

Accepted Manuscript

Bunched-beam Schottky monitoring in the LHC

M. Betz, O.R. Jones, T. Lefevre, M. Wendt

PII: S0168-9002(17)30945-2
DOI: <http://dx.doi.org/10.1016/j.nima.2017.08.045>
Reference: NIMA 60064

To appear in: *Nuclear Inst. and Methods in Physics Research, A*

Received date: 23 May 2017
Revised date: 21 August 2017
Accepted date: 25 August 2017

Please cite this article as: M. Betz, O.R. Jones, T. Lefevre, M. Wendt, Bunched-beam Schottky monitoring in the LHC, *Nuclear Inst. and Methods in Physics Research, A* (2017), <http://dx.doi.org/10.1016/j.nima.2017.08.045>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



Bunched-beam Schottky monitoring in the LHC

M. Betz^a, O.R. Jones^b, T. Lefevre^b, M. Wendt^b

^a*Lawrence Berkeley National Laboratory, Berkeley, California, USA*

^b*CERN, Geneva, Switzerland*

Abstract

The LHC Schottky beam diagnostic system has the potential to measure beam parameters such as tune, chromaticity and momentum spread in a non-invasive way. It does so by detecting beam induced signals at 4.8 GHz, and extracting the characteristic modulation caused by incoherent motion of particles within a bunch. The system was commissioned in 2011 and provided satisfactory beam-parameter measurements during LHC Run I for lead-ions. However, for protons its usability was substantially limited due to strong interfering common-mode signals. The system was upgraded with optimized traveling-wave pick-ups and an improved microwave signal path. Design and operational aspects for the complete system are shown, and the results from measurements with LHC beams in Run II are presented and discussed.

Keywords: elsarticle.cls, L^AT_EX, Elsevier, template
2010 MSC: 00-01, 99-00

1. Introduction

'Shot noise' was discovered in 1918 by W. Schottky during his experiments with vacuum tube audio amplifiers. While studying fluctuations in the tubes anode current, he discovered that with enough gain it was possible to distinguish the impulses from individual electrons arriving at the anode. He named this effect "shot noise" after the crackling sound he heard in his headphones. He later showed that the effect originates from the statistical fluctuation of current due to being made up of individual packets of finite charge [1].

In 1971 the same phenomenon was observed with proton beams in the Intersecting Storage Ring (ISR) at CERN [2]. It was discovered that the complex particle motion imprints valuable information about the state of beam and machine into the noise. This makes "Schottky signals" an important non-invasive diagnostic tool, in particular for storage rings.

Additionally, Schottky signals are also the foundation for stochastic cooling, a concept first published by Simon van der Meer in 1972 [3]. Schottky diagnostics and stochastic cooling techniques

became essential for every-day operation of accumulator rings, allowing many injections of long, low intensity bunches to be collected into a dense 'stack' of particles. Good examples are the \bar{p} accumulation schemes at CERN (ISR, SPS, AA, AC, AD, ELENA), Fermilab (Accumulator and Debuncher for $p\bar{p}$ collisions in the Tevatron), and BNL (RHIC).

Yet, using Schottky signals for the monitoring of bunched proton or antiproton beams with short, typically non-Gaussian particle distributions has turned out to be very challenging. A strong unwanted common mode signal content unexpectedly persists even at frequencies well above the first order high frequency cut-off of the beam spectrum. This typically leads to saturated amplifiers in the RF electronics and distortion problems. That is why attempts at bunched beam stochastic cooling have failed, both, in the SPS $p\bar{p}$ program at CERN (1990), and also in the Tevatron at Fermilab [2]. A successful cooling of bunched gold ions has been achieved at RHIC (BNL) in 2000 [4]. RHIC was profiting from the fact that the Schottky signal power for ion beams scales quadratically with the ionic charge number z . Therefore, gold ions with $z = 79$ provide almost 40 dB more signal power than a beam consisting of the same number of proton or antiprotons ($z = 1$).

Email address: mbetz@lbl.gov (M. Betz)

The Large Hadron Collider (LHC) at CERN has been equipped with a 4.8 GHz transverse Schottky system for beam diagnostics purposes, which was commissioned in 2011 as part of the U.S. LHC Accelerator Research Program (LARP). The LHC collides proton or lead-ion (Pb^{82+}) beams at a center-of-mass energy of up to 7+7 TeV.

The LHCs 400 MHz main RF frequency results in a 1 ns four-sigma bunch length and an intensity of $\approx 10^{11}$ protons [5]. These short, high intensity bunches cause strong coherent signal components in the Schottky spectrum, which can be up to 100 dB larger than those of the incoherent signals to be measured. The dynamic range requirements therefore pose a major challenge for bunched beam Schottky diagnostics in the LHC.

During LHC Run I (2010–2013) the practical usability of the LHC Schottky system was limited to lead ion beams, where the common-mode signals at the pick-up operational frequency are weaker (longer bunches), and the power in the incoherent signals gain with the square of the higher charge state ($z = 82$). A major overhaul of the Schottky beam pickups during the long shutdown 1 (LS1: 2013–2014), complemented by a redesign of the RF front-end signal processing finally enabled a reliable and reproducible detection of Schottky signals for bunched proton beams.

There are two important use-cases in which this system can support the everyday physics operation: One is the non-perturbative measurement of the machine chromaticity, the other is the measurement of the betatron tune. Chromaticity in the LHC is typically measured by modulating the RF frequency while observing the change in tune. However, for various reasons this technique is limited to low intensity beams, implying that chromaticity cannot be measured during physics production runs with high intensity beams [6]. The standard tune measurement system also suffers when the machine is filled up with high intensity bunches. While this system has a high sensitivity [7] its operational frequency overlaps with the transverse stability feedback system of the LHC, again, limiting the measurement capabilities for high intensity physics production beams. This limits the tune measurement to a small number of ‘probe bunches’, for which the transverse feedback is disabled. A reliable Schottky measurement system should overcome these issues while in addition providing a rich picture of the internal dynamics of the transverse beam motion of particles within a bunch, which should lead to a

better understanding of the machine.

This report starts with a high level overview of the LHC Schottky system, followed by a brief review of bunched beam Schottky theory (Section 2). A detailed look is then taken at the components which have been upgraded since 2013. This includes the overhaul of the beam pickups, the redesigned and optimized RF front-end, now equipped with new filters, and an analysis of the overall noise performance including the down-converter chain in the RF back-end. Several new algorithms, to process the acquired spectra and extract parameters such as tune and chromaticity are also discussed, along with their performance evaluated during dedicated machine development (MD) periods.

1.1. System overview

The LHC Schottky system consists of four identical channels, for the horizontal and vertical plane of each of the counter-rotating LHC beams. A schematic overview of the main components of one channel is shown in Fig. 1. The Schottky




Figure 1: High-level overview of the LHC Schottky pickup and signal processing components.

beam pickup (Section 3) consists of two symmetrically arranged, slot-coupled waveguides operating at 4.8 GHz, coupling to the beams coherent and incoherent signal contributions. Subtracting both output signals by means of a difference (Δ) hybrid reduces the large common mode (Σ) signal contributions, i.e. most of the unwanted coherent spectral components derived from the bunch intensity, enabling the detection of transverse low-level signals. A fast gate-switch (Section 4.2) allows to select one or more bunches and is followed by a narrow-band signal processing system, optimized for low-noise

and high dynamic range. In practice, the RF signal processing is divided in two parts. A 4.8 GHz ‘front-end’ gating and gain section (Section 4.3), which is mounted directly on top of the Schottky beam pickup, and a ‘back-end’ down-converter section (Section 5), located in a nearby underground alcove. Both RF sections are interconnected via ≈ 20 m long, high quality coaxial cables. After down-conversion in three stages, the resulting Intermediate Frequency (IF) signal represents a 15 kHz wide slice of the input spectrum centered around ≈ 4.8 GHz. A spectral analysis of this digitized signal reveals the ‘Schottky sidebands’ with their characteristic features, which are further processed by real-time data extraction algorithms (Section 6) to measure beam and machine parameters such as synchrotron and betatron tunes, chromaticity or momentum spread.

2. Theoretical Background

As shown in Fig. 2, Schottky signals can be separated into spectral components derived from both, the longitudinal and transverse motion of particles. Even though the LHC Schottky pickups have been designed to detect the fluctuations of the charge density in the transverse plane, they also have a residual sensitivity to longitudinal fluctuations, which explains why both components are visible in the measurement.

These can be further subdivided into a coherent part, which corresponds to the collective motion of the whole bunch and an incoherent part, which corresponds to the random movement of individual particles within the bunch. Qualitatively, the former results in strong sharp, well-defined peaks at characteristic frequencies in the spectrum, such as the revolution frequency f_{rev} or betatron frequency f_{β} , while the latter results in an array of lines merging into wide humps of lower amplitude around these peaks.

Extracting a beam parameter of interest requires an analysis of the measured Schottky spectra. See e.g. [8, 9, 10, 11]. The essentials are summarized as follows:

The chromaticity, Q' , is defined as the change in betatron tune $Q = Q_{\text{int}} + q$ (with Q_{int} and q being the integer and fractional tune respectively) for a relative change in momentum $\Delta p/p$ of the beam or an individual particle:

$$Q' = Q\xi = \frac{\Delta q}{\Delta p/p} \quad (1)$$

with $\xi = Q'/Q$ being the relative chromaticity. In practice, with many particles in the beam there is always a momentum spread $\Delta p/p$ present, which causes tune spread $\Delta Q = Q' \cdot \Delta p/p$. The ‘classical’ method to measure the chromaticity utilizes the fact that a change of the RF frequency changes the beam momentum, which leads to a change of the relative revolution frequency $\Delta f/f$ and therefore a change of the betatron tune proportionally to the given chromaticity:

$$\frac{\Delta Q}{Q} = \xi \frac{1}{\eta} \frac{\Delta f}{f} \quad (2)$$

with the slip factor $\eta = 1/\gamma^2 - \alpha$. For the LHC, the momentum compaction factor $\alpha = 1/\gamma_{\text{tr}}^2$ is given by gamma transition $\gamma_{\text{tr}} = 55.68$. The slip factor η relates a change of the relative momentum $\Delta p/p$ to a change of the relative frequency $\Delta f/f$:

$$\eta = \frac{\Delta f/f}{\Delta p/p} \quad (3)$$

where the frequency f can be the revolution frequency f_{rev} or the RF frequency f_{RF} .




