12 September 2025

Investigating linearity and producing monitoring plots with the calibration systems of the ATLAS Tile Calorimeter

Anikó Horváth[1,2]

Supervisors:
Antonio J. Gómez Delegido[3]
Henric Wilkens[3]

1: CERN Summer Student, 2: ELTE Eötvös Loránd University, 3: CERN

Abstract

In this project, I used two calibration methods of the Tile Calorimeter detector of ATLAS to achieve two goals: describe the linearity of different photomultipliers (PMTs) with the laser calibration system, and create monitoring plots for the charge injection calibration system. First, the different kind of PMTs present in the detector were separated into 3 groups: the legacy and a new model of PMTs connected to legacy electronics, and the new PMTs connected to new electronics (a prototype of the electronics that will be installed during the High Luminosity Large Hadron Collider (HL-LHC) upgrade). Using different quantities to describe the PMT response's deviation from linearity, it was observed that the HL-LHC demonstrator, a hybrid module equipped with prototypes of the HL-LHC electronics but with backward compatibility with the legacy system, showed more nonlinearity than the other observed PMTs. In the second part of the project, a toolkit was developed to make long-term monitoring plots of the calibration constant obtained from charge injection measurements. It creates a simple way to see the difference between the measured calibration constants and the ones stored in the Tile Calorimeter database for all $\sim$ 10000 channels. A webpage to display all of these plots was also prepared to ensure easy access to the monitoring tools.

Contents

1 Introduction

During my work as a Summer Student, I had a chance to familiarise myself with the ATLAS Tile Calorimeter and its calibration methods. ATLAS is one of the 4 major experiments in the LHC ring, and it aided in the discovery of the Higgs boson [1]. The detector consists of toroidal magnets and a cylindrical structure with close to 4π coverage in solid angle built up of concentric layers of different purpose detectors: the inner tracking detector closest to the interaction point and surrounded by a thin superconducting solenoid, followed by electromagnetic and hadronic calorimeters, and a muon spectrometer incorporating three large superconducting air-core toroidal magnets.

In the following sections, two calibration methods of the Tile Calorimeter (TileCal), will be used to investigate the linearity of different photomultipliers, and to create automatically generated plots for monitoring purposes. First, I will discuss the makeup of this detector in general, and then explain the different methods that are used to calibrate it during data taking periods. After that, I will present the methods and results of the project in detail.

1.1 The Tile Calorimeter

TileCal is the central hadronic calorimeter of the ATLAS experiment. It is a sampling calorimeter, composed of alternating layers of thin steel absorbers and plastic scintillators. It contributes to the energy measurements of jets, single hadrons, and taus and helps to determine missing transverse momentum with high precision [2].

It is divided into a long barrel (LB, $|\eta| < 1.0$) and two extended barrels (EB, $0.8 < |\eta| < 1.72$). In the azimuthal direction, it is separated into 64 modules ($\Delta\phi \approx 0.1$). To measure the energy deposited in the detector, the light produced in a collection of scintillator tiles is transported by two wavelength-shifting fibres per tile to photomultiplier tubes (PMTs). The grouping of fibres form approximately 5000 cells that are read out by 9852 PMTs. The signal is then shaped and digitised on-detector into 7 samples in two gains - with a signal amplification factor of ~ 64 - in the analog-to-digital converters (ADCs). After that, the data is transported off-detector for further processing.

1.2 Calibrating the Tile Calorimeter

In the ATLAS physics programme, a performance requirement of no more than 2% of non-linearity was set out for jets up to 4 Tev [3]. To achieve this goal, there are multiple

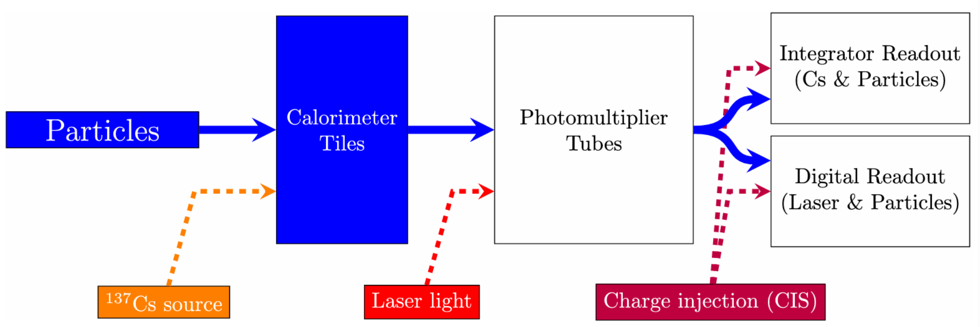

Figure 1: The different calibration systems of TileCal and their signal paths. Source: [2]

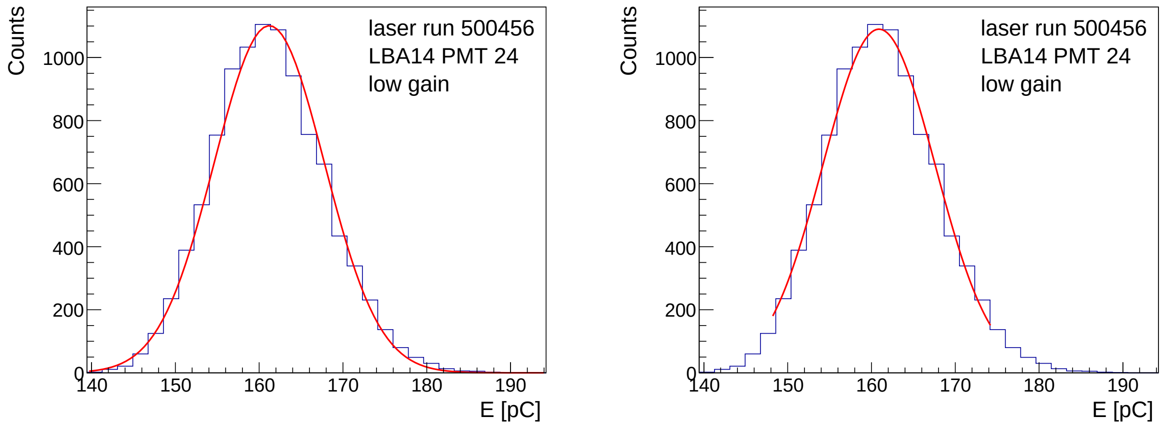

Figure 2: An example of the fits used to determine the response of the PMTs (left), and a refit on the $\pm 2\sigma$ interval (right) that was performed to prevent the tails of the distributions (which were not always perfectly Gaussian) from skewing the means.

systems built in to monitor the detector performance and calibrate the response if needed.

