

Effectiveness of Emittance Bumps in the NLC and USColdLC Main Linear Accelerators

P. TENENBAUM
LCC-NOTE-0138
May 6, 2004

Abstract

We report on a series of studies on the effectiveness of closed orbit bumps in the linacs of the NLC and the USColdLC. In the first study, emittance dilutions of a pure-wakefield or pure-dispersion character are introduced in each linac, and a set of emittance bumps is tested to determine the most effective bump location in the linac, and the net effectiveness. In the second study, a more realistic set of dilutions are introduced and the wakefield bumps used in the first study are applied.

1 Introduction

All of the designs for a linear collider main linac must contend with emittance dilutions from dispersion (due to beam-to-quad misalignments and RF structure tilts) and wakefields (due to beam-to-structure misalignments). Because the design emittances in the linear collider designs are so small (vertical normalized emittances of 20 nm are typical), it is generally understood that a local or quasi-local correction of the emittance dilutions is required. A number of such schemes have been considered, and a recent simulation study on their capabilities is available [5].

An alternate technique for reducing emittance growth in the linacs is to deliberately introduce a large offset between the beam and either the quadrupoles or the RF structures at a given location; this offset will generate dispersion and/or wakefield emittance dilutions which can be set in amplitude and betatron phase to cancel dilutions elsewhere in the linac. Such “emittance bumps” are generally considered a less desirable technique for emittance control, for several reasons:

- The bump amplitude and betatron phase can only be set by measuring the emittance on one or more profile monitors (typically laser wire scanners are assumed); the latter are at a small number of discrete locations in the linac, and therefore the resulting correction is highly nonlocal
- Changes in the betatron phase advance through the linac can cause the bump to go “out of phase” with the dilution sources and/or change the bump amplitude required; when this happens the bumps need to be retuned
- It is generally considered unlikely that a small number of bumps can effectively cancel the emittance growth in a linac which has not first been tuned via a local correction algorithm; prior to application of the local tuning algorithms the linac emittance growth is generally many times the incoming emittance.

For these reasons, it is generally expected that emittance bumps will be applied after local corrections have brought the emittance down into some sort of “capture tolerance” for the bumps.

In this Note we report on a series of studies of emittance bumps in the NLC and the USColdLC main linacs. The goal was to understand both relative effectiveness (“Is it true that bumps are more effective in a superconducting linac than in a normal-conducting one?”) and absolute effectiveness (“Just how much emittance growth can those things cancel?”).

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 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} .

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 $0.4 \mu\text{m}$ 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 the short range wakefield [4].

3 Sensitivity Study

In the first study, the initially-perfect linacs were misaligned in such a way as to generate either purely-dispersive or purely-wakefield emittance dilution. Then an appropriate set of bumps (dispersion-only or wake-only) were applied to each linac.

3.1 Generation of Dispersive Emittance Growth

The key to generating an emittance dilution which is caused solely by dispersion in the linac is to introduce a spectrum of errors which cause the beam to be misaligned with respect to the quads, yet aligned with respect to the RF structures.

In the case of the USColdLC, the dispersive emittance growth was generated by applying an RMS BPM-to-quad offset of $60\ \mu\text{m}$, and then using dipole steering correctors to steer the beam to the centers of the BPMs. Since the quads and RF structures are still perfectly aligned to the survey line, this results in the beam being misaligned by about $60\ \mu\text{m}$ with respect to the RF units and the quads; since the wakefields are weak in USColdLC, $60\ \mu\text{m}$ RMS beam-to-structure offset results in negligible wakefield emittance growth, while the combination of quad offsets and dipole deflection from the correctors results in a mean dispersive emittance growth of $20\ \text{nm}$ (averaged over 100 seeds).

In the case of the NLC, a much smaller BPM offset – $5.8\ \mu\text{m}$ rather than $60\ \mu\text{m}$ – generated $20\ \text{nm}$ growth in the normalized vertical emittance. In this case, the linac was steered flat via magnet movers rather than correctors, and the RF girders were then aligned via structure-BPM signals; this was necessary because, unlike the USColdLC, in the NLC the alignment tolerances for RF girders and quadrupoles are more nearly equal.