Figure 2: Schottky signals at ≈ 4.8 GHz, LHC ion run 2015. The x-axis has been normalized to show the revolution harmonic index n . Revolution harmonics are visible as narrow peaks at integer n , surrounded by the longitudinal Schottky noise bands (red). The upper and lower transverse Schottky sidebands are also visible (dark and light blue).

Due to the periodic nature of the bunch passage, the frequency domain representation of Schottky signals consists of an infinite series of revolution harmonics, at integer multiples of the revolution frequency ($nf_{\text{rev}} = n \cdot 11.245$ kHz). The incoherent motion of the particles is reflected in the bunched-beam spectrum as noise ‘humps’, created by a multitude of lines separated by the synchrotron frequency. The longitudinal band is centered on the revolution harmonic, whereas the upper and lower transverse sideband is located symmetrically

about the revolution harmonic at a distance corresponding to the betatron frequency. Fig. 2 illustrates this relationship with color-coding for a typical measured spectrum from the LHC 2015 ion run. The noise bands originate from amplitude and phase modulations due to the betatron and synchrotron motion of the particles respectively. Note that the structure shown in Fig. 2 repeats over the entire beam pickup bandwidth of ≈ 200 MHz.

2.1. Longitudinal Spectrum

The longitudinal Schottky time-domain signal of the i^{th} particle in a bunched beam at the n^{th} harmonic can be shown [8, 9, 11] to be given by

$$i_i(t) = ze f_{\text{rev}} + 2ze f_{\text{rev}} \text{Re} \left\{ \sum_{n=1}^{\infty} i_n \right\} \quad (4)$$

$$i_n = \sum_{p=-\infty}^{\infty} J_p(n\omega_{\text{rev}}\hat{\tau}_i) \exp j(n\omega_{\text{rev}}t + p\Omega_s t + p\Psi_i),$$

where $i_i(t)$ is the beam current contribution, measured at some location in the ring with an ideal pickup monitor, z is the charge state (for ions), e the elementary charge, n the revolution harmonic index, p the synchrotron satellite index, J_p the Bessel function of first kind of order p , $\omega_{\text{rev}} = 2\pi f_{\text{rev}}$ the angular revolution frequency in radian units, $\hat{\tau}_i$ the synchrotron oscillation amplitude, $\Omega_s = 2\pi f_s$ the angular synchrotron frequency and Ψ_i the initial synchrotron phase.

The synchrotron motion of each particle causes a modulation of its time-of-arrival at the pickup, resulting in a phase modulation (PM) of the beam current signal. In the presence of this PM, the spectrum of a single particle shows a fine-structure in terms of ‘synchrotron satellites’, which are lines spaced by the synchrotron frequency f_s , as illustrated in red in Fig. 3. Their magnitude depends on the value of the argument of the Bessel function ($n\omega_{\text{rev}}\hat{\tau}_i$). The magnitudes are also symmetric around $n f_{\text{rev}}$ due to the property: $J_p(x) = (-1)^p \cdot J_{-p}(x)$. The $(-1)^p$ factor does not impact the magnitude but results in a 180° phase shift for odd p .

With more than one particle in the bunched beam, the synchrotron amplitude and phase turn into random distribution functions, transforming the individual lines in Fig. 3 into bands. This has been illustrated in the upper plot in Fig. 4, where




Figure 3: Magnitude of frequency components of longitudinal (i_i) and transverse (d_i) Schottky signals for a single particle, taking its betatron and synchrotron motion into account.

Eq. 4 has been evaluated for 10^6 simulated particles. Each dot in the upper plot of Fig. 4 corresponds to the frequency and magnitude of one phasor as determined by Eq. 4 and color coded according to the synchrotron satellite index p . The lower plot in Fig. 4 shows a simulated spectrum analyzer measurement, which has been derived by sliding a 1 Hz wide rectangular window over the frequency range and calculating the vector-sum of all phasors which fall into this resolution bandwidth. The $n = 427727^{\text{th}}$ revolution harmonic at 4.809 GHz is shown. The longitudinal parameters loosely approximate the LHC nominal parameters at injection energy [5]. The synchrotron amplitudes of each particle ($\hat{\tau}_i$) follow a Gaussian distribution which is centered at zero and has a one-sigma width $\Delta\hat{\tau} = 0.25$ ns, which was derived from the 1 ns four-sigma nominal bunch length of the LHC. The synchrotron frequency of each particle has a non-linear relationship with its amplitude [12]. For the short bunches (with respect to the bucket length) of the LHC it can be approximated by

$$\Omega_s(\hat{\tau}) = \Omega_{s0} \left(1 - \frac{(\omega_{\text{RF}}\hat{\tau})^2}{16} \right), \quad (5)$$

where $\Omega_{s0} = 2\pi \cdot 61.8$ Hz is the zero-amplitude synchrotron frequency of the LHC at injection energy [5] and $\omega_{\text{RF}} = 2\pi \cdot 400$ MHz is its RF frequency. The initial synchrotron phase was set to uniformly random values in the range of $-\pi < \Psi_i \leq \pi$.

Except for the central line, the power of all synchrotron satellites adds up non-coherently due to their random phase Ψ_i and frequency f_s . The non-linear relationship between synchrotron oscillation amplitude and frequency gives rise to a spread in synchrotron frequency Δf_s . This turns the syn-




Figure 4: *Upper plot*: Monte Carlo simulation of a longitudinal Schottky spectrum at 4.809 GHz, visualizing Eq. 4 for 10^6 simulated particles with parameters typical for the LHC at injection energy. The points represent phasor magnitudes and have been color coded according to the order of the Bessel function p . *Lower plot*: Simulation of a spectrum analyzer measurement by using a sliding window for binning and vector-summing the points within a resolution bandwidth of 1 Hz.

chrotron satellites (excluding the line at $p = 0$) into horizontal bands of width $p\Delta f_s$. Furthermore, magnet current and RF fluctuations can lead to a spread in revolution frequency Δf_{rev} , which broadens all spectral features by $n\Delta f_{\text{rev}}$. The synchrotron bands are overlapping for $p\Delta f_s > f_s$ and $n\Delta f_{\text{rev}} > f_s$, which makes the structure in the longitudinal Schottky spectrum disappear for large values of n or p . This smoothing effect can be seen in Fig. 2 and in Fig. 4 at the lower and upper end of the spectrum. In this case, the momentum-spread of the beam is related to the width of the longitudinal hump [13, 14].

For the central line at 0 Hz in Fig. 4, p is zero and all points fall on the same place in frequency and phase, resulting in the strong coherent peak. It represents the Fourier component of the macroscopic beam current and its amplitude is given by the length and longitudinal shape of the bunch. In practice, these peaks have been observed over a wide range of frequencies (10^4 Hz $< f < 10^{10}$ Hz). On the lower end of this range, their amplitudes can be explained by the Fourier transform of the

bunch shape. At very high frequencies, the longitudinal particle distribution cannot fully explain the observed high signal levels of these revolution harmonics, but microwave effects like intrabunch instabilities may play a role, which have been observed in the Tevatron at FNAL, SPS at CERN, RHIC at Brookhaven, and HERA at DESY [10, 14, 15]. It is important to notice, for the LHC-beam with protons the amplitude of these revolution peaks can rise > 100 dB above the level of the transverse sideband humps to be measured! The design of a signal processing chain, which has to provide a sufficiently large instantaneous dynamic range without saturating from these large coherent harmonics, while still being sensitive to the low-level transverse sideband humps appeared to be one of the most critical engineering challenges of the LHC Schottky system.

2.2. Transverse Spectrum

Following the theory, the transverse (dipolar) time domain signal can be written as

$$d_i(t) = z e f_{\text{rev}} a_i \text{Re} \left\{ \sum_{n=1}^{\infty} d_n \right\}, \quad (6)$$

$$d_n = \sum_{p=-\infty}^{\infty} J_p(\Omega_{\beta} \hat{r}_i) \exp j [(\Omega_{\beta} + p\Omega_s)t + p\Psi_i + \varphi_i],$$

$$\Omega_{\beta} = (n \pm q_i)\omega_{\text{rev}},$$

where a_i is the amplitude and φ_i the initial phase of the betatron oscillation of each particle. The fractional incoherent tune of each particle is given by q_i . In case of ions, the charge state is $z > 1$, the LHC is typically operating with lead-ions of $z = 82$.

The Amplitude Modulation (AM) due to the betatron motion results in the transverse Schottky sidebands, centered at the frequencies $\Omega_{\beta} = 2\pi f_{\beta}$. The spectrum for a single particle has been illustrated in Fig. 3. For a large number of particles the signal needs to be averaged over the random phases Ψ_i and φ_i , and the lines turn into bands in a similar way as for the longitudinal case. Note that for this transverse case, due to the random betatron phase φ_i the power in the central line ($p = 0$) does not add coherently as in the longitudinal case. The width of the transverse sidebands is dominated by the spread of revolution frequencies Δf_{rev} and the spread of tunes Δq_i . In high energy machines, the latter is more significant and originates from a non-zero machine chromaticity.

