CBN 98-18

Simulation of Multibunch Longitudinal Instabilities Using the *OSCIL* Code

Daniel B. Fromowitz

Wilson Lab, Cornell University, Ithaca, NY 14853, USA
(e-mail: dbf5@cornell.edu)

December 15, 1998

A tracking simulation shows current thresholds near the values which are presently seen in CESR. Increases in the threshold are predicted each time a copper RF cavity is replaced with a niobium cavity given the dynamics included in the code. Growth rate and Fourier spectra of sample beams are presented.

Contents

1 Introduction

The *"oscil"* code [1] simulates longitudinal motion in the CESR storage ring. The *oscil* simulation is used to answer the question: Can the current thresholds measured in CESR be understood as a function of EM wake field modes in the RF cavities?

Section 2 introduces the physics incorporated into the simulation including how cavity wake fields influence the beam. Section 3 briefly mentions some basic phenomena that *oscil* correctly shows which serves as a check on the implementation of the physics in the code. Section 4 predicts CESR beam current thresholds during the course of the transition from five–cell copper cavities to single–cell superconducting cavities (part of the phase III upgrade), presents a calculation of longitudinal oscillation growth rates, and explains why certain spectral lines indicate an instability.

2 Dynamics

2.1 Path Length Change

The longitudinal displacement (long. disp.) and energy displacement (en. disp.) of the circulating particles are coupled as follows:

$$\Delta C = \frac{\alpha C(\Delta E)}{E} \tag{1}$$

where α is the momentum compaction factor, C is the circumference of the machine, ΔE is the en. disp., and E is the average energy. The definition of ΔE in eqn. 1 assumes that the ideal particle has a constant energy between cavities. However, the definition used in *oscil* allows ΔE to vary between cavities as explained in section 2.2. Using the initial value (ΔE_i) of ΔE at a particular cavity, the change in the long. disp. at the next cavity is

$$\Delta(\text{long. disp.}) = -\frac{\alpha L(\Delta E_i)}{E} + \frac{\alpha L^2 S}{2EC} \tag{2}$$

where Δ(long. disp.) is the change in the long. disp., L is the intercavity distance traveled, and S is the amount of synchrotron radiation per turn per particle. The change in sign of the first term is due to the difference between the long. disp. and the path length as shown in section 2.2. Note that the last term is a correction to take into account synchrotron radiation.

Details of the simulation may be found in Appendix A.

2.2 RF Drive

The change in the en. disp. due to the drive voltage waveform is modeled by the equation

$$\Delta(\text{en. disp.}) = eV_{RF}\sin(\omega_{RF}\cdot \text{long. disp.}/c) \tag{3}$$

where Δ(en. disp.) is the change in the en. disp, e is the charge of the positron, V_{RF} is the cavity drive voltage amplitude due to the klystron, ω_{RF} is the klystron drive angular frequency, and c is the speed of light.

The change in the long. disp. (as measured at a cavity) is negative the change in the path length. A long. disp. of zero means the particle will get no kick from the *drive* voltage.

The value ΔE is defined to be zero for a particle that begins with the design energy and emits no synchrotron radiation. Therefore, an actual particle will continually have its value of ΔE *fall* during intercavity flight. Subsequently, a real particle will usually acquire a negative ΔE by the time it reaches a cavity. Over time, eqns. 2 and 3 together determine the synchrotron motion of the particles.

2.3 Cavity Wake Fields

Wake fields in the cavities also couple the long. disp. and the en. disp. The energy kick given to a positron by mode n is [2]:

$$\Delta(\text{en. disp.})_n = -eq_n\omega_n\left(\frac{R}{Q}\right)_n cos(\omega_n t)\exp\left(\frac{-\omega_n t}{2Q_{L,n}}\right) \tag{4}$$

where q_n is the charge of the bunch that created mode n, and t is the time elapsed since mode n was created (related to the long. disp.). Corresponding to mode n are the frequency ω, total impedance R/Q, and loaded quality factor Q_L. The parasitic modes excited by previous passes of all particles must be considered. The total Δ(en. disp.) is the sum of all such wake fields, i.e.

$$\Delta(\text{en. disp.}) = \sum_n \Delta(\text{en. disp.})_n \tag{5}$$

where the summation is for all passes of all particles at all parasitic mode frequencies, *including* the fundamental mode.

The total impedance of all cavities in the fundamental mode is defined based on the following formula for the average cavity power P_{ave}:

$$P_{ave} = \frac{V_c^2}{2\left(\frac{R}{Q}\right)Q_L} \tag{6}$$

where V_c is the (time) average total cavity voltage amplitude and all quantities refer to the fundamental mode only.

2.4 Cavity Voltage Regulation

In order to simulate the RF cavities, it is necessary for *oscil* to keep the total cavity voltage amplitudes consistent with the CESR cavity drive regulation, i.e. the amplitudes should be independent of the beam current. (See section A.3 of Appendix A.) The cavity voltage phasor is a sum of the drive voltage phasor and the beam–induced wake field voltage phasor. Only the fundamental mode is observed in this cavity regulation process.

The drive voltage is near a perfect sinusoid. However, the wake-field voltage fluctuates due to the spacing and varying phases of the bunches in the beam. The regulation process takes into account two conditions that must be satisfied in each cavity. First, the sum of the real parts of the drive voltage phasor and the wake field phasor must be equal to the magnitude of the synchrotron radiation lost times the relative fraction of power delivered to beam by the cavity in question. This equivalence is a statement of the conservation of energy. This relationship is shown graphically in Figure 1. The cavity voltage due to the klystron is $\tilde{V}_{\text{kly}}$, and its magnitude is the V_{RF} from eqn. 3. The wake field phasor (for the fundamental mode) is $\tilde{V}_{bl}$ and takes into account the fundamental theorem of beam loading due to the present particle. $V_{\text{synch rad}}$ is the amount of synchrotron radiation for which the particular cavity must compensate.

Figure 1 only applies at the instant a particle passes through the cavity. The simulation calculates the real part of $\tilde{V}_{\text{kly}}$ from the other information given.

Using the (time) average total cavity voltage V_c, the (fundamental mode) wake field phasor $\tilde{V}_{\text{wake}}$ (which does not include the fundamental theorem of beam loading because only voltages between successive particle passes are considered), and the real part of $\tilde{V}_{\text{kly}}$ just calculated, the imaginary part of $\tilde{V}_{\text{kly}}$ is determined from the geometry seen in Figure 2.

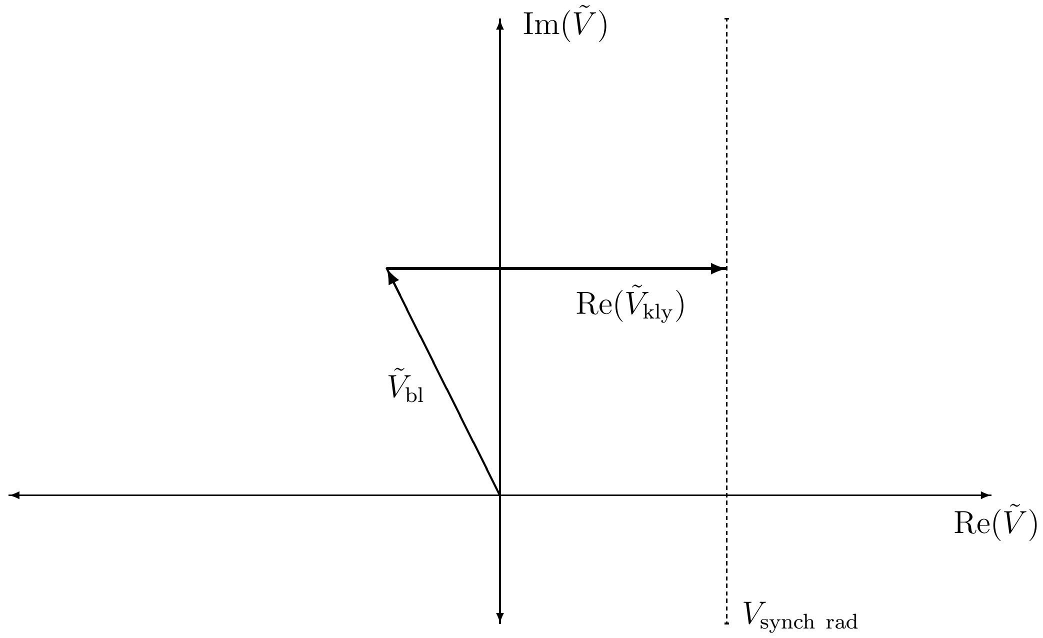

Figure 1: Calculation of the real part of the cavity voltage phasor due to the klystron ($\mathrm{Re}(\tilde{V}_{\mathrm{kly}})$) necessary to counteract the energy loss due to synchrotron radiation.

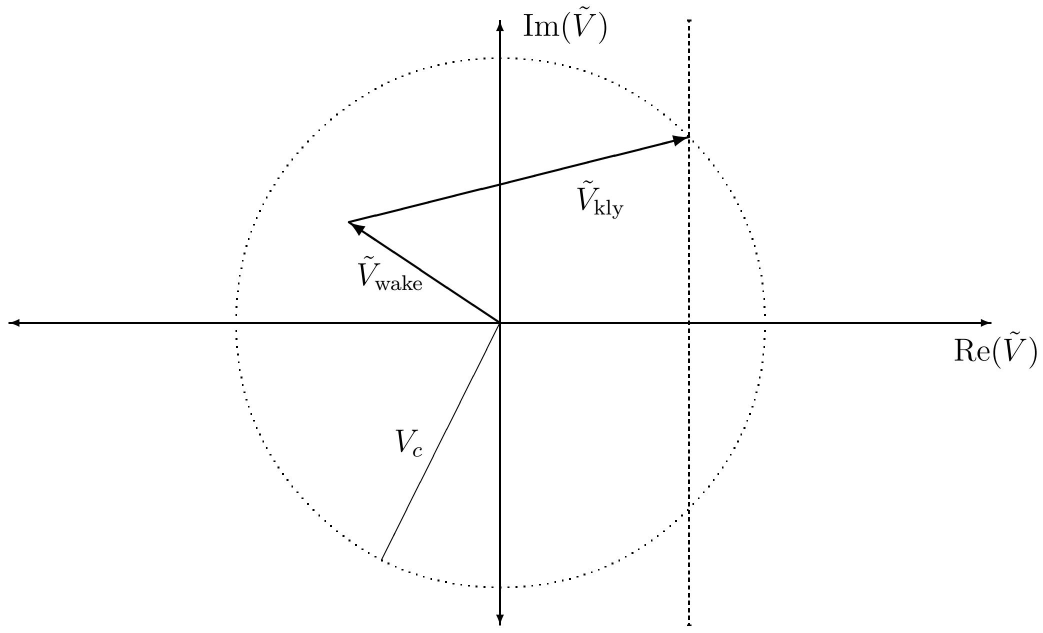

Figure 2: Calculation of the complete cavity voltage phasor due to the klystron ($\tilde{V}_{\mathrm{kly}}$).

Because the wake fields decrease in magnitude and rotate in phase between passes of successive particles, the phasor $\tilde{V}_{\mathrm{wake}}$ is the average wake field seen at the time the average particle traverses the cavity and at all subsequent multiples of the RF period thereafter.

The two phasors in Figure 2 must have a sum that lies on a circle of radius V_c. This relationship is the second condition that must be satisfied in the cavity regulation and is required because V_c in a real CESR cavity is kept near a constant value. The horizontal distance between the tail of the $\tilde{V}_{\mathrm{kly}}$ phasor and the vertical dotted line is the magnitude $\mathrm{Re}(\tilde{V}_{\mathrm{kly}})$.

The choice between the two intersection points of the circle and the phasor sum is made by making the sum (or $\tilde{V}_{\mathrm{kly}}$) have a positive slope in the complex plane for beam stability. As an example, if a particle has a high energy, it will reach the cavities later and later according to eqn. 2. The phasors in Figure 2 will rotate more clockwise and will therefore give less of a total energy boost to the particle that they would had $\tilde{V}_{\mathrm{kly}}$ been chosen to point to the lower intersection point. This is the situation we want.

3 Verified Phenomena

The *oscil* simulation reproduces key observed phenomena [1] in CESR including:

- The synchrotron frequency.
- The bunch length.
- The longitudinal radiation damping lifetime based on the amount of synchrotron radiation per turn per particle.

The simulation also shows characteristics of storage rings that are easy to verify theoretically [3] such as Robinson instability and damping.

4 Results

4.1 Longitudinal Oscillation Envelopes

A series of simulations were performed to find the current threshold as a function of ring and beam parameters. Most of the parameters were changed one at a time so that the effect of particular quantities could be understood.

First, it is necessary to show how the long. disp. depends on time. Often the envelope of the long. disp.—rather than all long. disp. data—provides relevant information.

A typical long. disp. envelope appears in Figure 3. The data were made from a phase II (4 Cu cavities) 9x2 configuration with a 42-ns (intratrain) interbunch spacing, pointlike bunches, 1.029 MeV/turn/particle of synchrotron radiation, and with the fundamental cavity mode impedance only, i.e. *no* higher–order modes). Although these conditions do not fully specify the beam (because of varying initial phase space distributions and cavity parameters such as voltage amplitude), Figure 3 does show a typical low–current long. disp. envelope corresponding to 72 mA of total beam current. Since each bunch is made of only one macroparticle, the long. disp. shown is by definition the center of mass of an analogous bunch in CESR. Notice that the beam is strongly damped.

