University of Groningen

Bachelor Thesis

Reproducibility of LHCb's VeloPix noise equalisation

Author: Borja Torrijos García

First supervisor: Dr. K.A.M. De Bruyn
Second supervisor: prof. dr. K.H.K.J. Jungmann

Contents

1 Introduction

We humans have an intrinsic curiosity that forces us to keep pushing questions to its limits. We can already see this in our childhoods, where we cannot escape the asking loop of: "and why is that?" even when we keep getting the answers. As physicists, we essentially do the same, just that usually with more transcendent questions, being one of the main and most naive ones what are things made of, what are the *building blocks* of matter.

In order to answer it we have built impressive things, not only experimental but theoretical. In the theoretical aspect, we started from really basic concepts such as the four elements, earth, fire, air and water, plus the ether, and we have gone all the way through bringing complex models, being the Standard Model, SM, the most established one nowadays. In this model we work with the most elementary particles known so far: quarks and leptons, with its corresponding six different species or flavours each; gauge bosons and the Higgs, and we try to describe their properties and how they interact with each other in order to create everything we have in our Universe. Note that for every particle there exist a counterpart, its antiparticle, with the same mass but opposite quantum numbers (opposite charge, for example).

The SM is already well understood and have been successfully tested since its current formulation in the 1970s along with the acceptance of quarks' existence [1]. Nonetheless, we now have consistent evidence that it is not the final theory of everything, as we still have unexplained observational facts we have to face within this model. Listing some of the main ones: We have gravity, which at least at our scale leads to significant interactions, is not included; it also explains the behaviour of just around the 5% of the Universe's energy content, being the rest attributed to dark matter and dark energy, where no established particles or interactions candidates have been found yet. Also, it predicts massless neutrinos, although it has been proved that they indeed have some mass as they oscillate between the three flavours. And finally, it does not have a consistent mechanism that could explain the matter-antimatter asymmetry in the Universe, which is made out mostly of matter while in principle there is not a trivial reason behind it.

Encouraged with those unsolved questions we move to the experimental aspect, and in order to find answers at this elementary level we need to be able to create and detect the mentioned elementary particles. We can think of it with the help of an analogy: In an effort to understand how a clock internally works if we see one for the first time, it would be hard to do it by just looking at how time passes on it. It is more useful (at least, as physicists) if we break it into pieces and reconstruct it. That way, we would understand which parts are involved and how they work with each other. That is what we do in particle colliders, we make two particle beams crash, creating a high-energy environment in an ephemeral time from which the elementary particles can emerge. After they are created, we need to detect and recognise them. We are able to do that in the detectors, where they interact with its components, letting us reconstruct their trajectories and identify which particle they were.

One of the most ambitious particle colliders is located at the European Organization for Nuclear Research, CERN, in Geneva, called the Large Hadron Collider, LHC. In particular, within the LHC we have one specialized experiment called LHCb, which focuses on the detection of one of the quarks' flavours, the b quark, also called bottom or beauty. This experiment aims specifically to answer the matter-antimatter asymmetry previously mentioned, looking for a possible slight asymmetry between particles and anti-particles in, for example, decay rates differences within some mesons, particles that contain quark-antiquark pairs, containing b quarks, such as the B mesons.

The whole accelerator has gone through a three years shutdown in order to make an upgrade, the goal is to work with even higher energies and achieve more collisions, and the LHCb is correspondingly also upgraded to allow a readout at 40 Mhz [2], [3], in order to do a complete read of the events and do a software-based selection of the collisions that will be stored for offline analysis. The shutdown has concluded last April 22, when the first collision since the start of the update period has successfully happened. Specifically, this thesis is focused on the very first layer of the LHCb detector, the vertex locator, VELO, which can be visualized in Figure 1.

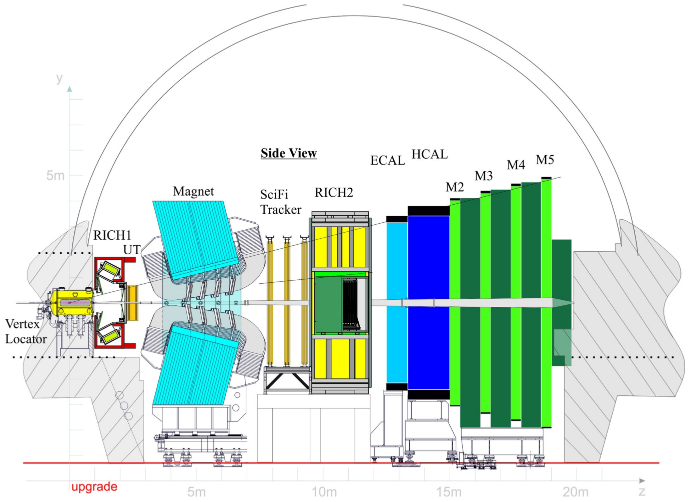

Figure 1: Schematic drawing of LHCb's detector made by CERN [2].

The VELO works detecting the particles that goes through it so we then reconstruct its trajectories, in particular the primary and secondary vertices (the point where the original particles collide and its first decay, respectively) as it is the first layer. Its new version has changed the previous silicon strips for hybrid pixel sensors [4] and has improved the technology overall, providing the quick data readout of all the events mentioned before and allowing a software-based trigger system. Also, the pixels reduce the occupancy, making it easier to reconstruct the paths and reduce the background contribution from collisions we are not interested in. A diagram of the upgraded VELO can be found in Figure 2.

