RF HEATING FROM WAKE LOSSES IN DIAGNOSTICS STRUCTURES

E. Métral, CERN, Geneva, Switzerland

Abstract

Heating of diagnostics structures (striplines, buttons, screen vessels, wire scanners, etc.) has been observed at many facilities with higher stored currents. Simulations of wake losses using 3D EM codes are regularly used to estimate the amount of power lost from the bunched beam but on its own this does not tell how much is radiated back into the beam pipe or transmitted into external ports and how much is actually being dissipated in the structure and where. This talk should introduce into the matter, summarise some of the observations at various facilities and illustrate what approaches of detailed simulations have been taken.

INTRODUCTION

As beam intensity increases, the beam can no longer be considered as a collection of non-interacting single particles: in addition to the "single-particle phenomena", "collective effects" become significant [1]. At low intensity a beam of charged particles moves around an accelerator under the Lorentz force produced by the "external" electromagnetic fields (from the guiding and focusing magnets, RF cavities, etc.). However, the charged particles also interact with their environment, inducing charges and currents in the surrounding structures, which create electromagnetic fields called wake fields (see Fig. 1). In the ultra-relativistic limit, causality dictates that there can be no electromagnetic field in front of the beam, which explains the term "wake". It is often useful to examine the frequency content of the wake field (a time domain quantity) by performing a Fourier transform on it. This leads to the concept of impedance (a frequency domain quantity), which represents, for the plane under consideration (longitudinal, horizontal or vertical), the force, integrated over the length of an element, from a "source" to a "test" particle, normalized by their charges. Impedances are complex functions of frequency and in general, the impedance in a given plane is a nonlinear function of the test and source transverse coordinates, but it is most of the time sufficient to consider only the linear terms.

The wake fields (or impedances) can influence the motion of trailing particles, in the longitudinal and in one or both transverse directions, leading to energy loss, beam instabilities, or producing undesirable secondary effects such as excessive heating of sensitive components at or near the chamber wall (called beam-induced RF heating). Therefore, in practice the elements of the vacuum chamber should be designed to minimise the self-generated (secondary) electromagnetic fields. For example, chambers with different cross-sections should be connected with tapered transitions; non necessary cavities should be avoided; bellows should preferably be separated from the beam by shielding; plates should be grounded or terminated to avoid reflections; poorly conductive materials should be coated with a thin layer of very good conductor (such as copper) when possible, etc. However, the issue with the diagnostics structures is that they are designed to couple to the beam!

In the case of the beam-induced RF heating of interest in this paper, it comes from the real part of the longitudinal impedance and the bunch length (and sometimes longitudinal profile) is the main parameter (once the bunch intensity and number of bunches have been fixed): usually, the longer the bunch, the better. Beam-induced RF heating has been observed in many places, as for instance recently in several CERN LHC components during the 2011 and 2012 runs when the bunch/beam intensity was increased and/or the bunch length reduced [2]. This caused beam dumps and delays for beam operation (and thus less integrated luminosity) as well as considerable damages for some equipment. This is why the rms bunch length was increased to ~ 9 cm in 2011 and ~ 10 cm in 2012, whereas the nominal value is 7.5 cm. The RF heating of some equipment is therefore worrisome for the future operation and it is closely followed up [3].

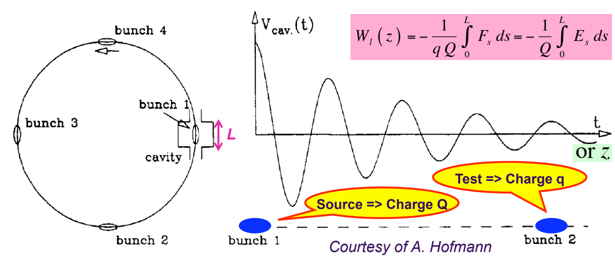

Figure 1: Illustration of the voltage induced by a bunch (bunch 1 in this example) going through a cavity. The bunch 2, at the distance z behind, will feel a longitudinal wake field given by the formula (more precisely, it is called wake function and it represents the response to a pulse of infinitely small length).

Section 1 reviews the main results related to beam-induced RF heating, while Section 2 is devoted to the summary and highlights of a one-day mini-workshop on "Simulation of Power Dissipation and Heating from Wake Losses in Accelerator Structures", which took place on 30/01/2013 at the Diamond Light Source (DLS) [4].

BEAM-INDUCED RF HEATING

Consider the case of M equi-spaced equi-populated bunches, which is a good approximation when the machine is almost full. In this case, the general formula

ISBN 978-3-95450-127-4

for the beam power loss (due to the interaction with the longitudinal impedance) can be written [5]

$$P_{loss} = M\, I_b^2\, Z_{loss}\,, \tag{1}$$

with

$$Z_{loss} = 2\, M \sum_{p=0}^{\infty} \mathrm{Re}\left[Z_l \left(p\, M\, \omega_0 \right) \right] \times \mathrm{PowerSpectrum}\left[p\, M\, \omega_0 \right], \tag{2}$$

where $I_b = N_b\, e\, f_0$ is the bunch current (with N_b the number of particles per bunch, e the elementary charge and f_0 the revolution frequency), $\omega_0 = 2\,\pi\, f_0$, Z_l the longitudinal impedance, and PowerSpectrum stands for the beam power spectrum.