What happens as the current is increased? At the higher current of 144 mA, the long. disp. amplitude grows exponentially in time up to about 0.11 s as seen in Figure 4. Then the beam "pulses" indefinitely. Note that what is seen is a "sloshing" back and forth of the

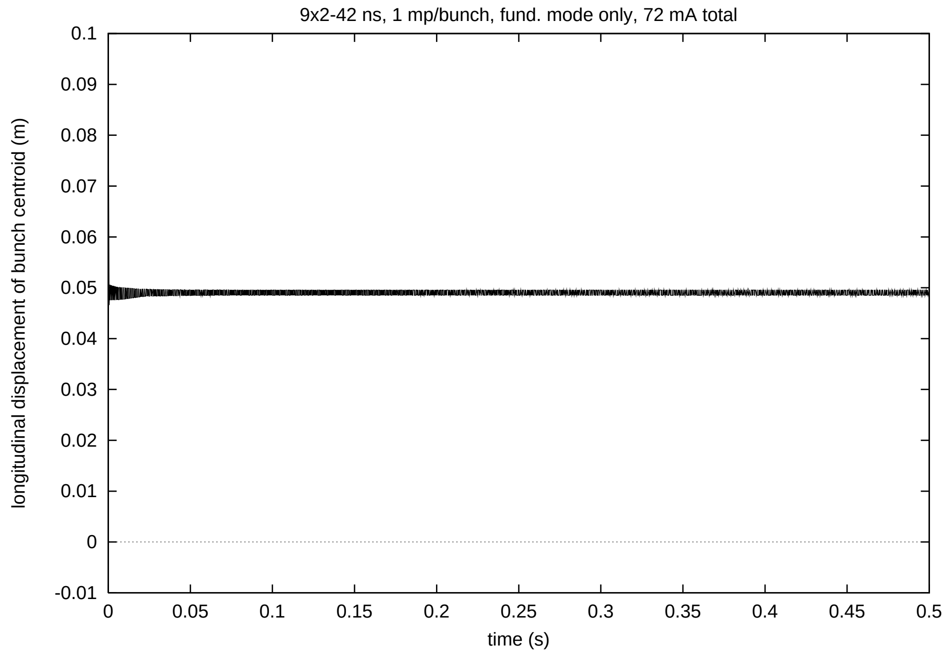

Figure 3: Envelope of longitudinal oscillations of 18 pointlike bunches at 72 mA total current. The interior of the envelope is darkened for clarity.

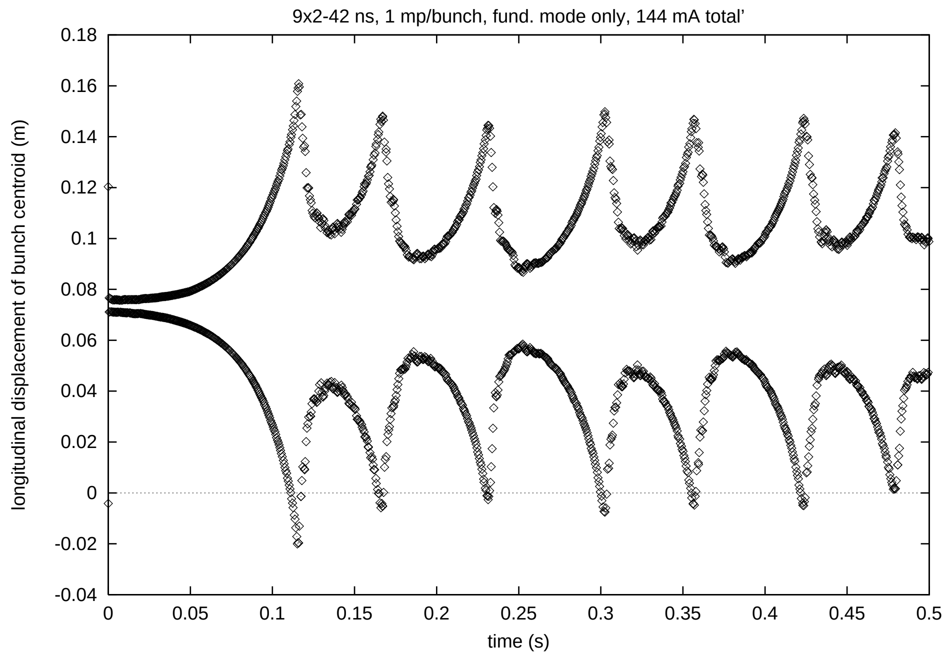

Figure 4: Longitudinal oscillation envelope at 144 mA.

bunch centroids and not of the longitudinal bunch size (since the bunches are pointlike). Similar pulses have been seen in the simulation ODYSSEUS [4].

At the even–higher current of 216 mA, the pulses vary chaotically in duration and height as seen in Figure 5. Remember that there are no higher–order mode impedances included; only the fundamental mode is modeled. At a slightly higher current, the beam is lost almost immediately, i.e. the long. disp. increases by several orders of magnitude. Note that the pulses begin sooner than in Figure 4 because the initial growth rate is higher when the current is higher.

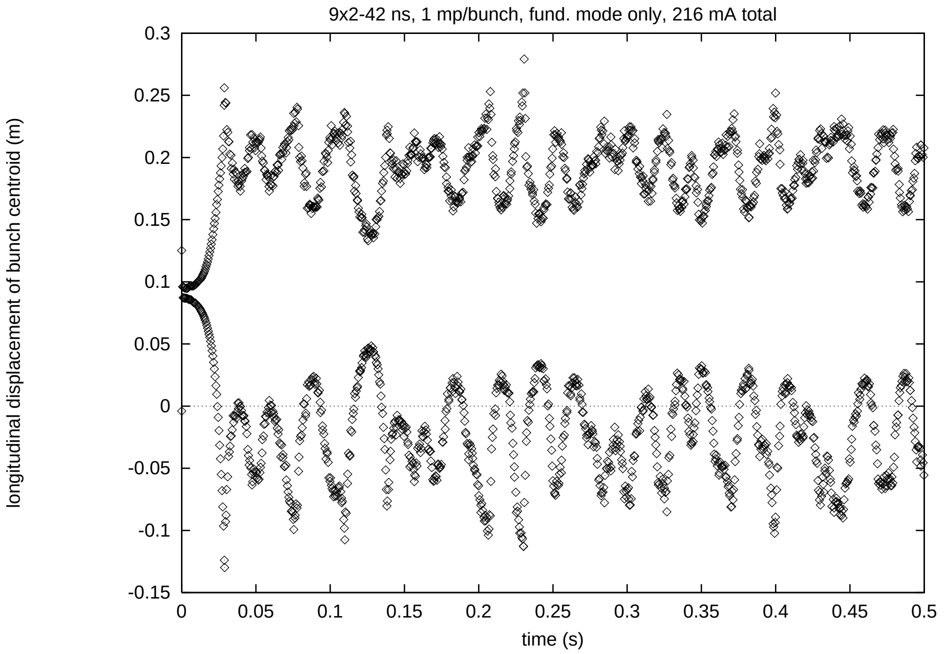

Figure 5: Longitudinal oscillation envelope at 216 mA.

Since bunches in storage rings are not made of only one particle but often around 10^{11} particles, we need to see what changes occur when the pointlike bunches are replaced with bunches containing many macroparticles. Perhaps Landau damping of the particles within a bunch can increase the current thresholds. *Oscil* treats each macroparticle as an independent entity. Therefore, the only interaction between macroparticles is indirect via the induced wake fields. With more than one macroparticle per bunch, the current threshold might be more accurately predicted.

Using 200 macroparticles per bunch, we get an envelope like that seen in Figure 6. Figure 6 is similar to the one macroparticle per bunch case in Figure 5, but the envelope does become smaller in places.

When plotting multiparticle bunches, the long. disp. is that of the center of mass, just as with pointlike bunches. Notice that the envelope is not always symmetric about its center. This effect is prominent around both 0.17 s and 0.47 s of Figure 6. The asymmetry is due to the cavity voltage regulation.

The bunch length may be found (for bunches containing more than one macroparticle) by computing twice the standard deviation of the displacement of each macroparticle from the bunch centroid. Figure 7 shows the bunch lengths that correspond to the data points in Figure 6. Note that the *envelope* of the bunch length is *not* shown; what is shown are the bunch lengths at the long. disp. envelope. Observe the correlations between the long. disp.

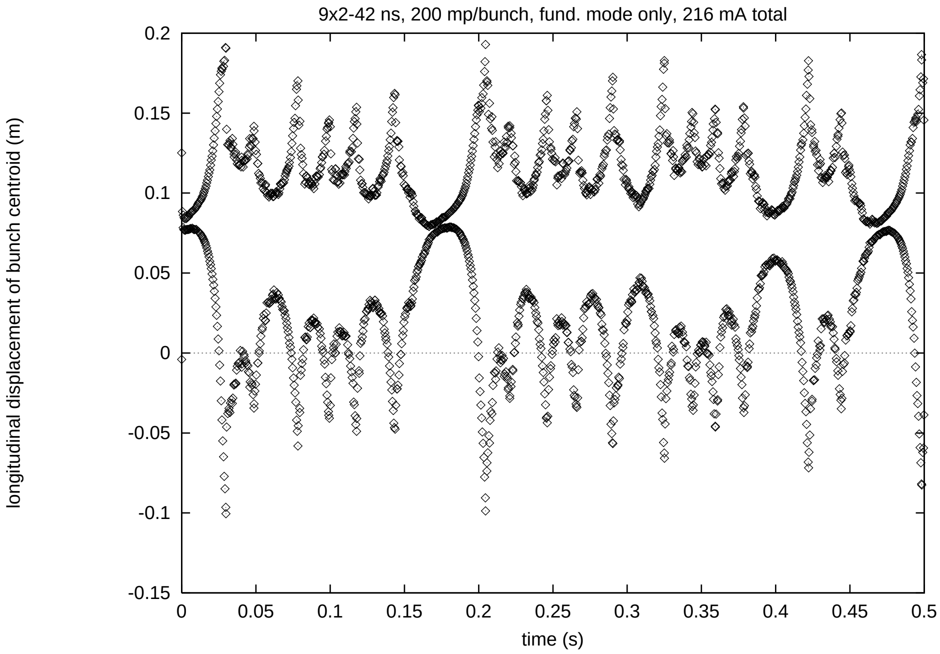

Figure 6: Longitudinal oscillation envelope at 216 mA for bunches made of 200 macroparticles each.

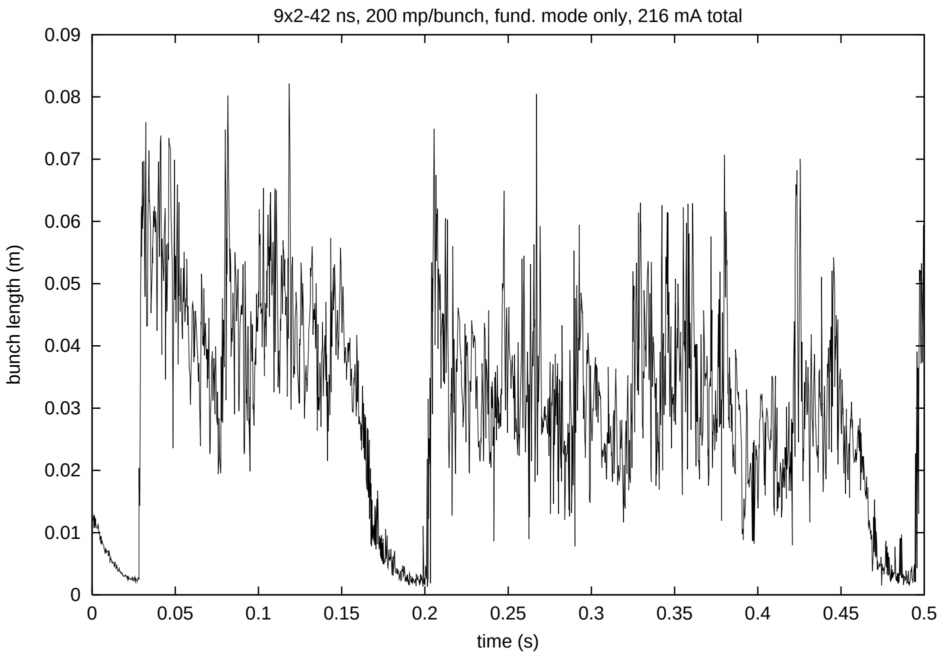

Figure 7: Bunch lengths along the longitudinal oscillation envelope shown in Figure 6.

and the bunch length when the long. disp. envelope is narrow. The correlation means that when there is not much longitudinal motion of the bunch centroid, the bunches are compact.

