

Radiation damage in LGAD detectors for HL-LHC

Anton Rindom Kristensen

Email: anton.rindom.kristensen@cern.ch

Supervisors: Jonathan Hollar and Ksenia Shchelina

(Dated: September 26, 2024)

SUMMARY

I worked with LGAD sensors, since they will be used to upgrade the PPS detector for the HL-LHC. I did experimental work by measuring IV and CV curves of trench-isolated LGADs, discovered that they break, broke them systematically and finally showed that they maybe don't break if all pixels are grounded. I also simulated NIEL curves in order to get a data point for 7TeV protons, discovering along the way that the current method might not be applicable for such high energies. I proposed a new method, that I still need to verify and cross check before claiming that it actually works. I also joined a bunch of researchers for beamtime and measured the timing performance of LGADs read out by ETROC chips.

All the code, data, plots, guides and logs can be found in [this GitHub repository](#). Don't hesitate to contact me!

CONTENTS

Summary	1
I. Introduction	2
II. LGAD Measurements	2
A. Probe station	2
B. Wire bonded	5
III. NIEL Simulations	7
IV. Testbeam	9
V. Outlook	11
A. LGADs	11
B. NIEL simulation	11
VI. Conclusion	11
VII. Acknowledgements	11
References	12

I. INTRODUCTION

The LHC will be upgraded after run 3 ends in the end of 2025. The upgrade will improve the luminosity by a factor of 5 to 7.5. This means that detectors will also receive significantly increased amounts of particles, which requires new technologies for various reasons. My project is specifically about the PPS sub-detector of the CMS detector.




Figure 1: The CMS with the PPS sub-detector. The PPS consists of several so called "Roman Pots" placed in various stations along the beam-line. Image shamelessly stolen from [2].

As can be seen on figure 1, the PPS has stations that are hundreds of metres away from the central CMS detector. The principle is as follows: Protons that scatter softly in the CMS interaction point will lose some of their energy. They will still follow along the beam pipe after the collision, but will be deflected more by the magnetic dipole field. This causes the lower energy protons to occupy a different spatial region from the main beam protons. By inserting detectors into this region, but keeping them away from the main beam, it is possible to measure the scattered protons.




Figure 2: The current timing detector, as well as the proximity to the main LHC beam. Most of the scattered protons will hit these diamond strips and generate a signal. The distance from the main beam to the closest strip is around 1.5mm. Image shamelessly stolen from [2].

As can be seen on figure 2, the sensors are very close to the main beam, and most of the scattered fluence will hit the sensors closest to the beam. This means that the sensors must be radiation hard, in order to

survive many years of operation in the LHC. The high-luminosity upgrade of LHC will worsen this radiation damage problem, as well as require better timing. For this reason, LGADs will be used as the sensor for the upgrade to the timing detectors of PPS.

Low Gain Avalanche Diodes, "LGAD", are a relatively new silicon detector technology [5]. The point is to have an intrinsic gain in the silicon, which the p^+ gain layer achieves, shown on figure 3. In this region the electric field is strong enough for charge carriers to multiply and create avalanches. A helpful side effect of the thin gain layer is that the drift time is very small, resulting in good timing performance.




Figure 3: Illustration of a standard LGAD (a) and a trench-isolated LGAD (b). The trench-isolation results in a much smaller insensitive region. Image shamelessly stolen from [6].

An unfortunate disadvantage of LGADs is the large inactive region between pixels. This inactive region is in the order of 0.1mm, and is significant when the pixels are only 1.3mm wide, and it makes it impossible to do fine grained tracking. An attempt at mitigating this is the trench-isolation as described in [6]. As can be seen on figure 3 the inactive region is much smaller. The original plan for my stay at CERN as a summer student was to measure the behaviour of irradiated trench-isolated LGADs.

II. LGAD MEASUREMENTS

A. Probe station

In building 28 at CERN there is a probe station which can perform IV and CV measurements on silicon sensors. The principle behind the measurements are relatively simple: Needles are physically manipulated, as shown on figures 4 and 5, so they touch various regions of interest on the chip. A bias voltage is then applied to

the silicon, and the resulting current and capacitance is read out by a source meter or LCR meter respectively.




Figure 4: A 2x2 trench isolated LGAD at the CERN probe station. Needles are touching the guard ring and a pad on one of the pixels.

Physically, what is expected to happen is that as the voltage is increased, the gain layer is gradually depleted. The current going through the silicon will quickly rise until the gain layer is fully depleted. Then it will rise slowly with voltage, until breakdown of the silicon, where the current will quickly rise again. Most of this can be seen on figure 6.




