Honours College
Individual Research Project

On the Temperature Dependence of
LHCb's Vertex Locator Pixel Noise

Fabian Gunnink
S4144775
May 18, 2022

Supervised by:
Kristof De Bruyn

Contents

1 Introduction

The Large Hadron Collider (LHC) is the world's largest particle accelerator, located near Geneva, Switzerland, and built at the European Organization for Nuclear Research (CERN) by an international collaboration of research institutes and universities from all over the world. One of the current experiments, LHCb, tries to resolve a major problem in physics: the matter-antimatter asymmetry problem [1, 2]. In everyday life we interact with objects of ordinary matter, but relativistic quantum mechanics predicted the existence of another type of matter. This antimatter is like a mirrored version of ordinary matter in the sense that the sign of physical charges is flipped, such as electric or color charge, while the mass is still the same.

By virtue of Einstein's Theory of Special Relativity, energy can be converted into matter-antimatter pairs of particles and hence it is thought that the Big Bang must have created equal amounts of both types of matter. However, the Universe seems to be built up from ordinary matter, apart from the trace amounts of antimatter created in natural processes and particle accelerators. This raises the question how the Universe came to prefer one type of matter over the other. The results from the LHCb experiment might shed new light on this problem by studying the decay of a specific particle to probe the validity of the current model of particle physics.

The detector used in the LHCb experiments consists of various detectors and trackers, as is visible in Figure 1. LHCb's innermost detector, surrounding the proton-proton collisions created by the LHC, is the Vertex Locator (VELO) [3, 4, 5]. The VELO is a silicon pixel detector, used to reconstruct the trajectory of charged particles created in the collisions. Its basic operational principle works on detecting a tiny current when charged particles traverse the silicon sensor and ionize the material, similar to a Charged-Coupled Device (CCD) now somewhat common in cameras and phones. The VELO contains 52 modules, each with 12 readout chips, referred to as ASICS, which all consist of a 256x256 grid of pixels. The pixels in these arrays always have an output current due to local thermal fluctuations and this is why a specific threshold needs to be set to differentiate between this constantly fluctuating noise and actual particle detections.

The VELO was already introduced in 2001 and has been operating since LHCb's first run back in 2010. After the second run ended in late 2018, the LHCb detector was upgraded during LHC's Long Shutdown 2, that ends in the first half of 2022. Of particular interest in this context is the VELO upgrade. The precision of many physics experiments near the end of run 2 were statistically limited due to the detector [6]. This is resolved by further boosting the luminosity, the number of collisions created in the detector per unit area and time, from $4 \times 10^{32}\,\mathrm{cm^{-1}s^{-1}}$ to $2 \times 10^{33}\,\mathrm{cm^{-1}s^{-1}}$, and by removing the previously used hardware trigger. The latter is replaced by a software trigger, upgrading the ASIC readout rate from 1 MHz to 40 MHz. Besides, the original VELO setup has its innermost strip positioned 8 mm

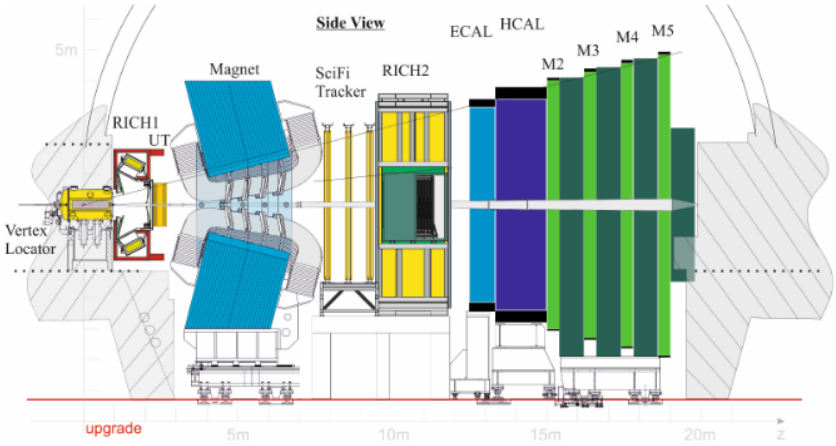

Figure 1: A schematic side-view of the upgraded LHCb detector. Source: [5].

from the central beam axis, which will be further reduced to 5.1 mm during the upgrade. Also worth mentioning is the fact that the original VELO was designed to operate under $-10°$C and that after the upgrade, the sensors should be kept under temperatures lower than $-20°$C.

One aspect that has been left untouched up till now, is the influence of temperature on the output signal. The temperature dependence in CCDs for example, has been studied over the past decades as it influences the so-called dark current: even in the absence of photons, an output signal can be measured (see for example [7, 8]). To increase the signal-to-noise ratio in such experiments, one should subtract this time and temperature dependent bias from the detected current. In much the same way, VELO is likely to have a similar temperature dependence that needs to be investigated. The goal of this project is hence to analyze the temperature dependence in VELO pixels and, if possible, find a reliable way to enhance the accuracy of future measurements.

2 The Equalisation Process

The previous section introduced the VELO detector setup and the background noise generated by the pixels. To minimize noise interfering with particle detections, these devices need to be carefully calibrated based on their position and sensitivity, a process known as equalisation. This section revolves around equalising the behaviour of the pixels within the ASICS, such that they register particles passing through at the highest possible efficiency.

2.1 Dealing with Background Noise

The main idea behind the VELO is that a particle produced in the experiment will locally ionize the medium it is travelling through, which leads to an increase in the output current of the pixels the particle passed through. Every ASIC contains a grid of 256x256 pixels, amounting to a total of 65536 sensor units, that can all separately produce a current. Simply measuring a non-zero current is however not the full story: even when no particles are travelling through the detector, the pixels will still produce a fluctuating background noise that carefully needs to be distinguished from the real particle measurements.

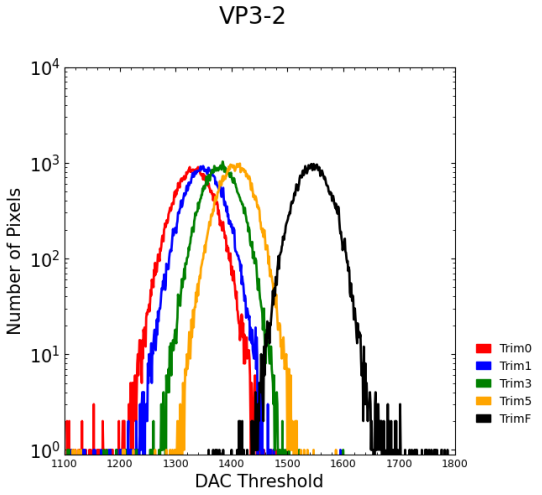

Figure 2: Number of pixels within one ASIC for a given mean noise position, measured in DAC. The different curves correspond to various trim levels and will be the topic of the next subsection.

