Analysing the current noise equalization in VeloPix

Author: Lucas Clandfield Galia

Supervisor: Dr. Kristof De Bruyn

Second supervisor: Dr. Julia Even

July, 2021

Contents

1   Introduction

Since the discovery of atoms we've wondered what lied beyond that. We're still studying the components of elementary particles. One way of achieving this is by making particles collide and seeing what results from it. This is currently being done at the Large Hadron Collider (LHC) in CERN. It accelerates protons and makes them collide to then study the particles that form them.

The LHC [9] is the world's largest and highest-energy particle colliders. It lies in a tunnel 27 km long in a circumference beneath the France-Switzerland border in Geneva. Since the end of 2018 it was shut down to implement upgrades and it has very recently got back online to start doing experiments again.

The collider has four different crossing points where the accelerated particles collide, these crossing points each come with their set of detectors to detect different phenomena and they entails different experiments. These experiments are: ATLAS, CMS, ALICE and LHCb. The experiment we'll be focusing on is the LHCb [5, 10], which we'll get into detail about in the next section.

A part of the LHCb is the Vertex Locator (VELO)[1, 6, 7] which is responsible for locating where the collision between protons occur. VELO has recently been upgraded [7] in the previous shutdown to a better version. Data from the upgraded VELO is read out by the new VeloPix ASICs. They consist of a grid of pixels where an electrical current gets measured every time a particle goes through a pixel. Every grid of pixels in the new detector are called VeloPix, we'll go into more detail about the structure of these grids in the next sections.

Pixels in VeloPix have background noise associated with them [2, 3], meaning that there is a residual current that gets measured whenever no particle is moving through the pixel. To eliminate this noise we want to implement a current threshold to consider a current spike as a relevant event. The main issue with this idea is that all pixels have different background noise distributions, leading to discrepancies, since the current that could be considered noise in a pixel could very well be considered an event on the next pixel.

The hardware limitations of the VeloPix makes it so we can only apply a global threshold, meaning that all pixels in a grid need to have the same threshold. To account for this we have to equalize the noise from all pixels into a similar level, this way we can safely implement our threshold and correctly distinguish events from background noise. This process is called equalization and will be the main concern of our research.

The equalization process is based of trimming all of our pixels into different trim levels. What trimming means is to add an extra current to all pixels, categorizing our pixels in 16 different trim levels depending on how much additional current we're providing to them.

Before equalization can take place, we need to know the noise value for each trim level. If we do 16 measurements setting all pixels to the different trim levels then we can just adjust every pixel to the trim in which it's noise comes as close as possible to the threshold. The problem with measuring all trim levels is that it is very time consuming. The current solution has been to measure the first and last trim level to then predict the other level's noise by linear interpolation.

In another thesis by Max Vos [2] he investigated how the noise from the trim levels don't really follow a linear trend but instead they follow a third degree polynomial. This has put in question the current equalization process and motivated this research. This research will explore how inaccurate is the current equalization working by comparing it to the ideal equalization process. We'll touch more on this in section 3, but by using scans for all trim levels we can determine the ideal trim for each pixel instead of needing to interpolate. We can compare this to the current equalization process to determine how inaccurate it actually is.

2 The detector

2.1 LHCb

The LHCb is a particle physics experiment at the Large Hadron Collider (LHC) at CERN (European Organization for Nuclear Research) that has been designed and built to make precise measurements of CP violation. [5, 10]

CP violation, is the violation of the combined conservation laws associated with a charge conjugation (C) and parity (P) by the weak force. Charge conjugation is an operation that transforms particles into anti particles. Parity is the reflection through the origin of the space coordinates of a particle, so the three space dimensions x,y and z become, respectively, -x,-y and -z. Parity conservation means that left and right and up and down are indistinguishable in this system. [4]

It studies principally the properties and decays of heavy particles that contain beauty and/or charm quarks created in proton-proton collisions at the LHC. Particles containing b quarks are of particular interest, since large CP violation is expected in certain of their decays. In Figure 1 we've included a schematic of the LHCb experiment.
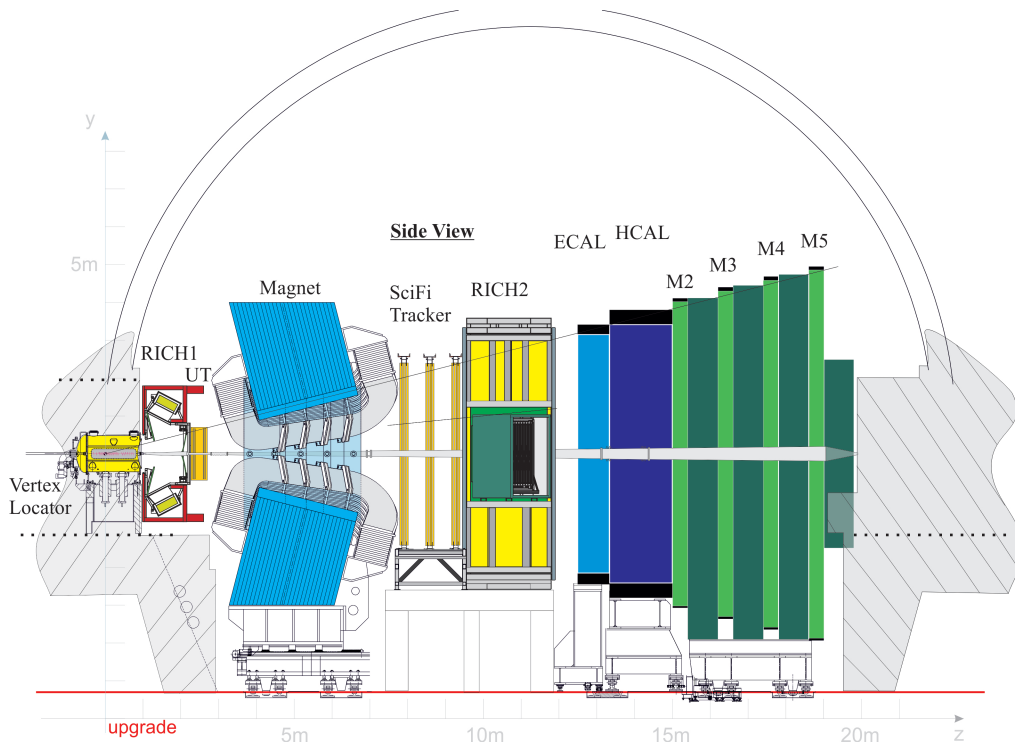

Figure 1: A schematic of the LHCb experiment [10]

The LHCb got upgraded in 2018, getting the front-end electronics of all sub-detectors replaced, included the Vertex Locator (VELO) detector readout and sensors chip. The Vertex Locator is the main tracking component of the LHCb detector.

