Main Linac II:

Transverse Dynamics

Peter Tenenbaum, SLAC

Reminder – The Linac RF Unit

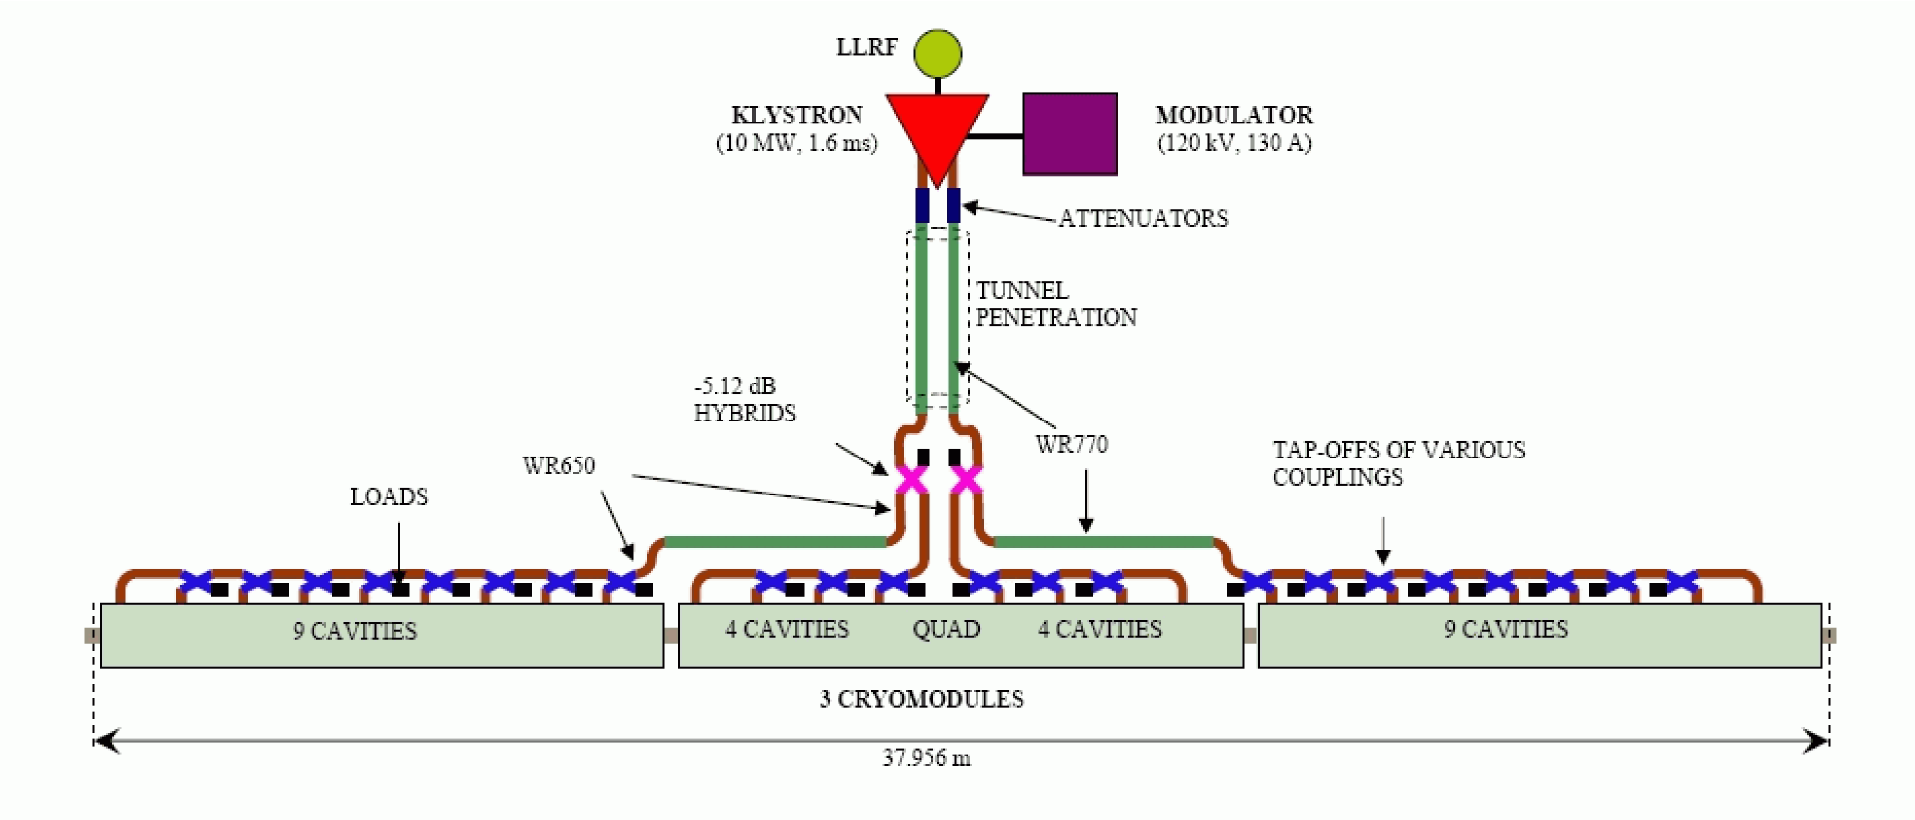

We use 278 of these in the positron linac, and 282 in the electron linac, plus additional in the sources and the bunch compressors.