Figure 5: A 5x5 trench isolated LGAD at the CERN probe station as seen through the microscope. Needles are touching the guard ring and a pad on one of the pixels. The other visible pads on the pixels are for bump-bonding and laser measurements.




Figure 6: *IV curves of pixels on a 5x5 trench isolated LGAD. Measurements done on various pixels. Both the current of the particular pixel (solid lines) and the total current of the silicon (dashed lines). Pixel 44 (blue) was the first pixel to be measured on this particular TI-LGAD. Notice the two different scales.*

Many of the conclusions from my measurements can be seen on figure 6. The first feature is the gain layer turn on. The current going through the individual pixels reaches a plateau at around 28V. Among other effects, the voltage, V_{gl} at where the gain layer is turned on becomes lower when the silicon is irradiated [7]. This means that V_{gl} becomes a way to measure the NIEL scaling coefficient of the sensor. I will go into more detail about this topic in III.

The most obvious feature of figure 6 is probably the blue line representing a measurement on pixel 44. This pixel was the first we measured on this particular LGAD, and we let the voltage go to high values, since we could measure on 2x2 trench isolated LGADs at high voltages without having too much total current going through the silicon. On the 5x5 trench isolated LGADs this turned out to not be possible. The pixel under measurement

would suddenly break, the current going through the pixel would increase by many orders of magnitude, which made measurements stop do to a compliance setting on the source meters which would stop measurements if too much current was supplied. The pixels that fail in this manner stay broken, meaning that successive measurements on the pixels will not show any IV curve behaviour, just a high current even at low voltages. This suggests the failure of the pixel has introduced an electrical short somewhere in the silicon.

This failure mode was unexpected, and has rendered all our 5x5 TI-LGADs with at least one broken pixel. Note that the measurement on pixel 44 of this sensor was likely done before any pixels had been broken. This can be seen from the total current graph of the first measurement (dashed blue line on figure 6). The total current rises sharply when the failure happens, and the successive measurements (dashed lines) have a much higher total current. Interestingly, some pixels would also become broken while doing measurements on other pixels. So the failure would sort of spread.

We also measured CV curves, but all of the measurements were done on on chips that had at least one broken pixel. The results for one of the trench isolated 5x5 LGADs are shown on figure 7.

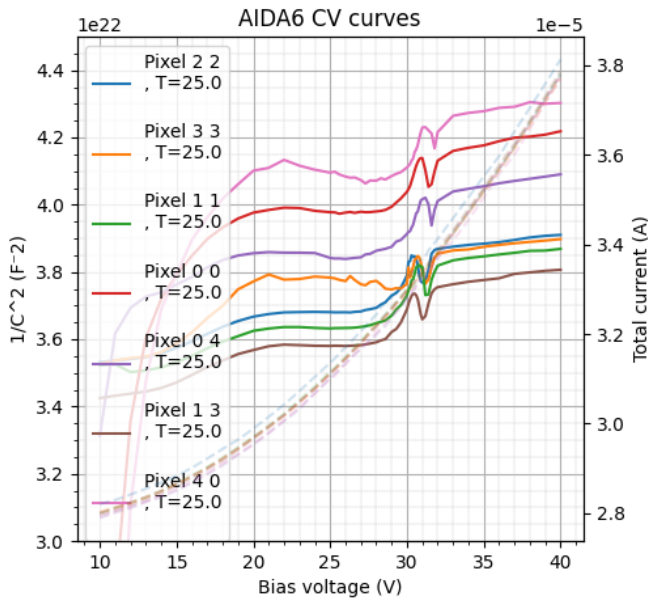


Figure 7: CV curves of pixels on a 5x5 trench isolated LGAD. Measurements done on various pixels.

One would expect the capacitance to fall (so $1/C^2$ would rise) when the gain layer turns on, which would be at 28V as read from figure 6. This is usually referred to as the "foot" of the CV curve, and can also be used to estimate the radiation damage coefficient of the sensor,

as it changes depending on received radiation dose. What is not expected is the "dip" at around 31V. It is uncertain whether this is a feature that comes from the broken pixel, or whether it is some quirk of the trench isolation. The 2x2 trench isolated LGADs didn't show this dip. That doesn't rule out that the dip could be from the trenches, especially since every pixel on a 2x2 TI-LGAD borders the guard ring. From figure 7 we can also see that every pixel has a different capacitance, which is plotted on figure 19.