3.2 Generation and Tuning of Dispersion Bumps

For each design (USColdLC and NLC), bumps were generated at 6 locations: main linac injection, 50 GeV, 100 GeV, 150 GeV, 200 GeV, and 250 GeV beam energies. At each location two bumps were generated, separated by approximately 90° in betatron phase. In the case of the USColdLC, the linac betatron phase advance is 60° per cell, so a closed bump could be generated by energizing a pair of correctors 3 cells apart, with appropriate scaling for the change in beam energy between the two. In the NLC main linac the phase advance per cell is not so regular, and so the bumps at injection, 50 GeV, and 100 GeV are “3-bumps,” in which 3 consecutive “D” quads are moved together to make a closed bump; for this reason, the two bumps at each of these points are not quite at opposite betatron phases. At the 150 GeV, 200 GeV, and 250 GeV points, the phase advance per cell is 90° , so pairs of quads could be moved to make a closed bump and the bumps are truly in opposite phases.

In both the USColdLC and NLC cases, at each bump amplitude the intervening RF structures are aligned to the beam with infinite precision; this is to prevent any wakefield effects from confusing the dispersion study. Note that while this is physically plausible in the NLC case, the USColdLC has no girder movers and cannot align its RF structures to the beam. This unrealistic method was adopted to minimize crosstalk between wakefield and dispersion effects in this study.

Both the NLC and the USColdLC designs call for a multiwire emittance measuring station downstream of the main linac, and the NLC has two additional stations within the linac. For this study, none of these diagnostic stations were used. The downstream stations were left out because they are technically within the beam delivery system, and their use would entail certain practical complications to the study. The intermediate NLC stations were not used because the USColdLC design has no such stations, therefore their use in the NLC would break the symmetry between the two linacs in this study. Instead the bumps were tuned by inserting wire scanners into the beamlines at appropriate locations near the end of each linac. In the case of the NLC it was possible to insert a pair of wire scanners at two consecutive “D” quads, separated by 90° in betatron phase. Because the USColdLC makes use of a 60° lattice, a somewhat less physically plausible method was used: a wire scanner is placed at a “D” quad, and at an “F” quad 90° away in betatron phase. Although such an arrangement would not be chosen in real life because the resolution of the small vertical

emittance would suffer, in simulation we can do this rather than go through the agony of designing and matching a proper emittance station onto the end of the linac. For each bump a wire scanner target was selected by determining which wire saw a larger fractional change for a given bump; the scan range was then set by the bump value needed to quadruple the extracted emittance, given an otherwise-perfect linac. This is the procedure that was used to tune emittances in the SLAC linac for SLC and FFTB. In all cases a wire scanner resolution of 2% of the measured beam size was assumed.

3.3 Results of Dispersion Study

Figure 1 shows the mean emittance growth in the USColdLC, averaged over 100 seeds, as a function of the bumps used to minimize the emittance growth from dispersion. In the absence of all bumps, a mean emittance growth of 22 nm is observed (somewhat larger than the target value of 20 nm); the injection bumps are the most effective and can cancel almost all of the emittance growth in almost every case; the remaining bumps become monotonically less effective, with the 250 GeV bumps leaving about 9 nm of emittance growth. In the USColdLC design the energy spread is much larger at injection than it is at any other point in the linac, so it is not too surprising that the injection bumps are the most effective – they are closer to the likely sources of emittance dilution than the downstream bumps are.

Figure 2 shows the mean emittance growth versus bump location in the NLC. The most effective bumps, at 150 GeV, only remove about half the emittance growth. A somewhat more effective approach is to use the wakefield bumps at 150 GeV in combination with another set of bumps. It turns out that the bumps at 50 GeV are the most effective for this; when both 50 GeV and 150 GeV bumps are used, the emittance growth can be reduced to under 5 nm, as shown in Figure 3.