2.2 VELO

The only point where the beams collide, and particles containing b and anti-b quarks are created, is inside the VELO sub-detector. The VELO's job is to measure the distance between the point where protons collide and the point where the B particles decay and to determine the collision point of the particles. The VELO can also detect the constituent particles of the B mesons after the decay has occured.

The upgraded VELO is a lightweight pixel detector operating in close proximity to the LHC beams. It has 624 ASICs (application specific integrated circuits). The ASICs are compiled into modules, each of the 52 modules has 4 sensors and each sensor is comprised of 3 VeloPIX ASICs each. This is important because later on we'll be analyzing different data points to study the noise of different ACICs we'll have them named based on their module, sensor, etc. So, for example, a data set of ours is labeled Module0-VP0-0. [1]

A schematic of the vertex locator set up is seen in Figure 2. Each yellow block is a module, made out of 4 sensors (two at the front and two at the back, which are not visible).
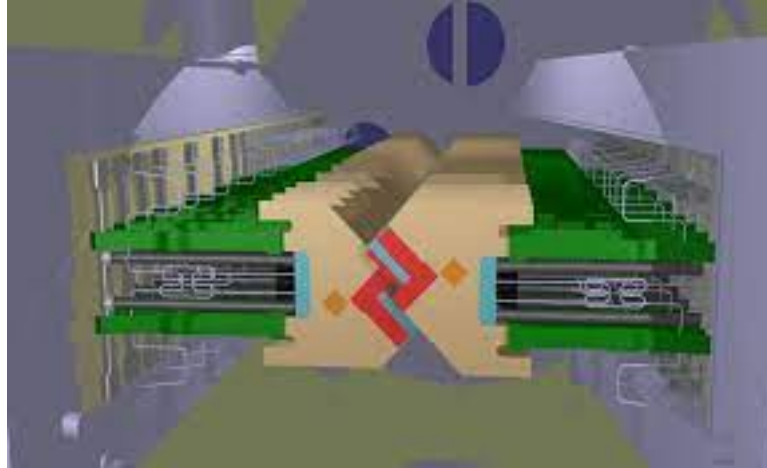

Figure 2: Schematic image of the VELOPix detector.

Each ASIC has a total of 256x256 pixels, with each pixel with a surface area of 55x55 μm^2.

Whenever a particle goes through a pixel a current gets emitted to mark this event, by looking all the different pixels activated we can get a trajectory of our particles and their collision point (the vertex). In Figure 3 we can see a schematic of this.
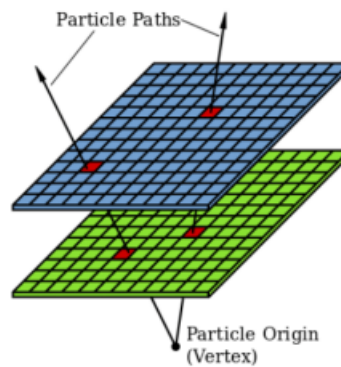

Figure 3: Schematic of the pixels in VeloPix and how they reconstruct particles trajectories

The current from the sensor is converted into the VeloPix into a digital signal. This means that in all subsequent plots and analysis, current will be given in an arbitrary unit called DAC.

As we've explained in the introduction, the main challenge is to equalize the noise of all pixels to succesfully set a global current threshold for all pixels.

3   Equalization process

We would want the detector to only show current whenever a particle goes through, since then we would be able to distinguish easily whenever the event happens. But every pixel has an associated background current that we'll call noise that is present at all times. This background noise is mainly due to thermal leakage. Thermal leakage is the effect where current charges vibrate within an electrical conductor.

The main issue to solve is that each pixel has a different sensitivity to background noise from manufacturing imperfections and since we can only set a global current threshold for all pixels in one of the ASICs we need to equalize all the pixel noise on it.

This is done by trimming our pixels into different trim levels. Trimming is a process in which we apply a baseline current to our pixels to archieve a uniform value of current noise for all pixels (this value is called target). The different trim levels (which are 16, labelled by hexadecimal values) equate to different baseline currents being applied to each pixels.

3.1   The current equalization process

The current equalization process can be summarized in the following steps for each pixel:

1. Two measurements for the noise mean of all pixels in the ASICs is performed. One with all the pixels set to the first trim and the other with all the pixels set to the last trim.

2. Taken a noise average of both measurements we calculate our target as an average of both means.

$$target = \frac{mean\_trim0 + mean\_trimF}{2}$$

3. We linearly interpolate the noise for trim0 to predict the trim that would give us a closer value to the target.

In Figure 4 we can see an example of noise distribution for a grid once it's all equalized.

We can see how the noise distribution gives a block distribution that is centered around the target as we were expecting. There are a couple of conditions that go into ignoring some pixels from our analysis, this is called pixel masking.

3.2   Pixel Masking

Since our detector is not perfect, there are a couple of mishaps that could happen when taking current measurements of our pixels. We'll try to mitigate the effects of these imperfections by masking pixels so that they don't affect the equalization process. The conditions to get our pixels masked are the following:

- A) Both the measurements from trim 0 and trim F are null. This means we have no data and cannot count in any information of this pixel.

- B) The measurement from trim 0 is null.

- C) The measurement from trim F is null. Both of these end up in the same practical reason as the first one but we'll put them in different categories for sake of comparison.

