Spatial and Temporal Beam Profiles for the LHC using Synchrotron Light

A. Jeff*[a,d], S. Bart Pedersen[a], A. Boccardi[a], E. Bravin[a], A.S. Fisher[b],
A. Guerrero Ollacarizqueta[a],
T. Lefèvre[a], A. Rabiller[a], C. P. Welsch[c,d]
[a]CERN, CH-1211 Geneva 23, Switzerland;
[b]SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA
94025, USA
[c]The Cockcroft Institute, Warrington, UK
[d]University of Liverpool, Liverpool, UK

Spatial and Temporal Beam Profiles for the LHC using Synchrotron Light

A. Jeff*[a,d], S. Bart Pedersen[a], A. Boccardi[a], E. Bravin[a], A.S. Fisher[b], A. Guerrero Ollacarizqueta[a],
T. Lefèvre[a], A. Rabiller[a], C. P. Welsch[c,d]
[a]CERN, CH-1211 Geneva 23, Switzerland;
[b]SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025, USA
[c]The Cockcroft Institute, Warrington, UK
[d]University of Liverpool, Liverpool, UK

ABSTRACT

Synchrotron radiation is emitted whenever a beam of charged particles passes though a magnetic field. The power emitted is strongly dependent on the relativistic Lorentz factor of the particles, which itself is proportional to the beam energy and inversely proportional to the particle rest mass. Thus, synchrotron radiation is usually associated with electron accelerators, which are commonly used as light sources. However the largest proton machines reach sufficiently high energies to make synchrotron light useful for diagnostic purposes.

The Large Hadron Collider at CERN will accelerate protons up to an energy of 7TeV. An optical arrangement has been made which focuses synchrotron light from two LHC magnets to image the cross-section of the beam. It is also planned to use this setup to produce a longitudinal profile of the beam by use of fast Single Photon Counting. This is complicated by the bunched nature of the beam which needs to be measured with a very large dynamic range.

In this contribution we present early experimental data of the transverse LHC beam profile together with a scheme for measuring the longitudinal profile with a time resolution of 50 ps. It includes the use of a gating regime to increase the dynamic range of the photon counter and a three-stage correction algorithm to compensate for the detector's deadtime, afterpulsing and pile-up effects.

Keywords: Synchrotron light, LHC, time-correlated single photon counting, Avalanche Photo Diode, fast gating

1. INTRODUCTION

When a charged particle is accelerated, it emits electromagnetic (EM) radiation. In a circular particle accelerator, centripetal acceleration is provided by the bending magnets, and the EM radiation which is thus emitted is called synchrotron radiation. In the particle's own frame of reference, the radiation has the well-known form of a Hertzian dipole. However, if the particle is moving with relativistic velocity, the shape of the radiation field is distorted when viewed in the laboratory frame of reference. Sideways and backwards components are eliminated, leaving a tight beam of radiation shining forwards i.e. in the direction of motion of the particle. It is this beam which is exploited in purpose-built accelerators known as Synchrotron Light Sources.

The total power radiated by a relativistic charged particle following a given curved path is proportional to γ^4, where $\gamma = \frac{1}{1 - v^2/c^2}$ is the relativistic Lorentz factor[1]. It can also be given by $\gamma = \frac{E}{m_0 c^2}$ where E is the total energy of the moving particle and $m_0 c^2$ is its rest mass energy. Thus, for a given energy the radiated power is inversely proportional to the fourth power of the particle's rest mass. This is why Synchrotron Light Sources are electron accelerators. The Large Hadron Collider (LHC) is however a proton accelerator. The amount of synchrotron radiation emitted is therefore less, by a factor of $(m_p/m_e)^4 \approx 10^{11}$, than would be emitted by an electron accelerator running at the same energy. Synchrotron radiation in proton accelerators has traditionally been ignored for this reason. However, at the unprecedented particle energies that the LHC will reach, it is no longer negligible. Instead, it can be a useful tool for beam diagnostics.

*adam.jeff@cern.ch; phone 0041 22 76 77179

2. PRODUCTION AND EXTRACTION OF LIGHT

2.1 Production of synchrotron light in the bending magnet

The Beam Synchrotron Radiation Telescope (BSRT) [2 & 3] is used in the LHC to monitor the beams using the synchrotron light which they produce. Due to the optics and sensors chosen, only the visible part of the synchrotron radiation spectrum is used. Since the LHC has two counter-rotating beams, two separate and symmetrical BSRTs have been installed, one attached to each beam pipe. The following description can apply to either.

The primary source of the synchrotron light is a 9m long superconducting bending magnet with a uniform vertical magnetic field. In order to keep the beam trajectory constant, the magnetic field strength is increased in proportion to the beam energy. The field strength in this magnet is varied from 0.25T to a maximum of 3.9T. The nature of the synchrotron light emitted in the bending magnet can be calculated analytically[4] or computed using a simulation code such as Synchrotron Radiation Workshop (SRW)[5]. Not only the power emitted by each proton but also the spectrum and angular distribution of the light depend on the γ of the protons i.e. on their energy. The light will therefore change during the LHC's accelerating cycle. The total light emitted is simply the sum of the light from individual protons. Thus the light intensity can be used as a proxy for proton density at different points along the beam.

