



UNIVERSIDADE DE SANTIAGO DE COMPOSTELA

Programa de Doutoramento en Investigación en Tecnoloxías da Información

Tesis doutoral

DEVELOPMENT OF ELECTRONICS FOR THE VELO UPGRADE DETECTOR.

Presentada por:

Antonio Fernández Prieto

Dirixida por:

Pablo Vázquez Regueiro

July 30, 2020



Pablo Vázquez Regueiro, Profesor contratado Doutor da Universidade de Santiago de Compostela na Área de Física de Partículas.

FAI CONSTAR:

Que a memoria titulada “**Development of electronics for the VELO upgrade detector.**” foi realizada por **Antonio Fernández Prieto** baixo a miña supervisión no departamento de física de partículas da Universidade de Santiago de Compostela. Así mesmo, expreso a miña conformidade co mesmo e autorizo a súa presentación ante a comisión examinadora do Programa de Doutoramento en Investigación en Tecnoloxías da Información, constituíndo así a Tese de Doutoramento para optar ao **Grao de Doutor en Enxeñaría**.

July 30, 2020

Asdo. Pablo Vázquez Regueiro

Asdo. Antonio Fernández Prieto

Dedicated to the memory of my father
Antonio Fernández Gómez

“Aut viam inveniam aut faciam.”

Hannibal Barca

“No es sabio el que sabe donde está el tesoro, sino el que trabaja y lo saca.”

Francisco de Quevedo

Acknowledgments

There are many people and institutions that I would like to thank for helping me. I hope not to leave anybody out, to all my sincere gratefulness.

I would like to start manifesting my gratitude to my supervisor Pablo Vázquez for his advice and patience, this thesis won't be possible without him.

I am pleased to be part of an extraordinary team as it is the VELO and LHCb group. It is extraordinary not only from the technical point of view but also on the personal level, as it was proven in the coronavirus crisis by creating the HEV [1]. I would also like to thank all VELO group. Special thanks to Karol Hennessy, who always helped me since the first moment I arrive at CERN. But, how can I forget the rest of the VELO people with whom we share a lot of work: Jan Buytaert, Paula Collins, Martin Van Beuzekom, Kristof De Bruyn, Dónal Murray and many others. I also want to thank the LHCb online CNRS and CPPM teams: Guillaume Vouters, Federico Alessio, Frederic Hachon, Paolo Durante, Luis Granado Cardoso and many others for their support and their prompt answers.

Last but not least, a special mention for the USC LHCb hardware team from whom I learned a lot, this is a well-balanced group where all the personal skills are complemented with the rest of the team: Abraham Gallas (he can get you anything), Pablo Vázquez (he will always think a new smart idea), Edgar Lemos (the tester and now the ambassador at CERN), Eliseo Pérez (Our “bonding man” and also a bit McGyver he only needs duct tape and cardboard to build a cooling system), Antonio Pazos (for whom there is no small footprint), Efrén Rodríguez (the rookie of the team) and Marcos Seco (even though he is not a hardware guy, he is always helping around and giving me more IPs).

I have to thank the Spanish Ministry of Economy and Competitiveness for funding my research.

ANTONIO FERNÁNDEZ PRIETO

Non quero esquecerme tampouco de todos os compañeiros cos que compartín o piso de Saint Genis Eliseo, Edgar, Julián, María, Antonio Romero, Miguel, Miriam, Carlos Vázquez, Óscar Boente, Bea, Pablo Baladrón e Asier. Mencionar tamén a todos os habituais das pausas para o café non citados antes: Juan, Cibrán, Nestor.

Sei que me queda moita xente no tinteiro coa que non tiveron o pracer de traballar pero grazas a todos os membros do IGFAE por axudarme sempre que o precisasei.

No terreo persoal quero agradecer principalmente a miña nai, Fina, que sempre me apoiou e a quen lle debo todo. Quero acordarme tamén dos meus amigos de Narón Alex, Iván e Adrián.

Antonio.

Contents

1	Introduction	1
1.1	Thesis content	2
2	The Large Hadron Collider and Large Hadron Collider beauty (LHCb) Experiment	3
2.1	The Large Hadron Collider (LHC)	3
2.1.1	LHC filling schemes	5
2.2	The LHCb Experiment	7
2.3	The LHCb upgrade	12
2.3.1	LHCb trigger and readout upgrade	13
2.3.2	LHCb on detector hardware upgrade	15
3	Vertex LOcator (VELO) upgrade	21
3.1	Planar silicon sensors	22
3.2	VeloPix Front-End (FE) Application Specific Integrated Circuit (ASIC)	23
3.2.1	VELO FE data framing	26
3.2.1.1	Pixel data	26
3.2.1.2	Special frames	27
	IDLE frame	28
	Sync frame	29
	Bunch ID frame	29
3.2.2	Timing and Fast Control (TFC) command response	29
3.3	On-detector electronics	30
3.3.1	VELO module	31

3.3.2	High-speed data tapes	32
3.3.3	Vacuum feedthrough board and Opto-Power Board	33
3.4	Cooling	34
3.5	Mechanics and RF foil	35
4	VELO upgrade Off-Detector Electronics	37
4.0.0.1	LHCb VELO data taking operation	39
4.1	PCIe40	41
4.2	LHCb VELO TELL40 firmware	43
4.2.1	Low Level Interface (LLI)	44
4.2.1.1	Intel Arria 10 Transceiver Native PHYsical layer (PHY) overview	44
4.2.1.2	VELO specific Physical Medium Attachment (PMA)	47
4.2.1.3	VELO specific Physical Coding Sublayer (PCS)	47
4.2.2	TFC	48
4.2.3	Pre-Router	51
4.2.4	Router	53
4.2.5	Post-Router	56
4.2.5.1	Data processing output format	59
4.2.6	Isolated Cluster Flagging (ICF)	61
4.2.7	Clustering	61
4.2.8	Event ID	62
4.2.9	MiniDAQ output (Multi Event Package (MEP))	63
4.2.9.1	Event packaging and multiple fragment packages	64
4.2.10	TELL40 error handling	65
4.2.10.1	LHCb framework error handling	65
4.2.10.2	VELO error handling	67
4.2.11	ASIC Identification	68
4.2.12	Monitoring	70
4.2.12.1	LLI Base Address Register (BAR) (BAR2)	70
4.2.12.2	General purpose BAR (BAR0)	72
4.2.13	Resources	73

4.2.14	Clocking	73
4.2.15	LHCb VELO bypass	75
4.2.16	Future developments	77
4.3	LHCb VELO control and timing distribution firmware (SOL40)	80
4.3.1	VeloPix Experiment Control System (ECS) link	82
4.3.1.1	VELO interface board BAR0	84
4.3.1.2	VELO FE transmission and reception operation	85
4.3.2	VeloPix TFC link	86
4.4	LHCb control software	87
5	Firmware development and hardware validation.	91
5.1	What is a Field Programmable Gate Array (FPGA)?	91
5.2	Firmware design flow	93
5.2.1	ModelSim/QuartaSim simulation	96
5.2.2	Intel/Altera Quartus design environment	96
5.2.2.1	Qsys	96
5.2.2.2	Chip Planner	97
5.2.2.3	Timequest	97
5.2.2.4	Hardware validation tools	97
5.2.3	Xilinx "ISE" and "Vivado" design environments	98
5.2.3.1	Intellectual Property core (IP) catalog and integrator	99
5.2.3.2	Timing tools	99
5.2.3.3	Hardware manager	99
5.3	Speedy PIXel Detector Readout (SPIDR) readout system	99
5.4	MiniDAQ 1 readout system	100
5.4.1	MiniDAQ server	102
5.4.2	MiniDAQ1 Data Acquisition (DAQ)	103
5.5	MiniDAQ 2 readout system	103
5.6	MiniDAQ 3 readout system	104
5.7	Hardware validation	105
5.8	VELO Tests with beam particles	115
6	Test benches for production quality assurance	119

6.1	High speed tapes test bench	120
6.2	High voltage tapes test bench	122
6.3	VeloPix Carrier board Quality Assurance (QA)	122
7	Conclusions and future work	125
8	Summary of the thesis	127
8.1	Context	127
8.2	LHCb upgrade	130
8.3	The LHCb VELO upgrade	133
8.3.1	Underground electronics	133
8.3.2	Electronics at the ground level	135
8.3.2.1	VELO control firmware	137
8.3.2.2	VELO readout firmware	137
8.3.3	QA firmware	138
8.3.3.1	Bypass firmware	138
8.3.3.2	Test-bench for High Speed and High Voltage tapes	139
9	Resumo da tese en galego	141
9.1	Contexto	141
9.2	Mellora do experimento LHCb	144
9.3	Mellora do detector de vértices do LHCb	147
9.3.1	Electrónica na caverna	148
9.3.2	Electrónica na superficie do experimento	150
9.3.2.1	Firmware de control do VELO	152
9.3.2.2	Firmware de adquisición de datos do VELO	152
9.3.3	Firmware para controis de calidade e desenvolvemento do hardware	153
9.3.3.1	Firmware de bypass	154
9.3.3.2	Banco de probas dos cables de alta velocidade e alta tensión	154
A	Old TELL40 output data format	155
A.1	Old data formats	155
	Bibliography	157

Acronyms	162
List of Figures	169
List of Tables	175

CHAPTER 1

INTRODUCTION

The Standard Model (SM) of particle physics is a theory that describes the electromagnetic, weak and strong interactions, explaining most of the phenomena observed in nature. The SM declares that the material in the universe is built of two types of particles: fermions and bosons. Fermions, which comprise all quarks and leptons, are the building blocks of matter and they follow the Pauli exclusion principle. On the other hand, bosons are the force carriers that act as the “glue” that holds matter together, they carry the electromagnetic, strong, and weak forces among fermions.

Unfortunately, the SM can not formulate some phenomena, such as: the gravitational interactions, why do the 3 flavour families (explained in the previous paragraph) exist?, why this structure? or how to explain the asymmetry between matter and antimatter in the universe, etc. This drives many New Physics models along the last years and also high energy physics experiments to explore these models.

In the area nearby Geneva (Switzerland) is the seat of the European Organization for Nuclear Research (CERN) that holds several high energy physics experiments. One of CERN facilities is the LHC, a 27 km circular accelerator designed to collide protons at a nominal centre of mass energy of $\sqrt{s} = 14\text{TeV}$. The LHC has 4 interaction points, located where the four experiments are built ATLAS, CMS, LHCb and ALICE.

This thesis is intended to be a guide that allows the understanding of the LHCb VELO upgrade electronics from both hardware and firmware point of view. It will also explain the integration of this system with the rest of the LHCb sub-detectors. The work presented in this thesis is the fruit of a large collaboration that goes beyond VELO collaboration and also

involves common developments for the whole LHCb experiment. However, this thesis does not cover other minor tasks carried out within the framework of VELO upgrade, but which were not the central tasks of the thesis: shifts on the LHCb experiment, total ionizing dose and single event effects tests.

Instituto Galego de Física de Altas Enerxías (IGFAE) group, to which I belong since the beginning of this thesis, is involved in the design of the back-end readout and control firmware, testing of the front-end ASIC, validation and integration of the different electronic components of the VELO system. IGFAE is also involved in the production and QA for the VeloPix carrier board used for the testing of the front-end ASIC, high speed and high voltage flexible tapes.

1.1 Thesis content

This thesis is organized as follows:

chapter 2 This chapter describes the LHC operation and the different filling schemes, which is important to understand and design the readout electronics of the LHCb VELO upgrade that will operate at the LHC rate. This chapter also describes the reason and the LHCb experiment, the weakness and needs of the original experiment that motivated a major upgrade during the LHC Long Shutdown 2 (LS2) that is taking place during 2019-2020.

chapter 3 This chapter shows a complete overview of the LHCb VELO upgrade that is going to be installed in the pit during LS2 of the LHC with a special emphasis on the electronics design.

chapter 4 This chapter describes the design of the electronics placed at the surface LHC intersection point 8, LHCb surface.

chapter 5 This chapter explains, chronologically, the different setups used for the design of the VELO upgrade electronics, culminating with a VELO test beam campaign in October 2018. This chapter also covers the methodology followed for the design of the firmware.

chapter 6 This chapter gathers all the setups designed at IGFAE to warranty the quality of the assembly components in which Santiago's group was involved.

CHAPTER 2

THE LARGE HADRON COLLIDER AND LHCb EXPERIMENT

2.1 The LHC

The world largest particle collider, LHC [2] is hosted by the CERN. It is placed in the surrounding area of Geneva (Switzerland). The LHC is a circular collider of 27 km built in the tunnels of the old LEP collider, placed between 45 m and 170 m under the ground level.

The particle acceleration mechanism is performed in several stages and it is supported by smaller accelerators. The LHC can accelerate different kinds of Ions, such as nuclei of lead, argon or xenon atoms. However, the most commonly used particles are the protons. In the case of protons, the acceleration mechanism starts by extracting them from a simple bottle of hydrogen gas. The extraction is performed by a duoplasmatron which consists of a metal cylinder with an electrical field applied, allowing the removal of the electrons present in the hydrogen. At this point, the protons leave the Duoplasmatron with 1.4% the speed of light (~ 4000 km/s) and enter in a series of accelerator stages: A 50 MeV LINear particle ACcelerator (LINAC), 1.4 GeV Proton Synchrotron Booster (PSB), 25 GeV Proton Synchrotron (PS), 450 GeV Super Proton Synchrotron (SPS).

Once the protons are accelerated at the SPS, they are finally transferred to the two beam pipes of the LHC, where they reach their maximum energy of 6.5 TeV per beam pipe and then they collide inside the four detectors present at the LHC ring (ALICE, ATLAS, CMS and LHCb) with an energy of 14 TeV at the centre of mass and a 30 MHz inelastic collision rate. The acceleration




Figure 2.1: LHC complex detector.

of the particles is achieved by using sixteen radio-frequency cavities distributed along the ring. However, the radio-frequency cavities accelerate the protons in a straight line. Thus, to maintain the protons, it is needed to bend the trajectory of the protons by using Niobium-Titanium (Nb-Ti) superconducting dipole magnets that have a nominal bending field of 8.65 T and operate below 1.9 K in static baths of pressurized super-fluid helium.

Each detector has its own physics target. While ATLAS [3] and CMS [4] are general purpose detectors, with a large physics programme, designed for the study of collisions that produce a high transverse momentum particles, the ALICE experiment [5] is devoted to heavy ion studies resulting from nucleus-nucleus collisions taken in the LHC lead special runs. Finally, the LHCb experiment [6], which is the frame of the content of this thesis (see section 2.2), was originally devoted to the CP violation and rare decays study. But, thanks to the excellent performance of the detector, the LHCb collaboration is able to deploy a larger physics program.

CERN experiments are not only the four shown above. There are multitude of small detectors placed along the LHC and the other pre-accelerators. The LHC also hosts TOTEM and LHCf experiments that focus on the forward particles and they are placed close to CMS and ATLAS re-




Figure 2.2: (Top left) ATLAS Experiment, (Top right) CMS experiment, (Bottom left) ALICE experiment, (bottom right) LHCb experiment.

spectively. MoEDAL, which is also placed in the LHC (around LHCb VELO), is used to search for a hypothetical particle called magnetic monopole. As shown in Figure 2.1, other experiments like a Radioactive Ion Beam facility (ISOLDE), neutron Time Of Flight (nTOF), fixed-target experiments (COMPASS, NA61, NA62, etc.) are placed along the rest of the accelerator complex. An Antiproton Decelerator (AD) is shown in Figure 2.1. It is used in combination with a new decelerator called ELENA to serve the experiments that study antimatter (ALPHA, ASACUSA, ATRAP and BASE).

2.1.1 LHC filling schemes

As it has been mentioned above, the LHC particle acceleration complex has multiple steps previous to the interaction of the two opposite beams. Due to the construction of the different accelerators, the beams are not a constant flow of accelerated particles, nor evenly distributed along time. The distribution of these particles with time is the so-called filling scheme of the LHC. The particles travelling the LHC are gathered in bunches, where the minimum separation

between the two colliding bunches is 25 ns.

These schemes have several constraints by the construction of the LHC:

- LHC clock that runs at $\sim 40\text{MHz}$ ¹.
- LHC ring length of 27.6km, which combined with the previous point, gives a theoretical maximum number of bunches present at the ring of 3564.
- The LHC injection also applies several constraints: a beam dump gap of at least $3\text{ }\mu\text{s}$

¹The actual LHC clock frequency can vary between 40.0784 and 40.0789MHz, but as a simplification, in this thesis, these clocks are referred as 40MHz or a multiple of it, for example, 160MHz clocks should be $\sim 160.3125\text{MHz}$.




Figure 2.3: 25 ns LHC Filling scheme.

to allow for the rise time of the dump “kicker” magnet², space between adjacent batches injected into the LHC must be greater than the rise time of the LHC injection “kicker” magnets ($0.95\mu\text{s}$) and the LHC injection “kicker” flat top cannot exceed $7.86\mu\text{s}$.

A filling scheme can vary with the particles present in the LHC [7] (75 ns, 43 bunch, Ions scheme, etc.), however, the most commonly used and also the most demanding in terms of frequency is the 25 ns filling scheme that uses 2808 of the 3564 slots, as can be seeing in Figure 2.3.

The coincidence of the beams in the four experiments could have four scenarios: beam-beam, beam-empty, empty-beam, empty-empty. Table 2.1 shows the number of beam-beam collisions on each interaction point per orbit³ for the 25 ns filling scheme. The different number of collisions is caused by the different gaps in the filling scheme. The number of collisions of ATLAS and CMS interaction points is coincident because they are built on opposite sides of the LHC ring, while ALICE and LHCb are not symmetric and thus they have a different number of collisions. The number of collisions delivered is maximized for ATLAS and CMS, giving an average rate of collisions $\sim 30\text{MHz}$.

Interaction Point (experiment)	Number of collisions
IP1 (ATLAS)	2808
IP2 (ALICE)	2736
IP5 (CMS)	2808
IP8 (LHCb)	2622

Table 2.1: Number of collisions for the 25 ns filling scheme in the four interaction points.

2.2 The LHCb Experiment

In the proton-proton collision mode, the LHC produces a huge amount of $b\bar{b}$ pairs per second in the forward and backward direction of the beams. The LHCb detector [6] was designed as

²This type of magnets are used in the injection and extraction stage because they minimize the beam loss and they are very precise in terms of trajectory, phase and space. These magnets are a type of dipoles that can quickly switch a particle beam between paths. At injection, they merge the circulating beam with the incoming one by appending the comming in particles after the circulating bunches. At dumping, it does the opposite by ejecting the entire bunch of protons.

³LHC orbit is a full loop; in other words, the entire filling scheme time.

a forward arm spectrometer (Figure 2.4), covering an angle between 10 and 300 mrad. LHCb complements the other LHC detectors, as it covers a pseudo-rapidity range⁴ between 2 and 5 (Figure 2.5).

The LHCb experiment has an excellent tracking resolution, particle identification and a flexible trigger. The tracking is performed by the VELO and four tracking stations distributed upstream and downstream the magnet. Particle identification is carried out with two Ring-Imaging Cherenkov Detectors (RICH) detectors, electric and hadronic calorimeters and a muon detector.

A dipole magnet that covers the forward acceptance of the detector was installed to measure the momentum of the charged particles. The magnet consists of two trapezoidal aluminium coils bent by 45° at the two transverse sides, arranged inside an iron window-frame yoke. The magnet gap is wedge-shaped in both, the vertical and horizontal planes, following the detector acceptance.

⁴Pseudo-rapidity describes the range of particle angles relative to the beam axis that a detector can measure. It has a value 0 for the plane perpendicular to the beam and infinite for the beam axis.




Figure 2.4: LHCb detector.




Figure 2.5: (Left) Angular distribution of b and $\bar{b}$. (Right) Illustration of the acceptance complementarity as a function of pseudo-rapidity between experiments at the LHC.

The flexible trigger system consists of two stages [8]. At first, a hardware level trigger uses information provided by the calorimeter and muon systems to filter 1 MHz of the 30 MHz of the colliding proton bunches. All the data that pass through the first trigger, are sent to a CPU farm where the events are reconstructed and selected. Thus, the LHCb output rate is 5 kHz. LHCb experiment data taking has been successfully performed during LHC Run 1 and 2, collecting nearly 9.2 fb^{-1} .

Although this thesis is made under the frame of the LHCb upgrade, it is necessary to describe the hardware of the original detector to justify the reasons behind the upgrade. The main systems are:

Vertex Locator (VELO) [9] is the sub-detector closest to the interaction point in the whole LHC, 8.2 mm during data taking⁵. VELO is used to measure the trajectory of the particles close to the interaction point in order to distinguish between primary and secondary vertices. VELO sub-detector consists of 42 semi-circular modules with two $300 \mu\text{s}$ thick silicon strip sensors arranged in radial shape (measuring r) and azimuthal (measuring ϕ).

⁵During beam injection or unstable beam, the sensors are moved away from the beam axis to prevent any damage.

Tracking stations consist of two sub-detectors with different technologies. The first one [10][11] is based on a large-surface silicon microstrip detector, which may itself be divided into two: four layers before and twelve after the magnet. In order to give an idea of the surface covered by this sub-detector, the station upstream the magnet covers a 1.7 m^2 , whereas the downstream station covers 4.2 m^2 . The second sub-detector is a gas tracking detector [12] that allows the measurement of the momentum of the particles covering a large area. It is built by grouping straw-tubes filled with a mixture of Argon and CO_2 . These straw tubes are arranged in three stations of four layers each, similar to the silicon trackers.

RICH [13] is the sub-detector capable of distinguishing the pions and kaons in a range between 2 and $100 \text{ GeV}/c$ by measuring the Cherenkov light. This sub-detector is divided into two stations, one before the dipole magnet between VELO and the first tracking station and other after the magnet between last tracking station and the calorimeter. When the particles travel through a very dense gas (C_4F_{10} or CF_4), faster than the speed of light traveling through this gas, they emit a cone of light that is reflected on mirrors, driving the light outside of the detector acceptance for digitalization on hybrid photon detectors.

Calorimeter [14] system is designed to identify and measure the energies of the electrons, photons and hadrons. It consists of several layers: the Scintillating Pad Detector (SPD), the pre-shower detector, the Electromagnetic Calorimeter (ECAL), and the scintillating tile iron plate Hadron Calorimeter (HCAL). The SPD determines whether particles hitting the calorimeter system are charged or neutral, while the pre-shower indicates the electromagnetic character of the particle (i.e. whether it is an electron or a photon). They are used at the trigger level in association with the ECAL to indicate the presence of electrons, photons, and neutral pions which are inductors of an event of interest for LHCb.

The working principle of the different layers is similar. It alternates layers of metal, to initiate a particle shower, with a scintillating material and a photo-detector to convert the shower into an electrical signal. Each calorimeter subsystem is configured with the appropriate metal layer, scintillating material, photo-detector and the dimensions for optimal detection of every type of particle. Thus, SPD and pre-shower use a 15 mm thick scintillating pads interspaced with a 2.5 times the radiation length lead converter. Then, the light is collected using wavelength shifting fibres and driven to a multi-anode photo-multiplier tube in the periphery of the sub-detector. ECAL uses lead plates to create the

particle shower and scintillating plates with a different cell size, depending on the position. The cell granularity is the same as for the SPD and pre-shower. The light detection is performed on photo-multiplier tubes with an individually regulated high voltage base. Finally, the HCAL uses iron plates interspaced with scintillating tiles arranged in parallel to the beam pipe. Like the ECAL, the inner and outer regions use different cell dimensions.

Muon detector [15] is located at the far-end of the LHCb. It comprises five stations (one in front of the calorimeter SPD and pre-shower and the remaining four downstream), the four latest stations are interleaved with iron filters.

Each one of the five stations is divided into four regions with respect to the radial distance to the beam axis, all with the same acceptance. Inner regions have better granularity due to their higher particle density. The granularity is likewise better in the horizontal plane, to give an accurate measurement of the track momentum and transverse-momentum.

The different granularity choice drives the use of two different technologies for the muon detection: the first one is used in the inner region of the first station and it is called triple Gas Electron Multiplier (triple-GEM) and the second one, used for the rest of the detector, is called Multi-Wire Proportional Chambers (MWPC).

Online system ensures the transfer of data from the FE electronics to permanent storage under known and controlled conditions. This includes not only the movement of the data themselves, but also the configuration of all operational parameters and the monitoring of these, as well as environmental parameters, such as temperatures or pressures. The online system must also ensure that all detector channels are properly synchronized with the LHC clock. It can be decomposed into three sub-systems: **Data AcQquisition (DAQ)** that transfers the data belonging to an event of a LHC bunch crossing from the detector FE to permanent storage, **Timing and Fast Control (TFC)** that drives all stages of the data readout between the FE and the final storage by distributing the synchronous clock, resets and fast control commands, and **Experiment Control System (ECS)** that controls and monitors the operation of the detector.




Figure 2.6: LHCb upgrade detector.

2.3 The LHCb upgrade

Results from LHCb analyses so far, have mostly been consistent with the SM. However, the experiment needs to collect more data to reduce the current statistical uncertainties, and therefore continue with the quest for physics Beyond the SM (BSM). In 2019-2020, during the LHC Long Shutdown 2 (LS2), LHCb is suffering an important upgrade (LHCb Upgrade Phase I) [16]. The detector operated in Run I and II was not able to handle all the luminosity⁶ that the collider could provide, thus the LHC was working with offset collisions in LHCb interaction point to reduce the delivered luminosity. This offset will be kept for the upgrade I but the upgraded detector will be designed to operate at 5 times the actual luminosity, up to $2 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$, expecting an integrated luminosity of 50fb^{-1} by the end of Run IV (2030). Moreover, the events will be fully-reconstructed at the LHC inelastic event rate using a fully software-based trigger. All LHCb sub-detectors will be upgraded as well as their readout system in order to sustain the new

⁶The luminosity measures the number of collisions produced per cm^2 and per second.

experimental conditions.

2.3.1 LHCb trigger and readout upgrade

The main limitations of the original LHCb experiment was that the collision rate had to be reduced to a fixed latency 1 MHz rate. The cause of this reduction was the hardware trigger of the experiment. The low acquisition rate causes inefficiencies in the trigger chain, especially for purely hadronic decays. Thus, from the readout point of view, the primary objective is to implement a trigger-less hardware detector. To reduce the storage costs, all sub-detectors must perform zero-suppression before propagating the data to the DAQ system.

The whole detector will be operated through the optical fibre radiation hard Versatile Link [17], differentiating two kind of links: full duplex for control and timing information and sim-




Figure 2.7: LHCb upgrade Readout architecture.

plex for data acquisition. Timing and control information (TFC/ECS) is driven on the same link using the GigaBit Transceiver (GBT) protocol [18]. On the other hand, data links will be read out using GBT for all sub-detectors but VELO that will use an specific protocol Gigabit Wireline Transmitter (GWT) [19] (more details on chapter 3).

All Back-End (BE) electronics are designed to control, synchronize and read out the experiment. They are placed in the surface area of LHCb and they share the same hardware, a board called PCIe40 connected to a server. The functionality of this board is determined by its firmware.

Figure 2.7 shows the architecture for the upgraded LHCb where all FE electronics record and transmit data synchronously at 40 MHz to the readout boards through the versatile links.




Figure 2.8: LHCb upgrade trigger diagram.

Muon and Calorimeter detectors transmit in parallel event information to the high level trigger software, which will take the physics accept or reject decision per LHC bunch crossing and propagate it through the TFC system. This system is based on two different types of boards: a readout supervisor board (S-ODIN) and an interface board (SOL40). The first one acts as an interface with the LHC, generating the TFC commands of the detector. The second one handles the distribution to all FEs of two kinds of signals: the previously mentioned TFC and the control/monitoring commands, ECS, from the multiple computers used as Supervisory Control And Data Acquisition (SCADA).

The flow of the data and its various transformations are summarized in Figure 2.8. In this architecture all FE electronics record and transmit data continuously to the readout boards, where each sub-detector pre-process the data in real time in a so-called “data processing” block. After the sub-detector processing, the combined data from all links of the readout board (TELL40) are passed as a new fragment to the data-processing block which can contain sub-detector specific processing. This fragment is passed on to the MEP building, which combines several fragments from consecutive bunch-crossings into an event-block and transfers them to a ring-buffer on the TELL40 server host. The Event Building software consists essentially of two processes: a readout-unit which is sending the blocks and a builder-unit which is collecting the blocks and creating multi-events (MEPs).

The data are driven from the sub-detectors FE via a multi-Tb/s readout network into a computing centre where events will be fully reconstructed prior to the lowest level trigger decision. The estimated data stored after the software trigger will be around 10Gb/s.

2.3.2 LHCb on detector hardware upgrade

From the detector point of view, all FE electronics will be replaced to cope with the new readout scheme, a new tracking system will be installed as well as an improvement in the radiation tolerance for the particle identification detectors.

The modifications to each sub-detectors are:

Vertex Locator: In the case of this sub-detector, it will suffer major changes [20]. The sensor technology will migrate from silicon strip to pixel that gives a better spatial resolution and can also be operated closer to the interaction point. This, combined with the increase in the delivered luminosity by a factor 5, dictates that the FE electronic and sensor design




Figure 2.9: CAD model of the LHCb VELO upgrade.

must be radiation tolerant and the sensor bias voltage will be increased from 500V to 1000V. The entire sub-detector will increase its power consumption around a factor 2, hence a novel cooling design was developed in order to cool the FE electronics and keep a low material budget. As the sensor technology is changed, the shape of the VELO stations will be different as well, therefore, a new radio frequency foil to separate VELO and LHC vacuum was designed. A detailed description of the system is given in chapter 3 and chapter 4.

Tracking stations: As the VELO case, the LHCb will have a brand new set of tracking stations [21]. The replacement of the tracking stations is needed because the original LHCb, operated during Run I and II, was not designed to be enough radiation hard for the integrated radiation dose of Runs I to IV. Moreover, the strips geometry present in the previous tracking stations leads to a significantly high occupancy under the new running conditions. The upgraded tracking system will consist of 2 different sub-detectors: The Upstream Tracker (UT) that will replace the old first tracking station between the RICH1 and the LHCb spectrometer magnet. After the LHCb spectrometer magnet, the tracking will be done by the Scintillating Fibre Tracker (SciFi) that will replace the silicon and gas tracking stations.




