnits = 10000.0*$f;
$shell->addFieldError("elements" => "^qfh$", "R" => 0.1,
                    "b" => [0.0, $dqInUnits]);
```

This introduces the fraction f by which quad strengths are to be incremented (in this case one percent) converts it into absolute units and increments all QFH elements.

Another smoothing is required next: ⑥

```
$shell->hsteer("adjusters" => "^kickh", "detectors" => "^bpmh");
$shell->vsteer("adjusters" => "^kickv", "detectors" => "^bpmv");
```

Finally the kicker strengths found in this smoothing are recovered: ⑦

```
$shell->getMagnetStrengths("elements" => "^kickh",
                          "order" => 0, 'attid' => 'b');
$shell->getMagnetStrengths("elements" => "^kickv",
                          "order" => 0, 'attid' => 'a');
```

The actual control file has more commands than have been exhibited, for example to perform iteration steps or to generate output data.

6. Numerical Output From the Simulation

In this section some output data from the simulation sketched in the previous section is given. (The UAL commands directing the calculation were shown but the commands causing the printout were not.) The next table was output after step ④ of the previous section.

q_num	el_index	kl1_bef	del_x	del_y	defl_x	kickh	defl_y	kickv
0	49	0.17187	0.00313	-0.00145	0.00145	-0.00146	-0.00240	-0.00266
1	54	0.17187	-0.01158	-0.01250				
2	105	0.17187	0.01003	0.01144	-0.00138	0.00139	0.00102	0.00120
3	110	0.17187	-0.00202	-0.00552				
4	164	0.17187	-0.00388	-0.00507	-0.00131	0.00130	-0.00209	-0.00217
5	169	0.17187	0.01151	-0.00712				
6	220	0.17187	-0.00077	0.01405	0.00222	-0.00223	0.00406	0.00404
7	225	0.17187	-0.01213	0.00958				
8	279	0.17187	0.00069	0.00882	-0.00008	0.00014	-0.00088	-0.00097
9	284	0.17187	-0.00023	-0.01394				
10	335	0.17187	-0.02179	-0.00303	0.00115	-0.00112	-0.00091	-0.00085
11	340	0.17187	0.01509	-0.00225				
12	394	0.17187	-0.00739	0.00697	0.00133	-0.00136	0.00115	0.00121
13	399	0.17187	-0.00033	-0.00028				
14	450	0.17187	0.00394	-0.00516	-0.00131	0.00129	0.00192	0.00190
15	455	0.17187	0.00365	0.01632				

This listing has a row for each of the 16 half-quadrupoles. Within TEAPOT, like MAD, the lattice is completely “flattened” into a sequential list of elements. The second column gives the indices of the QFH quads. The “kl1_bef” column gives the starting values of the quadrupole strengths (inverse focal lengths in inverse meters.) These entries were obtained using an instruction like ⑦. The columns “del_x” and “del_y” show the alignment errors introduced in step ③. (Of course these values would be unknown in actual operations.)

Because the two quad halves have been assigned uncorrelated random misalignments the closed orbit shift is proportional to the sum of these misalignments. It would be more realistic to correlate the misalignments applied to the separate halves, perhaps even modeling quad rotation as well. What has been done is equivalent to having applied both random translation and random rotation misalignments. The net horizontal deflections (in radians), (sum of displacements)*quad-strength, are given in column “defl_x”. The horizontal kicker strengths found in step ④ are given in column “kickh”. Even though the actual BPM calibration has not yet begun these numbers are clearly promising, since the kicks found (both horizontal and vertical) agree well with the deflections caused by the quad misalignments. But this determination was “easy” for `hsteer` and `vsteer` since the closed orbit shift (as well as being proportional to the misalignment) is proportional to the full quadrupole strength. In the next step the closed orbit shift is derated by the fractional increment factor f which is of order one percent.

The next table was output from UAL after the completion of step ⑦ .

