

A Longitudinal Density Monitor for the LHC



Thesis submitted in accordance with the requirements of the

University of Liverpool

for the degree of Doctor in Philosophy by

Adam Jeff

November 2012

Abstract

At the Large Hadron Collider (LHC), the world's largest and highest energy particle accelerator, ion bunches circulate in two counter-rotating beams and are brought into collision. Each bunch is confined within a bucket by the longitudinal focusing effect of the radio frequency (RF) cavities. The RF period is 2.5 ns, while the minimum bunch spacing is 25 ns. Thus, 9 out of every 10 buckets should be empty, as well as additional gaps to allow for the rise-time of injection and dump kickers. In practice, however, small numbers of particles can occupy these supposedly empty buckets, causing problems for machine protection and for the absolute calibration of the LHC's luminosity.

The Longitudinal Density Monitor (LDM) is a new monitor, designed to measure the longitudinal distribution of particles in the LHC with a sufficiently high dynamic range to quantify the relative particle population in the supposedly empty buckets. A non-interceptive measurement is made possible by the use of synchrotron radiation (SR). Single photon counting with an avalanche photo-diode operating in Geiger mode allows a very high dynamic range to be achieved despite the low levels of light available. The imperfect response of the avalanche photo-diode is compensated using a specially designed correction algorithm which reduces noise and distortion to a minimum.

This work presents the design, implementation and operation of the LDM. Signal correction methods are discussed with reference to the deadtime and afterpulsing of the avalanche photodiode, and the analysis of the LDM data for use in LHC luminosity calibration is explained. Experimental results with both proton and heavy ion beams are shown illustrating the LDM's exceptional performance, combining a high dynamic range of 10^5 with a 90 ps time resolution. Finally, a novel scheme to extend the dynamic range by several more orders of magnitude is presented.

Table of Contents

1. Introduction	9
1.1 The LHC and its Aims.....	9
1.1.1 Basic Principles of Radio Frequency Acceleration	10
1.1.2 Longitudinal Structure and Focusing	11
1.1.3 Equations of Longitudinal Motion	13
1.1.4 The Longitudinal Phase Space.....	16
1.1.5 Longitudinal Emittance and Bunch Length	19
1.1.6 LHC Filling Schemes.....	20
1.1.7 Injector Chain and Effect on Bunch Structure	21
1.2 Satellite and Ghost Bunches	23
1.2.1 Formation of Satellite and Ghost Bunches	24
1.2.2 Importance of Satellite and Ghost Bunches	27
1.3 Objectives and Specifications of the Longitudinal Density Monitor.....	28
Chapter Summary	28
2. Background	29
2.1 Interceptive.....	29
2.2 Electro-magnetic.....	29
2.2.1 Beam Current Transformers	29
2.2.2 Wall Current Monitor.....	31
2.2.3 Electro-optic.....	31
2.3 Optical.....	32
2.3.1 Synchrotron Radiation	32
2.3.2 Optical Transition Radiation.....	33
2.3.3 Optical Diffraction Radiation.....	33
2.3.4 Cherenkov Radiation.....	34
2.3.5 Smith-Purcell Radiation.....	34
2.3.6 Scintillators.....	35
2.4 Optical Profile Methods	35
2.4.1 Detector & Scope	35
2.4.2 Single Photon Counting.....	36

2.4.3 Streak Camera	36
2.4.4 Laser Mixing	37
2.4.5 Frequency-Resolved Optical Gating.....	38
2.4.6 Coherent Radiation Spectrum.....	40
2.4.7 Shot Noise Technique	40
2.5 Synchrotron Radiation	41
2.5.1 Dipoles	42
2.5.2 Edge Radiation	45
2.5.3 Undulators and Wigglers.....	45
2.5.4 SR from Protons and Ions.....	47
2.5.5 Time Structure of SR	48
Chapter Summary	51
3. Theory & Simulations.....	52
3.1 The Beam Synchrotron Radiation Telescopes	52
3.2 Arrangement of the Optical System	54
3.3 Synchrotron Radiation: Analytical Solution	59
3.4 Synchrotron Radiation: Simulations with SRW	60
3.4.1 Spectrum of Emitted SR	62
3.4.2 Intensity and Distribution of Emitted SR	65
3.4.3 Heat Load on the Extraction Mirror.....	70
3.5 Synchrotron Radiation: Simulations with SPECTRA	73
Chapter Summary	75
4. Single Photon Counting with Avalanche Photo-Diodes	76
4.1 Single Photon Counting.....	76
4.1.1 Principle of Single Photon Counting	76
4.1.2 Advantages and Disadvantages of SPC	77
4.1.3 Detectors for SPC	78
4.2 Avalanche Photo-Diodes.....	78
4.2.1 Principle of Photon Detection with Avalanche Photo-Diodes	78
4.2.2 Construction.....	79
4.2.3 Active and Passive Quenching	80
4.2.4 Deadtime.....	81

4.2.5 Afterpulsing.....	86
4.2.6 Dark Counts.....	87
4.2.7 Diffusion Tail	87
4.2.8 Gated APDs	87
4.3 Correction for the Detector Response.....	87
4.3.1 Principle of Deadtime Correction	87
4.4 Monte Carlo Simulation of the Detector Response and Signal Correction	89
4.4.1 Testing the Theoretical Detector Response at Different Count Rates.....	89
4.4.2 Testing the Signal Correction Algorithm - Finding the Optimal Count Rate	91
4.4.3 Testing the Signal Correction Algorithm – Deadtime Jitter	93
4.4.4 Testing the Signal Correction Algorithm - Pile-up Correction.....	94
4.4.5 Simulating the Response of a Gated Detector.....	94
4.5 Testing the Detector Response in the Laboratory	96
4.5.1 Instrument Response	96
4.5.2 Testing the Deadtime Effect	97
4.5.3 Afterpulses	98
4.5.4 Gating the Detector	103
Chapter Summary	106
5. Method	107
5.1 Overview of the System	107
5.2 Components.....	107
5.2.1 Avalanche Photo-Diode (APD)	107
5.2.2 Time-to-Digital Converter (TDC)	109
5.2.3 Beam Synchronous Timing System	111
5.2.4 Data Handling.....	111
5.3 Optical Layout	111
5.3.1 Free-space and Fiber Coupling.....	111
5.3.2 Fiber Coupling	111
5.3.3 Location of the LDM on the BSRT Optical Table	114
5.4 Software and Control.....	115
5.4.1 Filter Wheels	115
5.4.2 Translation Stages	116

5.4.3 Data Processing in the Front-End.....	117
5.4.4 Graphical User Interface (GUI).....	118
5.4.5 Logging.....	119
Chapter Summary	120
6. Results with Beam and Discussion.....	121
6.1 Example Profiles of the LHC Beams	121
6.1.2 Proton Beams.....	121
6.1.2 Lead Ion Beams.....	123
6.2 Analysis of the LDM Data.....	124
6.2.1 Effect of Signal Correction	124
6.2.2 Analysis Methods and Software Routines	125
6.2.3 Re-Binning.....	126
6.2.4 Calculation of Average Slots.....	127
6.2.5 Diffusion Tail	128
6.2.6 Baseline	129
6.2.7 Ghost & Satellite Fractions.....	130
6.2.8 Comparison with Results from the Experiments	133
6.2.9 Enhanced Satellites	135
6.3 Time Resolution	136
6.3.1 Sources of Timing Uncertainty.....	138
6.4 Dynamic Range & Uncertainties	140
6.4.1 Limiting Factors.....	140
6.4.2 Sources of Uncertainty.....	141
6.4.3 Dependency on Emittance.....	142
6.4.4 Modification of the Beamline	144
6.5 Beam Debunching.....	147
6.6 3D Bunch Shape Measurement	151
6.6.1 Method	151
6.6.2 Limitations.....	151
6.7 Relevant Results from Other Instruments	151
6.7.1 Synchrotron Light Intensity.....	151
6.7.2 Distribution of Light on the Extraction Mirror	152

Chapter Summary	154
7. Future Improvements	155
7.1 Improved Characterisation of the Detector Response	155
7.2 Software Improvements	155
7.3 Ultra-high Dynamic Range Method	155
Chapter Summary	160
8. Conclusion.....	161
8.1 Applicability to Other Machines	161
8.2 Conclusions	161
Acknowledgements.....	163
References	166

1. Introduction

1.1 The LHC and its Aims

The Large Hadron Collider (LHC) at CERN is the world's largest particle accelerator [1]. It occupies a tunnel 27 km long located outside Geneva and lying partly in Switzerland, partly in France. Particles accelerated in the LHC reach the highest energies ever achieved in a man-made accelerator. During the 2010-11 run protons were accelerated to an energy of 3.5 TeV [2], and 4 TeV was reached in 2012 [3]. The design energy, 7 TeV, is expected to be reached during 2014. In addition to protons, the LHC is also capable of accelerating heavy ions. In the first such run, fully stripped lead ions ($Z=82$, $A=208$) were used. These were accelerated up to 287 TeV, or 1.38 TeV per nucleon.

Two beams circulate in opposite directions. When they are brought into collision, the energy released can create new particles and resonances. The types and distributions of these particles, and those they decay into, are measured by hugely complex particle detectors. The resulting data will help to further our understanding of the fundamental physical laws of the Universe. The existence of theoretically predicted particles such as the Higgs boson will be proved or excluded, and new physics beyond the Standard Model could be observed.

The two beams travel in separate beam pipes which pass through different field regions of the same magnets. The magnets are arranged such that the two beams always feel the same magnitude of magnetic field, though in opposite directions, and thus the two beams must always circulate at the same energy. When the desired energy is reached the two beams can be steered into collision at designated points. By convention, the clockwise circulating beam is referred to as beam 1 and the anticlockwise circulating beam as beam 2.

Most of the LHC's circumference is devoted to dipole magnets, that is magnets in which only a vertical component of the magnetic field is present, whose job it is to bend the beams on a circular trajectory. The LHC is not, however, strictly circular. There are 8 straight sections without dipole magnets, at the centre of each of which is an interaction point (IP). Each IP has a different purpose, as summarised in Figure 1. Four are given over to the LHC's four flagship experiments: ATLAS [4] in point 1, ALICE [5] in point 2, CMS [6] in point 5 and LHCb [7] in point 8. IPs 3 and 7 are for collimation, the removal of particles whose trajectory and / or momentum is outside the tolerated values. The beam dump, a six ton block of graphite capable

of safely absorbing the entire beam in the event of a beam abort, is located at IP 6. Finally, the acceleration of the particles occurs at point 4.

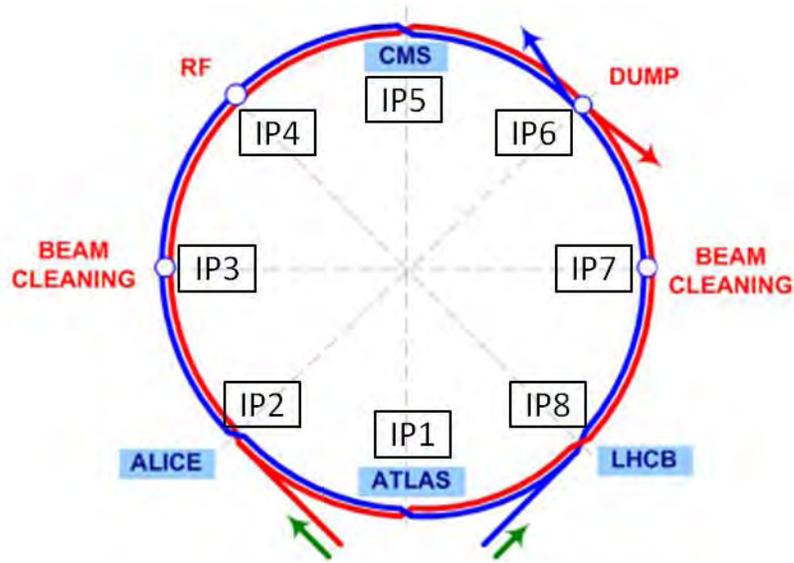


Figure 1. Layout of the interaction points (IP) of the LHC. From [8]

1.1.1 Basic Principles of Radio Frequency Acceleration

While magnetic fields are used to steer and focus the beams, electric fields are used to accelerate them. In order to allow continuous acceleration with a finite voltage, an alternating field is used. The field is generated by setting up a standing wave in a specially shaped resonating cavity. Since these waves are in the radio-wave part of the spectrum, this is known as radio frequency or RF acceleration. In the LHC, superconducting RF cavities are used, because of their high accelerating field and high stored energy, which minimises the effects of transient beam loading [9]. Each beam passes through eight single-cell cavities, each capable of delivering an accelerating potential of 2 MV. The LHC cavities oscillate at a frequency of 400.8 MHz. Some important parameters of the LHC RF system are shown in table 1.

Table 1. Key RF-related parameters of the LHC

	Protons	Lead Ions
Maximum voltage per cavity	2 MV	
Number of cavities (per beam)	8	
Revolution period	89 μ s	
Harmonic number	35,640	
RF period	2.5 ns	
Minimum bunch spacing	25 ns	100 ns
Maximum number of bunches	2,808	592

At the energies used in the LHC the beams are travelling at 99.9999% of the speed of light and it takes just 89 μ s for them to make one revolution. Because of special relativity, an increase in energy is manifested mainly as an increase in the relativistic gamma factor, not particle velocity. During acceleration of protons in the LHC from 450 GeV to 3.5 TeV, the revolution period changes by less than 2 parts per million (16 parts per million for lead ions).

1.1.2 Longitudinal Structure and Focusing

Since the field in the RF cavities oscillates, particles passing through the cavities may be accelerated or decelerated depending on the phase of the cavity when the particle arrives. A continuous stream of particles entering the cavity with the same energy will leave it with an energy modulation. In order to prevent a large increase in energy spread, particles circulate in bunches which are timed to pass the cavities at a certain phase. The centre of the bunch should reach the cavity at the same phase with each revolution. Thus, the RF frequency ω_{rf} must be an integer multiple of the particle revolution frequency ω_{rev} . This is the harmonic number h ,

$$\omega_{rf} = h\omega_{rev} \quad (1)$$

As well as accelerating the beam, the RF cavities provide a longitudinal focusing force which prevents the bunches from spreading out too far. The LHC operates above transition, meaning that a particle with higher energy has a lower revolution frequency, since the increase in its magnetic rigidity, and therefore in its path length, is greater than the increase in its speed. In

order to maintain longitudinal focusing, therefore, the RF should be synchronised such that the centre of a bunch arrives on the decreasing slope of the electric field. A particle with lower energy than the rest of the bunch arrives at the cavity earlier, when the field is stronger (more positive), and therefore receives a larger increase in energy, bringing it closer to the bunch average. Conversely, a particle with higher energy arrives later, feels a less positive or a negative field, and loses energy with respect to the bunch average [10].

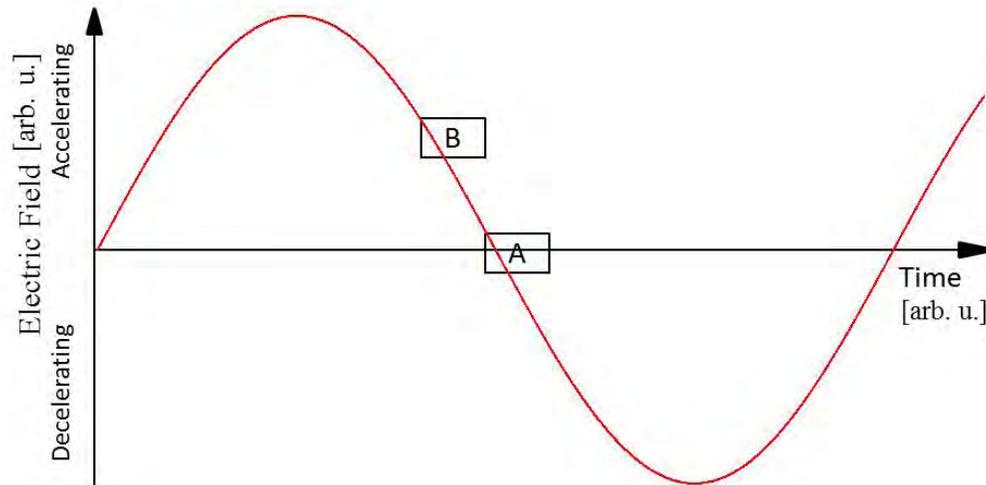


Figure 2. Electric field in an RF cavity against Time. A positive electric field indicates the field is in the direction of particle motion (accelerating), a negative field indicates the opposite direction (decelerating). A: position of bunch centre for longitudinal focusing without acceleration. B: position of bunch centre for longitudinal focusing with acceleration.

Because of this need for longitudinal focusing and to replace the small amount of energy lost by synchrotron radiation, the RF cavities are always powered during LHC operation, even when no acceleration is needed. When the LHC is maintaining a constant energy, the RF is synchronized so that the centre of the bunch arrives at the zero-crossing of the downwards slope of the electric field (position A in Figure 2). Thus, there is no increase in the average energy of the bunch. When acceleration is needed the bunch centre moves to the positive part of the downwards slope of the RF (position B).

1.1.3 Equations of Longitudinal Motion

Consider a bunch which is centred at the zero-crossing of the RF field (position A in Figure 2 above). A particle which crosses the cavity at this point will receive no acceleration or deceleration, and if its initial momentum is correct it will pass the cavity at the same point on its next revolution. This is known as the synchronous particle and the phase of the RF as the synchronous particle passes the centre of the cavity is the synchronous phase, φ_s .

Now consider a general particle which has a small deviation from the synchronous particle in momentum or in longitudinal position [11]. Its energy E , momentum p , revolution period τ and RF phase φ may be written as

$$E = E_s + \delta E \quad (2)$$

$$p = p_s + \delta p \quad (3)$$

$$\tau = \tau_s + \delta\tau \quad (4)$$

$$\varphi = \varphi_s + \delta\varphi \quad (5)$$

where the subscript s denotes the properties of the synchronous particle. Assuming that the particles pass through the RF cavity much faster than the RF period, the RF phase can be considered constant during the passage of the particle, and the energy gain per turn is given by

$$\Delta E = qV \sin \varphi \quad (6)$$

where V is the maximum RF voltage. From eq. (2), the change per turn in the deviation from the energy of the synchronous particle can be written as

$$\Delta(\delta E) = \Delta E - \Delta E_s = qV(\sin \varphi - \sin \varphi_s) \quad (7)$$

Assuming that δE changes in small steps then

$$\frac{d(\delta E)}{dt} \approx \frac{\Delta(\delta E)}{\tau_s} = \frac{qV}{2\pi} \omega_{rev} (\sin \varphi - \sin \varphi_s) \quad (8)$$

The momentum deviation is related to the energy deviation by

$$\frac{\delta p}{p} = \frac{1}{\beta^2} \frac{\delta E}{E} \quad (9)$$

where β is the relativistic beta, $\beta=v/c$. Similarly the difference in revolution period is given by

$$\frac{\delta \tau}{\tau} = -\eta \frac{\delta p}{p} \quad (10)$$

where η is known as the phase slip factor. The phase slip factor is derived from the transverse optics of the accelerator and takes into account both the change in speed of the particle and the change in path length due to its greater momentum. The phase slip factor depends on the energy of the beam and is negative for energies above transition.

Now,

$$\Delta \varphi = \Delta(\delta \varphi) \approx \frac{d(\delta \varphi)}{dt} \tau = \omega_{rf} \delta t \quad (11)$$

where t is the arrival time of the particle. Then

$$\frac{d(\delta \varphi)}{dt} = \omega_{rf} \frac{\delta t}{\tau} \quad (12)$$

Substituting from eqs. 10 and 9,

$$\frac{d(\delta \varphi)}{dt} = -\frac{\omega_{rf} \eta}{\beta^2} \frac{\delta E}{E} \quad (13)$$

Differentiating,

$$\frac{d^2(\delta \varphi)}{dt^2} = -\frac{\omega_{rf} \eta}{\beta^2 E} \frac{d(\delta E)}{dt} \quad (14)$$

Now substituting from eq. (8),

$$\frac{d^2(\delta \varphi)}{dt^2} = -\frac{\omega_{rf} \eta q V}{\beta^2 E 2\pi} \omega_{rev} (\sin \varphi - \sin \varphi_s) \quad (15)$$

From eq. (1),

$$\frac{d^2(\delta\varphi)}{dt^2} = -\frac{h\omega_{rev}^2\eta qV}{\beta^2 E 2\pi}(\sin\varphi - \sin\varphi_s) \quad (16)$$

Recalling from eq. (5) that $\varphi = \varphi_s + \delta\varphi$, a trigonometric identity is used to write

$$\sin\varphi = \sin(\varphi_s + \delta\varphi) = \sin\varphi_s \cos\delta\varphi + \cos\varphi_s \sin\delta\varphi \quad (17)$$

If the deviation of the particle from the synchronous phase is sufficiently small, the small angle approximations $\sin\delta\varphi = \delta\varphi$ and $\cos\delta\varphi = 1$ may be applied. Then eq. (16) can be written as

$$\frac{d^2(\delta\varphi)}{dt^2} + \Omega^2\delta\varphi = 0 \quad (18)$$

where a new parameter Ω has been introduced,

$$\Omega = \omega_{rev} \sqrt{\frac{h\eta qV \cos\varphi_s}{2\pi\beta^2 E}} \quad (19)$$

Equation (18) is the equation of simple harmonic motion, with Ω as the angular frequency of the oscillation. This longitudinal oscillation is known as *synchrotron oscillation* and Ω as the *synchrotron frequency*. It is also common to use the *synchrotron tune*,

$$Q_s = \frac{\Omega}{\omega_{rev}} \quad (20)$$

which is the proportion of a synchrotron oscillation completed in each revolution [12]. Since the RF acceleration is not performed uniformly around the ring but usually concentrated at a single point, it can be seen that the assumption of slow, adiabatic change made in eq. (8) is valid only if $Q_s \ll 1$. From eq. (19), this is true if $qV \ll E$, in other words if the maximum energy that can be transferred to a particle by the RF cavities in one pass is much less than the energy it already has. This is certainly true in all but the fastest-cycling synchrotrons. For the LHC, Q_s varies from 0.007 at injection to 0.002 at 7 TeV [13].

1.1.4 The Longitudinal Phase Space

It is conventional to define a longitudinal phase space in which the two coordinates are time (or phase) and momentum (or rate of change of phase). In such a phase space, eq. (18) describes an ellipse centred on the synchronous particle (Figure 3). Any particle which is not exactly synchronous will rotate about such an ellipse, though the size of the ellipse will be different for particles with greater or smaller deviation from the synchronous particle. For a synchrotron operated above transition, such as the LHC, the rotation is clockwise. In reality, the particle does not follow a smooth ellipse but moves in discrete steps, with the angle of each step being the inverse of the synchrotron tune. For particles with a larger deviation from the synchronous phase or momentum, the small angle approximations made in eq. (18) no longer hold, and the particle's path in phase space is not a true ellipse, as shown in figure 4.

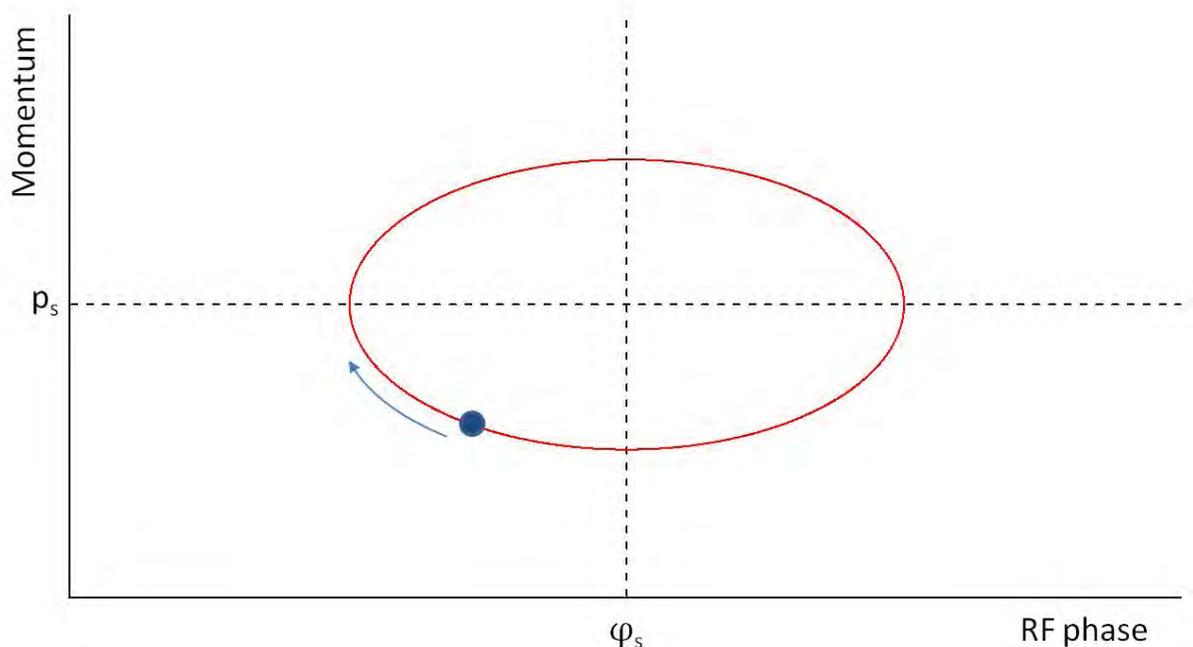


Figure 3. Momentum of a non-synchronous particle plotted against the phase of the RF as the particle passes the centre of the RF cavity. This representation is known as the longitudinal phase space and here shows the synchrotron oscillation of a single particle.

The phase-space region over which longitudinal focusing can keep particles contained is called a *bucket*, and the limit of the bucket in phase-space is known as the *separatrix*. Particles inside the separatrix will follow stable synchrotron oscillations.

If a particle's deviation from the synchronous particle is too large, the longitudinal focusing effect may not be sufficient to keep the particle bound to a single bucket. The particle is then known as *debunched*. Such a particle can still circulate in the accelerator but will not receive the same acceleration as other particles. A debunched particle's mean revolution frequency is different from that of the synchronous particle (a particle captured inside a bucket can instantaneously have a different revolution frequency, but its mean revolution frequency must be the same as that of the synchronous particle) and so it slowly drifts from one bucket to another. If its momentum is sufficiently different from that of the synchronous particle, it spends equal time in and out of phase with the RF, and receives no net acceleration.

The size of the bucket depends on the synchronous phase, as shown in Figure 4. The bucket is largest when the synchronous phase is at the zero-crossing of the RF, that is, when the beam is held at constant energy. This is known as the stationary bucket. Its acceptance stretches over the full RF period. In other words, a particle with momentum exactly equal to that of the synchronous particle will be captured in a bucket no matter what its initial longitudinal position. As the beam is accelerated, the synchronous phase changes and the separatrix becomes smaller. Thus, particles which were inside the bucket at injection can become debunched once the acceleration ramp begins. In the LHC, a 'radiation flash' is expected to start 18 seconds after the beginning of the ramp and last for 1 second as these debunched particles are lost [14].

If the difference in a particle's momentum from that of the synchronous particle is too large, the particle will be lost by hitting the beam aperture in a region with large dispersion. A specific 'momentum collimation' section in point 3 of the LHC is designed to collect such particles safely. The momentum acceptance of the LHC is approximately 6×10^{-4} , meaning that particles with a momentum deviation $\delta p > 6 \times 10^{-4} p_s$ will be lost. During acceleration, debunched particles are therefore lost almost immediately, but while the beam is kept at constant energy they can circulate for some time. Only a slow energy loss due to synchrotron radiation will eventually cause them to be lost. In the LHC at 7 TeV, the mean lifetime of a debunched proton is about 70 s, during which time it can drift through 5820 buckets ($\sim 1/6$ of the ring) [15]. At 450 GeV the lifetime is much longer, since much less synchrotron radiation is emitted, and debunched particles can drift around the whole ring.

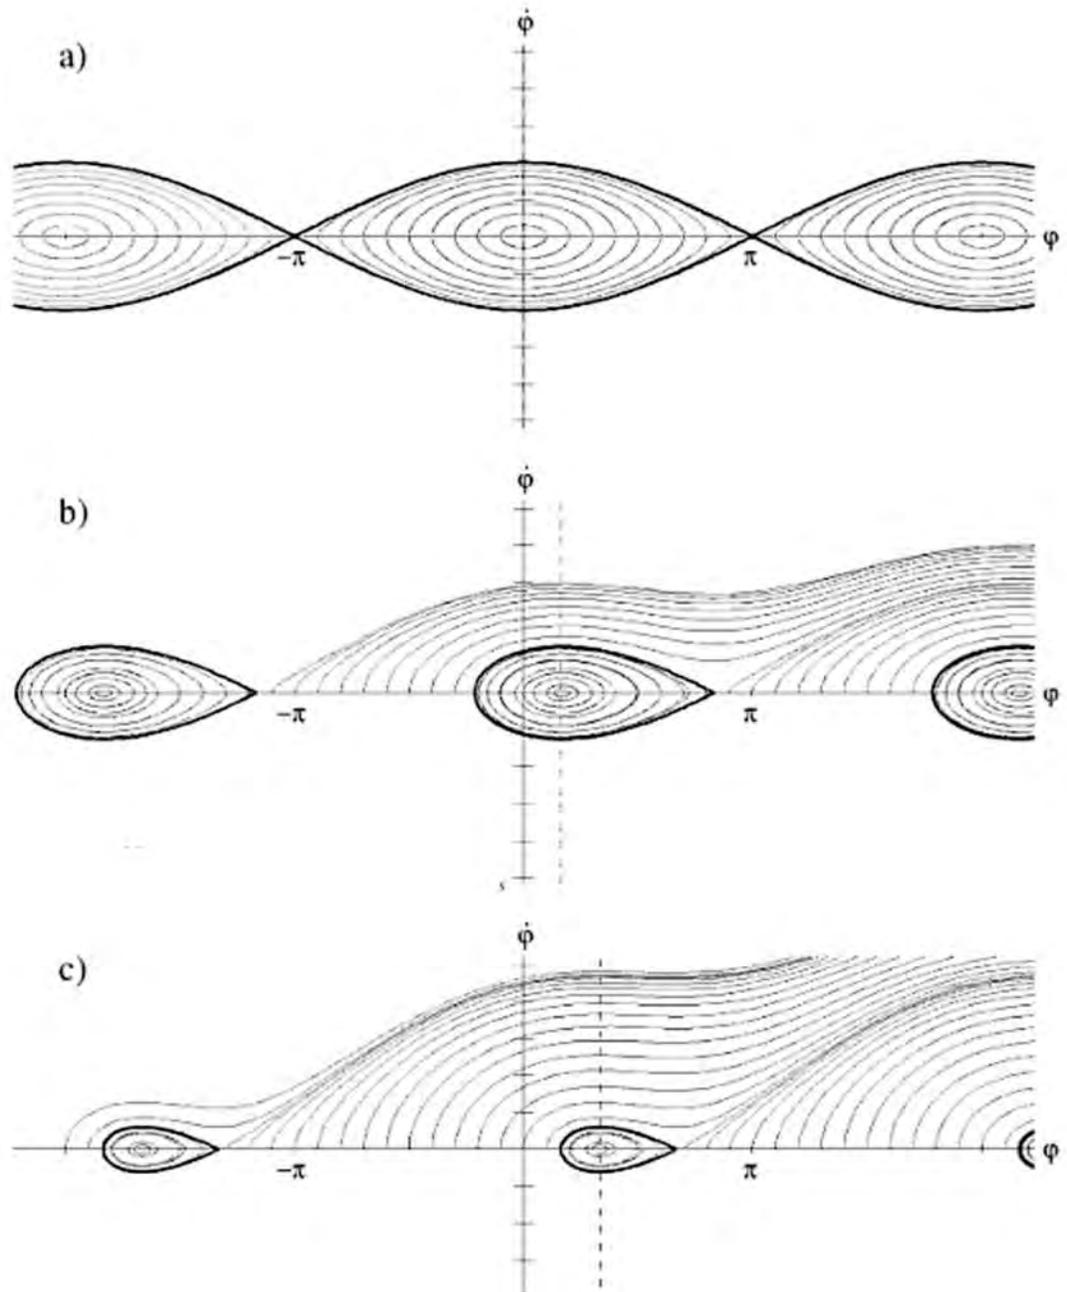


Figure 4. Longitudinal phase-space diagram plotting the rate of change of the particle's phase against the particle's phase. Possible particle tracks in phase space, for synchronous phase at the RF zero-crossing (a), 30° from the zero-crossing (b), and 60° from the zero crossing (c). The bold line marks the bucket separatrix, tracks outside the separatrix represent debunched particles. As the synchronous phase is moved further from the zero-crossing, more acceleration is provided, but the acceptance of the bucket is reduced. Reproduced from [11].

1.1.5 Longitudinal Emittance and Bunch Length

It has been shown above that the maximum area of a bucket in phase space is bounded by the separatrix. However, most of the particles in the bunch occupy an area much closer to the centre of the bucket. The area of the phase-space ellipse containing the bunch is called its *longitudinal emittance*. In reality, some very small proportion of the bunch population will always spread right to the separatrix. It is therefore necessary to define the longitudinal emittance as the area which contains a given fraction of the bunch, and refer for example to the '95% emittance' [16].

The longitudinal emittance is usually viewed as a conserved property. Various bunch manipulations can be carried out which change the phase-space shape of the bunch, for example by altering the RF voltage, but its phase-space area will stay the same [12]. In other words, if the bunch length is to be shortened, it is necessary to increase the momentum spread, and vice versa. However, at sufficiently high energy, synchrotron radiation has a slow damping effect which reduces the longitudinal emittance.

The projection of the phase-space map onto the time (or phase) axis is the longitudinal profile of the bunch. Since a few particles will be found at all longitudinal positions, it is again necessary to refer to a bunch length which contains a given proportion of the bunch population. Generally, the longitudinal bunch profile is approximated by a Gaussian,

$$\lambda(s) = \frac{N_b}{\sigma\sqrt{2\pi}} e^{-s^2/2\sigma^2} \quad (21)$$

where $\lambda(s)$ is the linear proton density, N_b is the number of protons in the bunch, and the bunch centre is taken as $s=0$. The bunch length is then given as a multiple of the Gaussian's σ . Table 2 gives the proportion of the bunch contained within given bunch lengths. The bunch profile is assumed to be symmetric and so the 2σ bunch length contains all the particles within $\pm 1\sigma$ of the peak. It should be noted that the 4σ or 95% bunch length is shorter than the projection onto the time axis of the 95% emittance, since the former includes particles which have a correct position but a large momentum deviation.

Table 2. Proportion of the bunch population contained in a given bunch length, assuming a Gaussian profile.

Bunch length definition	Proportion of the bunch contained
2σ	68%
4σ	95%
6σ	99.5%

In the LHC, a strong longitudinal instability was found to develop when the bunch length drops below 600 ps (4σ). In order to avoid this, a controlled longitudinal emittance blow-up is carried out during the ramp [17]. Parameters relating to longitudinal dynamics in the LHC are shown in table 3.

Table 3. Longitudinal parameters of the LHC.

	450 GeV	3.5 TeV	7 TeV
Synchrotron tune	4.9×10^{-3}	2.5×10^{-3}	2.1×10^{-3}
Bucket Area	1.7 eVs	5.4 eVs	7.62 eVs
Longitudinal emittance (95%)	1.0 eVs	1.6 eVs	1.0 → 2.5 eVs
Bunch length (4σ)	1.8 ns	1.2 ns	0.66 → 1.08 ns
Relative momentum spread	8.6×10^{-4}	5.4×10^{-4}	$1.3 \rightarrow 2.2 \times 10^{-4}$
Synchrotron radiation loss per turn	0.1 eV	440 eV	7 keV
Synchrotron radiation damping time	>20 years	384 hours	24 hours

1.1.6 LHC Filling Schemes

Particles are not injected evenly around the LHC rings. Instead, bunches are injected into specific buckets. Each bucket is given an identification number or BCID. The timing of ‘bucket 1’ is arbitrary, but it is synchronised for the two beams so that bucket n of beam 1 and bucket n of beam 2 will cross at interaction points 1 and 5. The list of buckets which are to be filled is known as the filling scheme.

In order to alleviate the effects of beam loading on the RF cavities and to reduce wake-field instabilities, bunches are spaced at a minimum of 25 ns. Each filled bucket is therefore followed by 9 empty ones. Each group of 10 buckets is then known as a *slot*. A filled slot contains 1 bunch, although the location of the filled bucket within the slot is subject to different definitions. Due to the constraints of the injection and extraction kickers, discussed below, some slots remain entirely empty.

1.1.7 Injector Chain and Effect on Bunch Structure

In order to keep the LHC RF frequency within a very small range and to reduce space-charge effects, particles are not accelerated from rest in the LHC. Rather, the injected particle beams are already highly relativistic. Injection energy is 450 GeV for protons. For heavy ions, the equivalent energy depends on their charge and mass. In 2010 fully stripped lead ions were used ($Z=82$, $A=208$) leading to an injection energy per nucleon of $450 \text{ GeV} \times Z/A = 177 \text{ GeV}$.

A sequence of pre-existing accelerators known as the LHC injector chain is used to accelerate the beam to this energy. The LHC injector chain for protons consists of the Linac 2, Proton Synchrotron Booster (PSB), Proton Synchrotron (PS) and Super Proton Synchrotron (SPS) [18] while for heavy ions Linac 3 and the Low Energy Ion Ring (LEIR) replace Linac 2 and the Booster respectively [19]. At the end of each stage, fast-rising kicker magnets are used to direct the beam to the next accelerator. To prevent any particles from receiving a partial kick, gaps must be left in the filling pattern so that each kicker can fire whilst it is empty of particles. Similarly, a fast kicker is used to redirect the LHC beam to the dump when it is no longer wanted, in case of equipment malfunction or if abnormal beam losses are detected. The gap left for the risetime of the LHC dump kicker is $3 \mu\text{s}$ long and is known as the Abort Gap. Figure 5 shows the injector chain while Figure 6 shows the constraints it imposes on the nominal LHC filling pattern.

The RF period and harmonic number is not equal in the various injectors (Table 4) and so a series of complicated bunch manipulations is necessary in order to pass the bunches up the injector chain without unacceptable losses [20]. The PSB consists of 4 stacked rings. Two of these are cycled twice, so that 6 bunches are injected into the PS [21]. The PS has a harmonic number of 7 at injection, leaving one empty bucket for injection & extraction kickers. A complex process of RF gymnastics [22] splits each bunch into 12 shorter bunches. The full PS ring is then injected up to four times into the SPS to make an SPS batch of up to 288 bunches.

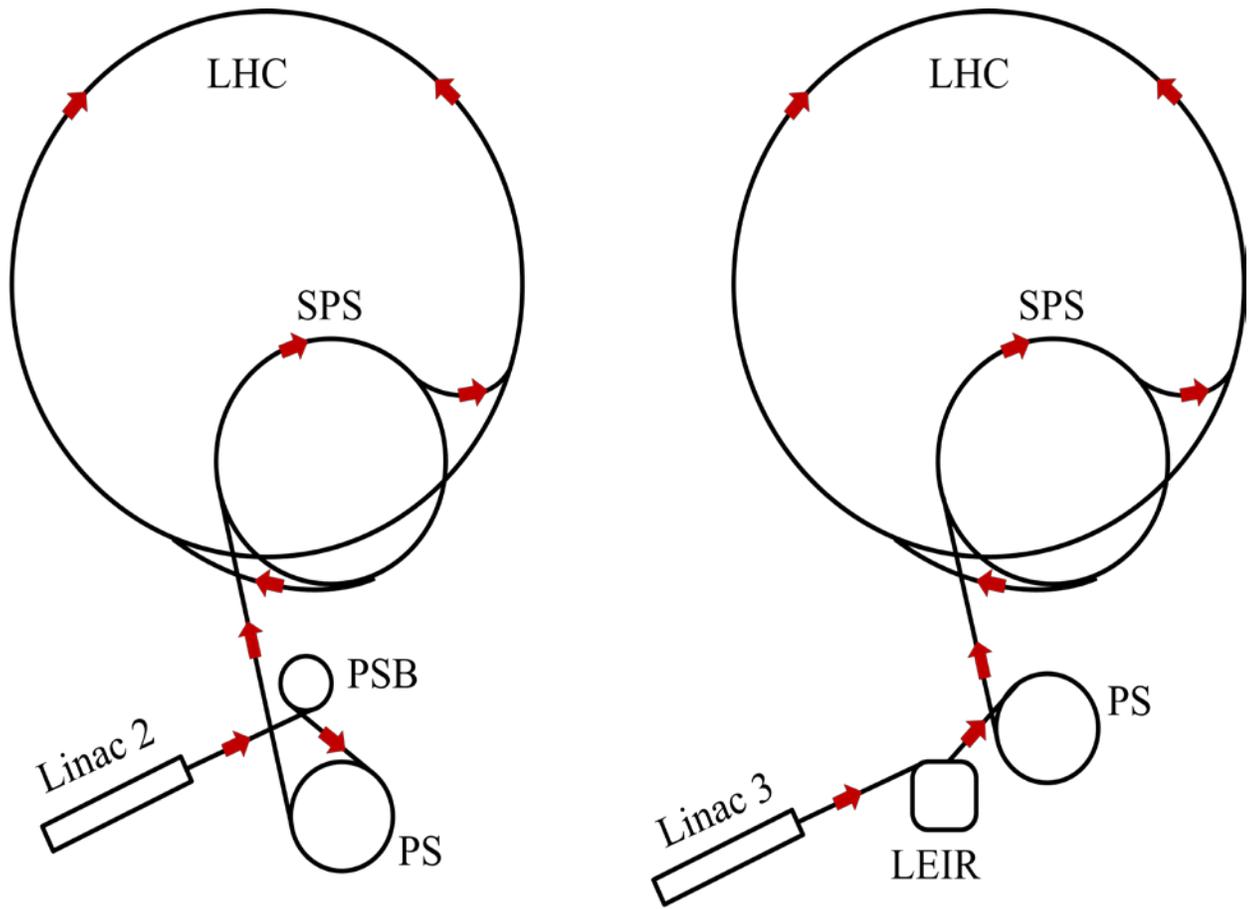


Figure 5. Schematic of the LHC injector chain for protons (left) and heavy ions (right). [1]

Table 4. Relevant parameters of the LHC injector chain for protons.

	PSB (four rings)	PS	SPS
Circumference (m)	157	628	6912
Kinetic energy (GeV)	0.05 → 1.4	1.4 → 26	26 → 450
RF frequency (MHz)	0.6 – 1.7	3.1, 9.3 – 10, 40	200.4
Harmonic number	1	7, 21, 84	4620
Number of bunches	1 (per ring)	6 → 72	144 – 288

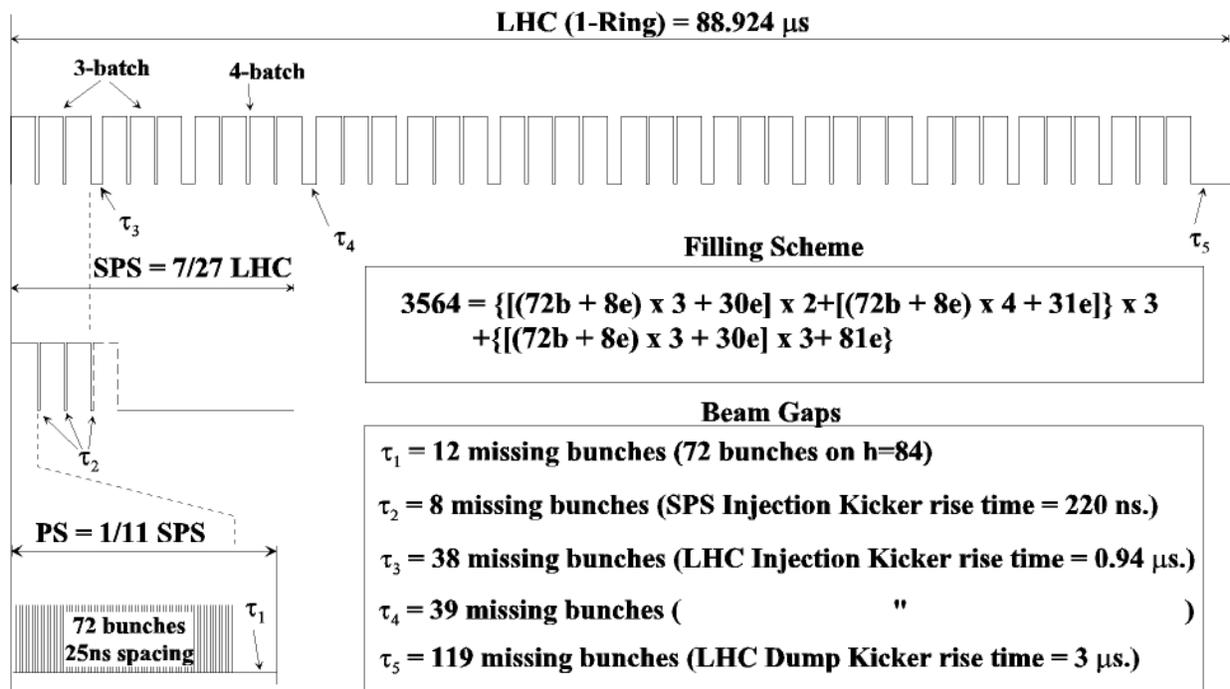


Figure 6. Ultimate LHC filling scheme for protons. Gaps are left to allow for injection and extraction in the LHC's injection chain, as well as to allow for the LHC dump kicker. [23]

1.2 Satellite and Ghost Bunches

It has been shown above that not all buckets in the LHC should contain a bunch. There are 35,640 buckets and a maximum of 2,808 of them should be filled. In practice, however, the nominally empty buckets can contain small numbers of particles.

Each filled slot contains 9 empty buckets and 1 filled one. Any particles occupying any of the 9 nominally empty buckets form *satellite bunches*. Particles in any of the buckets of an empty slot form *ghost bunches*. For the purpose of this definition, the slot is generally taken to be centred on the filled bucket. Thus, particles which are located >5 buckets from a nominal bunch are in ghost bunches, those outside the nominal bunch but ≤5 buckets away are in satellite bunches. Both ghost and satellite bunches are captured in an RF bucket and should not be confused with debunched beam.

The population of the ‘empty’ buckets is of course much less than that of the nominal bunches. One or more satellite bunches in each slot may contain up to 0.5% of the main bunch population. The population of the ghost bunches is generally smaller than that of the main bunches by a factor of at least 10^4 , but since there are far more empty buckets than filled ones, their total population can exceed 1% of the total beam charge.

1.2.1 Formation of Satellite and Ghost Bunches

Satellite and ghost bunches may already be present at injection into the LHC or they may form in the LHC itself. The largest satellite and ghost bunches are generally formed in the pre-injectors and are injected into the LHC along with the nominal bunches. In this case they can only be present inside the injection kicker window of one of the SPS → LHC injections and would not be present between SPS batches or in the LHC abort gap. The spacing of these large satellites and ghosts is an indication of their provenance, since the RF frequency and bunch spacing of each of the pre-injectors is different.

The bunch splitting process in the PS is of particular concern for the formation of satellites and ghosts. The nominal splitting procedure (Figure 7) is designed to produce bunches with 25 ns spacing. In 2010 and 2011, a filling scheme with 50 ns spacing has been mostly used for the LHC. This means that during the final two-way split, the PS must split the RF buckets but keep the bunch contained in only one of the two resulting buckets (‘rebucketing’) [24]. The RF gymnastics which are used to split bunches require the amplitudes of RF cavities operating at different frequencies to be carefully modulated. If there is a mismatch in any one of the RF voltages or in the precise timing of the modulations, particles can be captured in nominally empty buckets.

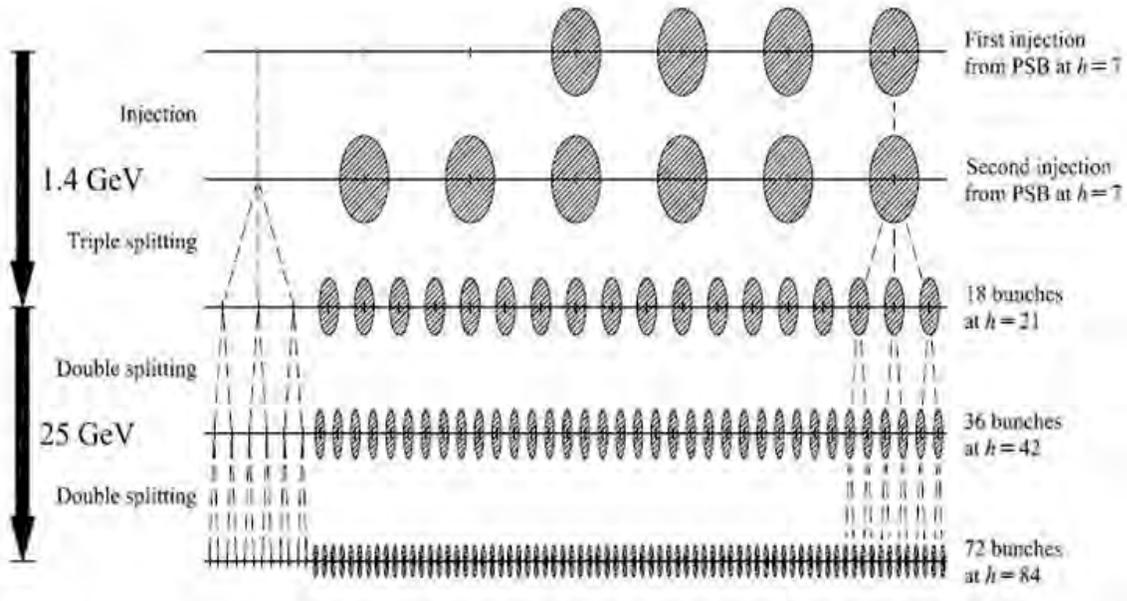


Figure 7. Bunch splitting scheme in the PS. Six long bunches are injected from the PSB and each is divided into 12 LHC-style bunches through three splitting procedures. From [20].

Even after splitting, the bunches in the PS are too long for the 200 MHz buckets in the SPS. Transfer of the beam from the PS to the SPS requires a non-adiabatic bunch compression [25]. The phase of the bunches in the PS must also be locked to that of the SPS RF before extraction. For the 50 ns scheme this is done during the final rebucketing [26].

Beams injected into the LHC are ultra-relativistic ($\beta \approx 1$) and an increase in energy does not lead to a change in speed. At the low energies of the PSB and PS, however, the speed changes considerably during acceleration, and the RF frequency must be tuned in order to keep the harmonic number constant during the ramp. This effect is even more pronounced for heavy ions. A different splitting scheme is employed in the PS for the LHC-type lead ion beam, involving a final rebucketing with a harmonic number of 169 [27] for a final bunch spacing after acceleration in the SPS of 500 ns.

The beam which is extracted from the PS has the correct structure for the LHC filling scheme so no bunch manipulations are necessary in the SPS. However, the bunches coming from the SPS's 200 MHz RF system are not well-matched for the LHC's 400 MHz buckets [28]. This leads to a small longitudinal emittance blow-up due to filamentation. In addition, charge may leak into neighbouring buckets, particularly if there is a voltage or phase mismatch [29]. A phase mismatch leads to a dipole oscillation, in which the bunch centre moves relative to the bucket

centre. A voltage mismatch means that the shape in longitudinal phase-space of the bunch does not match the shape of the bucket. The result is a quadrupole oscillation in which the bunch length fluctuates (Figure 8). This usually results in substantial losses from the bucket, which may form a large satellite on one or both sides of the main bunch, or may form debunched beam.

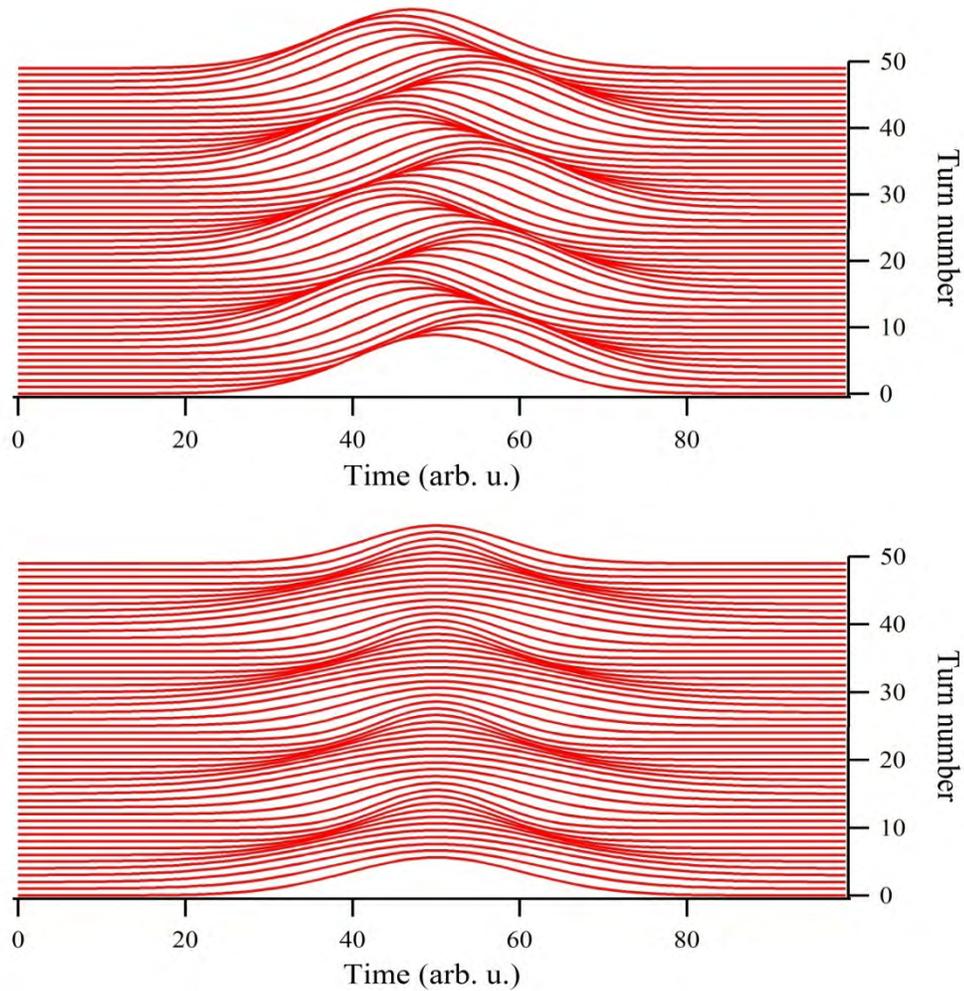


Figure 8. Longitudinal charge density against time, plotted for the same bunch over many turns. Simulated turn-by-turn bunch profiles illustrating mismatched injections. Top: dipole oscillations due to a phase mismatch. Bottom: quadrupole oscillations due to a voltage mismatch.

During operation with heavy ions, an additional mechanism increases the satellite bunch population. During injection of each batch the RF voltage is reduced from 7 MV to 3.5MV in order to improve the longitudinal matching and minimise capture losses. This causes the separatrix to shrink and some particles to leak from bunches already circulating in the machine.