Figure 2.10: (Left) Layout of the detection layers of the UT. Each rectangle represents a silicon sensor and diverse shadings indicate different sensor geometries. (Right) Sketch of a UT stave.

The UT sub-detector has four planes of silicon strip sensors as shown in Figure 2.10. The sensors have a finer granularity than the old ones to cope with the greater density of particles. The sensors strips are arranged vertically to increase the precision in the horizontal direction, which is the bending direction of the dipole magnet. The middle planes are angled by $\pm 5^\circ$ for stereo measurements. Each plane is built from units called “staves”. Staves placed far from the beam have fourteen sensor units (each alternatively mounted front and back) of a sandwich of foam and carbon fibre structure with embedded tubes acting as support and coolant. Central staves include four extra half-length sensor units in the region closest to the beam pipe. Each sensor unit contains one sensor, four or eight FE ASICs, a hybrid circuit and a stiffener. Each FE ASIC has 128 input channels with a 6 bit Analog to Digital Converter (ADC).

The SciFi sub-detector consists of three stations with four detector planes each, as shown in Figure 2.11. Each station is built from individual modules ($0.5\text{ m} \times 4.8\text{ m}$) comprising 8 2.4 m long fibre mats. The mats consist of 6 layers of densely packed scintillating




Figure 2.11: Schematic yz- and xy-view of one SciFi tracking station. It is composed out of 4 layers, two vertical (x) and two with a stereo angle of $\pm 5^\circ$ (u, v). Each layer is made of 10 or 12 individual fibre modules.

fibres with a diameter of $250\mu\text{m}$. The scintillation light is recorded with arrays silicon photo-multipliers, digitalized in a custom ASIC. To reduce the thermal noise of the silicon photo-multipliers are cooled down to -40°C .

RICH detectors: In order to allow operation of the RICH detector system at the upgrade luminosity, the first station of the RICH optical system is being redesigned: in particular, the focal length of the spherical mirrors will be increased from 2.7 m to 3.7 m to reduce the hit occupancy on the photo-detectors. The cooling system and the support mechanics will also be modified to allow stable operation and easier access to the detector during maintenance.

New photo-detectors (Multi-anode Photo Multiplier Tubes, MaPMTs) have been chosen and they will be read out using a 8-channels ASIC, named CLARO [22], designed to sustain a photon counting rate up to 40 MHz while minimizing the power consumption and the cross-talk. A 128 bit digital register allows selection of thresholds and attenuation values and provides features useful for testing and debugging. More detailed information in [23].




Figure 2.12: RICH1 has a vertical optical layout (y-axis), while RICH2 has a horizontal optical symmetry (x-axis). The two drawings are not to scale.

calorimeters: Calorimeters will only suffer minor changes. The main change is to rebuild the whole front end and back end electronics, along with the rest of the sub-detectors, to read out at the LHC frequency. Working at higher luminosity also forces the reduction of the Photo Multiplier Tubes (PMT) gain to ensure the lifetime during LHC Run 3. Finally, the PreShower and SPD will be removed since they are not used for the upgraded trigger.

Muon chambers: Muon system, as well as the rest of the particle identification system (RICH and calorimeters), will suffer only minor changes to read out the electronics at the LHC clock, will remove the first muon plane, and also additional shielding will be added around the beam pipe and in front of the second plane, M2, to reduce the occupancy generated in these regions by the escalation in fluence.

CHAPTER 3

VELO UPGRADE

The upgraded VELO sub-detector [20] will be a hybrid pixel detector consisting of 208 silicon sensors, each of those ball-bonded to 3 FE ASICs. The migration from the old strips technology into pixels is due to the better granularity and, therefore, tracking precision.

The closest sensor will be placed at 5.1 mm from the interaction point (3 mm closer to the beam than in the old sub-detector) within a secondary vacuum isolated from the LHC primary vacuum. Sensors are distributed in 26 stations along the beam in two opposite retractable (by 6 cm) halves as it is shown in Figure 3.1. Each station is divided up in 2 slices. Each of those includes the so-called module, containing the sensors and the FE electronics, and the auxiliary power and interconnection electronics. The absorption of the modules telescopic displacement is




Figure 3.1: Close-up view of the VELO sub-detector.

taken by the flexible data tapes and the powering cables.

The retraction of the modules keeps them far from the beam during the LHC injection, avoiding any damage in the sensors and FE. Once the beam is stable, the two VELO halves will close to the nominal beam distance to take data.

Based on Monte Carlo simulations, the highest occupancy ASICs will have pixel hit rates about 900 Mhits/s, which is equivalent to 520 Mhits/s in terms of cluster rates. This translates into a bit rate of ~ 16 Gb/s adding up to ~ 3 Tb/s data for the ~ 41 Mpixels of the whole VELO. One of the major challenge of the sensors and the readout chips is caused by their orthogonal arrangement to the beam pipe. The irradiation profile along the sensors is inversely proportional to the squared distance to the interaction point, with the greatest concentration of hits at the innermost sensor regions. Figure 3.2 shows an estimation of the VeloPix data rates in the hottest module¹, where the highest rate (red and yellow) is received in the region closest to the interaction point (pink dot).




Figure 3.2: VeloPix data rate [Gb/s].

3.1 Planar silicon sensors

The VELO upgrade will be made of hybrid pixels, consisting of a planar silicon sensor [24] connected to a three readout ASICs (section 3.2) by using flip-chip technology [25]. This technology allows the union between the sensor and the FE ASIC by using an array of small soldering bumps ($12 \mu\text{m}$ diameter) known as “bump bonds” shown in Figure 3.3.

The chosen planar silicon sensors are $200 \mu\text{m}$ thick Hamamatsu “n-on-p” optimized for electron collection in order to mitigate the timewalk effect². Sensor efficiency is higher than 99%,

¹The hottest module has the highest estimation of hits. This corresponds with the station number eight.

²This effect is caused by the different spectrum of the sensor collected charge, which generates a non-negligible spread in the time response at the input of the electronics pre-amplifier. Hence, hits with small collected charge generate signals with a slower rising edge, generating a late detection of the pixel hits.

reducing the risk of missing the first measurement points. The pitch of the sensor is compatible with the readout chip and with elongated pixels of $55 \times 137 \mu\text{m}^2$ in the region between ASICs. The sensor is built with a $450 \mu\text{m}$ guard ring and the deep reactive-ion etching needed for the foil clearance.

By the end of LHC Run IV (2030), the hottest sensor will be exposed to a fluence of $8 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$. Hence, they will operate at -20°C to protect the silicon from thermal runaway effects after irradiation. They also need to withstand bias voltage of 1000 V to provide enough signal at the end of their lifetime.




Figure 3.3: (Left) Hybrid pixel detector. (Right) VELO “Tile”: Sensor flip-chip to 3 VeloPix ASICs.

3.2 VeloPix FE ASIC

To read out the charge collected in the sensors, a new FE ASIC called VeloPix [26], based on the MediPix/TimePix family, has been designed to meet the very demanding readout rate, radiation

Technology	130 nm CMOS
Pixel matrix	256×256 pixels with a pixel size of $55 \mu\text{m} \times 55 \mu\text{m}$
Radiation hardness	4 MGy achieved using triplicated voting registers
Power consumption	$< 1.5 \text{ W/ASIC}$
Readout	Unsorted binary data driven at 40 MHz
Data throughput	20.52 Gb/s

Table 3.1: VeloPix ASIC specifications.

hardness and low power consumption requirements of the experiment (Table 3.1). A single sensor is bump-bonded to 3 VeloPix ASICs comprising a “Tile”. The VeloPix has to timestamp the collected data according to the 3564 bunch crossings of one LHC orbit, for that 12 bits are needed, but to optimize the data bandwidth, only 9 bits of tagging are transmitted. The full 12 bits are later reconstructed in the BE readout board (chapter 4).

The architecture of the ASIC is shown in Figure 3.4. The VeloPix transmits only zero suppressed binary data generated from time over threshold measurements in the pixel. On this architecture (from left to right on Figure 3.5) the signal enters in the pad from the bump bond connection to the sensor, it is amplified in the analog FE, digitized in the pixel level and grouped in clusters of 2 by 4 pixels (Super Pixel Packet (SPP)). The SPP data are vertically transmitted, with a clock rate of 160 MHz, over the matrix until they reach the End of Column block (EoC).




Figure 3.4: General architecture of VeloPix.

Then the data go to a central node where they are packed in groups of four SPPs, a bit of parity per SPP and a header of the full packet are added. The 128 bit frames are now sent into the ASIC transceiver where the data are scrambled following the polynomial $X^{30} + X + 1$ to maximize the transitions allowing a good data transmission³.




Figure 3.5: Schematic of the pixel and super pixel architecture.

With the aim of reducing the power consumption of the CERN standard serializer (GBT) and keeping in mind the radiation hardness of the ASIC, a readout link serializer block, named GWT [27] was designed. A single GBT serialiser and its corresponding high-speed Phase Locked Loop (PLL) requires ~ 300 mW whereas the GWT serializer uses a multi-phase delay locked-loop which consumes less than 100 mW for the equivalent functionality. The VeloPix ASICs are operated in vacuum and if we reduce the power footprint, the demands of the cooling system also decrease. Furthermore, the GWT serializer has a higher usable bandwidth of 5.13 Gb/s vs. 4.8 Gb/s for the GBT. This serializer works by feeding a 16 bit input round-robin multiplexer with a byte of data on both edges of a 320 MHz clock (8 times faster than LHC clock). The reading of this multiplexer is triggered by the signal generated on a multi-phase delay-locked

³The clock is transmitted embedded with the data and therefore a good clock recovery relies on these transitions.

loop, which generates 16 phases evenly spaced from the given 320MHz clock. The timestamp could also be encoded in gray code to reduce the digital activity in the ASIC and, therefore, reduce the power consumption.

The VeloPix has a slow control interface connected to the Experiment Control System (ECS) from which are accessible the configuration and monitoring registers. Full details about this interface can be found in section 4.3.

3.2.1 VELO FE data framing

The data generated by FE ASIC go to the BE board through optical fibre following the VELO-specific protocol GWT. The GWT output frame operates at the LHC clock frequency of 40MHz. Each frame consists of 128 bits, shown in Figure 3.6. This frame has a 4 bit header, always sending the binary value “1010”, that allows the word alignment of the data on the back end side (subsection 4.2.1). Moreover, four 30 bit SPPs are sent along with their corresponding parity bits. Depending on the content of each SPP, the data transferred could be: Pixel data or special frames.




Figure 3.6: GWT data frame.

3.2.1.1 Pixel data

When pixel data are transmitted, each SPP field follows the format shown on Figure 3.7. Pixel data can be distinguished from special frames because they have a hitmap field different from 0. From left, Most Significant Bit (MSB), to right, Less Significant Bit (LSB) the information transmitted is:

- 13 bits define the hitted Super Pixel (SP) address (upper 7 bit for the column and lower 6 bit the row). Pixel row = SP row x 4; Pixel column = SP column x 2.
- 9 bits of timestamp (Bunch Cross Identifier (BXID)) optionally encrypted in gray code.




Figure 3.7: SuperPixel Data format.

- 8 bits define the hit pixel in a super pixel. Pixels are geometrically ordered, as shown in Figure 3.8.




Figure 3.8: Hitmap pixel ordering.

The 9 bits of timestamp correspond to the lower nine bits of the full 12 bit LHC BXID. Special frames with the full 12 bit BXID can be sent from VeloPix, described in the following subsection.

3.2.1.2 Special frames

Valid SPPs are always indicated by a non-zero hitmap. All special frames, which do not contain actual event data, are indicated by zero hitmap. Also, valid SPPs don't have a fixed header at position [29:26], but all special control packets do in order to distinguish them more easily. Since special frames do not have to traverse the pixel columns, they do not suffer from the large latency variances of the SuperPixel hit data. Some small variances may remain, however, as the output to the BE is never guaranteed to be fixed latency. Table 3.2 shows the full list of special frames.

Header [29:26]	Configurable packet data[25:8]	Hitmap [7:0]	Packet type
0x0	6'h0, BX ID[11:0]	8'h00	BX ID
0x1	Configurable[17:0]	8'h00	TFC Sync
0x2	3'h0, chip ID[14:0]	8'h00	Chip ID
0x3	12'h0, TFC bits[5:0]	8'h00	TFC Align Mode
0x4	Configurable[17:0]	8'h00	Invalid (Idle)

Table 3.2: configurable FE SuperPixel Packet frame.

The special sub-frames are defined per SPP, so one can construct a complete GWT frame with four different special SPP sub-frames. The most useful would be with a BXID packet (0x0), a TFC Sync (0x1), a Chip ID packet (0x2), and a TFC Align Mode packet (0x3). The Sync mechanism is described below. The Chip ID is a unique VeloPix identifier (used to confirm fibre/cable mapping). The TFC Align Mode is used to identify which TFC bits which have been set in the TFC word.

IDLE frame

By default, when no valid hit data are arriving at the GWT serialiser, “Invalid packets” are sent with the SuperPixel header 0x4. See Figure 3.9. Note that unlike the GBT protocol, which has a different header for valid and invalid data, the GWT header remains fixed (binary “1010”) at all times, irrespective of idle or valid data. A mixture of idle and valid SPP sub-frames is allowed in a GWT frame. For example, if only one SPP arrives at the GWT serialiser in a 25 ns window, then the GWT frame will consist of three idle sub-frames and one valid data SPP sub-frame. The position of the valid sub-frame is not fixed and it is a function of the state of the end-of-column routing logic of the VeloPix (in short, the GWT frame is *not* filled SPP0 first, then SPP1 ...).




Figure 3.9: Idle frame.

Sync frame

When the VeloPix receives a TFC Sync message, it sends a corresponding TFC Sync frame. This is configured as at least one of the 4 available SPP slots. The TFC Sync frame would correspond to a header = 0x1 and configurable pattern 18h'2BABE by default. The remaining 8 bits are set to 0 to indicate that the packet is special as previously mentioned. The other SPPs can be configured to send VeloPix Chip ID, full 12 bit SPP as seen in Table 3.2.

The TFC Sync mechanism is required to time-align the FE and BE. The BE requires a sync frame to initiate valid data taking. Before the Sync frame, the BE cannot determine the correct allocation of a single SPP within an orbit, as only 9 of the 12 bits are transmitted. Subsequent Sync frames can be used as a cross-check of the FE to the BE. The 9 bit SPP imposes a maximum latency constraint on the output of the VeloPix. Any data delayed by more than 512 clock cycles, 9 bits, will be incorrectly timestamped at the BE. However, Monte-Carlo studies have shown that the fraction of VeloPix data exceeding this requirement is negligibly small.

Bunch ID frame

Bunch ID counter provides a global timing reference across the VeloPix. The counter runs at 40 MHz offering a maximum dynamic range of 9 bits for time stamping data packets. The full 12 bits of the counter can be queried with a TFC command Snapshot, or embedded into TFC Sync frames as described in the previous section.

3.2.2 TFC command response

All LHCb FE electronics receive TFC commands described in subsection 4.2.2. The transmission of these commands from the BE board is shown in subsection 4.3.1. VeloPix ASIC only reacts to the following commands:

BXID reset. This command is sent every turn of the LHC, it resets the 12 bit bunch counter.

FE reset. This command performs a reset of the VeloPix without touching the bunch counter.

It is asserted at the same time as the “Header only” and “BX Veto” commands. Moreover, a FE reset is always followed by a “Sync” command.

Sync. The VeloPix response is described above in subsection 3.2.2.

Calibration. Upon receiving a TFC calibration command, the VeloPix injects test-pulse data into the matrix. The exact configuration of which pixels are activated by the testpulses can be configured through the slow control interface (ECS).

Snapshot. This command is only be dealt over the ECS data path and GWT/DAQ data are unaffected. All counters are transferred to a set of snapshot registers on the current SPP and can be read via ECS and compared to the expected value. This command does not affect the data taking.

Shutter. This is not a LHCb TFC command, but this signal is transmitted to the VeloPix enabling the GWT data adquisition. A known feature of the VeloPix is that during the transitions of this signal, noise will be transmitted through the GWT links, and thus it is mandatory to reduce these transitions.

3.3 On-detector electronics

Figure 3.10 shows a VELO slice, while Figure 3.11 is its schematic. From right to left (in Figure 3.10), what can be seen is: a VELO module, the high-speed tapes and low and high voltage cables, the Vacuum Feedthrough Board (VFB) and the Opto-Power Board (OPB). The high-speed tapes [28] have low mass to reduce the radiation length in the detector acceptance and they are flexible enough to allow the detector opening during LHC beams filling without




Figure 3.10: VELO slice.

compromising signal integrity. High-speed transmission lines travel almost 1 m before reaching the optical converters. A feedthrough board is needed to make the electrical connection between vacuum and air-side electronics. All the electrical readout and control signals are converted into optical in the OPB and driven 300 m to the FPGA based BE board on the surface.

3.3.1 VELO module

A single module (Figure 3.12) is equipped with 4 sensor tiles on a silicon substrate, 2 per side, with microchannel evaporative CO₂ cooling (section 3.4). Every tile contains 3 VeloPix ASICs, with a peak module occupancy of 61 Gb/s. The variability in the hit distribution causes that the VeloPixes placed close to the interaction point receive more hits, thus the number of 5 Gb/s readout links on these ASICs is larger than the periphery ones. Monte Carlo simulations show the ASIC data rate of Figure 3.2. Hence, the nearest ASICs to the interaction point that has a data rate of 15 Gb/s uses 4 readout links, ASICs in the middle region of the module that expect an occupancy of 6 Gb/s use 2 links, and the periphery ASICs with less than 3 Gb/s of occupancy use only 1 link. Ten is the aggregated number of links of a single module. All of them share the




Figure 3.11: LHCb VELO electronics overview.

same off-detector readout board, TELL40 (section 4.2).

Each side of the substrate has a Giga Bit Transceiver ASIC (GBTx) ASIC [29] that controls and distributes the timing signals to the 6 VeloPix through a full-duplex link. The electronic design of the module was made to reduce the effect of the mismatch in the coefficient of thermal expansion and using the minimum mass in the acceptance of the sub-detector. Consequently, two kinds of Printed Circuit Board (PCB) interconnected by Kapton tapes were designed, one for hosting the GBTx ASIC and one for the VeloPix wire-bonding pads (Figure 3.12).




Figure 3.12: VELO module.

3.3.2 High-speed data tapes

All module readout and control signals are routed inside the vacuum using low mass high-speed flex tapes with Kapton microstrip technology in the acceptance region where low mass materials are critical for the detector performance (~ 5 cm). Four flexible PCB tapes built in stripline technology drive up to 7 links. The dielectric used for these flexible tapes has to be radiation tolerant and provide low dissipation loss for high frequency signals. Besides, high reliability and yield would be required for the impedance control. The best candidate material

found, Pyralux®AP Plus [30], was retained. It is also widely used in medical, aerospace and military applications.

3.3.3 Vacuum feedthrough board and Opto-Power Board

The VFB is placed at the end of the high speed tapes. This board has to deal with a pressure of $\sim 10^{-7}$ mbar on its inner side and the atmospheric pressure of the ambient air on the outer side. To help handling the difference of pressure the PCB was designed with edge metallization.

On the air-side of the vacuum feedthrough is connected to the OPB through a PCI express connector. The main functions of the OPB are: to provide low voltage to the FE electronics, to monitor the FE temperatures and to convert the signals from electrical into optical and vice-versa.

All the active components used on the OPB board are radiation tolerant: the FEASTMP DC/DC converters [31], the Versatile Twin Transmitter (VTTx) and Versatile Transceiver (VTRx) [17] electro-optical converters, the GBTx gigabit transceiver and the GBT Slow Control Adapter (GBT-SCA) ASIC [32] which is used for control and monitoring.




Figure 3.13: (Left) Cross section scheme for the microchannel cooling with a superposition of the VeloPix ASICs. (Right) Schematic of the microchannel attachment of the fluidic connector.



region and clearance to th

module contains 12 VeloX detectors. To prevent the prevention of the thermal damage caused by the irradiation suffered in the previous experiments, the operating temperature of -20°C was chosen. It is designed to fulfil the cooling requirements of the detector [33]. Figure 3.13 shows the layout of the detector consists of making circular pads on the silicon substrate. The pads have a diameter of $60\mu\text{m}$ and the main region is used for the boiling of the CO_2 , keeping the temperature below the boiling point of the liquid. It has been selected due to the compatibility with silicon ASICs and the low cost of the material budget.




Figure 3.15: Pie chart showing the contributions from different components to the total material budget in the upgraded VELO within the pseudorapidity angles 0.8° and 15.4° .

3.5 Mechanics and RF foil

The VELO sub-detector is a 1.1 m long piece surrounding the LHC point 8 collision point, divided into two movable halves facing each other. Figure 3.14 shows the aluminium foil of one of the halves. The subdetector must be placed in an isolated vacuum from the LHC primary vacuum to avoid compromise the accelerator operation. Thus, an isolation aluminium foil must withstand a differential pressure of 10 mbar between primary and secondary vacuum. The foil also shields the detector from radiofrequency pick-ups from LHC beams.

From the mechanical point of view, the aluminium foil for the upgrade I must accomplish several constraints:

- As it can be seen in Figure 3.15, the foil thickness has a significant effect on the impact parameter resolution. Hence, to reduce the multiple scattering effects, the thickness is reduced from $300\ \mu\text{m}$ in the old VELO to an uniform thickness of $250\ \mu\text{m}$, chemically etched down to $150\ \mu\text{m}$ in the region closest to the sensors.
- As the upgraded detector will operate closer to the beam (from 8.1 mm to 5.1 mm), the clearance of inner aperture of the foil must also be smaller, going from 5.5 mm to 3.5 mm.

- Inner clearance between the foil and the sensors must be $900\text{ }\mu\text{m}$.

CHAPTER 4

VELO UPGRADE OFF-DETECTOR ELECTRONICS

LHCb Off-Detector Electronics (ODE) comprises all the systems placed in the surface of LHC interaction point 8 that are responsible for control, timing and readout of the underground experiment. A new data centre [34] is being built to house all ODE within a cable distance around 300 m from the detector. All data transfer between the data centre and the detector is performed over optical links. Figure 4.1 shows the facilities of the new data centre.

The LHCb upgrade control and readout scheme was introduced already in section 2.3. The aim of this chapter is to go in-depth into the LHCb electronics on the surface, with an especial emphasis on VELO control, timing and DAQ readout boards.

Already explained in the detector overview, the event readout acquisition of the experiment is managed centrally from the readout supervisor board (also known as S-ODIN). Interface boards (SOL40) and Readout Boards (TELL40) work as slaves of S-ODIN.

Due to the centralized nature of the TFC system, the distribution of the signals presents some challenges in terms of timing distribution, clock recovery, jitter, synchronization and robustness of the system. TFC synchronization between a LHCb experiment and the accelerator is done centrally in the S-ODIN boards (Figure 4.2). To ease the detector operation, each sub-detector has one or several detector partitions controlled by a single S-ODIN board. Each S-ODIN is connected to the LHC independently, allowing better robustness. An extra board supervises all the experiment partitions. Each partition S-ODIN recovered the clock, cleaned and fanned it out. In addition to that, it has to monitor and control the system allowing to each partition to be



Figure 4.1: New LHCb data centre and the optical fibre cabling to the cavern.

aligned by itself. The TFC also has to be aligned with the LHC beam structure, as explained in subsection 2.1.1, with a tolerance below 100 ps.

As all communications are done via optical links, the transceivers in the cards that propagate the TFC information need to work in a way that it maintain the LHC clock quality transmission, phase and latency. All these signals have to reach ~ 2500 FE and 500 BE destinations.

The synchronicity between all the units that make up a partition is guaranteed over Passive Optical Network (PON) technology [35], where a passive optical splitter allows to reach many destinations from a single point of start without re-driving the stream. As the clock, timing and readout control signals are centrally generated and they are the same for all destinations, they can be distributed passively.

Figure 4.2 shows the architecture of a TFC partition, where the TFC information is centrally generated and synchronous with the LHC clock in S-ODIN. Data are sent over a timing-trigger and control PON splitter and connected to the different BE boards located in the surface of the experimental area (SOL40 and TELL40 boards). From the SOL40s, the ECS and TFC signals are distributed to the FE and the TELL40s collect FE data and make a real time pre-processing

of the sub-detector events before combining this data in the event building computers.




Figure 4.2: Architecture of the TFC system for a sub-detector partition

To implement the control and readout architecture of the LHCb experiment a common PCIe mezzanine board, called PCIe40 (Figure 4.4), was designed. Depending on the firmware flavour, it could be configured as S-ODIN, SOL40 or TELL40. This chapter describes the BE board hardware and firmware designs needed for operating the VELO sub-detector.

4.0.0.1 LHCb VELO data taking operation

The synchronization of all detector electronics is a fundamental task of the readout electronics. This is carried out by the TFC system from the readout supervisor board. Later in this chapter, there will be an in-depth description of the different synchronization commands that concerns the VELO FE control and the readout board, but in this item, we are going to describe the big picture for the synchronization and the starting of the data taking from VELO BE electronics point of view.




Figure 4.3: LHCb TFC starting sequence.

The LHCb data must be tagged with a unique identifier, BXID, that is restarted every 3564 40MHz clock cycles with the “BXID Reset” signal.

Starting the data taking requires a sequence with multiple steps as it is shown in Figure 4.3. At the beginning, the “shutter” for the data acquisition is unset to prevent useless FE data to be sent to the BE readout boards, for VELO case the shutter signal is asserted from a register on the interface boards. At this point, all the FE and BE Finite State Machines (FSMs) are driven to the initial state by sending their corresponding “Reset” commands, after that a configurable time-window with the FEs sending header only data¹ to configure the links to the BE and waiting for the BXID reset to guarantee that all FE counters are synchronous. Now, the synchronization between the VELO FEs and BE readout boards takes place. On one side, each FE receives the synchronization command generated in the S-ODIN board and responds with a specific TFC frame through the readout links (see subsection 3.2.2). On the other side, the readout board, which by default is on idle mode (not taking data) waiting for the FE specific frame to lock to its BXID. Afterwards, a programmable number of header only data, the shutter is open again to start physics data taking, triggering on the proton-proton events based on the filling scheme of the LHC.

¹Header only data mode allows the FE links to stay active without driving valid data to the BE readout board.




Figure 4.4: PCIe40 board fully populated.

4.1 PCIe40

The PCIe40 is a 14 layer PCB hosting an Intel Arria 10 FPGA (device model 10AX115S4F45E3SG) [36]. This PCB is meant to be used in the ODE of LHCb and ALICE experiments, but in this thesis, I will focus on the LHCb case. The board is populated with eight Multi-fibre Push On (MPO) optical connections with 12 links each, four transmitters and four receivers that handle up to 48 links running at 10Gb/s on each direction. The Opto electrical conversion is performed in the MiniPODs, 1 per MPO, and each link is connected to a single FPGA high speed input/output. The PCIe40 board is controlled and read out via two independent PCIexpress version 3 of 8 links each, with a maximum combined throughput of 128Gb/s and a sustained 108Gb/s [37]. Moreover, two 3.2Gb/s bidirectional timing-trigger and control PON links are mounted, allowing the interface of the timing signals of the experiment.

To improve the performance of the high speed data transmission on the FPGA, a very stable clock is needed, requiring the implementation of several external PLLs and Jitter cleaners (Si5345 [38]) that provide a clock with a jitter of 90 fs RMS, see Figure 4.5.

The functionality of the board is defined by its firmware. TELL40 and SOL40 firmwares are developed in an LHCb wide common framework with provision made for sub-detector specialization. The framework is focussed on the shared tasks for all sub-detectors like: input/output management, TFC interface, protocols handling, SCADA common developments, etc.

The FE data collected and processed in a single TELL40 board is propagated to the host computer through the PCIexpress connector. As it has two eight link interfaces, to reduce firmware complexity, two independent data-streams are instantiated, as it is shown for the VELO case in




Figure 4.5: PCIe40 clock distribution.

the Figure 4.6. Each of those framework data-streams has the following functionality blocks: GBT LLI that recovers the data from the optical fibre, error recovering block, a GBT decoding block, alignment of the different fibres according to their BXID information, TFC accept/reject of the data (and its associate communication with the system), sub-detector specific processing, identification of the events, event packaging and finally PCIeexpress interface.

The SOL40 board has a common development that handles TFC communication with S-ODIN, slow control communication with the host computer and FE (see section 4.3 and section 4.4).




Figure 4.6: VELO DAQ firmware architecture. Two parallel data streams are instantiated (each one associated with a PCIexpress interface).

4.2 LHCb VELO TELL40 firmware

The VELO readout board firmware scheme (Figure 4.6), flushed into the PCIe40 card (TELL40), is almost a full modification of the LHCb firmware framework. The rationale for adopting such a strategy is that the VeloPix ASIC is the only LHCb FE reading out unsorted data over a specific protocol, (section 3.2). Hence, a new firmware transceiver had to be designed to handle the VELO protocol. This protocol does not support forward error correction and, thus, it does not exist an equivalent to the framework GBT error correction block. Nevertheless, error detection is performed with a parity checker discarding corrupt data that do not match with their parity (more information in subsection 4.2.1.3). The main task of the readout firmware is sorting the BXID timestamp information, but also some auxiliary processing is performed by flagging isolated hits and clustering the data. Unfortunately, some of the common developments

are not useful for VELO. It is the case of BXID alignment block and TFC accept/reject block. The alignment and TFC processing are not possible as the FE data arrive unsorted, the solution adopted consists of: on one hand, a sorting block, so-called Router (subsection 4.2.4), that opens a time window of data sorting, big enough to do not need an alignment within a few clock cycles of different between input links. On the other hand, as the TFC processing information is related to the BXID, it must be performed after the Router. A small part of the TFC processing, related to the synchronization of the data, is performed before the Router.

Each VELO readout board can handle 20 DAQ links (as shown in Figure 4.6) and thus a single TELL40 board will handle the data from a module. That gives a total number of 52 boards needed for the entire VELO.


4.2.1 LLI

The LLI is the first layer of the firmware and it is responsible for recovering the data from the optical fibres, which is a customised and more extended version of the PHY of the OSI model. As the VELO FE ASIC uses the GWT serializer and protocol, a VELO FPGA LLI needs to be developed. For that, in this chapter the native tools of Intel Arria 10 are described in the first place and goes on showing how they were adapted for VELO.

