ADAPTIVE RF TRANSIENT REDUCTION FOR HIGH INTENSITY BEAMS WITH GAPS

J. Tückmantel, CERN, Geneva, Switzerland

Abstract

When a high-intensity beam with bunch-trains and gaps passes a cavity with a high-gain vector feedback enforcing a constant voltage, large transients appear, stressing the RF high power hardware and increasing the trip rate. By modulating the cavity voltage with a varying periodic waveform (set-function), the RF power can be made constant while still preserving the high feedback gain. The average cavity voltage is conserved but bunches have to settle at slightly shifted positions. A method is derived to obtain this set-function in practice while making no assumptions or measurements of the beam or RF parameters. Adiabatic iterations are made including the whole machine as an analog computing device, using all parameters as they are. A computer simulation shows the success of the method.

THE PROBLEM

An RF vector feedback (Fig. 1, top) with high gain is necessary in any high current synchrotron to prevent longitudinal coupled bunch instability due to the main impedance. For constant enforced voltage a high intensity beam with long bunch-trains and gaps induces large RF power transients. These require an increased installed RF power, more mains consumption and particularly stress all RF high power components, hence increasing the risk of trips with total beam loss. Therefore these transients should be largely reduced if this is compatible with other requirements of accelerator and detectors. A similar problem was encountered and handled at PEP II [1].

THE SET-FUNCTION DEFINITION

An RF system as in Fig. 1, top, is assumed. Bunches in coast may have diverse (equilibrium) shapes and charges, including zero in gaps. A first 'Gedanken-Experiment' is executed with this beam, but starting with an intensity scaled down so much that beam loading can be considered non-existent. Then cavity voltage and RF power are constant, bunches are at their nominal position. Furthermore, conditions remain unchanged when opening the feedback loop between the red and green triangle at (α) and injecting at the red triangle a constant drive wave 'd', identical to the (previous and present) signal 'm', the cavity probe signal minus the constant set-value V_0.

Now adiabatically bunch charges are scaled up again. To keep the cavity voltage on *average* at its nominal value, the constant drive 'd' may have to be adjusted; it is suitable, but not indispensable, to also detune the cavity for (average) reactive beam loading compensation. While scaling up the bunch charges, the cavity voltage will start dithering around its average and bunches will *slightly* drift away from their initial nominal position. Meanwhile the RF power remains perfectly constant along the beam revolution, up to fully re-scaled bunch charges.

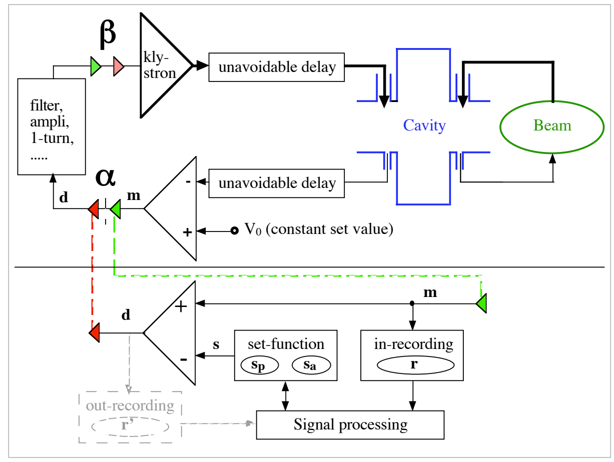

Fig. 1: Top: the principal blocks of a fast RF vector feedback. The cavity field is driven simultaneously by klystron and beam, the latter reacting itself to changes of the cavity field (coloured triangles to be ignored yet). Bottom: the 'Smoother': Incoming signals (green) can be recorded in a cyclic storage r matching one revolution period. A similar cyclic storage works as an *active* set-function s_a: synchronized to the beam, its contents turned out in a never-ending loop and fed into the comparator of unity gain. A *passive* set-function s_p allows hidden manipulations and, once ready, both set-functions can be swapped instantly. The dashed 'out-recording' may be used as an *alternative* to the 'in-recording'.

