EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

CERN – SL Division

CERN-SL-Note-2000-054 HRF

CERN-NUFACT Note 064

Simulation of the SPL SC RF system with Beam using "SPLinac"

J. Tückmantel

Abstract

The study of an H^- linac with superconducting RF at CERN, called SPL, requires examination of microphonics and Lorentz detuning of the superconducting cavities and their influence on the beam quality at the and of the linac. We report here on first studies using the program ‘SPLinac’

Geneva, Switzerland

November 2000

1. Introduction

At CERN a study is under way to make use of parts of the now decommissioned LEP2 superconducting RF system together with new hardware to establish an H^- linac capable to deliver 4 MW proton beam power for different final applications. The machine was initially designed as a proton linac, therefore the name SPL for Superconducting Proton Linac stuck. Details of this machine are given in [1].

One problem with superconducting linacs is the potential detuning of the cavities by external (microphonics) or internal (Lorentz-) forces thus modulating the resonance frequency, resulting in field variations that perturb the beam. To understand the influence of these perturbations, a simulation program, 'SPLinac' [2], was written and first simulations of the SPL machine with it are reported here.

2. Perturbations of the Cavity Fields and their Influence on the Beam

In a linac with particles of constant speed – highly relativistic ones – a local lack or excess of acceleration will not influence the acceleration efficiency in the following parts of the linac and the final particle energy will simply be different by this amount. Things are worse in SPL with particles of β significantly lower than unity. Here a local change of acceleration will modify not only the energy gained but also the speed of the particle. Therefore drift times to the subsequent cavities change, modifying phase angles and the acceleration in these cavities. A local perturbation modifies parameters globally.

The origins for field variations are multiple. Firstly, the resonant frequencies of the accelerating cavities are not stable due to external perturbations. Some are *coherent* for the whole machine due to a *common* origin – e.g. the cryogenic plant or ground motions – some are *incoherent* – e.g. *different* vacuum pumps on different cryostats. The loaded cavity bandwidth is, for SPL cavities, between 120 Hz and 180 Hz and even small mechanical perturbations can produce comparable shifts. Secondly, when the field is pulsed, the Lorentz force detunes the cavities with similar effect. Thirdly, the sudden onset of beam loading at the start of the beam pulse creates a fast transient of the cavity fields. In this context it is important to remember that beam loading is not uniform for all – otherwise equal – cavities since the transit time factor changes with the increasing speed of the particles from cavity to cavity, especially important in the low energy part.

To fight the consequences of these field variations a vector (sum-) feedback will be built around cavity and transmitter. However, this system is only capable of reducing the effect, not of eliminating it completely. The loop gain g is limited due to unavoidable delays and noise; we have chosen a value of 100 as a reasonable compromise. Then the fraction 1/g = 1% of any original perturbation remains. Furthermore, all fast transients are only compensated within a time-scale of the cavity filling time (in the order of 1 ms) divided by the loop gain g, i.e. within about 10 μs. The bunch position in phase space as seen at the end of the linac thus varies rapidly at the beginning of the beam pulse due to the beam loading.

2.1 Multi-Cavity Operation with one Transmitter

An additional difficulty arises in the sections β=0.8 and β=1 where, for cost reasons, a single klystron drives several cavities via a vector *sum* feedback. Mathematically the complex sum of individual cavity voltages, phase-adjusted for the particle's time of flight, expresses the total RF voltage as seen by the particle. Practically, such a RF vector sum is built up from the measured, phase corrected, individual cavity voltages to give a signal imaging the total RF voltage of the whole family of cavities. This signal is fed into the control loop as the 'measured cavity voltage'.

For particles with constant speed – i.e. β very close to 1 – the time of flight is invariant and the phase adjustments set into the summing device are always correct, thus the sum signal correctly images the total voltage independently of the individual components of the sum. However, for slower particles as in SPL this is no longer completely true. Any change of cavity voltage also influences the speed of the particle and thus the time of flight towards the following cavities. Therefore the phase adjustments in the summing device do not correspond to reality any more and the sum output signal does not image precisely the total voltage gain. However, the feedback loop will do its best to match this slightly wrong signal and the reference signal, thus inducing residual voltage errors depending on the individual cavity voltages.

2.2 Real Life System Errors

Finally there are practical imperfections in the real hardware, particularly perturbing when several cavities are driven by a common transmitter. Errors in the power splitting network result in cavities with unequal excitation in amplitude and phase. As we have seen above each deviation from the design voltage – for which the vector sum device is calibrated – leads to errors in the measured equivalent field sum. The scatter in the coupling strength of the main couplers causes a similar effect and additionally produces a non-uniform beam loading. Furthermore, the vector sum devices themselves will contain calibration errors; these errors lead the feedback system to stabilise a slightly wrong signal, not the true equivalent voltage, thus leaving residual errors.

Also the Lorentz constant k_L, expressing the static frequency shift per square of the accelerating field, may have a scatter over different cavities, important if one transmitter drives several cavities.

3. Short Description of the Simulation Program 'SPLinac'

To study the influence of all these perturbations and system errors on the beam – including the RF feedback loops and their real life limitations – the program 'SPLinac' [2] was written.

The machine is assumed to be composed of sections with different internal parameters, the cavities having β-dependent transit time factors. Fast RF vector (sum) feedback loops control one or several cavities driven by one transmitter. RF feedback loops include limited transmitter power and bandwidth and – if the option is chosen – loop-delay, as well as vector sum calibration errors, power splitting errors and scatter in the coupling strength to the cavities. Beam loading due to a pulsing beam can also be considered. For each cavity the externally driven vibrations are characterised by

their frequency, amplitude and phase. The amplitude of the vibration is not expressed by a mechanical quantity but directly by the induced resonant frequency sweep; thus the vibration amplitude ('stroke') is expressed in (RF-) Hz. For all cavities in one section we can either force the same amplitude or scatter it between zero and the maximum; similarly the phase can be fixed or scattered over 360º. Also the vibration frequency for one section can be defined to be identical for all cavities (coherent) or scattered (incoherent). Bunches can be defined in full longitudinal phase space representation or simply as a point bunch.

It is not intended that cumulative beam break-up is determined with this programme. Therefore to a very good approximation beam loading can be expressed over the pulse length by a continuous, averaged beam current at the nominal phase angle. Any microstructure and phase jitter is ignored; the inertia of the cavity stored energy averages it out in any case. Should the beam phase jitter become important, the validity of this beam loading calculation becomes invalid but the beam is lost then in any case.

To localise the origin of the most troublesome perturbations, the measured and ideal RF vector sum, individual (complex) cavity fields including beam loading, frequency deviations, and generator and reflected power can be plotted for each transmitter family. The programme offers a multitude of graphic displays including phase-space images along the linac - together with the jittering local RF voltage - or summary dot-map or overlay plot at the exit as well as all dynamic RF quantities of any transmitter or cavity. Precise numbers for voltages, powers, phases or phase-space points in plots can be obtained online at the cursor position in an 'info window'. All plots can be saved as standard graphic files, accessible to professional word-processors or graphic postprocessors.

The program is easy to use[1] starting by simply clicking the program icon or a document icon. It offers explanatory panels for entering or changing data or options. Default sets of data are saved automatically – e.g. together with a picture file – or on user request, keeping an easy track of the different run parameter sets during longer production runs with many different cases. Frames can be assembled to movies showing – off-line on any other computer - the dynamics of the system.

3.1 Simulation Error Creation

We want to compare pairs of machines under otherwise exactly the same conditions, where only one single (set of) parameter(s) is different. Therefore care has been taken while creating numerical random errors, that an error attributed to a certain object (e.g. the Q_{ext} of cavity 123) depends only on the random seed and the user-specified magnitude. It is created by exactly the same sequential call for the pseudo-random generator, independent of whether another option is activated or not.

4. Parameters of the Simulation

The relevant parameters as used are given in the following table

[1] It runs on any Macintosh with Power-PC processor (including G3 or G4) and MacOS8 or later

Table 1: The essential linac data used in the simulation

β_{nom}	(R/Q) [Ω]	Qext [E9]	Eacc [MV/m]	Vacc [MV]	φs[°]	cavities	cav/trans
0.52	60.7	2	3.5	3.1	-25	14*3=42	1
0.7	107.7	2.5	5	6	-20	8*4=32	1
0.8 (a)	191.8	3	9	15.3	-20	6*4=24	4
0.8 (b)	191.8	3	9	15.3	-15	6*4=24	4
1	234.4	2	7.5	12.8	-15	27*4=108	6

These data differ from the latest 'official' SPL data on one point. As demonstrated later, the chosen 'official' phase angle of –15° for the β=0.8 section can be slightly changed in the first half section in such a way that the beam quality is considerably increased. We have therefore split this section into two halves, running the first one at a phase angle of -20° and the second unchanged at the 'official' value of -15°. There are only 24 cavities affected – each with nominal 15.3 MV – and thus the total beam energy reduction – proportional to [cos(-15°)-cos(-20°)] – is less than 10 MeV.

(R/Q) and V_{acc} in Table 1 are expressed for the nominal β. Figure 1 shows the *relative* transit time factors as function of β for the four (distinct) groups of cavities. This function is normalised to 1 (horizontal red line) for the nominal β (vertical red line), the effective voltage for any other β is scaled by the shown function.

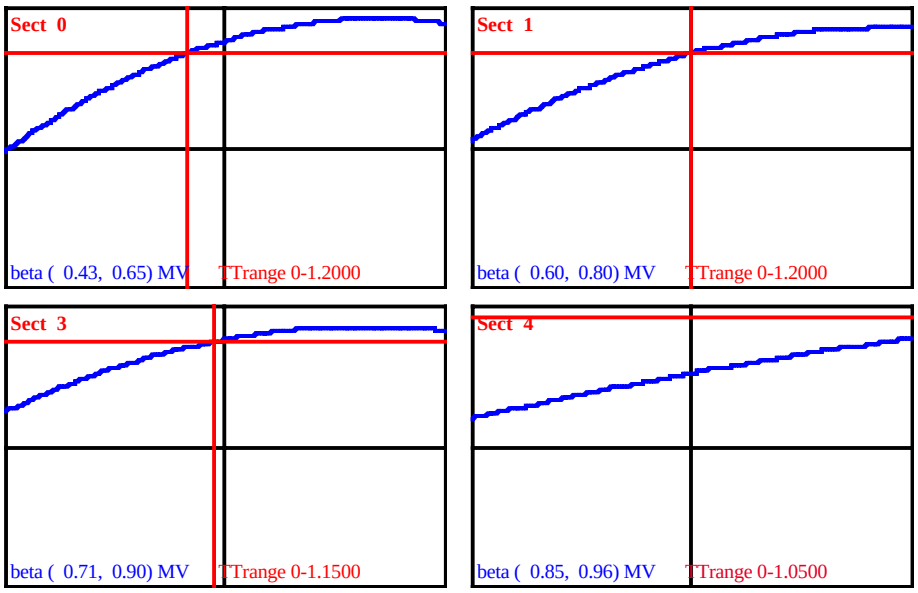

Figure 1: The (relative) transit time factors of the four sections. Top left: β=0.52 section from 0.43 to 0.65; top right: β=0.7 section from 0.60 to 0.80. Bottom left: β =0.8 sections from 0.71 to 0.90, bottom right: β =1 section from 0.85 to 0.96; vertical relative range: 0 to 1.2

To express the beam loading an averaged beam current of 11 mA (DC) is assumed over the pulse length.

5. Internal Parameters of an RF System Family

In the following we shall show the 'internal working' of the first transmitter family of the β=0.8 section ('transmitter 74') supplying four cavities ('cavities 74 to 77'. It is considered here as part of the 'first' machine with a resonant frequency swing of ±20 Hz, beam loading and RF system errors (for details see later). The corresponding phase space plot will be shown in Figure 17 (top left plot).

Figure 2 to Figure 5 show these data. On the left at t=0 the RF is switched on, and loading of the empty cavities to the nominal field starts. Figure 2 shows the RF power (red) rising nearly immediately to its maximum power. Figure 3 shows the cavity voltage real parts, i.e. in phase with the nominal voltage. Thin, light blue traces show all four cavity voltages rising (zero suppressed). The true sum voltage is shown by a thick dark blue trace and the measured and error prone sum voltage thick orange. The latter rises till the measured voltage approaches the set value (corresponding to the y=0 line). In parallel the reflected power in Figure 2 (magenta), goes down while the cavities are loading. When the measured sum voltage approaches the set value, the generator power reduces and approaches the reflected power; both do not perfectly join here – there are neither sc. cavity losses nor beam losses yet - since the cavities change their resonant frequency, the latter being traced in Figure 5.