RF Unit and Beam Parameters

Parameter	Value
# Cavities	26
Avg Gradient	31.5 MV/m
Total voltage	852 MV
Peak RF Power	10 MW
Pulse Length	1.6 msec
Fill time	0.6 msec
Length	38.0 m
Phase	5º

Parameter	Value
Bunches / train	2625
Trains / second	5
Bunch spacing	369 nsec
Beam current	9.0 mA
Bunch charge	3.2 nC
Initial energy	15 GeV
Initial energy spread	**1.5%**
RMS bunch length	0.3 mm
$\gamma\varepsilon_{x,y}$	**8 µm x 20 nm**

Emittance Preservation

What about the ILC linac is fundamentally new, above and beyond any existing electron linear accelerator?

1. Its enormous length (just a matter of scale – a big linac is just a lot of copies of a short linac)
2. Its enormous energy gain (same as the length, no big deal)
3. Its cost (OK, that's a big deal, but we aren't going to talk about that today)
4. ***The very small emittances which must be preserved (very big deal)***

The vertical emittance in the ILC is about the same as the smallest emittance ever produced at any light source in the world.

That emittance has to be preserved without dilution through > 10 km of linear accelerator.

What makes that so difficult?

Transverse Wakefields

In yesterday's lecture we discussed the fact that the beam can excite all of the monopole modes (accelerating / decelerating) with its "pancake" of co-moving fields; and that this can have an effect within a single bunch.

The beam can also excite all of the dipole modes (deflecting) with that same "pancake". The difference is that the beam needs to have a nonzero dipole moment wrt the cavity axis – ie, it needs to have a transverse offset!

Similar to the monopole wakefield, the short-range transverse wake can be approximated:

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

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

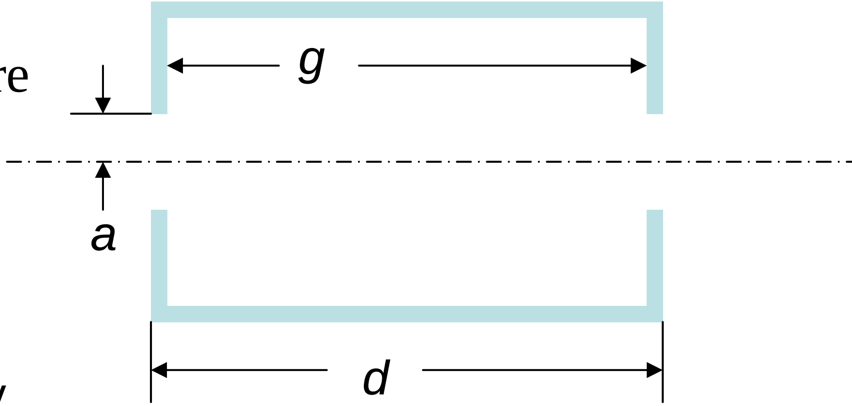

W in V/C/m^2 – need driving particle charge, cavity length, driving particle offset to get a voltage

Transverse Wakefield (2)

Properties of the transverse wakefield:

$W_\perp(0) = 0$ – a single particle does not deflect itself (different from monopole case)

The slope of $W_\perp$ at z=0 is given by: $W'_\perp(z=0) = \frac{2Zc}{\pi a^4}$

To first approximation, all that matters is the iris size, and it matters a lot! Note that, since the $W_\perp$ actually "rolls over" (ie, grows less than linearly) when the full expression is considered, a^{-4} is a bit of an overestimate of its strength – $a^{-3.8}$ is somewhat closer (not much of an improvement).

Since a is generally proportional to RF wavelength, low frequency cavities tend to have much weaker wakefields than high frequency cavities.

Beam Dynamics of the Transverse Short-Range Wakefield

Once again consider a 2-particle model with charge q and particle separation $2\sigma_z$

Trailing particle gets a kick given by $\Delta y'_2 = y_1 \frac{q}{2} L \frac{W_\perp\left(2\sigma_z\right)}{E_2}$

For a betatron oscillation down the linac, kicks on the trailing particle add coherently – when y_1 changes sign, so does y'_2, so $\Delta y'_2$ has same sign as y'_2 and so on…

Consider a beam with initial offset y_0 (both particles initially have the same offset). As it oscillates through the linac, the trailing particle gets a kick at each cavity given by:

$$\Delta y'_{2,c} = y_{1,c} \frac{q}{2} L \frac{W_\perp\left(2\sigma_z\right)}{E_c}$$

We can rewrite the first particle's offset at the cavity in terms of its initial offset:

$$\Delta y'_{2,c} = y_0 \sqrt{\frac{\beta_c}{\beta_0}} \sqrt{\frac{E_0}{E_c}} \cos \mu_c \frac{q}{2} L \frac{W_\perp\left(2\sigma_z\right)}{E_c}$$

SRWF BD (2)

Now project the motion of the trailing particle to the end of the linac – here we assume that the end of the linac has a 90º offset in betatron phase from the beginning, but that assumption just simplifies the math and does not affect the conclusion…

$$\Delta y_{2,f} = \Delta y'_{2.c} R_{34}\left(c \to f\right) = y_0 \beta_c \sqrt{\frac{\beta_f}{\beta_0}} \sqrt{\frac{E_0}{E_f}} \frac{qLW_\perp\left(2\sigma_z\right)}{2E_c} \cos^2 \mu_c$$

The final motion is given by the sum over all cavities, in which we can replace cos^2 with its mean value of ½:

$$y_{2,f} = y_0 \sqrt{\frac{\beta_f}{\beta_0}} \sqrt{\frac{E_0}{E_f}} \frac{qLW_\perp\left(2\sigma_z\right)}{4} \sum_{\text{cavities}} \frac{\beta_c}{E_c}$$

SRWF BD (3)

We want the resulting offset of the trailing particle wrt the leading one to be small compared to the size of the initial offset, y_0, when it is projected to the IP:

$$y_f = y_0 \sqrt{\frac{E_0}{E_f}} \sqrt{\frac{\beta_f}{\beta_0}}$$

Equivalently,

$$\frac{qLW_\perp\left(2\sigma_z\right)}{4} \sum_{\text{cavities}} \frac{\beta_c}{E_c} < 1$$

This provides a lower bound on the typical betatron functions which are acceptable in the linac. With q=3.2 nC, L = 1.04 m, $W_\perp(2\sigma_z) \approx 29{,}000$ GV/C/m^2, and $\Sigma(1/E_c) \approx 87.2$ GeV^{-1}, we find that we want

$$\left\langle \beta_c \right\rangle < 450 \text{ meters}$$

SRWF BD (4)

Of course, even if the beam goes perfectly down the axis of the accelerator, misaligned cavities can also cause wakefields which drive the tail's motion relative to the head.