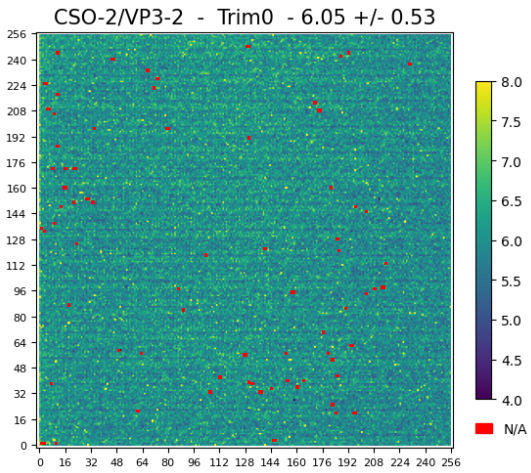

Figure 3: The width of the noise distribution plotted for each pixel in the 256x256 grid. Red dots are masked pixels and are discussed in the dedicated section.

This is in general not an easy task. One might naively assume that each pixel in the setup is identical, implying that its background noise output is concentrated

around the same equilibrium position. This is however not the case: even though all components should in principle be exact copies of one another, small deviations and inconsistencies might arise during the production process of the pixel. To quantify how much these pixels differ in their background output current, a noise distribution can be made. Such a plot gives information about the mean noise position, as well as its width around this average. Figure 2 shows the number of pixels that are present in some ASIC for a given mean noise position. The reason why this distribution looks more like a Gaussian than a delta function, is related to the dissimilarities between pixels, as quoted earlier. The theory behind the different trim levels shown is explained in the next subsection. Next to this, the width of the pixel's noise distribution can also be calculated. The resulting values are shown in the 256x256 pixel grid in Figure 3.

2.2 Using Trim Levels

Particle detections lead to a sudden spike in the output current. By introducing a threshold value for such an event, it becomes possible to differentiate between noise and the passing of a particle. While determining the exact position of this threshold, a delicate balance is sought between maximizing the particle detection probability, but simultaneously minimizing the chance that noise is incorrectly labeled as a measurement. One drawback of the current VELO setup is that this threshold can only be set for the whole pixel grid, which is why this will be referred to as the global threshold position. At first sight this seems problematic, since the mean noise position can differ substantially between pixels: for example, one pixel could constantly be firing since its mean noise position is close to the global threshold.

This nasty situation can luckily be resolved by making use of the pixel's so-called trim levels. A total of 16 levels, ranging from 0 to F in hexadecimal form, allows the mean noise position for every pixel to be shifted by some amount. This opens up the possibility to move all the pixel mean noise values individually to the optimal position based on the global threshold. The current setup hence uses the following strategy:

1. Measure the mean noise position for the lowest and highest possible trim level for every individual pixel and then average over the pixels;

2. Average these global mean noise positions over the two trim levels to get the global target, that each pixel should come close too;

3. Calculate per pixel the trim level needed to bring its mean noise position as close to the global target as possible;

4. Introduce a global threshold for particle detection based on the global target.

The calculation of the most suitable trim level in the third step currently relies on a simple linear interpolation formula between the two known endpoints. One

might wonder at this point, why the mean noise position for every trim level is not simply measured, such that the closest trim level no longer has to be calculated through interpolation. The answer to this question lies in the fact that measuring the noise distributions for all trim levels simply is too time-consuming. Recently, the linear interpolation scheme was shown to be more accurately described by a cubic interpolation method [9].

2.3 Masked Pixels

The main idea behind the equalisation process of the pixels is to shift their mean noise position as close to the global target as possible using the 16 available trim levels. For certain pixels, using even the outermost trim levels might not be enough to bring them to within a predefined range around the global target. If, for instance, the mean noise position of such a pixel lies too far above the global target, close to the global threshold, then this pixel might continuously give off a signal of particle detection. Along the same lines, a pixel whose mean noise position is far below the global threshold, even if the maximum trim level is used, perhaps never registers a particle when neighbouring pixels clearly detect one.

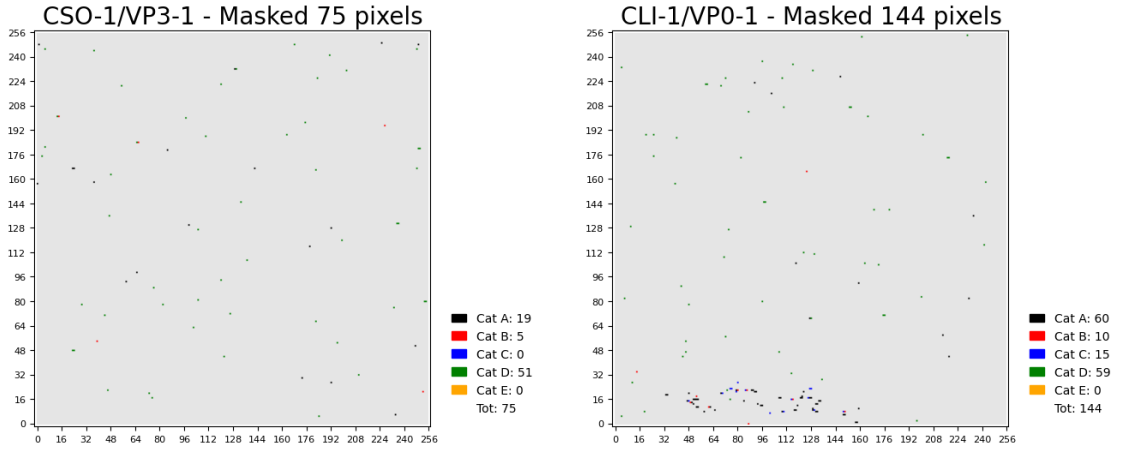

Figure 4: Two examples of a 256x256 matrix mask such that misbehaving pixels are not taken into account during further analysis. The categories correspond to the ones given in Table 1.