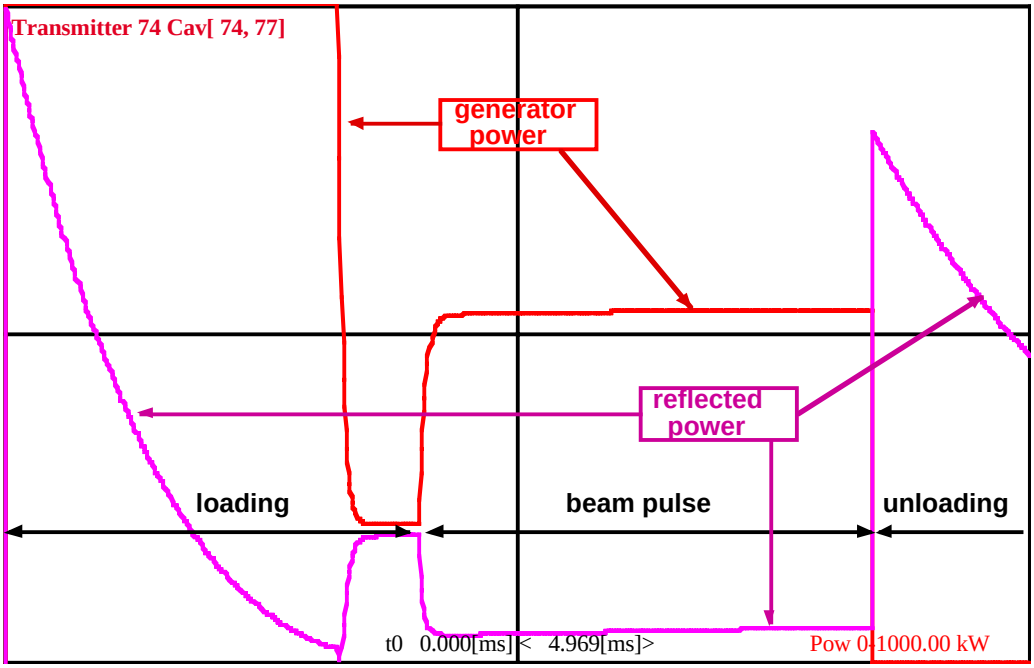

Figure 2: Transmitter (red) and reflected (magenta) power, horizontal time scale 5 ms, vertical power range 0-1MW.

The quadrature voltages are shown in Figure 4. When the *real* part of the measured sum voltage arrives at the set value, the *imaginary* part of the measured sum voltage (thick cyan) approaches also its set-value zero, the true voltage (thick dark green) stabilises slightly off due to the calibration errors in the vector sum. The imaginary parts of the four individual cavity voltages (light green) show different behaviour compared to the real parts – the imaginary voltages change with the first

order of the frequency change - depending on cavity tune status, even running out of scale, while the sum remains close to zero.

At t=2 ms the beam is injected and all real parts of the cavity voltages show a fast decrease, while the transmitter power shows a fast rise to supply the new power demand. Due to differences in transit time and phase angle, cavities are loaded differently by the beam. Therefore, despite the fact that the nominal vector *sum* is stabilised by the feedback system, individual cavity voltages drift with their natural filling time to the new equilibrium (four diverging light blue lines). Since the beam does not pass at the crest of the RF it induces also an imaginary voltage, the latter also drifting apart for the different cavities

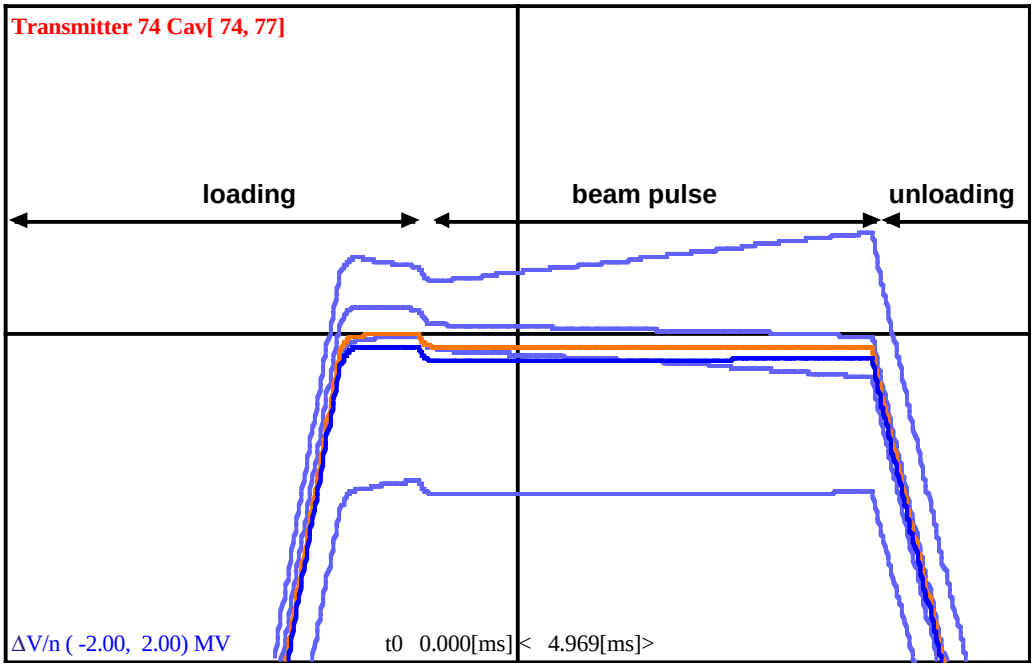

Figure 3: Cavity voltages and vector sum voltages real component (zero suppressed, the reference value is shifted to the centre of the y-axis). The scale for cavity voltages (light blue) is ±2MV, sum voltages ('true' in thick dark-blue, 'measured' – including VS errors – in thick orange) are displayed 'normalised per cavity', i.e. the total voltage scale is ±8MV then. Time scale 5 ms as Figure 2

At the end of the beam pulse (t=4.2 ms) the RF power is switched off instantly and the reflected power produces a power spike due to the sudden unloading of the cavities (Figure 2). All cavity voltages decay, approaching zero with the apparent exception of the imaginary voltages in Figure 4. The imaginary voltage increase is caused by the fact that the feedback is not active any more and the complex voltage vector starts to turn due to the detuning. The initially mainly real vector now turns towards the imaginary direction and thus increases this component in the beginning, even while the absolute value of the vector decays. In a plot of this function for a longer time, the imaginary part also finally decays to zero, as it should be.

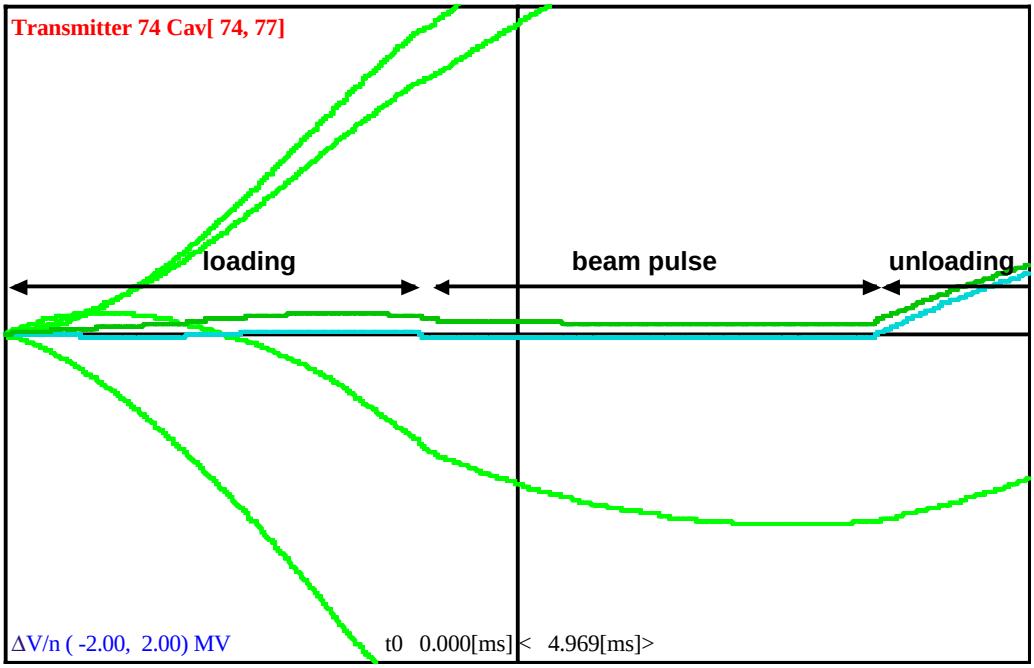

Figure 4: Cavity voltages and vector sum voltages imaginary component (quadrature). The scale for cavity voltages (light green) is ±2MV, sum voltages (in thick dark-green 'true', 'measured' with VS errors in thick cyan) are displayed 'normalised per cavity', i.e. ±8MV. Time scale 5 ms as Figure 2

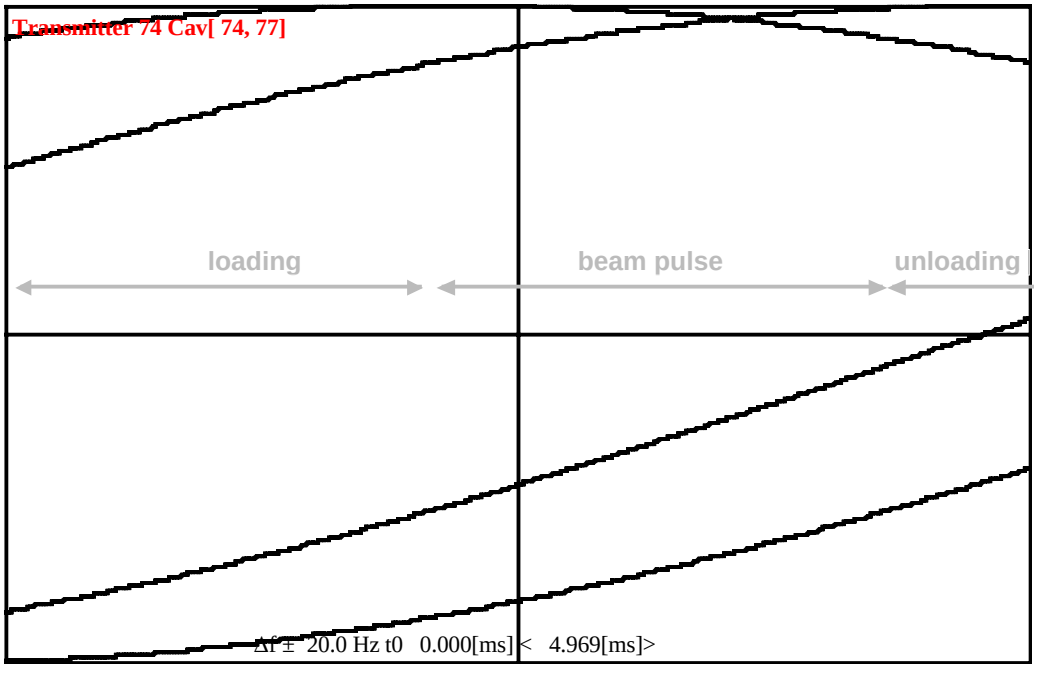

Figure 5: The drifting cavity resonant frequencies, ±20 Hz 'stroke', coherent, with random scattered phase. Time scale 5 ms as Figure 2

6. Number of Transmitters in the β=0.7 Section

We discussed in an early stage the question of how many cavities one transmitter could supply and it is evident that the lower β sections are the most sensitive ones. Therefore the β=0.5 section was already initially foreseen with one tetrode per cavity. The remaining open question was the procedure for the β=0.7 section, since the β=0.8 and β=1 sections are foreseen with LEP klystrons with several cavities per transmitter.

The nominal beam power per cavity for the β=0.7 section is about 65 kW; thus a LEP klystron could supply eight cavities easily, more would be too risky. If this set-up could not be used, a single tetrode per cavity would be another alternative. We now show results of a set of machines – of the same basic design as those used later – with one tetrode per cavity (this case is identical to Figure 17 in the next paragraph) and one LEP klystron per *eight* cavities. Shown are overlaid phase images at the end of the linac.

