Main Linac Single Bunch Emittance Preservation in the NLC and USColdLC Configurations

P. Tenenbaum

LCC-Note-0137
May 6, 2004

Abstract

We consider the emittance preservation issues in two possible linear collider main linac configurations: The Next Linear Collider (NLC), using 11.424 GHz copper accelerating structures at a loaded gradient of 52 MeV/meter, and USColdLC, using 1.3 GHz superconducting structures at a loaded gradient of 28 MeV/meter. Emittance dilution sources considered include misalignments of individual elements and groups of elements (girders or cryomodules), magnet strength errors, magnet rotation errors, and BPM offsets and resolution limits.

1 Introduction

A key operational issue in any future linear collider will be the preservation of the very small beam emittances during passage through the main linear accelerator. Sources of emittance dilution will include familiar sources such as quadrupole rotations and offsets, less familiar sources such as transverse wakefields from the accelerating structures, and rather esoteric effects such as RF deflections from pitched RF structures.

Over the years a variety of procedures for reducing emittance growth in linacs have been devised, simulated, tested, described, discussed, and dissected. In this note we consider the application of two such techniques to two candidate designs for a future, 500 GeV CM linear collider.

2 The Accelerator Designs

2.1 USColdLC

The superconducting linear accelerator design which was considered in this study was an adaptation of the TESLA TDR design [1] to the requirements of the parameters document of the US Linear Collider Steering Group [2]. Acceleration is provided by 9-cell structures at 1.3 GHz, with a fully-loaded gradient of 28 MeV/meter.

The linac cryogenic system is divided into cryostats, with 12 structures per cryostat. Superconducting quadrupole magnets are installed in alternate cryostats (ie, 50% of all cryostats have quads, 50% do not); the magnet optics is a FODO lattice with a betatron phase advance per cell of 60° in each plane. Each quadrupole has a cavity-style BPM and a vertical corrector magnet; horizontally-focusing quads also have a nearby horizontal corrector magnet. In sum, then, the USColdLC main linac lattice, for the purposes of this study, is quite similar in overall design to the first half of the TESLA TDR main linac, except that it is longer and has a higher accelerating gradient. The total number of quadrupoles is 356; there is an extra vertical corrector at the linac injection point, and an extra BPM at the linac extraction point. The short range wakefields used in this study are based on the TDR wakefields, although recent estimates suggest that the actual wakefield of the TESLA-style structures may be somewhat weaker [3]. Injection energy was 5 GeV, extraction energy was 250 GeV. Charge in the single bunch was 2×10^{10}.

The *ab initio* installation precision for quadrupoles, BPMs, RF structures, and cryostats are shown in Table 1. These values are generally consistent with the TESLA TDR values. Noteworthy exceptions include: the structure and cryostat pitch angles – since no tolerance was listed in the TDR, we have adopted a tolerance equal to the offset tolerance divided by the length of the object in question; and the BPM resolution – the 10 μm specification in the TDR is far worse than the state of the art, and the performance improvement in going from 10 μm to 1 μm was sufficiently large that we adopted the better resolution as our standard. Note that the aperture of the cavity BPM in the TDR is about 6 times as large as the aperture of the NLC X-band cavity BPM, the USCold bunch charge is about 2.7 times as large as in the NLC; given these parameters, the expected resolution of the X-band cavity BPM (0.4 μm) would scale to 0.9 μm, quite close to the value chosen for the USCold.

Table 1: Installation precision for USColdLC Main Linac components in the vertical plane.

Error	Tolerance	With Respect To...
Quad Offset	300 μm	Cryostat
Quad Rotation	300 μrad	Cryostat
BPM Offset	300 μm	Cryostat
BPM Resolution	1 μm	N/A
Structure Offset	300 μm	Cryostat
Structure Pitch	300 μrad	Cryostat
Cryostat Offset	200 μm	Survey Line
Cryostat Pitch	20 μrad	Survey Line

2.2 Next Linear Collider

The normal-conducting linear accelerator design considered in this study was the 2003 configuration of the Next Linear Collider (NLC). Acceleration is provided by 60 cm structures operating at X-Band (11.424 GHz), with a loaded gradient of 52 MeV/meter.

