ILC School Erice Linac II
Lecture Notes

Peter Tenenbaum
DRAFT October 3, 2007

Contents

1 Introduction

At the end of the previous linac lecture, we had defined the basic building block of the ILC's linear accelerator, the RF unit, shown schematically in Figure 1.

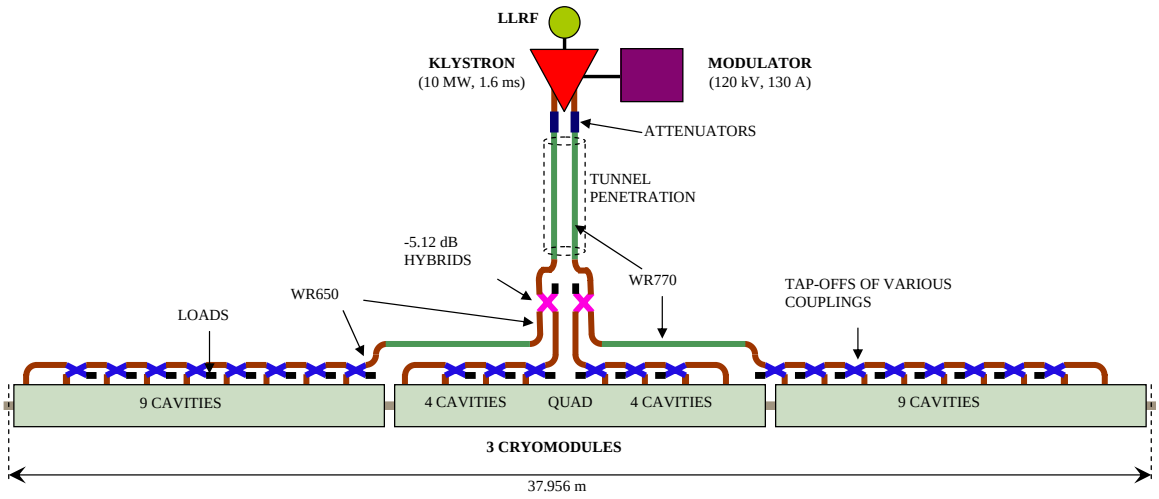

Figure 1: Schematic of an ILC RF unit.

Table 1 reviews the key parameters of the linac RF unit, and Table 2 reviews the key parameters of the beam which is injected into the linac.

Parameter	Value
# of Cavities	26
Average gradient	31.5 MV/m
Total voltage	852 MV
Peak RF Power	10 MW
RF Pulse Length	1.6 msec
RF fill time	600 μsec
Total length	38.0 m
RF-to-beam phase	5°

Table 1: Key parameters of the linac RF unit.

From the initial energy of the beam, the desired final energy of 250 GeV, and knowledge of the energy gain, phase offset, and single-bunch loading of the RF cavities, we know that the positron linac requires 278 RF units; from additional knowledge of the energy loss from the undulator in the electron linac which is used to produce positrons, we know that the electron linac requires 282 RF units.

The primary feature of the ILC main linac which distinguishes it from other linear accelerators (other than its tremendous length and vast final energy) is the requirement that the extremely small emittance of the beam injected into the linac be preserved over its full length. In order to understand the other issues which drive the design of the linac as a system, therefore, we must understand the beam dynamics effects which would tend to dilute this extremely small emittance.

Parameter	Value
# Bunches/train	2625
# Trains/second	5
Bunch spacing	369 nsec
Beam current	9.0 mA
Bunch charge	3.2 nC
Energy	15 GeV
Energy spread	1.5%
Bunch length	0.3 mm
$\gamma\epsilon_{x,y}$	< 9.0 μm (x), < 24 nm (y)

Table 2: Key parameters of the beam at entrance to the linac.

2 Transverse Wakefields

In the previous lecture on the linac, we saw that the phenomenon of beam loading could be understood in terms of the electron beam exciting modes in the accelerating cavity. In the long-term, only the modes with the highest Q values tend to survive (i.e., the fundamental accelerating mode), and therefore the effect can be modeled quantitatively with the R and Q values determined for the fundamental TM_{010} mode, while in the shorter term (such as the length of 1 bunch) a very large number of modes are present, interactions occur on a time scale which is short compared to the speed-of-light travel times about the RF structure, and a different model must be used to study beam loading.

We also saw in the previous lecture that RF cavities will support an infinite number and variety of modes. Accelerating modes, such as the TM_{010} mode (or "fundamental mode") are "monopole" modes: these modes can be excited by an external power source operating at the correct frequency, or by the passage of the beam anywhere in the cavity, and the modes contain longitudinally accelerating fields. There are also an entire spectrum of dipole modes: these modes are driven by any beam which has a dipole moment (i.e., nonzero charge and an offset from the symmetry axis of the cavity), and produce a deflecting field in the cavity. Like the longitudinal wakefields, the dipole modes produce effects on the single-bunch timescale and on the bunch-to-bunch timescale, and different methods are used to study the effects on these two very different time scales.

In this section, we will concentrate on the beam dynamics of single-bunch dipole modes, usually known as transverse wakefields. Later on we will consider the effects at the longer, inter-bunch timescale.

2.1 Single Bunch Wake Function

by analogy with the single-bunch beam loading (or longitudinal wakefield), it is possible to write down an approximate formula which estimates the deflecting voltage left behind by a point charge which is offset from the center of a cylindrically-symmetric accelerating cavity. A particle with charge q and offset x which passes through a cavity of length L at time $t = 0$ excites a time-dependent voltage $V(z) = W_\perp(z)Lqx$, where $z \equiv ct$ and

$$W_\perp(z) \approx \frac{4Zcs_0}{\pi a^4}\left[1 - \left(1 + \sqrt{\frac{z}{s_\perp}}\right)\exp\left(-\sqrt{\frac{z}{s_\perp}}\right)\right], \text{ where} \tag{1}$$

$$s_\perp \equiv 0.169 \frac{a^{1.79} g^{0.38}}{d^{1.17}},$$

where a is the aperture radius of the cavity, d is the cell length of the cells which make up the cavity, and g is the interior width of the cell (*i.e.*, the cell period minus the disc thickness) [1].

Since $W_\perp(0) = 0$, the leading particles in a bunch do not deflect themselves. This is qualitatively different behavior than what is observed for the longitudinal wakefield – we have seen that even a single electron is decelerated by its own longitudinal wake! We can also see that the slope of $W_\perp$ is nonzero at $z = 0$: $W'_\perp(0) \equiv \partial W_\perp / \partial z|_{z=0} \approx 2Zc/\pi a^4$. In other words, the deflecting field "takes off" much more rapidly as the size of the hole in the cavity is decreased. Since the size of the hole is very strongly tied to the fundamental-mode frequency of the cavity, the implication is that the strength of the head-tail effect is roughly proportional to the fourth power of the fundamental-mode frequency. In fact, the effect scales slightly less quickly than that, because while $W'_\perp(0)$ scales with the fourth power of a, the parameter $s_\perp$ scales with almost the second power of a, and it is this parameter which determines how fast the transverse wakefield "rolls over" (or becomes weaker than a strict linear dependence would dictate). The small compensating effect of aperture on the $s_\perp$ parameter leads to an approximate transverse wake dependence on the $a^{-3.8}$, rather than a^{-4}. Nonetheless, the relatively low frequency selected for the ILC cavities also leads to a relatively weak transverse wakefield compared to higher-frequency linacs (such as the 3 km SLAC linac, operating at 2.856 GHz).

Note that the sign of the transverse wakefield is always such that the tail of the beam is deflected in the same direction as the offset of the head, *i.e.*, if the bunch passes through the cavity above its symmetry axis, then the tail is kicked upwards.