In Figure 1, we can see the different calibration methods. With the Caesium system, the radioactive isotope is pushed through the detector pipes with a hydraulic system to monitor the full optical chain of the calorimeter, from the scintillating tiles to the wavelength shifting fibres and the PMTs. However, this calibration system employs a different readout that integrates the signals produced by the Cs source in a larger time window. The Cs calibration allows transporting the signals to the electromagnetic scale determined in test beam campaigns before the start of the LHC operations. This measurement takes significant time, so it is only done once a month.

In the charge injection system, two capacitors are used, and they are discharged to inject known amounts of charge into the TileCal front-end cards, that shape and amplify those signals. The signals are then digitised on-detector, mimicking the read out path of the physics signals. The system is used to monitor and calibrate the output of the ADCs, determining the conversion between charge (in pC) and the pulse amplitude in units of ADC counts. It is taken every 2-3 days.

The laser system uses a laser head positioned about 100 meters from the detector in the service cavern. The light produced by the 532 nm laser is transported to the detector via approximately 400 optical fibres to the PMTs in the detector [4]. The laser light is also routed to a set of photodiodes to monitor the light yield and normalise the signals measured in the detector. Calibration runs are taken on a daily basis.

2 Linearity investigation with the laser system

The motivation for this exploration was a study of the Run 2 data [4], where the linearity of the photomultipliers was investigated. While the focus there was more on monitoring the linearity in the detector over time, here the attention will be focused on the different types of photomultipliers present in the detector. For the analysis, laser runs taken in the second half of June 2025 with different settings will be evaluated.

In the last few years, some of the PMTs that were initially installed before the start of

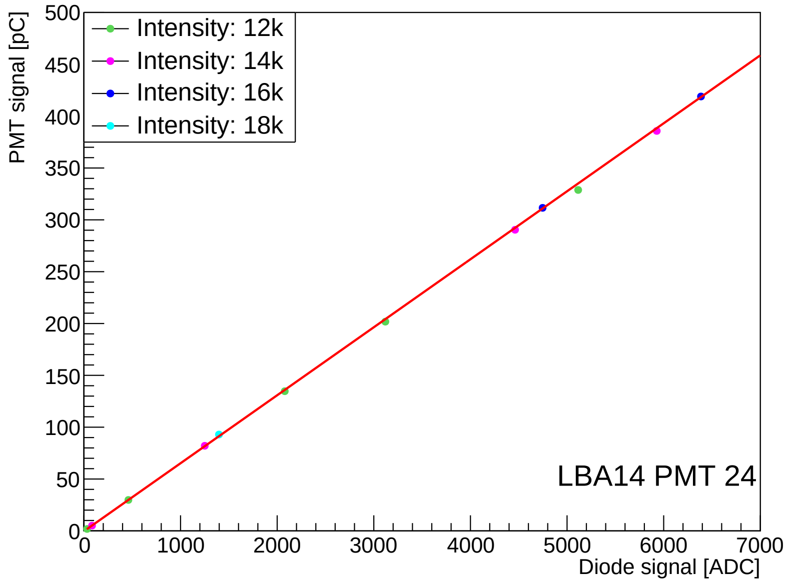

Figure 3: Example of the data points and line fit obtained by the analysis. The PMT response is plotted as a function of the reference diode signal.

the LHC operations (Hamamatsu model R7877) have been replaced with a newer model (Hamamatsu R11187). The replacement was necessary because the older PMTs' responses had degraded with accumulated anode charge, and the newer model exhibits less gain drift under the same conditions. This upgrade ensures more stable operation under the high integrated charges expected in future LHC runs with increased luminosity and pile-up.

During the High-Luminosity upgrade of LHC, about 10% of the PMTs will be replaced by the newer model, and the full TileCal readout system will be replaced to cope with the increased data flow, and to further tolerance against higher levels of radiation in the on-detector electronics [5, 6]. A prototype of these electronics were installed before the start of Run-3 in one of the modules (LBA14) as a demonstrator. All the PMTs present in the demonstrator are the new R11187 model.

In the different runs, the amount of laser light received was varied using the tools in the optics box. The intensity of the laser source itself can be changed, and a filter wheel is also part of the optical path, where the light attenuation can be adjusted by choosing different filters in the filter wheel (from 100% light transmission down to 0.3%). Using one of the diodes as a reference (Diode 6, which measures the laser light close from the fibre bundle that distributes it to the detector, and therefore close in optical path to the light used for calibration) and illuminating the PMTs with differing amounts of light, the linearity of the photomultipliers was tested.

