TSL Note-2001-51

Design Considerations for a Feedback System to Control Self-Bunching in Ion-Storage Rings

V. Ziemann

The Svedberg Laboratory, 75121 Uppsala, Sweden

(Dated: February 26, 2001)

Abstract

We discuss the feasibility of a feedback system to cure self-bunching of the electron-cooled coasting ion-beam in CELSIUS. Such a system may also aid stable operation of accumulator rings for future spallation neutron sources or heavy ion rings used for inertial fusion energy production.

PACS numbers: 29.27.Bd

I. INTRODUCTION

The electron-cooled ion-beam in ion storage rings such as CELSIUS frequently shows self-bunching which is typically interpreted as the beam operating outside the stability region [1] due to its high intensity and small energy spread. The source of this instability are wake fields or, equivalently, impedances. The mechanism of the instability is that the beam excites electromagnetic fields in the beam pipe that affect subsequent sections of the beam. These later sections of the beam in turn excite new fields and so on. This mechanism obviously forms a closed loop that can become unstable which is undesirable in nuclear physics ion-storage rings, because it creates background for the experiments and it will contribute to particle losses in high intensity accelerators thus increasing the radiation load.

The energy change $\Delta E(z)$ or wake field a particle at longitudinal position z experiences is determined by the longitudinal charge density $\rho(z')$ of the beam weighted by the wake potential $W(z-z')$ that depends on the distance $z-z'$ between the affected particle and the section of the beam that excites the field

$$\Delta E(z) \propto \int_{-\infty}^{z} \rho(z') W(z-z') dz' \ . \tag{1}$$

There are two major contributors to the wake potential in low energy ion storage rings [2]. First, the space charge impedance and second, contributions from resonator-like structures such as cavities and bellows. Even though the space charge impedance is purely capacitive and does not drive instabilities below transition we will include it in the simulation. The real “culprits” that make the beam unstable and cause self-bunching are the resonator-like wakes and we will focus on canceling their adverse effect. To accomplish this, note that the expression for the energy loss given in eq. 1 is almost identical to that of a finite-impulse response (FIR) filter [3], where the in-signal corresponds to the earlier charge density and the wake function $W(\Delta z)$ corresponds to the filter coefficients. It is intuitively appealing to try to compensate the adverse effect of the wake fields by a feedback system that records the beam intensity using the sum signal from a beam position monitor (BPM), then digitizes this signal at a high rate, passes the signal through a FIR filter, and feeds it back to the beam via an amplifier and a longitudinal kicker structure. We will show the feasibility and features of such a system in the remainder of this report.

Note that such an active feedback system bears a little resemblance to the “SPEAR capacitor”. discussed in ref. [4] where a device with a specially tailored impedance was introduced in SPEAR in order to compensate adverse effects of wake fields. This device had limited success, possibly because the artificial wake was fixed. On the other hand, in a digital system, many different configurations can be easily tested by downloading new filter coefficients.

In the following section we will discuss the simulation code `ltrak`, results from simulations which will define the parameters of the feedback system, and a short discussion of technical issues, followed by a concluding summary.

II. SIMULATION

In order to make the simulation as realistic as possible we utilize a multi-particle tracking code that typically propagates 10 000 particles in the longitudinal phase space. As phase space coordinates we choose the phase ϕ which is proportional to the arrival time at a given position in the ring and the relative energy offset with respect to a reference particle $\Delta p/p$. The unperturbed dynamics of the particles is given by

$$\begin{aligned} \phi_{n+1} &= \phi_n + \eta \left(\frac{\Delta p}{p}\right)_n \\ \left(\frac{\Delta p}{p}\right)_{n+1} &= \left(\frac{\Delta p}{p}\right)_n \end{aligned} \qquad (2)$$

where the subscript n denotes the turn number and $\eta = \alpha - 1/\gamma^2$ the phase slip factor. $\alpha = 1/\gamma_T^2$ is the momentum compaction factor, γ_T the transition energy, and γ the relativistic energy factor of the particle. Note that we use the convention that η is negative below transition. In the code the head of the distribution is on the left side, towards smaller phase values, i.e. the beam moves towards the left. In this way particles with positive momentum $\Delta p/p$ are moving toward the left, when operating below transition. We also see from eq. 2 that the energy of the unperturbed particle stays constant and particles with positive energy offset $\Delta p/p$ will move towards smaller phase values, provided that η is negative.

Figure 1 shows the main panel of the program `ltrak` which runs the simulation and displays the particle distribution as the simulation progresses. In the main window in the top left corner the longitudinal phase space is shown with phase between 0 and 2π in the horizontal direction and relative momentum $\Delta p/p$ in the range $\pm 10 \times 10^{-3}$ in the vertical direction. The window below the main shows the longitudinal distribution, i.e. the signal that the BPM sum signal would show. The window on the right shows the momentum distribution. The upper status bar visible in the lower right corner shows the relative number of surviving particles and the lower status bar shows the percentage of the execution time. The small window in the lower right corner displays the number of surviving particles and the bunching factor which is the absolute value of a Fourier harmonic that can be specified in the program.