2.2 Beam Dynamics of the Short-Range Wakefield

Since the transverse wakefield only becomes nonzero when the beam passes off-axis through an accelerating cavity, we can neglect its effect if the beam always passes on-axis through all cavities. Unfortunately, this is never going to be the case! there are two reasons why this is never the case: the beam has some initial offset passing through the linac (either a DC offset or an offset due to bunch-to-bunch or train-to-train position jitter), and the RF cavities in the linac are not perfectly aligned.

To understand these two cases, let us use a two-particle model of the beam: a leading particle with charge $q/2$, and a trailing particle with equal charge which follows the driving particle a distance $2\sigma_z$ behind. This is a reasonable model of a bunch with a Gaussian distribution of charge with RMS length σ_z.

2.2.1 Short-Range Wakefields and Betatron Oscillations

Let us assume that the beam initially has an offset y_0, and that both the leading and trailing particles have this offset. Let us further assume that the Twiss functions at the injection point are $\beta_y = \beta_0$, $\alpha_y = 0$, and that the beam energy at the injection point is E_0.

At the end of the linac, which we assume to be 90^d in betatron phase from beginning of the linac, the leading particle's position goes to zero, although its angle with respect to the beam axis is given by a simple out-of-phase betatron oscillation: $y'_{1,f} = y_0 \sqrt{E_0/E_f}/\sqrt{\beta_0 \beta_f}$. In the absence of wakefields (at very low charge, for example), the second particle's trajectory is the same as the first particle's. In the presence of wakefields, on the other hand, the second particle receives a kick at every cavity, proportional to the transverse offset of the first bunch at that cavity, and that kick leads to an additional betatron oscillation from the cavity to the end of the linac.

At a given cavity, the driving particle's position is given by $y_{c,1} = y_0\sqrt{\beta_c/\beta_0}\sqrt{E_0/E_c}\cos(\mu_c)$, where β_c, E_c, μ_c are the betatron function, energy, and betatron phase of the cavity of interest, respectively. The kick received by the trailing particle is given by:

$$\begin{aligned}\Delta y'_{2,c} &= \frac{y_{1,c}qL_cW_\perp(2\sigma_z)}{2E_c} \\ &= \frac{y_0qL_cW_\perp(2\sigma_z)}{2E_c}\sqrt{\beta_c/\beta_0}\sqrt{E_0/E_c}\cos(\mu_c).\end{aligned} \tag{2}$$

The resulting change in position at the end of the linac is given by

$$\begin{aligned}\Delta y_{2,f} &= \Delta y'_{2,c}R_{34}(c\to f) \\ &= y_0\beta_c\sqrt{\frac{\beta_f}{\beta_0}}\sqrt{\frac{E_0}{E_f}}\frac{qL_cW_\perp(2\sigma_z)}{2E_c}\cos^2\mu_c.\end{aligned} \tag{3}$$

If we sum over all cavities in the linac, we find:

$$y_{2,f} = \frac{y_0qL_cW_\perp(2\sigma_z)}{2}\sqrt{\frac{\beta_f}{\beta_0}}\sqrt{\frac{E_0}{E_f}}\sum_{\text{cavities}}\frac{\beta_c}{E_c}\cos^2\mu_c. \tag{4}$$

Is that a lot, or a little? Recall that the zero-charge betatron oscillation leads to a final angle for the particles of $y_0\sqrt{E_0/E_f}/\sqrt{\beta_0\beta_f}$. This is equivalent to a final position of $y_0\sqrt{E_0/E_f}\sqrt{\beta_f/\beta_0}$ (*i.e.*, this is the offset one would get at a point with the same beta function which is 90° away in betatron phase). We can thus rewrite the final position of the second particle:

$$y_{2,f} = y'_{1,f}\beta_f\frac{qL_cW_\perp(2\sigma_z)}{2}\sum_{\text{cavities}}\frac{\beta_c}{E_c}\cos^2\mu_c. \tag{5}$$

From this, we can see that the wakefield-driven motion of the second particle is large compared to the free betatron oscillation if:

$$\frac{qL_cW_\perp(2\sigma_z)}{4}\sum_{\text{cavities}}\frac{\beta_c}{E_c} > 1, \tag{6}$$

where we have substituted 1/2 for the mean value of $\cos^2\mu_c$. Since we want the wakefield-driven motion to be small compared to the free oscillation, we conclude that we want to reduce the typical betatron functions in the linac as much as possible, implying that strong focusing is desired.

2.2.2 Short-Range Wakefields and Cavity Misalignments

In the case of misalignments the situation is somewhat different. In this case, if the driving particle starts out on-axis (which we assume it does in this example), it arrives at the end of the linac on-axis. The trailing particle is another matter!

Consider a linac with a single misaligned cavity, with offset y_c from the nominal axis of alignment. At this location, the driving particle passes through the cavity off-axis, resulting in a wakefield kick to the trailing particle. The kick amplitude is given by:

$$\Delta y'_{2,c} = -\frac{y_cL_cqW_\perp(2\sigma_z)}{2E_c}. \tag{7}$$

At the end of the linac, the resulting kick to the beam is given by

$$\Delta y_{2,f} = -\frac{y_cL_cqW_\perp(2\sigma_z)}{2E_c}\sqrt{\beta_c\beta_f}\sqrt{\frac{E_c}{E_f}}\sin\Delta\mu_c, \tag{8}$$

where $\Delta\mu_c$ is the phase difference between the end of the linac and the cavity in question.

In the limit of uncorrelated misalignments, the kicks from the many RF cavities will tend to cancel one another out. Because the number of cavities is finite, the cancelation is imperfect, and we can compute the expected mean-squared offset of the trailing particle:

$$< y_{2,f}^2 >=< y_c^2 > \frac{\left(qL_cW_\perp(2\sigma_z)\right)^2\beta_f}{8E_f}\sum_{\text{cavities}}\frac{\beta_c}{E_c}. \tag{9}$$

Is that a lot or a little? We can compare the expected offset of the trailing particle to the betatron beam size at the end of the linac: $\sigma_y^2 = \epsilon_y\beta_f = \gamma\epsilon_y\beta_f m_e c^2/E_f$. After some boring mathematics, we find that our criterion – that we want the deflection of the tail to be small compared to the betatron size of the beam – can be written as:

$$< y_c^2 > \frac{\left(qL_cW_\perp(2\sigma_z)\right)^2}{8m_ec^2}\sum_{\text{cavities}}\frac{\beta_c}{E_c} < \gamma\epsilon_y. \tag{10}$$

At the limit, where the tail motion is equal to the beam size, the resulting emittance growth is 25%. This relation is derived in Appendix A.

2.3 Short-Range Wakefield: Implications

From the preceding discussion, we see that there are two criteria for achieving suitable dynamics in the presence of short-range transverse wakefields: stability of the beam in the case of a betatron oscillation, and stability of the beam in the case of misaligned cavities. These criteria are expressed in Equations 6 and 10. In both of these cases, the beam dynamics of short-range wakefields is improved by reducing the typical betatron functions in the linac.

How small does it have to get? Consider Equation 6. We can estimate the required betatron functions which are needed in the linac by using the ILC's nominal parameters (q = 3.2 nC, L_c = 1.04 m), and estimating the wakefield as $W_\perp(2\sigma_z) \approx 2\sigma_z \times 2Z_0c/4\pi a^4 = 2.9\times10^{13}\text{V/C/m}^2$. This approximation for the wakefield will give a somewhat pessimistic result, since it assumes that the wakefield rises linearly behind the trailing particle, whereas in real life it will have begun to turn over by $z = 2\sigma_z$; but this will allow a reasonable approximation of the results. We can replace the sum over $1/E_c$ with 87.2 GeV^{-1}, which is the correct value for a linac which accelerates from 15 GeV to 250 GeV in 280 RF units of 26 cavities each. Putting it all together, Equation 6 implies that the typical beta functions must be such that $0.0021\ \text{m}^{-1}\beta_c < 1$, or that betatron functions of under 450 meters are desirable in the main linac.