Now a 'Smoother', as described in the caption of Fig. 1, bottom, is added, first joining the two green triangles only; it records 'm', periodic* with the revolution period, as in-recording r. With a 'frozen' r the set-function s_a=r-d is made active. Then the 'Smoother' gives (red triangle) d'=m-s_a=m-(r-d). But m=r (periodicity) and hence d' is *constant* and equal to d, despite the dithering of the probe signal. s_a is the required set-function. Closing the loop by joining the red triangles is now completely transparent: the high gain is re-established with a *constant* RF power, (considerably) below previous peaks. The *average* cavity voltage is conserved but bunch positions have slightly drifted away from the nominal ones.

For LHC, even at ultimate beam intensity, these drifts in bunch position remain very small compared to the

* in reality with an open loop the LHC beam would get unstable

bunch length. Therefore the resulting change in proton-proton collision position and time is very small compared to the natural uncertainty: this RF manipulation remains imperceptible for the physics experiments [2].

It is obvious that this s_a is unique for the given boundary conditions; when these slightly change, e.g. by intensity loss in coast, small transients will reappear. Adapting s_a to new boundary conditions will cause tiny beam perturbations[†], hence this process should not run permanently but only rarely when considered worthwhile.

During *slow* energy ramping (as in LHC) the system is practically periodic in short term: an s_a can be iterated that is perfect for that instant. While ramping further, 'm' may start to deviate from the previous recording 'r' and small transients will start to show up. A new s_a may be iterated once in a while, hence also in slow ramp the RF power can be kept close to constant avoiding large power spikes.

In reality s_a cannot be determined as done above: a high intensity beam will go unstable when opening the loop. Since set-functions act inside a high gain loop, any manipulation error is amplified: a well-designed procedure chosen with the utmost care is necessary.

THE ADAPTIVE METHOD

In theory s_a could even be calculated and fed into the 'Smoother', provided the parameters of all bunches, RF system(s) and machine optics were *perfectly* known, either by assumption or measurement. However, any discrepancy to reality is amplified by the loop gain, making this a very difficult enterprise. Also other ideas relying on simultaneity at a reference-point encounter the problem of signal transmission properties to this point and need calibration.

For the previous open loop case the signal 'm' deduced as a constant term d was used successfully as a set-function providing constant RF power output. But this works only if previously bunches have been drifting adiabatically precisely to their new equilibrium position, compatible with constant RF power. This is not the case here and the activation of such a set-function, even adiabatically, would produce different transients but just as large.

To circumvent all these difficulties, we use the machine, RF system and beam at large as a sort of analog computer and iterate the set-function with it, embedding parameters as they truly are, simultaneity being intrinsic.

The first 'Gedanken-Experiment', just before closing the loop, can be considered as a feedback system with *zero* gain. This leads to the idea to first smooth the transients by slightly[‡] and adiabatically lowering the loop gain g. Then instantly g is switched back to the initial g_0 while *simultaneously* the set-function is modified such that the output of the 'Smoother' remains unchanged, conserving the smoothing of transients.

A second 'Gedanken-Experiment' deploys the same hardware with the 'Smoother' included in the closed loop, and full beam. The comparator is fed at (+) with 'm' and at (-) with the active set-function s_a, the high power chain then being driven by $d=m-s_a$. Initially s_a is set to zero, corresponding to the 'classical' system with $d_0=m_0$ with nominal loop-gain g_0, showing large RF transients.

Now the gain γ of the comparator, normally unity, is lowered *smoothly*, adiabatic for the beam, by a small amount to $\gamma=x<1$ (e.g. x=0.9), the loop gain being lowered by the same factor[§]. Bunches are drifting to slightly shifted positions, m_0 smoothly transforms to m_1 and the drive becomes $d_1=x\cdot(m_1-s_a)$, the loop-gain never being below $x\cdot g_0$, preventing any beam instability.