q_num	el_index	del_kick_x	-defl_x*f	del_kick_y	defl_y*f
0	49	-0.0000147	-0.0000145	-0.0000262	-0.0000240
1	105	0.0000141	0.0000138	0.0000118	0.0000102
2	164	0.0000132	0.0000131	-0.0000214	-0.0000209
3	220	-0.0000227	-0.0000222	0.0000399	0.0000406
4	279	0.0000015	0.0000008	-0.0000096	-0.0000088
5	335	-0.0000114	-0.0000115	-0.0000084	-0.0000091
6	394	-0.0000138	-0.0000133	0.0000119	0.0000115
7	450	0.0000131	0.0000131	0.0000187	0.0000192

There is one row for each quadrupole (two half-quads). The column labeled “del_kick_x”, the only new horizontal data in the table, gives the *change* in kicker strength that `hsteer` found was needed to re-smooth the orbit after the quad strengths had been trimmed. These entries are to be compared with entries in the column “-defl_x*f” which are the deflections caused by the quad trimming; they are calculated as the deflections from the previous table de-rated by factor f . The agreement between these two columns confirms the applicability of Eqs. (1.2).

The corresponding vertical entries also agree well. (The sign difference is an artifact of the vertical sign convention.) But the agreement is less good. This is probably what should have been expected since the quad family being studied is horizontally focusing so

the horizontal beta function is larger than the vertical beta function at locations of this particular family of quadrupoles.

7. Relaxation of Hardware Requirements

It has been emphasized that the alignment procedure relies on a null measurement. A quad deflection is nulled against an adjacent steering element. It follows that any particular quadrupole can only be determined if there is a steering element nearby. (Exactly *how near* has not been specified. For an example discussed below the error from this extrapolation must be held to less than 0.1 mm.) There was therefore no choice as to which adjustors had to be used and the condition $N_a = N_q$ was automatic. But it was not really necessary to use detectors adjacent to the quads being studied. Formally at least, any N_q could have been chosen and a unique solution would follow. But that is not the point; choosing $N_d \gg N_q$ is even better. The solution is still unique and the more detectors in use the more robust the measurement. If some of the quads being measured lack BPM's it really doesn't matter, except for the inconvenience of not having a secondary standard for beam centering at that location.

Another relaxation concerns the importance of accurate location of BPM's relative to quads. There is no importance whatsoever. The BPM's are only being used to establish the nulling, so their positions don't matter. This is much the same argument as was made in the previous paragraph.

Another relaxation relates to the actual operational procedure. So far it was stated that a set-up phase in which the QDF closed orbit was smoothed was to be followed by the actual calibration phase. But the set-up phase was really superfluous—as much as anything it was introduced as a pedagogical ruse to exercise and illustrate the simulation. Starting from any ragged orbit (though not so ragged as to bring nonlinear elements into the picture), after applying the systematic (factor f) quad trim, the only requirement is for the smoothing program to restore the closed orbit. The adjustor strengths needed to do this are just those that null the quad misalignment kinks.

8. Conclusions and Comments

An algorithm for centering the beam on all quadrupoles has been described. The algorithm is applicable even when multiple quad trims are powered from the same bus. Some effects that could potentially cause the algorithm to fail have been investigated. Sextupoles of strength needed to adjust chromaticities to zero have negligible effect. So also do the random and systematic magnetic field errors assumed in the only lattice file investigated.

Other effects, potentially more limiting, have not been investigated. A simulation such as this one could, however, anticipate the degree to which this calibration procedure would be reliable. Electronic noise and stability could be estimated and included. Also instrumentation issues such as the required least count precision of analog to digital conversion could be addressed.

For actual coding of the algorithm into the control system a certain amount of “house-keeping” would be needed. For robustness, allowances should be made for the occasional broken or missing BPM or kicker. Also correction will be needed to compensate for the effects of small relative longitudinal displacements of quads, BPM’s, and kickers. The sort of modification required is illustrated in Fig. 8.1. Because the local kicker K_i is slightly displaced longitudinally it cannot, by itself, completely cancel the effect of strength trimming of quad q_i . But a (relatively much smaller) compensating kick by kicker K_{i-1} can null the orbit everywhere outside the range from K_{i-1} to K_i . Determination of the offset x_i will still be dominated by the setting of K_i , but with a correction (smaller by a factor of order l/L) depending on K_{i-1} . Since the interpolation error in this procedure is itself proportional to the local slope, it is subject to being reduced in subsequent iterations.