After 3 seconds the RF voltage is increased again and the escaped particles can be recaptured in different buckets [30].

1.2.2 Importance of Satellite and Ghost Bunches

Satellite and ghost bunches are undesirable for a number of reasons. If the level of ghost bunches is sufficiently large, they can cause machine protection issues. Generally, particles in ghost bunches have the same momentum and transverse properties as those in the nominal bunches, and so circulate normally in the machine. However, they can cause heavy losses if they lie within the window of the next injection. In such a case they feel a large transverse force from the injection kicker and are sprayed onto the beam pipe. If the population of ghosts within the injection kicker window is significant, the large instantaneous losses force a beam dump.

Secondly, they cause problems in the calibration of other instruments, principally for the measurement of bunch current and thus luminosity. The luminosity of a pair of colliding bunches of equal transverse size is given by

$$\mathcal{L}_{pair} = f_{rev} \frac{F_g N_1 N_2}{4\pi\sigma_x\sigma_y} \quad (22)$$

where N_1 and N_2 are the populations of the two bunches, $4\pi\sigma_x\sigma_y$ is the effective cross-section of the bunch at the interaction point, and F_g is a geometric reduction factor which takes into account the crossing angle of the two bunches and the rapidly changing bunch size ('hourglass effect') [31]. The total luminosity $\mathcal{L} = \sum \mathcal{L}_{pair}$ is then dependent on the individual bunch population products and cannot be determined from the total populations of the two beams.

Absolute current in the LHC is measured by the DC beam current transformers (DCCT) [32]. The fast beam current transformers (FBCT) [33], which measure the bunch-by-bunch current, must be cross-calibrated with the DCCT and in order to do this an allowance must be made for the fraction of the beam which is stored in the ghost bunches, to which the FBCT is blind [34].

1.3 Objectives and Specifications of the Longitudinal Density Monitor

The LHC Longitudinal Density Monitor is designed to produce a longitudinal profile of the entire LHC ring. Since it must profile main bunches and satellite / ghost bunches at the same time, it must have a very high dynamic range of at least 3×10^4 . The target time resolution is 50 ps. Taking the nominal LHC bunch size of 1.15×10^{11} protons [35] and the minimum 4σ bunch length of 660 ps, the maximum longitudinal density assuming a Gaussian bunch is

$$\lambda_{max} = \frac{N_b}{\sigma\sqrt{2\pi}} \approx 3 \times 10^8 \text{ protons / ps} \quad (23)$$

using eq. (21) for the Gaussian distribution. The number of protons in the peak 50 ps bin is then a little under 1.5×10^{10} . For a dynamic range of 3×10^4 this gives a sensitivity of 5×10^5 protons per bin.

The functional specification [23] divides the task of longitudinally monitoring the LHC beams into three parts: ultra-high sensitivity monitoring of the particle population in the abort gap; high sensitivity, high resolution measurement of the whole ring; and standard sensitivity, fast repetition rate measurement of the main bunch parameters. The Abort Gap Monitor [36], [37] uses a photo-multiplier tube for the ultra-high sensitivity measurement of the abort gap. The measurement of the main bunch parameters (bunch length, bunch shape, relative population) with very short integration time is accomplished by the wall current monitor [38]. The LDM is then tasked with the high sensitivity measurement for the characterisation of the main bunch, main bunch tails, satellites and ghost bunches.

Chapter Summary

In this chapter the layout of the LHC and its injector complex was briefly outlined, and beam parameters relevant to the LDM were given. The concept of longitudinal focusing in RF cavities was explained, along with the way in which it splits the longitudinal phase space into buckets, which constrain the longitudinal position and momentum of particle bunches. This allowed the definition of *satellite* and *ghost* bunches, the measurement of which is a key motivation for the development of the LDM. Lastly, the specifications of the LDM were shown.

2. Background

Overview of Different Longitudinal Profile Methods

Many different ways of measuring the longitudinal properties of an accelerated particle beam have been invented. They can be roughly divided into those that detect the particle beam itself, those that detect its electric or magnetic field, and those that detect the radiation it produces. All have advantages and disadvantages, and the choice of an appropriate method must be made based on the parameters of the beam to be measured and the purpose of the measurements [39].

2.1 Interceptive

Some measurement methods involve directly intercepting the beam. Faraday cups [40] are usually used for current measurement but can give time-resolved measurements if they are connected to a fast oscilloscope. Time resolution is usually in the hundreds of picoseconds [41]. Better sensitivity is obtained using silicon detectors [42], which can have time resolution down to 150ps [43], but their limited radiation tolerance means their use is limited to low-intensity beams. Diamond detectors [44] are much more radiation hard and can have excellent time resolution, below 50ps [45]. Both silicon and diamond detectors are based on ionisation caused by a particle passing through the detector. Except for the lowest energies, the particle is not stopped in the detector and so measurement of the beam current must be based on assumptions on the beam energy.

Interceptive methods clearly cannot be used for on-line measurements of a circular accelerator as they destroy or heavily perturb the beam passing through them, such that the beams could no longer circulate.

2.2 Electro-magnetic

2.2.1 Beam Current Transformers

A beam current transformer (BCT) consists of a toroid of magnetic material placed around a non-conducting ceramic section of the beam pipe. This toroid couples to the magnetic field of the beam, which then acts as the primary of the transformer. A secondary winding is applied evenly around the toroid, and the current induced in this wire can be considered directly proportional to the beam current [46]. The ceramic gap prevents the image current from

passing through the BCT aperture and cancelling the magnetic field of the beam. An alternative path is arranged to allow the image current to pass outside the BCT.

This method is not sensitive to the DC component of the beam current. Such a transformer is therefore useful only with bunched beams. If the secondary coil is read out using a suitably fast Analogue to Digital Converter (ADC), the currents of individual bunches can be distinguished and this is known as a fast beam current transformer (FBCT). However, the time resolution is still limited by the impedance of the coils and reaches at best a few tens of ns [47], making it unsuitable for the measurement of ghost and satellite bunches in the LHC, which have a spacing of 2.5 ns.

The BCT principle can be adapted to measure DC current [48]. In this case, two toroids are needed, each with two separate sets of secondary windings. One set of these windings is attached to the output amplifier. The other is excited by a sine wave modulation with amplitude sufficient to drive the magnetic toroid into saturation at the peaks of the wave. The excitation windings of the two toroids are connected in series so that the current they carry is exactly the same, but are wound in opposite directions. The output windings are also connected in series, so that the resulting current from the two toroids, being excited in opposite directions, is zero. If, however, an additional current, the beam, passes through the toroids, one of the two will saturate for a longer period than the other, and this asymmetry will result in a current being induced in the output windings. This system is known as the DC current transformer (DCCT or DC-BCT).

In order to improve the linearity and dynamic range of the DCCT, it is usually operated as a zero-flux current transformer. A wire passes through the aperture of both toroids and carries a compensation current which cancels the effect of the beam. This compensation current is automatically adjusted so that no current is measured on the output windings, the size of the compensation current is then a direct measure of the beam current (Figure 9). A DCCT in this configuration is often the only instrument which allows direct current measurement, and is then used to calibrate the FBCT and all other current-measurement instruments. The current sensitivity is typically 1-2 μA , equivalent to one tenth of a pilot bunch in the LHC [48]. When a larger current is circulating the accuracy of the current measurement can be as low as 0.2% [49].

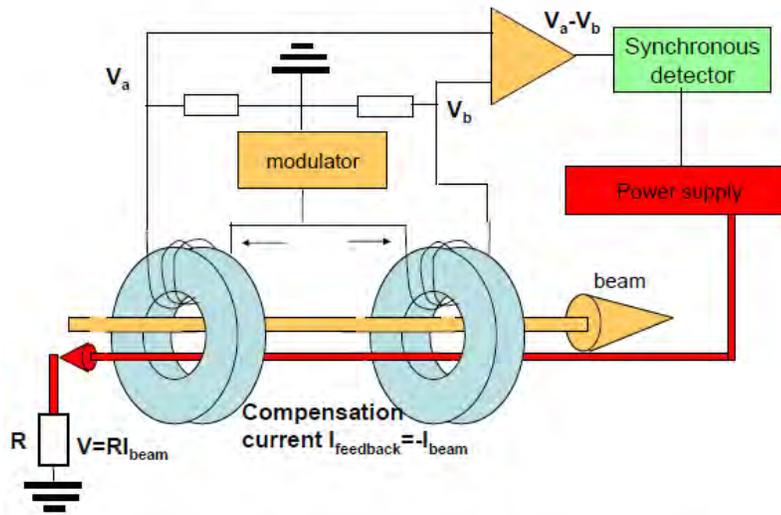


Figure 9. Schematic of a DC Current Transformer operating in zero-flux mode. From [50].

2.2.2 Wall Current Monitor

As the beam passes through the beam pipe, its electric field drags an equal but opposite ‘image current’ along the beam pipe. Wall Current Monitors (WCM) also use a non-conducting ceramic gap in the beam pipe to force this image current to find a new path. This is provided by a set of identical resistors evenly spaced around the gap. The voltage across these resistors is summed, in order to avoid a dependence on beam position, and is measured by a fast sampling oscilloscope [51]. The WCM can measure the beam current with a bandwidth of kHz to a few GHz. It can therefore be used for longitudinal profile measurement as well as bunch-by-bunch current measurement.

2.2.3 Electro-optic

Electro-optic (EO) techniques have been used to measure ultra-short electron bunches such as those used in FELs [52]. Time resolutions of a few tens of femtoseconds have been achieved [53]. The method uses a non-linear crystal such as gallium phosphide, which has the property of rotating the polarisation of light passing through it by an amount which is proportional to the electric field in the crystal. The crystal is placed inside the beam pipe where it sees the electric field of the beam. Crossed polarisers are fixed either side of the crystal, so that in the absence of any EO effect no light is transmitted. When an electric field is applied, however, the polarisation of the light is rotated between the two polarisers and some light is transmitted.

The intensity profile of light transmitted through the crystal is therefore the same as the bunch profile.

A spectral encoding technique is usually used to measure the intensity profile of the light with sufficiently high temporal resolution. A laser pulse is generated with a wavelength chirp, that is, with wavelength changing with time through the pulse. The laser can be located outside the beam pipe, and the pulse is passed through the crystal using thin windows and mirrors. The pulse is synchronised to pass through the crystal at the same time as the bunch. Because of the wavelength chirp, the bunch profile is then encoded in the spectrum of the light as well as in its time profile. A spectrometer is used to read out the profile.

2.3 Optical

Under certain circumstances, an accelerated charged particle can emit electro-magnetic (EM) radiation. This radiation, sometimes in the form of visible light, can be used to measure the properties of the emitting particle beam.

Some of the mechanisms which can be used to generate EM radiation are discussed below. Methods of profile measurement using this EM radiation are discussed in section 2.4.

2.3.1 Synchrotron Radiation

Synchrotron radiation (SR) is emitted when a charged particle moves on a curved path. It is therefore emitted in all circular accelerators, hence the name, but the radiated power is negligible unless the particles are ultra-relativistic.

Synchrotron radiation is an excellent source for optical diagnostics since it is non-destructive to the beam. It is emitted when the charged particle beam travels through the dipole magnets which are already present in a circular accelerator, and requires no further perturbation of the beam. Specific 'insertion' magnets may be added to enhance synchrotron radiation emission or change its spectrum, but if these are well designed their effect on the accelerator lattice will be negligible.

For this reason, synchrotron radiation has been chosen as the source for the Longitudinal Density Monitor, and is discussed in detail in section 2.5.

2.3.2 Optical Transition Radiation

Optical Transition Radiation (OTR) is emitted when a charged particle passes from one medium to another with a different dielectric constant [54], for example from vacuum into a metal foil. The intensity of the emitted radiation is proportional to the relativistic γ of the particle [55]. OTR is emitted in two discrete beams: forward OTR emitted as a hollow cone with opening angle $1/\gamma$ centred on the direction of motion of the particle, and backwards OTR as the specular reflection of the forwards OTR from the foil surface. The ‘hollow cone’ pattern of OTR opens the possibility of measuring very small beam sizes by comparing the intensity of the maximum to that of the central minimum [56].

Arrangements of multiple foils separated by inert spacers can amplify the OTR intensity by constructive interference. In addition, observation of the interference pattern of such a device can be used to extract the beam energy and angular divergence [57].

The foils used for OTR emission can be very thin, so OTR can be considered a non-destructive diagnostic method in single-pass machines i.e. Linacs and transfer lines. In a circular machine, however, the repeated passage through the OTR foil would cause an unacceptable emittance blow-up in the particle beam.

2.3.3 Optical Diffraction Radiation

Optical Diffraction Radiation (ODR) [58] is similar to OTR except that the particle is not required to traverse the foil. This gives it a great advantage over OTR in that it is non-interceptive and has a minimal effect on the particle beam. Diffraction radiation will be emitted if a relativistic charged particle passes within a distance h of a target,

$$h < \frac{\gamma\lambda}{2\pi} \quad (24)$$

where γ is the relativistic gamma of the particle and λ is the wavelength of the ODR observed. The direction and pattern of emission is similar to that of OTR [59].

When ODR is used for diagnostics, two slits or edges are usually used. The first blocks out upstream SR as well as producing ODR. If the two slits are located closer than the radiation formation length, interference between the forwards ODR of the first slit and the backwards ODR of the second can be observed [60] and used to infer the beam size and other beam properties.

2.3.4 Cherenkov Radiation

Cherenkov radiation [61] is emitted when a charged particle travels through a dielectric medium faster than the phase speed of light in that medium. The result is a wave of light which propagates forwards in a cone, similar to the acoustic shockwave caused by a supersonic object. The opening angle of the cone is proportional to the factor by which the speed of light is exceeded. The spectrum of Cherenkov radiation is given by the Frank-Tamm formula [62] and is roughly proportional to frequency. Thus, Cherenkov radiation appears blue-violet to the human eye. In reality the spectrum is continuous, but rising with frequency, and is usually strongest in the UV. Conservation of energy is respected because above some cut-off frequency, usually in the X-ray, the refractive index of the material drops below 1 and thus the phase speed of light cannot be exceeded, and therefore no Cherenkov light is emitted in that part of the spectrum.

Cherenkov radiation is widely used for particle detection and has been proposed for beam loss monitoring [63], but is less useful for beam profile measurement as it is destructive to the beam.

2.3.5 Smith-Purcell Radiation

Cherenkov radiation requires the emitting particle to be travelling faster than the *phase* speed of light in a medium. Unlike the actual propagation speed, the phase speed can be modified by use of a periodic structure. By placing a suitable periodic grating close to the beam, the particle can be made to emit at any arbitrary velocity. This is Smith-Purcell radiation [64].

A conductive grating is placed close to the beam in order to cause emission. The grating has a negligible effect on the beam trajectory, so Smith-Purcell radiation can be an excellent source for non-destructive diagnostics. The wavelength of emission is determined by the period of the grating and by the velocity of the beam, and is usually in the IR and microwave region. In addition, there is a strong dependence of wavelength on the observation angle, so that a Smith-Purcell grating can act as a kind of natural spectrometer.

Although direct observation of Smith-Purcell radiation could be used for diagnostics, the radiation is usually very weak. In addition, the intensity depends strongly on the distance between the gating and the beam, so beam losses on the gating can cause a substantial background. One experiment reports an emission intensity of 10^{-9} photons of Smith-Purcell

radiation per emitting electron, compared to 10^5 photons of OTR for electrons hitting the grating [65].

However, very short (sub-picosecond) particle bunches can emit Smith-Purcell radiation coherently. The radiation intensity is then multiplied by the number of particles in the emitting bunch [66]. Further, knowledge of the spectrum of this coherent radiation can be used to determine the longitudinal bunch profile [67]. However, the reconstruction of the profile from spectral information is not straightforward and without making assumptions about the bunch shape, a given spectrum cannot be shown to have arisen uniquely from a particular bunch profile.

2.3.6 Scintillators

Scintillation is the emission of light from a material which has been ionised by the passage or absorption of an ionising particle. The choice of material is important, since it must produce sufficient light, be transparent to the wavelengths of interest, and be resistant to radiation. Plastic scintillators are often used, since they are easily moulded for coupling into light guides or optical fibres. Doping of the material can be used to improve or adjust its light yield and emission wavelength.

The photon is emitted when the ionised atom recombines and drops to the ground state, not at the moment of ionisation. Thus, there is a variable delay between the particle arrival and the photon emission, known as the scintillator's decay time. The decay is exponential with a half-life of a few ns for plastic scintillators. Decay times under 1 ns have been reported for doped ZnO scintillators [68]. Since the scintillators decay time is the minimum time resolution of any scintillators-based instrument, scintillation is not usually a good source for longitudinal diagnostics.

2.4 Optical Profile Methods

2.4.1 Detector & Scope

The most direct way of measuring the longitudinal profile from some source of optical light is to use a photodetector such as a photo-multiplier tube (PMT) or a photodiode which is connected directly to a high-speed oscilloscope. Provided that sufficient optical power is available, a longitudinal profile can be obtained in a single shot. Diodes and oscilloscopes are

now available with over 20 GHz bandwidth. However, the dynamic range of this kind of measurement is limited, with a maximal signal-to-noise ratio of around 1,000 [69].

2.4.2 Single Photon Counting

Single photon counting (SPC) is a technique which allows a high dynamic range to be achieved even if the available optical intensity is very low. For this reason it has been chosen for the LDM, and is described in detail later.

2.4.3 Streak Camera

Longitudinal profiles with very high time resolution can be achieved using streak cameras [70]. The incident light passes through a slit and hits a photocathode where it releases electrons. The electrons are accelerated and focused by a series of high voltage electrodes, and hit a phosphor screen at the back of the streak tube. An additional pair of electrodes applies a transverse deflection. Applying a sawtooth voltage to these plates sweeps the electrons across the phosphor screen. The longitudinal profile is then displayed by the variation in intensity across the phosphor screen, which can be imaged by a normal CCD camera. The streak camera layout is shown schematically in Figure 10.

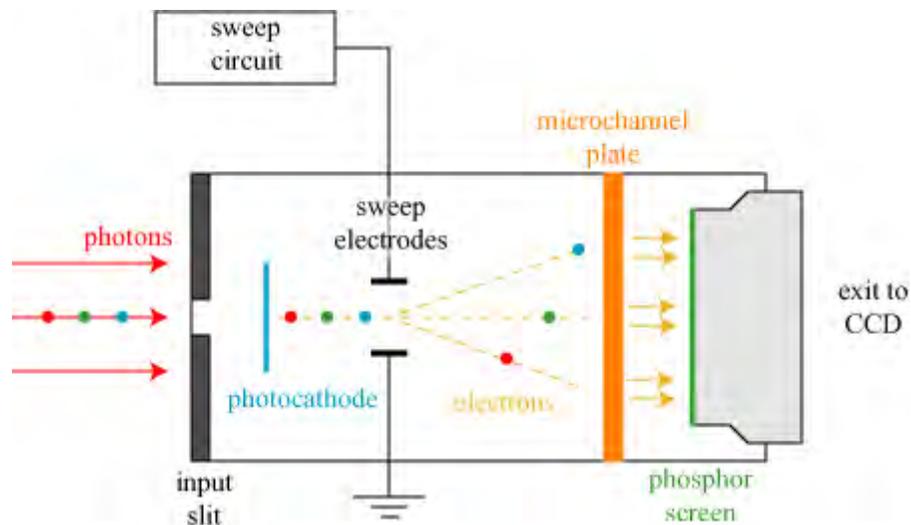


Figure 10. Simple streak camera. Incoming photons are converted to electrons at the photocathode, deflected by the sweep electrodes, intensified by the microchannel plate and then imaged on a phosphor screen. In more advanced streak cameras, the resolution and linearity is improved by the addition of a focusing magnetic field. From [71]

The sensitivity of the streak camera is not usually sufficient for a single-shot measurement so a ‘synchroscan’ mode is used. In this mode, the deflection voltage is carefully synchronised with either the bunch frequency or revolution frequency, allowing the profile of many bunches, or the same bunch over many turns, to be superimposed on the phosphor screen. The signal-to-noise ratio is thus much improved and can reach 10^4 with a long integration time and a bright source [72], but the dynamic range achievable by streak cameras is generally still low compared to other methods.

Another limitation is the record length. For a typical streak camera with a few ps resolution, the whole width of the screen corresponds to around 500ps [73]. The repetition rate is limited by the read-out time of the CCD camera, typically 20ms. The record length can be extended by operating a ‘dual-scan’ mode. A second set of deflecting plates applies a slower sweep in the direction perpendicular to the fast sweep. This allows a number of streaks to be painted across the phosphor screen. Even with this refinement, however, the record length is much shorter than the LHC revolution period of 89 μ s.

2.4.4 Laser Mixing

A longitudinal profile can be constructed by non-linear mixing of SR with light from a pulsed laser [74]. The SR and laser light are coincident on a non-linear crystal with a small crossing angle. Photons with a frequency equal to the sum of the SR and laser frequencies are then generated and travel with an angle between that of the SR and laser. These secondary photons are then detected by an APD or a PMT, which is protected with another monochromatic filter to prevent the signal being swamped by scattered photons from the much brighter primary beams.

By using a laser with a pulse length much shorter than the bunch length, and scanning the phase delay over the bunch using an optical delay line, the longitudinal profile of the bunch can be sampled with high resolution. Assuming that the laser intensity is stable from pulse to pulse, the intensity of secondary light measured is directly proportional to the bunch charge within the sampled slice.

The configuration of the crystal naturally selects SR photons from a narrow wavelength range for conversion. The conversion relies on a non-linear process, in which the number of output

(summed) photons is proportional to the product of the number of laser photons and the number of SR photons [75]. Thus, a powerful pulsed laser and an SR source with high spectral brightness at the chosen wavelength are needed to apply this technique.

An alternative laser-mixing technique applies the concept of spectral decoding as discussed in section 2.2.3 above. The pulse to be measured is passed through a monochromatic band-pass filter and then undergoes degenerate four-wave mixing (FWM) with a chirped laser pulse [76]. The FWM can take place in optical fibres or in specially designed silicon waveguides. The process generates a new component at a frequency given by $f = 2f_{laser} - f_{signal}$ which is known as the ‘idler’. The spectrum of the idler then contains the temporal information of the signal pulse. The idler is separated from the unmixed components using a band-pass filter and measured with a spectrometer.

This is sometimes called a ‘time lens’ and has been demonstrated in the measurement of laser signals, with resolution of a few hundred femtoseconds and record lengths of more than 100 ps [77]. The advantage over the pulsed-laser technique above is that single-shot measurements are possible. The time lens does not appear to have been used in accelerator diagnostics but there seems to be no fundamental obstacle to its application.

2.4.5 Frequency-Resolved Optical Gating

Frequency-Resolved Optical Gating (FROG) is now common in the measurement of ultra-short laser pulses. It is essentially a spectrally-resolved autocorrelation measurement [78] in which the pulse is split into two halves and then combined in some non-linear medium. One half of the pulse has a variable delay which is scanned across the non-delayed pulse to produce the autocorrelation. It is in this way similar to the pulsed-laser mixing technique above, except that the pulse is gated with itself. It is therefore possible to measure arbitrarily short pulses. Measuring the spectrum as well as the intensity of the autocorrelation enables the full information of the original pulse to be reconstructed (intensity vs time and spectrum vs time) unlike in traditional auto-correlation which not only cannot measure the spectral information but also suffers from ambiguities in the retrieval of the intensity trace.

The combination of the two halves of the pulse can involve changing its polarisation through the electronic Kerr effect [79]; self-diffraction in a non-linear medium [80]; or generation of a

second or third harmonic [81] (Figure 11). All of these are non-linear processes and therefore require high pulse powers. All are third-order processes apart from second-harmonic generation (SHG) which is second-order. SHG-FROG therefore produces considerably higher output intensity and has been used to measure pulses down to the pJ level. SHG-FROG does, however, contain more ambiguities than the other FROG methods, notably an ambiguity in the direction of time of the measured pulse. This can be resolved, however, by placing a thin piece of glass in the beamline in order to generate a trailing satellite which indicates the time-direction of the pulse [82].

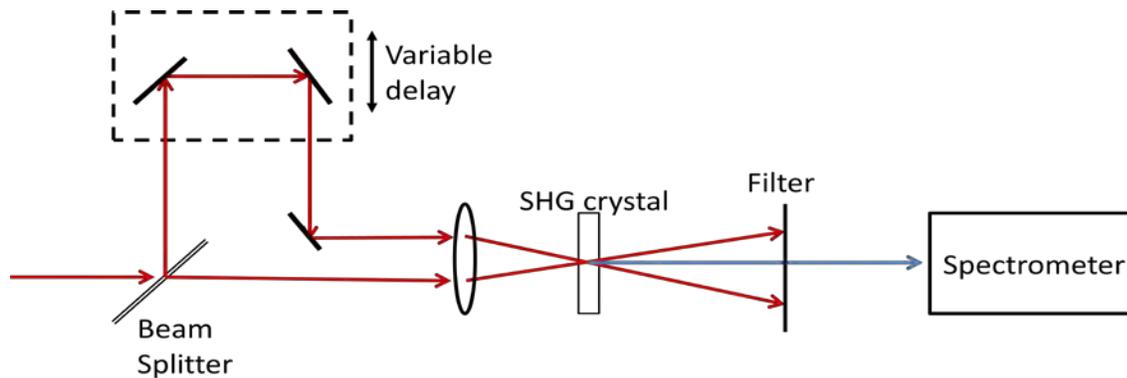


Figure 11. The pulse to be measured is split into two halves, one of which is subject to a variable delay. The Two pulses then co-propagate through a second-harmonic generation (SHG) crystal and a new pulse is generated with much lower power and a frequency twice that of the input pulse. Measurement of the intensity and spectrum of this new pulse as the variable delay is scanned allows reconstruction of the original pulse shape.

All of the above FROG methods involve scanning one copy of the pulse over the other over many repetitions, and therefore require the assumption that the pulse shape is constant. When this is not valid the FROG reconstruction can cause large and unquantifiable errors. Because the pulse is gated with itself, however, FROG is insensitive to arrival time jitter.

Single-shot FROG measurements are also possible by mapping the delay between the two pulse copies onto one of the transverse directions. The perpendicular transverse direction is then used for the spectrometer measurement. This spatial mapping of the delay can be achieved by focusing the pulse into a line and then crossing the two copies with a large angle. In the centre, the two pulses arrive at the same time, but moving away from the centre there is an increasing relative delay between the two (Figure 12).

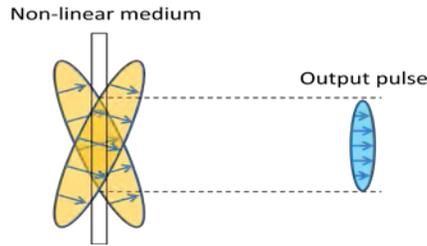


Figure 12. Illustration of a single-shot autocorrelation technique, in which the relative delay between two copies of the same pulse is mapped onto the transverse coordinate. This can be converted into a single-shot FROG measurement by using an imaging spectrometer with the dispersion in the perpendicular transverse direction.

FROG has been used to measure the output light from Free-Electron Lasers (FELs). Due to the high pulse powers required, however, it is not suitable for most optical diagnostics sources. Single-shot FROG has however been used in the measurement of signals from Optical Replica Synthesiser experiments [83], in which a high-power optical pulse with the same profile as the beam is generated by coherent synchrotron radiation.

2.4.6 Coherent Radiation Spectrum

Emission of light by several of the mechanisms described above (SR, OTR, ODR) can become coherent if the bunch length is smaller than the wavelength of light emitted [84]. Measurement of the radiation spectrum can then be used to reconstruct the bunch length. Reconstruction of the longitudinal profile is also possible by using the Kramers-Kronig technique [85] to find the phase of the bunch form factor, although some assumptions about the bunch shape have to be made [86] since the reconstructed profile is not unique. A very wide spectral measurement covering at least 4 octaves is necessary [87], since the coherent part of the radiation is usually contained in the IR or microwave region. Within the spectral range where the coherence condition is met, the coherent component of the radiation is several orders of magnitude stronger than the incoherent part.

2.4.7 Shot Noise Technique

Bunch length can be measured indirectly by looking at the fluctuation in the incoherent synchrotron light in a small bandwidth [88]. Although the spectrum of incoherent synchrotron radiation is typically broad and smooth, this is only observed by averaging over many bunches. By contrast, the single-shot spectrum is noisy due to photon shot-noise [89] and contains many

spikes whose width is inversely proportional to the length of the emitting bunch. Thus, measuring the spectrum would allow the bunch length to be measured. However, it is not easy to measure the synchrotron light spectrum with high accuracy in a single shot.

An equivalent measurement can be made by looking at the intensity of SR in a small bandwidth. Typically, a filter with 1nm bandwidth or less is used. Thus, an intense SR source is needed, so that sufficient power for single-shot measurement remains after passing through the filter. The intensity is measured for some 1000s of shots. The variance of these measurements is then a proxy for the 'spikiness' of the spectrum: a very spiky spectrum will show greater shot-to-shot variance. The variance can then be used to calculate the bunch length.

2.5 Synchrotron Radiation

According to classical electromagnetism, an accelerating charge will emit electromagnetic radiation. The power of the emitted radiation is given by the Larmor formula [90],

$$P = \frac{2}{3} \frac{q^2 \dot{v}^2}{c^3} \quad (25)$$

where q is the charge, \dot{v} its acceleration and c the speed of light. In a particle accelerator, charged particles are subject to two forms of acceleration: forwards acceleration in the RF cavities and transverse acceleration in the magnets. The intensity of EM radiation emitted by both these accelerations is, by equation (25), rather small. However, the Larmor formula is non-relativistic. When the radiating charges are moving at relativistic speeds the Larmor formula gives the radiated power as observed in the moving reference frame of the particle. To calculate the radiation observed in the stationary frame of the accelerator, a Lorentz transformation must be applied between the frame of reference of the particles and that of the accelerator. This has the effect of amplifying the radiation emitted in the direction of travel of the accelerated charge. This forward radiation is emitted when the charge experiences a transverse acceleration. Radiation from the forwards RF acceleration is emitted perpendicularly to the direction of motion and receives no Lorentz amplification. It is therefore negligible compared to the radiation from transverse acceleration, and is ignored from now on.

The amplified forward radiation is known as synchrotron radiation (SR) since it is usually associated with highly relativistic circular accelerators such as synchrotrons.

Synchrotron radiation was first discussed in 1944 by Isaak Pomeranchuk and Dmitri Iwanenko [91], although they did not name it at the time. It was first observed in 1947 by Floyd Haber, a technician at the General Electric synchrotron [92]. It was recognised by Frank Elder, Anatole Gurewitsch, Robert Langmuir, and Herb Pollock to be the same radiation predicted by Iwanenko and Pomeranchuk, and given the name 'synchrotron radiation' when they published the finding [93].

Today synchrotron radiation is much used as a high-intensity source of x-rays for experiments in material science, molecular biology and other fields, and accelerators known as synchrotron light sources are built for the express purpose of emitting it. It plays an important role in the physics of electron accelerators, where the power of synchrotron radiation which is emitted and must be replaced by the RF cavities is so large that it becomes a limiting factor for the maximum achievable energy of electron synchrotrons. It also causes a radiative damping effect which is the dominant factor in the small emittance of electron beams.

2.5.1 Dipoles

In most accelerators, the majority of the transverse acceleration occurs in bending magnets which have only vertical magnetic fields. The force on the particle is then entirely horizontal. In the frame of reference of the emitting particle, it emits as a simple Hertzian dipole. However, when the Lorentz transformation into the laboratory frame is carried out, the radiation emitted in the direction of motion of the particle is amplified. If the particle is ultra-relativistic this amplifying factor can be very large. The radiation emitted in other directions is not affected so the radiation also becomes collimated (see Figure 13). In general, for a very relativistic particle the radiation is contained within a cone of half-angle $1/\gamma$.

The SR power radiated by a particle of charge e being accelerated round a circular path of radius ρ is given by

$$P = \frac{e^2 c}{6\pi\epsilon_0(m_0c^2)^4} \frac{E^4}{\rho^2} \quad (26)$$

where ϵ_0 is the permittivity of free space, (m_0c^2) is the rest mass energy of the particle and E its total energy. Substituting $E = \gamma m_0c^2$ gives

$$P = \frac{e^2 c \gamma^4}{6\pi\epsilon_0 \rho^2} \quad (27)$$

The energy radiated in one complete revolution of period τ is then

$$\Delta E_{turn} = \int_0^\tau P dt = P \frac{2\pi\rho}{c} \quad (28)$$

where $2\pi\rho/c$ is the time taken to travel through the bending magnets. The time taken to traverse any straight sections is ignored since no radiation is emitted there. Substituting eq. (26) into (27),

$$\Delta E_{turn} = \frac{e^2 \gamma^4}{3\epsilon_0 \rho} \quad (29)$$

and the energy emitted in any bending magnet of bending angle α is given by

$$\Delta E_{BM} = \frac{\alpha e^2 \gamma^4}{2\pi 3\epsilon_0 \rho} \quad (30)$$

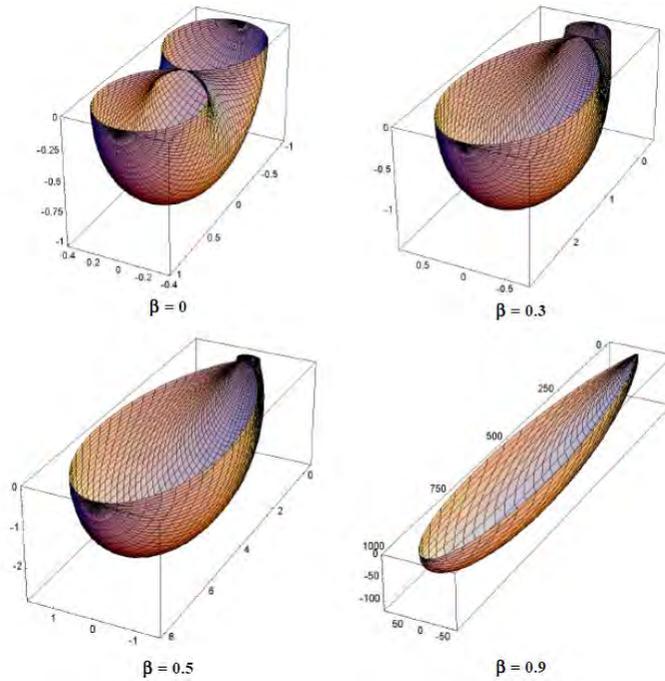


Figure 13. Distribution of synchrotron radiation emission for particles travelling at various velocities βc . From [12].

SR emitted by particles traversing a dipole field has a broad spectrum, which is again determined by the relativistic γ of the particles and by their bending radius ρ [94]. It is normal to characterise the radiation spectrum by its critical frequency,

$$\omega_c = \frac{3c\gamma^3}{2\rho} \quad (31)$$

The critical frequency divides the spectrum into two halves of equal power, that is, equal power is radiated above and below the critical frequency. The shape of the SR spectrum was derived by Schwinger [95]. He found that the number of photons within a spectral slice is

$$\frac{d\dot{N}}{d\varepsilon/\varepsilon} = \frac{P}{\omega_c \hbar} S\left(\frac{\omega}{\omega_c}\right) \quad (32)$$

where P is the radiated power defined in equation (27) and the spectral function S is given by

$$S(\xi) = \frac{9\sqrt{3}}{8\pi} \xi \int_{\xi}^{\infty} K_{\frac{5}{3}}(\xi) d\xi \quad (33)$$

where $K_{5/3}(\xi)$ is the modified Bessel function and $\xi = \omega/\omega_c$. An example spectrum is shown in Figure 14. The Schwinger spectrum is in fact an approximation of the real SR spectrum, and it is valid only with the assumptions that the dipole magnet is very long and that the observer is very far away. In particular, the long-wavelength part of the spectrum is not always accurately described [96].

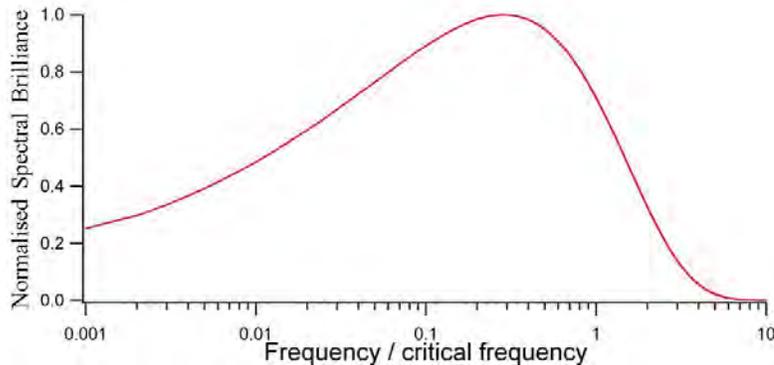


Figure 14. Spectral brilliance of SR against frequency. The spectrum of synchrotron radiation from a charged particle traversing a dipole magnet, normalised to the critical frequency, always has the same form. The critical frequency depends on the particle's relativistic γ and on the bending radius of the magnet.

2.5.2 Edge Radiation

The magnetic field of an ideal dipole magnet would be constant everywhere inside the magnet and zero outside. In practice, however, the field is not sharply contained inside the magnet. For a particle moving along the longitudinal axis of the magnet (i.e. the beam orbit) the field strength changes from zero to the nominal value over some length, generally of the order of a few mm. The length and shape of this 'edge field' varies depending on the magnet construction.

The spectrum of emitted SR can be viewed as the Fourier transform of the particle's motion in the laboratory reference frame [97]. If the time taken to traverse the magnetic field is very short, then the spectrum will be proportionately broad. A charged particle crossing the edge field will emit SR with a critical frequency given by eq. (31), as in the body of the dipole, but taking the value of the radius of curvature at that moment. Since the edge fields are crossed very quickly, however, the spectrum will be much broader than that shown in Figure 14, and indeed may be considered almost flat for wavelengths much higher than the critical frequency [98]. Although the total power of the edge radiation is negligible compared to the SR from the body of the dipole, the spectral brightness of edge radiation can be higher in some wavelength regions far from the critical wavelength.

2.5.3 Undulators and Wigglers

In addition to the dipoles which are necessarily present in a synchrotron, SR may be generated by specific 'insertion devices'. These generally take the form of an undulator, a sequence of short dipole fields of alternating polarity. The beam is deflected in alternate directions to give an undulating trajectory. If the integrated field strength in the two directions is equal then the total bending angle will be zero. If in addition the undulator is arranged so that the field sections at each end are of half the length of the others, the beam will exit the undulator with no transverse displacement. The undulator is then described as matched (Figure 15). A matched undulator can be inserted into the beamline without affecting the beam orbit.

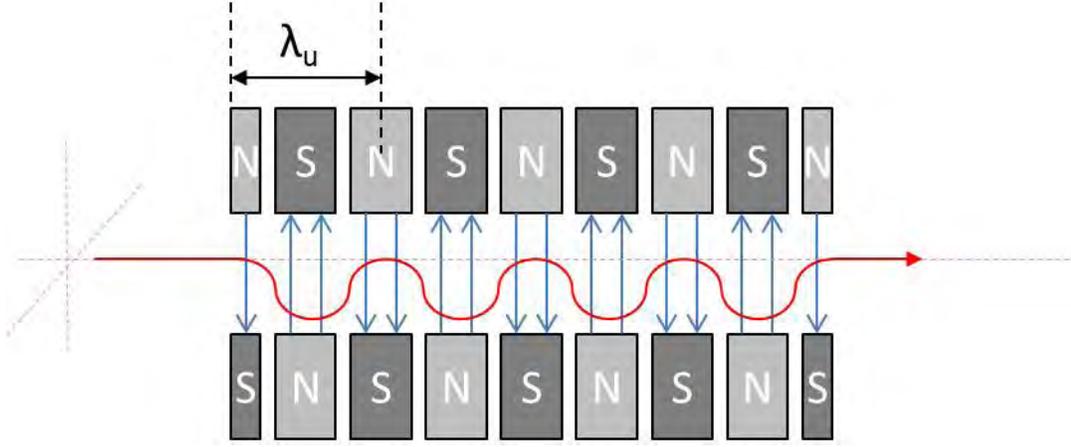


Figure 15. A matched undulator, which can be inserted into the beam without affecting the orbit.

The magnetic field along the beam orbit is generally sinusoidal, $B(s) = \tilde{B} \sin(k_u s)$, where $k_u = 2\pi/\lambda_u$. The dimensionless undulator parameter K is then defined as

$$K = \frac{\lambda_u e \tilde{B}}{2\pi m_0 c} \quad (34)$$

K is a measure of the strength of the undulator, more specifically of the deflection it causes. The maximum angle between the beam trajectory with and without the undulator in place is

$$\theta = \frac{K}{\gamma} \quad (35)$$

where γ is the relativistic gamma or Lorentz factor of the beam [99]. Since the opening angle of emitted SR is $1/\gamma$, an observer on-axis downstream of the undulator would see the undulator as a long continuous source provided that the trajectory deviation is always less than $1/\gamma$. From eq. (35), this requires $K \leq 1$. If $K > 1$ then $\theta > 1/\gamma$ and an on-axis observer would see a series of discrete sources. Such an undulator with $K > 1$ is called a Wiggler.

The periodic nature of the undulator gives rise to interference effects in the SR produced. The SR emitted by successive undulator periods will interfere constructively if

$$\lambda_{coh} = \frac{\lambda_u}{2\gamma^2} \left(1 + \frac{K^2}{2} + \gamma^2 \theta_{obs}^2 \right) \quad (36)$$

where Θ_{obs} is the angle at which the undulator radiation is observed. Eq. (36) is known as the coherence condition for undulator radiation. In an undulator with many periods, the interference effects dominate and near-monochromatic SR will be emitted. In addition to the wavelength given by eq. (36), SR is emitted at the harmonics $\lambda = \lambda_{\text{coh}}/n$. For an undulator, the intensity of the harmonics decreases with n while a wiggler may radiate preferentially at higher harmonics [100].

In an undulator with few periods the SR is not monochromatic but has a broad spectrum centred at λ_{coh} . The width of the spectral peak is inversely proportional to the number of undulator periods N .

The power emitted in the forward cone by a particle traversing an undulator [101] is given by

$$P = \frac{\pi e^2 c \gamma^2}{\epsilon_0 \lambda_u^2 N} \frac{K^2}{(1 + K^2/2)^2} \quad (37)$$

Strictly, this formula is valid only for a weak undulator, $K \ll 1$. An additional correction factor should be introduced to make the equation valid for any value of K [102]. Since the power is proportional to γ^2 , compared to γ^4 for SR from a dipole, undulators are particularly useful for enhancing SR emission at lower beam energies.

2.5.4 SR from Protons and Ions

From equation (26) it can be seen that, for a given particle energy and bend radius, the radiated power is proportional to $(1/m_0)^4$. Since $m_{\text{proton}}/m_{\text{electron}} = 1836$, a proton will emit more than 10^{13} times less synchrotron light than an electron with the same energy and trajectory!

For this reason, synchrotron light sources always accelerate electrons (or positrons). On the other hand, the energy lost to SR limits the use of circular electron accelerators in high energy physics. The LEP accelerator at CERN was probably the highest-energy electron synchrotron that will ever be built. It reached electron energies of 104.5 GeV, a figure largely limited by the inability of the RF cavities to supply any more than the 3.3 GeV radiated per turn by each electron. By contrast, the LHC is designed to reach 7000 GeV in the same tunnel, and will then radiate 6.7×10^{-6} GeV per proton per turn.

A fully stripped ion emits SR as a single charged particle, since all the charges are contained within a length much shorter than the wavelength of the SR. Then equation (26) is modified as

$$P = \frac{Z^2 e^2 c}{6\pi\epsilon_0(m_0c^2)^4} \frac{E^4}{\rho^2} \quad (38)$$

where Z is the charge number of the ion. An alternative description is that each proton within the ion is radiating like an individual proton. The neutrons, having no charge, do not emit SR. However, since all the protons within an ion are radiating within a distance much shorter than the SR wavelength, their radiation is coherent. That is why the radiation intensity scales with Z^2 rather than Z.

2.5.5 Time Structure of SR

Since the synchrotron light is emitted over an extended source area, the longitudinal distribution of the light is spread out compared to that of the emitting bunch. In order to quantify the time resolution of the synchrotron light compared to the emitting particles, two effects must be taken into account: the difference in speed of the particles and photons and the difference in path length.

SR is emitted with an opening angle of $1/\gamma$ and this determines the length of the extended source in a dipole magnet. As can be seen in Figure 16, photons emitted at A or B would both reach the detector from the same direction and would be indistinguishable. In practice, the finite acceptance of the detector would cause photons to be detected from an even greater source length. Considering only the theoretical minimum source length AB:

$$Path_{photons} = AB = \sqrt{2\rho^2 - 2\rho^2 \cos(2/\gamma)} \quad (39)$$

while

$$Path_{particles} = Arc_{AB} = \frac{2\rho}{\gamma} \quad (40)$$

The difference in speed must also be taken into account: photons travel at speed c while the particles travel with speed βc , where $\beta = \sqrt{1 - 1/\gamma^2}$

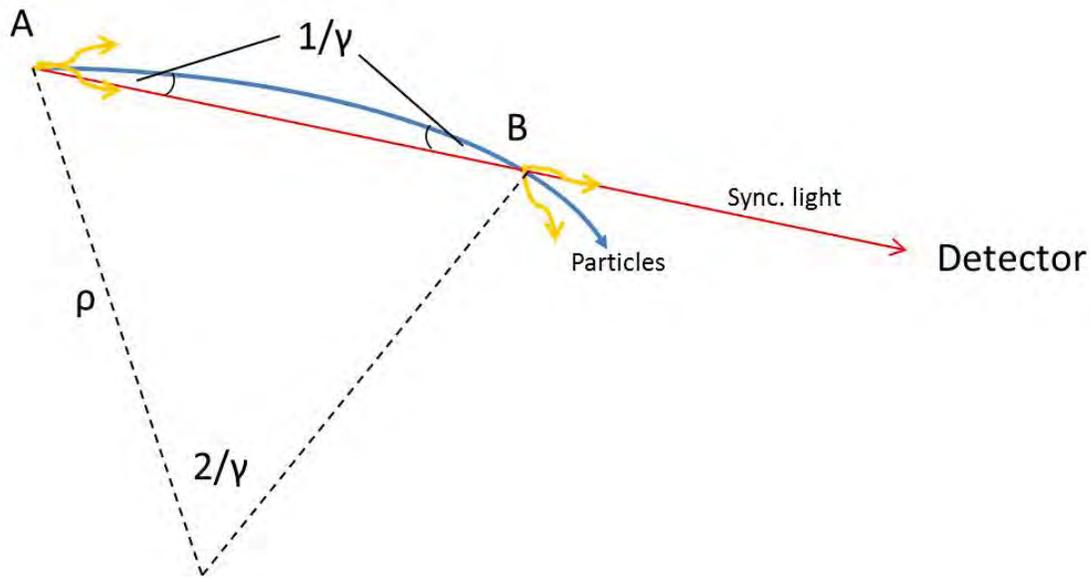


Figure 16. Diagram illustrating the minimum source length for synchrotron light in a dipole and the path length difference it implies (after [99]). A photon emitted at A will arrive at the detector before a photon emitted at B by the same particle.

Then the time to travel from A to B is

$$t_{photons} = \frac{\sqrt{2\rho^2 - 2\rho^2 \cos(2/\gamma)}}{c} \quad (41)$$

and

$$t_{particles} = \frac{2\rho/\gamma}{c\sqrt{1-1/\gamma^2}} \quad (42)$$

Then the time-of-flight difference is given by

$$\Delta t = \frac{\rho}{c} \left\{ \frac{2}{\gamma\sqrt{1-1/\gamma^2}} - \sqrt{2-2\cos(2/\gamma)} \right\} \quad (43)$$

For a strongly relativistic particle with large γ , the small angle approximation $\cos \theta = 1 - \theta^2/2$ can be applied. Eq. (42) then simplifies to

$$\Delta t = \frac{2\rho}{c\gamma} \left(\frac{1}{\beta} - 1 \right) \quad (44)$$

For almost all systems this is an extremely small number and so the SR can be taken to have the same time distribution as the emitting particles.

Similarly, for a particle traversing an undulator or wiggler, the longitudinal velocity v_s can be defined by simple application of Pythagoras' theorem as

$$v_s^2 = (\beta c)^2 + v_x^2 \quad (45)$$

where v_x is the transverse velocity due to the field of the undulator, given by

$$v_x = \beta c \frac{K}{\gamma} \sin(\omega_u t) \quad (46)$$

Averaging over a full period of the undulator,

$$\overline{v_x^2} = \frac{1}{2} \left(\frac{\beta c K}{\gamma} \right)^2 \quad (47)$$

then

$$\overline{v_s} = \beta c \sqrt{1 + \frac{K^2}{2\gamma^2}} \quad (48)$$

A detector on the axis of the undulator cannot distinguish photons emitted from the beginning or end of the undulator. For an undulator of length L_u these would have an arrival time difference of

$$\Delta t = \frac{L_u}{c} - \frac{L_u}{\overline{v_s}} \quad (49)$$

Assuming $\gamma \gg K$ then $\sqrt{1+x} = 1 + \frac{1}{2}x$ is a good approximation. Then

$$\Delta t = \frac{L_u}{c} \left(1 - \frac{4\gamma^2}{\beta(4\gamma^2 + K^2)} \right) \quad (50)$$

The value of K is less than 1 for an undulator, so it can be seen that for a strongly relativistic particle with $\gamma \gg 1$, Δt is again very small.

Chapter Summary

This chapter comprises a review of the most commonly-used longitudinal diagnostic techniques. Focusing especially on optical techniques, different mechanisms of light production are outlined, and then the methods by which the emitted light can itself be measured in order to provide information about the emitting beam. Since synchrotron radiation is to be used for the LDM, the theory and nature of synchrotron radiation is presented, including the differences between SR from different types of magnets: bending magnets, edge fields and undulators.

3. Theory & Simulations

3.1 The Beam Synchrotron Radiation Telescopes

The beam synchrotron radiation telescopes (BSRT) are installed at point 4 of the LHC. Each beam has a completely independent BSRT. The central part of point 4 contains the LHC's RF cavities and the two BSRTs are located to either side, approximately 100 m apart (Figure 17).

The LHC's two counter-rotating beams require magnetic fields acting in opposite directions. Each beam has its own beam pipe but the two pipes pass through the same magnet cold mass. Thus, the separation of the two beams is only 194 mm in the arcs [103]. However, the separation must be increased to 420 mm at point 4 so that separate RF cavities can be installed for each beam. This separation is achieved by a set of specially designed magnets in which the magnetic fields for the two beams are parallel, known as MBRS / D3 and MBRB / D4 dipoles. The D3 magnets are the main source of SR for the BSRT.

However, at the LHC's injection energy of 450 GeV, the D3 dipoles do not produce sufficient visible light for diagnostics to be possible. In order to allow SR diagnostics to be used at all LHC energies, a superconducting undulator was installed for each BSRT [104]. The undulator is attached to the upstream side of the D3 cryostat; each undulator affects only one beam. The layout of the magnets is shown in Figure 17. There is a gap of 937 mm between the undulator and the D3 dipole. Some key parameters of the D3 dipoles [105] and the undulators [106] are shown in Table 5 and Table 6. Although the undulator has a nominal peak field of 5 T, requiring 450 A of current, it has so far been operated with a current of 400 A and a corresponding peak field of 4.4 T.

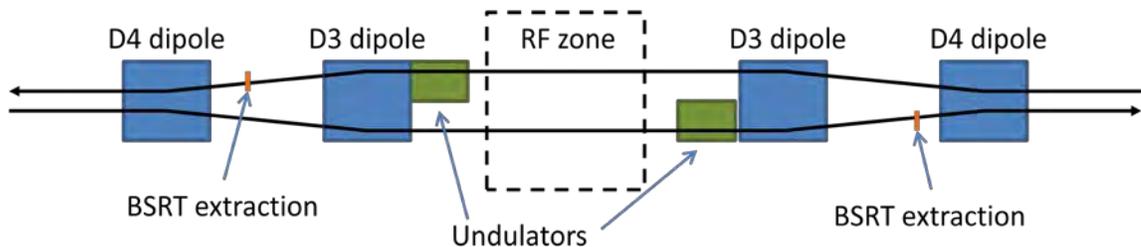


Figure 17. Layout of the separation dipoles in point 4 of the LHC. The undulators and D3 dipoles are the source of synchrotron radiation for the BSRT.

Table 5. Key parameters of the D3 dipoles.

Magnet length	9.45 m
Bending angle	1.58 mrad
Bending radius	6013 m
Nominal field strength (@ 7 TeV)	3.8 T
Nominal current	5750 A
Operating temperature	4.5 K

Table 6. Key parameters of the BSRT undulators.

Period length	280 mm
Number of periods	2
Nominal peak field strength on orbit	5 T
Nominal current	450 A
Operating temperature	4.2 K
Undulator parameter k	0.071

Synchrotron radiation (SR) is emitted in a narrow cone centred on the direction of motion of the emitting particle beam. Thus, light emitted in the undulator and the leading edge of the D3 initially travels along the same line as the proton beam. The proton beam is then deflected by the D3 dipole while the SR continues in a straight line. A small mirror is used to reflect the SR out of the beam pipe and onto the BSRT optical table. In order to prevent particle losses and damage to the mirror caused by beam halo hitting the mirror or its support, it must be located at least 20 mm from the nominal beam orbit, equivalent to 15 times the nominal horizontal beam size [107]. Since the D3 dipole separates the proton beam from the undulator SR by only

1.6 mrad, a long drift is needed before sufficient separation is reached. The extraction mirror is located 27 m downstream of the undulator, and is off-centred by 7 mm with respect to the undulator axis in order to keep it clear of the proton beam.

The extraction mirror is made of baked aluminium on a stainless steel support. Only the visible and near-IR component of the SR spectrum is reflected (see Figure 18). A stainless steel scraper protrudes 1 mm on the side closest to the beam orbit in order to protect the mirror in case of beam halo growth.