Synchrotron light is emitted in the forwards direction as a cone with a half-opening angle of $1/\gamma$, so as the particles are accelerated the beam becomes narrower. In addition, the critical frequency ω_c varies with γ^3. The critical frequency is defined here as the frequency which divides the spectrum into two halves of equal radiated power. It is close to but not the same as the peak frequency. When protons are injected into the LHC they have E=450GeV and γ=480, and the critical frequency is in the IR. Combined with the dependence of power on γ^4 as mentioned above, this means there is very little light emitted in the visible spectrum. Only when the energy is greater than 1TeV is there any significant visible light produced.

2.2 Use of an undulator magnet to produce additional synchrotron light

In order to allow the BSRT to be used at lower beam energies, an additional magnet has been installed just upstream of the bending magnet. This is an undulator magnet consisting of alternating regions of upwards and downwards magnetic field. The design of the magnet can be found here[6].This causes the particles to undulate but has no effect on average, i.e. the particles leave the magnet on the same trajectory regardless of whether the magnet is switched on or not. The undulations do however cause the emission of synchrotron light and in this case the emitted spectrum is dominated not by the critical frequency but by interference effects arising from the periodic structure of the undulator. This particular undulator has been designed to emit visible synchrotron light at proton energies below those at which the radiation from the bending magnet can be used.

2.3 Separating photons and protons

Since the synchrotron light is emitted in the direction of motion of the particle, the light and particle beams initially travel together. However the protons are bent in the magnetic field while the photons are not. As the protons are deviated the direction of the emitted light also changes, so that the beam is scanned across the inside of the beam pipe. The light emitted at the end of the magnet leaves with the same direction as the particles and so cannot be captured, however the light emitted by the undulator and at the start of the bending magnet has a small angular separation from the particle beam. Due to the very large radius of the LHC, the bending angle of this magnet is only 1.4mrad, and the extraction mirror must be located 21m downstream of the end of the magnet in order to allow it to be inserted without perturbing the beam. The light is then reflected out of the beam pipe through a fused-silica window.

2.4 Amount of light produced

The amount of light hitting the extraction mirror varies strongly with the particle energy. It can be computed using SRW, (Figure 1). This is in close agreement with calculated values[7]. It is at a minimum at the cross-over point around 1TeV proton energy, where the light from the undulator has passed into the UV but the light from the bending magnet still peaks in the IR.

Protons are injected into the LHC in bunches rather than as a continuous beam. Each bunch takes 89µs to complete a revolution. During the commissioning of the machine pilot bunches are injected which each contain around 5×10^9 protons. With a single pilot bunch in the machine we would thus have around 5×10^{13} protons per second passing through

the magnets. During normal running, however, 2808 full bunches will be injected each containing 1.7×10^{11} protons. We would then have about 5×10^{18} passes per second. The rate of useable photons varies from 5×10^7/s (single pilot bunch at 1TeV) to 2×10^{16}/s (full ring at 7TeV).
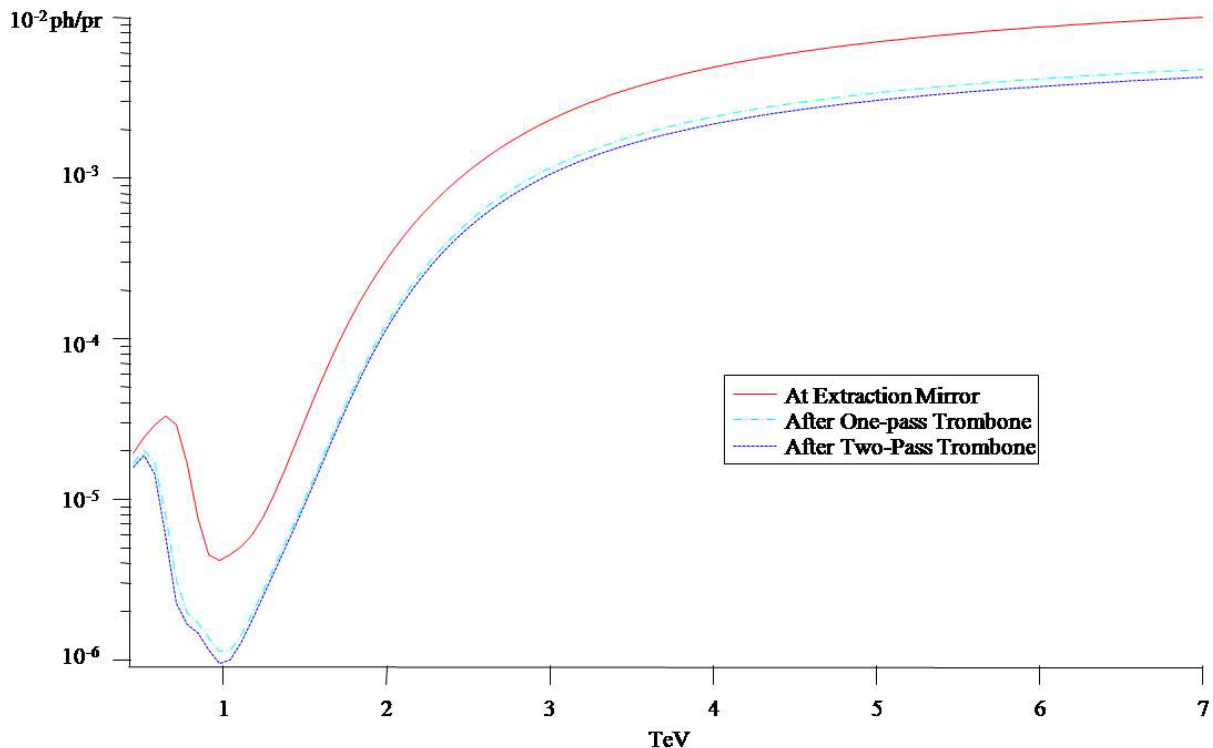