A simple but effective method to resolve this issue, is by simply ignoring any problematic pixels. Although this locally removes inconsistent behaviour, this comes at the cost of accuracy: after all, fewer pixels are available to confirm or deny the passing of a particle. The example in the paragraph above accounts for only one reason why a pixel should be masked, while there exist various other grounds for ignoring specific pixels. These masking categories are summarized in Table 1. Figure 4 visualizes the masked pixel categories in two different ASICS. Note how the

Category	Masking reason
A	Dead pixel: the pixel fails to respond to any sort of impulse;
B	Defective pixel: the pixel fails to respond to a scan on the lowest trim value;
C	Defective pixel: the pixel fails to respond to a scan on the highest trim value;
D	The pixel's mean noise position cannot be moved to the global threshold with the available 16 trim levels;
E	The pixel's mean noise position is outside the allowed range around the global threshold.

Table 1: The different causes as to why a pixel should be masked. Source: [10]

right-hand plot shows that the VP0-1 ASIC has an area near the bottom with a lot of dead pixels.

2.4 Manually Masked Regions

The previous subsection already highlighted that misbehaving pixels ought to be masked during a further numerical analysis, since their output values simply are nonsensical. Measuring artefacts in the noise scanning software can sometimes lead to similar circumstances, but instead of only a few deviating pixels, whole columns can be affected. This is for example visible in Figures 5 and 6. These regions are distinctly erroneous and significantly influence the width of the noise distribution and further calculated quantities. This issue is circumvented by inspecting each noise matrix by hand and manually masking any problematic regions.

Another possibility was to perform the measurements a second or even third time, since these measuring artefacts are not certain to arise. This would, however, take more time than when these strips are simply masked. We chose for the masking procedure since we were already constrained for time while trying to measure each of the 5 trim levels under varying external temperatures.

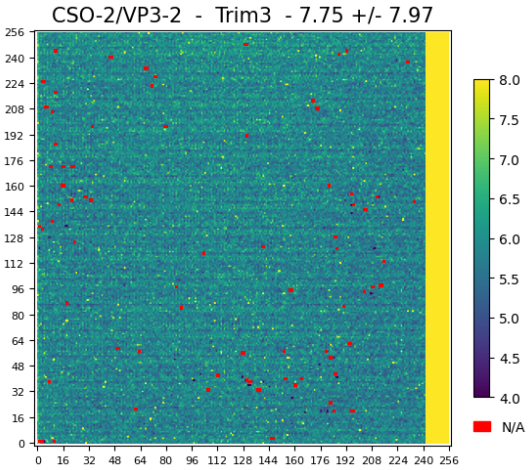

Figure 5: The width of the noise distribution plotted for each pixel in the 256x256 grid, similar to Figure 3. This time however, the columns with index ≥ 240 are clearly irrational and should hence be manually masked.

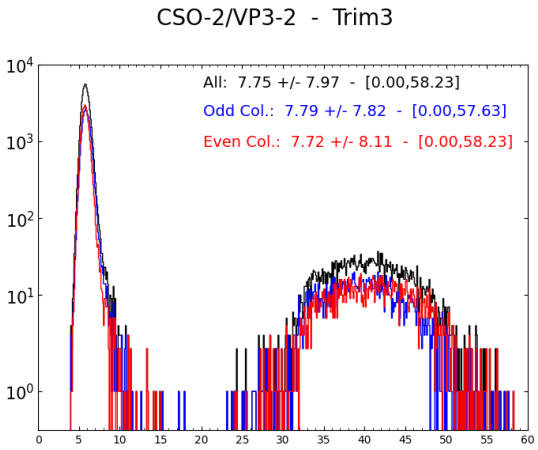

Figure 6: Data from Figure 5 binned in a histogram. The two floats in square brackets give the minimum and maximum values in the data range. For reference, the maximum value is for similar configurations at most ~ 20; those distributions miss the second, broader peak.

3 Temperature Dependence

Having explained the basics behind pixel equalisation and having introduced the necessary nomenclature, it is time to move to the objective of this report: analyzing the temperature dependence of the pixel noise distribution and the number of masked values. As was noted in a previous bachelor thesis [10], earlier measurements implied a temperature dependence on the trim 0 configuration, but not necessarily on higher trim settings. One of the reasons for writing this report was, therefore, to check to what extent this assumption is true. Instead of performing the analysis directly on all these pixels, the exploration of the detector's temperature dependence is first carried out based on how the global quantities behave in each grid. This is simply done to get acquainted with the available data without immediately getting overwhelmed by the dimensions of the problem.

The first subsection will hence focus on how the external conditions alter the global mean noise position and standard deviation for different trim levels. The latter is of particular interest: the width of the noise distribution around the mean essentially tells one how trustworthy a pixel really is. An equalized pixel with a very broad distribution could still lead to invalid detections since its width could surpass the global threshold. The next subsection improves on this approach by moving to the per-pixel basis, in which an identical technique is applied on each individual pixel noise distribution separately. The final subsection inspects how the number of masked pixels changes over temperature.

The data used in this project was acquired by K. De Bruyn. using the facilities and VELO device available at Nikhef, the Dutch National Institute for Subatomic Physics in Amsterdam. Noise distributions were calculated for 5 trim levels (0, 1, 3, 5 and F), under 5 different temperatures (18°, 8°, −12°, −21° and −27°). The uncertainty in the latter has for ease been set to zero. During data acquisition the temperatures were controlled using accompanying software, though the set values never corresponded with the measured temperatures. It is easy to imagine that this is the result of the detector coming into an equilibrium with on the one hand the cooling down mechanism, and on the other hand heat generated by reading-out the ASICs. This was of particular importance for the lowest temperatures. The conditions at which equilibrium would be reached were hard to predict, and led to the asymmetric spacing between the 5 temperatures mentioned. From the 12 ASICS in the setup, 9 were operational during the time of data taking.

3.1 Analyzing the Global Quantities

The temperature dependence is mapped for each of the 9 ASICS, but for brevity, only 1 ASIC will be actively followed in this section: VP0-2. Plots and statistics for the remaining 8 grids can be found in Appendices C & D. The first quantity that

will be looked at is the global mean and width of the noise distribution and their behaviour under a change in conditions. Using a simple Python3 script, the results were generated and plotted in Figures 7 and 8.

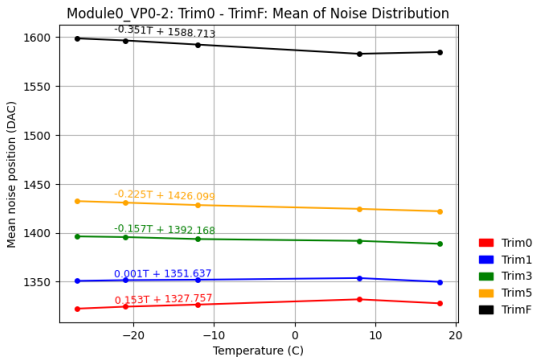

Figure 7: The global mean of the noise distribution for different trim levels as a function of temperature. The formula is the solution to the linear least squares fit of the form $y = aT + b$.