The power loss is always proportional to the square of the number of particles per bunch but depending on the shape of the impedance, it can be linear with the number of bunches (when the bunches are independent, i.e. for a sufficiently short-range wake field – or broad-band impedance – which does not couple the consecutive bunches) or proportional to the square of the number of bunches (when the bunches are not independent, i.e. for a sufficiently long-range wake field – or narrow-band impedance – which couples the consecutive bunches). These two extreme cases are discussed in detail below.

As concerns the beam (power) spectrum, examples of measurements performed in 2011 in the LHC before the ramp and in stable beams, which reveal interesting features, are shown in Fig. 2 (they correspond to a 4σ bunch length of ~ 1.2 ns). First, many peaks are spaced by ~ 20 MHz (as it is expected for the 50 ns bunch spacing beam used; the bunch frequency would be ~ 40 MHz for the nominal 25 ns beam) below an envelope which is decreasing with frequency until a certain value and which is then revealing a side lobe (and sometimes also others). This behaviour is exactly the one expected due to the finite length of the bunch (inside a finite bucket). To get a better feeling, let's consider four typical (theoretical) distributions, whose longitudinal profiles are represented in Fig. 3 [6]. The two extreme cases are, on one side the Gaussian distribution with infinite and smooth tails (therefore unrealistic) and on the other side the Water-Bag distribution with finite and sharp tails. The corresponding power spectra can be computed analytically and they are depicted in Fig. 4. It is clearly seen that only the (unrealistic) Gaussian distribution does not reveal side lobes due to the fact that the tails extend up to infinity. For all the other distributions (with finite lengths), sides lobes are revealed and the sharper the tails the higher the sides lobes. However, in the measurements of Fig. 2 the height of the first side lobe is at ~ - 35 or - 40 dB, whereas the theoretical distributions considered give higher values. This means that the real distribution must have smoother tails. Consider now a family of (finite) distributions, keeping the same half width at half height, depending on the parameter n (converging to a Gaussian distribution when n goes to infinity, see Fig. 5, similarly to what was done in Ref. [7]). The corresponding power spectra are depicted in Fig. 6. It is seen that the distribution with n = 3 should be a relatively good approximation (even if two side lobes are expected from theory in this case whereas only a large one was

