

The algorithm for FIR corrections of the VELO analogue links and its performance



Public Note

Issue: 1
Revision: 0

Reference: LHCb-2008-015
Created: July 3, 2007
Last modified: September 4, 2009

Prepared by: K. Akiba^a, J. Borel^b, J. Buytaert^c, L. Eklund^d, M. Gersabeck^d

^aIF-UFRJ, Brazil

^bEPFL, Switzerland

^cCERN, Switzerland

^dUniversity of Glasgow, UK

LHCb-2008-015
03/07/2007



Abstract

The data from the VELO front-end is sent to the ADCs on the read-out board over a serial analogue link. Due to imperfections in the link, inter-symbol cross talk occurs between adjacent time-bins in the transfer. This is corrected by an FIR filter implemented in the pre-processing FPGA located on the read-out board. This note reports on a method to determine the coefficients for the filter using data taken in-situ. Simulations are presented that show the performance of the method as it is implemented in the LHCb read-out board. The effectiveness of the algorithm is demonstrated by the improvements in tracking performance on beam test data it brings.

Contents

1	Introduction	2
2	Modelling the system	3
3	Algorithm to determine the FIR coefficients	4
4	Simulation results	5
5	Test beam results of the FIR filter corrections	6
6	Conclusion	8
7	References	8

List of Figures

1	The figure shows the analogue pulse as seen by the TELL1 ADC for 16 analogue links. The x-axis show the time-bucket number that correspond to a particular front-end channel. A calibration pulse is injected in channel 24 and the average signal height is determined for each of the 16 possible ADC sampling phases. The vertical lines indicate the sampling time judged to give optimal performance.	2
2	The figures show the cross talk measured as the ratio of the signal height in the adjacent bins to the signal height in the central (pulsed) bin. The histograms have one entry per analogue link and the left-hand plot show the cross talk to the previous bin and the right-hand plot show the cross talk to the following bin	3
3	Schematic view of model used for the system, where the distortion caused by the Beetle and the analogue link in seen as one unit and the digital filter as a second unit.	3
4	The left-hand plot shows the simulated ADC values before FIR corrections. The right-hand plot shows the residual signal distortion after corrections with float and 9-bit precision.	6
5	Schematic view of the read-out order of the different regions of a R measuring sensor. Note that increasing channel number means either increasing or decreasing radius depending on the position on the sensor.	7
6	Top: Mean of unbiased track residuals for the eight regions on the sensor. The direction of the shift in the mean corresponds to the direction in read-out order shown in Figure 5. Bottom: Same after application of the FIR filter. The larger error bars are due to lower statistics of the re-processed data.	7

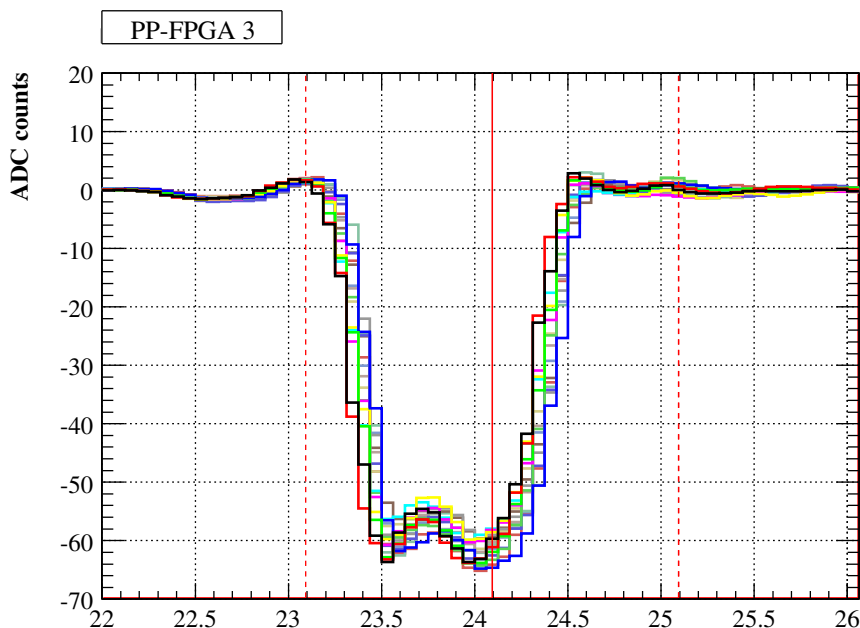


Figure 1 The figure shows the analogue pulse as seen by the TELL1 ADC for 16 analogue links. The x-axis show the time-bucket number that correspond to a particular front-end channel. A calibration pulse is injected in channel 24 and the average signal height is determined for each of the 16 possible ADC sampling phases. The vertical lines indicate the sampling time judged to give optimal performance.

