Single-Event Upsets in Photodiodes for Multi-Gb/s Data Transmission

Jan Troska, Alberto Jimenez Pacheco, Luis Amaral, Stefanos Dris, Daniel Ricci, Christophe Sigaud, François Vasey and Paschalis Vichoudis

CERN, 1211 Geneva 23, Switzerland

jan.troska@cern.ch

Abstract

A Single-Event Upset study has been carried out on PIN photodiodes from a range of manufacturers. A total of 22 devices of eleven types from six vendors were exposed to a beam of 63 MeV protons. The angle of incidence of the proton beam was varied between normal and grazing incidence for three data-rates (1.5, 2.0 and 2.5 Gb/s).

We report on the cross-sections measured as well as on the detailed statistics of the interactions that we measured using novel functionalities in a custom-designed Bit Error Rate Tester. We have observed upsets lasting for multiple bit periods and have measured, over a large range of input optical power, a small fraction of errors in which an upset causes a transmitted zero to be detected as a one at the receiver.

I. INTRODUCTION

Single Event effects have been widely documented to occur in the photodiodes typically used in modern high-speed serial communications [1, and references therein]. At CERN, we are currently designing the next generation of optical data transmission link for reading-out and controlling particle physics detectors to be operated at CERN's upgraded Large Hadron Collider (Super LHC). Such links will operate at multi-gigabit per second data-rates. The innermost regions of the detectors will encounter a radiation environment dominated by high-energy pions with a most-probable energy around 300 MeV, at fluxes of 10^6 – 10^8 particles/cm^2/s, depending upon position with respect to the beam (see Figure 1).

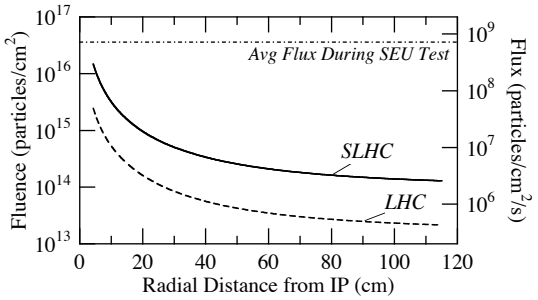

Figure 1: Extrapolation of expected fluxes for Inner Detectors from CMS data at 500 fb-1 (LHC) to 3000fb-1 (SLHC).

The control information flowing into the detectors from shielded control rooms is critical for maintaining the synchronization of the data-taking system, both internally and with respect to the bunched beams circulating in the SLHC. It is therefore of critical importance that this control information be transmitted error-free and, with the knowledge that Single Event Upsets (SEUs) will occur within a photodiode placed in such an environment, the use of Forward Error Correction (FEC) coding will be mandatory. Validation of any choice of FEC code depends upon a detailed knowledge of the statistics of the errors that are expected to be encountered and the test reported in this paper aims to gather that knowledge.

In order to gather as much information as possible, we performed a small survey of the radiation-response of several different devices. InGaAs PIN photodiodes operating at 1310nm, GaAs PIN photodiodes operating at 850nm were combined in this test with Receiver Optical Sub-Assemblies (ROSA) where the Transimpedance Amplifier (TIA) is mounted in the same TO-can as the photodiode. Again, both 1310nm InGaAs and 850nm GaAs ROSAs were included.

II. SEU TEST METHOD

A. Irradiation test setup

The irradiation was carried out at the PIF-NEB proton irradiation facility at the Paul Scherrer Institut (PSI), Villigen, Switzerland [2] using a 63 MeV proton beam. Every second the flux was measured by ionization chambers and its value stored in a file by the control software of the irradiation facility for later analysis.

The DUTs were mounted on a rotating axle that allowed the angle of incidence of the proton beam on the optoelectronic receivers to be varied between normal (0°) and grazing incidence (90°) by remote control from outside the irradiation bunker. Measurements were taken at 0°, 10°, 80° and 90°.

Data were generated inside the FPGA-based Bit Error Rate Tester (BERT) described below, that was sited below the beamline inside the irradiation bunker, but shielded with a combination of Aluminium and Polyethylene. Serial data was passed on to a laser driver and laser diode for conversion to an optical signal (see Figure 2). This signal passed through 25 m of optical fibre to the control room, where an optical attenuator and power meter were used to control and measure the amplitude of the light returning, via an optical splitter and another 25 m of optical fibre cable, to the DUTs in the irradiation bunker.