Figure 1. Average number of visible light photons emitted by each proton which passes through the magnets, as a function of proton energy. Upper line, all visible photons hitting the extraction mirror. Lower lines, visible photons available after passing through the window and reflection from 8 (dashed blue line) or 12 (black line) mirrors of the optics.

3. TRANSVERSE PROFILES

3.1 Optical arrangement

Synchrotron light exits the beam pipe through a fused-silica window and falls first on a steerable mirror which is used to correct for any misalignment. Two spherical mirrors are then used to focus the light and provide a demagnification of 0.3. Two optical density filter wheels providing a total attenuation of up to 10^5 are necessary to avoid saturation of the cameras given the very large variation of light intensity.

As described above, the source of synchrotron light varies with beam energy. At low energies the undulator is the main source; as the beam energy increases synchrotron light is emitted from the leading edge of the bending magnet and at higher energy still it is emitted uniformly through the bending magnet. For sharp imaging it is therefore necessary to introduce an optical delay line before the focusing mirrors. A translation stage allows an optical 'trombone' to be lengthened remotely; a second smaller trombone can be inserted to provide a further increase in path length by adding 4 more mirrors into the system. This allows the focus to be moved from the center of the undulator to some distance inside the bending magnet.

The light is finally split and imaged on two cameras. The main camera is a Proxicam HL4 S NIR with a red-enhanced S25 photocathode and an image intensifier. This produces an image of the beam cross-section averaged over the entire

ring. From the first results, the camera has the sensitivity to image a pilot bunch at injection energy in a single turn. A turn-by-turn and bunch-by-bunch imaging system is also currently being developed as an additional synchrotron light monitor. This device uses a gated image intensifier coupled to the Redlake HG-100k camera. The intensifier is an MCP with a minimum gate of 5ns. This is sufficiently fast to isolate single bunches. The intensifier can then be triggered to capture light from the same bunch each time it passes the camera, allowing an image to be built up over many turns.

4. ABORT GAP MONITOR

4.1 Purpose of the abort gap monitor

It is sometimes necessary to dump the particle beams circulating in the LHC, either because the beam has become unstable or simply to prepare for a fresh fill. The beam carries a large amount of energy and it cannot be allowed to hit the walls of the beam pipe or any other equipment. Instead, a special kicker magnet can be switched on which diverts the beam into a tangential beam line where it is absorbed in a huge graphite block. The kicker magnet takes 3µs to go from zero to nominal field. If a particle arrived during this time it would receive a partial kick and would hit the beam pipe or a downstream magnet. It is therefore important to limit the number of particles which can arrive in this period, and with this in mind the injection of bunches into the LHC is arranged with a 3µs long 'abort gap' which contains no bunches. Nonetheless some particles can escape from the bunches and drift into the gap, so the particle density in the gap must be carefully monitored to ensure it stays below a safe threshold.

4.2 Implementation

A beam splitter is installed after the first focusing mirror and diverts about 10% of the light towards the abort gap monitor (AGM). The AGM consists of a gated photo-multiplier tube which is synchronized with the abort gap. The abort gap is divided into 30 bins of 100ns each and the number of photons arriving in each bin is integrated over 100ms (~1100 passes of the abort gap). However, due to a limitation of the duty cycle of the PMT to 1%, each bin can only be scanned once in every 4 passes, thus 280 times in each integration period.

5. LONGITUDINAL DENSITY MONITOR

The third part of the BSRT will be the longitudinal density monitor (LDM), which is under development. The objective is to measure the variation of density with time, which is equivalent to a longitudinal profile of the beam. A scheme is presented here using single photon counting to fulfill this objective.

5.1 Bunch structure

In order to be successfully accelerated, particles in the LHC must arrive at the accelerating cavity when the RF field has the right phase. This means that particles must be injected not as a continuous beam but in discrete bunches. The RF operates at 400MHz so each of the 'RF buckets' which the particles may occupy is 2.5ns long. A bunch of particles is injected only into every 10[th] rf bucket, in other words the bunch separation is 25ns. There are additional gaps to allow for the rise time of injection magnets in the LHC and its feeder accelerators, as well as the abort gap discussed above, so that in total 2808 of the 35640 buckets will be filled.