2.1 Measuring laser response

The first step in analysing each run is looking at the seven signal samples of each event (for each PMT individually), and removing any event where the samples are all zero, or the ADC is saturated. After cleaning the events, the fit method results were extracted, which uses a pulse shape template to fit the 7 digitised samples and determine the amplitude of

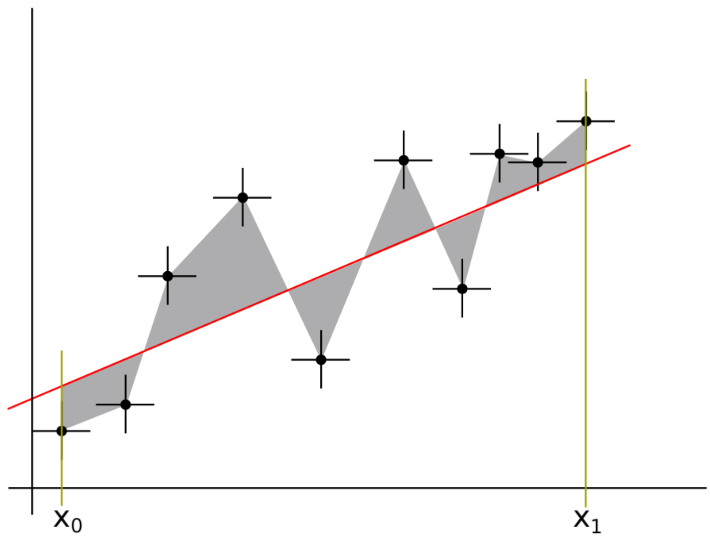

Figure 4: The area calculation for quantifying non-linearity. Source: [4]

the pulse in ADC units. The amplitude can be converted to pC using the CIS constant. The next step is fitting the charge distribution of the run with a Gaussian function. To have better precision, a refit was performed on the $\pm 2\sigma$ range of the first fit, as it can be seen in Figure 2. From this, we can extract the mean of the Gaussian (μ), and also the error of the mean:

$$\Delta\mu = \frac{\sigma}{\sqrt{N}},$$

where N is the number of events, and σ is the standard deviation of the Gaussian. Similarly, the value and error of the reference diode signal (in ADC counts) can also be evaluated for every run. If the readout electronics are linear, we expect the mean of the charge measured in the PMT to be proportional to the diode signal.

2.2 Quantifying linearity

Figure 3 shows an example of the fitted line obtained from the analysis in the previous section. The points at first glance do seem to fit on a line - their errors are too small to be seen due to the large number of events in the runs, which gives high precision on the determination of the mean of the distribution. Following the method employed in Reference [4], the line fit was performed iteratively to avoid high-energy non-linearities skewing the results in the low range.

To assess and compare their linearity, a quantitative description was needed. For that, I used 3 different methods: the first was a χ^2/NDF test, in the second I assigned a percentage of nonlinearity to every fit, and for the third one, I used the residuals of the fit. In the next sections, the method of the latter two quantifications is presented in more detail.

2.2.1 Non-linearity percentage

This is the method that was used by the Run-2 study. Calculating the area that is enclosed by connecting the data points with straight lines and intersecting them with the fitted line (shown in gray in Figure 4), and then dividing it by the area under the fitted line

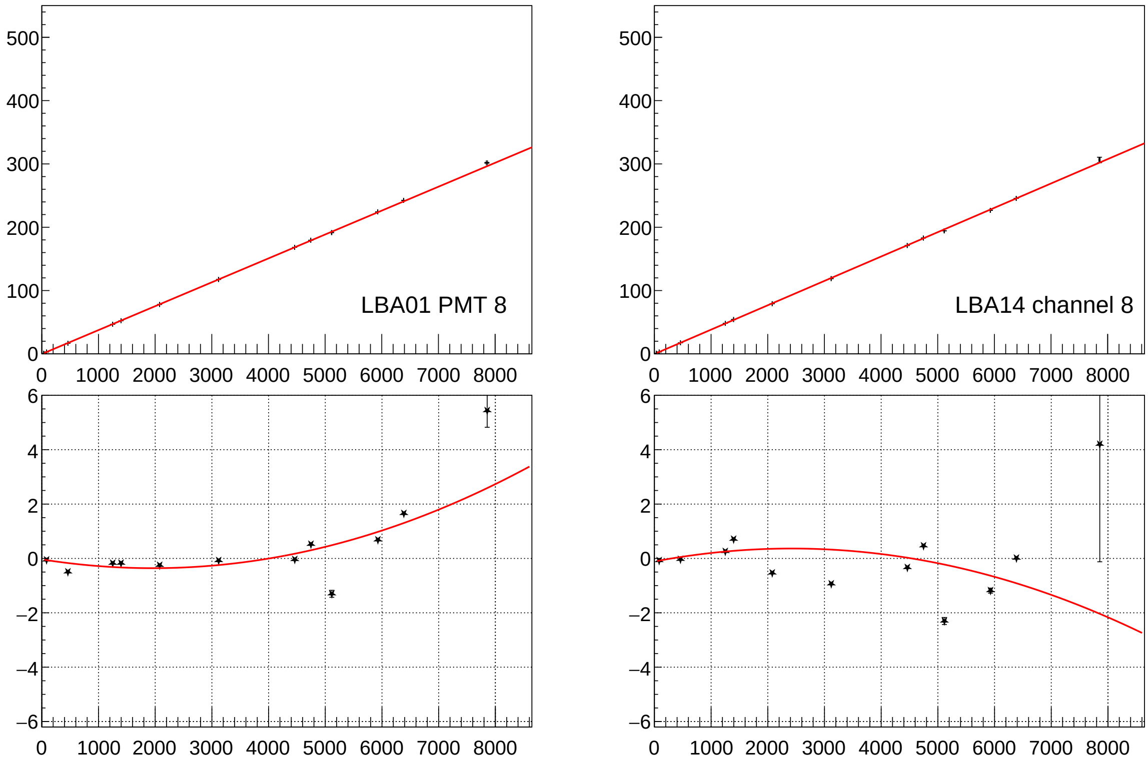

Figure 5: Two examples of the line fit (similar to the one shown in Figure 3) and the residuals polynomial fit. Left: a legacy PMT readout by legacy electronics, right: a PMT from the demonstrator module.

gives back a quantifier that describes the points' deviation from the line. This quantity, expressed as a percentage, will be used as one of the linearity estimates in this work. It can be less descriptive if the slope of the line is too small, which was not a problem for the PMTs we desire to evaluate.