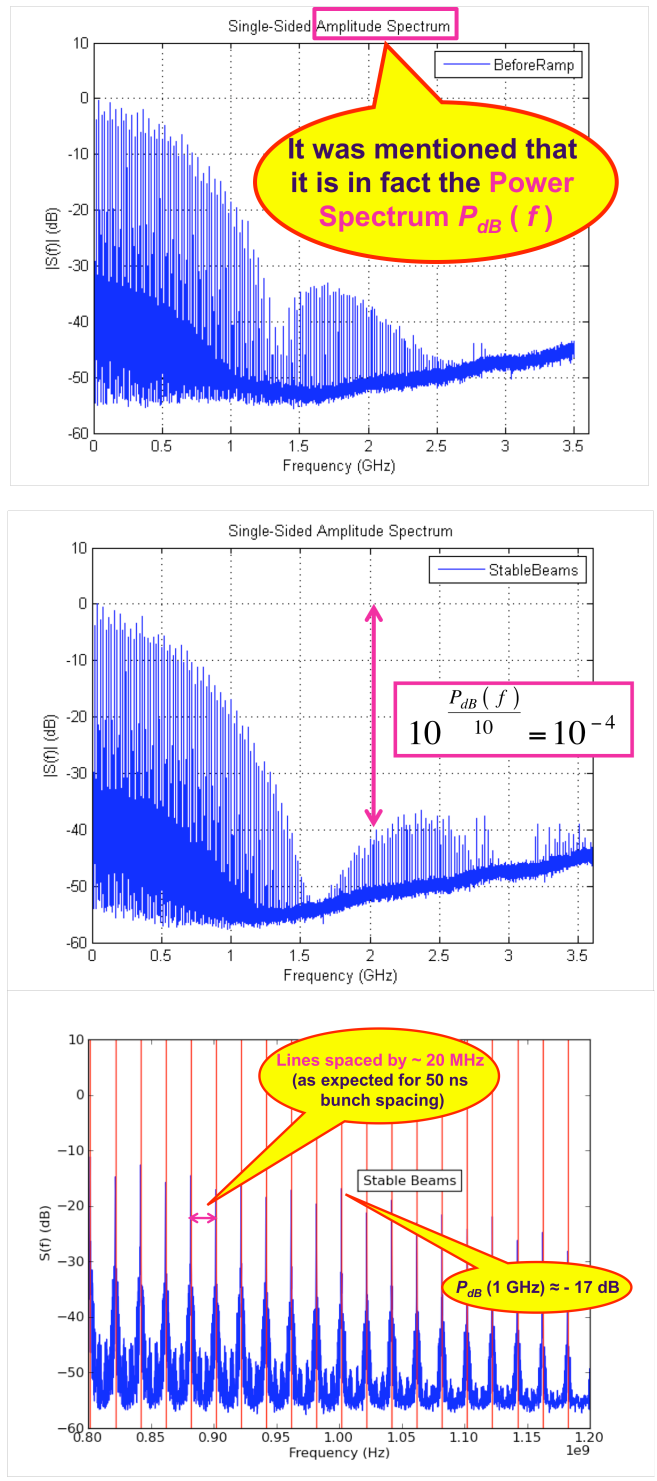

Figure 2: Power spectra measurements for LHC beam 1 on fill # 2261 (Courtesy of Themistoklis Mastoridis, Philippe Baudrenghien and Hugo Day).

ISBN 978-3-95450-127-4

measured). By taking the inverse Fourier Transform of the measured spectrum, the longitudinal profile of Fig. 7 has been obtained, which is consistent with the expected one from theory (with $n = 3$; see Fig. 5).

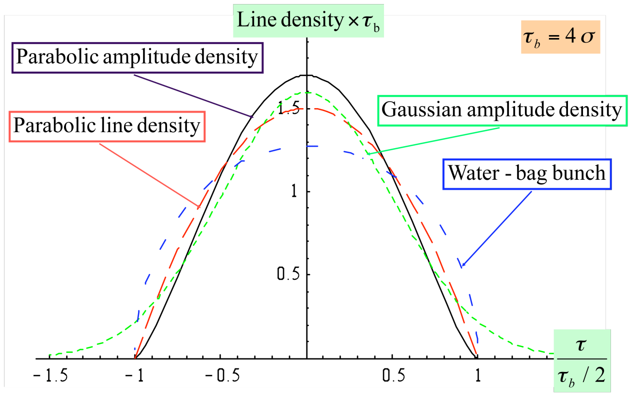

Figure 3: Longitudinal profiles of four typical (theoretical) distributions.

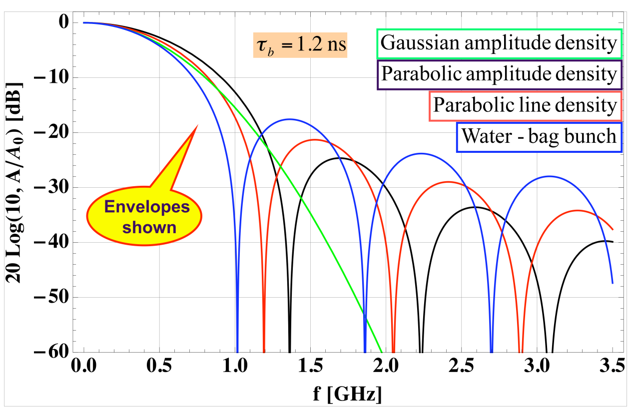

Figure 4: Power spectra corresponding to Fig. 3 (considering a full or 4σ bunch length of 1.2 ns).

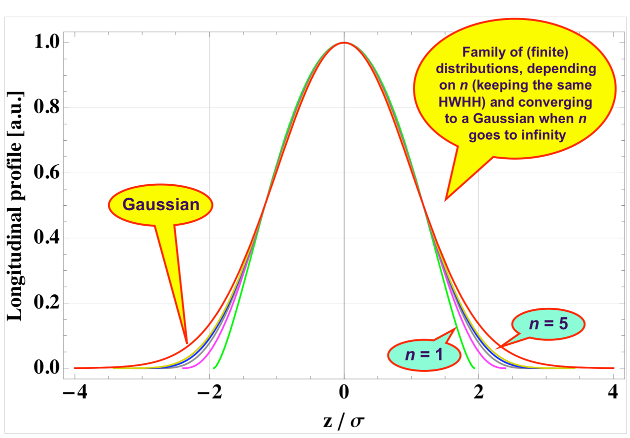

Figure 5: Family of (finite) distributions, keeping the same Half Width at Half Height (HWHH), depending on the parameter n, and converging to a Gaussian distribution when n goes to infinity.

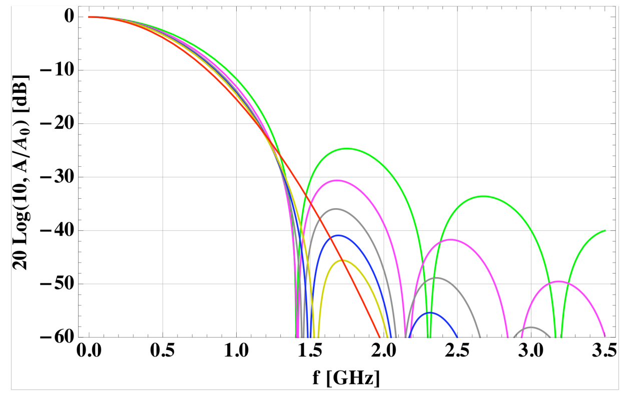

Figure 6: Power spectra corresponding to Fig. 5.