3.4 Generation of Wakefield Emittance Growth

Emittance growth which is dominated by wakefields is generated by a process complementary to the dispersive process: the beam needs to be misaligned in the structures but not the quadrupoles.

In the case of the NLC, this is accomplished by misaligning the RF girders with respect to the nominal survey line in an otherwise-perfect linac; an RMS girder misalignment of 12 μm generates a mean emittance growth of 20 nm, which is entirely due to wakefields.

The USColdLC case is somewhat more complicated: when the RF units are misaligned, their fringe-field focusing deflects the beam, resulting in beam-to-quad offsets; because the USColdLC is much more sensitive to beam-to-quad than it is to beam-to-RF offsets, simply misaligning the RF units actually results in more dispersive emittance growth than wakefield growth. This effect is counteracted by first misaligning the RF units and then steering to zero on the BPMs (which in this case are perfectly aligned with the quad centers). An RMS misalignment of the RF girders of 580 μm , while steering the beam to zero in the BPMs, leads to a mean emittance growth of 20 nm which is entirely due to wakefields.

3.5 Generation and Tuning of Wakefield Bumps

In general, the wakefield bump methodology is quite similar to the dispersion bump methodology: the bumps and wire scanners are at the same locations, the selection of a target wire and a range for each bump is the same. In this case, a “bump” consists of a single girder of RF units which are displaced from their nominal positions. This is a physically plausible approach in the NLC, while in the USCold it was done this way for simplicity and to ensure a fair comparison.

As an additional constraint, it was necessary for each bump in the USColdLC to include a dipole corrector. This corrector would be used to take out the steering from the RF-focusing of the offset structures; otherwise, each “wake bump” would generate more dispersive emittance growth than wakefield growth.

3.6 Results of Wakefield Study

Figure 4 shows the mean emittance growth, averaged over 100 seeds, as a function of the wakefield bumps used to minimize the emittance in the USColdLC. With the exception of the bumps at injection, all of the wakefield emittance bumps can reduce the mean emittance growth from 20 nm to about 3.5 nm.

Figure 5 shows the mean emittance growth, averaged over 100 seeds, for the NLC wakefield bump study. The remaining emittance growth is about 6.5 nm in the best case (150 GeV bumps), and is not as uniform as a function of bump location as what was observed for the USColdLC.

4 More Realistic Tuning Study

In the previous study, linacs were known to have either purely-dispersive emittance growth or purely wakefield emittance growth, and bumps appropriate to that condition were used for tuning. In this study a more complete spectrum of dilutions was applied, followed by local and semi-local correction schemes; both wakefield and dispersion bumps were then applied to the resulting, “tuned” systems.

4.1 Generating the Tuned Linacs

The tuned linacs were generated using the errors and tuning procedures described in [5]; in the case of USColdLC, the girder pitch angles, structure pitch angles, and structure offsets with respect to the girder were set to zero, rather than the nominal values. This resulted in a mean emittance growth of 5.1 nm for NLC and 7.4 nm for USColdLC.

4.2 Applying the Bumps

The USColdLC was tuned by applying the dispersion bumps at injection and the wake bumps at 50 GeV, in that order. The ranges and targets of the bumps were the same as in the sensitivity study.

The NLC was tuned by applying the dispersion bumps at 50 and 150 GeV, followed by the wake bumps at 200 GeV. The ranges and targets of the bumps were the same as in the sensitivity study.

4.3 Results of Tuning

In the case of the USColdLC, the tuning bumps reduce the mean emittance growth from 7.6 nm to 4.5 nm; the 90% CL is reduced from 15.9 nm to 8.0 nm. The distribution of emittance growth over the 100 seeds is shown in Figure 6.

In the case of the NLC, the tuning bumps reduce the mean emittance growth from 5.1 nm to 3.8 nm; the 90% CL is reduced from 8.4 nm to 5.6 nm. The distribution of emittance growth over the 100 seeds is shown in Figure 7.

4.4 Doubled Errors

Another case worth considering is one in which the errors listed above are doubled. This corresponds to a case in which the emittances are larger because tolerances are not met, and therefore the effectiveness of the bumps becomes more critical.