- D) The target is lower than the measurement done at trim 0 or the target is higher than the measurement done at trim F. This means that this pixel does not have the target in the interval between all possible 16 trims and we'll remove it.
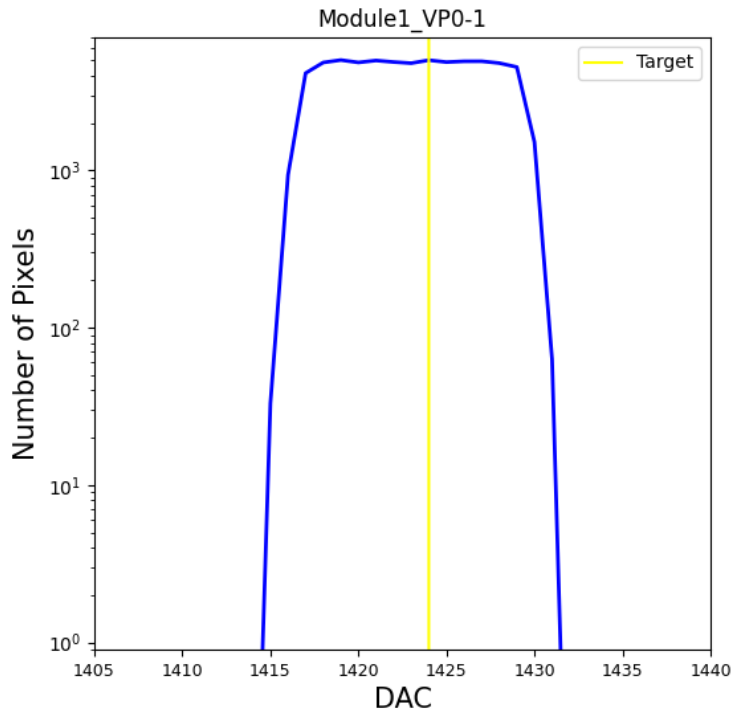
Figure 4: Noise distribution for all pixels in an ASIC

- E) The prediction falls to far from the target, in this case we'll apply this if the difference between the prediction and the target is more than 25.
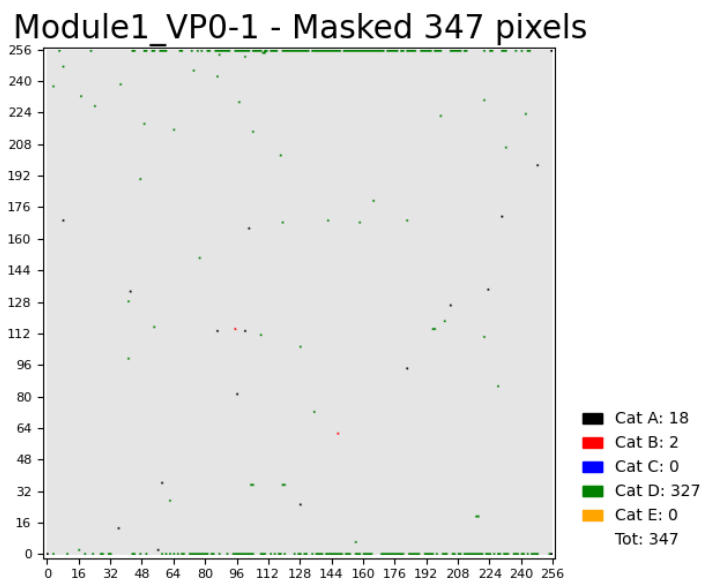
A map of the masking can be seen in Figure 5.

Figure 5: Colour map of all pixels categorized by the reason of their masking

As we can appreciate, most of these pixels get masked because the target is not between the lowest and the highest current value. And this happens mainly in the edges of the ASIC.

3.3 The ideal equalization process

For the sake of analysing how good of a job the current equalization process is doing we'll compare it to the ideal equalization process. We first need to stress why this isn't being used currently. A measurement on all trim levels is very time consuming, in the current procedure only two trim levels are being measured and the whole equalization process is being done based on them. This has proven to produce a good enough equalization for quality control during detector construction, but in another bachelor thesis by Max Vos [2], it was showed that trim levels follow a third degree polynomial instead of a linear dependence. Compared to a linear interpolation, which needs two inputs to interpolate, if we want to correctly interpolate a third degree polynomial we need four inputs. This means that we would need to measure all pixels in four different trim levels to implement it.

Based on this it begs the question, is it really worth it to start measuring four trim levels to make our equalization more precise? We'll bring this to an even more extreme case, we'll test the ideal equalization where we measure all trim levels and try to equalize on that. This will change the whole equalization process, so it now follows the following steps:

1. We measure a full scan of all pixels on the grid set to all the possible trim levels (from 0 to F).

2. We calculate the target as the average between the noise mean of trim 0 and the noise mean of trim F.*

3. For each pixel we find the trim level measurement that comes closest to the target. This will be the trim that we'll set for the pixel and the noise prediction will be the noise current measured at that trim level.

*Here we have calculated the target taking only the first and last levels. This has been done this way because choosing the correct target has been another topic that has proven to be more complicated that it'd seem. So, for sake to make our comparison easier we'll just calculate the target the same way we were doing it before.

For step 3 we can illustrate this as visible from Figure 6. In this figure we have plotted the noise measurement for all different trims along with the target for this equalization. In this specific example the pixel would get assigned the trim 4, since this is the reading that gets the closest to the target.
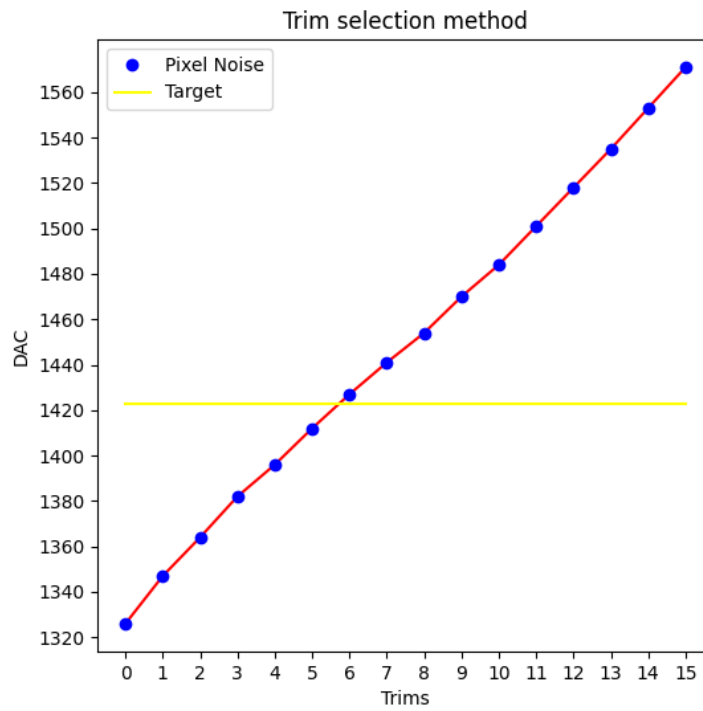
Figure 6: Example of the noise readings for a pixel in all trims along with the target

With this we'll compare the noise distribution and the trim assignation between these two approaches.