In the code the wake potential is represented as table with 256 entries that corresponds to the wake potential discretized over 256 points of a revolution. This requires that the wake has decayed over a revolution period or that the wake potential represents a “wrapped-around” version of the wake potential. The latter case is only a valid approximation if the distribution function changes slowly with respect to the revolution period, because when calculating the effect of a slice of beam that lies a few turns, say 3 turns back, we effectively use the distribution of the previous turn. Mostly the wakes are resonator like with impedances given by

$$Z(\omega)=\frac{R_s}{1+iQ(\omega/\omega_r-\omega_r/\omega)} \tag{3}$$

where R_s is the shunt impedance of the resonator, Q is its quality factor, and ω_r is the resonance frequency. Fourier transforming this expression leads to the following wake potential

$$W(\tau)=e^{-C\tau}\left(A\cos(D\tau)-B\sin(D\tau)\right) \tag{4}$$

with $A+iB=\omega_r R_s(1+i/\sqrt{4Q^2-1})/Q$ and $C-iD=\omega_r(1-i\sqrt{4Q^2-1})/2Q$. Note that the wake potential at the origin has the magnitude $\omega_r R_s/Q$ which translates into a energy loss of a delta-like charge distribution with charge Ne by $\Delta U=Ne\omega_r R_s/Q$.

The simulation of the space charge wake deserves some special attention. From ref. [1] we know that the longitudinal electric field due to space charge fields in a perfectly conducting beam pipe is given by

$$E_s=-\frac{e(1+2\ln b/a)}{4\pi\varepsilon_0\gamma^2}\,\frac{\partial\lambda}{\partial s} \tag{5}$$

where λ is the longitudinal charge density in particles/m, a is the transverse beam size and b is the beam pipe radius. $\varepsilon_0=1/Z_0c$ is the dielectric constant and Z_0 the impedance of free space (377 Ω). In the simulation we calculate the longitudinal charge distribution by binning the particles after every turn. The energy loss U_i for particles in bin number i due to space charge is then given by

$$\Delta U_i=Z_{\mathrm{sc}}\frac{I_{i+1}-I_i}{2\pi/M} \tag{6}$$

where $Z_{\mathrm{sc}} = Z_0(1 + 2\ln b/a)/2\beta\gamma^2$ is the space charge impedance, $M = 256$ is the number of longitudinal bins in one turn, and I_i is the current in bin number i. For a homogenous coasting beam we have $I_i = I_{\mathrm{total}}/M$ for all i. If we consider 48 MeV protons in CELSIUS the space charge impedance is about $10\,\mathrm{k}\Omega$ and if we assume that each of the macro-particles represents $1\,\mu\mathrm{A}$ we have $\Delta U \approx 0.4\,\mathrm{eV} \times (N_{i+1} - N_i)$. It is also instructive to calculate the Fourier-transform of the space charge wake where one obtains a function linear in frequency which means the the space charge wake acts like a high pass filter. As a matter of fact, this behaviour is not so surprising, because the the term $N_{i+1} - N_i$ looks like a discrete version of the derivative of a Dirac-delta function of which the Fourier transform is known to be a linear function in frequency.

In the simulation we have adjusted the space charge wake function and the resonator wake such that the contribution of the wake at the origin from space charge is ten times that of the resonator wake. Varying this ratio showed little influence on the growth rate. The resonator wake was chosen to be a $Q = 1$ resonator at the fundamental revolution harmonic with a strength such that the beam becomes unstable within 3000 turns in order to economize on the simulation time.

In the program it is possible to toggle a switch that causes the wake field to be displayed superimposed on the longitudinal distribution in the lower window on the main panel in `ltrak`. The zero is denoted by a horizontal line and the wake field autoscales to fill the vertical range. An example, which uses the same intitial distribution as that used in generating Fig. 1 is shown in Fig. 2. Note that the distribution is identical to that of Fig. 1, but superimposed is the energy gain (if above the horizontal line) or energy loss of the particles (if below the line). The maximum energy loss appears just behind (to the right) of the small peak in the middle of the graph. Note also, that the wake function is rather jagged which is a consequence of the space charge wake which acts as a high pass filter and makes the field more noisy. On the right hand side of Fig. 2 we show the distribution and wake after 3000 turns when using the same initial distribution, but having space charge fields turned off in the simulation. The wake is clearly much more smooth.

