

Data Acquisition in LHCb VELO Test beams

Niko Neufeld*CERN

E-mail: niko.neufeld@cern.ch

The 2006 testbeam campaign for the LHCb Vertex Locator was not only an important test for all sub-systems of this important detector. It was also the first test of the final LHCb data acquisition system. This paper will explain the LHCb data acquisition system and its major components. The correspondence in the DAQ setup for the testbeams will be emphasised. The final readout electronics of the Vertex Detector will also be discussed.

*The 16th International Workshop on Vertex detectors
September 23-28, 2007
Lake Placid, NY, USA*

*Speaker.

1. Introduction

LHCb's [1] Vertex Locator (*VeLo*) [2] must measure precise track coordinates close to the interaction region. The *VeLo* consists of 21 silicon stations each of which has modules consisting of strips at constant r and constant ϕ . This arrangement of the strips allows for very fast track-reconstruction, which is essential for use in LHCb's trigger. Each station is split into two half-stations, which can be moved in and out of the acceptance of the Large Hadron Collider (LHC). This is required, because the *VeLo* will be closer to the beam, than the nominal aperture of the latter during injection.

In total there are thus 84 modules to be read out. The radiation levels in LHCb are such, that any processing of the data, other than simple thresh-holding or digitization must be done far away from the detector in an area shielded from the radiation. In order to have the full information available for the optimal reconstruction of the data of each sensor, in the *VeLo* all strips are read-out over analogue links. The data are digitized upon reception.

While all detectors in LHCb acquire data at the nominal collision rate of the LHC, 40MHz, the data are transferred to the off-detector electronics at 1MHz only. The reduction in acceptance of events required is provided by LHCb's first level trigger, called L0-trigger. It uses information from the calorimeters and the muon-detectors to select events, characterised by tracks with high transverse momentum, p_t .

The data acquisition of the *VeLo* is complemented by a sophisticated control-system, responsible for accurate monitoring of the environmental conditions and the configuration and monitoring of the detector electronics.

Many individual components: sensors, readout-links, readout-cards, control-electronics, control-software, cooling and vacuum system, must work together to ensure optimal detector operation.

To ensure a successful integration, a testbeam campaign was undertaken in 2006, where in successive stages all essential components were brought together. This paper reports on this effort, which preceded the installation and commissioning of the *VeLo* in LHCb, which is currently (November 2007) ongoing and is foreseen to come to an end shortly in 2008. An emphasis will be put on the data acquisition, because the *VeLo* testbeams were at the same time the first occasion to use the hard- and software foreseen for the global LHCb data acquisition and control systems.

After a quick recapitulation of the architecture foreseen for LHCb's data acquisition and controls, the setup of the testbeam will be discussed, with an emphasis on the lessons learnt for the global commissioning of the *VeLo* and LHCb.

2. The LHCb Online system

Data acquisition and detector control are very closely integrated in LHCb. Wherever possible, the same technologies are re-used both in hardware and software. The data acquisition is based on the industry standard Ethernet. All connections are point-to-point, no path is shared. The control system uses a physically separated Ethernet network to avoid any interference between data and control traffic. The trigger has two levels, Level0, briefly described in the introduction, and the high-level trigger (*HLT*). The *HLT* is a collection of software algorithms, which run on a general purpose CPU. The high event rate of 1MHz necessitates a large number of processing units

(“cores”) and is currently estimated to be 6000 for the entire experiment. The data are pre-processed by the common LHCb read-out boards *TELL1* [3]. These boards, which will be described in some more detail later, send their data already on Ethernet links as IP-packets. Throughout the system data are *pushed*. Every station sends as soon as it is ready and it can assume that the receiving end can cope with the data. In case of a congestion every entity participating in the data-flow, except the routers in the network, which connect the TELL1-boards with the CPU servers with the processing units, can send a *throttle* signal to disable the trigger. The trigger decisions and in general all synchronous information is transmitted by the *Timing and Fast Control* system, which is based on the radiation-hard TTC technology developed for the LHC experiments [4]. When a throttle signal is received, the trigger is disabled. This gives the system time to recover from congestion, without biasing the trigger, because only event, which have never entered the system are thrown away. The events, which are already in the system, and might be responsible for causing the problem in the first place, will go through in any case.