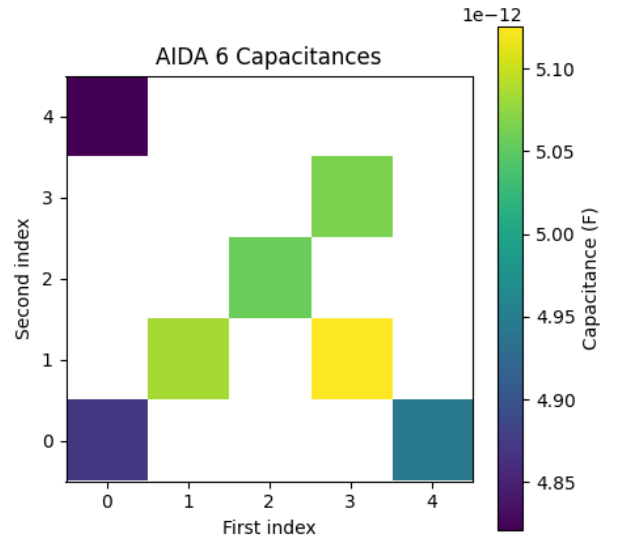


Figure 8: Map of the capacitance of the measured pixels.

It seems as if the pixels bordering the edge have a lower capacitance than the ones in the middle, but only by a few percentages. It might have an impact on the signal that the pixels will output to the readout electronics, giving a slightly larger timing uncertainty. But it is unclear whether this difference would still be there when the LGADs have no broken pixels, and also when all pixels are bump bonded to the ETROC readout chip.

As a side note, we also performed several measurements on other LGADs. They had been measured before, and nothing particularly new was discovered other than a weird "dip" in the IV curve of the 2x2 TI-LGADs just after the gain layer turn on.

B. Wire bonded

The working hypothesis from the measurements at the probe station was that individual pixels left floating could float to a high voltage, and then discharge that current in a way that left permanent damage somewhere in the silicon. In order to test this we wanted to measure on the

last 5x5 trench isolated LGAD we had at CERN, which was wire bonded to a little readout PCB. The wirebonding would allow us to ground all the pixels that were not to be measured, thus hopefully preventing this possible breakdown failure mode.

There was also a tiny hope that no pixels had been broken on the wire bonded TI-LGAD from the first preliminary measurements.



Figure 9: Wirebonded LGAD. Not every pixel has its own readout. Some of them are wire bonded to the same readout pad on the PCB.

In order to make measurements we rented two Keithley 6487 Picoammeters from the CERN electronics pool. I programmed a graphical user interface in python based on the PyQTGraph library. The program can scan through a list of voltages, and measure both the total and pad currents (with uncertainties!), as well as plot and save data and plots.

Since the Keithleys had limited current compliance functionality, I programmed a software compliance, where the picoammeter would slowly raise the voltage in small steps and measure whether it would hit the compliance limit entered by the user. The program, along with a guide on how to install and use it is available on my Github. The setup can be seen on figure 10.




Figure 10: Setup of the measurements on wire bonded LGADs. The LGAD is under the orange rag, inside a black plastic bag, because it is quite sensitive.

This setup is only able to measure IV curves, but could be modified with a function generator, a bias-T and an LCR meter, which would allow capacitance measurements also.

When we started measuring on the wire bonded TI-LGAD we quickly found that it unfortunately had a broken pixel already. This broken pixel was likely from the preliminary measurements where the TI-LGADs were all tested on pixel 00 when they were first received.

We still proceeded to do measurements on the TI-LGAD, the most significant shown in figure 11.

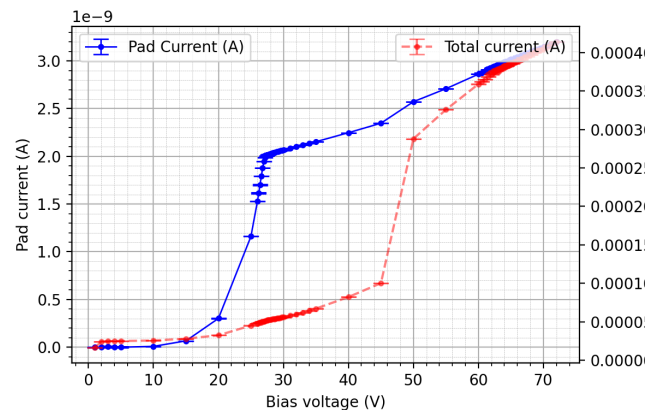


Figure 11: IV curve of a pixel on the wire bonded TI-LGAD. The silicon hit compliance after the measurement on 72V.