4 Results

4.1 Masking with the new process

The first thing we'll have a look at is how the masking of pixels has changed compared to our original equalization process. Where we have slightly changed our categories to be:

- A) Both the measurements from trim 0 and trim F are null. This means we have no data and cannot count in any information of this pixel.

- B) The target is lower than the measurement done at trim 0 or the target is higher than the measurement done at trim F. This means that this pixel does not have the target in the interval between all possible 16 trims and we'll remove it.

- C) The prediction falls to far from the target, in this case we'll apply this if the difference between the prediction and the target is more than 25.

We can see this masking in Figure 7. Comparing it to Figure 5, we can observe how most pixels are still being masked because the target falls out of the bounds of the pixel resolution. With this equalization method more pixels are being masked overall but overall we can appreciate how this is not significant enough to affect our interpretation of the results.
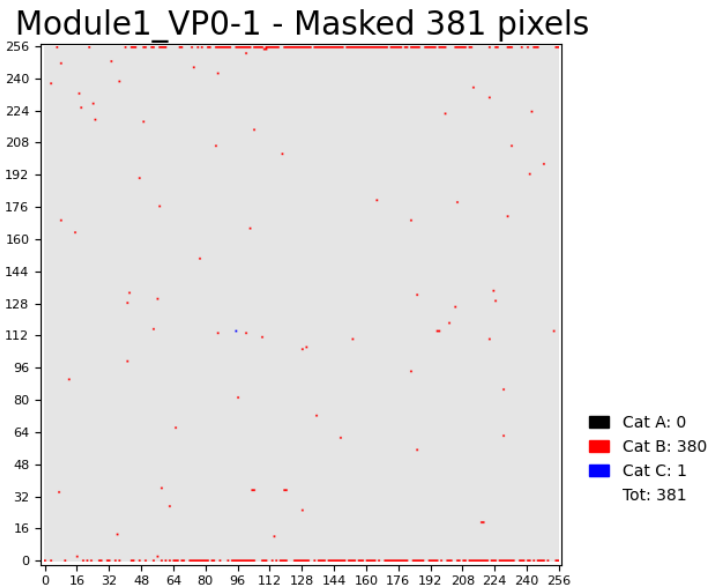

Figure 7: Masking of the pixels taking into account our new equalization process

4.2 Noise distributions

Next, and more importantly we'll compare how the noise distribution for the ASIC has changed after introducing this equalization. For this we'll use the same data set. A plot of this distribution can be seen along with the data from the original analysis in Figure 8.
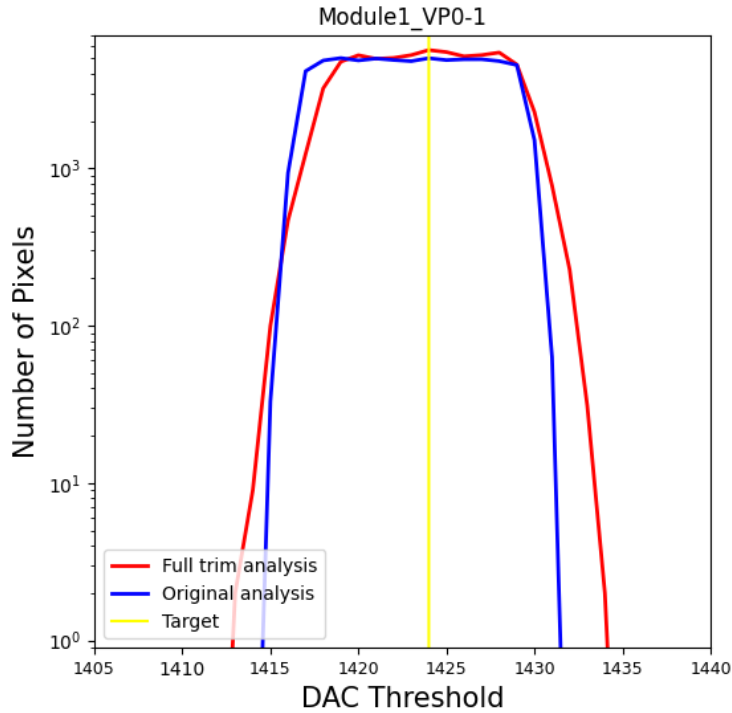
Figure 8: Comparison between the noise distribution of the original equalization and the ideal equalization using all trim levels

This distribution is similar to the Figure 4, we see a distribution resembling a block function centred around the target. The main difference between these two plots is that our new analysis has a wider distribution, meaning that there is a higher amount of pixels that steer away from the pixel. This means that the current equalization is failing to predict this wideness noise distribution. This also can be understood better by thinking that the current equalization is setting trims for the pixels and predicting the noise but the actual noise is given by the full trim analysis.

To make sure that these results are not isolated we'll make the same analysis for different data sets. For this we'll use 5 different data sets from different ASICs and study the difference between the original analysis and the analysis using all trims. Where we'll assume that results from a scan are going to be similar to a future scan, meaning that results are reproducible in different times, this is being investigated by a colleague Borja Torrijos[8]. The plots for these comparisons can be seen in Figure 9.