1 Introduction

The LHCb [1] Vertex Detector [2] consists of 42 double sided silicon micro-strip modules. Each module side has a sensor with 2048 strip read out by 16 Beetle chips [3]. The front-end pulse is sampled and stored in an analogue pipeline and events selected by the Level 0 trigger are transferred to the read-out boards (TELL1 [4]). Each Beetle chip has 4 serial analogue outputs where the data from 32 channels are transferred at a rate 40 M samples/s.

The read-out boards are located in the counting house, at approximately 65 m cable distance from the detector front-end. The signals in the serial links are amplified by analogue drivers that compensate the attenuation and frequency distortion in the cable. Due to imperfections in this compensation and component-to-component variations, the signal is still slightly distorted by the transfer. This mainly manifests itself as inter-symbol cross talk, where the signal in one time bin spills over to adjacent bins. The size of the cross talk can be measured by studying configurations where the whole signal is contained in one signal channel. This has been done in beam tests, considering orthogonal tracks traversing the sensor close to a strip centre, or with test pulses where a signal is injected in a single channel. The signal present in adjacent time-bins is then attributed to cross-talk. This is normally expressed in per cent of the signal height of the pulsed channel.

The size of the cross talk depends on the details of the implementation of the link and the sampling time of the analogue signal on the serial links. Values of up to 20% have been observed in beam tests and test set-ups. However with the well-tuned analogue links installed in the LHCb Experiment the cross talk values are considerably smaller.

The dependence on the sampling time is illustrated in Figure 1. The the plot is generated by pulsing one channel and accumulating a large number of events at each sample delay setting. As seen in the figure the cross talk is in general asymmetric and can be both positive and negative.

The cross talk for the final system was measured for the configuration used in the LHC injection tests in June 2009. Figure 2 shows the cross talk in the forward and backward direction for one detector half determined by calibration pulses, while using the sampling point indicated in Figure 1




Figure 2 The figures show the cross talk measured as the ratio of the signal height in the adjacent bins to the signal height in the central (pulsed) bin. The histograms have one entry per analogue link and the left-hand plot show the cross talk to the previous bin and the right-hand plot show the cross talk to the following bin




Figure 3 Schematic view of model used for the system, where the distortion caused by the Beetle and the analogue link is seen as one unit and the digital filter as a second unit.

The read-out board has ADC cards digitising the serial data and pre-processing FPGAs (PPFPGA) for on-line data processing. The PPFPGA performs pedestal subtraction, digital filtering, channel re-ordering, common mode noise subtraction and clustering before the data is sent over the read-out network to the trigger computing farm. The digital filter is implemented as a Finite Impulse Response filter (FIR) and its purpose is to compensate for the distortions that occur on the analogue link. The FIR filter is a well established digital signal processing algorithm and well suited for implementation in an FPGA. The specific challenge for the VELO application is to determine the coefficients used in the algorithm, in situ and for all 5632 analogue links in the system.

2 Modelling the system

In order to analyse the system, define the filter algorithm and determine the filter coefficients, a mathematical model of the system is needed. Figure 3 shows a schematic view of the model. Data from 32 channels are sent serially over one link and is viewed as a time series of numbers. The values $x[n]$ are the *true* values corresponding to the signal in the front-end of the chip. The values $w[n]$ are the values sampled by the ADC on the read-out board, distorted by the Beetles and by the analogue link. And the values $y[n]$ are the corrected values after the digital filter that are passed on to the next stage of processing in the PPFPGA. The two units in the system are characterised by their transfer functions, H for the Beetles and analogue links and G for the digital filter.

Assuming a linear and time invariant system (LTI), the transfer function H can be described in the time domain by the number series $h[n]$. The relation between the input and output of any time-discrete LTI system can be described in this way.