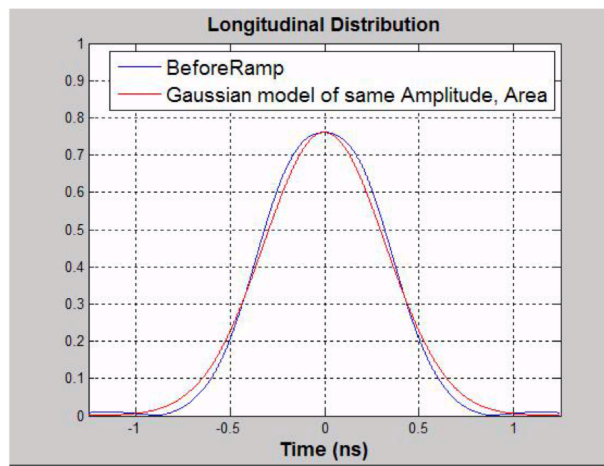

Figure 7: Longitudinal profile obtained by taking the inverse Fourier Transform of the measured spectrum before the ramp (see Fig. 2a; Courtesy of Themistoklis Mastoridis and Philippe Baudrenghien).

In the case of the DLS, for instance, the bunch length is much smaller than in the LHC and therefore the power spectrum extends to much high frequencies, as shown in Fig. 8.

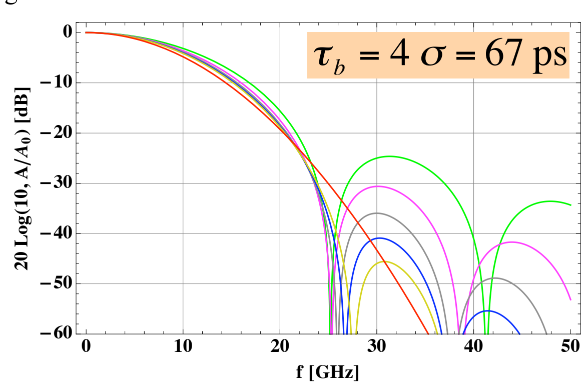

Figure 8: Power spectrum for a DLS bunch (5 mm rms, i.e. 67 ps at 4σ) assuming the bunch profiles of Fig. 5. To be compared to an LHC bunch (see Fig. 6).

ISBN 978-3-95450-127-4

In the case of a broad-band impedance, consider for instance the case of the resistive-wall impedance, and, as a numerical example, the particular case of the LHC beam screen (neglecting the holes, whose contribution has been estimated to be small in the past, and the longitudinal weld). Assuming a Gaussian longitudinal profile (other similar distributions would give more or less the same result in this case), the power loss (per unit of length) is given by

$$P_{loss/m} = \frac{1}{C}\,\Gamma\left(\frac{3}{4}\right)\frac{M}{b}\left(\frac{N_b\,e}{2\,\pi}\right)^2\sqrt{\frac{c\,\rho\,Z_0}{2}}\;\sigma_t^{-3/2}\,, \tag{3}$$

where C = 26658.883 m is the average LHC radius, Γ the Euler gamma function, b the beam screen half height (assumed to be 18.4 mm), c the speed of light, ρ the resistivity (assumed to be 7.7 10^{-10} Ωm for copper at 20 K and 7 TeV), Z_0 the free-space impedance and σ_t the rms bunch length (expressed in unit of time). Assuming the nominal LHC beam parameters (M = 2808, N_b = 1.15 10^{11} p/b and σ_t = 0.25 ns), Eq. (3) yields ~ 101 mW/m.

Consider now the case of a narrow resonance, describing a trapped mode due to the geometry. It is described by 3 parameters: (i) the resonance frequency, assumed to be here f_r = 1 GHz; (ii) a shunt impedance, assumed to be here R_l = 10 Ω; and (iii) a quality factor Q, whose value is scanned below. The impedance plots are represented in Fig. 9 together with the corresponding wake functions. It can be seen on this example that if the quality factor is bigger than ~ 100, only one line can be considered (the bunches are coupled and this is the total current which matters) whereas if the quality factor is smaller than ~ 20, then the bunches are not coupled. Indeed, in the case of a sharp resonance impedance (i.e. when $Q >> f_r\,/\,(\,2\,f_b\,)$ where f_b is the bunch frequency), the power loss is given by the simple formula (which is valid when $Q >> 1$ and $\Delta << 1$)

$$P_{loss} = R\,I^2 \times F \times G \tag{4}$$

with

$$F = 10^{\frac{P_{dB}(f_r)}{10}}\;,\quad G = \frac{\Delta^2}{\Delta^2 + \sin^2\left(\dfrac{\pi\,f_r}{f_b}\right)}\;,\quad \Delta = \frac{\pi\,f_r}{2\,Q\,f_b}\,, \tag{5}$$

where R = 2 R_l, i.e. using the Linac convention (Linac Ohms), $I = M\,I_b$ is the total beam current and $P_{dB}\,(\,f_r\,)$ is the beam power spectrum in dB at the resonance frequency f_r read from a power spectrum (computed or measured). The factor F describes the frequency dependence of the power loss, which depends on (i) the longitudinal bunch length, (ii) the longitudinal profile and (iii) the resonance frequency. It converges to 1 at zero (low) frequency (where it is the worst case) and it is