In this case, the deflection of the tail by a cavity with misalignment y_c is given by:

$$\Delta y'_{2,c} = \frac{y_c L q W_\perp\left(2\sigma_z\right)}{2E_c}$$

Once again we project the resulting deflection to the end of the linac:

$$\Delta y_{2,f} = \frac{y_c L q W_\perp\left(2\sigma_z\right)}{2E_c}\sqrt{\beta_c \beta_f}\sqrt{\frac{E_c}{E_f}}\sin\Delta\mu$$

In this case the contributions from the different cavities will add incoherently, assuming a normal-distributed set of cavity misalignments with zero mean. We can characterize the mean-squared offset of the trailing particle:

$$\left\langle y_{2,f}^2 \right\rangle = \left\langle y_c^2 \right\rangle \frac{\left[L q W_\perp\left(2\sigma_z\right)\right]^2 \beta_f}{8E_f}\sum_{\text{cavities}}\frac{\beta_c}{E_c}$$

SRWF BD (5)

How much misalignment is too much? We can compare $\langle y_{2,f}^2 \rangle$ with the square of the beam size at the end of the linac, and insist that the wakefield deflection should be small compared to the beam size. The beam size is given by:

$$\sigma_y^2 = \varepsilon_y \beta_f = \gamma \varepsilon_y \beta_f \frac{m_e c^2}{E_f}$$

So we can cancel a bunch of factors to find our criterion:

$$\left\langle y_c^2 \right\rangle \frac{\left[L q W_\perp \left(2\sigma_z \right) \right]^2}{8 m_e c^2} \sum_{\text{cavities}} \frac{\beta_c}{E_c} < \gamma \varepsilon_y$$

In the limit where the LHS and the RHS are equal, the resulting emittance growth is 25%.

For typical betatron functions of 450 meters, we can tolerate about 480 µm of RMS cavity misalignment based on this. For stronger focusing (smaller beta) the tolerance gets looser.

Pitched RF Cavities

In addition to a translational offset, the cavities can have rotational offsets in the yz plane – nonzero pitch misalignments.

The effect of this is to project the fundamental mode acceleration into the transverse plane, resulting in a vertical kick:

$$\Delta y'_c = \sin\psi_c \frac{V_c}{2E_c}\cos\left(z\frac{\omega}{c}+\varphi_c\right)$$

The factor of 2 in the denominator is due to the end-field focusing of the cavity’s fringe fields – it takes out exactly 50% of the effect of the cavity pitch.

There are two things that come out of this – time-dependent deflections (the head and the tail see different kicks) and dispersion (particles with different energy get different kicks).

Time-Dependent Kicks

The lowest-order dependence of the voltage on the particle arrival time is given by:

$$V'(z) = -V_c \frac{\omega}{c} \sin\left(z\frac{\omega}{c} + \varphi_c \right)$$

The increase in the RMS angular spread of a beam with RMS length σ_z is given by:

$$\Delta\sigma_{y'} = \sin\psi_c \frac{V_c}{2E_c} \frac{\omega\sigma_z}{c} \sin\varphi_c$$

Once again we can project this to the end of the linac and require that the resulting beam size growth be small compared to the nominal beam size:

$$\sigma_z^2 \frac{\omega^2}{c^2} \frac{\left\langle \psi_c^2 \right\rangle}{8} \frac{V_c^2}{m_e c^2} \sin^2\varphi_c \sum_{\text{cavities}} \frac{\beta_c}{E_c} < \gamma\varepsilon_y$$

Substituting numbers, we find that the time-dependent kick tolerance on the cavity pitch angle is about 2 milliradians. This does not seem like much of a problem for the ILC.

Dispersion from RF Kicks

When the beam gets kicked by a pitched RF cavity, the low-energy particles get a larger kick and the high-energy particles get a smaller kick. This results in a growth in the projected emittance.

However – we are going to ignore this effect! This is because there is a much bigger effect in the linac which is dispersive in nature.

When a beam undergoes a betatron oscillation, the low-energy particles and high-energy particles do not oscillate with the same frequency. This results in a growth in the beam size due to its energy spread.

How big an effect is this? For an initial offset of $n\sigma_{y,o}$, if allowed to propagate down an infinitely-long linac, the emittance growth at the end will be given by a factor of $\sqrt{1+n^2}-1$

For a 1 sigma oscillation, the ultimate emittance growth is 41%.

This suggests that we need to limit the oscillation to ~1 σ_f at the end of the linac. This limit turns out to be too restrictive, but we can use it any to get a rough and naïve estimate of the size of the problem.

Dispersion from RF Kicks

Doing the usual transformation of the centroid kick to the end of the linac, we find:

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

For ILC parameters, and a beta function of 450 meters, the resulting tolerance for the cavity pitch is 1.4 microradians.

Even if we have gotten our estimate too tight by a factor of 10, this is a problem!

Once again, a stronger lattice (smaller beta) helps.