The linac structures are installed on 2.5 meter girders, 4 structures to a girder. Quadrupoles are installed on separate supports, with 2 girders per quadrupole in the upstream end of the linac, 4 girders per quadrupole in the middle, and 6 girders per quadrupole at the end. Each quadrupole is provided with a BPM with 0.4 μm resolution at the design single bunch charge of 0.75×10^{10}, and is mounted on a magnet mover with 50 nm stepsize in both horizontal and vertical degrees of freedom; the BPM is assumed to be firmly attached to the quad, and moves with it when the quad's mover is actuated. There are 591 quadrupoles, and an equal number of BPMs; there are no correctors used for DC steering in the linac (there are several used by steering feedbacks, but those were not actuated in this study). The accelerating structures are provided with damping manifolds, and measurement of the wakefield power and phase in these manifolds allows them to function as structure-BPMs (S-BPMs). The girders are supported by remote-controlled movers with x, y, yaw, and pitch degrees of freedom. The short-range wakefields are those for the present structure design (H60VG3S17), based on the dimensions of the structure's cells and Bane's method for estimation of short-range wakes [4].

The *ab initio* installation precision for quadrupoles, BPMs, structures, and girders are shown in Table 2. A few of these figures are worth noting. First, the BPM-to-quad offset of 5 μm is not actually "*ab initio*", but rather the residual after quadrupole shunting (see section 3.1). The S-BPM resolution is actually thought to be about 1 μm, but in this case a poor resolution was used so that the resolution, averaged over the girder, will achieve the target accuracy limit in girder positioning (3.2 μm); the target was achieved in this way because the simulation program used in the preparation of this Note permits limited resolution of S-BPMs, but not fixed offsets in their readings.

Table 2: Installation precision for NLC Main Linac components in the vertical plane.

Error	Tolerance	With Respect To...
Quad Offset	50 μm	Survey Line
Quad Strength	0.25%	Design
Quad Rotation	300 μrad	Survey Line
BPM Offset	5 μm	Quad Center
BPM Resolution	0.4 μm	N/A
Structure Offset	25 μm	Girder
Structure Pitch	33 μrad	Girder
S-BPM Resolution	8.5 μm	N/A
Girder Offset	50 μm	Survey Line
Girder Pitch	15 μrad	Survey Line

3 Alignment and Steering Algorithms

3.1 Quadrupole Shunting

Quadrupole shunting is an extremely familiar and standard technique of accelerator alignment. In this procedure, the strength of a single quadrupole is varied by a substantial fraction (20% or even 50% are not uncommon) and the resulting change in beam orbit is measured on downstream BPMs. This allows the beam-to-quad offset to be estimated. If the quad contains a BPM (which all quads in this study do), then the offset of the BPM electrical center to the quad magnetic center can be estimated as well, and appropriate software changes can be implemented to give a zero reading when the beam is passing through the quad center.

The actual shunting procedure is excruciatingly dull and was not simulated in this study. Rather, the accuracy limit of the technique from quad-center motion during shunting was estimated analytically [5], and this limit was assumed to define the RMS BPM-to-quad offset after quad shunting. In the case of the NLC, studies of a prototype quadrupole has shown a total center variation of under 1 μm for a change in integrated strength of 20% [6]. This level of stability would lead to a systematic error in the quad shunting fit of 4 μm. At this time the ensemble performance of the NLC quads cannot be known (given only 1 prototype has been built so far). If the RMS (over the ensemble of quads) of the total center variation is 1 μm (ie, 68% of all quads have a total center shift of less than 1 μm when shunted 20%, but 32% of quads have more), then the RMS BPM-to-quad offset would be expected to be about 4 μm. If, on the other hand, the prototype represents a worst case, then it might be more sensible to model the distribution of quad stabilities with a square distribution, with 1 μm per 20% strength change as the edge of the distirbution. In this case the RMS BPM-to-quad offset would be reduced to 1.7 μm. In this study we assumed that BPM-to-quad offsets after shunting are Gaussian-distributed, with an RMS of 5.0 μm, allowing an extra level of conservatism over either of these expectations.