The data acquisition system is fairly large, consisting of over 300 TELL1 boards and over 1000 CPU servers. A large IP router provides the required connectivity. Each TELL1 board must be able to reach each server. A particular difficulty for commercial network equipment comes from the synchronous nature of the traffic. Whenever a trigger is sent to the TELL1 boards, they will process the event data, and send their event-fragments, practically at the same time to a pre-determined CPU server. Since all links can only provide the same throughput of 125MB/s this results in an enormous instantaneous overcommitment. Only high-end routers have the required buffer-space to buffer all packets and send them one after the another on to the destination¹. A schematic drawing of the data acquisition system is shown in Fig. 1.

The control system must configure, monitor and control all devices and electronics in the entire experiment. At the link-layer most equipment is either attached via Ethernet, these devices are in the radiation free part of the experimental area, or via the *SPECS* bus [6], which is a radiation and magnetic field tolerant bus, developed specifically for LHCb². The software for the experiment control system (*ECS*) is based on a commercial SCADA system, PVSS, and a generic framework developed jointly by the 4 large LHC collaborations [7].

3. The VeLo

The VeLo consists of 21 stations, each made up of two modules (upper and lower). Schematically shown in Fig. 2, consists of two silicon sensors, one of sensors at constant r , one of sensors at constant ϕ , two hybrids which contain most importantly the read-out chip (Beetle), the substrate, which provides mechanical support and thermal connection for cooling, the cooling block, which connects the substrate to the VeLo cooling system, and the paddles which allow precise movement and positioning of the modules. For the readout each module is connected to two of the standard LHCb readout-boards TELL1. The VeLo is special in that contrary to all the other sub-detectors of LHCb, it has been decided to readout the data over analogue cables and digitize them only at the input of the TELL1 boards, safe from the radiation close to the VeLo. The data need to be

¹Most routers will simply drop packets in such a situation. This is perfectly valid behaviour for Ethernet equipment, since Ethernet [5] is an unreliable protocol.

²There are, in particular for industrial equipment, also a few other field-buses in use, such as CAN

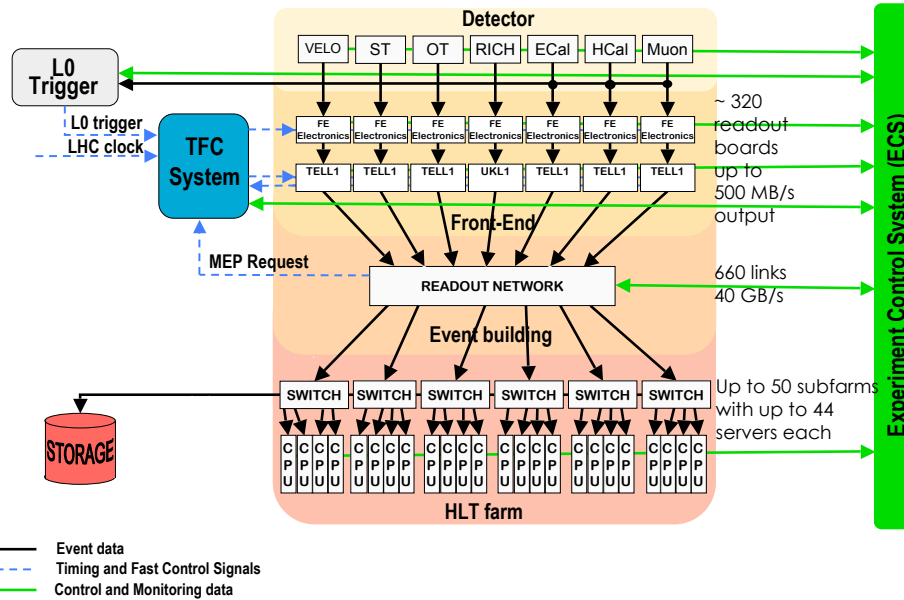


Figure 1: The LHCb Online system: data acquisition and control networks.

transported over $\sim 60\text{m}$ cables with minimum amplification and distortion. Testing this critical component was one of the important goals of the testbeam.