Note that Eq. 6 describes the motion of the individual particles for a ‘simplistic’ machine, i.e. particles follow only synchrotron and betatron motion, which sums to the incoherent transverse Schottky signal of the bunched beam. The particle motion in the LHC and other hadron ring accelerators is more complex and the bunch as ‘semi-rigid’ body follows forces, e.g. from instabilities (head-tail, electron-cloud), impedance effects, driven kicks (RF modulation, dipole kicks) etc. [8, 16]. In the Schottky spectra, these ‘coherent tune’ q_c effects lead to one or several peaks, super-positioned to the incoherent transverse sidebands, but not necessarily centered. Also non-linear or higher-order effects such as space-charge or beam-beam interactions can influence the shape and width of the transverse Schottky bands. For example, a beam with strong space-charge effects will show an asymmetric transverse Schottky sideband [17]. Its coherent tune line will not be centered on the incoherent noise humps due to a tune-shift effect. This has been discussed for an un-bunched ion beam in [18]. One method to determine the tune-shift due to these effects is to compare the results of the peak detection and curve-fitting method (as explained in Section 6), which measure coherent and incoherent tune respectively. In practice, with the nominal beam parameters of the LHC during physics runs, it has been found that the two tune measurements agree within $\approx 5\%$. Furthermore, no significant asymmetries have been observed, which would become apparent as errors in the gaussian curve-fitting procedure. Nonetheless, during dedicated Machine Development (MD) sessions with high brightness beams, strongly distorted spectra and large tune-shifts have been observed.




Figure 5: LHC Schottky spectrum at injection energy.

For a non-zero chromaticity Eq. 6 needs to be modified [9, 11, 16], the argument of the Bessel function should be replaced by

$$J_p [(\Omega_\beta - \Omega_\xi) \hat{r}_i], \quad (7)$$

$$\Omega_\xi = \frac{Q'}{\eta} \omega_{\text{rev}},$$

where Ω_ξ is the ‘chromatic frequency’. A non-zero chromaticity leads to an asymmetry between the lower and upper sideband humps, see Fig. 5. For large values of n , here $n = 427746$, the chromaticity can be simply evaluated from their relative widths:

$$Q' = \eta \left(n \frac{\Delta f_{\text{lsb}} - \Delta f_{\text{usb}}}{\Delta f_{\text{lsb}} + \Delta f_{\text{usb}}} + q \right) \quad (8)$$

In section 6, two different algorithms are discussed to extract the sideband widths Δf_{usb} and Δf_{lsb} from measured data. Once these parameters are known, Eq. 8 is used to determine the machines chromaticity.

3. Schottky Beam Pickup

The LHC Schottky beam pickups have been designed to be sensitive to the transverse fluctuations of the charge density within a bunch. They couple to the beam at a center frequency of $f_c \approx 4.8$ GHz, over a bandwidth of $BW_{3\text{dB}} \approx 200$ MHz. The choice of the operation frequency is a trade-off between avoiding strong coherent signals at lower, and overlapping Schottky bands at higher frequencies. Sufficient bandwidth is required for effective gating on single bunches separated by 25 ns, which excludes resonant style pickups.

The simplified cut-away view in Fig. 6 shows the basic elements of the beam pickup. It consists of two symmetrically arranged TE_{10} rectangular waveguides (type WR 187), which are placed on adjacent sides of the 60 mm square cross-section beam-pipe. A 0.2 mm thin slotted copper beryllium (CuBe) foil provides forward coupling to the Transverse Electromagnetic (TEM) field of the beam [19].

Due to limitations found during LHC Run I, the pickups were redesigned and upgraded during the long shutdown in 2014. In particular, the dimensions of the 270 rectangular slots were optimized through a semi-analytical approach [21], to obtain a tight coupling to the beam, while ensuring the phase velocity of the TE_{01} Δ -mode $v_p \approx c_0$.

Special care was taken on the four waveguide-to-coaxial transitions of the pickup. The downstream




Figure 6: Internals of a LHC Schottky pickup. Two rectangular waveguides are coupled electromagnetically through an array of slots to the rectangular beam-pipe. 4 Pumping ports allow to meet LHC vacuum requirements (blue arrows), while a mesh-grid shaped cover (light blue) keeps the beam impedance low. Source: [20]

pair acts as signal ports, while the upstream pair is used to feed a test tone signal for calibration and maintenance purposes. The geometry of these coupling pins, their position and the profile of the edge around the coaxial pin provide free parameters for minimizing the return loss. An initial configuration was found through simulations, then each of the 16 couplers were fine-tuned by hand [20]. This was performed using a measurement jig, terminating the waveguide and beam-pipe ports with microwave absorbing foam. Under these conditions, the only significant reflection originates from the coupling pin itself, allowing the optimization based on a S_{11} measurement from a Vector Network Analyzer (VNA). A typical set of measurement results is shown in Fig. 7, demonstrating that this technique provides sufficient sensitivity to discover a faulty coupling pin. One of the coaxial feed-troughs was recuperated from the 2010 pickup installation and suffered from an internal damage due to micro-cracks in its ceramic dielectric.




Figure 7: Measured return loss of several waveguide-to-coaxial transitions after fine-tuning. One transition showed significantly different characteristics due to a material defect, which was discovered by this method.

Before the 2014 upgrade, the pickups suffered from a mechanical issue related to the different thermal expansion coefficients between the aluminum bulk material and the CuBe foil used for the coupling slots. During the bake-out procedure, the two parts elongated to different values, leading to plastic deformation and a warping-effect of the foil. This problem was effectively resolved by machining the new pickups from copper. To ensure a reliable electrical contact to the foil, canted coil spring RF gaskets were incorporated along the entire length of the pickup.

To verify the electromagnetic characteristics of the pickups, their 4-port S-parameters were measured at 3 different points in time: After mechanical assembly in the laboratory (blue trace in Fig. 8), after bakeout and final vacuum leak test in the laboratory (orange trace) and after the final installation in the LHC tunnel (green trace). The parameter S_{31} , which is shown in Fig. 8, is especially sensitive to mechanical tolerances, as it measures the coupling between the two opposite plates of the pickup. For the first two measurements in the laboratory, during which the beam pipe ports have been terminated by microwave absorbing foam, the measurement results indicate no degradation of electrical properties. Nonetheless, after installation in




Figure 8: VNA measurement on one port-pair of the B1V Schottky pickup. Under laboratory conditions with absorbing foam in the beam ports, the coupling at 4.8 GHz is reasonably small (blue and orange trace). However, after installation (green trace) the coupling becomes non-negligible due to microwave modes and reflections at obstacles in the LHC beam-pipe.

the LHC accelerator beam-pipe, the S-parameters changed significantly. This can be explained by the absence of terminating foam on the beam-pipe ports. Due to their rather low waveguide cutoff frequency ($f_{TE01} = 2.5$ GHz), microwave energy can

propagate with low loss along the LHC beam-pipe with multiple reflections at various obstacles up and downstream of the pickup interfering with the measurement.

4. RF front-end

The RF front-end is mounted directly on top of the pickup. It provides common mode rejection, gating for bunch selection, filtering and amplification (Fig. 9). The signals at the inputs of the front-end still contain all spectral components over the full ≈ 200 MHz bandwidth of the pickup. Due to strong coherent components the peak power levels can be high and the active components of the front-end must be able to provide sufficient dynamic range to avoid excessive distortion of the weak underlying Schottky signals.

4.1. Compensation path

The unwanted longitudinal coherent bunch current harmonics at nf_{rev} are of a common mode nature and can therefore be effectively suppressed using a hybrid, which provides the difference of the two input signals at its Δ output. This works as long as amplitude and phase between the two inputs are well matched. Asymmetries caused by manufacturing tolerances, temperature drift related changes in the electrical length of the transmission-line components, and beam position offsets will degrade the common mode suppression. To compensate for these errors, one input of the hybrid has been equipped with a motorized attenuator and delay line. The other input has been equipped with a fixed attenuator and delay line, such that a relative adjustment range of ± 2 dB and ± 85 ps between the two hybrid inputs is achieved.

The sensitivity of the system has been determined by an AC simulation [22] in ‘Advanced Design System’ (ADS). This simulation determines the linearized frequency response, using the measured S-parameters of the actual hybrid as an input. The result is shown in Fig. 10 and indicates that in order to achieve ≈ 30 dB of common mode rejection, the delay line and attenuator must be adjusted with a resolution of $\Delta\tau_2 = 0.4$ ps in electrical length and 0.1 dB in amplitude.

While the installed components provide sufficient resolution, they suffer from non-linear step sizes and hysteresis effects of the internal mechanics, which requires some practice from the operator when carrying out manual adjustments.

Note that ultimately the phase error (φ) at the hybrid needs to be minimized, which depends on the electrical delay (τ_1 and τ_2) of the two (fixed and variable) input paths by

$$\varphi(f) = \text{mod}_{2\pi} [2\pi f(\tau_1 - \tau_2)]. \quad (9)$$

Only if $\tau_1 = \tau_2$ the phase error becomes independent of frequency so that the common mode rejection applies over a wide frequency range, as shown in Fig. 11. If this is not the case φ is frequency dependent, with good common mode rejection achieved only within narrow ranges, leading to periodically spaced notches at frequencies which are a multiple of $1/(\tau_1 - \tau_2)$. Note that such a ‘notched’ common mode rejection would only marginally reduce the integrated power at the front-end amplifier and hence must be avoided. Using only the 15 kHz final observation bandwidth of the down-mixed Schottky acquisition system, it would be impossible to distinguish between notched and wide-band common mode suppression. Hence the compensation path is tuned by monitoring the broadband Δ -signal at the output of the first amplifier on a Vector Signal Analyzer (VSA), adjusting phase and amplitude to achieve minimum integrated power over a bandwidth of ≈ 50 MHz.

During the Schottky upgrade, the electrical length of all critical cables in front of the hybrid were verified and re-matched. A typical VNA measurement is shown in Fig. 11 indicating that the upgraded system can provide up to 40 dB of common mode rejection over a bandwidth of > 1 GHz. This results in a significant reduction of unwanted signal power allowing for more gain to detect the weak Schottky signals before reaching saturation.

The sensitivity of the system to beam orbit variations was analyzed by moving the beam position by a known amount at the pick-up location, as shown in Fig. 12. A degradation in common mode rejection of ≈ 3 dB has been measured for a 0.1 mm change in beam orbit.

Despite all these improvements, not more than 20 dB of common mode rejection could actually be achieved with beam, compared to the 40 dB measured in the laboratory (Fig. 11). This poorer performance might be explained by microwave reflections from the beam pipe, mixing the common and differential mode signals. Due to their large difference in amplitude, even a weak mode mixing is significant, as this power cannot be rejected by the hybrid.




