GHOSTS AND SATELLITES MEASURED WITH THE LHC LDM

A. Jeff, CERN, Geneva, Switzerland & University of Liverpool, UK

Abstract

Accurate calibration of the LHC luminosity depends on knowledge of the size and distribution of ghost and satellite bunches. The very high dynamic range and high time resolution of the Longitudinal Density Monitor (LDM) allows satellites and ghosts to be clearly identified and quantified.

A short summary of the LDM design is given here, along with a discussion of how the ghost and satellite proportions are extracted from the LDM profiles.

INTRODUCTION

Each ring of the LHC is divided longitudinally into 3564 slots, only some of which contain a bunch. Each slot covers 10 RF buckets. In this paper, the slot is defined such that if the slot is filled, the bunch occupies the 5th bucket of the slot. The other 9 buckets would, in a perfect machine, be empty; in reality they can contain some charge and these are called satellite bunches. In an unfilled slot all 10 buckets should be empty; again in reality they contain some charge and these are defined as ghost bunches. These definitions are illustrated in fig 1 with an example profile from the heavy ion run 2010. Note that the bunch at -15 ns (six buckets before the main bunch) is counted as part of the ghost charge, since it is outside the slot, even though it has the same origin as the satellites.

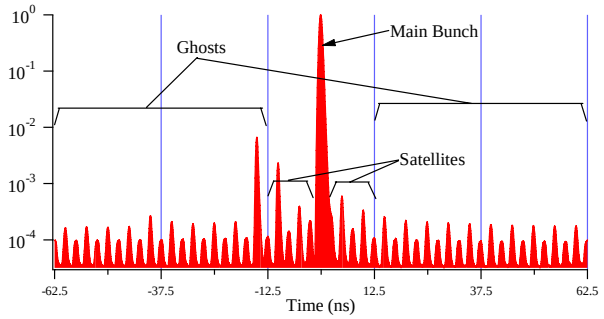

Figure 1. Definition of ghosts and satellites. Longitudinal profile in logarithmic scale.

It should be noted that both satellite and ghost bunches are captured in an RF bucket, and are thus accelerated and focused in the same way as the nominal bunches. It is also possible to have debunched beam in the machine, i.e. charge which is not captured by the RF and which therefore drifts around the ring. At 3.5 TeV, any debunched charge would spread round the whole ring. Debunched charge can exhibit a 400 MHz modulation in antiphase with the bunched charge, although as it spreads away from the source bunch this modulation is flattened [1]. Beam which is debunched at 450 GeV cannot survive the ramp.

The luminosity of a pair of colliding bunches depends on the product of the bunch currents. Relative bunch-by-bunch currents are measured by the fast BCT [2], and must then be normalized to the absolute current measurement performed by the DCCT [3]. The fast BCT integrates the charge in each filled slot. Its response is optimized for the main bunch by scanning the integrator phase; surrounding satellite bunches may or may not be integrated. Any slot which does not reach a certain threshold, i.e. an empty slot, is set as zero. In reality, these slots contain ghost charge. Therefore, the ghost charge proportion must be allowed for when normalizing the fast BCT to the DCCT, which integrates all charges circulating in the machine:

$$\sum Fast\ BCT + \sum Ghost\ \&\ debunched = DCCT$$

The satellite, ghost and debunched charges themselves give a negligible contribution to luminosity, since the bunch current product of two colliding satellite/ghost bunches is very much smaller than the product for two nominal bunches.

PRINCIPLE OF THE LONGITUDINAL DENSITY MONITOR

Synchrotron radiation (SR) has the same time structure as the emitting particle beam, so that making a longitudinal profile of the SR is the same as making a profile of the beam. The LHC is equipped with two sets of synchrotron light monitors, one for each beam, known as the BSRTs [4], located either side of the RF cavities at point 4. SR is emitted in a purpose-built undulator and in the D3 separation dipole immediately downstream. The wavelength and intensity of SR are strongly dependent on the beam energy. Simulations using the SRW code [5] show that the dominant source of visible synchrotron light changes during acceleration from the undulator to the edge field of the dipole and then the main body of the dipole.