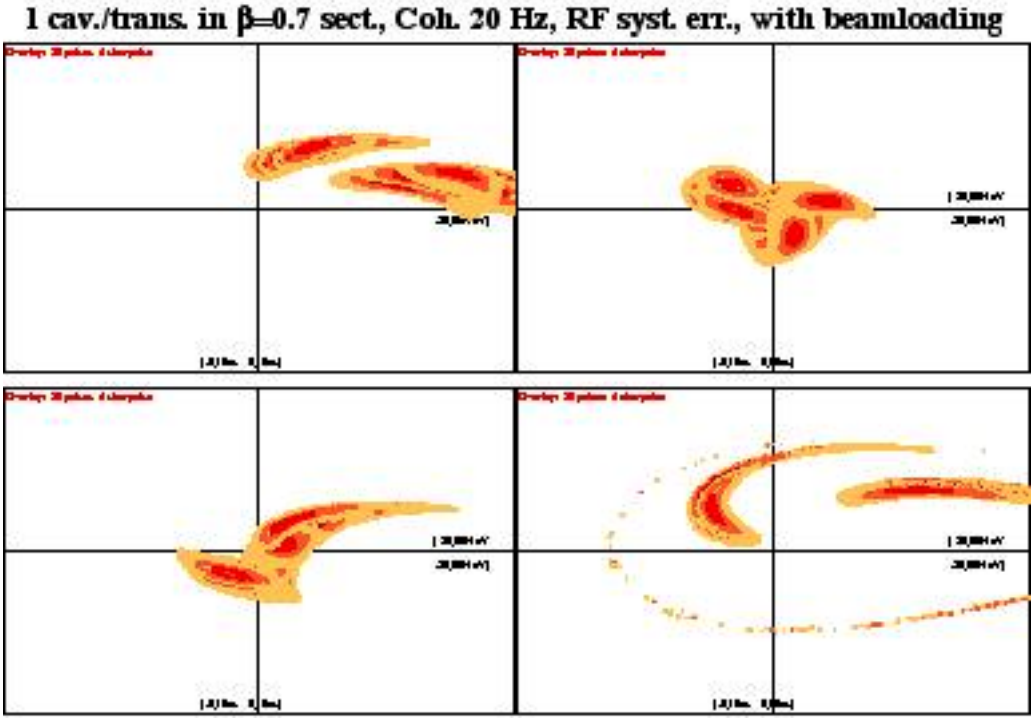

Figure 6: *One* cavity per transmitter (tetrode) in the β=0.7 section for four independent 'random' linacs. Overlaid longitudinal phase space images at the end of the linac for 20 pulses, five 'shots' per pulse, range ±100 ps (horizontal), ±20 MeV (vertical). For more details on the linac set-up see text in following paragraphs.

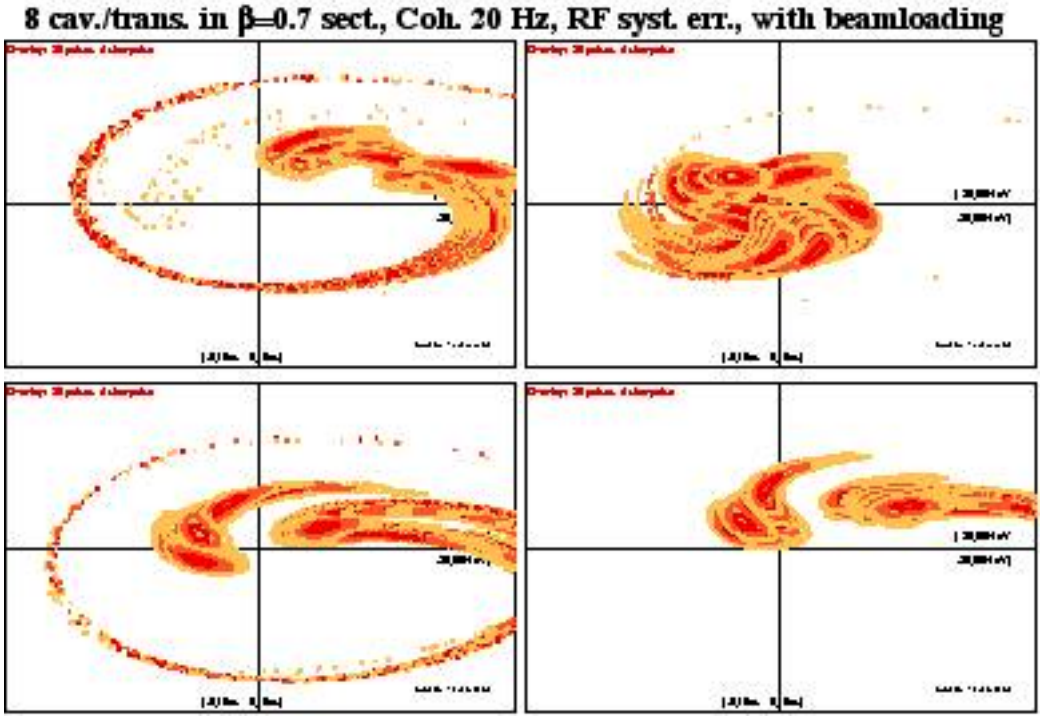

Figure 7: As Figure 6 but with *eight* cavities per transmitter (LEP klystron) in the β=0.7 section

We see immediately that the tetrode solution is much more stable and it was retained for SPL.

7. Examination of Microphonics

For the first and main part we have studied cavity microphonics – i.e. externally driven oscillations – and their influence on the beam. Two modes of microphonics are possible: coherent oscillations from a common source or incoherent oscillations driven by many different sources. Since we want to be sure that SPL can withstand the worst imaginable case, we examined the different configurations and from then on concentrated on the worst case.

For the coherent oscillations 'per definitionem' all cavities vibrate with the same frequency. To impose a maximum stroke we assume also that all cavities (of the section considered) have the same amplitude (resonance frequency stroke). However, if the phases are also the same, all cavities behave identically and a 'family' of one transmitter supplying several cavities is absolutely equivalent to individual transmitters for each of those cavities; we have constructed an 'easy case'. Therefore we have to assume that the phases are randomly scattered over 360º for the worst case.

The absolute value of the vibration frequency we use for probing is of no real importance provided the cavity field can follow the oscillation adiabatically. This is the case if the vibration frequency is somewhat smaller than the cavity bandwidth, the minimum being 120 Hz in SPL.

Nonetheless, a possible pitfall for the choice of this frequency exists. The vibration phase advance from pulse to pulse for a cavity vibrating at f_{vib} is given by $\Delta\phi=2\pi(f_{vib}-f_{rep})T_{rep}$, i.e. the k'th pulse passes the cavity at $\phi_k=\phi_1 + 2\pi(k-1)(f_{vib}-f_{rep})T_{rep}$. If now by 'bad luck' $(f_{vib}-f_{rep})T_{rep}$ were equal to 1, even if we simulate any number of pulses, we would always see the same phase; if $(f_{vib}-f_{rep})T_{rep}$ were 2/3 or 1/3, we would see only three phases 120° apart. In the coherent case this phase advance is unique for all cavities, i.e. for a badly chosen f_{vib} we might encounter the whole machine in only a very reduced number of states and miss a problem for other phases. Also for a ratio close to 'very rational numbers' we would only probe a very small total range. On the other hand this relation can be exploited to test 'all' phases in uniform steps with only few simulations. We have chosen the ratio[2] f_{vib} = 9/20 f_{rep} = 33.75 Hz, which lets the vibration phase scan exactly the full circle (in fact it describes 11 turns in phase) during the 20 pulses we will simulate.

For the incoherent case we have chosen vibration frequencies around the same value of 33.75 Hz but scattered by ±20 Hz. If one cavity vibrates 'very rationally' with respect to the repetition frequency, this does not matter, most do not. Also in reality this case can happen statistically. Oscillations are independent and in this spirit we have initially used a stroke scattered between zero and the maximum given value. Then results are – as to be expected – for the same (maximum) amplitude less perturbed than in the coherent case. To have the worst conditions we have also tried scattered frequencies but with the *same* amplitude. Then the beam perturbations for otherwise comparable conditions are very similar to the coherent case. To avoid an overwhelming flux of data (i.e. figures) without showing really new things, we have concentrated on the *coherent case* as defined above; it also allows the regular safe scan as shown above.

Figure 8 shows the RF behaviour of one family of four cavities and their transmitter for 20 pulses. The 20 individual plots for one pulse are overlaid in one frame – like a triggered storage scope – and the jitter in the parameters is well visible.

To examine the statistical scatter presented in such a machine, we have constructed four different machines, identical in theoretical design but with four different random seeds for error generation. As stated above, if an option is activated in a machine, all numerical error values are the same for all tests independent of otherwise different parameters settings. In all following figures (Figure 10 onwards), the plots shown in the top left corner correspond always to the same ('first') machine, the top right ones to another ('second') machine and so on.

We have fixed on two 'strokes', 20 Hz as the generally acceptable case – with rare exceptions which had to be 'trimmed' - and 40 Hz, often above the accepted limits of ±10 MeV at the and of the linac when considering the whole phase space occupation. However, to define a reference, we start with the cases without vibrations.

We simulate 20 'well distributed' beam pulses. The cavities do not move very much within the 2.2 ms of a beam pulse so that all bunches in the same pulse are nearly identical; residual cavity loading transients then change fields only slightly. However, beam loading modifies the cavity fields rapidly; the feedback with g=100

[2] any factor (n/20) would do provided n and 20 are mutually prime (i.e. n is not multiple of 2 or 5)

compensates only 99% of it. As we have said above, the e-time of the feedback is of the order of 10μs and therefore the beam sweeps at the beginning of the pulse with this time constant through the phase space. The first bunch sees no beam loading yet, following bunches sweep and once the new equilibrium state is established with the feedback e-time, later bunches then arrive all at about the same place. To probe the beam loading, each beam pulse is 'tested' by 5 reference bunches ('shots') equally distributed in time over the beam pulse.

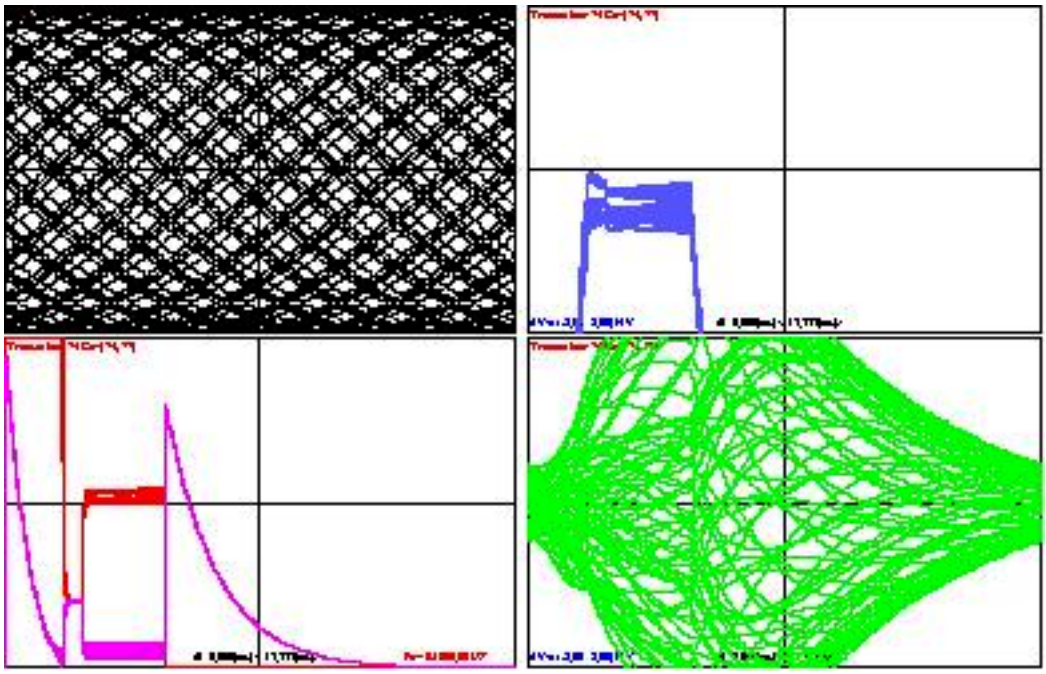

Figure 8: 'Overlay plot' of 20 pulses in a machine with 20Hz vibration stroke, beam loading and RF system errors (conditions as top left plot in Figure 17). It shows the cavity frequency drifts (top left), the real voltages (top right), imaginary voltages (bottom right) of the four cavities supplied by the first transmitter in the β=0.8 section ('number 74') and the power of the latter (bottom left); scales as in Figure 2 to Figure 5. The apparent regular frequency pattern has its origin in the choice of the test frequency as explained above, showing that the full phase range is probed. We get here an impression of the scatter of the cavity voltages in operation.

We have also included possible RF system errors. In the power distribution we have set a 5% error (in driving current amplitude) – resulting in unevenly excited cavities already without beam. For the scatter in Q_{ext} we have assumed 10%, hoping to choose the best candidates from the lot of the LEP2 cavities and paying special attention to the newly produced ones. This results in unevenly excited cavities without beam and different reaction under beam loading. These two errors correspond to a car that always pulls to, say, the left, but the driver can compensate a lot of it. However, the worst case is a partly blind driver, in our case a 5% field error in the vector sum, supplying erroneous information into the feedback system. Longitudinal cavity alignment errors are equivalent to phase errors and can thus be expressed in terms of the above errors. We have not included a special effect for them.

To avoid giving an enormous quantity of plots for 5 shots in 20 pulses (thus 100 phase space images per case) we have produced 'overlay plots'[3]. We have overlaid all macro-particles from all 'shots' of all pulses in the same frame showing the whole phase space area receiving particles during the operation of the linac, the density of 'colour' expressing the average density of particles during linac operation.