Figure 9: Block diagram of the LHC Schottky front-end electronics, installed directly on top of the pickups.




Figure 10: Contour plot showing the achievable common-mode rejection in dB as a function of attenuator and delay-line setting. Result of a simulation based on measured S-parameters for the hybrid and an ideal model for delay line and attenuator.

4.2. Fast gate

Behind the compensation path sits a fast gating circuit, which allows bunch selection and also acts to minimize the influence of beam pipe reflections, by opening the gate only during the ≈ 5 ns long

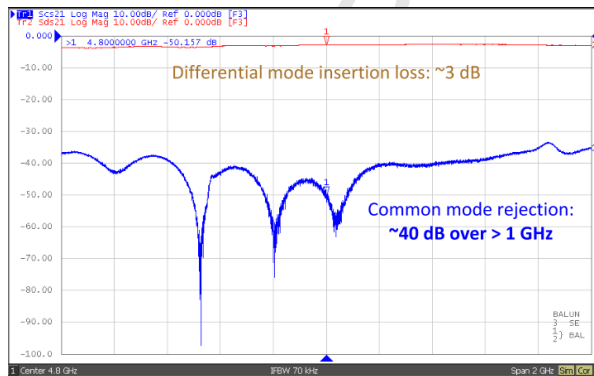


Figure 11: Measured common mode rejection of the compensation path at $f_{\text{center}} = 4.8$ GHz and $f_{\text{span}} = 2$ GHz, using a vector network analyzer with differential ports.




Figure 12: Measured beam orbit dependence of the common mode rejection circuit. The trace shows the power level of one revolution harmonic. Annotated is the change of beam orbit at each step.

bunch response signal. The gating circuit has been designed and built at CERN [23], and is based on the *HMC547LC3* chip from analog devices. The chip requires a control voltage of -5.5 V / $+0.5$ V, which is generated by a *THS3202* current feedback amplifier. The typical switching time in this configuration was measured to be in the order of ≈ 1 ns. By producing the circuit on a microwave substrate, less than 2.1 dB insertion loss at 4.8 GHz was achieved. To generate the trigger pulses for the gate, an FMC form-factor trigger generator was developed. It provides two independently programmable TTL output channels sharing clock and trigger inputs. Delay and pulsewidth can be controlled with sub-ns resolution while the signal path has been optimized for low jitter. The project documents are available on the CERN open hardware repository under the name *FMC-DEL-1NS-2CHA* [24].

4.3. Front-end design

For the Schottky microwave electronics, a trade-off between high sensitivity and low distortion had to be achieved. For instance, to achieve the former, there must be enough gain in the system, such that the weak transverse Schottky sidebands are

amplified above the inherent thermal noise of the active and passive microwave components along the chain. This is discussed in Section 5 by computing a component-wise signal to noise ratio.

On the other hand, the strong coherent signal components in combination with the wide bandwidth of the pickup leads to a sharp spike in instantaneous power level during a bunch passage. Hence gain must be distributed carefully along the chain to ensure that the active microwave components are only adding a minimum of unwanted but unavoidable distortion products to the signal. Distortion can result in an elevated noise floor and additional spectral components (spurs), which overshadow the weak transverse Schottky sidebands. In general, low distortion can be expected when operating an active component several dBs below its 1 dB compression point (P_{1dB}) [25].

To optimize the fronted, the maximum instantaneous power level in the time domain during a bunch passage has been determined from measurements with a fast (18 GHz bandwidth, 60 GS/s) oscilloscope directly at the Δ -output of the hybrid. The measured result is shown as trace ‘Hybrid (1)’ in Fig. 13. For nominal LHC bunches containing 10^{11} protons and a well-adjusted compensation path, a peak voltage of ≈ 2 V and peak power of $+19$ dBm was obtained.

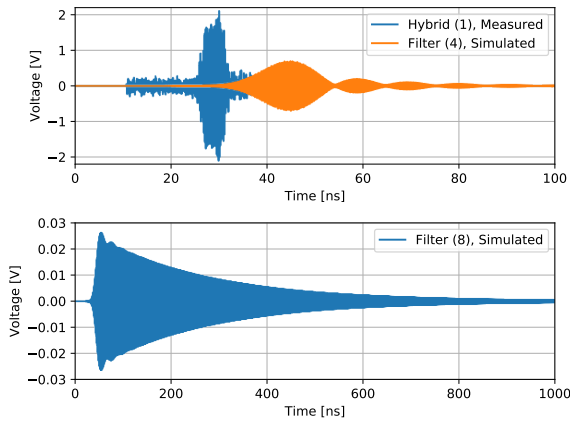


Figure 13: Measured (*Hybrid (1)*) and simulated (*Filter (4)*, *Filter (8)*) output voltage of front-end components during a single nominal bunch passage. Note how the peak voltage is dramatically reduced after the cavity filter due to its long impulse response.

Taking this initial data-point and the small-signal gain of each component into account, the peak power \hat{P}_{out} was calculated along the chain. The

result of this ‘cascade analysis’ is shown in Fig 14, where \hat{P}_{out} is compared to P_{1dB} . This metric was used for choosing a front-end arrangement which strikes a good compromise between sensitivity and low distortion.

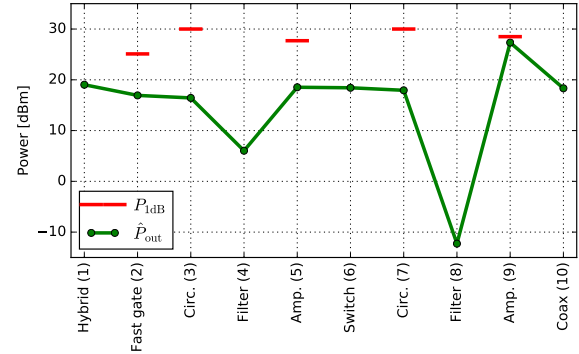


Figure 14: Maximum instantaneous peak power level at the output of each front-end component during a single nominal bunch passage (green) and its corresponding 1 dB compression point (red).

The final design consists of two bandpass filters, two isolators and two low noise amplifiers. Both filters were custom developments at CERN to meet their special requirements and are described in detail in sections 4.4 and 4.5.

Note that both filters provide a significant reduction in the instantaneous peak power, as their impulse response is spreading energy over time. The effect has been quantified by a transient simulation in ADS, based on models derived from measured S-parameter data. The resulting waveforms at the outputs of both filters are shown in Fig. 13 and indicate a reduction of the peak signal level of 10 dB by Filter (4) and 28 dB by Filter (8).

Filter (4) was optimized for minimum insertion loss and an impulse response within ≈ 25 ns. With a relative large 55 MHz bandwidth, it provides frequency pre-selection and prevents excessive peak power at the input of Amp. (5). This amplifier is of type *GRF5020* with modest amount of gain ($G=12$ dB) but high linearity (1 dB compression point $P_{1dB} > 27$ dBm) and low noise figure. Its function is to compensate for the losses of the following Filter (8), which has a bandwidth of only 1.4 MHz, drastically reducing the peak power during the bunch passage for all subsequent components. Amp. (9), which is placed after the narrow-band filter, provides high gain ($G=45$ dB) to overcome the losses of the coaxial cables from the tun-

nel to the alcove and to minimize the effect of any added noise of subsequent components downstream (see Section 5 for details). In the current configuration, it is the output of this amplifier, which is closest to $P_{1\text{dB}}$. This bottleneck will be addressed once Filter (8) has been silver plated (see Section 4.5). This is expected to double its Q factor and therefore reduce the peak power levels of all subsequent components by 3 dB. Isolators (terminated circulators) were added in front of both filters to absorb reflected stop-band signals. With Switch (6) a VSA can be connected to the signal path to observe wide-band (BW= 55 MHz) signals, which are useful to fine-tune the gating circuit and the compensation path.

4.4. Interdigital-line band-pass filter

A band-pass filter of moderate bandwidth is used as pre-selector, and helps to “clip” high common mode signal levels which tend to saturate the following LNA.

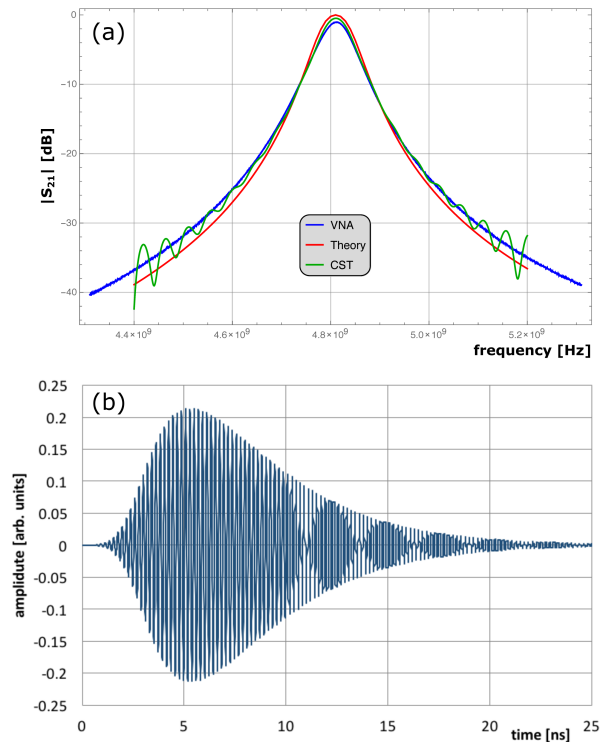


Figure 15: Characteristics of the interdigital band-pass filter (a) S_{21} response, measured (blue), analytical formula (red) and simulated (green). (b) Measured impulse response.