5.2 Specifications

The LDM should be capable of producing a longitudinal profile with a time resolution of 50ps. Two functionalities are requested. Firstly there should be a fast integration mode which can measure the bunch parameters (bunch length, density distribution) with an integration time of 1ms. Secondly, there should be a high-sensitivity mode which can produce a full longitudinal profile showing all the bunches with a sensitivity of 5×10^5 protons per 50ps bin. This implies a dynamic range of more than 30,000 compared to the maximum density at the centre of the bunch. Of particular importance is to measure the tails of the bunches and the nominally empty spaces between bunches[8].

5.3 Concept

Given the small amount of light available and the fast time resolution required, it is proposed to use time-correlated single photon counting to construct the profile. An avalanche photo diode (APD) operated in the Geiger mode is suitable for detecting single photons with high time resolution. It is proposed to install a second 10% beam splitter on the AGM line, thus diverting approximately 1% of the total light in the BSRT into the LDM. Two options exist for the location of the LDM.

The first is to locate it in the LHC tunnel next to the AGM. The electrical signals would then be carried by a cable to an electronics rack some 30m away where the time-to-digital converter (TDC) card would be located. This preserves the time structure of the optical pulse as much as possible. However, the LHC tunnel is a radiation zone, raising the possibility of radiation damage to the detector and/or of noise caused by ionizing radiation traversing the APD. It is known that exposure to radiation can increase the dark count rate of APDs[9] but no specific information is available on the radiation hardness of the APDs chosen for this application.

The second option is to couple the light into an optical fiber and transport it to the surface where both the APD and the TDC would be located. This avoids radiation damage to the detector and has the added advantage of making access much easier. Due to the running schedule of the LHC, access into the tunnel is possible only during brief technical stops, with the next major shutdown not scheduled until late 2011 or 2012. However, the fiber introduces dispersion which deteriorates the time resolution of the system. Single-mode fiber would have to be used to keep dispersion at an acceptable level over the few hundred meters to the surface, but this introduces problems with the coupling of the light into the fiber. While the coupling efficiency does not need to be large, the coupling must be done by remote control and it would be desirable to keep the coupling efficiency constant for calibration purposes. A multi-mode fiber would make this easier but destroys the time resolution of the signal.

5.4 Limitations of the APDs

APDs have two significant limitations which affect their use in this project. When used for single-photon counting they must be operated in the Geiger mode, in which the APD is biased above its breakdown voltage. When a photon is incident on the APD it generates an electron-hole pair which is self-amplified by the avalanche into a measurable pulse. If left unchecked the avalanche would continue until the APD overheated and was destroyed, so the APD must include a quenching circuit. This lowers the bias voltage below breakdown as soon as a photon is detected and restores it some time later when the avalanche has dissipated. This time is the deadtime, during which the APD is blind to further photons.
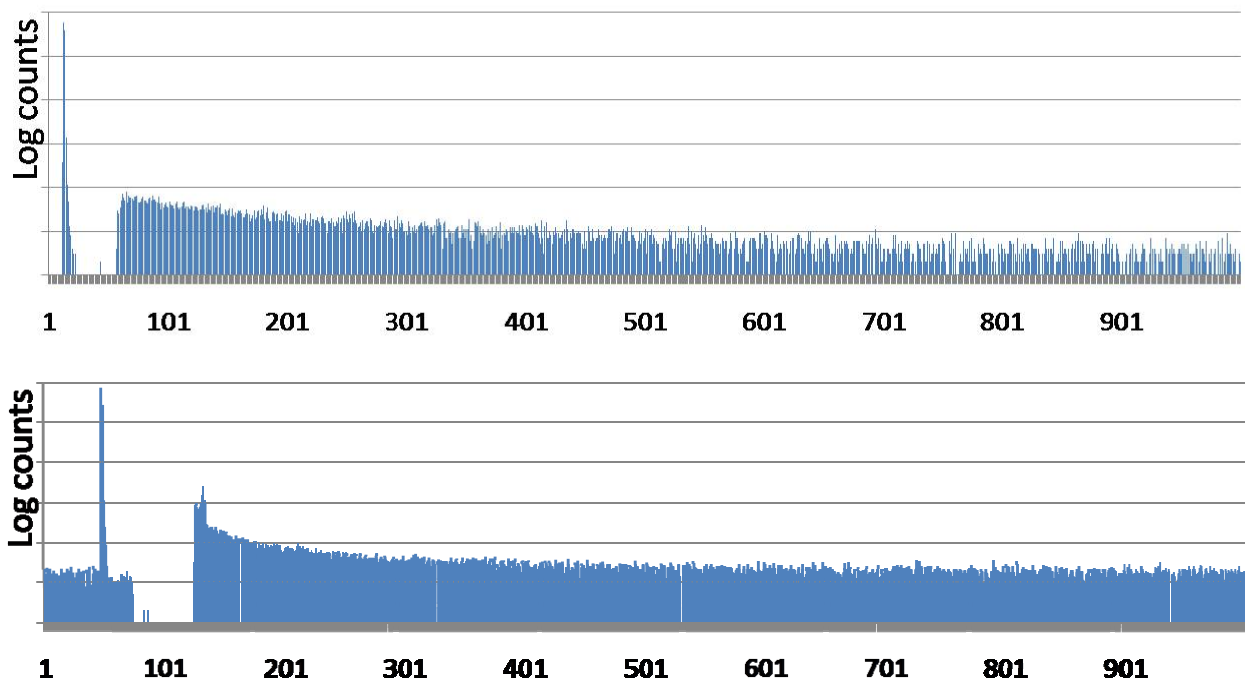