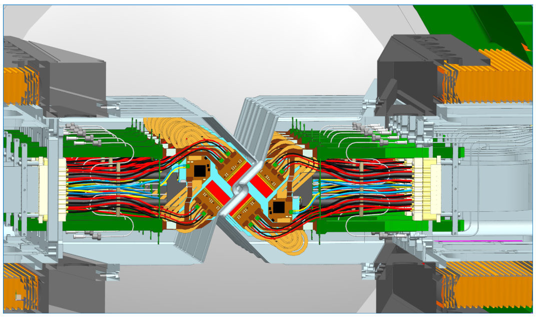

Figure 2: Schematic drawing of the upgraded VELO [4].

Apart from this background contribution, we have another phenomenon to deal with, the noise of the pixels, which is the signal they give when no particles are going through. The goal is then to determine the magnitude of that noise so we can discern which signals were indeed particles; To generalise this process to all the pixels consistently, as we will see in the next section, a process called equalisation is applied.

In order to apply the equalisation we first need to measure the noise of the pixels, which is made through a scan where a noise distribution for each of them is obtained and its noise mean is calculated; for simplicity, we will refer to this pixel noise mean as the noise position. The objective of this thesis is to study the equalisation discrepancies between different scans, which brings two major benefits. The main one is helping other analysis that are being made about other effects happening at the VELO, such as the noise temperature dependence [5] and more accurate treatment of the pixels noise [6], which have the open question of whether the systematic uncertainties due to this reproducibility could be significant in their results. And the second advantage is, if we get a consistent reproducibility, to be able to reduce the amount of scans (and their corresponding equalisation processes) needed over time, as it would mean that we just need to do them mainly to keep track of radiation damage, which in principle does not require scans that often, although it would need to be studied in more depth as being the first layer also means the radiation environment is much harsher.

2 Noise equalisation

In the VELO there are 52 modules, each with 12 ASICs, also referred to as VeloPix, and each VeloPix containing a grid of 256 × 256 pixels, so a total of more than 40 million pixels [7]. All of them cannot behave the exact same way, as production imperfections, location with respect to the electronics, radiation damage over time and some others factors have to be taken into account. Hence, every pixel have a different noise position, which means setting a single threshold, a value below which electronic signals are considered noise, will not be the best option.

In order to solve the noise differences, we set, apart from the threshold, individual trims for each of the pixels with sixteen possible levels, 0-15, which changes their noise position. For example, if we set a lower trim, the noise distribution higher than that level obtained in the scan will not be considered, so it will not contribute to the noise mean, decreasing it. This allows us trying to get every pixel to the same noise position so we can finally set a consistent threshold for all of them. But as the trim levels are discrete and the noise continuous we cannot get all the pixels to the exact same noise position; we need to calculate which trim fits each pixel the best to get them to a certain noise target, where we want all of them to be.

This process is the so-called equalisation, and we can see an example in the Figure 3, where the electronic signal in the x-axis is measured in DAC, an arbitrary current unit with 1 DAC defined as what VeloPix can measure. In this plot, the red and blue lines indicate the noise distribution (the noise position of every pixel) with a trim level of 0 and 15 (also referred as F), respectively, and the black one being the noise distribution after the equalisation is applied. The ideal scenario for the black line would be a Delta function, which would mean that after the equalisation all the pixels have the exact same noise position; but as mentioned the discreteness of the trim levels makes it impossible, giving a certain wideness around the target.

The default approach of calculating the target until now have been averaging the mean of the noise positions in the 0 and 15 trim distributions. Then we calculate the best trim level for each pixel so they come as close as possible to that target. This is what was applied in the Figure 3, and although more sophisticated ways are being studied with the use of more trim values [8], [9], that is not the research question addressed in this thesis. As we need to compare between equalisations, we just want to make sure that both were made through the same process. It is worth mentioning, as this would be a parameter that we can also compare between the scans, that there are pixels that cannot be equalised (such as dead pixels), the so-called masked pixels.

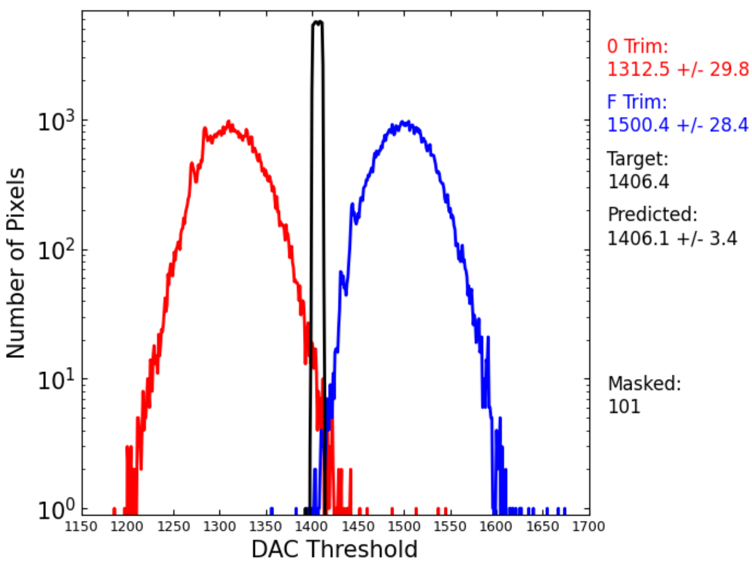

Figure 3: Equalisation of a VeloPix scan.

3 Reproducibility

To test the reproducibility of the equalisations between scans we will compare three parameters for a single ASIC: first the noise position of the pixels, then their best trim levels and how does this affect to the overall equalisation and finally the difference in masked pixels. In this section, the data sets that are going to be compared come from two scans that have been consecutively taken. In the next sections, we will extend this analysis to the different ASICs and scans separated in time.

