Beam-Based Alignment of the NLC Main Linac, Part Two: Dispersion-Free Steering and Dispersion Bumps

PETER TENENBAUM

LCC-NOTE-0071

10-Sep-2001

Abstract

The standard prescription for NLC main linac steering assumes that the RMS offset between a quad's magnetic center and the corresponding BPM's electrical center can be determined at the level of a micron. This is a fearsome requirement, and will be particularly difficult to achieve if hybrid iron/permanent magnet quads are used in the main linac. As an alternative, the Dispersion Free Steering (DFS) [1] algorithm is adapted to the NLC main linac environment; the DFS algorithm does not require knowledge of the quad-BPM offsets. The results of simulation studies of this adaptation are presented. In addition, the use of closed orbit bumps to globally correct dispersive emittance growth is considered. The studies indicate that DFS can be used successfully in the NLC main linac environment, and that dispersion bumps are a useful addition to the linac steering "toolbox," regardless of the main algorithm selected.

1 Introduction

It is well known that the principal sources of single-bunch emittance dilution in the NLC main linac will be wakefields from misaligned RF structures and dispersion from misaligned quadrupoles. For the present NLC parameters ($\gamma\epsilon_y = 2.0 \times 10^{-8}$ m.rad, $N = 0.75 \times 10^{10}$ particles, $\sigma_z = 110 \mu\text{m}$, $\sigma_E/E \approx 0.6\%$), RMS beam-to-structure misalignments at the level of $10 \mu\text{m}$ will cause an emittance increase of 25%, while RMS beam-to-quad misalignments at the level of $2 \mu\text{m}$ will cause a similar emittance increase [2]. Such tight tolerances on the beam-to-element offsets cannot be achieved by conventional survey and alignment techniques, but must rely on *beam-based alignment* (BBA). Note that in this context, we define "beam-based alignment" to mean *all* of the techniques which use measurements of beam parameters to estimate or infer beam-to-element offsets and determine the appropriate corrections. It is common to use this phrase to refer to a particular method of BBA, in which variation of quad strength is used to measure the beam-to-quad offset; in this Note, we shall explicitly refer to this as the "method of quadrupole variation."

Beam-based alignment of the NLC main linac is essential, and therefore the linac is well equipped with instrumentation and controls for this purpose. Figure 1 shows the main linac hardware which is intended for use in BBA: quadrupole and RF girder translation stages which allow elements to be repositioned during beam operation, and quadrupole and RF-structure BPMs which can accurately measure the beam position.

The procedure for aligning the RF structures to the beam is quite straightforward. Each structure includes a set of damping manifolds which are required for managing the long-range (bunch-to-bunch) transverse wakefields. By measuring the amplitude and phase of the RF power in one vertical and one horizontal manifold, the beam-to-structure offset can be directly determined. This measurement is unambiguous in that a "zero" reading from the BPM – zero power in the manifold – means that the beam is going through the structure with zero offset. Furthermore, the wakefield of a single structure is sufficiently weak that changing the transverse position of a few structures does not produce a significant change in the centroid trajectory downstream; this is in contrast to a quadrupole, which steers the beam when it is moved. Consequently, an RF girder can be aligned to the beam by simply measuring the beam position at the upstream and downstream

ends of each structure on the girder, fitting a straight line to the resulting measurements, and setting the girder translation stages to zero the average offset and slope of the BPM readings.

By contrast, the alignment of the quads to the beam is relatively difficult. One complication is that moving a quadrupole steers the beam, and so if one quad is moved to zero its offset the offsets in the remaining quads will change. A more difficult problem is that, unlike the RF structures, there is an ambiguity in the reading of the BPM associated with a quad: when the BPM reads zero, the beam may be centered in the BPM but probably not in the quad. The offset between the BPM and the quad is a function of the initial installation procedure, the cabling and electronics of the BPM, the variation of the quad's magnetic center position as a function of time, and many other variables. Thus, while the linac can easily be steered, via the quadrupole movers, to minimize the RMS BPM readings, this will not ensure that the RMS beam-to-quad offsets will be minimized.

