SIMULATIONS OF THE STATIC TUNING FOR THE TESLA LINEAR COLLIDER

D. Schulte, CERN, Geneva, Switzerland
N. Walker, DESY, Hamburg, Germany

Abstract

At the heart of the TESLA linear collider are the two 10 km long superconducting linacs. A linac is constructed from 858 cryomodules each containing 12 nine-cell 1.3 GHz superconducting cavities. 355 quadrupoles provide the necessary beam focusing. The advantages of low-frequency superconducting RF in terms of wakefield behaviour are well known, and the TESLA alignment tolerances are relatively loose. However, the effects of cavity tilts and their impact of the linac beam-based alignment algorithms have until recently not been fully investigated. In addition, the strong sensitivity to correlated emittance growth due to the high beam-beam disruption parameter makes it desirable to control the linac emittance down to a few percent. In this report we discuss various static tuning algorithms and present new simulation results. Discussions of the relative merits and applicability of the methods is also discussed.

INTRODUCTION

TESLA is a proposal for a e^+e^- linear collider with an initial centre of mass energy range up to 500 GeV, with a possible upgrade to 800 GeV, and a design luminosity of 3×10^{34} cm^{-2}s^{-1} (at 500 GeV) [1]. The linac technology is based on 1.3 GHz superconducting RF. Each 10 km linac is constructed from 858 cryomodules containing 12 nine-cell niobium cavities. The focusing (beam transport) is provide by a regular 60º FODO lattice with 355 quadrupoles. As within any of the proposed future linear colliders, a major challenge is the preservation of the tiny vertical emittance ($\gamma\epsilon_y$ ~ 2×10^{-8} m for TESLA) from the damping ring to the interaction point (IP), and in particular in the main linac. Table 1 lists the *luminosity* tolerances and the expected achievable *installation* tolerances. The luminosity tolerances indicate the level of alignment precision required to achieve (on average) the goal luminosity. The installation tolerances represent those believed achievable in the real machine. The relatively low frequency superconducting RF has the advantage of relatively weak wakefields, which is reflected in the corresponding cavity luminosity tolerance which can be achieved during construction and installation. The luminosity tolerances for the cavities are based on the TDR transverse wakefield [1] which has since be superseded by a more accurate calculation [2], indicating a transverse wakefield 30% weaker than previously thought; this should lead to a reduction of the wakefield induced* emittance growth by 60%.

Table 1: Component tolerances for the TESLA main linac. The *luminosity tolerances* are those random RMS values which result (on average) in the budgeted emittance growth after one-to-one linac steering. The *installation tolerances* are those accuracies to which the components are expected to be installed; these numbers are used in the simulations. Units are µm and µrad.

		reference axis	RMS tolerance
Luminosity tolerance			
BPM	offset	ref. line	25
cavities	offset	ref. line	500
	tilt	ref. line	300
Installation tolerance			
quadrupole	offset	cryomodule	300
structure	offset	cryomodule	300
	tilt	cryomodule	300
BPM	offset	cryomodule	200
	resolution	-	10
cryomodule	offsets	ref. line	200

Chromatic (dispersive) effects from offset quadrupoles and orbit deviation are the dominant source of emittance growth in TESLA; this is particularly true at the entrance of the linac where there is a large uncorrelated energy spread of ~3% resulting from the bunch compressor. The strength of the effect is reflected in the BPM offset luminosity tolerance of 25 µm (table 1); this should be compared to the expected installation tolerance of ~360 µm†.

In the past, so-called dispersion free steering (DFS) [4] has been extensively studied as the primary beam-based alignment method. However, in the results reported in [1] and [2], the effects of the cavity tilts were not included, as well as other 'realistic' errors such as beam jitter. A more complete study of the TESLA linac performance has since been made [5], where a complete set of realistic errors (including the cavity tilts) have been modelled using the LIAR code [6]: the average emittance growth was ~140% after DFS, which is several factors more than the previously reported results.

In this paper we will review new simulation results of beam-based alignment for the TESLA linac, including the new transverse wakefield potential, with specific emphasis on the ballistic alignment technique [7], where the average emittance growth is observed to be ~30%.

* It is actually misleading to talk about emittance growth from wakefields and dispersion separately, since the two are coupled.

† $\sqrt{200^2+300^2} \approx 360\mu$m

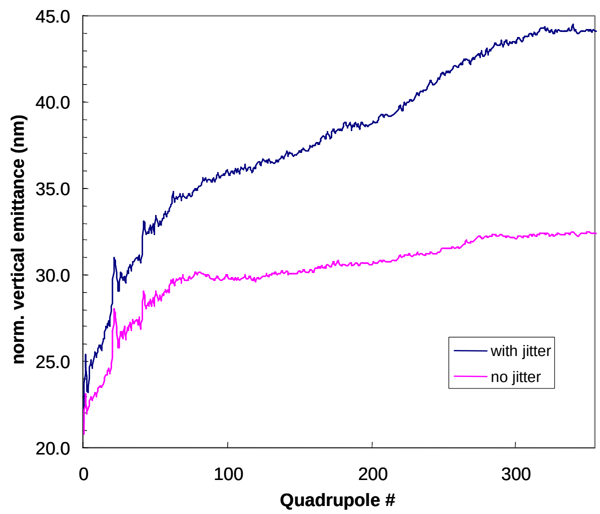

Figure 1: Results of simulating DFS on the TESLA linac (average over 100 seeds). The black line shows results with a stable; the magenta line shows the effects of ~1σ_y beam jitter. In both cases BPM noise was included (10 μm RMS).