A *Bessel-Thomson* (maximally-flat time delay) band-pass filter approximation was preferred to en-

sure no ringing of the time-domain response of the low-pass prototype (see Fig. 15(b)). The chosen relative bandwidth of $w = 0.15$ results in an impulse response which decays adiabatically to $<1\%$ within the 25 ns bunch repetition time. This allows the filter to be placed before or after the fast gate-switch. The filter was implemented as low-loss interdigital-line design based on parallel, $\lambda/4$ long rectangular bars between ground-planes. A 2-stage design with quarter-wave impedance transformers at the ports was found to be feasible. The normalized self C_k/ϵ and mutual $C_{k,k+1}/\epsilon$ coupling coefficients were computed following the design equations in [26]. From those coupling coefficients, a set of reasonable, and practically machinable cross-section dimensions for the bars were synthesized, while ensuring a high Q-value for the resonators. Based on the conformal mapping solution [27], the equations for the even and odd mode, parallel and fringe capacitances [28] were combined with all other related design equations in a script to compute the mechanical dimensions. A 3D numerical electro-magnetic computation was performed to verify and fine tune the design parameters, Fig. 15(a) compares the results of theory, 3D simulation and measurement for the frequency-domain transfer response $|S_{21}|(f)$.

4.5. Cavity band-pass filter

After gating, the bandwidth of the Schottky signals is reduced substantially using a narrow-band filter. This minimizes the integrated spectral power for subsequent components and, through the impulse response of the filter, spreads the remaining signal energy over a longer time span, hence reducing instantaneous peak signal power. A single-stage, cylindrical-like cavity band-pass filter operating on the TE_{011} mode was designed and optimized for this purpose. Comparatively high Q-factors are achieved for this mode as its surface current only flows along the circumference of the cavity. This design effectively avoids Ohmic losses caused by the contact resistance of removable parts, such as the top and bottom lids.

The mode chart in Fig. 16 (a) for an ideal cylindrical resonator with $f_{TE_{011}} \approx 4.8$ GHz (constant radius of the cavity, $r = 46.5$ mm) allows the dimension of the cavity height to be selected for retaining a good separation of the eigen-modes. Modifications on both end caps help to detune the degenerated TM_{111} -mode [29], and offer a simple solution to fine-tune $f_{TE_{011}}$, see Fig. 16(b). Using brass as the cavity body material results in a $Q_0 \approx 9300$. After

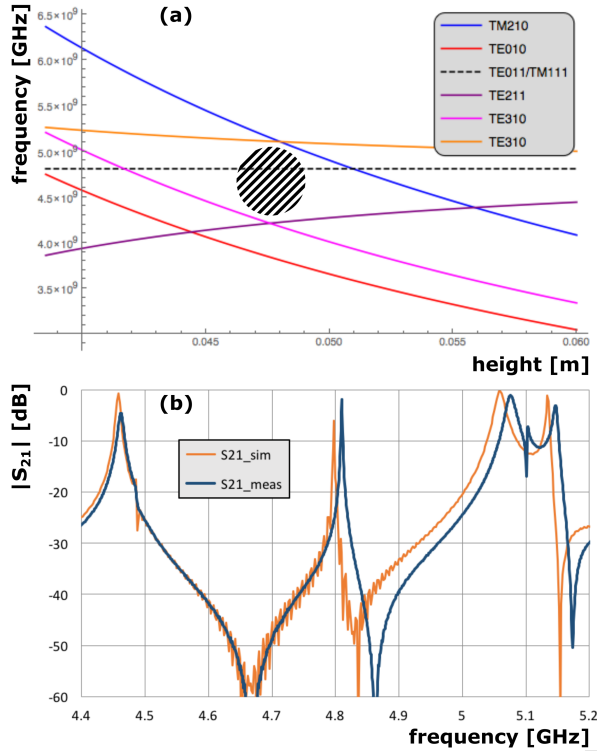


Figure 16: (a) Mode-chart of cylindrical cavity. Region with good mode separation is indicated. (b) Simulated (orange) and measured (blue) $|S_{21}|$ response of TE₀₁₁ filter with modified geometry.

optimization of the coupling loops for a minimum insertion loss of ≈ 2.2 dB the loaded quality factor, Q_L , was measured to be ≈ 3700 , which corresponds to a total signal decay time of $\approx 1 \mu s$. Preliminary tests with silver-plating of the cavity body showed an approximate doubling of Q_0 , as well as of the decay time, but makes the center-frequency more sensitive to variations of the ambient temperature in the accelerator tunnel. Nevertheless, this option is considered as a possibility for a future upgrade.

5. Backend

The front-end is connected by ≈ 20 m of 7/8" corrugated *Heliax* coaxial cable to the backend, located in a service gallery next to the LHC tunnel. A block diagram of the backend components is shown in Fig. 17. A variable attenuator at the input adjusts the incoming signal level, while an electronic gate is used to significantly reduce the average noise power. A cascade of 3 mixer stages is used to down-convert the signals to an intermediate frequency in the audio range.

As the transverse Schottky signal is essentially a pulsed signal of relatively low duty cycle and the background thermal noise is continuously present, the gating circuit makes a substantial improvement in the signal to noise ratio possible. To quantify its effect, a basic analytical system noise analysis has been carried out. In the model, the output noise power density N_{out} (in [W/Hz]) of an amplifier consists of amplified input noise plus an additional noise contribution due to its own internal noise:

$$\begin{aligned} S_{out} &= S_{in}G, \\ N_{out} &= N_{in}G + N_{add}, \\ N_{add} &= T_0(F - 1)kG, \\ F &= 10^{NF/10}, \end{aligned} \quad (10)$$

where S_{in} , N_{in} , S_{out} and N_{out} are the signal and noise power densities at its in and output, G is its linear small signal gain and N_{add} is the additional noise, referred to its output. The latter can be derived from its noise figure NF (in dB) or linear noise factor F , the reference temperature $T_0 = 290$ K and Boltzmann's constant k . For passive devices like attenuators, filters or splitters, the same equations can be applied with $G < 1$ and $F = 1/G$.

The output noise power density of the noise reduction gate is:

$$N'_{out} = D(N_{out}) + (1 - D)kT, \quad (11)$$

where $D = t_{On}/(t_{On} + t_{Off})$ is the gating duty cycle and $T \approx T_0$ the noise temperature of its termination resistor in the 'off' state. Note that the gate is typically driven with $D = 10^{-2}$, limited by the fact that the on-time must be sufficiently long to pass the entire filtered bunch signal response without clipping it ($t_{On} \approx 1 \mu s$), whereas the off-time is given by the time between a bunch passage ($t_{Off} \approx 1/f_{rev} \approx 90 \mu s$). For this case, the average noise power at its output is reduced by ≈ 20 dB, while the signal power is only slightly reduced due to the insertion loss of ≈ 2 dB. Hence, a significant increase of the Signal to Noise Ratio (SNR) is achieved after the gate. There is a trade-off to be maintained between reducing the gating duty cycle to reduce the noise contribution of the amplifier and reducing the bandwidth of the cavity filter, resulting in a longer ringing time, to make the system more robust to high level common mode signal content.

Using Eq. 10 and Eq. 11, the signal and noise power density were calculated for the output of




Figure 17: Block diagram of the backend components.

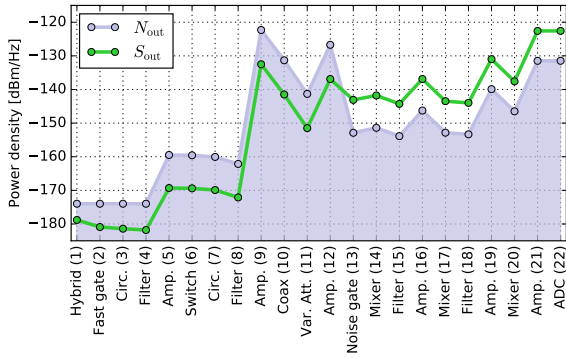


Figure 18: Calculated signal (green) and noise (blue) power density levels at the output of each component. The variable attenuator was set to 10 dB. It is clearly visible how 'Noise gate(13)' reduces noise power substantially while barely affecting signal power.




Figure 19: Calculated signal to noise ratio at the output of each component. Note how only the first 5 components contribute most of the noise in the system.

each component along the chain. The signal power density of the transverse Schottky sidebands was set to $S_0 = -176$ dBm/Hz, which is based on an estimate [19] for a single bunch of a nominal LHC physics fill at injection energy and 2 mm rad emittance. The noise power density was set to $N_0 \approx -174$ dBm/Hz, which corresponds to a 50 Ω system operating at room temperature. The results

are shown in Fig. 18.

Note how the transverse Schottky signal only emerges from the background noise floor at the output of the *Noise gate* (13). Furthermore, the power increase after the high gain *Amp.* (9) is clearly visible, such that the additional noise of the following *backend* components does not have a significant influence on the SNR.

The SNR corresponds to the distance between the green and blue trace, which has been plotted in Fig. 19. The SNR is significantly enhanced after the noise reduction gate due to Eq. 11 ($D = 0.01$). The components in front of the first gain stage and the first gain stage itself dominate the noise contribution, while the noise of subsequent components is almost negligible.

Nevertheless, the components up to and including the first amplifier do significantly diminish the SNR. With the current system, this has been accepted as a necessary trade-off in favor of low-distortion operation. However, a new generation of microwave devices based on Gallium Nitride semiconductors [30, 31] is becoming commercially available. These, in combination with the silver coated front-end filters, have the potential to significantly improve the sensitivity and low distortion operation of the Schottky diagnostic system. The SNR at the very beginning of the chain could be further improved by using movable electrodes to eliminate the insertion losses of an electronic compensation path.

After the noise reduction gate, the signal is down-converted in 3 stages. The first mixer is driven by a Local Oscillator (LO) which is harmonically derived from the accelerator RF frequency utilizing a frequency multiplier.

$$f_{LO1} = 11f_{RF} \approx 4.4 \text{ GHz.} \quad (12)$$

This ensures the best possible phase noise performance. Alternatively, LO1 can be driven by a frequency synthesizer, which gives the flexibility to

choose a different operating frequency. The first Intermediate Frequency ($IF_1 = f_{RF} \approx 400$ MHz) is down-converted to $IF_2 = 21.4$ MHz as this is a standard frequency for narrow-band crystal filters. After another down-conversion stage to $IF_3 = f_{rev} \approx 11.2$ kHz, the signal is digitized with a 24-bit digitizer at a sample-rate of $4f_{rev}$. All local oscillators and the ADC clock are locked to the same RF frequency driving the accelerating cavities. This ensures no frequency drifts in the resulting spectra, even during the energy ramp where the RF frequency changes.