2.2.2 Residual evaluation

This method is similar to an evaluation that is used at the TileCal test bench used to test the PMTs before installing them on the detector. After performing the linear fit described previously, the residuals of the data points with respect to the best-fit line are extracted and a second degree polynomial is fitted on them. An example of the fits with the legacy electronics (left) and the new electronics (right) can be seen in Figure 5.

Since the laser light is transported by many different fibres, the light received by individual PMTs, and consequently the charge range of the linear fits can vary (the reference signal is the same for all of them). To avoid biases in the test that could be introduced from quantifying the non-linearity of the PMTs at different operational regimes, only PMTs that test the range 300-500 pC were chosen, and the polynomial fits were extrapolated to measure the PMT non-linearity always at the same operational value. The extrapolation was done by evaluating the linear fit at 600 pC, and finding the corresponding value of the quadratic fit on the residuals. This means that for lines with smaller slopes, the extrapolation will be further than for lines with larger slopes.

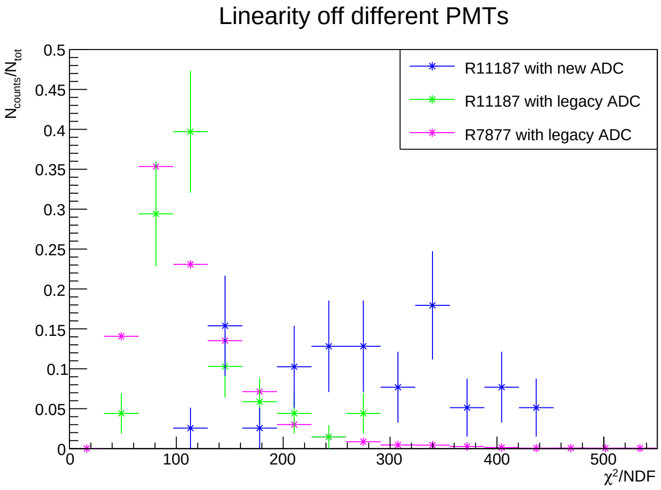

Figure 6: Comparing the distribution of different types of PMTs (normalised to unity) for the χ^2 test.

2.3 Results

In this section, I will describe the results of the different kinds of analyses. The three categories of PMTs to evaluate are the following:

- Legacy photomultipliers readout by legacy electronics,
- New photomultipliers readout by legacy electronics,
- New photomultipliers readout by new electronics (in the demonstrator module).

To compare them while taking into account the differing numbers of PMTs in each category, all distributions are shown normalised to 1. As we can see in Figure 6, both models of PMTs perform in a similar manner when connected to the same electronics. The large absolute values of the χ^2 can be attributed to the small errors of the mean (coming

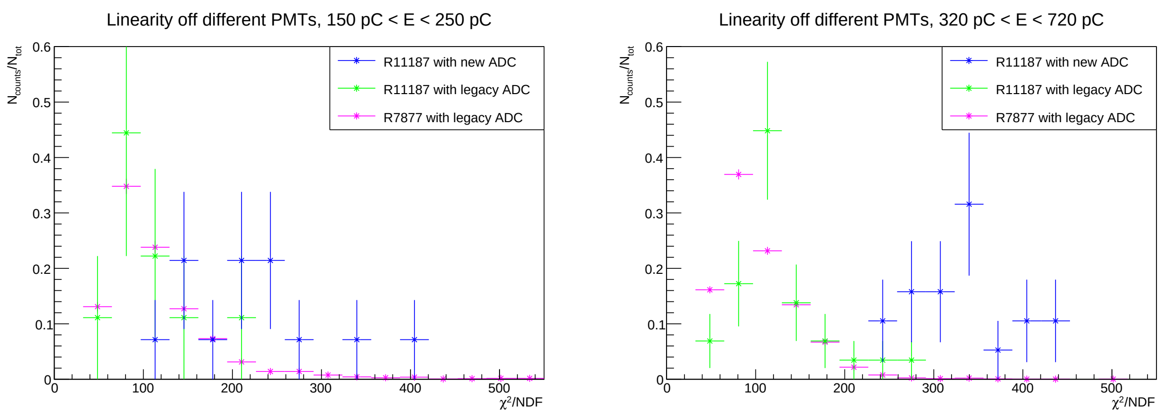

Figure 7: PMTs separated according to the charge range covered in the χ^2 test. The distributions on the left show the linearity test results for PMTs producing pulses in the 150–250 pC range, while on the right are the results for 320-720 pC.

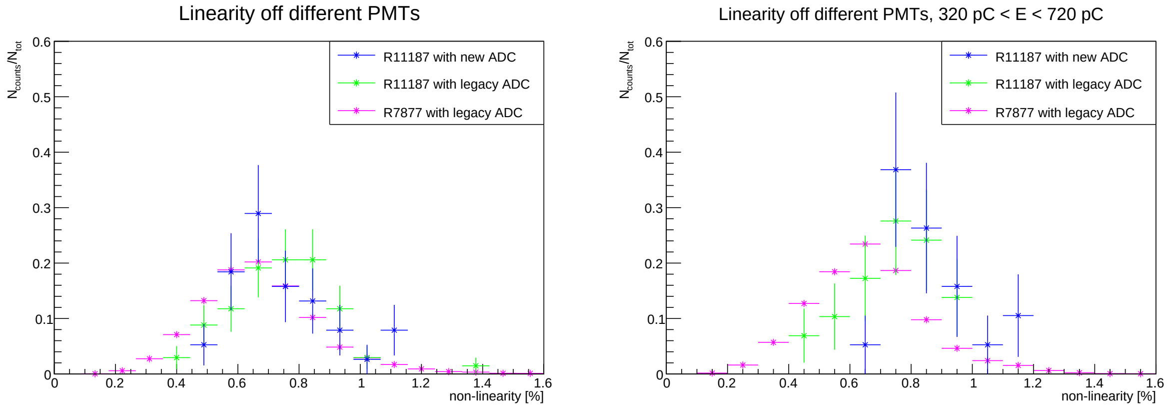

Figure 8: Distribution of non-linearity for all PMTs (left) and the ones in the higher energy range (right).

from many events in each run). However, we can clearly see that the demonstrator module (in blue) shows a larger non-linearity than the other two categories.