Figure 9, left, shows the perfect machine; there are no vibrations, no system errors and no beam loading effects; Figure 9, right, includes beam loading; the sweep during onset of the beam loading is clearly visible with the 400 'shots' we have simulated. For later plots we will only use 5 'shots' per pulse and there may be two clearly distinct 'blotches', the one generally closer to the centre corresponding to the first bunch (first 'shot') without beam loading. Since for 5 'shots' the next reference bunch comes 2.2ms/4=550μs later, the new equilibrium with beam loading is established for a long time already and all further bunches see this new field and cluster in a different 'blotch'. In reality the first few bunches of each pulse would sweep from the first 'blotch' to the second one, all following bunches arriving there, as in Figure 9, right.

Figure 10 shows a quiet machine without any vibration and no beam loading but with RF system errors. The linac parameters are exactly the official ones. We see that in at least one machine bunches are already well distorted without any vibration. If we apply now the minor modification – -15° instead of -20° phase angle in the first half of the β=0.8 section – we get under otherwise identical conditions the plots of Figure 11. A similar comparison can be done between Figure 12 and Figure 13, as for Figure 10 and Figure 11 but this time with beam loading. We see that bunches behave much better with this small modification and in the following simulations we will keep it.

Figure 14 shows an otherwise perfect machine; only 20 Hz vibrations are present. Bunches arrive in full phase space size nicely within the design limit for all four machines. Figure 15 shows the same conditions with beam loading, the two 'blotches' are hard to distinguish since each blotch is large enough to partly overlap the other one. Also bunches here arrive nicely within the design limit.

Figure 16 and Figure 17 show similar cases but including RF system errors. Some tails of the bunches are outside the ±10MeV agreed acceptance range. However, by a static machine adjustment the 'arrival area' can be centred to stay within the design limit except for the fourth machine, bottom right. This machine had to be trimmed; i.e. RF system errors had to be reduced to better approach the situation without errors in Figure 14 and Figure 15.

Figure 18 presents exactly the same case as Figure 17 but showing only the centroids. We see that the impression is much more optimistic, i.e. it is recommended that not only should the centroids be looked at but also the full bunch as in the other figures.

The following Figure 19 to Figure 23 present exactly the same collection with the only difference that a stroke of 40 Hz (instead of 20 Hz) was applied. Without system errors (Figure 19, Figure 20) bunches still arrive within the limits. However, for the other cases several plots show large tails outside the design limit. The

[3] here a movie would do an even better job

'blotches' are also so extended around the centre that a static shift as in the 20 Hz case cannot be applied any more for these cases. Since the cases without system errors in Figure 19 and Figure 20 are also very close to the limits, it would be nearly impossible to 'trim' the machine so perfectly as to hit the design specifications. As already said above, the centroids in Figure 23 look much more optimistic than the reality of the corresponding full bunch figures in Figure 22.

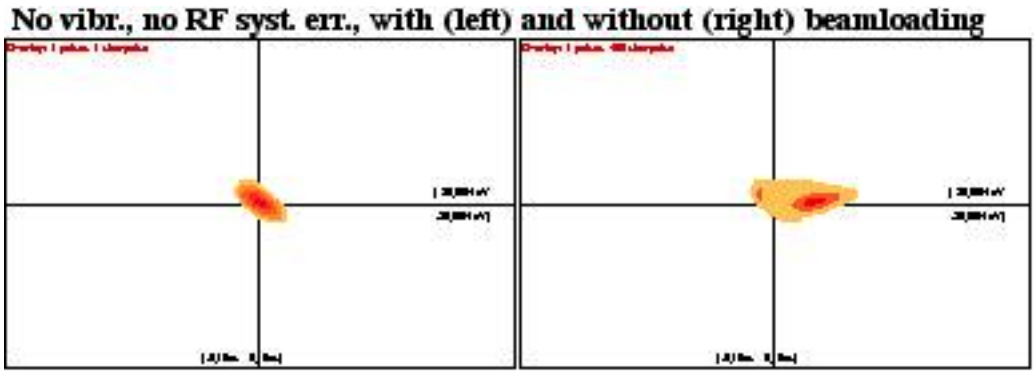

Figure 9: There is only one ideal machine without vibrations and without system errors. Left without beam loading (one shot is sufficient in this case), right with beam loading (400 shots), the sweep is clearly visible; range ±100 ps (horizontal), ±20 MeV (vertical).

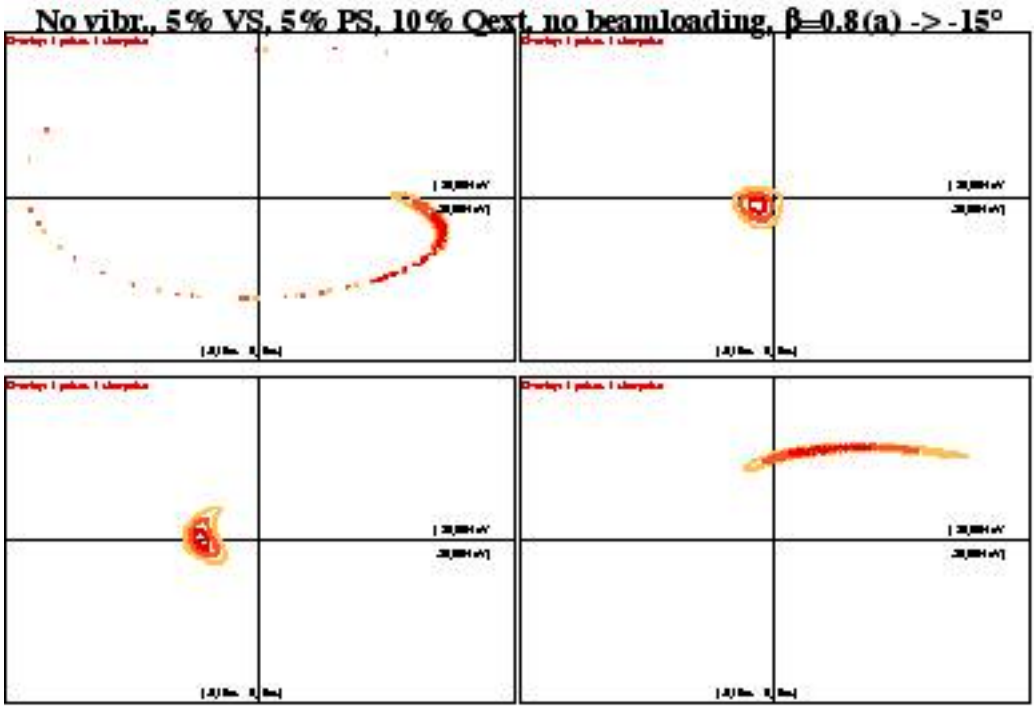

Figure 10: The quiet machine, no vibrations, no beam loading but system errors, with the 'official' phase angle –15° for the whole β=0.8 section; one shot is sufficient in this case. Range ±100 ps, ±20 MeV.

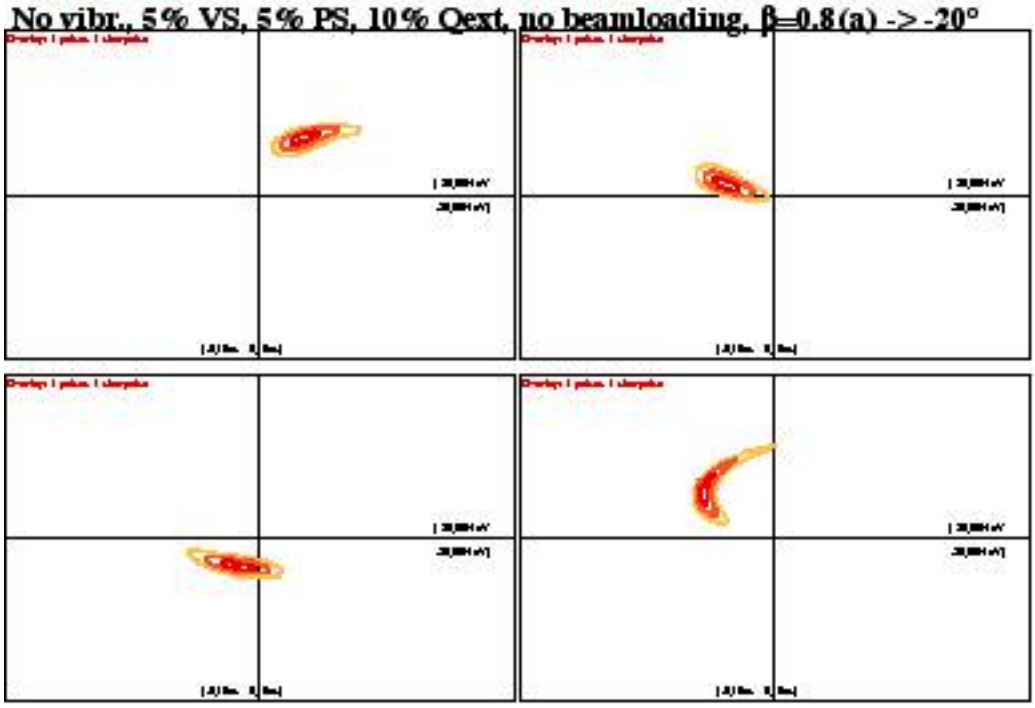

Figure 11: The same machine as in Figure 10 but with modified phase angle –20º for the first half of the β=0.8 section; range ±100 ps, ±20 MeV.

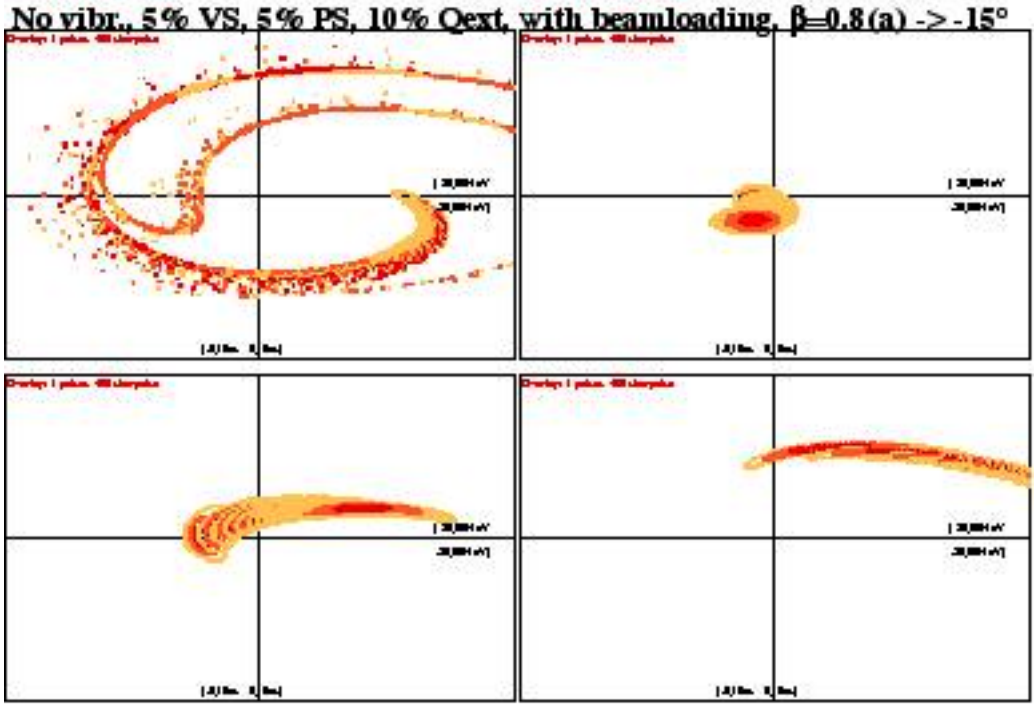

Figure 12: The quiet machine, no vibrations but with beam loading (400 shots) and system errors, with the 'official' phase angle –15º for the whole β=0.8 section. In the bottom right machine, most bunches arrive outside the plot to the right; range ±100 ps, ±20 MeV.