4.2.1.1 Intel Arria 10 Transceiver Native PHY overview

The FPGA mounted in the PCIe40 board has 96 transceiver channels, distributed in 16 transceiver “Banks” of 6 links each, distributed as 50% input and 50% output. Figure 4.7 is an overview of the main blocks of any FPGA transceiver. It shows the transmitting and receiving side of the transceiver. For the transmitter side, the parallel data is encoded and scrambled according to the desired protocol in the PCS block, serialized in the PMA and finally leaves the FPGA through the high speed transmitter pin. The incoming serial data arrive in a first place to the Clock Data Recovery (CDR) where the clock from the link is recovered and the data grouped in a parallel bus in the de-serializer, then the data are driven through the PCS block for de-scrambling and decoding the data transmission protocol. Each block of Figure 4.7 is described below:

Physical Medium Attachment (PMA). Inside the receiver PMA, it is possible to differentiate three main blocks as shown in Figure 4.8.



Notes:

(1) The FPGA Fabric - PCS and PCS-PMA interface widths are configurable.

Figure 4.7: This figure shows a full duplex Arria 10 transceiver channel. (Top) Transmitter path. (Bottom) The incoming serial data.

- The “Receiver Buffer” is the analog part of the incoming link, which allows the equalization of the links. It is possible to tune the common mode of the signal, the value of the termination resistance, the DC and AC gain of the input Continuous Time Linear Equalizer (CTLE) filter. It is also possible to add gain to the analog signal with the Variable Gain Amplifier (VGA). Finally, a decision feedback equalization can be applied to the signal instead of the CTLE filter with the advantage of a better signal to noise ratio because it stores an historical of the previous data received.
- The CDR block works as a PLL with two input clocks feeding two independent phase detectors, a first phase detector with a very stable reference clock and the second one with the serial input data, which allows the CDR to have two different working modes: Lock To Reference (LTR) and Lock To Data (LTD). The output of those phase detectors is multiplexed and it feeds a charge pump and a loop filter that provides a stable voltage to a voltage controlled oscillator generating the output clock of the CDR. A second output clock is generated with a given factor slower than the serial clock, this clock is called “Parallel clock”. During normal behaviour,




Figure 4.8: Scheme of Intel PMA transceiver

the CDR starts in LTR mode, allowing all the rest of the CDR electronics to work. Once a signal is detected in the link, the CDR output clock is locked, with a stable frequency and the output clock in phase with the reference clock, then it switches to LTD mode. Finally, in case something goes wrong with the incoming data and the CDR loses the lock in LTD mode, the LTR mode will be automatically switched back.

- The de-serializer takes the serial input data from the receiver buffer block and clocks the data using the CDR serial clock. The parallel data are clocked using the CDR low-speed parallel clock. The de-serializer forwards the data to the receiver PCS.

For the transmitter side the mechanism is simpler. A serializer only needs a parallel and a serial clock in sync, similar to the receiving side. The serial data goes directly into a transmitter buffer for the analog circuit, which comprises a programmable differential output voltage, programmable pre-emphasis circuit and internal termination circuit.

Transceiver PCS. This block implements the encoding/decoding mechanisms of the most commonly used transmission protocols: 8b/10b used for Gigabit Ethernet data, 64b/66b used up to 100 Gigabit Ethernet, SATA, PCIe version 3, etc. This PCS also allows the bypass of the data from the PMA for a custom made PCS.

4.2.1.2 VELO specific PMA

The analog parameters for the receiver buffer should be optimized each time the physical link is modified². Nevertheless, the termination of the link is always fixed to 100 Ω . To provide a reference in the analog parameters optimization, for the test setup we decided to tune this parameters manually and, thus, we bypass the VGA and decision feedback equalization³, optimizing the AC gain to compensate the swing loss of the incoming link. The CTLE is configured in “full bandwidth high gain” mode, which allows data rates up to 17.4Gb/s.

As it was mentioned previously, every transceiver Bank shares a common reference clock. In the case of the PCIe40 board, this clock is provided by an external PLL and jitter cleaner (see section 4.1). This reference clock is set by design in the PCIe40 board to 6 times the LHC clock, an “optimal” frequency for the CERN GBT protocol that is used in the rest of the experiment. This frequency is “optimal” because the GBT protocol uses a 120 bits word running at the LHC clock, this gives PMA de-serialized output word of 20 bits. In the specific case of VELO, whose GWT protocol data width is 128 bits, it does not match for a clock 6 times the LHC. Hence, the solution adopted consisted in converting the external reference clock in an FPGA PLL to 4 times the LHC clock and to using a de-serialized width of 32 bits. Providing a very stable PLL clock is critical for the performance of the readout, this requires an extra step in the start up operation to calibrate the re-clocking.

4.2.1.3 VELO specific PCS

As it was mentioned in section 3.2, the VeloPix ASIC has its own data transmission protocol and therefore its own scrambler based on the polynomial function, $X^{30} + X + 1$. This custom protocol demands a dedicated PCS for the BE board. That is possible thanks to the bypass option of the native FPGA PCS and the development of a custom block.

The first stage of the VELO specific PCS is a word aligner block. Its function is to pile the 32 bit input word from the PMA. It looks for the header of the GWT frame (subsection 3.2.1) and, when more than a configurable number of consecutive frames with the header in the right

²This is the case of specific developments required for different test setups (These setups are described from section 5.4 to section 5.6). Also, due to the irradiation of the fibres during the operation of the detector, these parameters should be optimized.

³VGA and decision feedback equalization are anyhow accessible from registers for future optimization (if needed). After testing these PMA blocks with the MiniDAQ setup, no significant improvement was found and thus we decided to bypass them.

position are found, it declares that the alignment is locked and sends the data to the next stage⁴. At this point, it is time to look for a transmission problem by checking the parity bit of each SPP in the GWT frame. An error in the parity causes the ruled-out of the packet and an error counter increase. Valid SPPs are now de-scrambled, following the same polynomial as the VeloPix ASIC encoding.

In order to keep the link in a ready to receive data state, the FE ASIC sends idle GWT frames (section 3.2.1.2). These frames provide a balanced signal in the optical transmission and a proper clock recovery in the FPGA CDR⁵. However, from the data process point of view, these idle frames are useless and thus after the de-scrambling a zero suppression block was created to propagate only meaningful data. Along this whole chain, a series of registers monitor and configure different stages of the chain like the GWT frame inversion, the required number of clocks for locking and unlocking the link and parity errors, etc.

For the transmitter side, a basic PCS was developed in order to loop-back test functionalities. Even with the bypass of the native PHY transceiver for the Arria 10 we can make use of the internal FPGA memory and some embedded features to test the quality of the link by generating and checking patterns like PseudoRandom Binary Sequence (PRBS) 7, 15 or 31.

4.2.2 TFC

The BE readout board receives timing information from the readout supervisor board once per bunch crossing and stores it in an internal Random Access Memory (RAM) (TFC buffer). Unlike the LHCb common framework TFC processing, The VELO's TFC (Fig. Figure 4.9) only triggers the data acquisition when the Pre-Router receives a certain number of synchronization commands from the FE. Once the synchronization command is received, its BXID information is extracted and sent to the Post-Router in order to start the data acquisition sequence. In parallel, other kind of information of the link like the BXID, the Chip ID, and the status of the VeloPix ASICs are gathered and sent downstream to the data path (in parallel to the Router) and stored in dedicated registers. The Post-Router collects the data from the Pre-Router and stores them in a First In First Out (FIFO) memory. At the same time, a synchronization flag from

⁴By default, the constant values to declare the link as locked or unlocked are the same as CERN GBT FPGA transceiver: 23 consecutive words are requested to set the link and 4 of the last 60 ones to unset it.

⁵To recover the sampling clock, receiver needs a reference clock of approximately same frequency. To generate the recovered clock, the receiver needs to phase align the reference clock to the transitions on the incoming data stream. Hence, it is essential to feed the CDR with a well balanced signal in terms of transitions.

the Pre-Router is used to trigger the reading of the TFC buffer RAM. This buffer stores 3564 64 bit words with the timing information (Table 4.1). The Post-Router reads the TFC buffer in groups of 512 words once every $12.8 \mu\text{s}$, starting from the BXID number received from the Pre-Router. From the content of the TFC buffer, the Post-Router is responsible to manage the “Header only”, “BXID VETO” and “Trigger” commands, creating a 512 bit mask that enables the reading of the RAMs that stores the data in the Router (subsection 4.2.4). The TFC word is also synchronously propagated with the data.

Bit	63-52	51	50	49	48-43	42-36
TFC Data	BXID(11..0)	External trigger	MEP accept	Stop run	Step run counter	Run number
35-34	33-18	17-14	13-10		9	8 7
Bunch cross type	MEP destination address	trigger type	Calibration type	Synchronization command	Snapshot	Trigger
6	5	4	3	2	1	0
Bunch count veto	Non zero suppress mode	Header only	TELL40 reset	FE reset	Event ID reset	BXID reset

Table 4.1: 64 bit TFC word. It contains all the BXID commands distributed across the experiment that keeps the data taking synchronous.

The VELO data processing reacts to the following TFC commands:

BXID Reset. This command ensures that the synchronicity of the system is maintained. It resets bunch counters once per orbit . This reset is locally aligned to flag the crossing of bunch 0 (up to 3563) of Beam 1 at the LHCb interaction point as BXID=0. In the case of VELO BE firmware, it acts in both Router and Post-Router.

Synchronization (Sync). The previously described synchronization mechanism requires the TELL40 board to detect the synchronization identifier from the VeloPix to adjust its internal counters and to start the data taking (VELO FE response to a Sync command is described in subsection 3.2.2). The TFC Sync mechanism is required to time-align the FE and BE. The BE requires a synchronization frame to initiate data taking. Before the Sync sub-frame, the BE cannot determine the correct orbit of the BXID. Subsequent, Sync sub-frames can be used as a cross-check of the FE to the BE. The 9 bits BXID imposes a maximum latency constraint on the output of the VeloPix. Any data delayed by more than 512 clock cycles (2^{9bit}) will be incorrectly timestamped at the BE. The adjustment of the 3 bits, which are omitted in the data transmission, can only be made by the Sync command. It is processed at the Pre-Router level.

Header only. When the FE electronics receive, for a given BXID, a header only TFC command they send exclusively header information in the data corresponding to this BXID. This feature allows to run the system in an idle state, for instance, during the configuration or reset of the readout electronics or the Event Filter Farm, when the run is stopped or for special running modes which require a lower rate of events such as in the non zero suppression readout mode described below, or a time alignment event mode in which consecutive events are processed together.

For the case of VELO BE firmware, this signal acts in the Post-Router disabling the reading of the SPP RAM containing the event information (subsection 4.2.5).

The Header Only is handled identically to the bunch crossing VETO. These commands are ignored by the VeloPix so they have to be processed by the TELL40. The Router processing is required to determine the full BXID of each SPP and therefore, the VETO must come in the Post-Router stage of the data processing. All the latency requirements of standard SPPs will apply to the bunch crossing VETO and Header Only data (more information about the BE board latency in subsection 4.2.4 and subsection 4.2.5). A single empty Event ID block (256 bit) is propagated synchronously with the BXID and other metadata.

Bunch crossing VETO. When the FE electronics receive, for a given BXID, a VETO TFC command they send exclusively header information in the data corresponding to this BXID. The signal exclusively based on the LHC Bunch Filling Schemes, is used to reject the empty bunch crossings.

Calibration. Upon receiving a TFC calibration command, the VeloPix injects test-pulses into the matrix. The configuration of which pixels are activated by the test pulses can be configured over ECS. The VELO firmware makes no special treatment of calibration data and simply treats them in the same way as regular “physics” data.

Snapshot. It triggers all status and monitoring counters in the FE electronics to be sampled in ECS registers. The information may then be read via the control system. This allows to gather a consistent and instantaneous picture of the system across sub-detectors.

For the case of VELO BE firmware, this signal will save in registers all the counters present in the Router and Post-Router blocks.

This command is only processed over the ECS data path while data are unaffected. All counters are transferred to a set of snapshot registers and can be read via ECS.

Non zero suppression. This command is ignored and the TELL40 will send data as standard “physics”. VeloPix non zero suppression data can only be read over the ECS data path.

Trigger. Since the FE is data driven/self-triggered, Trigger is meaningless in this context. The TELL40 will accept and propagate all data where the Trigger bit is set to 1, and increase the Event ID number by 1.

TELL40 reset or BE reset. With the reception of this signal, the VELO TELL40 firmware clears all the operational logic involved in the data processing, formatting and transmission, such as state machines, data buffers and their write/read pointers, and operational counters and status bits. However, the slow control and configuration registers remain intact.

FE reset. This command is ignored by the TELL40 firmware. Although it is typically accompanied by a “BE Reset” signal which is intended to reset all of the TELL40 blocks including the data processing. For the VELO firmware, buffers will be cleared and the firmware will be set into its pre-synchronization state (i.e., awaiting the first synchronization signal).

4.2.3 Pre-Router

The Pre-Router is the first block of the subdetector specific data process and its functions are: to trigger the data taking mechanism, to change the clock domain from 40MHz to 160MHz (the main data processing clock), to decode the gray BXID (see section 3.2) and to flag the ASIC where the data comes from (subsection 4.2.11).

The triggering of the data taking is performed with a sniffer code on each input link. Once it detects a certain number of Sync frames (typically ten), it declares the link as locked. When this state is asserted, the Pre-Router captures the VeloPix full identification and stores it in a register. It captures the full BXID, aligned for the ten links (or the number of active links if they are less than ten) and sends it to the Post-Router in a 256 bit frame, more information about this frame in subsubsection 4.2.5.1.




Figure 4.9: VELO Data processing scheme.

The clock domain transition is performed by an FSM acting in a similar way as a FIFO memory⁶. This FSM stores and propagates each SPP, according to its associated Data Valid (DV)⁷.

4.2.4 Router

The main processing function of the VELO BE readout boards is performed in the Router and consists in sorting the VeloPix data by its timestamp. This is mandatory as the VeloPix sends unsorted data.

Different architectures were studied, one of the major constrain comes from the fact that the readout board (PCIe40) has two independent channels in the PCI express V3, handling a sustained throughput of 50 Gb/s. Consequently, two parallel data processing streams are designed, addressing ten FE links each. With this architecture, each of the two Routers have to sort 10 links in parallel, running each of them with a maximum throughput of 5.13 Gb/s.

One of the key parts of a complex FPGA design, like the Router, is determining the working frequency. Higher clock speeds allow faster computing with the same resources, but in FPGAs, the maximum frequency has a physical floating limit given by multiple factors. The first one is the technology of the FPGA, in the case of Intel Arria 10 is a 20nm TSMC technology (in particular the model 10AX115S4F45E3SG) and allows working frequencies higher than 500MHz. However, the biggest constraints for the clock speed are given by the design itself, increasing the resources usage, a bad design architecture or a bad clock distribution reduce drastically the working frequency.

An additional constraint for the frequency is given by the input GWT frame decoded on the LLI (subsubsection 4.2.1.3), to make the data processing easier the clock should be a multiple of 4 times the LHC clock, that gives us two possibilities 160MHz or 320MHz. That said, and as the project demands a great amount of FPGA resources, a 160MHz clock is selected as the data processing clock.

The 9 bits sorting algorithm is divided in 2 different conceptual sorters. The first one handles the 4 MSB and is designed to allow the maximum throughput without losing data, this architecture is not used for the whole Router because it consumes a huge amount of memory blocks. The second one will sort the 5 LSB and stores them into memory with the minimum latency.

⁶The FIFO memory is the optimal way of bridging clock domains, but due to the GWT framing this is not possible and, thus, a bit of extra logic is needed.

⁷The format of the data between the LLI and the Pre-Router contains the four SPPs of a GWT frame in the range of bits 0-119 and a DV of each SPP in the range between 120-123.

- **4 MSB block:** This block is formed by combining two basic bit classifiers (Figure 4.10) in a matrix structure. The matrix has 4 columns (1 per sorted bit), the number of rows increases exponentially with the number of columns. Four columns is the minimum number that allows the combined sorting of the ten input links⁸. The first column widens the 10 input links into 16 outputs by using two 2-input 2-output blocks and six 1-input 2-outputs blocks in parallel. The following columns are only made up by 2 input output blocks. Each of the sixteen outputs of 4 MSB block contains the data sorted from all the ten input links combined. This is the required condition to implement the next stage of the Router.




Figure 4.10: (left) 1 input 2 outputs block sorts 1 of the bits in the frame without any memory. (Right) 2 input 2 output block sorts 2 links at a time and consumes 4 small FIFO memories.

- **5 LSB block:** This entity consists of two parts. The first one uses 5 stages of 1 input 2 output blocks to sort the 5 remaining bits of each one of the 16 input from the previous block. The second stage gathers the data in 512 RAM memory segments, each storing data from a certain event, until the Post-Router reads it (subsection 4.2.5). A 512 words event counter takes care of how many packets of data are stored in every segment. This counter is temporarily stored in a RAM memory.

Simulations show that the unsorted data arriving at the data processing will follow the latency curve of Figure 4.11. This allows the Router to open 16 time windows, 1 per connection between 4 MSB and 5 LSB blocks, where a specific BXID arrives. This allows the absorption of the VeloPix latency and delay of the input fibres.

In order to optimize these windows and the data transfer from the Router to the Post-Router, a so-called "Swinging Buffer" memory structure is implemented. This swinging buffer memory architecture duplicates the size of the RAMs at the end of the Router. When the Router is

⁸For a certain number of columns n , 2^n links can be sorted, thus for sorting 10 links 4 columns are needed.




Figure 4.11: Internal latency of the VeloPix ASIC.

writing in one half of the RAMs, the Post-Router is reading the other half and every $12.8 \mu\text{s}$ these halves are swapped⁹. The time window of $12.8 \mu\text{s}$ ($2^9 \times 25 \text{ ns}$) is given by the 9 bit timestamp transmitted by the VeloPix. This is the maximum time the router can wait for packets coming from a certain BXID, before start mixing them with those coming from the BXID + 512.

To design of the size of these RAMs, without the swinging buffer, it is necessary to consider that with a higher number of stored words the data loss is reduced but also could create timing issues in the FPGA and prevent the possibility of further processing in the FPGA. Table 4.2 shows a comparison of the events stored in RAMs with respect to the data loss and the FPGA memory needed. In this case, the simulations show that for 512 events stored on each segment of the RAM the event loss will be kept below 1.5×10^{-7} per datastream of a TELL40 board, which is acceptable within the requirements as events with more hits will be candidates for trigger rejection.

The Router is designed to collect data spread in time, that is why TFC signals are barely implemented in this block. The only TFC signal needed is the “BXID Reset” that will reset the internal counters and restart all the internal FSMs.

⁹The swap of the buffer is synchronous with the BXID reset command.

RAM segment size	RAM overflows	Memory resources
7 bit (128 word)	2×10^{-3}	6%
8 bit (256 word)	2×10^{-5}	9%
9 bit (512 word)	$< 1.5 \times 10^{-7}$	15%

Table 4.2: This table shows the increase in memory resources due to the number of words stored in the full Router. Ram overflows represent the data estimated data loss in the hottest module.

4.2.5 Post-Router

The main function of the Post-Router is to read the data from the Router RAMs, to format them with the Intel Avalon streaming interface (Avalon-ST) [39] and with the LHCb Run 3 data format described in subsection 4.2.5.1. Nonetheless, this is not the only function of the Post-Router. It is also responsible for processing the TFC information associated to the BE board, therefore to either discard or accept good events. Moreover, it is also responsible for the propagation of the TFC synchronization information and the errors that may occur in the Pre-Router, Router and Post-Router blocks as it is only at this point where the full 12 bit BXID is fully reconstructed. The layout of the Post-Router can be seen in Figure 4.12. It works as follows:

- The TFC synchronization flag from the Pre-Router triggers the data acquisition of the VELO data processing. The Pre-Router sends to the Post-Router two different signals: the first one is asserted while a sync command is detected at the input of the Pre-Router and the second one is a 256 bit bus containing the 10 link synchronization frames received from the VeloPix and following the data format shown in Table 4.4.

The first signal is used to trigger the data acquisition (in the “TFC Control” and “RAM control” blocks of Figure 4.12) and as a write condition for storing the second signal in a FIFO memory. The data will be read out of this memory at its corresponding timeslot in sync with the TFC information.

- Once the TFC synchronization flag is asserted and just 200 ns before each Router event RAM swaps the swinging buffer, an FSM reads the content of the TFC buffering RAM that stores the information of the 32 events of the Router that are intended to be read. The same FSM looks at bits 4 and 6 of the TFC information, “VETO” and “header only” commands respectively (See subsection 4.2.2) and create a mask with the discarded events. Besides

that, only events with bit 7 (“Trigger”) set to 1 are propagated and the FIFO memory containing the sync information is only sent downstream if bit 9 (“sync command”) is asserted.

- Sixteen identical blocks (“RAM control”) are responsible for reading the Router output RAMs (each one contains 32 Events). Each one of these blocks is responsible for generating the reading signals for the RAM that stores the number of events. Once the algorithm of this block knows the number of events to read it will check the mask described in the previous bullet point and if the reading is allowed, it will proceed to generate the reading signals of the RAM that store the data.

At this point, the doubt may arise from the reason why the event RAM is read if a particular event is going to be discarded. The answer is simply because reading the event RAM does not waste more than a clock cycle and it simplifies the algorithm.

This block also acts as a master for the entities described in the following items, generating the signals that identify the transitions between events and the control signals used in the following entities.




Figure 4.12: Architecture of the Post-Router.

- As all the communication signals between the Router and Post-Router are performed sequentially by a single bus, the address request is multiplexed on the Post-Router side. On the Router side, the data from the 16 RAMs converge into a multiplexer and are sent to the Post-Router by a single output.
- Errors and warnings, described in subsection 4.2.10.2, originated in Pre-Router, Router and the Post-Router itself are driven to the Post-Router “Error Hub” block that packets them with the error output format described in subsection 4.2.5.1 and sends them downstream at its specific BXID.
- Both TFC buffering and Router data streams converge in the same output entity that synchronizes the output signals and format the data as explained in subsection 4.2.5.1. Moreover, the upcoming 384 bit data bus from the Router is split in two and inflated up to 256 bit output word.

The latency of the Post-Router processing depends on the occupancy of the current and previous events on each of the sixteen blocks that read the Router RAMs. The maximum time to read a single RAM segment with an event of 512 SPPs is 128 clock cycles or 800ns.




Figure 4.13: Data processing block output data interface.

4.2.5.1 Data processing output format

The output format described in this section, is generated in the Post-Router and all the downstream blocks must fulfil this format. The signals described in Figure 4.13 build the interface between the VELO specific data processing and the LHCb firmware framework. It follows the Intel standard Avalon-ST. The Avalon data stream protocol forces the 256 bit PCIe data bus to be sent synchronously with the so-called “TELL40 output clock” which is nothing more than the PCIe clock running at 250MHz. The Avalon-ST standard also requires the transmission of a DV signal whenever a data signal is being transmitted, Start Of Frame (SOF) flagging the beginning of the transmission and End Of Frame (EOF) flagging the end of the transmission.

255	224	192	160	128	96	64	32	0	
VELO Specific Data Out (x8)									Event 1
VELO Specific Data Out (1 < N < 8)					Empty				
VELO Specific Data Out (x8)									Event 2
VELO Specific Data Out (1 < N < 8)					Empty				

31	30	29	25	24	23	22	14	13	8	7	0
ICF		Padded data			Sens. ID		Sensor SP column		Sensor SP row		Hitmap

Table 4.3: (Top) VELO TELL40 output format. (Bottom) VELO Specific Data Output.

In parallel with Avalon-ST data and synchronous with the first data transmitted for each event (when Start Of Frame (SOF) is set to one), different LHCb standard signals are asserted: the full BXID is transmitted to the Event-ID block, event associated TFC data recovered from the TFC buffering block, FSIZE signal that expresses in bytes the length of the event transmitted and, finally, also an FTYPE field that encodes the event type and which will be used downstream in the CPUs to decode the fragment, this FTYPE field could also flag errors in the data.

The output format is shown in Table 4.3. Data are sent as 256 bit words composed of N times 32 bit in a VELO specific sub-frames (also 32 bit), where N is the number of SPPs. Output frames that are not a multiple of 256 bit are padded to the word size. The SPPs sub-frames are modified from their original form coming from the VeloPix. The main differences are that at the output, the BXID information is contained in the global header (along with the length of the

255-250	249-225	224-200	199-175	174-150	149-125	124-100	99-75	74-50	49-25	24-0
Pad data	Sync9	Sync8	Sync7	Sync6	Sync5	Sync4	Sync3	Sync2	Sync1	Sync0

25	22 21	16 15	12 11	0
Padded data	VPx TFC info	Chip ID	BXID	

Table 4.4: (Top) VELO TELL40 TFC synchronization output format (Bottom) VELO Specific "Sync" frame.

whole output frame).

Each Post-Router VELO Specific output sub-frame, contains the SP hitmap (8 bit) and the address of the SP in the sensor (15 bit), the Sensor ID (2 bit), the ICF bit and padded data to meet the 32 bit word boundary.

In addition to the regular physics data described above, the Post-Router block handles two extra special packets that as physics data follows the interface shown in Figure 4.13, but with a different format in the field "FE Data":

- Whenever a TFC synchronization occurs, the information propagated downstream will has an output shown in Table 4.4 and could be identified downstream by checking at bit 9 of TFC data. The VeloPix TFC info propagates the status of the VeloPix at the time of the synchronization with information like wether the data acquisition is active or not, the VeloPix is in a reset state, counters are being reset to its pre-loaded value, the ASIC is saving the value of the internal counters and, finally, the VeloPix is being calibrated with test pulses.
- In case of error or warning in the Pre-Router, Router or Post-Router. The information sent downstream has the format shown in Table 4.5 with the FTYPE equal to the hexadecimal value 3C. The 32 bit error encoding is not completely decided yet. Nonetheless, the 4 MSB are reserved to indicate the origin of the error and the following 28 bits indicate the type of error. The source of the error is labelled as 1 for the generic data processing, 2 for the Pre-Router, 3 for the Router, 4 for Post-Router, 5 ICF and 6 Cluster block.

Different versions of the data format were implemented before reaching the format shown here, those are explained in detail in Appendix A.

255	32 31	28 27	0
Pad data	Error source	Error code	

Table 4.5: FE data for the case of VELO data processing error or warning.

4.2.6 ICF

A search is performed to determine if a SP has a neighbour or it is completely isolated (i.e., not part of a greater cluster). A bit is asserted on the output of each SPP to indicate if it is isolated (1=isolated, 0=unknown). The ICF is shown to have a positive impact on the CPU processing of the data in the high-level trigger farm.

Different ways of implementing this algorithm were studied. In the first approach, the idea was to use one dimension (by column) sorting for a given BXID after the Post-Router, the reason of a single dimension sorting is because sorting by row and column was too expensive in terms of FPGA resources. Once the hits were sorted, they can be compared in a quest for adjacent hits. This architecture has multiple disadvantages like: great latency due to the sorting algorithm, a great number of parallel processes to maintain the throughput, large quantity of FPGA resources, not all events can be flagged (need of a cutoff).

The bubble sorting algorithm approach for the ICF is the typical software approach. Trying a hardware approach could be more efficient from the FPGA point of view. This consists of converting the incoming data from regular binary data to reflected binary gray code, in a single clock cycle, and then measuring the hamming distance¹⁰. One of the properties of the reflected gray code is that the hamming distance between two consecutive numbers is always 1. A measured Hamming distance of 1 is a necessary, but not sufficient condition. In a positive case then a second stage is needed to guarantee that the two pixels in question are neighbours.

4.2.7 Clustering

A real time cluster algorithm for FPGA is being implemented, releasing the computing resources and time from the CPU farm. The design of the clustering is being made keeping a relatively small amount of FPGA resources to make it fit in the PCIe40 board, but there is also a backup solution based on a dedicated FPGA. It has been also considered the transmission of clusters in-

¹⁰Hamming distance is a metric for comparing two binary data strings. While comparing two binary strings of equal length, Hamming distance is the number of bit positions in which the two bits are different.

stead of SP, which will allow a reduction of the bandwidth around $\sim 24\%$. Latest developments were tested with SPs from the most populated sensor, using 1000 minimum bias events and measuring throughputs of 37.9MHz, more than enough for the 30MHz inelastic rate of LHC (see subsection 2.1.1).

The clustering firmware, in the first place, resolves SPs without neighbours with a lookup table filter. The remaining SPs are filled into small matrices of 10 by 12 pixels where every pixel of every matrix is analyzed looking for neighbours patterns in a parallel way, if a pattern is found, a cluster candidate is declared, which will be resolved with a second lookup table.

4.2.8 Event ID

The Event ID block is associating a 64 bit unique tag called event ID to each Event-Fragment received from the Post-Router/ICF/Cluster.

One of the challenges of the TELL40/Event builder communication is that there is a significant number of TELL40 in the experiment very close to the PCIe interface limit of 100 Gb/s network. Hence, in order to save bandwidth, the TELL40 card does not transmit the Event ID and BXID information, as they can be reconstructed later. Event ID and BXID data are substituted by monotonic counters that are increased by steps of one for each accepted event. The not propagation of the Event ID and BXID information forces the FPGA to detect the corruption or de-synchronization of the data, thus the Event ID block compares the BXID data of VELO data processing with the TFC information.

As it can be seen in Figure 4.14, in comparison with Figure 4.13, Event ID block replaces the BXID field by the “Event ID”.

It also performs several checks to ensure the correct continuity of a run:

Event-Fragment accepted trigger checking. Each event-fragment of the same orbit and same BXID in all the TELL40 needs to have the same Event ID tag in order to be associated together correctly in the event builder. The event ID block must ensure that each trigger accepted by the TFC system results in an Event-Fragment at the output of the Data Processing block.

Input interface checking. The event ID block is also checking the correct use of the data interface between the Data Processing block and the Event ID block. If the checking mecha-




Figure 4.14: Event ID block output data interface.

nism detects a wrong behaviour, it will turn the Event ID block status into error and stops any activity.