It is interesting to observe the behaviour of the bunching factor as the instability develops in the small windows at the bottom right of the main panel. As can be seen in Fig. 1 it shows exponential growth before saturating in the last quarter of the simulation. The growth rate can be analyzed with the aid of a feature hidden in the `Special` menu that is labelled

`Bunching Factor Growth rate.` Selecting that point will directly produce Fig. 3 in which the logarithm of the bunching factor is plotted as a function of the turn number together with a linear least squares fit in a selected region. The growthrate which is 451 turns in this case is displayed at the top of the graph. This feature will be used exhaustively later when analyzing the performance of a feedback system.

In the model we introduce a feedback system by using the sum signal of the BPM which corresponds to the horizontal projection of the distribution and passing it through a FIR digital filter whose filter coefficients can be specified by writing values to an array. The resulting kicks are multiplied by a gain factor and written to a ring-buffer that keeps track of the feedback kicks for last 8 turns with 256 kicks per turn. By picking kicks from the ring-buffer after a variable time delay and applying them to the particles it is possible to simulate the effect of a FIR-based feedback system on the beam. In a real implementation a buffer depth corresponding to two turns should suffice in order to allow adjusting the delay time over the range of a full revolution period.

In the same way as for the wake potential it is possible to specify the filter coefficients in `ltrak`, where the user can specify low, high, bandpass, or bandstop filters, load coefficients from file or copy the wake potential. Other features include digitizing the filter in time, which means applying a zero-order-hold on a data value for a number of samples. Moreover digitizing in amplitude is also implemented in order simulate finite word lengths if the filter is implemented on a fixed point DSP or programmable logic.

III. RESULTS

Initial tests with various filter types such as band or lowpass were rather unsuccessful to reduce the growth rate of the instability. Finally we realized that we could compensate the effect of the wake as given in eq. 1 by *using the wake function as filter coefficients for the FIR filter* and adjust the gain of the feedback in such a way that the kick applied to the beam is equal to that from the wake thus creating an artificial wake that compensates the real one. Figure 4 shows the growth rate of the instability as a function of the feedback gain when the filter coefficients are chosen to be the same as the coefficients that define the wake function. In Fig. 5 we show the kicks that the beam gets from the wake and from the feedback as a function of position inside the ring. Observe that the shape of the wake kicks is closely followed by that of the feedback kicks, except for the opposite sign. In this simulation the gain of the feedback was adjusted to provide kicks half the strength of that from the wake potential. This can be verified in Fig. 5 where we display the peak-to-peak values of kicks coming from the wake and from the feedback. The peak-to-peak value for the feedback is $1.37\,10^{-5}$ and from the wake is about twice as large, namely $3.07\,10^{-5}$. Obviously the wake is much more noisy due to the influence of the space charge component which is not included in the feedback coefficients.

In the previous section we applied the feedback kicks a multiple of the revolution time after we applied the wake kicks. In practice we do not know when and where in the ring the wake kicks hit the beam and in a practical implementation we need to control the delay between picking up the BPM data and kicking the beam. This delay parameter will then have to be optimized experimentally. In order to investigate the sensitivity of the feedback on this parameter we run the `ltrak` code and vary the delay time at which the feedback kicks are taken from the ringbuffer. In Fig. 6 we show the growth rate as a function of the delay time, represented here as a phase offset where 360 degree correspond to the full revolution period and negative phase means longer delay. Zero degree means that the feedback kicks are applied an integer number of revolution periods after the wake kicks. For unexplained reasons the growth-rate has an asymmetric behaviour. However, we can observe that a delay that is a little too large (negative phase) will cause the growthrate to increase dramatically, whereas a decrease of the delay by 35 degree will still prevent the startup of the instability. In any case we have a window of about 40 degrees in which the feedback is functioning well which means that the tolerance for this parameter is rather loose. For CELSIUS at injection energy the revolution time is about $1\,\mu$s and the tolerance for the delay is thus on the order of 100 ns. Furthermore, we found that shifting the delay time by multiples of the revolution time does not significantly change the behaviour of the feedback system. This can be attributed to the fact that the longitudinal beam distribution changes rather slowly on the time scale of several 100 turns whereas the delay time is varied over a few turns, only.

So far we always employed filter coefficients that are given as floating point numbers. In a real system where the filter calculations have to be made at a rate of a few tens of MHz we can gain some speed by reducing the word length in the calculation. We simulate this by digitizing the filter coefficients in steps of powers of two and found that only the average energy loss is changed which can be easily compensated by adding a constant to the filter coefficents. The growth rate, on the other hand, is not affected even if the filter coefficients are digitized in 8 steps, that is 3 bit. Using a low number of bits makes it possible to use a look-up table approach for the multiplications needed in the FIR filter calculations, provided the number of bits used in the digitization of the BPM signal is also small. However, in order to be on the safe side we will try to use 6 or 8 bits in the ADC and filter coefficients which will also make matching the dynamic range of the BPM and the ADC less critical.