A (stable) measurement m_1 for one turn is frozen as $r\equiv m_1$ and the passive set-function $s_p=r\cdot(1-x)+x\cdot s_a$ is determined with it. Then *simultaneously* γ is set back to unity and s_p is made active. For the signal m_2 the new drive is $d_2=(m_2-m_1)+x\cdot(m_1-s_a)=(m_2-m_1)+d_1$. If $m_2\equiv m_1$ is true also $d_2\equiv d_1$ holds: the switching cannot be detected outside the 'Smoother', everything runs as before. After such a step all transients are reduced, corresponding to $g=x\cdot g_0$, but the full loop gain g_0 is recovered.

Instead of instant switching, the 'return path' to $s_a \to s_p$ and $\gamma\to 1$ could be executed slowly, even consecutively, but always adiabatically[**] enough to avoid any beam perturbation: the final state will be the same. Then one complete step would consist of four parts: 1) $\gamma\to x$; 2) stabilize, determine s_p; 3) $s_a \to s_p$; 4) $\gamma \to 1$.

In a sequence of such steps one ends by ramping from $\gamma=x$ to $\gamma=1$ while the next step starts by ramping from $\gamma=1$ to $\gamma=x$. Dropping this useless double operation yields the new sequence (• $\gamma\to x$); •stabilize, determine s_p; • $s_a \to s_p$; • stabilize, determine s_p; • $s_a \to s_p$; …. and so on.

The first unique ramping $\gamma\to x$ can even be left out: one can imagine that starting with an even higher gain $g_0'=g_0/x$ it was already done. This even economizes on the (difficult) hardware for a smooth gain ramping.

One last point remains: each such step also reduces the apparent gain g for the average cavity voltage $\langle V\rangle=V_0\cdot g/(1+g)$ by x, letting also $\langle V\rangle$ converge to zero. To prevent this, the 'Smoother' always has to preserve the average; this is done by shifting $\langle s_p\rangle$ to zero at each step, i.e. replacing s_p as expressed above by $s_p - \langle s_p\rangle$.

After n such steps transients will correspond to a gain of $g=x^n\cdot g_0$, finally converging to zero, while g_0 is recovered after each iteration step, $\langle V\rangle$ being conserved.

Alternative hardware options

There are two hardware alternatives, possibly handier for certain designs. First, as sketched in Fig. 1, bottom, instead of in-recording r the out-recording r' might be used. Then s_p has to be defined as $s_p = s_a+r'\cdot(1/x -1)$ with, as above, subtraction of its average to preserve $\langle V\rangle$.

[†] as small as desired by correspondingly reducing the adaptation speed

[‡] by far remaining within the loop gain range assuring a stable beam

[§] strictly true only for a perfectly linear chain (no important difference)

[**] at the same time avoiding problems of precision and simultaneity

Second, the ‘Smoother’ might be installed at (β) instead of (α). To prove this, the same chain of arguments as above for (α) has to be followed. The set-functions at (β) and (α) are probably different from each other.

MULTIPLE DIFFERENT CAVITIES

Till now the RF system was treated like a single cavity. In reality different RF systems may work in parallel, as the 400 MHz and future 200 MHz systems in LHC, but even ‘identical’ cavities operate at slightly different parameters due to calibration uncertainties (V_{acc}), setting differences ($\Delta\omega$, Q_{ext}) and manufacturing scatter (R/Q). Due to these differences and the high gain, each cavity needs its own set-function adapted to its precise parameters.

For any multiple-cavity system, the first ‘Gedanken-Experiment’ can be repeated in operating *all* cavities with open loop while bunch charges are scaled up again adiabatically. Once this is done, all cavity loops can be closed as described above, such that each cavity is controlled again by its *individual* high gain feedback system while receiving constant RF power. This argument proves that also here there is a (unique) set of set-functions for the given boundary conditions.

This set can be iterated similarly to the unique cavity case. When iterating cavity set-functions one-by-one, the beam cannot directly approach its final position, as defined above. Therefore it is much more efficient to iterate all cavities in parallel. Since all changes are executed adiabatically, this parallelism does not mean that all operations have to be *perfectly* synchronized nor use the same reduction factor x.