The fibre transmission is different from fibre to fibre, which implies that PMTs illuminated by different fibres are excited in different charge ranges. In the long barrel, two fibres transmit the light to each module. In LBA14, one of the fibres transmits more light, and as a consequence, the PMTs can be separated into a high- and a low-energy group. Figure 7 shows that we see somewhat more consistency with the rest of the PMTs in the low-energy range, while in the high-energy one, the demonstrator distinctly shows worse non-linearity. Based on this, we can say that the non-linear effect is larger in the high-energy range.

To avoid the problem of different fibres exciting different charge ranges, we take advantage of the fact that the same fibres that transmit light to LBA14 also transmit it to LBC14, which gives an opportunity of direct comparison of the entire readout chain. As seen in Figure 9, the performance of the demonstrator also shows more non-linearity even when subjected to the same light.

In Figure 8, we can see the results of the non-linearity percentage test. On the left side,

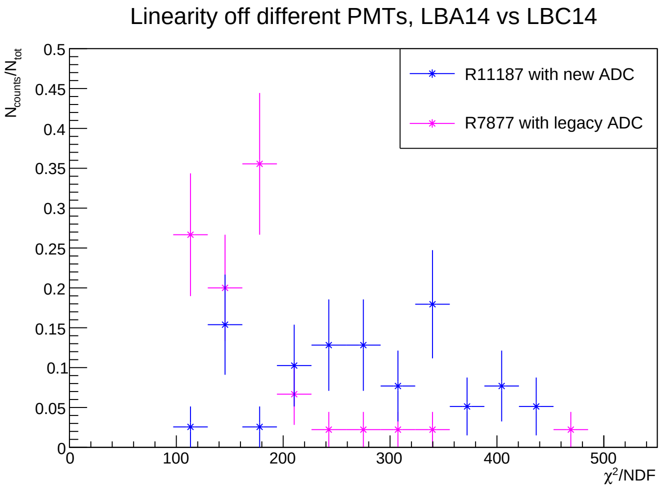

Figure 9: Comparing the distribution of PMTs in LBA14 and LBC14, two modules that receive light by the same fibres.

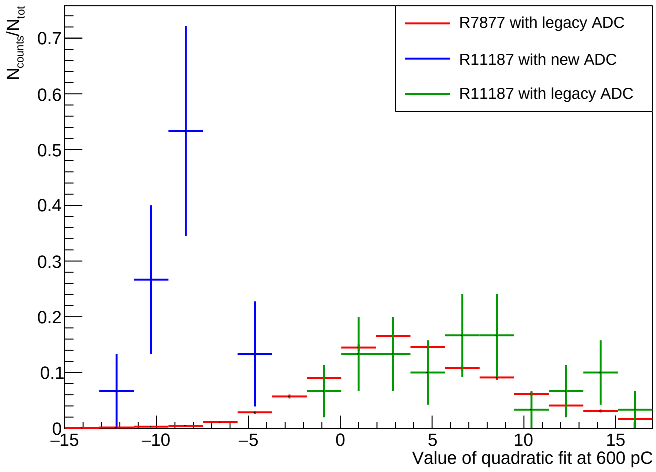

Figure 10: Distribution of PMTs in the residual test.

no clear difference can be seen between the three categories. However, when considering only the PMTs that are illuminated with larger light yields, the distribution of new PMTs and electronics does show a tendency towards less linearity (the mean of the distribution is larger than the legacy system's), as seen in the previous test.

To investigate further, a third test where a second degree polynomial fits the residuals between the data points and the best-fit line resulting from the linearity fit was also done. Figure 10 shows the result of the extrapolation of the polynomial fit to 600 pC of measured charge. As we can see, the two different kinds of PMTs again show good overlap when readout by legacy electronics, with a tendency towards a positive, upward curve of the parabola. The demonstrator module tends to have larger, and solely negative values of the fit. This indicates a downward curving parabola with larger deviation from linearity as we increase the energy to 600 pC. We can also see in Figure 5 that the demonstrators' points are scattered and do not adjust well to a parabola, implying that this method cannot give a quantitative description of the non-linearity in the demonstrator. However, it can still affirm that the behaviour is different in the demonstrator than in the legacy and upgrade PMTs readout by the legacy electronics.

Based on the tests, the type of PMT installed did not cause any significant difference in the linearity when readout by the same electronics. In the case of the demonstrator module, the linearity seems to be worse, especially as we go towards higher ranges of energies tested. The origin of the worse non-linearity seen in the demonstrator module is not clear. Among other possibilities, it could be caused by additional noise in the prototype electronics. Another possibility that should be investigated is usage of a non-linearity correction on the low gain ADCs that was derived during the first years of data taking for the legacy electronic. The drop in response in the legacy low gain ADC around 40 pC occurs when the high gain branch of the analogue electronics begins to saturate. The two branches are connected across a common source, and when the high gain saturates, small shifts of its input offset voltage can propagate to the other branch. It is unclear whether this effect is present in the prototypes for the HL-LHC electronics that are installed in the demonstrator, and this should be investigated in the future.