Next we will discuss the effect of sampling and processing the signals at a lower speed. In the preceeding sections we always used 256 sample points per turn for the beam intensity, the wake, and the filter coefficients. This would correspond to using a 256 MHz data processing system when running at a revolution frequency of 1 MHz. This exceeds the bandwidth of today's ADC and filters and one thus has to settle for slower processing speeds. In the code `ltrak` we have implemented sampling and processing at a lower rate by using only every n-th sample of the beam intensity signal and also updating the feedback-kick ringbuffer only every n-th sample which is then held constant for n samples. An example of the kicks applied to the beam after 3000 turns for the case where only 16 kicks per turn were applied is shown in Fig. 7. Observe that the feedback kicks come in steps which are 16 bins long whereas the kicks from the wake field are quasi continuous.

Judging the performance degradation due to a slower feedback system using the bunching factor as a figure of merit proved difficult. In Fig. 8, we therefore choose to show phase space after 3000 turns for feedback systems operating a different sampling speed. In the top left corner we show phase space when the beam was kicked 128 times per turn (the kick is held constant for two consecutive beam slices as described above) and reduced the number of kicks per turn in steps of factors of two down to 4 kicks/turn. In the top row we have 128, 64, and 32 kicks per turn, and in the bottom row 16, 8, and 4 kicks per turn. We observe from the longitudinal profile – the projection in the lower part of the respective pictures – that the bunching gets progressively worse as we reduce the number of kicks per turn, but also that the momentum spread – the width of the vertical projection on the right of the respective pictures – increases. It appears that the distributions in the top row are tolerable, but in the bottom row are not, because the energy spread increases dramatically. In any case, we see when comparing to Fig. 1 that the feedback is beneficial in all cases with the exeption of the last two in the bottom row.

When experimenting with different parameters we found that we could improve the performance of the feedback system by decreasing the delay-time or, equivalently increasing the phase, c.f. Fig. 6, by about 35 degrees. The pictures with reduced delay time corresponding to those displayed in Fig. 8 are shown in Fig. 9. We observe that self-bunching is somewhat reduced and that the momentum spread in the first picture in the second row is considerably smaller compared to that in Fig. 8. This means that we can operate a feedback system which only samples the beam at a rate which is 16 to 32 times the revolution frequency. In CELSIUS at injection energy this means we can operate the feedback at about 30 MHz. This relieves the demand on the ADC and the FIR filter drastically.

During the investigation described in the previous paragraph we observed that the average energy applied to the beam by the feedback is considerable and the beam's distribution moves up or down in phase space. This behaviour can, however, be compensated by subtracting or adding a constant to the filter coefficients.

The required power can be estimated from the amplitude of the kick applied to the beam in order to compensate the growth of the instability for a given bunching factor. In the simulation the growth rate of the un-compensated instability is about 500 turns or about 0.5 ms. The peak-to-peak value of the wake kicks is about $\Delta p/p \approx 10^{-6}$ when the bunching factor is 1 %. This translates to an energy change of about ± 1 kV delivered to the beam. If the growth-rate of the instability is smaller the required voltage is reduced proportionally. This also translates into a reduced current that can be stored without self-bunching.

Next we investigate the effect of feedback amplifier power saturation on the behavior of the feedback system. To this end we use the standard simulation scenario and limit the maximum kick that the feedback can deliver to the beam. We find that the feedback can prevent the instability from starting as long as fluctuations are small. Once a fluctuation has created a sufficiently large inhomogeneity the feedback is, however, unable to contain it. The details are a function of the instability's growth-rate and the amount of fluctuations in the beam.

A practical difficulty comes from the inability to accurately know the shape of the wake function. In a real implementation we have to guess the shape which invariably leads to differences between the real wake that drives the instability and the filter coefficients that constitute the artificial wake. We are thus interested in the deterioration of the performance of the feedback system when there is a mismatch between the frequency or Q-value of the wake and the feedback. Since the wake functions for different frequencies (harmonics) look rather similar we tested to correct a wake on the first harmonic with a feedback filter operating on the second harmonic, but that did not work at all. So, as a rule the harmonic of the wake and the feedback have to agree. In a next step we try to understand the performance of the feedback system if there is a mismatch of the $Q-$value between the wake and the feedback system. To this avail we set up the wake to be a $Q = 5$ wake at the first harmonic, add the space charge wake and then try to compensate with a feedback system that operates at the first harmonic but vary the $Q-$value between 1 and 10. In almost all cases is it possible to compensate the instability almost perfectly. Only when we have a feedback system with $Q = 1$ do we see considerable bunching on the second harmonic. The first harmonic is properly damped here as well. This is not so surprising, because the $Q = 1$ resonator is rather broadband and even when it cancels the instability on the first harmonic does it contain sufficient spectral power on the second to drive an instability there.