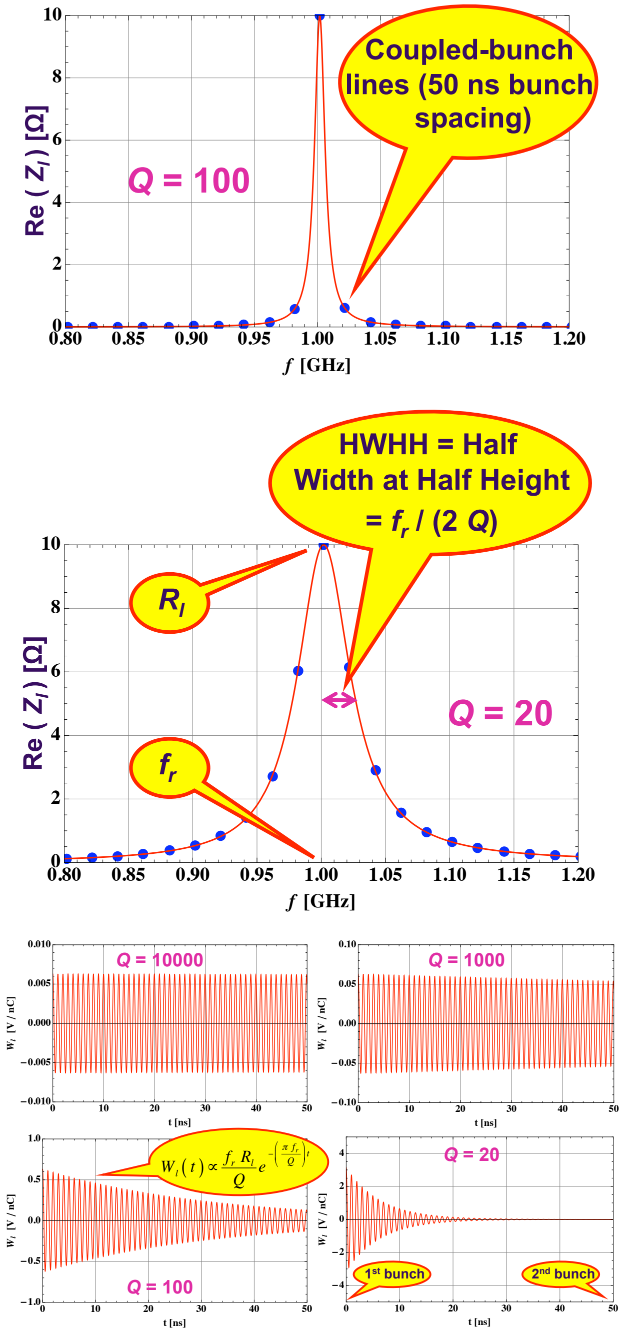

Figure 9: Impedance plots for a resonance with the resonance frequence f_r = 1 GHz, shunt impedance R_l = 10 Ω, for different values of quality factors Q and corresponding wake functions.

between 0 and 1 for any frequency. For a Gaussian bunch, the factor F is given by Exp[- ($2\,\pi\,f_r\,\sigma_t$) ^{2}]. The factor G describes the off-resonance effect [5]: if the resonance falls exactly on an harmonic of the bunch frequency (i.e. on resonance), it is equal to 1, otherwise it is between 0 and 1 (as will be seen in Fig. 11).

Assuming a total beam current of 1 A (the nominal LHC value is 0.58 A) and considering the theoretical longitudinal bunch spectrum of Fig. 10 (left) for an rms bunch length of 9 cm (similar to the LHC case in 2011), a

ISBN 978-3-95450-127-4

sharp resonance R_l = 5 kΩ (usual typical values are between few hundreds and few tens of thousands Ohms) at 1.4 GHz (i.e. on resonance) would therefore generate a power loss of 1 W. However, this result is very sensitive to the bunch length. It can be seen for instance from Fig. 10 (right), that dividing the bunch length by 2, i.e. going from 9 cm rms to 4.5 cm, would increase the power loss by a factor ~ 2000, i.e. going from 1 W to 2 kW! Therefore, any (major) bunch length reduction should be considered with great care. For completeness, the off-resonance effect was also studied in Fig. 11, where it can be seen that the power loss rapidly decreases with the frequency offset for high-Q resonances, which could be a useful knob. The problem is that in practice this offset is usually not known with sufficient precision.