The relation between the number series $w[n]$, $x[n]$ and $h[n]$ is

$$w[n] = \sum_{k=-\infty}^{k=\infty} h[k] \times x[k - n]. \quad (1)$$

The series $h[n]$ is called the impulse response, since it can be measured by sending in a single impulse $\delta[n]$ (called the Kronecker delta function) and observing the output. Hence if

$$x[n] = \delta[n] = \begin{cases} 1, & \text{if } n = 0 \\ 0, & \text{if } n \neq 0 \end{cases} \quad (2)$$

the output $w[n]$ is identical to the the impulse response $h[n]$. This is performed in-situ by injecting a signal in a single channel at the front-end. In practise, this can be done in two different ways. Either by using the built-in calibration pulse facility in the Beetle chip with an isolated channel pulsed. Or by using reconstructed tracks that pass sufficiently close to the centre of the strip to ensure that no charge sharing occurred. Both methods have been used successfully in VELO test beams.

Similarly to the way the Beetles and the analogue links are modelled, a digital filter is described by its transfer function G . In the time domain it defined by the series $g[n]$ and the relation between the input and output of the filter is

$$y[n] = \sum_{k=-\infty}^{k=\infty} g[k] \times w[k - n]. \quad (3)$$

The transfer function of the whole system from the front-end input to the output of the digital filter is found by combining the transfer functions G and H . In order for the filter to completely compensate for the distortion of the analogue link, the transfer function of the combined system should be the unity function. In the time domain the unity function is given by the Kronecker delta $\delta[n]$. Hence combining Equation 1 and 3 and equating it to the $\delta[n]$ gives

$$\delta[n] = \sum_{k=-\infty}^{k=\infty} g[k] \times h[k - n]. \quad (4)$$

This is the relation between the measured number series $h[n]$ and the number series $g[n]$ for the filter. Section 3 describes how this equation can be solved in the context of an FIR filter.

3 Algorithm to determine the FIR coefficients

The algorithm of how to apply the general digital filter is given by Equation 3. Because of the infinite sum it is clearly not possible to implement this in practice. A class of filters that approximate the general function in an easily implemented way is called Finite Impulse Response (FIR) filters. The name comes from the fact that the filters impulse response, $g[n]$, has a finite length and hence the infinite sum is truncated. The number of non-zero element in $g[n]$ is called the order of the filter. Hence for a M^{th} order symmetric FIR filter, where M is an odd number, the algorithm is

$$y[n] = \sum_{k=-\frac{M-1}{2}}^{k=\frac{M-1}{2}} g[k] \times w[k - n]. \quad (5)$$

This approximation will also reduce Equation 4 to a finite sum of M terms.

$$\delta[n] = \sum_{k=-\frac{M-1}{2}}^{k=\frac{M-1}{2}} g[k] \times h[k - n]. \quad (6)$$

It can be considered as an infinite set of equations, one for each n with the $g[n]$ as M unknowns. Assuming that the inter-symbol cross talk extends a finite distance from the original signal gives further simplification. In other words the series $h[n]$ also has only N non-zero elements, where N is

an odd number. Combining these assumptions, Equation 4 is reduced to a set of $M + N - 1$ equations with M unknowns [5].

$$\delta[n] = \sum_{k=-\frac{M-1}{2}}^{k=\frac{M-1}{2}} g[k] \times h[k-n] \quad \text{where} \quad \delta[n] = \begin{cases} 0, & \text{if } -\frac{M+N}{2} + 1 \leq n < 0 \\ 1, & \text{if } n = 0 \\ 0, & \text{if } 0 < n \leq \frac{M+N}{2} - 1 \end{cases} \quad (7)$$

Equation 7 can be expressed as a matrix equation, $AX = Y$, if each value of n is considered as one row and the terms of the sum in ascending k are considered as columns. In this view, the right hand side Y and the unknowns X are matrices of dimension $(M + N - 1) \times 1$

$$X = \begin{pmatrix} g[-\frac{M-1}{2}] \\ \vdots \\ g[-1] \\ g[0] \\ g[1] \\ \vdots \\ g[\frac{M-1}{2}] \end{pmatrix} \quad Y = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \quad (8)$$