We can see that for all of our data sets the behaviour is similar to that of Figure 8, where our new analysis is wider than the current equalization process. So again, we can safely say that the current equalization is failing to predict a wider noise distribution, which could be a problem if that difference could interfere with the threshold.
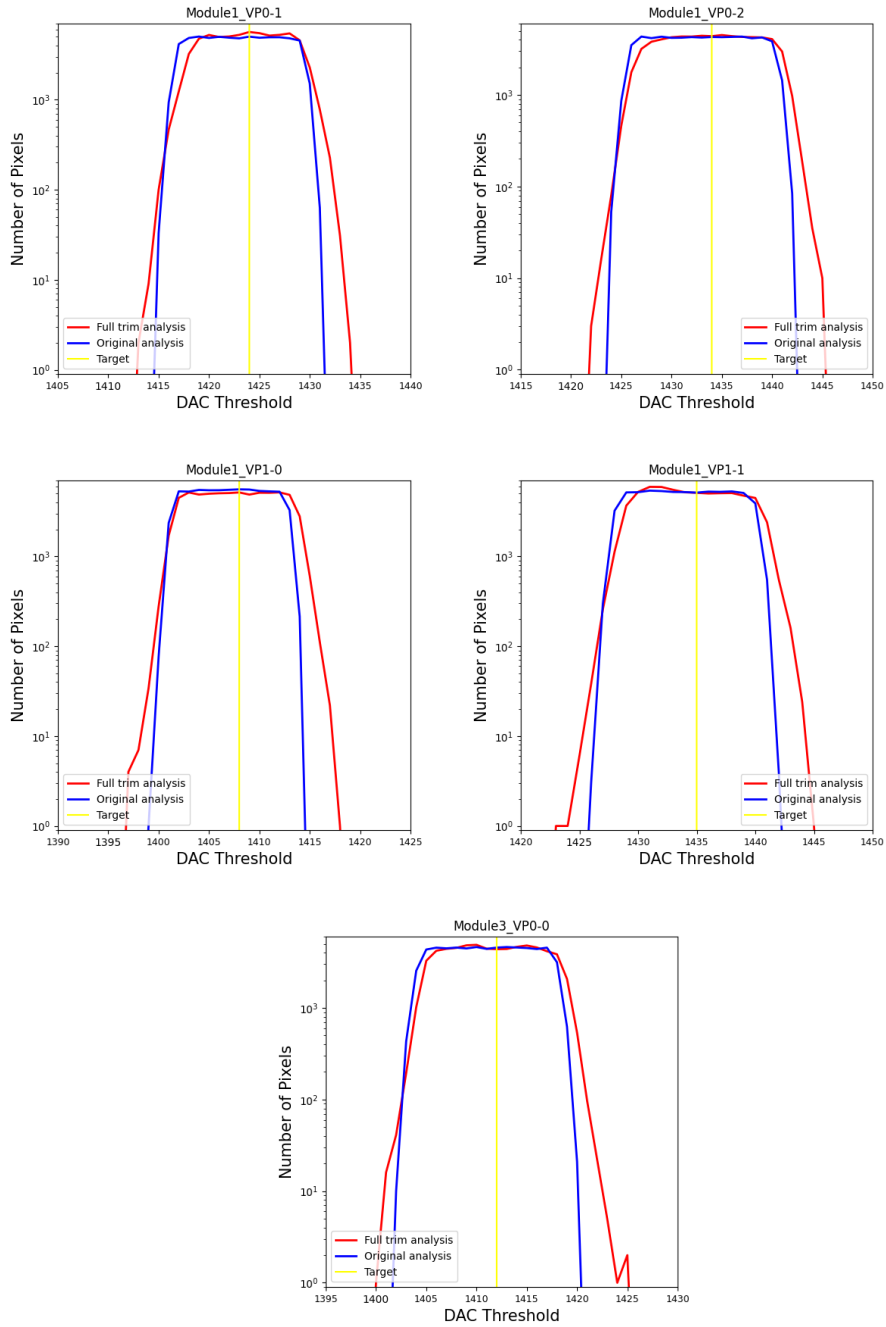
Figure 9: Comparison of the noise distribution between the original equalization and the equalization using all the trims

4.3 Trim differences

Next thing we'll have a look at is how trims are assigned to the pixels. We'll display for each pixel the difference between the trim assigned with the full analysis and the trim assigned with the linear interpolation. This map can be seen in the Figure 10.
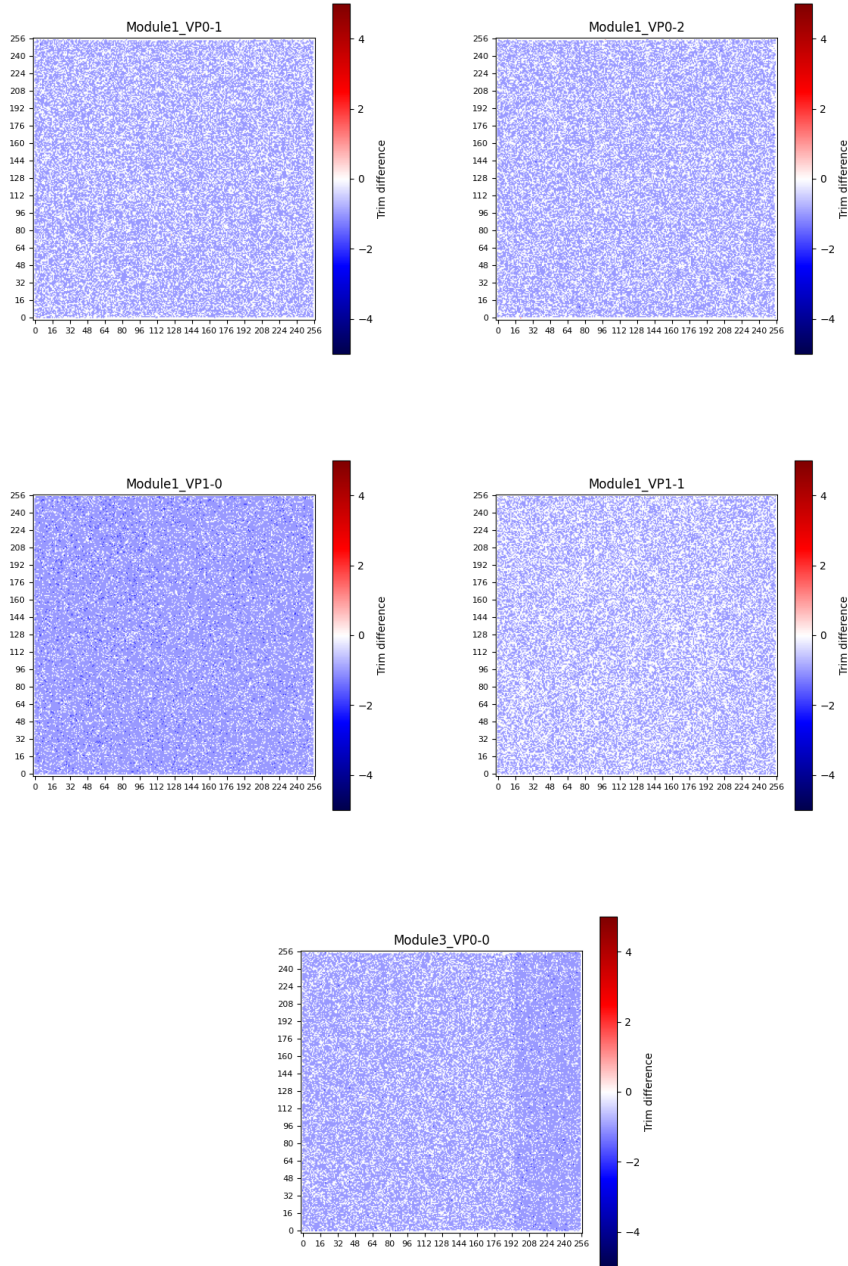
Figure 10: Colour maps of the trim differences between the original analysis and the analysis of all the trims. Trim difference = Full Trims analysis - Original analysis