4.2 Current Thresholds

While plots of the long. disp. show differences in growth rates and behavior of beams with different parameters, we are interested in specifically locating the current thresholds. Two types of current thresholds are described and plotted. The lower–current threshold may be described by comparing Figures 3 and 4. In Figure 3, the longitudinal oscillations damp and remain damped. At the higher current in Figure 4, the oscillations first increase in magnitude then pulse indefinitely. There is a threshold between these two conditions, called the "damping-to-pulses" threshold in this report. Below the "damping-to-pulses" threshold, the oscillations always damp. Above the "damping-to-pulses" threshold, the oscillations always grow and then pulse.

The higher–current threshold is the boundary between rapid (and rapidly–varying) pulses as seen in Figure 5 and a higher current when the beam cannot be contained. This threshold is therefore coined the "pulses-to-divergence" threshold. Between the currents corresponding to the "damping-to-pulses" and "pulses-to-divergence" thresholds, the oscillations pulse. Above the "pulses-to-divergence" threshold, the oscillation growth rate is so high that particles of the beam leave their original RF bucket and are lost.

These two thresholds are plotted for a variety of cases, mostly using a 9x2-42 ns beam. In each of the remaining plots in this section, the current limits are plotted versus the number of SRF (niobium) cavities (with the total number of cavities constant at four). Thus zero SRF cavities represents the phase II case as in CESR until mid–1997, and one SRF cavity is the beginning of phase III upgrade as in CESR starting late 1997. (The full phase III scenario with four SRF cavities is not plotted for a reason explained in section A.3. of Appendix A.) The subsequent threshold plots allow one to see how the thresholds change as the phase III upgrade progresses, if the only parameters in the ring were the ones used in the particular simulation and plot.

Since the Cu and Nb cavities have different properties and thus affect the beam dynamics differently, it is helpful to recall the differences between the two type of cavities. See Table 1. These values are those used in the *oscil* simulation and are approximately the values in CESR.

Table 1: Differences in normal and superconducting cavities.

Quantity	Normal (Cu) Cavity	Superconducting (Nb) Cavity
fund. Q_L	6×10^3	2×10^5
fund. R/Q	426.762 Ω/m	145 Ω/m
# of cells	5	1
V_c	~ 1.5 MV	1.9 MV
detuning offset	$-5°$	$-30°$

The average cavity voltage V_c for the copper cavities varies with the number of SRF cavities based on readings seen in the CESR control room. The "detuning offset" is the detuning value at zero current that is added to the fundamental mode detuning value if the cavities are automatically detuned to minimize reflected power back to the klystron. See section A.4 of Appendix A for further information.

Furthermore, when a higher–order mode (HOM) is included in this report, all cavities have the same mode frequency, even if some cavities are copper and some are SRF. The

HOM frequency used is about 1400 MHz because the impedance is high for this mode in both types of cavities [5]. However, the values for the R/Q and Q_L are allowed to vary between the copper and SRF cavities. In CESR, it is impossible to "turn off" certain modes, but in this simulation, the effects of particular modes may be studied by including only those modes in question.

The sum of the cavity voltage amplitudes for all cavities used in these simulations is shown in Table 2. The total number of cavities is always four.

Table 2: Sum of time–averaged, fundamental-mode, total cavity voltage amplitudes.

Number of SRF Cavities	$\sum_{i=1}^{4} V_{c,i}$ (MV)
0	5.92
1	6.32
2	6.82
3	7.22
4	7.60

The amount of synchrotron radiation varies depending on whether the CHESS wigglers are open or closed. One would probably expect higher thresholds with the wigglers closed because there would be more radiation and therefore more damping. With the wigglers open, each particle loses 1.029 MeV per turn due to synchrotron radiation. With the wigglers closed, the figure rises to 1.158 MeV [6]. Figure 8 shows the damping-to-pulses thresholds for pointlike bunches with the fundamental mode (and no HOM's) for both wigglers open and wigglers closed. Figure 9 shows the analogous pulses-to-divergence thresholds. Note that the thresholds increase with the number of cavities that are superconducting.

As may be expected, the damping-to-pulses threshold (Figure 8) is increased when there is more synchrotron radiation (more damping). (Compare the upper and lower data sets.) However, the pulses-to-divergence threshold *decreases* with increasing radiation, but the effect is not significant. This effect draws attention to the complexity of the wake fields and the cavity regulation.

The error bars seen in the threshold plots exist for three reasons. First, they result from the finite resolution of bunch currents (usually integers) in successive simulations. Second, they also indicate a spread in threshold depending on initial conditions (long. disp. and en. disp.). Last, at high currents, some beams may diverge, but slightly higher–current beams may pulse, meaning that the pulses-to-divergence threshold is not well–defined.

When a HOM with a frequency of 1398.8 MHz is added to the cavities, the two thresholds change as seen in Figures 10 and 11. The effects of opening or closing the wigglers has similar effects as compared to when only the fundamental mode exists.

Comparing Figures 8 and 10, we see that the damping-to-pulses thresholds are significantly decreased by the addition of the HOM. This effect shows that the HOM is a "parasitic" mode, whose perturbation decreases the beam stability. However, the pulses-to-divergence thresholds are greatly *increased* by the addition of the HOM as seen in Figures 9 and 11. This effect is caused by nonlinearity of the cavity voltage caused by the HOM and its resulting Landau damping among the various bunches in the ring.

It is known from cell probe measurements that the HOM frequencies vary from cell to cell. Also, temperature variations affect the frequencies within a particular cell. In Figure 12, the damping-to-pulses thresholds are shown for two different HOM frequencies, one of 1398.8 MHz, and other of 1399.8 MHz—a change of only 1 MHz. There is a huge difference in these thresholds. Increasing the HOM frequency by 1 MHz increases the thresholds around

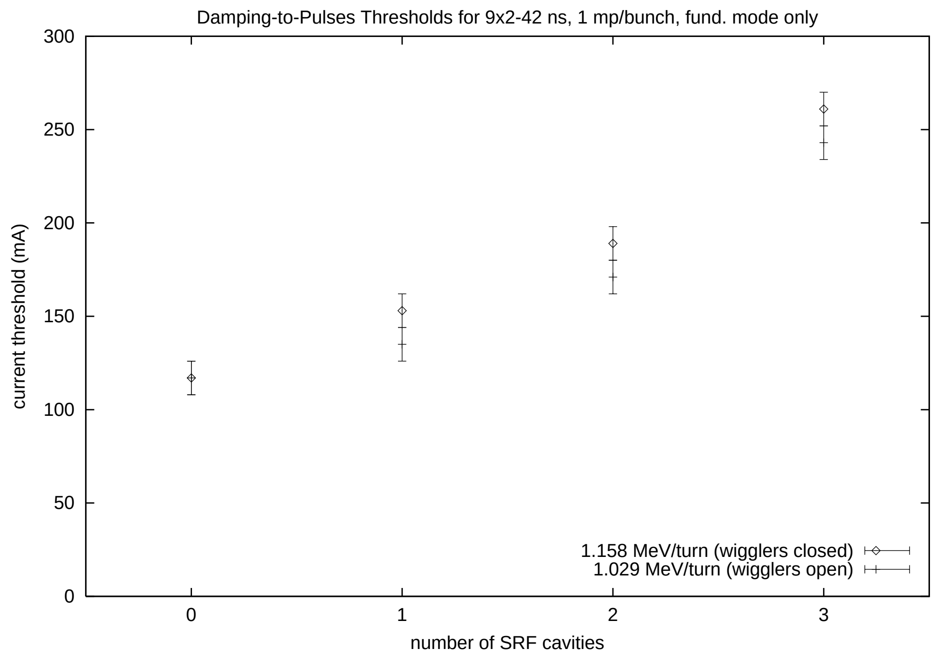

Figure 8: Damping-to-pulses current thresholds for the fundamental mode alone as a "function" (two different data sets) of the amount of synchrotron radiation.

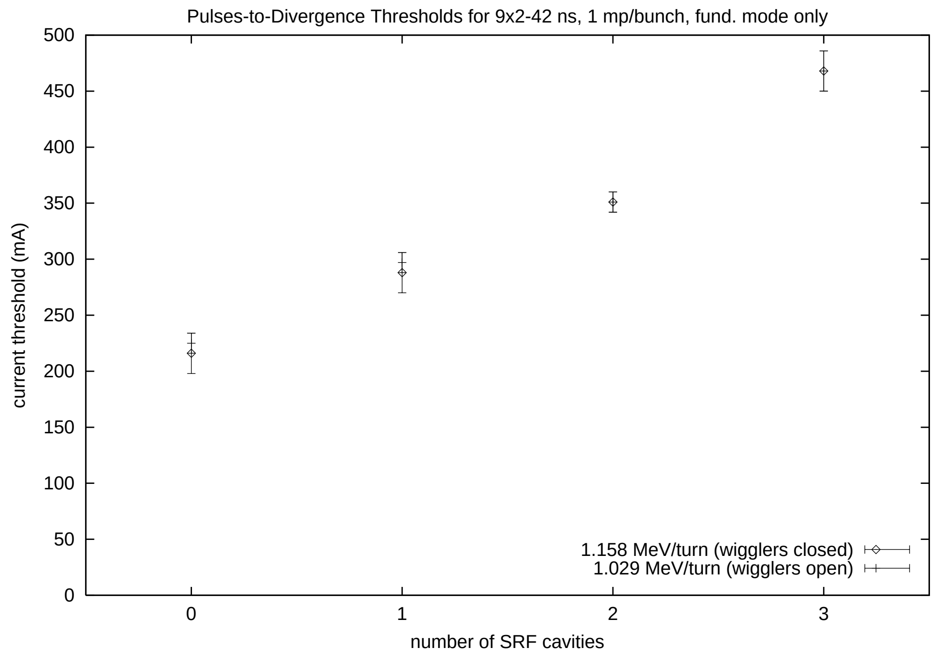

Figure 9: Pulses-to-divergence current thresholds for the fundamental mode alone as a function of the amount of synchrotron radiation.

Figure 10: Damping-to-pulses thresholds for the fundamental mode and a HOM as a function of synchrotron radiation.

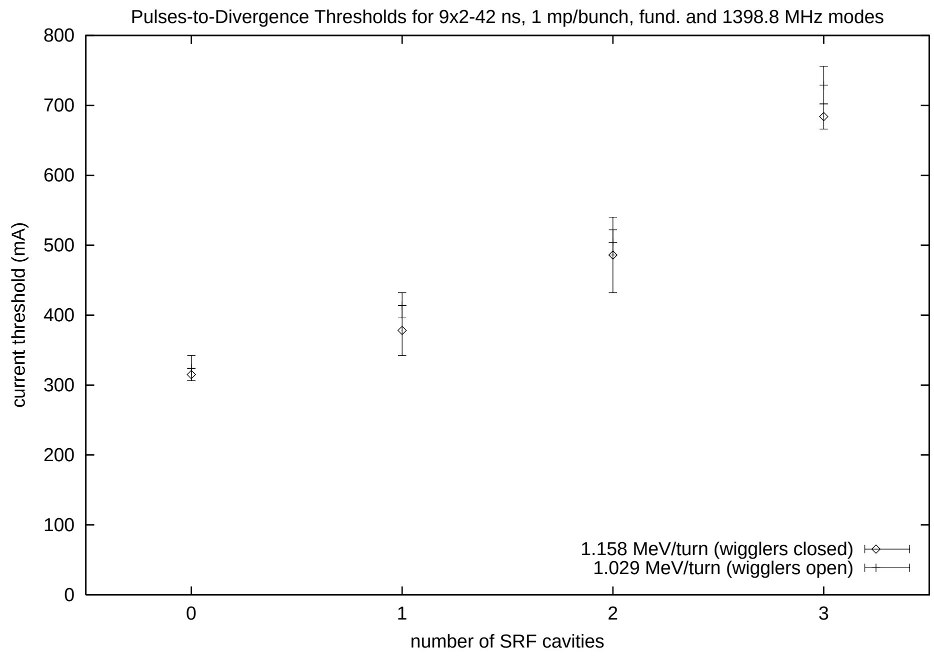

Figure 11: Pulses-to-divergence thresholds for the fundamental mode and a HOM as a function of synchrotron radiation.

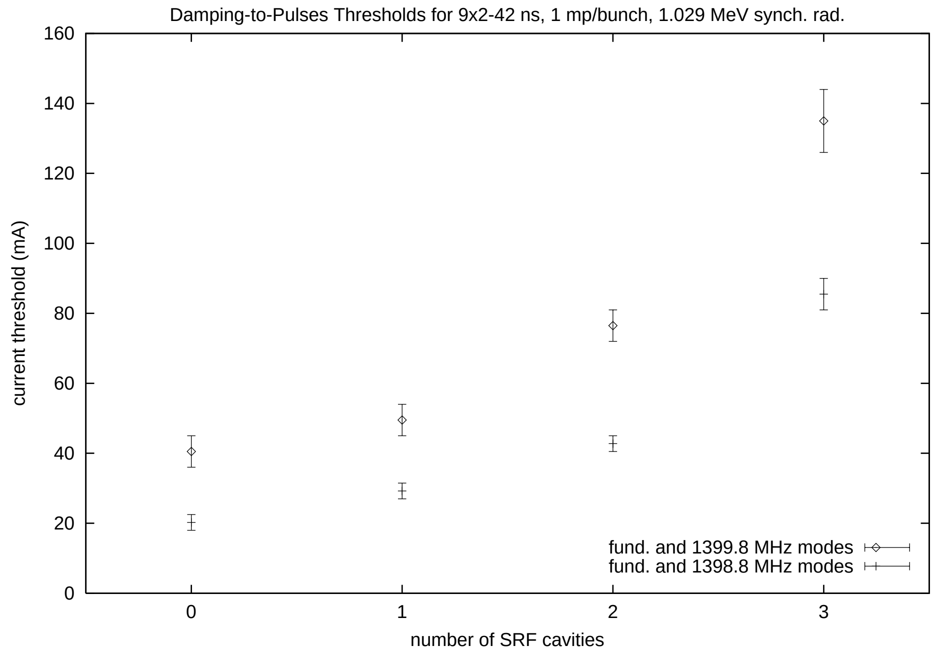

Figure 12: Damping-to-pulses thresholds as a function of HOM frequency.