From measurements at the probe station we had seen that the pixels would break when the bias voltage was around 65V. So for the first measurement on the wire bonded TI-LGAD we only went to 70V. And it didn't break! So it seemed that grounding the other pixels (which is how TI-LGADs would be implemented with the ETROC readout chip) prevents the failure mode we had seen. The measurement on figure 11 is the second high voltage measurement. I made my setup scan up to 100V, but with the software compliance turned on. After 72V

the setup hit the software compliance and stopped. Afterwards successive measurements didn't show that the measured pixel (or any other pixel for that matter) was broken in the same way that we had seen at the probe station. But the total current going through the chip was now much higher, to the point that we couldn't do anymore IV curves without hitting compliance at 0.5mV.

That's where we got to on the measurement part during my stay at CERN. We had found a weird failure mode when pixels were not grounded, at 65V, that was repeatable. We showed that you can avoid this failure mode if grounding all the pixels. Personally I suspect that the new failure mode we saw in this last measurement was due to the already broken pixel (through which most of the current was probably going) broke even more. More testing needs to be done though, and I believe the next step is to get our hands on some fresh TI-LGADs to redo measurements, this time with no broken pixels. If those don't break, then the irradiation program can then continue as originally planned.

III. NIEL SIMULATIONS

Doing an irradiation test will give a coefficient that describes how much the performance of a sensor will deteriorate per amount of radiation it receives. But without any conversion it only really works for the type of particle that it was irradiated with, and at that specific energy. That's where the NIEL (Non-Ionizing Energy Loss) scaling hypothesis is useful. The core principle is that the damage in a silicon sensor scales with the amount of non-ionizing energy that is deposited in the sensor. By finding the damage coefficient from protons at one energy, one can then calculate the damage that protons at another energy will cause. Normally the NIEL value is scaled to the value of neutrons at 1MeV, which is 95MeV mb.

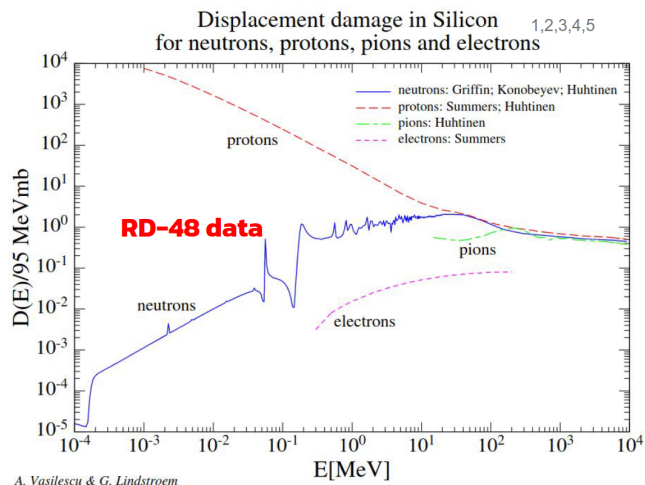


Figure 12: NIEL curve for protons, neutrons, pions and electrons in silicon. Taken from [4]

This is particularly helpful for the PPS sensors, which are hit by protons with energies above 6TeV. There is no irradiation facility that can irradiate at such high energies (closest is HiRadMat at cern with 450GeV), so the damage coefficient must be scaled in order to estimate radiation damage. There is an analytical approach to calculating the NIEL value, based on [3]:

$$NIEL(E_0) = \frac{N_A}{A[g]} \sum_i \int_{E_{min}}^{E_{max}} Q(E) E \left(\frac{d\sigma}{dE} \right)_i dE \quad (1)$$

E_0 is the incident energy of the particle, so the x-axis on figure 12. N_A is Avogadro's number, A is the atomic mass, so around 28 for silicon, $[g]$ is the unit of gram, i denotes every reaction that the incident particle can participate in, $Q(E)$ is the Lindhard partition function and $\left(\frac{d\sigma}{dE} \right)_i$ is the differential cross section of the i 'th reaction.