Misaligned Quadrupoles

Misaligned quadrupoles are very similar to pitched RF cavities – they deflect the beam, causing emittance growth through the *chromatic filamentation* of the beam oscillating down the linac.

We can skip directly to the (intermediate) result for cavity misalignments, based on our experience with other effects:

$$\frac{\sigma_q^2}{2m_e c^2} \sum_{\text{quads}} \left(K_q L_q\right)^2 \beta_q E_q < \gamma \varepsilon_y$$

Notice that in this case, the biggest problem is the quads at the high-energy end of the linac. That's because the focusing strength of the quads is fixed, but the geometric emittance at the high-energy end is small; so the same kick angle is a larger fraction of the beam angular divergence at the high-energy end.

For the cavities, the low-energy end is more problematic because the low-energy beam gets kicked more by a given voltage than a high-energy beam would be.

Misaligned Quads (2)

Let us not forget that K_qL_q and β_q are correlated – to get the beta functions down you need stronger quads! When we include that correlation, we see a somewhat different relation:

$$\frac{\sigma_q^2}{2m_e c^2}\frac{1}{\beta_t}\frac{1}{\cos^2\left(\mu_{\mathrm{cell}}/2\right)}\sum_{\mathrm{quads}} E_q < \gamma\varepsilon_y$$

This relation is still not entirely satisfactory, since β_t and μ_{cell} are correlated, but it does show that for quad misalignments we want to have the lattice as *weak* as possible.

As we saw earlier, for the ILC that's about 450 meter beta functions, corresponding to a phase advance per cell of 10º.

For those parameters, this relation tells us that we need to hold quad misalignments to about 0.5 μm. If we want a stronger lattice, then it gets worse.

The ILC Main Linac Lattice

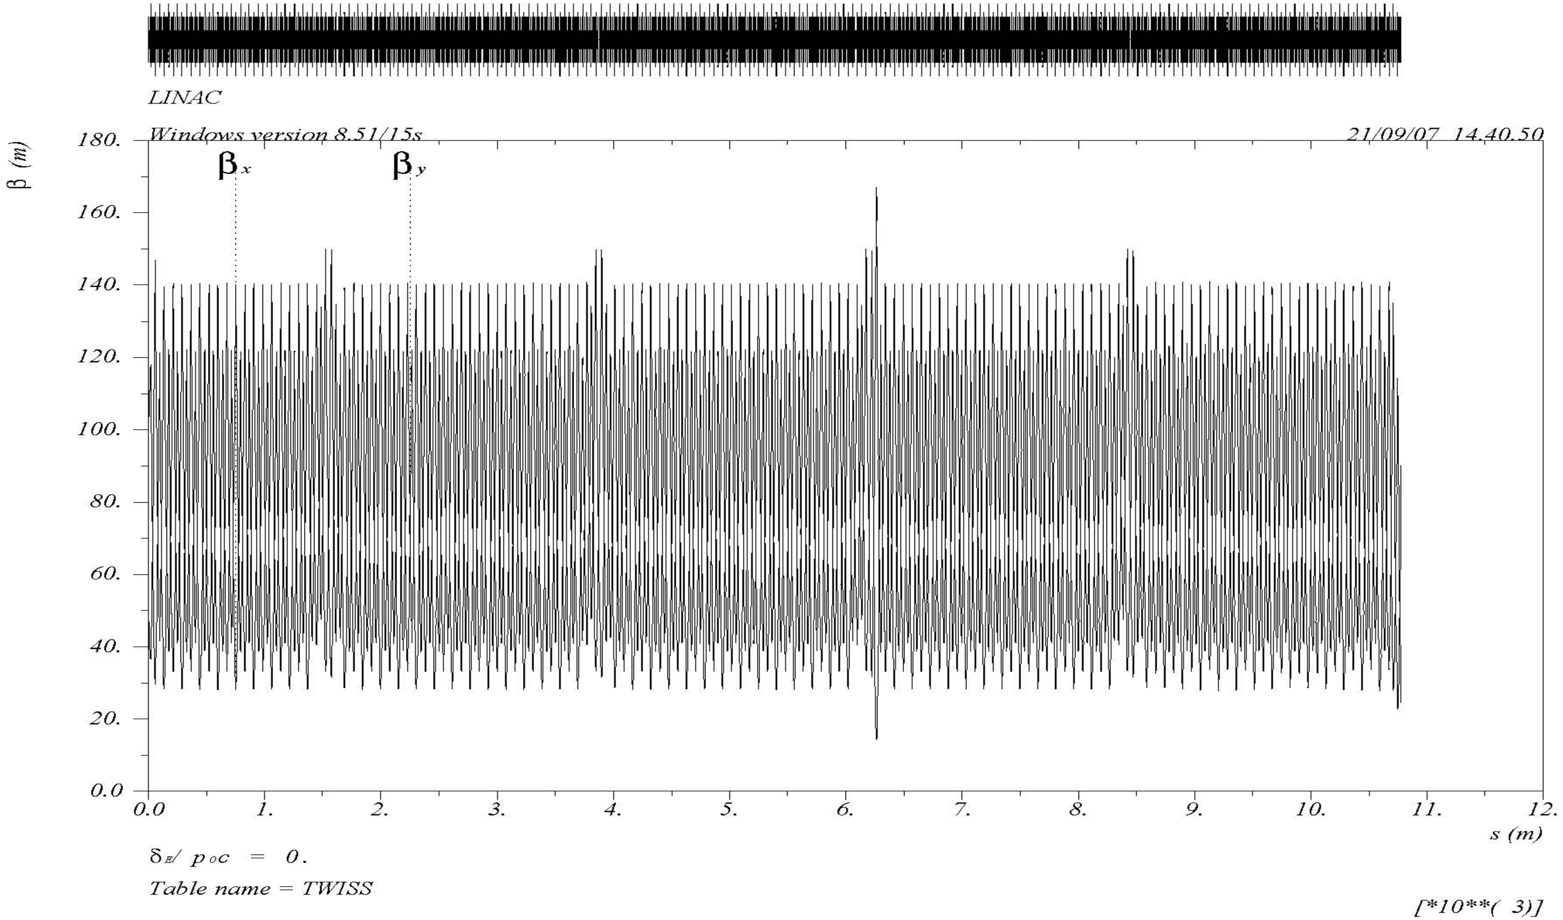

Linac Lattice (2)