BEAM-BASED ALIGNMENT METHODS

Dispersion Free Steering (DFS)

As its name implies, DFS attempts to find a trajectory which minimises the dispersion generated. This is achieved by finding an orbit which minimises the difference orbit when either the beam energy or the lattice strength is modified. In the absence of cavity misalignments or other 'external fields' which can kick the orbit between the quadrupoles, there is no conceptual difference between changing the quadrupole strengths and modifying the beam energy. However, when we consider the effects of cavity steering, then it is important to use the energy method in order to get a 'true' dispersion measurement (merely changing the quadrupole strengths does not 'measure' the dispersive kicks from the cavities).

The method adopted in [5] divided the linac into a number of overlapping sections or *bins* (typically 20 quadrupoles with an overlap of 10). DFS is then applied to each bin in turn. For each bin, the upstream energy was varied by ~20% and the difference orbit measured. With the additional knowledge of the optics an orbit can be found which minimises the measured difference.

DFS suffers from several problems:

- In the presence of BPM noise, the exact solutions typically result in an absolute orbit which exhibits a long wavelength bowing with very large unrealistic amplitudes. This tendency must be compensated by constraining the resulting orbit to be within the expected absolute alignment error of the BPMs.
- The upstream energy is modified by changing either phase or amplitude of the upstream cavities. In the presence of cavity alignment errors (particularly tilts), this causes a steering of the beam which – if not corrected for – will confuse the measurement of the 'dispersive' orbit.
- Similarly incoming random beam jitter must be compensated, or averaged away.

Figure 1 shows the results of applying DFS to the TESLA linac using the simulation code PLACET [7]. Specifically, the plot shows the effects of upstream beam jitter on the results (upper curve); in this case two BPMs at the entrance to each bin where used to fit out the incoming oscillation. The degree to which this can be achieved depends on the BPM resolution assumed (10 μm RMS). The results indicate that the jitter is responsible for a factor ~3/2 in emittance growth, and would appear to explain some of the discrepancy between the TDR results and those reported in [5].

Ballistic Alignment (BA)

In ballistic alignment [8], a reference line is established by first turning off all the components (quadrupoles, cavities, correctors) within a given section, and then allowing the beam to 'coast' through that section. Assuming that there are no other fields which can influence the beam, the BPM readings define a straight line (to within the BPM resolution) to which the beam is re-steered after components are restored to their original values. As with DFS the linac is divided into bins, but unlike DFS they are not overlapped. To prevent the corrected orbit 'walking away', the ballistic BPM readings are first linearly corrected to arbitrarily zero the last BPM in the current bin; this BPM acts as a pivot location or node, where the straight ballistic reference lines are allowed to have an angle. Figure 2 shows schematically a section of linac after ballistic alignment. For the case of perfect BPMs (zero resolution), the remaining dispersive kicks are simply given by the angles at the section boundaries (α_i). The effect of BPM noise is indicated by the red dotted line in figure 1.

The main advantages of BA over DFS are:

- since a single pulse can be used to determine the ballistic trajectory, the method is relatively insensitive to upstream beam jitter;
- there is no 'fit' in the process which may be numerically unstable (as is the tendency in DFS);
- notwithstanding remnant fields, the components are either off, or at their nominal settings, so the motion of the magnetic centres (for example) with excitation is not a problem.

The main disadvantage is the control of the beam during the ballistic measurement. Turning off 7 FODO cells and the associated RF causes a large β-beat, and a large energy error (ΔE ~ 7 GeV, corresponding to ~300% in the worst case at the beginning of the linac). If left uncorrected, the downstream orbit would also perform large amplitude oscillations which could damage the linac. TESLA does have the advantage of relative large apertures (70 mm diameter), and in practise, the alignment would probably start with all the RF off and the entire lattice scaled to the initial energy of 4.6 GeV. As each section is ballistically aligned (and the RF in that section turned on), the remaining downstream linac will be scaled to the new energy. In addition feedback will be used to maintain the downstream orbit. Further simulations of the control of the beam are needed to determine the maximum allowed length of ballistic section – particularly in the first linac sections where the effects are likely to be most dramatic.

Another possible issue which also affects DFS performance is that of BPM resolution. Initial beam-based alignment is likely to be at reduced beam power to protect the linac from possible damage; this means either a significant reduction in the number of bunches or a reduction in single bunch charge or (most likely) both. This may have a significant impact on the BPM performance. Reduction of the bunch charge will also affect the wakefields, although in TESLA this is probably not an issue.