Having discussed the requirements and limitations of such a feedback system we will briefly collect the relevant parameters that would be used in such a system for CELSIUS at injection energy with a revolution frequency around 1 MHz.

- 30 MHz sample rate.
- 6 or 8 bit data word width.
- 60 filter taps (2 times sample rate over revolution frequency)
- a few 100 V peak amplitude.

At peak energy in CELSIUS the revolution frequency is about tripled such that the sample rate and the number of taps is also tripled. Since self-bunching mostly plagues low-energy beams we will focus mainly on the low-energy scenario laid out above and will discuss technical implementation issues in the next section.

IV. PRACTICAL ISSUES

The central item in the implementation of the feedback system is the FIR-filter that serves as an artificial wake field and must operate at very high sample rates. There are recent developments in the field of programmable logic devices that make FIR filters operating above 100 MHz with filter lengths of 256 taps, 16 bit word length possible [5]. As shown above we need only 60 taps and word lengths of less that 8 bit at speeds of about 30 MHz such that using a FPGA prototype board with on-board ADC and DAC will be suitable [6]. These boards are originally designed for the telecommunication market, but can equally well be applied in accelerator based research. In CELSIUS the bandwidth of the front end of the BPM operates at bandwidths well over 50 MHz such that we can directly use these signals. There is a longitudinal kicker structure in the ring that can be used to apply the kicks to the beam. The available voltage is about ± 60 V without problems and ± 120 V with some extra effort.

One of the first experiments with such a feedback system will be to use it as an artificial wake to make a low intensity beam unstable and observe the growth-rate of the instability. It is important to check that the correct harmonic is excited and whether the growth-rate scales with feedback gain in the manner expected from theory [1]. In this way the feedback system acts as an artificial wake function that can be used to investigate the accuracy of simulation codes for instabilities similar to that discussed in ref. [2].

Once the expected behavior of the feedback system is established one can use it to counteract self-bunching that occurs at higher beam currents. First one has to observe at which harmonic self-bunching is most strongly visible by observing the zero-span signal of a BPM sum signal on a spectrum analyzer at different harmonics and determine which harmonic is worst. From the data of zero-span power versus time we can determine the growth-rate of the instability. From a snapshot of the longitudinal distribution, visible on a digital oscilloscope, we can deduce the bunching factor. Taken together this information can be used to estimate the strength R/Q of the wake which in turn will determine the gain of the feedback system such that the magnitude of the applied kick is of the same order of magnitude. Since the magnitude of the $Q-$value is uncritical we would choose it to be about 5. Together with the knowledge of the dominant harmonic this will allow us to specify the filter coefficients. The last parameter to set is the delay time of the feedback signal which can be adjusted to minimize the occurrence of self-bunching. Subsequent iterations of the parameters may be used to improve the situation further. Having stabilized one harmonic one can repeat the procedure on the next stronger self-bunching harmonic.

V. CONCLUSION

We have determined the design parameters for a longitudinal feedback system in order to limit self-bunching instabilities in low-energy ion storage rings such as CELSIUS. Such a feedback system is based on creating a digital artificial wake using an ADC, a fast FIR filter based on programmable logic and a DAC. A BPM sum signal is fed into such a system and the output voltage is applied to an amplifier feeding a longitudinal kicker structure. We found that sample rates of about 30 MHZ, data word lengths of 6 or 8 bits and filter length of 60 taps are sufficient to stabilize the beam at injection energy in CELSIUS. We found that choosing the filter coefficients identical to those of the wake potential led to near perfect cancellation of self-bunching. A major drawback is that the wake potential is normally not known. In order to cure this deficiency we outlined a procedure that allows to roughly determine the filter coefficients experimentally. If such a tuning procedure worked satisfactorily it could certainly be automatized to make the feedback self-adapting which would also constitute a measurement of the wake potential. But that remains to be seen after we have secured funds to build and test such a system.

REFERENCES

[1] A. Hofmann, "Single Beam Collective Phenomena: Longitudinal," *In *Erice 1976, Proceedings, International School Of Particle Accelerators (CERN 77-13)*, Geneva 1977, 139-174.*

[2] G. Rumolo, et.al., "Comparison between theory and simulation for longitudinal instabilities of coasting beams," Nuclear Instruments and Methods, **A415,** 411, 1998.

[3] E. Ifeachor, B. Jervis, "Digital Signal Processing," Addison Wesley, 1993.

[4] K. L. Bane *et al.*, "Bunch Length And Impedance Measurements In Spear," SLAC-PUB-4620 *Presented at 1st European Particle Accelerator Conf., Rome, Italy, Jun 7-11, 1988.*

[5] In the XtremeDSP section on `www.xilinx.com`.