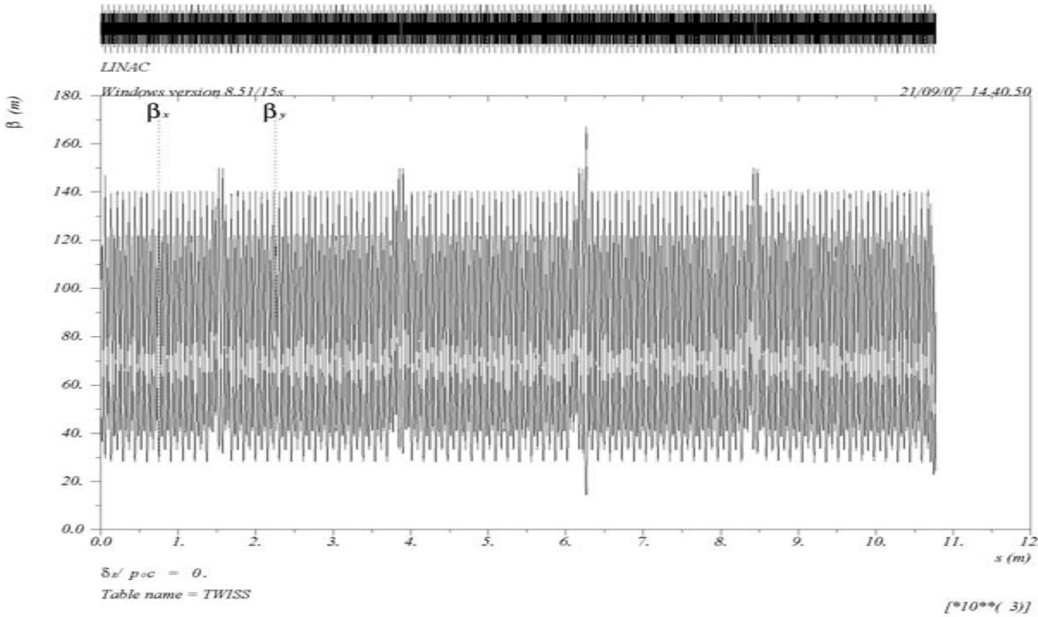

Phase advance per cell = 75º x, 60º y

Typical β ~ 80 meters

Beta "bumps" are at points where the quad spacing is not uniform – needed to accommodate cryo fluid distribution system

Linac is actually curved in vertical (follows Earth's curvature) so there's small vertical dispersion throughout the linac

What sort of tolerances do we get for this optics?

Cavity misalignment – around 1.4 mm

Cavity pitch (time-dependent kick) – around 4.6 mrad

Cavity pitch (dispersion) – around 3.3 μrad

Quad misalignment – around 0.2 μm

At least it's clear where the problems are…

Steering and Alignment

The main problems are all due to the kicks to the central trajectory, so orbit correction schemes can in principle correct these problems!

Simplest correction: use the dipole correctors to steer to the BPM centers – make the BPMs read zero. Since the number of correctors = the number of BPMs, this is called “one-to-one” correction (also sometimes 1:1 or 121).

If the BPMs are perfectly aligned to a straight line, then this will reduce the emittance growth to almost zero.

If the BPMs are misaligned, it will leave in some emittance growth.

Have we traded one impossible problem (aligning the cavities and quads) for another (aligning the BPMs)?

1:1 Steering

Imagine that all the BPMs are perfectly aligned except for 1. We will use the correctors upstream and downstream of the BPM, and the one at the BPM, to put in an orbit bump to get the beam through the BPM center.

The resulting orbit error at the last corrector for an off-energy particle is:

$$\Delta y' = \delta y_{\mathrm{BPM}} K_q L_q$$

We can compute the alignment tolerance by assuming uncorrelated BPM offsets, and requiring that the offset of an off-energy particle be comparable to or smaller than the nominal beam size. We can assume that at any point the RMS energy spread is given by the adiabatically-damped energy spread from linac injection (ie, that the linac adds no net energy spread, which is about right for the ILC)

$$\sigma^2_{\mathrm{BPM}} \frac{\sigma^2_{\delta,0} E_0^2}{2 m_e c^2} \frac{1}{\beta_t} \frac{1}{\cos^2\left(\mu_{\mathrm{cell}} / 2\right)} \sum_{\mathrm{quads}} \frac{1}{E_q} < \gamma \varepsilon_y$$

Weak focusing is still good, but now the low-energy end (where energy spread is largest) is worse than the high-energy end. This criterion sets the BPM alignment tolerance to about 85 μm – much better than without steering, but still too tight!

Kick Minimization (KM)

Consider a linac where the beam follows a perfectly straight trajectory (which would eliminate emittance growth from dispersion), where the quads are misaligned (but not the cavities), and the BPMs are aligned to the quads (but not the survey axis).

The beam passes off-axis through the quads (since they are misaligned), but it isn’t kicked because the correctors at the quads exactly cancel the kicks from the quad misalignments.