4. The VeLo testbeam campaign in 2006

The VeLo testbeam campaign in 2006 is called Alignment Challenge and Detector Commissioning (ACDC), reflecting the dual use of the three periods. The first stage in April was not actually a testbeam, but an integration test in the lab.

It was followed by the second stage, ACDC2, which was taking place at CERN in the H8 testbeam line, which is part of the SuperProtonSynchrotron (SPS). H8 provides various kind of secondary particle beams and also primary protons. The particles used for ACDC were 180 GeV pions and 400 GeV protons. At this stage 3 modules of the VeLo were ready as well as most of the final components used for control and DAQ. In total this testbeam lasted for 18 days. Over 1 TB of data was acquired and lots of useful lessons learnt. The full length readout cables were used and found to perform very well.

The final test-beam took place in November of the same year. ACDC3 had one half of an entire VeLo half in use, in total 10 modules. All other components, such as cooling modules, power supplies were taken from the final production batches. The software in the FPGAs and the receiving computers was a prototype but contained most of the important features. The rest of this section focuses on this important integration test.




Figure 2: Schematic drawing of a VeLo module. Explanations are in the text.

4.1 A “slice” of the VeLo

“Slice”-test has been a very popular word in the LHC community. It reflects the great modularity of the complex detector designs in the LHC experiments. The VeLo is no exception. While the vacuum tank and main mechanical support is of course shared, as are the cooling system and power-supplies, each half-station is a largely independent system, with its own set of components for detection, data acquisition and controls. The VeLo slice is shown in Fig. 3. Starting from the right the sensor, hybrid and the readout chips, can be seen, followed by the two caption cables transporting the data through the vacuum tank to the repeater board. The repeater board amplifies the analogue data and sends them on towards the TELL1 board. It also sends control and timing signals, as well as the high and low voltage to the hybrids. There is no control-logic on the repeater boards, this is provided by the *control board*. The control board is connected to the TFC system and the SPECS bus, which provides the interface to the Experiment Control System. High and low voltage are also provided, from the counting house, to the repeater board. Temperature monitoring is done by the *temperature board*, which is based on an ELMB [8]. This ELMB is controlled via the CAN bus.

The setup is like it is currently being installed in the LHCb experimental area, with one ex-

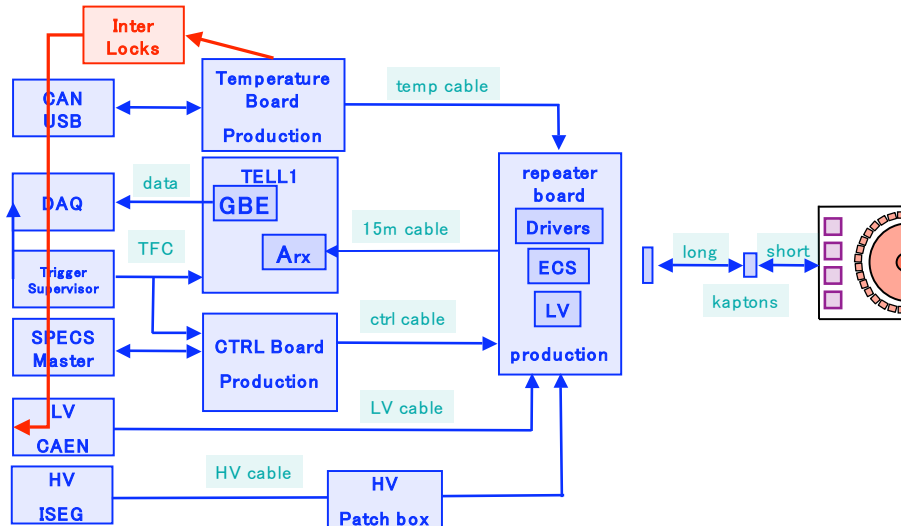


Figure 3: A “slice” of the VeLo. Explanation can be found in the text.

ception. The data cables are only 15 m long, because there were not enough of these sophisticated and expensive cables available for the test-beam. However the final length cables (60m) have been previously tested in ACDC2.