[6] `www.gvassociates.com`

LIST OF FIGURES

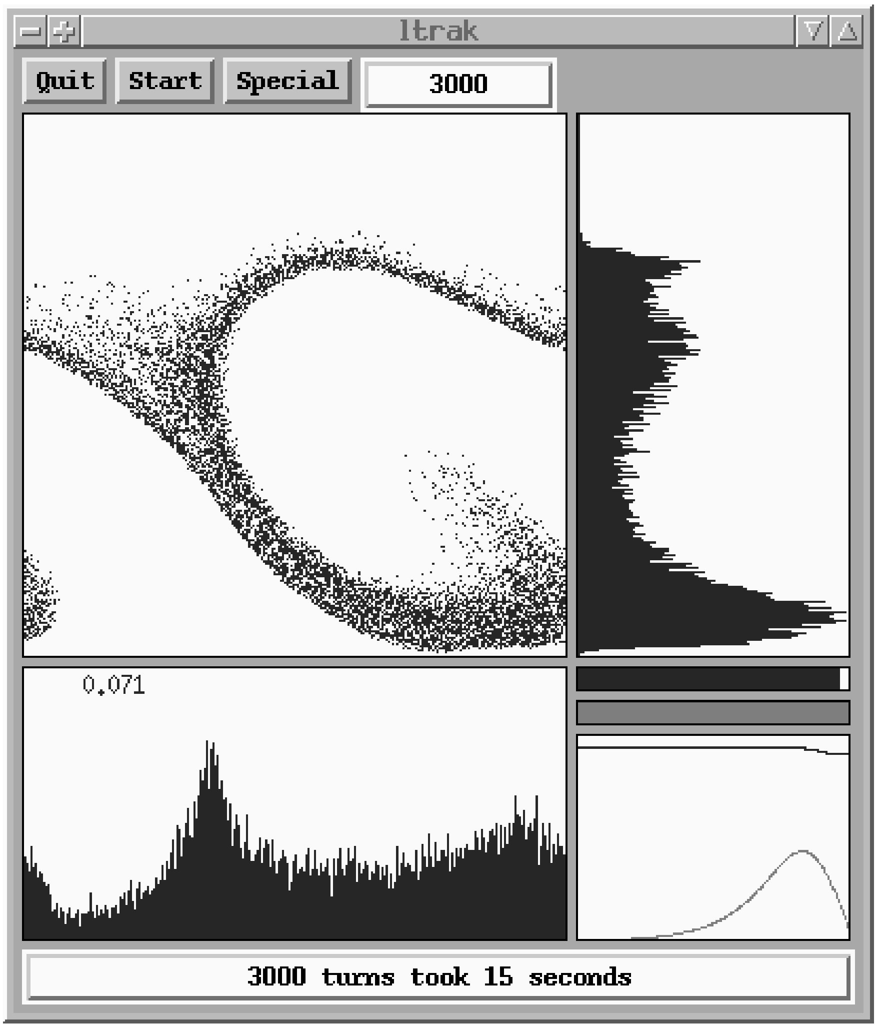

FIG. 1: Output of the simulation code `ltrak` shows the longitudinal phase space and projections after 3000 turns. See the text for details.

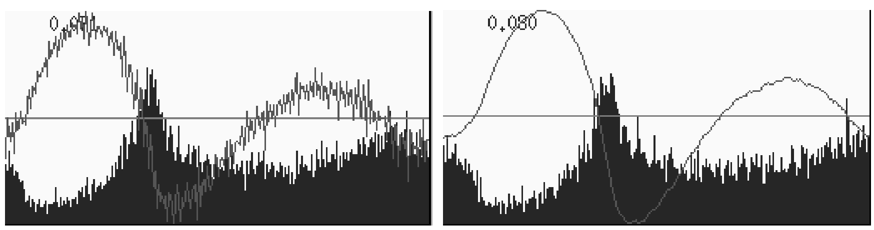

FIG. 2: Longitudinal distribution and the corresponding wake field shown on the left in Fig. 1. The right plot shows the same simulation with space charge turned off.

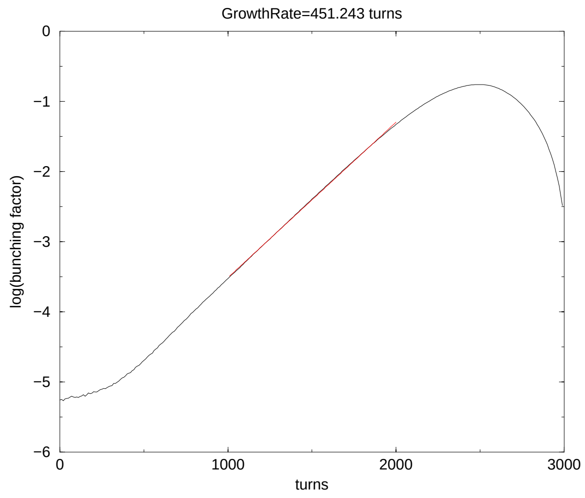

FIG. 3: The logarithm of the bunching factor as a function of the turn number. A linear least squares fit to the slope in a finite region yields the growth rate of the the mode that is unstable.