6. Beam measurements and parameter extraction




Figure 20: Typical LHC Schottky spectra for protons at injection (yellow, 450 GeV) and collision energy (purple, 6500 GeV). The area under the transverse sidebands scales with geometric emittance, explaining the drastic difference between the two traces.

As shown in Fig. 20, the transverse sidebands change in their appearance as the beam is accelerated from injection (450 GeV) to collision energy (6.5 TeV). The power in the sidebands reduces with the decrease in geometric emittance, which is proportional to the relativistic gamma of the beam. An ≈ 12 dB reduction in total power is therefore to be expected between signals at injection and at top energy, while the overall width decreases due to a factor 3 reduction in the momentum spread. The peaks of the incoherent Schottky humps are therefore expected to be 7 dB lower at top energy as compared to injection energy.

In order to measure the machine tune, i.e. the difference in frequency between the center of the transverse Schottky sideband and the revolution frequency, two complimentary parameter-extraction algorithms have been developed. The *peak detection* algorithm relies on the sharp coherent peaks

on top of the transverse sidebands. These peaks are usually visible with all proton physics beams, independent of beam energy, making it possible to measure the coherent fractional tune during most machine conditions.

The other algorithm ignores these peaks and relies on curve-fitting to extract amplitude, width and position of the two incoherent transverse Schottky sideband humps. From the position and width of both sidebands, the incoherent fractional tune and the chromaticity can be calculated. This latter algorithm requires relatively broad transverse sidebands with > 10 dB signal to noise ratio, currently limiting its application to beams at injection energies.

6.1. Peak-detection algorithm

This method is used to measure the transverse coherent fractional tune q , as distance between a revolution line and a coherent peak of a transverse sideband. A typically acquired spectrum spans over $2f_{rev}$ as shown in Fig. 21. There are 3 revolution harmonics at -1, 0 and +1, but only the middle one (0) is visible due to the crystal IF filters bandwidth of 15 kHz. As all clocks in the down-converter chain, including the ADC clock, are locked to f_{rev} , the revolution harmonics will always be at these fixed positions in the spectrum. Furthermore there are 4 transverse sidebands visible. The algorithm to determine their distance to a revolution harmonic and therefore the tune q consists of the following steps [32]:

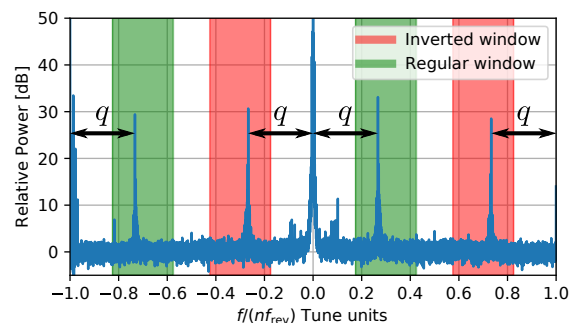


Figure 21: The coherent tune q can be measured as the distance between a revolution harmonic and a transverse sideband as indicated. Note that there are 3 revolution harmonics at -1, 0 and 1. To improve accuracy, data from 4 windows at predefined and fixed positions from these revolution harmonics is overlaid and averaged as shown in Fig. 22. The 2 inverted windows (red) need to be mirrored around the horizontal axis beforehand.




Figure 22: Processing steps of the peak detected tune extractor. For the filtering step, a second order Savitzky-Golay moving average filter with a window length of 13 bins has been used. The parameters have been chosen to match the filter impulse response to the shape of the tune peak, hence maximizing its noise rejection ability while preserving the shape of the peak [33].

Preprocessing To remove amplitude-ripple, a reference spectrum is acquired without beam (or alternatively, with front-end gating in its Off state) and subtracted from all subsequent measurements, resulting in a flat baseline.

Windowing Frequency windows are defined with a fixed position and width around the four transverse sidebands. The frequency axis of the ‘Inverted windows’ needs to be mirrored.

Overlapping Since the Schottky sidebands are repeated around each revolution harmonic, all four windows can be overlapped and averaged, improving the signal to noise ratio, resulting in the black trace of Fig. 22.

Peak detection A Savitzky-Golay [33] moving average filter is applied to minimize the random fluctuations due to the remaining noise in the spectrum while preserving the shape and amplitude of the tune peak. Its two parameters, window width and filter order, are manually tuned for specific beam conditions, such that the filters impulse response approximates the shape of the peak in the spectrum. This maximizes noise rejection. The betatron tune is then calculated based on the position of the global maximum of the filtered peak.

The algorithm has been implemented as Root-based C code and runs on a front-end computer, providing the extracted tune values to the CERN control system in real-time. Its typical performance




Figure 23: Schottky peak detected tune over time compared to the traditional BBQ tune measurement [34]. The BBQ data was binned and averaged over 30 s intervals to match the spectral averaging of the Schottky monitor.

is compared with the traditional LHC tune measurement system ‘Base Band Tune’ (BBQ) [34] in Fig. 23. At 10:45, the LHC changes to ‘collision optics’ and the resulting tune jump is clearly visible and agrees well in absolute terms with the standard system. Note that the BBQ system operates with a measurement rate of 10 Hz (light gray trace), leading to a stronger variance in the result compared to the Schottky system, for which spectra were averaged over 30 s for each measurement (red trace). To allow a fair comparison, BBQ data was binned in 30 second intervals and averaged in post-processing (dark gray trace). During dedicated machine measurement sessions the Schottky system proved to be sensitive enough to detect tune changes in the order of 10^{-4} without any external excitation of the beam.

6.2. Curve-fitting algorithm

To extract the data from the incoherent transverse Schottky signal sideband ‘humps’ a curve-fitting algorithm is applied to the data. This is performed in the following way [32, 35]:

Preprocessing Baseline subtraction as described in Section 6.1.

Sideband splitting The upper and lower Schottky sidebands are extracted from a given frequency range and processed separately. The center position and width of the ranges are user defined constants.

Optional envelope detection In case the raw spectrum shows strong synchrotron peaks (as in Fig. 26), only the envelope function is of

interest, and a peak detection algorithm is applied to extract any local maxima.

Masking A ‘deadband’ is centered on each transverse Schottky hump, where strong coherent peaks are often present. All data within the deadband is removed. The gray dots in Fig 24 show a typical result at this point of processing.

Gaussian curve-fit A Gaussian-like function, Eq. 13, is used to fit each sideband. The `curve_fit` function of the python package `Scipy` is used to carry out a robust non-linear regression [36].

Parameter extraction Frequency offset, amplitude and width are then extracted from the Gaussian fits of both sidebands and are used to compute the relevant beam and machine parameters following Eq. 14 and Eq. 8.

Note that the masking step is critical to discriminate between the coherent and incoherent part of the spectrum. Only the latter provides information from which the chromaticity or the incoherent tune can be extracted. On the other hand, coherent signals, which have been observed as strong sharp peaks close to the center of the incoherent humps can lead to errors in the curve-fit procedure and hence must be removed beforehand. A deadband is used for that purpose. Its position and width are derived from f_β and Δf of the previous curve-fit result, centering the deadband on the Gaussian-like function and setting its width to 1 standard deviation. This is effective under the assumption that the machine tune and chromaticity change significantly slower than the Schottky measurement rate. If the previous curve-fit result is not available, the deadband parameters can be determined from a dedicated iteration through the curve-fit algorithm in which the masking step is skipped. Data points falling in the deadband are simply removed and no interpolation is applied. Therefore the following processing steps must be able to handle the resulting gaps in the data-set. For that reason no smoothing filter was applied, which – as opposed to the Gaussian curve-fit – would require uniformly sampled data.

The Gaussian-like function which is used for the curve-fit of each side-band is

$$y = A \cdot \exp\left(-0.5 \left|\frac{(x - f_\beta)}{\Delta f}\right|^{\text{ORD}}\right), \quad (13)$$

where A is the amplitude, f_β is the frequency offset and Δf the width of the hump. The parameter ORD is a user defined constant, which modifies the slope of the Gaussian curve as shown in Fig. 24. It has been empirically tuned to match the shape of the transverse sideband. Good results were achieved for $\text{ORD} = 3$, as indicated in Fig. 25, where the Root Mean Square (RMS) error between model function and data-points has been used as a metric for a series of curve fits. Note that it is required for the model function of lower and upper sideband to share the same value of ORD, otherwise calculating their difference in width (in Eq. 8) may lead to wrong results. For that reason, ORD has not been included as a free parameter in the the curve-fit procedure.

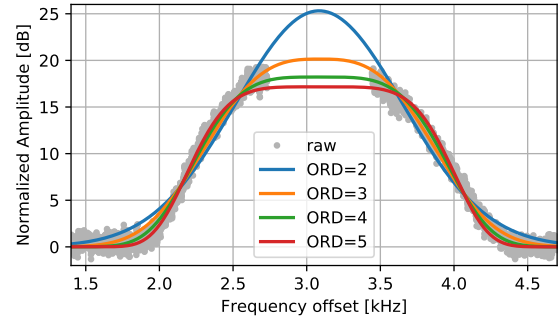


Figure 24: The user-defined ‘ORD’ parameter defines the steepness of the gaussian-like model function. Here the raw data points of a transverse Schottky hump (gray points) and the result of 4 curve fits with different ORD values are overlaid. Note how the slope of $\text{ORD}=2$ undershoots and $\text{ORD}=5$ overshoots the data.




Figure 25: RMS-errors for a series of curve-fits with different ‘ORD’ parameters. Schottky data taken during MD1767 at injection energy with several different chromaticity set-points.

After the parameters of the Gaussian fits from the lower and upper sideband (identified by subscript l and u) are extracted, the fractional tune q is given by:

$$q = \frac{f_{\beta u} - f_{\beta l}}{2f_{\text{rev}}}. \quad (14)$$

The chromaticity Q' is then computed following Eq. 8. To illustrate the required accuracy: for the LHC ($n \cdot \eta \approx 136$) a 1 % width-difference corresponds to ≈ 1 unit change in chromaticity. Hence these measurements are demanding with regards to the accuracy of the curve-fit procedure and the signal to noise ratio of the transverse sidebands.