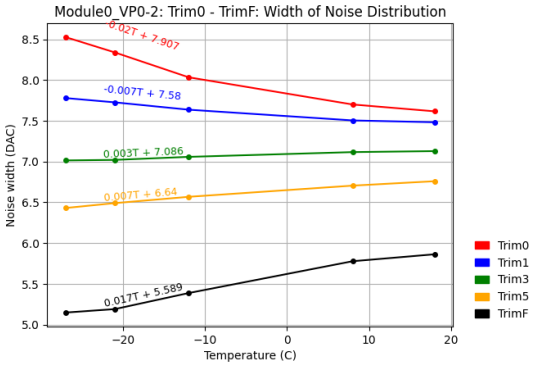

Figure 8: The global width of the noise distribution for different trim levels as a function of temperature. The formula is the solution to the linear least squares fit of the form $y = aT + b$.

From these figures, one notes that the mean noise position slowly changes in approximately a linear fashion, but remains fairly stable overall. A linear least squares fitting procedure with a degree-1 polynomial resulted in the formulas shown on the graphs. For a positive slope coefficient, the mean of the noise distribution moves to higher values at higher temperatures, and vice versa for a negative slope.

The width of the noise distribution steadily increases or decreases, depending on the trim level under consideration. Once again, a positive slope coefficient means that the distribution increases in width towards the higher temperature regime, while the opposite holds true for a negative coefficient. An increase in the width when the temperature drops to operating conditions is clearly strongest for the trim 0 configuration.

These figures already give us a first answer to one of the questions posed in the previous sections, whether the temperature dependence is only apparent in the trim 0 configuration or not. It is true that the effect is strongest in this trim setting, judging by the magnitude of the slope coefficient for both the mean and width of the noise distribution, but is certainly non-zero for the other trim levels. The trim F, for instance, is influenced by temperature in an almost equal but opposite fashion relative to trim 0. Any trim level in between, is somewhere between these two maxima. Note, however, that out of the 16 trim levels, only trim 0, 1 and possibly 2 have a negative slope coefficient for the width of the noise distribution for VP0-2, while the remaining 13 trims presumably have a positive slope. There is hence an

asymmetric distribution between these symmetric bounds. The clear separation of lines visible in these figures are certainly not apparent in other ASICS, as the reader can confirm by looking into Appendix C & D.

3.2 Analyzing Per-Pixel Quantities

Investigating only the behaviour of the average quantities obscures part of the details on how these observables in reality change. A follow-up on the previous approach is hence to look at every pixel individually and see how this one in particular acts under these conditions. As a start, one can take a randomly chosen pixel and see how its mean and width in the noise distribution are behaving. These results are visible in Figures 9 to 12.

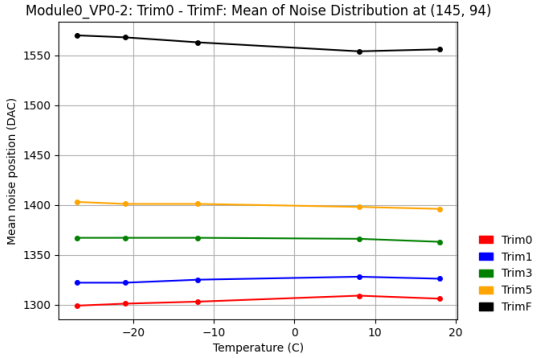

Figure 9: The mean of the noise distribution for pixel (145, 94) for different trim levels as a function of temperature.

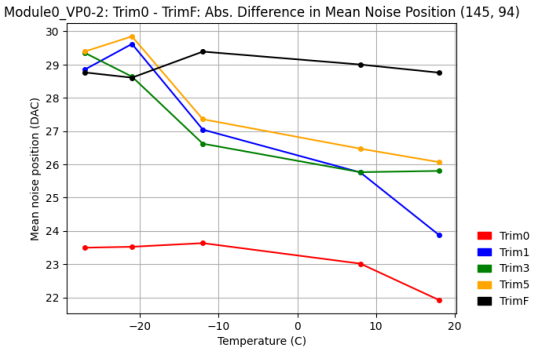

Figure 10: The absolute difference between the global and pixel (145, 94) mean noise position for different trim levels as a function of temperature.

The mean noise position for this randomly selected pixel in Figure 9 follows the same overall pattern as the global mean noise position from Figure 7. By subtracting the contribution of the latter, what remains is the absolute difference between this sample pixel and the average behaviour and is plotted in Figure 10. Such a plot gives an indication of how the pixel acts relative to the global values. All the values being positive implies that this pixel's mean noise position, or for that matter probably the whole noise distribution, is shifted towards higher DACs. Also note how a match between the pixel and global behaviour up to a constant shift would result in a straight horizontal line for every trim level.

Two identical plots are given in Figures 11 and 12, but this time for the width of the noise distribution and how its magnitude compares to the global behaviour from Figure 8. Recall that a positive absolute difference in this case means that this pixel's noise width is wider than the average of all pixels. Note in particular that

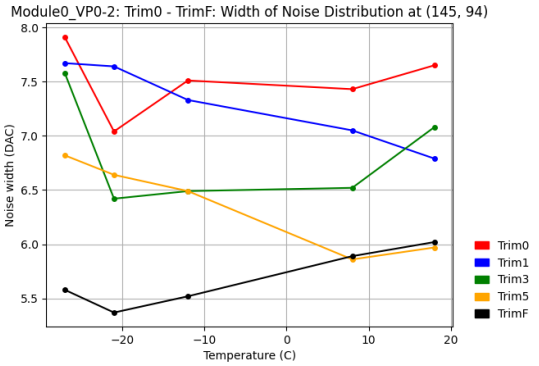

Figure 11: The width of the noise distribution for pixel (145, 94) for different trim levels as a function of temperature.