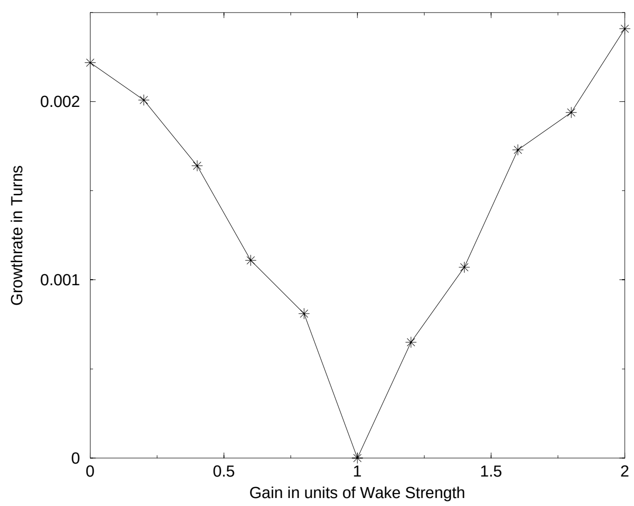

FIG. 4: Growth-rate of the bunching factor as a function of the gain of the feedback system. The filter coefficients are identical to the wake function ($Q = 1$ resonator at first harmonic plus space charge). Note the perfect cancellation when the gain is identical to the strength of the wake.

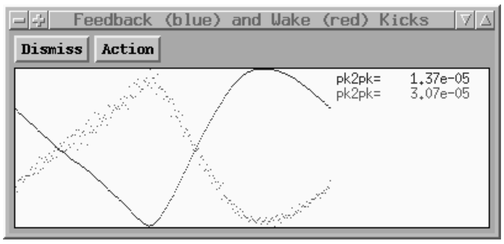

FIG. 5: The kicks from the wake and the feedback. Note that the feedback kicks are just the negative of the wake kicks.

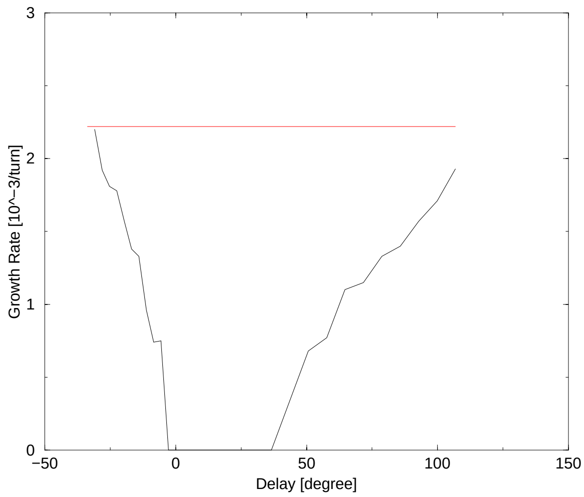

FIG. 6: The growth-rate of the bunching factor as a function of the time delay given as phase (360 degree correspond to the revolution time). The horizontal line denotes the growth-rate in the absence of feedback.

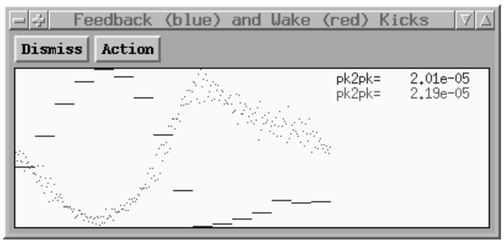

FIG. 7: The kicks from the wake field and the feedback system for the case where the feedback kicks the beam 16 times per turn only. These are the kicks corresponding to the first picture in the second row in Fig. 8.