From Figure 10 we can see that the trims are being over calculated in the original approach, with most pixels being assigned a higher trim value of what they should be. We can also appreciate the data set for Module3_VP0-0 has a slight gradient of trim difference for the right most pixels, meaning that the right most pixels tend to have their trim over calculated by a higher amount compared to the rest of them. This is odd, since it only occurs for one of the five ASICs. We'll assume to be a slight hardware feature that makes the pixels in this area harder to linearly interpolate.

We can also visualize this information by inputting the trim differences into an histogram, these series of histograms are displayed in Figure 11.
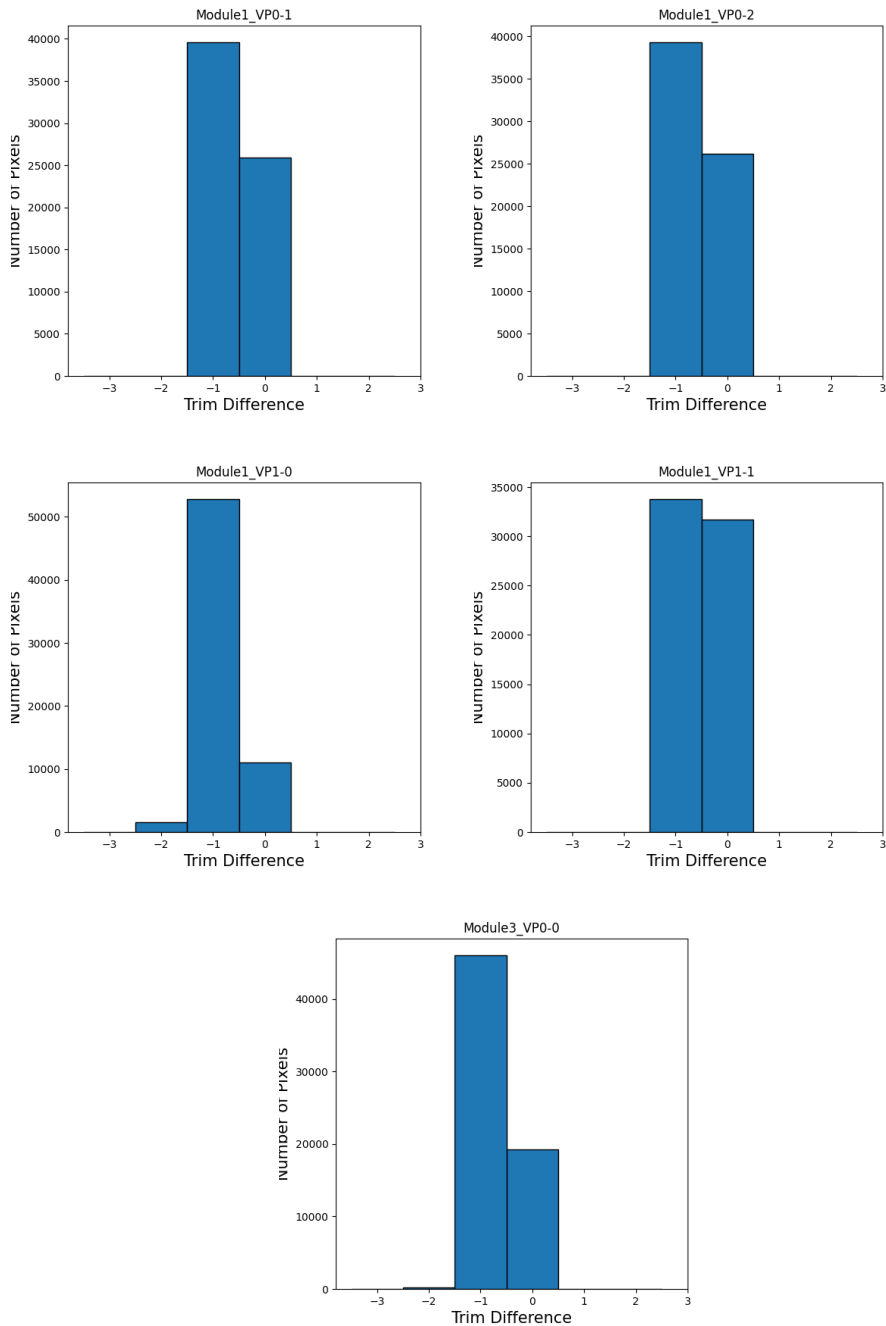
Figure 11: Histograms of the trim differences between the original analysis and the analysis of all the trims. Trim difference = Full Trims analysis - Original analysis

These histograms confirm again the information we had on the maps (in Figure 10), most pixels are being over estimated and the second most likely scenario is that pixels are being attributed the same trim values. To reiterate, we would expect that al of the pixels have the same trim level attributed to them but we can see how the current equalization routine is over calculating trims in general.