Figure 2. Response of the id100 (top) and MPD (bottom) detectors. The deadtime can clearly be seen after the main pulse, as can the increase in false counts at the end of the deadtime (afterpulsing). Each bin is 1ns, which is the smallest supported by the acquisition card.

In addition, the avalanche does not always properly dissipate. Charge carriers can be trapped in imperfections in the APD. At the end of the deadtime, when the voltage is restored, they generate a new avalanche which is indistinguishable from one caused by a new photon. This is known as afterpulsing and can be a source of noise in the measurement[10].

Two APDs have been chosen for the LDM. One is an id100 single-photon detection module from idQuantique. The second is a photon detection module from MPD. The response of both APDs was characterized using a pulsed laser (Figure 2). The time resolution was limited to 1ns by the acquisition card; both devices claim a time resolution better than 50ps. Attempts were made to stop outside light entering the system but nonetheless a background remains.

The response of both detectors is similar. The MPD module has a longer deadtime and also has a shoulder after the main pulse, which might have been caused by poor matching of the output to the acquisition card. Both detectors have a similar degree of afterpulsing. However, in the id100 these are spread over a long time, with the count rate not returning to the background level until approximately 15μs after the initial pulse. In contrast the majority of afterpulses in the MPD module occur immediately after the deadtime, with a faster tail-off to the background rate.

5.5 Distortion of the signal due to deadtime and its correction

Because the signal is time-varying, as well as limiting the maximum count rate the deadtime also causes a distortion of the signal [11]. The detector is more likely to be ready ('up') to receive a photon at the start of the bunch than at the end, since the detector will be in its deadtime ('down') if a photon has been received at any time earlier in the bunch. This causes a skewing of the signal towards the front of the bunch, which is more pronounced the higher the photon arrival rate. In traditional TCSPC applications the integration time is not important so the arrival rate can be made sufficiently low that the distortion is negligible. However in this case the integration time is limited and so we must accept the distortion and correct for it.

To this end a Monte Carlo simulation was made[11] using C++ to quantify the extent of the distortion and to test a correction algorithm. The number of photons seen by the detector is histogrammed over many turns of the LHC. To calculate the number of photons C_i counted by an ideal detector (i.e. one with no deadtime) in a particular bin i, the number of photons counted by our (non-ideal) detector is summed over all the bins i-d to i-1, where d is the deadtime of the detector divided by the bin width. Dividing this by the number of turns that integrated over gives the probability that the detector was down during bin i. Then

$$C_i = \frac{x_i}{P(up)_i} = x_i \frac{N}{N - \sum_{j=i-d}^{i-1} x_j}$$

where x_i is the number of counts given by the non-ideal detector in bin i over N turns.

A second correction can be applied to account for afterpulsing. The probability of an afterpulse occurring after a given number of bins a can be measured in the lab, and the signal can then be corrected by subtracting from bin i a number of counts based on the number of counts in bin i-a. However, since the afterpulses are spread over a long period, the adjustment per bin is small and a substantial amount of noise is introduced.

The third correction accounts for pile-up. The corrected signal C is not directly proportional to the particle density since even an 'ideal' photon counter operates in a binary mode. Each bin either contains a photon or does not. This leads to an underestimation of the signal since if more than one photon arrives during one bin, only one is counted. The effect only becomes significant at high photon arrival rates since the chance of two photons arriving in the same bin is otherwise negligible.

In order to correct for this, the probability of two photons arriving at the same time must be estimated. The emission of synchrotron radiation is a stochastic effect involving a very large number of particles, each with a very small chance of emitting a photon within the acceptance of the detector. The photons can thus be considered to have a Poissonian distribution. If the number of counts expected in a particular bin (given the proton density) is λ, then the probability of having a given number k of photons emitted is

$$p(k) = \frac{\lambda^k e^{-\lambda}}{k!}$$

If *C/N* is the probability that the ideal detector would see at least one photon then