3.1 Noise comparison

We first calculate the noise position differences of all the pixels and show it in both the whole grid and in a histogram-like plot, obtaining Figures 4 and 5. It has been calculated subtracting the first scan to the second one, so a positive difference means the second scan have a higher noise mean. We see in the grid that the differences go up to around 4 DAC, and in the histogram that most of the pixels are within a 2 DAC difference. This is a good first result taking into account we mentioned that 1 DAC is the resolution of the VeloPix. Particularly, in the grid we also notice that there is not an obvious pattern or clustered bigger differences, which would have meant that there is a more specific reason that affects certain areas and not just «random» fluctuations within pixels because of the reasons we discussed.

In Figure 5 the ideal scenario would be a Delta function centered at zero, meaning that there is no difference nor fluctuations between the scans. What we obtain is a Gaussian-like distribution instead, which is what we expect from random fluctuations. The key results are that, on average, there is no difference between the scans as the mean is zero, and the fluctuations, indicated by the width of the Gaussian, have a magnitude of 0.6 DAC, which is again a really good result as it is lower than the resolution. Hence, the differences we obtain are not significant, but we still need to check how they affect the equalisation result in the next section.

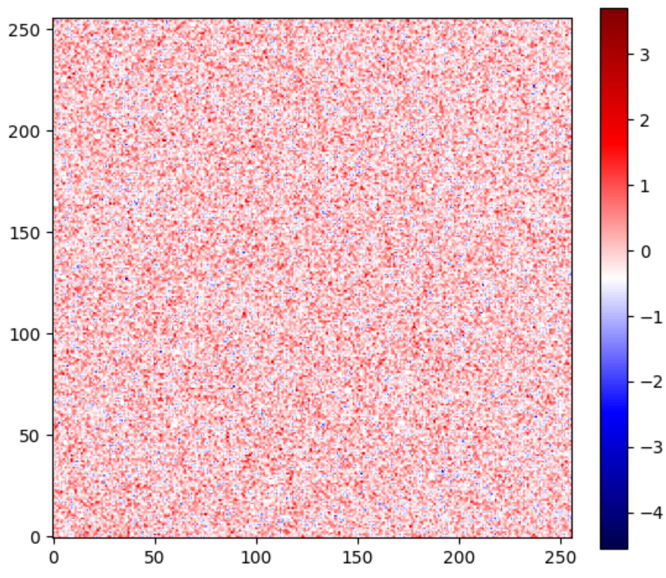

Figure 4: Noise mean difference between the scans for every pixel, plotted in a 256×256 grid, where the values are calculated in DAC.

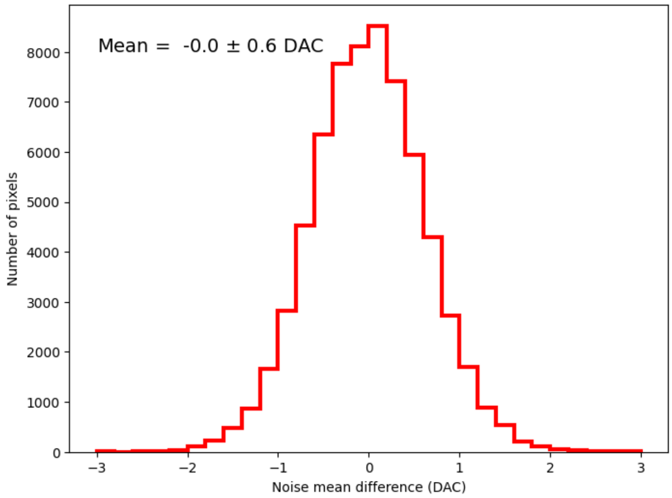

Figure 5: Noise mean difference between the scans for every pixel, plotted in a 1D histogram.

3.2 Trims comparison

Now we analyse how this noise fluctuations affect the calculated best trim levels, and which effect does this possible changes in the trims have on the equalisation. Comparing the grid of trim levels we obtain Figure 6:

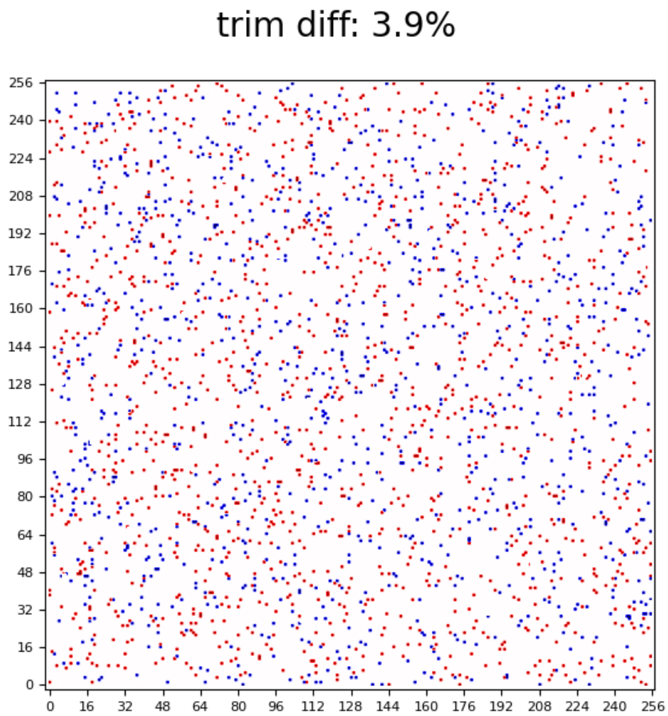

Figure 6: Difference of the best trim level to equalise each pixel between the scans, with red points being +1, blue -1, and white 0 difference.