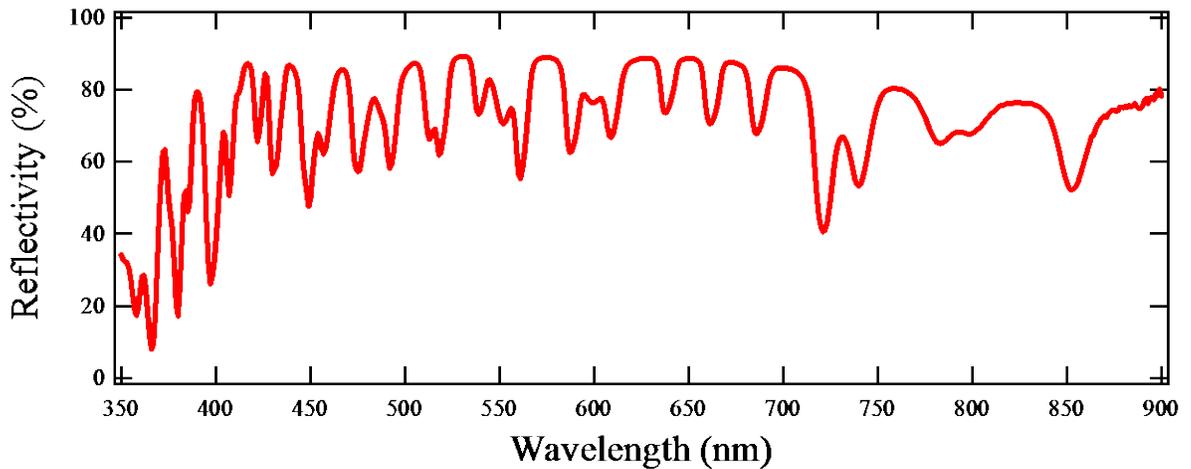


Figure 18. Reflectivity of the extraction mirror against wavelength. The extraction mirror is made of baked aluminium and located inside the beam pipe. From manufacturer's specifications [108].

3.2 Arrangement of the Optical System

The SR incident on the extraction mirror is reflected vertically downwards and exits the beam pipe through a fused silica window. The fused silica is 20mm thick and has excellent transmission across the visible and near-IR spectrum (Figure 19). Since the SR beam passes perpendicularly through the window, there is no dispersion, but the different velocity of light of different wavelengths in the silica leads to a transit time difference given by

$$\Delta t = \frac{n_1 \times 0.02\text{m}}{c} - \frac{n_2 \times 0.02\text{m}}{c} \quad (51)$$

The fused silica has a refractive index varying between 1.48779 and 1.45250 for light between 300 nm and 850 nm [109]. Substituting into eq. (51) gives $\Delta t = 2.4$ ps, which is negligible compared to the time resolution of the LDM detector.

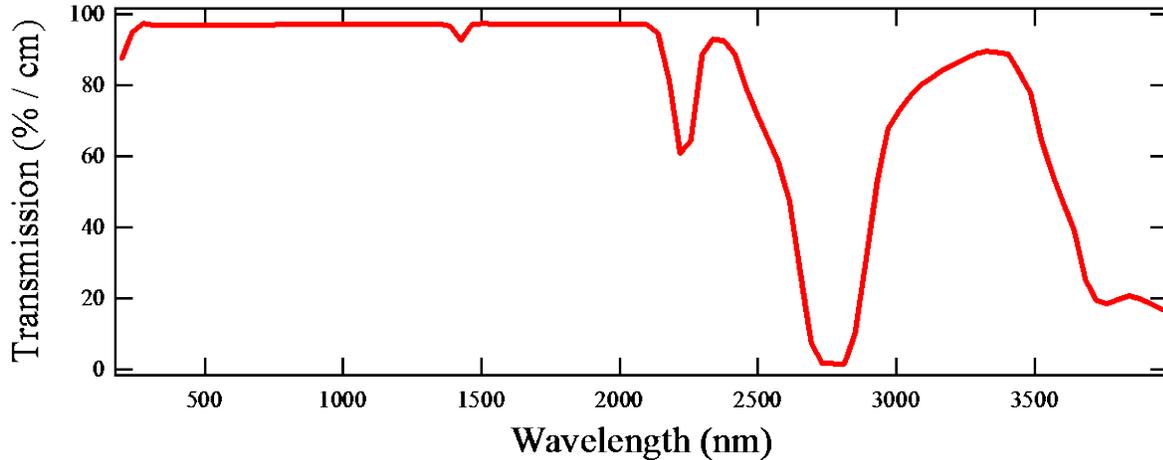


Figure 19. Transmission of the beam-pipe window against wavelength. The synchrotron light exits the beam pipe through a fused silica window. From manufacturer's specifications [110]. Note change of horizontal scale.

Directly below the extraction mirror is a 2-axis steering mirror which allows the SR beam to be aligned with the rest of the BSRT optical components. The mirror is remotely controlled and is usually operated in an auto-steering mode which keeps the beam spot at a fixed position on the BSRT camera.

The layout of the BSRT optical table [111] is shown in Figure 20 and schematically in Figure 21. The optical table is located below the beam pipe. The SR beam first enters a variable optical delay line or 'trombone'. This allows the optical path to be increased by up to 3m (Figure 22), in order to move the focus of the system from the undulator to the beginning of the D3 dipole. The first stage of the trombone consists of 4 mirrors, all angled at 45° to the main axis. Two of the mirrors are on a 750 mm translation stage, allowing the path length to be changed by 1.5 m. The second stage or 'small trombone' has a further 4 mirrors also at 45° . A smaller translation stage can insert two of these into the light path, causing the light to travel a further 1.5 m. The SR beam is then focused by a spherical mirror F1 of focal length 4.0 m. A beam splitter then reflects 85% of the light onward to a second spherical mirror, F2, of focal length 0.75 m. The remaining 15% is transmitted by the splitter and is used for the Abort Gap Monitor

[37]. After F2 a second splitter separates 8% of the remaining light for the LDM. The rest is used for transverse imaging [112]. Each instrument is equipped with its own set of neutral density filter wheels, allowing the SR intensity to be controlled independently. All of the mirrors on the BSRT optical table are coated with protected silver (Figure 23). The reflective spectra of these mirrors can be combined with that of the extraction mirror and the silica window, leading to a combined transmission spectrum for the BSRT which is shown in Figure 24.

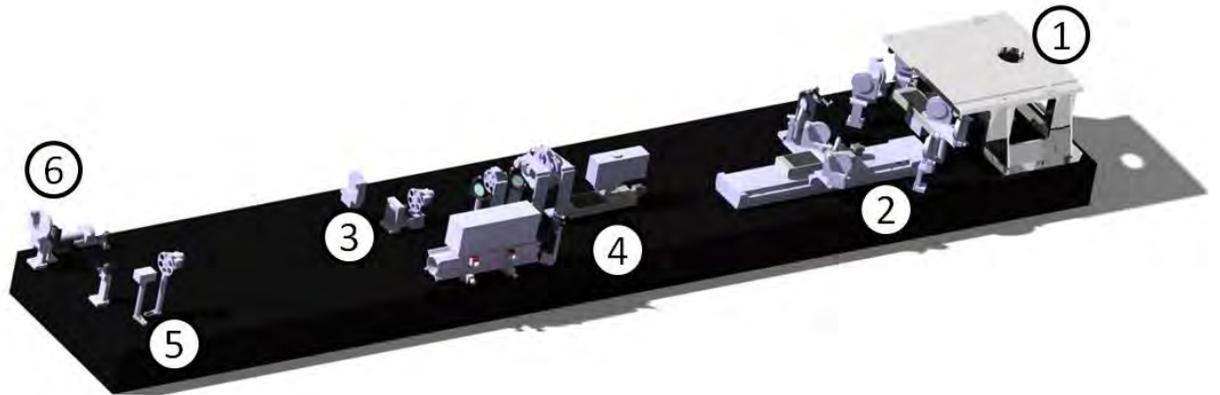


Figure 20. Layout of the Beam Synchrotron Radiation Telescope (BSRT). The BSRT is located below the beampipe and the synchrotron light is directed down through the support (1) onto a motorised alignment mirror. It passes through a variable delay line or ‘trombone’ (2) which is used to move the focus onto the dominant source of synchrotron light. The light is then split between the abort gap monitor (3), the transverse profile cameras (4) and the longitudinal density monitor (5). A calibration line (6) can be substituted for the synchrotron light.

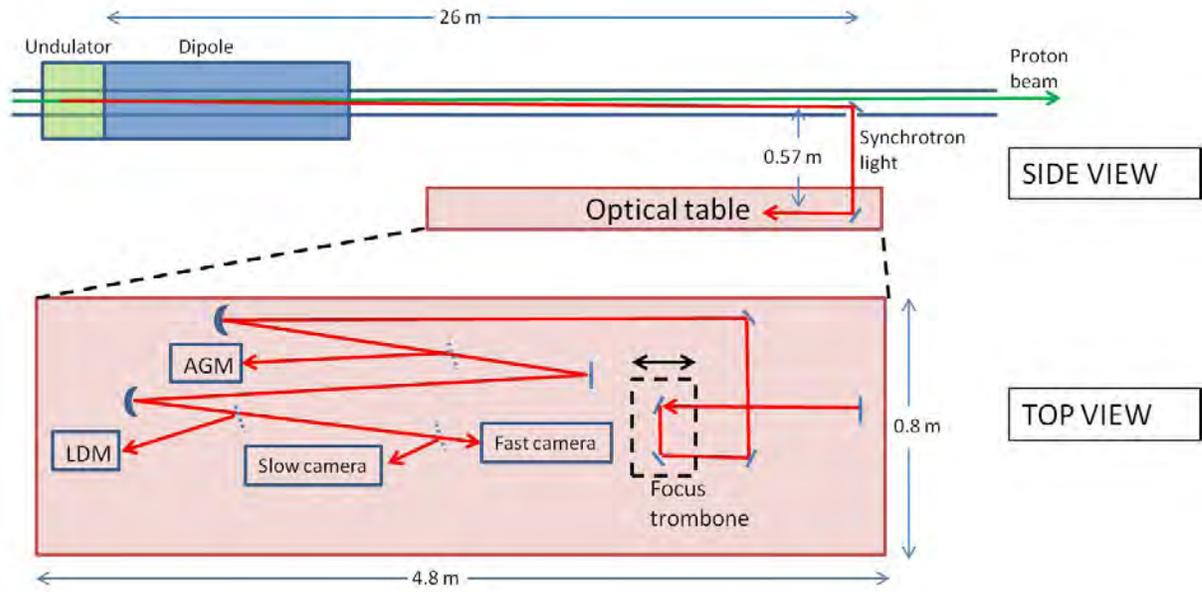


Figure 21. Simplified schematic of the Beam Synchrotron Radiation Telescope (BSRT). Above, side view. Below, top view of the optical table. Some components are omitted for clarity, e.g. filter wheels and the small trombone.

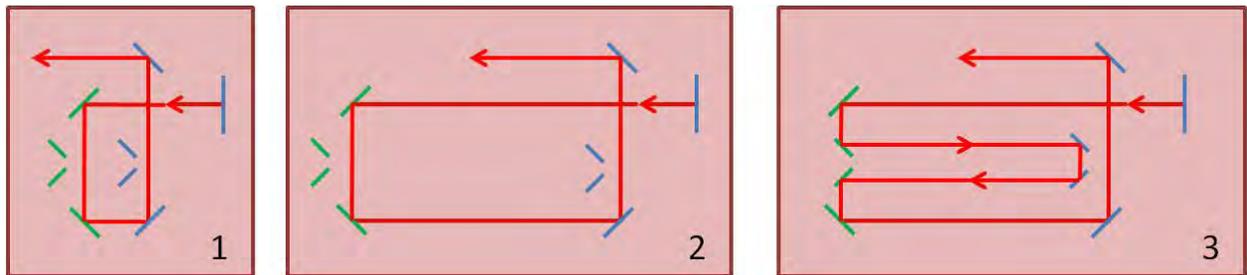


Figure 22. The optical trombone controls the focus of the BSRT by altering the path length of the SR before it arrives at the first focusing mirror. (1) Minimum extension. (2) Medium extension, +1.5m. (3) Maximum extension using the small trombone, +3m.

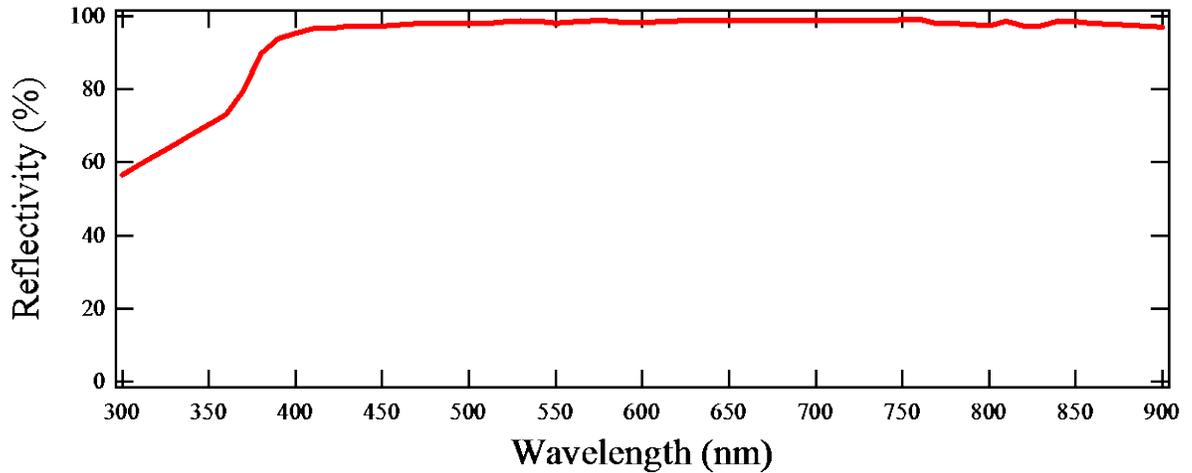


Figure 23. Reflectivity of the protected-silver mirrors against wavelength. Twelve such mirrors are used on the BSRT optical table. From manufacturer's specifications [108].

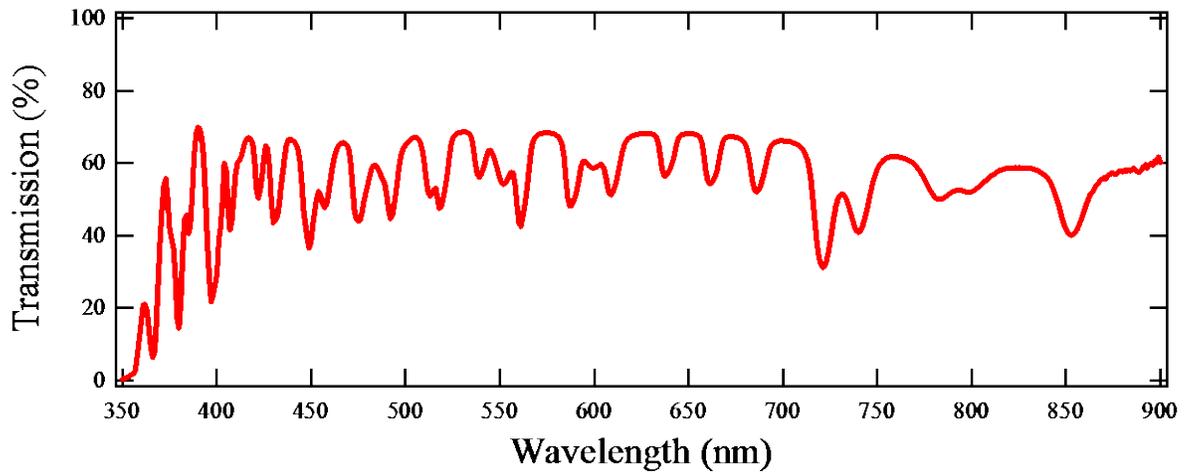


Figure 24. Transmission of the BSRT optical system against wavelength, summing all the effects from the extraction mirror to the LDM detector, assuming that the small trombone is not inserted.

The LDM thus receives around 7% of the total SR available. The LDM detector is located in the image plane of F2. The detector is mounted on an x-z translation stage allowing its position to be controlled with 1.25 μm absolute precision. The active area of the detector has a diameter of only 50 μm , while the beam spot is of the order of a few hundred μm , leading to a substantial coupling loss.

Two filter wheels are mounted on the LDM line, allowing the intensity at the LDM to be reduced by up to a factor of 2×10^7 . The content of the filter wheels is shown in Table 7.

Table 7. Neutral density filters on the LDM line.

Wheel 1		Wheel 2	
Optical density	Transmission	Optical Density	Transmission
Empty	1	Empty	1
2	0.01	0.1	0.79
3	10^{-3}	0.3	0.5
4	10^{-4}	0.6	0.25
5	10^{-5}	1	0.1
6	10^{-6}	1.3	0.05

3.3 Synchrotron Radiation: Analytical Solution

Formulae for calculating the intensity and spectrum of SR from a constant magnetic field are well known and were presented in chapter 2. Similarly, equations are available for infinitely long undulators and wigglers, and these provide a very good approximation for undulators with many periods.

However, analytical solutions for the radiation from the edge field of a dipole and from a short undulator involve some approximations. These types of SR are better modelled by a simulation code which treats each short segment of magnetic field separately. The magnetic field within these short segments may be taken as constant and the SR emission from each segment is then well-modelled with no further approximations. The SR from each segment is propagated together in order to take account of interference effects.

3.4 Synchrotron Radiation: Simulations with SRW

Synchrotron Radiation Workshop (SRW) [113] is probably the most complete SR simulation code. It was developed by O. Chubar and P. Elleaume [114] at the European Synchrotron Radiation Facility (ESRF) and is primarily aimed at simulations for electron storage rings used as synchrotron light sources. Nonetheless it can be used to simulate SR from protons and ions with a couple of small adjustments.

The SR emitted by a moving charge is determined only by its charge, its speed (through the Lorentz factor γ) and its trajectory. Thus, a proton and an electron with the same γ moving with the same radius of curvature will emit SR identically (albeit the direction of the EM field of the SR would be opposite at any point, due to the opposite charge of the proton and electron).

However, SRW requires the beam energy to be input. Instead of entering the true energy of the proton beam, the energy which an electron would have if it had the same γ as our protons must be entered. Thus,

$$E_{effective} = E_{true} \frac{m_e}{m_p} \quad (52)$$

where m_e and m_p are the electron and proton masses respectively. SRW is then ‘tricked’ into calculating the correct γ . A similar trick is used to get the correct trajectory from the magnetic field,

$$B_{effective} = B_{true} \frac{m_e}{m_p} \quad (53)$$

since an electron with energy $E_{effective}$ passing through a magnetic field $B_{effective}$ would follow the same trajectory as a proton with E_{true} passing through a field B_{true} . With the correct γ and trajectory SRW will now simulate the SR from a proton beam correctly.

The magnetic field defined in the SRW model of the BSRT system is shown in Figure 25. The origin of the coordinate system is the point where the beam enters the D3 dipole, and the longitudinal axis follows the direction of the beam orbit at this point. The nominal undulator current of 450 A and peak field on orbit of 5 T is assumed.

Two approximations are made. Firstly, although the full undulator field map is available, it is closely matched by SRW’s in-built sinusoidal undulator function and this is used instead. Use of

the in-built function was found to give better results than the arbitrary field map, as the finite number of points in the field map is interpreted as a series of discontinuities by SRW. Secondly, the precise nature of the dipole edge field is not known. The D3 dipole was produced by Brookhaven National Laboratory. Before installation into the LHC its higher-order field components were measured but a longitudinal field map was not made. The edge field is modelled using SRW's default edge field function and is defined as increasing from 10% to 90% of maximum over 56 mm. The true edge width is not known but this is suggested as a reasonable approximation given the magnet field and gap height. A small variation of this simulation parameter does not significantly influence the results. SRW models the edge field with a sigmoid function,

$$B(s) = \frac{B_{max}}{1 + e^{-s/k}} \quad (54)$$

where k is a constant which is calculated by SRW to define the steepness of the sigmoid in accordance with the stated 10-90% length of the edge field. A \cos^2 shaped edge field has also been suggested, but this model has not been implemented.

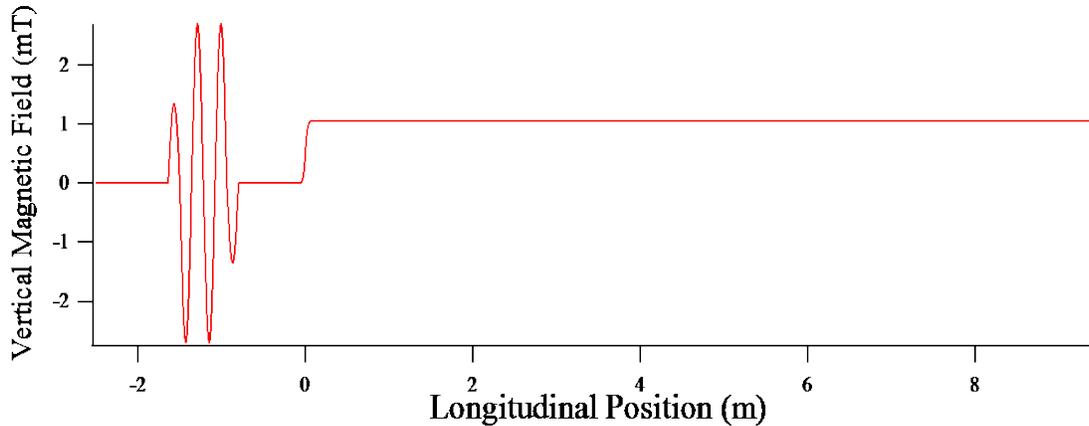


Figure 25. Vertical magnetic field against longitudinal position. Magnetic field definition for the SRW model of the BSRT system, for beam energy 3.5 TeV. The field strength is 1/1836 its real value, due to the need to model a proton beam with code designed for electrons.

SRW can compute the beam trajectory from the input parameters, and this is a useful check to ensure the scaling factors have been applied correctly. An example trajectory is shown in Figure 26 and Figure 27. The magnetic field in the dipole is increased during acceleration of the

beam so as to keep the orbit constant. The field of the undulator is however kept constant. Thus, it introduces a greater deviation in the trajectory at lower beam energies.

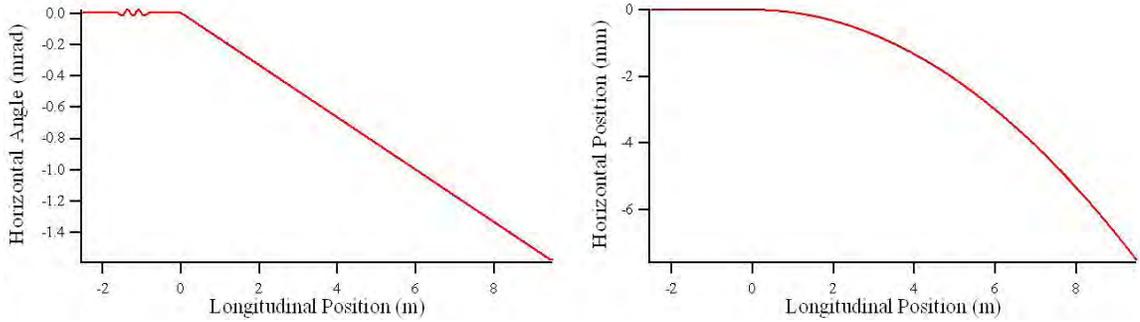


Figure 26. Trajectory of a 3.5 TeV proton beam through the BSRT undulator and dipole. Horizontal angle against longitudinal position (left) and horizontal position against longitudinal position (right). Computed using SRW with appropriate scaling factors for the beam energy and magnetic field strength. The origin of the coordinate system is at the entrance of the D3 dipole. The change in horizontal position due to the undulator is too small to be seen on this scale and is reproduced in Figure 27 below.

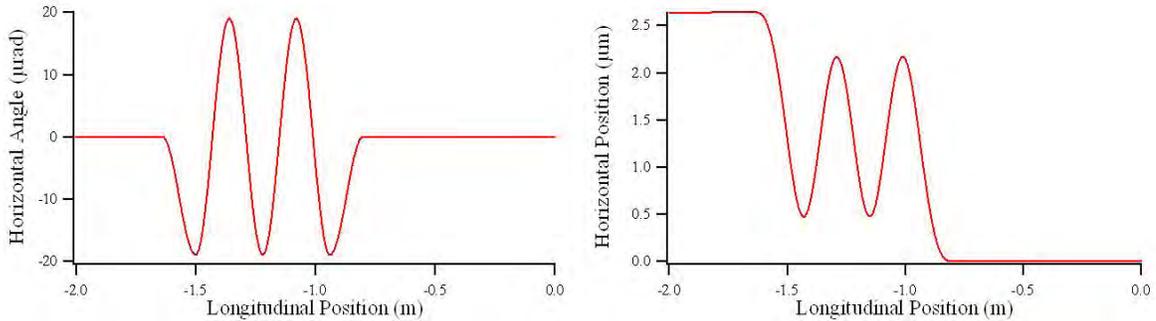


Figure 27. Trajectory of a 3.5 TeV proton beam through the BSRT undulator. Horizontal angle against longitudinal position (left) and horizontal position against longitudinal position (right). Computed using SRW with appropriate scaling factors for the beam energy and magnetic field strength. The origin of the coordinate system is at the entrance of the D3 dipole. The field integral of the undulator is zero, so the entrance and exit of the beam are parallel, but they are not collinear because the undulator poles are not matched.

3.4.1 Spectrum of Emitted SR

The central frequencies of the SR emitted by the dipole and the undulator can easily be calculated analytically using the formulae presented in chapter 2. They are shown below in Table 8 for protons and Table 9 for lead ions. The spectrum of SR emitted by a dipole magnet is broad. Since the undulator has only two periods it also has a rather broad spectrum. Thus,

even when both are centred outside the visible there is still substantial visible light available. In addition, the wavelength of the undulator radiation peak depends on the angle of emission. The undulator wavelength shown in Tables 8 and 9 is only the peak wavelength for on-axis observation. Longer wavelengths are observed away from the emission axis.

Table 8. Wavelengths of SR emission from a proton beam at various energies. The wavelength given for the dipole is the wavelength equivalent to the critical frequency, which is not the spectral peak. Critical frequency was discussed in ch. 2. For the undulator the peak wavelength on-axis is given.

	Dipole	Undulator
450 GeV	230 μm	610 nm
3.5 TeV	485 nm	10 nm
7 TeV	60 nm	2.5 nm

Table 9I. Central wavelengths of SR emission from a lead ion beam at various energies. The wavelength given for the dipole is the wavelength equivalent to the critical frequency, which is not the spectral peak. Critical frequency was discussed in ch. 2. For the undulator the peak wavelength on-axis is given.

	Dipole	Undulator
177 GeV/u	3.7 mm	3.9 μm
1.38 TeV/u	7.9 μm	64 nm
2.76 TeV/u	990 nm	16 nm

At some beam energies both the dipole and the undulator are centred outside the visible range, and in these cases radiation from the edge field of the dipole can be significant. Although the total power radiated from the edge is always much smaller than that from the dipole centre, it can produce more visible light, since its spectrum is different to that of the dipole centre.

For a full treatment of the SR spectrum at different beam energies, taking into account the edge radiation and the spatial dependence of the undulator radiation spectrum, an SRW

simulation has been used. The observation window is defined at the extraction mirror. In fact, not all of the light which is reflected from the extraction mirror is gathered into the telescope: a spectral cut is added by the transmission of the fused-silica beam-pipe window and by the reflectivity of the BSRT mirrors, while a spatial cut may be applied if the light beam spills outside the mirrors, for example due to incorrect alignment. Since the precise alignment of the BSRT with respect to the SR axis is not known, this spatial cut cannot be implemented in the simulation. This effect should be small since even the smallest mirror is several times larger than the beam size, so that even if the beam spot is not well centred little light should be lost.

The spectrum of SR calculated by SRW and integrated over the area of the extraction mirror is shown below for protons (Figure 28) and for lead ions (Figure 29). The transmission spectrum of the BSRT, shown in Figure 24 above, is then applied to these spectra in order to estimate the spectrum of SR arriving at the LDM detector, or at the transverse profiling cameras. In order to focus the cameras on the dominant SR source at any beam energy, the small optical trombone is inserted during beam acceleration. Thus, at 450 GeV the SR is reflected from 8 protected silver mirrors, but at 3.5 TeV and 7 TeV it is reflected from 12 mirrors and the transmission is slightly less. After folding with the BSRT transmission spectrum for the appropriate number of mirrors, the spectra are shown below for protons (Figure 30) and lead ions (Figure 31).

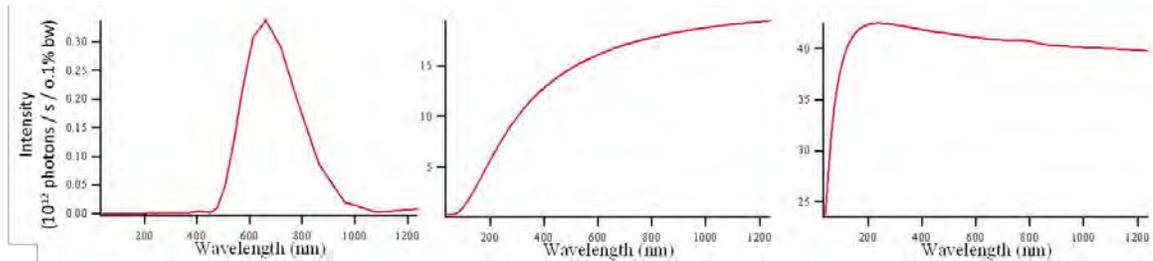


Figure 28. Intensity against wavelength. Spectrum of SR incident on the extraction mirror, for a proton beam at 450 GeV (left), 3.5 TeV (centre) and 7 TeV (right).

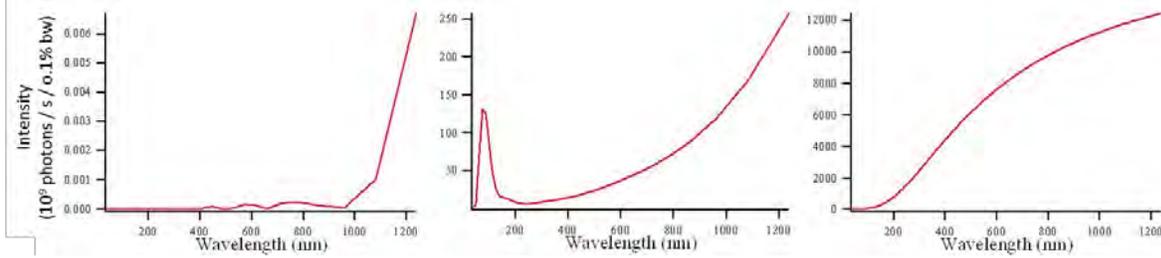


Figure 29. Intensity against wavelength. Spectrum of SR incident on the extraction mirror, for a lead ion beam at 177 GeV per nucleon (left), 1.38 TeV per nucleon (centre) and 2.76 TeV per nucleon (right).

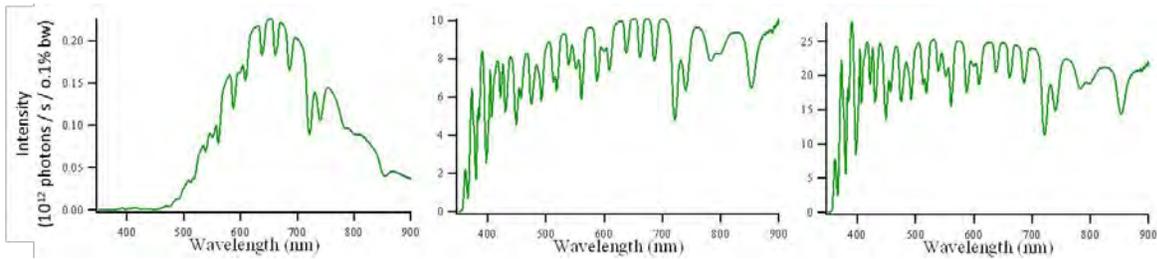


Figure 30. Intensity against wavelength. Spectrum of SR incident on the LDM detector, for a proton beam at 450 GeV (left), 3.5 TeV (centre) and 7 TeV (right).

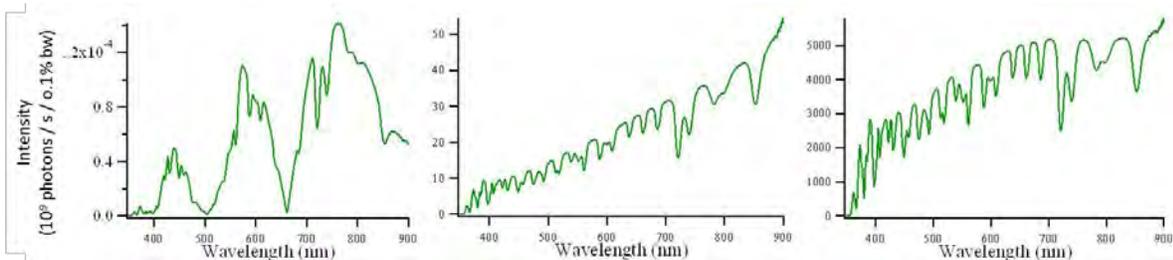


Figure 31. Intensity against wavelength. Spectrum of SR incident on the LDM detector, for a lead ion beam at 177 GeV per nucleon (left), 1.38 TeV per nucleon (centre) and 2.76 TeV per nucleon (right).

3.4.2 Intensity and Distribution of Emitted SR

The emitted SR power is strongly dependent on the energy of the beam. As was shown in ch. 2, the total power radiated whilst a particle travels through a dipole field is proportional to γ^4 . For the undulator, the radiated power is proportional to γ^2 . Figure 32 shows that the undulator then radiates at the same rate as the dipole at around 3 TeV with a proton beam. The undulator radiates more power than the dipole at lower energies, and at all energies for the lead ion beam.

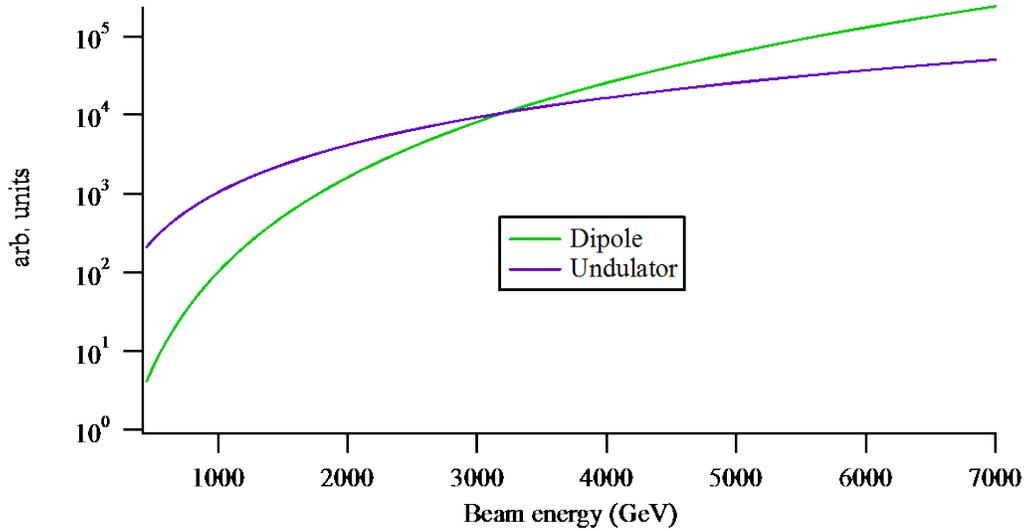


Figure 32. Total power radiated from the dipole and the undulator against beam energy, for a proton beam.

However, this radiated energy is spread over a broad spectrum. The LDM and the other BSRT instruments are sensitive only to visible light. Since the SR spectrum is also dependent on the beam energy, the variation of visible SR power with beam energy becomes more complex. Again, an SRW simulation was used, and the number of SR photons integrated over the area of the extraction mirror and over the visible range of photon energies. The results are shown in Figure 33. For protons, the undulator is centred in the visible at injection energy (450 GeV) but quickly passes into the UV. A minimum is seen at around 1 TeV where the undulator is emitting SR mostly in the UV while the dipole is still mostly in the IR. Above 1.2 TeV the SR from the dipole edge starts to enter the visible and the visible SR intensity increases. For lead ions, very little visible SR is produced at injection because both the dipole and the undulator are centred in the IR. The visible SR intensity increases strongly up to around 510 GeV per nucleon (equivalent to 1.3 TeV for protons) as the undulator radiation enters the visible range. There is then a dip around 870 GeV per nucleon (equivalent to 2.2 TeV for protons) as the undulator radiates in the UV while the SR from the dipole is still centred in the IR. The issue was also investigated in [111] without using a simulation code, and similar results were obtained as shown in Figure 34.

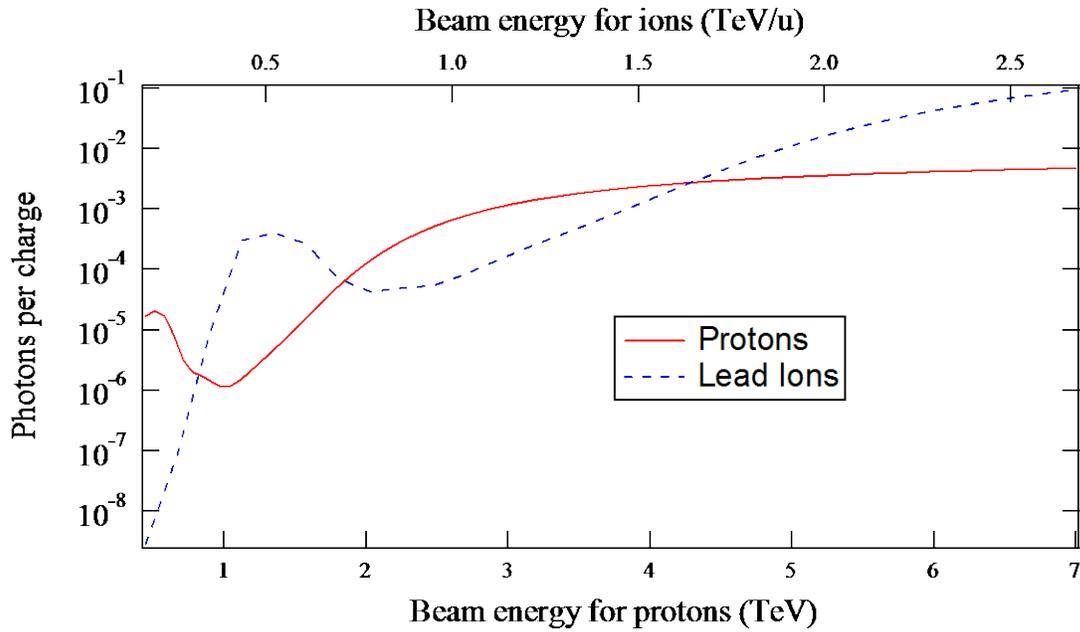


Figure 33. Visible SR intensity against beam energy, for protons and lead ions. From a simulation using SRW.

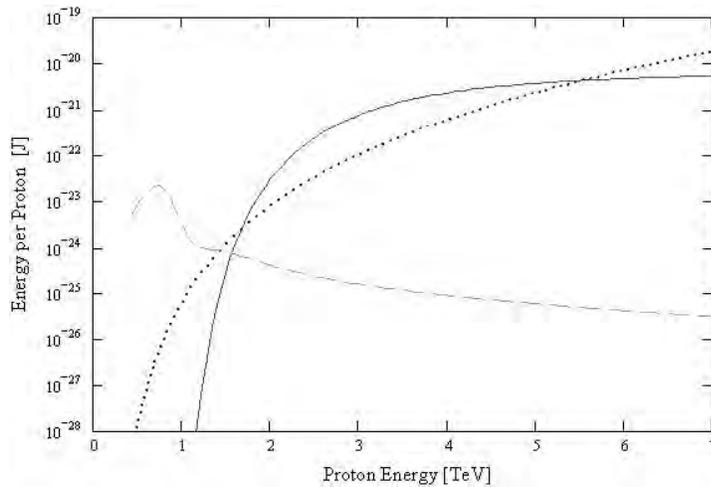


Figure 34. Radiated energy per proton collected by the extraction mirror, from the dipole centre (solid line), the dipole edge (dotted line) and the undulator (dashed line), for wavelengths between 200 nm and 900 nm. Reproduced from [111].

Synchrotron light is emitted mainly in the direction of motion of the emitting particle. The precise distribution, however, depends strongly on the beam energy. In addition, at many points during the beam ramp the SR spectrum peaks outside the visible range, and the spatial distribution of visible synchrotron light can be substantially different from the total SR distribution.

The origin of our coordinate system is the point where the beam enters the dipole and the longitudinal axis is collinear with the beam at this point. Thus SR emitted on axis from the undulator or at the entrance of the dipole will hit the extraction mirror at (0,0). This is not the centre of the mirror. SR from the dipole shows up as a streak from (0,0) to the edge of the mirror closest to the centre of the beam pipe. This is because of a ‘searchlight’ effect: as the particles are bent through the dipole, the beam of SR is swept across the extraction mirror. Only the first 3 m of the dipole contribute SR to the BSRT. After 3 m the SR beam misses the mirror.

The evolution of the SR distribution on the extraction mirror during acceleration of a proton beam can be seen in Figure 35. Here the SR has been integrated over the visible range (350 nm – 800 nm). At 450 GeV the undulator is the dominant source and the SR peak is at (0,0). At 1 TeV the undulator is still the dominant source but the on-axis undulator radiation is in the UV; the visible undulator radiation is seen as a hollow ring. The edge radiation is starting to be visible as a spot at (0,0). At 1.67 TeV the edge radiation is the dominant source. Above 2 TeV the SR from the body of the dipole is dominant.

The equivalent distributions for an accelerating lead ion beam are shown in Figure 36. The undulator is the dominant source at energies up to around 1 TeV per nucleon (equivalent to 2.5 TeV for protons). The dipole edge radiation then dominates until around 2 TeV / nucleon (equivalent to 5 TeV for protons), above which the SR from the dipole body is stronger.

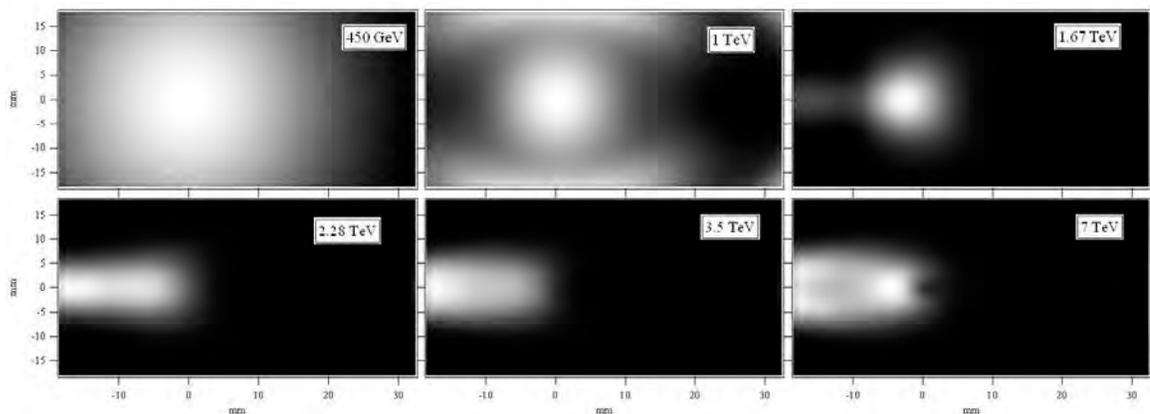


Figure 35. Spatial distribution of visible SR on the extraction mirror for a proton beam at various energies, from a simulation using SRW.

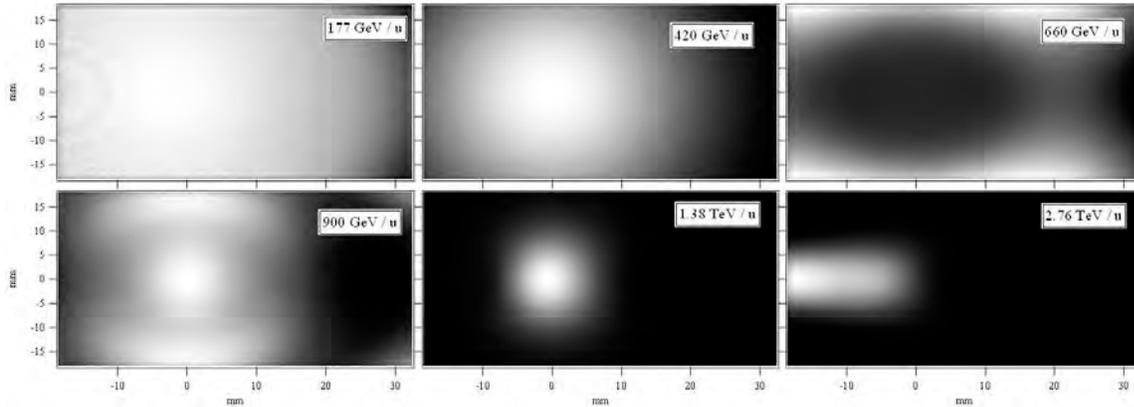


Figure 36. Spatial distribution of visible SR on the extraction mirror for a lead ion beam at various energies, from a simulation using SRW.

In the case of undulator radiation, the wavelength is strongly dependent on the angle of observation. Radiation is emitted on-axis at the undulator coherence wavelength. Other wavelengths are emitted in a hollow cone with opening angle proportional to the difference between the observed wavelength and the coherence wavelength. This can be clearly seen in Figure 37, for a proton beam at 450 GeV with a coherence wavelength of 610 nm. For SR from the dipole, the peak intensity is always on-axis but the opening angle of the emission cone is wider for longer wavelengths, as shown in Figure 38 for protons at 3.5 TeV.

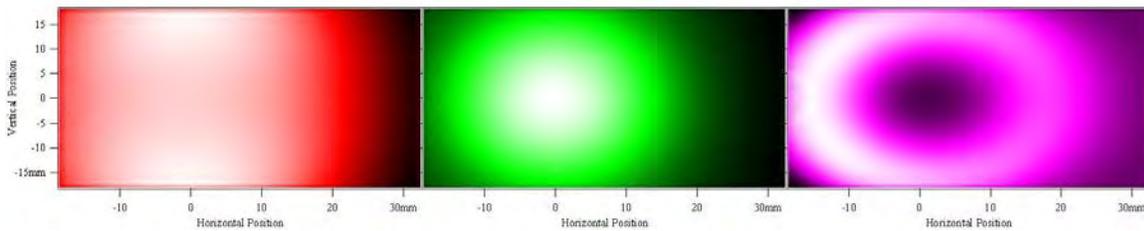


Figure 37. Distribution of SR on the extraction mirror for a 450 GeV proton beam, from a simulation using SRW. Left 700nm. Centre 550 nm. Right 400 nm.

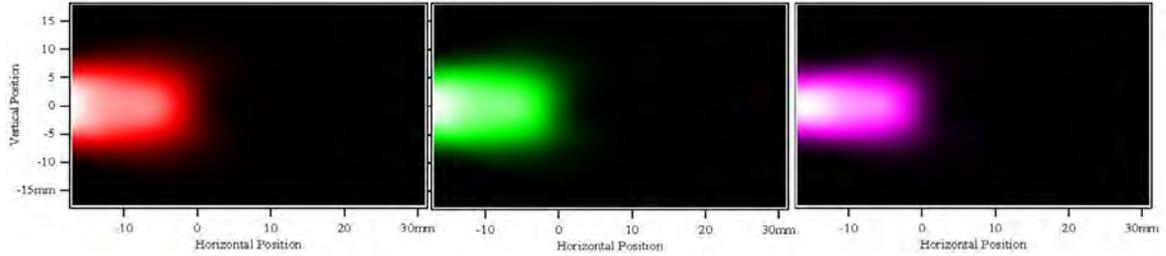


Figure 38. Distribution of SR on the extraction mirror for a 3.5 TeV proton beam, from a simulation using SRW. Left 700nm. Centre 550 nm. Right 400 nm.

SR is predominantly polarised in the plane of the particle motion, horizontal in this case. In fact, an observer in this plane will see only transverse particle motion in the plane, so the SR is fully horizontally polarised. An observer above or below this plane, however, would see an elliptical particle motion. The polarisation of the SR would then have a small vertical component, as shown in Figure 39.

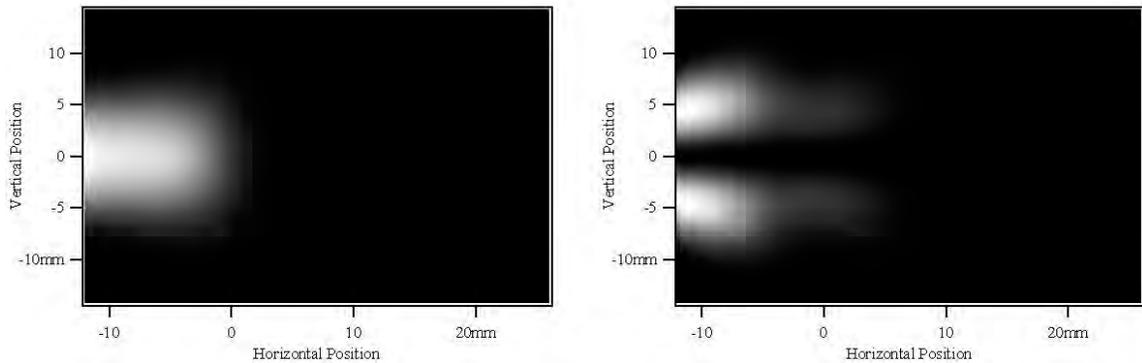


Figure 39. Distribution of SR from protons at 3.5 TeV, from a simulation using SRW. Left, horizontal polarisation. Right, vertical polarisation. The vertically polarised component constitutes only 11% of the total intensity.

3.4.3 Heat Load on the Extraction Mirror

The reflective surface of the BSRT extraction mirror, which is used to direct the SR out of the beam pipe, is made of baked Aluminium and its reflectivity decreases sharply at wavelengths below 400 nm. Shorter wavelengths are absorbed and cause a local heating of the mirror. At higher beam energies, the SR striking the mirror has a substantial UV and x-ray component. A simulation was carried out using SRW in order to estimate whether the SR power would be

great enough to cause a temporary distortion of the image due to uneven heating and expansion of the mirror, or even permanent damage to the reflective coating.

SRW offers the option 'compute power density' which allows the total SR power, integrated over all wavelengths, to be calculated at any position. In this case the upper and lower bounds of photon energy specified in the 'radiation sampling' box are ignored. The total power calculated in this way is a good approximation for the power deposition on the mirror above ~5 TeV, when a very large proportion of the SR power is contained in short wavelengths. An example of such a calculation is shown in figure 40.

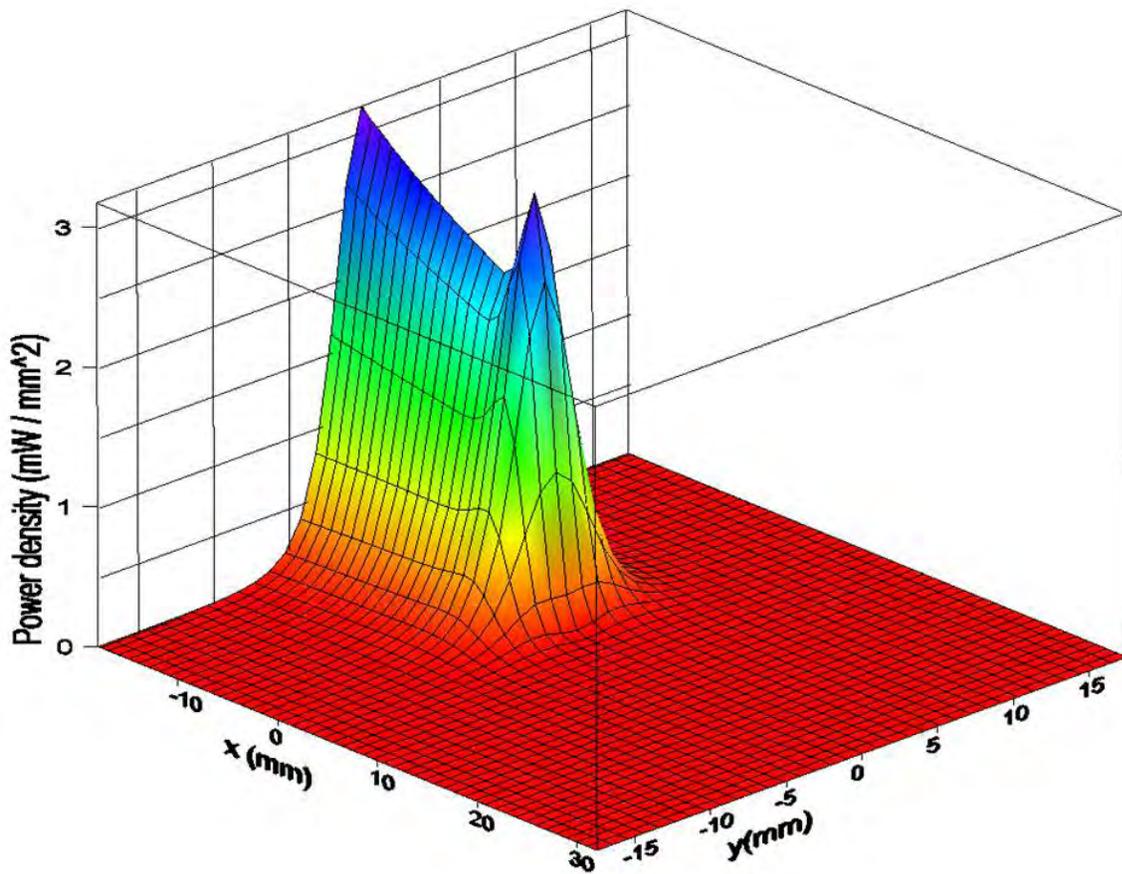


Figure 40. SR power distribution on the extraction mirror, arbitrary units. Simulation using SRW of the undulator and dipole for a proton beam at 7 TeV.

At lower beam energies, the SR at wavelengths longer than 400 nm is not negligible. This SR is reflected down to the BSRT and therefore does not contribute to heating of the mirror. The total SR power density is then not a good approximation for the heat load on the mirror.

Instead, the full electric field calculation must be used, and the results integrated over the wavelength range of interest.

Care must be taken to specify an appropriate integration step for the electric field calculation. Reducing the step size increases dramatically both the time and the memory required to run the calculation. However, if the step is too large, spurious radiation may be predicted at short wavelengths. It is therefore necessary to use a much shorter step size when running simulations of SR power, which depend on accurate simulation of SR in the x-ray region, than when simulating visible light for the BSRT imaging system. The independence of the results on the integration step size is a necessary (although by no means sufficient!) condition for the reliability of the simulation. Figure 41 shows the heat load on the mirror for protons at 7 TeV. The simulation was repeated for 3 different values of the integration step size and differences in the results were negligible.

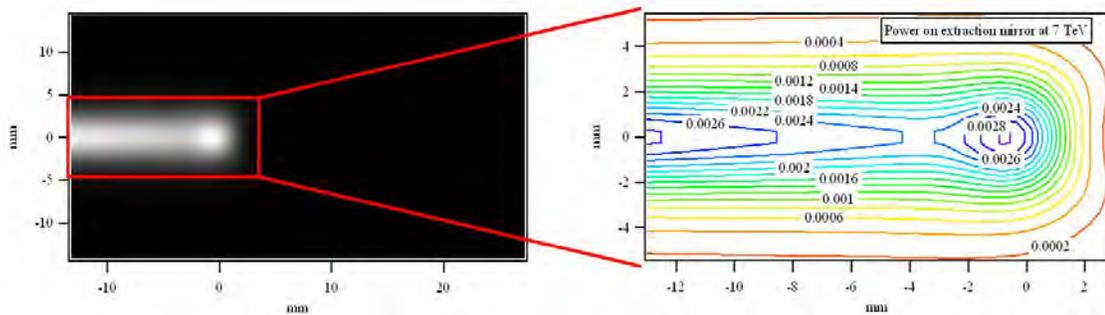


Figure 41. Heat load on the extraction mirror for protons at 7 TeV, W/mm^2 . Simulation with SRW, integrated over the wavelength range 5 nm - 400 nm.