$$p(0) = 1 - C/N = e^{-\lambda}$$

The expected number of photons is then given by

$$\lambda = -\ln\left(1 - C/N\right)$$

which is now directly proportional to the proton density in that bin. As expected this is approximately equal to *C/N* for (*C/N*)<<1. Figure 3 shows that the correction is effective even when the deadtime is much longer than the bunch separation.
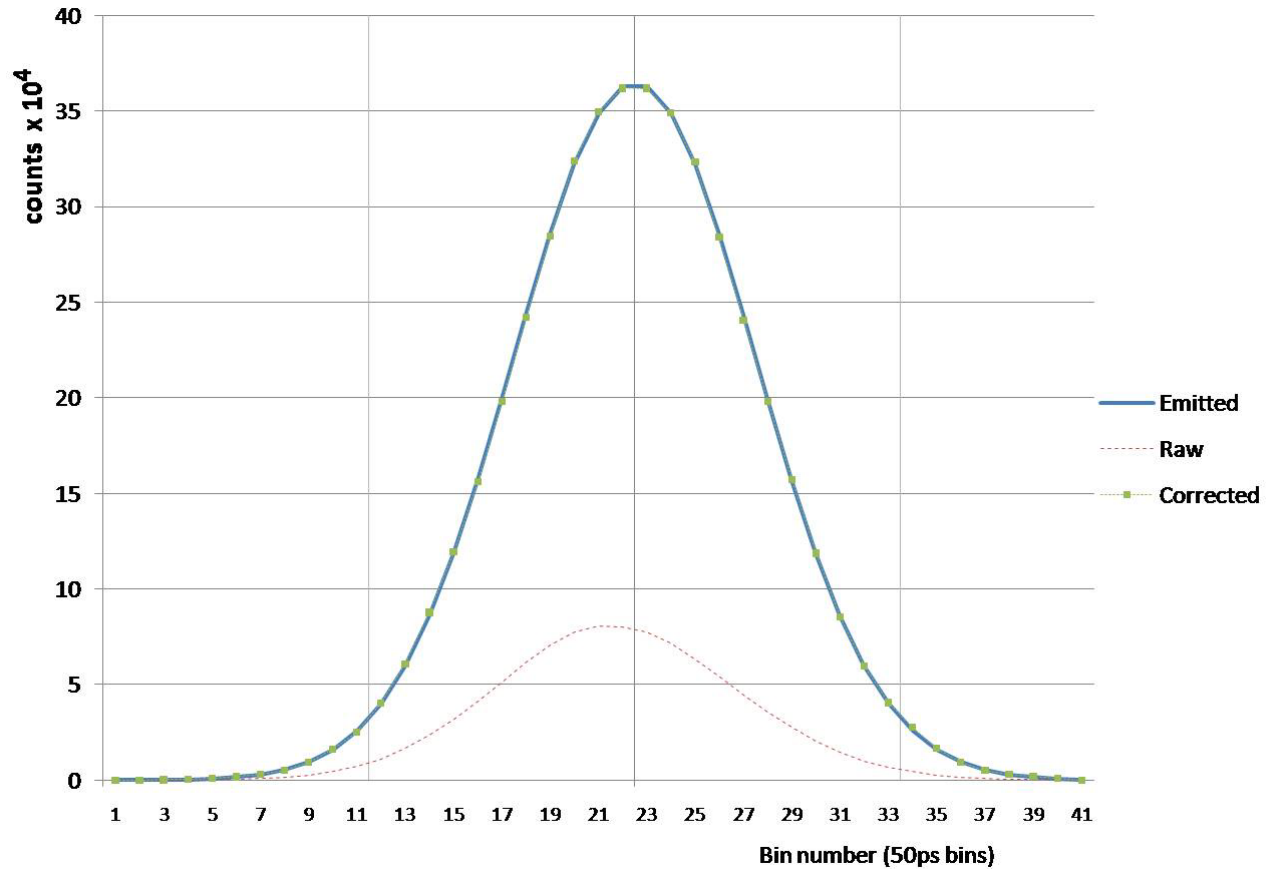

Figure 3. Simulated performance of the APD and correction algorithm. Each bin is 50ps. The photon arrival rate is 0.5 per bunch and the deadtime is 195ns. The centre of the bunch is marked to show the skew of the raw APD signal towards the front of the bunch. After correction the signal exactly matches the true number of photons emitted.

5.6 Limit to the dynamic range of the detector

The correction algorithm described above is effective up to a photon arrival rate of a few photons per bin. When the arrival rate becomes very high, however, the probability of the detector still being up in the later part of the bunch is almost zero. In this case no information is available about the end of the bunch and correction is impossible. The detector can then be said to be saturated. This sets a limit to the amount of light that can be allowed to hit the detector.

In order to achieve the desired sensitivity in the tails and ghost bunches, however, it would be desirable for the detector to capture as much light as possible. The specifications[8] state that the system should be sensitive to a ghost bunch of 5×10^5 protons. In order for just one photon from this ghost bunch to reach the detector during a 10s integration, the arrival rate must be 3 photons per full bunch. There is then a large probability that the detector would be down anyway

when the photon from the ghost bunch arrives. In practice, of course, more than one photon would be needed to stand out above the noise level, and the arrival rate from the full bunch would then need to be above saturation of the detector.