The signals from the photodiodes require amplification in order to be sent over coaxial cables to the shielded Bit Error Rate Tester (BERT). Combined TIA/Limiting Amplifiers from Maxim Semiconductor (MAX3866) were mounted in very close proximity to the photodiodes on the test board. The electrical signals from the ROSAs were further amplified using a Limiting Amplifier (LA), also from Maxim Semiconductor (MAX3748B). The amplifiers were shielded from the pro-

ton beam by 6.5 mm of brass, sufficient to stop 60MeV protons. In addition, each set of eight DUTs was accompanied by two reference photodiodes and TIA/LAs that were also shielded. These references, one SM and one MM, were provided to measure any possible noise induced by external sources within the irradiation bunker.

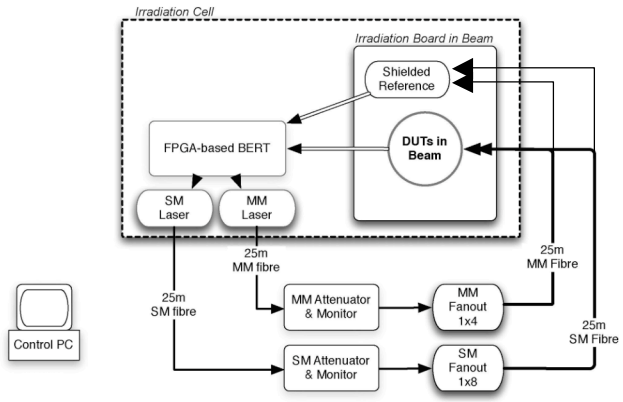

Figure 2: Setup for the proton irradiation test.

B. FPGA-Based Bit Error Rate Tester

A custom BERT was implemented (see Figure 3) based upon the Transceiver Signal Integrity Development board available from Altera for the Stratix II GX family of FPGAs, which include embedded high-speed transceivers capable of operating at data-rates up to 6.375 Gb/s [3].

The primary testing goal of measuring error statistics was achieved through the use of an error log memory that could hold up to 8K 20bit words. For every received word, the XOR of the transmitted and received data is evaluated and if one or more bit errors are encountered this error pattern is stored in the memory along with a timestamp for later analysis. In addition, basic Bit Error and Word Error counters were implemented.

A second memory of 8K 20bit words contained the pattern cyclically sent by the transmitter. In our case this memory was filled with random data, 8B/10B encoded offline for line-balancing and with commas inserted every 64 words to aid synchronization in the receiver. The received data were compared "as-is", i.e. not decoded, to measure the raw BER due to SEU errors only and not errors due to decoding problems.

Firmware was developed that would allow operation at the three data rates used in the test (1.5, 2.0 and 2.5 Gb/s) by simply supplying a different frequency base clock to the FPGA.

Ten optoelectronics receivers were simultaneously tested on each irradiation board and all of the boards contained both SM and MM devices. Each FPGA supports a maximum of four full speed electrical transceivers, and therefore three FPGA platforms were required for our test. On each FPGA a maximum of one transmitter channel was active, because only one SM and one MM laser source were employed in our setup. Each FPGA board can work independently or can be connected with a second one in master-slave mode.

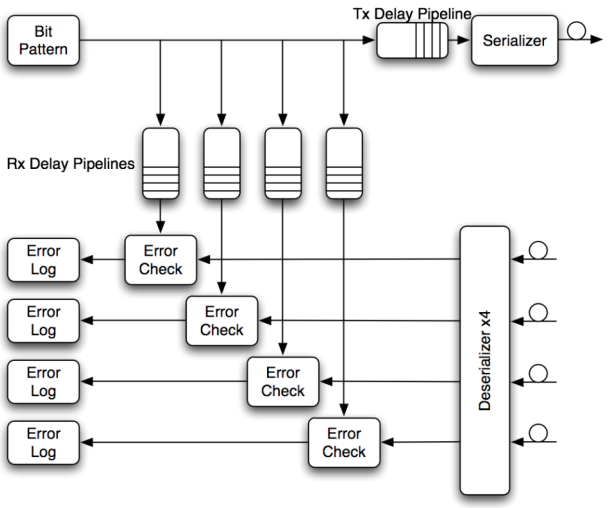

Figure 3: Simplified diagram of the BERT implemented in firmware.