Q describes how much of the energy will be deposited as non-ionizing energy. It has been fitted from data by [1]. The energies in [1] go up to 100keV, and it is assumed that all the energy of the created recoil particles will get deposited in the material. This assumption breaks down for very high energy recoil products, or for very thin sensors. Since the particles that hit the PPS sensors are very high energy, and the sensors (especially the gain layer of the LGADs) are thin, I started looking into the validity of this assumption, and tried to simulate it. Originally I just wanted a data point at 7TeV, but I ended up focusing more on the validity of the Lindhard approach, and on coming up with a new method that would be more suitable for simulating damage at very high energies.

To simulate NIEL I followed the approach of [4], using Geant4 and a 1mm x 1mm x 0.1mm silicon geometry. The first simulation I did was a test to see the cross sections of each recoil product, and to calculate how much NIEL damage they cause. The result on figure can be seen on figure 13.

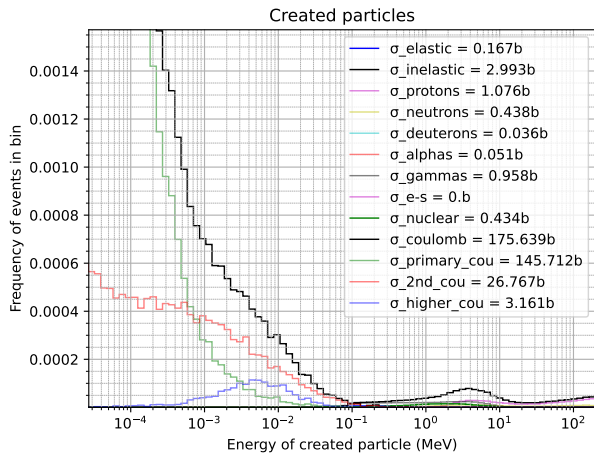


Figure 13: Cross sections of various recoil particles generated from incident protons with 200MeV. The coulomb and elastically scattered particles are silicon atoms.

These cross sections fit well with those from [4], so I felt relatively safe that I was simulating the same basic physics. I then calculated the damage that each of these recoils corresponds to, which is shown on figure 14.

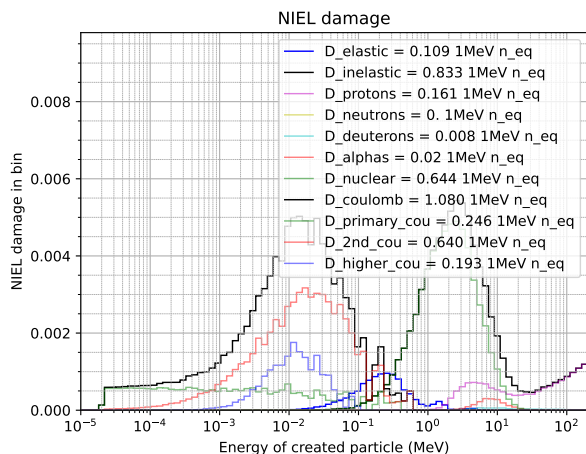


Figure 14: Damage done by various recoil particles generated from incident protons with 200MeV.

I could write a lot about these two plots, along with the similar plots for 7TeV. But the main point is that the approach of equation (1) counts a 100MeV proton that has been generated from an inelastic scattering as if all its 100MeV energy would be deposited in the silicon. From my simulation I could see that this was not the case, and in fact that such highly energetic particles would mostly just escape the volume, rarely depositing much energy. This effect is worse at 7TeV, where in general the recoil products are smaller particles, and tend to have higher energies.

These problems led me to consider an alternative approach to calculating the damage. I was inspired by the fact that my simulation could show me not only the damage due to the coulomb scattering from the primary protons, but also the damage done by coulomb scattering from the recoil products of inelastic events. My idea was to not use equation (1) directly on the product of inelastic events, but let the products propagate through the volume and coulomb scatter silicon atoms on their way. I would then use (1) to calculate the damage from these coulomb scattered atoms, which is much more valid since their energy is in the range of the fit from [1].

I implemented such a routine Geant4 and simulated protons at various energies to see the difference in the NIEL curve produced by the standard Lindhard approach and by my approach of looking only at the coulomb and elastically scattered silicon atoms. The results are shown on figure 15.