If we now consider Equation 10, and take into consideration that $\gamma\epsilon_y \approx 20$ nm, we find that for a 450 meter typical betatron function, an RMS cavity misalignment of 480 μm is required.

Bear in mind that these are upper limits for the quantities under consideration. In point of fact, it is not acceptable for the wakefield contribution to the final beam motion to be equal to the betatron contribution, nor is it acceptable for the emittance growth from transverse wakefields to equal the initial emittance. However, these estimates do make clear that, from the perspective of wakefield control, the betatron functions want to be small, but they do not have to be incredibly small; and the RF cavities want to be well-aligned, but they do not have to be incredibly well aligned.

3 Pitched RF Cavities

In the preceding Section, we considered RF cavities which were misaligned by a pure translation. Another possible RF cavity misalignment is a pitched cavity: the longitudinal center of the cavity

is perfectly aligned to the nominal trajectory, but the cavity has a net angle with respect to the nominal trajectory.

To first order, a pitched RF cavity does not result in any beam dynamical effects from wakefields – the kicks in the upstream half of the cavity are compensated by the kicks in the downstream half. In the next approximation there is a small effect from wakefields in this case because the beam energy is low in the upstream end of the cavity and high in the downstream end, so the compensation is not exact. As you might expect, this is a small enough effect that we can neglect it.

A more important effect is that the pitched cavity couples the very strong fundamental mode of the cavity into the transverse plane – the cavity's main accelerating field, at 31.5 MV/m, produces a slight vertical acceleration as well as the large longitudinal one. This results in two important beam dynamics effects.

3.1 Basic Formalism

Consider a bunch which is accelerated in an RF cavity with frequency ω and voltage V_c. The voltage as a function of z position in the bunch is given by:

$$V(z) = V_c \cos(z\frac{\omega}{c} + \phi_c), \tag{11}$$

where ϕ_c is the cavity's RF phase offset. If the cavity has a pitch angle ψ_c with respect to the survey line, then there is a resulting deflection experienced by particles passing through the cavity:

$$\begin{aligned} \Delta y_c'(z) &= \sin\psi_c \frac{V(z)}{E_c} \\ &= \sin\psi_c \frac{V_c}{E_c} \cos(z\frac{\omega}{c} + \phi_c). \end{aligned} \tag{12}$$

One important subtlety is that the previous result does not take into account the cavity fringe fields. At each iris within the cavity, there are radial electric fields and azimuthal magnetic fields which deflect particles which pass off-center through the irises. For a pitched cavity, it can be shown that these fringe fields have a focusing effect which is proportional to the deflection from the fundamental-mode deflection, and which acts in the opposite direction, with the result that the actual deflection is only 1/2 of what would be predicted from considering only the longitudinal field of the pitched cavity [2]. Therefore, the net effect is:

$$\Delta y_c'(z) = \frac{\sin\psi_c}{2} \frac{V_c}{E_c} \cos(z\frac{\omega}{c} + \phi_c). \tag{13}$$

Applying the logic of the betatron oscillation, we find that at the end of the linac, the resulting deflection is given by:

$$\Delta y_f(z) = \frac{\sin\psi_c}{2} \frac{V_c}{E_c} \cos(z\frac{\omega}{c} + \phi_c) \sqrt{\beta_f \beta_c} \sqrt{\frac{E_c}{E_f}} \sin\Delta\mu_c. \tag{14}$$

3.2 Time-Dependent RF Kicks

The first interesting phenomenon is that the head and tail of the bunch, which arrive at different times, receive different deflections and thus travel on different trajectories to the end of the linac. The voltage slope as a function of z is given by:

$$V'(z) = -V_c \sin(z\frac{\omega}{c} + \phi_c)\frac{\omega}{c}. \tag{15}$$

At $z = 0$, the RMS deflection over the length of the bunch can be approximated as $V'(z = 0)\sigma_z$, or

$$\sqrt{<(\Delta y_c')^2>} \approx \sigma_z \frac{\sin\psi_c}{2}\frac{V_c}{E_c}\sin\phi_c\frac{\omega}{c}. \tag{16}$$

This results in a spread in final positions of the particles given by:

$$\sqrt{<y_f^2>} = \sigma_z \frac{\sin\psi_c}{2}\frac{V_c}{E_c}\sin\phi_c\frac{\omega}{c}\sqrt{\beta_f\beta_c}\sqrt{\frac{E_c}{E_f}}\sin\Delta\mu_c. \tag{17}$$

As we did for the case of misaligned cavities, we can recast this as a statistical ensemble of pitched cavities, estimate the expected mean-squared spread in final particle positions, and compare to the nominal beam size, with the requirement that the former quantity be smaller than the latter. This leads to the requirement that:

$$\sigma_z^2\frac{\omega^2}{c^2}\frac{<\psi_c^2>}{8}\frac{V_c^2}{m_e c^2}\sin^2\phi_c\sum_{\text{cavities}}\frac{\beta_c}{E_c} < \gamma\epsilon_y. \tag{18}$$

One might anticipate that the emittance growth effect of the time-dependent RF kicks is a small one – after all, the beam rides close to the RF crest, so $\sin\phi_c$ is small, and the bunch is short, so σ_z is also small. And indeed, one finds that with 450 m betatron functions, the RMS cavity pitch angle needed to double the emittance from the time-dependent RF kick effect is on the order of 2 milliradians. In the ILC bunch compressors, by comparison, one might expect this to be a large effect – the beam runs close to the zero-crossing, and the bunch length is much larger.

3.3 Dispersive Centroid Kick

A much more important effect comes from the fact that not all the particles in the bunch have the exact same energy. The RMS energy spread at the entrance to the linac is about 1.5%, and this value adiabatically damps down to 0.09% at the end of the linac. There is an additional contribution from the short-range wakefield, and another contribution from quantum excitation of the synchrotron radiation emitted in the positron production wiggler in the electron linac, but for now we will neglect these contributions.

When a beam with a finite energy spread undergoes a betatron oscillation in a FODO lattice, the particles with different energy oscillate with different betatron wavelengths due to the chromaticity of the lattice. As a result, the particles with low and high energy no longer oscillate coherently with one another, leading to emittance growth. As a general rule, it can be shown that a beam which is launched down an infinitely-long FODO lattice with an initial offset of $n\sigma_y$ and normalized emittance $\gamma\epsilon_y$ will reach the (infinitely-distant!) end of the lattice with a final offset of zero and a final normalized emittance of $\gamma\epsilon_y\sqrt{1+n^2}$. Thus, a 1 σ_y initial offset translates to a 41% growth in normalized emittance.

The result above suggests that we need to limit cavity pitches such that the resulting offset of the beam at the end of the linac, calculated from linear optics, is less than about 1 $\sigma_{y,f}$. This is the right order of magnitude, but it is actually more restrictive than is necessary. This is for several reasons. First, the linac lattice is not infinite, so the breakdown in coherence of the oscillation between the particles of different energies (known as *filamentation*) is not yet complete. Second, unlike the case of an initial betatron oscillation, the oscillation described here actually builds up gradually as the sum of individual cavity kicks; thus, a 1 σ_y final offset, driven by thousands of small cavity kicks, is not as severe, from emittance growth point of view, as a 1 σ_y initial offset in a perfectly-aligned linac. Even better, as the kick amplitude is growing, the RMS energy spread is

adiabatically damping, leading to reduced sensitivity of the emittance to the betatron oscillation (or, if you prefer, as the energy spread decreases the number of betatron oscillations needed to achieve filamentation increases). Finally, in real life the trajectory of the beam centroid through the accelerator is something that we can measure and control, via BPMs and correctors; thus, ultimately it is the performance of the BPMs and correctors which will set the level of emittance growth from sources such as the cavity pitch.