The total heat load integrated over the mirror surface is 11.5 times higher for a proton beam at 7 TeV compared to a beam at 3.5 TeV. This heating causes the mirror surface to expand unevenly and thus bend, which in turn will be seen as movement and distortion of the image produced on the BSRT transverse profile cameras. Since the synchrotron light travels a further 10.5 – 13.5 m from the extraction mirror to the plane of the camera and the LDM, depending on the position of the trombone, a small change in the planarity of the mirror can cause significant movement in the camera plane. If the heat deposition is sufficiently fast it could in addition cause permanent damage to the mirror coating. One possible mitigation strategy would be to turn off the undulators once they are no longer needed, i.e. once the beam has been accelerated past 1.5 TeV. However ramping up or down the undulators, as with any of

the LHC's superconducting magnets, carries a small risk of triggering the quench protection system (QPS) which would cause a beam dump.

The simulation was repeated with only the dipole magnetic field present and then with only the undulator. The contribution to the mirror heating by the undulator is only 10-15% of the total, as shown in figure 42. Based on these calculations it was decided that the undulators should not be routinely ramped down during operation at 7 TeV. If there is a problem of image distortion due to heating of the mirror, it would not be solved by switching off the undulators.

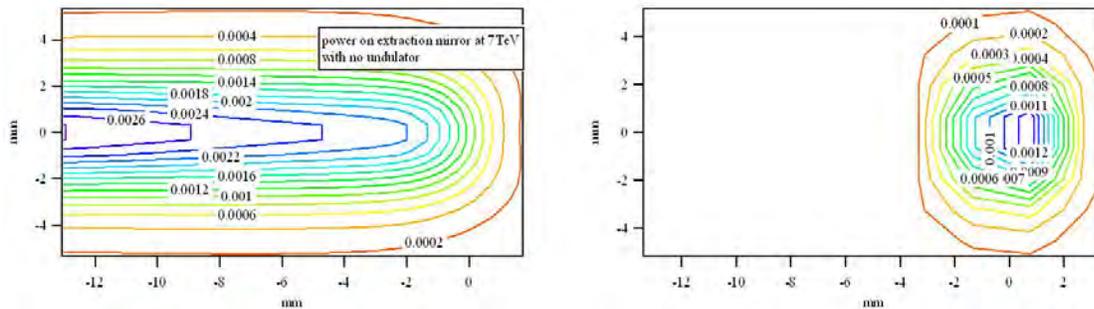


Figure 42. Heat load on the central part of the extraction mirror for protons at 7 TeV, W/mm^2 . Simulation with SRW, integrated over the wavelength range 5 nm - 400 nm. Left: contribution from the dipole, including dipole edge radiation. Right: contribution from the undulator.

3.5 Synchrotron Radiation: Simulations with SPECTRA

SPECTRA [115] is an alternative synchrotron radiation code developed at Spring-8 in Japan. Like SRW, it is designed for use with electron accelerators but can easily be used with protons or ions by dividing the beam energy and magnetic field strength by an appropriate factor, as given by eq. (52) and (53) above.

The principle used by SPECTRA for the calculation of synchrotron radiation is the same as SRW except that a far-field approximation is applied. This is valid providing that the SR is observed at a distance from the point of emission which is much greater than the wavelength of the radiation. The far-field approximation allows the calculation to be carried out much more quickly.

In addition, SPECTRA contains preset functions for the calculation of undulator radiation. However, these are based on the long undulator equation. In the case of the BSRT undulator

this is not appropriate so the undulator radiation must be modelled by inputting a full field map of the undulator.

SRW operates as a plug-in to the Igor environment [116], a commercially available mathematical analysis and graphing package, while SPECTRA is a stand-alone program. While this of course makes the installation easier, the program is less flexible as it cannot be managed through the Igor macros. SPECTRA does include the option to automatically perform repeated calculations while varying one parameter. However this is not always helpful as it is frequently necessary to vary two parameters at the same time. For example, simulating the change of SR emission during acceleration of the beam requires both the beam energy and the dipole field strength to be varied. This process cannot be automated in SPECTRA.

SRW has been used as the primary tool for SR simulations in this work due to its greater flexibility. Nonetheless, it is useful to cross-check some of its results with SPECTRA simulations. No substantial disagreement has been found between the two codes for any of the cases examined. Figure 43 shows a comparison of the results from SPECTRA and SRW for the SR intensity distribution on the BSRT extraction mirror for a 450 GeV proton beam, with agreement to better than 10% at all positions. The differences in the results are probably caused by the different means of defining the undulator field in the two codes.

Other codes for the simulation of SR are available, for example Program UR [117] and Zgoubi [118]. The former was found to be overly simplified for the present case, while the latter is more concerned with the effects of synchrotron radiation on the particle optics.

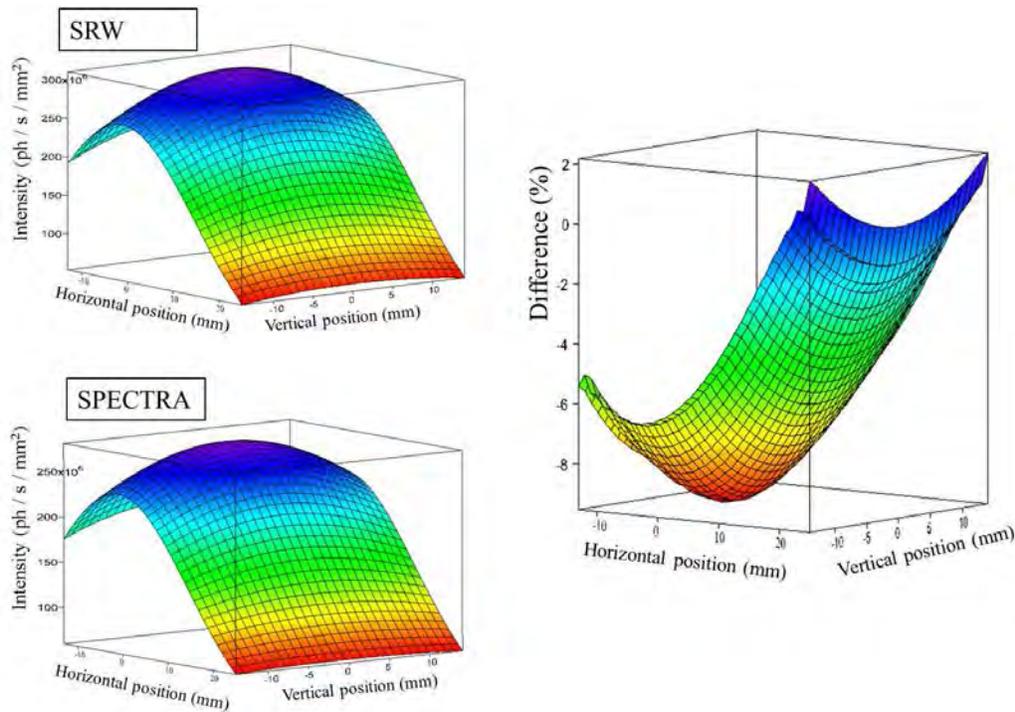


Figure 43. Visible SR intensity distribution on the BSRT extraction mirror, for protons at 450 GeV, simulated using SRW and SPECTRA. Right, difference between the two results as a percentage of the maximum bin.

Chapter Summary

The Beam Synchrotron Radiation Telescope, which gathers synchrotron light for the LDM, has been presented. The three key sources of synchrotron radiation in the BSRT (undulator, dipole and edge) are discussed and their relative importance is quantified for different beam energies. The synchrotron radiation code SRW has been used to perform extensive simulations of the emission from these magnets, and results for the spectral and spatial distribution are shown. The number of visible light photons collected on the extraction mirror varies strongly with the beam energy. For proton beams, minimum intensity is reached around 1.2 TeV, but the light is still sufficient for photon counting measurements. For lead ion beams, insufficient light is available at the injection energy of 450 Z GeV, but measurements become possible above 1 Z TeV. Heating of the extraction mirror due to undulator radiation should not be significant.

4. Single Photon Counting with Avalanche Photo-Diodes

4.1 Single Photon Counting

4.1.1 Principle of Single Photon Counting

Whether it is viewed as a wave or a particle, it is known that EM radiation can only be emitted in discrete quanta of energy [119], known as photons. Thus, the detection of a single photon represents the maximum sensitivity that any detector can achieve.

Such sensitivity can be achieved by detectors which have a sufficiently high internal gain. The most common types are photo-multiplier tubes (PMTs) and avalanche photo-diodes (APDs). The group of techniques which uses these detectors to record individual photon-detection events is known as single photon counting (SPC). SPC is an inherently digital technique, since at any time either a photon is detected or it is not. Generally, the arrival of multiple photons simultaneously cannot be processed in SPC. Analogue techniques which use a single-photon-sensitive detector but integrate its output over some time window are not considered SPC.

Various forms of SPC exist [120]. Steady-state SPC simply counts the photons and works out their arrival rate, in order to measure the average intensity of a very weak light source. Gated SPC counts only those photons which arrive within a time window of interest. When used with a periodic signal, the gate may be shifted slightly at each repetition, so that the waveform of the periodic signal is sampled over a large number of cycles.

A more efficient way of recording the waveform of a periodic signal is time-correlated single photon counting (TCSPC). TCSPC records the arrival times of photons with respect to the start of the signal. By accumulating the arrival times over many cycles and building up a histogram [121], the waveform is accurately reconstructed. The more measurement cycles which are accumulated into the histogram, the better the signal-to-noise ratio (SNR) of the reconstruction (Figure 44). For a detector which gives no false counts,

$$SNR = \sqrt{N} \quad (55)$$

where N is the number of photons in each bin of the histogram. It is clear that the smaller the bin width of the histogram (i.e. the higher the bandwidth of the measurement), the lower the SNR will be for the same accumulation time.

If two or more photons arrive at the same time or faster than the TCSPC system can process them, they will be recorded as a single photon. It is therefore important that the light intensity is kept sufficiently low as to make this improbable. Stronger light signals may be attenuated in order to fulfil this condition. The probability of a photon arriving in any given bin on any cycle is then much less than one, so that the number of cycles needed is much larger than the square of the SNR required.

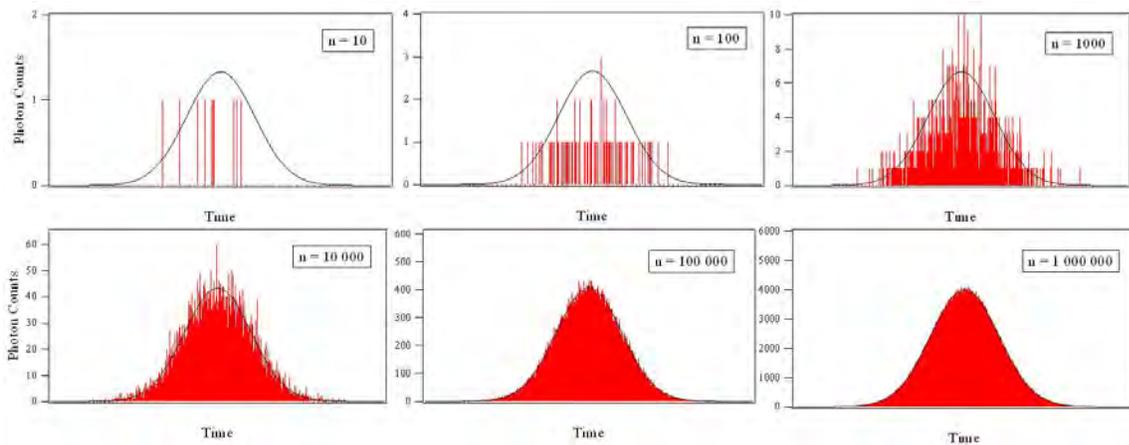


Figure 44. Histogram of photon counts against time. Time-correlated single photon counting (TCSPC) becomes more accurate as the number of cycles is increased. Here, a simulation of TCSPC on a Gaussian pulse is shown for a number of counted photons n from ten to one million.

TCSPC can be used with any signal which can be repeated many times. It is widely used in biochemistry to measure fluorescence lifetimes; in materials science for the characterisation of crystal imperfections through positron lifetime experiments; and in a variety of remote sensing applications such as laser range finding.

4.1.2 Advantages and Disadvantages of SPC

Digital techniques like TCSPC can achieve a much higher dynamic range than analogue measurements, which are generally limited by electronic noise.

On the other hand, TCSPC is slow, requiring millions of cycles for a precise measurement. It is therefore insensitive to fast bunch shape changes such as the quadrupole oscillations that may follow a badly matched injection.

4.1.3 Detectors for SPC

Any detector which has a sufficient internal gain for the associated electronics to be able to pick up the signal of a single photon may be used for TCSPC. Suitable detectors [4] [5] include photomultiplier tubes (PMTs), microchannel PMTs (MCPs), avalanche photodiodes (APDs), superconducting single photon detectors [124], and photoconductive switches [125]. The choice must be based on the required time resolution, sensitivity, noise level, and budget. The reasons for the choice of APDs for use in the LDM will be detailed below.

4.2 Avalanche Photo-Diodes

4.2.1 Principle of Photon Detection with Avalanche Photo-Diodes

An avalanche photo-diode (APD) is a solid-state device which detects photons by their ability to release photoelectrons at a semiconductor junction. A bias voltage is applied over a slab of semiconductor containing positive and negative doped regions. A high internal electric field is then generated at the junction of these regions, such that a photoelectron in this field is accelerated sufficiently to generate a further electron-hole pair by impact ionisation [122]. These new electrons are in turn accelerated and generate further pairs, so that an avalanche occurs.

In a normal APD, the bias voltage is below the breakdown voltage of the semiconductor, and the holes do not gain enough energy to initiate ionisation. The avalanche thus spreads only in one direction, from the p to the n region, and stops once it leaves the high-field junction region. The multiplication is usually no more than 200 electron-hole pairs from one original photoelectron.

However, if the bias voltage is above breakdown, both the electrons and the holes gain enough energy for further ionisation. The avalanche is then self-sustaining as electrons and holes travel

backwards and forwards in a positive feedback effect. The APD is then said to be in Geiger mode in analogy to the Geiger counter used for radiation detection.

APDs designed to be operated in the Geiger mode (G-APDs) were first developed by McIntyre [126] and Haitz [127]. The multiplication factor of such a self-sustaining avalanche can be extremely high. Gains of up to 10^9 have been reported [128], although 10^5 - 10^7 is more usual [129]. This is sufficient to be detected by a sensitive discriminator or amplifier, and thus the device can be used as a single photon detector. For this reason the G-APD is sometimes known as a single-photon avalanche diode (SPAD).

4.2.2 Construction

Since the G-APD is biased above the breakdown voltage, any free electrons will initiate an avalanche. The semiconductor used must therefore have a sufficiently high band-gap to reduce the appearance of thermally generated free electrons. The bandgap must however be smaller than the energy of the photons to be detected.

Silicon is the most common material for visible-light APDs [130]. Silicon has a bandgap of 1.11 eV at room temperature [131], equivalent to a photon of wavelength $\sim 1 \mu\text{m}$, and is thus ideal for detection of visible light. Due to its prevalence in computing and electronics, silicon has been extremely well-studied, and many manufacturing techniques and facilities exist. In addition, the APD and its associated readout circuits can be built into the same chip.

Other materials have been used mostly for photon detection in the infra-red. This is of particular interest for the telecommunications industry and in quantum key distribution (QKD). Among other materials, Germanium [132] and Indium Gallium Arsenide (InGaAs) [133] have been used for IR-sensitive APDs. InGaAs in particular has shown promising results for high count-rate photon detection [134]. However, the lower band-gap of such materials means that they suffer from high thermally-generated noise rates and must be cooled, sometimes to cryogenic temperatures [135].

The diode is manufactured by doping the chosen semiconductor with carefully selected impurities. Certain regions or layers of the semiconductor are n-doped, by adding a dopant which increases the availability of conduction electrons, while another region or layer is p-doped, i.e. the added dopant increases the availability of positively-charged electron holes. A high electric field will develop at the junction of these regions, and this is the crucial

multiplication region of the APD, where the run-away avalanche occurs. Many different arrangements have been used (Figure 45). In order to achieve a good time resolution, it is important that the doped regions be arranged such that the propagation time of avalanches is uniform and independent of the site of the photoelectron.

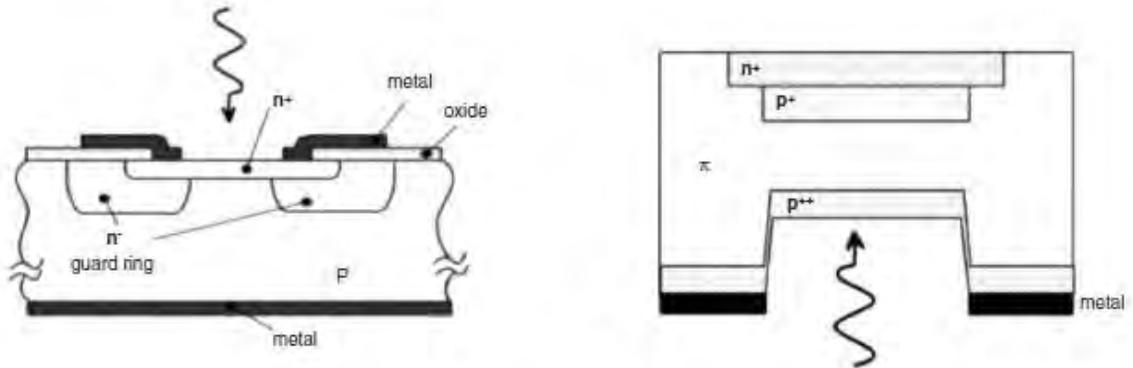


Figure 45. Possible arrangements of doped regions in a Geiger-mode APD. Left, the planar or epitaxial type, originally developed by Haitz. Right, the ‘reach-through’ type originally developed by McIntyre. From [129].

4.2.3 Active and Passive Quenching

In a Geiger-mode APD the avalanches are naturally self-sustaining. The avalanche must be stopped or ‘quenched’ by the reduction of the electric field strength below breakdown, or the heat generated would damage the material. Two methods have been developed: passive quenching and active quenching.

Passive quenching means that a high-ohm resistor is connected in parallel to the G-APD [136] so that the current generated by an avalanche creates a back-voltage across the resistor, which reduces the bias voltage across the APD below the breakdown voltage. This simple solution is usually employed in multi-pixel G-APD devices (often called silicon photomultipliers or SiPMs). However, the time taken to quench the avalanche is relatively long, and devices employing passive quenching tend to have relatively poor time resolution and long deadtime.

In active quenching, an electronic circuit is employed to sense the onset of the avalanche and reduce the voltage below breakdown [137]. Initially external to the G-APD, such active

quenching circuits can now be incorporated into the APD chip [138]. Because the circuit can sense the avalanche and quench it before it becomes very large, active quenching APDs tend to have shorter deadtime. In addition, the active quenching circuit can double as a timing circuit [139]. Timing resolutions as low as 20 ps have been reported for G-APDs built on this principle [140].

4.2.4 Deadtime

The avalanche occurring in a Geiger-mode APD is externally quenched by the lowering of the bias voltage. The voltage must be kept below the breakdown level for long enough to allow the avalanche to fully dissipate, i.e. for all the generated charge carriers to have recombined. Whilst the voltage level is lowered, a new self-sustaining avalanche cannot occur, and the detector is effectively blind to any arriving photons. This period is known as the deadtime.

In a passive quenching APD, the deadtime is determined approximately by the value of the resistance used, and by the parasitic capacitance of the circuit. In the case of active quenching, the deadtime can be set to the desired value by careful arrangement of the active quenching circuit, and is sometimes programmable. This generally allows the deadtime to be minimised. If it is made too short, however, the avalanche might not have fully dissipated and remaining charge carriers would immediately start a new avalanche.

In traditional TCSPC applications, the photon arrival rate is kept so low that the probability of another photon arriving within the deadtime is negligible. For the LDM, however, it is desirable to keep the acquisition time to a minimum and this requires a higher count rate to be used. The effect of the detector's deadtime on the acquired signal is then significant.

It is clear that the deadtime of the detector reduces the measured count rate, since any photons which arrive during the deadtime are not counted. If the signal is time-modulated then in addition the signal shape will be distorted. The deadtime of the APD used for the LDM is 77 ns, which is much longer than the bunch length, and therefore only one photon can be detected from each bunch on each turn. If a photon is detected in the earlier part of the bunch then the APD will be blind to any further photons. The reduction in count rate is thus larger for times later in the bunch. The bunch profile is skewed towards the front of the bunch, as well as being reduced in amplitude [141].

Consider an isolated bunch with a bunch length shorter than the deadtime. Let $p(t)$ be the probability of a photon being received at time t , and $a(t)$ be the probability of the detector being available i.e. not in deadtime. Then

$$x(t) = a(t) p(t) \quad (56)$$

where $x(t)$ is the probability of a photon being detected at time t . Here it is assumed that the detector has a photon detection efficiency (PDE) of 1, i.e. all arriving photons are detected if the detector is available. In reality, the PDE will be less than 1, but this does not change the analysis except that $p(t)$ would then be the probability that a photon is received multiplied by the PDE.

If any photon is detected then the detector will not be available, so

$$\frac{da(t)}{dt} = -a(t) p(t) \quad (57)$$

Integrating,

$$\int_0^t \frac{1}{a(t)} da(t) = - \int_0^t p(t) dt \quad (58)$$

Now,

$$\int_0^t \frac{1}{a(t)} da(t) = \ln a(t) - \ln a(0) \quad (59)$$

At the beginning of the isolated bunch the detector has not received any photons so $a(0)=1$. Then

$$a(t) = e^{-\int_0^t p(t) dt} \quad (60)$$

If the total probability of receiving a photon from the bunch is much less than one, this can be written as

$$a(t) = 1 - \int_0^t p(t) dt \quad (61)$$

The effect of the deadtime is more pronounced the higher the count rate. The effect of deadtime on a single isolated bunch is illustrated in Figure 46 for different total photon probabilities.

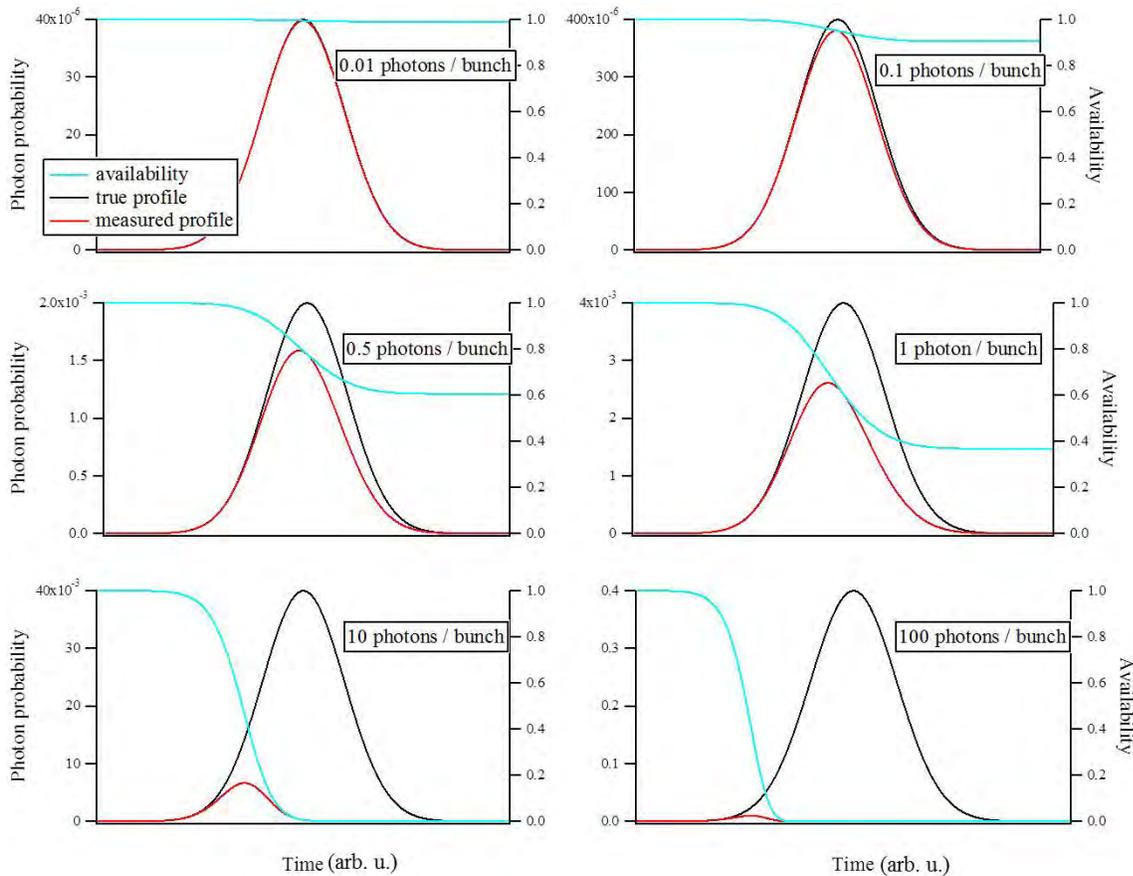


Figure 46. Photon probability against time (red and black lines, left axis), detector availability against time (blue line, right axis). Effect of deadtime on the measured profile of a single isolated bunch, for different photon arrival rates.

Since one of the principal uses of the LDM is the quantification of satellite bunches, the effect of the deadtime on the detection of a small bunch close to the main bunch should be considered. It can be seen from Figure 46 that the availability of the detector is lowest immediately after the bunch. The LDM is therefore less sensitive to trailing satellites (those arriving immediately after the main bunch) than leading ones (those arriving immediately before). The photon arrival rate should therefore be chosen so as to maximise the sensitivity to the trailing satellites. It can be seen that if the photon rate is too high, the detector will almost always be in deadtime when the satellite passes, while if it is too low then very few photons from the satellite will arrive.

Consider a satellite immediately following an isolated bunch. Let M be the ratio of the main bunch population to the satellite population, so that the satellite emits p photons per turn and the main bunch emits Mp . Then the probability of a count from the satellite bunch is

$$x = pa = p e^{-Mp} \quad (62)$$

where a is the availability of the detector at the start of the satellite bunch. The probability of receiving two photons from the satellite on the same turn is very small, so that the deadtime effects within the satellite can be neglected. As can be seen in Figure 47, x is maximised when $Mp=1$. The maximum count rate from a trailing satellite bunch is then

$$x = \frac{1}{eM} \quad (63)$$

It should be noted that here $Mp=1$ is the number of photons arriving at the detector from the main bunch.

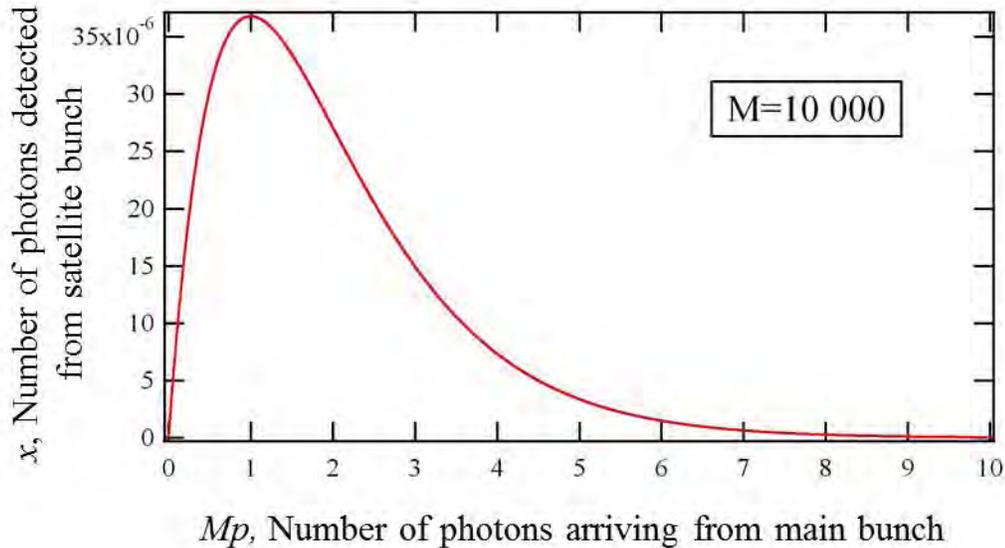


Figure 47. Probability of a count from a trailing satellite against number of photons emitted from the main bunch. M is the ratio of the main bunch to satellite bunch population.

The APD's deadtime of 77 ns is longer than the bunch separation in some LHC filling schemes, so the bunches cannot always be considered isolated. It is necessary to consider the effect of deadtime from one bunch to the other. If $P_{(b)}$ is the total probability of a photon being received from bunch b ,

$$P_{(b)} = \int p(t)dt \quad (64)$$

then the probability of a photon being detected during bunch b is

$$X_{(b)} = A_{(b)}(1 - e^{-P_{(b)}}) \quad (65)$$

where $A_{(b)}$ is the availability of the detector at the start of bunch b . Here the term in brackets is derived from equation (60) remembering that the availability was defined as the probability that a photon has not been detected. It can equally be derived from Poissonian statistics, since a very large number of charged particles are involved, each with a very small probability of emitting a photon. The Poisson distribution [142] is given by

$$f(k, \lambda) = \frac{\lambda^k e^{-\lambda}}{k!} \quad (66)$$

where $f(k, \lambda)$ is the probability of exactly k events occurring if the expected number of events is λ . Then the probability of at least one event occurring is

$$1 - f(0, \lambda) = 1 - e^{-\lambda} \quad (67)$$

Let us assume that the bunch separation is less than the deadtime but greater than half the deadtime, as is the case for the 50ns filling scheme in the LHC. Then the deadtime due to bunch b affects bunch $b+1$ but not bunch $b+2$. The availability at the start of any bunch is then

$$A_{(b)} = 1 - X_{(b-1)} \quad (68)$$

Equation (65) can then be re-written as

$$X_{(b)} = (1 - X_{(b-1)})(1 - e^{-P_{(b)}}) \quad (69)$$

If the count rate from the first bunch is large, $X(1) > 0.5$, then $X(2)$ must be less than 0.5, and X oscillates between high and low values. This behaviour is illustrated in Figure 48 for various values of P . If the bunch separation is less than half the deadtime, oscillations with longer period can be observed.

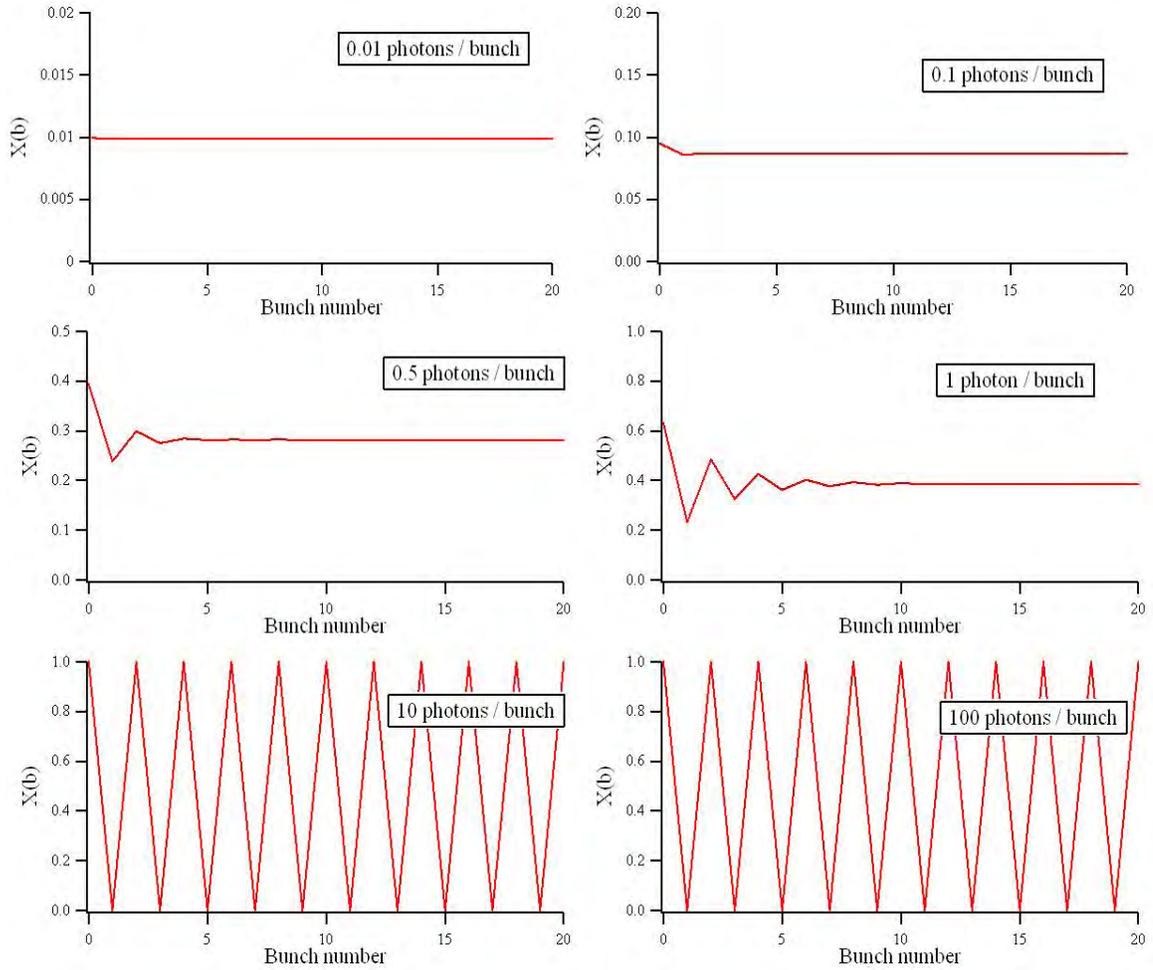


Figure 48. Detected photons per bunch against bunch number, shown for different photon arrival rates. Here the bunch separation is less than the deadtime but greater than half the deadtime, as is the case for the 50ns filling scheme in the LHC.

4.2.5 Afterpulsing

During an avalanche, some charge carriers can become trapped in energy states within the band-gap of the semiconductor [143]. The presence of impurities strongly increases this probability. The trapped carriers are later released by thermal excitation [144], with a trapping lifetime which is generally of the order of hundreds of ns. If the bias voltage is above the breakdown level at the time when the charge carrier is released, it will cause an avalanche. Such an avalanche is known as an afterpulse and is indistinguishable from a photon-caused avalanche. Thus, afterpulsing causes false photon counts which are correlated with the true photon counts.

4.2.6 Dark Counts

Electrons in the APD may be released by thermal excitation, creating an avalanche. These ‘dark counts’ occur even in the absence of illumination and are uncorrelated with the true photon counts. The dark count rate depends on the temperature and on the bandgap of the semiconductor used. Infra-red APDs made of materials such as InGaAs can have dark count rates in the 100s of kHz, while a cooled silicon APD can have a rate as low as tens of dark counts /second.

4.2.7 Diffusion Tail

Most avalanches propagate quickly and uniformly, such that the output pulse can have low time jitter. 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 full-width at half maximum (FWHM) time resolution of the APD is usually quoted, which ignores the diffusion tail effect. However, for high-dynamic range measurement the diffusion tail becomes important.

4.2.8 Gated APDs

A Geiger-mode avalanche can only be produced if the APD is biased above the breakdown voltage. Thus, it is in principle possible to turn the APD off and on very quickly, by altering the bias voltage. This is known as gating, and gated APDs are commonly used in some applications. In telecommunications and quantum-key distribution, the approximate arrival time of the photon is often known, and the APD can be gated on only at this time, in order to reduce the dark counts. In fluorescence lifetime imaging, on the other hand, the APD is usually gated off around the arrival time of the excitation pulse, so that it does not contaminate the measurement of the fluorescence decay.

4.3 Correction for the Detector Response

4.3.1 Principle of Deadtime Correction

As shown above, the deadtime and afterpulsing of the detector lead to a distortion of the measured signal. However, the true profile can be recovered if a suitable correction is applied. The principle of deadtime correction [145] is based on calculating the availability of the

detector for any given bin. The availability is given by the probability that no photon was detected for one deadtime prior to the bin of interest. If τ is the deadtime in bins, then the availability of the detector in bin i is

$$a_i = 1 - \sum_{j=i-\tau}^{i-1} \frac{x_j}{N} \quad (70)$$

where x_j is the number of counts recorded in bin j after acquiring for N turns. The probability of receiving a photon in bin i was then

$$p_i = \frac{x_i}{a_i N} \quad (71)$$

since $a_i N$ is the number of turns on which the detector was available. The number of photons counted over N turns by a detector with no deadtime would then have been

$$c_i = \frac{x_i}{a_i} = \frac{x_i}{1 - \sum_{j=i-\tau}^{i-1} \frac{x_j}{N}} \quad (72)$$

This correction does not take into account 'pile-up', or the possibility that two photons arrive in the same bin. If the bins are short then the probability of two simultaneous photons is very small. However, it may not be negligible for the peak bins of each bunch. A second-order correction can be carried out to account for pile-up by again using the Poisson distribution. From equation (67),

$$p_i = 1 - e^{-C_i/N} \quad (73)$$

where C_i is the number of counts by a detector with no deadtime or pile-up.

Combining equations (71) and (73) gives

$$C_i = -N \ln(1 - p_i) = -N \ln\left(1 - \frac{x_i}{a_i N}\right) \quad (74)$$

4.4 Monte Carlo Simulation of the Detector Response and Signal Correction

4.4.1 Testing the Theoretical Detector Response at Different Count Rates

The effect of deadtime on the measured pulse shape can be calculated analytically, as was shown above. However, in order to cross-check the calculations and allow them to be extended to cover different situations (e.g. deadtime jitter, bunches of different sizes, complex filling schemes, etc.), a Monte Carlo simulation was constructed [141].

The probability of a photon arriving in any bin is defined by the filling scheme, bunch shape and the average count rate. From this probability, a lookup table of cumulative probabilities is constructed, where the cumulative probability is the probability for any given bin that no photon has arrived since time $t=0$ (defined as the start of LHC bucket 1). In reality, $t=0$ has no special significance, as the detector is always on and the acquisition can start at any point in the ring. Nonetheless, it is a reasonable approximation to begin the simulation at this point, with the detector available, since bucket 1 is immediately preceded by the abort gap. Since the abort gap is much longer than the deadtime of the detector and contains very few particles, the availability of the detector at $t=0$ is close to 1.

The arrival time of the first photon is then randomised by generating a random number q between 0 and 1 and comparing this to the table. The bin for which q is just smaller than the cumulative probability is the arrival time of the photon. This arrival time is recorded.

The bin is then incremented by the deadtime of the detector, since no photons can be detected in the deadtime. The deadtime can be a constant or can have a randomised jitter. A new random number q' is generated and is multiplied by the cumulative probability in the bin where the simulation lies. This is again compared to the table to find the arrival time of the next photon.

The last (smallest) number in the lookup table is the probability that no photon is detected for a whole turn, p_0 . If q' is smaller than p_0 , then the bin counter is returned to zero and q' is divided by p_0 , moving the simulation on by one turn. The lookup then continues as before.

These steps are repeated as many times as necessary in order to generate a histogram of arrival times. The correction algorithm described above is then applied to this histogram, and

the result is compared to the input photon probability distribution. An example is shown in Figure 49. The results can be compared to those achieved analytically and shown in Figure 46. However it should be noted that in Figure 49 the ultimate LHC filling scheme with 25 ns bunch separation is simulated, so that the bunches cannot be considered isolated.

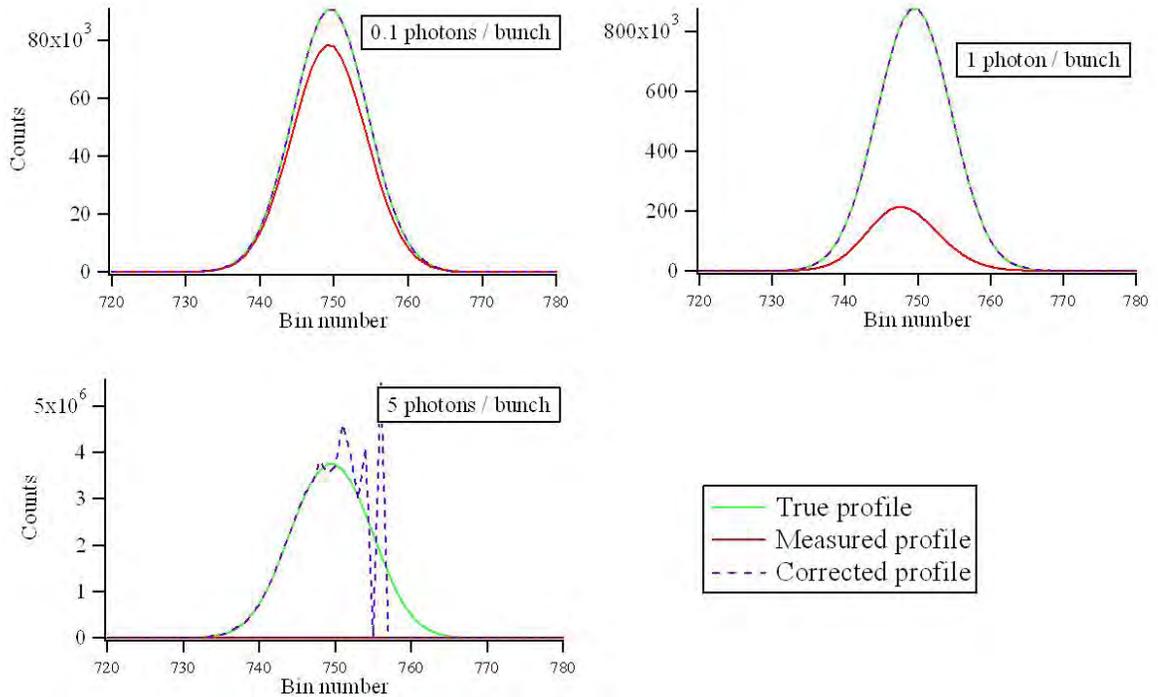


Figure 49. Photon counts against time. Results of the Monte Carlo simulation for different photon arrival rates. Bunch separation is 25ns, detector deadtime 45 ns, acquisition time 1000 seconds (1.12×10^7 turns).

Since the 25 ns bunch spacing of the nominal LHC filling scheme is shorter than the detector deadtime, the deadtime due to photons detected during any bunch affects the following bunch. This causes an oscillation in the measured bunch charge, as shown in Figure 50. The simulation results are in agreement with the calculations shown in figure 48, and the correction algorithm is able to accurately reconstruct the true bunch-by-bunch population.

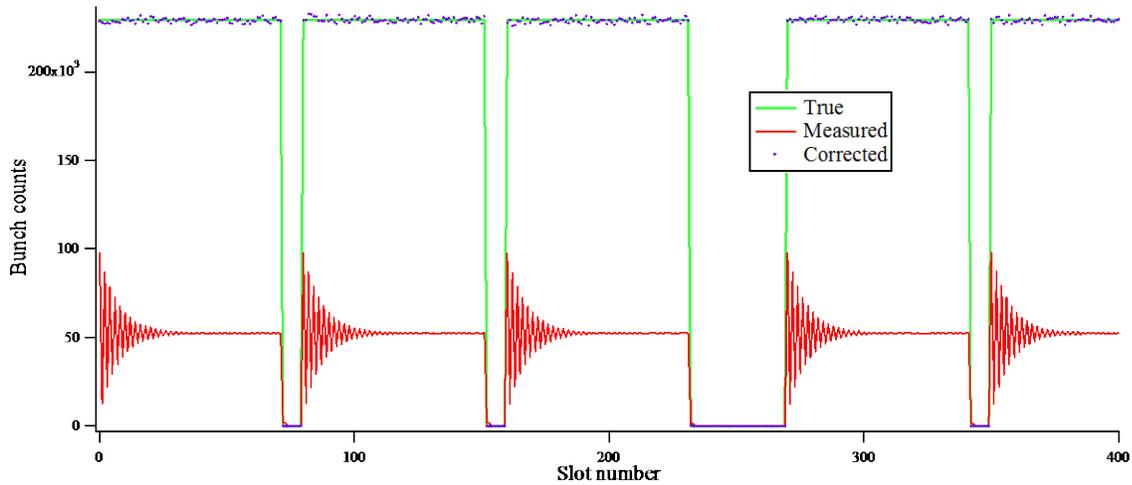


Figure 50. Photons detected per bunch against bunch number. Results of the Monte Carlo simulation showing counts integrated over each bunch for the ultimate LHC filling scheme with 25ns separation. Arrival rate 2 photons per bunch per turn, acquisition time 10 sec, detector deadtime 45 ns. The raw measurement shows a strong oscillatory behaviour since the bunch separation is smaller than the detector deadtime, but after correction the bunch currents closely match the true values.

4.4.2 Testing the Signal Correction Algorithm - Finding the Optimal Count Rate

In order to investigate the optimal count rate for the LDM system, a series of simulations were run, varying the photon arrival probability. In this case, a single isolated bunch is used. The RMS error of the histogram is found, and normalised to the total number of counts (Figure 51). As expected, the error in the raw measurement increases strongly with the count rate. The error of the corrected histogram is much smaller, and reaches a minimum around 1.4 photons per bunch per turn.

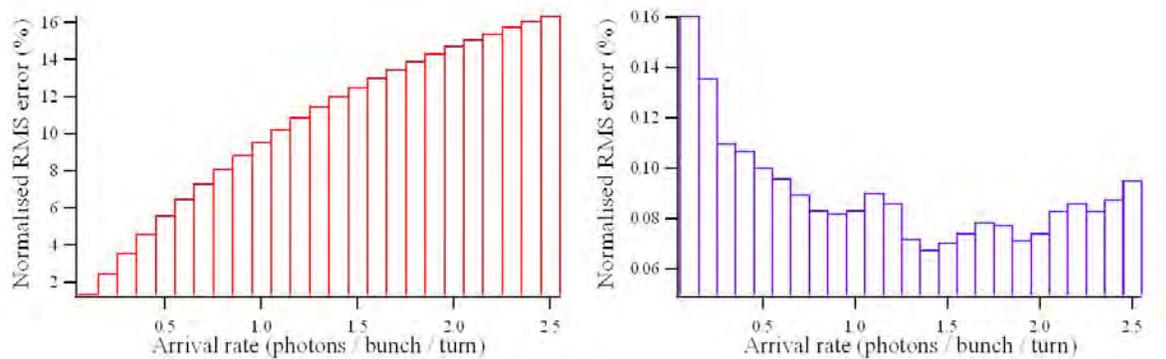


Figure 51. RMS error of the histogram, normalised to the total counts, against photon arrival rate. From the Monte Carlo simulation of a single isolated bunch. Left: raw measurement. Right: after correction.

It was derived above that the optimum count rate for visibility of a trailing satellite is 1 photon / bunch / turn. In this case, the optimum is found to be higher since the average error across the whole bunch was considered. The leading edge of the bunch profile is more accurate with a higher count rate, since the effect of deadtime is negligible there. The simulation shows that 1.4 photons / bunch / turn is the optimal trade-off between the error on the leading edge (dominated by statistical error) and trailing edge (dominated by the deadtime effect).

Since one of the main purposes of the LDM is to measure the satellite ratio, the simulation was run once more, this time using an isolated bunch with a trailing satellite. The satellite population was set at 10^{-3} of the main bunch population. The satellite ratio was calculated for each histogram and compared to the true ratio (Figure 52). As the count rate is increased, the underestimation of the trailing satellite size from the raw measurement increases, as the satellite is in the deadtime of the main bunch counts. The ratio calculated from the corrected histogram is much more accurate. As predicted, the error is minimised for a photon arrival rate close to 1 photon per bunch per turn.

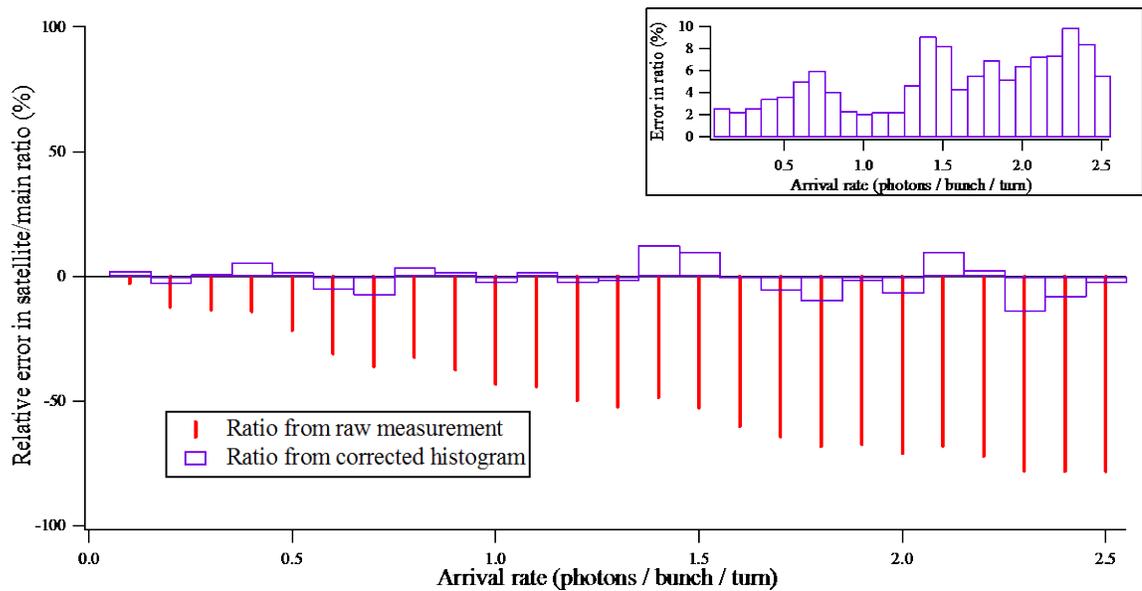


Figure 52. Error in the estimation of the satellite / main ratio for a trailing satellite, against photon arrival rate. From the Monte Carlo simulation with a single isolated bunch and acquisition time of 300 s. Inset: magnitude of the error in the ratio calculated from the corrected histogram.

4.4.3 Testing the Signal Correction Algorithm – Deadtime Jitter

In reality, the deadtime of the detector is not always the same, but is subject to some jitter. The deadtime jitter was incorporated into the Monte Carlo simulation by randomising the deadtime after each photon between the limit $deadtime_{min}$ and $deadtime_{max}$.

One effect of the deadtime jitter is to strongly increase the uncertainty of the correction algorithm when the mean deadtime is close to a multiple of the bunch separation (Figure 53). In the presence of deadtime jitter, the exact deadtime after a particular count is not known. The correction factor for any bin i therefore has an uncertainty proportional to the number of counts from bin $i - deadtime_{max}$ to bin $i - deadtime_{min}$. The uncertainty is large if a bunch falls between these limits. The uncertainty in the corrected histogram is proportional to the number of counts in bin i as well as to the uncertainty in the correction factor. Therefore, a large uncertainty in the corrected histogram will result if there is a bunch at i (large number of counts in i) and between $i - deadtime_{max}$ and $i - deadtime_{min}$ (large uncertainty in the correction factor). This can be the case only if the bunch separation is between $deadtime_{max}$ and $deadtime_{min}$.

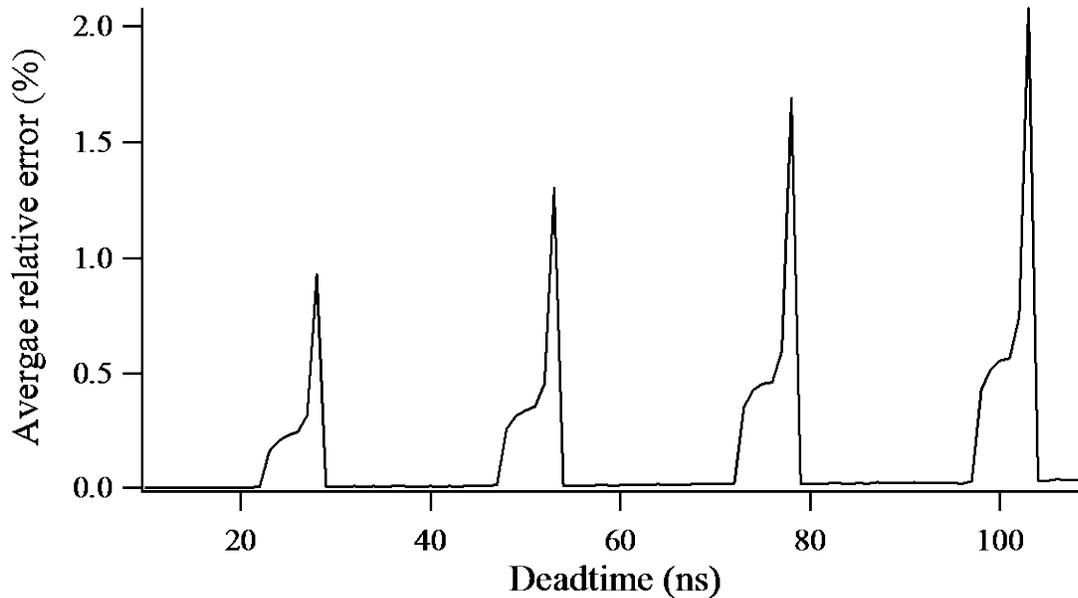


Figure 53. Relative error in the corrected histogram against detector deadtime. Results of the Monte Carlo simulation showing the average relative error of all bins in the corrected histogram versus detector deadtime. Ultimate LHC filling scheme with 25ns bunch separation. Arrival rate 1 photon per bunch per turn, acquisition time 10 sec, deadtime jitter +/- 2.5ns.

4.4.4 Testing the Signal Correction Algorithm - Pile-up Correction

The Monte Carlo simulation also identified the need for pile-up correction. Use of the simple deadtime correction algorithm given in equation (72) led to a consistent underestimate of the bunch current. The need for the additional correction for pile-up shown in equation (74) is illustrated in Figure 54. The pile-up correction becomes negligible at low count rates, but is significant at the count rates employed in the LDM.