In the case of the USColdLC, we did not assume that all of the quads would be shunted. This is due to two factors. First, the USColdLC quads are superconducting, and thus a shunting procedure might take a very long time to implement in the real machine. Second, there is no experimental basis for estimating the stability of the magnetic center as a function of excitation current in superconducting magnets; therefore we could not estimate the residual BPM-to-quad offsets in the USColdLC after shunting. The one place in the USColdLC in which we assumed that the BPM-to-quad offsets would be measured and corrected with greater accuracy is in the first 7 quads in the linac. The emittance growth in the linac is rather sensitive to the element alignment in this region, due to the large energy spread and low beam energy; also, because these magnets are the furthest upstream, they cannot be aligned via dispersion-free steering (see section 3.5). For this reason, it was assumed that the first 7 BPMs in USColdLC would be aligned to their neighboring quads, and that the accuracy limit would be comparable to the limit in the NLC quads, but scaled by the ratio of their apertures. This led to a 30 μm RMS BPM-to-quad offset in the first 7 quads of the USColdLC linac.

3.2 USColdLC Launch Steering

Once the BPM-to-quad offsets in the first 7 quads of the USColdLC have been eliminated, as described above, this region (henceforth known as the launch region) is steered as follows: first, all RF structures in the launch region are switched off to eliminate RF kicks from pitched structures or cryostats; the beam is then transported through the launch, and the BPM readings are extracted; this information is used to estimate the quad offsets with respect to the survey line, and corrector settings are then computed which ideally would result in a straight trajectory of the beam through the launch region (effectively "moving" the quad magnetic centers, via the dipole correctors, to the survey line). The orbit after setting the corrector magnets constitutes a reference or "gold" orbit for the launch.

Once the gold orbit for the launch is determined, the launch region RF units are restored and the orbit is re-steered to the gold orbit. In this way the effect of RF kicks in the launch region is cancelled.

3.3 Steering to BPM Centers

Both the USColdLC and the NLC linacs are steered to minimize RMS BPM readings, but using quite different techniques.

In the case of the USColdLC, the BPM transverse positions are fixed. It is therefore sufficient to use the correctors to steer the beam to the center of each BPM. This is done in segments, with approximately 50 BPMs and correctors per segment, thus seven segments in all. The BPMs in the launch region are not steered. Each segment is steered twice before the next segment is steered.

In the case of the NLC, the BPMs move with the quadrupoles, so some additional refinements are necessary. Specifically, the steering algorithm simultaneously seeks to constrain the BPM readings after steering and the RMS change in quadrupole position (recall that the NLC linac has no correctors, and instead the magnet movers are used to steer the beam). The ratio of the BPM constraint to the motion constraint is 1:10, since the expected BPM-to-quad offset is 5 μm and the anticipated quad offset from the survey line is 50 μm. Steering is done in segments, with approximately 50 quads per segment; also, the overlap between segments is 50% (ie, the quad in the center of segment 1 is the quad at the beginning of segment 2), so the total number of segments is 24. After a segment's quads are moved, RF girder alignment of the RF structures in that segment is performed (see below). Each segment is steered 3 times before the next segment is steered.

3.4 RF Girder Alignment

After the quads in an NLC linac alignment segment have converged on their new positions, the readings of the S-BPMs on all RF structures in the segment are determined. For each girder, the change in position and angle needed to null the average offset and average slope of the BPMs is calculated, and the girder is moved by that amount. All girders in a segment are aligned simultaneously, not sequentially; it is assumed that the change in position at one girder due to motions of an upstream girder are negligible.

3.5 Dispersion Free Steering

Both the USColdLC and the NLC linacs make use of a steering algorithm designed to minimize dispersion while simultaneously constraining the correction to eliminate long-wavelength misalignments [7]. The algorithms have certain features in common, and certain implementation-dependent features.

Both the USColdLC and the NLC divide the linac into alignment segments. In each segment two orbits are measured using the BPMs: the first orbit is under nominal conditions, while the second orbit is performed with some RF stations upstream of the segment switched off. No RF stations were switched off within the segment which was being DF-steered: this allowed the technique to measure and correct dispersion in a given segment while the segment was in its operating condition (ie, all quad strengths and energy gains equal to nominal, only the incoming beam is changed), and in particular allowed dispersion from pitched RF structures to be included in the correction; the disadvantage of this choice was that the acceleration in the segment would adiabatically damp the fractional energy change along a segment, so that the resolution of the dispersion in downstream BPMs is reduced compared to upstream ones. The maximum energy variation for a given segment is 20% of the nominal beam energy at the upstream end of the segment, or 18 GeV, whichever is smaller; the limits are designed to simultaneously prevent the beam energy change from exceeding the bandwidth of the linac optics, and to prevent the beam energy change from exceeding the bandwidth of the post-linac, pre-BDS extraction system (assumed to be about 10% of final beam energy, or 25 GeV). In both cases a few BPMs upstream of the first quad to be fitted are included in the measurement, so that any change in the incoming orbit can be separately fitted and subtracted from the dispersion effect in the alignment segment. In both cases, the first few quads and BPMs cannot be aligned by DFS because it is impossible to vary the energy sufficiently upstream of a certain point in the linac.