60–100% in this case! A 44 kelvin temperature change is required to produce such a 1 MHz shift in a CESR copper cavity [7]. This would be a large temperature change, but even moderate temperature fluctuations would shift the thresholds by some amount.

In Figure 13, the pulses-to-divergence threshold is also shown to vary enormously when the HOM frequency is changed by 1 MHz. However, note that while the damping-to-pulses threshold greatly *increased*, increasing the HOM frequency by 1 MHz *decreased* the pulses-to-divergence threshold by about 30–40%! The large effect that HOM frequency variation has on the thresholds has important ramifications. In CESR, slightly different shapes of the cavities, including effects due to fabrication imperfections, cell probes, and stretching due to temperature increases, cause different HOM frequencies to prevail in different cells. As the temperature changes, even the frequencies in a given cell will vary. Therefore, the current thresholds will be hard to duplicate because they will change in time. This effect is consistent with actual CESR threshold measurements that change over time.

Figure 14 shows that the damping-to-pulses thresholds do not change when breaking up a pointlike bunch into multiple macroparticles. Nevertheless, the pulses-to-divergence thresholds do increase when breaking up the bunches as presented in Figure 15.

This increase would be due to Landau damping of the individual macroparticles within a bunch as suggested in section 4.1.

The macroparticle–number threshold effects are similar when also including a HOM with frequency 1399.9 MHz. The damping-to-pulses and pulses-to-divergence thresholds may be found in Figures 16 and 17, respectively. It appears that the effect of modeling bunches with more particles is independent of which modes are included.

Measurements in CESR show that a beam in the 9x2 configuration has different thresholds depending on the interbunch spacing. 42-ns spacing has a higher threshold than 14-ns spacing, for example. In Figures 18 and 19, the damping-to-pulses and pulses-to-divergence thresholds are seen to be higher in the 42-ns simulation than in the 14-ns simulation, in accordance with CESR data.

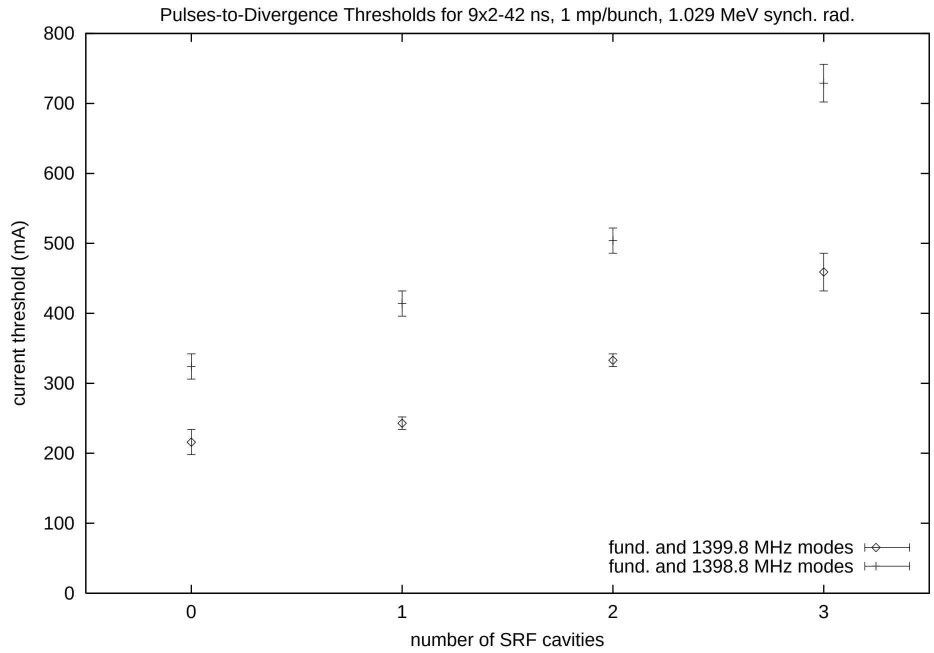

Figure 13: Pulses-to-divergence thresholds as a function of HOM frequency.

4.3 Interpretation of Current Thresholds

There is a vast difference between the damping-to-pulses and pulses-to-divergence threshold levels. The correct interpretation of the two different thresholds is not completely clear. However, the damping-to-pulses threshold may represent the current where a longitudinal instability might begin to be noticed while the pulses-to-divergence threshold may indicate when the instability is so strong that particles are lost from the beam. A threshold of merit would probably involve the pulses-to-divergence threshold with many macroparticles in a bunch and at least one HOM. An example is the upper set of data points in Figure 17. Keep in mind that changing the HOM frequencies (including adding more HOM's) can drastically change the threshold as shown in Figures 12 and 13.

Note that in all of these beam current threshold plots, the thresholds increase as we go from no SRF cavities (phase II) to one SRF cavity (beginning of the phase III upgrade) although the increase is not always drastic. However, the actual thresholds did not appear to significantly increase in CESR when the first SRF cavity was installed. The two most significant effects associated with replacing a copper cavity with a niobium cavity are a decrease in the total impedance and an increase in the total cavity voltage amplitude V_c (see Table 1). The replacement definitely does decrease the total impedance. However, V_c does not increase by necessity. As shown in Table 2, the sum of the V_c's around the simulated ring are 5.92 MV for no SRF cavities and 6.32 MV for one SRF cavity. (The actual CESR value for the one–SRF–cavity case was 6.61 MV (at 200 mA) in November 1997.) However, since the installation of the SRF cavity, the total V_c for all CESR cavities has decreased; the value was only 5.91 MV (at 200 mA) on September 23, 1998. This value is *lower* than the simulated phase II value!

Recalculating the one–SRF–cavity case with a total ring voltage amplitude of 5.91 MV showed damping-to-pulses thresholds that were the same or very slightly lower than the previously graphed results at 6.32 MV. However, the pulses-to-divergence thresholds are

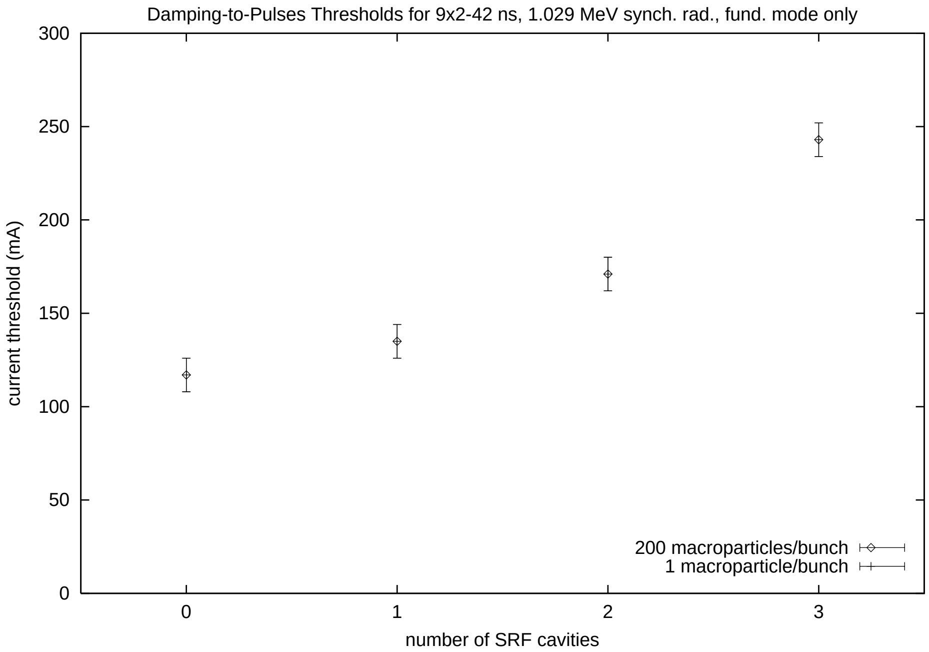

Figure 14: Damping-to-pulses thresholds for the fundamental mode alone as a function of the number of macroparticles per bunch.

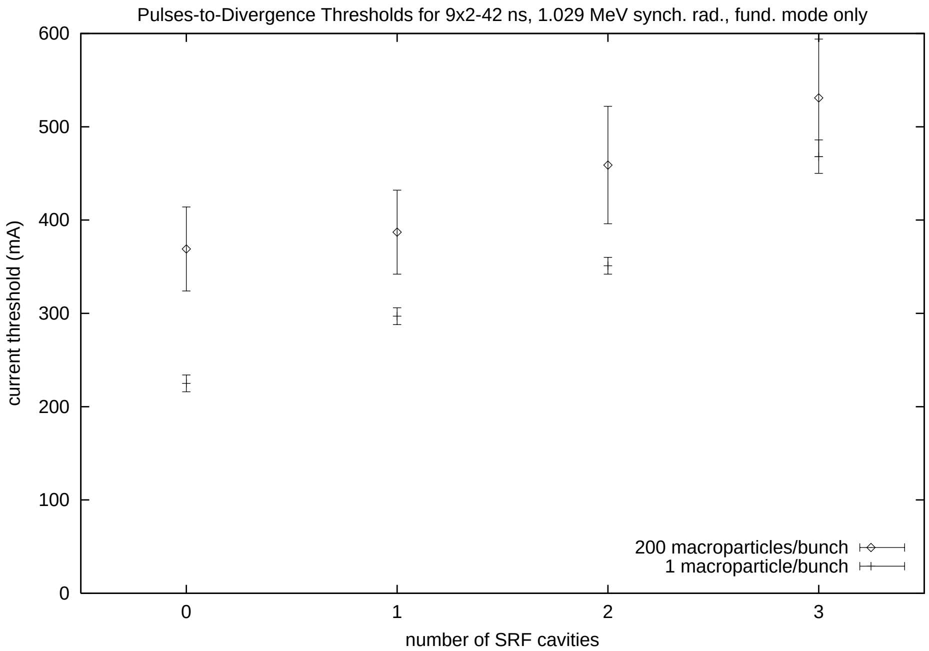

Figure 15: Pulses-to-divergence thresholds for the fundamental mode alone as a function of the number of macroparticles per bunch.

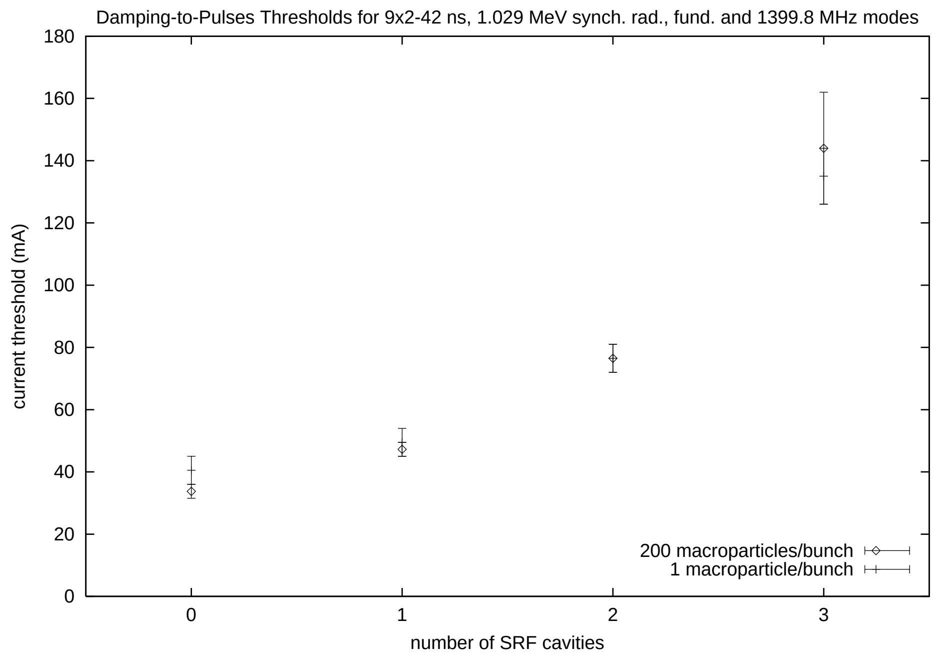

Figure 16: Damping-to-pulses thresholds for the fundamental mode and a HOM as a function of the number of macroparticles per bunch.