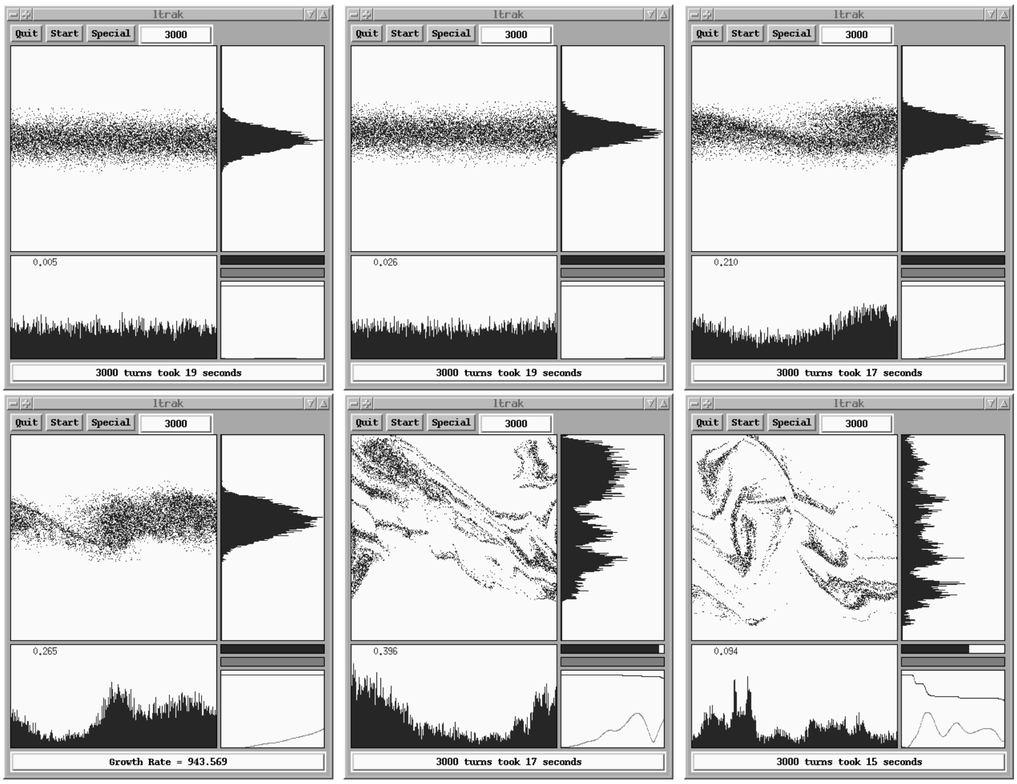

FIG. 8: The individual pictures show the longitudinal phase space after 3000 turns with 10000 particles with wake and feedback. The feedback uses from 128 kicks/turn (top left), 64 kicks/turn to the right down to 4 kicks/turn (bottom right) in power-of-two steps. Compare with the corresponding system without feedback in Fig. 1. See the text for details.

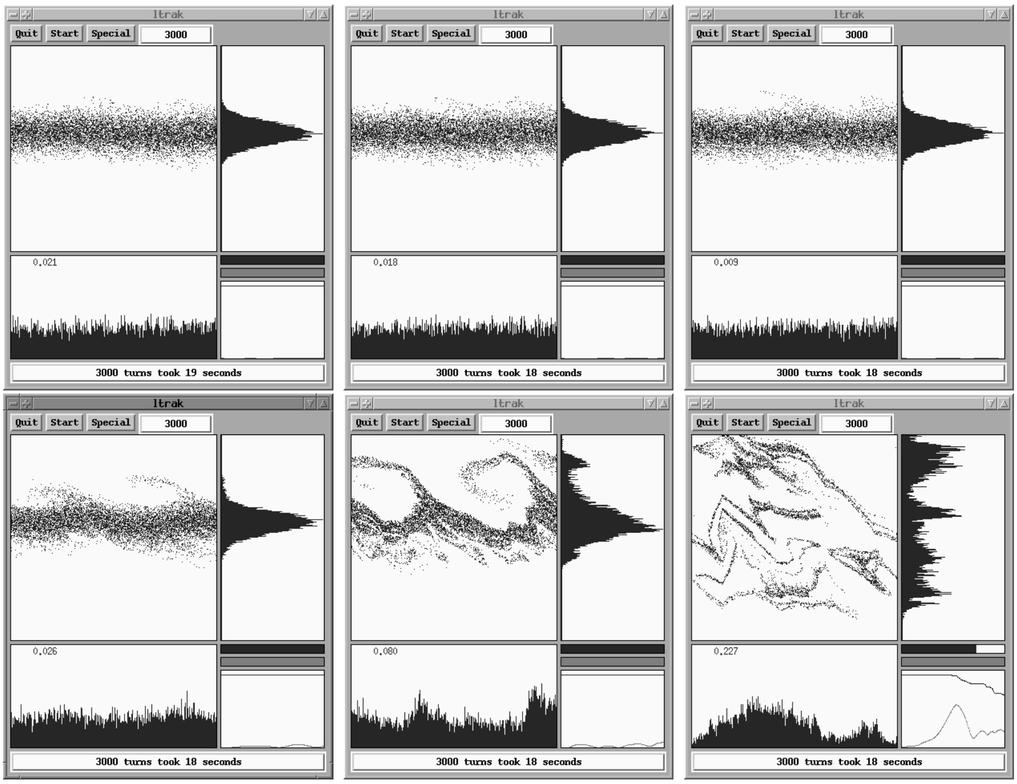

FIG. 9: The pictures corresponding to those in Fig. 8 where the feedback delay was reduced by a tenth of a revolution ($\approx 35^o$.)