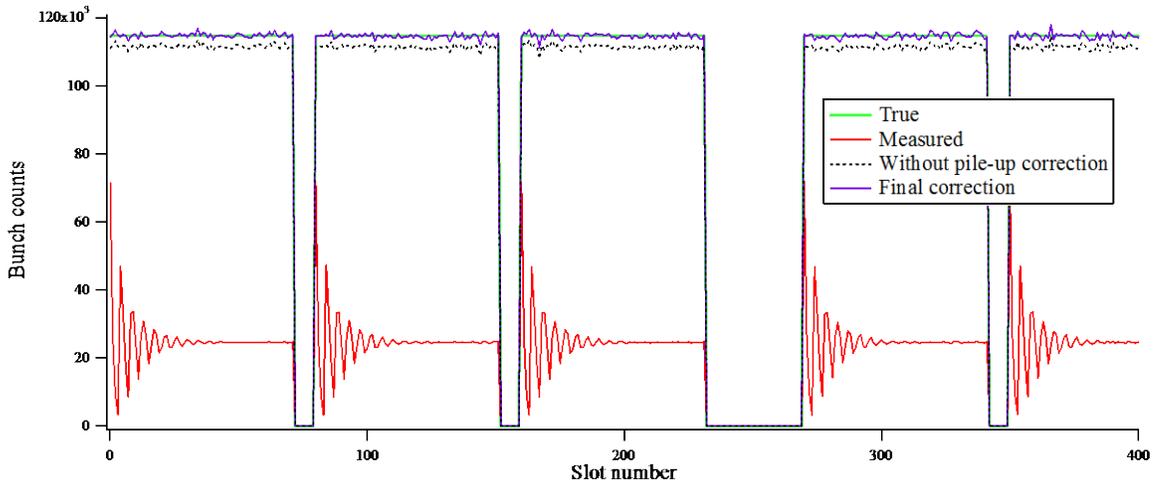


Figure 54. Photons detected per bunch against bunch number. Results of the Monte Carlo simulation illustrating the importance of pile-up correction. Counts integrated over each bunch. Simulation for the ultimate LHC filling scheme with bunch separation 25 ns, detector deadtime 77 ns, arrival rate 1 photon / bunch / turn. Without pile-up correction, the bunch current is consistently underestimated.

4.4.5 Simulating the Response of a Gated Detector

In order to optimise the measurement of the profile of the leading satellite bunches, a large count rate should be used. If the bunches are isolated, then the count rate should be chosen so that 1 photon arrives from the leading satellite(s). However, in this case the count rate for the main bunch would be very large, hundreds or thousands of photons depending on the relative size of the satellite(s). As shown in Figure 46 and Figure 49 above, the main bunch profile would be very strongly distorted, so that the correction algorithm could no longer restore the true bunch shape. Furthermore, the availability at the end of the main bunch would be close to zero, so that the trailing satellites could not be seen.

This problem could be bypassed if the detector is gated, i.e. switched off during the passage of the main bunch. In this case, an optimal count rate could be achieved for both the leading and trailing satellites. The availability of the detector would remain high because it would not be sensitive to photons from the main bunch.

A modification of the Monte Carlo program allows the simulation of the response of a gated detector. In this simplified case the detector is assumed to have zero detection efficiency when the gate is OFF, and to switch instantly between gate ON and OFF. The latter assumption, in particular, is not physically realistic. However, provided that the real switching time is smaller than the delay between the main bunch and the closest satellite, this should not affect the principle. The results of this simulation are shown in Figure 55. As expected, gating OFF the detector during the passage of the main bunch allows both the leading and trailing satellites to be seen.

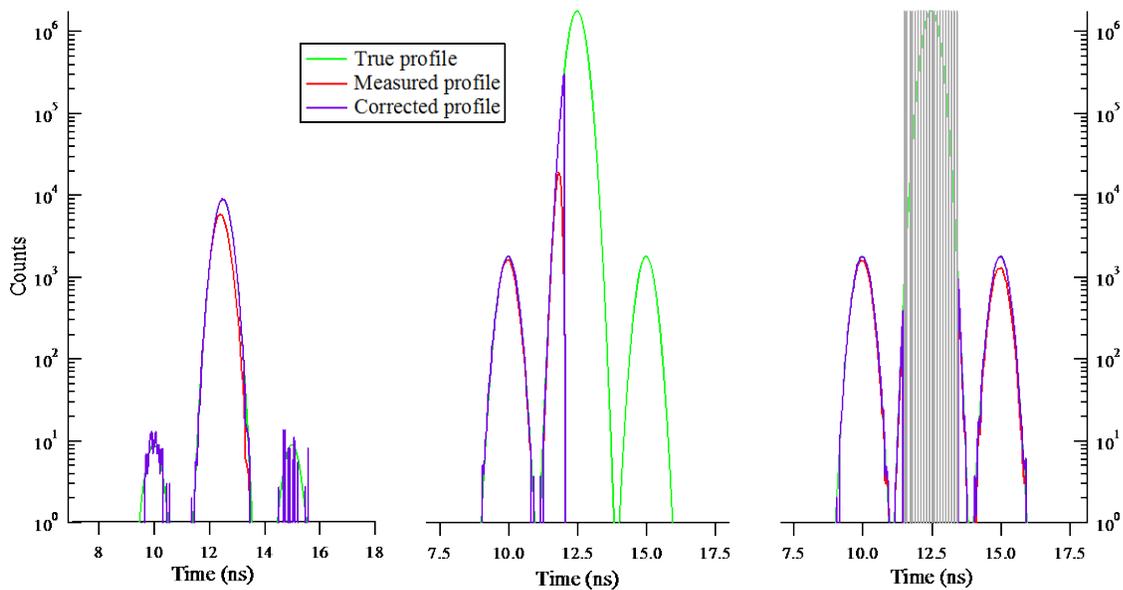


Figure 55. Photon counts against time. Results of the Monte Carlo simulation with a gated detector. An isolated bunch is centred at 12.5 ns. Satellites are located at 10 ns and 15 ns, each 10^{-3} of the main bunch. The detector deadtime is 77 ns and the acquisition time is 10 s (1.12×10^5 turns). Left, the count rate is 1 photon / turn from the main bunch, the satellite profiles are very noisy. Centre, the count rate is 250 photons / turn from the main bunch, and the detector is blind to the trailing satellite. Right, the count rate is 250 photons / turn from the main bunch but the detector is gated OFF for the central part of the bunch, and both the satellites can be measured.

4.5 Testing the Detector Response in the Laboratory

4.5.1 Instrument Response

The instrument response of the LDM detector is shown in Figure 56. This is not the electrical 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 time-to-digital converter and integrated over 10^8 cycles of a pulsed laser. The diffusion tail, the deadtime and afterpulsing are marked.

Only the shape of the response, not the magnitude, was studied. That is to say that the photon detection efficiency (PDE) of the APD was not measured. The PDE can be measured either by using a calibrated source or by using pairs of photons generated through parametric downconversion [146]. However, it is not necessary to know the PDE in order to perform the signal correction. Only the detected photons are of interest, photons which do not cause an avalanche do not influence the measurement.

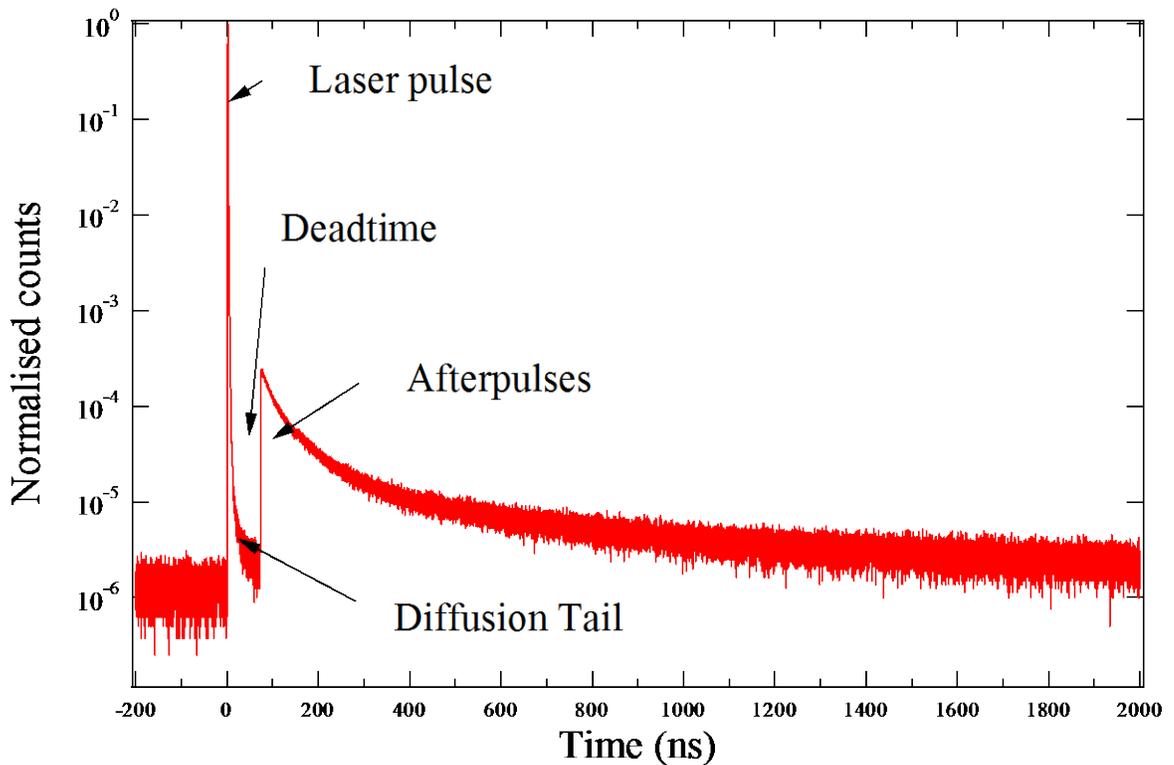


Figure 56. Photon counts against time, normalised to the maximum bin. Photon-counting response histogram of the APD, measured in the lab for 10^8 cycles of a pulsed laser.

4.5.2 Testing the Deadtime Effect

The principle of pulse-shape distortion due to deadtime and the effectiveness of the correction algorithm were tested in a laboratory experiment. The set-up is shown in Figure 57. The pulsed LED used was the PDL-800B driver with a PLS head [147] from PicoQuant. The pulse length is around 1 ns, making it possible to profile the LED pulse with the LDM detector. First, a neutral-density filter was inserted between the LED and the detector in order to reduce the photon arrival rate to less than 1 per 1000 pulses. The probability of receiving two photons from the same pulse is then very small, and the distortion due to deadtime is therefore negligible. The histogram of arrivals thus gives the ‘true’ pulse shape. Next the filter was replaced by a filter with less attenuation, so that the arrival rate was higher. As expected, the deadtime effect becomes significant and the pulse shape is distorted (Figure 58). The measured profile is both flattened and shifted towards earlier times. However, by applying the correction algorithm described above the true pulse shape is restored. The procedure was repeated with a range of neutral density filters. In order to keep the time-of-flight constant, the distance between the LED and the detector was changed in each case, to compensate the thickness of the glass filter. The correction fails when the arrival rate is so high that the probability of the detector being available at the end of the pulse is negligible.

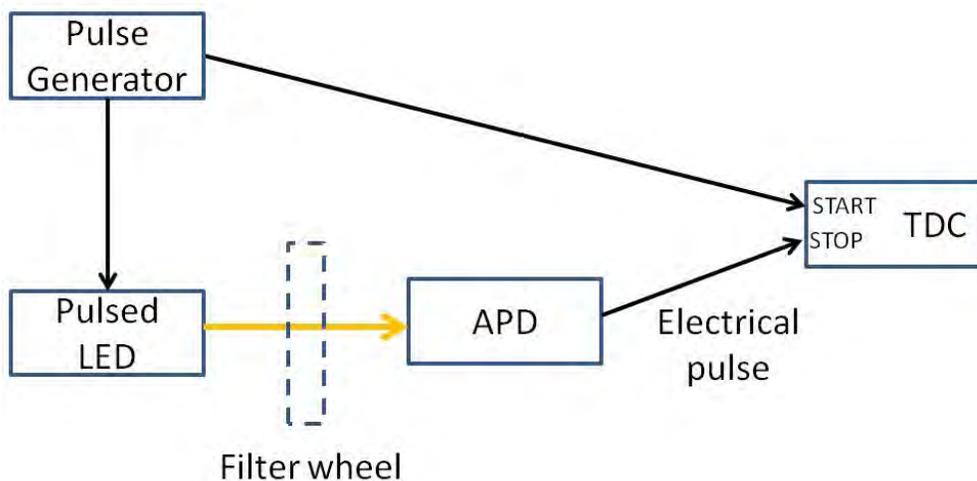


Figure 57. Laboratory set-up used to test the pulse-shape distortion due to deadtime and the algorithm used to correct for it.

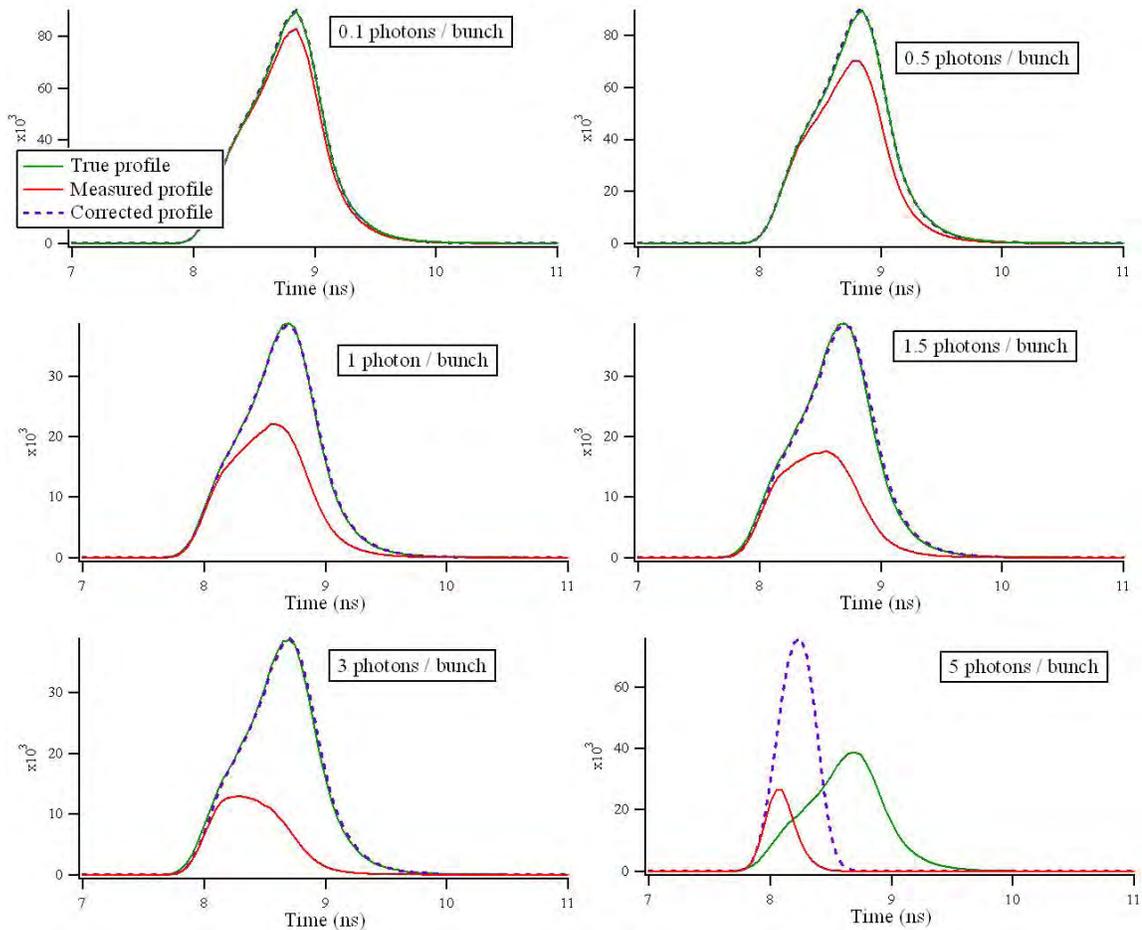


Figure 58. Photon counts against time. Pulse shape of a pulsed laser measured in the laboratory. The ‘true’ profile is in fact the normalised profile of a strongly attenuated pulse with a photon arrival rate of less than 10^{-3} , so that the effect of deadtime is negligible. The measured pulse shape is strongly distorted. However, after correction the true pulse shape is restored. However, when the photon arrival rate is so high that no photons are received from the later part of the bunch, reconstruction of the true bunch profile is impossible (lower right). The cut-off arrival rate above which incorrect reconstruction occurs will depend on the integration time.

4.5.3 Afterpulses

Although the afterpulsing behaviour of some APDs has been characterised very thoroughly [148], no comprehensive model of APD afterpulsing exists. While the afterpulse time distribution is invariably modelled by a series of exponential decays, the number of exponentials needed varies between different devices. In the LDM case the afterpulsing was found to be adequately fitted by a sum of three exponentials, as shown in Figure 59. The fit may converge to a number of quite different parameter sets, depending on the starting values chosen.

Significant variation is observed between three detectors of the same model, so that a different set of parameters is found for each of the LDM detectors.

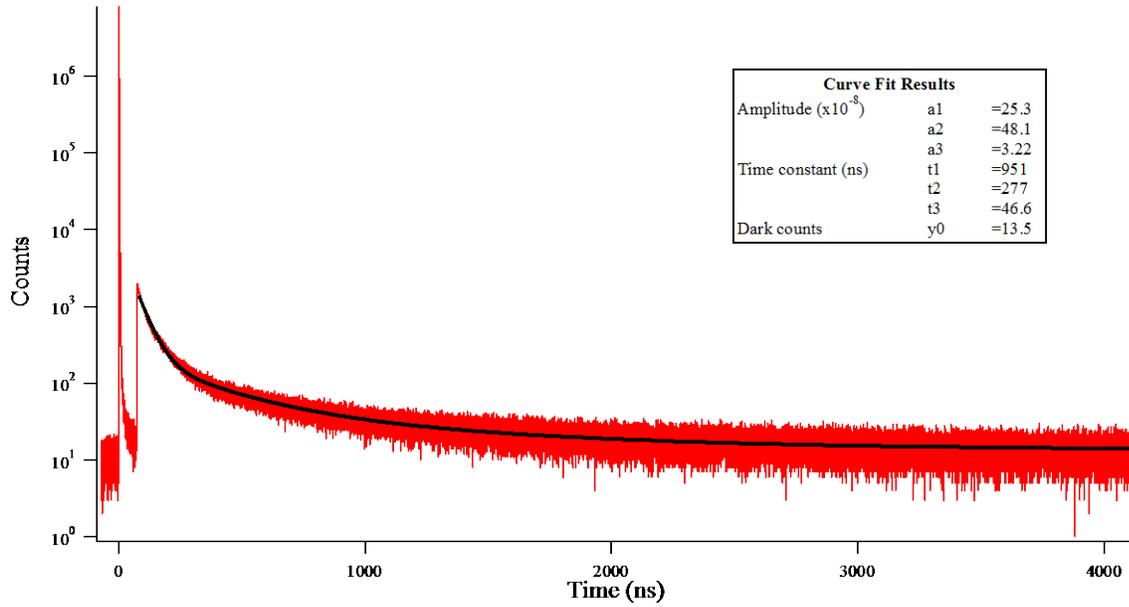


Figure 59. APD counts against time. A short pulse of photons is incident on the detector at $t=0$. All other counts are afterpulses. The decay of the afterpulsing probability is fitted by the sum of three exponentials. The parameter space contains many local minima, so the fit parameters found depend strongly on the starting values. The fit parameters for this case are shown in the box, for the fitted function $y = a_1 e^{t/t_1} + a_2 e^{t/t_2} + a_3 e^{t/t_3} + y_0$.

Afterpulses account for approximately 3% of the counts integrated over the whole histogram. It should be noted that an afterpulse is indistinguishable from any other avalanche, and so can itself generate an afterpulse. Thus, fitting a sum of exponentials to the afterpulses observed in the photon-counting response will result in an overestimate of the afterpulsing probability distribution and of the integrated afterpulsing probability. If R_{ap} is the total probability of an afterpulse occurring after any avalanche, then the ratio of the total number of observed afterpulses to the original number of photon-induced avalanches will be

$$R_{obs} = R_{ap} + R_{ap}^2 + R_{ap}^3 + \dots \quad (75)$$

It can be shown that

$$\sum_{n=0}^{\infty} x^n = \frac{1}{1-x} \quad (76)$$

Re-arranging the last two equations,

$$R_{ap} = \frac{R_{obs}}{R_{obs} + 1} \quad (77)$$

Thus, the applied afterpulse correction settings are not those given by the fit shown above but are adjusted to account for afterpulses caused by afterpulses. An example of afterpulse correction is shown in Figure 60.

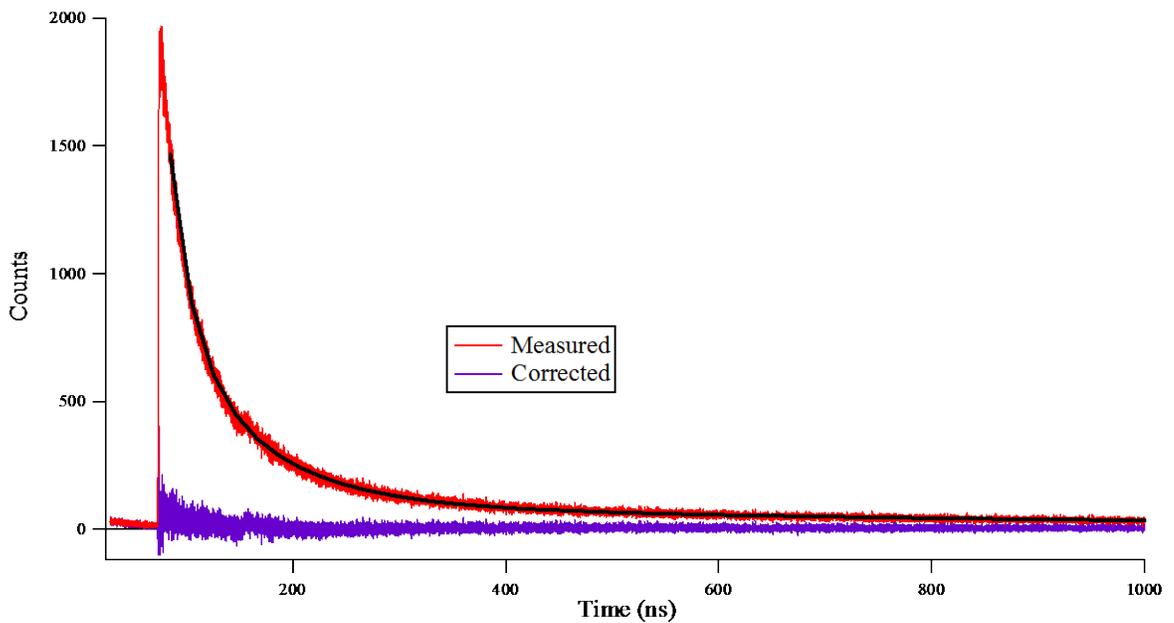


Figure 60. APD counts against time. Afterpulses are corrected by using three infinite impulse response filters (IIRs). The parameters of the IIRs are taken from the fit shown in Figure 59 and adjusted following equation (77).

The decay time of the afterpulses may depend on the internal temperature of the detector [149], since the trapped charge carriers are predominantly released by thermal excitation. The detector used for the LDM uses a Peltier cooler to keep the APD at a constant temperature. However, at high count rates the cooling power may not be sufficient and the internal temperature may rise.

In addition, the direct effect of count rate on afterpulsing is not well understood. Two models are suggested in the literature, depending on the effect of subsequent avalanches on the trapped charge carriers [150][151]. First, any trapped charge carriers are cleared if another

avalanche occurs (the 'reset' model). In this case the afterpulsing distribution depends only on the most recent avalanche, and the ratio of afterpulses to photon-induced avalanches depends on the count rate. Alternatively, subsequent avalanches have no de-trapping effect (the 'additive' model) and therefore the afterpulse probability distribution is the sum of the distributions due to each avalanche. In this case, the ratio of afterpulses to photon-induced avalanches is independent of count rate. The reset model would apply if there are only one or few trapping sites, while the additive model would apply if there are many potential trapping sites, only a small proportion of which are filled during an avalanche.

In order to determine the most suitable model for the LDM detector a measurement was carried out in the lab using a pulsed laser. The PDL-800B driver was used with an LDH-series pulsed diode laser head [147]. The laser driver was triggered externally using a signal generator. First, the signal generator produced a pair of pulses 100 ns apart, with a repetition rate of 1 kHz. Secondly, a single pulse was used, also with a repetition rate of 1 kHz. The histogram of the single pulse was duplicated, one copy was shifted by 100 ns and the two copies added together. This composite response of two single pulses closely matched the twin pulse response (Figure 61), showing that the detector follows the additive model of afterpulsing.

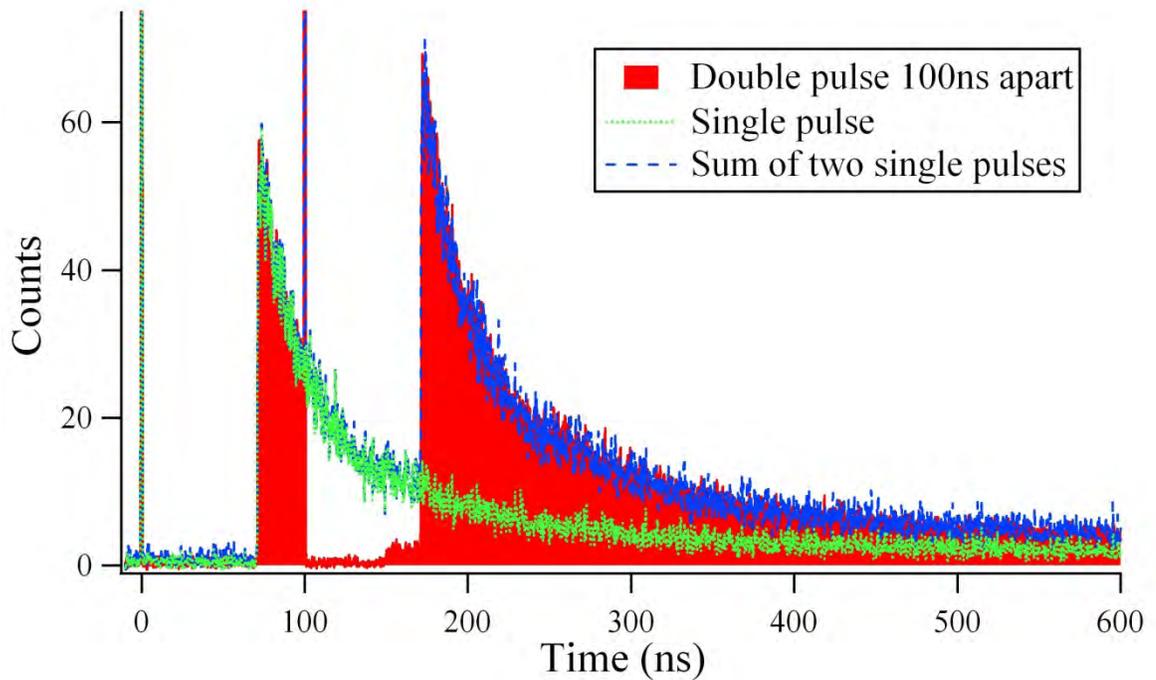


Figure 61. APD counts against time. Afterpulses are additive, i.e. subsequent avalanches do not release the trapped carriers. The solid red histogram is the response to a pair of laser pulses 100 ns apart. The peak of the laser pulses is at 100,000 counts (not shown). The green line is the response to a single pulse of the same magnitude. The blue line is the sum of the green and a duplicate of the green shifted by 100 ns.

It has been suggested that the afterpulsing rate can be reduced by periodically switching off the detector [152] but this could not be replicated with the LDM detector. The detector was operated in free-running mode, both with and without illumination, and no change in the afterpulsing or dark count rate was observed (Figure 62). A technique for reducing afterpulsing by photoionising the trapped charge carriers using a pulse of light of wavelength longer than the detection threshold of the APD has been suggested [153], but has not been investigated in the LDM case.

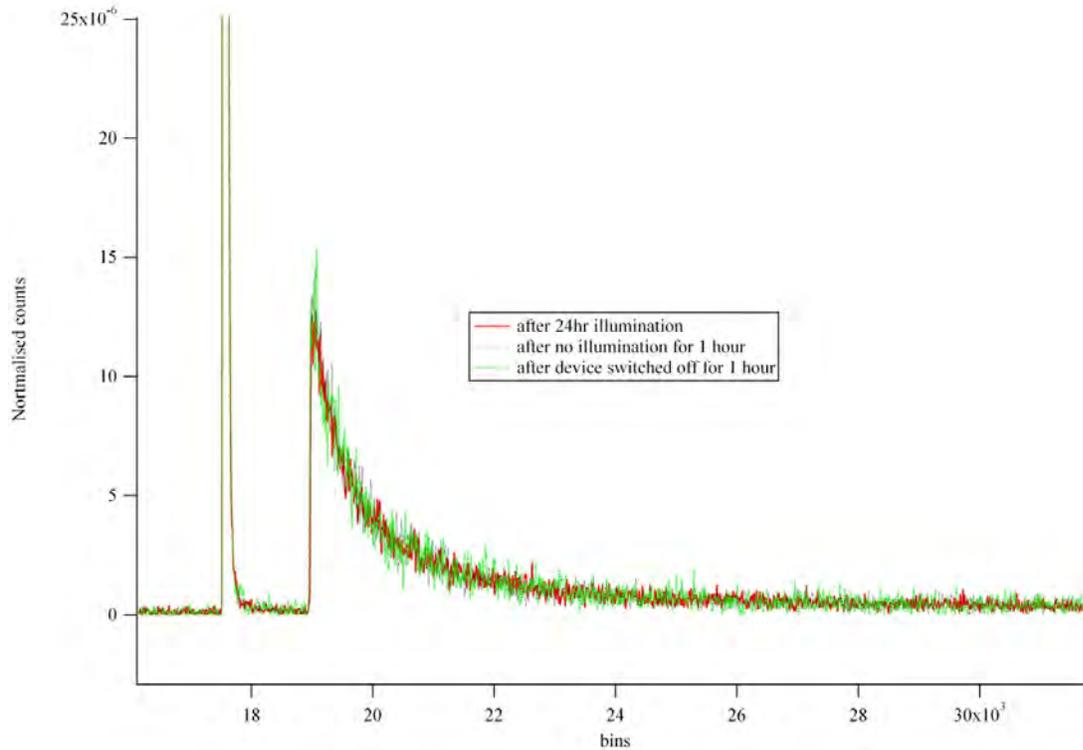


Figure 62. APD counts against time. Afterpulsing decay curves, normalised to the integral of the main peak, for the detector under different circumstances. The afterpulsing is not affected by stopping the illumination for one hour or by turning the detector off and on, so the three curves are indistinguishable.

4.5.4 Gating the Detector

One of the two APDs tested for use in the LDM, the PDM module [154] can be operated in gated mode [155]. A TTL signal is used to control the APD. When the TTL input is high, the APD operates normally. When it is low, the bias voltage across the APD is reduced and the APD is insensitive to incoming photons. Gating of the detector could be used to increase the dynamic range of the measurement, as discussed in 4.4.5. However, the detector would need to be gated off for every passing bunch. The required repetition rate of up to 40 MHz is extremely demanding.

The gated-mode operation of the PDM was tested in the laboratory. First, the module was gated with a low repetition rate. A signal generator was used to gate the module off for a short period. The timing of the pulsed laser was then adjusted so that the pulse arrives when the detector is off or on. It can be seen from Figure 63 that the pulse is successfully masked when

the detector is off, and that the pulse is unaffected if it arrives after the detector is switched back on. The application of the gate pulse itself causes some false counts, seen as two small peaks at 0 and 80 ns. The size of the false peak at 80ns, when the APD is switched back on, increases if the APD is strongly illuminated during the gated-off period (Figure 64). However, the effect is relatively small.

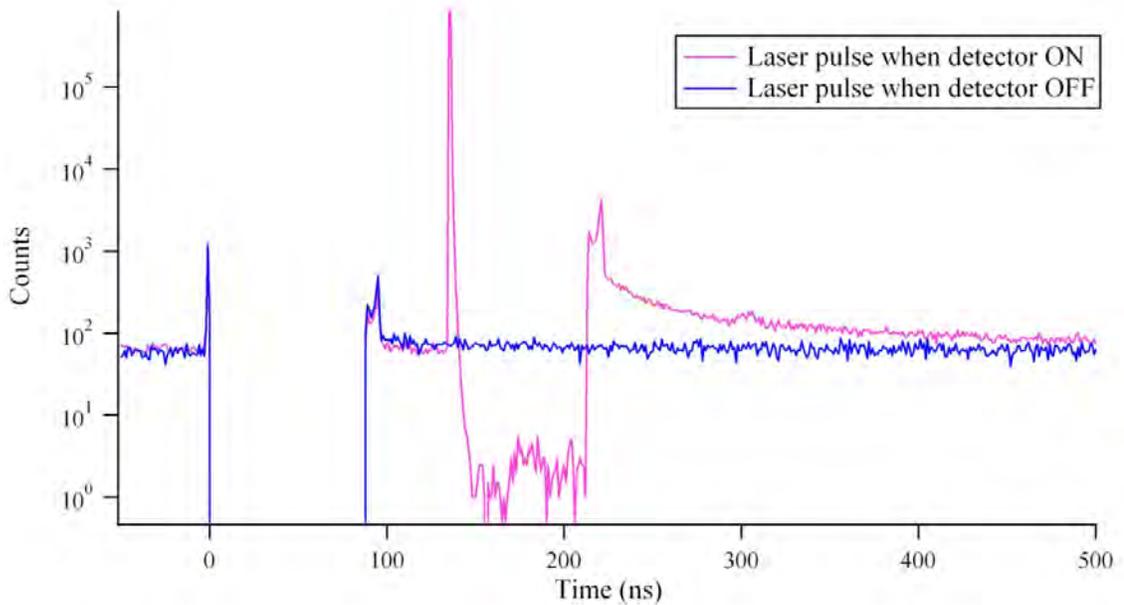


Figure 63. APD counts against time. Response of the gated APD with a gate length of 80 ns at low repetition rate of 10 kHz. When the laser pulse is timed to coincide with the gate OFF it is completely masked. Some background light was present in order to show the position of the gate.

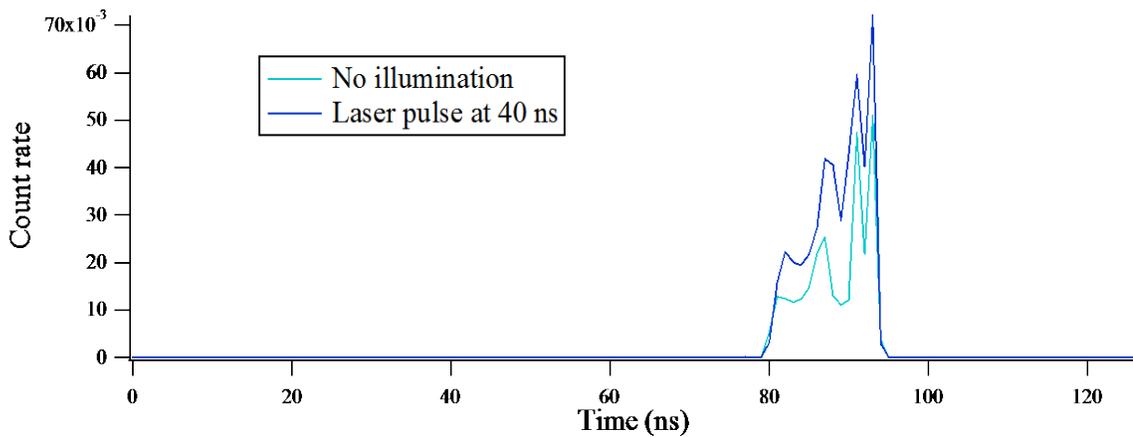


Figure 64. APD counts against time. False peak at the end of the gate-off period. The peak increases if the APD is illuminated while it is gated off.

The repetition rate was then increased to reflect the requirements for operation in the LHC. The pulsed laser was operated at 20 MHz, equivalent to the 50 ns LHC filling scheme, and the gate was turned on for 15 ns in each gap. The size of the false peaks then increases dramatically, as shown in Figure 65. When the laser pulse is weak (approximately 0.2 photons / pulse) the false peaks are of the same magnitude as the pulse. When the pulse intensity is increased to >10 photons per pulse, as would be required in the LDM case, the false peak becomes so large that the APD is always in deadtime for the rest of the gate. It is therefore concluded that gated operation of the APD is not suitable for the high-dynamic range scheme discussed in 4.4.5. An alternative scheme to optically gate the SR before arrival at the detector is discussed in chapter 7.

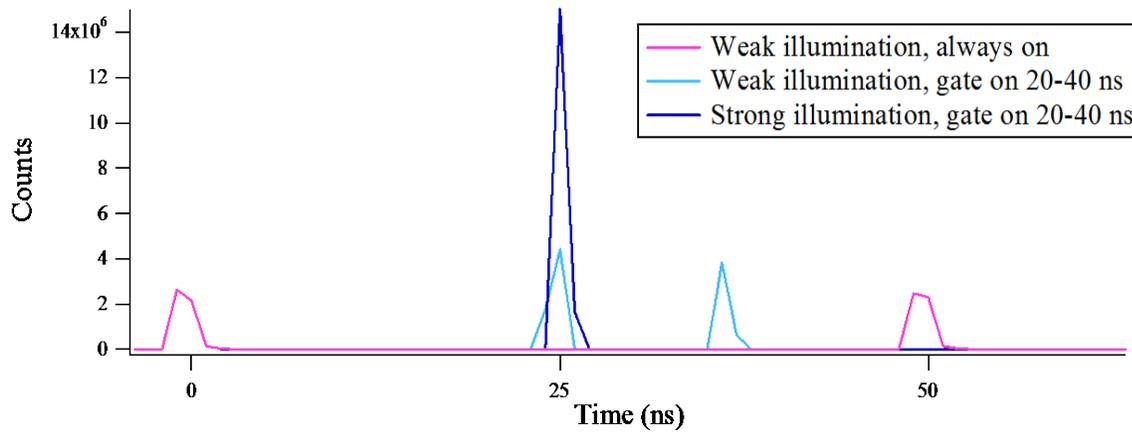


Figure 65. APD counts against time. Response of the APD when gated at a very high repetition rate.

Chapter Summary

In this chapter, Single Photon Counting is explained. The Avalanche Photo-diode is presented, along with the concepts of deadtime and afterpulsing. The effects of deadtime and afterpulsing on the LDM measurement are tested using Monte Carlo simulations, which are later benchmarked against laboratory measurements. A procedure for the statistical elimination of the deadtime bias is explained, and successfully tested against both simulated and measured profiles.

Simulations show that the use of a gated detector, which is switched off so as not to detect photons coming from the main bunches, would allow very high dynamic range measurements of the satellite and ghost bunches. However, laboratory testing of gated APDs shows that they cannot be operated at sufficient speed to apply this technique in the LHC.

5. Method

5.1 Overview of the System

The Longitudinal Density Monitor (LDM) is a photon-counting system which uses an avalanche photodiode (APD) operated in the Geiger mode to detect photons of synchrotron radiation. The setup is illustrated schematically in Figure 66. The SR source and the optical setup used to guide the light onto the APD have been described in chapter 3.

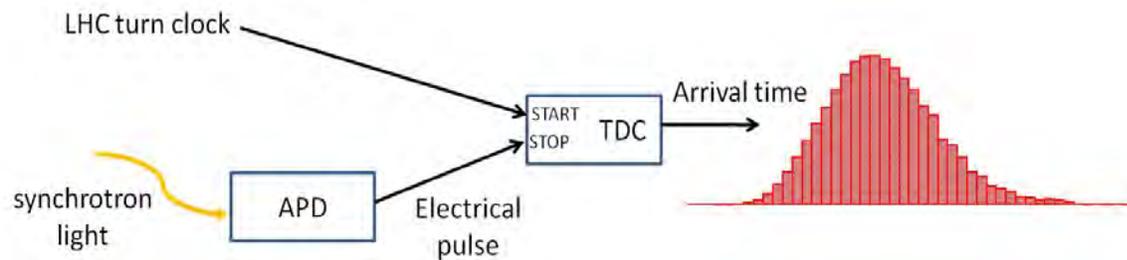


Figure 66. Schematic of the LDM showing the avalanche photo-diode (APD) and the time-to-digital converter (TDC).

Photons of SR are incident on the APD. If a photon is detected, an avalanche occurs in the APD and this is detected by the active quenching circuit. The APD module's built-in electronics then produce an output pulse of standard amplitude and shape. The output uses NIM logic levels [156].

The time-to-digital converter (TDC) [157] receives this pulse and calculated the difference in arrival time between the APD pulse and the LHC turn clock. A multi-stop TDC is used, so that the arrival times of many photons can be measured for each turn clock signal.

A histogram of arrival times is created. After many turns, this histogram approaches the longitudinal profile of the beam.

5.2 Components

5.2.1 Avalanche Photo-Diode (APD)

Only two brands of APD on the market had a sufficient time resolution for the LDM, the PDM from Micro Photon Devices [154] and the id100 from ID Quantique [158]. Table 10 shows that

while the characteristics of both devices are similar, the id100 has some advantages, most notably a smaller deadtime.

Table 10. Manufacturer's specifications of the two APDs [158][154].

	Id100	PDM
Photon Detection Efficiency (averaged over visible range)	20 %	35 %
Deadtime	45 ns	77 ns
Time resolution (FWHM)	40 ps	50 ps
Active area diameter	50 μm	50 μm
Dark count rate	20 Hz	250 Hz
Afterpulse probability	3 %	3 %

Both devices were tested in the lab and in operation in the LHC, with the PDM initially installed on beam 2 in October 2010 and the id100 on beam 1 in March 2011. In addition, a test was carried out with a fibre-coupled id100 detector which is described below.

Two issues were identified with the response of the id100 detector. Firstly, its diffusion tail was considerably longer than that of the PDM. Thus, although it has a slightly narrower time resolution by the FWHM measure, the time response was much less suitable for the purposes of the LDM. While the diffusion tail of the PDM blinds the detector to the bucket following the main bunch, that of the id100 blinds it for at least 4 buckets, meaning that none of the trailing satellites were visible. The much larger diffusion tail is clearly seen with the logarithmic scale of Figure 67.

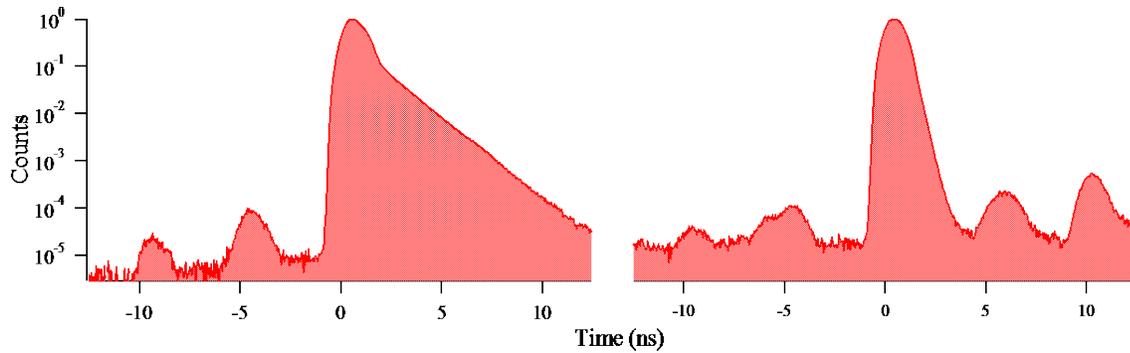


Figure 67. APD counts against time. Comparison of the time response of the two detectors. Left: Beam 1 longitudinal profile, measured with the id100. Right: Beam 2 longitudinal profile, measured with the PDM.

The second problem arises when the pulse repetition rate is similar to the deadtime of the id100. When the operating voltage is restored to the APD, there is a slight overshoot. Any photons arriving during this overshoot phase have a higher avalanche probability and the avalanches propagate faster. This is particularly problematic for the LDM case since the deadtime of the id100 is approximately 45 ns, which is close to the 50 ns bunch separation in the standard LHC filling scheme. When the id100 was used during this filling scheme, a false peak appears in the profile around 1 ns before the second and subsequent bunch of each train. This is caused by the shorter propagation delay of some of the avalanches from the main bunch.

Because of these two problems, the PDM was chosen as the detector for the LDM, and the id100 detector which had been installed on beam 1 was replaced with a PDM in August 2011.

5.2.2 Time-to-Digital Converter (TDC)

The TDC used is an Acquiris TC890 (also called Agilent U1051A) [159]. The TC890 is a multi-stop TDC so that the time stamps of many STOP pulses can be given relative to each START. For each STOP pulse, the time difference between it and the last START is recorded. The TC890 has 6 STOP channels, so that both the LDMs (corresponding to the two LHC beams) are recorded in the same TDC.

The START pulse is provided at the revolution frequency by the LHC Beam Synchronous Timing (BST) system [160]. The BST optical signal is converted into an electrical pulse by a BOBR

module [161]. The STOP signals are the electrical pulses generated by the PDM. The minimum separation of two STOPS is 15 ns; since this is less than the deadtime of the PDM it has no effect on the measurements. A minimum separation of 15 ns between STOP and START pulses is also required. In fact, some STOP pulses arriving within +/- 15 ns of the START pulse are counted, but some are lost. This means that the LDM has a reduced sensitivity for a 30 ns period in each turn. In order to reduce the impact of this 'blind spot', the START pulse should ideally arrive during the Abort Gap. However, since the same START pulse is used for both beams and they are not in phase at IP4 where the LDM is located, this condition cannot be met for both beams. The START pulse is located in the Abort Gap for beam 2, and in bucket 26,664 for beam 1.

The minimum bin width of the arrival time histogram is set by the value of the least significant bit (LSB) in the TDC time-stamp. The TC890 has a LSB value of 50 ps. The maximum time difference is 10.48 ms. This is never reached in the LDM case since the LHC revolution period is $\sim 89 \mu\text{s}$.

Each input channel has an impedance of 50Ω and is equipped with a voltage comparator. The threshold can be set between +/- 1.5 V. The common START channel is set to +1 V, while the two STOP channels are set to -0.4 V. In each case this is approximately half the pulse height. The rising edge of the pulse is steepest in this region, leading to the minimum time jitter.

The maximum count rate is limited mainly by the data read-out rate. The TC890 is equipped with two 8 Mbyte internal memory buffers which operate in ping-pong mode, i.e. one buffer is read out while the other records events; the buffers are switched when the recording buffer is almost full. Each hit is recorded as a 4-byte word, so that each buffer can contain 2×10^6 events. Timestamps can be written to the buffer at a rate of 5×10^7 events per second per channel. However this rate can be sustained only for a short time, otherwise the recording buffer will fill up before the other buffer has been read out. In this case the TDC is frozen until the buffer can be fully read out, and some events are lost.

The TC890 is built in Compact PCI (cPCI) [162] architecture and data is streamed to the computer via the cPCI bus. This allows a data throughput of approximately 100 MB / s. The maximum average count rate is thus limited to 25×10^6 events / s.

5.2.3 Beam Synchronous Timing System

The Beam Synchronous Timing (BST) system [160] is responsible for distributing the LHC turn and bunch clocks. The clocks are distributed by fibre-optic link and converted locally to an electrical pulse which is sent to the common START channel of the TDC. The optical fibre also carries a data stream which is interleaved with the clock signals and which gives information on the state of the machine.

5.2.4 Data Handling

The front-end CPU for the LDM is located in the cPCI crate alongside the TDC. It is responsible for control of the TDC and processing of the data. The CPU is a PP-712-083 produced by Concurrent Technologies [163] running the CERN Scientific LINUX operating system.

5.3 Optical Layout

5.3.1 Free-space and Fibre Coupling

Due to uncertainty over the effect of radiation on the APDs, and over the level of radiation to be expected at the BSRT location, it was initially planned to locate the APDs outside the tunnel. SR would be coupled into an optical fibre and guided to a surface laboratory or to a radiation-free underground area, such as the US45 electronics hall.

5.3.2 Fibre Coupling

A fibre-coupled APD was purchased, allowing light from a connectorised fibre to be easily coupled into the APD with a high efficiency. However, coupling of the free-space SR beam in the BSRT to an optical fibre is difficult, especially since the alignment procedure must be accomplished by remote control.

All optical fibres suffer from dispersion, i.e. a travelling time difference for photons traversing the fibre [164]. This clearly impacts the achievable time resolution of the LDM if a fibre is used. Three main kinds of dispersion occur.

Chromatic dispersion is caused by the different speed of photons of different wavelengths in the medium of the fibre. All fibres exhibit chromatic dispersion to varying degrees. Although fibres have been designed which reduce chromatic dispersion to a minimum over a certain

wavelength range [165], or cancel it with a section of fibre of opposite dispersion [166], none work over the very broad spectrum of the BSRT.

Modal dispersion is caused by the different path lengths of photons which enter the fibre with different angle or position. Such photons are then said to be travelling in different 'modes'. If the fibre core is sufficiently narrow, only one mode is possible, and there is then no modal dispersion. This is known as single-mode fibre (SMF). If the core is wide enough that more than one mode can propagate (multi-mode fibre, MMF) then modal dispersion will occur. It can be mitigated by the use of graded index (GRIN) fibre, in which the travel time difference between different modes is minimised because the difference in path length is partly compensated by the change in the speed of light in different parts of the fibre. These three types of fibre are illustrated in Figure 68. Recently, exotic types of fibres based on photonic crystal lattices have been demonstrated to have dispersion as low as 0.6 ps /nm km [165].

Finally polarisation dispersion occurs if the cross-section of the fibre is not circular. Light which is polarised in different directions with respect to the fibre will propagate with different speed. An initially circular fibre can exhibit polarisation dispersion if it is tightly coiled, so that the fibre core is stressed in one plane.

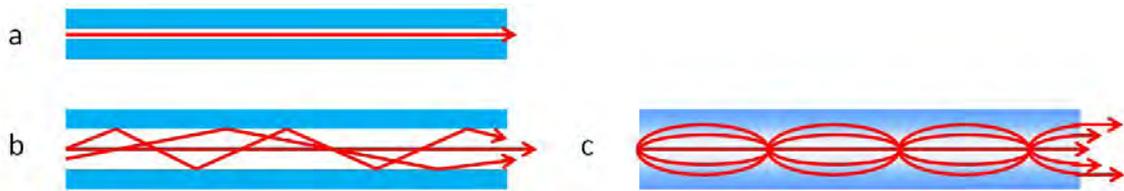


Figure 68. Modes of propagation in optical fibres. a) in single-mode fibre (SMF) only one mode is possible. b) in multi-mode fibre (MMF) light can propagate in different modes depending on the angle and position with which it enters the fibre. Light in different modes travels down the fibre with different effective velocities. c) in graded index MMF different modes exist but their effective velocities are closer.

The larger aperture of MMF makes the coupling of the SR into the fibre easier and more efficient. However, it also worsens the time resolution, due to the presence of modal dispersion.

Due to the difficulty and delay in pulling new fibres, an existing fibre was used to carry the signal from the BSRT to US45, the radiation-free underground electronics hall close to the LHC

tunnel. The fibre was a multimode fibre with a 100 μm core, and a length of approximately 50 m. A GRIN lens with an aperture of 1 cm was used to couple the SR into the fibre. The lens was mounted on a fixed support on the BSRT optical table, and the steerable mirror of the BSRT was used to align the SR beam with the lens. This produced a rather poor coupling efficiency, as the SR beam can only be steered in two axes. However, enough light was gathered to demonstrate the principle of the LDM.

A profile measured using the fibre-coupled detector is shown in Figure 69. It can be seen that the bunch shape is extremely distorted. The dispersion spreads over a range of over 15 ns.

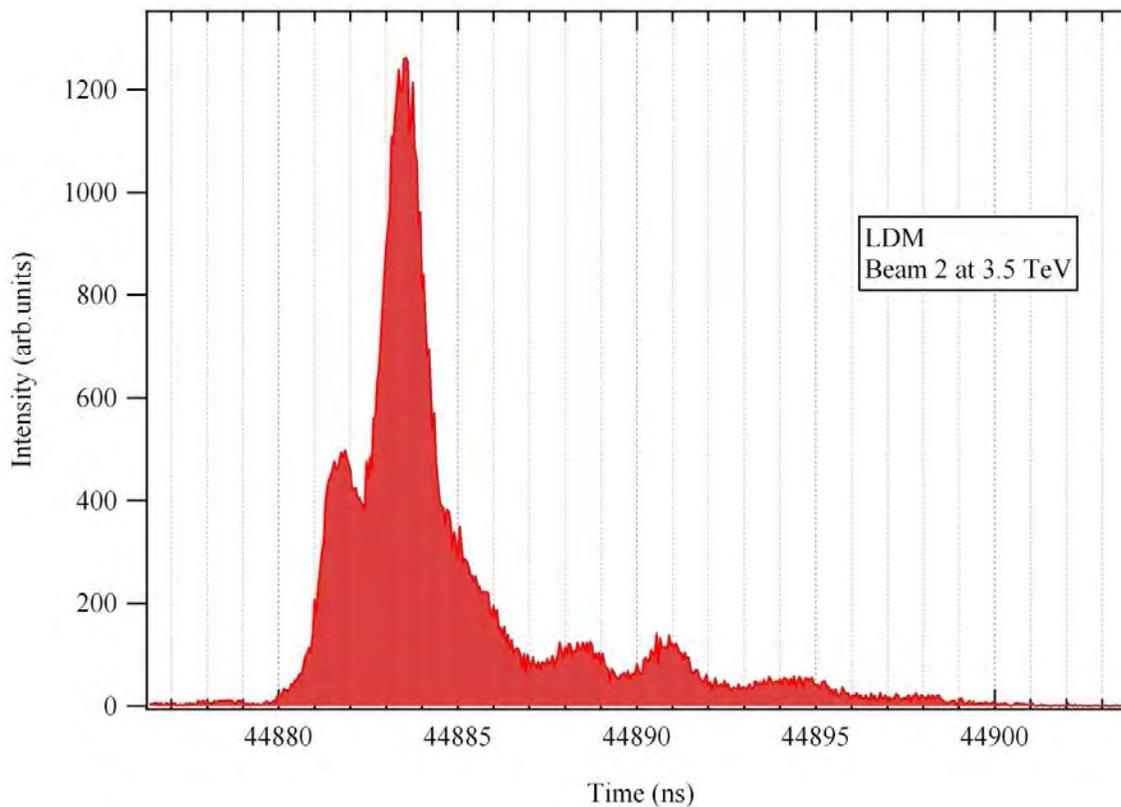


Figure 69. APD counts against time. An example of a bunch profile measured using the fibre-coupled prototype of the LDM. The bunch shape is distorted by dispersion in the optical fibre.

Problems with dispersion would probably have been overcome by the use of single-mode fibre, although this would present additional challenges for alignment. However, the radiation level in the tunnel was found to be low enough for the APD to be located in the tunnel and it was

thus unnecessary to continue with the fibre-optic coupling tests. The APD was located on the BSRT optical table as shown below, and the free-space SR beam is incident directly on the APD.