4.4   Noise differences

One last aspect we'll study about this difference is the noise difference from pixel to pixel. This means that, for every pixel we'll take the predicted noise given by the full analysis and the noise prediction given by the linear interpolation. We are going to calculate the difference between both

for all pixels and display them. As we have seen in Figure 9, the noise distribution with both approaches is similar, with the new equalization turning into a wider distribution for the noise. We'll plot a maps and histograms as we have done with the trim differences to see if the noise difference has any tendency. The map of the noise differences can be seen in Figure 12.
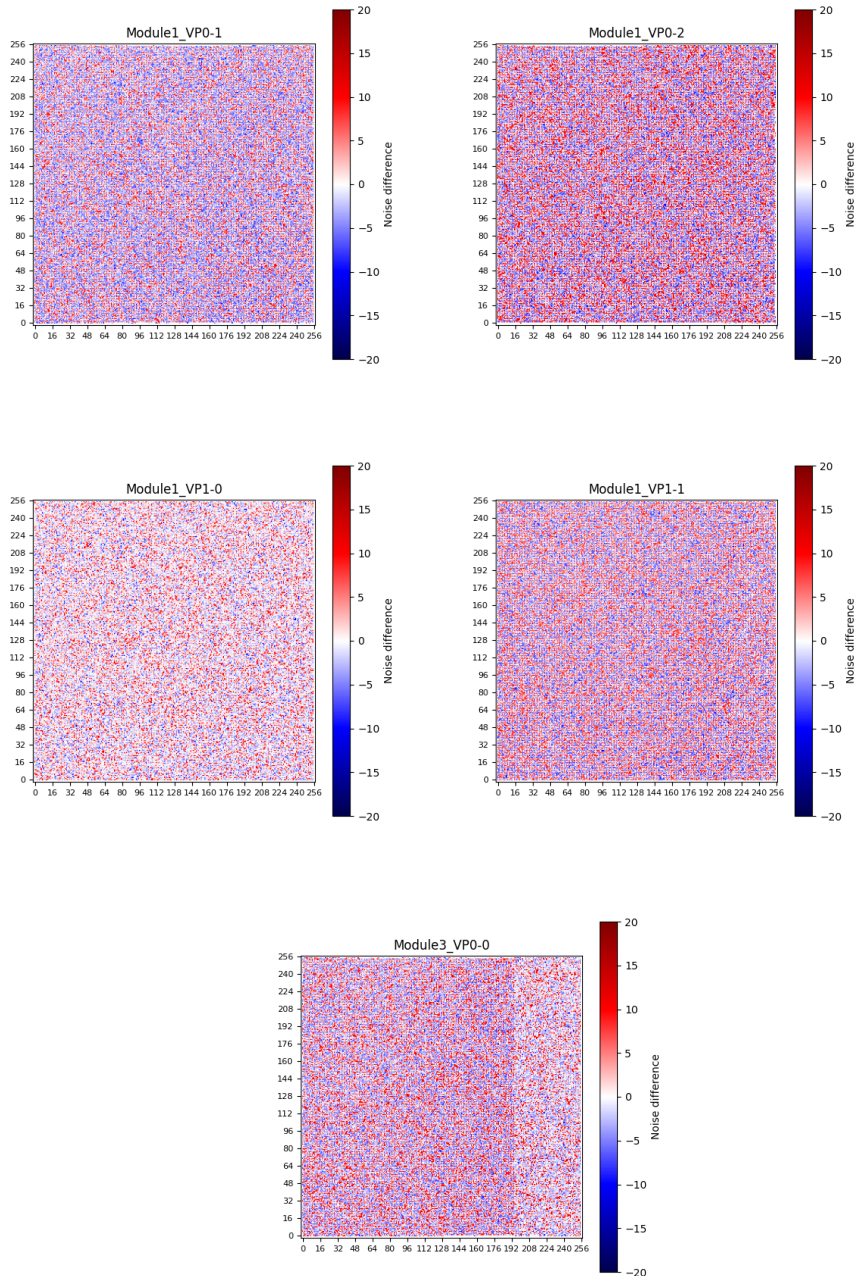

Figure 12: Colour maps of the noise differences between the original analysis and the analysis of all the trims. Noise difference = Full Trims noise - Original noise

We can appreciate from Figure 12 how the noise difference between pixels is random, where some pixels are being over calculated compared to the ideal equalization while other pixels ate being under calculated. This is in line with the noise distributions in Figure 9, where both distributions are roughly symmetric so it is expected that noise differences from pixel to pixel should be close to zero or should average out to be zero in total.

Another aspect we can appreciate from these maps is that the ASIC Module3_VP0-0 has again the same gradient as we've seen with the trim difference in Figure 10. Where the pixels that have higher over estimated trim values end up being the pixels with a lower noise difference.