The emittance ε is proportional to the area below the transverse sidebands. For absolute emittance measurements, the system gain of front-end and down-converter chain needs to be known and long-term drifts need to be eliminated by a regular calibration procedure. So far, this has not been studied in detail with the LHC Schottky system.

Note that the curve-fit algorithm suffers from the effect of coherent lines and artifacts on top of the incoherent transverse sideband humps. They are not taken into account by the ‘model function’ of Eq. 13 and hence must be treated as *outliers*. While the masking step with a deadband is effective at removing most of these outliers, some of them may remain if they occur far from the center of the hump. Hence an additional method was identified for making the fitting procedure more accurate and stable in the presence of a few outliers.

A traditional least-squares regression minimizes the squared error between model function and data-points. The larger the distance between the expected model function value and an outlier, the stronger is its impact on the parameters, which results in biasing errors. This can be avoided by a robust least squares algorithm [37]. The weighting metric is modified by an $\arctan()$ function, softly clipping the maximum value of the residuals and hence limiting the impact of outliers on the curve fit result. The effectiveness of this method is demonstrated in Fig. 26.

6.3. Threshold algorithm

A disadvantage of the previously described data-extraction method is its high computational cost. A simpler algorithm has been tested, which does not rely on curve-fitting and hence can be more easily implemented on a front-end computer. As shown in Fig. 27, it is based on smoothing the data (in




Figure 26: Data processing steps with the curve-fit method. Gray: Raw spectral data, synchrotron peaks are clearly visible. Blue: Selected data points after envelope detection. The ‘Masking’ step has been skipped for the sake of demonstrating the effectiveness of the robust fit method with challenging data. Green: Fit-result of a least squares fit. Red: result of a robust least squares fit.

the same fashion as the peak detection algorithm in Section 6.1) and using a simple threshold detection of the transverse sideband humps to measure their difference in width. As this method determines the two points on each side of the hump, where the amplitude value crosses a certain threshold, e.g., $\text{THR} = 3$ dB above the noise-floor, it is not affected by the coherent peaks around the center of the hump. For that reason, no data masking step is necessary.

Choosing a low threshold value allows weaker Schottky spectra to be measured and avoids synchrotron satellites and other irregularities (as shown in Fig. 27) from disturbing the measurement. Furthermore, THR should be at least 2 dB above the average noise floor to make it unlikely for the filtered random fluctuations to generate a threshold crossing. Within these guidelines, the chromaticity results are insensitive to the choice of threshold-value, as long as the same is chosen for both sidebands. This is demonstrated in Fig. 28.

6.4. Experimental results

Schottky chromaticity measurements were benchmarked during a dedicated machine development session, of which the results are shown in Fig. 29. The LHC sextupole magnets were used to achieve certain chromaticity set-points in the range of $0 < Q' < 20$, shown by the black

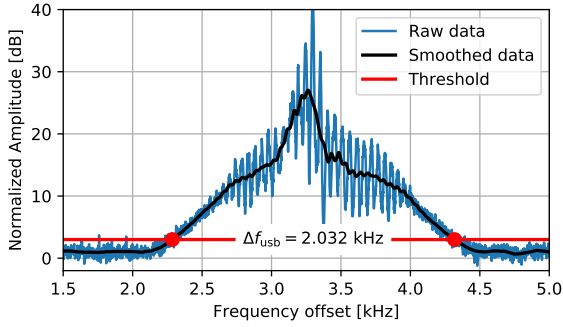


Figure 27: Data processing with the ‘threshold’ method. The raw Schottky data of the upper transverse sideband is shown in blue. The moving average filtered data in black. The threshold value, $\text{THR} = 3$ dB above the average noise floor, is indicated as red line. The two red points indicate the threshold crossing, their horizontal distance determines the measured width of the hump (Δf_{usb}). The same algorithm is applied to the lower transverse sideband. Equation 8 is used to get Q' .

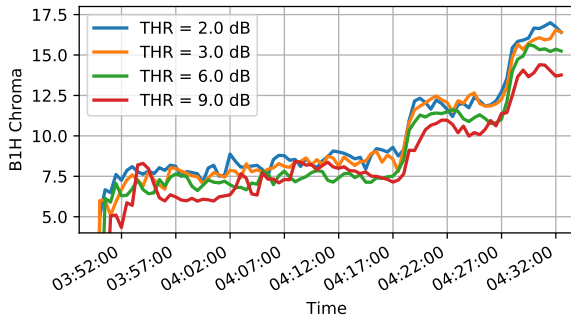


Figure 28: Results of chromaticity extraction by the threshold method for different THR parameters. Schottky data taken during MD1767 at injection energy.

curve in Fig. 29. Chromaticity measurements, using the described curve-fitting (blue trace) and threshold (orange trace) extraction algorithms, were compared with the traditional beam energy modulation method [6] (red points).

Note that from 3:45 to 4:45 the machine was filled with a single bunch of nominal intensity (10^{11}). From 4:40 onwards, the machine was filled with 5 bunch-trains of 96 nominal bunches each and an additional train of 12 nominal bunches. For machine safety reasons, the traditional momentum modulation method could not be applied in the latter configuration. Hence there are no ‘Reference-points’ during that timespan in Fig. 29. Furthermore, the gaps in the traces correspond to times when the

LHC filling pattern or the front-end gating were adjusted and no Schottky signals could be acquired.

Both parameter extraction methods proved to be satisfactory at capturing relative changes in Q' . The absolute values agree better with the reference measurements than with the set-points, which can be attributed to the fact that the system setting the chromaticity is not well calibrated. During the time from 4:45 to 6:00 the chromaticity set-point was held constant. Nonetheless, both extraction methods reveal a slight value drift, which the authors believe to be a real change of machine parameter. Furthermore, the Schottky spectral averaging time constant of 30 s becomes apparent during the large chromaticity set-point steps at 4:32 and 6:35.

To get a quantifiable comparison between the two extraction methods, the RMS error between the reference-points (red-dots) and the extracted chromaticity results from matching instances in time were calculated. This was done for the whole dataset, consisting of 4 planes (beam 1 and 2, horizontal and vertical plane), resulting in 15 valid measurement points in the timespan of 4:15 to 4:50. Furthermore the standard deviation of the extraction results was calculated. For this, a time window from 3:50 to 4:15 was chosen, where the chromaticity set-point is constant. The results are shown in Table 1 and indicate that the curve-fit method provides lower RMS-error, meaning its results agree closer with the ones from the traditional RF modulation method. Furthermore the lower standard deviation indicates that the curve-fit results are less noisy compared to the threshold method. The slightly less accurate and more noisy results of the threshold method can be explained by considering that it extracts information from only two smoothed data points of each transverse sideband (the actual threshold crossing points), whereas the curve-fit method takes almost all of the data-points into account. This tends to better suppress the inherent random fluctuations and improves robustness.

Both measurement methods suffer from systematic uncertainties. The coherent peaks around the center of the transverse humps cause biasing errors in the curve-fit procedure. Hence these data-points must be treated as outliers and have been removed before the curve-fit by a deadband. Furthermore, a curve-fit algorithm which is more robust against the impact of strong outliers, compared to the classical least-squares regression, has been chosen. In practice, these measures prove to be sufficient for




Figure 29: Result of chromaticity extraction with the curve-fitting (blue) and threshold (orange) method. The uncalibrated set-point (black trace) and reference measurements by the traditional RF modulation method (red dots) are also shown. Protons (beam 1, horizontal plane) at injection energy during MD1767.

Method	RMS-error	Std. dev.
Curve-fitting	2.07	0.44
Threshold	4.42	0.52

Table 1: To compare the two data extraction methods, the absolute RMS-error (in chromaticity units) has been measured between extracted chromaticity and 15 reference points. Furthermore, the standard deviation (also in chromaticity units) of the 60 extracted chromaticity values during 3:50 to 4:15 has been determined.

a useful chromaticity measurement during nominal physics operation at injection energy. While the threshold method provides slightly less accurate and more noisy results, it is not affected by the coherent peaks around the center of the transverse sidebands. Nonetheless it is sensitive to spurs and artifacts around the edges, which might lead to a wrong threshold point being detected. Both methods suffer from low SNR and very small difference in width of the sideband humps at collision energy (7 TeV) with protons, causing a large uncertainty in the measured chromaticity result. Furthermore, under these conditions, no good method has been found yet to distinguish between the coherent and incoherent part of the spectrum, as both features are of about the same width and overlaying each other (see the purple trace in Fig. 20).

During the 2015 and 2016 runs, the amplitude of the transverse Schottky sideband was typically 30 dB above noise level for ions (see Fig. 30), 15 dB for protons at injection energy (450 GeV, see Fig. 26) and 5 dB for protons at 7 TeV (Fig. 20).

The excellent results with ions in Fig. 30 are due to the fact that the power density of the Schot-

tky bands scale quadratically with charge number z . Hence signal power is much increased for lead-ions ($z = 82$) compared to protons ($z = 1$). Furthermore, the figure illustrates how geometric beam emittance, which is reduced at higher energies, relates to an reduction in area under the incoherent Schottky bands. Furthermore a shift in tune is apparent, shifting the center of the transverse sidebands outwards.




Figure 30: Schottky spectra of a LHC lead-ion bunch at 450 GeV (black) and 6370 GeV (purple).

7. Conclusion

The upgraded LHC transverse Schottky system delivers improved signal quality, enabling the non-invasive, bunch gated measurement of beam parameters, such as betatron tune and chromaticity for protons and ions beams. It is anticipated that this system will in the future allow a regular on-line measurement of chromaticity with high intensity proton physics production beams, something which is not currently possible using existing methods. Some operational challenges still remain, such as producing reliable measurements during the energy

ramp when the proton bunches are intentionally excited longitudinally to maintain a constant bunch length. While the strong unwanted coherent signals remain a challenge for the front-end electronics, new microwave amplifier developments based on wide-bandgap semiconductors like Gallium Nitride are a promising match for the high dynamic range requirements and have the capability to further improve the sensitivity and low distortion operation of the Schottky signal processing chain.