In the USColdLC algorithm, the correction is applied by dipole steering magnets. The correction is weighted to simultaneously minimize the measured dispersion and the RMS value of the BPM readings; the weight ratio of the two constraints is about $\sqrt{2}$: 300, the ratio of the BPM noise (remember that 2 orbits are measured, with 1 μm noise on each) to the expected offset of the BPM with respect to the cryomodule. Three BPMs upstream of each segment are used for fitting the incoming trajectory. When the alignment of the first segment is performed, all RF upstream of the fourth quad is switched off, and the alignment segment begins at the seventh quad; in this case, the energy variation exceeds that specified by the 20%/18 GeV rule. Each alignment segment contains about 40 quads, correctors, and BPMs; the segments overlap by 50% because this was found to be an effective combination in previous studies [8]. Each segment is DF steered twice before the next segment is steered.

In the NLC algorithm, the correction is applied by changing the vertical positions of the quadrupole magnets, via their magnet movers. Since the BPMs move with the quads, constraints on the absolute BPM readings will generally not be effective; instead, the measured dispersion and the RMS change in magnet positions are simultaneously constrained, with relative weights of

0.57 : 5, this being the ratio of the BPM noise in the dispersion measurement to the expected quad motion requirement. After each segment's quads are moved, the segment's RF girders are aligned. Each segment contains about 50 quads, and the first segment begins at quadrupole number eight in the linac. Two upstream BPMs are used per segment to extract incoming orbit changes, and each segment is aligned 4 times before the next segment is aligned.

4 Results of Simulations

For each of the two designs considered here, 100 sets of random misalignments and errors were applied to each, and the tuning algorithms as described above were unleashed. This led to an ensemble of 100 misaligned, tuned machines for each design. The ensembles were analyzed statistically after minimizing the BPM readings, and again after DFS.

4.1 USColdLC Results

After steering to minimize BPM readings, the mean emittance growth of the USColdLC main linac was 461 nm. Out of the 100 seeds, 90% achieved emittance growth of under 984 nm.

The DFS algorithm was then applied to each "steered-flat" linac. This led to a mean emittance growth of 10.7 nm, and a 90% confidence level of 21.2 nm.

The emittance growth noted above is much larger than has been observed in previous studies of DFS on superconducting linacs with low emittance beams [8]. Part of the discrepancy is that the energy variation used in this study is limited to 20% of the incoming beam energy, or 18 GeV, whichever is less for a given segment, while in previous studies the variation was over 50%. In addition, in this study the RF structures in an alignment segment are not switched off; this means that, while the energy change at the first quad may be as much as 20%, it is lower in subsequent quadrupoles, with a corresponding reduction in the resolution at the downstream end of the segment (see Section 3.5 for an explanation of the reasoning behind leaving all RF structures switched on in an alignment segment) In earlier studies, the fractional energy change was a constant along the segment. This was often obtained by changing the quad strengths; such an approach was rejected here because it would lead to motions of the quad center which would confuse the algorithm. This study also includes quadrupole rotations, which is believed to be responsible for about 2 nm of the mean emittance growth.

A third difference between this study and earlier ones is the inclusion of pitched RF structures and cryostats. A pitched RF element has some of its longitudinal electric field rotated into the beam's transverse degrees of freedom, resulting in a deflection to the beam. This deflection generates dispersion, and also results in a deflection which changes when the RF unit is switched off. In order to understand the role these deflections play in the emittance growth, we repeated the steering and alignment simulations with 100 seeds in which both girder and structure pitches were set to zero; structure offsets with respect to the girder were also zeroed for this study, but these offsets do not appear to contribute more than about 1 nm emittance growth.