1.1 The Technique of Quadrupole Variation

One method of addressing the ambiguity described above is to first measure the offset between each quad's magnetic center and the corresponding BPM's electrical center. This allows accurate determination of the BPM reading that corresponds to a beam-to-quad misalignment of zero. In this model, quad alignment is a two-step process: first the BPM-to-quad offset is determined, and then this knowledge is incorporated into the steering algorithm and the quads are moved to minimize the true beam-to-quad RMS offset.

The most popular method for measuring the BPM-to-quad offset is the method of quadrupole variation, shown schematically in Figure 2. A beam which enters a quad with an offset x_q receives a deflection $x' = K_q x_q$, where K_q is the inverse focal length of the quad. By varying the quad strength and measuring the downstream deflection, the value of x_q can be deduced; the BPM-to-quad offset is the difference between the beam-to-quad offset and the quadrupole's BPM reading.

The use of quadrupole variation followed by quad-mover steering to minimize the RMS beam-to-quad offset has several advantages. Chief amongst these is that the correction of the orbit is quite local: the beam-to-quad offset is minimized at every quad, and therefore the resulting correction is quite stable. The quadrupole variation technique is a nulling technique: because zero deflection of the beam indicates zero offset in the quad, which is the desired condition, the measurement is relatively insensitive to small errors in the knowledge of the quad strength, BPM scale factors, transport matrices, etc. There are, however, several disadvantages as well. For example, the quadrupole variation technique is quite invasive: it cannot be performed while the beams are colliding, since the beam orbit is changed and the change in focusing ruins the small beam sizes required at the IP. This in turn implies that the BPM-to-quad offsets must be stable, so that a single measurement of the offsets by this technique can be valid over many days, weeks, or months of operation and linac steering. This implies that the BPM electronics and mechanical installation must be quite stable, and that the quadrupole center position must be similarly robust.

In addition to these difficulties, there are important systematic effects which can fool the quad-variation measurement technique. The particulars of the systematics are different for electromagnet quads and various forms of permanent magnet quads. All of these systematics can cause an error in the fitted beam-to-quad offset which in turn becomes an error in the quad-to-BPM offset. Because of these issues, quadrupole variation may not be adequate to achieve micrometer-level accuracy in the measurement of quad-to-BPM offsets. Appendix A discusses the particulars of these systematic effects in greater detail.

2 Dispersion Free Steering

Because of the difficulties inherent in measuring the quad-to-BPM offset accurately, it is logical to seek a means of minimizing the dispersive emittance growth in the NLC main linac which does not require knowledge of the quad-to-BPM offsets. Such a technique was proposed in the early 1990's by Raubenheimer and Ruth [1]. In this technique, the matching of the beam energy to the quadrupole lattice is varied, and the resulting change in the orbit through the linac is measured. The change in orbit as a function of beam-lattice energy mismatch is the dispersion; thus, this technique directly measures the dispersion and permits a steering solution to be formulated which zeroes the dispersion at all points. Such "Dispersion-Free Steering" (DFS) relies upon the BPM resolution but not the quad-to-BPM accuracy.

Early tests of the DFS algorithm in the SLAC linac were not successful in reducing the end-linac transverse emittance. In these tests, the quadrupoles in the linac were scaled to emulate energy variation; it is now suspected that the quad centers may have moved during scaling, which would have led to systematic errors in locating the dispersion-free orbit. Subsequently, the DFS algorithm has been used successfully in many venues, including the SLAC linac; typically, successful DFS implementations actually vary the beam energy rather than the scaling of the lattice. This has led to consideration of DFS via energy gain variation as a possible technique for correcting the NLC main linac trajectory.

2.1 Algorithm Choices

Choices which must be made in the algorithm include the manner and degree of energy variation used, the longitudinal segmentation of the linac, and the relative balance of the algorithm towards minimization of dispersion and minimization of RMS quad motion.

2.1.1 Energy Variation

The issues related to quad center motion made clear that the NLC's DFS algorithm should operate by varying the linac energy gain, not the quad strength. This also enhances the speed of operation in the real machine: the energy gain of each RF girder can be varied on a pulse-by-pulse basis by changing the relative phases of the klystrons in each 8-pack.