3 Monitoring plots for the charge injection system

The stability of the response of a given readout channel to a calibration system is a good indicator of a well-behaving channel. The scrutiny of channels with unstable responses has proven to be an effective way to identify pathological channels that should be treated or masked to maintain a good quality on the data. For this purpose, there already existed a webpage with monitoring plots for the laser calibration system: https://tilecal.web.cern.ch/Tucs/Laser/.

The goal of this work was to create something similar for the charge injection system (CIS). The TileCal Universal Calibration Software (TUCS) enviroment was used for this. It consists of a data structure and processors called workers that operate on it. A macro is a python code that tells TUCS the order in which to run the workers. For example, in the macro I created (`CIS_stability.py`), the charge injection runs in the given time frame are selected, the bad channel flags and calibration values are read from the database, and then my worker for calculating differences and making individual module plots (`do_cis_stability_plot.py`) runs last.

Some key arguments that can be given to the macro (other than the generally used ones):

- percentage: Configures whether the plot shows the relative (in %) or absolute difference with respect to reference point. If it is False, two axes are plotted for high gain and low gain in case of absolute difference. This allows the differences in the two

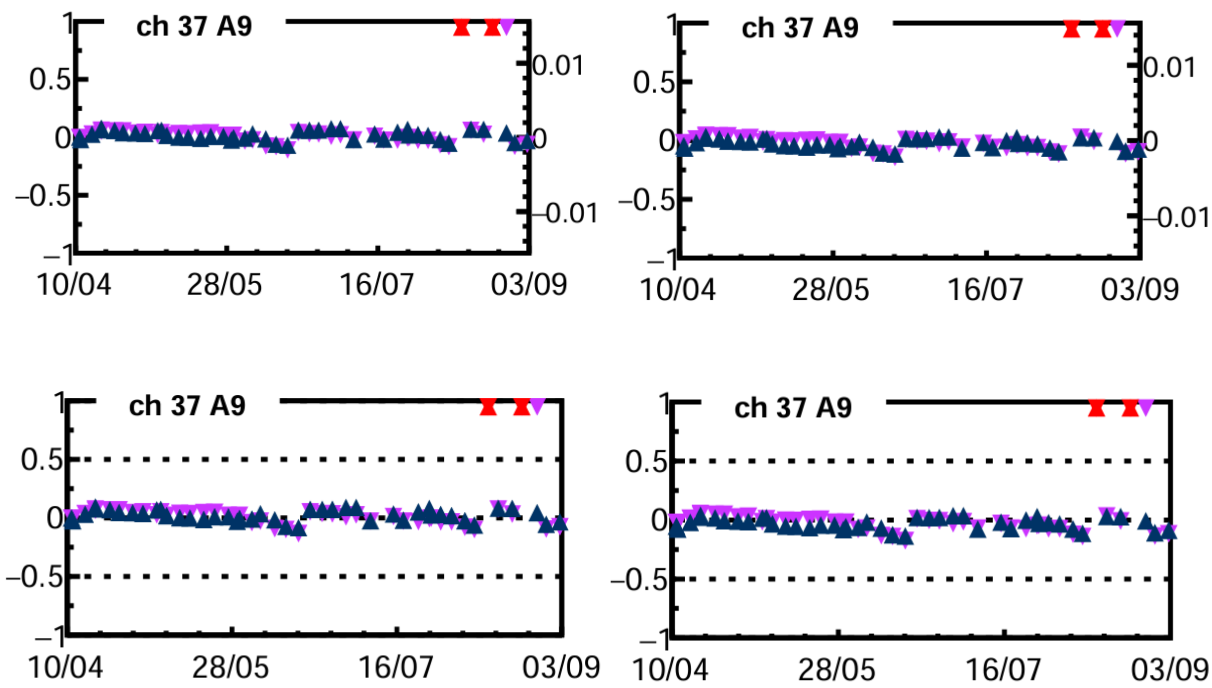

Figure 11: The different options for plotting. Triangles pointing up correspond to high gain, and ones pointing down to low gain measurements. The left column compares to the first measurement, while for the right column, the reference is the correct IOV database value. The top row shows absolute differences, with the left axis corresponding to high gain and right axis to low gain. The bottom row shows the relative differences in percentage.

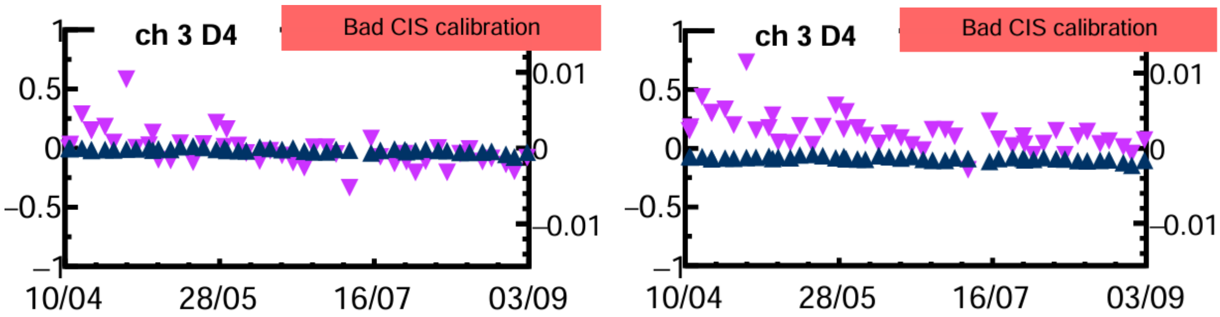

Figure 12: The two options depicting absolute differences for channel 3 of the EBC13 module. On the left plot, the reference point is the first measurement, and on the right side it is the database value. The channel is flagged for Bad CIS calibration in the database.

gains — one of which is approximately 64 times larger — to be visually comparable, with the low-gain axis scaled by a factor of $\frac{1}{64}$ relative to the high-gain axis.