5.7 Solution using two APDs

To overcome this limitation it is proposed to use two APDs. The synchrotron light would again be split. One branch would go to a free-running (i.e. always on) APD via a controllable attenuator. This detector would be capable of measuring the core of the bunch but would be blind to the much fainter signals of the tails and ghosts. The second branch would go to a fast gated APD. This APD would be switched off during the arrival of photons from the bunch and would thus be able to receive much more light from the tails without becoming saturated (figure 4).
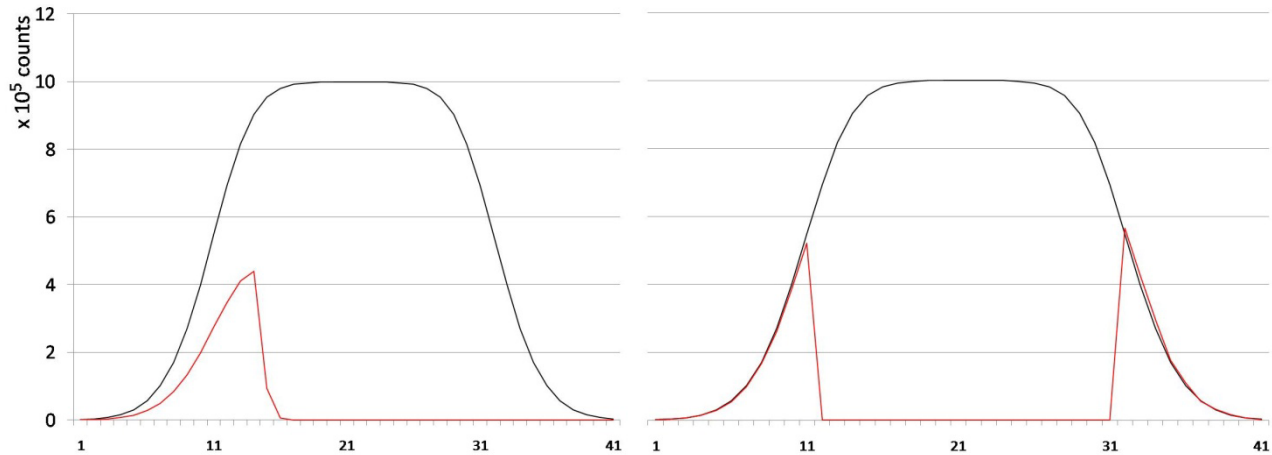

Figure 4. APD signal after correction compared to the true bunch shape, from simulation. Arrival rate is 90 photons per bunch, each bin is 50ps and integration time is 10s. Left, the detector is saturated and correction is unable to restore the true signal. Right, the detector is gated off in the central part of the bunch; the correction algorithm is then able to accurately portray the tails.

5.8 Gating of the detector

The performance of the MPD detector in gated mode was tested. At low repetition frequency, the gate operates well, with pulses arriving when the detector is gated OFF being completely hidden (figure 5). The gate switching speed is just a few nanoseconds. The application of the gate pulse causes some spurious counts at the beginning and end of the gate, but this is not so frequent as to limit the use of the detector. There appears to be no limit to the duty cycle of the detector.
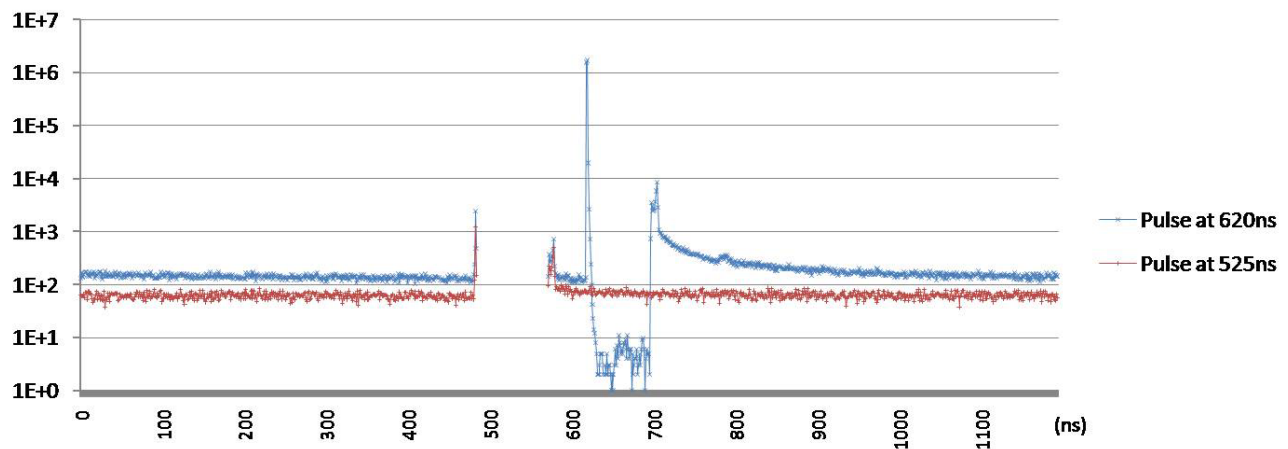

Figure 5. Performance of the MPD module in gated mode. The module is gated off for 80ns. If the pulse occurs when the detector is on, the response is normal; if it occurs when the detector is off it is completely hidden.