Again, we cannot discern a clear pattern in the differences, with no apparent clusters and an arguably random scattering and amount of both positive and negative values. We also see that the highest differences are just of one trim level, with 3.9% of the pixels having changed. Now we need to understand whether both noise and trim changes affects significantly the equalisation result. In Figure 7 we plot the equalisation result for the second scan in black, with the blue and red lines indicating the pixels that have changed their trim by +1 and -1 respectively. Also, the first scan's equalisation result is plotted in green, although it is hard to discern as it is underneath the black one. This overlap indicates that the equalisations are indeed compatible.

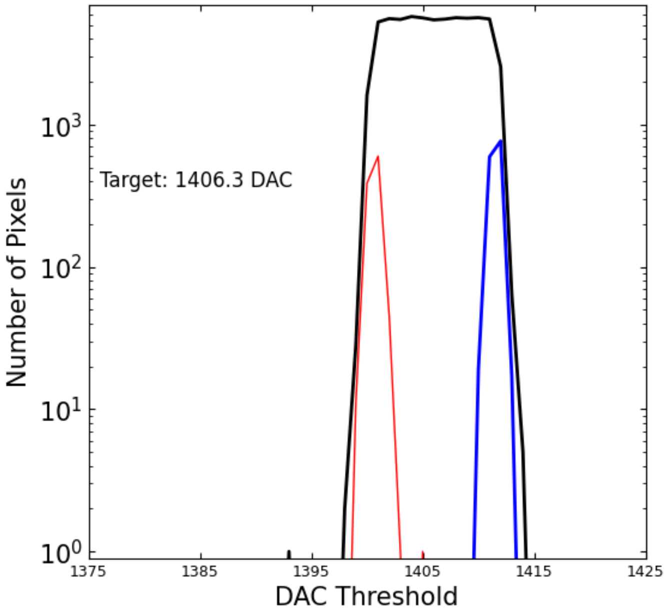

Figure 7: Black line: Equalisation result of the second scan, with the blue distribution being the pixels that increased their trim +1 and the red ones -1. Green: Equalisation result of the first scan.

Our hypothesis to explain why there are a significant amount of pixels that have changed their trims but it is not affecting the overall equalisation was that the ones that have changed were already around the middle of two possible trim levels so a small change in the noise could have swapped them from one trim to the other, giving a similar distance to the target. A diagram of this behaviour can be found in Figure 8.

We have confirmed this hypothesis with what we obtained in Figure 7 for the changed pixels, as the ones that increased by one are now grouped on the right, meaning that they were originally on the left border and so with a compatible distance to the target, as we discussed with the diagram. Also, the exact but opposite case occurs to the ones that decreased the trim by one, and as we have similar amount of both, the small discrepancies between the distances to the target due to the jumping cancels each other out, resulting in a compatible equalisation. It is worth noting that the y-axis, which counts the number of pixels, is in a logarithmic scale, so the peaks for the ones that changed their trims are narrower than what we can observe here, meaning that the jumping effect is even more pronounced.

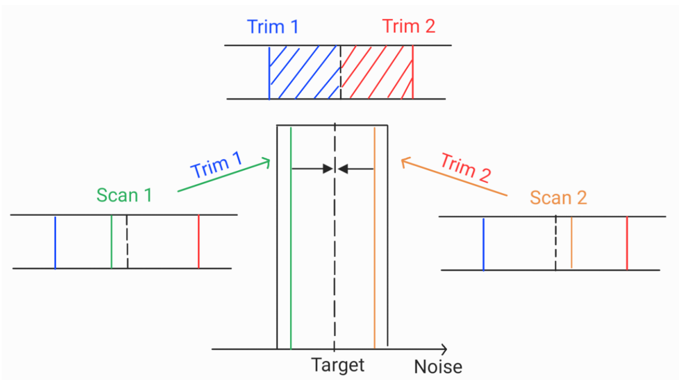

Figure 8: Jumping process diagram.

And finally, we also check whether is viable to equalise the second scan using the trim values calculated in the first, so we also achieve the second goal mentioned in the introduction of being able to «skip» some scans and do them less often just for radiation damage tracking. The result is shown in Figure 9, where we obtain an almost perfect overlap between the equalisations, also giving the same achieved target result.

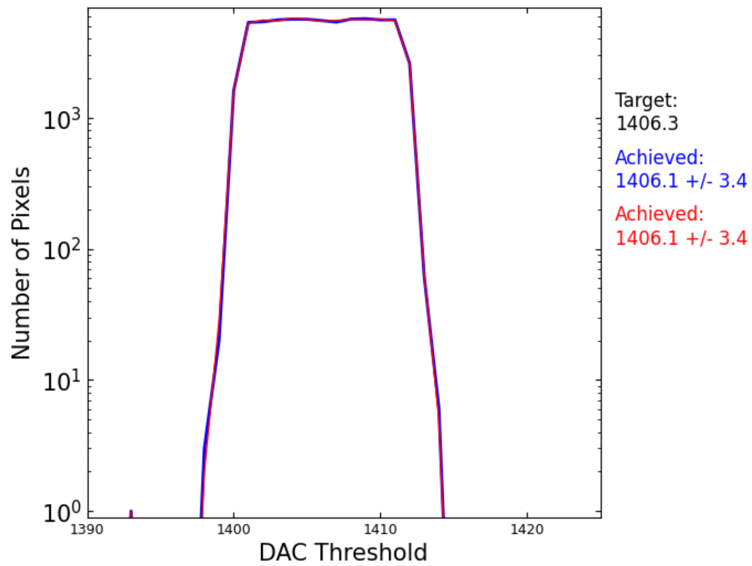

Figure 9: Equalisation result of the second scan, where the blue distribution uses the first scan equalisation's trim levels, and the red one uses its own equalisation.