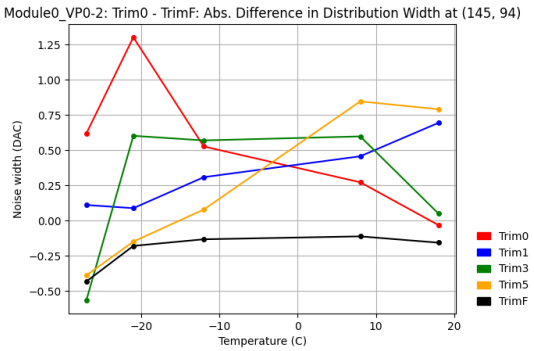

Figure 12: The absolute difference between the global and pixel (145, 94) noise distribution width for different trim levels as a function of temperature.

these trends are a lot bumpier than those visible in Figure 8 and no longer follow the somewhat linear global behaviour. However, the pattern found for the global width is still recognizable in Figure 11: trim F has a positive slope and trim 0 a negative one. For this specific pixel also trim 3 and 5 have a clear negative slope, which differs from the global quantities.

A quick way to visualize and compare how each pixel acts under a change in conditions, one can take a linear least-squares fit through the different datapoints to generate a first-order model for the temperature dependence of the pixels mean noise position and noise distribution width. Examples of these for both the mean noise position and the width of the noise distribution for a single pixel is visible in Figures 13 and 14. Note the meaning of the two parameters in the fitting function $y = aT + b$: the coefficient a is directly related to the temperature T and gives an indication of how the pixel behaves under varying circumstances. The parameter b, on the other hand, is not of particular interest since it simply quantifies a temperature-independent offset due to the specifics of the pixel under consideration.

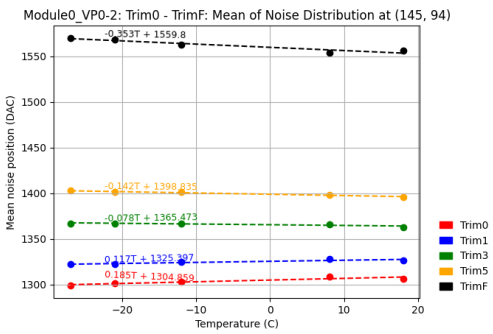

Figure 13: The mean of the noise distribution for pixel (145, 94) for different trim levels as a function of temperature, fitted with linear functions.

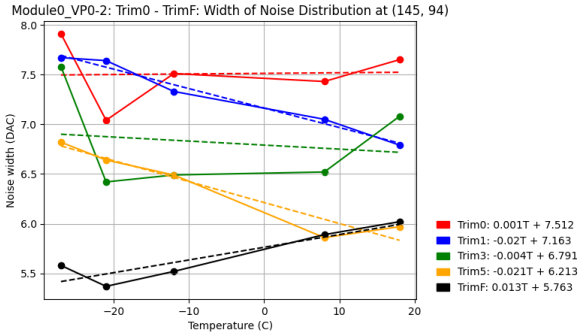

Figure 14: The width of the noise distribution for pixel (145, 94) for different trim levels as a function of temperature, fitted with linear functions.

The parameters resulting from this fitting process belong to a specific pixel, so when this analysis is performed on the whole grid, a histogram can be made in which pixels are grouped by their coefficients. As was noted in the previous paragraph, only the slope parameter will be looked at in more detail since the offset value is irrelevant for now. The results are visible in Figures 15 to 18 for the outermost trim levels.

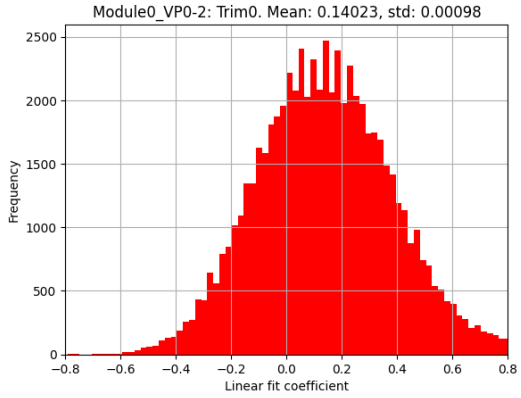

Figure 15: Histogram in the slope parameter for a linear fit through the mean noise positions over temperature in trim 0, for all pixels in the ASIC.

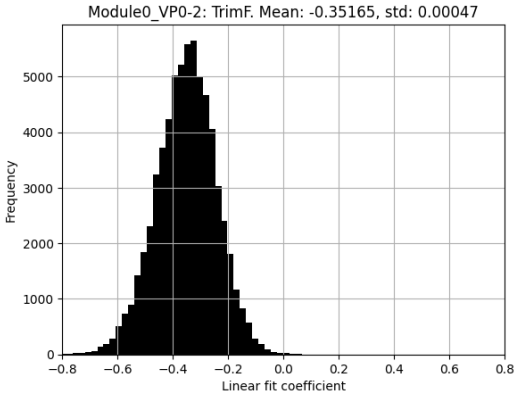

Figure 16: Histogram in the slope parameter for a linear fit through the mean noise positions over temperature in trim F, for all pixels in the ASIC.

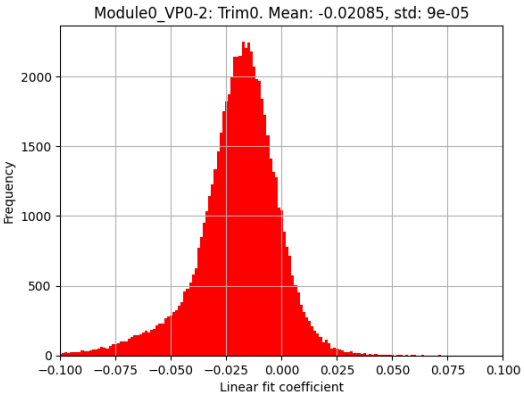

Figure 17: Histogram in the slope parameter for a linear fit through the width of the noise distribution over temperature in trim 0, for all pixels in the ASIC.

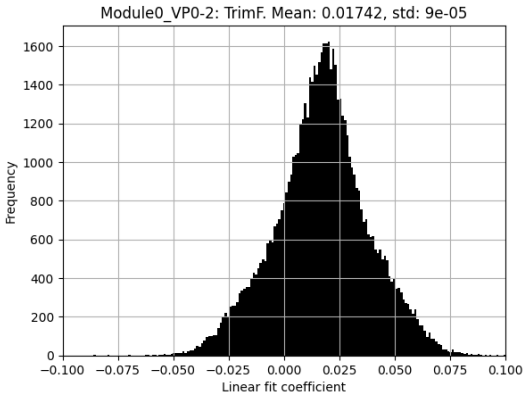

Figure 18: Histogram in the slope parameter for a linear fit through the width of the noise distribution over temperature in trim F, for all pixels in the ASIC.