Again, we can also visualize this information in histograms to have a better idea of the overall trend for the noise difference. These histograms can be seen in Figure 13.
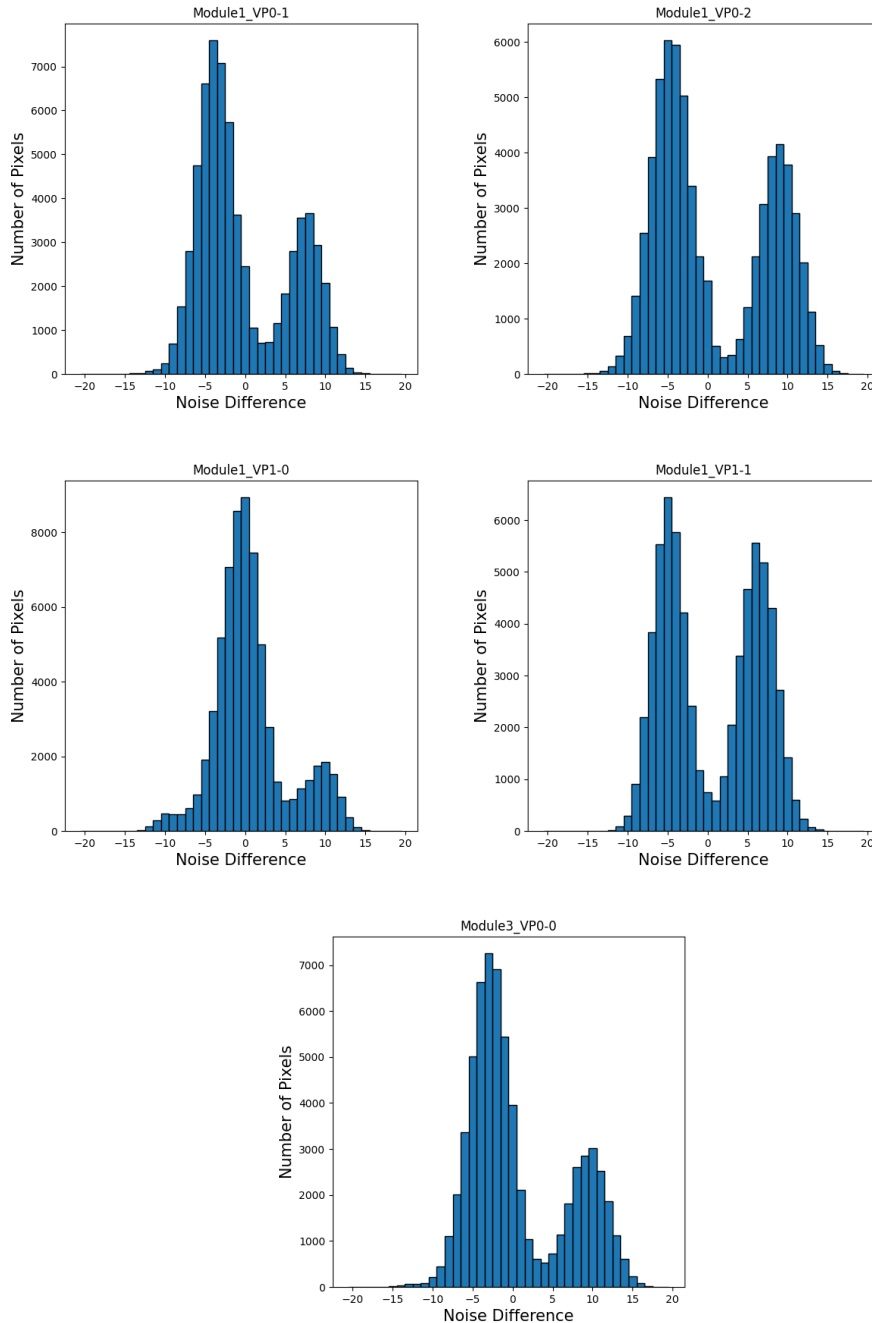

Figure 13: Histograms of the noise differences between the original analysis and the analysis of all the trims. Noise difference = Full Trims noise - Original noise

From this plots we can extract some extra information, but it looks inadvisable to draw closing conclusions out of it since we don't see a clear repeating pattern. The most notable feature from our histograms in Figure 13 is how the noise difference valleys, meaning that there is a very low chance for the pixel to have the same noise values for both approaches. The next pattern we notice

is how there's a higher tendency for the current equalization process to give lower noise than our equalization process. This can be seen by the left peak being higher than the right peak in Figure 13. Another notable difference is that the ASIC that seems to show the least amount of noise difference on average (Module1_VP1-0) is the same ASIC that shows the highest trim difference (from Figure 10 or 11). This again is confusing since we would expect that the pixels that get assigned the same trim as the ideal equalization process should get the same amount of noise, but this apparently isn't the case.

5   Conclusion

The equalization process is a very important process under the analysis of our data. Without proper equalization, the threshold is set incorrectly and this can lead to problems; for example, some peaks of the background noise could be misunderstood as an event if the threshold is set too low but an event could be misunderstood as background noise if the threshold is set too high.

We've discussed how the current equalization process works by using the 2 outer trims and linearly interpolate between them to predict the ideal trim for each pixel. In previous bachelor thesis by Senne Baker [3] and Max Vos it has been discussed how this equalization process is on a flawed basis. Max Vos showed that the trim level noise means follow a third degree polynomial instead of a linear tendency. Senne Baker investigated in using two arbitrary trims instead of performing the linear interpolation with the two outer trims, this went to show how the use of trim 1 ends in better results than the use of trim 0.

Both of these investigations have put in question the equalization process and ow precise it really has been. In these thesis we've compared it to the ideal idea of equalization, where we have data readings of noise for every pixel set at every different trim level. This way we can determine that the ideal trim for a pixel is the trim that results in a noise reading closest to the target. Comparing this approach to the equalization process currently used will give us an idea at precise the latter is.

From this analysis we have seen a couple of features of the current equalization process. First of all, we've seen that the currently used noise distribution fails to predict a correct noise distribution, with the current one being the thinner distribution. This, as we've said, could end up in a big problem if the threshold is set closely to the tail end of the distribution and that slight miscalculation could start confusing noise as events. This is arguably the most important result from the equalization process, and it should be taken into account to make sure we're not making this mistake.

We've also studied how this process assigns trims to our pixels and we have demonstrated how it is overestimating the trim all pixels should have, this is supported by Figure 10 and 11. The implications of this are also unsure, specially considering what we've just discussed of the current process having such a thin distribution. But it is something we should keep in mind for future investigations.

In conclusion, with a lack of further evidence, the current equalization process has proven to be as precise as we need it to be but the miscalculation of the distribution should be something to keep in mind for further investigations. There are a some anomalies about it that require a further look into and we would benefit from gaining an understanding from them.

6 References

[1] Poikela, T. & Gaspari, M. & Plosila, Juha & Westerlund, Tomi & Ballabriga, Rafael & Buytaert, Jan & Campbell, Michael & Llopart, X. & Wyllie, K. & Gromov, V. & van Beuzekom, Martin & Zivkovic, V.. (2015). VeloPix: the pixel ASIC for the LHCb upgrade. Journal of Instrumentation. 10. C01057-C01057. 10.1088/1748-0221/10/01/C01057.

[2] M. Vos. Predicting the noise position for unmeasured trim levels. June 2021. Bachelor thesis.

[3] S. Baker. The VeloPix upgrade for the LHCb experiment. June 2021. Bachelor thesis.

[4] https://www.britannica.com/science/CP-violation

[5] Jeroen J. van Hunen. The LHCb tracking system. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 572(1):149–153, 2007

[6] LHCb Collaboration. LHCb VELO Upgrade Technical Design Report. Technical report, Nov 2013

[7] L. Diamante. Transforming LHCb: What's in store for the next two years? Feb 2019.

[8] B. Torrijos. Reproducibility of LHCb's VeloPix noise equalisation. July 2021. Bachelor thesis

[9] LHC Collaboration. Facts and figures about the LHC.

[10] LHCb Collaboration. The LHCb detector at the LHC. August 2008