5.3.3 Location of the LDM on the BSRT Optical Table

The BSRT optical system is shown schematically in Figure 70. The focusing is achieved with two spherical mirrors. The abort gap monitor is located at the intermediate image plane, while both the cameras and the LDM are located at the final image.

An uncoated pellicle beam splitter reflects around 10% of the light incident on it towards the LDM detector module. The transmitted 90% goes to the transverse profile cameras. The filter wheels, diffuser and final focus lens (described below in section 6.4.4) are located between the beam splitter and the detector module. The final part of this beamline is shown in Figure 71.

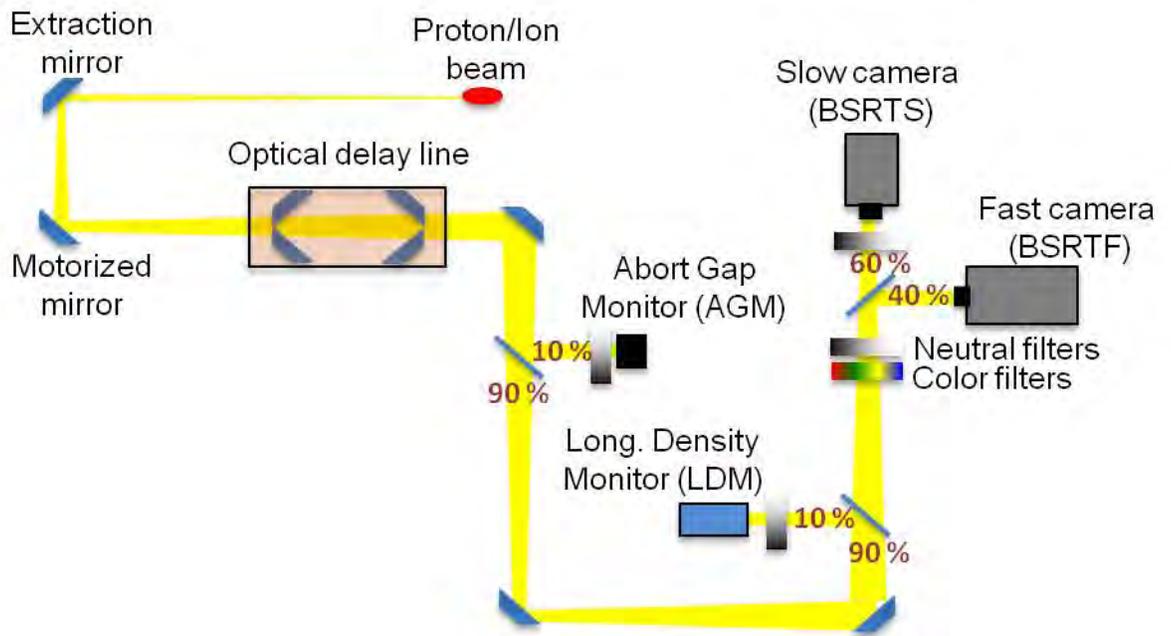


Figure 70. Schematic of the BSRT optical system showing the distribution of light to the various instruments.

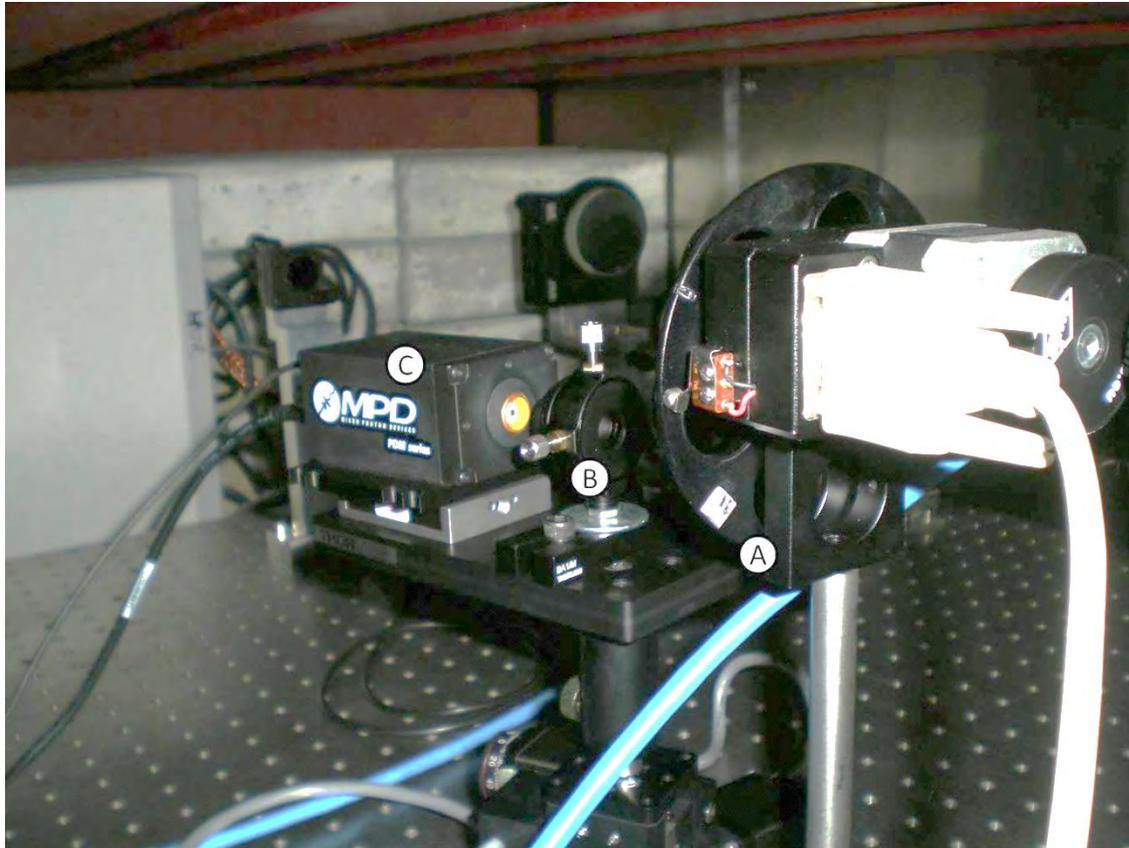


Figure 71. Final part of the LDM beamline, showing the second filter wheel (A), the final focus lens (B, described below in section 6.4.4) and the detector module (C). The filter wheels are tilted slightly off the perpendicular with respect to the light path, in order to prevent parasitic reflections arriving at the detector.

5.4 Software and Control

The LDM has been integrated into the CERN Front-End Software Architecture (FESA) [167]. The FESA server runs when the LDM front-end computer is booted. User applications may then be run on any machine as and when required, and obtain data from the FESA server. Since it runs whenever the LDM is operational, automatic controls should also be run by the FESA server. This functionality is under development, and automatic controls are currently operated by the expert application.

5.4.1 Filter Wheels

Two remote-controlled neutral density filter wheels are placed on the LDM line between the beam-splitter and the APD. The attenuation of the SR reaching the detector can thus be

controlled independently of the other BSRT instruments. The filter wheel position is automatically changed so that the count rate per bunch is kept within a set range at all times.

5.4.2 Translation Stages

The LDM is mounted on two translation stages for horizontal and vertical movement. The horizontal stage has a range of 25mm and a resolution of 1.25 μm , while the vertical stage has a range of 13 mm and a resolution of 0.02 μm . This allows extremely precise alignment of the LDM on the beam spot in order to achieve maximum coupling efficiency.

During operation, the beam spot moves in the plane of the detector. The movement is thought to be caused by heating & deformation of the extraction mirror, and by air currents caused by local heating in the BSRT. The extraction mirror is heated by absorption of short-wavelength synchrotron radiation and by image currents caused by the passing beam charge. Investigations are currently under way to determine the relative magnitude of these two effects. The result is a slow movement of the beam spot during the fill, as can be seen in Figure 72. In addition, a smaller and faster random movement of the beam spot is observed. This is thought to be caused by air inside the BSRT. The SR photons travel through approximately 12 m between the viewport and the LDM (15 m if the trombone is fully extended) and refraction caused by temperature variation of the air along this path can cause significant movement of the beam spot.

The first mirror of the BSRT can be remotely steered, and this is used to keep the beam spot within +/- 1 mm of the nominal position on the BSRT camera. The LDM is located the same distance from the beam splitter as the camera, so that movement of the spot on the camera is equivalent to movement on the LDM. Fine alignment is carried out using the LDM translation stages. A Gaussian fit is made to the BSRT profiles, and the LDM translation stages are moved in order to keep the LDM at the equivalent location of the centre of the fit.

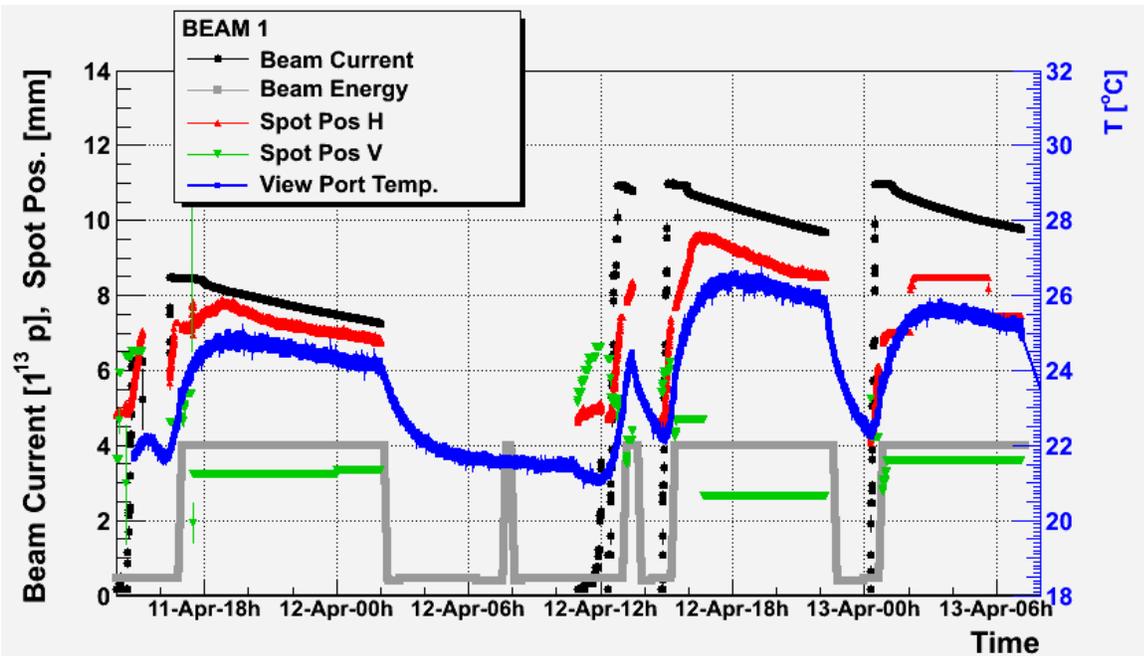


Figure 72. Horizontal and vertical position of the synchrotron radiation spot (red and green line respectively), beam current (black line), beam energy (grey line) and temperature of the beampipe window (blue line) against time. There is a correlation between temperature recorded on the beampipe window (the closest available position to the BSRT extraction mirror) and the position of the beam spot (corrected to assume no steering of the first BSRT mirror). Plot by F. Roncarolo.

5.4.3 Data Processing in the Front-End

The FESA server is responsible for reading the photon arrival times from the TDC, constructing a histogram, and performing corrections to the data. The acquisition time and the notification time can be specified, as a default 5 minutes and 10 seconds respectively. Counts are then added to the histogram continuously. At the end of each notification period the correction algorithm is applied. At the end of the acquisition time the histogram is copied to a new array, called the 'aged histogram', and the current histogram is cleared.

A re-binning of the histogram is also carried out on the front-end, in order to reduce the data transfer rate when the full histogram is not needed. A 'bucket' histogram is created by re-binning into 35640 bins of approximately 2.5 ns each, with the correct slot phase and width calculated from each aged histogram. A 'bunch' histogram is created by decimating the bucket histogram to keep only the central bin of each slot, i.e. the one which could contain a bunch.

Applications running on the CERN technical network may then subscribe to receive the full, bucket or bunch histogram with or without correction at the 'notification' rate, and the full aged histogram at the 'acquisition' rate.

5.4.4 Graphical User Interface (GUI)

The GUI allows the LDM settings and motors to be controlled, and visualises the LDM data. The GUI does not interact directly with the TDC, but only passes settings and receives data from the FESA server. Two LDM GUIs exist, an operational GUI which follows the standard CERN application model, and an expert GUI which gives greater flexibility for expert users (Figure 73 and 74 respectively).

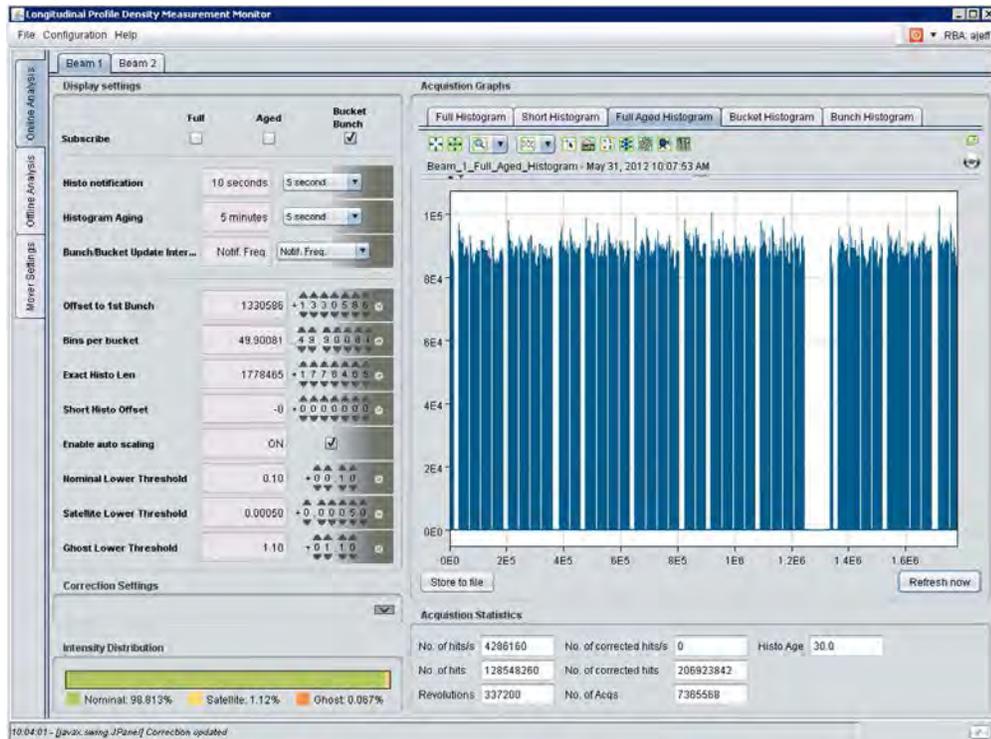


Figure 73. Screenshot of the standard LDM GUI. The full histogram or the re-binned 'bucket' and 'bunch' histograms can be viewed, and FE variables can be set.

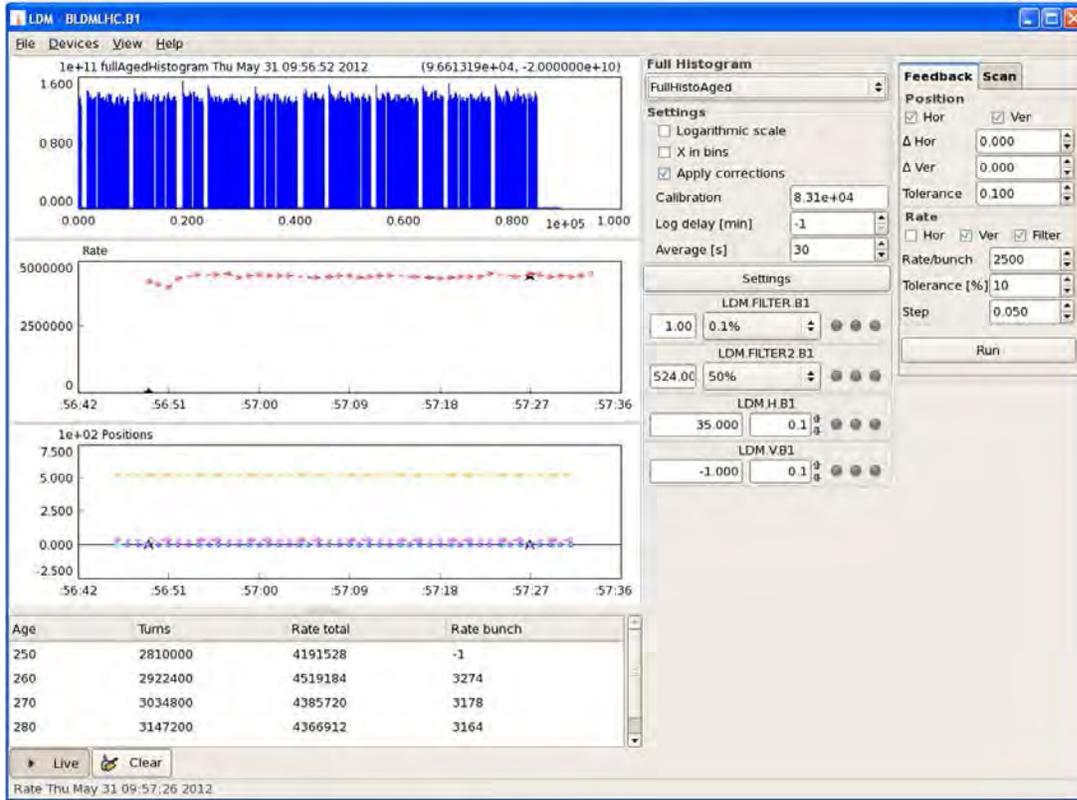


Figure 74. Screenshot of the LDM expert application. As well as viewing the various histograms, the histogram can be re-binned in bins of arbitrary width, bunch currents can be compared with the fast BCT measurements, and the LDM auto-steering and auto-filter parameters can be adjusted.

5.4.5 Logging

LDM data is logged whenever the LHC Beam Presence flag is ON. The logging frequency is defined in the FESA. As default, the bucket histogram and the positions of the LDM and the filter wheels are logged every 10 seconds in the CERN Measurements Database (MDB). The full histogram is only logged when it is aged, as a default every 5 minutes. The raw histogram is recorded, so that the correction parameters can be adjusted off-line. The histogram is too large for the MDB, and the histograms are saved directly into the CERN file system using the Self-Describing Data-Set (SDDS) format [168].

Chapter Summary

In this chapter the optical and electronic layout of the LDM system has been presented. The choice of components has been explained with reference to the specifications of the LDM, and the reasons for the choice of a free-space over a fibre-coupled layout have been given. The software used for the control and read-out of the LDM has been illustrated.

6. Results with Beam and Discussion

6.1 Example Profiles of the LHC Beams

The LDM was active from October 2010 for beam 2 and from May 2011 for beam 1. It took data for beam studies and luminosity calibration fills throughout this period. In addition, since August 2011 profiles have been automatically logged for all fills.

6.1.2 Proton Beams

When the LHC is filled with protons the LDM can measure longitudinal profiles at 450 GeV and at top energy (3.5 TeV in 2010 and 2011, 4 TeV in 2012). Although sufficient light is available at intermediate energies, conditions during the energy ramp change too quickly for the LDM profile to be meaningful.

An example LDM profile is shown in Figure 75. The PS and SPS batch structure can be seen, as well as the Abort Gap at the far right of the plot. The full LDM profile contains 1.8 million bins, which clearly cannot be portrayed on a single plot. Thus individual bunches cannot be distinguished in this figure, but are revealed by zooming on the first 3.5 μs of the profile (Figure 76). In addition, ghost and satellite bunches are too small to be seen on a linear vertical scale which shows the main bunches. They are revealed in Figure 76 using a logarithmic vertical scale. Finally, individual satellite bunches are revealed in Figure 77 by increasing the zoom level.

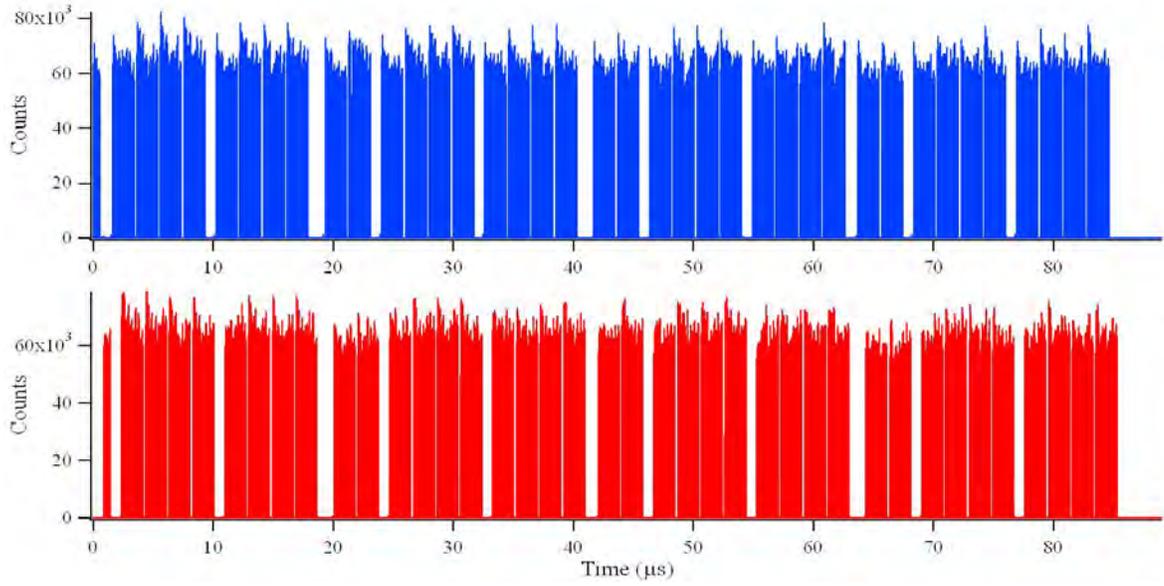


Figure 75. APD counts against time. An example beam profile taken in October 2011 with protons at 3.5 TeV. Top, beam 1. Bottom, beam 2.

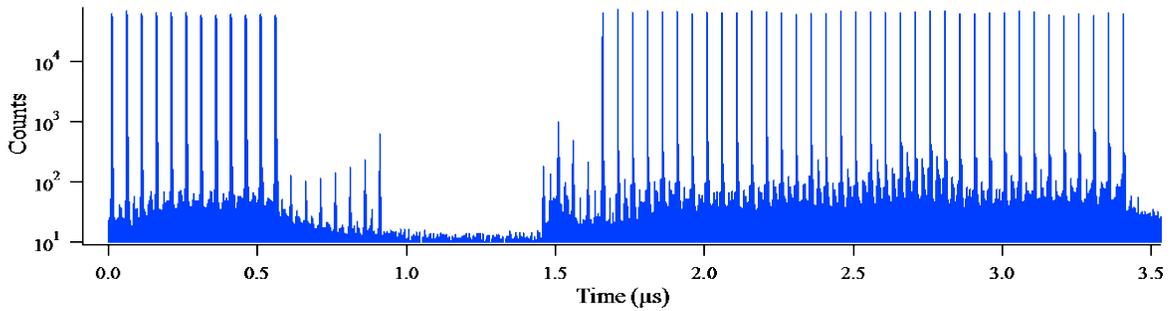


Figure 76. APD counts against time. Zoom on the first two batches of the beam 1 profile shown in Figure 75, log scale. Taken with protons at 3.5 TeV.

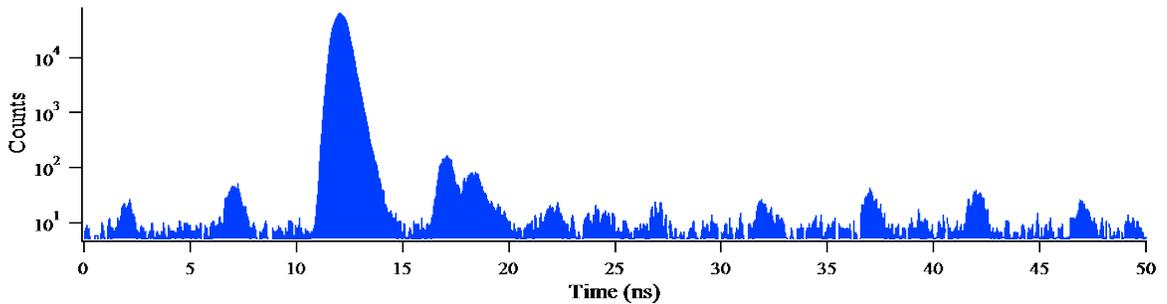


Figure 77. APD counts against time. Zoom on the first bunch of the beam 1 profile shown in Figure 75, log scale. Protons at 3.5 TeV. The additional peak at 16 ns is an accidental reflection within the instrument optics which has since been eliminated.

6.1.2 Lead Ion Beams

The LHC operated with beams of fully stripped lead ions (Pb^{82+}) in November-December 2010 and November-December 2011. Heavy ions produce insufficient synchrotron light at injection energy so LDM profiles could only be measured at top energy. The top energy for lead ions in the LHC was $3.5 \times 82 = 287$ TeV, or 1.38 TeV per nucleon. For ease of comparison with proton runs, this is usually expressed as 3.5 Z TeV.

Example LDM profiles with lead ions are shown in Figure 78. Again, the batch structure can be seen, this time with the rising intensity within each batch which is typical in lead ion fills and is caused by losses to circulating bunches during the build-up of each batch. Further detail is revealed by the logarithmic plots zoomed on the first batch and the first bunch (Figure 79 and Figure 80).

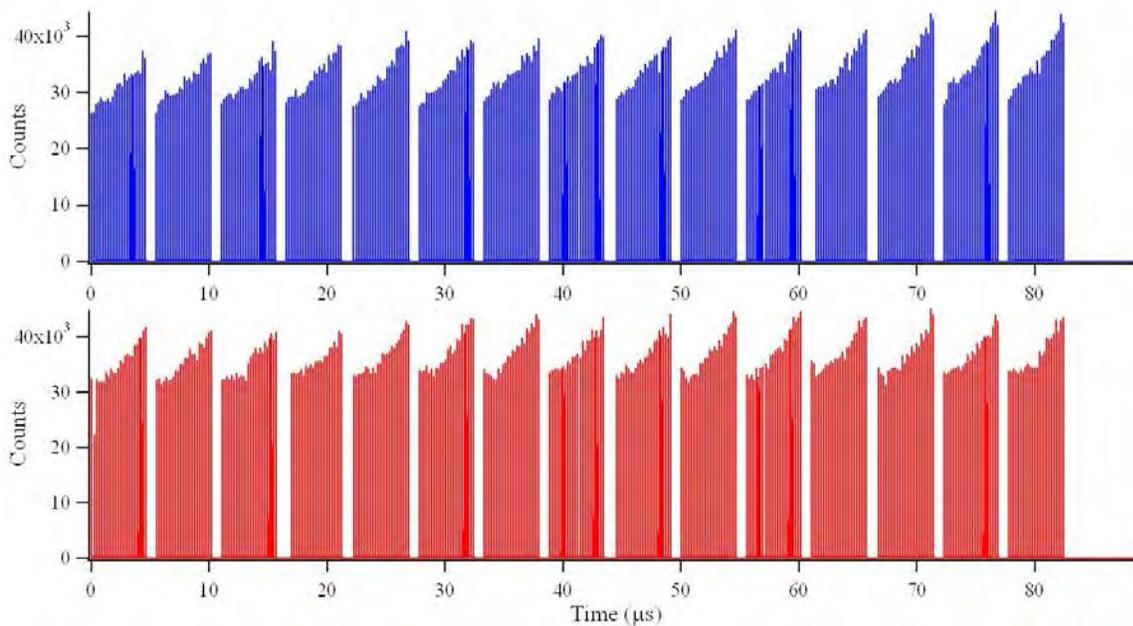


Figure 78. APD counts against time. An example beam profile taken in November 2011 with lead ions at 3.5 Z TeV. Top, beam 1. Bottom, beam 2.

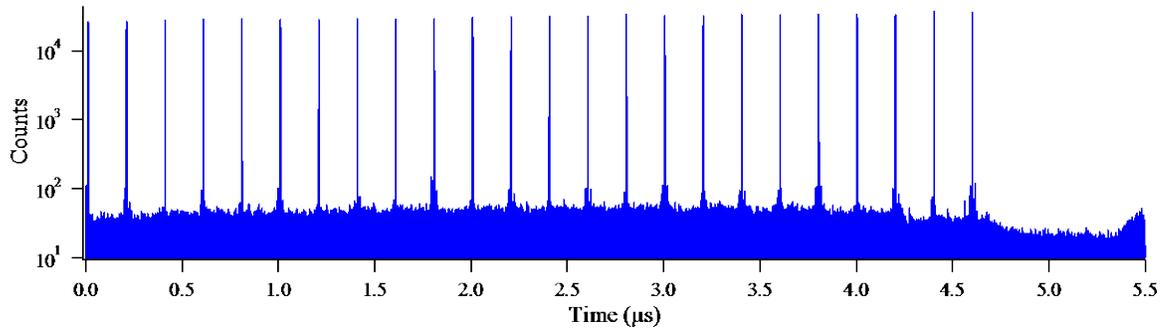


Figure 79. APD counts against time. Zoom on the first batch of the beam 1 profile shown in Figure 78, log scale. Taken with lead ions at 3.5 Z TeV.

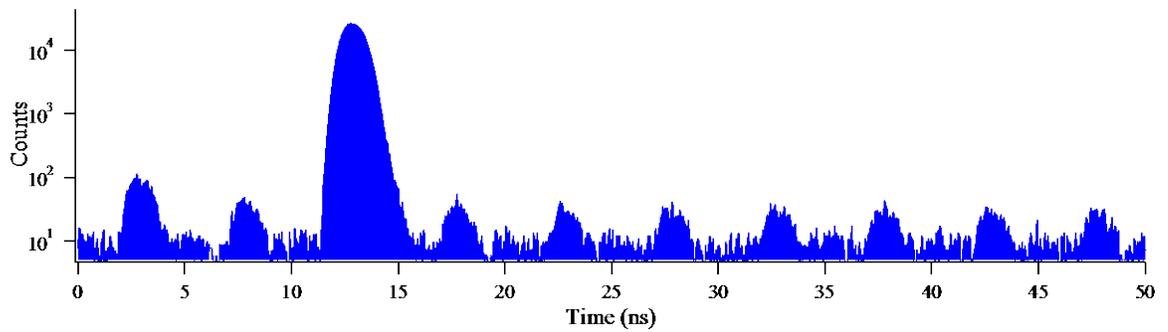


Figure 80. APD counts against time. Zoom on the first bunch of the beam 1 profile shown in Figure 78, log scale. Taken with lead ions at 3.5 Z TeV.

6.2 Analysis of the LDM Data

6.2.1 Effect of Signal Correction

The importance of the correction algorithm is illustrated in Figure 81. It can be seen that without correction being applied, the calculation of ghost and satellite fractions would be extremely inaccurate.

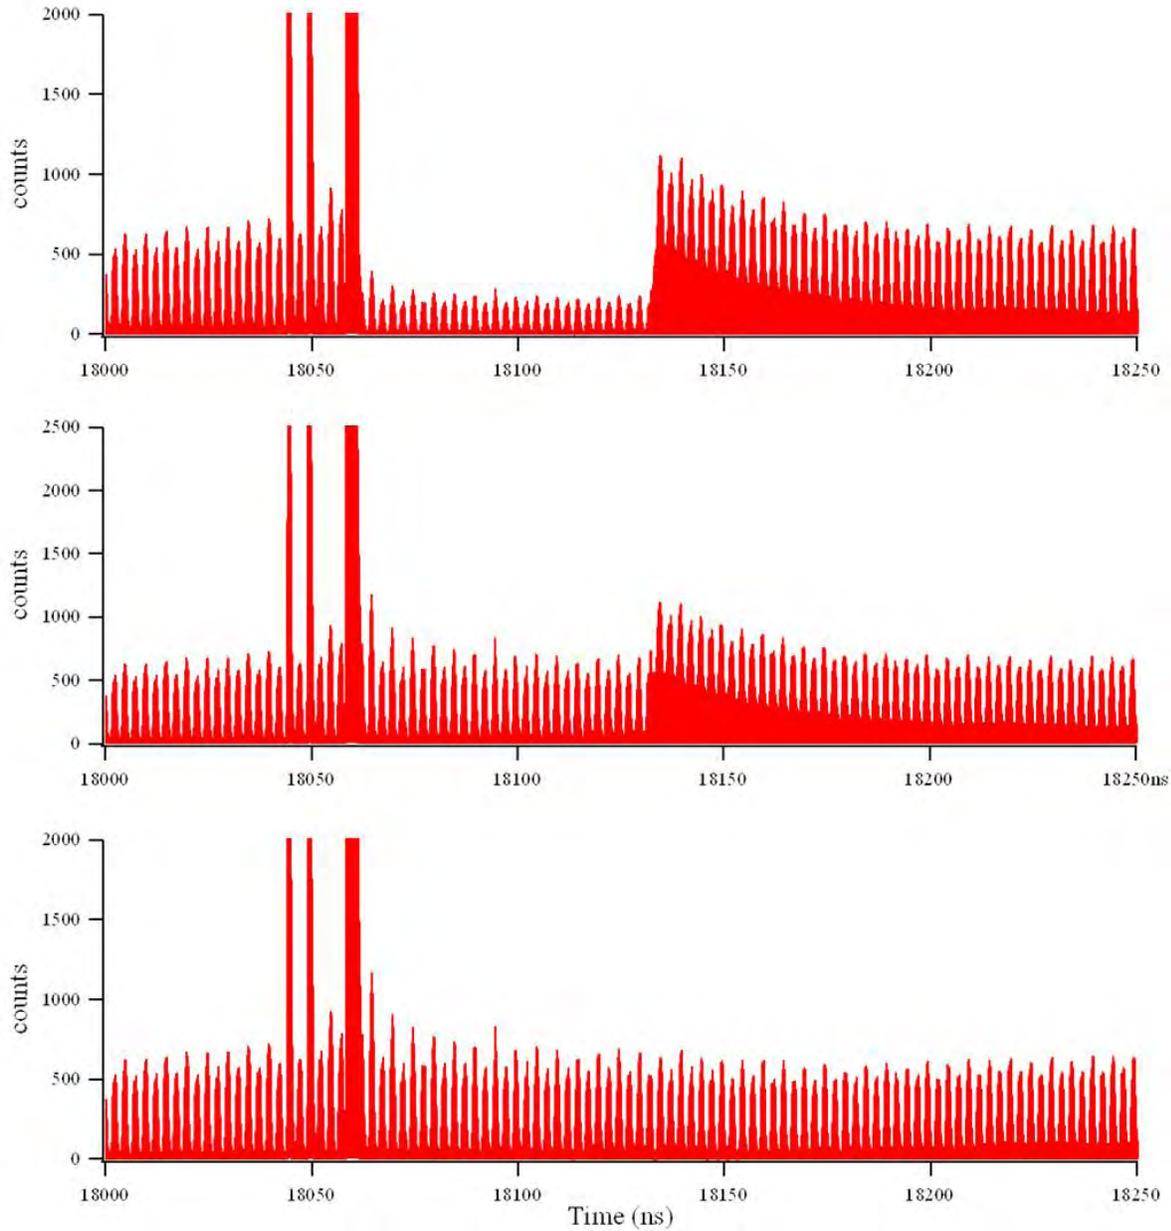


Figure 81. APD counts against time. The importance of correction in obtaining an accurate longitudinal profile. The main bunch is at 18060 ns and has a peak at 10^6 counts; two satellites preceding it are also out of range on this graph. Top, without correction. Centre, corrected for deadtime effects. Bottom, corrected for deadtime and afterpulsing.

6.2.2 Analysis Methods and Software Routines

There is an arbitrary delay in the turn clock provided to the LDM and in the LDM acquisition chain, so that the first bin of the histogram does not correspond to bucket 1. The corrected histogram must first be re-phased to match the standard bucket definitions. This requires that

the phase of bin 1 and the bucket width (or the slot width, since each slot contains exactly 10 buckets) be known.

The exact slot width 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. The first and last bunches in the histogram are found by looking for bins which are higher than half the histogram maximum, stepping through the histogram first forwards from bin 1 and then backwards from the last bin. A Gaussian fit is made for the first and last bunch so that the bunch centre can be found. Knowing the two bunch centres and the fact that each slot is approximately 499 bins wide, the exact slot width can be calculated by

$$slot\ width = \frac{centre_2 - centre_1}{round\left(\frac{centre_2 - centre_1}{499}\right)} \quad (78)$$

where *round* is rounding to the nearest integer. Knowing the slot width and the position of any one bunch centre, the bin corresponding to the official LHC bucket 1 can be calculated, since its approximate position is already known to within less than one slot. The true bunch width is always sufficiently close to 499 that equation 78 works for any two chosen bunch centres.

6.2.3 Re-Binning

Once the bucket period and phase have been calculated, it is simple to re-bin the histogram into bins of bucket or slot width. For example, in order to compare bunch currents with those measured by the fast BCT, re-binning into 25 ns slots is used, since the fast BCT response integrates over approximately a full slot, including any neighbouring satellites [169]. An example of the bunch populations calculated in this way and compared to those measured by the fast BCT is shown in Figure 82. Since there is no constant calibration factor for the LDM, this method is only useful for measuring relative bunch populations, which can then be normalised to the BCT measurements.

For other purposes, rebinning into buckets is more useful. This 'bucket histogram' can be used to identify significant satellites and estimate their population relative to the main bunch. However, it ignores the LDM baseline and thus should not be used to calculate the overall ghost and satellite fractions.

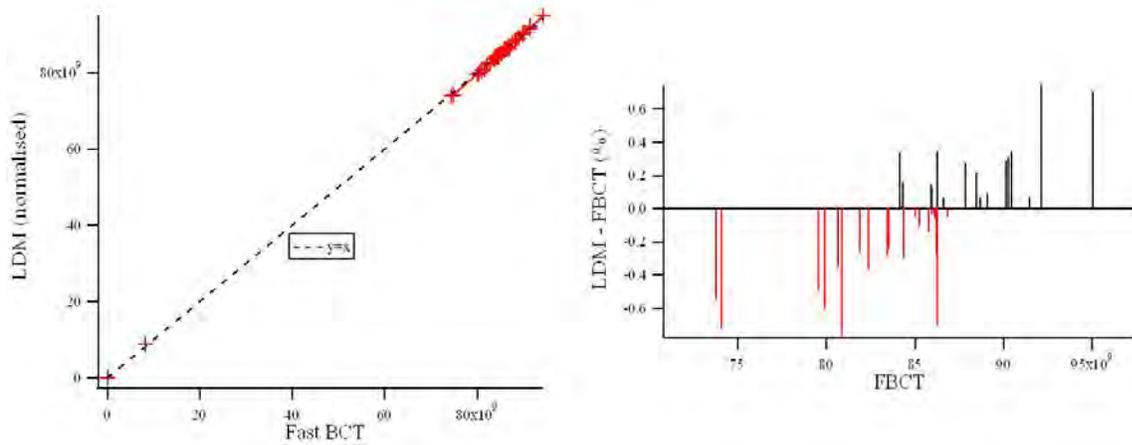


Figure 82. Relative bunch currents calculated by re-binning the LDM histogram into slot-width bins, against bunch currents measured by the fast beam current transformer (FBCT). A small non-linearity can be seen when the residuals are plotted (right). It is not clear whether this originates from the LDM or the FBCT.

6.2.4 Calculation of Average Slots

Knowing the slot width and phase also allows the 'average' slot to be calculated. A new array of 500 bins is created and each slot is added to it bin by bin. Making an average in this way greatly reduces the noise in the measurement (Figure 83).

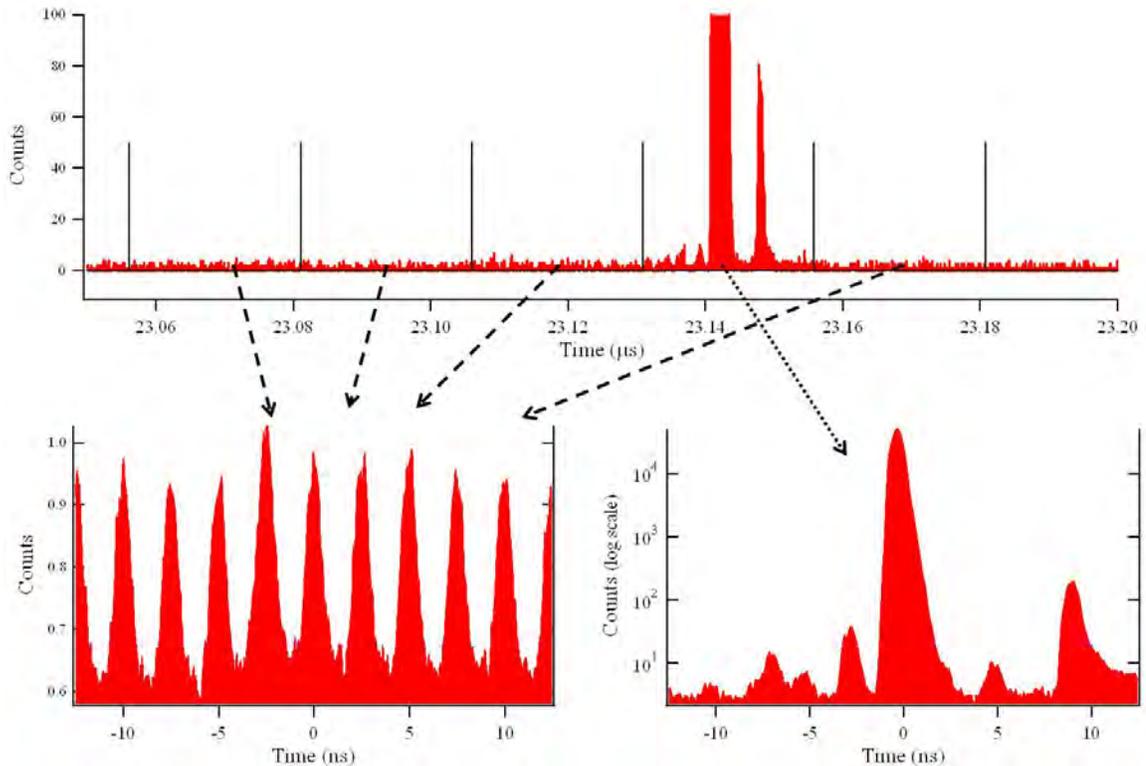


Figure 83. APD counts against time. Construction of average slots from the full ring histogram. Top, section of the full ring histogram with the slot separators marked (main bunch out of range). Bottom left, average empty slot. Bottom right, average full slot (shown in log scale to make main bunch visible).

In order to differentiate between satellites and ghosts, two averages are made, the average of all empty slots and the average of all full slots. The bunch threshold is usually set at 5%, meaning that all slots with at least 5% of the most-populated slot population are considered full. During van der Meer scans a filling scheme in which only some bunches collide in each of the experiments is generally used. In this case, separate averages can be made over any given subset of bunches.

6.2.5 Diffusion Tail

The diffusion tail of the APD causes a small proportion of counts from the main bunch to be counted much later (see 4.2.7). It can be seen in Figure 84 that the bucket immediately following the main bunch is swamped by delayed counts from the main bunch. The LDM is effectively blind to this bucket, and its contents are discarded during calculation of the ghost and satellite fractions.

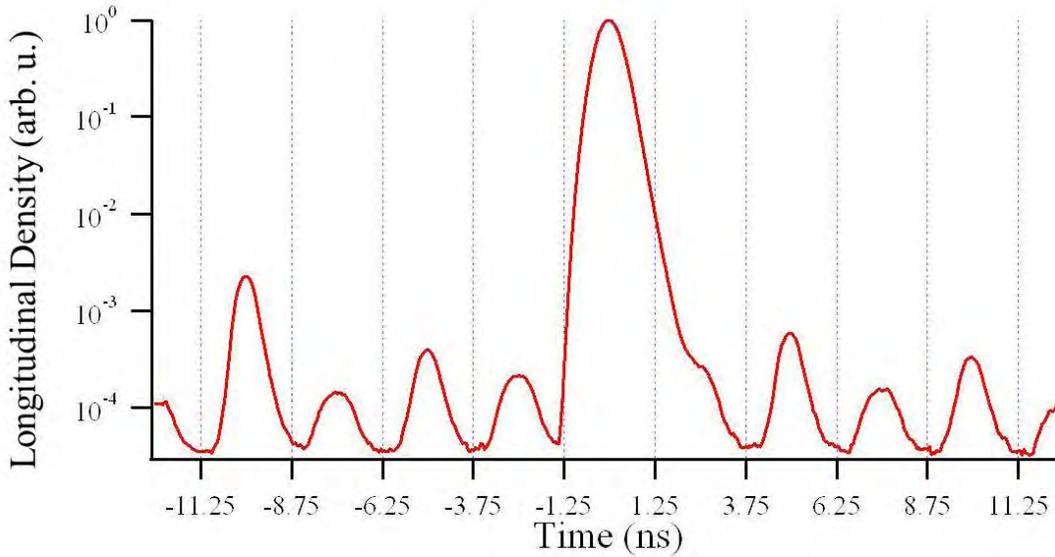


Figure 84. APD counts against time. The diffusion tail of the APD means bucket following the main bunch is swamped by delayed counts. Here it can be seen that a small satellite is present at 2.5 ns, but it cannot be accurately quantified.

6.2.6 Baseline

If all beam in the machine is bunched (i.e. all particles are longitudinally captured in an RF bucket), zero counts would be expected at the junction of two buckets (the bucket separatrix). However, it can be seen in Figure 83 that average slot histograms do not reach zero. This noise baseline could be caused by APD dark counts, by uncorrected afterpulses or by stray light in the BSRT optical setup.

The noise baseline is generally between 0 and 3 counts per bin for a normal 5 minute acquisition. The baseline level is not constant, but varies between fills, especially if the filling scheme is different. In the case of afterpulse over-correction, the baseline can also be negative.

For the calculation of bunch currents, and even for the calculation of the relative population of large satellites, the contribution of the baseline is negligible. However, for the quantification of the overall ghost and satellite fractions, the noise baseline is very significant and must be subtracted to obtain an accurate fraction.

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, generally only the average baseline is of interest for calculating the ghost

and satellite fractions. A slot-by-slot baseline can be calculated, but due to statistical noise this increases the uncertainty in the determination of the baseline.

In order to calculate the average baseline, the mean of the ten bins corresponding to the bucket separatrices in the average empty slot histogram is taken. The bucket separatrices are calculated using the slot phase and width, not by finding the minima of the histogram, as this makes the baseline less sensitive to noise.

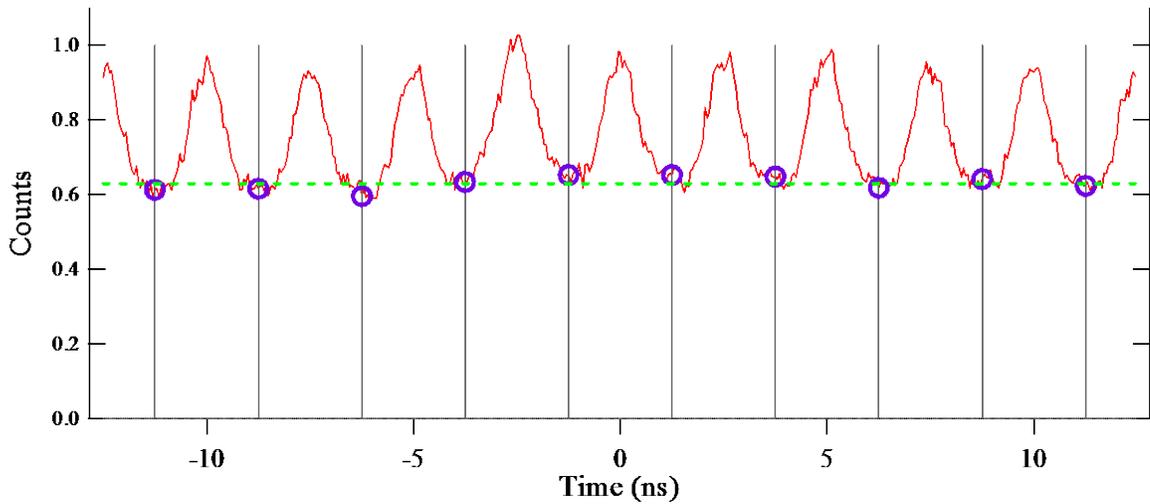


Figure 85. APD counts against time. Calculation of the average noise baseline from the average empty slot histogram. The average of the value at each of the bucket separatrices is taken.

6.2.7 Ghost & Satellite Fractions

Once the baseline level has been calculated, it is simple to calculate the ghost and satellite fractions. The baseline is subtracted from the average slot histogram, so that

$$Ghost\ fraction = \frac{e \int (H_{empty} - b)}{e \int (H_{empty} - b) + f \int (H_{Full} - b)} \quad (79)$$

where H is the average slot histogram, b is the noise baseline, e is the number of empty slots and f is the number of full slots.

To calculate the satellite fraction, the integral of the average full slot histogram without the main bunch or the following bucket is calculated. This is then multiplied by 9/8 in order to compensate for the bucket which is discarded:

$$Satellite\ fraction = \frac{9}{8} \times \frac{f \left[\int_{-12.5}^{-1.25} (H_{full} - b) + \int_{3.75}^{12.5} (H_{full} - b) \right]}{e \int (H_{empty} - b) + f \int (H_{Full} - b)} \quad (80)$$

where the limits of the integrals are time expressed in ns, assuming that the main bunch is centred at $t=0$.

6.2.8 Bunch Current Normalisation

The measurement of the absolute luminosity of the LHC is of great interest to the experiments since it allows the cross-section of different reaction processes to be determined. Knowledge of the absolute cross-sections can be used to constrain QCD models of proton-proton interactions and potentially to detect new physics beyond the standard model [170].

The absolute luminosity calibration for the LHC interaction points is carried out principally through van der Meer scans [171]. During the van der Meer scan one of the beams is slowly displaced with respect to the other, while the reaction rate of one or more common interactions is measured by the experiments. By scanning the beams across each other in this way, the contribution of beam size and shape in the luminosity calibration is eliminated [172]. The absolute luminosity can then be determined provided that the number of particles in the colliding bunches is known.

An alternative method involves direct measurement of the beam shape and overlap at the interaction point by vertex imaging [173]. However, the bunch current product must still be known in order to determine the luminosity.

Bunch-by-bunch relative current measurements in the LHC are made by the fast BCT, which is then normalised to the absolute total current measured by the DCCT, as described in 1.2.2. The DCCT measures all charges circulating in the machine, while the fast BCT is blind to ghost bunches, and this must be taken account of in the normalisation. In addition, the bunch current measured by the fast BCT is the integral of the current measured over one slot, and thus includes the satellite bunches. These satellite bunches contribute very little to the luminosity since they either cross the main bunch of the opposite beam away from the centre of the experiment or, depending on the crossing-angle, do not cross the main bunch at all.

It is thus essential to measure the population of the ghost and satellite bunches during the van der Meer scans. The uncertainty in this measurement is a major contribution to the total uncertainty in the LHC luminosity. For the luminosity determination using the earliest LHC van

der Meer scans in 2010, it was estimated that the uncertainty in the bunch current measurement was twice as large as all other sources of uncertainty put together [170]. In this case, the bunch current uncertainty was dominated by the DCCT total current measurement, and not by the ghost charge contribution. However, due to the dramatic improvement of the DCCT accuracy [49] and the use of more bunches during the van der Meer scans, the uncertainty in the ghost charge is now of the same order as the uncertainty in the total current [174]. In the lead ion fills it is the dominant contribution to the bunch current uncertainty.

Since its commissioning the LDM has been in use for all van der Meer scans. The measured ghost and satellite fractions are shown in Table 11 and Table 12 respectively. The values shown are averaged over all bunches and over the whole period of stable beams. The errors shown are estimates of the uncertainty at the 68% confidence level. The sources of uncertainty are discussed in 6.4.2 below.

Prior to the deployment of the LDM, the ghost and satellite charge was estimated using data from the experiments. However, the addition of LDM data has been used to reduce the uncertainty in the ghost and satellite charge and thus improve the overall determination of the absolute LHC luminosity [175].

Table 11. Percentage of beam population classified as ghosts, for fills used for van der Meer scans. Ghosts are defined as charge outside filled slots. The van der Meer scans in December 2011 were split over two fills.

	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	0.2	(+0.2 / -0.04)	0.4	(+0.4 / -0.1)
Oct 2011	0.7	(+0.7 / -0.2)	0.7	(+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 12. Percentage of beam population classified as satellites, for fills used for van der Meer scans. Satellites are defined as charge in the filled slots but outside the filled bucket. The van der Meer scans in December 2011 were split over two fills.

	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)	0.1	(+/- 0.03)
Oct 2011	0.2	(+0.06 / -0.08)	0.6	(+0.2 / -0.1)
Dec 2011	0.4	(+/- 0.1)	0.3	(+0.1 / -0.08)
Dec 2011	0.3	(+0.08 / -0.06)	0.3	(+0.08 / -0.06)

6.2.9 Comparison with Results from the Experiments.

No other LHC instrument has the sensitivity to cross-calibrate the LDM satellite measurements. However, the satellite bunches generate collisions in the interaction points and this can be detected by the LHC experiments. Satellite-Main collisions, that is collisions caused by the crossing of a satellite from one beam with the main bunch of the other, occur off-centre in the detectors due to the different timing of the satellite bunches. The proportion of collisions at different locations within the detector can therefore be used to determine the relative satellite population. Satellite-Satellite collisions occur both centred and off-centred, but with negligible rate, since the probability of a collision is proportional to the product of the populations of the two colliding bunches.

The location of collisions within the experiment can be determined by direct imaging of the luminous region (vertex reconstruction) or by comparing the arrival time of collision products at detectors at each end of the experiment. In ATLAS the latter has been applied using the Minimum Bias Trigger Scintillators (MBTS) [176] and is shown in figure 86 for a lead ion fill in November 2010. This technique can only be applied to the heavy-ion run since the high pile-up in the detectors during normal proton physics makes it impossible to determine the location of

any one collision using this arrival-time difference technique. The distribution of collisions at the interaction points is equivalent to the convolution of the longitudinal profiles of the two beams, averaged over all colliding bunches. At this time, the LDM was only operational on one beam, so the convolution of the LDM profile was performed with the same profile time-reversed, that is assuming that the profile of both beams is the same. The comparison is shown in Figure 86.