From a pedagogical point of view, it might be best to again summarize what these histograms truly mean. The first subsection of this chapter analyzed the behaviour of the pixels in a rather crude way: one takes the mean noise position for the whole assemble of pixels and then calculated how this value changes over a broad temperature range. The fitting process gives a formula that tells how the global noise position changes over temperature. The same is done for the width of the noise distribution. It is important to note that this glosses over the details of what the pixels are actually saying. The approach of this subsection is to look first at how each individual pixel behaves under temperature and only then compare the pixels to each other in a histogram. The mean of the histogram should be relatively close to the slopes in Figures 7 and 8, and indeed they do, as the reader can check using Table 2. In a sense, the operation of averaging and fitting is reversed between these two sections.

Trim Level	**Mean (global)**	**Mean (pixel hist.)**	**Width (global)**	**Width (pixel hist.)**
0	0.153	0.140	-0.020	-0.021
1	0.001	-0.008	-0.007	-0.009
3	-0.157	-0.163	0.003	0.001
5	-0.225	-0.227	0.007	0.007
F	-0.351	-0.352	0.017	0.017

Table 2: Comparing the slope values in the global quantities to those determined using the pixel histograms, specifically for VP0-2.

The histograms in Figures 15 and 16 summarize what the pixel mean noise positions do under a change in temperature for trim levels 0 and F, respectively. Most of the pixels in the trim 0 histogram have a positive slope, meaning that these shift to lower mean noise positions for lower temperatures. The trim F level, curiously, has almost no pixels with a positive coefficient, implying that essentially all pixels have their mean noise, possibly the whole noise distribution, shifted towards higher values when the temperature drops. Also note how the distribution for the trim F configuration is a lot narrower compared to the trim 0 setting. This implies that the pixels in this trim behave more evenly compared to, for example, the trim 0 setting. Figures 17 and 18 are similar to the other plots mentioned, but instead of being related to the mean noise position, they refer to the width of the noise distribution. A negative coefficient means that the distribution gets wider for lower temperatures and vice versa for a positive parameter.

The sample mean and standard deviation can be determined from the histograms. These are plotted categorically in Figure 19 and 20. Comparing the average behaviour of the outermost trim levels, we can conclude the following for VP0-2:

1. For trim 0, the mean noise position is shifted towards *lower* DACs for lower temperatures, while the width of the noise distribution *increases*.

2. For trim F, the mean noise position is shifted towards *higher* DACs for lower temperatures, while the width of the noise distribution *decreases*.

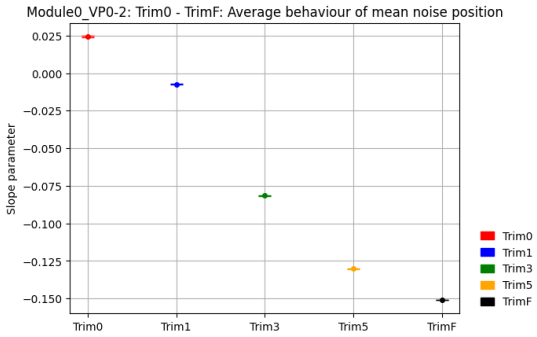

Figure 19: Average and standard deviation of the histograms in the mean noise position, plotted categorically for the examined trim levels in VP0-2.

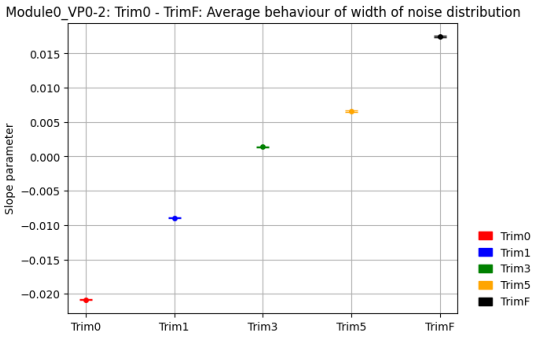

Figure 20: Average and standard deviation of the histograms in the width of the noise distribution, plotted categorically for the examined trim levels in VP0-2.

3.3 VP0-2: A Deceptive Example

It turns out that VP0-2 was special after all. Comparing the outcomes of different ASICS, as can be found in Appendix I & J, shows that VP0-2 is the strongest example of an ASIC for which a higher trim level corresponds to an increasingly

lower value of the average mean noise position slope. For completeness' sake, the final plots for VP3-0 have also been added in Figures 21 and 22 to make a clear distinction between the pixels of both ASICS. To reiterate, the remaining 6 ASICS have a closer resemblance to VP0-2.

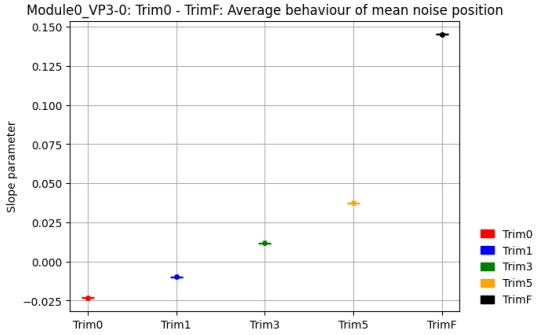

Figure 21: Average and standard deviation of the histograms in the mean noise position, plotted categorically for the examined trim levels in VP3-0.

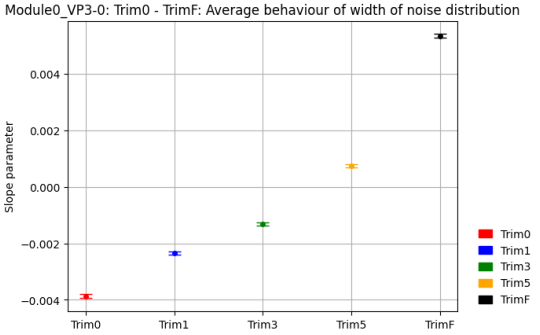

Figure 22: Average and standard deviation of the histograms in the width of the noise distribution, plotted categorically for the examined trim levels in VP3-0.

Based on Figures 21 and 22, we conclude the following for VP3-0:

1. For trim 0, the mean noise position is shifted towards *higher* DACs for lower temperatures, while the width of the noise distribution *increases*.

2. For trim F, the mean noise position is shifted towards *lower* DACs for lower temperatures, while the width of the noise distribution *decreases*.