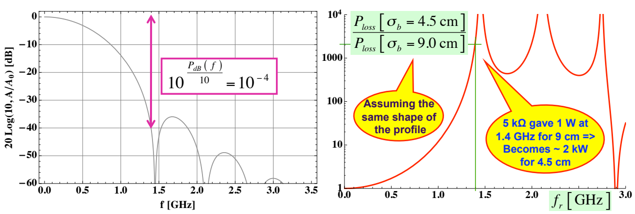

Figure 10: Theoretical (see Fig. 6 with n = 3) longitudinal bunch spectrum (left) for the case of a LHC bunch in 2011 (9 cm rms bunch length) and power loss increase for the case of a bunch two times shorter (4.5 cm rms) assuming the same shape.

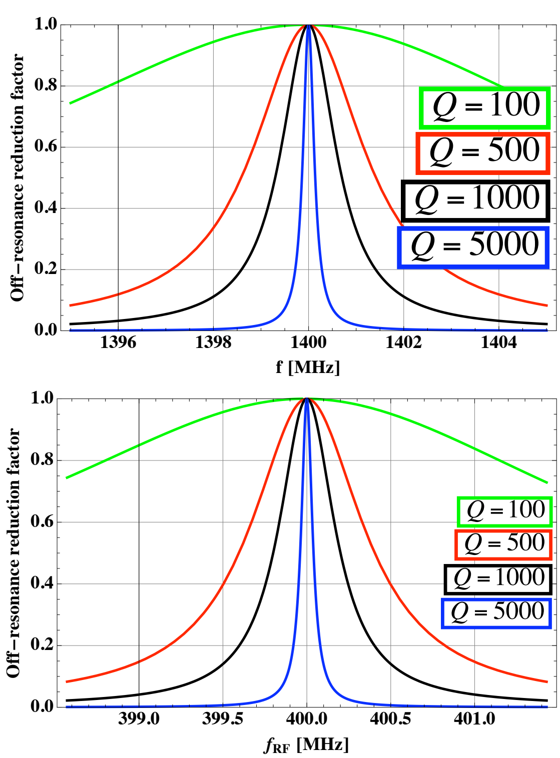

Figure 11: Off-resonance effect on the power loss, applied to the case of Fig. 10, vs. (a) the resonance frequency and (b) the RF frequency (assuming the resonance frequency at 1.4 GHz).

The usual solutions to avoid beam-induced RF heating are the following, depending on the situation:

i) Increase the distance between the beam and the equipment.

ii) Coat with a good conductor if the heating is predominantly due to resistive losses and not geometric losses.

iii) Close large volumes (which could lead to resonances at low frequency) and add a smooth transition. This is why beam screens and RF fingers are installed.

iv) Put some ferrite with high Curie temperature and good vacuum properties (close to the maximum of the magnetic field of the mode and not seen directly by the beam) or other damping materials. Adding a material with losses (the type of ferrite should be optimized depending on the mode frequency), the width of the resonance will increase (the impedance will become broader) and the (maximum) impedance will decrease by the same amount. The power loss will therefore be (much) smaller. However, the ferrite will then have to absorb the remaining power. Even if much smaller, the heating of the ferrite can still be a problem if the temperature reached is above the Curie point, or is above the maximum temperature allowed by the device. To cool the ferrite one should try and improve the thermal conduction from the ferrite as most of the time only radiation is used (given the general brittleness of the ferrite it is difficult to apply a big contact force).

v) Improve the subsequent heat transfer:

- Convection: there is none in vacuum.
- Radiation: usually, the temperature is already quite high for the radiation to be efficient. One should therefore try and improve the emissivities of surrounding materials.
- Conduction: good contacts and thermal conductivity are needed.
- Active cooling: the LHC strategy (for instance) was to water cool all the near beam equipment.

vi) Try and design an All Modes Damper (AMD) if possible, to remove the heat as much as possible to an external load outside vacuum, where it can be more easily cooled away. This can also work together with a damping ferrite.

vii) Increase the bunch length. The longitudinal distribution can also play a very important role for some devices, and it should be kept under tight control.

viii) Install temperature monitoring on critical devices to avoid possible damages.

Following some issues with RF fingers on some LHC equipment in 2011 (as observed also before in other machines), a task force was set up during 2012 to review the design of all the components of the LHC equipped with RF fingers. The lessons learnt and the mitigation measures for the CERN LHC equipment with RF Fingers were reported in Ref. [8]. It is worth mentioning that for all the cases studied, no problem with impedance was revealed for conforming RF fingers. But the top priority for the future should be to try and reach robust mechanical designs to keep the contacts of all the RF fingers and to do a very careful installation (92 non-

ISBN 978-3-95450-127-4

conformities were revealed in 2012 after an X-ray campaign).

HIGHLIGHTS FROM A MINI-WORKSHOP AT DLS ON 30/01/2013 [4]

A one-day mini-workshop on "Simulation of Power Dissipation and Heating from Wake Losses in Accelerator Structures" took place on 30/01/2013 at the DLS [4]. It was organized by G. Rehm and the motivation/worry was that diagnostics systems are designed to couple to the beam, which can lead to large amounts of energy being lost from the beam. With the currents settings, a power loss of 189 W was estimated for the striplines. With the plan settings (to go to higher currents and shorter bunch lengths) this power loss would increase to 313W. With these huge amounts of power lost by the beam, it is important to study in detail how much of the power removed from the beam is radiated back into the beam pipe or transmitted into external ports (where present) and how much is actually being dissipated in the structure, and where. The final question to answer being: what is the impact of the dissipated power in terms of deformation, stresses or potential damage?