Due to the small bending angle of the dipole, only 27m downstream of the dipole are the SR and particle beams sufficiently separated to extract the synchrotron light. The SR is reflected downwards through a silica window onto an optical table where it is focused by two spherical mirrors. A beam splitter reflects approximately 7% of the available light onto the LDM detector. The LDM line is equipped with two filter wheels allowing the attenuation to be controlled independently of the other BSRT instruments, and the LDM is mounted on translation stages for horizontal and vertical alignment.

The LDM uses time-correlated single-photon counting (TCSPC) [6] to produce a longitudinal profile. Each bunch passes the LDM once every turn and a photon could be detected from anywhere in the bunch, or not at all. In order to make a meaningful bunch profile the data

have to be collected over many turns. Incoming photons are detected by a silicon avalanche photo-diode (APD). The APD is operated in Geiger mode, so that each photon (if detected) produces a measurable electrical pulse. A time-to-digital converter (TDC) then time-stamps the arrival of the pulse relative to the LHC turn clock. The histogram of arrival times which is built up is then the longitudinal beam profile (fig 2). The LDM system is described in more detail in [7].

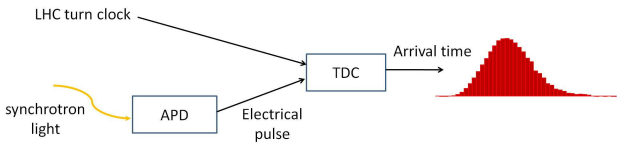

Figure 2. Schematic of the LDM system.

The longer the acquisition, the more counts are added to the histogram and the higher the dynamic range of the measurement. An integration time of 10 s (~10^5 turns) allows the shape, length and relative populations of the main bunches to be measured. Quantifying the satellite and ghost fractions requires an integration time of at least 5 minutes (~3×10^6 turns). A dynamic range of better than 10^5 has been reached with an even longer integration time.

BEHAVIOUR OF THE DETECTOR

In an APD operated in Geiger mode each photoelectron causes a self-sustaining avalanche. The rising edge of the avalanche is detected by an active quenching circuit. This circuit reduces the bias voltage across the APD in order to quench the avalanche, and also shapes an output timing pulse. The shape and magnitude of the output pulse is independent of the size of the avalanche. An avalanche caused by two co-incident photons is indistinguishable from a single photon avalanche.

Most avalanches propagate very quickly, and the timing of the output pulse has a FWHM jitter of just 40 ps. However, if the initial photoelectron lies in the so-called diffusion region of the APD, it must first travel into the high-field junction region, before the avalanche begins. Thus, a small proportion of the APD counts have a substantially longer delay. This is the 'diffusion tail' of the APD response.

The operating voltage of the APD is restored after a fixed time, during which the avalanche will have dissipated. This is known as the deadtime of the detector, which is 77 ns for the LDM detector. During the deadtime, the detector is blind to any further photons.

During an avalanche, charge carriers may become trapped in impurities in the silicon. These trapped carriers will be released some time later, when they can cause a new avalanche. This is known as afterpulsing and is the dominant source of noise for the LDM detector.

The APD used in the LDM is the Photon Detection Module (PDM) from Micro Photon Devices [8]. Its response is shown in fig 3, showing the diffusion tail, the deadtime and afterpulsing. This is not the output pulse shape of the device, which is irrelevant in a photon-counting system. Instead it is the histogram of photon counts time-stamped by the TDC and integrated over 10^8 cycles of the pulsed laser.

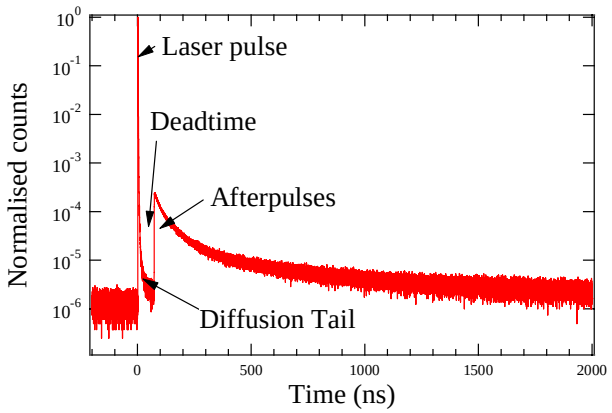

Figure 3. Photon-counting response histogram of the APD, measured in the lab for 10^8 cycles of a pulsed laser.