COMPUTER SIMULATIONS

The above algorithm with different refinements was incorporated into the program CYCLOPS [3] and simulations were done. As example, Fig. 2a shows the initial state before, Fig. 2b the final state after adaptation of a set-function, the success of the adaptation is apparent. The bunch energy deviation (red bands) stays close to zero, as it should be in equilibrium, and the equilibrium time/phase (black bands) have found their new equilibrium.

A more detailed theoretical analysis with enlarged scope, refinements to the (simulated) execution and further simulations and phase space representation of bunches can be found in [4]

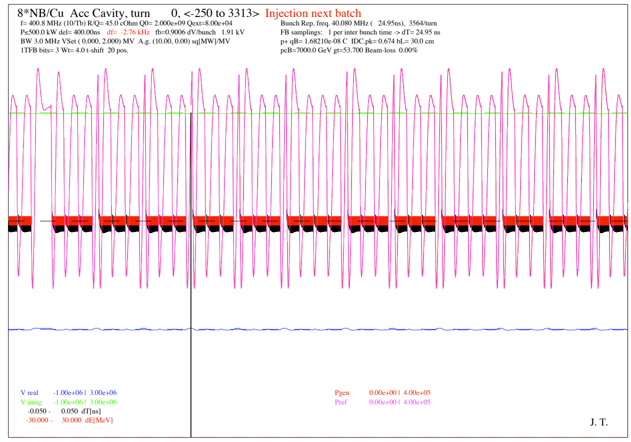

Fig 2a: One revolution in LHC (89 μs) before adaptation: large transients. Red (incident, partly covered by pink reflected trace), pink (reflected) RF power (scale 0-400 kW); green (I), blue (Q) accelerating voltage components (-1 to +3 MV); red bands: bunch energy deviation (scale ±30 MeV); black bands: bunch equilibrium phase (expressed in time, scale ±50 ps). The beam dump gap (close to the left end, no bands) and the kicker-gaps are clearly visible.

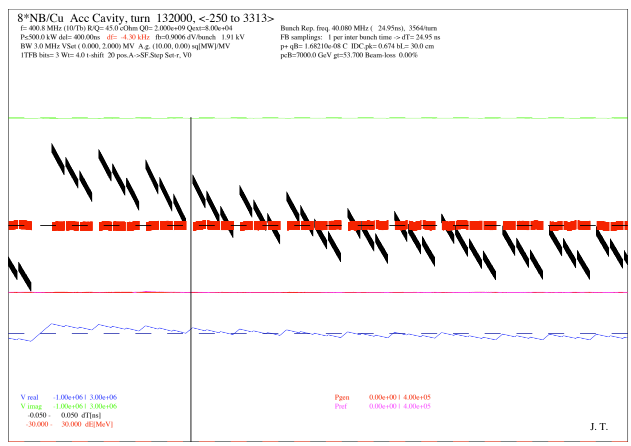

Fig 2b: As Fig. 2a but after adaptation for 132000 turns (about 11 s): nearly constant RF power. Bunch positions now shifted as already shown in [2], energy deviation from equilibrium ±0.4 MeV. Cavity tuning drifts from half-detuning at –2.76 kHz to –4.32 kHz.

REFERENCES

[1] W. Ross et al., Gap Voltage Feed-Forward Module for PEP-II Low Level RF System, SLAC-Pub-7478 and Proc. PAC97, 12-16 May 1997, Vancouver, BC, Canada.

[2] J. Tückmantel, The LHC Beam with Suppressed Transients, CERN-AB-Note-2004-022. http://documents.cern.ch/cgi-bin/setlink?base=internal&categ=ab&id=ab-note-2004-022

[3] J. Tückmantel, Realistic RF system and Beam Simulation in Real Time for a Synchrotron, Proc. HEACC 2001, Tsukuba, Japan, 26-30 March 2001. http://conference.kek.jp/heacc2001/pdf/p2hc07.pdf (The 2005 program CYCLOPS is a considerably improved version, e.g. allowing full phase-space bunch representation)

[4] J. Tückmantel, Theory and Simulations for Adaptive RF Transient Reduction for High Intensity Beams with Gaps, to be published.