Another concern is the possibility that RF deflections in the structures could cause systematic errors in the DFS algorithm. RF deflections are caused by structures which are pitched relative to the beam angle, by asymmetric fields in the input and output couplers, and other effects. These deflections would vary as the energy gains of the structures are altered, and the changing RF deflections must be properly accounted for in the algorithm. In order to manage the RF deflections, the beam energy gain is changed in a region of the linac, and dispersion is measured downstream of that region. In the area in which the dispersion is being measured, the beam energy is reduced because of the adjustment of energy gain, but the structures in the measurement area maintain their design energy gains. Thus, the change in RF deflections become a change in the incoming orbit of the beam, which can be eliminated by steering the beam into the measurement segment or by fitting the incoming oscillation.

2.1.2 Linac Segmentation

Because the energy gain in each measurement segment is left at its nominal value, the change in beam energy from the design is maximum at the upstream end of the segment and decreases along the segment. For example, a change in beam energy by 20% at the 10 GeV point is 2 GeV;

at the end of the linac, this is less than 1% of the beam energy, and the relative variation is so small that dispersive errors will be extremely difficult to resolve. Thus, the linac must be divided into measurement segments longitudinally. For this study, we have divided the linac into 25 such segments, each of which contains an equal number of quadrupoles. This means that the segments in the upstream end of the linac are only 33% as long as those in the downstream end, due to the change in quad spacing down the beamline.

2.1.3 Range of Energy Variation

The optimal results are obtained if the beam energy is varied as much as possible. The practical limitations on the range of energy variation come from optical stability and the energy acceptance of the post-linac energy collimation system. In these studies, we have limited the energy variation at the upstream end of a measurement segment to 20% of the design energy or 35 GeV (7% of the final energy for the 500 GeV main linac), whichever is less.

The 8 GeV injection energy of the main linac cannot be easily varied due to the dynamics of the bunch compressors. A variation of 20% of the design energy implies that the first quad which can be in a measurement segment is at the 10 GeV point in the main linac. The quads upstream of the 10 GeV point must be aligned by other means. For the purposes of this study it is assumed that they are perfectly aligned to the beam *ab initio*.

2.1.4 Fit Constraints

Beam-based alignment techniques are typically very sensitive to short-wavelength misalignments, but relatively insensitive to long-wavelength misalignments. While long-wavelength misalignments have little effect on the beam quality, the practical limitations of the vacuum chamber aperture, magnet mover ranges, etc. dictate that alignment solutions that improve the short-wavelength alignment but degrade the long-wavelength alignment must be suppressed. This limitation was encountered in the design of the main linac's steering algorithms as well as the DFS context [6].

DFS is typical of beam-based alignment schemes in that the long-wavelength misalignments which are generated by the algorithm must be minimized. The DFS algorithm used for the NLC main linac suppresses long-wavelength misalignments via a "soft constraint:" the RMS dispersion measured by the BPMs and the RMS motion of the magnet movers in a measurement region are simultaneously minimized.

2.1.5 Summation of the Algorithm

The DFS algorithm used for the NLC main linac is as follows:

- Subdivide the 10 GeV – 500 GeV region of the linac into 25 measurement segments with equal numbers of quads
- For each segment, determine how many upstream RF structures must be switched off to vary the energy at the upstream end of the segment by 20% of the design or 35 GeV, whichever is smaller
- Switch off the necessary structures
- Steer the incoming beam orbit to the position/angle which was present at the design energy
- Measure the change in BPM readings throughout the segment and in the first BPM downstream of the segment

- Determine the magnet moves which simultaneously minimize the RMS dispersion and the RMS magnet motion, with the former weighted 250 times as important as the latter
- Apply the fitted moves and perform RF girder alignment in the segment
- Iterate each segment 4 times before going on to the next segment.

2.2 Simulation of DFS

The algorithm described above was added to LIAR (LInear Accelerator Research code) [5]. Simulation was performed using the NLC 2000 main linac optics. These optics are similar to the 2001 design, but with several distinct differences:

- The 2000 linac was shorter than the 2001 design, and as a consequence of this the 2000 linac contains approximately 750 quads, rather than NLC2001's 820.
- Each RF girder contains 3 RF structures of 1.8 m length, rather than NLC2001's 6 structures of 0.9 m length.

It is not expected that the differences above will yield dramatically different results.