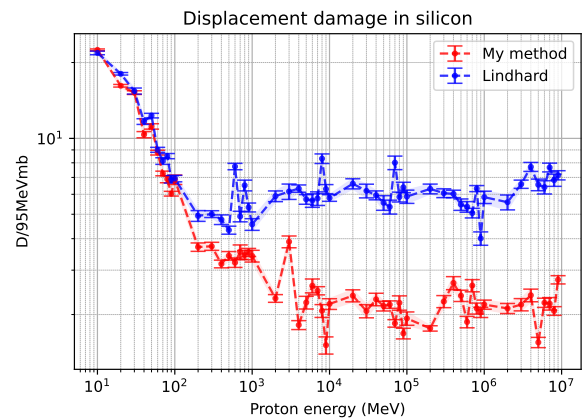


Figure 15: NIEL curve of the standard Lindhard method versus my method. The uncertainties shown are statistical, and found by performing several simulations and reporting the mean of the simulations' result. The statistical uncertainties are likely overshadowed by systematic errors.

The two approaches seem to agree well up to 100MeV. This makes sense, because in this region the damage is dominated by coulomb and elastic scattering. Then when inelastic scattering starts to have an effect, the two methods diverge. It is also worth noting that the values I calculate are too high compared to 12. I should be overcounting the damage, because according to TRIM simulations from [4] the value of Q of silicon recoils is about 20% to 30% higher than it should be, depending on the energy. I am also counting damage due to protons that are generated from inelastic scattering, which was not done in [4], I believe.

Care should be taken not to force new results to agree with old results, though. It can be argued that what I calculate in figure 15 and what is shown on 12 is not actually the exact same quantity conceptually. 12 shows

a property of silicon in general, while my simulations on figure 15 will depend on the geometry of the sensor, so it is more a property of my detector and its dimensions.

In order to study the behaviour of my method I did a few simulations. The first was to see whether a larger volume of silicon would make the two methods converge. This would be the case, since then the very high energy products from inelastic events would actually deliver all their energy to the material. The results are shown on figure 16.

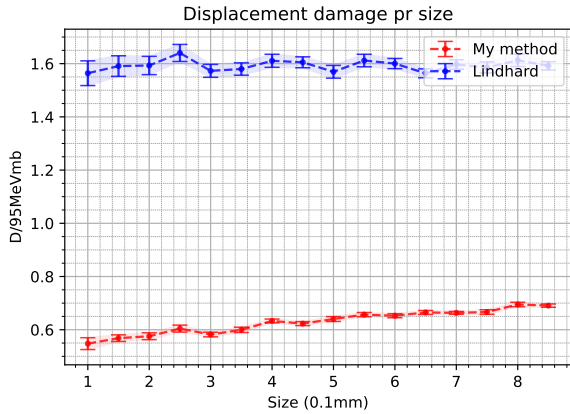


Figure 16: Simulation investigating the impact of sensor size of my method for protons at 7TeV. The time to simulate scales badly with size, and this simulation took around 3 days to run on my laptop, because of the high statistics.

As expected the damage increases as the size of the silicon increases. I am not certain that it should actually converge to the Lindhard method as the size goes to infinity. That would only be the case if the Q is correctly calculated for high energy product, which I know it not to be. The fit from [1] only uses data up to 100keV, and mostly for heavier ions, not protons, deuterons, tritons and alpha which make up the main body of the hadronic products at 7TeV. Furthermore the fit parameters from [1] do not have uncertainties on it, so it is difficult to quantify how well the fit works at high energies.

Since the damage done to a silicon sensor seems to slightly vary by size, I wanted to investigate whether the damage is uniformly distributed inside a sensor. This is particularly interesting for LGADs, since they are not just a piece of uniform silicon. If an LGAD was placed in a beam such that the gain layer is the last part that the proton travels through, then it could be imagined that the gain layer would receive more damage than if placed the other way around. My main argument for this is that the products from an inelastic shower tend to move in the same direction as the incident proton. So the probability of being hit by a shower product will increase the further you go into the material, since the incident proton then has more chances to create an inelastic scattering event.

I simulated the damage as a function of depth, and the results are shown on figure 17.




Figure 17: Simulation investigating the distribution of radiation damage by 7TeV protons along the depth of the sensor.

From this I would hesitate to conclude anything based on the uncertainties. There is a small downwards trend, but another run of the simulation showed a slight upwards trend. My takeaway from this simulation is that higher statistics might show some definite trend, but that the effect would probably be small. The effect should scale with the size of the sensor though.

This is the point that I have currently gotten to at the end of my stay at CERN. I plan to do more work on this, which will be described V.

IV. TESTBEAM

During my stay at CERN the group also had a week of beamtime at the SPS North Area. At the North Area protons from SPS (400GeV) are shot into targets, which create secondary beams. It is these secondary beam which we used for our tests. Interestingly, this secondary beam is not particularly well controlled, and is just described as "mainly consisting of different types of hadrons". Fortunately that does not impact our tests much as all these hadrons behave as Minimum Ionizing Particles (MIPs).