In principle, the transverse wakefields will define the 'straightness' of the ballistic trajectory. In CLIC simulations [8] the ballistic process is generally iterated to converge to the final desired orbit. The current simulations suggest that no such iteration is required for TESLA, as the wakefield effects on the ballistic orbit are small.

The performance of the ballistic method depends on:

- The size of the node angles α_i, the typical size of which depend on the expected BPM offsets, and the length of the ballistic sections (L_b);
- The resolution of the BPMs.

For the current simulations, bins of 14 quadrupoles were taken (7 cells, or a total phase advance[‡] of $7\pi/3$) corresponding to a length of ~410 m in the low energy section of the linac. A BPM single-shot resolution of 10 μm was also used.

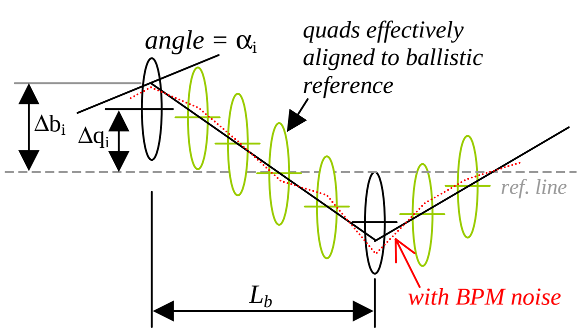

Figure 2: The result of ballistic alignment: the quadrupoles are aligned on straight line segments between bin boundaries (nodes). The RMS offset of the nodes is simply given by the RMS BPM offset with respect to the reference line (~360 μm for TESLA). The solid lines represent the best case for perfect BPMs (no noise); the red dotted line indicates the effect of finite BPM resolution.

Figure 3 shows the results of applying ballistic alignment to the TESLA linac, with all the installation

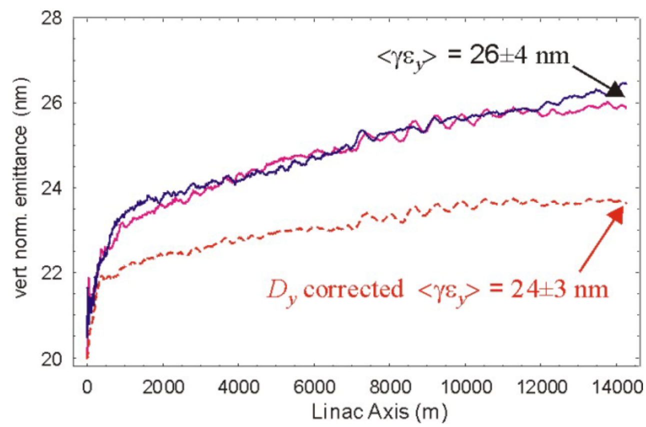

Figure 3: Results of simulations of ballistic alignment for the TESLA linac (averaged over 100 seeds). The two solid lines show independent results from the simulation codes MERLIN [9] (red) and PLACET [7] (blue). The red dotted line shows the emittance after the linear energy correlations are removed.

errors in table 1, including 10 μm RMS random BPM noise and an initial random beam jitter of ~5 μm ($\sigma_y/2$) RMS. The average emittance achieved as 26±4 nm, corresponding to an emittance growth of 30%; ~85% of the machines produced an emittance growth of 50% or less. Note that half of this emittance growth comes from the first ~750 m or ~12 GeV of the linac where the energy spread is large due to bunch compression.

A large fraction of the 'projected' emittance shown in the solid lines in figure 3 is correlated with energy. In principle, the linear correlations $\langle y\delta\rangle$ and $\langle y'\delta\rangle$ can be removed with use of orbit bumps in the linac, or the dispersion correction available in the beam delivery system (BDS). The green dotted lines in figures 3 and 4 show the emittance after this correction: the average linac exit (dispersion corrected) emittance is now ~18%. Applying two such corrections (for example, two bumps at the correct phase as suggested in [2]) may improve performance since some fraction of the correlation corrected at the first bump would otherwise filament (decohere) by the end of the linac. The study of realistic bumps has still to be made.

REFERENCES

[1] TESLA Technical Design Report, TESLA-Report 2001-13 (2001)
[2] V. Tsakanov, TESLA-99-21 (1999).
[3] I. Zagorodnov and T. Weiland, these proceedings (2003).
[4] T. Raubenheimer and R. D. Ruth, NIM **A302**:191 (1991).
[5] P. Tenenbaum, R. Brinkmann, V. Tsakanov, Proc. EPAC 02 (2002).
[6] R. Assmann *et al*, SLAC/AP-103 (1991)
[7] http://dschulte.home.cern.ch/dschulte/placet.html
[8] T. Raubenheimer and D. Schulte, PAC 99 (1999)
[9] http://www.desy.de/~merlin

‡ care must be taken to avoid multiples of π.