and A is a matrix of dimension $(M + N - 1) \times M$ given by

$$A = \begin{pmatrix} h[\frac{N-1}{2}] & 0 & \dots & 0 & \dots & 0 & 0 \\ h[\frac{N-1}{2} - 1] & h[\frac{N-1}{2}] & \dots & 0 & \dots & 0 & 0 \\ \vdots & \vdots & & \vdots & & \vdots & \vdots \\ h[-\frac{N-1}{2} + 1] & h[-\frac{N-1}{2} + 2] & \dots & h[1] & \dots & h[\frac{N-1}{2}] & 0 \\ h[-\frac{N-1}{2}] & h[-\frac{N-1}{2} + 1] & \dots & h[0] & \dots & h[\frac{N-1}{2} - 1] & h[\frac{N-1}{2}] \\ 0 & h[-\frac{N-1}{2}] & \dots & h[-1] & \dots & h[\frac{N-1}{2} - 2] & h[\frac{N-1}{2} - 1] \\ \vdots & \vdots & & \vdots & & \vdots & \vdots \\ 0 & 0 & \dots & 0 & \dots & h[-\frac{N-1}{2}] & h[-\frac{N-1}{2} + 1] \\ 0 & 0 & \dots & 0 & \dots & 0 & h[-\frac{N-1}{2}] \end{pmatrix}. \quad (9)$$

As for all over-constrained systems of equations it can be solved by the least square method, where

$$X = (A^T A)^{-1} A^T Y. \quad (10)$$

Hence by measuring the cross talk of a signal injected on a single channel the impulse response $h[n]$ of the analogue link can be determined. These values are used to form the matrices described in Equation 8 and 9, and then solved using Equation 10.

4 Simulation results

The FIR filter is implemented in the PFPFGA on the LHCb read-out board, where the resource usage puts constraints on the precision of the algorithm. The current implementation is a 5th order filter where the coefficients are stored with 9-bit precision. Simulations were performed to estimate the influence of this implementation compared to the ideal case.

Five channels are simulated where the central channel is assigned a value corresponding to the expected signal size of 32 ADC counts. Typical values of +8% in the forward direction and -5% in the backward direction were added to the adjacent bins, corresponding to the values shown in Figure 2. The channels two steps away are assigned the cross talk values squared. Even if these values are very small, it is important to include them in the calculation of the coefficients. The length of the non-zero impulse response, called N in Section 3, determines the width of the region taken into account in the




Figure 4 The left-hand plot shows the simulated ADC values before FIR corrections. The right-hand plot shows the residual signal distortion after corrections with float and 9-bit precision.

minimisation (Equation 10). Hence keeping N to at least the order of the filter (M) ensures that the signal does not leak out of the minimised region after corrections.

The simulation is run with 10000 events and a Gaussian noise with a sigma of 1.5 ADC counts is added to all channels. The corrections are applied for each event in two different variants to compare the performance. One with full float precision and one with the 9-bit precision as implemented in the TELL1.

Figure 4 shows the the simulated signal before corrections and the residual distortion after corrections. The remainder is negligible when the calculations are done in float precision, but there is a small distortion of the signal with 9-bit precision. This comes from the rounding errors introduced by the bit truncations.

5 Test beam results of the FIR filter corrections

To evaluate the performance of the suggested digital filtering method, it was used to correct the data from the test beams in August and November 2006. This verifies both the method of determining the filter coefficients, as described in Section 3, and the application of the FIR filter on real data. The FIR correction of the beam test data is done off-line with a bit-perfect emulation of the processing occurring on-line in the FPGA. The filter applied was a 5th order filter calculated with nine bits precision.

The VELO test beam in November 2006 used a non-standard analogue link. For practical reasons the link consisted of 15m cables instead of 65m with analogue drivers modified to partially compensate for this difference. In this non-optimised setup the measured cross talk was in the order of +5% in the forward direction and -5% in the backward direction.

The main influence of the cross talk on the tracking performance comes from the skew it introduces in the reconstructed track position in the sensor. The signal induced by a particle traversing the sensor will be shared between several strips, and the track position is interpolated between the strip centres depending on the relative signals sizes of the strips in the cluster. Moreover it will depend on track parameters such as the track angle and its inter-strip position. Hence, the cross talk will shift the mean of the track residuals and depending on the exact nature of the cross-talk it may also degrade the resolution.

As seen in Figure 2, the cross-talk is normally asymmetric. The read-out order of the VELO R measuring strips depends on the location on the sensor, as shown in Figure 5. The shift in reconstructed track position will hence be shifted towards either larger or smaller radius, depending on the position. This effect is visible in Figure 6. It shows the mean of the track residuals from the VELO beam test in November 2006, with the different regions annotated with the same numbering as in Figure 5. The shift track residuals before FIR corrections correspond to the difference in read-out order on the R measuring sensor. The data was re-processed using an algorithm that emulates the 5th order FIR that is implemented in the pre-processing FPGA, applying coefficients determined by the method described in Section 3. Figure 6 shows how the effect vanishes after the filtering.