C. Devices tested

Devices were selected based upon current availability from six vendors. One device type (Man. 4, Mod. 1) was included to provide a comparison with previous work carried out at CERN [4] at lower data-rates. Table 1 shows the devices tested and some relevant parameters.

The devices were arranged across three test boards that were exposed in succession to the proton beam.

Family	Wave-length (nm)	Device Type (# tested)	Active diameter (μm)	Responsivity measured
PIN	1310	Man. 1, Mod. 1 (2)	30	0.7 A/W
PIN	1310	Man. 1, Mod. 2 (2)	60	0.8 A/W
PIN	1310	Man. 1, Mod. 3 (1)	60	0.75 A/W
PIN	1310	Man. 1, Mod. 4 (2)	80	0.8 A/W
PIN	1310	Man. 2, Mod. 1 (2)	60	0.8 A/W
PIN	1310	Man. 3, Mod. 1 (2)	—	0.75 A/W
PIN	1310	Man. 4, Mod. 1 (3)	80	0.8 A/W
PIN	850	Man. 5, Mod. 1 (3)	100	0.6 A/W
PIN	850	Man. 6, Mod. 1 (2)	90	0.6 A/W
ROSA	1310	Man. 6, Mod. 2 (2)	65	3.0 $mV_{pp}/\mu W$
ROSA	850	Man. 6, Mod. 3 (2)	90	2.2 $mV_{pp}/\mu W$

(Man.: Manufacturer; Mod.: Model)

Table 1: Devices Tested. 1310nm devices are Single-mode, 850nm devices are Multi-mode.

III. Results: Overall Trends

For every combination of the selected data rates and angles, attenuation scans monitoring the BER were systematically performed both with the beam on and off, to be able to distinguish in every case errors caused by protons from those due to electrical and environmental noise.

As an example, we show the effect that turning on the beam has on the BER performance of two devices in Figure 4. In this figure and throughout the paper, the Optical Modulation Amplitude (OMA) is measured at the input of the optoelectronic receivers. When the beam is on, the range of OMA can be divided in two regions, one where performance is dominated by noise and one where it is dominated by radiation in-

duced errors. The almost perfect matching in the noise dominated region between the plots with beam on and off shows the good reproducibility of the results.

In the following, most results will be presented in terms of the Bit Error Cross Section, defined as the quotient between the number of bit errors occurring during the testing time and the accumulated fluence. This cross section is only defined and presented in the SEU dominated region.

Plots similar to that of Figure 4, comparing the BER with beam off and on, were analysed for the reference photodiodes at all incidence angles. These have shown that the shielding was not working perfectly at all angles; specifically, some upsets could be observed in the reference photodiodes near grazing incidence.

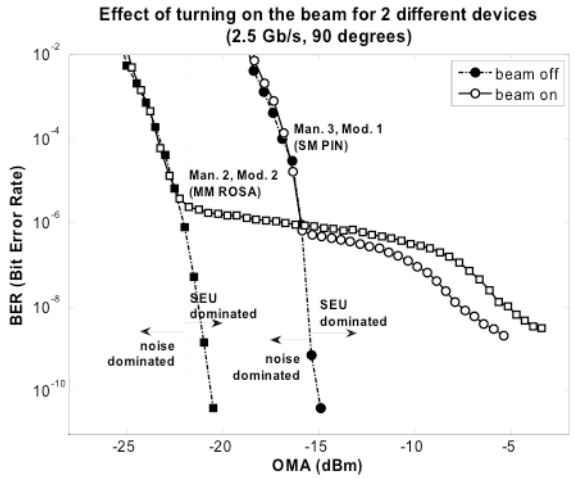

Figure 4: Illustration of the effect of the beam on two different devices, a SM photodiode (marker '•') and a MM ROSA (marker '■').