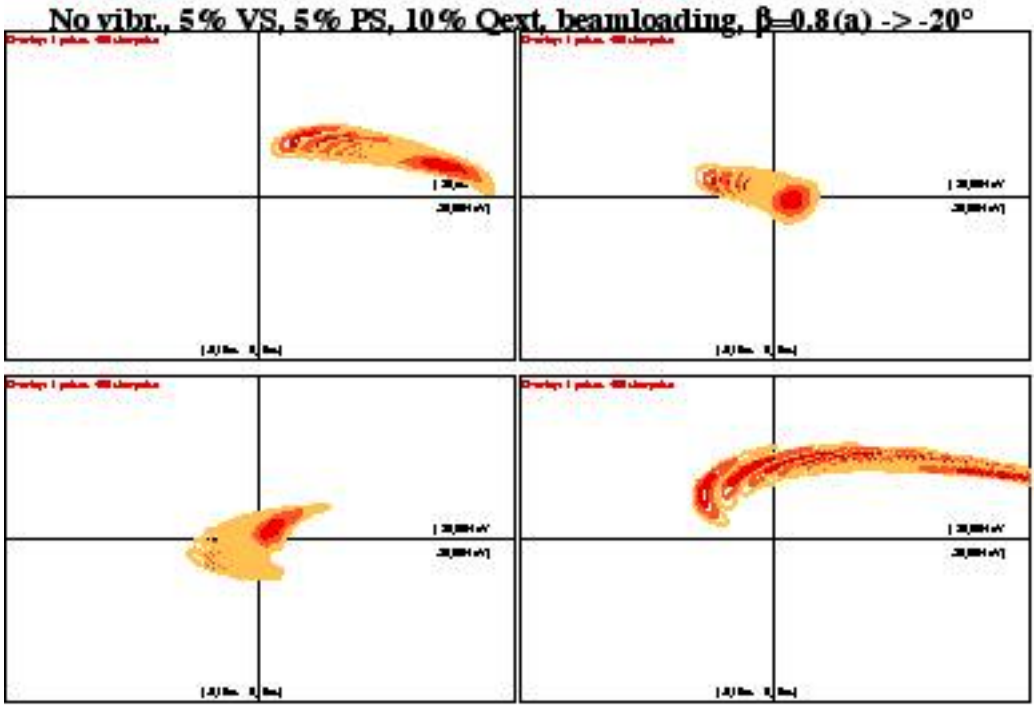

Figure 13: The same machine as in Figure 12 but with modified phase angle –20º for the first half of the β=0.8 section. In the bottom right machine, many bunches arrive outside the plot to the right side, but less pronounced than in Figure 12.

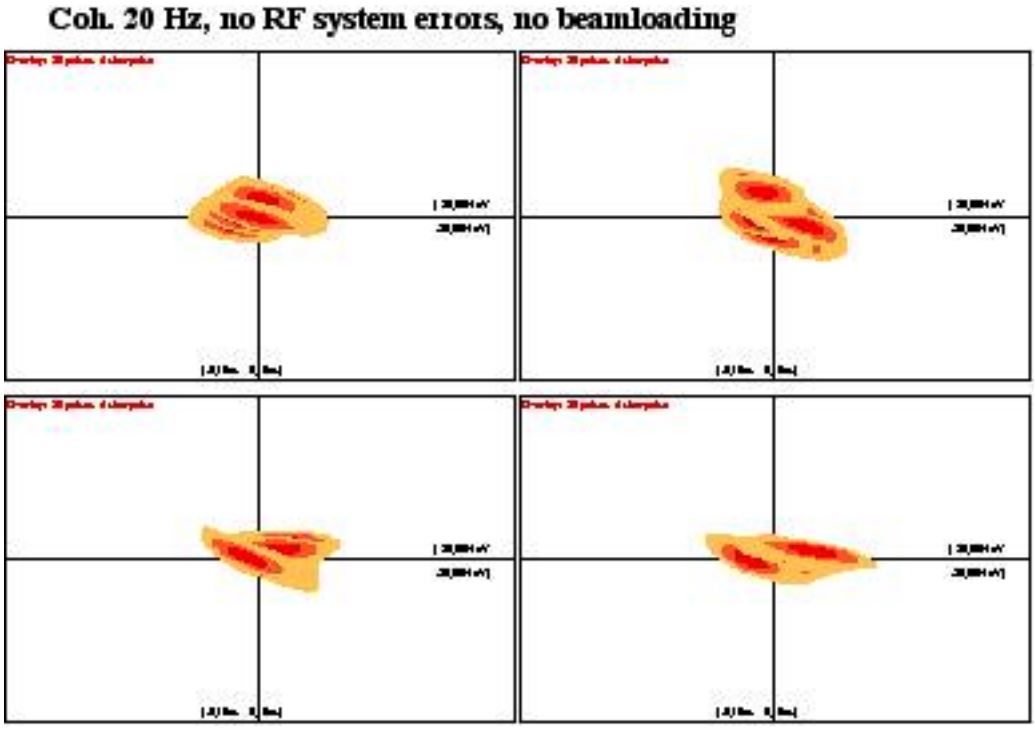

Figure 14: 20 Hz stroke vibrations, but no system errors nor beam loading; 20 pulses, 5 ‘shots’ per pulse, range ±100 ps, ±20 MeV.

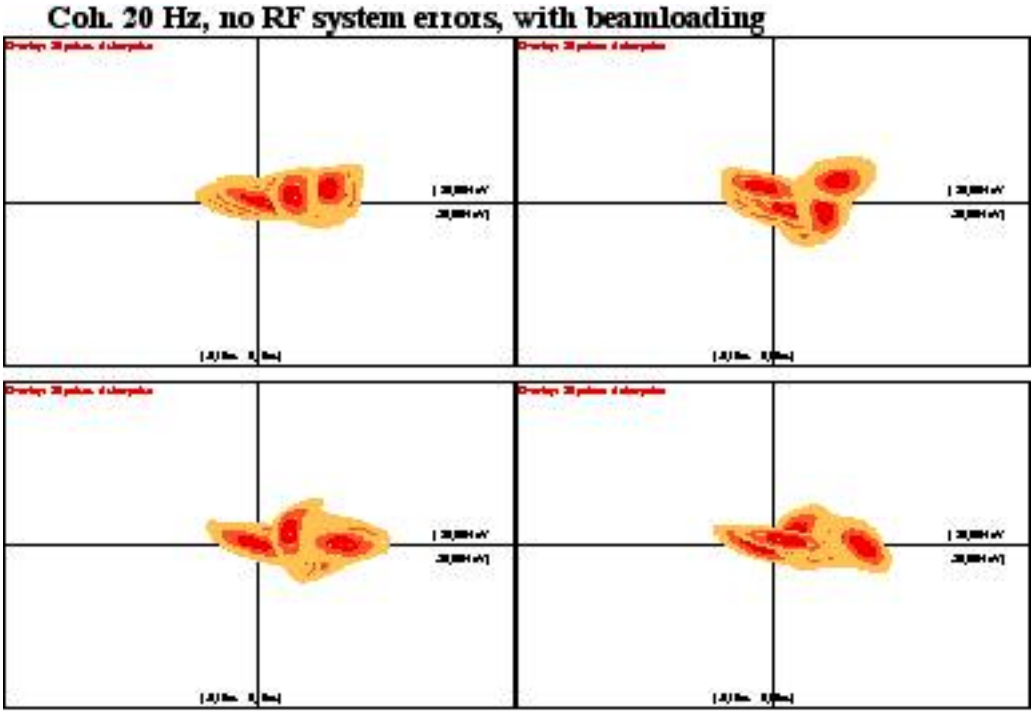

Figure 15: as Figure 14 but with beam loading

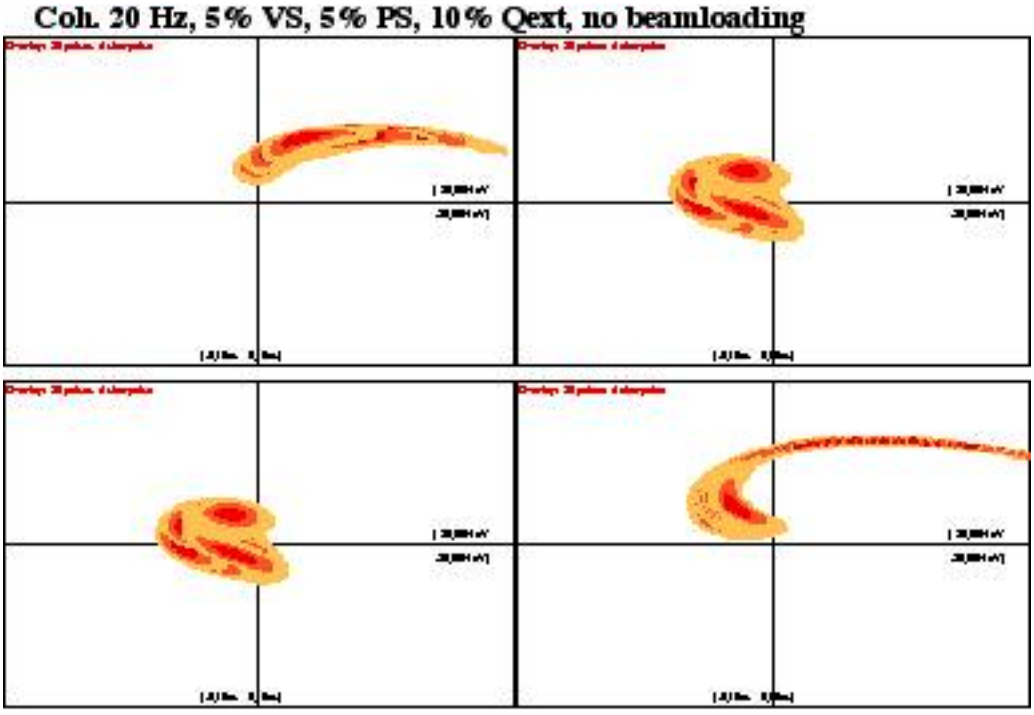

Figure 16: 20 Hz stroke vibrations and system errors but no beam loading; 20 pulses, 5 ‘shots’ per pulse, range ±100 ps, ±20 MeV.

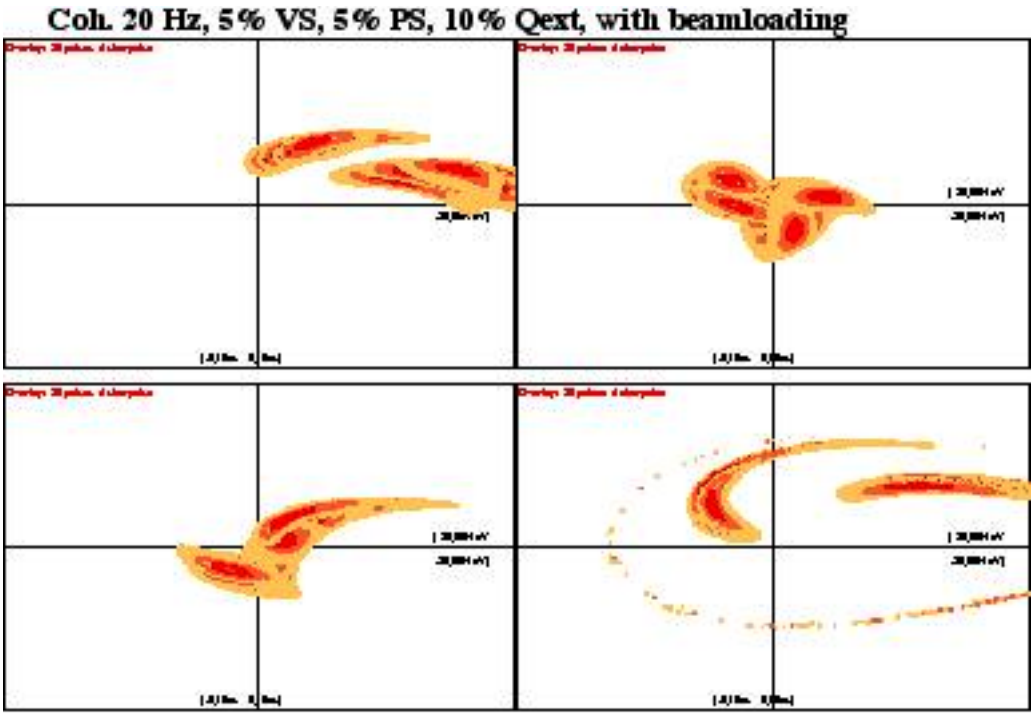

Figure 17: As Figure 16 but with beam loading (worst case)