Since the BPMs are aligned to the quad centers, there is a correlation between the BPM readings and the corrector settings. Specifically, at each quad package,

$$y_{\mathrm{BPMreading}} = -\frac{\theta_{\mathrm{corrector}}}{K_q L_q}$$

If we steer to achieve the criterion above at each package, then there will be no net kicks, no dispersion growth.

KM (2)

In practice, KM solution to steering is unstable! Imagine that the BPM-to-quad alignment is not exact, but has some small error. This will lead to an error in setting the corrector at that location, which leads to a big orbit error at the next quad, which leads to a big corrector setting there, and so on.

Need to simultaneously constrain both the absolute BPM readings and the kicks:

$$\chi^2 = \underbrace{\sum_{\text{BPMs}} \frac{y^2_{\text{BPMreading}}}{\sigma^2_{y,\text{BPM}}}}_{\text{1:1 steering}} + \underbrace{\sum_{\text{BPMs}} \frac{\left(y_{\text{BPMreading}} + \theta / K_q L_q\right)^2}{\sigma^2_{y,\text{BTQ}}}}_{\text{Kick Minimization term}}$$

If the BPM-to-quad alignment is poor, KM doesn't do much. How do we make it good?

Quad Shunting

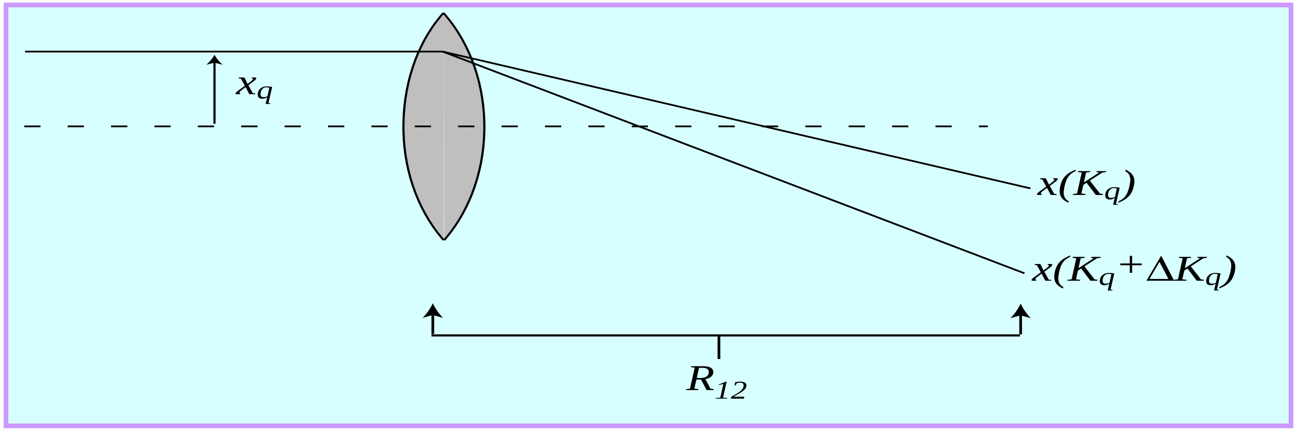

Varying strength in quad with beam off-axis changes the orbit – can estimate beam-to-quad offset. BPM-to-quad offset easy to get at that point.

Ultimate limit: quad center can move when quad strength is changed:

$$\Delta y_{\text{fit}} = \Delta y_{\text{center}} \left(\frac{K_q}{dK_q} + 1 \right)$$

For good quads (iron-core, small bore), estimated that we can get Δy_{fit} down to a few μm. For current-dominated SC quads, don't know.

Cavity Pitches in KM

Assumption that quads are the sole source of orbit kicks is implicit in KM algorithm design.

What about cavity pitches? Each corrector needs to steer out effect of pitched cavities between the corrector and the next BPM. Those pitches are unknown, so correlation between BPM reading and corrector is weakened – need to change chisq expression.

Research into this is ongoing…

Dispersion Free Steering (DFS)

This method makes a direct measurement of dispersion and applies a semi-local correction to cancel it.

Dispersion measurement: best method is to vary the energy gain in the linac (turn RF units on/off, rephase, lower the gradient) – really looks at what you are interested in (ie, the change in orbit with energy). Be careful – pitched RF cavities will change the orbit when they are adjusted to change the energy!

Alternately, can vary the quad/corrector strength (mismatches beam energy to lattice, equivalent to varying energy). This method only looks at dispersion from quads and correctors (neglects pitched cavities, stray fields, etc), and can introduce systematic errors (the quad centers can move when you vary them, as we discussed earlier).

Since there is vertical dispersion, need to actually include this in the model – so it’s more like “dispersion-matched steering”. Implies that BPM scale factors can have an impact on the result (not a huge one).

DFS (2)

As with KM, just looking for a set of corrector settings which zero the measured dispersion is not stable in presence of errors – need to simultaneously minimize the absolute orbit on BPMs and the dispersion:

$$\chi^2 = \sum_{\mathrm{BPMs}} \frac{y^2_{\mathrm{BPMreading}}}{\sigma^2_{y,\mathrm{BPM}}} + \sum_{\mathrm{BPMs}} \frac{\left[\Delta y_{\mathrm{meas}} - \Delta y_{\mathrm{model}}\left(\vec{\theta}\right)\right]^2}{2\sigma^2_{\mathrm{res}}}$$

1:1 Steering

DF portion – relies on BPM resolution

Ballistic Alignment (BA)

Emittance growth is associated with electromagnetic fields in beamline components. What happens if we just get rid of those fields – turn off the magnets and the cavities?

Beam should follow a straight line!

Measure BPM readings of the “straight line” trajectory, then turn everything on and steer back to that orbit.