A. Device Families

We compare the Bit Error Cross Section of every model used in our test under a common set of conditions, 2.5 Gb/s and grazing incidence (90°), in Figure 5. For a given value of OMA the difference in cross section among devices spans more than two orders of magnitude, but the plots for all models exhibit the same general shape. The variety of active diameters, packaging materials and manufacturing processes among devices from different manufacturers makes general trends difficult to observe.

Nevertheless, photodiode Model 1 from Manufacturer 1, which happens to have the smallest active diameter among all the devices tested (30 microns), stands out as remarkably better than the rest. Since the path of the protons through the active volume is minimised, so is the BER, especially the contribution due to direct ionization.

The ROSAs do not rank among the devices with worse performance (especially for the single mode case), even though in these devices we are observing the combined effects of SEUs in the photodiode and in the unshielded TIA (The LA is shielded).

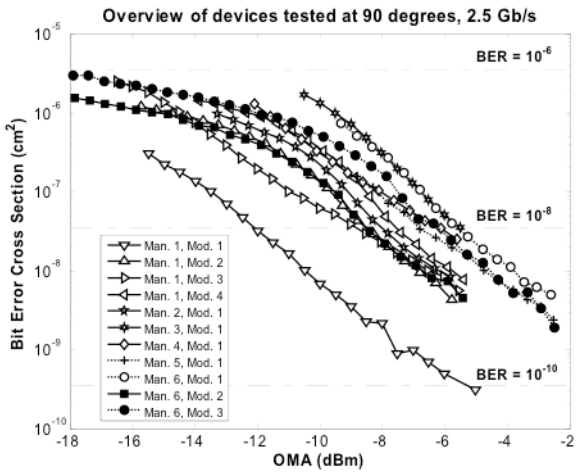

Figure 5: Overview of results for all device models at 2.5 Gb/s and 90°. Continuous lines correspond to SM devices, dotted lines to MM; open markers correspond to photodiodes, solid markers correspond to ROSAs. The three horizontal references for BER = 10^{-6}, 10^{-8}, 10^{-10} assume a perfectly constant flux of 8×10^{8} p/cm^2/s.

B. Angular dependence

We confirm the observation made in previous tests by other authors [1] and our own team [4]: that the maximum of the cross section as a function of the incidence angle occurs near 90° (grazing incidence) and it is minimum for 0° (normal incidence), as shown in Figure 6. This is expected, as 90° corresponds to the longest ionizing path of the protons through the active volume of the photodiode.

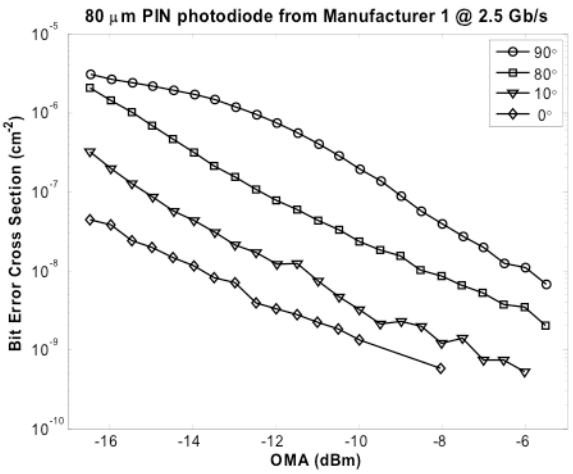

Figure 6: Bit error cross section as a function of received optical power for an 80 μm photodiode from Manufacturer 1 tested at 2.5 Gb/s and different angles.

We were expecting a very selective peak of the cross section around 90°, as shown for instance in [1] or [4]. However that is not exactly what we observe in Figure 6, where the plots corresponding to 0° and 10° should be closer to one another, and closer to that of 80°, to agree with this expectation. This deviation from the expected behaviour could be explained by partial shadowing of the DUTs by the optical fibres and connectors attached to them, which could degrade the energy of the beam for angles near normal incidence.

C. Data Rate dependence

For a given device and angle, when plotting the BER as a function of the received optical power for the three data rates measured, we observe that - within the limits of the experimental error - they superimpose almost perfectly, indicating that there is no dependence of the SEU induced BER on data rate (see Figure 7).