The beam conditions selected were based on the 2001 design values: a bunch charge of 0.75×10^{10} , RMS bunch length of $110 \mu\text{m}$, initial energy and uncorrelated energy spread of 8 GeV and 1.6%, respectively. The initial emittances were the design DR-extraction values ($\gamma\epsilon_{x,y} = 3.0 \times 0.02 \text{ mm.mrad}$). The linac phasing was selected to give an RMS correlated (BNS) energy spread of 0.6% to 0.7%, which was reduced to 0.25% at the end of the linac.

RF structure RMS misalignments with respect to the girder axis were $15 \mu\text{m}$ and $33 \mu\text{rad}$; RMS girder misalignments with respect to the survey line were $50 \mu\text{m}$. Quad misalignments with respect to the survey line were $50 \mu\text{m}$.

The RF structure BPMs were assumed to have a resolution of $5.0 \mu\text{m}$. Two sets of quadrupole BPM resolutions were used: $0.3 \mu\text{m}$ and $1.0 \mu\text{m}$. The former value represents the design pulse-by-pulse resolution of the Q-BPMs; the latter figure includes expected errors in the beamline model and other effects which can degrade the "effective" resolution in applications that require fitting trajectories to a model.

Ten seeds were used in each simulation. In the case of $0.3 \mu\text{m}$ BPM resolution, the end-linac emittance dilution after DFS averaged 20%. In the case of $1.0 \mu\text{m}$ BPM resolution, the emittance dilution after DFS averaged 80%, with results ranging from 43% to 205%. Since the linac's total emittance dilution budget is 40%, the $0.3 \mu\text{m}$ BPM resolution could be acceptable but $1.0 \mu\text{m}$ "effective" resolution would not be acceptable.

3 Dispersion Bumps

An additional tool for emittance control in the NLC main linac is direct use of closed orbit bumps to generate dispersion; this dispersion can then be used to cancel other sources of dispersion in the same phase. Such techniques were used widely in the SLC.

The number of bumps required for the linac is a function of the beam energy spread and the optics. Because of the chromaticity of the linac optics, particles with different energies will oscillate with slightly different frequencies; at some point, particles in the bunch which initially were oscillating coherently will be 90 degrees out of phase with one another; at this point, dispersion bumps cannot eliminate the emittance growth due to the oscillations. The phase advance per cell in the NLC main linac is typically around 90 degrees, and therefore the chromatic phase advance

$\xi \equiv d\nu/d\delta$ is around $-1/\pi$. The condition for decoherence within a bunch can be approximately written:

$$N_{\text{cell}}\sigma_{\delta}\xi \approx -0.25. \quad (1)$$

An RMS energy spread of 1% implies that $\sigma_{\delta}\xi \approx 0.003$. The decoherence condition then implies that a pair of bumps every 100 cells, or 2 per 200 quads, should be sufficient. In the interest of conservatism, we used a total of 14 bumps: one pair at injection and 6 pairs spaced approximately equally in betatron phase advance down the linac to the end of the line.

Each bump was implemented as a “three-bump,” in which one magnet introduced the bump and two magnets downstream closed it again. In the interest of flexibility, the bumps were modelled with quadrupoles on ultra-high-resolution movers; in real life, weak dipole correctors would be used. In order to limit the impact of wakefields, after the bump quads were set to their new positions the intervening RF girders were realigned to the beam trajectory.

Each bump was tuned against the beam sizes at the nearest downstream diagnostic region. The tuning algorithm was designed to accurately mimic the actual tuning experience of the SLC:

- For each orbit bump knob, the downstream beam size monitor with the best response was empirically determined by tuning the bump to a large value and observing the beam size at each monitor.
- During tuning, each knob in turn would be scanned through 5 values, and the beam size at the optimal downstream monitor would be measured at each value.
- The curve of beam-size squared versus knob value was fitted to a parabola; the extremum of the parabola corresponds to the optimal knob setting.
- The bump would be set to its fitted optimum, and the next bump in the series would be tuned.

For the purposes of simulation, the beam size measurements were assumed to be perfect – without resolution limits.