CORRECTION FOR DEADTIME AND AFTERPULSING

The deadtime of the detector introduces a distortion to the measured bunch shape. If a photon is detected from the earlier part of the bunch, the detector will be blind to any further photons, so that the measured pulse is shifted towards earlier time, as well as reduced in magnitude. The strength of this deadtime effect is strongly dependent on the count rate.

We can define the availability a_i of the detector in any bin i as the probability that it was not in deadtime, i.e. that no photon had been detected in the previous 77 ns. Then

$$a_i = 1 - \sum_{j=i-\tau}^{i-1} \frac{x_j}{N}$$

where x_j is the number of counts in bin j after an acquisition of N turns.

The decreasing availability of the detector through the bunch can be seen in fig 4, along with its effect on the measured bunch shape. Here the bunches are assumed to be isolated, in other words separated by more than 77 ns, which is usually the case during van der Meer scans.

It can be seen that for large counting rates the availability approaches zero and no information can then be gained about the later part of the bunch, or any trailing satellites. A very low counting rate means that the distortion due to the deadtime is negligible. However, in this case the acquisition time needed in order to build up the histogram to a given dynamic range would be very large. An optimal compromise is reached around 0.5 photons / bunch / turn. In this regime the distortion due to deadtime is significant, but a correction can be applied in order to restore the true bunch shape.

The probability that a photon was received in bin i is

$$p_i = \frac{x_i}{a_i N}$$

and we can therefore apply a correction to calculate the number of counts that would have been measured by a detector with no deadtime. The availability for each bin is calculated from the measured histogram using a running sum [9].

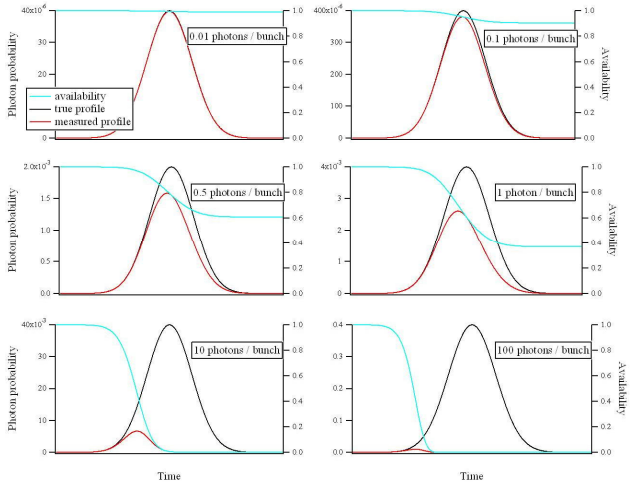

Figure 4. Effect of the detector deadtime on the measured bunch shape, for different numbers of photons arriving per turn.

A second-order correction accounts for the possibility of two photons arriving in the same bin. The probability of this occurring is small but not negligible for the peak bins of each bunch. Using Poissonian statistics, we can write

$$p_i = 1 - e^{-C_i/N}$$

where p_i is the probability of seeing at least one photon in bin i and C_i is the true number of photons arriving in N turns. Equating the previous two equations,

$$C_i = -\mathrm{N} \ln\left(1 - \frac{x_i}{a_i N}\right)$$

The majority of the noise present in the histogram is caused by afterpulsing. The probability of an afterpulse at some time after any count is modelled by a sum of multiple exponentials, with time constants between 100 ns and a few µs. Trapped charge carriers are not released by the next avalanche, so that the afterpulsing probability due to different photon counts is additive. The afterpulsing can then be statistically eliminated by subtracting a series of infinite impulse responses (IIR) filters. The parameters of the IIRs are found by experiment. The best-fitting parameters vary slightly according to the LHC filling scheme and LDM count rate, probably due to a change in the internal temperature of the detector. In practice, therefore, afterpulsing is not totally eliminated, but is very strongly reduced by use of the IIRs.

CALCULATION OF THE GHOST AND SATELLITE FRACTIONS

The corrected histogram must first be separated into slots. Each slot is approximately 499 bins long. The exact length changes slightly between 450 GeV and flat top, as does the LDM's phase with respect to the turn clock, due to the change in RF frequency and in the synchrotron light source, respectively. The exact phase and width are therefore calculated each time, by using the center of Gaussians fitted to the first and last bunch in the histogram.