The plan was to test the timing performance of LGADs along with the readout electronics. Researchers from Fermilab brought the "Poor Guy", which is a suitcase containing a setup with four planes of LGAD sensors along with ETROCs and other electronics for readout.




Figure 18: *Me standing next to the Poor Guy, which I had aligned to the beam.*

We spent the first few days setting up the equipment, where I helped with various practical tasks, the power supplies, and the actual alignment of the suitcase to the beam. Once everything was ready and the first tests were done, we started a long series of measurements. Each measurement was about 4 hours of beam on the detectors, which meant sometimes coming in late and starting a new measurement. The measurements were scanning through a parameter space of different voltages on the LGADs, and different thresholds of the ETROC readout chip.

There was already scripts for doing the entire process, from data acquisition to final analysis. I didn't look too deeply into it, but it basically works like this: All hits on all pixels are stored and logged. From these hits (which are timestamped precisely by the ETROC), tracks are calculated from 3 of 4 detector boards. These tracks are then matched with hits on the last board, which is called the Device Under Test (DUT). This matching will create a distribution of time differences, the width of which gives the time resolution of the DUT. This way we could measure the time resolution of various settings. The results of one such measurement is shown on figure 19.

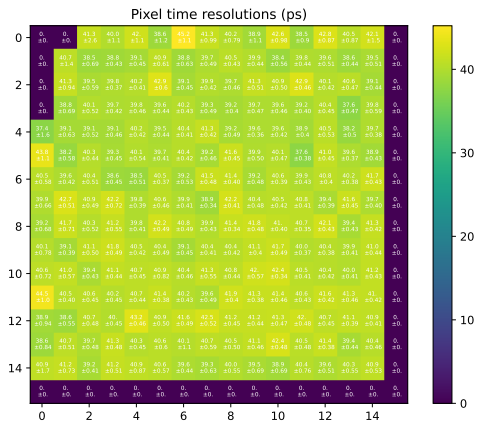


Figure 19: *Map of the pixels and their time resolutions.*

As can be seen, a few pixels on the edges did not have enough good tracks to calculate the time resolution from. The analysis script then binned these time resolutions and fit a Gaussian in order to get value for the time resolution of the entire board. I thought it was slightly silly to bin the data, when maximum-likelihood fitting can just fit directly on the PDF, so I also tried bootstrapping data according to the mean and uncertainties reported and then ML fitting a Gaussian to that. I also did the weighted mean for normally distributed data. Interestingly enough, the three different methods gave slightly different answers for the spread of the time resolution. 1.7ps, 1.3ps and 1.2ps respectively.

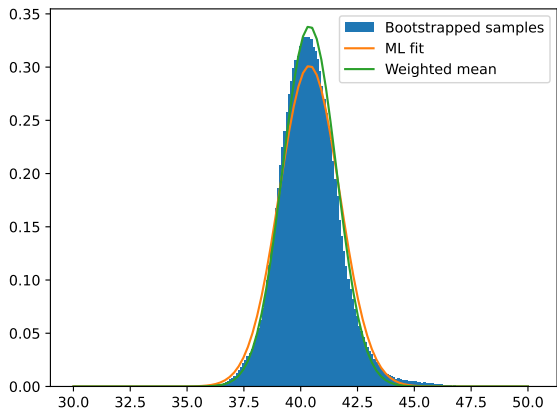


Figure 20: *Bootstrapped time resolution data, and ML-fit and weighted means result.*

It looks like the data is not entirely normally distributed, which makes sense since some of the pixels can be close to going into breakdown mode, which changes the timing performance. Regardless, I don't feel like it's particularly significant to use one or the other method to calculate the

timing performance, as long as the same method is used consistently.

Just as a side note, the best timing resolution we found was about 35 picoseconds, which I find absolutely incredible.

V. OUTLOOK

A. LGADs

The next step seems to be to get a fresh production of TI-LGADs, so that they can be tested. It would technically be possible to irradiate the broken TI-LGADs and try to characterize them, but any conclusion drawn would be questionable, since any effect could be caused by the pixels being broken. There is also the matter of whether the total current would rise enough to make IV curves impossible.

Since the 6 TI-LGADs were all from the same wafer, there is a probability that it was just a production error that will be fixed by the production of a new wafer. But if it is not a production issue, but a feature of the trenches, then it might be that any future measurements on large TI-LGADs need all pixels to be wire bonded and grounded.