This is known as ballistic alignment.

Issues with BA:

1. Getting the beam down all or part of the linac with everything turned off
2. Did everything really turn off, or are there remnant fields?
3. What about stray fields? What’s the impact of the Earth’s field on the beam over 10 km, especially with the beam at low energy (15 GeV) all that way? Etc.

Global Corrections

Emittance dilution results from uncontrolled buildup of small effects (wake kicks, dispersion, etc) from unknown sources throughout the linac.

We can deliberately introduce an “error” – dispersion, for example – in a controlled way. As we saw, a closed orbit bump thru a quad introduces dispersion, so we can do that.

If the dispersion introduced with the bump is in phase with some of the emittance growth sources, it can either correct them or make them worse; so vary the amplitude of the dispersion bump and measure the emittance. Use 2 bumps, 90 degrees apart, and you can in principle correct all the emittance growth from dispersion.

In practice, dispersive emittance growth can filament through the linac, so the bump can only correct errors which are within a few betatron wavelengths of the bump.

Other effects limit effective correction range of bumps, so some more local (steering) solution is needed first.

Wake bumps are possible in principle, but are much more complicated to design.

Long-Range Transverse Wakefields

Transverse wakefields are dipole-mode equivalents of longitudinal wakefields – caused by beam excitation of the various eigenmodes of the cavity.

Over a short (intra-bunch) time, there are a near-infinite number of transverse modes which have a transient excitement which deflects the tail of the beam.

Over a longer (inter-bunch) period, most of the very high-frequency modes die off (remember that the cavity resistance α f^2). But many modes at lower frequencies are still present, esp. frequencies around the fundamental (1.3 GHz).

Result: trailing bunches can be deflected by fields left behind by leading bunches.

If the Q for the low-frequency dipole modes is as high as the Q for the fundamental mode (10^6 – 10^{10}) then the field sticks around for the whole bunch train – big trouble!

Higher Order Modes

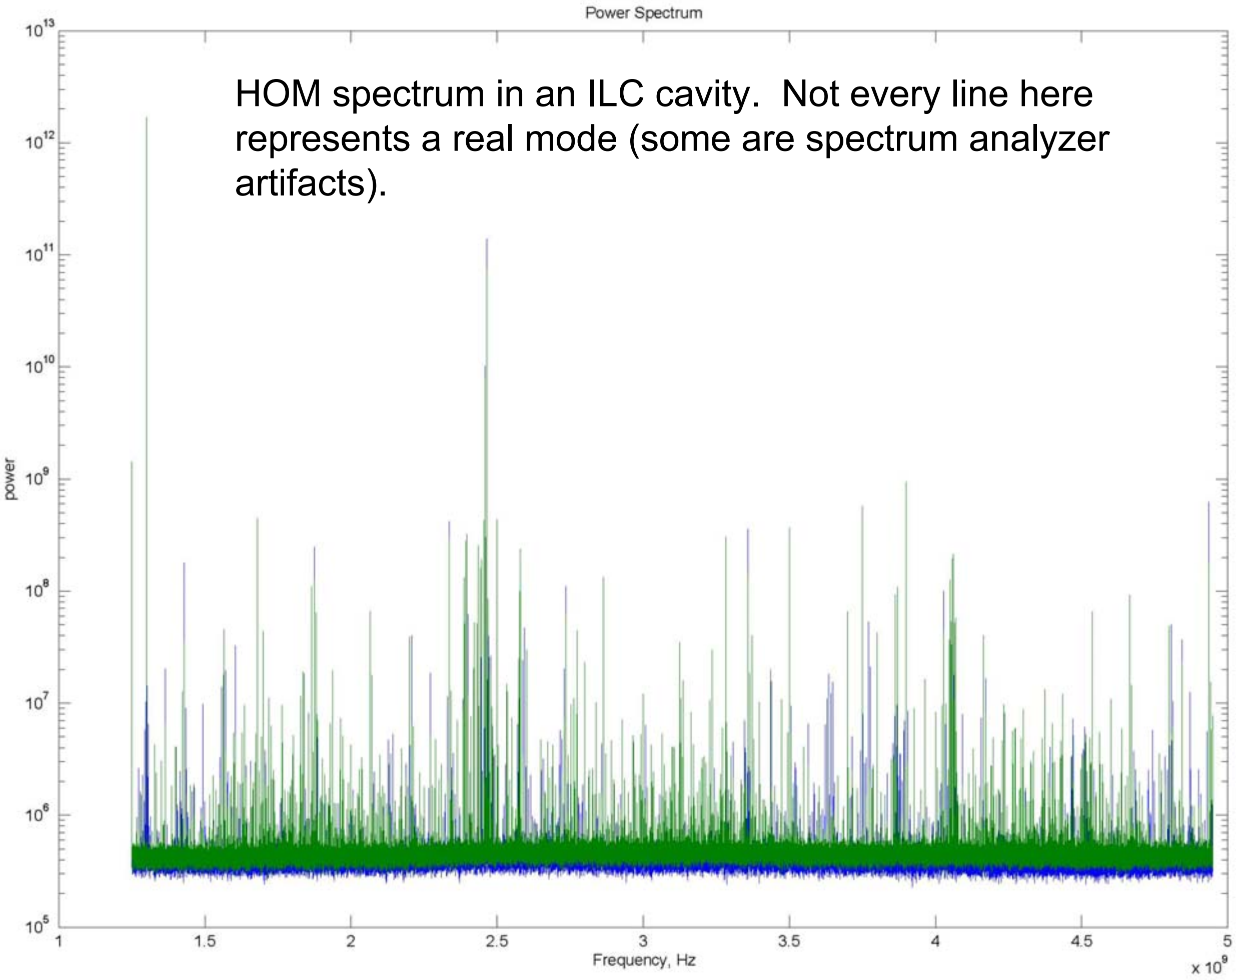

HOM spectrum in an ILC cavity. Not every line here represents a real mode (some are spectrum analyzer artifacts).