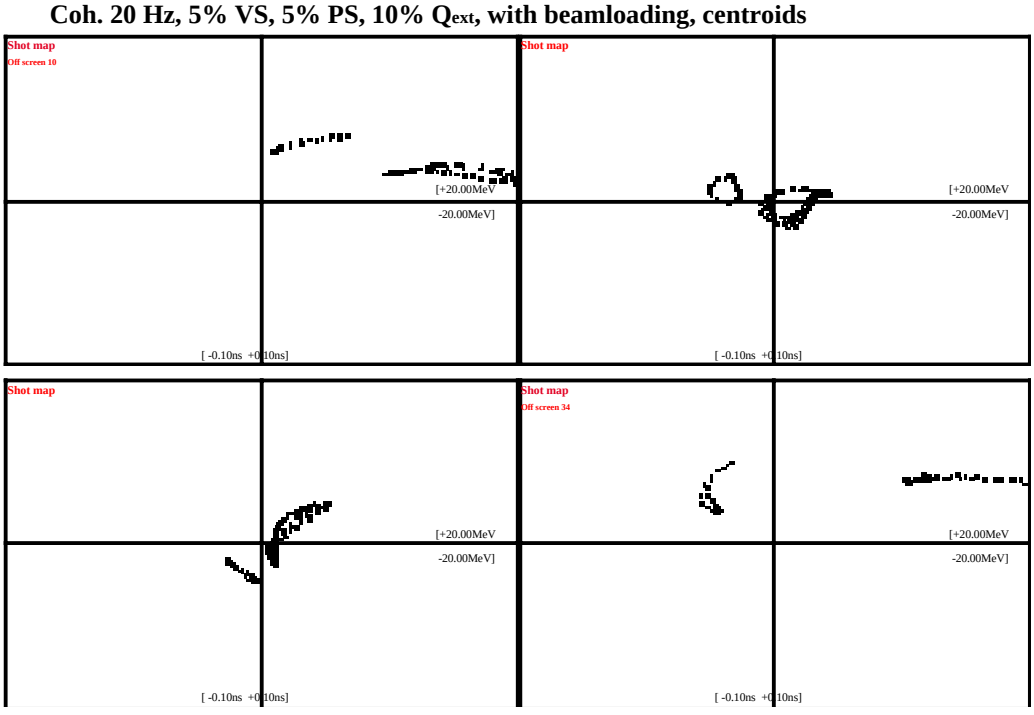

Figure 18: As Figure 17 but showing only the centroids, looking much smoother…

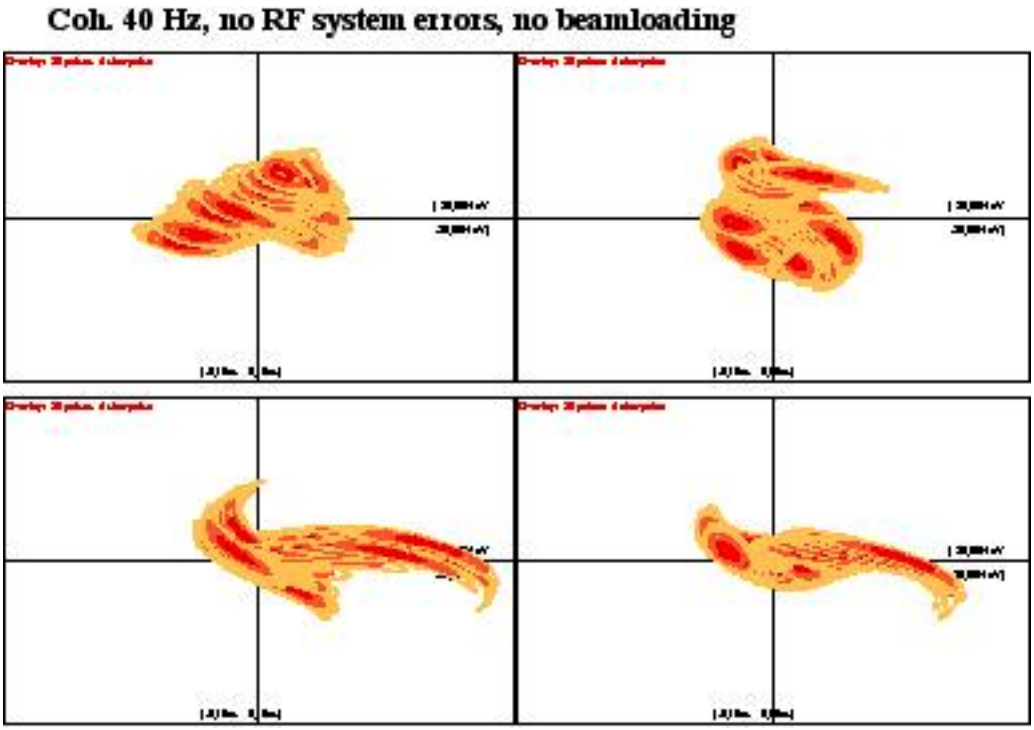

Figure 19: 40 Hz stroke vibrations, but no system errors, no beam loading; 20 pulses, 5 ‘shots’ per pulse, range ±100 ps, ±20 MeV.

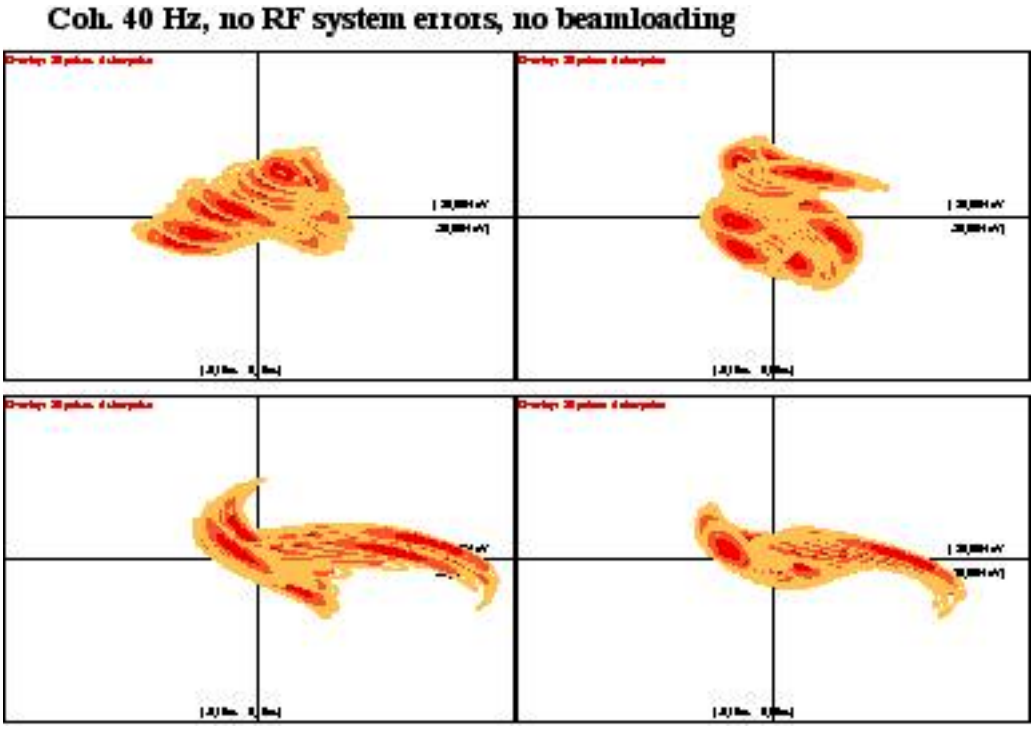

Figure 20: As Figure 19 but with beam loading

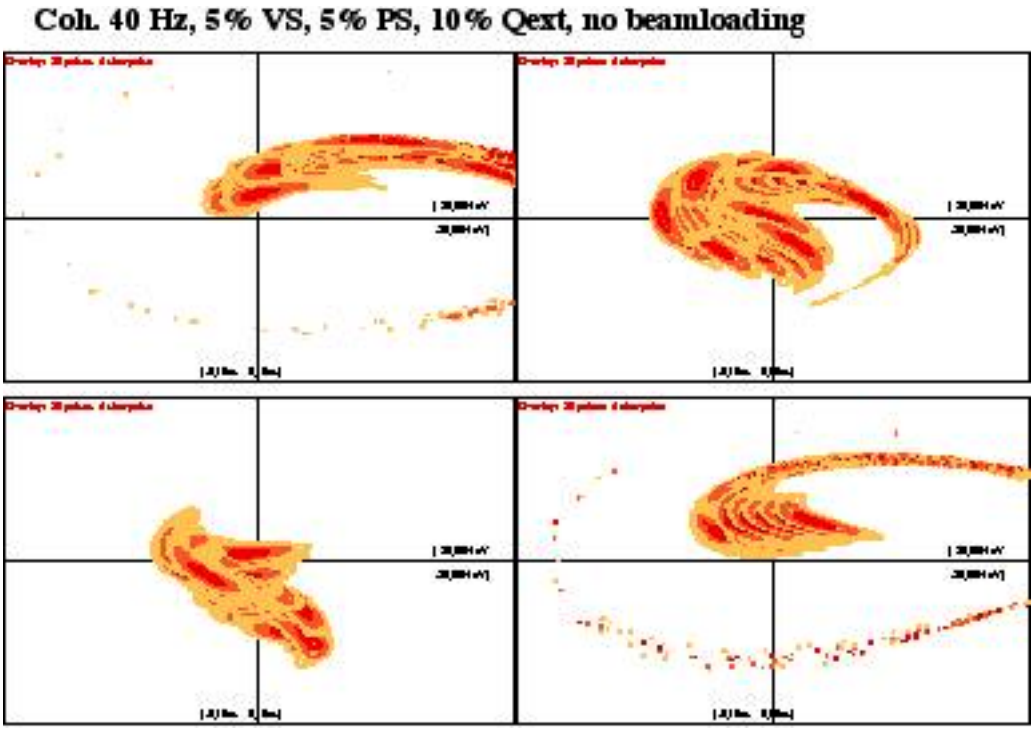

Figure 21: 40 Hz stroke vibrations and system errors, but no beam loading; 20 pulses, 5 ‘shots’ per pulse, range ±100 ps, ±20 MeV.

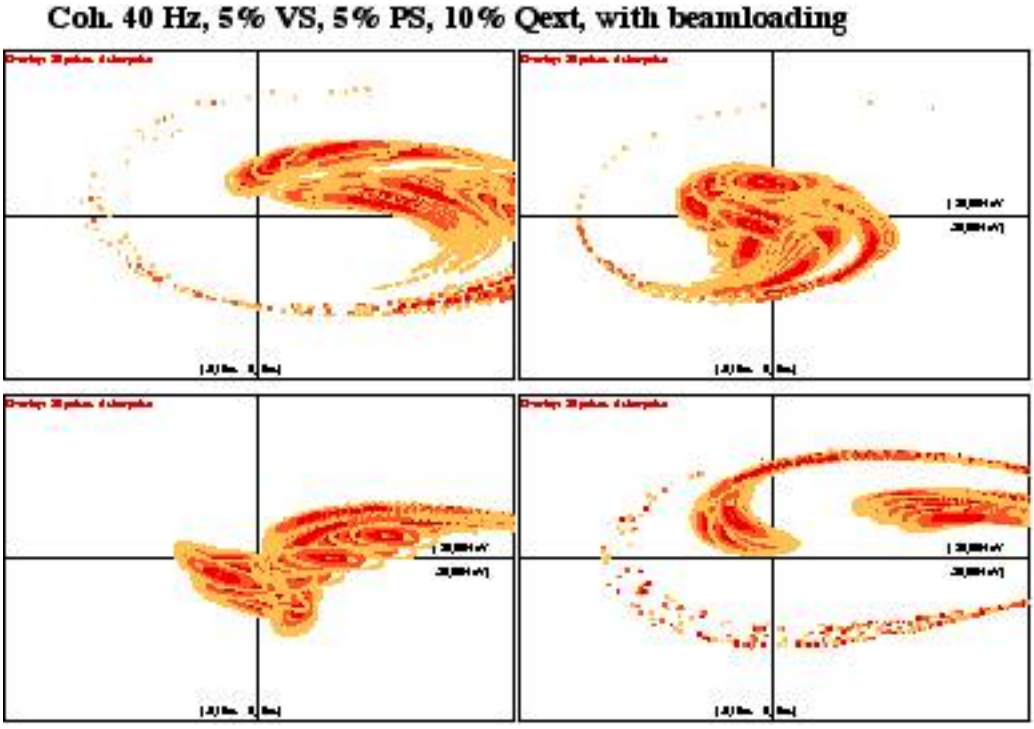

Figure 22: As Figure 21 but with beam loading (worst case)

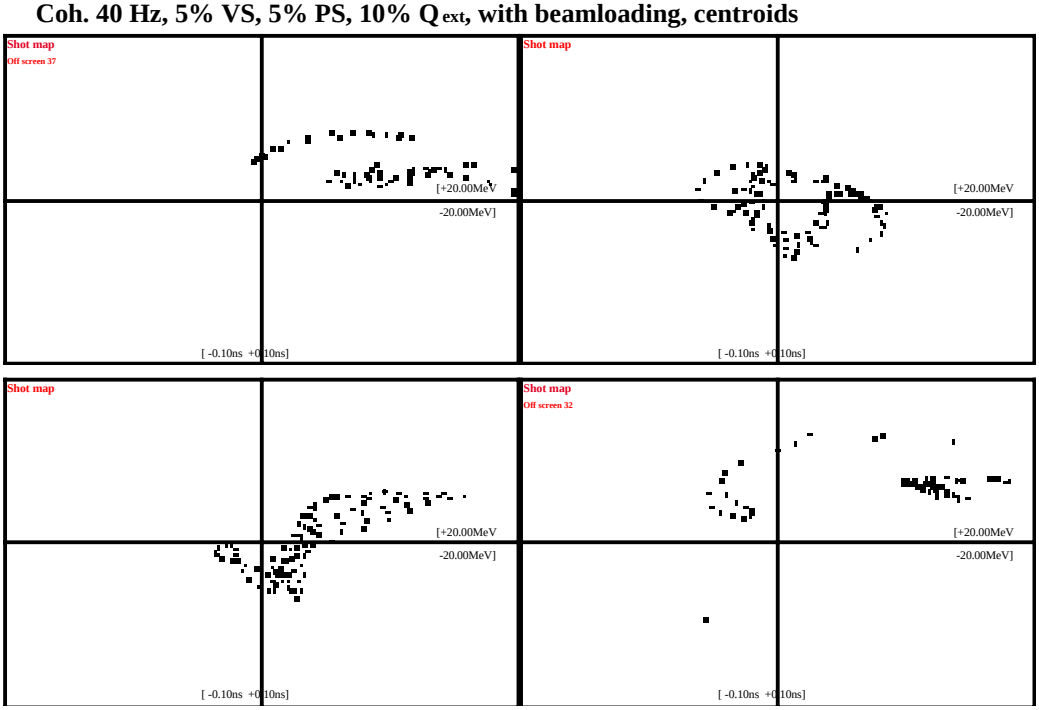

Figure 23: The centroids of Figure 22, looking much smoother…

6. First Examination of Lorentz detuning

When the RF field in a cavity rises, it creates forces on the cavity surface, which detune the latter; this effect is called 'Lorentz detuning'. It was intensively studied e.g. in the context of the TESLA cavities (1.3 GHz 9-cell cavities with stiffener rings, gradients of 25 MV/m or even more) [3] and solutions were proposed [4][5]. Since the LEP2 cavities were designed for a coasting beam with constant accelerating field, no particular features against Lorentz detuning are incorporated in their design.