Next, 'average slot' histograms are created by adding the slots together bin by bin. Separate histograms are created for the filled slots and the empty slots. Examples are shown in fig 5 and 6 respectively.

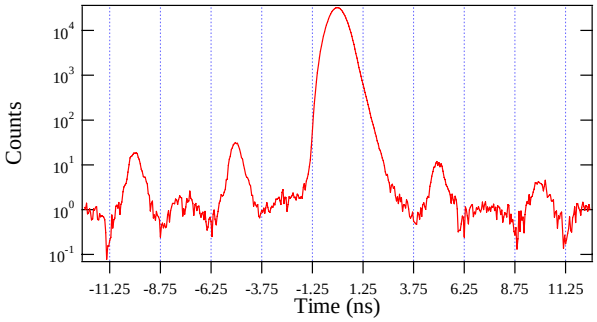

Figure 5. Averaged filled slot. From the beam 1 longitudinal profile during lead ion fill 2292 (November 2011)

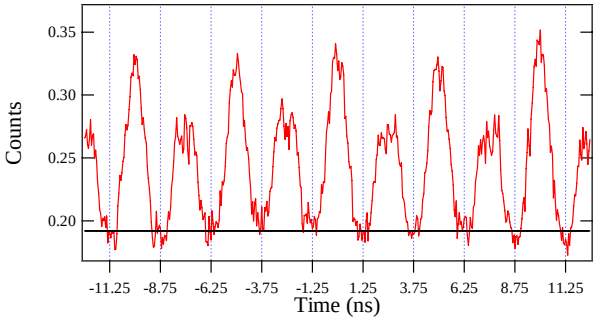

Figure 6. Averaged empty slot. From the beam 1 longitudinal profile during lead ion fill 2292 (November 2011). The baseline level is shown.

Under the assumption that there is no debunched beam, we would expect there to be zero counts at the junction of two buckets (the bucket separatrix) but this is not the case. Even allowing for the LDM time resolution, and in particular the diffusion tail of the APD, we would expect the value at the separatrix to be only 1% of the preceding bunch peak, as shown for the main bunch in figure 4**.** In fact the value at the separatrix is found to be much higher than this. This noise baseline should be removed before calculation of the ghost and satellite fraction.

In order to calculate the baseline, we take the mean of the ten bins corresponding to the bucket separatrices. The bucket separatrices are calculated slot phase and width, not by finding the minimum of the average slot, as this makes the baseline less sensitive to noise.

The true baseline is not constant around the ring. Since it arises mostly from uncorrected afterpulsing, it is usually higher towards the end of a train, and closer to

zero far from any bunches. However, only the average baseline is of interest for calculating the ghost and satellite fractions.

After subtracting the baseline, the ghost and satellite proportions are calculated.

It can be seen in figure 4 that the diffusion tail of the APD means that the bucket immediately following the main bunch is swamped by delayed counts from the main bunch. Any satellite in this bucket cannot be distinguished. This bucket is not counted in the satellite proportion, nor are the separatrices each side of the main bunch counted in the baseline calculation.

RESULTS FOR VAN DER MEER SCANS

For accurate luminosity calibration, it is particularly important to know the bunch currents during the van der Meer scans, and this requires knowledge of the charge fraction outside the main bunches. The LDM was in use during all van der Meer scans since its installation. The results are shown below. Where multiple LDM datasets are available, the results are averaged over the duration of the van der Meer scans, and the satellites are averaged over all main bunches. The satellite fraction averaged over a subset of bunches, e.g. only bunches colliding at a given IP, can also be calculated from the LDM profile.

Table I. Ghost charge measured by the LDM during van der Meer scans, defined as the percentage of beam charge outside filled slots.