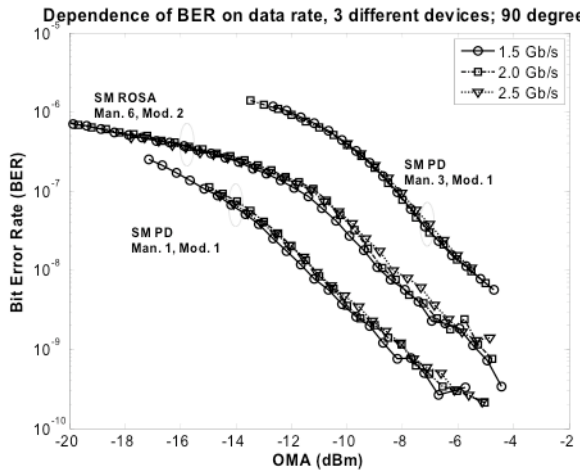

Figure 7: Bit error cross section as a function of received optical power for three different devices at 2.5, 2.0 and 1.5 Gb/s; grazing incidence.

With this in mind, using the relationship below:

$$\mathrm{BER} = \frac{\sigma \times \overline{\mathrm{flux}}}{\mathrm{data\_rate}} \qquad (1)$$

between BER, Bit Error Cross Section (σ), data rate and average flux, and taking into account the good stability of the proton beam used, we conclude that there is a linear dependence of the cross section with data rate, at least in the range from 1.5 to 2.5 Gb/s. This contrasts with results presented by other researchers [5], who have shown a linear relationship at low data rates, but a greater than linear dependence at the higher end of their measurement range (1.2 Gb/s).

IV. Results: Error Log Analysis

The error logging mechanism implemented in the custom FPGA BER Tester allow us to obtain a very detailed analysis of the error statistics: burst length histograms, Error Free Interval (EFI) histograms, pattern dependence, as well as correlation of errors with the transmitted pattern.

In order to characterise an error burst, not only is its length important, but also the value of the Error Free Threshold (EFT) used in the analysis [6][7], defined as the maximum number of successive correct bits allowed inside a burst. The value of the EFT must be carefully selected, examining simultaneously its effect both on the EFI histogram and on the burst histogram. We have selected an EFT of 10 bits, implying that any two bit-errors separated by 10 or less correct bits are considered part of the same burst and have obtained satisfactory results.

A. Error Classification

We have used the results from the error log analysis to classify the errors according to the following criteria:

- **Error length**. We distinguish between single (isolated) errors and burst errors. For reason that will become clear later, we further subdivide burst errors in short (length between 2 and 20 bits) and long bursts (length over 20 bits).
- **Fraction of 0-to-1 errors**. We can correlate the logged error patterns with the transmitted sequence to find out what fraction of the bit errors are due to sent 0's being mistaken at the receiver by 1's, and vice-versa.
- **Burst occupancy** (sometimes also termed burst density): is computed as the number of bits that were actually flipped in a burst divided by the length of that burst.

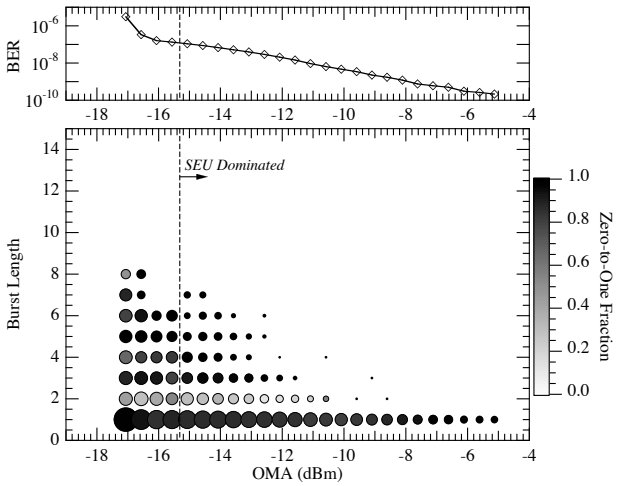

Figure 8: Two-dimensional burst histogram for a 30 µm SM photodiode (Man.1 Mod.1) tested at 2.5 Gb/s and grazing incidence. EFT = 10 bits.