Nonetheless, we can use the "naive" estimate of the orbit deflections from cavity pitches, in the absence of correction, to get a sense for the scale of the problem. Starting from Equation 13, replacing $\sin\psi$ with ψ and $\cos\phi$ with 1, we find:

$$\frac{\sigma_\psi^2}{8}\frac{V_c^2}{m_e c^2}\sum_{\text{cavities}}\frac{\beta_c}{E_c} < \gamma\epsilon_y. \tag{19}$$

Assuming a typical β_c of 450 meters, this would yield a tolerance on the pitch angle ψ of about 1.4 microradians. While this estimate is both crude and "naive," it demonstrates that the scale of cavity tolerances which can be tolerated without correction is too small to achieve at installation, that the tolerance becomes looser for stronger focusing, but that the emittance problem from pitched cavities is too severe to have any hope of correcting it by simply reducing β_c.

4 Quadrupole Misalignments

Like the cavity pitch error, misaligned quadrupoles lead to emittance growth from dispersion, since a misaligned quadrupole deflects the beam and the deflection is inversely proportional to the energy of the particle. We can recast Equation 13, replacing the deflection from the pitched RF cavity with the deflection from the misaligned quadrupole:

$$\Delta y_q' = -y_q K_q L_q, \tag{20}$$

where y_q is the quad offset with respect to the accelerator axis, K_q is the normalized quad strength, $K_q = (dB_q/dy)/B\rho$, the sign of K_q is such that $K_q > 0$ is horizontally focusing and vertically defocusing, and $B\rho$ is the magnetic rigidity, $B\rho = E_q/c$. The resultant deflection at the end of the linac is given by:

$$\Delta y_f = -y_q K_q L_q \sqrt{\beta_f \beta_q}\sqrt{\frac{E_q}{E_f}}\sin\Delta\mu_q. \tag{21}$$

We can make the same requirement here that we did for pitched cavities, specifically that the total deflection at the end of the linac should be comparable to the nominal beam size. This leads us to the naive criterion for emittance growth:

$$\frac{\sigma_{y,q}^2}{2m_e c^2}\sum_{\text{quads}}(K_q L_q)^2\beta_q E_q < \gamma\epsilon_y. \tag{22}$$

So far, this looks a lot like the expressions we have already seen, and seems to make similar implications. However, this is not the case, and the reason it is different should be obvious: that β_q, the typical betatron function at the quadrupoles, is not independent of $K_q L_q$, the focusing strength of the quadrupoles. We can make this explicit as follows: in a FODO lattice, the maximum betatron function β^+ and the minimum beta function β^- can be written in terms of $K_q L_q$, the quad-to-quad spacing L_{qq}, and the phase advance per cell μ_{cell}, as follows [3]:

$$\beta^\pm = \frac{1 \pm \sqrt{1-\cos^2(\mu_{\text{cell}}/2)}}{\cos(\mu_{\text{cell}}/2)K_q L_q}. \tag{23}$$

We can get an approximate sense of the scale of the problem by replacing β_q in Equation 22 with the mean of β^+ and β^-, and solving for $K_q L_q$ in terms of the typical beta function and the cell phase advance:

$$K_q L_q = \frac{1}{\beta_t} \frac{1}{\cos(\mu_{\text{cell}}/2)}. \tag{24}$$

Now our criterion for emittance preservation becomes the following:

$$\frac{\sigma_{y,q}^2}{2m_e c^2} \frac{1}{\beta_t} \frac{1}{\cos^2(\mu_{\text{cell}}/2)} \sum_{\text{quads}} E_q < \gamma \epsilon_y. \tag{25}$$

Equation 25 is not entirely satisfying – for one thing, the parameters β_t and μ_{cell} are implicitly correlated because the quad-to-quad spacing is fixed in the ILC linac to 1 quad per 3 cryomodules (or about 1 quad per 40 meters). Nonetheless, it does show us that, unlike the other issues we have studied to date, the dispersive emittance growth from quad misalignments benefits from weaker focusing – weaker focusing increases β_t and drives $\cos^2(\mu_{\text{cell}}/2)$ towards its maximum value of 1. Also of interest is that the impact of misaligned quads is stronger at the high-energy end of the linac, whereas the impact of pitched or misaligned cavities is stronger at the low-energy end of the linac.

To make this more concrete, let us consider again the linac with 450 meter typical betatron functions and 280 quadrupoles. With a 40 meter quad spacing, the phase advance per cell is about 10° (see Appendix B to see how to do this calculation), the sum of E_q is 37,100 GeV, and Equation 25 shows that a quad alignment with respect to the survey line of better than 0.5 μm is needed to limit the emittance growth due to quad misalignments. Again, this neglects the effects of steering corrections, which reduce the emittance growth from dispersion, and is a substantial overestimate due to the crude way in which the calculation was performed.

5 The ILC Main Linac Lattice

So much for the theory of single-bunch emittance preservation in a linear accelerator – how is the ILC main linac designed to address these issues?

Figure 2 shows the betatron functions of the positron main linac. The phase advance per cell is 75° in the horizontal, and 60° in the vertical, with typical betatron functions of 80 meters. As shown in the discussion of Section 2.3, this is much smaller than the maximum tolerable betatron function, which is around 450 meters. Therefore, we expect that emittance growth from wakefields in a betatron oscillation will not be much of an issue for the ILC main linac.

Review of Equations 10, 18, 19, and 22 suggest that the "naive" tolerances will be approximately as follows:

- Cavity misalignments: 1.4 mm RMS
- Cavity pitch (time-dependent kick): 4.6 mrad RMS
- Cavity pitch (dispersion): 3.3 μrad RMS
- Quad misalignments: 0.2 μm RMS.

As we can see, with these parameters it should be straightforward to limit the impact of cavity misalignments and time-dependent kicks through careful construction and installation, since these