Truncation mechanism when PCIe interface is not ready. A truncation mechanism has been implemented to avoid stopping the run when the PCIe interface is not ready. Event-Fragments are truncated to 1 PCIe frame and the FTYPE field is flagged as DAQ error. The not ready information is provided by the downstream block (MEP) to help to prevent filling up the downstream FIFOs. Truncation occurrence is monitored to detect any strange behaviour of a PCIe interface.

4.2.9 MiniDAQ output (MEP)

Once the data from all FE optical links have been aggregated onto a single output stream and tagged with a globally unique event identifier, the final stream is transmitted to a readout computer called event builder. The output collected at this point for the different versions of the readout demonstrator (MiniDAQ) is used for preliminary analysis and to validate the dataflow. However, in the final system, the distributed event-building algorithm exchanges and synchronize these data between all readout units to assemble full events out of partial fragments. Due to

the development history of the TELL40 data acquisition board, two different output interfaces are used depending on the target hardware:

- On the AMC40 board (MiniDAQ1), the output component of the firmware implements a bidirectional UDP/IP network stack over 10GBASE-R [40] optical Ethernet, data is received in the readout computer over simplex Berkeley sockets API [41].
- On the PCIe40 board (MiniDAQ 2, 3 and TELL40), the output component is implemented as a streaming DMA engine [42] that coordinates PCI-Express transactions directly with the target CPU, data is received in the user application using a memory mapped interface.

4.2.9.1 Event packaging and multiple fragment packages

Each one of the TELL40 outputs communicates with the event builder CPUs using arrays of fragments. The fragments transmitted by a TELL40 always corresponds to a monotonically increasing sequence of event IDs as explained in the previous section. Hence, it is not necessary to transmit the 64 bit event ID of every fragment. For a given packet of fragments from a TELL40, only the first event ID is transmitted in the multi-fragment header, together with the number of those in the packet, also known as the packing factor. The packing factor must be set centrally by the ECS to the same value on all TELL40.

The CPU event-building process consists of collecting fragments from each of the active TELL40 in a partition for each event. To avoid unnecessary re-formatting of the data, the event-building process is simply a concatenation of the fragments starting with the same event ID. This choice has two important consequences:

- All data sources in a partition are required to transmit fragments starting and ending at the same event ID boundaries.
- Fragments belonging to the same event are not contiguously stored in memory.

The concatenation of the fragments starting with the same event ID from all the data sources in a partition is called MEP. A multiple-event packet comprises a fixed-size part that is 64-bits long, and a variable-size part that contains the offsets to every fragment in the packet (see Figure 4.15).




Figure 4.15: The Multi Event Package (MEP).

4.2.10 TELL40 error handling

The error handling code for VELO is not yet implemented, and the description of this section is likely to be modified in a future. For each block of the TELL40 firmware, a list of potential problems has been catalogued according to two types of problems:

- Error, which is a critical problem and requires to stop the data acquisition. It is monitored in an error register of the firmware and it is accessible through BAR.
- Warning is a minor problem that does not require to stop the data acquisition but should be flagged in a register and monitor in the experiment SCADA.

4.2.10.1 LHCb framework error handling

VELO's firmware is significantly different from the standard LHCb code, but it still shares a few blocks. One of the priorities for the custom blocks is to be as compatible as possible with the

common developments (this will save time during the control software developments). Some of the generic errors imported or adapted to VELO's firmware are:

PLL lock. For the case of losing the lock in the PLL that generates the 160MHz clock used in the data processing block a critical error is raised.

TFC buffer errors. The TFC buffer design, described in subsection 4.2.2, considers the following errors **during physics data taking** that would stop the run: Orbit length different than 3564 clock cycles, wrong BXID, non-monotonic jump on BXID and error trying to obtain the TFC information from S-ODIN either at initialization or during the run.

If the **system is not taking physics data** and any of the previous conditions appears, these potential error will trigger a warning instead of an error.

Latency error between TFC data and FE-fragments. By default, there is a random latency between the FE data and the TFC information from S-ODIN. Two potential problems could appear at this stage: First, if the FE data arrives at the TELL40 in a way that, after being processed for the VELO Router, the data from S-ODIN is still not yet present in the TFC buffer, it will cause a critical error. Second, as the TFC buffer is designed to store one LHC orbit, a critical error can be caused if the data arriving at the Post-Router have a latency greater than one orbit of difference with respect to TFC buffer.

Event ID errors. One of the main tasks of the Event ID block, described in subsection 4.2.8, is to ensure the correct continuity of a run, if any of the multiple items that guarantee is not fulfilled, an error is raised:

- The Event ID block must ensure that each trigger accepted by the TFC system results in an event fragment at the output of the data processing block. If an event fragment is missing, the Event ID block must create a fake one with the associated Event ID and assert a warning. If ten consecutive event fragments are missing, the checking mechanism will turn the Event ID block status into error and will stop any activity. Besides, if the latency between TFC information and the event fragment received is too high, the TFC information buffer might be full turning the Event ID block status into error and stop any activity.

- The Event ID block is also checking the correct use of the data interface between data processing and the Event ID blocks: two consecutive SOF or two consecutive EOF are forbidden and, thus, bringing the system to the state of error.
- The Event ID block error register is accessible from a register with information of the error.

4.2.10.2 VELO error handling

VELO has identified a series of potential errors that compromise the data taking performance:

Link Errors. These errors are sourced in the LLI GWT bank. Link errors can be identified by the GWT header locking mechanism. The GWT_RX bank has a header search and lock routine. A search for the header pattern (“1010”) in the incoming serial stream is performed. Once a header is found, successive headers in the same position (every 128 bits) are counted. If N successive headers are found (typically 23), the “header lock” signal is asserted. Header lock is de-asserted if greater than M incorrect headers are found (typically 4) in a given frame window of headers (typically 60). Warnings can be placed on the number of incorrect headers, either if it is greater than zero or less than the “header lock lost” threshold (M) during running conditions. An error can be raised if the “header lock” is de-asserted during a run.

GWT Parity errors. Apart from header errors, the only error checking that can be performed on incoming GWT frames of the 4 parity bits (one per SPP). These will be monitored with counter registers which will be incremented with each parity error. Warning and Error thresholds will be put in place to indicate the course of action to be taken.

BXID errors. Given that the SPPs arrive in the TELL40 out of time order and they are truncated to 9 bits, detecting all BXID errors is practically unfeasible. BXID “out of order” errors cannot be determined nor corrected. However, it may still be possible to monitor some BXID error signatures such as “forbidden BXIDs” if they are in the range (3564-4095); these values should never be found and indicate a synchronization or timing issue. An error would be raised in this case.

Some BXID latency measurements may be possible. Measuring the latency should produce a histogram with a characteristic shape during standard data taking. Deviation from

this characteristic shape should indicate a problem that could be monitored in the online monitoring farm. The feasibility of implementing such latency measurements is to be determined and it is not guaranteed.

Post-Router and Event ID blocks also check for any de-synchronization between the TFC data and the data propagated downstream of the Router, if that is true, it will be flagged as an error and the run will be stopped.

ChipID errors. ChipIDs will be fixed, one per link during the commissioning of the detector and a change of the ChipID indicates that a fibre has been moved between the OPB VTTx output and the TELL40 input, or that there is a problem with the VeloPix ASIC. An error would be raised in this case.

Synchronization errors. As described in subsection 4.0.0.1, after the initial synchronization between the VeloPix and the TELL40, successive Sync messages can be used as a cross-check of the synchronization of the FE and BE. An error would be raised if this is not the case. Similarly, the relative latency between neighbouring links on the TELL40 is not expected to deviate from each other by more than one or two clock cycles (given that fibre lengths are expected to be approximately the same from one OPB to one TELL40). A large relative latency difference between links should raise an error.

FIFO overflow errors. FIFOs are used extensively in the VELO firmware and form an essential element of the BXID Router. Similarly, most (if not all) data processing blocks have input and/or output FIFO. Where possible FIFOs have been implemented with “full” and “almost full” ports. “Full” FIFOs indicate a loss of data. Low tolerance for full FIFOs is expected (but not zero). It is expected that above a certain rate of full FIFO assertions would raise an error. “Almost Full” FIFO assertions should raise a warning.

4.2.11 ASIC Identification

All FE ASICs send a unique 15 bits identifier when they receive the TFC synchronization command (see subsection 4.2.2). All these FE identifiers are stored in a register based on the BE board input link. These registers are located in the Pre-Router (subsection 4.2.3). In order to reduce the bandwidth, the identifier propagated downstream allows to distinguish a single sensor in a BE board (as the BE board is already identified later on).

Four bits are needed to identify the 12 VeloPixes present in a module and connected to a given BE readout board. However, it is possible to go a bit further with this identification and save computing resources for offline processing but especially for the software trigger, where the computing time is critical. Instead of arranging the ASIC ID consecutively, they will be




Figure 4.16: Four chip ID bits in the data processing according to the FE ASIC of origin. Two MSB (sensor ID) in red and two LSB in black inside each VeloPix.

labeled as shown in Figure 4.16. The three ASICs of the same tile always label, from left to right, the two LSB of chip ID as “00”, ”01” and “10”. The two MSB identify the sensor, but as the front hybrid and back hybrid will go to each datastream (see Figure 4.6), the higher bit is linked to the datastream (0 for front hybrid and 1 for back hybrid) and the next bit identifies the sensor in the front or back hybrids (“0” for the one connected to the left of the GBTx and ”1” to the right one).

This particular way of naming the VeloPix has some advantages:

- The two LSB naming allows the creation of a SPP sensor column (0 to 383) instead of the VeloPix SPP column (0 to 127), without making a new code and therefore saving precious

FPGA resources. This was a requirement from the computing team that will optimize the resources needed for processing the data on a CPU.

- Fixing the MSB to the datastream identifier allows for not sending it through the Router and aggregating it later in the Post-Router. This saves 4% of the Router memory resources.

4.2.12 Monitoring

The monitoring of the PCIe40 board is performed via the PCI configuration space. That allows up to six 32 bit base BAR. PCIe40 uses two of these 32 bit BAR creating two address domains in the FPGA. The first domain maps the monitoring and configuration address for the GBT and GWT LLI, this domain has access to the transceivers, MiniPODs and some delicate components of the system that could compromise the behaviour of the card, thus the access to this BAR is restricted only to experts in the system. A second domain maps the general purpose registers (See section 4.3 for SOL40 control and timing board), in the VELO readout board it controls and monitors the data processing, event ID and MEP blocks.

4.2.12.1 LLI BAR (BAR2)

The design of the BAR2 scheme and address map is strongly bonded, like the rest of the LLI, to the specific FPGA. This causes 3 different LLI designs to each of the versions of the MiniDAQ (see chapter 5), nonetheless MiniDAQ 2 and 3 share almost entirely the same BAR2 scheme.

- The first MiniDAQ demonstrator (section 5.4) built for the LHCb readout was based on an Altera/Intel Stratix V FPGA built-in an Advanced Mezzanine Card (AMC) board [43], this implies that the BAR runs with a clock of 125MHz. This means that all control signals must be re-synchronize from the 40MHz or 160MHz clock of the system to 125 MHz. This BAR allows several functionalities:
 - BAR2 provides a mechanism for tuning and monitoring the PMA and PCS (subsection 4.2.1) of the Intel Stratix V transceiver such as: check the status of the different steps, reset the PMA or PCS, create a loopback in the serialized data, invert the polarity of the links and the endianness of the generic PCS.

- A mechanism to inject GWT data on the transmitter links to test the firmware in a loopback mode. The injection works by feeding data into a RAM and controls the injection through the BAR2 bus.
 - VELO specific PCS described in subsection 4.2.1, has the possibility of monitoring and reset the parity errors in the data transmission, the number of headers received consecutively. It can also enable/disable the scrambler for either transmission or reception and invert the endianness of the GWT word.
 - A series of pattern generators and checkers were included on each link for both transmission and reception. They can be configured and monitored through BAR2 registers for sending and checking a 5.13GHz clock or data patterns like PRBS7, PRBS15, PRBS23, PRBS31.
- Version 2 of MiniDAQ demonstrator (section 5.5) was built as the final readout system with an Intel Arria 10 FPGA [44], causing several changes on the LLI. The PCI connection, in this case, is routed via PCIe V3 connector and the BAR2 clock is now at 40MHz, like most of the firmware. Most of the functionalities described for MiniDAQ1 are still valid for this version, nonetheless there are some differences:
 - The GWT transceiver allows meticulous configuration of the analog parameters of the PMA and the standard PCS for both transmission and reception. As most of the settings need to be optimized for the reception links coming from the VeloPix, this item is focused on the reception PMA and the standard PCS. The analog parameters of the PMA described in subsection 4.2.1, can be modified via BAR2. The polarity of the links can be also changed, but it has to be the same for the whole transceiver bank, unlike MiniDAQ 1 transceiver that allowed changing the polarity per link.

One of the different features of the Arria 10 transceiver in respect of Stratix V, is the embedded pattern generator and checker in combination with embedded accumulators that allow an easier way of testing the links.

 - As explained before, the VELO sub-detector has its own transceiver and LLI which differ from the standard GBT ones even in the reference clock. In order to use the same readout board (PCIe40), which does not has the possibility of generating a clock with the GWT reference frequency (160MHz), a cascade fractional PLL is instantiated to generate such reference clock from the 240MHz clock provided

by the board (see Figure 4.5). Besides that, each transceiver bank has a fractional PLL that generates the reference clock (running at 2.565 GHz) for the transmitter channels of the transceiver.

Both fractional PLLs, mentioned above, need to be calibrated over BAR2, after turning on or loading the firmware into the FPGA in order to lock to its reference clock.

- The VELO specific PCS for MiniDAQ 2 and 3 includes new monitoring features in comparison with MiniDAQ 1, for example: counters for the GWT word aligner lock or the GWT signal valid.
- From the BAR point of view, the final PCIe40 readout board does not differ from the MiniDAQ 2 board, except that the transceiver transmitters are not mounted.

4.2.12.2 General purpose BAR (BAR0)

A second BAR structure, similar to BAR2, has roots in the rest of the BE board code. It allows the configuration and monitoring of the system. In the case of the VELO BE readout boards, each block of the system has its range of address, as shown in Table 4.6.

Data processing block	Range of address (hex)
Data processing top wrapper	from "B00000" to "B0FFFF"
Pre-Router	from "B10000" to "B1FFFF"
Router	from "B20000" to "B2FFFF"
Post-Router	from "B30000" to "B3FFFF"
ICF	from "B40000" to "B4FFFF"
Clustering	from "B50000" to "B5FFFF"

Table 4.6: VELO data processing address map.

Inside each block, an specific code allows the monitoring of the readout chain and the modification of the configuration signals. This work is still under development and, therefore, this section will need revision in the future.

Data processing top wrapper. The top level of the architecture contains all VELO specific code besides the LLI (covered by BAR2). The control and monitoring functions of this block are: to enable/disable each of the ten input links and to select the output data to be propagated downstream from the Post-Router, ICF or Clustering.

One of the next steps to be implemented is the integration of the latency monitor sniffer at the input of the Router (see subsection 4.2.16). The control and histogram registers will be allocated in this wrapper addresses region.

Pre-Router. It allows to enable or disable: the gray encoding block, TFC synchronization mechanism (in case that we want to accept all data from the FE without being triggered by the sync) and finally modify the number of TFC synchronization consecutive frames to trigger the data acquisition. It also allows the monitoring of the unique 15 bits ASIC identifier.

Router. The monitoring of the occupancy of the FIFO memories.

Post-Router. It monitors the output FIFOs, the synchronicity of the VELO data processing and the typology of the errors that may occur on the Pre-Router, Router and Post-Router itself.

4.2.13 Resources

The resources of the actual design of the firmware are 41% for logic utilization and 72% for Intel memory blocks (M20k). Those numbers do not include neither the ICF nor the Clustering blocks, by adding the the Clustering they are expected to increase to 73% and 88% respectively.

Figure 4.17 shows the details of a current compilation of the VELO readout firmware.

4.2.14 Clocking

The GWT transceivers require an input CDR reference clock of 160MHz. In the LLI, each GWT receiver produces a separate 160MHz clock to match the data coming from that receiver.

The GWT data are crossed to the global 40MHz clock domain before being passed to the data processing. In the data processing, a separate 160MHz clock is generated from the global 40. Most of the Router logic on this block operates on individual SPPs rather than the GWT frame. Since there are 4 SPP slots per GWT frame, then the clock must be four times the 40 MHz clock for the GWT frame.

At the output of the data processing, the data (and metadata) cross to the 250MHz domain and the data is formatted to the 256 bit TELL40 output format. ICF and Clustering blocks

Fitter Resource Usage Summary				
<<Filter>>				
	Resource	Usage	%	
1	Logic utilization (ALMs needed / total ALMs on device)	182,272 / 427,200	43 %	
2	> ALMs needed [=A-B+C]	182,272		
3				
4	Difficulty packing design	Low		
5				
6	▼ Total LABs: partially or completely used	27,884 / 42,720	65 %	
1	-- Logic LABs	27,807		
2	-- Memory LABs (up to half of total LABs)	77		
7				
8	▼ Combinational ALUT usage for logic	211,536		
1	-- 7 input functions	632		
2	-- 6 input functions	95,201		
3	-- 5 input functions	22,807		
4	-- 4 input functions	17,282		
5	-- <=3 input functions	75,614		
9	▼ Memory ALUT usage	1,412		
1	-- 64-address deep	0		
2	-- 32-address deep	1,412		
10				
11				
12	▼ Dedicated logic registers	317,778		
1	▼ -- By type:			
1	-- Primary logic registers	272,357 / 854,400	32 %	
2	-- Secondary logic registers	45,421 / 854,400	5 %	
2	▼ -- By function:			
1	-- Design implementation registers	310,576		
2	-- Routing optimization registers	7,202		
13				
14	Virtual pins	36		
15	▼ I/O pins	408 / 960	43 %	
1	-- Clock pins	31 / 50	62 %	
2	-- Dedicated input pins	100 / 155	65 %	
16				
17	M20K blocks	1,869 / 2,713	69 %	
18	Total MLAB memory bits	8,128		
19	Total block memory bits	21,240,512 / 55,562,240	38 %	
20	Total block memory implementation bits	38,277,120 / 55,562,240	69 %	
21				
22	▼ Total DSP Blocks	0 / 1,518	0 %	
1	-- Total Fixed Point DSP Blocks	0		
2	-- Total Floating Point DSP Blocks	0		
23				
24	IOPLLs	6 / 16	38 %	
25	FPLLs	9 / 32	28 %	
26	▼ Global signals	33		
1	-- Global clocks	29 / 32	91 %	
2	-- Regional clocks	3 / 16	19 %	

Figure 4.17: TELL40 resources in detail without ICF and clustering.

process the data at 250MHz. Figure 4.18 shows the different clock domains present in the VELO TELL40 board.




Figure 4.18: Clock domains of the VELO TELL40.

4.2.15 LHCb VELO bypass

The firmware described in the previous sections is very complex and requires a lot of manpower to remain functional and reliable. However, the development and production of the hardware can not be delayed until a fully functional firmware is developed and tested. This forces an implementation of a parallel design for the DAQ firmware, allowing the test of the hardware on lab, production QA and in tests with particle beams (See section 5.8). Additionally, the particle beams tests require synchronization to external references as the TimePix3 telescope [45].

Figure 4.19 shows the structure of the bypass firmware. It only shares with the final design the VELO LLI shown in subsection 4.2.1 and the standard transceiver for the PCIeexpress.




Figure 4.19: Structure of VELO bypass firmware.

The data processing code for this DAQ and testing firmware is a completely new design. It triggers the reading mechanism of a link when ten consecutive TFC synchronization commands are received. After that, it packets the frames with the format of Table 4.7 and stores them in a FIFO memory. The data stored on individual FIFOs per link are read out whenever one of the links is receiving data. In the particular case that more than one link is receiving data, they are read out with higher priority to the FIFO with greater occupancy. The output is driven directly to the PCIe interface.

The TimePix 3 telescope has the possibility of being triggered by an external signal (typically two scintillators, one upstream and one downstream the telescope arms). For this test beam, the

Bit	255-192	191-180	179-160	159	158	157	156	155-124	123-120	119-0
Data	64 bit counter	BXID	Data length	Start of Run	Trigger	Data flow	Sync	Info	Link Number	GWT frame

Table 4.7: VELO bypass firmware output format.

TimePix 3 telescope works as a slave to the DAQ firmware. The S-ODIN, which controls the readout chain, unsets a TFC command called “Stop run” (Table 4.1) during the data taking time window. By creating a falling edge detector, we implemented a trigger signal to the telescope. As a summary for the synchronization, when the run is started from the SCADA used for data taking, the S-ODIN unsets the stop run and sends the synchronization command to the FE. Consequently, the bypass firmware sends a trigger to the TimePix 3 telescope to start recording data. At the same time, the DAQ waits for ten sync commands to start the data taking and propagate all the data subsequently received.

The content of the 256 bit data frame sent downstream of the bypass firmware (Table 4.7) is:

- 64 bit counter that will flag with a common timestamp for the whole data process the time of the data arrival into the FPGA. During test beam operation, the UTC time is preloaded into this counter and it starts to count from thereon.
- BXID represents a full orbit bunch count ID between 0 and 3563.
- Data length is a constant value of 32 bytes for the bypass, it is given by the 256 bit output width.
- Start of Run flags when the operator presses the button to start a new run from the SCADA.
- Trigger, this flag forwards the trigger signal from the readout supervisor board.
- Data flow indicates which of the two parallel data processed sends the data.
- Sync indicates if the bypass receives a synchronization command from the VeloPix ASIC.
- Info field allows the operator to introduce a short comment about the run in the data.
- Link number identifies the fibre from which the data received.
- GWT frame is the raw VeloPix data after the LLI.

4.2.16 Future developments

Several designs with extra functionalities have been either discarded or will be implemented if there is enough room in the FPGA (see subsection 4.2.13), after integrating the main blocks explained in previous chapters. This section gathers some functionalities that were implemented

and not integrated with the final design for lack of time/resources and some extra functionalities or upgrades to the final design:

- The development of procedures that reacts to the large variety of errors that may occur in the TELL40 board. The subsection 4.2.10 describes the tools that VELO is implementing for this purpose.
- A BXID latency monitoring block was developed to prove the different parts of VELO DAQ. This block was interesting especially in the case of the Router.

Latency monitor works by sampling each one of the 10 Router inputs or the 16 links that connect the Router 4 MSB and the 5 LSB blocks. This data is compared with an internal counter synchronized with the TFC “BXID reset” command. A histogram of the latency is stored in a RAM memory with a configurable number of bins (between 16 and 2048 in powers of 2) to reduce the memory resources of the FPGA. Figure 4.20 shows an example of the latency sampling a VeloPix link at the input of the Router and the combined histogram at the interconnect links between Router 4 MSB and 5 LSB blocks.




Figure 4.20: This figure shows the SPPs latency at the input of the Router (in blue) and at the moment in which the data is stored in memory (in red). This latency is measured with the data processing clock, 4 times faster than the LHC.

Latency monitor integration in the final data readout path was put aside because the latency in the Router is not a limiting factor since the “swinging buffer” has been implemented, as it maximizes the latency window in which the Router accepts data. Moreover, the resources are currently prioritised to implement a clustering mechanism in the PCIe40. Despite this, the latency monitor implementation is not discarded between the Pre-Router and Router. It is convenient to measure the latency of each link to early detect a bad behaviour in any FE ASIC and more importantly, to detect a VeloPix noisy pixel that needs to be masked avoiding the overflow of the Router RAMs.

- Post-Router optimization. With the actual design of the Post-Router, see subsection 4.2.5, the maximum throughput is given by the bottleneck of the continuous reading iteration of the Router RAMs. On this design, each segment of the SPP RAM is read out at a rate of 16 SPPs per reading. This is not a bottleneck by itself. However, if we want to maximize the performance of the Router, based on Monte Carlo simulations that predict a mean datastream occupancy of 24 SPPs per BXID, it is possible to improve the actual design by increasing the width of the bus, which interconnects the Router and Post-Router, from 16 SPPs to 32 SPPs. Hence, the majority of the events in the hottest module would be addressed in a single 160 MHz clock cycle and moving the bottleneck downstream to the Post-Router FIFOs that inject the data in the PCIeexpress, nevertheless this is not avoidable. The mitigation of this is only possible if we reduce the death times of this FIFOs, and, thus, the data flow will be more continuous and robust. The necessary firmware changes consist of: resizing the interconnection buses between Router and Post-Router, changing the FSMs that take care of the SPP RAM reading and change the FSMs that handle the building of the output format.
- The necessity of reducing the cost in computing resources forces the possibility of adding new features to the FPGA real time processing. One of the multiple possibilities that is still under study and could be integrated with future versions of the VELO firmware is based on the Router “swinging buffer” architecture. “Swinging buffer” changes every $12.8 \mu s$ on each of the 16 internal links of the Router that connect MSB and LSB blocks. Based on Monte Carlo simulations for the majority of the events, which has occupancies lower than 50%, there is around $7 \mu s$ of dead time that could be used for further processing. The idea is to instantiate 512 small blocks (1 per BXID) that could handle a limited number of

SPPs per BXID, but that will cover a large fraction of the events, as explained the mean datastream occupancy is around 24 SPP/BXID, therefore with 32 SPP/BXID (the next power of 2), will cover most of the events. For larger events, the sorting algorithm will be bypassed and sent out the data unsorted and flagged. The expression 4.1 shows the memory resources needed to implement this sorting. As the FPGA memory resources are the limiting factor for the design, the number of bits to sort and the maximum SPPs to process must be chosen taking a compromise with the memory of the rest of the design.

$$Mem = 512 * SPP_{max} * 24 * 2^{N_{sortbit}} \quad (4.1)$$

Where:

Mem = Memory resources in bits.

SPP_{max} = Maximum number of SPP/BXID that will be sorted.

$N_{sortbit}$ = Number of bits to be sorted.

4.3 LHCb VELO control and timing distribution firmware (SOL40)

The VELO control and timing firmware, which is loaded to the interface boards (SOL40), is a variation of the standard upgrade LHCb framework. It is modified to include the communication with the VeloPix FE ASICs both to write and read the configuration commands and to distribute the TFC information for keeping the sub-detector synchronous with the rest of the experiment. As it can be seen in Figure 4.21, a single VELO slice needs three control links: one for handling the OPB and monitoring the FE temperatures and analog signals and two additional links for controlling 6 VeloPix ASICs each. As a PCIe40 board configured as SOL40 has 48 bi-directional links, four boards are needed for the whole VELO.

GBT bit	Downlink (SOL40 to OPB)	Uplink (OPB to SOL40)
32-33	GBT-SCA1 data	GBT-SCA1 data
80-81	GBT-SCA0 data	GBT-SCA0 data

Table 4.8: This table shows the selection bit distribution of the OPB GBT word.




Figure 4.21: Control and Timing firmware architecture for a VELO slice.

The link connected to the OPB uses the CERN standard GBT-SCA protocol over the GBT standard link. GBT-SCA is connected to the optical link GBTx ASIC using a full duplex transmission protocol, based on the HDLC standard (ISO/IEC 13239:2002) [46]. On the other hand, each VeloPix has two Scalable Low-Voltage Signaling (SLVS-400) [47] links connected to a GBTx: a full duplex Serial Peripheral Interface (SPI) 80Mb/s link for the slow control and a simplex 320Mb/s link for timing. The VeloPix ASIC was designed allowing the possibility of connecting all these links for different ASICs in the same bus, but, in order to avoid any further problem, a more robust point to point connection to individual e-ports of a GBTx ASIC was chosen.

The CERN radiation hard link transmission protocol, GBT [18], is based on a 4.8Gb/s transmission link. It has several configuration modes depending on its use. In the VELO case, which needs a very reliable control link and where the data transmission is not a priority, it uses the forward error correction mode that reduces the usable data bandwidth to 3.2Gb/s, and, thus, an 80 bit word is available for data transmission. Table 4.8 shows the selection bits used for the OPB GBT-SCA communication. Similarly, Table 4.9 shows the selection bits used for the

transmission of ECS and TFC data to the VeloPix ASICs.

4.3.1 VeloPix ECS link

The SLVS-400 SPI protocol embedded in the 80Mb/s full-duplex link that is used for GBTx and VeloPix ASICs communication is shown in Figure 4.22.

The VeloPix always has a slave role in the bus. By default, it is always looking for the “start of frame” or Header (hexadecimal value E8). Once the header is found, it looks at the Chip address and if it matches with its pre-programmed address or it is all 0¹¹, it checks whether the requested action is to write or to read and reacts accordingly. For reading actions, the VeloPix replies with the value of the address and 16 bit of status word, and for writing, only the status word is replied. The status word contains information about the performance of the ASIC like:

¹¹Chip Address 0 is used for the broadcast mode, which means that all VeloPix in the bus must reply.

GBT bit	Downlink (SOL40 to FE)	Uplink (FE to SOL40)
0-1	ECS data for VeloPix 5	ECS data from VeloPix 5
2-3	Reset VeloPix 5	Not used
4-5	ECS data for VeloPix 4	ECS data from VeloPix 4
8-9	Reset VeloPix 4	Not used
0-1	ECS data for VeloPix 3	ECS data from VeloPix 3
10-11	Reset VeloPix 3	Not used
16-17	ECS data for VeloPix 2	ECS data from VeloPix 2
18-19	Reset VeloPix 2	Not used
20-21	ECS data for VeloPix 1	ECS data from VeloPix 1
22-23	Reset VeloPix 1	Not used
24-25	ECS data for VeloPix 0	ECS data from VeloPix 0
26-27	Reset VeloPix 0	Not used
32-39	TFC data for VeloPix 5	Not used
40-47	TFC data for VeloPix 4	Not used
48-55	TFC data for VeloPix 3	Not used
56-63	TFC data for VeloPix 2	Not used
64-71	TFC data for VeloPix 1	Not used
72-79	TFC data for VeloPix 0	Not used

Table 4.9: This table shows the selection bit distribution of the FE GBT word for the 6 VeloPixes controlled from a single GBTx. This word is distributed at a 40MHz rate, thus 2 bit slots corresponds to 80Mb/s FE links and 8 bit slots to 320Mb/s FE links.




Figure 4.22: VeloPix ECS protocol layer.