- fromdatabase: Decides the reference point chosen. If the flag is set to true, the reference for each point is the database constant for the correct interval of validity (IOV). If the flag is set to false, the reference point is the first measured constant in the selected time frame.

The plots resulting from the 4 possible combination of settings in one channel of the module LBC20 are shown in Figure 11. The axes are the same for all plots, so for any values that would fall outside the range, they are stuck to the top edge of each plot. The ones coloured red correspond to measurements with the "Fail likely calib." CIS flag, defined within Tucs, and are points that fall out of 6.25% range of the detector average.

We can see that all plots look very similar in Figure 11. This is explained by the fact that high gain constants tend to be around 80, and $1/80 = 0.0125 \approx 1\%$. The differences can be spotted when analysing a channel with larger variation, as shown in Figure 12. In the macro, the argument "limit" can be given to the worker, which sets the range of the plots, either in percentage or in absolute value of the high gain. The channels are typically recalibrated when the difference between the measured CIS constant and the one stored in the database is larger than 0.5%. A drift of 1% already indicates abnormal behaviour in the front-end electronics, so the default value of the limit parameter was chosen to be 1.

The layout in each module follows ordering by channel number, as it can be seen in the Appendix (Figures 14 and 13). Empty regions correspond to channels where no PMT is instrumented. For channels that are flagged in the bad channel database, a red box is attached to the top right corner, displaying the problem(s) that are labelled in the bad channel database.

3.1 Making the webpage

For displaying the plots with percentages of difference from the database values, a webpage was made: https://tilecal.web.cern.ch/Tucs/CIS/. It contains the PDFs for the partitions, the demonstrator, and also all modules together in one file. Alongside these, channel specific plots are also made for either channels flagged as bad in the database, or ones that have sizeable variation. These are made with already available workers, which produce a plot for each channel and gain, and show the absolute value of the CIS constant as a function of time, among some other useful information.

An automated, recurring task that executes a specified command to produce the plots and a script to update the html source code of the webpage at a pre-defined schedule, managed by the cron daemon, is executed daily on the TileCal calibration workstation. The plots produced in the webpage can serve several purposes, for example for data quality, so that the team responsible for guaranteeing the integrity of the data can quickly check the integrity of the channels, or to check if the variation of laser response in coming from a change in the ADCs.

4 Conclusions

In this work, the linearity of different photomultipliers was tested. Using multiple evaluation methods, I could conclude that the choice of PMT model did not show any significant difference, meaning that the resilience against aging does not seem to degrade the performance in terms of linearity. However, the demonstrator module consistently performed worse in the test. There is a possibility that this might have to do with a correction implemented for the legacy electronics applied to the entire detector, but determining this needs further investigation.

A webpage was also made for monitoring the charge injection constants, similar to the page already used for the laser system. It gets updated daily, and can be used to quickly check the performance of all channels in the detector, so it provides another option alongside the previously used CIS plots, where one can investigate issues deeper channel by channel.

Acknowledgements

I would like to thank my supervisors Antonio and Henric for their constant support during my time at CERN, and their insightful comments that helped me in moving the project forward. I also want to thank my friends, family and partner in encouraging me to apply to the Summer Student Programme, and the friends I made during my time here.

References

[1] ATLAS Collaboration et al. Observation of a new particle in the search for the standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716:1–29, 2012.

[2] ATLAS Collaboration. Operation and performance of the ATLAS tile calorimeter in LHC Run 2. *European Physical Journal C*, 84:1313, 2024.

[3] ATLAS Collaboration, G. Aad, B. Abbott, et al. Readiness of the ATLAS Tile Calorimeter for LHC collisions. *European Physical Journal C*, 70:1193–1236, 2010.

[4] M.N. et. al. Agaras. Laser calibration of the ATLAS Tile Calorimeter during LHC Run 2. *Journal of Instrumentation*, 18(06):P06023, jun 2023.

[5] Sandra Leone on behalf of the ATLAS Tile Calorimeter System. Upgrade of ATLAS hadronic Tile Calorimeter for the High Luminosity LHC. *Journal of Instrumentation*, 19(01):C01036, jan 2024.

[6] Technical Design Report for the Phase-II Upgrade of the ATLAS Tile Calorimeter.