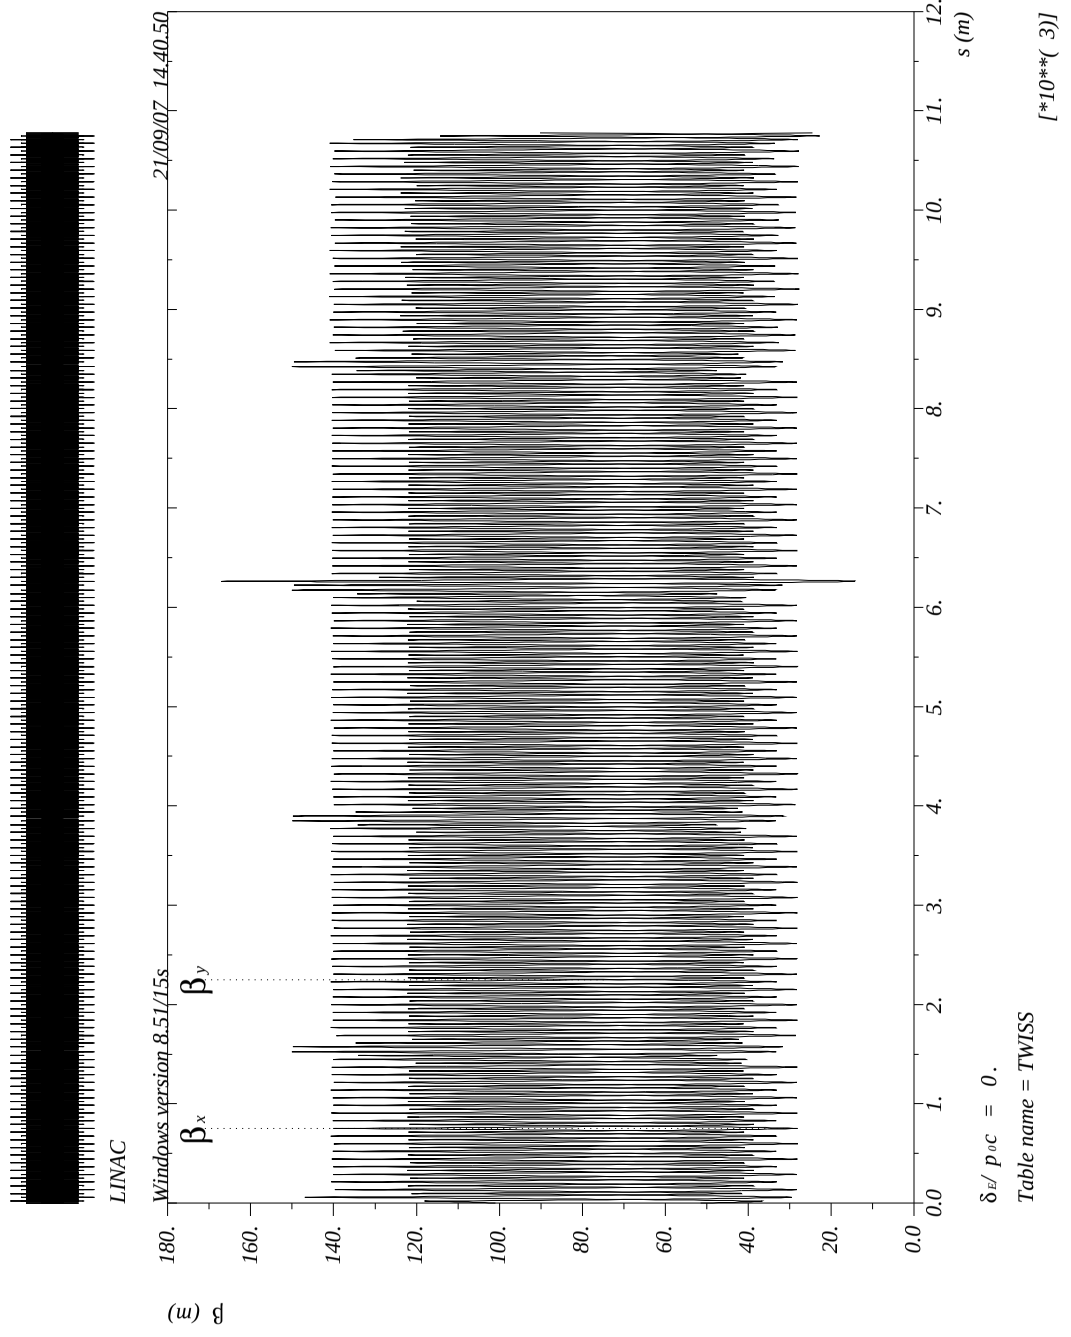

Figure 2: Betatron functions of the ILC positron main linac.

tolerances are quite loose. On the other hand, the tolerances on growth from dispersion are extremely tight, even if the actual tolerances which must be achieved are several times larger than the ones shown here.

One might wonder whether it makes more sense to try and increase the beta functions in the linac, and relax the focusing, to try and loosen the quad misalignment tolerance in particular. The answer is that it does not make sense. As we have seen, there is an upper limit on the betatron function in the ILC linac, which is set by wakefields acting on a jittering beam. This limit is only a few times larger than what was ultimately chosen, and would only result in an incremental improvement in the required alignment tolerance (from 0.2 μm to 0.5 μm, for example). This is too small an improvement to be worth making, given that the achievable vertical alignment of an accelerator is certainly no better than 50 μm, and for a superconducting system is probably closer to 300 μm, due to the movement and change in dimensions of the components as they are cooled from room temperature to 2 K. Much better, therefore, to strengthen the beta functions enough to correct the wakefield effects, which are otherwise difficult to fix by beam-based tuning, and leave the quad misalignments (and the cavity pitches) with tight tolerances that can be corrected later via beam-based tuning.

5.1 Aperiodicity of the Lattice

As Figure 2 shows, the betatron functions are not perfectly regular – there are periodic variations ("beats") in the optical functions. These are necessary due to the engineering requirements of such a large cryogenic system. The ILC linac cryogenic system is broken into 2.5 km long subunits, each of which is fed by a cryogenic plant on the surface. At the junction between these subunits is a room-temperature drift and diagnostic section which disrupts the regularity of the linac layout. In addition, the cryogenic fluids are distributed to cryomodules in segments of about 160 meters, and therefore every 160 meters is an additional interruption to the regularity of the linac.

5.2 Vertical Curvature of the Linac

The "linear" in linear accelerator implies that the main accelerating component of the ILC should be constructed in a straight line. In fact, the linear collider is not designed to be straight, but rather to have a small curvature in the vertical plane. This curvature causes the linac to follow a gravitational equipotential: to a good approximation, then, the vertical radius of curvature of the linac is about the same as the mean radius of the Earth (typically taken to be 6370 km).

The purpose of this curvature is twofold. First, it vastly simplifies the distribution of superfluid helium if there is no gravitational force pulling the helium from one point in the accelerator to another point. Second, it allows the ILC tunnel to lie at a near-constant depth below the surface. In the absence of this curvature, the center of the site would need to be much deeper than the far ends. As an additional benefit, this method of construction means that the final length of the ILC does not need to be known in advance, implying that an energy upgrade by "digging out" from the center becomes feasible.

The vertical curvature is achieved by exciting the vertical steering correctors at each quadrupole. Figure **??** (which is missing in this draft!) shows the resulting dispersion function, including all of the "beating" due to lattice aperiodicity. The dispersion spikes at the beginning and end are places where the correctors have been tuned to introduce and cancel the necessary vertical dispersion, since the matched dispersion in the RTML and BDS injection are both zero.

6 Steering, Alignment, and Emittance Preservation

As discussed in the previous section, the ILC main linac lattice design is acceptable on the topic of emittance dilution from wakefields and time-dependent transverse kicks from pitched cavities, but is very far from acceptable in terms of emittance growth from steering errors, most notably pitched cavities and misaligned quadrupoles. For these errors, the typical tolerances are on the order of micrometers and microradians, while the achievable installation tolerances for superconducting components is probably closer to 300 μm or 300 μrad.

Fortunately, the ILC main linac has vertical steering dipoles at each quad, and a beam position monitor (BPM) at each quad as well. This means that it is possible to correct the steering errors which drive emittance growth in the linac, provided that one has a proper plan to take advantage of these devices. We shall discuss several possible plans (or algorithms, if you prefer) for steering in this section.

6.1 One-to-One Steering Correction

The simplest correction one could imagine applying to the linac is to simply use the correctors to steer the beam such that all of the BPMs read zero. Since the number of BPMs is equal to the number of correctors, this is often known as "one-to-one" steering. Because the number of measurements equals the number of corrections, the system has a unique solution.

It is interesting to note that, in the limit where the BPMs are perfectly aligned to the nominal accelerator axis, one-to-one steering will reduce the emittance growth almost to zero no matter how badly the other components are misaligned. This is because any kick is corrected at the nearest BPM; as a result, the misaligned quads, pitched cavities, and corrector magnets form a complicated chicane, and the off-energy and on-energy trajectories will always meet at each BPM in that case. Thus, the performance of one-to-one steering depends mainly on the alignment of the BPMs themselves, and only to a very small degree on the alignment of the quads and cavities.

What happens if the BPMs are not perfectly aligned to the accelerator axis? Or, to change the sense of the question, how well do the BPMs need to be aligned to the accelerator axis? To get a sense of this, we will follow the approach of Ruth [5], considering an accelerator in which everything is initially perfectly aligned except for one BPM, which has a misalignment y_{BPM}. The correctors immediately upstream of the BPM, immediately downstream of the BPM, and at the BPM will all be excited to produce a closed bump such that the misaligned BPM reads zero, along with all the other BPMs in the beamline. To good approximation, the nominal deflection angles of the upstream and downstream correctors are given by:

$$\theta_{-1,1} = \frac{y_{\mathrm{BPM}}}{L_{qq}}. \tag{26}$$

The total kick angle at the BPM must obviously be $-2y_{\mathrm{BPM}}/L_{qq}$, but some of this deflection is provided by the fact that the beam is passing off-axis through the quadrupole at the location of the misaligned BPM, leading to a kick angle of $\theta_q = y_{\mathrm{BPM}} K_q L_q$; therefore, the middle corrector has a kick angle of $\theta_0 = -2y_{\mathrm{BPM}}/L_{qq} - y_{\mathrm{BPM}} K_q L_q$. Simple arithmetic will demonstrate that, at the location of the downstream corrector, both the position and the angle of the centroid are zero.