3.3 Mask comparison

Finally, we do a similar analysis as with the trim levels but with the masked pixels, just to check that there are not significant changes here either, as the masked pixels do not appear in the noise analysis. Comparing the mask matrices of both scans to see which ones have changed we obtain Figure 10, and we see that just 2 pixels have changed, which is a negilible compared to the total number of pixels we have in one VeloPix.

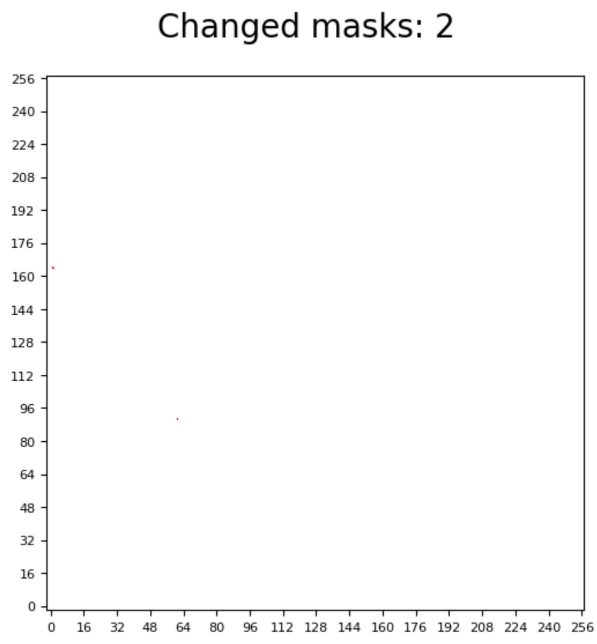

Figure 10: Masked pixels that have changed between the scans.

4 Extension to different ASICs

Now we extend the analysis for all the ASICs in the module we are working on, so we see if we keep obtaining good results. We created a table with the key values we discussed for the purpose of having an overall comparison between them, obtaining Table 1, where the ASICs labels indicate in which tail they are and its location on that tail, respectively. Also, the *Trim increase rate* column indicates, within the pixels that changed the trim, the percentage of the ones that have increased it.

Module	**Mean difference (DAC)**	**Trim change (%)**	**Trim increase rate (%)**	**Masks changed**
0-0	0.0 kk 0.6	3.9	57.2	2
0-1	0.1 kk 0.6	7.1	99.1	3
0-2	-0.1 kk 0.7	2.8	56.6	0
1-0	0.1 kk 0.7	3.3	41.7	1
1-1	0.1 kk 0.6	3.1	35.0	1
1-2	0.1 kk 0.6	3.3	47.8	2
2-0	-0.1 kk 0.7	3.5	61.8	3
2-1	0.0 kk 0.7	3.5	49.0	5
2-2	-0.1 kk 0.7	4.1	73.2	9
3-0	0.0 kk 0.6	3.5	60.3	4
3-1	0.0 kk 0.7	7.2	0.5	3
3-2	0.2 kk 0.6	3.3	71.3	3

Table 1: Comparison of the reproducibility key values between the ASICs of the N020 module.

From the table, we first see that the means are very similar, all compatible with zero and with almost the same uncertainty, so we get the same conclusion as with the first ASIC. Same goes for the masks, where although it gets up to 9, it is still non significant compared to the total amount of pixels. In the column *Trim change* two results stand out, the ones for the 0-1 and 3-1 ASICs, that we are going to analyse individually. For those ASICs we also see that the trim change is very asymmetric, with most of the trims increasing for the 0-1 and the opposite for the 3-1. Some others have also a pretty asymmetric change such as the 2-2, but we will see that even the highest asymmetries in the previous two will not be significant, so we will not worry about those.

Let us check then that the reproducibility in those is still acceptable. In Figure 11 we plot the noise scan for the 0-1 ASIC, where the down-left part stands out, with some sort of cluster of masked pixels and different noise levels from what we see in the rest of the grid. This means some kind of damaged pixels area, but does this affect our reproducibility significantly? We first plot the trim change matrix, Figure 12, where no clustered changes appear, so the damaged area is not the reason behind the higher and asymmetric trim change. This is a possible open question for future analysis, but we will focus whether this affects our reproducibility. For that purpose, we show the equalisation result again in Figure 13.

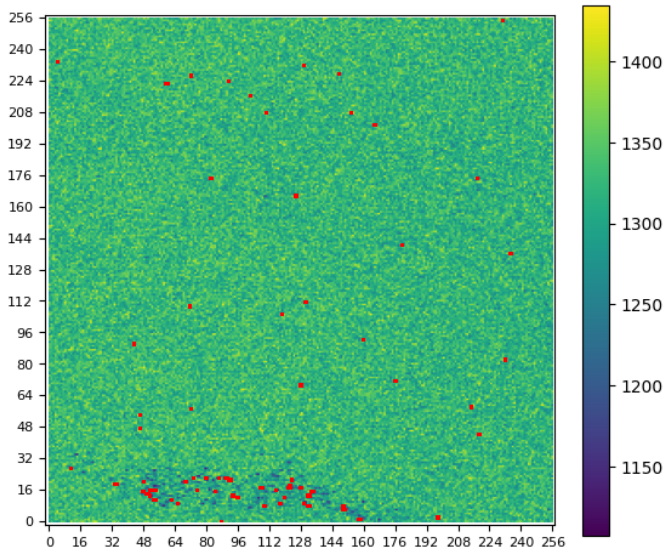

Figure 11: Noise mean of the first scan for every pixel in ASIC 0-1, plotted in a 256 × 256 grid, where the values are calculated in DAC.