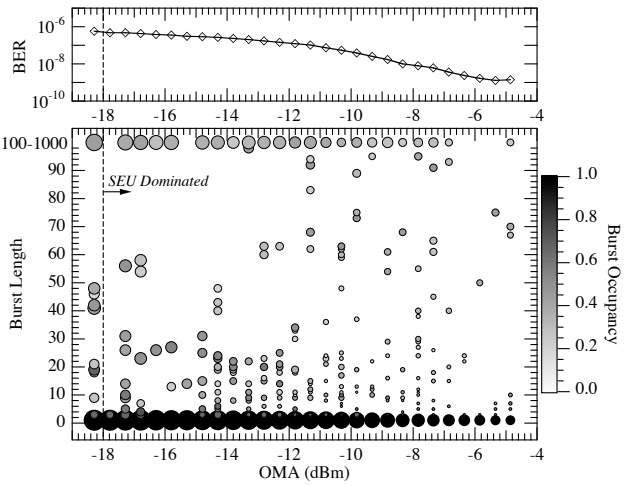

Figure 9: Two-dimensional burst histogram for a SM ROSA (Man.6 Dev.2) tested at 2.5 Gb/s and grazing incidence. EFT = 10 bits.

For a given set of test conditions (data rate, angle), a convenient way to simultaneously visualize the burst histograms of a device for the complete range of attenuation values tested is what we call a two-dimensional colour-coded burst histo-

gram, for which we show two examples in Figure 8 (for a photodiode) and Figure 9 (for a ROSA). A circle is represented at point (x, y) if one or more bursts of length y bits are present in the 1D-burst histogram for a received power level of x (dBm). The size of the circle is logarithmically proportional to the BER contribution due to all bursts of length y at this power level. The colour of the circle gives information about the average 0-to-1 fraction (Figure 8) or burst occupancy (Figure 9), following the colour scale shown to the right of the figure. The upper part of each plot shows the dependence of the total BER on OMA for comparison to data shown previously.

Making a global classification of the errors required the careful examination of this kind of 2-D histogram for all devices and test conditions, but the two examples presented here are representative of the general behaviour of photodiodes and ROSAs, respectively.

We can classify the errors induced by SEU in three groups:

1. **Single errors**: this is by far the most frequent type of error. Independently of the device, data rate, angle or power level, the bin of length 1 dominates all burst length histograms. Almost all single errors are due to 0-to-1 bit flips. However, even for very low attenuation values, for which the probability of a noise-induced error is virtually zero, 1-to-0 bit flips still occur at the level of a few per cent (photodiodes) or a few per mille (ROSAs) as shown in Figures 10 & 11.

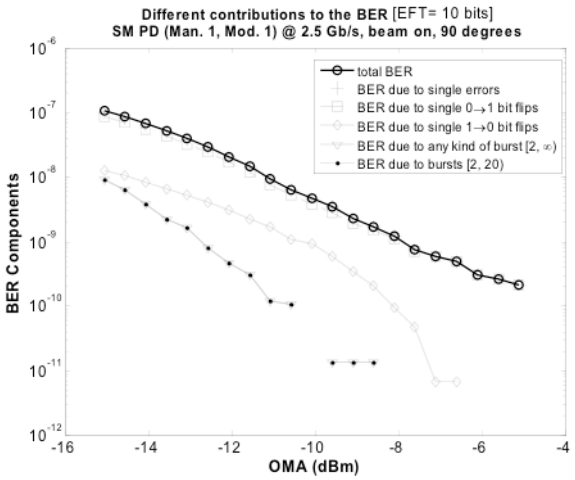

Figure 10: Different contributions to the BER for a 30 μm SM photodiode tested at 2.5 Gb/s and grazing incidence.

It is generally reported in the literature that SEUs in photodiodes can only be produced as 0-to-1 transitions. However, in our test we are observing bit errors at a global system level, and these 1-to-0 transitions appear in fact as the convolved response of the photodiode and the TIA/LA to a particle strike in the photodiode.