GWT DLL lock, single event upsets detected, DAQ enabled, test pulse enabled, address set ok, chip address ok, etc. The only status information checked in the firmware is the matching of the address and the chip address. If those are set to 1, it means that the address and the chip are present and the writing/reading operation was successful.

From the BE FPGA point of view, the code is very simple, as all the processing is performed on a CPU. The CPU interface is done via the PCIe v3 BAR0 of 32 bits. This means that for each VeloPix there are two 32b registers, one for reading and one for writing.

The data going from the CPU to the VeloPix are received in the FPGA in 32 bit words and stored in a FIFO memory. Once the full word is received in the FPGA, a trigger command serializes the full content of the FIFO and sends it to the VeloPix, as it is shown in Figure 4.23.




Figure 4.23: Timing diagram of the signals sent from VELO SOL40 firmware and the GBTx deserialized word sent to the VeloPix ASIC. 0 and 1 are the 2 bit word sent out by the downstream FIFO.

The up-link from the VeloPix is always being checked by FPGA's code looking for the header (E8) that flags the start of transmission. Once this is found and the chip address is checked, it triggers the reading of the VeloPix data and stores it on a deserializing FIFO. This FIFO will wait for a word of 1536 bit, which is the maximum number of bits that the VeloPix can reply, and close the input shutter from the VeloPix. After that, the CPU can read the data from the FPGA through the 32 bit BAR0.

4.3.1.1 VELO interface board BAR0

VELO is the only sub-detector that has its own dedicated communication protocol with its FE ASICs. This requires the integration of this protocol into the LHCb framework in order to control all sub-detector FEs with the same software tools.

To operate and monitor a single side of a VELO module (1 GBT link), 36 registers are needed (6 per ASIC) with the following functionalities:

- A 32 bit register where the data are transmitted. Whenever a new word is written in this register, VELO's code forwards it to a FIFO, where it is temporally stored until a trigger is received.
- A second register called "go-VELO" acts as a trigger. When a rising edge is detected the data stored in the FIFO is serialized and sent to the GBTx in the module.
- Changing the transmission FIFO content is possible thanks to a "flush" register. Whenever it has a 1 written in the LSB, it erases the data already stored.
- Each FE ASIC reset pin is directly connected to the LSB of a register. By setting it to "1" the ASIC can be sent to a reset state.
- VeloPix TFC data acquisition (subsection 4.3.2) can be enabled by opening the shutter and closing the shutter. This opening will be done in normal operation before the start of collision in the LHC because the transitions of the shutter cause fluctuations in the power distribution of the ASIC and thus inject noise packet transmission through the GWT links.
- Similarly to the configuration data sent through the down-link to the VeloPix, the ECS response, arriving from the VeloPix ASICs in the up-link, is deserialized and temporarily

stored into a FIFO. At the time that 1536 bits are received, the received word is read through a specific BAR0 register.

- A 24 bit register is used to monitor the status of the 12 VELO FIFOs. Each one of these memories is monitored via its full and empty signals.
- Each down-link and up-link has its packet counter for the sent or received serialized packets.
- Similar to the previous item, the up-link has implemented a counter of the packets detected as erroneous. The erroneous packets are declared if the up-link receives data at the same time as the down-link is transmitting (abnormal behaviour).
- A BAR0 register is dedicated for resetting all counters.
- Two registers allow to predict and monitor the header response of the VeloPix ECS frame, and if it does not match with what is expected, it does not propagate the data and declare them as erroneous.

4.3.1.2 VELO FE transmission and reception operation

This section describes the sequence of instructions that the firmware expects from the host computer to write and read a register into a given FE ASIC:

Writing a register

- Flush the link FIFOs (down and up-link) of the VeloPix that we want to write.
- Generate the full word that we want to transmit to the VeloPix following the sequence described in Figure 4.22.
- Split the word in smaller 32 bit words and transmit them to the firmware through the appropriate BAR0 address with big-endian ordering.
- Once the word is fully transmitted to the firmware, set the associated "go-VELO" trigger to 1 and then 0 to start the transmission to the ASIC.

- Wait until the data is fully transmitted, the VeloPix replies with the acknowledgement signal. Optionally, once the receiving side has an extra packet detected, or the empty signal of the associated FIFO changes to 0, it is possible to read it and check the ASIC status of the written register.

Reading a register

- Flush the link FIFOs of the VeloPix that we want to write.
- Generate the full word that we want to transmit to the VeloPix as shown in Figure 4.22.
- Split the word in smaller 32 bit words and transmit them to the firmware through the right BAR0 address with big-endian ordering.
- Once the word is fully transmitted to the firmware, set the associated "go-VELO" trigger to 1 and then to 0 to start the transmission to the ASIC.
- Wait until the up-link counter increases and then start reading the VeloPix register in 32 bit frames (again in big-endian ordering).

4.3.2 VeloPix TFC link

The VeloPix TFC is used only for critical signals that must be perfectly aligned in time. These signals are generated every bunch cross (25 ns). The signals are sent over a serial connection based on a 320 MHz clock. The clock allows 8 bits to be transferred but only 7 bits are transferred in the case of VELO. One of the bits is used for flagging the start of the 8 data bits transmission in synchronization with the LHC clock. Figure 4.24 illustrates the data transmission on the firmware side and how they are received by the VeloPix after being deserialized by the GBTx ASIC. The other bits transmitted to the VeloPix are: FE reset, BXID-reset, synchronization command (sync), calibration, snapshot and shutter. The first five signals are defined by the LHCb TFC system, while the shutter signal is specifically built for the VeloPix. The LHCb TFC shutter can not be used directly here because it can have multiple transitions during data acquisition and it is known that transitions of the VeloPix shutter create perturbations on its readout links for several clock cycles. Hence, minimizing shutter operations is a must. This




Figure 4.24: TFC signals on firmware and between module GBTx and VeloPix ASIC. Where: 0 (BXID reset), 1 (FE reset), 2 (Test pulse injection (calibration)), 3 (Snapshot), 4 (TFC sync), 5 (Shutter), 6 (Not used), 7 (Alignment, always 1).

shutter fed to the VeloPix is created with an AND gate between a signal controlled by ECS (acting as enable) and a signal provided by a “SR Flip-flop” set with the rising edge of the LHCb shutter and reset with the “FE Reset” signal.

From the BE board firmware point of view, the changes are minimal with respect to the rest of the LHCb sub-detectors as the protocol for TFC is common and it was only needed to reorder the positions of each TFC bits and create the shutter specific command.

4.4 LHCb control software

Although this is not the main subject of this thesis, it is advisable to give a few guidelines about the software that interconnects with the firmware registers previously explained. With the aim of easily and reliably control and configure all LHCb sub-detectors, a SCADA system (Based on Siemens WinCC OA) [48] is being developed, integrating all sub-detectors. Due to the custom nature of the different sub-detectors, it requires the implementation of a framework capable of handling the communication with the FE devices via all the available field-buses, providing a link to the ECS and abstracting the description of the hardware and model it into the data structures used by WinCC OA.

Figure 4.25 is a depiction of the typical architecture of the control system for the FE electronics devices of the experiment through the interface boards. On the software side, the framework for the control system is composed of 3 main components:

- A GBT Server – runs on the host PC that holds the FPGA board.




Figure 4.25: Architecture of the LHCb control system for the FE electronic boards.

- A GBT Client – runs on the controls PC.
- A hardware abstraction tool – installed on the controls PC.

All LHC experiments use as their SCADA system the software WinCC OA. To reduce development duplication, the Joint Controls Project (JCOP) [49] was created to provide common controls solutions for the big four experiments. JCOP generates a framework for the creation of components, WinCC OA packages containing all the required user panels, libraries, scripts and other software that can be easily installed and distributed. The two later components of the control system software for the FE electronic devices (GBT server and client) are developed as framework components.

The GBT server implements the lower level communication with the FPGA board devices and firmware, the master GBT, and the GBT-SCA chips. It also implements the connection to the control system, acting as the bridge between the Control System and the hardware devices and firmware. For the local FPGA board, the GbtServ can control and monitor devices using the I2C and SPI protocols, it can communicate with the FPGA firmware using the local bus (typically PCIe). For the remote FE devices, which in general are connected via the GBT-SCA chipset, the server is based on Distributed Information Management System (DIM) [50] that can easily be interfaced from WinCC OA. DIM is a communication system for distributed/mixed environments and it provides a network transparent inter-process communication layer. It is based on the server/client paradigm and it uses the concept of publishing/subscribing services and commands. The GbtServ provides the connection to the control system as it implements services and commands via DIM for all the available field-buses, both to the local FPGA board

and remote GBT-SCA chips, which are then available to be subscribed to the GBT client in the Control System. The Control System abstraction tool can also send a description of the connected hardware, with all the relevant settings and configurations to the GbtServ, and the GbtServ exposes DIM services and commands based on these descriptions.

CHAPTER 5

FIRMWARE DEVELOPMENT AND HARDWARE VALIDATION.

This chapter is intended to explain, in first place, the firmware design workflow, followed by a chronological design overview of the different versions of the hardware demonstrator for the LHCb VELO upgrade.

5.1 What is a FPGA?

A Field Programmable Gate Array (FPGA) is a semiconductor device based on a matrix of programmable logic cells that can be freely interconnected and configured, providing high flexibility in terms of hardware. The architecture of the logic cells depends on the manufacturer and even on the device family, but the philosophy could be based on multiplexers (for example Microsemi FPGAs) or memory based (Xilinx and Intel FPGAs). As new Intel and Xilinx devices use memory based logic units, I will only explain this type. Figure 5.1 shows a basic example of memory based logic cell where the Lookup Table (LUT) contains the programmed logic, a D-type flip-flop latches the output of the lookup table and a multiplexer selects between the flip-flop or LUT outputs. This cells could also be improved by including adders for handling the carry in/carry out bit with adjacent cells, this allows to perform fast arithmetic operations. Moreover, extra flip-flops could be added to improve timing performance.

The main difference from ASICs lies in their capability of being re-programmed by changing the interconnection of the logic cells, on the other hand, ASICs are custom manufactured.




Figure 5.1: FPGA logic cell.

The programmable logic cells can be configured on a wide range of logic functions, since logic gates (AND, OR, XOR, etc.) to complex combinational functions (shift registers, multiplexers, counters, etc.). Besides that, FPGAs also have built in memory blocks, a wide range of communication transceivers that allow multiple standards or even customize them, etc.

Digital electronic designs, like FPGA design, are based on Hardware Description Language (HDL) like Verilog [51] or Very High Speed Integrated Circuit Hardware Description Language (VHDL) [52]. These designs differ strongly from CPU programming. This design flow is cover in the next section. However, nowadays the use of OpenCL [53] compilers is spreading for complex data processing applications using new (and large in resources) FPGAs. The language OpenCL is well extended in GPU programming, hence it allows parallel programming, and it is more suited for developers used to work with high level programming tools.

FPGAs have several advantages like: they allow real time critical applications as they have no software built in, massively parallelized as they are like a white paper hardware-wise, they allow building ad-hoc systems improving the efficiency, short time development than an ASIC, etc. Unfortunately, they also have some disadvantages like: require a high knowledge of HDL programming and hardware related issues (timing, power consumption, etc.), large time development than a CPU based system. Table 5.1 helps to fit FPGA technology by showing a comparison with some of the most common silicon technologies.

Criteria	CPU	GPU	FPGA	ASIC
Processing power	Moderate	High	Very High	Highest
Flexibility	Highest	Medium	Medium	Very low
Parallelization	Low	High	Very high	Highest
Latency	High	Very high	Low	Lowest
Input/outputs	Very Low	None	Very High	Highest
Implementation and validation time	Low	Medium	High	Highest
Power consumption	Highest	Medium	Very Low	Lowest
Cost	Low	Medium	High	Highest

Table 5.1: Comparison of the most popular silicon technologies.

5.2 Firmware design flow

Figure 5.2 shows the generic FPGA design flow, valid for all designs explained in this thesis, which uses two vendors Intel/Altera and Xilinx. To give some context, these two vendors are the most important players in the FPGA business, having a combined market share around 90%.

Before starting any design, it is necessary to create a conceptual model of the architecture that can go from a simple paper draw to a very detailed schematic. It is better not to skip this step and take the time needed because it will be the baseline for a good design.

Once the structure is conceptually defined, it is required to identify all the IP needed for the design (transceiver, PLL, memories, etc.) and configure them in vendors IP catalog under Quartus environment for Intel/Altera or ISE/Vivado for Xilinx FPGAs. These IP components are usually paid and, thus, it has to be considered when designing the architecture of the code.

Configured IP cores generate a bunch of HDL and Tool Command Language (Tcl) files, defining the entity for the hardware implementation and simulation. Generated IP HDL code is usually implemented in Verilog. An HDL wrapper to the entire IP is also included, which can be configured to the user language (in our case VHDL).

With the HDL design complete, or a significant part (in case of big projects) a behavioural simulation is performed. For the work presented here, simulations were performed using ModelSim/Questasim tools¹. On those, the most common simulation consists on creating a wrapper for the code under test, feeding the inputs with expected input signals, either generated on an HDL code or by reading text files, and checking the output or intermediate signals. In case of

¹ModelSim/Questasim was chosen because it is a more powerful tool for complex projects compare to the simulation tools provided by FPGA vendors. It works with projects developed on different languages, scripting and it is faster than most of the equivalent platforms.




Figure 5.2: FPGA design flow.

bugs in the HDL code, it can be re-designed. The full VELO firmware readout chain is tested by injecting monte-carlo data from files and matching the output with this data. Small blocks are tested with custom data generators.

Once the designer is satisfied with the behaviour of the HDL code, to continue with the implementation in real hardware, it is necessary to apply some constraints to the synthesization of the code by setting the input/output pins in the hardware, the signals and clocks distributed along with the FPGA and defining the location and routing of the code (Something similar, although with some important differences to the PCB routing). For Intel/Altera FPGAs, these are set in Synopsis Design Constraints (SDC) files and on Xilinx devices Xilinx Design Constraints (XDC) files. Constraints files are written in a proprietary language.

Now, going back to the vendor's design environment (Quartus/ISE/Vivado), the synthesis and optimization can be performed. If the size of the design matches within the FPGA, the routing and placement of the code can be launched as well as the timing analyzer. In case of not succeed in the synthesis, a change in the code is required. For timing errors, they should be analyzed and, depending on the typology of the error, it could be solved by applying new constraints or changing the HDL code. The important concepts about timing constraints and possible failures are related to setup and hold times of registers and how they, combined with the propagation and combinational circuit delays, determine how fast a system could run. The setup time is the amount of time, before a clock edge, required for the data to be stable at the flip-flop input. In contrast, hold time is the minimum amount of time required for the input of the flip-flop to be stable after a clock edge.

The last verification step, previous to the approval of the firmware version, is to test the design in the hardware via vendors tools that act through Joint Test Action Group (JTAG) programming and debugging bus. An alternative way of testing the code in hardware, when the design is more mature, is by implementing and monitoring user BAR registers.

If the design is approved, then the generation of the Electrically Erasable Programmable Read-Only Memory (EEPROM) memory bit file concludes the design process.

The following items describe the simulation software used, and the environment for each FPGA vendor family and its specificities.

5.2.1 ModelSim/QuestaSim simulation

The Questa Advanced Simulator (QuestaSim) is a multi-language HDL simulation environment by Mentor Graphics for simulation of hardware description languages such as VHDL, Verilog and SystemC, and includes a built-in C debugger. Questasim can be used independently, or in conjunction with Intel Quartus Prime, Xilinx ISE/Vivado.

The LHCb collaboration designed an entire framework that emulates the behaviour of the whole firmware, with space for each sub-detector to test their developments. Furthermore, a test-bench using GitLab repository is being developed, with the possibility to inject data into the data processing block from a FE emulator or a text file. Besides that, the VELO sub-detector, due to the large modifications of the common project, developed local test-benches for each functional block.

5.2.2 Intel/Altera Quartus design environment

Quartus is a design software for Intel and former Altera FPGAs that comprises design and Intel IP cores configuration, analysis, synthesis, timing analysis and simulation of HDL designs. For the designs involved in this work, several versions of the environment were used, from the first versions in Quartus 15.1 standard edition to the latest using Quartus 19.1 Professional edition. The reasons behind these changes in version are, in one way or another, related to the passage of time: Intel Arria 10 FPGA is only supported in Quartus 16.1 or newer, improvements in design tools and compiler engine.

Quartus environment encompasses different designs and debugging tools widely used for the LHCb design, the more used ones are: Qsys, Chip Planner, Timequest, Signaltap, Transceiver toolkit.

5.2.2.1 Qsys

Qsys is a platform designer tool created to ease the interconnection of IPs with user functions and subsystems. Qsys platform is especially handy in large projects thanks to its hierarchical framework. It is also very flexible with respect to the design methods, allowing HDL languages, block-based entries, schematic and black boxes (empty wrapper with code that will be filled later on in the design flow).

5.2.2.2 Chip Planner

The Chip Planner provides a visual display of FPGA resources. The Chip Planner can show logic placement, logic lock regions, relative resource usage, detailed routing information, fan-in and fan-out connections between nodes, timing paths between registers, delay estimates for paths, and routing congestion information. It can also make changes in the assignment, such as creating and deleting resource assignments, and also perform post-compilation changes creating, moving, and deleting logic cells and Input/Output (I/O) atoms.

Finally, Chip Planner allows the viewing and creation of assignments for a design floor plan, perform power and design analyses, and implement last-minute changes. With the Chip Planner and Resource Property Editor, connections between resources could be changed and make post-compilation changes to the properties of logic cells, I/O elements, PLLs, and RAM and digital signal processing blocks.

5.2.2.3 Timequest

The TimeQuest timing analyzer is a powerful ASIC-style timing analysis tool that validates the timing performance of all logic in a design using industry standard constraint, analysis, and reporting methodology. The TimeQuest timing analyzer includes support for SDC. Additionally, the TimeQuest timing analyzer includes an extensive Tcl scripting. This tool was widely used, in the case of the work exposed here, for applying timing constraints to the VELO specific LLI and data processing blocks.

5.2.2.4 Hardware validation tools

The Quartus design environment has several hardware validation tools like Logic Analyzer, SignalProbe, System Console, Bus Analyzer, etc. But either because they were not available for the FPGA families, we used and/or they were not suited for the test we defined, they were not used. Hence, this chapter will only describe those tools actually used in the project.

- The Quartus SignalTap Embedded Logic Analyzer is a tool that allows the designer to capture signals from internal FPGA nodes while the device is running, giving non-intrusive access to signals from internal device nodes. The Embedded Logic Analyzer megafunction works within the Quartus environment for design development, debugging, and verification. Quartus SignalTap allows the configuration of the signals that the designer/tester

would like to probe, previous to the synthesis. Once the Embedded Logic Analyzer mega-function is configured and compiled, it can be downloaded into the device with the code under test and debugged it via the JTAG blaster communications cable. Any change in the trigger to the Embedded Logic Analyzer can be made without recompiling the device design. The Quartus software displays data acquired by the Embedded Logic Analyzer as waveforms.

- The Transceiver Toolkit in the Quartus environment allows a quick test for the functionality of the transceiver channels and helps to improve the signal integrity of the transceiver links. In today's high-speed interfaces, stringent Bit Error Rate (BER) requirements are not easy to meet and debug. The Transceiver Toolkit in the Quartus software allows checking and improving the signal integrity of transceiver links on the device before the completion of the final design, saving time and helping in finding the best PMA (subsection 4.2.1) settings for the high-speed interfaces.

The underlying framework for the Transceiver Toolkit is the System Console. The System Console performs low-level hardware debugging of the design. The System Console provides read and write access to the IP cores instantiated in the design. The System Console is used for the initial bring-up of the PCB and low-level testing. The Transceiver Toolkit works at run-time and real-time configuration, including performing high-speed link tests for the transceivers in the devices.

5.2.3 Xilinx "ISE" and "Vivado" design environments

Xilinx ISE or Vivado are the design suites for Xilinx FPGAs. In 2013, Xilinx abandoned the ISE denomination for the environment in detriment of Vivado, thus in this chapter only Vivado is going to be described. This environment was used for all VC707 board [54] developments, and in particular the QA test-benches explained in chapter 6 using version 2014.4 of Vivado design suite.

Similar to Quartus environment, Vivado includes several analogous tools under its own naming. For simplification, this chapter describes only those used in the design process of the production test-benches, which are small projects and do not require the usage of most Vivado design tools.

5.2.3.1 IP catalog and integrator

The Xilinx Vivado design suite IP integrator tool, in combination with the IP catalog, is the equivalent to Intel Quartus Qsys. They allow the creation of complex design systems by configuring, instantiating and interconnecting IP cores with user's code. They allow working in both graphical and Tcl interfaces.

5.2.3.2 Timing tools

After Synthesis is being executed, several reports are available like: timing, clock networks, clock interaction, pulse width. Report timing summary combines two criteria: the design is fully constrained and it meets timing. Vivado allows applying closure techniques at the implementation level, such as: refining the timing constraints, floorplanning, re-routing, modify the implemented logic and optimise the implementation tools to strive to meet timing. All these changes will be saved in XDC files to help the implementation engine in future compilations.

5.2.3.3 Hardware manager

Vivado hardware manager is a hardware validation tool that allows the debugging of the code after the bitstream is being loaded in the device. The hardware manager supports graphical and Tcl interfaces, the tools used for the testing explained in section 6.1 and section 6.3 were Logic Analyzer and Serial I/O Analyzer.

The logic analyzer, allows the monitoring and driving of internal and pins of the FPGA.

Serial I/O allows the debugging and optimization of the FPGA transceivers. It provides built-in BER measurements on multiple channels, performs 1D/2D eye scans, and adjusts transceiver parameters in real time while the I/O channels are working.

5.3 SPIDR readout system

The SPIDR readout system (Figure 5.3) was designed to exploit the full capability of TimePix3 family of ASICs (TimePix 3, MediPix 3 and VeloPix). The SPIDR design is based on the Xilinx development board VC707 [54] connected to a custom made carrier board, specifically designed for the ASIC under test validation, through the FMC connector of the VC707.




Figure 5.3: SPIDR readout system.

The SPIDR system was widely used for the verification of the VeloPix ASIC design (section 3.2) and the quality assurance of the ASIC in the silicon wafers. The SPIDR system is controlled over a 1 Gb/s ethernet interface that access a set of registers in the LEON3 CPU built in the FPGA [55], these registers control both the FPGA and the VeloPix ASIC. The slow control readout (ECS running at 80 Mb/s) is also handled through the 1 Gb/s link. For the high speed readout, the 4 GWT links of the ASIC running at 5.13 Gb/s, are connected to a high speed FPGA GTX (Gigabit Transceiver) capable of running up to 12.5 Gb/s. The FPGA recovers the link, removes the empty GWT frames and sends the data over 4 10Gb/s Ethernet connections making use of the IEEE 802.3 Standard Jumbo Packet [56].

5.4 MiniDAQ 1 readout system

The MiniDAQ 1 readout system, shown in Figure 5.4, was the first prototype for the LHCb BE board used until the end of 2017. The MiniDAQ 1 system consists of a PCB hosting an Altera Stratix V GX FPGA and 36 bi-directional optical links connected over an Advanced Mezzanine Card (AMC) to a credit card PC² that controls the FPGA and FE systems. The control, ECS data taking and filesystem of the credit card PC are done via 1 GbEth link. The communication with the FE is performed through 24 bi-directional optical links that could be configured either as control or DAQ. Twelve extra links are used as 10 GbEth to read out the acquired data from

²A credit card PC is the name given to a computer with a very small size, typically with similar dimensions to those of a credit card.




Figure 5.4: MiniDAQ 1 connected to the first prototypes of the VELO electronics.

the FPGA to a computer.

In the case of VELO sub-detector, this setup has never been fully exploited as it was used only for development and testing the firmware and the hardware prototypes.

In the first place, and because the LHCb software was still under development, a VELO specific server running in the credit card PC of the MiniDAQ 1 was used to test the SOL40 control firmware (a detailed description can be found in subsection 5.4.1). This VELO server solution emulated the SPIDR functionality and could run with up to 3 VeloPix in parallel, it

was halfway solution between the SPIDR approach and the final experiment firmware, allowing the control of the full chain and the use of the computer VeloPix Apps used with the SPIDR system. The MiniDAQ 1 was used mostly with the first version of the on-detector electronics (see Figure 5.4) which could only handle up to 3 VeloPix ASICs. The first prototype of VELO on-detector electronics setup was used for validating the hardware design and optimizing the auxiliary ASICs (GBTx, GBT-SCA, Giga Bit Laser Driver (GBLD), VTRx, VTTx).

A prototype version of the VELO LLI firmware, described in subsection 4.2.1, was developed for this setup. The MiniDAQ 1 LLI had its own Transceiver and therefore a FPGA specific PMA receiver and transmitter. Nevertheless, the PCS block was tested and developed with this setup and it is common for the GWT protocol, hence easily migrated to the final PCIe40 board. The VELO PCS also included the capability to perform BER, giving the possibility of testing the GWT FE link. By using this functionality, automatic phase scans of the VeloPix were performed giving a BER of 1.3×10^{-15} with 325 m of optical transmission (more information about the readout link optimization later on in this chapter).

The VELO data processing firmware was not tested with this setup as all the development effort was turned into MiniDAQ 2 setup, which will be described in the next section.

5.4.1 MiniDAQ server

This server, embedded in the credit card PC, establishes a TCP communication between the software application and the firmware associated to a single VeloPix. Hence, three of these kind of servers are needed for controlling the system shown in Figure 5.4 and six of them for handling all the potential VeloPix controlled through a single GBT link.

Each server runs on a specific port with a base number of 50000 plus the VeloPix under control number³ (within the dynamic or private ports defined by IANA [57]). This server architecture is neither escalable to the full sub-detector nor easily integrated with the rest of the experiment. Hence, this was a temporary solution before migrating to the final LHCb standard server.

The data coming to the server have 3 fields: the type of command, the length of the data and the data itself. The command field allows controlling the firmware in several ways, all commands handled by the server are listed below.

³The naming of the VeloPix inside a module is always from lower number to higher number in the associated GBTx e-port, in other words, if we look at the module having as a reference the GBTx ASIC from left to right.

- SOL40 access (subsubsection 4.3.1.1).
- TELL40 access (subsubsection 4.2.12.2).
- LLI registers access (subsubsection 4.2.12.1).
- Writing and reading to a VeloPix register (subsubsection 4.3.1.1).

5.4.2 MiniDAQ1 DAQ

Even though the MiniDAQ 1 demonstrator has never had a fully working version of the readout board firmware, it was used for development purposes until it was abandoned in detriment of the MiniDAQ 2 and the final readout structure. The limited resources of the Intel Stratix V FPGA forced the MiniDAQ 1 readout firmware to have several specificities by comparison with the final readout structure shown in Figure 4.6. The MiniDAQ 1 firmware only handles one of the two parallel datastreams of the final system. Besides this, the DAQ output is sent through a 10 GbEth Small Form-factor Pluggable (SFP) connector instead of the PCIe v3 of the final setup.

5.5 MiniDAQ 2 readout system

The MiniDAQ 2 readout system (See Figure 5.5) was born as a development tool for the different teams involved in the LHCb upgrade to test their FE setups and to develop the software and firmware needed for the upgrade. The MiniDAQ 2 consists in the first version, of the PCIe40 card (described in section 4.1) in an ASUS ESC 4000 G3 server PC.

The MiniDAQ 2 was available at the beginning of 2018. Therefore, from that moment it took over from the MiniDAQ 1.

The VELO control firmware didn't change with this new hardware. However, while the server was not migrated to the LHCb ECS, the communication with the software was done through a similar server as the one used with the MiniDAQ 1. As well as the server all apps of the SPIDR systems were migrated to the experiment software.

Concerning the VELO readout board firmware, a new LLI transceiver is generated, see section 4.2.1. For the validation of the LLI, as taking real data will flood us with useless data, a scrambled constant value is chosen to be sent from the VeloPix. That allows the debugging of the VELO ad-hoc PCS (already described in subsubsection 4.2.1.3), in the first place, using




Figure 5.5: MiniDAQ 2 setup.

the Intel hardware debugging tool, Signaltap, and later by registers. Once the LLI was working and stable, the next aspiration is to have a system able to take raw data from the FE either in tests with beam particles or in the lab for the qualification of the assembly modules. This is the reason why a VELO GWT data bypass was developed as explained in subsection 4.2.15. The LHCb framework had already implemented a bypass mechanism but it didn't fulfil the VELO requirements as it can only handle GBT data and bypass 1 link at a time. On the other hand, the VELO bypass handles up to 20 GWT links at the same time, synchronizes the data and timestamps it. This bypass can handle a maximum throughput of 16Gb/s. The drawback is that it can not be integrated on the experiment common firmware due to the high cost in terms of resources.

5.6 MiniDAQ 3 readout system

This setup is used in parallel with the MiniDAQ 2 and they are almost identical, but for the MiniDAQ 3 using the version 2 of the PCIe40 board. After extensive tests on the PCIe40_v1, some changes were made to the new version (v2) of the board:

- Rearranging the power mezzanine boards with new DC/DC components and better cool-

ing air flow.

- Removal of the PCIe40 bridge to reduce costs, moving from a 16 lanes PCIe to 2 parallel 8 lanes PCIe.
- A second SFP connector for reducing the number of TFC/ECS modules and cost.
- New PLLs to improve the jitter filtering and phase.
- Implementation of an Intelligent Platform Management Interface (IPMI) monitoring to check the hardware status of the board.
- An EEPROM to track the board during production.
- Routing optimization by using staggered bias.
- Power sequencing improving to ease maintenance and increase the longevity of the module.
- New better heat sink.
- New USB Blaster to increase the programming speed by a factor 4.

Most of the hardware improvements explained above involve firmware changes. Most of them are common for the whole LHCb collaboration, but in the particular case of the VELO sub-detector, a new GWT LLI is needed due to the FPGA pinout changes and the new jitter cleaner.

The MiniDAQ 3 will also be dedicated to test the final readout system, hardware and firmware wise. It was not available for the VELO group until 2020 and thus these studies are not included in this thesis.