Now consider the trajectory of an off-energy particle. The three correctors each produce a deflection which differs from the nominal one by a factor of $1-\delta$. In the case of the deflection from passing off-axis through the quad, the deflection differs from nominal by a factor of $(1-\delta)^2$: one factor of $1-\delta$ comes from the fact that the quad strength, $K_q L_q$, is larger for low-energy particles, and one comes from the fact that the position offset of the particles, driven by the kick angle of

the first corrector, is also larger for low-energy particles. Adding the four terms and neglecting contributions of order δ^2, we find that an off-energy particle has an angle error from this system given by:

$$\Delta y' = -\delta y_{\rm BPM} K_q L_q. \tag{27}$$

Computing the effect from the ensemble of BPMs, and requiring that the resulting displacement of the off-energy particles be smaller than the beam size at the end of the linac, we find the following criterion:

$$\sigma_{y,\rm BPM}^2 \frac{\sigma_\delta^2 E_0^2}{2 m_e c^2} \frac{1}{\beta_t} \frac{1}{\cos^2(\mu{\rm cell}/2)} \sum_{\rm quads} \frac{1}{E_q} < \gamma \epsilon_y. \tag{28}$$

As was the case for our unsteered linac, we find that the linac emittance preservation benefits from a weaker lattice. Unlike in that previous case, we find here that the low-energy end of the linac has a bigger impact on the overall emittance dilution. Substituting numbers for the current linac design, we find that the BPM alignment tolerance is about 85 μm. This is clearly much better than the alignment tolerances in the absence of steering, but is still about a factor of 4 tighter than what is likely to be achieved at installation. Thus, we are forced to rely on somewhat more complex methods for steering the linac to limit emittance growth.

6.2 Quad Shunting

Since a quad deflects a beam which passes off-axis through it, one can estimate the beam-to-quad distance by varying the quad strength and observing the resulting change in orbit downstream of the quad. The precision of this technique is set by the maximum change in quad strength which is achievable while still transporting the beam, by the number, location, and resolution of the downstream BPMs which perform the measurement, and by the number of beam pulses which can be measured at each quad strength setting. Ultimately, there is a systematic limitation to the achievable accuracy: this comes from the fact that the quad center can move when the quad strength is varied, which results in a deflection which mimics the deflection of an off-axis beam passing through a quad with varying strength. The systematic error can be roughly estimated as [6]:

$$\Delta x_{\rm QSfit} = \Delta x_{\rm center} \left(\frac{K_q}{dK_q} + 1 \right), \tag{29}$$

where K_q/dK_q is the reciprocal of the fractional strength change used in the quad shunting procedure and $\Delta x_{\rm center}$ is the resulting motion of the quad center. For a system in which the quad strength is reduced by 20%, this implies that the error in finding the beam-to-quad offset will be four times as large as the unwanted motion of the quad center.

Once the beam-to-quad offsets are known, the BPM-to-quad offsets can be computed for each BPM/quad package in the linac. This information can then be put to use in steering algorithms such as Kick Minimization.

6.3 Kick Minimization

After Quad Shunting, the linac BPMs are well-aligned to the linac quads, but neither the quads nor the BPMs are well-aligned to the nominal accelerator axis. We can make use of that correlation in a subtle and fascinating way:

Imagine that, by accident or miracle, the linac gets steered in such a way that the beam is exactly on the design axis for the entire length of the accelerator. In the absence of the kicks from the pitched cavities, what we would find is that the beam goes off-center through the misaligned

quads; the correctors at each quad are set such that the corrector kick exactly cancels the kick from the nearby misaligned quad; and the BPM readings show a non-zero orbit at each BPM, since the BPMs are perfectly aligned to the quads but the quads are misaligned to the beam. Putting all this together, it means that there is a correlation between the BPM reading at each quad and the corrector setting at each quad, and the correlation is a byproduct of the good alignment between the BPM and the quad.

Quantifying the issue, in the case of the accidentally-perfect alignment we expect that, at each quad/BPM/corrector location

$$y_{\mathrm{BPM,reading}} = -\frac{\theta_{\mathrm{corrector}}}{K_q L_q}. \tag{30}$$

In this case, the net kick at each BPM is locally corrected, thus the name of the algorithm: "Kick Minimization" (KM) [7].

In principle, we can imagine a steering solution in which we seek to zero the quantity $y_{\mathrm{BPM,reading}} + \theta_{\mathrm{corrector}}/K_q L_q$ at all quad/BPM/corrector locations simultaneously. In practice, such a solution is unstable in the presence of errors, since an erroneously-large value of $y_{\mathrm{BPM,reading}}$ can be "balanced" by a large value of $\theta_{\mathrm{corrector}}$, leading to a large offset at the next BPM and therefore a large corrector value at that location, and so on – a solution which tends to grow along the linac. To provide stability against such errors, it is helpful to weakly constrain the absolute reading at each BPM, so that such a growth is prohibited. This results in an overconstrained problem: the number of constraints, considering both the absolute orbit constraint and the constraint on the BPM readings plus the local corrector settings, is twice as large as the number of parameters (in this case, corrector settings). Such problems are solved by forming an appropriately weighted χ^2 and minimizing it as a function of the parameters. In this case,

$$\chi^2 = \sum_{\mathrm{BPMs}} \frac{y^2_{\mathrm{BPM,reading}}}{\sigma^2_{y,BPM}} + \sum_{\mathrm{BPMs}} \frac{(y_{\mathrm{BPM,reading}} + \theta_{\mathrm{corrector}}/K_q L_q)^2}{\sigma^2_{y,BtQ}}, \tag{31}$$

where $\sigma_{y,BPM}$ is the RMS misalignment of the BPM with respect to the nominal accelerator axis, and $\sigma_{y,BtQ}$ is the RMS misalignment of the BPMs to the quads, based on the expected precision and accuracy of the quad shunting technique.

In simulation studies, the minimization above is very effective in finding an orbit which achieves a small emittance growth, provided that cavity pitch angles are neglected. When cavity pitches are included, the optimal setting of each corrector changes: in addition to correcting the kick from the nearest quad, each corrector needs to steer out the deflections from the cavities between the corrector and the next downstream BPM. Since these deflections are not known, this modification has the effect of increasing the denominator of the second term in the χ^2. Studies of such advanced steering methods are ongoing.

6.4 Dispersion Free Steering

Another option for correcting the trajectory to limit emittance growth from dispersion is to directly measure the dispersion and to derive a set of corrector setttings which minimize it. This is known as "Dispersion Free Steering" (DFS) [8].

More generally, the key to DFS is to generate a mismatch between the beam energy and the optics of the lattice. This can be done by varying the actual beam energy (via changing the RF system parameters) or by varying the magnet strengths, though this can introduce the same strength-dependent variation in the quad center position which is a limitation for quad shunting.