Figure 11 shows that the BER due to single errors of the type 1-to-0 is several orders of magnitude below that of single 0-to-1 for ROSAs, whereas in photodiodes it is only about one order of magnitude below (Figure 10). This different behaviour might be explained by the use of different amplifiers in the two types of devices (The ROSAs include a TIA 7770 from Vitesse Semiconductors integrated with the photodiode in the TO-can).

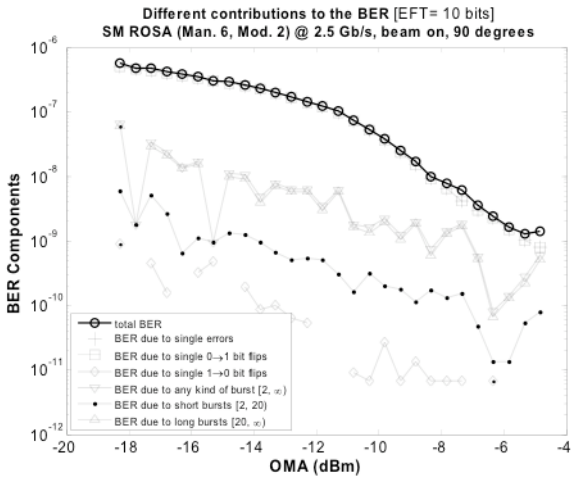

Figure 11: Different contributions to the BER for a SM ROSA tested at 2.5 Gb/s and grazing incidence.

2. **Short bursts (2-20 bits)**: For this type of error the conclusions differ slightly between ROSAs and PINs:
   - For photodiodes there is a strong correlation between the optical power at the receiver input and the occurrence of this type of bursts: the lower the power, the higher the number of bursts, and also the more important their contribution to the total BER. With respect to the 0-to-1 fraction, similarly to what happened for single errors, it is very close to 1. Many of the photodiodes show an anomaly by which all short bursts have a high 0-to-1 fraction except bursts of length 2. This is for example the case for the SM device in Figure 8, for which most double errors are in fact pairs of 1's mistaken in the receiver as pairs of 0's.
   - For ROSAs, there is also a high correlation of the short bursts with the power level, but it is more irregular and, contrary to photodiodes, a few short bursts are still present for very high values of the received power. With respect to the 0-to-1 fraction, in the ROSAs it is higher than for photodiodes, almost exactly equal to 1. There is no anomaly affecting the 0-to-1 fraction for double errors.

   It is especially interesting that for both type of devices the occupancy of these bursts is close to 100% (for double errors it is exactly 100%). This is the first direct measurement of multiple-bit bursts in photodiodes reported in the literature.

   We have also observed, both for photodiodes and ROSAs, that the contribution of short bursts to the total BER is strongly correlated with angle: short bursts occur more often at 90° than at 80°, and more often at 80° than at 10°; they almost disappear when the angle approaches 0°.
3. **Long bursts (length > 20 bits)**, which are almost ex-

clusively present in the ROSAs. The main characteristic of this type of error is that the burst occupancy is low, around 30-40%. The 0-to-1 fraction is very close to 1, as was also the case for single errors and short bursts in the ROSAs.

Long bursts are to some degree correlated with the received power level, as shown by the plot of their contribution to the total BER in Figure 11. However, a quick look at Figure 9 also reveals that long bursts can basically occur for any value of OMA.

We have also observed some dependence on angle: many fewer long bursts occur at 0° than at other angles.

In Figure 9 we see that the distribution of burst lengths is more or less continuous, without a gap between short and long bursts. Short bursts actually occur in ROSAs more frequently than long ones, but even so the contribution of the long ones to the total BER is much more important (Figure 11). This is because, despite their low occupancy, long bursts can last up to a few hundred bits. In contrast, bursts longer than 10 bits are virtually absent in photodiodes

As for the reference devices, we almost exclusively observed single errors, and a few short bursts at very low power levels, induced by noise. The very few short bursts occurring at higher power levels can be explained by the fact that the shielding was not working perfectly for angles near grazing incidence, as mentioned in Section III.

B. Hypothesis for the origin of bursts

We hypothesise that the long bursts occurring in the ROSAs are due to upsets taking place in the TIA, rather than in the photodiode. The fact that long bursts are almost exclusively present in the ROSAs, where the TIA cannot be shielded, supports this hypothesis; on the other hand, the very few long bursts that appear on photodiodes could still share the same origin because the shielding was not completely effective.

