Performance of the LHCb Detector During the LHC Proton Runs 2010 - 2012

Richard Jacobsson, *Member, IEEE,* on behalf of the LHCb Collaboration

LHCb-PROC-2013-030
23/05/2013

***Abstract*–The LHCb experiment at CERN's Large Hadron Collider (LHC) searches for New Physics through precision measurements in the domain of heavy flavour physics, exploiting in particular the large B hadron production. After a short introduction to the requirements of the LHCb detector, this paper reviews the operational strategy during the first three years of data taking. The focus is on the detector performance together with a description of several fundamental system developments which emerged in this period and which allowed LHCb to venture well beyond its design parameters and to extend the physics program. With the inclusion of a solid charm physics program and electroweak and soft QCD measurements in the forward direction, LHCb has established itself as an excellent forward general purpose detector at the LHC. Running at twice the design luminosity, LHCb has been able to collect an integrated luminosity of more than 3 fb^{-1} at an operational inefficiency of less than 4% in the first three years.**

I. Introduction

THE LHCb experiment [1] is located at one of the four interaction points on the Large Hadron Collider (LHC) at CERN. Its prime discovery potential for New Physics lies in measuring the effects of new physics in CP violation and in rare decays. Deviations from the Standard Model predictions are expected to manifest themselves most visibly in processes which are strongly suppressed in the Standard Model, such as those which are predominantly mediated by loop diagrams and involving flavor changing neutral currents. The virtual effects in the loop diagrams allow LHCb to increase the sensitivity to new physics in a mass range far beyond the center-of-mass energy scale accessible in direct searches at the LHC. In that respect the LHCb experiment is complementary to the searches by the Atlas and the CMS experiments.

The initial aim of the LHCb experiment was to perform precision measurements in the very large and diverse domain of processes involving B hadrons. In the proton-proton collision mode, the LHC accelerator is to a large extent a heavy flavor factory producing over 100.000 $b\bar{b}$–pairs every second at the LHCb design luminosity of 2x10^{32} cm^{-2}s^{-1}, giving the experiment access to all quasi-stable b-flavored hadrons.

Together with the progressive commissioning strategy of the LHC accelerator in the first years of operation and thanks to a very good detector, flexible trigger, high data acquisition capacity and powerful offline processing, LHCb has been able to expand the physics program. The charm production cross-section is 20 times larger than the beauty cross-section, and CP violation and mixing in the charm sector are both generally considered to be small in the Standard Model but have not been fully explored.

Below follows an account of the main challenges and highlights, and the evolution of the LHCb operational strategy which laid the foundation for a large number of world-best physics results in the first three years of operation.

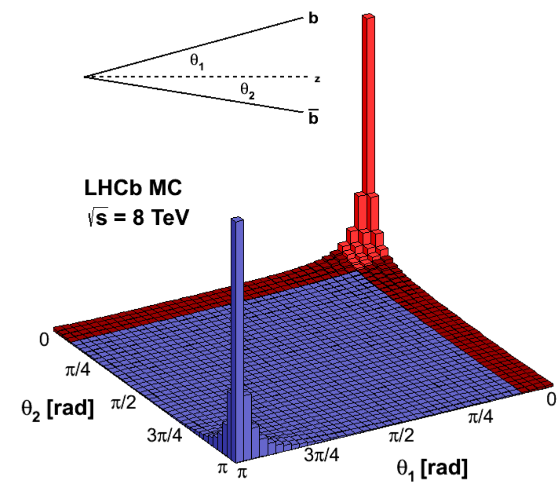

Fig. 1: The distribution of the polar angles of the **b**-quark and the $\bar{\mathbf{b}}$-quark as they are produced at the LHC.

II. LHCb Detector and Design Operating Conditions

At the LHC, the initial state partons have generally different longitudinal momenta leading to a strong boost of the $b\bar{b}$- and the $c\bar{c}$-pair. Fig. 1 shows the distribution of the polar angles of the b and the $\bar{b}$-quarks. Consequently the resulting pair of B or alternatively D hadrons appears in the same hemisphere. This opens the possibility of using the decay products of the accompanying hadron to tag the identity of the hadron decay of interest by detecting the final states in a reduced fiducial volume around the beam pipe with relatively small angle.