However, when tested with a faster gate repetition rate such as would be needed for this application, the detector did not perform as expected. If the gate is switched on immediately after a light pulse is received, a count will be generated as the gate is switched on. This could be because the incident light causes the generation of electron-hole pairs in the APD even if the voltage is below breakdown; if the operating voltage is restored before these have recombined, an avalanche will occur. When the gate-on signal followed the laser pulse by some 10ns, a false avalanche occurred in nearly 100% of cases, and the detector was therefore unable to detect any real photons in the gate, which was shorter than the deadtime of the detector.

5.9 Optical gating

Due to the problems encountered with the fast gating of the detector, the possibility of gating the light itself is being considered. One method would be to couple all the LDM's light into an optical fiber and then use a telecommunications-style optical switch to direct it down one of two fibers to the two detectors. However, very fast switches are hard to find outside the narrow telecommunications bands. Also, this method would exclude cross-calibration of the two detectors since there would be no time when both detectors receive light; indeed there would be a small blind spot in the profile due to the switching time. Nonetheless this could be a simple and cost-effective solution.

The second method would be to have a beam splitter, one beam going to the 'free' APD via an attenuator and the other going to the 'gated' APD via an optical gate. The gate modulation could be provided by a Pockels cell, but this would be difficult to implement due to the requirement for high voltage pulses. There are additional problems with wavelength dependency and the limited extinction ratio of the cell[12].

A better option might be to use an electro-optic (EO) deflector[13]. EO deflectors are essentially achromatic and the deflection is 100% effective. A suitable optical gate can be made by placing the deflector and the APD on opposite sides of a converging lens, each at a distance of 2*f*. Collimated synchrotron light entering the deflector then reaches the APD regardless of deflection. A small mask is placed on one side of the lens to block the central line. This arrangement avoids the need for a pulse generator; instead it requires only a sinusoidal voltage at half the bunch repetition frequency, synchronized so that the light hits the masked line whenever a bunch is passing.

6.   CONCLUSION

Synchrotron light is a useful tool for diagnostics in the LHC. It can be used to determine both spatial (transverse) and temporal (longitudinal) profiles of the particle beam. Two detectors have already been installed and have produced useful results during the short 2009 run of the LHC. A third detector for full longitudinal profiles is under development.

The scheme presented here would allow longitudinal profiles to be measured with a time resolution of 50ps and a very high dynamic range. An integration time of 1ms is sufficient to determine the bunch parameters, if it is accepted to average over all bunches. Integration over 10s can produce a full profile showing all individual bunches, the bunch tails and any ghost bunches. The main difficulty is the need for a very fast gating scheme to prevent the detector from being saturated.

ACKNOWLEDGEMENTS

REFERENCES

[1]. **Wille, K.,** [The physics of particle accelerators: an introduction]*, Oxford University Press, 2000.

[2]. **Ponce, L., Jung, R and Meot, F.** "*LHC proton beam diagnostics using synchrotron radiation."* CERN-2004-007. (2004)

[3]. **Fisher, A.S., Goldblatt, A. and Lefevre, T.** *"The LHC synchrotron light monitors",* Basel, Switzerland : DIPAC, p.164 (2009).

[4]. **Hofmann, A.** [The Physics of Synchrotron Radiation], Cambridge University Press, 2004.

[5]. **Chubar, O. and Elleaume, P.** *"Accurate and efficient computation of synchrotron radiation in the near field region".* EPAC. pp. 1177-9. (1998).

[6]. **Maccaferri, R., et al.** *"The 5-T superconducting undulator for the LHC synchrotron radiation profile monitor",* Lucerne, Switzerland : EPAC, (2004).

[7] **Fisher, A.S.** *"The LHC Synchrotron-Light Telescope (BSRT) and Abort-Gap Monitor(BSRA), Using Dipole and Undulator Radiation, for both Protons and Ions",* CERN project note (2009)

[8]. **Fischer, C.** *"High sensitivity measurement of the longitudinal distribution of the lhc beams"*: CERN, LHC-B-ES-0005.00 (2003).

[9]. **Sanchez Majos, S. et al.** *"Noise and radiation damage in silicon photomultipliers exposed to electromagnetic and hadronic radiation",* Nuclear Instruments and Methods in Physics Research A, Vol. 602 (2009).

[10]. **Renker, D. and Lorenz, E.** *"Advances in solid state photon detectors",* JInst, Vol. 4 (2009)

[11]. **Bravin, E.** *"Dead time effect on single photon counting for the longitudinal density monitor of LHC",* CERN, AB-Note-2006-002 BI (2006)

[12]. **Becker, K., Bruce Johnson, J and Edwards, G.** *"Broadband Pockels cell and driver for a mark III-type free electron laser",* Rev. Sci. Instrum., Vol. 65 (1994)

[13]. **Steve Payne, Leysop Ltd.,.** private communication.