Another fact that backs up this theory is that the median length of long bursts, when expressed in absolute time units (ns, rather than bit periods), turns out to be fairly independent of the data rate. It takes values around 50-60 ns. These events are very long compared to the speed of the TIA, so probably the errors are not due to hits in the signal path but to hits in other nodes of the circuit with much longer time constants.

In contrast, we favour the hypothesis that short bursts are indeed related to upsets in the photodiodes, since we observed a very marked dependence on the received power level and on the incidence angle, the occupancy of these bursts is quite high and the majority of the bit flips correspond to 0's turning into 1's.

V. Conclusions

Results of an ambitious SEU test with protons of a large selection of PIN photodiodes and ROSAs operating at high data rates have been presented. Tests at various incidence angles have confirmed that the highest error cross sections are obtained for angles near grazing incidence. The SEU induced BER turned out to be independent of the data rate in the measurement range, from 1.5 to 2.5 Gbps.

The use of a custom BER tester allowed us to obtain detailed statistics of the error events. For instance, isolated errors in which a transmitted 1 is detected as a 0 at the receiver have been observed at power levels where they cannot have been induced by electrical noise.

We have also shown that multiple bit errors can occur in optoelectronic receivers. Short error bursts spanning up to a dozen bits, were observed in the photodiodes and longer bursts, up to a few hundred bits in length, have been measured in the ROSAs. To the best of our knowledge, this kind of behaviour, where an SEU can upset several successive bits, has not been previously reported in SEU tests performed with photodiodes at other data rates. Short bursts could be originated by upsets in the photodiodes, but long bursts in the ROSAs are most probably related to proton hits in the unshielded TIA. In either case, burst errors will have to be mitigated using FEC coding in future optical links to be used inside Super LHC detector systems. The detailed statistics collected during this test will prove essential in the design and validation of an appropriate FEC scheme.

VI. Acknowledgments

The authors would like to thank Dr. Wojtek Hajdas for his help during irradiation at PSI. We would also like to acknowledge the fruitful discussions held with Dr. Federico Faccio, Dr. Paulo Moreira, Dr. Jorgen Christiansen, Dr. Philippe Farthouat, Dr. Alessandro Marchioro and Mr. Csaba Soos at CERN.

VII. References

[1] P. W. Marshall, P. T. Wiley, R. N. Prusia, G. D. Rash, H. Kim, and K. A. LaBel, "Proton-induced BitError Studies in a 10Gb/s Fiber Optic Link," *IEEE Trans. Nucl. Sci.*, Vol. 51, No. 5, pp. 2736–39, 2004.

[2] W. Hajdas, A. Zehnder, L. Adams, and B. Nickson, "The proton irradiation facility at the Paul Scherrer institute," *Nucl. Instrum. Methods Phys. Res. B*, vol. 113, p. 54, 1996.

[3] Altera "Stratix II GX EP2SGX90 Transceiver Signal Integrity Development Board: Reference Manual". Available online: http://www.altera.com/literature/manual/rm_si_bd_2sgx90.pdf

[4] F. Faccio, G. Berger, K. Gill, M. Huhtinen, A. Marchioro, P. Moreira, and F. Vasey, "Single-Event Upset Tests of an 80-Mb/s Optical Receiver", *IEEE Trans. Nucl. Sci.*, Vol. 48, No. 5, pp.1700-1707, 2001.

[5] C. J. Marshall, P. W. Marshall, M. A. Carts, R. A. Reed, S. Baier and K. A. Label, "Characterization of Transient Error Cross Sections in High Speed Commercial Fiber Optic Data Links", in Proc. RADECS 2001.

[6] E. A Newcombe and S. Pasupathy, "Error rate monitoring for digital communications," *Proc. IEEE*, vol. 70, no. 8, pp. 805–828, Aug. 1982.

[7] D. J. Kennedy and, M. B. Nakhla, "Burst error characterization of FEC coded digital channels," International Journal of Satellite Communications, Vol. 10, No. 5, pp. 243-250, 1992.