As with KM, the DFS solution is naturally unstable in the presence of errors, and must be weakly constrained by simultaneously minimizing the measured orbit:

$$\chi^2 = \sum_{\text{BPMs}} \frac{y^2_{\text{BPM,reading}}}{\sigma^2_{y,BPM}} + \sum_{\text{BPMs}} \frac{[\Delta y_{\text{meas}} - \Delta y_{\text{model}}(\vec{\theta})]^2}{2\sigma^2_{\text{res}}}, \tag{32}$$

where σ_{res} is the BPM resolution, Δy_{meas} is the measured change in BPM readings when the energy match is varied, and $\Delta y_{\text{model}}(\vec{\theta})$ is the expected change in the orbit, which is a function of the corrector settings, $\vec{\theta}$. Depending on the details of how the DFS algorithm is configured, it can be designed to properly manage the kicks from pitched RF cavities as well as the kicks from misaligned quads or other effects, since it can be made into a direct measurement of the local dispersion. Note that, in the case of a vertically-curved lattice, the model Δy must include the effects of the design dispersion function. In this case, "dispersion *free* steering" is not an accurate name, and the technique is more often referred to as "dispersion matched steering" or simply "dispersion steering." Finally, it is important to note that in the presence of non-zero design dispersion, the technique is not a classical "nulling" technique, which means that the BPM scale factor becomes important as a potential source of systematic error.

6.5 Ballistic Alignment

In this technique, a straight-line, non-deflected beam-based reference is obtained by switching off all magnets and RF cavities in the linac, or a portion of it, and measuring the BPM readings on all BPMs. The devices are restored, and the beam is re-steered to the readings which were obtained with all devices switched off. This is known as "ballistic alignment" (BA) [9].

The simplicity of BA is hard to beat, and it obviously manages misaligned quads and pitched cavities simultaneously. The main limitations to this are systematic: the degree of certainty that the devices are all truly off, with no remnant fields; the effect of stray electromagnetic fields on the nominally-ballistic initial orbit; potential difficulties in transporting the beam through long sections of deactivated beamline.

6.6 Dispersion Bumps

In Section 6.1, we saw that introducing a closed bump around a quadrupole produces a nonzero dispersion. We can make use of this to deliberately introduce a given amplitude and phase of dispersion, in order to globally cancel dispersion from one or more unknown sources. The resulting "dispersion bump" is varied in size until the emittance measured at the end of the linac is minimized.

Emittance bumps have limited effectiveness due to filamentation and high-order dispersion effects. Thus, they are typically used as an "afterburner" applied to a system which has been corrected with a good, but not perfect, steering solution.

7 Long-Range Transverse Wakefields

In Section 2, we reviewed the beam dynamics of short-range transverse wakefields, which act within a single bunch – deflecting the tail based on the transverse position of the head. The short-range wakefield is the sum of the transient excitation of an infinite number of deflecting modes, similar to the way in which the short-range longitudinal wakefield is the sum of an infinite number of accelerating modes.

As we have seen previously, the Q of a superconducting accelerating cavity is proportional to ω^{-2}. This means that, for extremely high frequencies, the Q is very low and the mode quickly

damps away. Nonetheless, for a large number of low-frequency modes, typically modes which in the same frequency regime as the fundamental accelerating mode (1.3 GHz), the Q can remain quite high for the deflecting modes. This means that the modes which are excited by one bunch can remain in a cavity when the next bunch arrives. These are the higher-order modes (HOMs), or long range transverse wakefields, of an accelerating cavity.

Because the HOMs can have high Q values, and very limited damping between bunches, it is possible for the HOMs in a cavity to be excited by all of the bunches in a train. The excitation is generally not "resonant," that is, the frequency of the HOMs is generally not harmonically related to the bunch interval, so the HOMs which are driven by a bunch are generally different in phase from the HOMs which are left in the cavity by previous bunches. Nonetheless, the amplitude of the HOM can build up over time, resulting in a deflection of late bunches which is much larger than the deflection of early bunches. In addition, a bunch train which is executing a betatron oscillation will excite HOMs in all of the cavities in the linac, leading to a deflection amplitude that grows with S along the accelerator as well as growing from bunch to bunch within a train. If, at the end of the linac, the deflection of trailing bunches exceeds the amplitude of the initial betatron oscillation, the dread "beam break-up instability" (BBU) can occur.

7.1 Managing Higher Order Modes

The techniques for managing higher-order modes are well-understood and have been for several decades. The fundamental techniques are *damping* and *detuning*. We will briefly describe each of these techniques.

7.1.1 Damping

The most damaging higher order modes, in terms of beam dynamics, are the ones which have the lowest frequencies: in the case of the ILC cavities, modes in the frequency band from 1.3 GHz to 4.0 GHz are typically of greatest concern. Ideally, we would like the excitation of these modes which is generated by one bunch to vanish by the time the next bunch arrives 369 nanoseconds later. In order to reduce the stored energy in these modes to $1/e$ of their peak values in this time, the Q values for the modes must be around 5×10^3.

As we have seen, the wall Q for modes in this frequency range is likely to be vastly larger than the desired 5×10^3 level, due to the inconvenient (in this context) fact that the cavity walls are in a superconducting state. Fortunately, we can take a lesson from the fundamental mode, in which the wall Q is around 2×10^{10} but the Q with the fundamental mode coupler is closer to 3×10^6. This suggests that we should, essentially, cut a hole in the cavity which the HOM power can leak out of, and then provide it with some sort of cooled load which can safely absorb it. In fact, this is more or less exactly what is done – a set of *higher order mode couplers* are attached to the cavity to reduce the loaded Q of the dipole modes to something closer to the desired level. An additional constraint on the HOM couplers is that they should reject the fundamental mode at 1.3 GHz, *i.e.*, the fundamental mode should not be coupled to, or damped, any more than is necessary.

The ILC design incorporates HOM couplers at the upstream and downstream end of each cavity. In addition, there is a HOM load incorporated into the vacuum system near each quadrupole in the linac, which serves to absorb HOM power which is at frequencies high enough to propagate freely down the linac (*i.e.*, frequencies above the cutoff of the cavity irises and the beam pipe). The HOM couplers do not achieve the full measure of damping which would be required to produce one e-folding between bunch arrivals: the Q values for the lowest modes are typically between 10^4 and 10^5 [10]. Figure 3 shows the spectrum of RF power which is removed from an ILC cavity via the cavity's HOM couplers. Note that the 1.3 GHz fundamental mode line can be seen in this

plot; however, while the excitation from the fundamental mode is larger than the excitation from the HOMs, the actual coupling of the coupler to the fundamental is quite small, as required; the excitation from the fundamental is so large because the fundamental mode has vastly more stored energy than the HOMs.

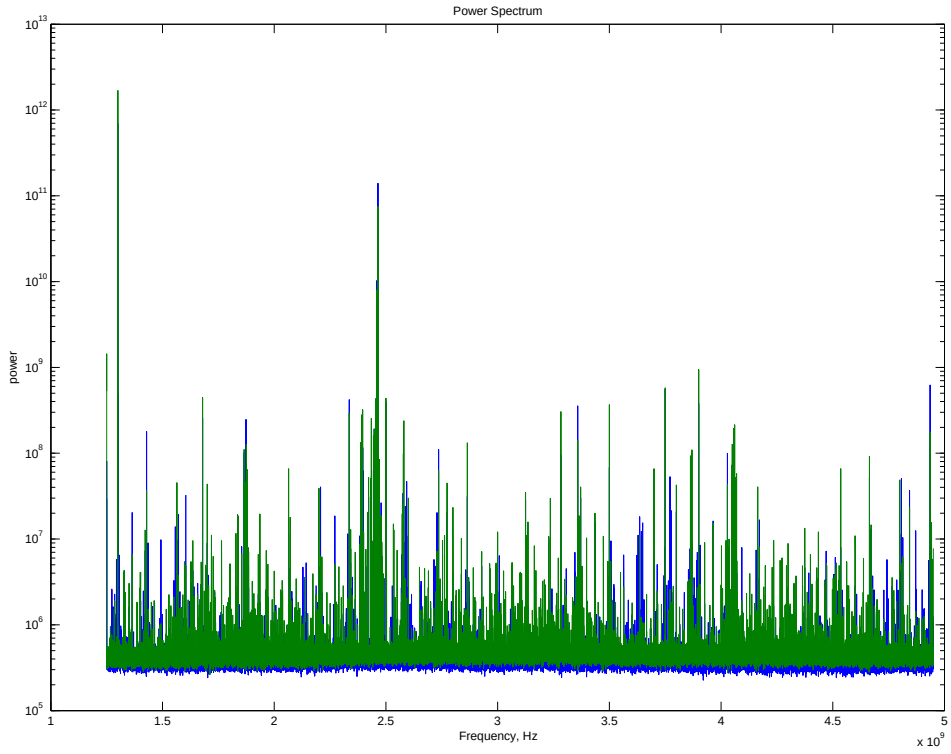

Figure 3: Spectrum of RF power detected in the HOM couplers of an ILC RF cavity. Not all of the lines represent real HOMs, as there are a number of instrumentation artifacts in the signal. Figure courtesy of S. Molloy.