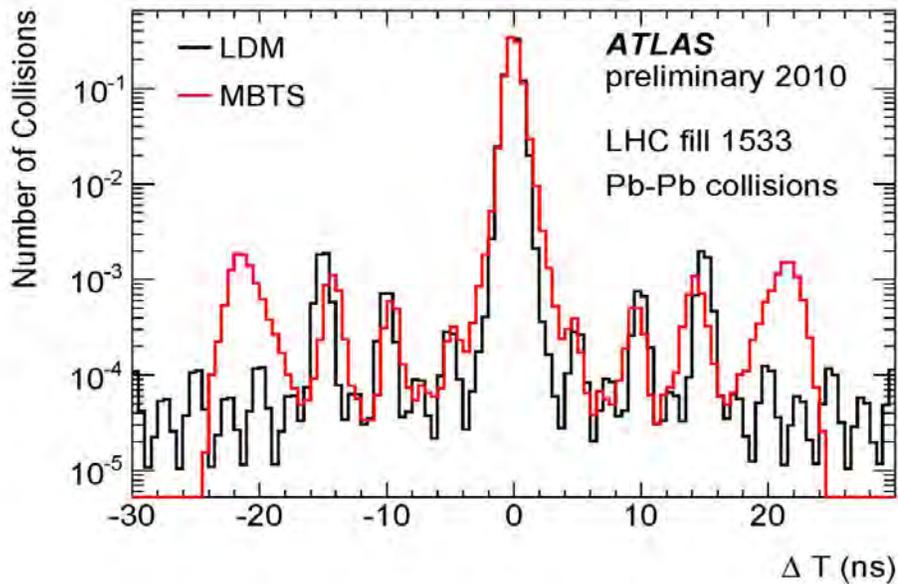


Figure 86. Comparison of satellites measured by the LDM and the MBTS. The red histogram shows the difference between the arrival time of hits in the Minimum Bias Trigger Scintillators (MBTS) on the two sides of ATLAS, requiring that at least 14 of the 16 counters be hit. This is shown for LHC fill 1533 during the 2010 Heavy-Ion run. The main peak corresponds to in-time collisions; the small peaks are interpreted as collisions of a main bunch from one beam with satellite bunches from the other beam. At the edges ($|\Delta T| > 18\text{ns}$) there is a large contamination from beam-halo events which masks any satellite collisions. Particles produced upstream of the detector which travel parallel to the beamline from one side of the MBTS to the other side have a transit time of about 21ns since the distance between the two MBTS detectors is 7.2m. The black histogram shows a measurement made by the LHC Longitudinal Density Monitor (LDM), with an integration time of 15 minutes during the same fill. Only the beam-2 LDM was available at this time and so the black histogram shows the convolution of the beam-2 longitudinal profile with the same profile reversed in time, as a proxy for the beam-1 profile. The longitudinal distribution of collisions in the detector mirrors the convolution of the longitudinal profiles of the two beams.

6.2.10 Enhanced Satellites

In October 2011 a new kind of LHC filling scheme was trialed, in which satellites spaced at 25 ns from the nominal bunches were deliberately enhanced by adjustment of the bunch-splitting process in the LHC injector chain [177]. The filling scheme is arranged such that main bunches collide with main bunches in the high-luminosity experiments (ATLAS, CMS, LHCb) while main bunches collide with enhanced satellites in ALICE, which requires much lower luminosity. By tuning the enhancement of the satellites, the ALICE luminosity can thus be tuned independently of the luminosity at the other interaction points [178].

The enhanced satellites were measured by the LDM in order to verify the correct tuning of the modified bunch-splitting process. The results are shown in Figure 87 and Figure 88.

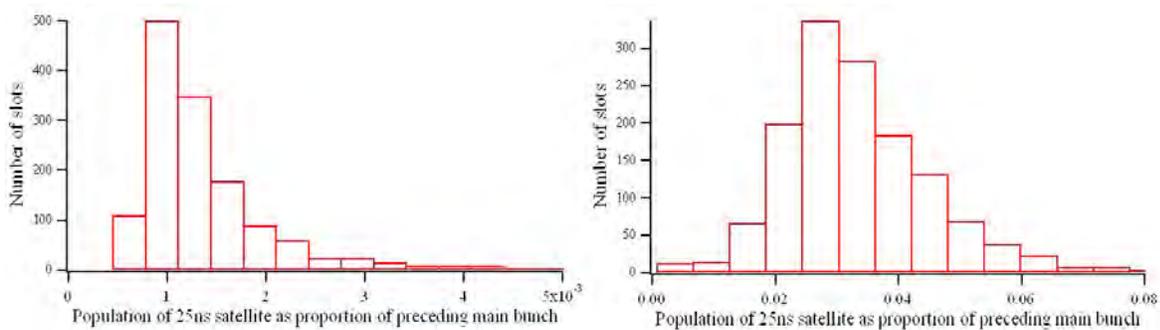


Figure 87. Distribution of relative satellite populations in fills with and without enhanced satellites. Fill 2219 (left) used the normal bunch splitting scheme and had only spontaneously occurring satellites, while fill 2261 (right) used the modified scheme with enhanced satellites. The mean spontaneous 25 ns satellite population was 1.4×10^{-3} of the preceding main bunch, while for enhanced satellites it was 3.3×10^{-2} .

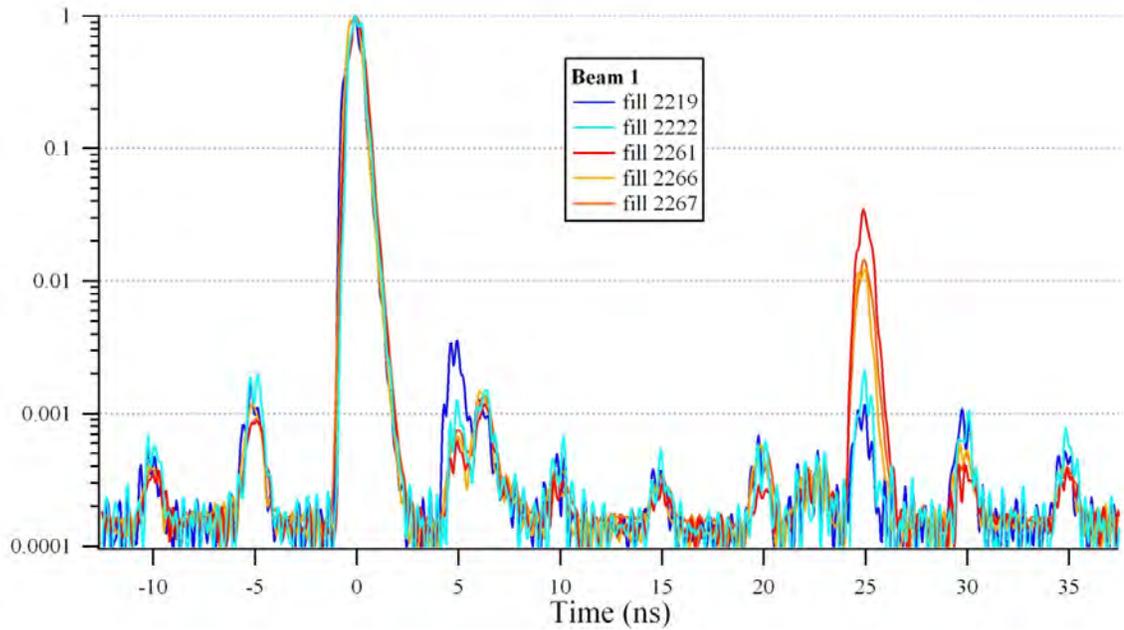


Figure 88. APD counts against time. Longitudinal profile of the filled slot and following slot, showing the effect of satellite enhancement in the LHC injector chain. Fills 2219 and 2222 used the normal bunch splitting scheme and had only spontaneously occurring satellites, while fills 2261, 2266 and 2267 used the modified scheme with enhanced satellites.

6.3 Time Resolution

An example of a single-bunch profile is shown in Figure 89. The LDM's time resolution is sufficient to determine the bunch shape and length. In this case, the bunch is Gaussian. The larger tail on the right-hand side (trailing edge) is in fact caused by the diffusion tail of the APD.

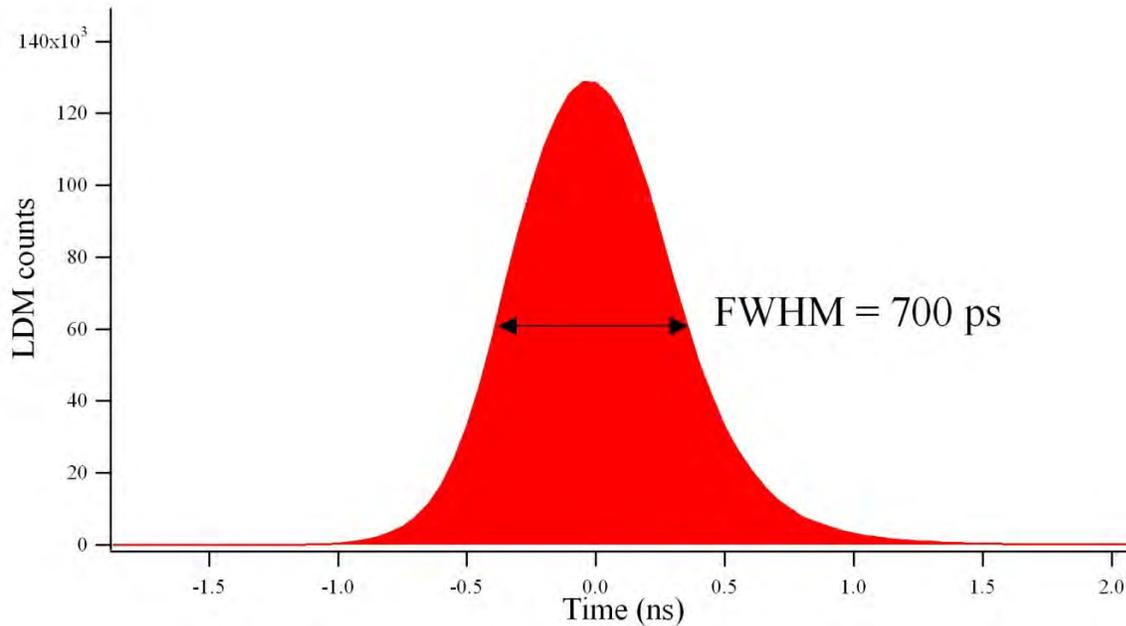


Figure 89. APD counts against time. Single bunch from a longitudinal profile obtained during an LHC lead-ion fill at 3.5 Z TeV. The spacing of main bunches was 500 ns, integration time was 250 s. Corrections for detector deadtime and afterpulsing have been applied.

Since the profile obtained by the LDM covers the entire ring, all the individual bunch lengths can be determined at the same time. In this case, a high dynamic range is not necessary and an integration time of 10 seconds is sufficient. A Gaussian fit is applied to each bunch. The bunch lengths thus obtained are compared with those measured by the LHC wall current monitor, sometimes known as the Beam Quality Monitor, in Figure 90. The LDM resolution of $\sigma=90$ ps is subtracted in quadrature from the σ given by the Gaussian fit. The wall current monitor provides a single measurement taken every 5 s and digitised at 8 GSamples / s. The FWHM bunch length is calculated by interpolation of samples lying above and below half maximum. This is converted to σ assuming that the bunch is Gaussian and a subtraction is made to account for dispersion in the pickup and cables. The agreement is very close around the whole ring.

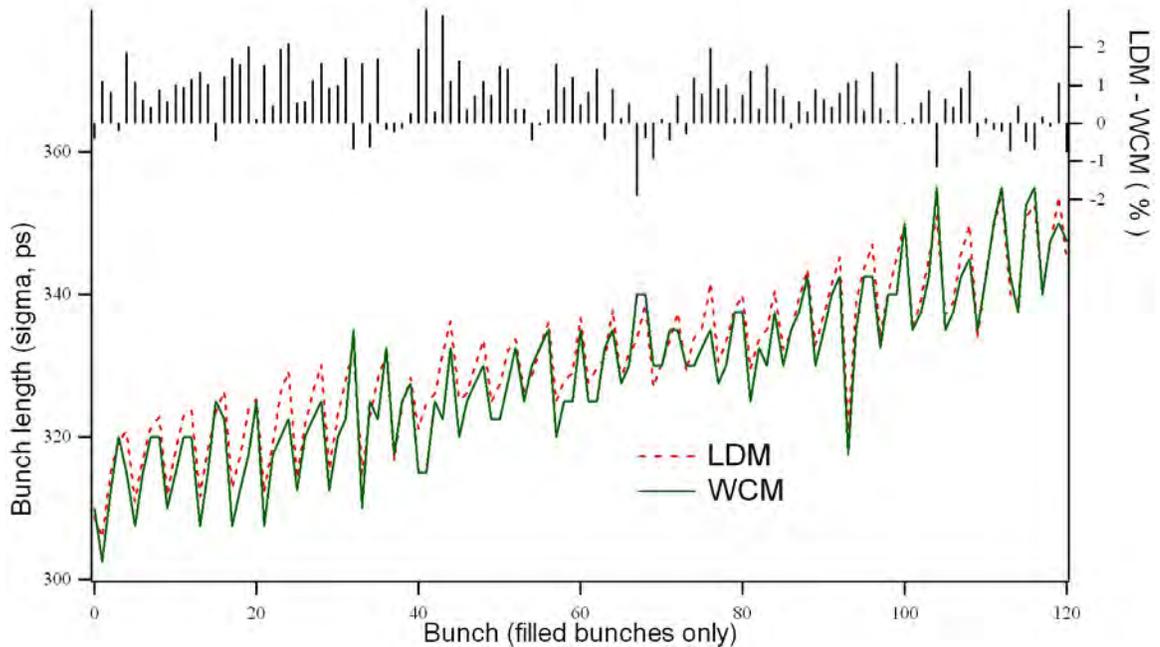


Figure 90. Bunch lengths measured by the LDM and the wall current monitor (WCM). The ring contained 121 bunches of lead ions at 3.5 Z TeV with minimum spacing of 500ns. The LDM jitter of 90ps is subtracted in quadrature from the measured bunch length. Below and left axis, measured bunch lengths. Above and right axis, difference in the two measurements.

The LDM has a time resolution of approximately $\sigma=90$ ps. The uncertainty in the photon arrival time can have three components: uncertainty in the turn clock, uncertainty in the photon detection, and uncertainty in the time-stamping.

6.3.1 Sources of Timing Uncertainty

The largest source of uncertainty in time is the jitter of the START pulse provided to the TDC. Since the TDC measures the arrival time of each photon relative to the most recent START pulse, jitter in the START pulse is equivalent to jitter in the arrival time. The START pulse is given by the LHC turn clock. The turn clock is part of the Beam Synchronous Timing (BST) system [160] and is distributed to all points of the LHC via fibre-optic links. The signal carries the turn clock, bunch clock, and a short message giving information on the status of the machine. A separate fibre carries the information for each of the two beams. A Beam OBServation receiver (BOBr) card receives the optical signal and produces an electrical pulse

which is sent to the START input of the TDC. This electrical pulse has a jitter of approximately $\sigma=75$ ps, as shown in Figure 91.

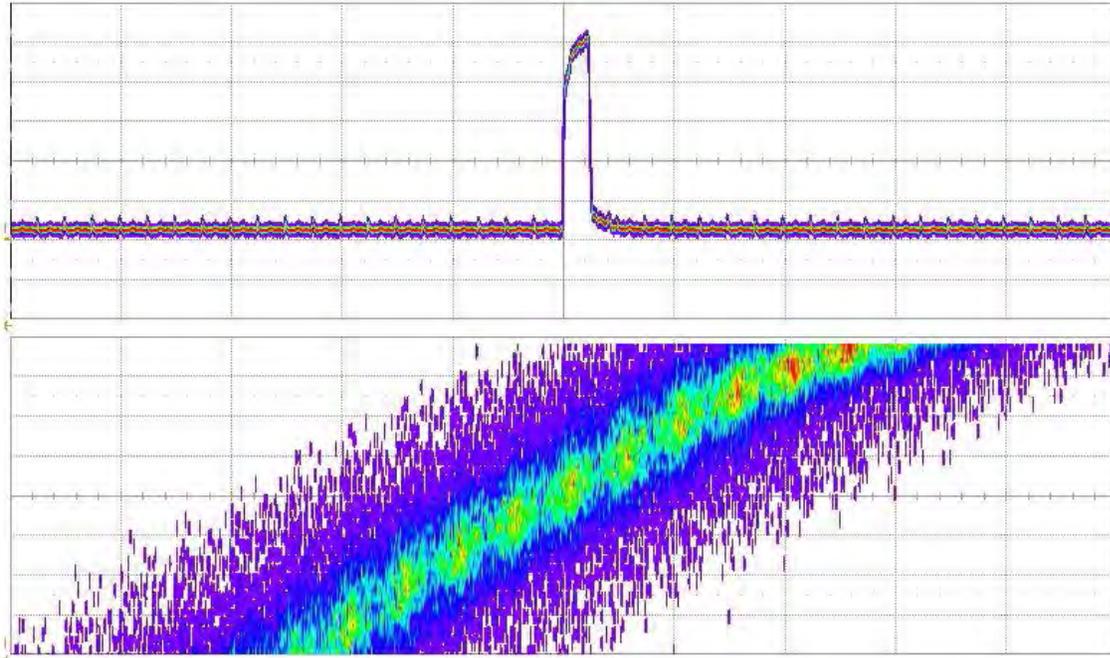


Figure 91. Oscilloscope trace of the turn clock signal produced by the BOBr card. Upper plot, 100 ns per division. Lower zoom, 200 ps per division. Making a cut at 1.5 V showed the jitter to be roughly Gaussian with $\sigma=75$ ps.

The resolution of the PDM detector is quoted as 50 ps FWHM [154]. Assuming that the jitter is Gaussian this gives $\sigma=21$ ps. The TC890 has a time resolution of 50ps [159]. This in turn includes the internal clock jitter, TDC nonlinearity and the quantisation error due to binning. All the errors are independent, thus

$$\sigma_{LDM} = \sqrt{\sigma_{BOBr}^2 + \sigma_{PDM}^2 + \sigma_{TDC}^2} = 92 \text{ ps} \quad (81)$$

Other sources of timing jitter have been considered and found to be negligible. For example, the transit time spread of photons of different wavelength through the fused silica beam-pipe window was calculated in chapter 3 as 2.4ps. The spread in time resolution of the SR relative to the emitting photons, due to the geometry of the emission, was examined in chapter 2. It is greatest for lower proton energies. At 450 GeV, following equations (43) and (49), the time spread in the undulator is

$$\Delta t = \frac{L_u}{c} \left(1 - \frac{4\gamma^2}{\beta(4\gamma^2 + K^2)} \right) = 80 \text{ fs} \quad (82)$$

and in the dipole is

$$\Delta t = \frac{2\rho}{c\gamma} \left(\frac{1}{\beta} - 1 \right) = 360 \text{ fs} \quad (83)$$

Even for lead ions at injection energy of 177 GeV per nucleon, the time spread is only 6 ps. Lead ions at injection do not emit enough SR for use of the LDM so the time spread for all cases where the LDM can be used is less than 6 ps.

6.4 Dynamic Range & Uncertainties

6.4.1 Limiting Factors

The dynamic range (DR) of an instrument is the ratio of the largest to the smallest signal that it can measure simultaneously. Because the LDM must measure very small satellite and ghost bunches at the same time as nominal bunches, it needs to have a very large DR. The DR of the LDM is principally limited by two factors: shot noise and false counts.

Since photon counting is a quantised system, it is subject to shot noise. The expectation value of the number of photons counted is directly proportional to the proton density in any given bin. However, photons arrive stochastically, and the number of photons counted is only an approximation of the expectation value. However, as more photons are added to the histogram, this approximation becomes more accurate. Thus, the shot noise limit is really a limit of the feasible acquisition time. Increasing the acquisition time increases the DR.

False counts occur either randomly spread around the ring (dark counts) or correlated to the true photon counts (afterpulses). If the detector response is well characterised then both the dark counts and the afterpulses can be statistically removed. Since the false counts are themselves subject to shot noise, the removal will be more accurate the longer the acquisition time.

6.4.2 Sources of Uncertainty

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. All the errors are taken to be independent, and are thus added in quadrature to give the total uncertainty, with the debunched beam uncertainty acting only in the positive direction.

Table 13. Sources of error in the LDM ghost / satellite measurement

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

The statistical uncertainty (or shot noise) in any bin is given by the square root of the number of counts in that bin [120]. This is true regardless of any re-binning. Similarly, the baseline uncertainty is given by the effect of the statistical error in the bins used to calculate the baseline.

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. An upper limit can also be set by comparing LDM data to results from the

Abort Gap Monitor. 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 population measured by the LDM has been shown to have a dependence on the transverse beam size. 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, 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 that of the main bunches. Indeed, since space-charge effects are negligible at the high energies of the LHC, most satellite and ghost bunches are not expected to have systematically different emittance to the main bunches. However, some satellites can be formed early in the LHC injection chain, where space charge effects are significant. These satellites might be expected to have a smaller emittance than the main bunches, since their smaller population produces smaller space-charge forces.

6.4.3 Dependence on Emittance

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 centred 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 centre, 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 measurements, since the alignment of the LDM independently of the BSRT was only possible after the installation of additional translation stages in August 2011.

This effect was investigated during an MD carried out in May 2011 with groups of large- and small-emittance bunches [179]. Two trains of 12 bunches were injected into the LHC. During the injection of the second train, a screen was inserted into the transfer line in order to blow up the emittance. The result was that the second train had an emittance about 80% larger than

normal in both planes, giving a transverse beam size about 3 times larger. The measured population was indeed considerably affected (Figure 92). On this occasion, the LDM detector was positioned off-centre, so that the bunches with larger emittance gave a larger signal. The correlation can clearly be seen in Figure 93.

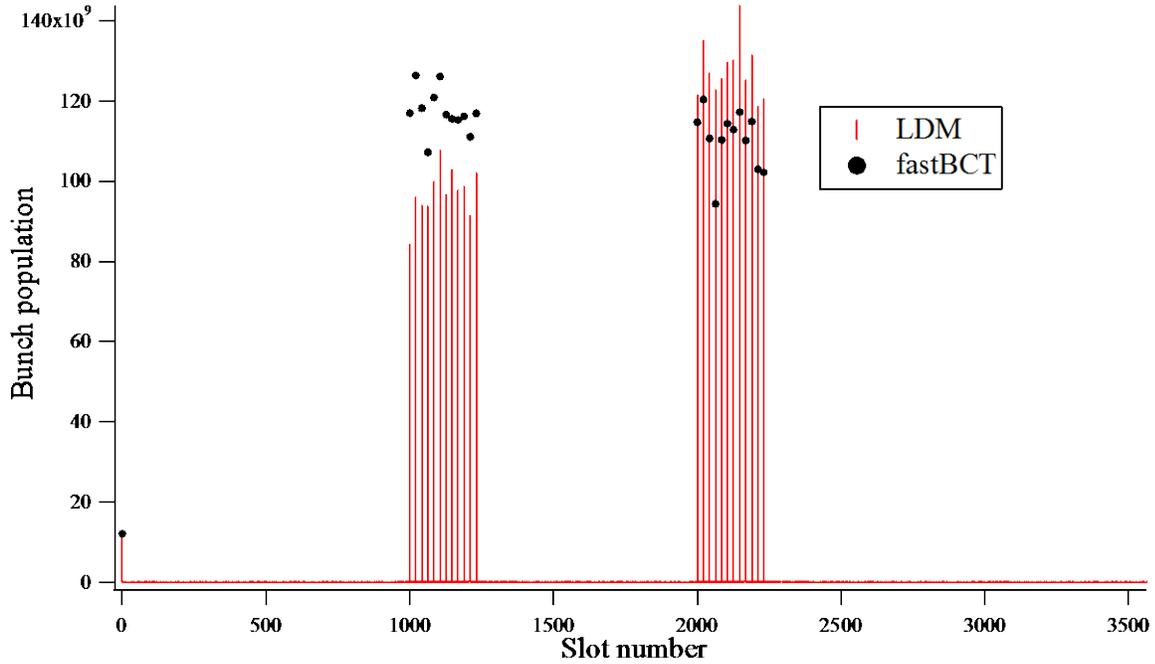


Figure 92. Normalised bunch population against bunch number. Comparison between the relative bunch populations measured by the fast Beam Current Transformer and the LDM, during an MD with bunches of different transverse size. The bunch population measured by the LDM is sensitive to the transverse emittance of the bunches. The second 12-bunch train had significantly larger emittance than the first.

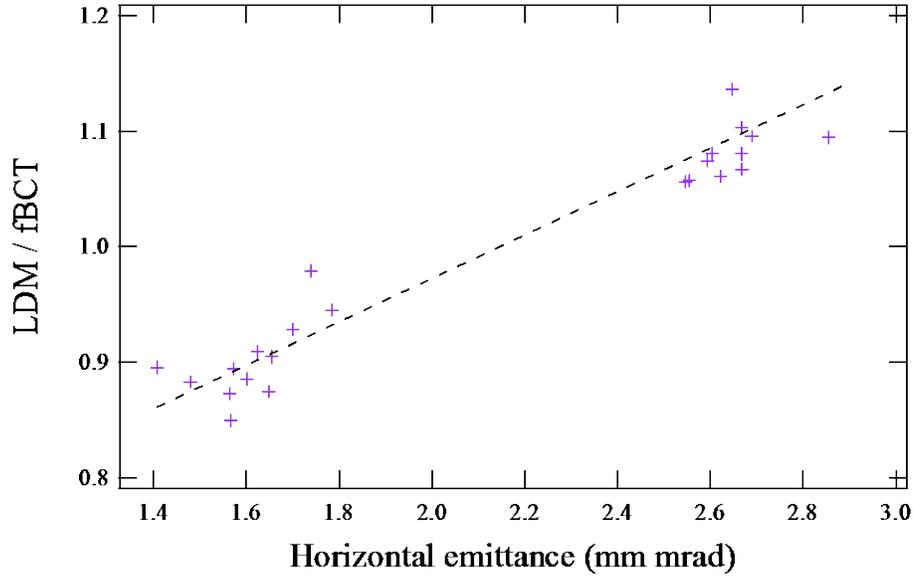


Figure 93. Ratio of the relative bunch populations measured by the LDM and the fast Beam Current Transformer against horizontal beam emittance. Measured during an MD with bunches of different transverse size. The bunch population measured by the LDM is sensitive to the transverse emittance of the bunches. A linear fit is shown.

6.4.4 Modification of the Beamline

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 94. A diffuser with a Gaussian point spread function (PSF) scatters the incident light. This produces a beam spot with size

$$\sigma_{After\ Diffuser}^2 = \sigma_{Original}^2 + \sigma_{PSF}^2 \quad (84)$$

The diffuser is chosen with $\sigma_{PSF} \gg \sigma_{Original}$, so that the dependence on the original beam size of the spot size after the diffuser is negligible. Thus, the measured population no longer depends on the bunch emittance. The large size of the beam spot after the diffuser causes an unacceptable loss of coupling efficiency, so a lens is used to increase the amount of light captured.

A simulation of the coupling efficiency with and without the diffuser was carried out and the results are shown in Figure 95. In the original configuration, the relative bunch population would be measured inaccurately if the bunches had different transverse size (middle row) or

position (bottom row). With the new layout, however, the relative bunch population is measured correctly regardless of transverse size or position.

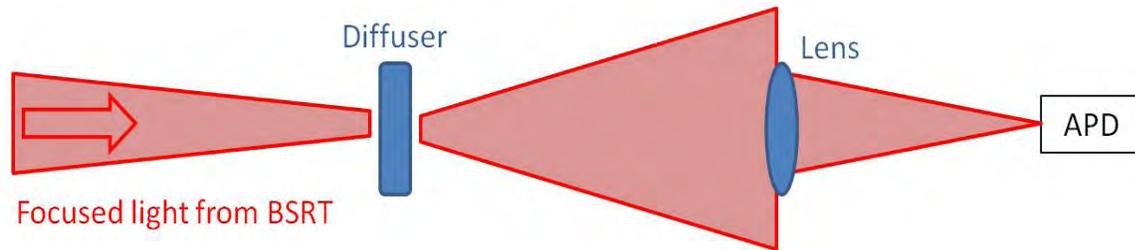


Figure 94. Modified layout of the final part of the LDM beamline. Instead of being focused directly onto the APD, the SR is first scattered through a diffuser and then re-focused by an achromat lens.

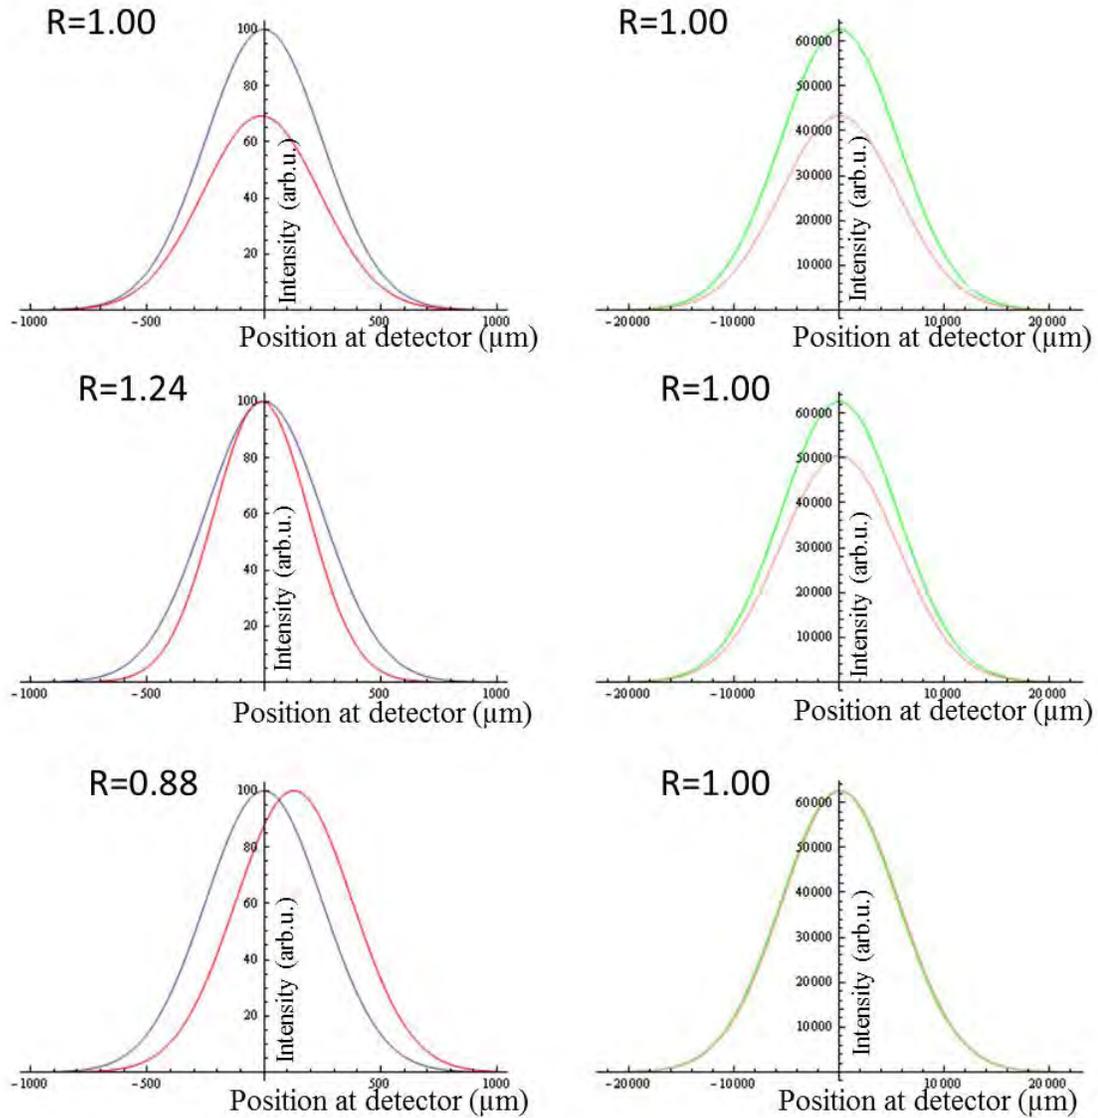


Figure 95. Light intensity against transverse position at the detector plane. Simulation of the effect of transverse beam size on the measured population. Left, original setup. Right, new setup with diffuser and lens. R is the ratio of the measured to the true population of the bunch in red, an accurate measurement thus gives $R=1$. Top, red bunch has smaller population. Centre, red bunch has smaller transverse emittance. Bottom, red bunch has different transverse position.

A second MD session was carried out in order to validate the effect of the diffuser. In this case, a large emittance spread was achieved by using the transverse dampers to selectively blow up the emittance of chosen bunches in one or both planes. Figure 96 shows that as expected, the dependence on emittance is completely eliminated.

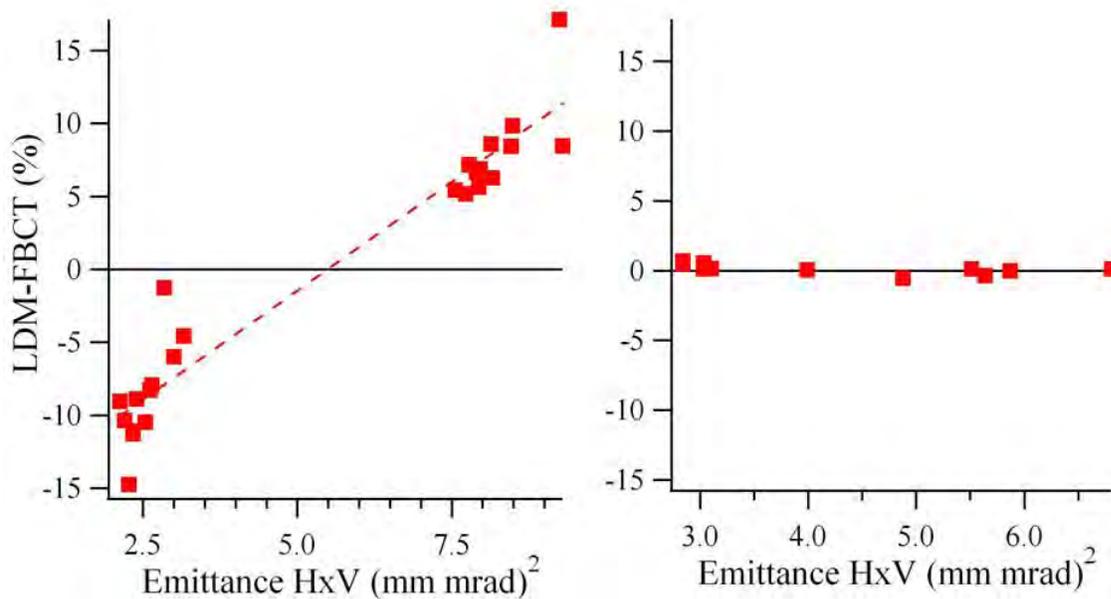


Figure 96. Difference in bunch current measured by the LDM and the Fast BCT against beam emittance, before (left) and after (right) installation of the diffuser and lens.

6.5 Beam Debunching

Two tests of beam debunching have been carried out at the LHC. Due to the nature of single photon counting, the LDM is not ideally suited to observation of rapidly changing events. Nonetheless, some useful measurements can be made. The most important is to cross-check that the LDM total count rate remains the same during the debunching, which was found to be the case within the statistical uncertainties which follow from the short acquisition times. This is a necessary condition to show that the deadtime correction algorithm is working properly. The count rate starts to drop only about 5 minutes after the debunching, when some particles have lost sufficient energy to be lost on the momentum cleaning collimators.

The following results were recorded during a machine development (MD) session on the 14th of March 2011. There were three nominal bunches and one pilot present, and the beam was at 3.5 TeV.

The RF system was turned off ($T=0$) and the loss of longitudinal focusing immediately caused the bunches to spread out (Figure 97). In addition, the energy radiated as synchrotron radiation was no longer replaced by the RF so the mean energy of the particles began to fall.

Since the LHC operates above transition, a loss of energy causes the bunch centres to move earlier in the ring. The rate of bunch centre movement allows the rate of energy loss due to synchrotron radiation to be estimated. The growth in bunch length allows the initial energy spread to be estimated.

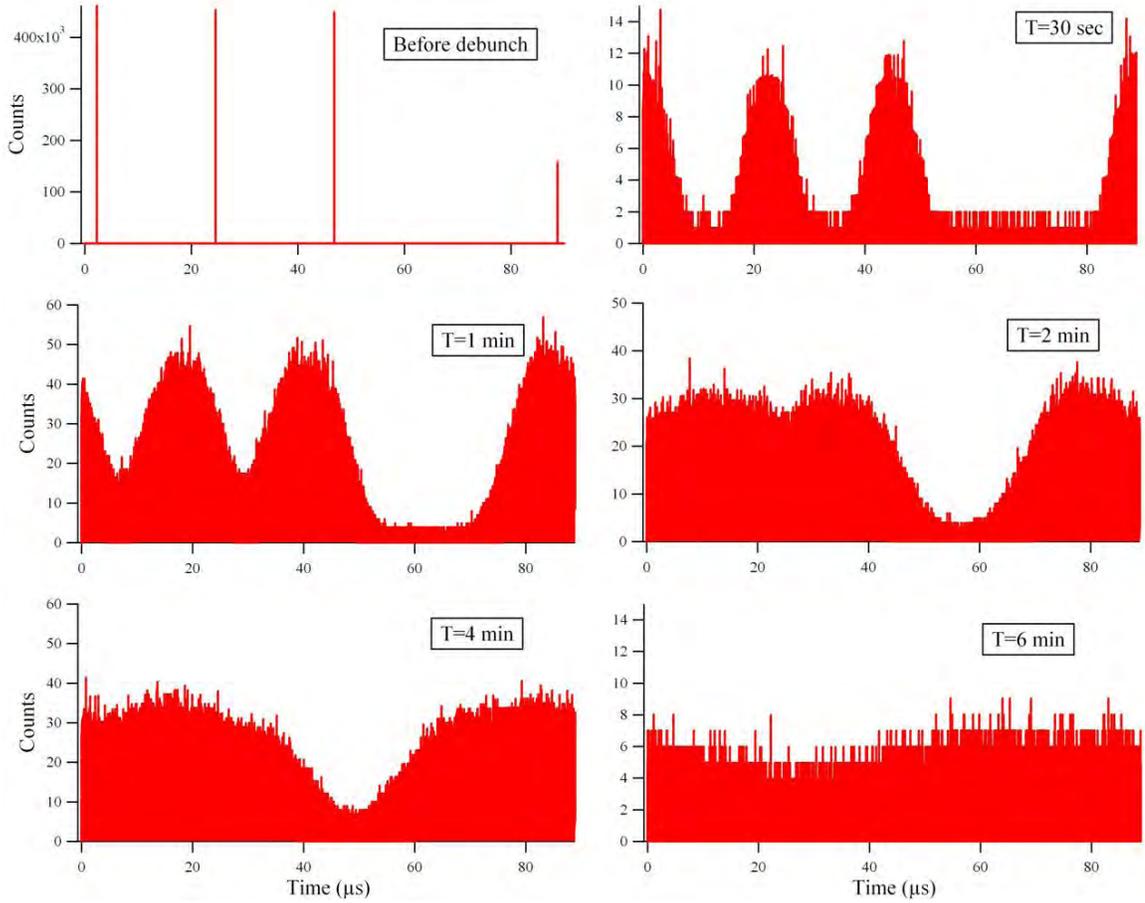


Figure 97. APD counts against time. LDM measurements of beam debunching after the RF was turned off at 3.5 TeV.

The slip factor η is defined by the LHC lattice and the beam energy. At 3.5 TeV it is equal to -0.00032. The slip factor relates the change in particle energy to the change in revolution frequency,

$$d\omega/\omega = \eta dE/E \quad (85)$$

The slip factor is negative because the LHC operates above transition energy. Thus, an increase in particle energy leads to a decrease in the revolution frequency, since the higher-energy particle follows a longer path. The rate of change is then

$$\frac{1}{\omega} \frac{d\omega}{dT} = \frac{\eta}{E} \frac{dE}{dT} \quad (86)$$

Re-arranging,

$$\alpha = \frac{\eta\omega}{E} \frac{dE}{dT} \quad (87)$$

where $\alpha = \frac{d\omega}{dT}$ is the rate of circular acceleration.

The phase shift after some time T is given by

$$\Delta\phi = \Delta\omega T + \frac{1}{2}\alpha T^2 \quad (88)$$

where $\Delta\omega$ is the initial difference in revolution frequency between a given particle and that of an on-momentum particle. It can be seen that the first term relates to the spreading out of the bunch due to the initial energy spread while the second term relates to the precession of the bunch centroid due to radiation of energy. Substituting eqs. (85) and (87) into eq. (88),

$$\Delta\phi = T \omega_{rev} \eta \frac{dE}{E} + \frac{1}{2} T^2 \omega_{rev} \frac{\eta}{E} \frac{dE}{dt} \quad (89)$$

The phase shift is given by

$$\Delta\phi = 2\pi \frac{\Delta t}{t_{rev}} \quad (90)$$

where Δt is the shift in the arrival time of the particle. Eq. (89) can thus be rewritten as

$$\Delta t = T \eta \frac{dE}{E} + \frac{1}{2} T^2 \frac{\eta}{E} \frac{dE}{dt} \quad (91)$$

By fitting Gaussians to each bunch in the profiles shown in Figure 97, the energy spread of the beam before debunching and the rate of energy loss in the coasting beam can thus be estimated. In practice, the pilot and the main bunch are close together and join in the first 30

seconds, so that the fits do not perform well. The results for the second and third bunches are shown in Table 14.

Table 14. Results of a Gaussian fit on two of the bunches in the profiles shown in Figure 97. The RF was turned off at $T=0$.

T (sec)	Bunch centre (μs)		Bunch σ (ns)	
	2 nd bunch	3 rd bunch	2 nd bunch	3 rd bunch
0	24.5	46.8	0.13	0.13
30	22.6	44.9	3876	3201
60	18.1	40.4	6681	4538
120	14.4	33.3	no fit	12480

Ignoring the movement of the bunch centre and setting the second term of eq. (91) to zero, the energy spread before debunching is calculated as

$$\frac{dE}{E} = -6.5 \times 10^{-6} / 60 \times \eta = 3.3 \times 10^{-4} \quad (92)$$

which is close to the expected value.

Now looking only at the movement of the bunch centre, the first term of eq. (91) becomes zero and thus

$$\frac{dE}{dt} = \frac{2\Delta t E}{T^2 \eta} = 8.3 \text{ MeV/s} = 740 \text{ eV/turn} \quad (93)$$

This is larger than the expected rate of synchrotron radiation, which was calculated at 440 eV per turn. However, the beam can lose energy in other ways, notably due to the impedance of the beam pipe and other accelerator components, so that it is not surprising to find a higher rate of energy loss than that expected from SR alone.

6.6 3D Bunch Shape Measurement

6.6.1 Method

The LDM detector is located at the image plane of the BSRT. Since the active area of the APD is only 50 μm in diameter, only a part of the image is sampled. Since the image size produced by the BSRT is between 100 and 500 μm , only part of the image is sampled. While this creates an undesirable dependence of coupling efficiency on the transverse size of the beam (discussed above), it also opens the possibility of obtaining a 3-dimensional beam profile.

The APD is mounted on remote translation stages. It can thus be scanned horizontally and vertically over the image plane. Combining the longitudinal profiles obtained at each position allows a 3-dimensional beam profile to be constructed.

6.6.2 Limitations

Measurement of the 3D profile requires considerable time. For a longitudinal profile of the nominal bunches, an acquisition of around 10 seconds is sufficient. Since the filters can be adjusted between measurements, acquisitions in the beam tails, where the light level is much lower, need not take any longer than acquisition of the bunch centre. Allowing a few seconds for the movement of the stepper motor, a 10x10 scan requires 25 minutes. The profile is only valid if the beam distribution is constant over this period, although gradual beam losses which do not change the beam shape can be compensated by cross-referencing with the beam current transformers.

6.7 Relevant Results from Other Instruments

6.7.1 Synchrotron Light Intensity

The LHC Abort Gap Monitor (AGM) [37] is located on the BSRT optical table and receives about 10% of the available light. The AGM consists of a gated photomultiplier tube (PMT) and the associated acquisition electronics. The PMT is gated on for 3 μs and the signal is acquired in bins of 100 ns. In normal operation, the PMT gain is high and the 3 μs gate coincides with the LHC abort gap, which must remain empty.

However, the delay of the AGM gate with respect to the turn clock can be changed. By centering the gate on the first bunch, the relative intensity of SR during the ramp can be measured (Figure 98). The gain of the PMT is set lower than in normal operation, as there are many more particles in the gate period. The gain is automatically adjusted in steps to allow for the changing intensity during the ramp. The agreement with simulations carried out using SRW (presented in section 3.4) is quite good. The discrepancy, especially for protons between 450 and 1200 GeV, may be ascribed to the different spectral response of the AGM PMT compared to the LDM detector.

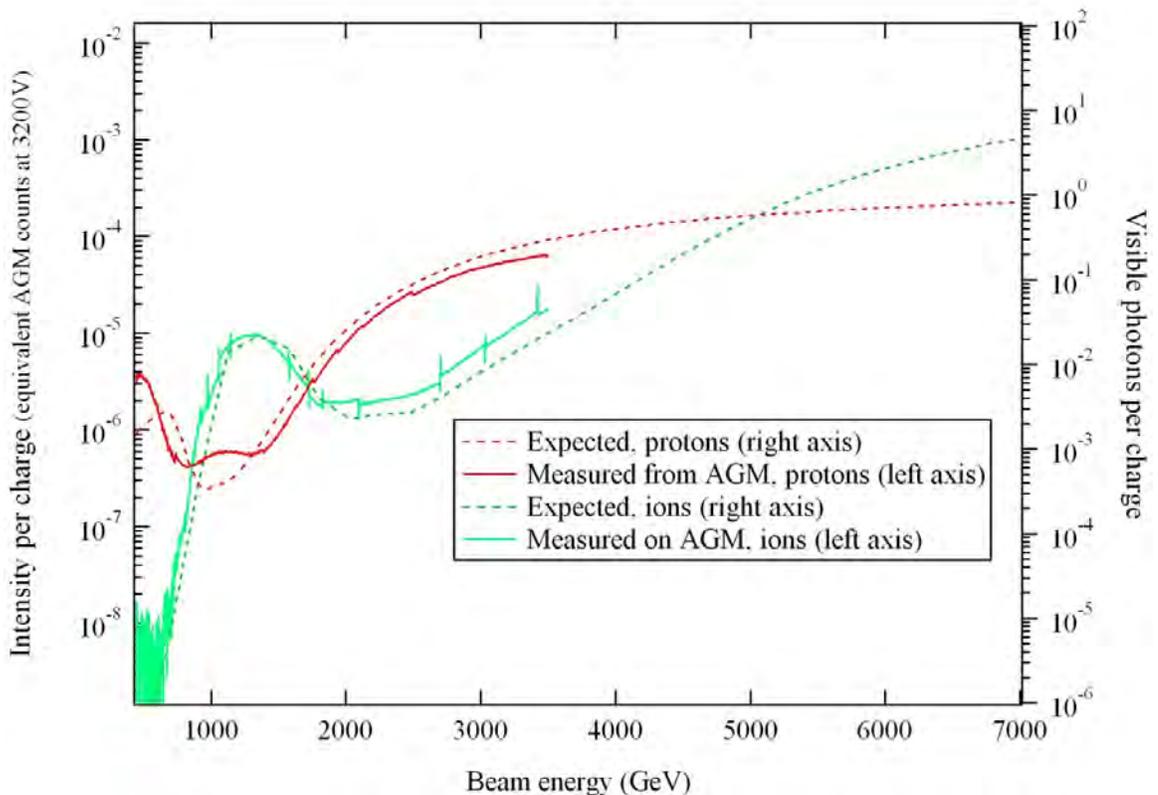


Figure 98. Intensity of visible synchrotron light collected by the BSRT against beam energy. Comparison of intensity measured by the Abort Gap Monitor (AGM) with expected values from simulation.

6.7.2 Distribution of Light on the Extraction Mirror

During the January 2012 shutdown an additional camera was placed on the BSRT optical table with the intention to image the distribution of light on the extraction mirror. The camera is located close to the entrance mirror of the BSRT optical system and the focal plane of the camera is set at the extraction mirror. The camera is placed off-axis in order to avoid the SR beam hitting the CCD directly. Instead light which is scattered from the extraction mirror with a

small angle is used. This approach is imperfect since it adds an angular dependence to the measured distribution. However, it is sufficient for a good approximation of the SR distribution on the extraction mirror, which is shown in Figure 99 for different beam energies. The agreement with the simulations shown in 3.4.2 is good, although the on-axis beam spot appears closer to the left hand edge of the mirror, indicating that the alignment of the extraction mirror is not quite as used in the simulations.

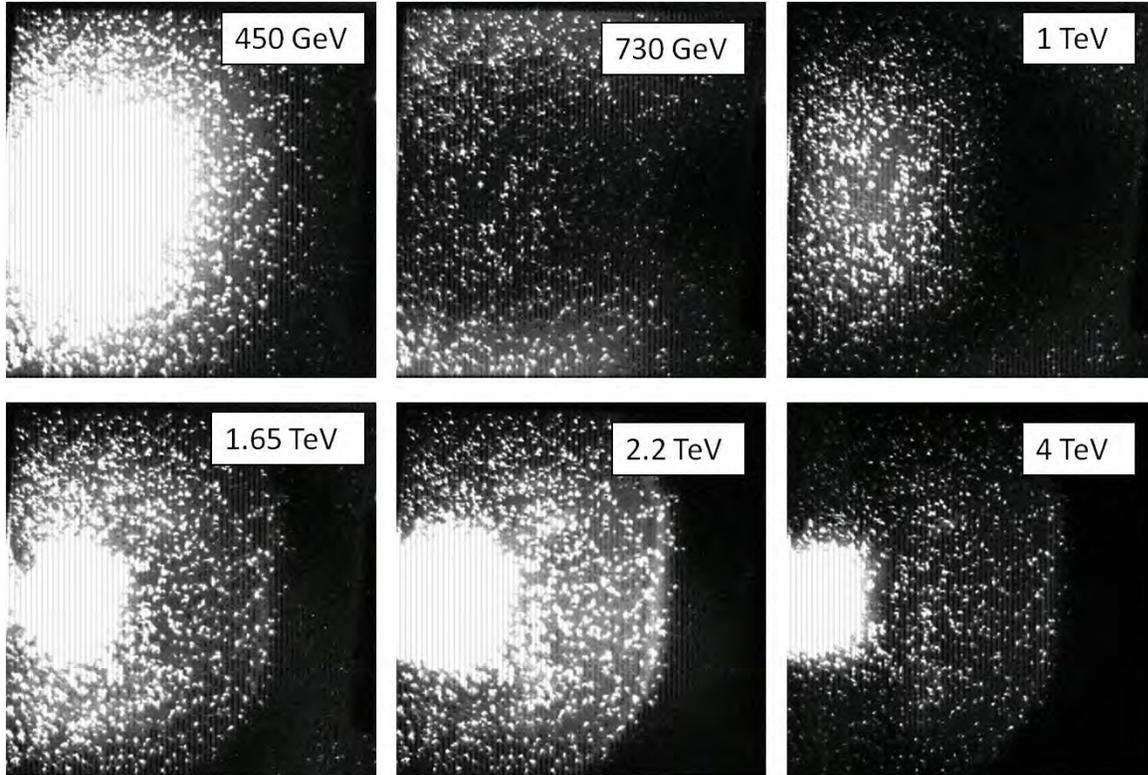


Figure 99. Images of the extraction mirror at different beam energies.

Chapter Summary

Examples of longitudinal profiles measured with the LDM have been shown to illustrate the capabilities of the instrument. In particular, the LDM's dynamic range has been demonstrated to exceed 10^5 within an integration time of approximately 30 minutes. The importance of effective correction for the effects of the APD's deadtime and afterpulsing in achieving this extremely high dynamic range has been shown.

The LDM's time resolution has been estimated at 90 ps, and the sources of timing jitter have been identified. At the same time, the LDM has a record length of 89 μ s, equal to a full LHC revolution, giving rise to a longitudinal profile histogram of 1.8 million bins. Statistical methods for dealing with this data have been explained, allowing the LDM results to be compared with data from other instruments and from the LHC experiments. In all cases, there has been excellent agreement with other measurements.

The original arrangement of the LDM optics was found to introduce an undesirable dependence of the longitudinal measurement on the transverse bunch size. The introduction of an optical diffuser and an additional lens into the optical system has been proven to eliminate this dependence.

7. Future Improvements

7.1 Improved Characterisation of the Detector Response

Two aspects of the detector's response require more detailed study. Firstly the dependence of the afterpulsing on the internal temperature of the detector, and thus on the counting rate. At present, different fit parameters for the afterpulsing correction are applied by hand for Van der Meer scan fills, which have a much smaller number of bunches than normal physics fills and thus a lower average count rate. A better solution would be to have a table of correction parameters against different count rates, and for the relevant parameters to be applied automatically by the front-end software.

Secondly, the diffusion tail of the APD should be modelled and a correction algorithm established. This would allow the measurement of satellite bunches occurring in the RF bucket immediately following the main bunch, which is presently not possible.

7.2 Software Improvements

Establishing accurate numbers for the satellite and ghost charge fractions currently requires an offline analysis. However, the procedure is now well understood and could be carried out automatically. Efforts are under way to move the analysis to the front-end server. This would allow the satellite and ghost charge fractions and the average slot arrays to be logged automatically at the end of each LDM acquisition. As well as being more time-efficient, having these numbers in the standard CERN measurement database would make it much easier for users to access the LDM data.

7.3 Ultra-high Dynamic Range Method

The dynamic range which can be achieved by a photon counting system is in principle only limited by the statistical shot noise, which is equal to the square root of the number of counts in the bins of interest. In order to increase the dynamic range, either the integration time or the count rate should be increased. However, the integration time is limited to the period over which beam conditions can be considered constant, while the deadtime of the detector limits the count rate at which the detector can be run without losing counts from the end of the

bunch and the trailing satellites. In practice, furthermore, other factors such as the imperfect subtraction of afterpulses reduce the dynamic range.

The dynamic range could be increased considerably by the use of two detectors, one operating in gated mode. This detector would be blind to the main bunches, and would thus be able to accept many more photons from the ghost and satellite bunches. The Monte Carlo simulation of such a system was shown in 4.4.5. However, as has been shown in chapter 4.5.4, the APD cannot be gated at sufficient speed.

An alternative method is to gate the light before it is incident on the detector. A number of methods exist by which light can be switched at very high speeds [180]. For example, optical switches based on a miniaturised Mach-Zehnder interferometer [181] are routinely used in the telecommunications industry and achieve switching times as low as 8 ps. However, the switching is only effective for a single wavelength, making this solution unsuitable for broadband light sources such as those available in the BSRT. Ultrafast non-linear interferometers such as that illustrated in [182] suffer the same limitation. Movable mirrors based on microelectromechanical systems (MEMS) clearly do not have any wavelength dependence, but cannot achieve the switching speed required, and their reliance on moving parts makes them insufficiently robust.