Appendix: Plots for an LB and EB module

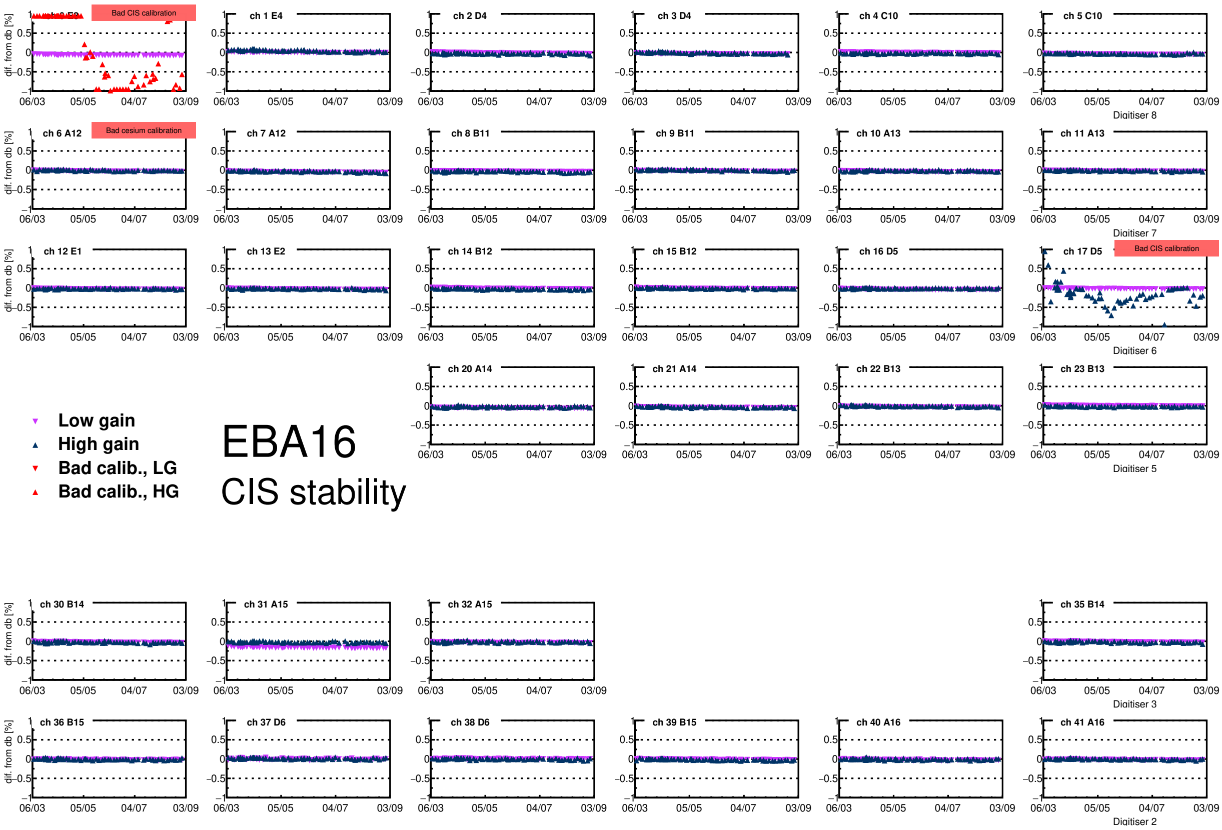

Figure 13: Channel layout of a module in an extended barrel module.

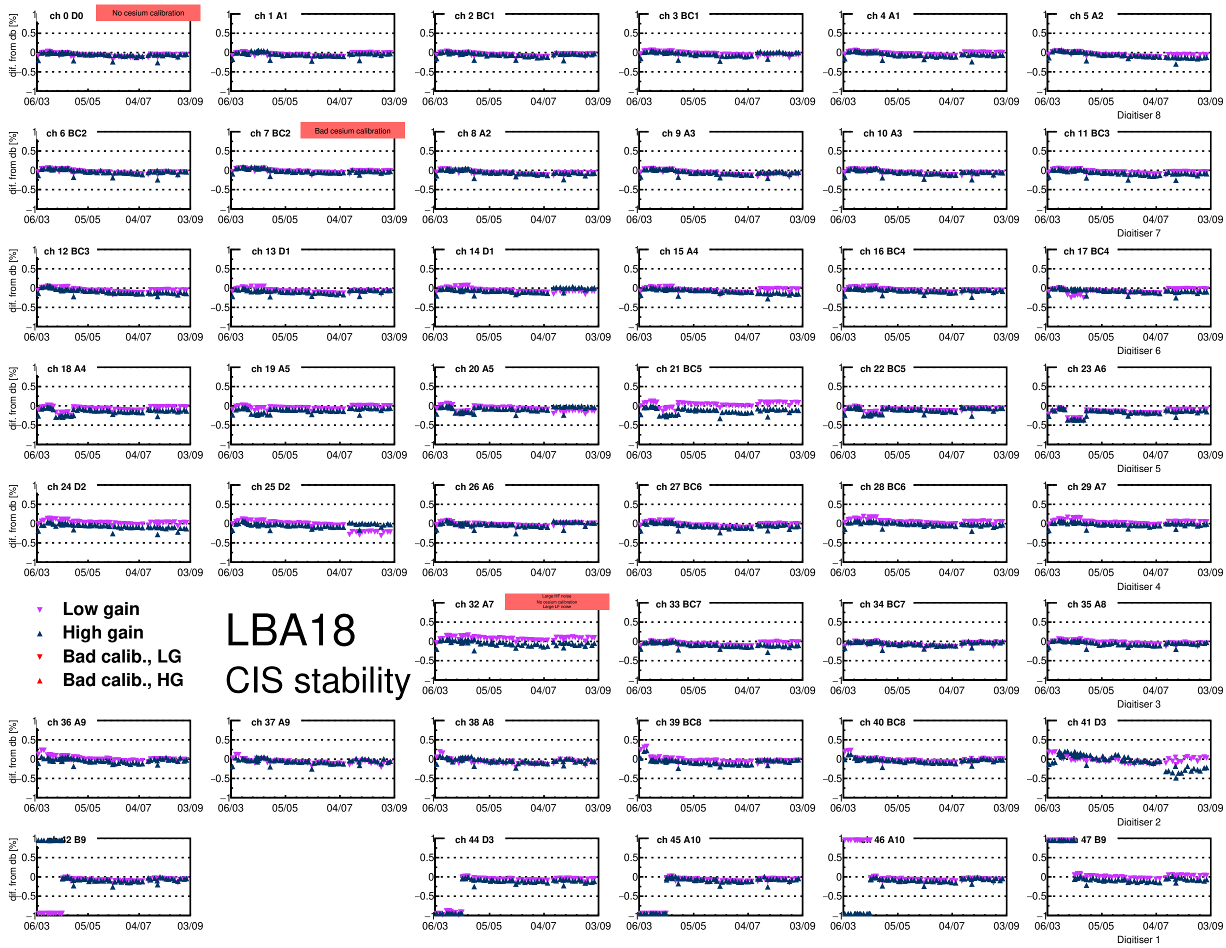

Figure 14: Channel layout of a module in a long barrel module.