4.2 Trigger

The H8 beam-line provides of course a rather different beam-structure than the LHC will. A dedicated trigger has been developed. Its output has been connected to the generic input of the Readout-Supervisor³. The scintillator setup is shown in Fig. 4. The trigger rates, which could be achieved are shown in Table. 1

4.3 Data acquisition in ACDC3

The data acquisition system for ACDC3 is shown in Fig. 5. It is maybe not obvious, when comparing with Fig. 1 showing the full LHCb DAQ, but Fig. 5 is actually a small version (about 5%) of the final system. All elements are there: 12 TELL1 read-out boards, each connected with one of their 4 Gigabit links. A powerful router, in this case an HP5412 and 2 data acquisition PCs. Most of the time only one PC was used, because the data were analyzed on the fly. Having several PCs, while technically possible, would require to collect the results from several live-processing

³Normally the Readout-Supervisor will get a trigger decision from a dedicated decision unit, which combines input from all high p_t trigger systems.

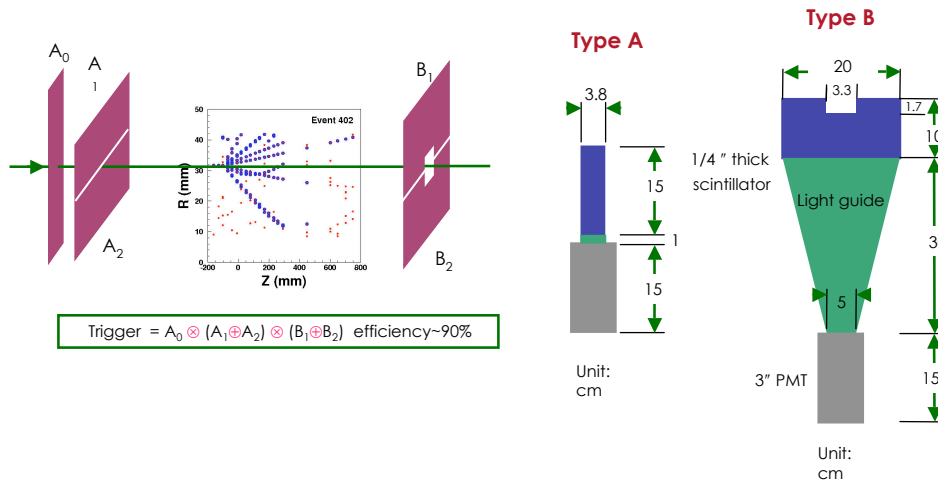


Figure 4: The scintillator trigger setup for ACDC2 and ACDC3

trigger type	rate / spill	scintillator use
straight track	6k - 9k	upstream scintillators in AND
detector interaction	5k	upstream scintillators in AND and-ed with the OR of B_1 and B_2 .
target interaction	1k	upstream scintillators in AND and-ed with B_2 and an additional scintillator
random trigger	10 kHz	

Table 1: Trigger rates for different trigger types, where 1 spill is about 4 s and the time between spills about 12 s.

jobs into a coherent output. This software framework was not ready for the testbeam. It is currently under test. So in order to have *all* data on a single PC, the trigger rate had to be limited to 2kHz. This rate results from the rather large event size of 60kB(!)⁴ from the twelve boards. In order to have a perfect understanding also of the internal FPGA processing of the TELL1s, tested for the first time as well under real beam conditions, the zero-suppressed as well as the non zero-suppressed data were read out for each trigger.

It is interesting, that a classical problem of all push-architectures manifested itself during the

⁴This should be compared with the average *total* event-size of LHCb, which is only 35 kB.