It is interesting to note that, for a mode with a frequency of 4 GHz and a Q of 10^5, the e-folding time is 4 microseconds. This implies that in 36 microseconds, the deflecting voltage left by the first bunch has fallen to 1% of its maximum value – in essence, the deflecting field from the first bunch is lost within a few tens of microseconds. Since the bunch train is about 1 millisecond long, we expect that, when the HOM damping is included, the distortions to the bunch train trajectories from HOMs will reach a "steady state" within 5-10% of the bunch train length.

7.1.2 Detuning

As described above, the damping of the HOMs provided by the HOM coupler is not sufficient to remove all the HOM power in a time comparable to the bunch spacing, and so prior to the onset of steady-state deflections the beam dynamics effects of the HOMs can be substantial. Since it is impractical to add additional damping, another technique must be used to limit the multi-bunch emittance growth during the first 50-100 microseconds of the bunch train.

One option is to detune the HOMs in the various cavities with respect to one another. HOM-driven beam break-up is caused by resonant excitation: a bunch train receives exactly the same wakefield kick in each of a large number of cavities, leading to a buildup of small kicks to a large total deflection. This can be mitigated by adjusting the frequencies of the HOMs such that they are slightly different in each cavity. When summed over all cavities, the resulting kicks will quickly tend to cancel out rather than add coherently.

It is expected that in the ILC cavities, errors in the cavity construction will lead to frequency shifts from nominal, with an RMS of about 10^{-3} [11]. Since the frequencies of interest are at the GHz level, the cumulative "line width" of each of the detuned HOMs, when considered over the length of the linac, is around 1 MHz. This implies a fall-off of the wakefield in a time comparable to 1 microsecond, which is sufficient to limit the amount of multi-bunch emittance growth in the 50-100 microseconds prior to the onset of steady-state.

Given that we expect to achieve a fall-off time of 1 microsecond using natural detuning, what is the purpose of adding damping as well, given that damping does not take full effect until after some tens of microseconds? The damping is required because the detuned wakefields will recohere after some time – that is, the deflections of trailing bunches would be reduced due to detuning, and then would suddenly shoot back up to an unacceptable level as the detuned modes got back into phase with one another. Addition of damping ensures that the recohered wakefields will be at a sufficiently low amplitude to prevent unacceptable multi-bunch emittance dilution.

A Computing the Emittance Growth Due to Wake-Driven Motion of the Bunch Tail

In Section 2, we set a limit on the permitted misalignments of the RF cavities such that the tail motion driven by transverse wakefields is equal to the nominal beam size. How much actual emittance growth does that correspond to?

Consider a two-slice model, as used in Section 2, with one leading and one trailing slice, where each slice has full transverse coordinates (x, x', y, y'), but each slice has zero length and zero energy spread. Each slice contains N particles. Transverse wakefields do not deflect the leading slice, but the trailing slice is deflected by $\sigma_{y,f}$ at the end of the linac. We would like to estimate the resulting beam size and emittance growth, so we should start by estimating the mean-squared beam size:

$$\sigma_y^2 = < y^2 > - < y >^2 . \tag{33}$$

Let us start by estimating $< y^2 >$ for this ensemble. We can write this as:

$$< y^2 >= \frac{1}{N}\left[\sum y_1^2 + \sum y_2^2\right], \tag{34}$$

where y_1 and y_2 are variables representing the positions of particles in the leading and trailing slices, respectively. We can replace y_1 with a Gaussian-distributed variable with RMS width σ_y; we can also replace y_2 with $\sigma_y + \Delta y_2$, where the first term is the offset of the slice and the term Δy_2 is a Gaussian-distributed variable with RMS width σ_y. This implies that

$$< y^2 >= \frac{1}{N}\left[\frac{N}{2}\sigma_y^2 + \sum(\sigma_y + \Delta y_2)^2\right]. \tag{35}$$

Expanding the sum, the cross-term between σ_y and Δy_2 vanishes, leaving:

$$< y^2 >= \frac{1}{2}\sigma_y^2 + \frac{1}{N}\sum(\sigma_y + \Delta y_2)^2 = \frac{3}{2}\sigma_y^2. \tag{36}$$

Meanwhile, the value of $< y >$ is trivially $\sigma_y/2$, therefore:

$$\sigma_y^2 = \frac{3}{2}\sigma_y^2 - \frac{1}{4}\sigma_y^2 = \frac{5}{4}\sigma_y^2. \tag{37}$$

If we assume that the dilution of the beam size is uniformly distributed to all betatron phases, then both the σ_y^2 and $\sigma_{y'}^2$ are increased by the same factor of 25%, leading to an emittance growth of 25% from the 1 σ_y and 1 σ_y' deflection of the tail.

B Calculating the Phase Advance per Cell from β_t and L_{qq}

Given the quad-to-quad spacing $L_q q$ and the typical betatron function value β_t, how does one compute the resulting phase advance per cell? To solve this we follow the approach of Wiedemann [4].

Start with a few definitions: $b \equiv \beta_t / L_{qq}$, and $\kappa \equiv 1/(K_q L_q L_{qq})$. The phase advance relation is

$$\sin(\mu_{\text{cell}}/2) = 1/\kappa, \tag{38}$$

and the relation between b and κ is:

$$b = \frac{\kappa^2}{\sqrt{\kappa^2 - 1}}. \tag{39}$$

This allows us to write and solve a quadratic equation for κ^2, which has as its solution:

$$\kappa^2 = \frac{b^2 + \sqrt{b^4 - 4b^2}}{2}. \tag{40}$$

This allows solution for κ, which leads directly to a solution to μ_{cell} via Equation 38.

References

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

[2] R. Miller, private communication.

[3] K. Brown, R. Servranckx, "First and Second Order Charged Particle Optics," 54 (1984).

[4] H. Wiedemann, *Particle Accelerator Physics*, Vol. 1, 187-188 (1993).

[5] R.D. Ruth, "Emittance Preservation in Linear Colliders," SLAC-PUB-4436 (1987).

[6] P. Tenenbaum, T.O. Raubenheimer, "Resolution and Systematic Limitations in Beam Based Alignment," PRST-AB 3:052801 (2000).

[7] K. Kubo, "Emittance Dilution Due to Misalignment of Quads and Cavities of ILC Main Linac,"
`http://lcdev.kek.jp/~kkubo/reports/MainLinac-simulation/lcsimu-20050310.pdf`
(2005).

[8] T.O. Raubenheimer and R.D. Ruth, "A Dispersion Free Correction Technique for Linear Colliders," NIM **A302** 191 (1991).

[9] T.O. Raubenheimer, D. Schulte, "The Ballistic Alignment Method," Proceedings of Pac 1999, 3441 (1999).

[10] *TESLA TDR*, Vol. II, 72 (2001).

[11] N. Baboi and R. Brinkmann, "Higher Order Mode Effects and Multi-Bunch Orbit Stability in the TESLA Main Linac," TESLA-2000-28 (2000).