3.1 Results of DFS + Bump Tuning

The total linac tuning package of DFS followed by emittance bumps was simulated for 10 linac seeds; as with the DFS alone, two BPM resolutions were used, but all other conditions were held constant as described above. In the simulations with $0.3 \mu\text{m}$ BPM resolution, the post-bump end-linac emittance growth averaged 8%, while with $1.0 \mu\text{m}$ BPM resolution the post-bump end-linac emittance growth averaged 30%, with values ranging from 16% to 51%. In both cases, the emittance dilution is reduced by roughly a factor of 3, and in both cases the dilution is acceptably below the linac emittance budget, although the case with poor effective BPM resolution is marginal.

4 Conclusions

Dispersion Free Steering and global dispersion bumps are techniques to reduce the emittance dilution in the NLC main linac which do not require that the quad-to-BPM offsets be known with high accuracy. Simulations indicate that DFS alone will reduce the emittance dilution in the linac to acceptable levels if the expected BPM resolution is achieved. If the expected resolution is not achieved, or if other effects dilute the “effective resolution” of the BPMs, global dispersion bumps can be used as an “afterburner” to DFS to further reduce the emittance at the end of the linac.

5 Acknowledgements

The commentary and ideas of J. Galayda, R.K. Jobe, T. Raubenheimer, A. Ringwall, M. Ross, C. Spencer, and J. Tanabe were invaluable in the preparation of this Note.

A Quad Variation and its Discontents

The main linac quadrupole architecture has been the subject of considerable thought. At this time, there are three candidate designs for the main linac quad: a conventional electromagnet with a dedicated power supply; an iron-core magnet excited by permanent magnet blocks which can be adjusted via movable tuners in each pole-piece; an iron-core magnet excited by permanent magnet blocks which can be adjusted by counter-rotating longitudinal segments of the magnet (thus changing the normal-quad component while, to first-order, keeping the skew-quad component at zero). All three of these magnets are subject to systematic errors which can influence the quad-variation procedure for measuring the beam-to-quad offset.

The most prominent effect which can degrade the accuracy of the quad-variation technique is movement of the magnetic center during the quad-variation process. This effect has been studied in some detail [3]. If the quadrupole strength is varied in steps of dK_q , and on step j the center has shifted by an amount x_j from its position at nominal strength, then the systematic error in the fitted center position is given by:

$$\Delta x = \frac{\sum_j j^2 x_j}{\sum_j j^2} + \frac{K_q}{dK_q} \frac{\sum_j j x_j}{\sum_j j^2}. \quad (2)$$

The second term on the LHS of Equation 2 is quite disturbing, because it indicates that RMS center motions x_j are amplified by a factor K_q/dK_q into the fit error: for a quadrupole change of 20%, a tolerance of 1 μm in the fit error translates to a tolerance of about 0.2 μm in the center motion during shunting. Such a tight tolerance would be quite difficult to demonstrate, much less achieve. In a conventional electromagnet, there are a number of potential sources of magnetic-center motion during strength variation: variable ground currents in the different poles, varying mechanical stresses, fabrication or assembly errors, thermal changes, etc.

A related systematic error can occur in a hybrid iron/PM quad with counter-rotating slices. Consider a single longitudinal segment of such a magnet, and assume that the axis of rotation is not the same as the magnetic neutral axis. Let the horizontal distance between the magnetic neutral axis and the rotation axis be Δx_{rot} . Let us assume that the magnet is originally at full strength (zero rotation angle), and is reduced by an amount dK_q in a single step to measure the beam-to-quad offset. Since the normal-quad strength of the segment is given by $K_q = K_q(\text{max}) \cos 2\theta$, and the change in the vertical position of the neutral axis is given by $\Delta y = \Delta x_{\text{rot}} \sin \theta$, Equation 2 can be used to estimate the systematic fit error caused by such a defect:

$$\Delta y_{\text{fit}} \approx \Delta x_{\text{rot}} \left(\theta + \frac{1}{2|\theta|} \right). \quad (3)$$

Note that in Equation 3 we have replaced sine and cosine expressions with their small-angle approximations. Equation 3 implies that the error in the rotation axis positioning must be at least several times better than the tolerance on the BBA fit's systematic error.

A hybrid iron/PM quad with rotating tuners is subject to errors which are similar to those described above. In addition, such magnets can suffer from construction defects which introduce a shift in the position of the neutral axis that is undetectable by the quad-variation technique.