Twenty-three people attended this workshop where eight talks were given. Several machines were discussed, whose relevant parameters are summarized in Table 1, where the ratio between the incoherent power loss (i.e. neglecting the coherent effects between the bunches) and the loss factor is also given, with

$$k_{loss} = \int ds \, W_l(s) \, \lambda(s) \tag{6}$$

$$\frac{P_{loss}^{Incoh}[\mathrm{W}]}{k_{loss}[\mathrm{V/pC}]} = M \, Q[\mathrm{nC}]^2 \, f_0[\mathrm{kHz}] \, 10^{-3}, \tag{7}$$

where $W_l(s)$ is the monopole longitudinal wake potential, $\lambda(s)$ the normalized bunch charge density and Q the bunch charge. A comparison between the different longitudinal beam power spectra is shown in Fig. 12, revealing the frequency ranges of interest for the different machines.

Table 1: Main parameters of the different machines discussed during the DLS workshop [4] and conversion factor between the incoherent power loss (i.e. neglecting the coherent effects between the bunches) and the loss factor (see Eqs. (6) and (7)).

	M	$Q = N_b e$ [nC]	f_0 [kHz]	I_{beam} [A]	W / (V / pC)	σ_z [mm]
ALBA	448	0.8	1118.6	0.4	319	4.6
SOLEIL	416	1.3	844.5	0.44	551	6
DLS	900	1.0	533.8	0.5	520	4
NSLS	1080	1.2	378.8	0.5	611	4.5
PETRA-III	40	19.2	130.1	0.1	1921	13
LHC	2808	18.4	11.2	0.58	10691	75.5
PEP-II	1700	12.9	136.3	3	38838	8

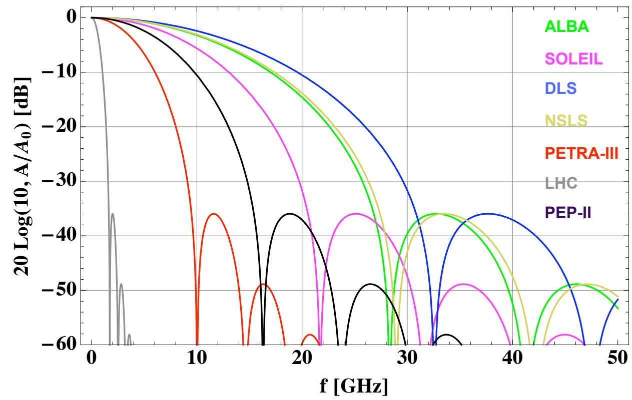

Figure 12: Comparison between the different longitudinal beam power spectra, revealing the frequency ranges of interest for the different machines.

A dedicated experiment at DLS, called COLDDIAG (= COLD vacuum chamber for DIAGnostics), was also discussed, whose aims are threefold: (i) measure the beam heat load on a cold bore simulating the liner of superconducting Insertion Devices (IDs) with different operating conditions; (ii) gain a deeper understanding in the beam heat load mechanisms and (iii) study the influence of the cryosorbed gas layer on the beam heat load. The measurements and their analysis is work in progress and will continue both on the experimental and theoretical sides.

The methods applied by the different presenters and/or their colleagues were discussed (time-domain wake field simulations to get the loss factor and time domain long-range wake field simulations; eigenmode simulations to identify critical modes; identification of dangerous regions from the different modes; thermal simulations using all power distributed according to the mode to get the temperature distribution; damping of some critical modes; etc.) and some useful analytical computations were reviewed: (i) power loss (incoherent and coherent); (ii) quality factor and loss factor for the different parts of a device; (iii) loss factor for an off-resonance line; (iv) loss factor for single-bunch and multi-bunch operations; (v) analytical impedance expressions of step transitions, tapers, surface roughness, resistive wall effect. It was reminded that the beam-induced RF heating of a machine can sometimes impose more stringent requirements on the vacuum chamber structures than those from beam instabilities and several examples of heating and damages were presented: melted materials (RF fingers), BPM buttons falling down, mirror and support of a synchrotron light monitor damaged, high detector background, wake fields outside the beam chamber, operation delays, beam dumps, etc.

Several codes were used to estimate the power loss (and sometimes the repartition and subsequent temperature increase), such as (i) GdFidL (1^{st} 2 talks and 8^{th} one); (ii) CST (3^{rd}, 4^{th}, 6^{th}, 7^{th} and 8^{th} talks), with also some thermal studies (input power and field distribution to get temperature distribution); (iii) MAFIA, NOVO and

ISBN 978-3-95450-127-4

Omega3P (5th talk); (iv) ECHO (8th talk) and (v) ANSYS for the temperature and stress distribution.