A Pockels cell [183] can be used as a light shutter if it is placed between crossed polarisers. When a voltage is applied to the cell, the polarisation of light passing through it is rotated, allowing some light to pass the second polariser. This method has been demonstrated for synchrotron light diagnostics at the Spring-8 synchrotron [184]. However, the shutter has a limited repetition frequency due to heating of the Pockels cell. Opening the shutter with a low duty cycle would increase the required integration time and thus negate the benefits of the two-detector system. In principle, a series of shutters could be used to overcome the repetition frequency limit, but such a system would be expensive and would require very careful optical alignment.

The solution proposed is to use an electro-optic (EO) deflector [185] as an optical switch. The deflector would be located several metres upstream of the APD, such that the deflection is translated into a horizontal displacement. EO deflectors can give an angular deflection of several mrad with a relatively low voltage and high frequency [186]. A displacement of up to 1

cm might then be expected close to the APD. Since this is much larger than the horizontal spot size, a large extinction ratio should be possible.

A lab test of the deflector-based optical gate has been carried out. A custom-made EO deflector was manufactured by Leysop [187]. Two deflecting crystals are combined in series to obtain a large angular deflection. A resonant circuit tuned to 20 MHz is used to modulate the deflector, allowing a relatively small input voltage to be used. A modulation of 20 V should produce a deflection of 6 mrad [188]. However, for this lab test a suitable driver was not available, and only 5 V could be produced at the required frequency, leading to a maximum deflection of 1.5 mrad. Nonetheless, this allowed the principle to be shown.

A low-power laser beam was shone through the deflector and onto the APD, which was located at a distance of 2 m from the deflector. The signal generator produced a 5 V, 20 MHz sine wave which was used to drive the deflector. It also produced a synchronisation signal which was used to trigger the TDC. The output of the APD was connected to the TDC STOP channel and the photon-counting histogram was built up. An example is given in Figure 100. An extinction ratio of 20:1 was achieved. This was mostly limited by diffraction. Since the aperture of the deflector is small, it causes significant diffraction of the laser beam as it passes through. Diffraction rings were visible at the APD. When the main spot is deflected off the APD, the diffraction ring might still hit it, meaning that the transmission does not drop to zero. This problem could be overcome by using a deflector with a larger clear aperture. In addition, using a specially-built generator to produce a 20 V, 20 MHz wave would allow a greater deflection to be achieved and thus a better extinction ratio.

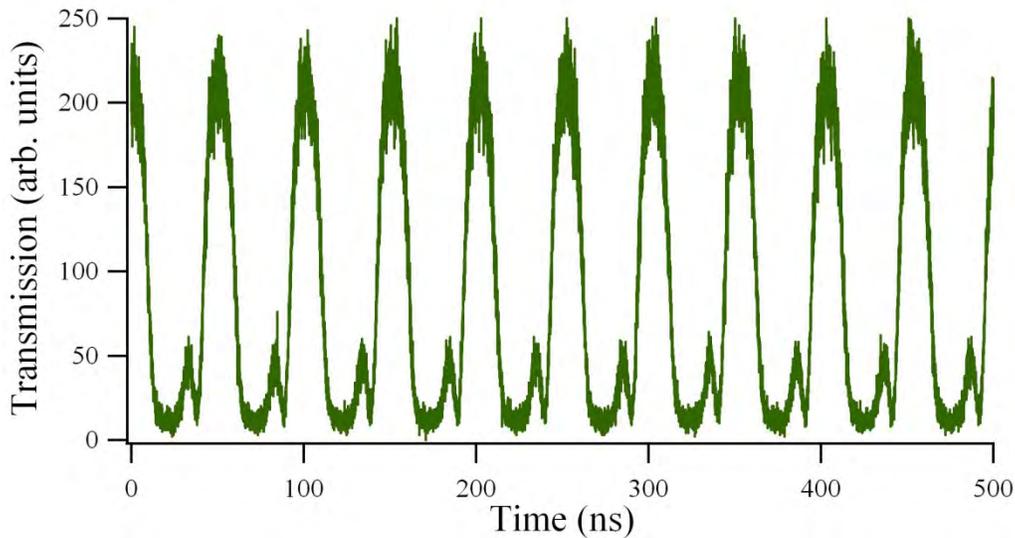


Figure 100. APD counts against time. Optical gating using the electro-optic deflector. An extinction ratio of 20:1 is achieved.

Installing a mask across which the light beam is scanned would allow the transmission to be varied with any waveform required [189]. In the LDM case, it would be desirable to have high transmission for most of the time, with zero transmission for a period of 2.5 ns coinciding with each nominal LHC bunch. In this case, a lens with a diameter greater than the maximum deflection would be placed in front of the APD, such that all of the light is gathered onto the active area of the APD regardless of the deflection given. A mask consisting of a single opaque line would then be placed in front of the lens (Figure 101). As the beam is swept across the mask, the light is blocked for a short period. If the deflector is modulated with a sinusoidal voltage, then the sweep speed will be greatest at the centre of the pattern, i.e. the position with zero deflection. A faster switching speed can be achieved by placing the blocking mask at this position, since the switching time is the size of the beam spot divided by the speed at which the spot is swept across the mask plane. Since the spot is swept across the centre twice in each cycle, it is sufficient to operate the deflector at half the LHC bunch frequency.

An example of the high duty cycle that can be achieved using this method is shown in Figure 102. Once again, diffraction rings around the beam spot limit the extinction ratio and cause additional structure in the waveform.

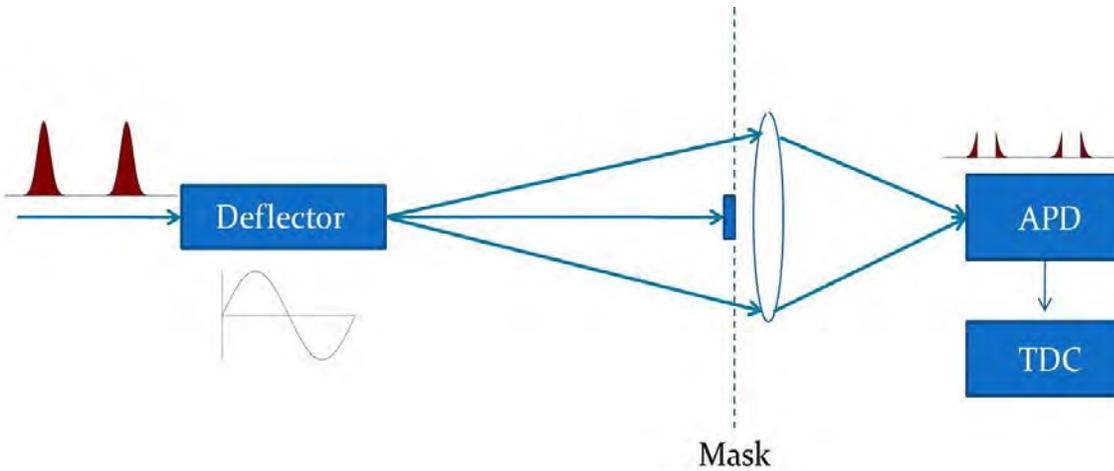


Figure 101. Layout of the proposed deflector-based optical gate. The beam spot is swept across a mask which blocks the light coming from the nominal LHC bunches. A lens then gathers all the remaining light onto the active area of the APD.

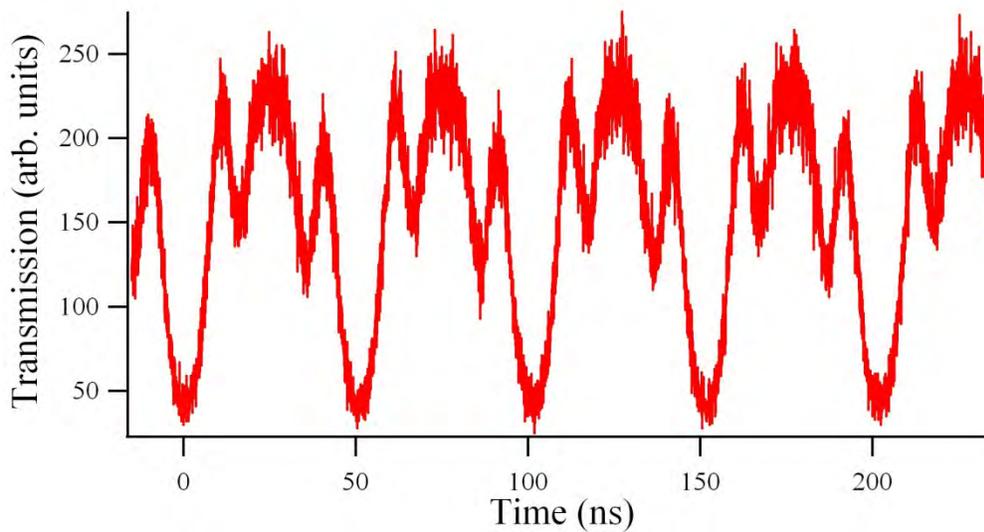


Figure 102. APD counts against time. Adjusting the position of the mask allows different transmission waveforms to be created. In this example, the transmission peak is broadened and the system is strongly attenuating for only a short period, as would be required for the LHC 50 ns filling scheme.

Such a system could dramatically increase the dynamic range. Equivalently, the desired dynamic range could be achieved in a much shorter time, allowing for example to check the quality of a fill before proceeding to the energy ramp. In principle, if the gate operated with a 100% extinction ratio, the dynamic range could be doubled. In practice it would be preferable for the ranges of the gated and free-running detectors to overlap, in order to allow cross-

calibration, and the improvement in dynamic range would then be slightly less than a factor of two.

Provided that sufficient light is available, any number of detectors could be employed, each with different gating schemes and attenuation rates, to achieve an arbitrarily large dynamic range. For example, one detector would be free-running, the second sensitive to everything but the main bunches, the third could be gated off during all filled slots, and a fourth gated on only during the Abort Gap. The splitting ratio and/or attenuations would be adjusted so that each detector operates at the optimum average count rate of 1 photon per deadtime period. Only when there is insufficient light available to achieve this optimum count rate is it futile to add further detectors.

Chapter Summary

Further potential improvements of the LDM system, which could not be completed within the time-frame of the project, are outlined. In particular, a novel scheme for increasing the dynamic range with the use of an optically gated detector is presented, and a suitable ultra-fast optical gating scheme is suggested. A prototype of this scheme has been tested in the laboratory and showed promising results.

8. Conclusion

8.1 Importance of the LDM in LHC luminosity calibration

The LDM plays an important part in the van der Meer scan procedure which is used for absolute luminosity calibration by the LHC experiments. The uncertainty in the ghost charge is an important source of uncertainty in the LHC luminosity, due to the need to normalise the relative bunch-by-bunch currents to the measured total circulating current. When the LHC is operated with lead ions the ghost charge fraction is larger, and the uncertainty in the measurement of the ghost charge is then the dominant uncertainty in the determination of the bunch currents.

While the ghost and satellite charge can be estimated from an analysis of collision data from the experiments, the use of the LDM to complement and in some cases replace these analyses has led to an improvement in the luminosity calibration of the LHC.

8.2 Applicability to Other Machines

The photon counting method can be applied in many cases where a high-dynamic range longitudinal profile is required. The source of radiation may vary, as synchrotron light diagnostics are limited to electron machines and hadron machines at the highest energies. Similarly, other detection techniques could be used, such as X-ray or charged-particle detection. Nonetheless, the principles would be the same and the techniques of signal correction developed for the LDM can be directly applied to any counting system.

The main limitation is the time resolution of such a system. In machines with much shorter bunch lengths, i.e. most electron machines, a true longitudinal profile may not be possible. However, bunch purity measurements could still be made provided that the bunch separation is considerably larger than the time resolution.

8.3 Conclusions

It has been shown that a photon-counting method can produce high-resolution longitudinal beam profiles of the LHC beams. The method is suitable for the low intensity synchrotron light

found in high energy hadron accelerators and has been demonstrated with both protons and heavy ions. Alongside the other instruments of the BSRT, this is the first time that synchrotron light from heavy ions has been used in beam diagnostics.

A high count rate, and consequently a shorter integration time, can be used provided that suitable correction is applied. The dynamic range of the system is largely limited by the integration time and by afterpulsing in the APDs. Correction of the signal for this effect substantially increases the dynamic range. A further improvement in dynamic range could be achieved if an optical gating could be applied to the signal, and a scheme is proposed to implement this using electro-optic deflection.

The small active area of the APDs caused difficulties with coupling stability and emittance dependence, but this has been eliminated by the addition of an optical diffuser.

An acquisition time of 5 minutes is sufficient to characterise the structure of ghost and satellite bunches and to quantify the fraction of the beam in ghosts and satellites. Bunch-by-bunch measurement of all bunches in the machine is made simultaneously. A dynamic range greater than 10^5 can be achieved by extending the integration time to 30 minutes.

The longitudinal density monitor has been proven against established instruments such as the beam current transformers and wall current monitors. However, it is more sensitive than these instruments and has a higher dynamic range. The population of satellites close to the colliding bunches has also been confirmed by data from the experiments.

The LDM is now an established part of the LHC toolbox and is widely used in LHC operation. The LDM is the primary tool for the quantification of ghost and satellite bunches in the LHC, which is essential for measurement of the colliding charge. In addition, it has been used by the operators to diagnose injection problems and to tune the bunch splitting procedures during 'enhanced satellite' operation.

Glossary

- AGM:** Abort Gap Monitor, an instrument which measures the particle population in the LHC abort gap by detecting synchrotron radiation with a gated photo-multiplier tube.
- APD:** Avalanche Photo-Diode, a solid-state detector with large internal amplification for the measurement of low-intensity light. When operated in the Geiger mode, APDs can be sensitive to single photons.
- BSRT:** Beam Synchrotron Radiation Telescope, the optical arrangement which collects and focuses synchrotron light from the LHC beams for transverse and longitudinal profile measurement.
- DCCT:** DC Current Transformer, the primary instrument used for measuring the absolute total beam current in the LHC.
- DR:** Dynamic Range, the ratio between the largest and smallest signal that can be measured by a given instrument.
- FBCT:** Fast Beam Current Transformer, the primary instrument used to measure the population of bunches in the LHC.
- FWHM:** Full Width at Half Maximum, a measure of the width of a signal.
- LDM:** Longitudinal Density Monitor, a new LHC instrument described in this work.
- LHC:** Large Hadron Collider, the world's highest-energy particle accelerator, located at the CERN particle physics laboratory.
- PDM:** Photon Detection Module, a commercial Geiger-mode APD module produced by Micro Photon Devices.
- PMT:** Photo-Multiplier Tube, a device for detecting light with large internal amplification, which can be sensitive down to the single-photon level.
- PS:** Proton Synchrotron, an accelerator at CERN which forms part of the LHC injector chain.

- PSB:** Proton Synchrotron Booster, an accelerator at CERN which forms part of the LHC injector chain for protons.
- RF:** Radio-Frequency
- SPC:** Single Photon Counting, a digital light measurement technique in which individual photons are detected.
- SPS:** Super Proton Synchrotron, an accelerator at CERN which forms part of the LHC injector chain.
- SR:** Synchrotron Radiation, electromagnetic radiation which is emitted when ultra-relativistic charged particles are transversally deflected.
- SRW:** Synchrotron Radiation Workshop, a software tool used to simulate the emission of SR.
- TCSPC:** Time-Correlated Single Photon Counting, a technique for measuring a repetitive optical signal by building up a histogram of the arrival times of individual photons.
- TDC:** Time-to-Digital Converter, a device for recording the time delay between two signals.

Acknowledgements

This thesis, and the Longitudinal Density Monitor itself, would have been impossible without the support of my supervisors at CERN and at the University of Liverpool. I would like to give my sincere thanks to Carsten Welsch, Andrea Boccardi, Enrico Bravin and Thibaut Lefevre.

Aurelie Goldblatt gave invaluable help in the installation of the LDM. Federico Roncarolo helped me to understand and use the BSRT. Alan Fisher helped to design the BSRT optical table and find a space for the LDM. Jean-Jacques Gras supported the LDM with ideas and analysis. Software for the LDM was written by Enrico Bravin, Maxim Andersen and Serkan Bozyigit.

Many other CERN colleagues, both in the BI group and at the CCC, gave help and advice without which it would have been much harder to complete this project.

My thanks to the members of the Bunch Current Normalisation working group, who have provided valuable feedback as well as data from the experiments for cross-calibration of the LDM, especially Massimiliano Ferro-Luzzi, Martino Gagliardi, Beate Heinemann, Gabriel Anders and Jaap Panman.

I thank all the members of DITANET and the Quasar Group, who have given such valuable advice, friendship and entertainment.

Finally, a big thank you to my wife, my friends and all the members of the CERN canoe club and the Randonnée section of the CERN ski club, for keeping me happy and sane throughout this whole process.

The work leading to this thesis was made possible by a DITANET Marie Curie fellowship, funded by the E.U. under contract PITN-GA-2008-215080.

References

- [1] O. Bruning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole, and P. Proudlock, “LHC design report”, CERN, Geneva, (2004)
- [2] S. Myers, “Large Hadron Collider commissioning and first operation,” *Phil. Trans. Royal Society A* **370** (2012) pp. 859–875.
- [3] CERN Press release, “LHC physics data taking gets underway at new record collision energy of 8TeV”, <http://press.web.cern.ch/press/PressReleases/Releases2012/PR10.12E.html>. [Accessed: 21 Aug 2012].
- [4] ATLAS Collaboration, “The ATLAS experiment at the CERN large hadron collider”, *JINST* **3** (2008) S08003.
- [5] ALICE Collaboration, “The ALICE experiment at the CERN LHC”, *JINST* **8** (2008) S08002.
- [6] S. Chatrchyan and CMS Collaboration, “The CMS experiment at the CERN LHC”, *JINST* **3** (2008) S08004.
- [7] LHCb Collaboration, “The LHCb Detector at the LHC”, *JINST* **3** (2008) S08005.
- [8] LHC Machine Outreach team, www.cern.ch [Accessed: 09/01/2013].
- [9] D. Boussard and T. Linnecar, “The LHC superconducting RF system”, LHC Project Report 316, CERN, Geneva (1999)
- [10] P. J. Bryant and K. Johnsen, *The principles of circular accelerators and storage rings* (Cambridge: Cambridge University Press, 1993)
- [11] M. Conte and W. W. MacKay, *An Introduction to the Physics of Particle Accelerators* (Singapore: World Scientific, 2008)
- [12] K. Wille, *The Physics of Particle Accelerators: an Introduction* (Oxford University Press, 1996)
- [13] M. A. Furman and W. C. Turner, “Beam-beam simulations for separated beams in the LHC”, CBP Note 350, LBNL, Berkeley (2000)
- [14] E. Shaposhnikova, “Abort gap cleaning and the RF system”, LHC performance workshop, Chamonix (2003) pp. 182–187.
- [15] E. Shaposhnikova, “Longitudinal motion of uncaptured particles in the LHC at 7 TeV”, LHC Project Note 338, CERN, Geneva (2004)
- [16] M. J. Syphers, “Some Notes on Longitudinal Emittance”, FNAL (2002), <http://home.fnal.gov/~syphers/Accelerators/tevPapers/LongEmitt.pdf> [Accessed: 16 October 2012]

- [17] P. Baudrenghien, A. Butterworth, M. Jaussi, T. Mastoridis, G. Papotti, E. Shaposhnikova, and J. Tuckmantel, “Longitudinal emittance blow-up in the LHC”, Proc. IPAC, San Sebastian (2011) pp. 1819–1821.
- [18] M. Benedikt, R. Cappi, M. Chanel, R. Garoby, M. Giovannozzi, S. Hancock, M. Martini, E. Métral, G. Metral, K. Schindl, and J.-L. Vallet, “Performance of the LHC pre-injectors”, Proc. HEACC, Tsukuba (2001)
- [19] A. Beuret, J. Borburgh, A. Blas, H. Burkhardt, C. Carli, M. Chanel, A. Fowler, M. Gourber-Pace, S. Hancock, and M. Hourican, “The LHC lead injector chain”, Proc. EPAC, Lucerne (2004) pp. 1153–1155.
- [20] H. Damerou, “Creation and Storage of Long and Flat Bunches in the LHC”, PhD Thesis, TU Darmstadt (2005)
- [21] M. E. Angoletta, A. Blas, A. Butterworth, A. Findlay, P. M. Leinonen, J. C. Molendijk, F. Pedersen, J. Sanchez-Quesada, and M. Schokker, “CERN’s PS Booster LLRF renovation: plans and initial beam tests”, Proc. IPAC (2010) pp. 1461–1463.
- [22] R. Garoby, “Bunch merging and splitting techniques in the injectors for high energy hadron colliders”, Proc. HEACC, Dubna (1998)
- [23] C. Fischer, “High sensitivity measurement of the longitudinal distribution of the LHC beams”, LHC-B-ES-0005, CERN, Geneva (2003)
- [24] H. Damerou, S. Hancock, and M. Schokker, “Longitudinal performance with high-density beams for the LHC in the CERN PS”, Proc. HB2010, Moschach (2010)
- [25] R. Steerenberg, G. Arduini, E. Benedetto, A. Blas, W. Höfle, E. Métral, M. Morvillo, C. Rossi, and G. Rumolo, “Nominal LHC beam instability observations in the CERN Proton Synchrotron”, Proc. PAC, Albuquerque (2007) pp. 4222–4224.
- [26] H. Damerou, S. Hancock, and M. Schokker, “Recent improvements of the RF beam control for LHC-type beams in the CERN PS”, Proc. IPAC, Kyoto (2010) pp. 2716–2718.
- [27] D. Manglunki, M. E. Angoletta, P. Baudrenghien, G. Bellodi, A. Blas, T. Bohl, C. Carli, E. Carlier, S. C. Cave, M. Chanel, K. Cornelis, H. Damerou, A. Findlay, S. Gilardoni, S. Hancock, J. M. Jowett, D. Küchler, S. Maury, E. Métral, S. Pasinelli, M. Schokker, G. Tranquille, B. Vanderpe, J. Wenninger, and U. Wehrle, “Ions for LHC: performance of the injector chain”, Proc. IPAC, San Sebastian (2011) pp. 2529–2531.
- [28] J. Tuckmantel, “The SPS / LHC longitudinal interface”, 9th LEP-SPS Performance Workshop, Chamonix (1999)
- [29] P. Baudrenghien, “Low-level RF”, CERN Accelerator School, Ebeltoft (2010) pp. 341–367.

- [30] J. M. Jowett, P. Baudrengnien, R. Bruce, C. Carli, D. Manglunki, T. Mertens, and D. Wollmann, "Heavy ions in 2011 and beyond", LHC Performance Workshop, Chamonix (2011) pp. 254 – 263.
- [31] K. Nakamura et al. (Particle Data Group), "2011 Review of Particle Physics: Accelerator Physics of Colliders", J. Phys. G. **37** (2010) 075021
- [32] P. Odier, M. Ludwig, and S. Thoulet, "The DCCT for the LHC Beam Intensity Measurement", Proc. DIPAC, Basel (2009)
- [33] 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, Santa Fe (2010)
- [34] J. Gras, C. Barschel, D. Belohrad, M. Ludwig, and P. Odier, "Optimization of the LHC beam current transformers for accurate luminosity determination", Proc. IPAC, San Sebastian (2011) pp. 1395–1397.
- [35] O. Bruning, "LHC challenges and upgrade options", J. Phys.: Conf. Ser. **110**, 11 (2008) p. 112002
- [36] J. Byrd, S. D. Santis, P. Denes, M. Placidi, W. Turner, and M. Zolotarev, "Design of an Abort Gap Monitor for the Large Hadron Collider", CBP Tech Note 329 (2005)
- [37] T. Lefevre, S. Bart Pedersen, A. Boccardi, E. Bravin, A. Jeff, A. Goldblatt, F. Roncarolo, and A. S. Fisher, "First Operation of the Abort Gap Monitors for LHC", Proc. IPAC, Kyoto (2010) pp. 2863–2865.
- [38] T. Bohl and J. F. Malo, "The APWL Wideband Wall Current Monitor", CERN-BE-2009-006, CERN, Geneva (2009)
- [39] C.P. Welsch, 2nd DITANET Topical Workshop on Longitudinal Profile Measurements, <http://indico.cern.ch/conferenceDisplay.py?confId=93401> [Accessed: 16 October 2012]
- [40] S. R. Mohanty, H. Bhuyan, N. K. Neog, R. K. Rout, and E. Hotta, "Development of Multi Faraday Cup Assembly for Ion Beam Measurements from a Low Energy Plasma Focus Device", Japanese Journal of Applied Physics, **44**, 7A (2005) pp. 5199–5205
- [41] D. Egger, O. Mete, M. Csatari, A. Dabrowski, S. Dobert, T. Lefevre, M. Olvegard, and M. Petrarca, "Performance Of The Time Resolved Spectrometer For The 5 MeV Photo-Injector PHIN", Proc. DIPAC, Hamburg (2011)
- [42] P. Finocchiaro, "Particle detectors for low intensity ion beam diagnostics", AIP Conference Proceedings, **862**, 1 (1999) pp. 862–865

- [43] M. Fiorini, V. Carassiti, A. Ceccucci, E. Cortina, A. Cotta Ramusino, G. Dellacasa, S. Garbolino, P. Jarron, J. Kaplon, A. Kluge, A. Mapelli, F. Marchetto, E. Martin, S. Martoiu, G. Mazza, M. Morel, M. Noy, G. Nuessle, F. Petrucci, P. Riedler, G. Aglieri Rinella, A. Rivetti, and S. Tiuraniemi, “The Gigatracker: An ultra-fast and low-mass silicon pixel detector for the NA62 experiment”, *Nucl. Instrum. and Meth. A*, **628**, 1, (2011) pp. 292–295
- [44] E. Griesmayer, B. Dehning, D. Dobos, E. Effinger and H. Pernegger “Diamond detectors as beam monitors”, *Proc. BIW, Santa Fe* (2010)
- [45] F. Schirru, B. S. N. Singh, L. Scruton, M. A. Bentley, S. P. Fox, A. Lohstroh, P. J. Sellin, A. Banu, M. McCleskey, B. T. Roeder, E. Simmons, A. A. Alharbi, L. Trache, M. Freer, and D. Parker, “Development of large area polycrystalline diamond detectors for fast timing application of high-energy heavy-ion beams”, *JInst*, **7**, 5 (2012) P05005
- [46] K. B. Unser, “Toroidal AC and DC current transformers for beam intensity measurements”, *Atomkernenergie Kerntechnik*, **47**, 1 (1985)
- [47] P. Forck, “Measurement of Beam Current”, *Joint Universities Accelerator School, Archamps* (2010)
- [48] P. Odier, “DCCT technology review”, *Proc. CARE-HHH-ABI workshop, Lyon* (2004) CARE-Conf-2004-023-HHH
- [49] C. Barschel, M. Ferro-Luzzi, J.-J. Gras, M. Ludwig, P. Odier, S. Thoulet, “Results of the LHC DCCT Calibration Studies”, *CERN-ATS-Note-2012-026-PERF*, CERN, Geneva (2012)
- [50] U. Raich, “Beam Diagnostics”, *CERN Accelerator School, Varna* (2010)
- [51] R. Capi, G. Cyvoct, J. Durand, M. Ruetten, and E. Schulte, “Single shot longitudinal shape measurements of nanosecond particle bunches”, *Proc. PAC, Washington* (1987) p. 643.
- [52] I. Wilke, a. MacLeod, W. Gillespie, G. Berden, G. Knippels, and A. van der Meer, “Single-Shot Electron-Beam Bunch Length Measurements”, *Phys. Rev. Letters*, **88**, 12 (2002)
- [53] B. Steffen, V. Arsov, G. Berden, W. Gillespie, S. Jamison, a. MacLeod, a. van der Meer, P. Phillips, H. Schlarb, B. Schmidt, and P. Schmußer, “Electro-optic time profile monitors for femtosecond electron bunches at the soft x-ray free-electron laser FLASH”, *Phys. Rev. ST-AB*, **12**, 3 (2009) pp. 1–16
- [54] P. Goldsmith and J. V. Jelley, “Optical Transition Radiation from protons entering metal surfaces”, *Philosophical Magazine*, **4**, 43 (1959) pp. 836–844
- [55] G. M. Garibyan, “Radiation from a Charged Particle Passing Through a Layered Medium”, *Journal of Experimental and Theoretical Physics*, **6** (1958) p. 1079

- [56] A. Aryshev, S. T. Boogert, D. Howell, P. Karataev, N. Terunuma, and J. Urakawa, "A novel method for sub-micrometer transverse electron beam size measurements using optical transition radiation", *J. Phys.: Conf. Ser.* **236** (2010) p. 012008
- [57] L. Wartski, S. Roland, J. Lasalle, M. Bolore, and G. Filippi, "Interference phenomenon in optical transition radiation and its application to particle beam diagnostics and multiple-scattering measurements", *Journal of Applied Physics*, **46**, 8 (1975) p. 3644
- [58] Y. N. Dnestrovskiy and D. P. Kostomarov, "Radiation of ultra-relativistic charges moving across round aperture in flat screen", *Doklady Akademii Nauk SSSR*, **4** (1959)
- [59] A. P. Potylitsyn, "Transition radiation and diffraction radiation: Similarities and differences", *Nucl. Instrum. Meth. B*, **145** (1998) pp. 169–179
- [60] A. Cianchi, M. Castellano, L. Catani, E. Chiadroni, K. Honkavaara, and G. Kube, "Nonintercepting electron beam size monitor using optical diffraction radiation interference", *Phys. Rev. ST-AB*, **14**, 10 (2011)
- [61] P. A. Cherenkov, "Visible emission of clean liquids by action of γ radiation", *Doklady Akademii Nauk SSSR*, **2**, 451 (1934)
- [62] I. Frank and I. Tamm, "Coherent visible radiation of fast electrons passing through matter", *Doklady Akademii Nauk SSSR*, **14** (1937) pp. 109–114
- [63] S. Mallows, E. B. Holzer, J. V. Hoorne, A. Mechev, and C. Welsch, "Requirements of a beam loss monitoring system for the CLIC two beam modules", *Proc. IPAC, San Sebastian* (2011) pp. 2385–2387
- [64] S. J. Smith and E. M. Purcell, "Visible Light from Localized Surface Charges Moving across a Grating", *Phys. Rev.*, **92**, 4 (1953) p. 1069
- [65] G. Kube, H. Backe, W. Lauth and H. Schope, "Smith-Purcell radiation in view of particle beam diagnostics", *Proc. DIPAC, Mainz* (2003)
- [66] Y. Shibata, S. Hasebe, K. Ishi, S. Ono, M. Ikezawa, T. Nakazato, M. Oyamada, S. Urasawa, T. Takahashi, T. Matsuyama, K. Kobayashi, and Y. Fujita, "Coherent Smith-Purcell radiation in the millimeter-wave region from a short-bunch beam of relativistic electrons", *Phys. Rev. E*, **57**, 1 (1998) pp. 1061–1074
- [67] R. Bartolini, C. Clarke, N. Delerue, G. Doucas, and A. Reichold, "Electron bunch profile reconstruction in the few fs regime using coherent Smith-Purcell radiation", *JInst*, **7**, 1 (2012) P01009
- [68] E. D. Bourret-Courchesne, S. E. Derenzo, and M. J. Weber, "Development of ZnO:Ga as an ultra-fast scintillator", *Nucl. Instrum. and Meth. A*, **601**, 3 (2009) pp. 358–363

- [69] H. Matsuda, A. Miura, H. Irie, S. Tanaka, K. Ito, S. Fujisaki, T. Toyonaka, H. Takahashi, H. Chiba, S. Irikura, R. Takeyari, and T. Harada, "High-sensitivity and wide-dynamic-range 10 Gbit/s APD/preamplifier optical receiver module", *Electronics Letters*, **38**, 13 (2002) p. 650
- [70] K. Scheidt, "Review of streak cameras for accelerators: features, applications and results", *Proc. EPAC, Vienna* (2000) p. 182.
- [71] www.hamamatsu.com [Accessed: 09/01/2013].
- [72] C. Bonté, M. Harmand, F. Dorchies, S. Magnan, V. Pitre, J.-C. Kieffer, P. Audebert, and J.-P. Geindre, "High dynamic range streak camera for subpicosecond time-resolved x-ray spectroscopy", *Rev. Sci. Instrum.*, **78**, 4 (2007) p. 043503
- [73] T. Nordlund, *Topics in Fluorescence Spectroscopy*, vol. 1 (Boston: Kluwer Academic Publishers, 2002)
- [74] J.-F. Beche, J. Byrd, S. De Santis, P. Denes, M. Placidi, W. Turner, and M. Zolotorev, "Measurement of the Beam Longitudinal Profile in a Storage Ring by Non-Linear Laser Mixing", *Proc. BIW, Knoxville* (2004) pp. 112–119.
- [75] J. Zhou, N. Park, K. J. Vahala, M. A. Newkirk, and B. I. Miller, "Four-wave mixing wavelength conversion efficiency in semiconductor traveling-wave amplifiers measured to 65 nm of wavelength shift", *Photonics Technology Letters*, **6**, 8 (1994) p. 984
- [76] R. Salem, M. A. Foster, A. C. Turner-Foster, D. F. Geraghty, M. Lipson, and A. L. Gaeta, "High-speed optical sampling using a silicon-chip temporal magnifier", *Optics express*, **17**, 6 (2009) pp. 4324–9
- [77] M. A. Foster, R. Salem, D. F. Geraghty, A. C. Turner-Foster, M. Lipson, and A. L. Gaeta, "Silicon-chip-based ultrafast optical oscilloscope", *Nature*, **456** (2008) pp. 81–85
- [78] R. Trebino, *Frequency-resolved optical gating: the measurement of ultra-short optical pulses* (Boston: Kluwer, 2000)
- [79] J. Kerr, "A new relation between electricity and light: Dielectrified media birefringent", *Philosophical Magazine*, **50**, 332 (1875) pp. 337–348
- [80] V. Vinetskiĭ and N. Kukhtarev, "Dynamic self-diffraction of coherent light beams", *Soviet Physics Uspekhi*, **22**, 9 (1979) p. 742
- [81] K. DeLong, R. Trebino, J. Hunter, and W. White, "Frequency-resolved optical gating with the use of second-harmonic generation", *JOSA B*, **11**, 11 (1994) pp. 2206-2215
- [82] G. Taft, A. Rundquist, M. Murnane, I. Christov, H. Kapteyn, K. DeLong, D. Fittinghof, M. Krumbugel, J. Sweetser, and R. Trebino, "Measurement of 10-fs laser pulses", *Selected Topics in Quantum Electronics*, **2**, 3 (1996) pp. 575–585

- [83] G. Angelova-Hamberg, V. Ziemann, P. V. D. Meulen, P. Salen, M. Larsson, H. Schlarb, J. Bodewadt, A. Winter, F. Lohl, E. Saldin, E. Schneidmiller, M. Yurkov, S. Khan, and A. Meseck, “Recent results from the optical replica synthesizer experiment in FLASH”, Proc. DIPAC, Basel (2009)
- [84] T. Nakazato, M. Oyamada, N. Niimura, S. Urasawa, O. Konno, A. Kagaya, R. Kato, T. Kamiyama, Y. Torizuka, T. Nanba, Y. Kondo, Y. Shibata, K. Ishi, T. Ohsaka, and M. Ikezawa, “Observation of Coherent Synchrotron Radiation”, Phys. Rev. Letters, **63**, 12 (1989) pp. 1245–1248
- [85] K. Lekomtsev, G. Blair, G. Boorman, P. Karataev, M. Micheler, R. Corsini, and T. Lefevre, “Coherent Diffraction Radiation as a tool for non-invasive bunch length diagnostics, theory”, DITANET Workshop on Longitudinal Beam Profile Measurements, Daresbury (2010)
- [86] S. Wesch, C. Behrens, E. Hass, and B. Schmidt, “Fast Bunch Profile Monitoring with THz Spectroscopy of Coherent Radiation at FLASH”, Proc. BIW, Newport News (2012)
- [87] M. Tordeux, M. Labat, F. Dohou, and O. Marcouillé, “Bunch length measurements from the incoherent synchrotron radiation fluctuation at SOLEIL”, Proc. IPAC, San Sebastian (2011)
- [88] F. Sannibale, G. Stupakov, M. Zolotarev, D. Filippetto, and L. Jägerhofer, “Absolute bunch length measurements by incoherent radiation fluctuation analysis”, Phys. Rev. ST-AB, **12**, 3 (2009)
- [89] J. J. Larmor, “A Dynamical Theory of the Electric and Luminiferous Medium,” Phil. Trans. Royal Society A, **190** (1897) pp. 205–493
- [90] D. Iwanenko and I. Pomeranchuk, “On the maximal energy attainable in a betatron”, Phys. Rev., **65** (1944) p. 353
- [91] V. A. Bordovitsyn, *Synchrotron Radiation Theory and its development* (Singapore: World Scientific, 1999)
- [92] F. Elder, A. Gurewitsch, R. Langmuir, and H. Pollock, “Radiation from electrons in a synchrotron”, Phys. Rev., **71** (1947) pp. 829–830
- [93] H. Wiedemann, *Particle Accelerator Physics* (Berlin Heidelberg: Springer, 2007)
- [94] J. Schwinger, “On the classical radiation of accelerated electrons”, Phys Rev, **75**, 12 (1949) pp. 1912–1925
- [95] G. Geloni, E. Saldin, E. Schneidmiller, and M. Yurkov, “On limitations of Schwinger formulae for coherent synchrotron radiation produced by an electron bunch moving along an arc of a circle”, Nucl. Instrum. and Meth. A, **528**, 1 (2004), pp. 520–524

- [96] R. Coisson, "Angular-spectral distribution and polarization of synchrotron radiation from a 'short' magnet", *Phys. Rev. A*, **20**, 2 (1979)
- [97] R. A. Bosch, "Edge radiation in an electron storage ring", *Il Nuovo Cimento D*, **20**, 4 (1998) pp. 483–493
- [98] K. Wille, "Synchrotron Radiation", Joint Universities Accelerator School, Archamps (2010)
- [99] D. Attwood, *Soft X-rays and Extreme Ultraviolet Radiation* (Cambridge University Press, 2007)
- [100] A. Hoffmann, *The Physics of Synchrotron Radiation* (Cambridge University Press, 2004)
- [101] K.-J. Kim, "Characteristics of synchrotron radiation", *AIP Conference Proceedings*, **184**, 1 (1989) pp. 565–632
- [102] L. Evans and P. Bryant, "LHC machine", *JINST*, **3** (2008) S08001
- [103] L. Ponce, R. Jung, and F. Méot, "LHC proton beam diagnostics using synchrotron radiation", CERN-2004-007, CERN, Geneva (2004)
- [104] BNL Magnet Group, "Design Report for the interaction region dipoles and the RF region dipoles", BNL, Upton, New York (1999)
- [105] R. Maccaferri, S. Bettoni, D. Tommasini, and W. Venturini Delsolaro, "Manufacture and Test of the Prototype 5T Superconducting Undulator for the LHC Synchrotron Radiation Profile Monitor", LHC Project Report 894, CERN, Geneva (2005)
- [106] R. Jung, P. Komorowski, L. Ponce, and D. Tommasini, "The LHC 450GeV to 7 TeV synchrotron radiation profile monitor using a superconducting undulator", *Proc. BIW*, Upton, New York (2002)
- [107] www.polymicro.com [Accessed: 16 October 2012][108] www.cvimellesgriot.com [Accessed: 09/01/2013]
- [109] <http://refractiveindex.info/> [Accessed: 09/01/2013]
- [110] www.corning.com [Accessed: 09/01/2013]
- [111] A. Fisher, "The LHC Synchrotron-Light Telescope (BSRT) and Abort-Gap Monitor (BSRA), Using Dipole and Undulator Radiation, for both Protons and Ions", LHC Performance Note 014, CERN, Geneva (2009)

- [112] T. Lefevre, E. Bravin, G. Burtin, A. Fisher, A. Guerrero, A. Jeff, A. Rabiller, and F. Roncarolo, “First Beam Measurements with the LHC Synchrotron Light Monitors”, Proc. IPAC, Kyoto (2010) pp. 1104–1106
- [113] <http://www.esrf.eu/Accelerators/Groups/InsertionDevices/Software/SRW> [Accessed: 16 October 2012]
- [114] O. Chubar and P. Elleaume, “Accurate and efficient computation of synchrotron radiation in the near field region”, Proc. EPAC, Stockholm (1998) p. 1177.
- [115] T. Tanaka and H. Kitamura, “Recent Progress of the Synchrotron Radiation Calculation Code SPECTRA”, AIP Conference Proceedings, **879** (2007) pp. 355-358
- [116] Igor Pro Manual (version 6), www.wavemetrics.com/products/igorpro/manual.htm [Accessed: 16 October 2012]
- [117] R. Dejus and a Luccio, “Program UR: General purpose code for synchrotron radiation calculations”, Nucl. Instrum. and Meth. A, **347**, 1–3 (1994) pp. 61–66
- [118] F. Meot, “The ray-tracing code Zgoubi”, Nucl. Instrum. and Meth. A, **427**, 1–2 (1999) pp. 353–356
- [119] A. Einstein, “Über einen die erzeugung und verwandlung des liches betreffenden heuristischen gesichtspunkt”, Annalen der Physik, **17** (1905) pp. 132–148
- [120] W. Becker, *Advanced Time-correlated Single Photon counting techniques* (Berlin Heidelberg: Springer, 2005)
- [121] S. Cova, M. Bertolaccini, and C. Bussolati, “The measurement of luminescence waveforms by single-photon techniques”, physica status solidi (a), **18**, 1 (1973) pp. 11–62
- [122] D. Renker and E. Lorenz, “Advances in solid state photon detectors”, JInst, **4**, 04 (2009)
- [123] R. Thew, N. Curtz, P. Eraerds, N. Walenta, J. Gautier, E. Koller, J. Zhang, N. Gisin, and H. Zbinden, “Approaches to single photon detection”, Nucl. Instrum. and Meth. A, **610**, 1 (2009)
- [124] G. N. Gol’tsman, O. Okunev, G. Chulkova, A. Lipatov, A. Semenov, K. Smirnov, C. Williams, B. Voronov, A. Dzardanov, and R. Sobolewski, “Picosecond superconducting single-photon optical detector”, Applied Physics Letters **2**, **79**, 6 (2001) pp. 705–707
- [125] X. Zheng, Y. Xu, R. Sobolewski, R. Adam, M. Mikulics, M. Siegel, and P. Kordoš, “Femtosecond response of a free-standing LT-GaAs photoconductive switch”, Appl. Opt, **42** (2003) pp. 1726–1731
- [126] R. J. McIntyre, “Theory of Microplasma Instability in Silicon”, Journal of Applied Physics, **32**, 6 (1961) pp. 983–996

- [127] R. H. Haitz, "Model for the Electrical Behavior of a Microplasma", *Journal of Applied Physics*, **35**, 5 (1964) pp. 1370–1377
- [128] I. Prochazka, "Semiconducting single photon detectors: the state of the art", *physica status solidi (c)*, **2**, 5 (2005) pp. 1524–1532
- [129] A. D. Renker, "Geiger-mode avalanche photodiodes, history, properties and problems", *Nucl. Instrum. and Meth. A*, **567**,1 (2006) pp. 48–56
- [130] B. Jalali and S. Fathpour, "Silicon Photonics", *Journal of Lightwave Technology*, **24**, 12 (2006) pp. 4600–4615
- [131] B. G. Streetman and S. Banerjee, *Solid state electronic devices*, 5th ed. (New Jersey: Prentice Hall, 2000) p. 524.
- [132] P. C. M. Owens, J. G. Rarity, P. R. Tapster, D. Knight, and P. D. Townsend, "Photon counting with passively quenched germanium avalanche photodiodes", *Applied Optics*, **33**, 30 (1994) p. 6895
- [133] N. Namekata, S. Adachi, and S. Inoue, "1.5 GHz single-photon detection at telecommunication wavelengths using sinusoidally gated InGaAs / InP avalanche photodiode", *Optics Express*, **17**, 8 (2009) pp. 1827–1829
- [134] A. Dixon, J. Dynes, Z. Yuan, A. Sharpe, A. Bennett, and A. Shields, "Ultrashort dead time of photon-counting InGaAs avalanche photodiodes", *Applied Physics Letters*, **94** (2009) p. 231113
- [135] D. Dravins, D. Faria, and B. Nilsson, "Avalanche diodes as photon-counting detectors in astronomical photometry", *Proc. SPIE* (2000) pp. 298–307
- [136] M. M. Hayat, M. A. Itzler, D. A. Ramirez, and G.J. Rees, "Model for passive quenching of SPADs", *Proc. SPIE* (2010) p. 76082B.
- [137] M. Stipčević, "Active quenching circuit for single-photon detection with Geiger mode avalanche photodiodes", *Applied Optics*, **48**, 9 (2009)
- [138] A. Gallivanoni, I. Rech, D. Resnati, M. Ghioni, and S. Cova, "Monolithic active quenching and picosecond timing circuit suitable for large-area single-photon avalanche diodes", *Optics Express*, **14** (2006) pp. 5021–5030
- [139] F. Zappa, S. Tisa, A. Gulinatti, A. Gallivanoni, and S. Cova, "Complete single-photon counting and timing module in a microchip", *Optics letters*, **30**, 11 (2005) pp. 1327–1329
- [140] S. Cova, A. Lacaita, M. Ghioni, G. Ripamonti, and T. A. Louis, "20-ps timing resolution with single-photon avalanche diodes", *Rev. Sci. Instrum.*, **60**, 6 (1989)

- [141] E. Bravin, “Dead time effects on single photon counting for the longitudinal density monitor of LHC”, CERN-AB-Note-2006-002, CERN, Geneva (2006)
- [142] S. D. Poisson, *Probabilité des jugements en matière criminelle et en matière civile, précédées des règles générales du calcul des probabilités* (Paris: Bachelier, 1837)
- [143] S. Cova, A. Lacaita, and G. Ripamonti, “Trapping phenomena in avalanche photodiodes on nanosecond scale”, *Electron Device Letters*, **12**, 12 (1991) pp. 685–687
- [144] K. E. Jensen, P. I. Hopman, E. K. Duerr, E. a. Dauler, J. P. Donnelly, S. H. Groves, L. J. Mahoney, K. a. McIntosh, K. M. Molvar, A. Napoleone, D. C. Oakley, S. Verghese, C. J. Vineis, and R. D. Younger, “Afterpulsing in Geiger-mode avalanche photodiodes for 1.06 μm wavelength”, *Applied Physics Letters*, **88**, 13 (2006) p. 133503
- [145] M. Höbel and J. Ricka, “Dead-time and afterpulsing correction in multiphoton timing with nonideal detectors”, *Rev. Sci. Instrum.*, **65**, 7 (1994) p. 2326
- [146] S. Polyakov, M. Ware, and A. Migdall, “High-accuracy calibration of photon-counting detectors”, *Proc. SPIE*, **6372**, 301 (2006) 63720J
- [147] www.picoquant.com [Accessed: 16 October 2012]
- [148] J. Zhang, R. Thew, J. Gautier, N. Gisin, and H. Zbinden, “Comprehensive Characterization of InGaAs – InP Avalanche Photodiodes at 1550 nm With an Active Quenching ASIC”, *Journal of Quantum Electronics*, **45**, 7 (2009) pp. 792–799
- [149] X. Jiang, M. A. Itzler, R. Ben-michael, K. Slomkowski, M. A. Krainak, S. Wu, and X. Sun, “Afterpulsing Effects in Free-Running InGaAsP Single-Photon Avalanche Diodes”, *Journal of Quantum Electronics*, **44**, 1 (2008) pp. 3–11
- [150] Y. Kang, H. X. Lu, Y. H. Lo, D. S. Bethune, and W. P. Risk, “Dark count probability and quantum efficiency of avalanche photodiodes for single-photon detection”, *Applied Physics Letters*, **83**, 14 (2003) p. 2955
- [151] W. J. Kindt and H. W. Van Zeijl, “Modelling and fabrication of Geiger mode avalanche photodiodes”, *IEEE Transactions on Nuclear Science*, **45**, 3 (1998) pp. 715–719
- [152] B. X. Yang, W. E. Norum, S. Shoaf, and J. Stevens, “Bunch-by-bunch diagnostics at the APS using time- correlated single-photon counting techniques”, *Proc. BIW, Santa Fe* (2010) p. 238.
- [153] M. Krainak, “Photoionization of trapped carriers in avalanche photodiodes to reduce afterpulsing during Geiger-mode photon counting”, *Proc. CLEO, Baltimore* (2005) pp. 588–590.
- [154] www.microphotondevices.com/products_pdm.asp [Accessed: 18 October 2012]

- [155] G. Ribordy, J.-D. Gautier, H. Zbinden, and N. Gisin, "Performance of InGaAs/InP Avalanche Photodiodes as Gated-Mode Photon Counters", *Applied Optics*, **37**, 12 (1998) p. 2272
- [156] U. S. NIM Committee, "Standard NIM Instrumentation System", Springfield, Virginia (1990)
- [157] F. Pozar, "The Time-to-Digital converter," *Nuclear Instruments and Methods*, **74** (1969) pp. 315–321
- [158] www.idquantique.com [Accessed: 18 October 2012]
- [159] www.agilent.com [Accessed: 18 October 2012]
- [160] B. G. Taylor, "Timing Distribution at the LHC", 8th Workshop on Electronics for LHC Experiments, Colmar (2002)
- [161] D. Dominguez, J. Gras, J. Lewis, J. Savioz, J. Serrano, and F. Ballester, "An FPGA based multiprocessing CPU for beam synchronous timing in CERN's SPS and LHC", *Proc. ICALEPCS*, Gyeongju (2003)
- [162] www.picmg.org [Accessed: 18 October 2012]
- [163] www.gocct.com/ [Accessed: 18 October 2012]
- [164] G. Keiser, *Optical Fiber Communications, Encyclopedia of Telecommunications* (Hoboken: John Wiley & Sons, Inc., 2003)
- [165] W. Reeves, J. Knight, P. Russell, and P. Roberts, "Demonstration of ultra-flattened dispersion in photonic crystal fibers", *Optics express*, **10**, 14 (2002) pp. 609–13
- [166] T. A. Birks, D. Mogilevtsev, J. C. Knight, and P. St. J. Russell, "Dispersion compensation using single-material fibers", *IEEE Photonics Technology Letters*, **11**, 6 (1999) pp. 674–676
- [167] A. Guerrero, J.-J. Gras, J.-L. Nougaret, M. Ludwig, M. Arruat, and S. Jackson, "CERN Front-End Software Architecture for accelerator controls", *Proc. ICALEPCS*, Gyeongju (2003) pp. 342–344
- [168] C. Roderick, R. Billen, and M. G. Pace, "The CERN Accelerator Measurement Database: On the Road to Federation", *Proc. ICALEPCS*, Grenoble (2011)
- [169] G. Anders, N. Bacchetta, V. Balagura, C. Barschel, D. Belohrad, H. Burkhardt, M. Ferro-Luzzi, C. Gabaldon, M. Gagliardi, J.-J. Gras, P. Hopchev, A. Jeff, W. Kozanecki, M. Ludwig, D. Marlow, K. Oyama, J. Panman, S. White, and A. Zuranski, "Study of the relative LHC bunch populations for luminosity calibration", *CERN-ATS-Note-2012-028-PERF*, CERN, Geneva (2012)

- [170] M. Ferro-Luzzi, “Determination of the luminosity at the LHC experiments”, 35th International Conference on High Energy Physics, Paris (2010)[171] S. van der Meer, “Calibration of the effective beam height in the ISR”, CERN-ISR-PO-68-31, CERN, Geneva (1968)
- [172] V. Balagura, “Notes on Van der Meer scan for absolute luminosity measurement”, Nucl. Instrum. and Meth. A, **654**, 1 (2011) pp. 634–638
- [173] M. Ferro-Luzzi, “Proposal for an absolute luminosity determination in colliding beam experiments using vertex detection of beam-gas interactions”, CERN-PH-EP/2005-023, CERN, Geneva (2005)
- [174] ATLAS Collaboration, “Updated Luminosity Determination in pp Collisions at 7 Tev using the ALAS detector”, ATLAS-CONF-2011-011, CERN, Geneva (2011)
- [175] A. Alici, G. Anders, N. Bacchetta, V. Balagura, C. Barschel, E. Bravin, H. Burkhardt, M. Gagliardi, J. J. Gras, B. Heinemann, P. Hopchev, M. Huhtinen, A. Jeff, W. Kozanecki, S. P. Griso, J. Panman, S. White and A. Zuranski, “Study of the LHC ghost charge and satellite bunches for luminosity calibration”, CERN-ATS-Note-2012-029-PERF, CERN, Geneva (2012)
- [176] T. Martin, “Development and online operation of minimum bias triggers in ATLAS”, JInst, **5**, 12 (2010) C12051
- [177] M. Ferro-Luzzi, “LHC 2011 operation - as viewed from the experiments”, LHC performance workshop, Chamonix (2012)
- [178] I. Kralik, “Enhanced satellite analysis in ALICE for 2011 dedicated fills”, ALICE Note, CERN, Geneva (2012)
- [179] E. Bravin, A. Boccardi, B. Dehning, J. Emery, T. Lefevre, J. Gras, A. Jeff, A. Rabiller, F. Roncarolo, M. Sapinski, V. Kain, M. Pojer, H. Bartosik, and B. Salvant, “LHC Transvers Profile Monitors studies (MD on May 6th, 2011)”, CERN-ATS-Note-2011-049 MD, CERN, Geneva (2011)
- [180] X. Ma and G.-S. Kuo, “Optical switching technology comparison: optical MEMS vs. other technologies”, IEEE Optical Communications, **41**, 11 (2003) pp. 16–23
- [181] S. Nakamura, “Experimental investigation on high-speed switching characteristics of a novel symmetric Mach–Zehnder all-optical switch”, Applied physics letters, **65**, 3 (1994) pp283-285
- [182] K. Hall and K. Rauschenbach, “100-Gbit/s bitwise logic”, Optics letters, **23**, 16 (1998) pp. 1271–1273
- [183] F. Pockels, “Über den Einfluss des electro-statischen Feldes auf das optische Verhalten piezoelektrischer Kristalle”, Abh. Gottingen Ges. Wiss. Math. Physik., **39** (1893) pp. 1–204

- [184] K. Tamura, "Development of a Fast Light Shutter for Beam Diagnostics", Spring-8 Annual Report, Harima (1998)
- [185] M. Gottlieb, C. L. M. Ireland, and J. M. Ley, *Electro-optic and acousto-optic scanning and deflection* (New York: Marcel Dekker, 1983)
- [186] K. Nakamura, J. Miyazu, M. Sasaura, and K. Fujiura, "Wide-angle, low-voltage electro-optic beam deflection based on space-charge-controlled mode of electrical conduction in $\text{KTa}_{1-x}\text{Nb}_x\text{O}_3$ ", *Applied Physics Letters*, **89**, 13 (2006)
- [187] "www.leysop.com." [Accessed : 18 October 2012]
- [188] S. Payne, Private communication (2009)
- [189] J. Z. Jun Kang, Wei Zhang, Hui Wei, Shaohe Chen, "Pulse reshaping by using a liquid crystal spatial light modulator and deflector for producing a specific waveform", *Chinese Optics Letters*, **4**, 3 (2006) p. 184