Consider a hybrid quadrupole which is mechanically and magnetically perfect: in such a magnet, the neutral axis of the magnet is identical to its mechanical symmetry axis, and the rotating tuners are perfectly balanced such that when they rotate the quad strength is varied without introducing dipole or other fields. In such a case, a beam on the neutral axis experiences no deflection when the quad is at its nominal strength and it experiences no deflection when the quad is varied in strength.

Now consider the quad in Figure 3. This is a mechanically- and magnetically-perfect quad, except that one of the energizing blocks of PM material is removed from between poles 1 and 2. We can analyze this situation using the formalism of Halbach [4] for magnet errors. Pole 1, which is “positive,” experiences a fractional change in excitation $-\epsilon_b$, while pole 2 experiences a fractional change ϵ_b , where $\epsilon_b > 0$. This introduces a dipole field which can be compared to the quadrupole field at the pole-tip:

$$\frac{H_1^*}{H_2^*} = \frac{1}{2} \Delta C_1 \exp(-i\theta) = \sum_{\text{poles}} 0.199i\epsilon_{\text{pole}} \exp(-i\theta), \quad (4)$$

where θ is the pole-angle with respect to the horizontal axis (in this case, $\pi/4, 3\pi/4, 5\pi/4, 7\pi/4$). In this case the sine-like terms in the complex exponentials cancel, and the ratio of H_1^*/H_2^* is given by $-0.199\sqrt{2}\epsilon_b$. If the tuner elements are now rotated to increase the strength of the quad by a fraction ϵ_q , the ratio of the dipole field to the *design* quad field becomes:

$$\begin{aligned} \frac{H_1^*}{H_2^*} &= \sum_{\text{poles}} 0.199i\epsilon_{\text{pole}} \exp(-i\theta) \\ &= 0.199i[(-\epsilon_b + \epsilon_q) \exp(-i\pi/4) + (\epsilon_b - \epsilon_q) \exp(-3i\pi/4) \\ &\quad + (\epsilon_q) \exp(-5i\pi/4) + (-\epsilon_q) \exp(-7i\pi/4)]. \end{aligned} \quad (5)$$

In the expression above, the terms dependent on ϵ_q cancel; the remaining terms are identical to those for the quad at normal strength. This means that the dipole field at the quad’s mechanical axis is nonzero but does not vary as the quad strength varies. Thus, the quad variation technique would not reveal in this case that the beam was not positioned on the magnetic neutral axis of the magnet.

The discussion above suggests that conventional electromagnets, tuner-type adjustable hybrids, and rotating-segment adjustable hybrids are all unlikely to meet the tolerances required if we wish to make the quad-variation technique our sole means of measuring the quad-to-BPM offset to the required accuracy.

References

- [1] T. Raubenheimer and R.D. Ruth, “A Dispersion Free Trajectory Technique for Linear Colliders,” *Nuclear Instruments and Methods* **A302**:191-208 (1991).
- [2] Snowmass 2001 Report, Linear Collider Working Group M3 Summary, in preparation.
- [3] P. Tenenbaum and T.O. Raubenheimer, “Resolution and Systematic Limitations in Beam-Based Alignment,” *Physical Review Special Topics – Accelerators and Beams* **3**:052801 (2000).
- [4] K. Halbach, “First Order Perturbation Effects in Iron-Dominated Two-Dimensional Symmetrical Multipoles,” *Nuclear Instruments and Methods* **74**:147-164 (1969).
- [5] R. Assmann *et al*, “LIAR, A Computer Program for the Modeling and Simulation of High Performance Linacs,” SLAC/AP-103 (1997).

- [6] P. Tenenbaum, "Beam Based Alignment of the NLC Main Linac, Part One: Single Bunch Comparative Study of Three Algorithms," LCC-Note-0013 (1999).

Figure 1: NLC main linac equipment related to beam-based alignment and steering.

Figure 2: Principle of quadrupole-variation technique of beam-based alignment.

Figure 3: Mechanically perfect hybrid quadrupole with one weak block of PM material. The weak block will cause a shift in the magnetic center, relative to the mechanical center, which cannot be detected by using the tuners to vary the quad strength.