It was already noticed during the PEP-II time that large amounts of energy can travel through the beam pipe, but it was not quantified. At that time, they had to found rapid fixes: some designs were changed to avoid sources of wake losses where possible, ceramic tiles were added to localise the losses and water-cooling was also added to almost everything.

A. Morgan discussed the approach currently used at the DLS (see Fig. 13) to try and disentangle the different power loss contributions [9]. Applying this approach to their striplines and BPMs with the current settings, satisfactory results were obtained. Indeed, subsequent thermal simulations were found to be in relatively good agreement with real world data of temperatures.

Finally, a homework was also proposed by the workshop organisers (before the workshop) on a simplified version of their stripline, with a single bunch of 1 nC and an rms bunch length of $\sigma = 5$ mm [9]. In this case, the loss factor is given by $k_{loss} = 858$ mV/pC and the power loss repartition is the following: (i) 11% down the beam pipe; (ii) 84 % into the signal ports and (iii) 5% left in the structure. The losses in the structure are about a quarter in the striplines and about three quarters in the vessel.

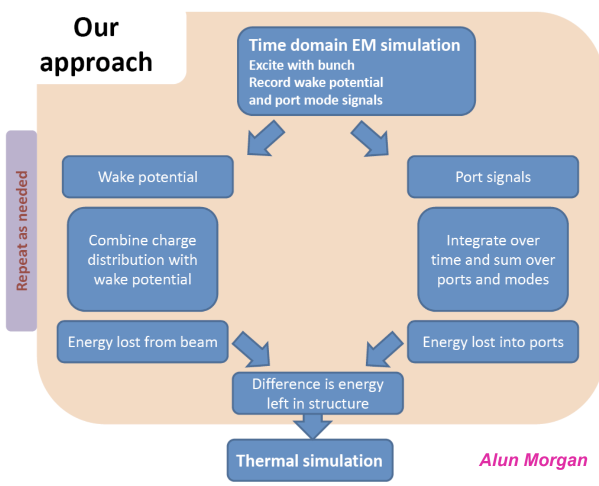

Figure 13: Approach being currently applied at the DLS to study the beam-induced RF heating [9].

CONCLUSIONS

Beam-induced RF heating is a very important mechanism, which can limit the performance of an accelerator and which can impose sometimes more stringent requirements on the vacuum chamber structures than those from beam instabilities. Several examples of heating, damages, beam dumps or delays have been observed in many machines, and therefore it should be treated with great care.

One of the main point raised at the recent mini-workshop which took place at the beginning of the year at DLS [4], is that for several structures, such as striplines, a large fraction of the power is sent down the beam pipe. This will/could act as an additional heat load on nearby structures. The fact that large amounts of energy can travel through the beam pipe was already discussed in the past but it was never quantified. How can this be correctly taken into account to accurately localize all the power losses? Do we need to simulate also the adjacent structures? Integrated calculation of dissipated power distribution from wake losses is not available at the moment already for a short-range wake field (i.e. single bunch) and the most important and critical case is a long-range wake field with a train of bunches to simulate the coherent case. Some approximated methods are used at the moment but the post-processing analysis is quite complicated and time consuming, and one often needs many simulation runs. None of the demonstrated methods discussed during the DLS workshop [4] is fully consistent, and these methods can be used only as a first step. More discussions should therefore take place with the code developers to tackle these different challenges.

ACKNOWLEDGEMENTS

Many thanks to all the people working at CERN on the important issue of beam-induced RF heating and to the organisers and participants to the mini-workshop at the DLS at the beginning of the year [4], for the very fruitful discussions.

REFERENCES

[1] A.W. Chao, Physics of Collective Beam Instabilities in High Energy Accelerators, New York: Wiley, 371 p, 1993.
[2] E. Métral et al., “Beam-Induced Heating / Bunch Length / RF and Lessons for 2012”, Proc. of the LHC Performance workshop, Chamonix, France, 06-10/02/2012.
[3] B. Salvant et al., “Beam Induced RF Heating”, Proc. of the LHC Beam Operation workshop, Evian, France, 17-20/12/2012.
[4] http://www.diamond.ac.uk/Home/Events/Past_events/Simulation-of-Power-Dissipation---Heating-from-Wake-Losses.html.
[5] M. Furman et al., “Energy Loss of Bunched Beams in SSC RF Cavities”, SSC-116, March 1987.
[6] J.L. Laclare, “Bunched Beam Coherent Instabilities”, CERN Accelerator School, Oxford, CERN 87-03, p. 264, 1987.
[7] E. Métral and A. Verdier, “Stability Diagrams for Landau Damping with a Beam Collimated at an Arbitrary Number of Sigmas”, CERN-AB-2004-019-ABP.
[8] E. Métral et al., “Lessons Learnt and Mitigation Measures for the CERN LHC Equipment with RF Fingers”, Proc. of the IPAC’13 conference, Shanghai, China, 12-17 May, 2013.
[9] A. Morgan, “Analysis of time domain wake potential and port signals for calculation of radiated and dissipated power due to wake losses”, presented at the DLS workshop [4].

ISBN 978-3-95450-127-4