The prerequisite to explore fully flavor physics at the LHC consists of collecting a very large statistics of B and D final states with the help of a fast but flexible trigger with high sensitivity to a very large variety of different final states. Due to the complexity of the final states the trigger needs to perform a level of reconstruction in order to identify the signal vertex signatures with high efficiency against a background of inelastic events which is several orders of magnitude larger. The key detector requirements consist of

Manuscript received November 16, 2012.
Richard Jacobsson is with CERN, 1211 Geneva 23, Switzerland (telephone: +41-22-767 3619, e-mail: Richard.Jacobsson@cern.ch).

- excellent vertex and impact parameter resolution in order to resolve the complex vertex signatures,
- excellent proper time resolution to resolve very fast oscillations
- high track reconstruction efficiency,
- a dipole magnet with precisely known field for high precision momentum measurement,
- very good mass resolution to reject combinatorial background,
- and extremely good π, K, e, γ, μ identification over a wide momentum range for the flavor tagging, and the reconstruction and identification of the particular heavy flavor decay,

all of which have to span the full forward acceptance.

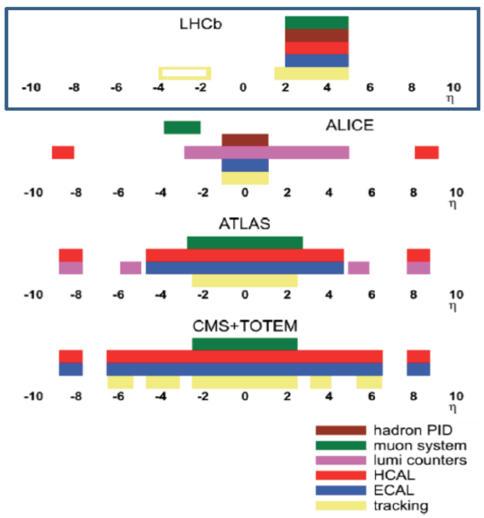

Fig. 2: Illustration of the unique acceptance of the LHCb detector in pseudo-rapidity as compared to the other LHC experiments.

Fig. 2 shows the unique coverage of the LHCb instrumentation as compared to the other LHC experiments. The LHCb experiment has been conceived as a 20 m long single-arm forward spectrometer with special emphasis on particle identification (Fig. 3). The detector acceptance spans the polar angles 15 mrad to 300 mrad in the horizontal bending plane of the spectrometer magnet and 250 mrad in the vertical non-bending plane, equivalent to a pseudo-rapidity of about $2 < \eta < 5$. While this corresponds to only 4% of the solid angle, it includes ~40% of the $b\bar{b}$-pair production cross-section. Below is a short summary of the main detector features. A complete detailed description of LHCb may be found in [1].

Vertexing is ensured by the VErtex LOcator (VELO). It consists of 2 x 21 half-moon silicon strip sensors organized perpendicularly to the beam over a distance of about one meter around the LHCb interaction point. Each sensor has a radius of 44 mm and provides both an R and a φ measurement. The inner edge of the sensitive area of the sensor is at 7.5 mm from the beam axis in the data taking position. For reasons of protection the VELO is a movable device which is moved out by 30 mm during the injection and preparation for the next LHC physics fill. As soon as the beams are in stable collision mode, the VELO is closed around the beam to the data taking position determined by a real-time reconstruction of the luminous region. In this way the data-taking position is reproducible to ~5 μm.

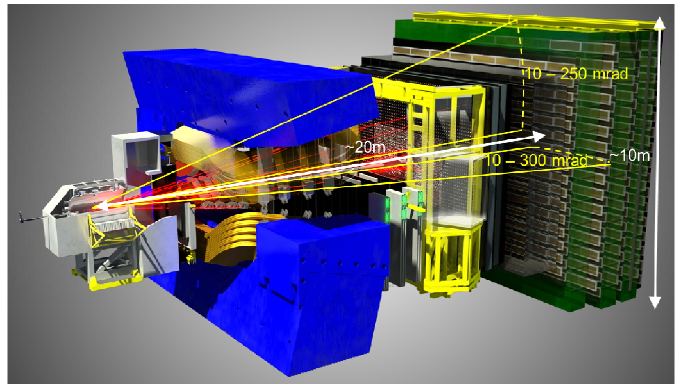

Fig. 3: The LHCb single-arm forward spectrometer detector at the LHC.