The procedure that has been described has to be applied successively to all the quad families in the ring. (If there are no trim circuits whatsoever one can (at least in principle) apply the method by using the main quad bus or buses.) Once all quads displacements have been determined the information is available for a final application of the smoothing algorithms to put the closed orbit through the effective centers of every quadrupole.

The measurable signal is proportional to the fractional quad trim f . Since the accuracy of the method is sure to be limited by BPM irreproducibility one will always be tempted to increase f . But the tolerable value of f is limited by the tolerable level of lattice




Figure 8.1: Displacement x_i is being inferred by the strength of kicker K_i needed to null the effect of changing quad q_i by fractional amount f . Because kicker K_i is not quite superimposed on the quad, a small deflection by neighbouring kicker K_{i-1} is needed to null the orbit everywhere else in the ring.

distortion. M. Blaskiewicz points out that trimming pairs of families simultaneously may allow a large increase in tolerable f value if appropriate trim strengths and polarities are chosen.

Finally it can be asked “does this procedure yield the best possible closed orbit?” Any answer to this question has to be qualified since factors other than those considered (such as the need to avoid obstacles) may enter. Also it would not be unprecedented for the BPM’s to be so erratic that they could not even produce reliable null settings. Let us ignore such possibilities. With this qualifications there is an interpretation of the question above such that the answer is “yes!” The reason accelerators are high-strung and tempermental is not that they have many-many bending magnets but that they have many-many lenses. If the closed orbit goes through the effective centers of all quads then the optics will agree well with the design optics and will be relatively insensitive to quadrupole deviations from their design global positions. This means that (like the BPM positioning) the quad positioning is also not critical *provided one steers the beam through quad centers*.

The previous paragraph may be regarded as too glib. A more conventional specification would be to require the orbit to be coplanar to some accuracy, say ± 0.5 mm. In this circumstance the ultimate achievable coplanarity will be limited by the accuracy with which effective quad centers lie on the same horizontal plane. With customary survey methods it should be possible to position the geometric centers with an accuracy of 0.1 mm, and it should be possible to maintain comparable accuracy for the relation between geometric

and effective quad centers. If these conditions have been met, the procedure advocated in this note could be useful provided its accuracy is comparable. This would require the precision given by Eq. (4.2) not to exceed 0.1 mm. Once all these things have been achieved it would surely be vertical steering caused by bending magnet roll errors that would limit the coplanarity. If these errors are sufficiently small an orbit coplanar to better than ± 0.5 mm would seem to be achievable.

There is one important situation in which steering through quad centers is not optimal. Suppose it is critical that the closed orbit be optimized for the preservation of spin orientation. Unwanted spin precession occurs due to vertical bends, be they due to vertical displacement in quads, to roll errors in dipole magnets, or to vertical steering correctors. In this case it is not optimal for the closed orbit to pass through quad centers. Precession due to steering in a quad can be cancelled in the adjacent steering element, provided, of course, that there is no intervening bend. If there were no roll errors in bending magnets and all effective quad centers were coplanar (neither of which is realistic) the optimal orbit would pass through quad centers, but the realistically optimal orbit has to be a compromise among the different vertical steering contributions. Any practical scheme for achieving this compromise would rely on the absolute BPM accuracy.

One way the null method of this note could contribute would be to measure the alignments of BPM electrical centers relative to quadrupole effective centers, since these precisions determine the accuracy of practical optimization schemes. An even more optimistic scenario assumes that all bending magnet roll errors have been eliminated but the effective quad centers have unknown coplanarity. Then the orbit best preserving polarization could have appreciable bending in individual quads provided the bending was immediately canceled by the adjacent steering corrector. In this case it would be advantageous for the (bend free) separation between quad and steering corrector to be large, not small, so vertical orbit excursions could be corrected with small deflections.

We wish to acknowledge useful conversations with M. Blaskiewicz, W. Fischer, Y. Lee, D. Raparia (who suggested the investigation), and J. Wei. We make no claim of originality for anything in the report (other than the simulation tools.) The ideas are all pretty obvious and have no doubt been published in the past. For example a recent paper by D. Barber et al. *Application of a Beam Based Alignment Technique for Optimizing*

The Electron Spin Polarization at HERA, EPAC96, Sitges/Spain (1996) describes, with numerous references, procedures that appear to be similar to the ones considered in this report.