For the USColdLC with doubled errors, the mean emittance growth after local tuning is 27.9 nm, which is reduced by bumps to 16.9 nm; the 90% CL is 56.2 and 29.9 nm, respectively. The distribution of emittance growth over the 100 seeds is shown in Figure 8.

For the NLC with doubled errors, the mean emittance growth after local tuning is 19.8 nm, which is reduced by bumps to 15.0 nm; the 90% CL is 31.0 and 23.3 nm, respectively. The distribution of emittance growth over the 100 seeds is shown in Figure 9.

5 Conclusions

Two studies on the effectiveness of emittance bumps in the USColdLC and NLC main linacs were performed. In the first study, pure emittance growths were carefully generated and appropriate bumps (either dispersion or wake type) were tuned; in the second study, a more realistic spectrum of errors and suite of local tuning algorithms was applied and the bumps were used to try and “clean up” the remaining emittance growth. In all cases the tuning was performed in an operationally-realistic manner, by varying the bump amplitude and seeking to minimize the projected beam size on a downstream profile monitor with limited resolution.

In the case of the USColdLC the bumps were found to be very effective against purely dispersive or purely wakefield emittance growth: in the former case a single pair of bumps could essentially eliminate all growth, while in the latter case a single pair of bumps could typically eliminate 5/6 of all growth. The NLC bumps were less effective: a single pair of dispersion bumps could remove about half the dispersive emittance growth, while a single pair of wake bumps could eliminate about 2/3 of all wakefield emittance growth. Using two pairs of dispersion bumps allowed 75% of all dispersive emittance growth to be cancelled.

In a study more representative of the real world, neither machine saw as much effectiveness of its emittance bumps. In the USColdLC, 60% of the emittance growth in the linac remained after applying one pair of dispersion bumps and one pair of wake bumps. In the NLC, 75% of emittance growth remained after applying two sets of dispersion bumps and one set of wake bumps. These ratios of pre-bump and post-bump mean emittances were preserved when the errors in the two designs were doubled. Note that in no case was even half of the linac emittance growth eliminated by the use of bumps.

Note that it is possible to imagine and simulate a system in which either more exotic instrumentation or a more byzantine tuning methodology is used to extract a more complete picture of the emittance dilution (the proportions due to dispersion versus wakefields, the betatron phase of the source, its approximate location in the beamline, etc). Practically speaking, minimizing the projected beam size as a function of bump amplitude is a technique which has been demonstrated, and more complicated, subtle, and sophisticated systems are often difficult to implement and use in an operating accelerator.

6 Acknowledgements

Without the strong encouragement of T.O. Raubenheimer, this Note would not have been possible.

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, "Main Linac Single Bunch Emittance Preservation in the NLC and USColdLC Configurations," LCC-Note-0137 (2004).




Figure 1: Mean emittance growth in USColdLC linac with dispersive errors and dispersion bumps.




Figure 2: Mean emittance growth in NLC linac with dispersive errors and dispersion bumps.




Figure 3: Mean emittance growth in NLC linac with dispersive errors and dispersion bumps; in this case the 150 GeV bumps are combined with one other pair of bumps.




Figure 4: Mean emittance growth in USColdLC linac with wakefield errors and wakefield bumps.




Figure 5: Mean emittance growth in NLC linac with wakefield errors and wakefield bumps.




Figure 6: Distribution of emittance growth before and after applying bumps to the USCold linac with design errors. Mean growth is indicated by red lines, 90% CL by green lines.




Figure 7: Distribution of emittance growth before and after applying bumps to the USWarm linac with design errors. Mean growth is indicated by red lines, 90% CL by green lines.




Figure 8: Distribution of emittance growth before and after applying bumps to the USCold linac with doubled errors. Mean growth is indicated by red lines, 90% CL by green lines.




Figure 9: Distribution of emittance growth before and after applying bumps to the USWarm linac with doubled errors. Mean growth is indicated by red lines, 90% CL by green lines.