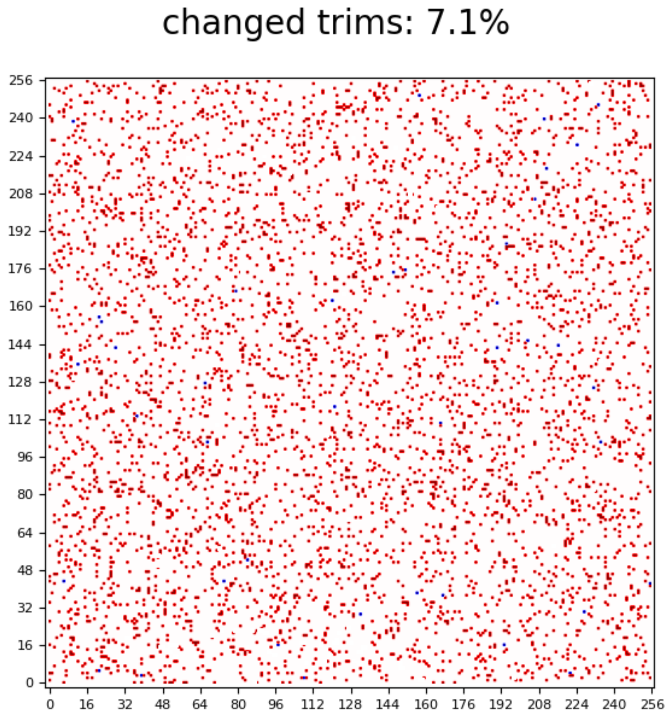

Figure 12: Difference of the best trim level to equalise each pixel between the scans for ASIC 0-1, with red points being +1, blue -1, and white 0 difference.

In Figure 13 we indeed see that the equalisation is worse than in the case of the first ASIC, as the green line does not perfectly overlaps the black one, but we still observe this jumping process where the changed trims are grouped in the borders of the main distribution peak, resulting again in a compatible equalisation, as the displacement of the black distribution is just of a few DAC. We also see some minor peaks within the blue distribution a bit far from the borders, but as previously mentioned the scale is logarithmic, so they can be neglected as they represent just a few pixels.

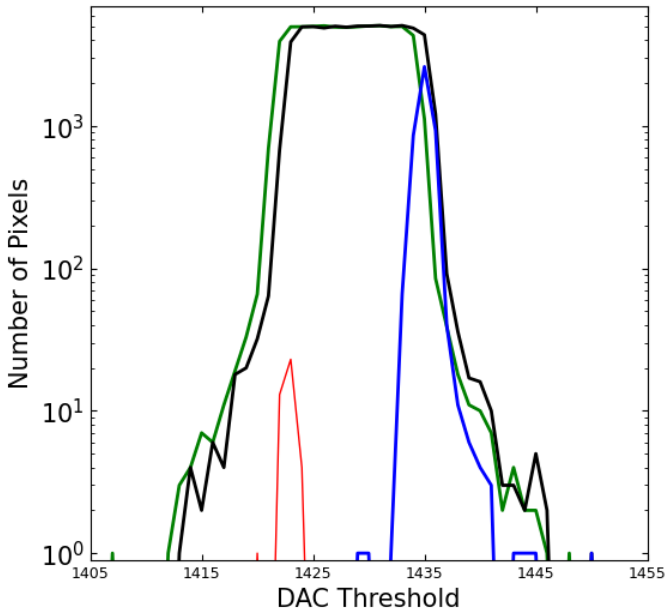

Figure 13: Black line: Equalisation result of the second scan for 0-1 ASIC, with the blue distribution being the pixels that increased their trim +1 and the red ones -1. Green: Equalisation result of the first scan.

Moving to the 3-1 ASIC, we make the same analysis, but now the noise distributions and differences look really similar to the ones we obtained in the first ASIC. If we plot the trim change, Figure 14, we see almost no red dots, which means as mentioned that almost all of the pixels that have changed their trim have decreased it. This is again an open question, as we cannot see any reason behind it from the parameters we have studied here. We then check again if the equalisations are still compatible, in Figure 15, and again the jumping process is happening and the difference between them are not relevant.

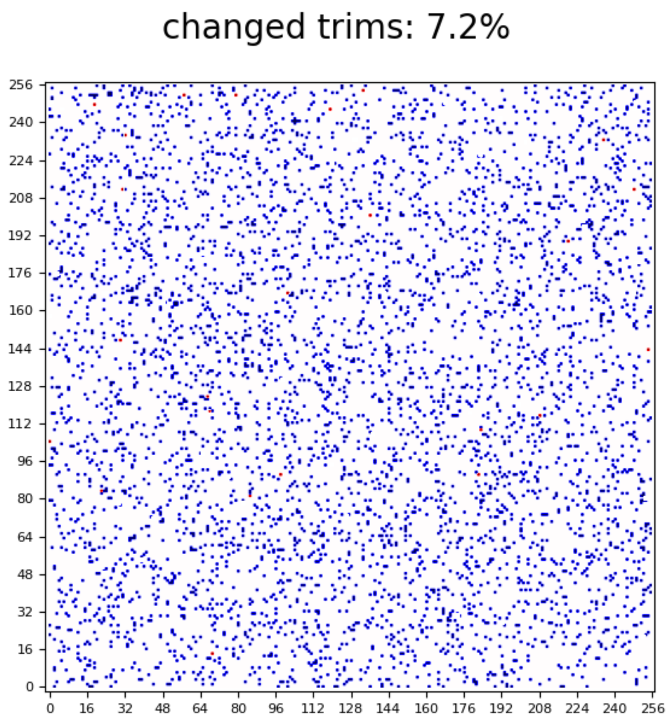

Figure 14: Difference of the best trim level to equalize each pixel between the scans for ASIC 3-1, with red points being +1 and blue -1 difference.