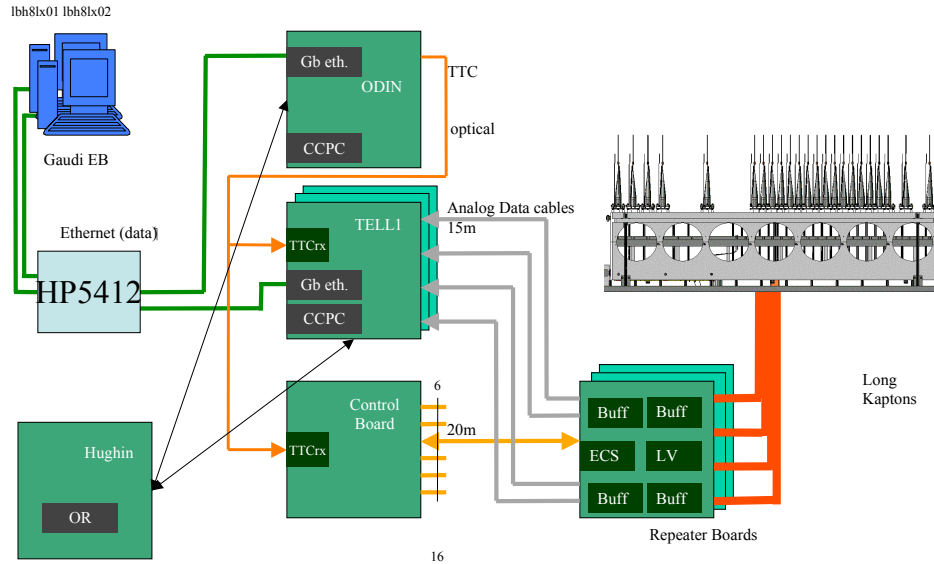


Figure 5: The data acquisition for ACDC3

first day of running. Since all TELL1s send at the same time, the switch has to buffer a lot of data, because it receives data on 12 links but can send them out only on a single link. The router which was originally used is a HP3400, which is completely loss-less at wire-speed for all its 48 ports under *random* traffic. This so-called full-mesh test is a standard mark of quality for any high-end router. The bursty nature of the data acquisition traffic is however very unfavourable for any router. And it required a much large router, with much more buffer space to cope with this type of traffic. This problem is the main reason, why sometimes a more complex *pull* protocol is chosen, where the destination PCs request data one-by-one from each source.

After this and a host of smaller configuration problems had been fixed, data acquisition went smoothly. In total 60 Million events or about 3.5TB were acquired. In a DAQ only test, zero-suppressed data alone could be read at up to 100kHz, which gives a lot of confidence that with the final system, which will be 30 times larger, the required factor 10 in readout speed can be achieved.

The main limitation was the speed at which data could be written to local mass-storage and from there copied to tape. This chain, which had to run on commodity hardware was limited to 30MB/s⁵. This rate was high, because no software trigger was applied in addition, but rather all data were stored.

⁵Again, this must be compared with the needs of the full-size LHCb data acquisition, which are about 80MB/s storage speed in total.

4.4 Eventbuilding software

The event-building software is build on the standard LHCb software framework Gaudi [9]. Since Gaudi is an object-oriented framework, it is sufficient to re-implement a several services for the Online world [10]. These concern mostly the I/O and the logging. At the high rates of data processing, it is very important to avoid any unnecessary memory to memory copies. This is achieved by a shared memory architecture. In LHCb data belonging to several triggers are packed together into one packet, in order to reduce the overheads due to the data transport. The eventbuilder is the first algorithm. It verifies the syntactical structure of the data and populates the shared memory. A system of buffers accessed under the publish / subscribe paradigm feeds data, strictly speaking references to data, into consumer algorithms, which put their results back into the shared memory. The architecture is illustrated in Fig. 6 The overall overhead for data formatting