Tracking is provided by one station before the magnet consisting of 4 layers of silicon strip modules, and 3 stations behind the magnet consisting of silicon strip detectors in the central high-occupancy region and 4 layers of straw tubes per station towards the outside.

The π/K separation is provided by two RICH detectors with three different radiators to cover the full momentum range of the B hadron decay products between 2 – 100 GeV/c. Electron and photon identification is mainly based on the balance of the energy deposited in the calorimeter system and the track momentum, and the matching between the corrected barycentre position of the calorimeter cluster and the extrapolated track impact point. The calorimeters provide a fast and simple first level trigger information in terms of the local clusters with the highest transverse electromagnetic and hadronic energy.

Muon tracking and identification is provided by five stations of multi-wire proportional chambers interleaved with iron filter walls. The first muon station is located before the calorimeters and has gas electron multiplier chambers in the central high occupancy region. The muon chambers in the five stations have projective geometry in order to provide a fast first level trigger based on the muon candidates with the highest transverse momentum.

The need for reconstruction to achieve high trigger efficiency has driven the design of the trigger and the readout architecture. Consequently, LHCb opted for a relatively simple and inclusive first level trigger (L0) logic based on transverse energy and momentum of leptons, photons, hadrons and muons as measured by the calorimeters and the muon detector system, and full event readout at 1 MHz to an event filter farm consisting of a very large number of commercial multi-core PCs. The event filter farm performs the main event selection by a software High-Level Trigger (HLT) with access to all detector information. The initial design specified an output event rate to storage of 200 Hz.

The LHC accelerator was designed to deliver a luminosity of $10^{34} cm^{-2} s^{-1}$ through a very large number of interactions per crossing. However, flavour precision physics relies on resolving properly the vertex structure. Event pileup significantly complicates the b-decay vertex reconstruction and flavour tagging, and increases the combinatorial

pileup veto condition in the trigger was dropped. However, a fundamental but extremely challenging turn point in the operational strategy of LHCb came when the LHC commissioning changed strategy in June 2010 from

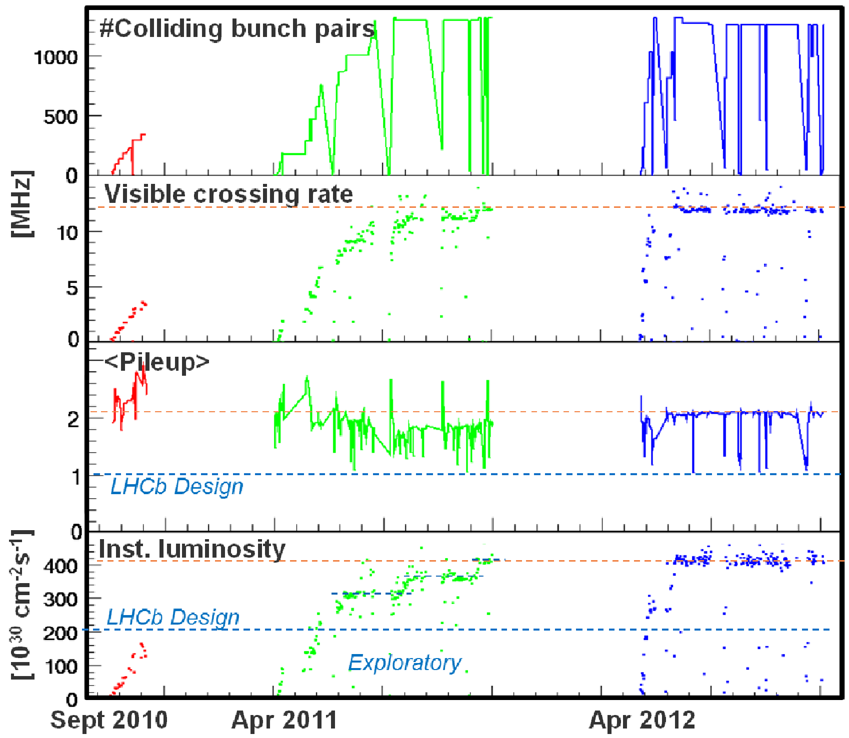

Fig. 4: LHCb operating conditions between 2010 and 2012 in terms of the number of colliding bunches, the rate of visible bunch crossings which is seen by the first level trigger, the average event pileup per visible crossing, and the instantaneous luminosity.