In this last set of simulations, the mean emittance growth from steering flat was 432 nm, with 90% confidence level at 910 nm growth. The mean emittance growth from DFS was 7.6 nm, with 90% confidence level at 15.9 nm growth.

4.2 NLC Results

After steering to minimize BPM readings, the mean emittance growth of the NLC main linac was 14.0 nm. Out of the 100 seeds, 90% achieved emittance growth of under 25.0 nm.

The DFS algorithm was then applied to each "steered-flat" linac. This led to a mean emittance growth of 5.1 nm, and a 90% confidence level of 8.3 nm.

5 Conclusions

The principal sources of emittance dilution in the main linac of a linear collider are misalignment of the quadrupole magnets and the RF accelerating elements. We have studied the effectiveness of two of these techniques – steering to minimize BPM readings, and dispersion free steering – in the case of a superconducting, low-frequency linac and the case of a normal-conducting, high-frequency linac. The studies were not precisely matched at the algorithmic level because of significant differences in the basic designs of the two machines (presence or absence of remote-controlled translation stages in the design, presence or absence of RF structure BPMs, detailed properties of superconducting quadrupoles versus conventional iron-dominated electromagnet quadrupoles), nor was such an algorithmic match a goal of this study; rather, the goal was to understand the performance of the two algorithms when the design-specific changes to the algorithms were included.

In the case of the warm linac, quad shunting allowed the BPM-to-quad offsets to be reduced to a level of 5 μm RMS, far smaller than the hundreds of micrometers achieved in the cold linac via mechanical survey and alignment. Largely because of this difference the emittance growth in the warm linac was only 14 nm after steering to minimize BPM readings, compared to 461 nm in the cold. Similarly, after applying a DFS algorithm the mean emittance growth in the warm was 5.1 nm, compared to 10.7 nm in the cold.

One of the major contributors to the emittance growth in the cold case was RF structure and girder pitch angles; such angles are measured and corrected in the warm machine by its S-BPM system and the girder movers which support the RF units. When RF pitches were suppressed in the cold design, "steer-flat" emittance growth was reduced to 430 nm and DFS emittance growth was reduced to 7.6 nm.

It is interesting to note that sensitivity studies performed on the cold machine tentatively suggest that the DFS algorithm described here actually eliminates almost all of the dispersive emittance growth in the linac if RF element pitches are suppressed. Of the 7.6 nm remaining emittance growth, as much as 6.5 nm may be due to wakefields from the cryostat misalignments and coupling from rolled quadrupoles (4.4 nm and 2.1 nm contributions, respectively). By contrast, the 5.1 nm of emittance growth in the NLC still contains a substantial contribution of dispersion (over 2 nm), while about 3 nm of total growth arise from wakefields and quad rotations.

6 Acknowledgements

The many helpful contributions and encouragement of C. Adolphsen, R. Brinkmann, D. Burke, T.O. Raubenheimer, and V. Tsakanov are gratefully acknowledged.

References

[1] The TESLA Collaboration, *TESLA Technical Design Report* (DESY, 2001).

[2] U.S. Linear Collider Steering Group, *US Linear Collider Technology Options Study,* `http://www.slac.stanford.edu/xorg/accelops/` (2004).

[3] D. Schulte, N. Walker, "Simulations of the Static Tuning for the TESLA Linear Collider," Proceedings *PAC-2003*.

[4] K. Bane, “Short-Range Dipole Wakefields in Accelerating Structures for the NLC,” LCC-Note-0116 (2003).

[5] P. Tenenbaum and T.O. Raubenheimer, “Resolution and Systematic Limitations in Beam-Based Quadrupole Alignment,” *Physical Review Special Topics – Accelerators and Beams* (2000).

[6] C.E. Rago, C.M. Spencer, Z. Wolf, G. Yocky, “High Reliability Prototype Quadrupole for the Next Linear Collider,” *IEEE Transactions on Applied Superconductivity* (2002).

[7] T. Raubenheimer, R.D. Ruth, “A Dispersion Free Trajectory Correction Technique for Linear Colliders,” *Nuclear Instruments and Methods A* (1990).

[8] R. Brinkmann, V. Tsakanov, “Emittance Preservation in TESLA,” (2001).