8. Acknowledgement

The authors are grateful for the practical hints and assistance from T. Levens, F. Caspers, M. Gasior, J. Fullerton, M. Hamani and T. Bogey. Thanks to the BE department management for support.

- [1] W. Schottky, Über spontane Stromschwankungen in verschiedenen Elektrizitätsleitern, *Ann. Phys.* 57 (1918) 541–567.
URL <http://doi.org/10.1002/andp.19183622304>
- [2] F. Caspers, D. Möhl, History of stochastic beam cooling and its application in many different projects, *Eur. Phys. J. H* 36 (2012) 601–632.
URL <http://cds.cern.ch/record/1447119>
- [3] S. van der Meer, Stochastic damping of betatron oscillations in the ISR, *Tech. Rep. ISR-PO-72-31*, CERN, Geneva (1972).
URL <http://cds.cern.ch/record/312939>
- [4] J. M. Brennan, M. Blaskiewicz, Stochastic cooling in RHIC, *Proc. of PAC09*, Vancouver, B.C., Canada (2009) 1910–1914.
URL <https://accelconf.web.cern.ch/accelconf/pac2009/papers/we3gri01.pdf>
- [5] O. S. Brüning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole, P. Proudlock, LHC Design Report, CERN Yellow Reports: Monographs, CERN, Geneva, 2004.
URL <https://cds.cern.ch/record/782076>
- [6] K. Fuchsberger, G. H. Hemelsoet, LHC online chromaticity measurement - experience after one year of operation, *Proc. of IBIC*, Barcelona, Spain (2016) 20–23.
URL <http://accelconf.web.cern.ch/AccelConf/ibic2016/papers/mob104.pdf>
- [7] A. Boccardi, M. Gasior, R. Jones, K. Kasinski, R. Steinhagen, The FPGA-based continuous FFT tune measurement system for the LHC and its test at the CERN SPS, *Proc. of PAC07*, Albuquerque, New Mexico, USA (2007) 4204–4206.
URL <http://cds.cern.ch/record/1055568>
- [8] S. Chattopadhyay, Some fundamental aspects of fluctuations and coherence in charged-particle beams in storage rings, *Proc. of Summer School on Physics of High-Energy Particle Accelerators*, Upton, New York, USA (1983) 486–641.
URL <http://doi.org/10.1063/1.35175>
- [9] S. van der Meer, Diagnostics with Schottky noise, *Proc. of 3rd US-CERN School on Particle Accelerators*, Anacapri, Italy (1988) 423–433.
URL <https://cds.cern.ch/record/195311?ln=en>
- [10] M. Blaskiewicz, J. M. Brennan, Bunched beam stochastic cooling in a collider, *Phys. Rev. ST Accel. Beams* 10 (2007) 061001.
URL <http://doi.org/10.1103/PhysRevSTAB.10.061001>
- [11] D. Boussard, Schottky noise and beam transfer function diagnostics, *Proc. of CERN Accelerator School: Accelerator Physics*, Berlin, Germany (1989) 90–92.
URL <http://cds.cern.ch/record/302475>
- [12] J. F. Esteban Müller, Longitudinal intensity effects in the CERN Large Hadron Collider, Ph.D. thesis, école polytechnique fédérale de Lausanne (2016).
URL <http://infoscience.epfl.ch/record/219008>
- [13] E. Shaposhnikova, T. Bohl, T. Linnecar, Longitudinal peak detected Schottky spectrum, *Proc. of HB2010*, Morschach, Switzerland (2010) 363–367.
URL <http://accelconf.web.cern.ch/Accelconf/HB2010/papers/tuo1c04.pdf>
- [14] D. A. Goldberg, G. R. Lambertson, A high-frequency Schottky detector for use in the Tevatron, *Proc. of American Institute of Physics* 229 (1991) 225–234.
URL <http://doi.org/10.1063/1.40735>
- [15] R. J. Pasquinelli, A. Jansson, Microwave Schottky diagnostic systems for the Fermilab Tevatron, Recycler, and CERN Large Hadron Collider, *Phys. Rev. ST Accel. Beams* 14 (2011) 072803.
URL <http://doi.org/10.1103/PhysRevSTAB.14.072803>
- [16] M. E. Castro Carballo, Transverse diagnostics for high energy hadron colliders, Ph.D. thesis, Universidade de Santiago de Compostela (2007).
URL <http://cds.cern.ch/record/1103499>
- [17] D. V. Pestrikov, Dipole coherent oscillations and fluctuations of a coasting ion beam with strong space charge, *Nucl. Instr. Meth. Phys. Res.* 562 (2006) 65–75.
URL <http://doi.org/10.1016/j.nima.2006.03.006>
- [18] S. Paret, V. Kornilov, O. Boine-Frankenheim, T. Weiland, Transverse Schottky and beam transfer function measurements in space charge affected coasting ion beams, *Phys. Rev. ST Accel. Beams* 13 (2010) 022802.
URL <http://doi.org/10.1103/PhysRevSTAB.13.022802>
- [19] T. Kroyer, F. Caspers, J. M. Jimenez, O. R. Jones, C. Vuitton, T. W. Hamerla, A. Jansson, J. Misek, R. J. Pasquinelli, P. Seifrid, D. Sun, The 4.8 GHz LHC Schottky pick-up system, *Proc. of PAC07*, Albuquerque, New Mexico, USA (2007) 4174–4176.
URL <http://accelconf.web.cern.ch/Accelconf/p07/PAPERS/FRPMN068.PDF>
- [20] M. Ehret, Untersuchung eines breitbandigen Koaxial-Hohlleiterübergangs für einen Schottky-Strahl diagnosedetektor, Master's thesis, Offenburg University of Applied Sciences (2013).
URL <http://cds.cern.ch/record/1642139>
- [21] D. McGinnis, Slotted wavguide slow-wave stochastic cooling arrays, *Proc. of PAC99*, New York, USA (1999) 1713–1715.
URL <https://accelconf.web.cern.ch/accelconf/p99/PAPERS/TUP100.PDF>
- [22] Keysight, Technical documentation for AC simulations, accessed: 2017-08-10.
URL <http://literature.cdn.keysight.com/litweb/pdf/ads2005a/pdf/cktsimac.pdf>

- [23] M. Wendt, M. Betz, O. R. Jones, T. Lefevre, T. Levens, Improvements to the LHC Schottky monitors, Proc. of IBIC16, Barcelona, Spain (2016) 453–456.
URL <http://accelconf.web.cern.ch/AccelConf/ibic2016/papers/tupg46.pdf>
- [24] M. Betz, M. Wendt, T. Levens, J. Pospisil, CERN open hardware repository - FMC DEL 1 ns 2 cha delay module, accessed: 2017-08-10.
URL <http://www.ohwr.org/projects/fmc-del-1ns-2cha/wiki>
- [25] C. Rauscher, Fundamentals of spectrum analysis, Rhode & Schwarz, 2005.
- [26] G. Matthaei, L. Young, E. Jones, Microwave filters, impedance-matching networks, and coupling structures, McGraw-Hill Book Company, 1964.
- [27] J. D. Cockroft, The effect of curved boundaries on the distribution of electrical stress round conductors, J. IEE 66 (1926) 385–409.
URL <http://doi.org/10.1049/jiee-1.1928.0033>
- [28] W. J. Getsinger, Coupled rectangular bars between parallel plates, IRE Trans. Microw. Theory Tech. 10 (1962) 65–72.
URL <http://doi.org/10.1109/TMTT.1962.1125447>
- [29] B. Yassini, M. Yu, B. Keats, A ka-band fully tunable cavity filter, IEEE Trans. Microw. Theory Tech. 60 (2012) 4002–4012.
URL <http://doi.org/10.1109/TMTT.2012.2224367>
- [30] K. W. Kobayayashi, Y. Chen, I. Smorchkova, B. Heying, W. B. Luo, W. Sutton, M. Wojtowicz, A. Oki, A cool, sub-0.2 dB, ultra-low noise gallium nitride multi-octave MMIC LNA-PA with 2-watt output power, Proc. of CSICS08, Portland, Oregon, USA (2008) 1–4.
URL <http://doi.org/10.1109/CSICS.2008.54>
- [31] O. Axelsson, K. Andersson, Highly linear gallium nitride MMIC LNAs, Proc. of CSICS12, La Jolla, California, USA (2012) 1–4.
URL <http://doi.org/10.1109/CSICS.2012.6340109>
- [32] O. Chanon, Schottky signal analysis: tune and chromaticity computation, Tech. Rep. CERN-STUDENTS-Note-2016-177, CERN, Geneva (2016).
URL <http://cds.cern.ch/record/2212811?ln=en>
- [33] R. W. Schafer, What Is a Savitzky-Golay filter?, IEEE Signal Processing Magazine 28 (2011) 111–117.
URL <http://doi.org/10.1109/MSP.2011.941097>
- [34] M. Gasior, S. Jackson, R. Steinhagen, Advancements in the base-band-tune and chromaticity instrumentation and diagnostics systems during LHC's first year of operation, Proc. of DIPAC11, Hamburg, Germany (2011) 476–478.
URL <https://cds.cern.ch/record/1372193>
- [35] M. Betz, M. Wendt, T. Lefevre, Summary of LHC MD377: Schottky pick-up, Tech. Rep. CERN-ACC-NOTE-2015-0039, CERN, Geneva (2015).
URL <http://cds.cern.ch/record/2110648>
- [36] The Scipy community, Technical documentation for `scipy.optimize.curve_fit`, accessed: 2017-08-10.
URL https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.curve_fit.html
- [37] B. Triggs, P. F. McLauchlan, R. I. Hartley, A. W. Fitzgibbon, Bundle adjustment - a modern synthesis, Proc. of Vision Algorithms 99, Corfu, Greece (1999) 298–372.
URL http://doi.org/10.1007/3-540-44480-7_21