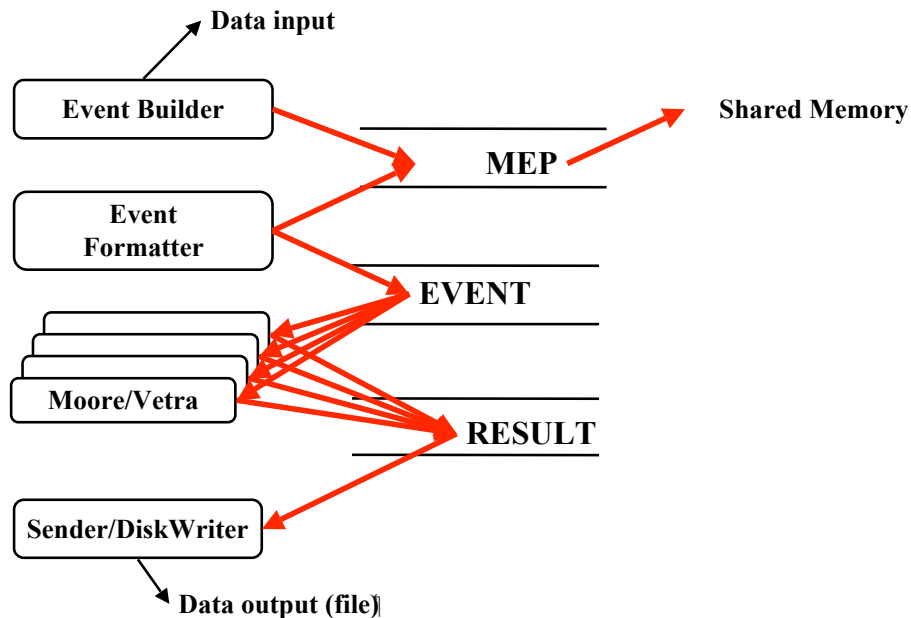


Figure 6: The layered architecture of algorithms in the LHCb Online software framework GaudiOnline. Data are copied only once into the shared memory as soon as they are received from the network. From then on only references are passed between algorithms.

and transport has been measured to be approximately 6% of a dual-processor, 2.8GHz server with 2 GB of RAM.

For the VeLo a specific software package, VETRA, has been developed, which takes into account the different environment, beam-structure and detector setup in the testbeam. All these are quite different compared to what the VeLo will have to deal with in LHCb. The standard algorithms developed for the LHCb High Level Trigger are therefore not suitable for the testbeam

and a custom software was needed. It will however also be used in LHCb for calibration and test-runs and in particular whenever the VeLo is not in its standard position for physics, but moving in or out of the beam.

5. Conclusion

ACDC3 was the last testbeam for and with the LHCb VeLo. In all important aspects it can be considered a resounding success. All major design decisions proved to be correct and the various sub-components worked well together. It was also a very important validation of the soft- and hardware architecture of the LHCb Online system. A host of small and not-so-small problems could be worked out, when FPGA firmware, network transport and eventbuilding software worked together for the first time. The overall system performance from the point of view of the DAQ was very satisfactory and there is every reason to be confident that the system will work well, when it goes into operation in May 2008.

Acknowledgments

The author wishes to thank to entire VeLo team for the excellent collaboration during the testbeam periods. In particular and in no specific order: Kazuyoshi Akiba, Paula Collins, Pavel Jalocha, Marina Artuso, Ray Mountain, JC Wang for the material they provided for this paper and the talk it is based on. I would also like to thank my colleagues from the LHCb Online team, in particular Artur Barczyk and Sai Suman Cherukuwada.

References

- [1] LHCb Collaboration, *LHCb Reoptimized Detector: Design and Performance* (CERN, Geneva, Switzerland, 2003), No. CERN-LHCC-2003-030.
- [2] LHCb Collaboration, *LHCb Vertex Locator Technical Design Report* (CERN, Geneva, Switzerland, 2001), No. CERN-LHCC-2001-.
- [3] A. Bay *et al.*, *Nuclear Instruments and Methods* **A560**, 494 (2006).
- [4] Timing, Trigger and Control (TTC) Systems for the LHC, 2007, <http://ttc.web.cern.ch/TTC/intro.html>.
- [5] *IEEE 802.3 LAN/MAN CSMA/CD Access Method*, edited by IEEE (IEEE, Piscataway, N.J., 2005).
- [6] *SPECS : A Serial Protocol for the Experiment Control System of LHCb*, LAL, 2007, <https://lhcb.lal.in2p3.fr/Specs/Documentation/Specs4.0.pdf>.
- [7] The JCOP Framework Project, <http://itcobe.web.cern.ch/itcobe/Projects/Framework/welcome.html>.
- [8] *Embedded Local Monitor Board with the ATmega128L processor*, NIKHEF, <http://elmb.web.cern.ch/ELMB/ELMBhome.html>.
- [9] G. Barrand *et al.*, in *Proc. CHEP* (Science Press New York Ltd., New York, 2001).
- [10] S. S. Cherukuwada *et al.*, in *Proc. International Conference on Computing in High Energy and Nuclear Physics* (IOP, Victoria, BC, 2007).