Figure 5 Schematic view of the read-out order of the different regions of a R measuring sensor. Note that increasing channel number means either increasing or decreasing radius depending on the position on the sensor.

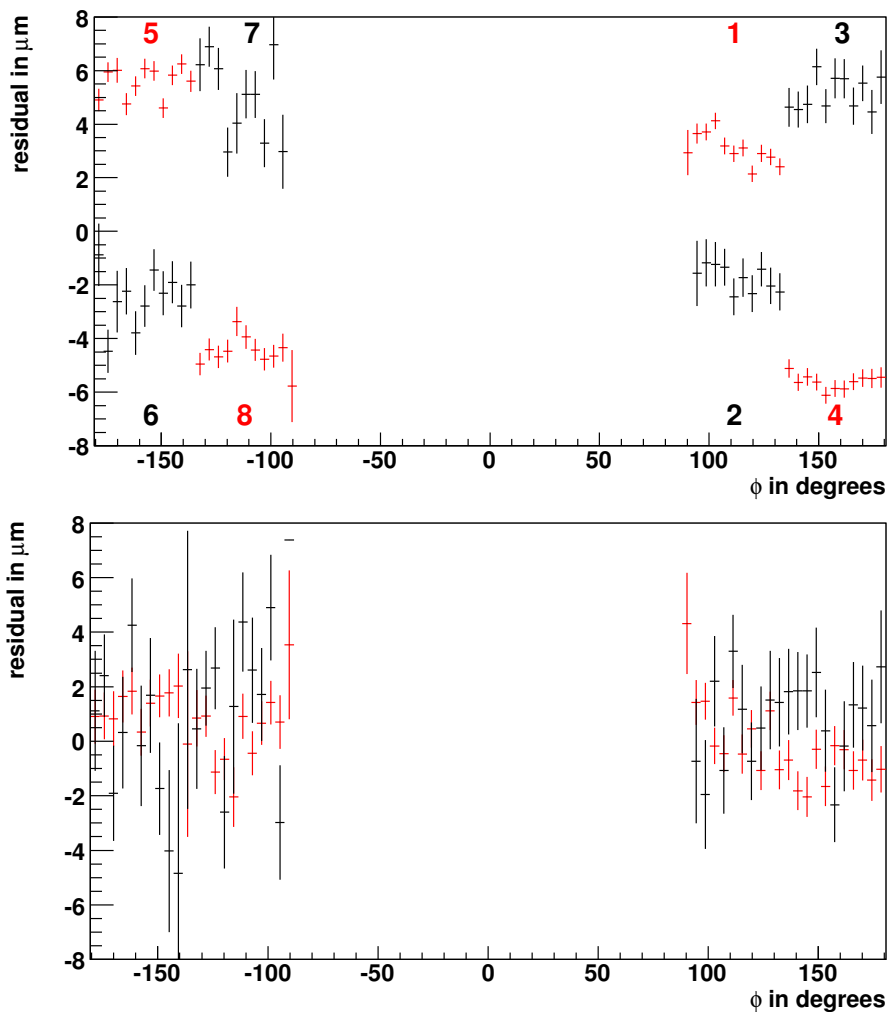


Figure 6 Top: Mean of unbiased track residuals for the eight regions on the sensor. The direction of the shift in the mean corresponds to the direction in read-out order shown in Figure 5. Bottom: Same after application of the FIR filter. The larger error bars are due to lower statistics of the re-processed data.

6 Conclusion

A method has been developed to determine the coefficients for the FIR corrections performed in the VELO read out boards. The method is robust and can be applied in-situ to the large number of analogue links present in the VELO. Simulations show that even with the approximations made the residual signal distortion is small. And when applying these corrections to beam test data, a measurable improvement in tracking performance is achieved.

7 References

- [1] "LHCb Technical Design Report, Reoptimized Detector Design and Performance," CERN/LHCC, 2003-030 (2003)
- [2] "LHCb VELO Technical Design Report," CERN/LHCC 2001-011 (2001)
- [3] S. Loechner, M. Schmelling, "The Beetle Reference Manual," LHCb Note LHCb-2005-105
- [4] G. Haefeli, A. Bay, A. Gong, H. Gong, M. Muecke, N. Neufeld and O. Schneider, "The LHCb DAQ interface board TELL1," Nucl. Instrum. Meth. A **560** (2006) 494.
- [5] M. Unser and M. Eden, "FIR approximations of inverse filters and perfect reconstruction filter banks," Signal Processing **36** (1994) 163-174