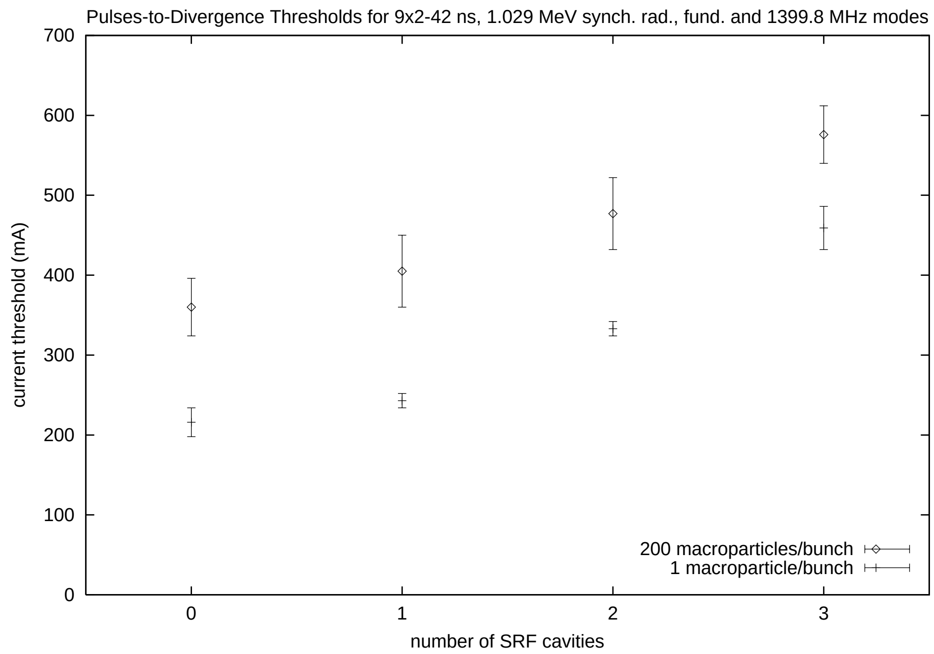

Figure 17: Pulses-to-divergence thresholds for the fundamental mode and a HOM as a function of the number of macroparticles per bunch.

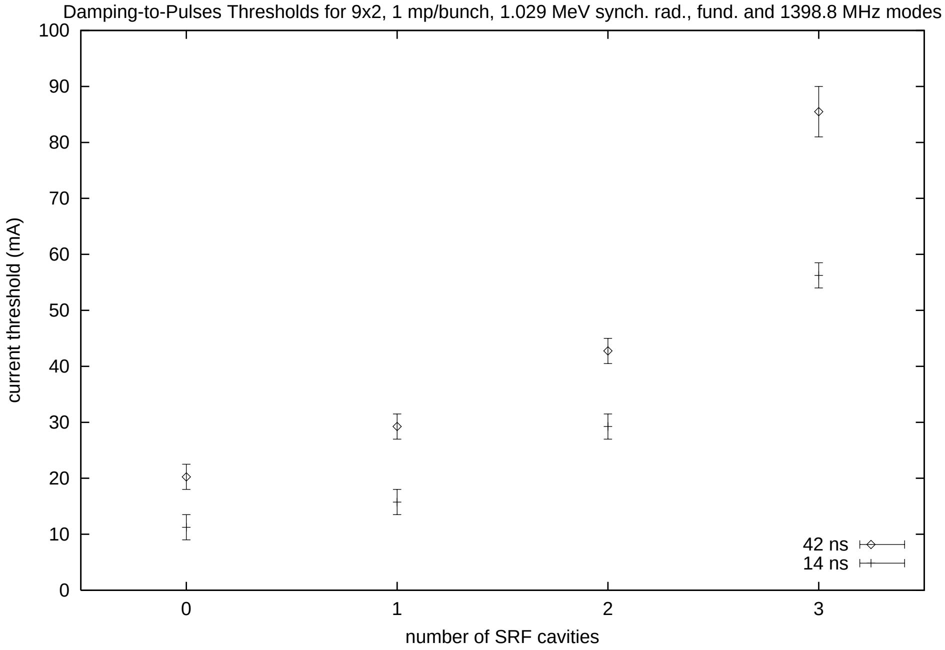

Figure 18: Damping-to-pulses thresholds as a function of interbunch spacing.

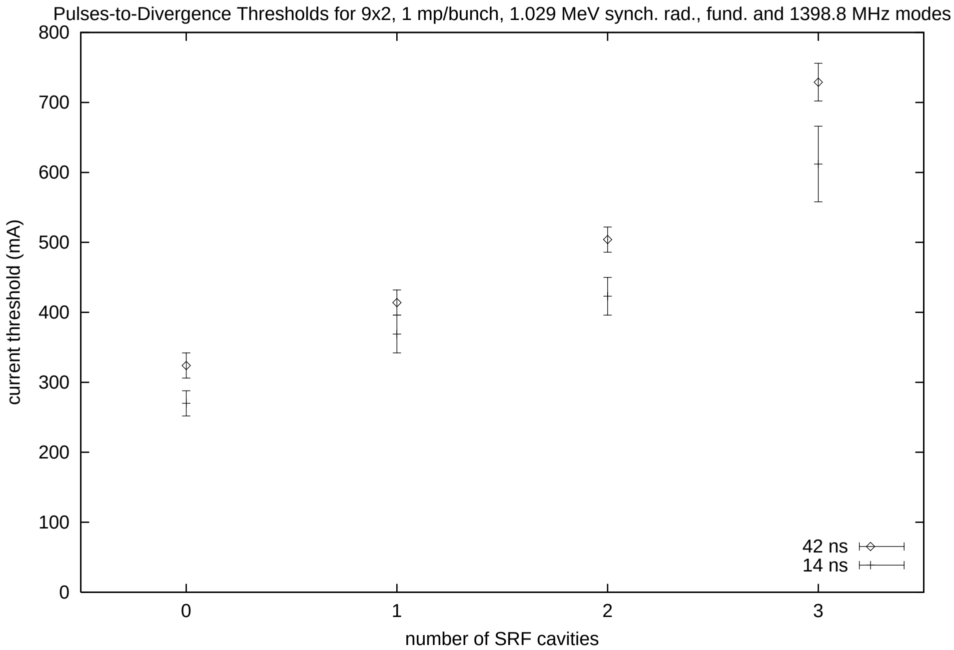

Figure 19: Pulses-to-divergence thresholds as a function of interbunch spacing.

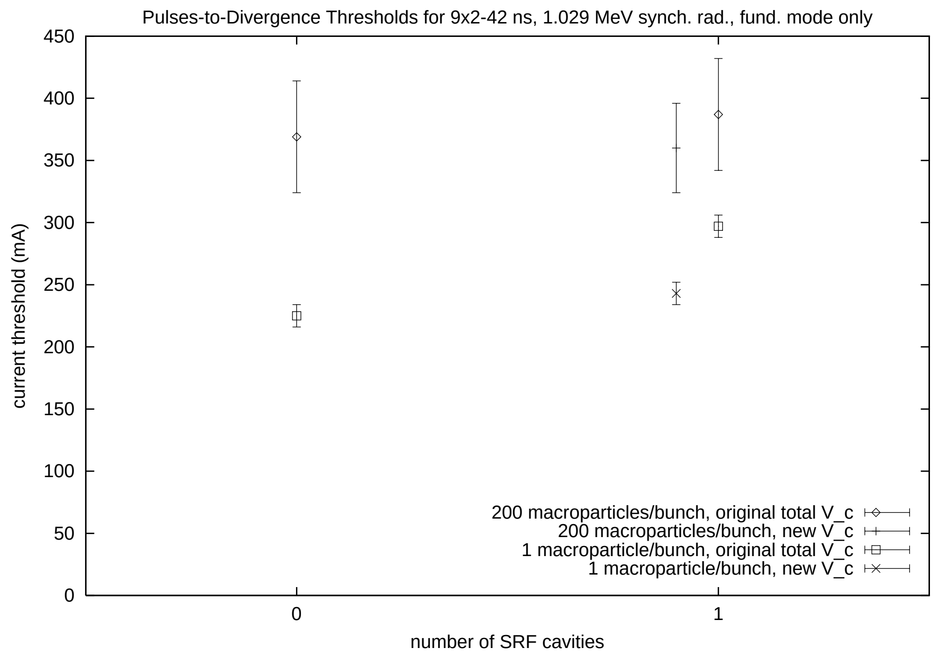

Figure 20: Pulses-to-divergence thresholds for the fundamental mode alone as a function of the number of macroparticles per bunch and the total ring cavity voltage. The one–SRF–cavity results at the new voltage of 5.91 MV are offset horizontally for clarity.

noticeably lower with the lower total voltage. Figure 20 reproduces part of Figure 15 and includes the new one–SRF–cavity case.

Notice that the one–SRF–cavity threshold is reduced to around the value in phase II as the total ring voltage is lowered. Figure 21, which duplicates part of Figure 17, adds a HOM and also shows the lower–voltage results. The data in Figure 21 also show that lowering the total voltage will lower the threshold if other parameters are held constant.

A final comparison of the different one–SRF–cavity voltages is presented in Figure 22 where different interbunch spacings are examined.

Figures 20–22 demonstrate that replacing copper cavities with superconducting cavities (and therefore lowering the total impedance) may not be enough to raise the current thresholds. According to the *oscil* simulation, the voltage should be raised on each SRF cavity—without lowering the voltage of the other cavities—in order to significantly raise the thresholds as the CESR phase III upgrade is completed.

Although only a 9x2 beam has been simulated so far, other train and bunch configurations, with varying interbunch spacings, may also be studied to determine their current thresholds.

4.4 Growth Rates

A key quantity used to describe instabilities is the growth rate. Can we determine why certain growth rates are seen in CESR and the *oscil* simulation?

Note that the growth rate may not always be clear–cut. Consider that a beam with initial positive growth rate must either be quickly lost or have the growth turned to damping. Therefore, the growth rates must change in time for beams that are high–current but contained. Such examples are seen in Figures 5 and 6, where there is a jumble of growing

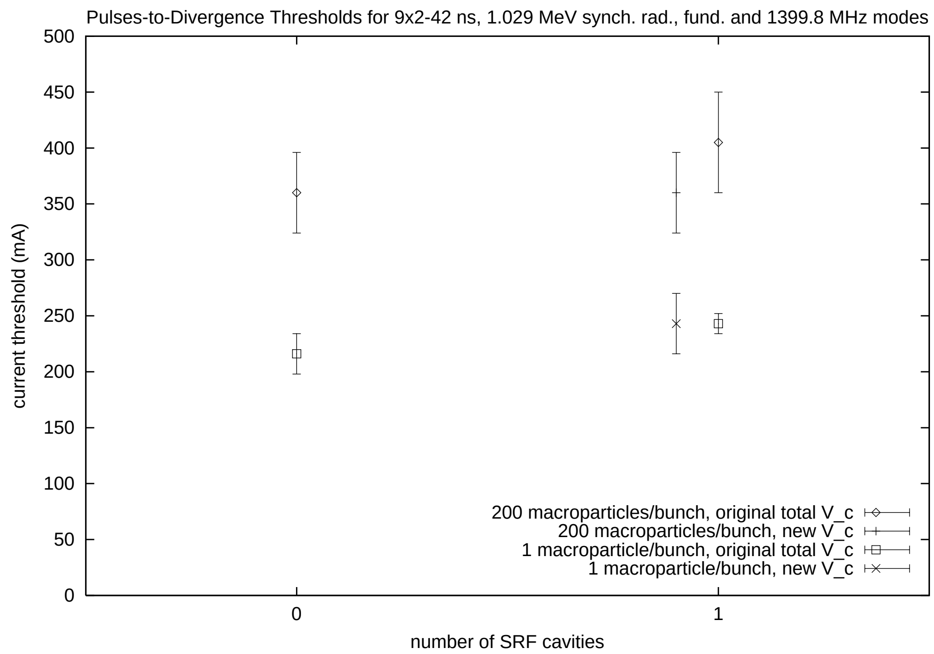

Figure 21: Pulses-to-divergence thresholds for the fundamental mode and a HOM as a function of the number of macroparticles per bunch and the total ring cavity voltage. The one–SRF–cavity results at the new voltage of 5.91 MV are offset horizontally for clarity.

and damping motions. Nevertheless, growth rates may be calculated for the very beginning of the beam's existence. The theory for the growth rates rests in work cited and created by Chao [3].

According to Chao, the shift in the complex longitudinal mode frequency for a beam of M *evenly–spaced* bunches, each with a "water–bag" distribution, is

$$\Omega^{(l)} - l\omega_s = i\frac{MNe^2\alpha\omega_0}{\pi\omega_s T_0 E\hat{\tau}^2} l \sum_{p=-\infty}^{\infty} \frac{Z(\omega')}{\omega'} J_l^2(\omega'\hat{\tau}) \tag{7}$$

where

$$\omega' = Mp\omega_0 + \mu\omega_0 + \Omega \quad , \tag{8}$$

N is the number of electrons (or positrons) per bunch, e is the charge of the positron, ω_0 is the angular revolution frequency, ω_s is the angular synchrotron frequency, T_0 is the revolution period, l is the longitudinal mode number, μ is the multibunch longitudinal mode index, Z is the total complex impedance for the fundamental mode for all cavities, J_l is the Bessel function of the first kind of order l, and $\hat{\tau}$ is the initial time displacement amplitude of the particles in the water–bag bunch.