5.7 Hardware validation

Determining that the hardware design is compliant with the requirements and making all the pieces work together within the full complexity of the system (described in chapter 3) requires planning and sequencing. These are common for all versions of the hardware and this section describes the followed procedure:

Device validation. A first step could be to ensure that every component by itself works as expected, discarding malfunctioning devices or PCBs with production or design failures. The usual way of performing this is with an exhaustive checking of the hardware under the microscope and by probing with a multimeter.

Due to the very specific design of the transmission lines, special care were taken in their validation. A Vector Network Analyzer (VNA) in combination with a PRBS generator were used for the transmission losses measurements (see Figure 5.6) and a Time-Domain Reflectometer (TDR) was used for the impedance characterization (see Figure 5.7). These tools facilitated the diagnosis of a great (but still acceptable) degradation in the transmission of the signal. In order to improve the signal integrity, a CTLE circuit (high-pass circuit with the transfer function showed in Figure 5.8) was implemented in all the transmission lines, at the end of the line for the data acquisition and at the beginning for the control.

Moreover, the CTLE circuit adds an inductor and a resistor to compensate the impedance




Figure 5.6: Data transmission loss, measured in the final prototypes, for all the FE data acquisition links. Test performed in the copper lines from including high-speed data tapes, VFB and OPB.




Figure 5.7: Characteristic impedance (final design) per link for all the FE links.

of the following component in the chain.

By design, the Nyquist Frequency of the CTLE is 2.4 GHz for the control links and 2.56 GHz for the data readout links, with a gain of -2.1 dB and -1.6 dB respectively. The DC gain is -9.5 dB (control) and -8.5 dB (data). The effect of the CTLE circuit can be




Figure 5.8: Schematic of the CTLE circuit.

seen in Figure 5.6, it is represented by a ridge with its maximum around the working data rate (~ 2.5 Gb/s).

Board validation. The next step consists of powering on the OPB, configuring the GBTx and connect it with the GBT-SCA. In the default configuration of the GBTx placed in the OPB, the clock is recovered from the input data and the communication with the GBT-SCA placed in the OPB is given by the bit position in the GBT frame (see section 4.3).

GBTx ASICs are, by default, not configured, therefore they must be configured via their I2C bus just after being powered on. Prototype boards of the VELO sub-detector provide I2C access through an onboard connector, while for the final design this configuration will be fused into the ASICs before assembly.

OPB optimization. All the integrated circuits of the OPB are switched on/off through the GBT-SCAs over General Purpose Input Outputs (GPIOs) (DC/DC converters, VTRx and VTTx) and configured over I2C (GBLD and electro/optical converters). The GBLD




Figure 5.9: Optimisation of the control link transmission running at 4.8 Gb/s.

ASICs present in the hardware have two different functionalities: they are mounted the optical converters (VTTx or VTRx) and they need to be tuned to amplify the electrical signal within the working range of the laser, this optimization was done by maximizing the BER after the optical conversion. The GBLDs are also used at the beginning of the module control links, making a pre-emphasis of the signal to guarantee the signal integrity after travelling almost one meter over copper. Figure 5.9 shows the eye diagrams of the module control link. Starting from the VTRx, where the optical signal is converted into electrical in the down-link, amplified in the GBLD and filtered in the CTLE before travelling through the flexible tapes until reaching the GBTx placed in the module. A similar optimization happens on the up-link.

Data link validation in loop-back mode. At this point, with the control down-link of the module configured, the next step is the configuration of the GBTx located in the module to recover the data from the link and distribute them to the 3 FE ASICs (first version of the hardware), or 6 FE ASICs (second and final version of the hardware).

As the VeloPix ECS protocol is completely different from the rest of the FEs of the experiment, the testing and configuration of the GBTx was more exhaustive. Before explaining these tests, let me open a side note here to briefly explain how the GBTx high speed data path works (more detailed information can be found in the manual [58]). Figure 5.10 ex-




Figure 5.10: GBTx data path block diagram.

plains its behaviour. The top/bottom rows of the figure are for the receiving/transmitting data paths. The first stage, for the receiving path, is the deserializing block, then a series of frame swappers align the word, forward error correction can be applied (on demand), data is de-scrambled from the optical transmission and sent out of the chip. The bottom row shows the opposite data path. Three different loop-back paths can be selected: after the frame swapper, after the forward error correction decoder and after the de-scrambler.

To test the transmission between the BE board and the GBTx placed in the module, the first step is to create an internal GBTx loop-back after the de-scrambler. The quality of the link was proved to work by sending a test message from the server (explained in subsection 5.4.1), down to the module and back to the server in a reliable way.

Once the quality of the link has been tested and verifying that what arrives to the MiniDAQ server is what has been sent to the GBTx, it is time to move forward and exchange the internal loop-back for an external one with a PCB connecting the GBTx ECS transmitter and the receiver. Previous tests can now be repeated while probing at the same time with a scope, as can be seen in Figure 5.11, to check that the data fed to the VeloPix complies with the requirements of voltage, timing and frequency. The same procedure is performed for the eports with the TFC commands and clocks.

This step allowed us to find design and hardware errors in the first prototype versions, as a swap in the ECS differential lines being sent to the VeloPix, without damaging rare and valuables ASICs.




Figure 5.11: Capture of the ECS data sent to the VeloPix.

Low level communication and validation of the FE. At this point, the system is ready to communicate with the different VeloPix ASICs. The first approach was to read a single register that, by default, is not zero and compare it with the expected value, then to modify it and read it back again.

High level communication and validation of the FE. The following natural step is to perform high level communication with the FE ASIC, by applying more complex functions as equalizing or scanning the noise and taking data over the ECS links. Unfortunately, no high level software applications were available for these purposes at this point. One of the main premises for the design of previously mentioned MiniDAQ server was to make it as compatible as possible with the software applications, already developed for the SPIDR setup. To do so, a series of software libraries were created. They communicate with the MiniDAQ server from the control computer and “fake” the functionalities of the SPIDR system not yet present in the final readout (Carrier board temperatures, voltages, etc.) to avoid the crashing of the software applications.

Software applications validation. The software developments explained in the previous subsection do not avoid the modification of the software applications. Nonetheless, they significantly simplify the migration of the software applications. The Migration of the IP address and port of the client (control computer)/server(SPIDR, MiniDAQ 1 or 2). The SPIDR setup used the IP address “192.168.1.1” (server) and “192.168.1.10” (client). Meanwhile, MiniDAQ 1 uses “192.168.1.11” for the client and MiniDAQ 2 uses the local-host as the control computer is embedded in the same computer as the server. MiniDAQ setups allow multiple VeloPix handling, not included for SPIDR, requiring the modification of the software applications to specify the VeloPix to handle. This is done by editing the communication port from the SPIDR unique port (50000) to the 6 communication ports of MiniDAQ setups (50000 to 50005).

FE operation with software applications. With the software applications in place, it is time to use them to equalize the VeloPix and take data through the ECS links. These are important steps because they allow to validate that the VeloPix responds correctly, masking the noisy pixels that will flood the readout links and degrade its performance. Figure 5.12 shows in




Figure 5.12: (Top) VeloPix noise counts vs global threshold before and after equalization. (Bottom) VeloPix matrix noise scan.

the top side the noise counts vs global threshold before and after a VeloPix equalization⁴, in the bottom side, it also shows the data acquired in absence of any particle source, noise scan. With no pixels giving hits because of the noise and with a very small number of pixel masked, we can continue with the tuning of the readout chain.

Readout link optimization. The first step for testing the performance of the readout links consisted of configuring the VeloPix transmitter to send PRBS data⁵. Probing with a high-speed real time scope at the input of the OPB optical converter allows measuring the

⁴The VeloPix ASIC has two parameters used for the equalization: a 4 bits digital to analog converter on each pixel and a global threshold for the chip. The equalization of the ASIC consists on performing two count-over-threshold or noise scans over the global threshold, one with pixel converters set to their minimum value (left small red Gaussian) and one to their maximum (right small blue Gaussian). Then pixel converters are adjusted to equalize the noise level of all pixels to produce the middle narrow curve.

⁵It is recommended to equalize and mask noisy pixels before validate the GWT links to prevent the serializer to be stucked with useless data.




Figure 5.13: BER of the GWT links with respect to the phase of the 320MHz clock. BERs below 10^{-12} is considered acceptable, as they represent less than 1 error every 200s.




Figure 5.14: VeloPix BER vs digital/analog power supply (measured on the ASIC). Tests were done on VeloPix v1, MiniDAQ v1, and first prototypes of the on-detector electronics.

quality of the transmission, which, by default, is not good enough. Improvements of the data transmission were done by optimizing the VeloPix GWT serializer and improving the link transmission with a CTLE circuit.

For the VeloPix, the synchronicity between the clock distributed to the matrix and the clock used to serialize the GWT (at 320MHz) frame causes an instability in the data transmission. It is known that changing the phase of the 320MHz clock helps to improve the serializer performance⁶. Thus, a scan of the 128 possible phases is required. As finding the best phase among all the possibilities is a long and tedious task, a software application was implemented for the MiniDAQ systems. It communicates with the VeloPix and with a custom made PRBS checker embedded on the LLI for MiniDAQ 1 or the transceiver for MiniDAQ 2.

The working phase was found to be for all VeloPixes about 40 as shown in Figure 5.13. Nonetheless, the working point degrades with temperature and supply voltage. Lower temperature and higher supply voltage reduce the BER. As the working temperature is already low, it was decided to increase the supply voltage measured on chip from 1.2 V

⁶The final version of the VeloPix ASIC improved the stability of the 320MHz clock and the design of the serializer. Unfortunately, the optimization in the phase is still needed for maximizing the performance of the link.

to 1.3 V. Figure 5.14 shows that for proper behaviour of the high speed links (BER below 10^{-12}), 1.3 V is needed for digital and analog power supplies. The improvements on the transmission line are based on a CTLE circuit with a Nyquist on 2.56 GHz. Moreover, in the second version of the VeloPix Hybrids two links behaved worse than expected due to an asymmetry in the displacement of the capacitors nearby the transmission line.

Firmware optimization. The following step is the optimization of the different firmware components of the BE board (section 4.2), but as these are specific to each version of the MiniDAQ systems they are described in their corresponding sections (section 5.4 for MiniDAQ 1, section 5.5 for MiniDAQ 2 and section 5.6 for the final PCIe40 board).

System integration. Finally, most of these steps are automatized and integrated into the experiment control project, JCOP, by migrating the MiniDAQ server to the experiment server, migrating the SPIDR software applications, etc. These are all software tasks, on which I am not involved, therefore, it is not covered in the subject of this thesis.

5.8 VELO Tests with beam particles

In October 2018, in the last opportunity for tests with beam particles of the LHC run 2 at the SPS North Area beam facility, the LHCb collaboration planned a large test campaign. In the particular case of the VELO sub-detector, it was a great challenge due to the fact that the first module prototypes had a very tight production schedule. Moreover, the readout board firmware design was far from being ready in time. Nonetheless, it was possible to attend it because the production sites that managed to build 1 just on time and 2 more along the beam campaign. A specific firmware specifically designed for this test beam (subsection 4.2.15) was later used in the production sites for QA of the modules.

Even though the SPS hit rate is not comparable with the expected particle rate of the upgraded VELO (3-4 orders of magnitude smaller), the goal of the tests was to validate the whole readout chain of a slice of the detector. To do so, 114 GB of data were taken over the GWT links synchronously with a TimePix 3 telescope [45] allowing the measurement of the same tracks in both systems and checking the synchronization between planes with less than 1 ns.




Figure 5.15: Picture of the 3 VELO slices during the test beam.

Different tests were performed over the control links to the three planes under test:

- Sensor bias voltage scan.
- VeloPix analog gain, threshold and time-over-threshold scans.
- VeloPix GWT clock-phase-scan.
- Time-of-flight alignment between slices.

First beam tracks of the VELO upgrade were recorded (see Figure 5.16). In addition, each of the three planes under test was correlated in time and space with the TimePix 3 telescope.




Figure 5.16: Preliminary VELO telescope data.

CHAPTER 6

TEST BENCHES FOR PRODUCTION QUALITY ASSURANCE

The IGFAE group contributes to the production of the phase I upgrade of the LHCb VELO by: mounting the high speed flexible data tapes and the high voltage tapes described in subsection 3.3.1, designing and mounting the carrier board for testing the VeloPix ASIC and validating/commissioning the hardware. Once the high speed flexible data tapes are mounted, it is mandatory to guaranty that they fulfil the quality requirements of data transmission performance and support the operational voltage before being assembled in the detector. A test bench setup was designed for high speed and high voltage tapes (Figure 6.1).

6.1 High speed tapes test bench




Figure 6.1: High speed flex tapes test bench.

This setup was designed based on the premise of having an automatic way of validating the performance of the FE data transmission links and the unipolar links used for the FE temperature sensors and analog output of the VeloPix. It is based on the same evaluation board, Xilinx VC707, used by the SPIDR system for the VeloPix validation. The firmware developed for this purpose takes an external clock 4 times the LHC (160 MHz) and it uses it as a reference to generate the emulated VeloPix data, at 5.13 Gb/s. The emulated data are transmitted in loopback through the data tapes. Moreover, we performed a scan in the phase and voltage of the emulated GWT signal that produces a pseudo eye diagram for each link. On the other hand, unipolar links are tested by sending a LVCMOS18 signal in loopback for each link. On the receiving side, continuity is checked as well as any shortcut with adjacent traces.

The alignment of the connector in the tape is crucial. A small displacement of the connector can prevent the other far end of the tape for being connected. To measure it two cameras are deployed, one on each end of the connector measuring the displacement in a millimetre PCB ruler as it is shown on the left side of Figure 6.2.




Figure 6.2: (Left) Alignment pictures showing the displacement of the cable when the opposite far end is connected. (Right) Example of an eye diagram generated for every single link of all the cables.

At the end of the test, an automatic report is generated with all the production and test data ready to be integrated into the VELO production database. Along with this report, eye diagrams and alignment photos are saved into the database. The automation of these tests is done via two layers of code. A high level layer is written in a linux bash script, allowing the interface with the operator of the test, collecting the information needed for the test (serial number, type of tape, batch, alignment photographs, production report) and producing test output files. A low level layer is written in a Tcl script, it takes the information introduced by the operator in the top layer, interacts with the firmware performing the actual hardware tests and saves the eye diagram information of each link.

The QA acceptance criterion for these tapes is based on the measured misalignment less than 4mm, the eye diagram of all the links of the cable is within a 30% of the mean value per link, and the unipolar links transmission is correct.

6.2 High voltage tapes test bench

For the validation of the high voltage tapes, the approach is similar to the high speed tapes. A test bench was design based on two high voltage power supplies (Keithley 2410) and a PCB with several LEDs enabling a simple diagnostic method to find shortcuts or open circuits. Once the tape has passed the continuity test, the LED board is removed to perform the leak test by applying 1000 V during 20 seconds.




Figure 6.3: (Left) Continuity test setup. (Right) Open circuit for 1000V test setup.

The automation of these test is again divided into two layers. An user interface layer, written in linux bash shell, gathers the information of the tape (serial number, kind of tape, batch, production report, etc.) and generates the output file. A custom made library, also written in bash, interacts with the high voltage power supplies and sends the information provided on the top layer and returns the current consumption.

6.3 VeloPix Carrier board QA

It was IGFAE's responsibility to design and produce the FMC carrier board compatible with the SPIDR system, used for the validation of the VeloPix ASICs, see section 5.3. To guarantee the quality of the carrier boards, a specific firmware for the VC707 board was developed. This firmware emulates the signals that control the VeloPix and the different components present in the carrier board like the DC/DC converter control signals, the cooling fan Pulse Width Modulated (PWM) control signals, the DC/DC that provides the VeloPix fuse levels and the mounting LEDs. The control of the firmware was made through Vivado hardware manager logic analyzer,

more information is shown in subsection 5.2.3.3.

All the functionalities of the carrier board were tested at IGFAE's probe station by probing in the bonding paths of the PCB, as shown in Figure 6.4.




Figure 6.4: Carrier board test setup.

CHAPTER 7

CONCLUSIONS AND FUTURE WORK

The work described in this thesis shows the design of the Back-End firmware and the validation and tune-up for the different parts of on detector electronics hardware. This design is framed under the phase I upgrade of the LHCb Vertex Locator sub-detector that is taking place during the LHC LS2. Some minor contributions developing specific firmware for testing and Quality Assurance during production are also shown.

The overall level of compliance with the objectives was very high, as the hardware was probed to be operational and has been constructed. From the Back-End firmware point of view, the control and timing firmware for VELO is operational, fully tested in the lab and in tests with beam particles at CERN SPS facilities. It was tested with several versions of the hardware and, now, it is ready for the commissioning of the system. On the other hand, several versions of the DAQ firmware were developed for testing the on-detector electronics and for the QA of the production. The experiment DAQ firmware was created and simulated. However, given the limited manpower resources and as some firmware modifications may improve the downstream computer resources, it is still under development. We are currently finalizing the implementation of the basic firmware and performing extensive simulations and hardware tests (on PCIe40 board) that will result in the basic firmware for the physics data taking and control.

New functionalities are going to be developed in the TELL40 and released in future versions of the firmware. These versions will include a larger firmware, monitoring, errors and warnings handling, an optimization in the Post-router for handling higher data rates and possibly extra bit sorting.

ANTONIO FERNÁNDEZ PRIETO

Finally, all the firmware shown in this thesis needs to be commissioned within the rest of the experiment.

CHAPTER 8

SUMMARY OF THE THESIS

This PhD thesis summarizes the the tasks I devoted to design and verification of the electronics for the LHCb VERTex LOcator (VELO), on which I am involved as an IGFAE member since december 2015.

8.1 Context

The Standard Model (SM) of particle physics is the frame within it is possible to explain the structure of the relationships between all the fundamental particles. It is the most reliable model that we know. It asserts that all the matter present in the universe is built of elementary particles called fermions, which interact through fields. The particles associated with the interacting fields are called bosons. The SM leaves some unanswered questions like how it is possible the asymmetry between matter and antimatter. During a small fraction of time, right after the Big Bang, there was an equilibrium between matter and antimatter that disappeared when the universe expanded and cool down. This phenomena is only possible under CP violation conditions and it understood within the SM. Notwithstanding, the magnitude predicted by this model is way smaller than the observed in the universe. Multitude of theories arise in consequence. Prove or disprove those with the help of accelerators and particle detectors is one of the tasks of experimental physics.

The world largest particle accelerator (LHC) is placed in CERN facilities under the French-Swiss border nearby Geneva. The LHC collides two proton beams most of the time, but there are some short special runs with lead and in the future with oxygen too. The proton acceleration

process starts on a hydrogen bottle. First, the hydrogen passes through a strong electric field that removes the electrons. After that, the protons are injected in a series of linear accelerators and synchrotrons (PSB, PS, SPS) before being injected to the LHC.

This collider has four interaction points, where the two proton beams collide, with four detectors: ATLAS, CMS, ALICE and LHCb. The three first have a barrel shape around the interaction point with an almost cylindrical detecting region. Two of the detectors are general purpose, ATLAS and CMS, or in other words, they are focus on a quest for new physics through the production of new particles. ALICE, on the other hand, studies the particles generated from heavy ions collisions (mostly lead). Finally, LHCb, which is the frame of this thesis, was originally designed to investigate the small differences between matter and antimatter through the study of b quarks. In spite of that, this experiment showed an extraordinary performance during “Runs” I and II of the LHC (2010-2018) expanding the physics studies to a general purpose detector in the forward region. Unlike the other experiments, the LHCb is built to measure only a small pseudo-conical region on which the interaction point is located in the vertex of the cone and the beam flows through its axis in opposite directions. The angular opening of this detector is 17° .

All particle physics detectors placed in circular colliders have several similarities in terms of the systems that conform them. From now on, I will make a brief description, from the engineering point of view, of how the sub-particles generated in proton collisions are detected and using as an example the original design of LHCb experiment. The main goals of these detectors are identifying the trajectory, energy and type of the particles. Each of these goals are achieved in different sub-systems that are spatially arranged situating the systems with less material close to the interaction point and, thus, minimizing the interference between systems. Besides that, it is common to use a magnet able to bend the trajectory of the electrically charged particles. For example, LHCb uses a dipole magnet that generates a uniform magnetic field of 4 Tm bending the particles vertically.

Let us delve into each of the detection systems. In first place (starting from the interaction point), we find the tracking detectors, responsible for measuring the points where the particles pass through and reconstruct the trajectory of the particle. This type of sub-detectors has several design constraints: they have to minimize the material budget to not interfere with the downstream detectors, they must have high efficiency, precision and granularity, allowing the tracing of the particles trajectory and avoiding pileup. The most common technologies are silicon sen-

sors (either pixels or strips) or scintillating fibres. In the case of LHCb, it combines two tracking stations located upstream and downstream the dipole magnet along with a characteristic vertex detector (VELO). The VELO sub-detector was designed to improve the study of the bottom quarks, which requires a high accuracy differentiating the origin of the reconstructed trajectories are either proton-proton collisions, what we call primary vertices, or the short decay of particle generated in a primary vertex, what we call secondary vertices. The VELO sub-detector consists of a series of stations perpendicular to the beam axis surrounding the interaction point. They are divided in two retractable halves that can be separated from the beam to avoid damages in the detective region when the LHC has no stable beams.

Another peculiarity of the LHCb experiment is its pion and kaon identification system, based on measuring the Cherenkov light rings¹ on the so-called RICH detectors. The Cherenkov rings of light are reflected in mirrors that drive them outside the detector acceptance, where they are converted into electrical signals in photo-detectors. The LHCb experiment has two RICH stations, one between the VELO sub-detector and the first tracking station, and the second one between the last tracking station and the first calorimeter.

The identification of electrons and hadrons, as well as the measurement of their energies, are tasks for the calorimeters. The technology used for those purposes consists of a series of sandwiches that alternates metallic layers (usually iron or lead) with scintillating plastics. The metallic layers stop the particles at a certain energy depending on the depth and creates a “shower” of ultraviolet light in the scintillating plastics, which is captured in photo-multipliers outside the acceptance.

The last systems are dedicated to measure muons, the heaviest charged particles detected in particle physics experiments. Muons are able to pass through all the previous systems. The concept of these sub-detectors rely on stopping them with iron layers able to create a shower similar to calorimeters.

Capturing the very rare and interesting events from the physics point of view, requires the generation of a huge amount of collisions per second, in the LHC case the peak beam crossing rate is 40MHz. It creates a huge amount of data in all systems described above, not feasible to be stored on disk. Choosing the events to be stored on disk is the task of the so-called “trigger” system. The LHCb experiment originally used a two-layer trigger system. The first layer was a

¹Cherenkov light is produced when charged particles passes through a certain medium (as very dense gas, C₄F₁₀) faster than the speed of light, then it causes a conical light spectrum. The RICH measures the projection of this cone.

hardware limit that processed the data in real time of VELO, calorimeters and muon detectors and the second layer selected the interesting events by running software algorithms in a CPU farm. With this trigger the data stored in disk was reduced by a factor 3200.

8.2 LHCb upgrade




Figure 8.1: Longitudinal section of the LHCb upgrade experiment.

Even though the LHCb experiment had an excellent performance during LHC Runs I and II (2010-2018), where it acquired 9.2 fb^{-1} of proton-proton collisions, most measurements are limited by its statistical uncertainty. A major upgrade is undergoing during 2019 and 2020 to solve the weaknesses of the experiment (see Figure 8.1): new VELO, new tracking system, and new trigger and readout systems to increase the data acquisition. Starting with the trigger and readout systems, in the new architecture the hardware trigger is going to be removed and therefore, the FE electronics will transmit the data to the BE at the collision rate (up to 40MHz) and the trigger will be entirely software, allowing more flexibility. As a result, all the sub-detector

electronics need to be redesigned. With the aim of minimizing the developments, a common design is done to standardize all systems. The common effort is focused on the data acquisition, control and synchronicity of the experiment. The other improvements of each detector are described below:

Vertex Locator (VELO) The upgraded VELO needs to fulfil several requirements: being radiation tolerant in the new LHC environment, improve the granularity for a robust track reconstruction, minimizing the material in the region closest to the interaction point to optimize the performance of the detector. However, the VELO sub-detector will still maintain the two retractable halves surrounding the interaction point, but in this case the number of stations will be increased from 21 to 26. Moreover, the minimal sensor distance to the collisions is reduced down to 5.1 mm and they will be only apart from the collisions by an aluminum foil of $150\text{ }\mu\text{m}$ that separates the primary vacuum of the LHC machine from the secondary VELO vacuum. Each station is split in two slices, those slices are made up by a module, in the region closest to the interaction point, four sensors (two per side) are located in an L shape around the beam. Each sensor is built with an “n-on-p” technology and it is bump-bonded to three custom design ASICs called VeloPix derived from TimePix family. The control of these VeloPixes is done directly through a high speed link by the GBTx ASIC, also present in the module. In order to dissipate the heat generated in the module, a novel cooling system was designed. It consists on building capillaries between two bonded silicon wafers used as a substrate to the module. These capillaries increase their section underneath the ASICs, allowing the expansion of the liquid coolant (CO_2), which is close to the gas transition state. The change of phase into gas removes the heat and keeps the module at stable temperature (-20°C at the working pressure).

The communication of the modules with the outside is done through flexible tapes to make possible their telescopic displacement, an interface PCB, designed to support the differential pressure between the secondary vacuum and the atmospheric pressure, and a service PCB, which takes care of the module power supply and electro optical conversion.

Ring-Imaging Cherenkov Detectors (RICH) The optical systems of the RICH subdetector need to be upgraded due to the increase in luminosity. In particular, the first station will improve the focal distance of its mirrors from 2.7 m to 3.7 m and thus will reduce the

occupancy of the photo-detectors. The cooling system will be improved with the aim of easing the access to the detector for maintenance.

Tracking system The new LHCb design will maintain two tracking stations, UT and SciFi but with a new design.

The UT will be placed between the first RICH station and the dipole magnet. It will increase the resolution of the momentum of the particles as well as radiation tolerance. The UT sub-detector has four stations based on silicon strips. The technology of the sensors, as well as the dimensions vary with the position within the sub-detector: the innermost region, with higher irradiation, uses “n-on-p” technology with small strips. The outer region will have “p-on-n” technology with larger pitch due to the lower radiation.

The SciFi sub-detector will be placed after the dipole magnet. It has 12 layers allowing a reconstruction of the tracks and the measurement of the momentum. Each layer is built by scintillating fibres of 2.4m long and 250 μm of diameter. They are arranged in vertical with a depth of 6 fibres per layer. The electrical conversion is done at the far ends of the fibers by silicon photo-multipliers (SiPM).

Calorimeters This system will only suffer minor changes. The photo-multiplier tubes will reduce their gain to maintain the same anode current and minimize the degradation.

Muon detector This system will not suffer any upgrade, besides the LHCb common upgrade in terms of readout electronics.

Online and trigger systems Even though this is not a system installed in the cavern of the experiment, it is advisable to describe the system called “online” that comprises the common developments in terms of experiment control and data acquisition.

Figure 8.2 shows the experiment DAQ and control systems. On this architecture, all sub-detectors shown above (FE) transmit the acquired data at a 40MHz rate to the surface computer “farm” over optical fibres. The farm tasks are controlling and acquiring all sub-detectors data. The data received are processed in real time on dedicated FPGAs with dedicated sub-detector’s firmware. FPGAs are connected to a PC via a PCIe express at 100Gb/s and from there on the data follow a standard format for all LHCb sub-detectors. All the computers of the farm are interconnected through a 100Gb/s network, allowing a distributed trigger processing that decides the events to be stored.




Figure 8.2: LHCb readout system.

8.3 The LHCb VELO upgrade

The work shown on this thesis is divided in two main blocks. Firstly, the validation of the hardware design of several prototypes for the VELO upgrade, finding errors, improving the design and finding the optimal configuration for the different ASICs. Secondly, the design and implementation of the FPGA control, distribution and DAQ firmwares as well as testing the different FPGA models used for prototyping.

8.3.1 Underground electronics

This section describes, from the electronics point of view, the VELO sub-detector slice (Figure 8.3). The electronic design was made by other members of the collaboration and my contri-

bution was the validation and setting up of prototypes and system.

The VELO module can be seen on Figure 8.3. Starting from the LHC beam, we find the tile sensors, Hamamatsu “n-on-p” silicon pixels optimized for the LHC environment and bonded to three VeloPix ASICs through flip chip 3D encapsulating technology. The communication of these tiles with the rest of the system is done through wire bonds to an auxiliary PCB called hybrid. These, in turn, transfer to the VeloPix the control and timing signals from a third hybrid hosting a GBTx ASIC and transfer from the VeloPix the acquired data outside the module, over twenty 5.13 Gb/s links.

All the communications with the module (control and DAQ) are driven through: flexible PCBs of ~ 50 cm that absorb the displacement of the module, a Vacuum Feedthrough Board (VFB) and an Opto-Power Board (OPB) that sends the data out of the cavern electronics.

Put to work a slice requires multitude of steps:

1. At the arrival of the prototypes, it is necessary a quick visual inspection of the different PCBs, as well as checking for shorts or open connections in different parts, with the help of a multimeter. Besides that, the integrity of the high speed transmission lines was tested with the help of VNA and TDR.
2. With the prototype assembled, we start up the different parts of the system step by step. In the first place the OPB, controlled over a GBTx acting and monitoring all the active components of the board. An optimization of all the optical links over the GBT-SCA is needed in order to minimize transmission errors.




Figure 8.3: (Left) VELO “slice” cross section. (Right) Module zoom.

3. Before starting the communication with the VeloPix, it is necessary to configure the GBTx of the module to propagate the control and TFC signals at the right speed and ports.
4. At this point it is possible to write and read VeloPix registers over the control link, allowing the equalization and data acquisition and testing the behaviour of the chip.
5. The optimization of the GWT transmission is required before taking data over them. It is attained by the combination of three factors: first and the most important, adjusting the phase of the clock that serializes the GWT word. Second, implementing a high pass filter (CTLE) to remove the external noise on the link. Third, amplifying the signal before converting into optical. With this tuning, a BER of 10^{-12} is achieved over 400m of optical fibre.