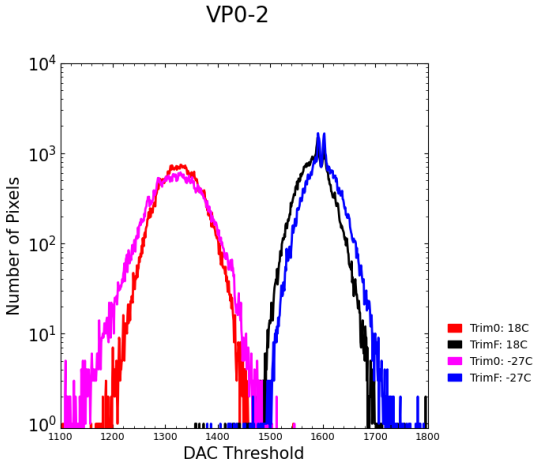

Figure 23: Number of pixels in VP0-2 for a given mean noise position, at 18°C and −27°C, for the 0 and F trim.

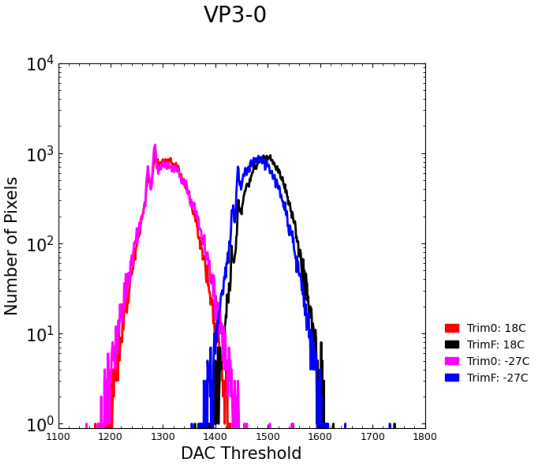

Figure 24: Number of pixels in VP3-0 for a given mean noise position, at 18°C and −27°C, for the 0 and F trim.

These statements can be matched to the original equalisation plots for both ASICS, which show the number of pixels for a given DAC threshold. To allow for a clear comparison, only the trim 0 and F curves are shown at 18°C and −27°C. From Figure 23 for VP0-2, one notes that:

1. For trim 0, the distribution becomes wider and moves to lower DAC values for the lower temperature. This should result in a positive slope coefficient for the mean noise position.

2. For trim F, the distribution stays constant in width but moves to higher DAC values for the lower temperature. This should result in a negative slope coefficient for the mean noise position.

These observations are in compliance with the results from Figure 19.

From Figure 24 for VP0-2, one notes that:

1. For trim 0, the distribution stays constant in width but moves to slightly higher DAC values for the lower temperature. This should result in a barely positive slope coefficient for the mean noise position.

2. For trim F, the distribution stays constant in width but moves to lower DAC values for the lower temperature. This should result in a positive slope coefficient for the mean noise position.

These observations are in compliance with the results from Figure 21.

In which direction the mean of the noise distribution moves, as we have seen, differs per ASIC. The width of the noise distribution over the measured trim values is from this perspective more consistent: a higher trim level attains a higher slope coefficient, with a few exceptions. With trim F, the width always decreases at lower temperatures. VP0-0 being the exception, this width only increases at these temperatures for trim 0. Of course, it is surely possible for two distribution to have their means moving in opposite directions and have their widths increasing at the same rate. It does, however, raise the question why VP0-2 differs then in this aspect from the other readout modules, since they should all have more or less the same overall behaviour.

3.4 Influence on Masked Pixels

As stated in the previous section, there are a total of five reasons why a pixel could be marked as improper. These are masked to prevent them from introducing unwanted pixel behaviour under a further analysis. Table 1 showed these five different categories and the specific reason as to when a pixel would belong to this group. A few of these categories could in principle be influenced by the temperature of the

detector. This subsection will therefore explore the number of masked pixels per category as a function of temperature and what this implies for the sensitivity of VELO.

By the nature of the cause, the masking categories A, B and C are not likely to show an explicit temperature dependence. A dead or defective pixel is simply working improperly from the start. Categories D and E, on the other hand, depend on the pixel's mean noise position and, as was already shown in the previous subsection, this quantity does change as a function of temperature. It would not come as a surprise if either of these have a different number of masked pixels under varying conditions. Figure 25 visualizes for two ASICS, VP3-0 and VP3-2, how the number of masked values per category changes over temperature. Note that these plots have a logarithmic scale on the y-axis with base 2, and that for the plot on the left for VP3-0, categories B and C are hidden behind category E at a constant level of 0.

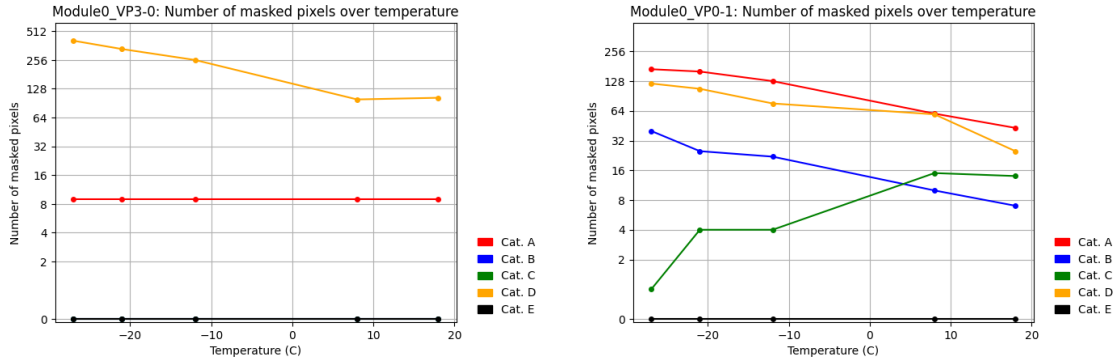

Figure 25: The number of masked pixels per category as a function of temperature for two different ASICS.

The absolute difference between the number of masked pixels at $-27°$C and $18°$C is also summarized in Table 3. These values should give a gross indication of whether a category is temperature dependent or not. Categories that have values close to zero in every row, are unlikely to be influenced by temperature. For some unknown reason, VP0-1 behaves differently compared to the other ASICS: its number of dead pixels, surprisingly, skyrockets as temperature goes down. This behaviour is not, or at most of limited significance, apparent in the others. Apart from this inconsistency, the earlier hypothesis that categories A, B and C should not be heavily influenced by external conditions are confirmed.