The value l represents the longitudinal structure within a bunch and is an integer that satisfies $0 \leq l \leq N$. Chao illustrates that $l = 0$ represents a constant bunch shape and no centroid motion, $l = 1$ is also a constant bunch shape but with sinusoidal motion of the centroid, and $l = 2$ is sinusoidal motion of two halves of the bunch which necessitate a change in the bunch shape.

The value μ indicates the relative phase between successive bunches. The criteria for μ is that it be an integer satisfying $0 \leq \mu \leq M - 1$. In mode μ, each bunch is $\mu\pi/2$ radians ahead of the previous bunch in longitudinal phase.

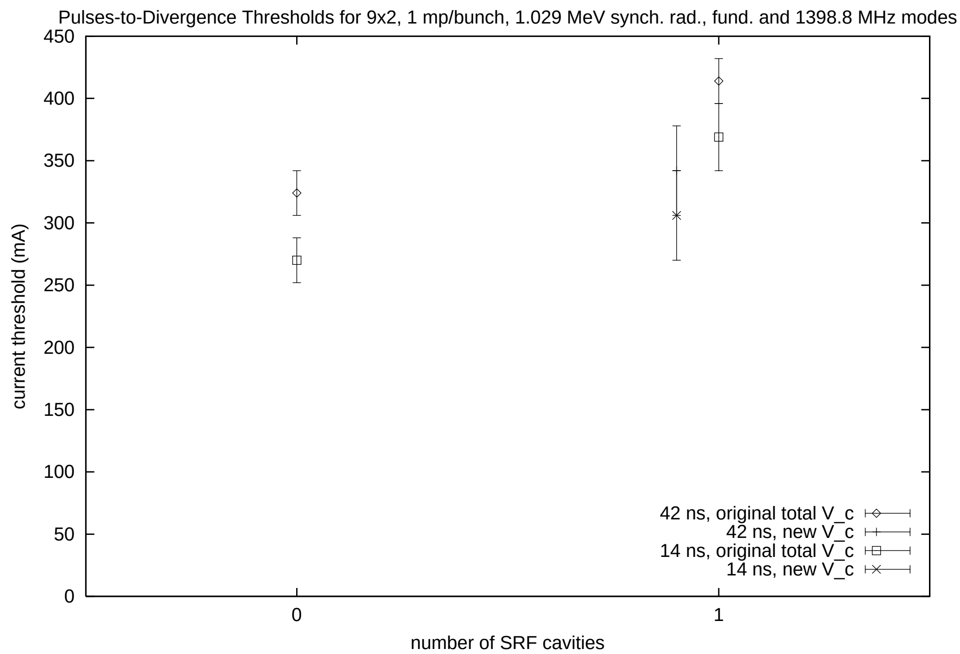

Figure 22: Pulses-to-divergence thresholds as a function of interbunch spacing and the total ring cavity voltage. The one–SRF–cavity results at the new voltage of 5.91 MV are offset horizontally for clarity.

Consider pointlike bunches ($N = 1$) even though the water–bag distribution would not be satisfied. This simplification restricts the only growth or damping rates to be associated with $l = 1$. Because the impedance is sharply–peaked near the RF frequency $h\omega_0$, only values of the Bessel function J_1 near the argument $h\omega_0$ need be considered. If the bunches are considered short (or small long. disp. for pointlike bunches) as compared to the RF wavelength, then the following approximation may be used: $J_1(x) = x/2$. Low currents are generally required to keep the long. disp. small. These approximations show that

$$\Omega^{(1)} - \omega_s = i\frac{MNe^2\alpha\omega_0}{4\pi\omega_s T_0 E}\sum_{p=-\infty}^{\infty}\omega' Z(\omega') \quad . \tag{9}$$

Note that the variable $\hat{\tau}$ has been eliminated.

Taking the imaginary part of eqn. 9 gives the growth rate. Performing this operation while remembering that only $\omega' \approx \pm h\omega_0$ gives a significant contribution to the impedance, we get:

$$\tau^{-1} = \frac{MNe^2\alpha h\omega_0^2}{4\pi\omega_s T_0 E}\left(\sum_{p=1}^{\infty}\mathrm{Re}[Z((Mp+\mu)\omega_0+\Omega)] - \sum_{p=-1}^{-\infty}\mathrm{Re}[Z((Mp+\mu)\omega_0+\Omega)]\right) \quad . \tag{10}$$

Substitute IT_0/e for MN (where I is the total beam current), $2\pi/T_0$ for ω_0, and use the zero order approximation for Ω ($\approx \omega_s$) to get:

$$\tau^{-1} = \frac{\alpha I e h\omega_0}{2ET_0\omega_s}\sum_{p=1}^{\infty}(\mathrm{Re}[Z((Mp+\mu)\omega_0+\omega_s)] - \mathrm{Re}[Z((Mp-\mu)\omega_0-\omega_s)]) \tag{11}$$

remembering that the real part of the impedance is an even function.

Equation 11 is the general formula (based on Chao's work) for the growth rate of mode μ for pointlike bunches assuming low currents. The beam current enters directly as I and indirectly in Z via eqn. 18 in section A.4 of Appendix A. Equation 11 assumes a single RF cavity so that V_c and $(R/Q)_{tot}$ from eqn. 18 must be the total values for *all* CESR cavities.

As Chao mentions, if there is only one bunch ($M = 1$ and $\mu = 0$), then the formula for the Robinson growth rate is obtained:

$$\tau^{-1} = \frac{\alpha I e h \omega_0}{2ET_0\omega_s}(\mathrm{Re}[Z(h\omega_0 + \omega_s)] - \mathrm{Re}[Z(h\omega_0 - \omega_s)]) \tag{12}$$

since only the synchrotron sidebands of harmonics near $\pm h\omega_0$ contribute significantly to the impedance. (A negative growth rate indicates positive damping.)

In general, there are M different (approximate) eigenmodes of interbunch phases corresponding to the M different values of μ. In addition to these growth (or damping) rates, there is a damping rate of about 76 s^{-1} due to synchrotron radiation (with the wigglers open). Adding the synchrotron damping rate to the μ-mode rates gives a total growth rate for each μ-mode. Because of the magnitude of the impedance near $\pm h\omega_0$, it is only necessary to consider the values of p that are near these two importance harmonics.

Figure 3 showed a 72-mA beam that damps rapidly. Figure 23 shows the early time behavior of the oscillations in detail. Examining the plot shows that the damping rate is about -6200 s^{-1}.

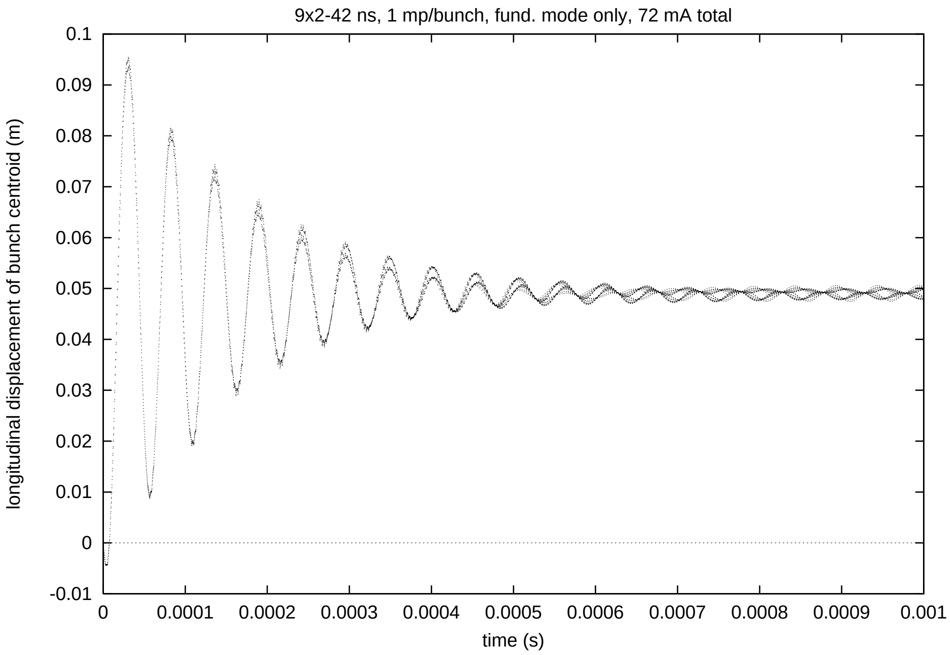

Figure 23: Longitudinal oscillations of 18 pointlike bunches at 72 mA total current. This figure corresponds to the early part of Figure 3.

The -6200 s^{-1} damping rate is close to (although not exactly) the theoretical value of about -4906 s^{-1} using eqn. 11 (see Appendix B).

Figure 4 shows some initial damping which is more easily seen in Figure 24. The damping rate is seen to be about -10,000 s^{-1} from Figure 24. Note that the damping is not perfectly exponential. This is due to the periodic cavity voltage regulation—see section A.3 of Appendix A. Hence the theoretical value should be lower than the approximate value averaged over most of the peaks seen in Figure 24. Equation 11 does give a higher rate for the $\mu = 0$ mode—about -12,666 s^{-1}. The rate could be calculated from only the first one or two oscillations of Figure 24 to show a better agreement of the experiment and theory, but the Bessel function approximation used in eqn. 9 is not accurate when the long. disp. amplitude is so high. Still, determining the growth rate from the third and fourth oscillations alone gives a rate of about -10,600 s^{-1}.

The highest value of μ ("20" when $M = 21$) gives the only growth in eqn. 11—at 22.84 s^{-1}. This value is very close to value of 25.9 s^{-1} seen as the first rise in Figure 4.

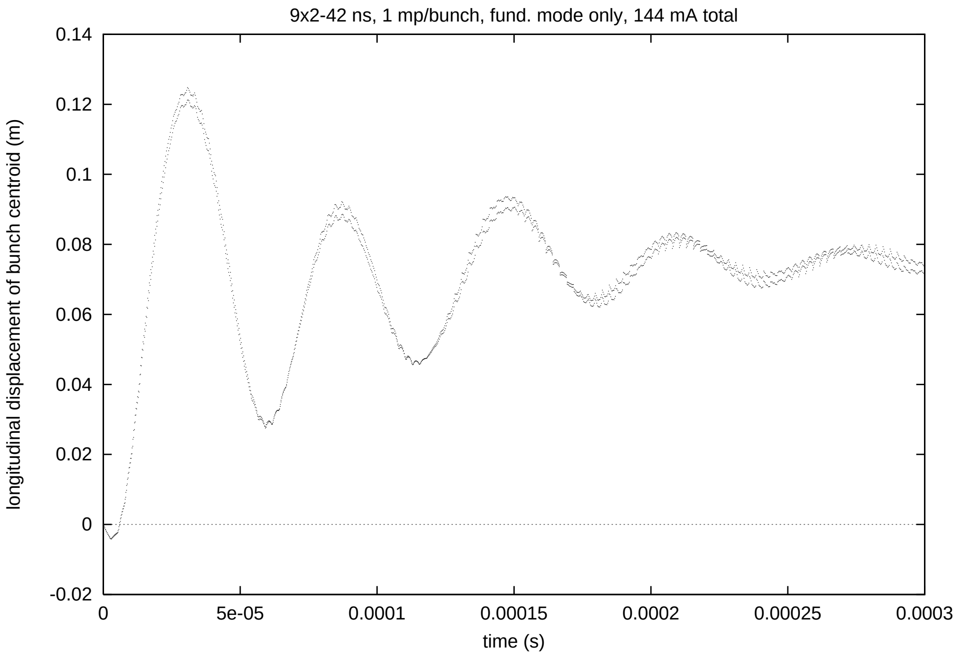

Figure 24: Longitudinal oscillations of 18 pointlike bunches at 144 mA total current. This figure corresponds to the early part of Figure 4.

Similarly, Figure 25 shows the early part of the data of Figure 5. The damping rate from Figure 25 is about -10,400 s^{-1} when using an exponential damping approximation for all of the peaks, but when considering only the third and fourth cycles, the rate is about –10,800 s^{-1} which is slightly closer to the eqn. 11 value of -15,217 s^{-1}. The first rise in Figure 5 shows a growth rate of 154 s^{-1} which is close to the eqn. 11 value of 150 s^{-1}.