Ponderomotive oscillations – also driven by Lorentz detuning – troubled the LEP2 RF system. This effect was analysed mathematically [6] and the Lorentz detuning constant was measured[4] to be in the range of k_L=–1 to –2Hz/(MV/m)^2. Using this value together with the measured bandwidth of the main mechanical resonance [7] (at about 100 Hz), indicating a Q_{mech} of about 10-20, gave a rough match between theory [6] and the observed onset of the oscillations in LEP. However, since a *precise* value of k_L is much more important here, this measurement should be repeated carefully for all cavity types foreseen.

In the β=0.8 cavities excited up to 9 MV/m (we assume the same constant k_L as for the LEP cavities, since they are very similar) the detuning for a constant excitation from zero to maximum would be about160 Hz. In the linac the excitation is present only for about 3 ms (averaged RF-up time) within 13.3 ms (repetition time), but the

[4] Measured were the field E, the He bath pressure p and the resonant frequency f while ramping the field slowly up and down many times. The Lorentz constant k_L was determined from the 3 parameter fit for the function $f = f_o + \alpha\ p + k_L\ E^2$

excitation repeats periodically and resonant enhancement is possible if the cavity resonant frequency is close to a multiple of the machine repetition rate. To estimate the effect we can make a simple model of a resonator excited by a 'rectangular force' between 0 and T_1 ('RF up') and oscillating free from T_1 to the repetition time T ('RF off'). We use directly the frequency stroke – not the mechanical deformation – as variable x

$$\ddot{x} + 2\sigma\dot{x} + \omega_o^2 x = \begin{cases} \omega_o^2 F & 0 \le t \le T_1 \\ 0 & T_1 \le t < T \end{cases}$$

with $\sigma = \omega_o/(2Q_{mech})$. We see that if we were to apply a constant F forever, the inhomogeneous asymptotic solution for $t \to \infty$ - where $\ddot{x}$ and $\dot{x}$ vanish – is simply $x_\infty = F$; thus we can determine directly $F = k_L E_{nom}^2$, and in our case we can write $F = -160\ \ Hz$. It is easy to show that the oscillation can be described as follows (ϖ is here the *complex* free mechanical resonant frequency $\varpi \approx \omega_o + i\sigma$, also containing the attenuation)

$$x_a(t) = F\left(1 - \frac{1-\exp(i\varpi(T-T_1))}{(1-\exp(i\varpi(T)))}\exp(i\varpi t)\right) \quad \{range\ \ a: 0 \le t \le T_1\}$$

$$x_b(t) = -F\left(\frac{1-\exp(-i\varpi T_1)}{(1-\exp(i\varpi(T)))}\right)\exp(i\varpi t) \quad \{range\ \ b: T_1 \le t \le T\}$$

We can plot this function for our case, adding a vertical line where the pulse starts and ends and calculating the frequency drift during this period as shown in Figure 24.

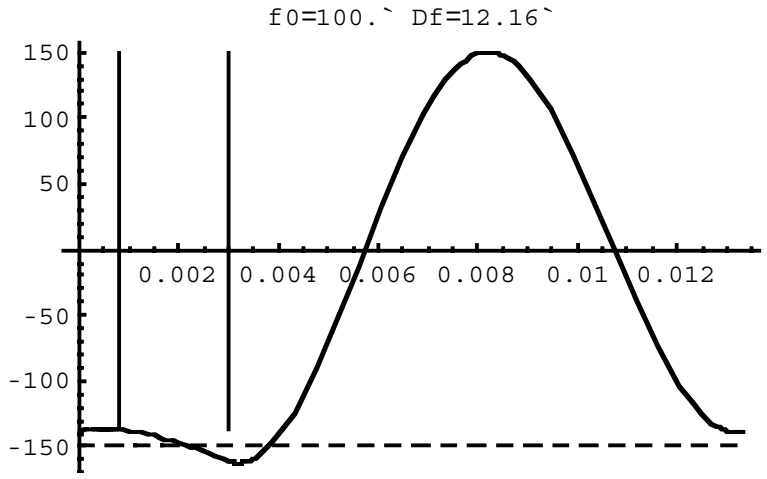

Figure 24: Model calculations: Oscillation of the resonant frequency under Lorentz force detuning with the 'rectangular force'; conditions as described in the text. Vertical scale in (RF-) Hz, horizontal time from 0 to 13.3 ms

We see that the total frequency swing is about ±150 Hz despite the short excitation of 3 ms only; the frequency moves by about 12Hz while the beam pulse is running.

The same case can be simulated closer to the truth with SPLinac. We have assumed that *only one* (not four as in the 'official' design) *cavity is driven by one transmitter* and all external vibrations have been 'frozen'. To reduce the power needed the cavity was statically pretuned by +200 Hz in such a way that the oscillating resonant frequency just passes the machine RF frequency when 'RF up' is required for cavity loading and the beam pulse. In Figure 25 two consecutive pulses are shown after the transients have died out (about 30 pulses) and a dynamic equilibrium has been established.

The cavity frequency goes up to a maximum of +165 Hz and down to a minimum of –236 Hz. The frequency swing is between –175 Hz and –235 Hz while the pulse is present, i.e. a stroke of about 60 Hz while the pulse passes.

We have also made a beam simulation, assuming that the whole machine has RF system errors and beam loading but 'freezing' all external vibrations. Figure 26 shows the overlay phase space plot at the end of the linac; large parts are out of specification.

We have simulated the same case but connecting *two* cavities to one transmitter. The plot similar to Figure 25 is not repetitive at all, i.e. there is no stable dynamic equilibrium any more. Higher or lower excitation of the previous pulse changes the stroke of the following pulse in a non-linear manner, resulting in a sort of 'chaotic system'. Figure 27 shows the frequency movement of these two cavities between pulse 30 – where the transients in the single cavity case have clearly died out – and pulse 60; voltages were not displayed to avoid overloading of the plot.

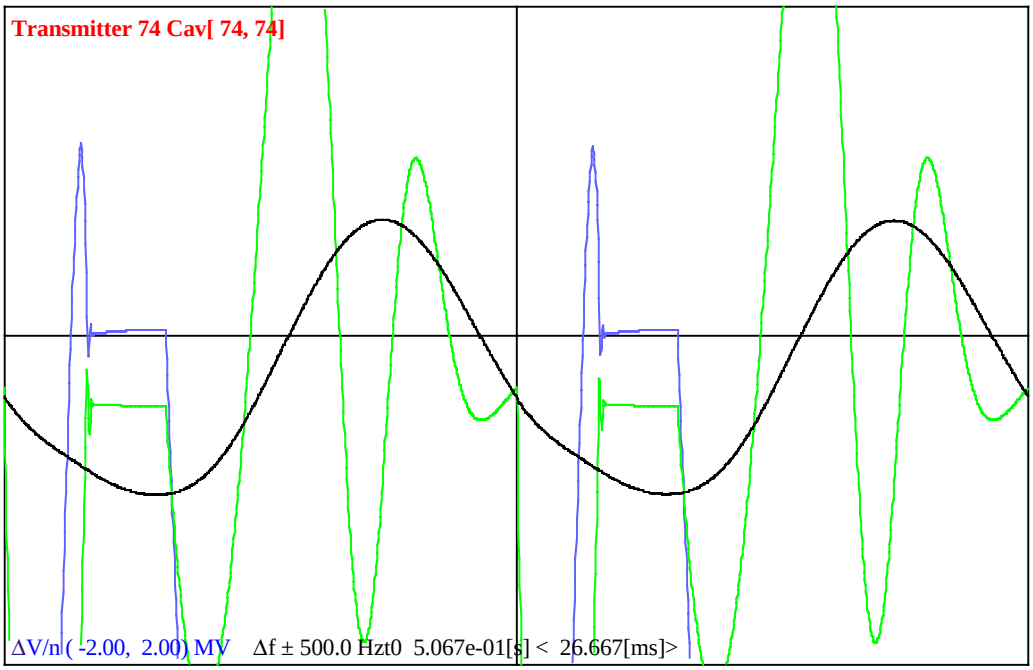

Figure 25: Two consecutive pulses once the dynamic equilibrium established (after about 30 pulses they become identical). Shown in black is the cavity frequency detuning (range ±500 Hz, the cavity bandwidth being 120 Hz), in light blue the real voltage, in light green the imaginary voltage ('phase'). Assumed are one cavity per transmitter and a static detuning of +200 Hz. Vertical scale ±2MV, horizontal scale 0-26.6 ms.

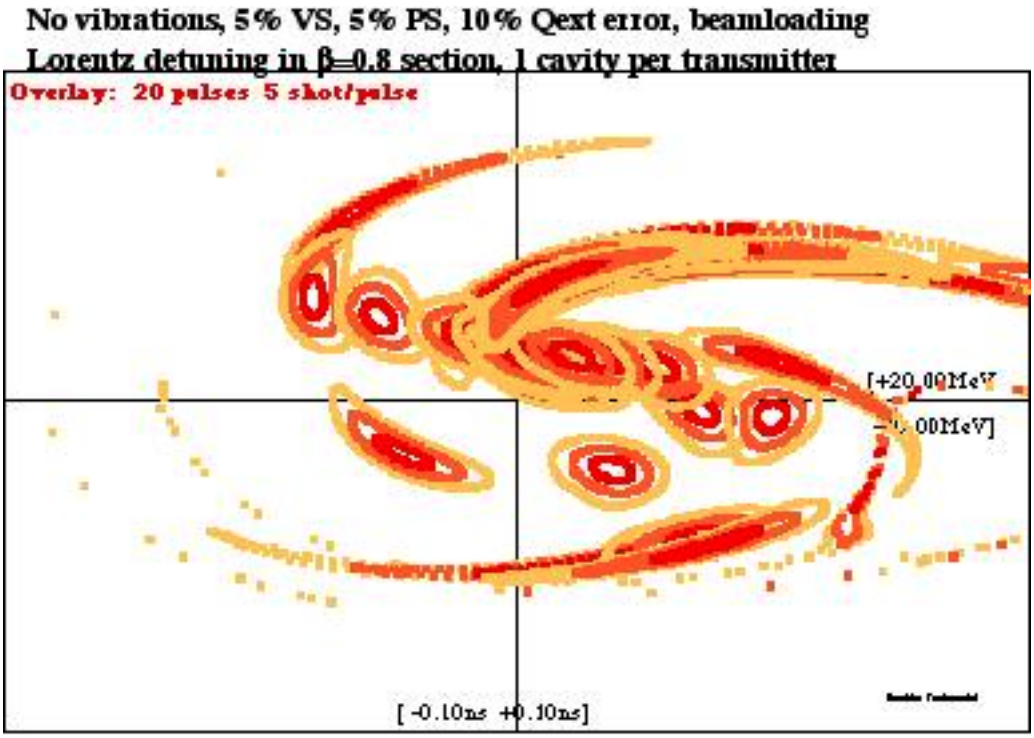

Figure 26: No vibrations but RF system errors in the whole machine and beam loading. Lorentz detuning considered in the β=0.8 section only (highest gradients) assuming only one cavity per transmitter there. Pulses are lost; 20 pulses, 5 ‘shots’ per pulse, range ±100 ps, ±20 MeV.