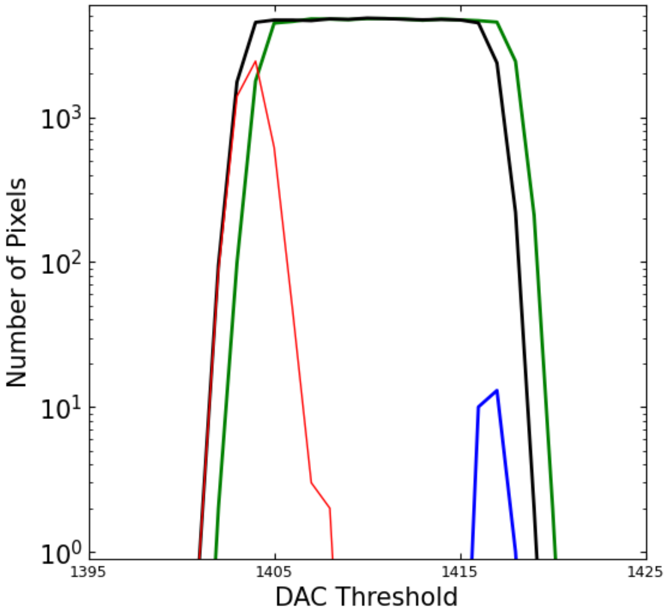

Figure 15: Equalisation result of the second scan for ASIC 3-1, with the blue distribution being the pixels that increased its trim +1 and the red ones -1.

As an additional check, we also show what happens to the second equalisation when using the first equalisation's parameters as we did in the first ASIC, as now the jumping process is still taking place but the big asymmetry between the increased and decreased could be relevant here. The reason behind the possible discrepancy is that the jumping process switch the pixels and puts them in an almost same distance to the target as if we used the original trim level, but the small differences canceled each other out in Figure 9 as we had similar amounts of increased and decreased pixels, resulting in a compatible equalisation. In this case, the unbalance between the two cases could end up in a discrepancy of the equalisation as those differences would be accumulated, depending on how small the individual differences are. The result we obtain is in Figure 16, where we indeed see a shift between the distributions, but it is just of 1.0 DAC, so again perfectly compatible with the resolution. Also, between the blue distribution and the target, which is what really tells us about how good the equalisation still is, there is again a non significant difference of 1.2 DAC.

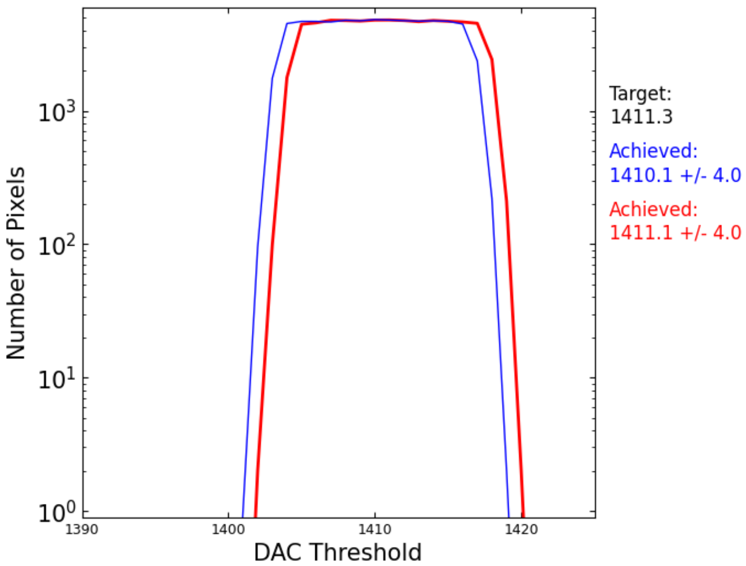

Figure 16: Equalisation result of second scan, where the blue distribution uses the first scan equalisation trim levels, and the red one uses its own equalisation, done for ASIC 3-1.

5 Extension to scans over time

Finally, we compared the reproducibility found between the first and a consecutive second scan as a reference with the one found between this first scan and one made some months ago. This will tell us how much this time difference affects our reproducibility, as the detector would have gone, for example, through several reconfigurations and reassembles.

5.1 Same conditions

First, we used a scan which was made under the same temperature condition, $18°C$, and we obtained Figure 17. We see that the red distribution, the one made with the different time scan, is indeed worse as we could expect from the consecuences of the reconfigurations and the possible damage over time, as it is a bit shifted from zero towards the positive values. Nonetheless, we see that the difference is just of 0.4 DAC, which as we discussed is not significant; even with the different ASICs in the previous version we obtained those differences. And the same goes for the width, being almost the same, just 0.1 DAC difference.

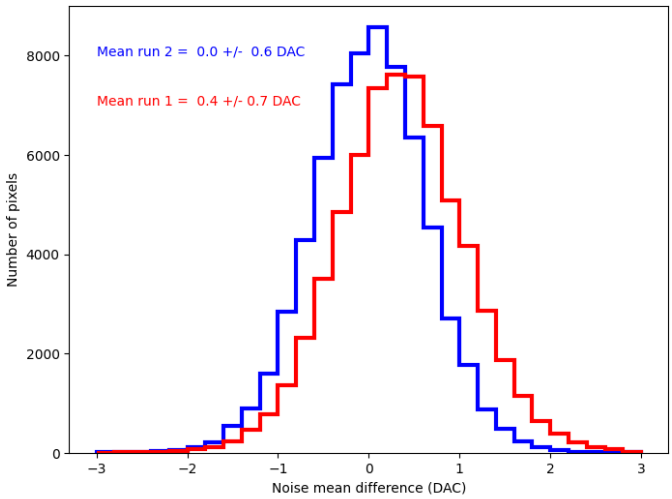

Figure 17: Noise mean difference between the scans | blue: consecutive scans, red: different time scans, same temperature.