8.3.2 Electronics at the ground level

The data center from where the underground experiment is controlled and readout is located in the ground level close to the detector. Figure 8.4 shows graphically the control and synchronicity of the experiment. Maintaining the detector aligned is vital for the operation of the detector. Hence, a hierarchical design was implemented, where the sole communication point with the LHC is the readout supervisor board (S-ODIN) which receives the LHC clock and distributes it to the whole experiment. This board also contains the filling scheme of the beams colliding in the LHCb. Subsequently, S-ODIN applies and distributes the trigger signal for the events where there are no collisions through the TFC system along with other signals associated with a BXID. The distribution of the TFC commands and the clock to the system is performed by the interface boards (SOL40) along with the Experiment Control System (ECS) commands. ECS signals allow the configuration and monitoring of the detector over a series of SCADAs and computers connected to the SOL40. On a different path, the data read out of the cavern are processed, in real time and create the events by the data acquisition boards (TELL40).

The VELO sub-detector has several design constraints, given the fact that it has the closest sensors to the interaction point in the entire LHC, it must reduce the material budget in the active region, keeping the sensors and electronics at -20°C and read out at 40MHz. These constraints compel a more challenging design than other subdetectors of the experiment: sending the data out of the FE ASIC unsorted in time, using a different high speed link serializer (GWT) that




Figure 8.4: LHCb readout architecture.

reduces the power consumption, controlling the VeloPix directly over the high speed control link (GBTx).

Those constraints have implications in the surface electronics design. From the SOL40 point of view, a small modification in the LHCb framework is needed, allowing the communication with the VeloPix over the GBTx link. Besides that, from the TELL40 point of view, it requires almost a new firmware to handle the VeloPix unsorted data.

8.3.2.1 VELO control firmware

The control signals provided to the VeloPix have two different natures: configure and monitoring signals (ECS) and fast control and timing signals (TFC). Even though both type of signals share the same link, the first ones have a bi-directional link running at 80Mb/s while the second ones a unidirectional link (towards the VeloPix) running at 320Mb/s. The communication bus between the SOL40 board and the GBTx uses an 80 bits word running at 40MHz. Hence, each ECS signal uses 2 bits and the TFC 8 bits. The associated firmware that packets this bus was design specifically for the VELO sub-detector.

8.3.2.2 VELO readout firmware

The only common development of the VELO sub-detector with the LHCb firmware framework are the PCIexpress and the block packaging, the rest of the firmware is custom made.

Each VELO TELL40 is designed to match with a module, and thus reading 20 GWT input links. The Low Level Interface (LLI) handles the optical link recovery, decoding of the data and transmission error identification and, finally, bridges the data to the FPGA processing clock domain.

The real time data processing consists basically in two tasks: the main task is to sort the VeloPix data by its timestamp, the second optional task (it will only be implemented if the main tasks is small enough to fit extra processing in the FPGA) consists on de-allocating resources from the computer farm and allocating those in the FPGA, a good candidate is the clustering of the hits.

The sorting task is performed in the so-called “Router”, which is an algorithm capable of sorting and mixing the data of half a module in real time (50Gb/s of throughput). The Router algorithm is divided in two blocks. The first one, takes care of the 4 MSB of the received GWT BXID, and it uses a network of pseudo decision trees on which each node of the three is a FIFO memory. This network absorbs the high throughput at the same time that the data are being sorted. The second block, which handles the remaining 5 bits of the GWT BXID, changes the way of sorting the bits to reduce the memory resources. It is based on lookup tables that assigns the timestamp with the address of one of the sixteen RAM memories where the sorted data are temporarily stored. Unfortunately, the Router, by its own, is not able to process all the data, and thus two conditioning and processing stages (Pre-Router and Post-Router) are placed upstream

and downstream of the Router. The Pre-Router takes care of flagging the data according to the VeloPix of origin, identifies the start of transmission commands and alerts the Router and Post-Router to start the data taking and finally re-clock the data from the 40MHz FPGA clock domain to the VELO data processing clock of 160MHz. On the other hand, the Post-Router manages the data taking of the VELO data processing. It processes the TFC commands, completes the BXID with 3 extra bits to avoid mix-up of the data² and formats the output following the Intel standard Avalon-ST.

There is also the possibility of doing extra data processing in the FPGA like flagging the isolated clusters and clustering the data.

In terms of FPGA resources, the current figures are 41% for logic utilization and 72% for memory. With these numbers in mind, there is some extra space where to place some of the additional functionalities mentioned above. The estimations, with the clustering included, are expected to increase to 73% for logic and 88% for memory. The estimations for the block that identifies isolated hits in its original version used many resources, but now with a conceptual change (flagging by using the hamming distance properties of the Gray encoding) will reduce dramatically the impact in terms of resources.

8.3.3 QA firmware

8.3.3.1 Bypass firmware

The development of the VELO firmware is a complex and endurance task. However, the hardware developments need to be tested and validated both in the lab and in production. A simplified version of the firmware was designed for these matters, it bypasses the data recovered from the fibre directly from the LLI to the PCIexpress to be saved into disk. This firmware has a dedicated FIFO memory per input link, as well as a management system that decides which link is being propagated to the output according to the occupancy of each FIFO. The firmware was also used in a test with particle beams in the SPS test facilities at CERN, for that it was modified to be synchronized with an external telescope used as a reference.

²Flagging and entire LHC orbit requires 12 bits. VeloPix only transmits 9 bits of data to save data bandwidth, thus the TELL40 must reconstruct the remaining bits.

8.3.3.2 Test-bench for High Speed and High Voltage tapes

One of the tasks of IGFAE in the LHCb upgrade was the assembly of the high speed and high voltage tapes placed between the VFB and the module. Two test benches were designed for testing the quality of the tapes using a FPGA Xilinx Virtex 7 evaluation board.

CHAPTER 9

RESUMO DA TESE EN GALEGO

Esta tese abrangue o deseño e verificación da electrónica necesaria para a mellora do detector de vértices (VELO) do experimento LHCb no que estou involucrado como membro do Instituto Galego de Física de Altas Enerxías (IGFAE) dende decembro de 2015.

9.1 Contexto

O Modelo Estándar (SM) da Física de Partículas é o marco que explica a estrutura e relacións de todas as partículas fundamentais máis fiable ata o momento. O SM afirma que a materia do universo está composta de partículas elementais chamadas fermións, que interactúan por medio de campos. As partículas asociadas ás interaccións dos campos son os chamados bosóns. Esta teoría deixa algunhas cuestións sen resolver completamente, como é o caso da asimetría entre a materia e a antimateria. Despois do Big-Bang, durante un curto espazo de tempo mantívose un equilibrio entre materia e antimateria. Cando o universo comezou a expandirse e arrefriou pasou por unha serie de cambios na súa composición. As partículas adquiriron as súas masas características e un fenómeno ocorreu que diferenciou a materia da antimateria, causando asimetría entre as dúas. Este fenómeno, só é posible baixo as condicións de violación da Carga-Paridade (CP), como recolle o SM. Non obstante, a magnitude da violación CP predita non é suficiente para explicar, dominio de materia sobre antimateria no universo actual. Múltiples teorías de nova física xorden para intentar explicar este e outros fenómenos. Unha das labores principais da física experimental é probar ou desmentir estas teorías coa axuda de experimentos deseñados expresamente con ese fin.

O maior acelerador de partículas do mundo (LHC) atópase nas instalacións do CERN na fronteira franco-suíza nas inmediacións de Xenebra (Suíza). Nel fanse colidir dous feixes de protóns a maior parte do tempo, aínda que en ocasións tamén se fan colidir outras partículas, por exemplo chumbo e nun futuro oxíxeno. O proceso de aceleración e colisións no caso dos protóns comeza dende unha simple botella de hidróxeno en estado gasoso. O hidróxeno faise pasar a través dun cilindro metálico no que se aplica un campo magnético capaz de eliminar os electróns. A continuación, os protóns circulan a través dunha serie de aceleradores lineais e sincrotróns (PSB, PS, SPS) antes de ser inxectados no LHC.

Neste colisionador atópanse 4 puntos de interacción, nos que se fan chocar os dous feixes de protóns, con senllos detectores: ATLAS, CMS, ALICE e LHCb. Os tres primeiros teñen unha forma de barril arredor do punto de interacción cunha detección espacial practicamente cilíndrica. ATLAS e CMS son detectores de propósito xeral, e dicir, pescudan nova física a través do estudo da produción de novas partículas. Por outro lado, ALICE está especializado en detectar as partículas provintes de colidir ións pesados no LHC (principalmente colidindo dous feixes de chumbo). Por último, está LHCb [16], que é o experimento para o que se realiza esta tese, foi deseñado orixinalmente para investigar as pequenas diferenzas que existen entre materia e antimateria a través do estudo dos quarks abaixo, aínda que debido o extraordinario comportamento do detector durante os “Runs” I e II do LHC (2010-2018) expandiuse o seu campo de estudo a un detector de propósito xeral. A diferenza dos outros tres detectores do LHC, que teñen forma de barril arredor do punto de colisión dos feixes, o LHCb so se centra nunha pequena rexión seudo cónica na cal o punto de interacción estaría situado no vértice do mesmo e os feixes de partículas viaxan na dirección do eixo de dito cono e en sentidos opostos. Este detector ten unha cobertura angular de 17° .

Todos os detectores de física de partículas situados en colisionadores de partículas teñen certas similitudes respecto á disposición dos distintos sistemas que os conforman. A continuación, fago unha pequena radiografía dende o punto de vista enxeñeril de como se detectan as partículas xeradas polo esmagamento de protóns usando como exemplo o LHCb orixinal. O obxectivo dun detector de partículas situado nun acelerador consiste en identificar a traxectoria, a tipoloxía e a enerxía das mesmas. Para acadar cada un destes fins, deséñanse diferentes sub-sistemas que normalmente teñen unha disposición espacial que sitúa os sistemas con menor materia na rexión próxima o punto de colisións e os que teñen maior material, sitúanse afastados. Isto permite que os primeiros non resten eficiencia ó resto de sistemas. A maiores, soe utilizarse un imán para

modificar a traxectoria das partículas cargadas electricamente, por exemplo o LHCb utiliza un imán dipolar capaz de xerar un campo magnético uniforme de 4 Tm que curva as partículas no eixo vertical.

Profundando en cada un dos diferentes sistemas de detección, en primeiro lugar, comezando dende o punto de colisión, están os detectores de trazas que miden os puntos a través dos que pasan partículas cargadas e reconstrúen a súa traxectoria. Estes detectores deben minimizar a materia, deixando pasar o maior número de subpartículas posible, deben contar cunha alta precisión e granularidade para permitir posicionar e distinguir as trazadas o mellor posible sen sufrir unha errónea reconstrución. As tecnoloxías comunmente utilizadas neste tipo de sistemas son os sensores de silicio (en forma de tiras ou píxels) ou as fibras escintiladoras. O LHCb combina dúas estacións de trazado, unha antes do imán e outra despois, xunto con un detector de vértices (VELO) característico de este experimento. O VELO deseñouse para mellorar o estudo dos quarks “abaixo”, que requiren unha gran precisión a hora de diferenciar os vértices xerados nas colisións protón-protón (vértices primarios) ou xerados a partir de rápidos decaementos das partículas xeradas nos vértices primarios (vértices secundarios). O VELO componse dunha serie de estacións arredor do punto de interacción dispostas perpendicularmente ó feixe. Estas divídense en dúas metades que se poden afastar ou acercar ó punto de interacción cón fin de non danar o detector cando os feixes do LHC non son estables.

Outro dos factores diferenciais do LHCb é o seu sistema de identificación de partículas (pións e kaóns) a través da medición dos aneis de luz Cherenkov¹ nos detectores chamados RICH que reflexan os aneis nunha serie de espellos e convírtenos en sinais eléctricos en fotodetectores. O LHCb conta con dous detectores RICH un entre o VELO e o primeiro detector de trazas e o segundo entre o último detector de trazas e o primeiro calorímetro.

A identificación de electróns e hadróns, así como a medición da súa enerxía faise nos chamados calorímetros. A tecnoloxía utilizada alterna capas metálicas (chumbo ou ferro) con plásticos escintiladores. Nas capas metálicas páranse as partículas xerando unha cascada ou “shower” de luz ultravioleta nos plásticos escintiladores que se captura posteriormente en tubos fotomultiplicadores.

Por último, atópanse os detectores de muóns, que son as partículas cargadas máis pesadas detectadas nos experimentos. Estas son capaces de atravesar todo o experimento sen deterse

¹ A luz Cherenkov prodúcese cando partículas cargadas atravesan un determinado medio (no caso do LHCb é un gas moi denso, C_4F_{10}) no cal viaxan máis rápido que a velocidade da luz. Isto provoca un espectro de luz cónico cuxa proxección se mide para a identificación de partículas.

nos anteriores sistemas. Co obxectivo de parar os muóns sitúanse pranchas de ferro capaces de iniciar unha cascada intercaladas con detectores de gas que capturan a cascada de xeito similar ó calorímetro.

Para poder capturar eventos pouco comúns, requírese un gran número de colisións por segundo, no caso lo LHC a frecuencia de entrecruzamento dos feixes é de 40MHz de pico. Isto provoca que o longo de todos os sistemas descritos anteriormente prodúcese unha cantidade inxente de datos, que non son viables de gardar en disco. Para filtrar esta gran cantidade de datos, utilízase o chamado sistema de disparo ou "trigger". No caso do LHCb orixinal este era un sistema hardware que procesaba os datos en tempo real do detector de vértices, calorímetro e detector de muóns e selecciona os eventos interesantes dende o punto de vista físico e gádaos en disco. Este filtrado dos datos permite reducir o número de datos unhas 3200 veces.

9.2 Mellora do experimento LHCb




Figure 9.1: Sección lonxitudinal do LHCb mellorado.

A pesar do excelente comportamento do experimento LHCb durante os Run I e II do LHC (2010-2018) onde adquiriu 9.2 fb^{-1} de colisións protón-protón, moitas medidas tomadas no detector están limitadas por unha gran incerteza estadística. Co obxectivo de aumentar o rendemento do detector, ó longo de 2019-2020, estase a mellorar o experimento. Os cambios céntranse principalmente en: mellorar o detector de vértices, o sistema de trazado do experimento e por último, o sistema actual de disparo e poder así adquirir así un maior número de datos. Antes de pasar a enumerar as melloras de cada un dos sistemas que forman o experimento, cabe mencionar que se abandona o anterior sistema de disparo e adquisición de datos por un sistema sen disparo hardware, e dicir toda a electrónica do detector envía os datos a frecuencia coa que teñen lugar as colisións (40MHz) e o disparo será, polo tanto, puramente software, o que permite unha maior flexibilidade. Como consecuencia, a electrónica de todos os sub-detectores precisa ser redeseñada. Para facilitar a integración do conxunto do detector, realízanse unha serie de deseños e estándares comúns para o conxunto de subsistemas no eido da transmisión de datos, control dos sub-detectores, sincronía do conxunto do experimento así como a adquisición de datos. A continuación, enuméranse os cambios desta mellora:

Detector de vértices (VELO) As especificacións requiridas para o novo VELO son: resistente a radiación no entorno do LHC (moi superior á actual), alta granularidade que permita a reconstrución das trazas de maneira robusta, minimizar o material na zona máis próxima o punto de interacción para optimizar o rendemento do detector.

O novo deseño do VELO seguirá constando de dúas metades retráctiles arredor do punto de colisión. O número de estacións presentes en cada metade verase incrementado de 21 a 26. Ademais, a distancia con respecto ó punto de interacción redúcese ata os 5.1 mm e só as separa do punto de colisión unha lámina de aluminio de $150 \mu\text{m}$ que separa o baleiro primario do LHC e o secundario do VELO. Cada unha das estacións presentes nunha metade do detector (chamadas "slices") consta dun módulo na rexión máis próxima ó punto de interacción con 4 sensores de silicio "n-on-p" (dous por cara) distribuídos en forma de L arredor do feixe, estes sensores están unidos mediante a tecnoloxía de "bump-bonding" a uns ASICs deseñados a medida chamados VeloPix, baseados na familia TimePix. O control dos VeloPixes lévase a cabo dende un link de alta velocidade directamente a través dun ASIC propietario do CERN chamado GBTx, situado no módulo tamén. Para refrixerar o módulo, recorreuse a un novidoso sistema que consiste en

fabricar uns capilares entre as dúas obleas de silicio que serven como sustento do módulo. Ditos capilares teñen unha maior sección xusto antes de pasar baixo as zonas de emisión de calor, isto xunto con uns quentadores a entrada do módulo que levan o refrixerante (CO_2) líquido ó borde do estado gasoso, unha vez este gas pasa a través do ensanche no silicio vólvese gasoso. Mantendo o refrixerante no punto de transición líquido/gas a temperatura de refrixeración, mantense constante en todo o módulo (en torno a -20°C á presión de traballo).

Os módulos comunícanse co exterior a través de cables flexibles que permiten o movemento telescópico do mesmo, unha PCB especificamente deseñada para soportar a diferenza de presións entre o baleiro secundario do LHCb e o aire e por último outra PCB que convirte todas as sinais a fibra óptica e as envía a superficie.

Detectores de luz cherenkov (RICH) Co fin de soportar o incremento na luminosidade do LHCb, as dúas estacións de detectores RICH mellorarán os seus sistemas ópticos. En particular, a primeira estación RICH vai mellorar a distancia focal dos seus espellos, incrementándoa dende 2.7 m ata 3.7 m e reducindo así a ocupación dos foto-detectores. O sistema de refrixeración tamén será mellorado co fin de permitir un acceso máis sinxelo para o mantemento así como un funcionamento máis estable.

Sistema de trazado A mellora do experimento manterá dúas estacións de trazado, UT e SciFi, pero con un deseño completamente novo.

O sub-detector UT estará situado entre a primeira estación de RICH e o imán dipolar. O motivo polo que se precisa mellorar este detector é principalmente aumentar a resolución no momento das partículas co incremento da luminosidade así como soportar a radiación no novo entorno. UT estará composto por catro estacións baseadas en tiras de silicio. A tecnoloxía dos sensores, así como as súas dimensións dependen da localización: na parte máis próxima ó feixe (con maior radiación) utilízase tecnoloxía “n-on-p” con tiras curtas e na parte externa utilízanse tiras máis separadas con tecnoloxía “p-on-n” xa que a radiación é menor.

O SciFi situarase despois do imán, consta de 12 capas que permiten a reconstrución das trazas e a medida do momento das partículas. Cada capa está composta por fibras escintiladoras de plástico de 2.4 m de longo e $250\ \mu\text{m}$ de diámetro, dispostas en posición

vertical e con unha profundidade de 6 fibras por capa, o que da un grosor total de 1.5 mm. A lectura destas fibras faise dende os extremos con fotomultiplicadores de silicio (SiPM).

Calorímetros Nos calorímetros os cambios serán menores. Reducirase a ganancia dos tubos foto-multiplicadores para manter a mesma corrente no ánodo dos mesmos e reducir a degradación.

Detector de muóns Non sufrirá ningún cambio máis alá da electrónica necesaria para ler o detector a 40 MHz.

Sistema de disparo e “Online” Pese a non ser un sistema instalado na caverna do experimento, cómpre describir o sistema chamado “online” que se encarga do desenvolvemento da electrónica de control e adquisición de datos común a todo o experimento e situada fisicamente na superficie do experimento.

A figura 9.2 mostra o sistema de control e adquisición de datos do experimento. Nesta arquitectura toda a electrónica dos diferentes sub-detectores expostos previamente (Front-End (FE)) transmiten os datos adquiridos a 40 MHz a través de fibra óptica a superficie do experimento onde se atopa a “granxa” de ordenadores dende a que se controla e adquiren os datos do experimento. Nesta granxa, os datos recibidos pola fibra óptica son procesados en tempo real por tarxetas dedicadas, baseadas en Field Programmable Gate Array (FPGA) nas que cada un dos sub-detectores realiza un pre-procesado dos datos. Estas tarxetas comunícanse cun ordenador a través dun porto PCIe a 100 Gb/s. A partir deste punto, os datos seguen un formato estándar para todo o experimento. Todos estes ordenadores que reciben datos dos diferentes ordenadores están interconectados a través dunha rede de 100 Gb/s. Neste punto faise un primeiro procesado dos datos, decidindo cales se gardan en disco para ser analizados máis tarde.

9.3 Mellora do detector de vértices do LHCb

O traballo aquí exposto divídese en dous grandes bloques. O primeiro consiste en verificar o correcto funcionamento dos diferentes prototipos hardware deseñados para o experimento, atopar erros de deseño e configurar adecuadamente os distintos ASICs do sub-detector. O segundo bloque consiste en deseñar e implementar o firmware de control, distribución das sinais de con-




Figure 9.2: Diagrama do sistema de adquisición de datos do LHCb.

trol rápido e adquisición de datos así como probar o seu funcionamento en diferentes versións das FPGAs utilizadas.

9.3.1 Electrónica na caverna

Dende o punto de vista da electrónica, nesta sección descríbese unha das unidades mínimas anteriormente descritas “Slice” (Figura 9.3). O deseño da electrónica foi realizado por outros compañeiros da colaboración do VELO. A miña participación referente á electrónica limitouse principalmente a validar e poñer en funcionamento o hardware.

Na figura 9.3, pódese observar o módulo que se situará a 5.1 mm do punto de interaccións do LHCb. Tomando como referencia o punto de interacción, atopámonos en primeiro lugar os

chamados “tiles” que están formados por un sensor de Hamamatsu “n-on-p” optimizados para o entorno do LHC unidos a tres VeloPix mediante a tecnoloxía de empaquetado 3D de chips chamada “flip chip”. A comunicación dos tiles co exterior realízase mediante a soldadura por “wire bonding” a unha PCB auxiliar chamada híbrido. Ditos híbridos transfiren aos VeloPix as sinais de control procedentes dun tercer híbrido que sustenta o ASIC encargado da xestión do link de control de alta velocidade (GBTx). En paralelo os VeloPix dun “slice” envían os datos adquiridos a través de vinte links a unha velocidade de transmisión de ata 5.13 Gb/s.

Todas as sinais de alta velocidade que comunican o módulo co exterior (control e adquisición de datos) circulan a través de: uns PCBs flexibles de aproximadamente 50cm de longo que absorben o desprazamento telescópico do módulo, unha PCB que absorbe a diferenza de presión entre o baleiro no que se atopa o módulo e o ambiente (VFB) e finalmente unha última PCB que as transforma de sinais eléctricas a ópticas (OPB) e as envía a electrónica de control e adquisición de datos da superficie do experimento.

A posta en funcionamento das diferentes partes que forman unha “slice” require multitude de pasos:

1. Unha vez recibidos os primeiros prototipos, fíxose unha rápida inspección visual das diferentes PCBs, así como, unha procura de de curto-circuitos ou camiños non conectados coa axuda dun multímetro. A continuación, inspeccionouse a integridade das sinais a través das liñas de transmisión de alta velocidade coa axuda dun Vector Network Analyzer




Figure 9.3: (Esquerda) Sección transversal de unha “slice” do VELO. (Dereita) Zoom do módulo.

(VNA) e un Time-Domain Reflectometer (TDR).

2. Unha vez ensamblado o prototipo, comezase a alimentalo pola OPB. Esta PCB conta con un GBTx a través do que se controlan e monitorizan todos os compoñentes activos que a conforman. É preciso optimizar dende a OPB os links de control que se transmiten ata o módulo, evitando así erros de transmisión.
3. Para comezar a operar o módulo, e preciso configurar os GBTx que envían as sinais de control e sincronía aos distintos VeloPixes.
4. Neste momento, seremos capaces de acceder a cada VeloPix a través do link de control para ecualizalo, adquirir datos e comprobar que funciona correctamente.
5. Antes de tomar datos a alta taxa cómpre optimizar o link de transmisión. Isto conseguiuuse pola combinación de tres factores: o primeiro e máis importante consiste en axustar a fase entre do reloxo co que se xeran os datos no link de transmisión, en segundo lugar implementouse un filtro paso alto (CTLE) para eliminar ruído exterior no link e por último amplificando o sinal antes de convertelo a fibra óptica. Esta mellora no link permite obter un Bit Error Rate (BER), medido despois de transmitir as sinais a través de 400 m de fibra óptica, menor que 10^{-12} .

9.3.2 Electrónica na superficie do experimento

Na superficie do LHCb sitúase un centro de datos dende o que se operará (control e adquisición de datos) o detector. A figura 9.4 mostra un esquema simplificado do sistema de control e sincronización do LHCb. Unha labor vital para obter os datos de forma correcta é manter todo o detector síncrono en todo momento. Para elo deseñouse o sistema xerárquico da figura 9.4, no que o único punto de comunicación co LHC é a tarxeta supervisora (S-ODIN) que toma o reloxo de 40 MHz do LHC e distribúeo a todo o detector. Nesta tarxeta tamén se mantén control do esquema de recheo dos feixes que colisionan no LHCb, en consecuencia a tarxeta supervisora aplica e distribúe as sinais de disparo para os eventos nos que hai colisión e que interesa procesar (sistema de control rápido, TFC). Tanto os reloxos coma estas sinais de TFC son transmitidas á electrónica da caverna mediante as tarxetas de interface (SOL40) e control así como as tarxetas de adquisición de datos (TELL40) que procesan os datos e crean os eventos en tempo real. A configuración, control e monitorización do experimento realízanse a través de




Figure 9.4: Arquitectura do sistema de control e sincronización do LHCb.

unha serie de Supervisory Control And Data Acquisition (SCADA) e ordenadores de control conectados mediante rexistros as tarxetas SOL40.

O feito de que o VELO ten os sensores situados o máis cerca da zona de interacción de todo o LHC, sumado a necesidade de reducir o material na zona próxima as colisións, manter os sensores a -20°C e a velocidade de lectura a 40MHz forzan ó deseño dun subdetector con características moi diferentes do resto cuxas singularidades son: envío de datos desordenados temporalmente dende o VeloPix, uso dun serializador e protocolo propietarios nos links de adquisición de datos do VeloPix, control e sincronización dos VeloPix directamente dende un link de alta velocidade (GBTx).

Dende o punto de vista da electrónica da superficie, foi preciso en primeiro lugar facer pequenas adaptacións no código da SOL40 para que este aceptase o protocolo de control do

VeloPix a través do GBTx. Ademais, a TELL40 precisa un firmware practicamente novo para permitir a lectura do protocolo propietario do VeloPix e para ordenar os datos temporalmente.

9.3.2.1 Firmware de control do VELO

As sinais de control que se deben proporcionar o VeloPix son de dúas naturezas posibles: sinais de configuración e monitorización (ECS) e sinais de sincronía (TFC). Pese a que ámbalas dúas comparten o mesmo link, a primeira transmite os datos bidireccionalmente a 80Mb/s e a segunda transmite datos de forma unidireccional (cara o VeloPix) a 320Mb/s. O bus entre a SOL40 e o GBTx é de 80 bits con unha frecuencia de 40MHz, polo que as sinais de control e sincronía utilizan 2 bits (80Mb/s) ou 8 bits (320Mb/s) respectivamente.

9.3.2.2 Firmware de adquisición de datos do VELO

O firmware de adquisición de datos comparte co firmware común de LHCb unicamente o interface PCIexpress e o bloque identificador de eventos.

Cada TELL40 recibe os datos unicamente dun módulo. A primeira etapa firmware prodúcese no chamado interface de baixo nivel (LLI) que se encarga de recuperar os datos da fibra óptica, descodificalos e identificar posibles erros de transmisión² e transferir os datos ó dominio do procesado de datos da FPGA.

A seguinte etapa é o procesado de datos que ten dúas labores: a tarefa fundamental é ordenar temporalmente os datos que chegan dos VeloPix, a segunda opcional (só se implementará si a tarefa fundamental deixa espazo libre suficiente na FPGA) consiste en liberar recursos dos previamente asignados as granxas de ordenadores.

A labor fundamental do procesado de datos lévase a cabo no chamado “Router” que é un algoritmo capaz de ordenar e mesturar os datos da metade dun módulo en tempo real a unha tasa de ata 50Gb/s. O algoritmo do Router está dividido en dous grandes bloques. O primeiro utiliza unha rede de seudo arbores de decisión entrelazados, nos que cada un dos nodos están compostos por memorias FIFO que permiten a absorción e ordenamento dos datos a gran velocidade segundo os catro bits máis significativos (dos nove bits de marca temporal asignados no VeloPix, BXID). O segundo bloque, encargado de ordenar os cinco bits menos significativos cambia a filosofía de ordenamento, para reducir o uso de memorias, baseándose en táboas de

²O protocolo de transmisión de datos do VeloPix non permite a recuperación de erros.

procura mediante as que se asigna a dirección correspondente en unha das dezaseis memorias RAM nas que se almacenan temporalmente os datos xa ordenados. Desafortunadamente, o Router por si mesmo non é capaz de facer o procesado de datos completo, polo que son precisas dúas etapas de acondicionamento e procesado auxiliares ó Router chamadas Pre-Router e Post-Router. O Pre-Router, como o seu nome indica sitúase augas arriba do Router, é o encargado de: marcar os datos respecto ó VeloPix de orixe, identificar os comandos de inicio da adquisición de datos e alertar ó Router e Post-Router en consecuencia, así como actuar coma ponte entre o dominio de reloxo da FPGA (40MHz) e o procesado de datos do VELO (160MHz). O Post-Router encárgase da xestión do procesado de datos. É dicir, lee as memorias que conteñen os datos ordenados polo Router de acordo a os datos de disparo do LHCb (só se propagan os datos declarados como válidos), aplica o formato de Intel Avalon-ST, procesa as sinais de sincronía específicas do VELO na TELL40 e completa o marcado temporal dos datos que non aplica o VeloPix³.

Opcionalmente e co fin de optimizar o posterior procesado dos datos en ordenadores, os datos serán marcados como impactos illados así como a identificación de datos que impactaron en zonas adxacentes.

Segundo as últimas estimacións os recursos necesarios para a implementación do firmware básico necesario son do orde dun 41% dos recursos lóxicos da FPGA e un 72% de bloques de memoria. Con estes resultados a vista, queda un pequeno espazo onde implementar as funcionalidades extra mencionadas anteriormente, o bloque de identificación de impactos adxacentes estímase que incrementa os recursos ata un 73% de recursos lóxicos e un 88% de bloques de memoria. Por último, o bloque de identificación de impactos illados na súa versión orixinal consumía gran cantidade de recursos, pero un cambio conceptual do deseño que pasará a basearse en identificación por código Gray permitirá un limitado impacto en recursos deste bloque.