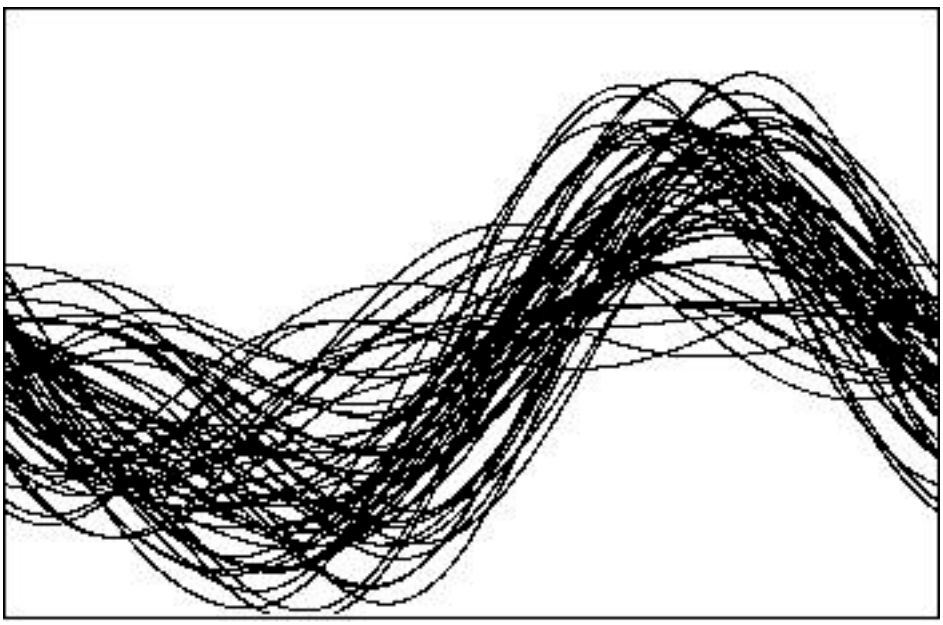

Figure 27: Frequency stroke for Lorentz force driven cavities as Figure 25 but with two cavities per transmitter. There is no true stable dynamic equilibrium any more. (Pulse 30 to 60; range ±500 Hz)

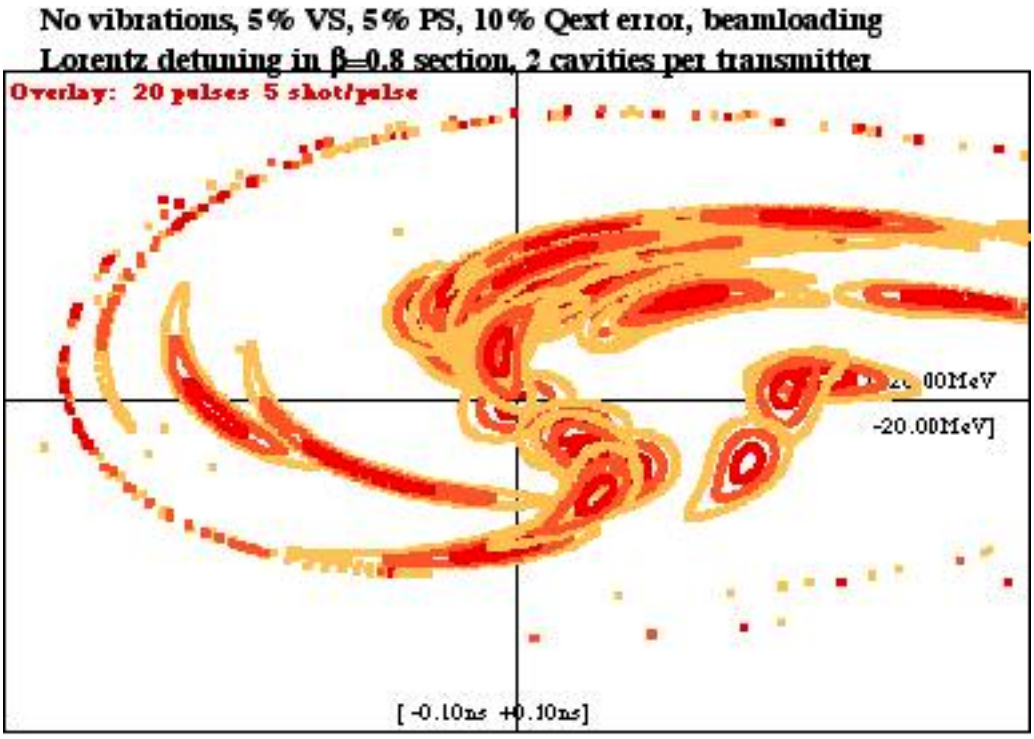

Figure 28: As Figure 26 but two cavities per transmitter in the β=0.8 section; several pulses are lost

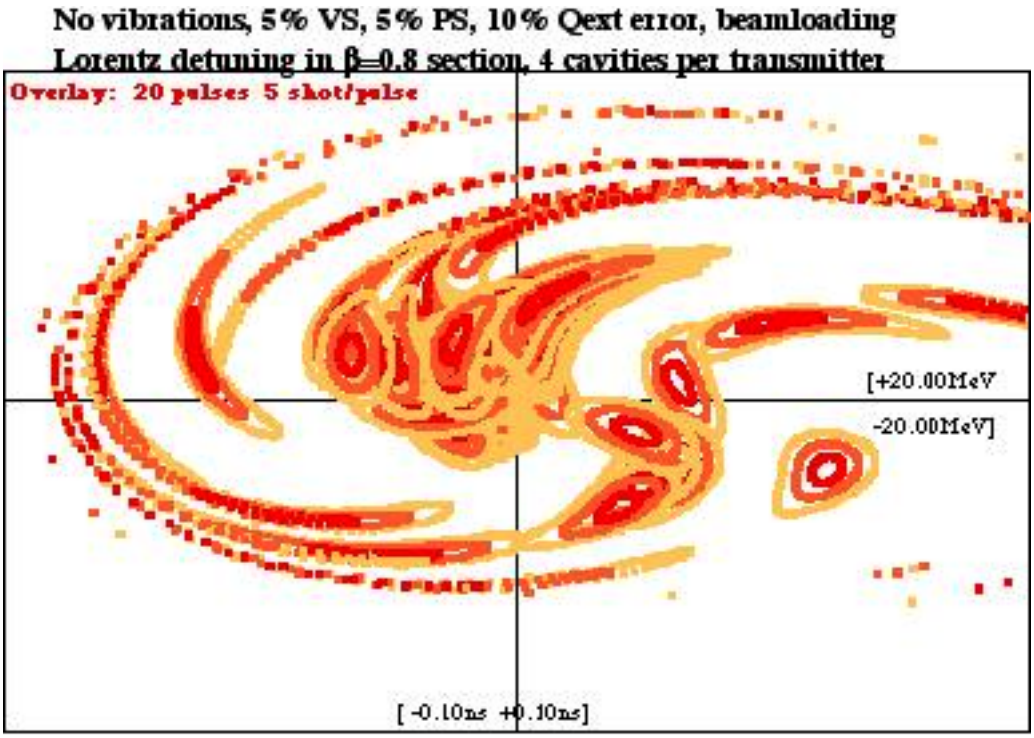

Figure 29: As Figure 26 and Figure 28 but four cavities per transmitter in the β=0.8 section (the 'official' case); more pulses are lost.

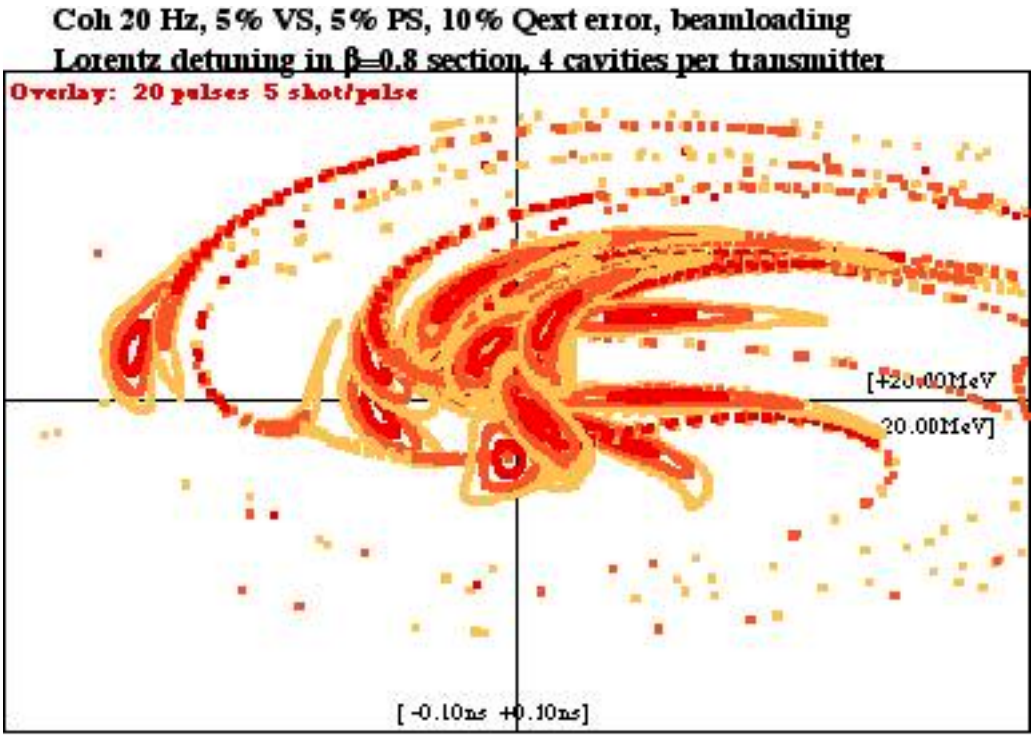

Figure 30: As Figure 29 but with 20Hz coherent oscillations in the whole machine (most realistic case). About half of the pulses are lost.

Under these conditions the beam quality is very bad and pulses are lost before they arrive at the end of the machine. We have compared the same machine ('top left' in previous plots) without any vibrations but RF system errors and beam loading with Lorentz detuning 'allowed' in the whole section β=0.8 for different conditions. One cavity per transmitter is shown in Figure 26, two cavities per transmitter in Figure 28 and four cavities per transmitter in Figure 29 while additionally coherent oscillations of 20 Hz stroke are present in Figure 30. These plots show that other methods have to be incorporated (e.g. feedforward) and tests with real cavities have to be done to demonstrate the feasibility of these methods.

7. Conclusion

Concerning microphonics we have seen that for a 'stroke' of 20 Hz nearly all machines show an acceptable situation. For the only case having larger deviations one should be able to approach the other cases by improvement of the RF system errors and stay perfectly within limits. The cases with 40 Hz stroke and a perfect RF system are still just in the limits. With realistic RF system errors the centroids stay acceptable but parts of the bunch tails are out of limits and it looks nearly impossible to trim the RF system so that all tails will be within bounds. Therefore it depends on the initial density distribution in phase space and the power capability of the collimator system if the situation is to remain acceptable; i.e. the 40 Hz case cannot be assured in every case.

Concerning Lorentz detuning the situation is not so clear-cut. For the case of a single transmitter per cavity the voltage is repetitive after a short transient of a few pulses so that static detuning lets the cavity oscillate through the nominal linac RF frequency while the RF is needed to load the cavity and sustain the beam pulse. Simulating beam passage in this situation – even with 'frozen' external vibrations – shows that bunches have a scatter in phase space larger than the allowed limits. Nevertheless, feedforward techniques as studied for TESLA might improve the situation.

However, when connecting two cavities to one transmitter (four or six are planned) and working with fast pulsing rate[5] (75 Hz) the oscillations do not approach a stable dynamic equilibrium (see Figure 27). Thus feedforward techniques – relying on prediction – cannot be copied straightforwardly from the 'one cavity per transmitter case', and further studies will be necessary.

Therefore, microphonics seems to be the lesser problem, provided the externally driven oscillations can be kept below a stroke of, say, 30 Hz, but the effect of Lorentz detuning needs further study. As a first step the Lorentz constant k_L should be determined more precisely, then a single LEP2 cavity has to be pulsed, followed by more than one cavity per klystron and finally also the fully equipped new cavities have to be tested in the same way.

References

[1] 'Conceptual Design of a High Intensity Superconducting H^- Linac at CERN', ed. A. Lombardi and M. Vretenar, CERN 2000-0012 ('Yellow Report')

[2] J. Tückmantel, 'SPLinac', A Program to Simulate SC Linac RF Systems with Beam, CERN-SL-Note-2000-053HRF

[3] B. Aune et al., Superconducting TESLA Cavities, Phys. Rev. Spec. Topics: Accel. And Beams, Vol 3, 092001 (2000)

[4] A. Mosnier, J.M. Tessier, Field Stabilisation Studies for TESLA, TESLA 94–16

[5] S. Simrock, Advances in RF control for high gradients, proc. IX Workshop on RF Superconductivity, 1-6 November 1999, Santa Fe (to be published)

[6] J. Tückmantel, Analysis of Ponderomotive Oscillations for the LEP sc. Cavities, SL Note 95-119(RF) and CERN LEP2 Note 95-36

[7] D. Boussard, P. Brown and J. Tückmantel, Electroacoustic Instabilities in the LEP2 Superconducting Cavities, CERN SL 95-81(RF)

[5] at TESLA with 5 or 10 Hz repetition rate the mechanical oscillation amplitudes are already strongly reduced till the next pulse arrives