In order to predict how much the sensors will be damaged inside the roman pots during run 4 I still need to set up a Geant4 simulation of the entire geometry. I am planning to do this after I have returned home from my stay at CERN. I can use both the standard Lindhard approach and my own more empirical approach to see what the differences are. The simulation should hopefully give a decent idea of which exact pixels on which boards will be the first to become unusable due to radiation damage, and give a prediction of how long it will take. If lucky the simulation will allow some engineering that could help ameliorate the radiation damage.

B. NIEL simulation

I feel like my approach is a good first step towards simulating NIEL damage done by very high energy protons, but I still need to verify that Geant4 is actually doing physics correctly. I am planning to use the TRIM software to cross check the Geant4 single coulomb scattering physics list, since that's the core of my simulations. There are also different Geant4 physics lists for low energy physics which it could be interesting to look at and see if they handle low energy recoil silicon atoms differently.

If the coulomb scattering physics is ok, I would like to see if I could get Geant4 to do some analysis of clustering. There is some suspicion, as mentioned by [4], that

clustered damage is different from non-clustered damage, and I suspect that the clustering is caused by higher energy silicon recoils (usually from elastic scattering instead of coulomb). It is not particularly relevant for radiation damage of LGADs, since any removed dopant will deteriorate performance. But it is interesting for a general idea of radiation damage in silicon, and definitely applicable to SiPMs for example, where the dark count caused by radiation damage can be a large problem.

VI. CONCLUSION

I worked with LGAD sensors, since they will be used to upgrade the PPS detector for the HL-LHC. I did experimental work by measuring IV and CV curves of trench-isolated LGADs, discovered that they break, broke them systematically and finally showed that they maybe don't break if all pixels are grounded. I also simulated NIEL curves in order to get a data point for 7TeV protons, discovering along the way that the current method might not be applicable for such high energies. I proposed a new method, that I still need to verify and cross check before claiming that it actually works. I also joined a bunch of researchers for beamtime and measured the timing performance of LGADs read out by ETROC chips.

VII. ACKNOWLEDGEMENTS

I want to thank Jonathan Hollar and Ksenia Shchelina for giving me the opportunity to come here, and to let me pursue the topics that caught my interest. I hope I can keep contributing to PPS so that I can pay back some of the gratitude I feel. I also want to thank Mohsen Khakzad, Cristóvão Silva and Francisco Garcia. It's been a joy doing experiments with you, and I really value the discussions we had and the advice you gave me. I hope I won't forget it!

I also want to thank the danish taxpayers for sponsoring my trip, and the CERN summer student team for making it possible.

-
- [1] A. Akkerman and J. Barak. New Partition Factor Calculations for Evaluating the Damage of Low Energy Ions in Silicon. *IEEE Transactions on Nuclear Science*, 53(6):3667–3674, December 2006. Conference Name: IEEE Transactions on Nuclear Science.
- [2] Edoardo Bossini. The CMS Precision Proton Spectrometer timing system: performance in Run 2, future upgrades and sensor radiation hardness studies. *Journal of Instrumentation*, 15(05):C05054–C05054, May 2020. arXiv:2004.11068 [physics].
- [3] J Lindhard. INTEGRAL EQUATIONS GOVERNING RADIATION EFFECTS (NOTES ON ATOMIC COLLISIONS, III).
- [4] V Maulerova-Subert, I Dawson, E Garutti, and M Moll. NIEL(non-ionizing energy loss).
- [5] N. Moffat, R. Bates, M. Bullough, L. Flores, D. Maneuski, L. Simon, N. Tartoni, F. Doherty, and J. Ashby. Low Gain Avalanche Detectors (LGAD) for particle physics and synchrotron applications. *Journal of Instrumentation*, 13(03):C03014–C03014, March 2018.
- [6] G. Paternoster, G. Borghi, M. Boscardin, N. Cartiglia, M. Ferrero, F. Ficorella, F. Siviero, A. Gola, and P. Bellutti. Trench-Isolated Low Gain Avalanche Diodes (TI-LGADs). *IEEE Electron Device Letters*, 41(6):884–887, June 2020.
- [7] Esteban Curras Rivera, Alessandro La Rosa, Michael Moll, and Fasih Zareef. Gain layer degradation study after neutron and proton irradiations in Low Gain Avalanche Diodes, June 2023. arXiv:2306.11760 [hep-ex, physics:physics].