Managing HOMs

Most obvious approach to controlling the multi-bunch emittance growth from HOMs = damping – reducing the Q to less dangerous levels.

Damping down to Q ~ $10^3 \rightarrow$ mode damps between bunches! Ideal! But tough to do…

Compromise: cavity damps lowest-frequency modes to Q ~ $10^3 – 10^5$. Wakes "last" for 100 or so bunch times (or fewer).

Implies that wake kick enters "steady state" after about 100 bunches – bunches far behind bunch 1 get kicks, but they all get the same kick (no multi-bunch emittance growth).

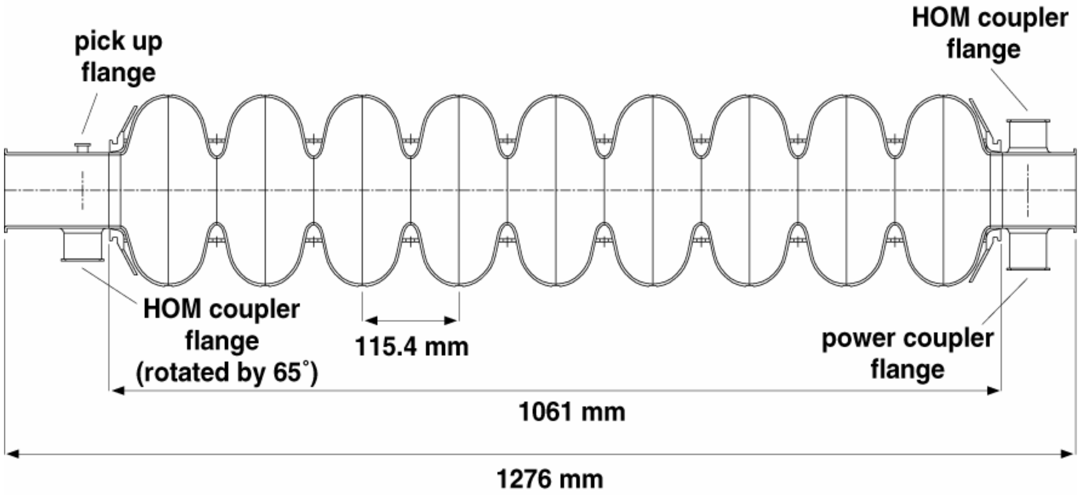

Many dangerous modes are at frequencies which are not cut off by cavity irises – they can propagate down the linac! Additional "broadband" dampers at every quad in the linac.

Managing HOMs (2)

HOMs are mainly dangerous because their effects are *coherent* – if bunch 100 gets a big kick in cavity 1, it also gets it in every cavity which is at the same betatron phase as cavity 1.

Additional tool – make the wakes different at all the cavities via *detuning*.

When summed over all cavities, detuning has the effect of “broadening” the HOM lines

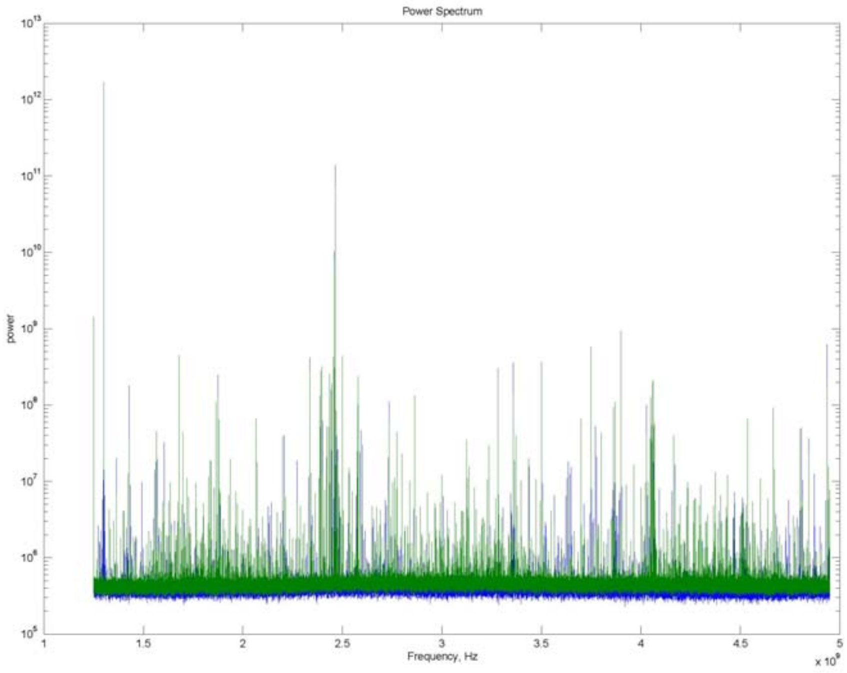

Fourier transform of a broad line = sharp spike in time domain, so total deflection of trailing bunches reduced.

For finite # of modes, “recoherence” effect – detuned modes get back into phase with each other after some time.

Fortunately, with $Q < 10^5$, wakes are pretty weak by the time they recohere.

More on Detuning

How much detuning is needed? Studies indicate that ~10^{-3} detuning should be enough (ie, around 2 MHz frequency errors for 2 GHz HOMs)

In theory, should be able to get that from natural construction errors!

In practice, cavities may be too reproducible!

This was a problem at the SLAC linac in the early 1960's – there was not enough frequency spread in the initial linac design to produce the necessary HOM detuning.