9.3.3 Firmware para controis de calidade e desenvolvemento do hardware

³O BXID preciso para marcar os datos xerados no LHC é de 12 bits. O VeloPix só transfere 9 bits para aforrar ancho de banda, polo tanto a TELL40 ten que reconstruír os 3 bits restantes.

9.3.3.1 Firmware de bypass

O desenvolvemento do firmware final é unha labor longa e complexa, polo que para probar os deseños hardware e para verificar a produción deseñouse unha versión moito máis sinxela do firmware que permite enviar os datos dende o LLI directamente a o interface PCIexpress para ser gardados en disco. Este firmware conta con unha memoria FIFO por cada canle de entrada e un sistema de xestión que as valeira dándolle prioridade en todo momento a memoria mais chea. Este firmware foi utilizado tamén en test con feixe de protóns no SPS do CERN, para elo introduciuse un sistema de sincronización con un telescopio usado como referencia.

9.3.3.2 Banco de probas dos cables de alta velocidade e alta tensión

Unha das labores do Instituto Galego de Física de Altas Enerxías (IGFAE) no eido da mellora do LHCb consiste no montaxe dos cables flexibles que transfiren a alta tensión e as sinais de alta velocidade entre o VFB e o módulo . Para verificar que ditos cables cumpran coas especificacións deseñáronse senllos bancos de probas baseados na FPGA de Xilinx Virtex 7.

APPENDIX A

OLD TELL40 OUTPUT DATA FORMAT

During the development of the TELL40, several data processing output formats were studied before agreeing the actual data format. This appendix explains the 10 GbEth data format designed for the MiniDAQ1 system, but has never been used, as the design was abandoned in favour of MiniDAQ2. At the same time, the old data format used in the first versions of the readout firmware is shown, but it was abandoned as well as it was not taking advantage of all the potential bandwidth.

A.1 Old data formats

Two versions of the Post-Router output format were studied and implemented. In these versions the Avalon-ST interface only propagates downstream a 256 bit data word for PCIe or 64 bit for the 10 GbEth interface of MiniDAQ1¹ along with SOF, EOF and DV. No separate TFC and event header were propagated, this information was included in the common Event ID block.

Table A.1 shows the first version of the data format, where the 64 MSB of the first output word of an event contain a global header with the TFC information followed by a header that identifies the BXID and the event length. The global header field was left empty in VELO Post-Router and it was filled by the Event ID block. VELO data was also slightly different, as it maintained a similar bit structure to the VeloPix SPP data, just adding the ICF flag and removing the BXID of the sorted data.

¹The 64 bit word is just the 256 bit chopped.




Table A.1: (Top) VELO TELL40 output format (Bottom) VELO Specific Data Output.

There was also a second version of the output format that sent the 64 bits of the global header on a separate bus in order to increase the bandwidth by 2 SPPs in the first Avalon-ST transaction.

Bibliography

- [1] J. Buytaert, “The HEV Ventilator Proposal,” Tech. Rep. EP-Tech-Note-2020-001, CERN, Geneva, Apr 2020. Full author list will be added at subsequent submission.
- [2] L. Evans and P. Bryant, “LHC machine,” *Journal of Instrumentation* 3, no. 8, p. S08001, 2008.
- [3] T. A. Collaboration, “The ATLAS experiment at the CERN Large Hadron Collider,” *Journal of Instrumentation* 3, no. 8, p. S08003, 2008.
- [4] T. C. Collaboration, “The CMS experiment at the CERN LHC,” *Journal of Instrumentation* 3, no. 8, p. S08004, 2008.
- [5] T. A. Collaboration, “The ALICE experiment at the CERN LHC,” *Journal of Instrumentation* 3, no. 8, p. S08002, 2008.
- [6] T. LHCb Collaboration, “The LHCb detector at the LHC,” *Journal of Instrumentation* 3, no. 8, p. S08005, 2008.
- [7] R. Bailey and P. Collier, “Standard Filling Schemes for Various LHC Operation Modes,” Tech. Rep. LHC-PROJECT-NOTE-323, CERN, Geneva, Sep 2003.
- [8] L. Collaboration, “LHCb trigger system: Technical design report,” *Technical Design Report LHCb*, CERN, 2003.
- [9] L. Collaboration, “LHCb VELO (VERTex LOcator): Technical design report,” *Technical Design Report LHCb*, CERN, 2003.
- [10] L. Collaboration, “Layout and expected performance of the LHCb TT Station: Technical design report,” *LHCb Note*, CERN, 2003.

- [11] L. Collaboration, “LHCb inner tracker: Technical design report,” *Technical Design Report LHCb, CERN*, 2002.
- [12] L. Collaboration, “LHCb outer tracker: Technical design report,” *Technical Design Report LHCb, CERN*, 2001.
- [13] L. Collaboration, “LHCb RICH: Technical design report,” *Technical Design Report LHCb, CERN*, 2000.
- [14] L. Collaboration, “LHCb calorimeters: Technical design report,” *Technical Design Report LHCb, CERN*, 2000.
- [15] L. Collaboration, “LHCb muon system: Technical design report,” *Technical Design Report LHCb, CERN*, 2001.
- [16] T. L. collaboration, *Framework TDR for the LHCb upgrade*, 2012.
- [17] L. A. et al., “The versatile link, a common project for super-lhc,” *Journal of Instrumentation*, 2009.
- [18] P. Moreira, R. Ballabriga, S. Baron, S. Bonacini, O. Cobanoglu, F. Faccio, T. Fedorov, R. Francisco, P. Gui, P. Hartin, K. Kloukinas, X. Llopart, A. Marchioro, C. Paillard, N. Pinilla, K. Wyllie, and B. Yu, “The GBT Project,” 2009.
- [19] *Development of a low power 5.12 Gbps data serializer and wireline transmitter circuit for the VeloPix chip*, 2015.
- [20] T. L. Collaboration, *LHCb VELO Upgrade Technical Design Report*, 2013.
- [21] T. L. collaboration, *LHCb Tracker Upgrade Technical Design Report*, 2014.
- [22] M. Baszczyk, P. Carniti, L. Cassina, A. C. Ramusino, P. Dorosz, M. Fiorini, C. Gotti, W. Kucewicz, R. Malaguti, and G. Pessina, “CLARO: an ASIC for high rate single photon counting with multi-anode photomultipliers,” *Journal of Instrumentation*, vol. 12, pp. P08019–P08019, aug 2017.
- [23] T. L. Collaboration, *LHCb Tracker Upgrade Technical Design Report*, 2014.

- [24] E. Buchanan, “The lhcb vertex locator (velo) pixel detector upgrade,” *Journal of Instrumentation*, vol. 12, pp. C01013–C01013, 01 2017.
- [25] J. H. Lau, *Flip Chip Technologies*. McGraw-Hill, 1995.
- [26] T. Poikela, R. Ballabriga, J. Buytaert, X. Llopart, W. Wong, M. Campbell, K. Wyllie, M. van Beuzekom, J. Schipper, S. Miryala, and V. Gromov, “The VeloPix ASIC,” *Journal of Instrumentation*, vol. 12, pp. C01070–C01070, jan 2017.
- [27] V. Gromov, V. Zivkovic, M. van Beuzekom, X. Llopart, T. Poikela, J. Buytaert, M. D. Gaspari, M. Campbell, and K. Wyllie, “Development of a low power 5.12 gbps data serializer and wireline transmitter circuit for the VeloPix chip,” *Journal of Instrumentation*, vol. 10, pp. C01054–C01054, jan 2015.
- [28] E. L. et al., “Production, measurement and simulation of a low mass flex cable for multi gigabit/s readout for the LHCb VELO upgrade,” 2012. pp.C01018.
- [29] P. M. et al, “The gbt project,”
- [30] DuPontTM, *Pyralux®AP-Plus all-polyimide thick copper-clad laminates*, 2009.
- [31] B. A. et al., “Custom dc/dc converters for distributing power in slhc trackers,” *Journal of Instrumentation*, 2009.
- [32] A. C. et al, “The gbt-sca, a radiation tolerant asic for detector control and monitoring applications in hep experiments,” *Journal of Instrumentation*, 2015.
- [33] O. de Aguiar Francisco et al, “Evaporative co₂ microchannel cooling for the lhcb velo pixel upgrade,” *Journal of Instrumentation*, 2015.
- [34] L. Brarda, B. Jost, D. Lacarrere, R. Lindner, N. Neufeld, L. Roy, and E. Thomas, “A new data-centre for the LHCb experiment,” *J. Phys. Conf. Ser.*, vol. 396, p. 012009, 2012.
- [35] J. M. et al, “Trigger and timing distributions using the ttc-pon and gbt bridge connection in alice for the lhcb run 3 upgrade,” 2018.
- [36] M. B. et al, “A pcie gen3 based readout for the lhcb upgrade,” *Journal of Physics: Conference Series*, 2014. Conf. Ser.513 012023.

- [37] P. Durante, N. Neufeld, R. Schwemmer, G. Balbi, and U. Marconi, “100gbps pci-express readout for the lhcb upgrade,” *Journal of Instrumentation*, vol. 10, pp. C04018–C04018, 04 2015.
- [38] S. Labs, *Si5345A*, 2019.
- [39] Intel, *Avalon® Interface Specifications*, 2019.
- [40] “IEEE standard for information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications,” standard, International Organization for Standardization, Geneva, CH, Dec. 2005.
- [41] M. T. Jones, *BSD Sockets Programming From a Multi-Language Perspective*. USA: Charles River Media, Inc., 1st ed., 2003.
- [42] “Streaming DMA accelerator functional unit (AFU) user guide,” standard, Intel, Dec. 2018.
- [43] Intel, *V-Series Transceiver PHY IP Core UserGuide*, 2019.
- [44] Intel, *Intel® Arria® 10 Transceiver PHY User Guide*, 2019.
- [45] K. Akiba, M. van Beuzekom, H. Boterenbrood, E. Buchanan, J. Buytaert, W. Byczynski, X. C. Vidal, P. Collins, E. Dall’Occo, A. D. Suárez, R. Dumps, T. Evans, V. F. Lima, A. G. Torreira, J. G. Pardiñas, B. van der Heijden, C. Hombach, M. John, S. Kulis, X. L. Cudie, F. Marinho, E. Price, S. Richards, P. R. Perez, D. Saunders, Å. S. Folkestad, H. Schindler, F. Schreuder, H. Snoek, P. Tsopelas, J. Velthuis, M. V. Diaz, and M. Williams, “LHCb VELO timepix3 telescope,” *Journal of Instrumentation*, vol. 14, pp. P05026–P05026, may 2019.
- [46] “ISO/IEC 13239 information technology — telecommunications and information exchange between systems — High-Level Data Link control (HDLC) procedures. HDLC standard,” 2002.
- [47] J. Standard, “Scalable low-voltage signaling slvs-400,” *JESD8-13*, October 2001.

- [48] E. ETM Professional Control GmbH, “Simatic wincc open architecture,” *Eisenstadt, Austria*.
- [49] “Jcop framework configuration db,” *website*:
<https://wikis.web.cern.ch/wikis/display/EN/JCOP+Framework+Configuration+Database>.
- [50] C. Gaspar, “Dim - a distributed information management system for the delphi experiment at cern,” *IEEE Real Time Conference*, 1993.
- [51] I. DASC, *IEEE standard Verilog hardware description language*, 2001.
- [52] I. DASC, *IEEE Standard VHDL Language Reference Manual*, 1993.
- [53] J. E. Stone, D. Gohara, and G. Shi, “Opencl: A parallel programming standard for heterogeneous computing systems,” *Computing in Science Engineering*, vol. 12, no. 3, pp. 66–73, 2010.
- [54] Xilinx, *VC707 Evaluation Board for the Virtex-7 FPGA*. Xilinx.
- [55] A. Kchaou, W. El Hadj Youssef, and R. Tourki, “Performance evaluation of multicore leon3 processor,” in *2015 World Symposium on Computer Networks and Information Security (WSCNIS)*, pp. 1–4, 2015.
- [56] C. A. M. E. A. N. F. . I. Q. Dell, Cisco, “Ethernet jumbo frames,” *Ethernet Alliance*, 2009.
- [57] I. E. T. F. (IETF), *Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry*, 2011.
- [58] K. W. P. Moreira, J. Christiansen, *GBTX MANUAL*, 2018.

Acronyms

CERN	European Organization for Nuclear Research. ix, 1, 3, 4, 25, 47, 81, 82, 91, 125, 127, 138, 141, 145, 154
LHC	Large Hadron Collider. xi, 1–7, 9, 11–19, 21–23, 25–27, 29, 30, 33, 34, 37, 38, 40, 47, 50, 53, 62, 66, 79, 85, 86, 88, 115, 125, 127–131, 134, 135, 138, 142–144, 149–151, 153, 171, 173
LHCb	Large Hadron Collider beauty. ix, xi–xiv, 2–5, 7–17, 19, 30, 34, 37–41, 43, 48, 49, 56, 59, 66, 70, 75, 80, 84, 86–88, 91, 101–103, 115, 119, 125, 127–133, 135–137, 139, 141–144, 146–148, 150–154, 171–175
RICH	Ring-Imaging Cherenkov Detectors. 8, 10, 16, 18, 19, 129, 131, 132, 143, 146, 171
VELO	VERtex LOcator. ix, xi–xiv, 1, 2, 8–10, 13, 16, 21–23, 25–27, 29–35, 37, 39–41, 43–53, 55–57, 59–63, 65–73, 75–87, 89, 91, 95, 96, 100–105, 107, 115–117, 119, 121, 125, 127–131, 133–135, 137, 138, 141, 143, 145, 148, 151–153, 155, 156, 171–174, 177
ASIC	Application Specific Integrated Circuit. xi, 2, 17, 18, 21–26, 29, 31, 33, 34, 43, 44, 47, 48, 51, 55, 60, 68, 69, 73, 77, 78, 80–87, 91–93, 97, 99, 100, 102, 107, 109–111, 113, 114, 122, 131, 133, 135, 145, 147, 149, 155, 156, 172–174, 177
BAR	Base Address Register. xii, xiii, 65, 70–73, 83–86, 95
BE	Back-End. 13, 24, 26, 27, 29, 30, 38–40, 47–51, 53, 56, 68, 69, 72, 83, 87, 100, 110, 115, 125, 130
BER	Bit Error Rate. 98, 99, 102, 109, 113–115, 135, 150, 174

BXID	Bunch Cross Identifier. 27, 29, 40, 42–44, 48–51, 54–56, 58–62, 66–68, 77–80, 86, 87, 135, 137, 138, 152, 153, 156, 173, 177
CDR	Clock Data Recovery. 45, 46, 48, 73
CTLE	Continuous Time Linear Equalizer. 45, 47, 106, 108, 109, 114, 115, 135, 150, 174
DAQ	Data AcQquisition. xiii, 11, 13, 30, 37, 43, 44, 62, 63, 76–78, 83, 100, 103, 125, 132–134, 172
DIM	Distributed Information Management System. 88, 89
DV	Data Valid. 51, 59, 156
ECS	Experiment Control System. xiii, 11, 13, 15, 26, 29, 30, 38, 50, 64, 82, 83, 85, 87, 100, 103, 105, 109–111, 135, 137, 152, 158, 173, 174
EEPROM	Electrically Erasable Programmable Read-Only Memory. 95, 105
EOF	End Of Frame. 59, 67, 156
FE	Front-End. xi, xiii, 11, 14–17, 21–24, 26, 28, 29, 33, 38–44, 48–51, 53, 60, 61, 64, 66, 68, 69, 73, 77, 78, 80, 82, 84–88, 96, 100, 102, 103, 109–111, 120, 130, 135, 147, 155–157, 173, 177
FIFO	First In First Out. 48, 51, 54, 57, 63, 68, 73, 76, 77, 79, 83–86, 137, 138, 152, 154, 173
FPGA	Field Programmable Gate Array. xiii, 15, 30, 41, 44, 45, 47, 48, 53, 55, 61, 62, 70–72, 77, 78, 80, 83, 84, 88, 89, 91–100, 102, 103, 105, 132, 133, 137–139, 146–148, 152–154, 157, 173
FSM	Finite State Machine. 40, 51, 56, 57, 79
GBLD	Giga Bit Laser Driver. 102, 107, 109
GBT	GigaBit Transceiver. 13, 25, 28, 42, 43, 47, 70, 72, 81, 82, 84, 88, 89, 102, 104, 107, 177
GBT-SCA	GBT Slow Control Adapter. 33, 81, 82, 88, 89, 102, 107, 134

GBTx	Giga Bit Transceiver ASIC. 31–33, 81, 82, 84, 86, 102, 107, 109, 110, 131, 134–137, 145, 149–152, 177
GWT	Gigabit Wireline Transmitter. 13, 25–28, 30, 43, 44, 47, 48, 51, 53, 67, 70–73, 75, 77, 78, 83, 85, 100, 102, 104, 105, 113, 114, 135, 137, 172, 174
HDL	Hardware Description Language. 92, 93, 95, 96
I/O	Input/Output. 97, 99
ICF	Isolated Cluster Flagging. xii, 59–62, 72–75, 155, 156, 173
IGFAE	Instituto Galego de Física de Altas Enerxías. x, 2, 119, 122, 123, 127, 139, 141, 154
IP	Intellectual Property core. 93, 96, 98, 99
JCOP	Joint Controls Project. 88, 115
JTAG	Joint Test Action Group. 95, 98
LLI	Low Level Interface. xii, 42, 44, 51, 53, 67, 70–73, 76, 78, 102, 103, 105, 114, 137, 138, 152, 154
LS2	Long Shutdown 2. 2, 11, 125
LSB	Less Significant Bit. 26, 53, 54, 69, 70, 78, 80, 84, 173
LTD	Lock To Data. 45, 46
LTR	Lock To Reference. 45, 46
MEP	Multi Event Package. xii, 15, 49, 63–65, 70, 173
MSB	Most Significant Bit. 26, 53, 54, 60, 69, 70, 78, 80, 137, 156, 173
ODE	Off-Detector Electronics. xii, 37, 39, 41, 43, 45, 47, 49, 51, 53, 55, 57, 59, 61, 63, 65, 67, 69, 71, 73, 75, 77, 79, 81, 83, 85, 87, 89
OPB	Opto-Power Board. 30, 33, 68, 80–82, 106, 107, 113, 134, 149, 150, 174, 177

PCB	Printed Circuit Board. 31–33, 41, 95, 98, 100, 105, 110, 120, 122, 123, 131, 134, 146, 149, 150
PCS	Physical Coding Sublayer. xii, 45–48, 71, 72, 102, 103
PHY	PHYsical layer. xii, 44, 45, 48
PLL	Phase Locked Loop. 25, 41, 45, 47, 66, 72, 93, 97, 105
PMA	Physical Medium Attachment. xii, 45–47, 71, 98, 102, 172
PMT	Photo Multiplier Tubes. 18
PON	Passive Optical Network. 38, 41
PRBS	PseudoRandom Binary Sequence. 48, 71, 106, 113, 114
PS	Proton Synchrotron. 3, 128, 142
PSB	Proton Synchrotron Booster. 3, 128, 142
QA	Quality Assurance. xiv, 2, 76, 98, 115, 121, 122, 125, 138
RAM	Random Access Memory. 48–50, 54–58, 71, 78, 79, 97, 137, 138, 153
SCADA	Supervisory Control And Data Acquisition. 15, 41, 66, 77, 87, 88, 135, 151
SciFi	Scintillating Fibre Tracker. 16, 17, 132, 146
SDC	Synopsis Design Constraints. 95, 97
SFP	Small Form-factor Pluggable. 103, 105
SLVS-400	Scalable Low-Voltage Signaling. 81, 82
SM	Standard Model. 1, 11, 127, 141
SOF	Start Of Frame. 58, 59, 67, 156
SP	Super Pixel. 27, 59–62
SPD	Scintillating Pad Detector. 10, 11, 19
SPI	Serial Peripheral Interface. 81, 82, 88
SPIDR	Speedy PIXel Detector Readout. xiii, 99–101, 103, 111, 120, 122, 173
SPP	Super Pixel Packet. 24–30, 47, 50, 51, 58, 59, 61, 67, 70, 75, 79, 80, 156, 173
SPS	Super Proton Synchrotron. 3, 115, 125, 128, 138, 142, 154
Tcl	Tool Command Language. 93, 97, 99, 121

TDR	Time-Domain Reflectometer. 106, 134, 150
TFC	Timing and Fast Control. xi–xiii, 11, 13, 14, 27–30, 37–42, 44, 48–50, 55–60, 62, 63, 66–69, 73, 76–78, 80, 82, 85–87, 105, 110, 135, 137, 138, 150, 152, 156, 172, 173, 177
UT	Upstream Tracker. 16, 132, 146
VFB	Vacuum Feedthrough Board. 30, 33, 106, 134, 139, 149, 154, 174
VGA	Variable Gain Amplifier. 45, 47
VHDL	Very High Speed Integrated Circuit Hardware Description Language. 92, 93, 96
VNA	Vector Network Analyzer. 106, 134, 150
VTRx	Versatile Transceiver. 33, 102, 107, 109
VTTx	Versatile Twin Transmitter. 33, 68, 102, 107
XDC	Xilinx Design Constraints. 95, 99

List of Figures

Fig. 2.1	LHC complex detector.	4
Fig. 2.2	(Top left) ATLAS Experiment, (Top right) CMS experiment, (Bottom left) ALICE experiment, (bottom right) LHCb experiment.	5
Fig. 2.3	25 ns LHC Filling scheme.	6
Fig. 2.4	LHCb detector.	8
Fig. 2.5	(Left) Angular distribution of b and $\bar{b}$. (Right) Illustration of the acceptance complementarity as a function of pseudo-rapidity between experiments at the LHC.	9
Fig. 2.6	LHCb upgrade detector.	12
Fig. 2.7	LHCb upgrade Readout architecture.	13
Fig. 2.8	LHCb upgrade trigger diagram.	14
Fig. 2.9	CAD model of the LHCb VELO upgrade.	16
Fig. 2.10	(Left) Layout of the detection layers of the UT. Each rectangle represents a silicon sensor and diverse shadings indicate different sensor geometries. (Right) Sketch of a UT stave.	17
Fig. 2.11	Schematic yz- and xy-view of one SciFi tracking station. It is composed out of 4 layers, two vertical (x) and two with a stereo angle of $\pm 5^\circ$ (u, v). Each layer is made of 10 or 12 individual fibre modules.	18
Fig. 2.12	RICH1 has a vertical optical layout (y -axis), while RICH2 has a horizontal optical symmetry (x -axis). The two drawings are not to scale.	19
Fig. 3.1	Close-up view of the VELO sub-detector.	21
Fig. 3.2	VeloPix data rate [Gb/s].	22

Fig. 3.3	(Left) Hybrid pixel detector. (Right) VELO “Tile”: Sensor flip-chip to 3 VeloPix ASICs.	23
Fig. 3.4	General architecture of VeloPix.	24
Fig. 3.5	Schematic of the pixel and super pixel architecture.	25
Fig. 3.6	GWT data frame.	26
Fig. 3.7	SuperPixel Data format.	27
Fig. 3.8	Hitmap pixel ordering.	27
Fig. 3.9	Idle frame.	28
Fig. 3.10	VELO slice.	30
Fig. 3.11	LHCb VELO electronics overview.	31
Fig. 3.12	VELO module.	32
Fig. 3.13	(Left) Cross section scheme for the microchannel cooling with a superposition of the VeloPix ASICs. (Right) Schematic of the microchannel attachment of the fluidic connector.	33
Fig. 3.14	(Left) Draw of the edge region and clearance to the module. (Right) Picture of the LHCb VELO RF foil.	34
Fig. 3.15	Pie chart showing the contributions from different components to the total material budget in the upgraded VELO within the pseudorapidity angles 0.8° and 15.4°	35
Fig. 4.1	New LHCb data centre and the optical fibre cabling to the cavern.	38
Fig. 4.2	Architecture of the TFC system for a sub-detector partition	39
Fig. 4.3	LHCb TFC starting sequence.	40
Fig. 4.4	PCIe40 board fully populated.	41
Fig. 4.5	PCIe40 clock distribution.	42
Fig. 4.6	VELO DAQ firmware architecture. Two parallel data streams are instantiated (each one associated with a PCIexpress interface).	43
Fig. 4.7	This figure shows a full duplex Arria 10 transceiver channel. (Top) Transmitter path. (Bottom) The incoming serial data.	45
Fig. 4.8	Scheme of Intel PMA transceiver	46
Fig. 4.9	VELO Data processing scheme.	52

Fig. 4.10 (left) 1 input 2 outputs block sorts 1 of the bits in the frame without any memory. (Right) 2 input 2 output block sorts 2 links at a time and consumes 4 small FIFO memories.	54
Fig. 4.11 Internal latency of the VeloPix ASIC.	55
Fig. 4.12 Architecture of the Post-Router.	57
Fig. 4.13 Data processing block output data interface.	58
Fig. 4.14 Event ID block output data interface.	63
Fig. 4.15 The Multi Event Package (MEP).	65
Fig. 4.16 Four chip ID bits in the data processing according to the FE ASIC of origin. Two MSB (sensor ID) in red and two LSB in black inside each VeloPix.	69
Fig. 4.17 TELL40 resources in detail without ICF and clustering.	74
Fig. 4.18 Clock domains of the VELO TELL40.	75
Fig. 4.19 Structure of VELO bypass firmware.	76
Fig. 4.20 This figure shows the SPPs latency at the input of the Router (in blue) and at the moment in which the data is stored in memory (in red). This latency is measured with the data processing clock, 4 times faster than the LHC.	78
Fig. 4.21 Control and Timing firmware architecture for a VELO slice.	81
Fig. 4.22 VeloPix ECS protocol layer.	83
Fig. 4.23 Timming diagram of the signals sent from VELO SOL40 firmware and the GBTx deserialized word sent to the VeloPix ASIC. 0 and 1 are the 2 bit word sent out by the down-stream FIFO.	83
Fig. 4.24 TFC signals on firmware and between module GBTx and VeloPix ASIC. Where: 0 (BXID reset), 1 (FE reset), 2 (Test pulse injection (calibration)), 3 (Snapshot), 4 (TFC sync), 5 (Shutter), 6 (Not used), 7 (Alignment, always 1).	87
Fig. 4.25 Architecture of the LHCb control system for the FE electronic boards.	88
Fig. 5.1 FPGA logic cell.	92
Fig. 5.2 FPGA design flow.	94
Fig. 5.3 SPIDR readout system.	100
Fig. 5.4 MiniDAQ 1 connected to the first prototypes of the VELO electronics.	101
Fig. 5.5 MiniDAQ 2 setup.	104

Fig. 5.6	Data transmission loss, measured in the final prototypes, for all the FE data acquisition links. Test performed in the copper lines from including high-speed data tapes, VFB and OPB.	106
Fig. 5.7	Characteristic impedance (final design) per link for all the FE links.	107
Fig. 5.8	Schematic of the CTLE circuit.	107
Fig. 5.9	Optimisation of the control link transmission running at 4.8 Gb/s.	108
Fig. 5.10	GBTx data path block diagram.	109
Fig. 5.11	Capture of the ECS data sent to the VeloPix.	110
Fig. 5.12	(Top) VeloPix noise counts vs global threshold before and after equalization. (Bottom) VeloPix matrix noise scan.	112
Fig. 5.13	BER of the GWT links with respect to the phase of the 320 MHz clock. BERs below 10^{-12} is considered acceptable, as they represent less than 1 error every 200 s.	113
Fig. 5.14	VeloPix BER vs digital/analog power supply (measured on the ASIC). Tests were done on VeloPix v1, MiniDAQ v1, and first prototypes of the on-detector electronics.	114
Fig. 5.15	Picture of the 3 VELO slices during the test beam.	116
Fig. 5.16	Preliminary VELO telescope data.	117
Fig. 6.1	High speed flex tapes test bench.	120
Fig. 6.2	(Left) Alignment pictures showing the displacement of the cable when the opposite far end is connected. (Right) Example of an eye diagram generated for every single link of all the cables.	121
Fig. 6.3	(Left) Continuity test setup. (Right) Open circuit for 1000V test setup.	122
Fig. 6.4	Carrier board test setup.	123
Fig. 8.1	Longitudinal section of the LHCb upgrade experiment.	130
Fig. 8.2	LHCb readout system.	133
Fig. 8.3	(Left) VELO “slice” cross section. (Right) Module zoom.	134
Fig. 8.4	LHCb readout architecture.	136
Fig. 9.1	Sección lonxitudinal do LHCb mellorado.	144
Fig. 9.2	Diagrama do sistema de adquisición de datos do LHCb.	148

Fig. 9.3	(Esquerda) Sección transversal de unha “slice” do VELO. (Dereita) Zoom do módulo.	149
Fig. 9.4	Arquitectura do sistema de control e sincronización do LHCb.	151

List of Tables

Tab. 2.1	Number of collisions for the 25 ns filling scheme in the four interaction points.	7
Tab. 3.1	VeloPix ASIC specifications.	23
Tab. 3.2	configurable FE SuperPixel Packet frame.	28
Tab. 4.1	64 bit TFC word. It contains all the BXID commands distributed across the experiment that keeps the data taking synchronous.	49
Tab. 4.2	This table shows the increase in memory resources due to the number of words stored in the full Router. Ram overflows represent the data estimated data loss in the hottest module.	56
Tab. 4.3	(Top) VELO TELL40 output format. (Bottom) VELO Specific Data Output.	59
Tab. 4.4	(Top) VELO TELL40 TFC synchronization output format (Bottom) VELO Specific "Sync" frame.	60
Tab. 4.5	FE data for the case of VELO data processing error or warning.	61
Tab. 4.6	VELO data processing address map.	72
Tab. 4.7	VELO bypass firmware output format.	76
Tab. 4.8	This table shows the selection bit distribution of the OPB GBT word.	80
Tab. 4.9	This table shows the selection bit distribution of the FE GBT word for the 6 VeloPixes controlled from a single GBTx. This word is distributed at a 40 MHz rate, thus 2 bit slots corresponds to 80 Mb/s FE links and 8 bit slots to 320 Mb/s FE links.	82
Tab. 5.1	Comparison of the most popular silicon technologies.	93
Tab. A.1	(Top) VELO TELL40 output format (Bottom) VELO Specific Data Output.	156