4.5 Interpretation of Growth Rates

The theory provided by Chao does a good job of calculating the growth and damping rates of the initial longitudinal oscillations. In a multibunch case after a damping-to-pulses threshold, the oscillations become large and the Bessel function approximation used in eqn. 11 is no longer valid. The rates become jumbled and may even reverse sign. A chaotic beam may even arise as seen in Figures 5 and 6. Still, the damping-to-pulses threshold may be approximated by using eqn. 11 to find when the highest-μ mode changes from negative to positive (see Appendix B). The pulses-to-divergence threshold may be approximated when the growth

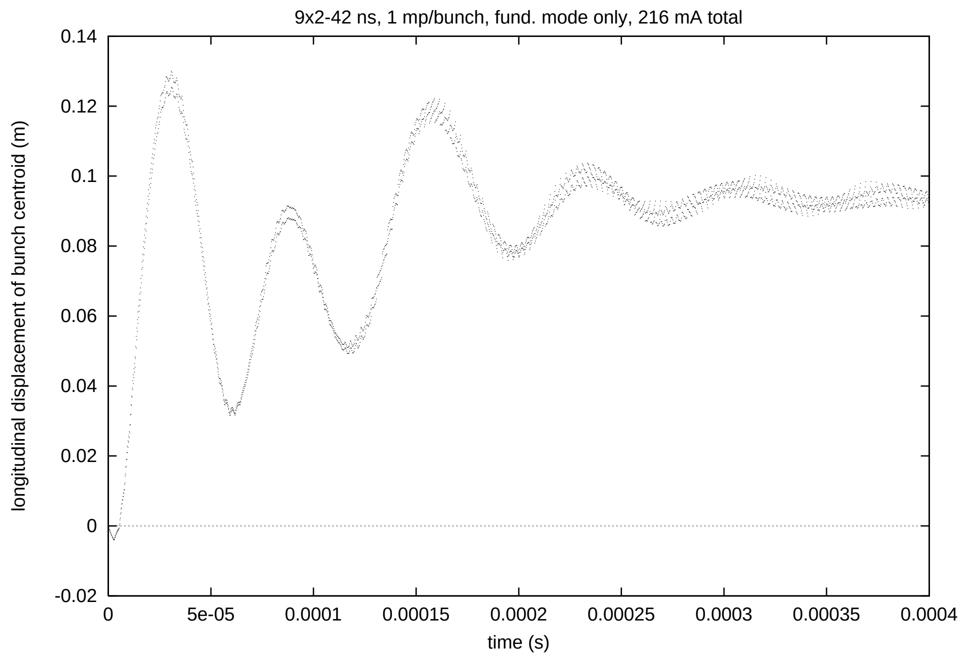

Figure 25: Longitudinal oscillations of 18 pointlike bunches at 216 mA total current. This figure corresponds to the early part of Figure 5.

rate for the highest-μ mode is greater than a certain value—perhaps somewhere around 250 s^{-1}. This rate is about the highest rate seen in the *oscil* simulation for beams that are not lost.

The initial damping, as seen in Figures 23, 24, and 25, results from the initial conditions—both of the beam parameters and the fact that the cavities begin with no wake fields.

The initial growth (not damping) rates described in section 4.4 may be explained using a simplified formula that resembles the Robinson growth rate formula of eqn. 12. Using eqn. 11 with different values of μ shows that given the parameters of CESR, the only growth stems from the multibunch mode with the largest allowable value of μ. Therefore, for that particular value of μ, we know that $(Mp' + \mu)\omega_0$ is near $\pm h\omega_0$ where the p' are the most significant values of p in the summation in eqn. 10 (*not* eqn. 11). We also know $\mu = M - 1$ and then:

$$Mp' + (M-1) \approx \pm h \quad . \tag{13}$$

Solving for p' we get

$$p' \approx \pm\frac{h}{M} - 1 + \frac{1}{M} \quad , \tag{14}$$

but since we know that p', h, and M are integers and that M divides h, we are left with the equality

$$p' = \pm\frac{h}{M} - 1 \quad . \tag{15}$$

Finally, we find:

$$Mp' + \mu = \pm h - 1 \quad . \tag{16}$$

The reduced formula for the initial growth (not damping) rate of a multibunch beam is

$$\tau^{-1} = \frac{\alpha I e h \omega_0}{2ET_0\omega_s}(\mathrm{Re}[Z((h-1)\omega_0 + \omega_s)] - \mathrm{Re}[Z((h+1)\omega_0 - \omega_s)]) \quad . \tag{17}$$

This formula is identical to eqn. 12 except that instead of using the impedances at the two synchrotron sidebands nearest $h\omega_0$, the next two closest synchrotron sidebands are used—those on the $h\omega_0$ side of the two nearest revolution harmonics [1]. It is predominantly only these two sidebands that determine the initial growth rate of a multibunch beam.

4.6 Fourier Spectra

The two sidebands in eqn. 17 may be observed in Fourier spectra of output from *oscil.* The Fourier transform of the one millisecond of long. disp. data from Figure 23 (72 mA) is shown in Figure 26. The $h\omega_0$ (ω_{RF}) harmonic and its sidebands are seen. The $(h \pm 1)\omega_0 \mp \omega_s$ sidebands are the peaks just on the inside (closer to the center of the plot) of the indicated revolution harmonics.

At 216 mA, the first millisecond of data is past the end of the graph in Figure 25, but the oscillations still do not grow in that time interval. The spectrum of that data is found in Figure 27.

The key feature to notice is that the sidebands that contribute to growth have significantly increased in magnitude. Also, the sidebands on the other side of indicated revolution harmonics have decreased in magnitude.

When there are a superposition of nearly–identical frequencies in any system, the difference of the frequencies is usually observed as a beat. Because the $h\omega_0$ and $(h \pm 1)\omega_0 \mp \omega_s$ sidebands are both seen in the Fourier spectra and have similar frequencies, an $\omega_0 - \omega_s$ line may also be seen in the spectra. Figure 28 shows the low frequency spectrum at 72 mA. Increasing the current to 216 mA produces the spectrum in Figure 29.

As with Figures 26 and 27, the current increase shows an increase in a certain spectral line. In this case it is an increase of the $\omega_0 - \omega_s$ (370.1 kHz) line with increasing current. (The ω_s and $3\omega_0$ lines—due to basic synchrotron oscillations and the symmetry of groups of three CESR trains, respectively—are also strong in these spectra.)

The increase in magnitude of the specified spectral lines in these figures indicate that the beam is becoming more unstable as the current increases.

The zero of long. disp. corresponds to the time when a hypothetical, *non–radiating* particle enters the first cavity. Therefore, the Fourier transforms in Figures 26–29 are equivalent to monitoring the long. disp. at periodic time intervals.

However, it may be desired to look at the beam current when the bunches actually arrive at the first cavity. In this case, the particles always have the same longitudinal position when observed, but the charge of each bunch may still be measured, and the times of the measurements will not be periodic. The Fourier transform of the charge at the arrival time of each bunch is similar to Figures 26 and 27 when the spectrum is observed near the drive frequency. The problem is that at low frequencies not much data is observable. Figure 30 is analogous to Figure 29 with the difference being that the charge (instead of the long. disp.—see Appendix B) is transformed. As with all the other spectra, the first 1 millisecond of data is transformed. The only discernible features are the $3\omega_0$ line and the DC component. The $\omega_0 - \omega_s$ line is not seen.

5 Conclusion

The *oscil* tracking code allows rough calculations of current thresholds based on a variety of beam and cavity parameters including, but not limited to, interbunch spacing and synchrotron radiation levels. A strong dependence of the threshold on HOM frequency is seen in Figures 12 and 13 which confirms qualitatively the observed dependence on cavity temperature and other factors. Nevertheless, if realistic longitudinal cavity modes are included, current thresholds similar to those seen in CESR may be obtained as shown in Figure 17. A more accurate threshold might be simulated by including many HOM's.

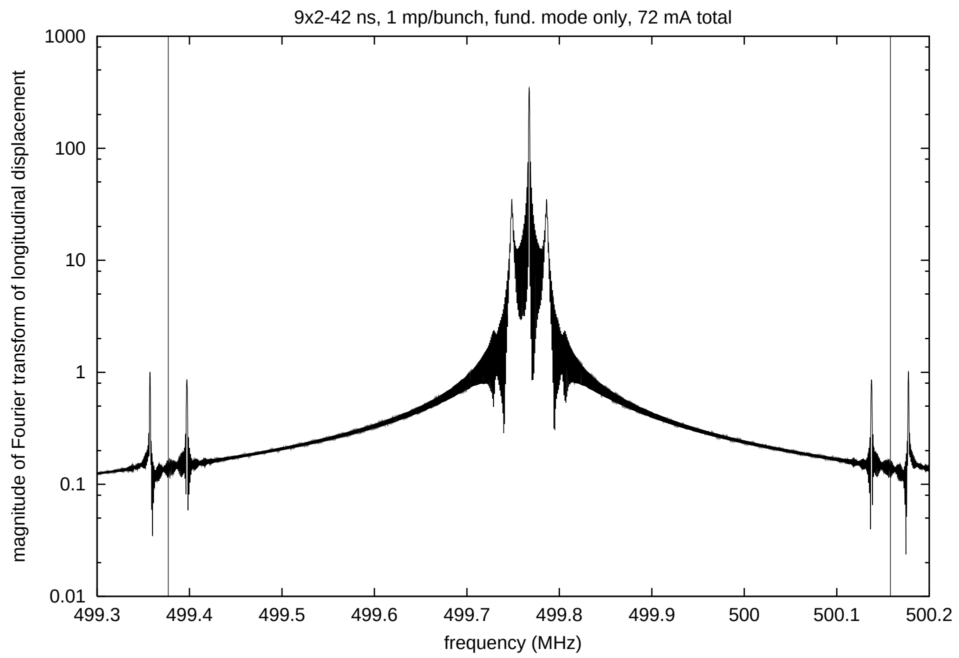

Figure 26: Fourier Transform of the longitudinal displacements from Figure 23. The beam current is 72 mA. The vertical lines indicate the $(h \pm 1)\omega_0$ revolution harmonics.

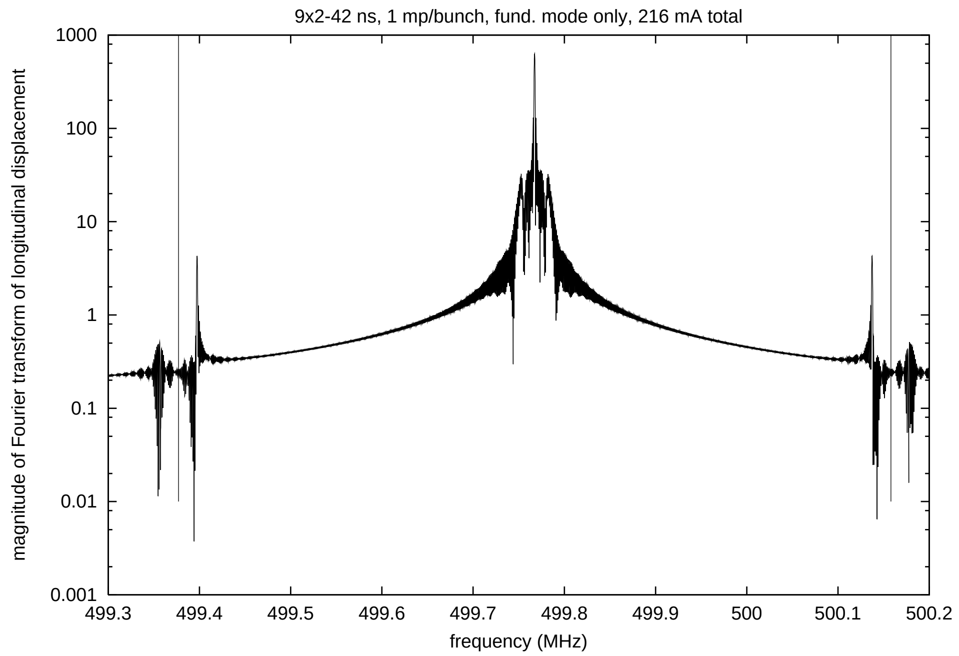

Figure 27: Fourier Transform of the longitudinal displacements from Figure 25. The beam current is 216 mA. The vertical lines indicate the $(h \pm 1)\omega_0$ revolution harmonics.

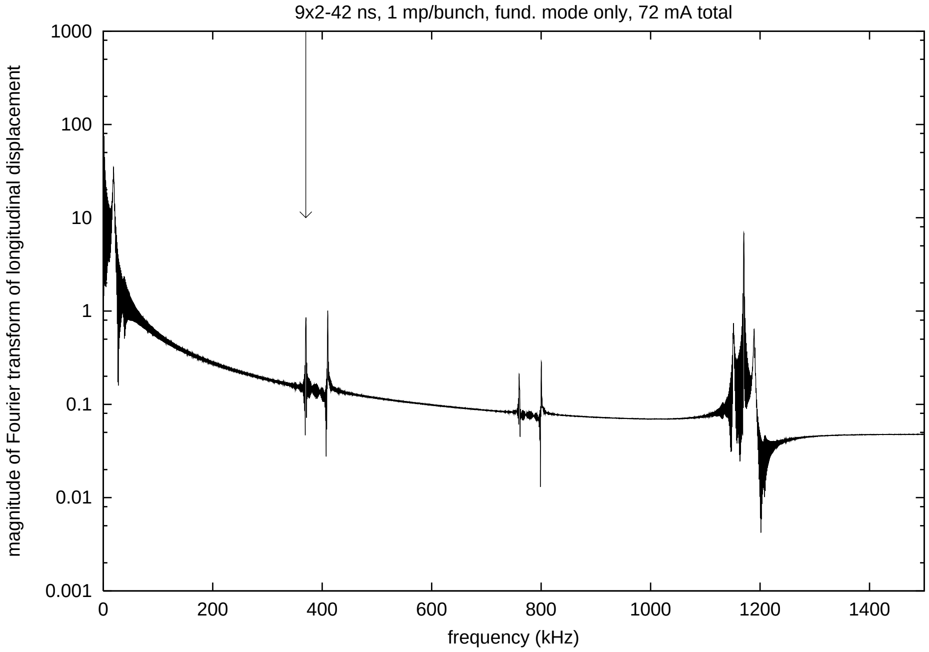

Figure 28: Fourier Transform of the longitudinal displacements from Figure 23. The beam current is 72 mA. The arrow is at the position of $\omega_0 - \omega_s \approx 370.1$ kHz.