	Beam 1		**Beam 2**	
Nov 2010	n/a		2.5	+2.5 / -0.6
March 2011	n/a		0.6	+0.6 / -0.2
May 2011	0.18	+0.18 / -0.04	0.40	+0.4 / -0.1
Oct 2011	0.69	+0.7 / -0.2	0.71	+0.7 / -0.2
Dec 2011	3.1	+3.1 / -0.8	2.8	+2.8 / -0.7
Dec 2011	2.1	+2.1 / -0.5	2.3	+2.3 / -0.6

Table II. Satellite charge measured by the LDM during van der Meer scans, defined as the percentage of beam charge inside the filled slots but outside the filled bucket.

	Beam 1		**Beam 2**	
Nov 2010	n/a		0.4	+/- 0.1
March 2011	n/a		0.1	+0.05/-0.03
May 2011	0.02	+0.03 / -0.01	0.12	+/- 0.03
Oct 2011	0.21	+0.06 / -0.08	0.57	+0.17 / -0.12
Dec 2011	0.38	+/- 0.1	0.34	+0.1 / -0.08
Dec 2011	0.28	+0.08 / -0.06	0.29	+0.08 / -0.06

The LDM was installed on beam 2 in October 2010 and on beam 1 in April 2011. Initially, a different brand of detector was used on beam 1, but the results were not as good so it was replaced with a detector of the same type as beam 2 in August 2011. Since August 2011 the LDM profiles have been logged automatically whenever there is beam in the machine. There are two entries for December 2011 as the van der Meer scans were carried out over two fills, 2335 and 2337.

Shown below are the main sources of error in the LDM ghost / satellite charge estimation. The values given are typical uncertainties at the 68% confidence level, for an integration time of 5 minutes, but the actual uncertainty depends on the number of counts in the histogram, which varies between acquisitions. The relative error on the satellite fraction is smaller since the satellites themselves are larger.

Table III. Sources of error in the LDM ghost / satellite measurement

	Error for ghosts	Error for satellites
Statistical	10 %	5 %
Baseline	12 %	3 %
Emittance	20 %	20 %
Debunched beam	100 %	25 %
	-25% / +100%	**-20% / +30%**

Since the baseline of the LDM histogram is set on the assumption that there is no debunched beam, the LDM can at present only measure the bunched beam component. The assumption that there is no debunched beam is compatible with beam-gas data, which shows zero counts at the bucket separatrix. However, due to limited statistics this can only provide an upper limit to debunched beam, and the presence of debunched beam is not necessarily constant across fills. To allow for debunched beam an arbitrary error of +100% is assigned to the ghost population, which is then taken to mean all charge outside the filled slots.

The active area of the APD has a diameter of only 50 μm. The beam spot produced by the synchrotron light telescope is roughly Gaussian with a sigma between 100 and 500 μm depending on the emittance of the beam. This reduces the coupling efficiency since only a fraction of the beam spot can be sampled. In addition, it creates a dependence of the coupling efficiency (and therefore the measured beam population) on the transverse size of the bunch, and therefore on the bunch emittance. If the detector is centered in the beam spot, then bunches with larger emittance will appear to have lower population. Conversely, if the detector position is away from the beam center, bunches with larger emittance will appear to have a larger population. The alignment of the LDM detector with respect to the beam spot is not known for some of the van der Meer datasets, since the alignment of the LDM independently of the BSRT was only possible after the installation of additional translation stages in August 2011.

If the transverse emittance of the satellite / ghost bunches is systematically different from that of the main bunches, this will result in an incorrect estimation of their population. Based on investigation of this effect during an MD with groups of large- and small-emittance bunches [10], a +/- 50% difference in emittance between ghosts /

satellites and main bunches would lead to a 20% difference in the measured population. This is a very conservative estimate of the uncertainty, since when the transverse size of the satellites has been able to be measured it has been within 10% of the main bunches.

OUTLOOK FOR 2012

The optical line for the beam 2 LDM has been modified during the winter stop 2011/12 in order to reduce to a minimum the dependence on transverse beam size. This will be repeated on beam 1 if the results are positive. The new setup is shown in figure 7. A diffuser with a Gaussian point spread function (PSF) scatters the incident light. This produces a beam spot with size

$$\sigma_{After\ Diffuser}{}^2 = \sigma_{Original}{}^2 + \sigma_{PSF}{}^2$$

The diffuser is chosen with $\sigma_{PSF} >> \sigma_{Original}$, so that the dependence of the spot size after the diffuser on the original beam size is negligible. The large size of the beam spot causes an unacceptable loss of coupling efficiency, so a lens is used to increase the amount of light captured.