Category E is irrelevant for the present data: it attains a value of 0 for every possible configuration for the whole temperature range measured. The most interesting part of this analysis, is of course the behaviour of pixels in category D. As suspected, these heavily rely on temperature through their dependence on the mean noise position. The trend visible in each of the ASICS confirms that the number of masked pixels in this category is higher for lower temperatures. This in itself might be surprising:

ASIC	**Cat. A**	**Cat. B**	**Cat. C**	**Cat. D**	**Cat. E**
VP0-0	1	1	0	291	0
VP0-1	126	33	-13	96	0
VP0-2	-4	26	0	24	0
VP1-0	-4	5	0	171	0
VP1-1	0	0	0	96	0
VP1-2	-1	0	0	98	0
VP3-0	0	0	0	304	0
VP3-1	-1	2	-1	124	0
VP3-2	14	13	-1	71	0

Table 3: The absolute change in the number of masked pixel between the temperature range endpoints at $-27°$C and $18°$C. A positive value hence means that there were more pixels masked at $-27°$C than at $18°$C.

the VELO is set to operate under conditions of $-30°$C to minimize thermal noise within the pixels, yet, Table 3 shows that the number of fully functioning pixels decreases by $0.15 - 0.45\%$ at these conditions compared to room temperature.

4 Discussion and Conclusions

This project aimed to give insight in how a VeloPix ASIC acts in different temperature regimes and tried to quantify this behaviour through fitting and comparing. It is clear that this device has a non-trivial temperature dependence. The behaviour of the global mean noise position and the global width of the noise distribution already showed the peculiar behaviour in some ASICS, paving the way for a more in-depth analysis on a per-pixel basis.

After fitting the mean noise position data for every pixel individually and comparing the resulting histograms in the slope values between different trim levels using Figures 19 and 21, it is clear that higher trim settings have a more positive linear fit coefficient, except for VP0-2. The trim 0 setting itself has a negative or positive slope, based on the ASIC under consideration, and the trim F setting only pushes these numbers higher. In other words, a higher trim level means that the pixel is more prone to acquire a lower mean noise position as the temperature drops. Consequently, the number of pixels that were masked by category D starts to change. Pixels masked at room temperature can suddenly become equalizable at lower temperatures due to their noise distribution being shifted, but it turns out that there are far more pixels undergoing the opposite reaction: they used to be reliable pixels, but the changing temperature altered their noise distribution so much, that they suddenly become unequalizable and need to be masked.

At the same time, the width of the noise distribution also changes as the temperature drops, as is visible in plots like in Figures 20 and 22. A positive slope value in these diagrams means that for lower temperatures the noise distribution becomes narrower, and vice versa for negative values. It is clear that for all ASICS, these values become more positive for higher trim levels. So, in other words, a trim level F functioning under low temperatures has a relatively narrow distribution. This is beneficial for the detector itself: a narrower noise distribution minimizes the chance that a random fluctuation is accidentally measured as a particle passing through. If this were to happen to all pixels, then the difference between the global mean noise position and the global threshold could be further decreased, increasing the sensitivity of the detector. Note however that this is only true for trim levels having a positive slope coefficient; the trim 0, for example, has a negative value in each of the ASICS, meaning that the reverse logic applies.

As has been argued, there are both pros and cons of working in the low temperature regime when considering pixel behaviour only. On the one hand, it could increase the sensitivity of the detector at a high trim level when the detector is set to operate under $-30°\text{C}$. On the other hand, this comes at the direct cost of fully functioning pixels, reducing the overall sensor sensitivity. It comes down to finding a delicate balance between these two opposing effects. The second effect was estimated to lower the efficiency of the detector by a factor of $0.15 - 0.45\%$. The first effect

is, however, hard to quantify, since it depends on how much the global threshold is moved towards the global target. Moreover, the main reason why the detector is currently kept under conditions of $-30°$C is to minimize leakage current due to radiation damage.

This report tried to give a straightforward overview of the behaviour of the pixels within the VELO setup. In several steps in the analysis, a linear fitting procedure was carried out to quantify the change in the noise distributions between pixels of the same ASIC. If the use of such a simple formula was allowed in hindsight, is still to be discussed. As an example, Figure 8 shows the global width of the noise distribution, while Figure 11 shows the same quantity but then for a specific pixel. The global values seem to be well-behaved and, at first sight, a linear fit could suffice. When looking at the behaviour of a specific pixel however, it becomes clear that a lot of details in the global quantities were averaged over. The bumpy and unpredictable lines sometimes have clear outliers and seem to defy a simple linear fit. However, one has to keep in mind that no low-order polynomial or other basic mathematical function are a proper fit for the data. Moreover, the degree of a higher-order function would need to be suppressed in order not to overfit the six datapoints or clash with Occam's Razor. A linear fit might be naive, but still grabs the most important information in a single and simple formula.

One might argue that the data points in, for example, Figure 11 lack error bars: if these were large enough, then this might still allow for a linear fit. Quantifying this inaccuracy could be done by performing the same measurements under alike conditions, but due to a lack of time, this part of the procedure was omitted. In hindsight, it would have been better to include this as part of the experiment, since we now lack an independent method of determining the size of the error bars. The results as shown in this work are not completely sidelined, however, since they should still capture the gist of what the pixel behaviour looks like.

In conclusion, the analysis posed in this work gives a first-order picture of how external conditions influence the sensitivity and further behaviour of pixels in the VeloPix. This story is by no means complete: the exact reason as to why this particular behaviour surfaces was not explored due to a lack of background in engineering. Still, the author hopes this report could perhaps be of interest towards those working on upgrades of the VELO detector and LHCb experiment.

References

[1] The LHCb Collaboration. The LHCb Detector at the LHC. *Journal of Instrumentation*, 3:S08005–S08005, 08 2008.

[2] LHCb. https://lhcb-outreach.web.cern.ch/.

[3] LHCb Collaboration. Lhcb velo (vertex locator) : Technical design report. *Geneva*, CERN/LHCC 2001-011, 2001.

[4] J Libby. Velo: the lhcb vertex detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 494:113–119, 11 2002.

[5] The LHCb Collaboration. Lhcb velo upgrade technical design report. *Geneva*, CERN-LHCC-2013-021, 2013.

[6] Stefano De Capua. The lhcb velo upgrade, 2018.

[7] Ralf Widenhorn, Morley M Blouke, Alexander Weber, Armin Rest, and Erik Bodegom. Temperature dependence of dark current in a ccd, 2002.

[8] Ralf Widenhorn, Ines Hartwig, Justin C. Dunlap, and Erik Bodegom. Measurements of dark current in a ccd imager during light exposures. *SPIE Proceedings*, 02 2008.

[9] Max Vos. *An Accurate Prediction of LHCb's Vertex Locator Pixel Noise*, 2021.

[10] Senne Bakker. *The VeloPix Upgrade for the LHCb Experiment*, 2021.