This is in principle enough for verifying that they are compatible, but we also checked the rest of the analysis and obtained indeed good results, with a trim change of 6.8%, 10 masks changed, and the jumping behaviour in the equalisation with the peaks of the pixels that changed trim in its borders, everything according to what we have been obtaining so far.

5.2 Different conditions crosscheck

Finally, as a crosscheck of all our analysis and code, we decided to also use a scan that have been made at the same time as the one in the previous subsection, but with a totally different temperature, $-27°C$, which should give widely different results because of the temperature dependence of the detector [5]. Doing the same comparison we obtain Figure 18.

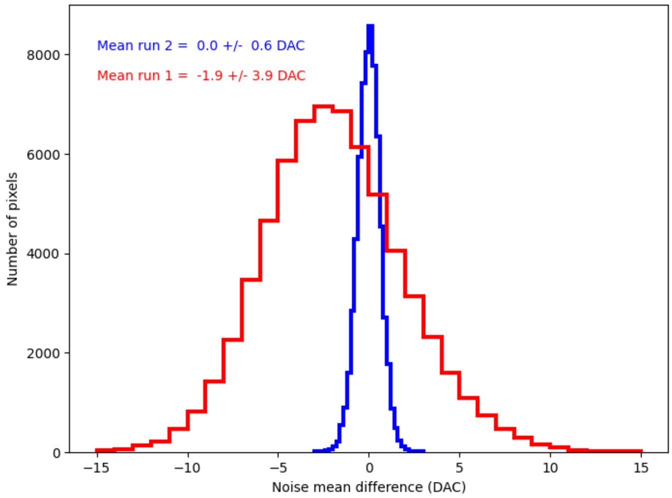

Figure 18: Noise mean difference between the scans | blue: consecutive scans, at $18°C$, red: different time scans, one at $-27°C$.

In this plot we in fact see a clear discrepancy between the distributions, where the one made with the different temperature, the red one, is at -1.9 DAC and also with a higher width, 3.9 DAC. And although the differences in the mean are not that high in comparison to the resolution, we have maximum differences of up to 40 DAC, a trim change of 33.9%, and trim changes of up to 4, so we clearly see a much worse reproducibility, as expected.

6 Conclusions and outlooks

In this thesis, the reproducibility between the equalisations of different scans of the upgraded VELO has been put to the test, comparing the key parameters for the pixels in the VeloPix: noise, trims and masks. And looking at how their changes affect the equalisation result.

We first looked at a single ASIC and compared two scans taken consecutively, obtaining non significant noise differences compared to the resolution of 1 DAC, also with a zero average difference and a negilible fluctuation. Furthermore, we obtained a change of 3.9% on the trims, checking that this did not result in a discrepancy between the equalisations, succesfully explained by the hypothesised jumping process where the pixels switch from one trim to another but still ending with a similar distance to the equalisation target. We then conclude here that the equalisations were perfectly compatible with small random fluctuations.

Later, we extended the analysis to the different ASICs within the module, and we found similar results between almost all of them, with two exemptions that we have studied individually and that have not resulted in a significant discrepancy between the equalisations either. We leave here the open question of why do those two ASICs have a higher and asymmetric trim change, that can be analysed in future studies.

And finally, we have looked at what happens when we compare two scans taken at different times, obtaining a slightly worse but compatible equalisation. Also, as a crosscheck we compared it to one made with a wide different temperature, obtaining an expected significantly higher discrepancy. This time analysis means that the possible discrepancies appearing when, for example, reassembling the modules to a different location, are not due to the reproducibility but to some new factors appearing in the new location.

So as a general conclusion, the reproducibility was obtained in all the ASICs and between scans that have been taken both back to back and with some months difference. This results can be applied to the other studies about the VELO where the uncertainty due to the reproducibility can now be neglected, such as in the temperature dependance study [5] where the fluctuation peaks need now another explanation. Also, when more advanced equalisation calculations are established, the reproducibility can again be put to the test to check if it still holds.

7 Acknowledgements

I would mainly like to thanks my first supervisor, Kristof, for helping me through all the process from the very beginning, with always quick responses and good discussions where he always said what I needed to hear to understand everything and keep going. Also, wanted to thanks the people from the VSI team, that were always kind to all the new students in the weekly meetings we had.

A The code

Repository for the python code used to analyse the already decoded data:

https://github.com/Mortrar

References

[1] R. Mann. An Introduction to Particle Physics and the Standard Model. CRC Press. 2010.

[2] LHCb Collaboration. LHCb Tracker Upgrade Technical Design Report. 2014.

[3] F. Alessio. A 40 MHz Trigger-free Readout Architecture for the LHCb Experiment. Proceedings of the Topical Workshop on Electronics for Particle Physics, pp.514-519, 2009.

[4] E. Buchanan. The LHCb Vertex Locator (VELO) Pixel Detector Upgrade. Journal of Instrumentation, JINST 12 C01013, 2017.

[5] F. Gunnink. On the Temperature Dependence of LHCb's Vertex Locator Pixel Noise. RUG Honours college's individual research project. May, 2022.

[6] M. Vos. An accurate prediction of LHCb's Vertex Locator pixel noise. RUG bachelor thesis. June, 2021.

[7] F. Prieto. Validation of the Front-End Electronics and Firmware for LHCb Vertex Locator, PoS, TWEPP-17, 2018.

[8] S. Bakker. The automation and extension of the equalisation process for the pixel noise matrix. June, 2021.

[9] L. Clandfield. Analyzing the current noise equalization in VeloPix. RUG bachelor thesis. July, 2022.