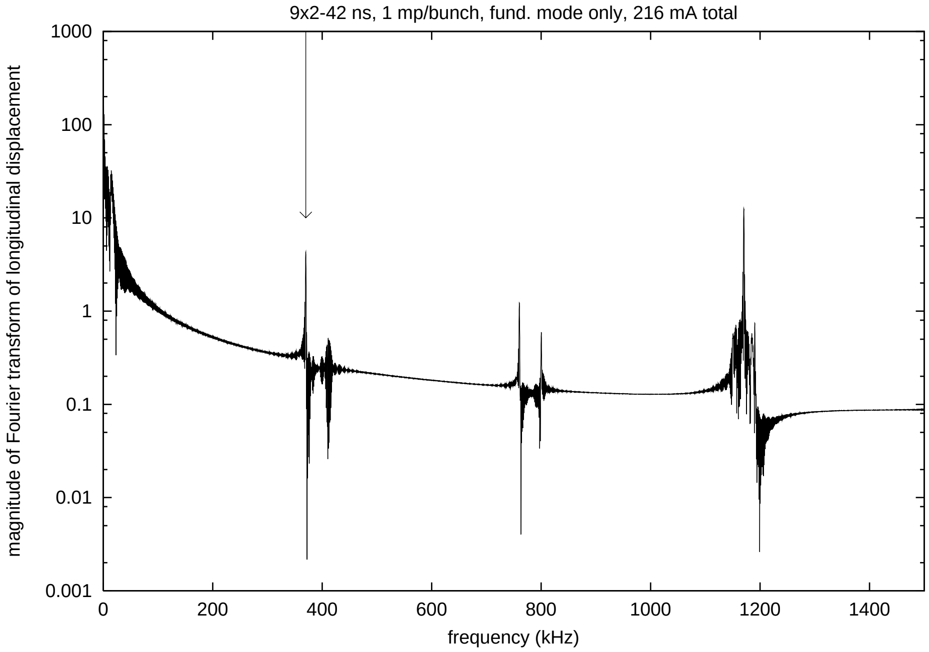

Figure 29: Fourier Transform of the longitudinal displacements from Figure 25. The beam current is 216 mA. The arrow is at the position of $\omega_0 - \omega_s \approx 370.1$ kHz.

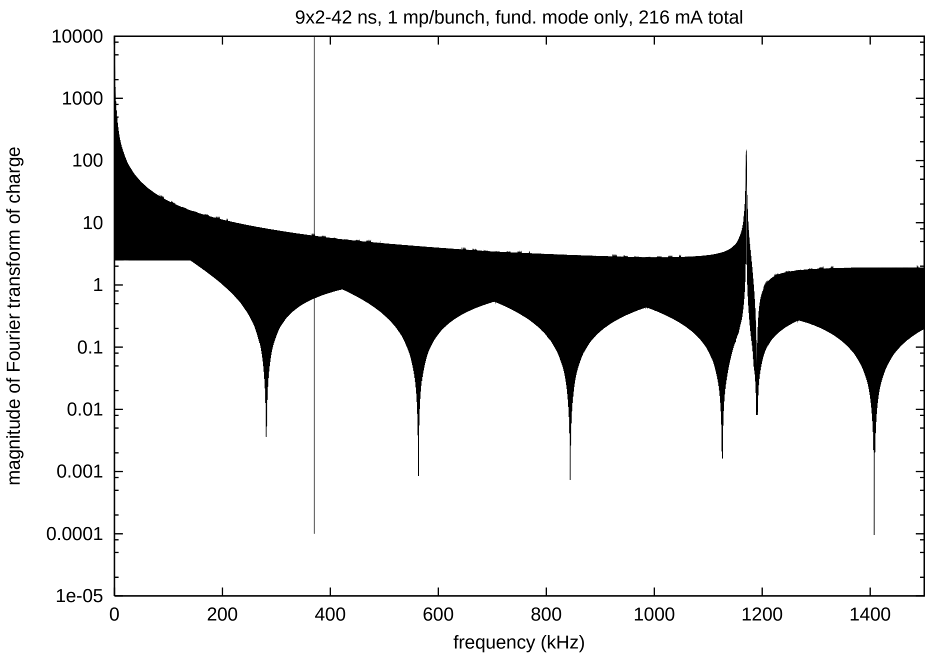

Figure 30: Fourier Transform of the charge from Figure 25. The beam current is 216 mA. The vertical line is at the position of $\omega_0 - \omega_s$.

An increase in the beam current threshold as the phase III upgrade is completed is predicted as long as the total ring cavity voltage amplitude is increased as well.

The initial damping (and growth for multibunch cases) rates are computed and may be predicted from theory. Although the rates are known for low currents and may be used to find the damping-to-pulses thresholds, at high currents the theory is not applicable, and therefore the *oscil* simulation is needed to find the maximum (pulses-to-divergence) thresholds. The theory explains why growth of the 370 kHz line indicates the onset of an instability.

6 Acknowledgments

I wish to thank Zipi Greenwald and Sasha Temnykh for data on Cu cavity modes and Sergey Belomestnykh for data on both Cu and Nb cavity modes. For discussions on bunch shapes and modeling, I thank Mike Billing, Cheng–Yang Tan, Edwin Anderson, and Kern Ormond. Also, I appreciate input from Zipi Greenwald regarding the dynamical equations. Last, but not least, no paper written by a graduate student is complete without an acknowledgment of the research adviser. Therefore, I wish to thank David Rubin for his suggestions, guidance, and motivations which have greatly helped this work.

A Appendix: *Oscil* Simulation Details

Detailed information concerning how to execute the *oscil* program is located on the online CESR Documentation Software pages under the entry "OSCIL." Additional information regarding the particulars of the code is given below.

A.1 Simplifications

The following simplifications are made in *oscil*:

1. There is only one beam present, specifically a positron beam based on the order in which the RF cavities are traversed.
2. The beam starts with its full current rather than having the current run up gradually.
3. Each multicelled cavity is treated as a single cell. To represent the location in the ring, each cavity is considered as having zero length and is placed at the position of the center of the real CESR cavity.
4. There are no explicit optics. All optics (except for the RF cavities) manifest themselves via synchrotron radiation (synch. rad.) and the momentum compaction factor α.
5. α applies to any segment of the ring rather than just to the entire ring.
6. Synch. rad. losses are continuous rather than discrete.
7. There is no transverse vertical motion. There is no transverse horizontal motion except for that due to dispersion.
8. There are no space–charge effects.

A.2 Adjustable Parameters

The simulation allows for a wide range of variables so that various conditions in CESR may be modeled. The variables that may be easily changed (via an input file) are:

1. The number of trains—either one, three, or nine trains may be used.
2. The number of bunches per train—by specifying the occupied/vacant status of each of the first five slots.
3. The interslot spacing—may be set to either 14 or 10 ns.
4. The number of "macroparticles" per bunch. The previous version of *oscil* could only use one macroparticle per bunch, i.e. a pointlike bunch.
5. The bunch current—applies to every bunch.
6. The initial en. disp.—applies to every macroparticle.
7. The initial standard deviation of the long. disp. within a bunch—applies to every bunch. The initial shape of the bunch is purposely rectangular to allow *oscil* to determine the bunch shape itself.
8. The amount of synchrotron radiation in MeV/turn/particle.

9. The longitudinal cavity modes—may be defined separately for each cavity. (There are no transverse modes.) Each mode is defined by the center frequency f (or ω), the impedance per unit length R/Q, and the loaded quality factor Q_L. Additionally, the number of cells in the particular cavity is needed to determine the *total* impedance in the cavity. The fundamental mode may be detuned automatically to minimize the reflected power from each cavity to the klystron, if desired.

10. The zero–current tuning angle offset per cavity—defined separately for each cavity.

11. The average total cavity voltage for the fundamental mode—defined separately for each cavity.

12. The relative (as compared with the values from the other cavities) net power to beam per cavity.

13. The number of revolutions for *oscil* to execute.

A.3 Cavity Regulation

The klystron–supplied drive voltage amplitude is regulated every 50 turns so that the total cavity voltage amplitude is the right value as explained in section 2.4.

Because of the cavity regulation as diagrammed in Figures 1 and 2, it is impossible (with the present code) to find the thresholds for phase III (four SRF cavities). For all threshold plots in this report, pointlike bunches in phase III conditions were strongly damped up to about 1400 mA total current, and for 200–macroparticle bunches, up to 1600 mA.

However, raising the current slightly more than the said current causes the beam to diverge. The reason is that at the start of the simulation, the cavities are devoid of wake fields. Thus the initial drive voltage phasor is vastly different than the phasor required after the wake fields have been built up, and the beam is lost almost immediately.

This problem would not occur if the currents were run up slowly from zero, but that would require recoding of several parts of *oscil*. It is, however, safe to say that *oscil* shows that the simulated phase III threshold should more than 1500 mA by probably a factor of two or more (as long as there no current–limiting effects that are not included in *oscil*).

A.4 Other Details

The simulation can track many particles for many turns in a reasonable amount of time. Choosing to output data pertaining only to the *envelope* of the long. disp. decreases the execution time and keeps output files to manageable sizes. The envelopes for the long. disp. plots in this report were created by printing only the maximum displacement and minimum displacement that occurred for all particles in all bunches during each 200–turn interval.

Throughout this report, the fundamental mode has been always automatically detuned. The change in the fundamental mode frequency $\Delta\omega$ for a particular cavity may be expressed as

$$\Delta\omega = -\omega_{RF}\frac{I\cos(\phi_s)}{V_c}\left(\frac{R}{Q}\right)_{tot} + (\text{offset}) \tag{18}$$

where ϕ_s is the synchronous phase angle, and $(R/Q)_{tot}$ is the total impedance for the cavity for the fundamental mode.

Two sets of initial conditions were used for the threshold plots shown in section 4.2. One set started with a standard deviation of macroparticle long. disp. of 0.01 m and an en. disp. of 3.0 MeV, and the other set began with the values being 0.02 m and 0 MeV. The long. disp. standard deviation does not apply to pointlike bunches.

B Appendix: Theoretical Growth Rates

Equation 11 with $M = 21$ gives the theoretical growth rates for each multibunch mode (including synchrotron radiation damping) that are seen in Table 3.

Table 3: Theoretical Growth rates including radiation damping for 21 evenly–spaced bunches. The current is 72 mA and the conditions are those used for Figure 23.

μ (multibunch) mode index)	growth rate (s^{-1})	μ (multibunch) mode index)	growth rate (s^{-1})
0	-4906.19	10	-75.7488
1	-93.6241	11	-75.742
2	-78.1094	12	-75.7342
3	-76.4283	13	-75.7231
4	-76.0153	14	-75.705
5	-75.878	15	-75.6718
6	-75.8152	16	-75.6117
7	-75.7842	17	-75.4518
8	-75.767	18	-74.9845
9	-75.7564	19	-72.978
		20	-51.4568

As is seen in the table, all the rates are negative indicating damping (for this particular current). The zero mode always shows the strongest damping. Had there been no synchrotron radiation, then some modes would show growth.

The reason that *18* bunches were not used in Table 3 is that the bunches must be evenly spaced. In other words, M must divide h. The only acceptable values for M are 1, 3, 7, 21, 61, 183, 427, and 1281. The use of any acceptable value greater than or equal to 7 produces the same rates for the highest μ-modes while the rate for the zero mode is the same for all of the acceptable M values. The value 21 for M is used because it is near the number of bunches often used in CESR and the *oscil* simulation and because it is one of the lower possible values for M.

When $M = 18$ is used in eqn. 11, the rates bear no resemblance to those in Table 3, the reason being that 18 bunches cannot both be evenly–spaced in CESR be near the null of the RF drive waveform. With 18 bunches, eqn. 8 could not be satisfied with any set of integers M, p, and μ.

When the current is increased, the only multibunch mode that shows growth is the one with the highest index—20 in the case of Table 3 with a higher current. Synchrotron radiation suppresses other potential growth modes given the parameters of CESR.

For the data in Figure 30, the charge is observed when each bunch reaches the first cavity. Since all bunches in *oscil* have the same charge, the charge is considered to be unity. Using unity instead of the actual charge merely changes the scale of a Fourier transform that is not normalized. Not using the actual bunch charges also speeds the transform computation.

References

[1] D. Fromowitz. CON 97-10. 1997.

[2] D. L. Rubin, J. Graber, W. Hartung, J. Kirchgessner, D. Moffat, H. Padamsee, J. Sears, and Q. S. Shu. CLNS 89/909. 1989.

[3] A. W. Chao. "Coherent Instabilities of a Relativistic Bunched Beam," *Physics of High Energy Particle Accelerators* **105**, 1983.

[4] E. B. Anderson. ODYSSEUS simulation. See CBN 97-33.

[5] Sergey Belomestnykh. Unpublished SRF cavity data.

[6] Don Bilderback. Private communication.

[7] Sasha Temnykh. Private communication.