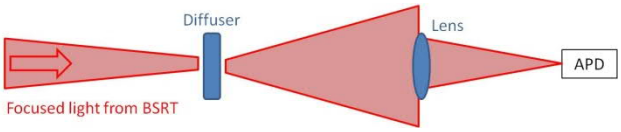

Figure 7. Modification of the optical line to eliminate dependence on transverse beam size.

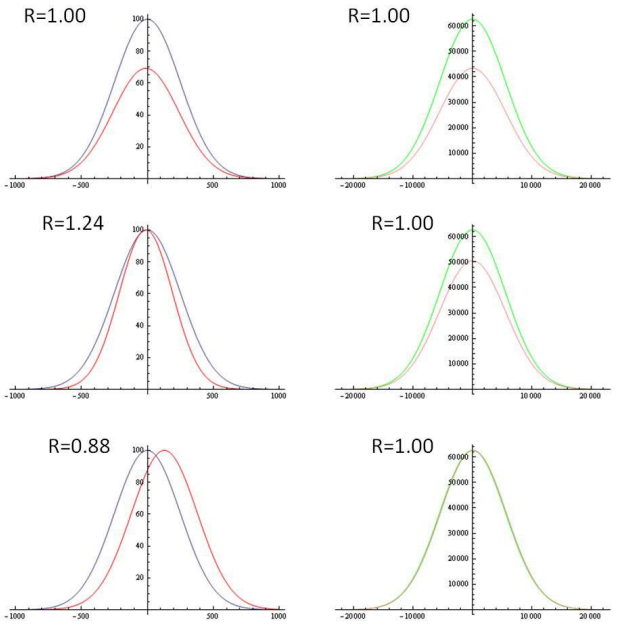

Figure 8. Simulation of the effect of transverse beam size on the measured population. Left, original setup. Right, new setup with diffuser and lens. *R* is the ratio of the measured to the true population of the bunch in red.

A simulation of the effect of the diffuser on the measured population is shown in figure 8. In the original setup (left) the relative population of a bunch with different transverse beam size or position is incorrectly measured. After the diffuser, however, there is negligible difference in the transverse profiles and no effect on the measured populations.

CONCLUSION

The LHC longitudinal density monitor can produce longitudinal beam profiles with a dynamic range of more than 10^5. It is used during van der Meer scans to quantify the ghost and satellite fraction, in conjunction with data from the experiments.

The accurate determination of ghost and satellite populations from the LDM relies on correction for the effects of deadtime and afterpulsing, and on subtraction of the noise baseline. The baseline can be effectively subtracted, but only with the assumption that there is no debunched beam in the machine.

The principal sources of uncertainty have been identified and efforts are underway to reduce them.

REFERENCES

[1] E. Shaposhnikova, "Longitudinal motion of uncaptured particles in the LHC at 7 TeV", LHC Project Note 338 (2004)

[2] D. Belohrad, L. K. Jensen, O. R. Jones, M. Ludwig, and J.-J. Savioz, "The LHC Fast BCT system: A comparison of Design Parameters with Initial Performance", Proc. BIW (2010) p269.

[3] C. Barschel, "Results of the LHC DCCT Calibration Studies", these proceedings.

[4] T. Lefevre et al., "First Beam Measurements with the LHC Synchrotron Light Monitors", Proc. IPAC (2010) pp. 1104-1106.

[5] O. Chubar and P. Elleaume, "Accurate and efficient computation of synchrotron radiation in the near field region", Proc. EPAC (1998) p1177.

[6] W. Becker, *Advanced Time-correlated Single Photon counting techniques* (Berlin Heidelberg: Springer, 2005).

[7] A. Jeff et al., "First results of the LHC longitudinal density monitor", Nucl. Instrum. & Meth. A, **659**, 11 (2011).

[8] www.microphotondevices.com

[9] E. Bravin, "Dead time effects on single photon counting for the longitudinal density monitor of LHC", CERN-AB-2006-017-BI (2006).

[10] E. Bravin et al., "LHC Transverse Profile Monitors studies (MD on May 6th, 2011)", CERN-ATS-Note-2011-049 MD (2011).