background. The increased detector occupancy also leads to excessive reconstruction times in the High-Level Trigger. As a consequence, the LHCb experiment was designed to take data at a luminosity of $2x10^{32}cm^{-2}s^{-1}$ with ~2600 bunches at 25ns spacing corresponding to an average number of visible interactions per crossing of ~0.4 in order to maximize the single interaction triggers. To further limit event pileup to one interaction per recorded bunch crossing, the LHCb detector includes two additional upstream silicon strip sensors in the VELO to provide a fast pileup veto signal in the L0 trigger.

In order to run at two orders of magnitude lower luminosity than the LHC design required in the initial plan a defocusing at the LHCb interaction point and a "progressive" trigger along with the luminosity decay in each fill.

III. Operating Conditions and Strategy 2010 - 2012

Fig. 4 shows the evolution of the LHCb operating conditions between 2010 and up to mid-October 2012 in terms of the number of colliding bunches at the LHCb interaction point, the rate of visible crossings to be dealt with by the trigger, the average event pileup per visible bunch crossing, and the instantaneous luminosity. As the number of bunches and the bunch intensity was expected to remain limited in 2010, the same beam focusing as the Atlas and CMS experiment was applied at the LHCb interaction point and the

commissioning many bunches with low intensity to rather commissioning first nominal intensity per bunch. The average event pileup in LHCb quickly reached as high as three. The sub-detectors and the readout system performed extremely well and the reconstruction was much more robust than anticipated in these conditions of high occupancy. On the offline side, physics analyses demonstrated that event pileup could be handled largely thanks to well separated primary vertices and the simple fact that the B and D events of interest are typically accompanied by the much higher rate of visible minimum bias events with lower average multiplicities. The main limitation came from the High-Level Trigger and the offline reconstruction which suffered from excessively long processing times associated with bunch crossings with very large event pileup.

In order to cope, three developments of major significance emerged in a very short time, the first of which was a change of the trigger strategy with the introduction of global event cuts on hit multiplicities to reject events with very large event pileup. Secondly, the CPU capacity of the event filter farm was tripled. Thirdly, the concept of an LHCb-driven real-time luminosity control based on adjusting the beam transversal overlap at the LHCb interaction point was proposed and tested in July 2010 [2]. These are described more in detail in Section

IV together with the detector performance. Other subsequent measures to manage the performance with the higher event pileup and higher luminosity included regular sub-detector ageing and calibration scans with beam, and automated end-of-fill calibration.

Encouraged by the high performance demonstrated by the sub-detectors, and the potential of the operational improvements, the operational strategy continued to evolve in 2010 – 2012.

A. 2010 Run

While 2010 was a "test drive" for the LHCb sub-detectors, the validation of the trigger and the offline processing concept was of fundamental importance. With the change of the LHC commissioning strategy, the first year of operation became characterized by an exploration of the LHCb performance in the event pileup domain.

In the very first phase of the commissioning with a very low number of bunches and low bunch intensity, LHCb focused on a minimum bias physics program very much profiting from an increase of the event storage rate from the design of 200 Hz to 2 kHz. Studies of strangeness production and of anti-particle/particle production ratios in the forward region are examples of interesting early measurements. In the intermediate phase with nominal bunch intensities but low number of bunches, LHCb was able to maintain lower than nominal trigger thresholds and hence boost the trigger efficiencies for hadronic B decays up to ~75%. This phase also allowed exploring the LHCb physics potential in charm physics by reserving a significant portion of the bandwidth for charm triggers with corresponding trigger efficiency of up to ~40% for prompt D hadrons. As the LHC luminosity increased with increasing number of bunches in the second half of the 2010 running year, trigger conditions were gradually tightened in order to exploit at maximum the available readout bandwidth and CPU power in the HLT farm. As shown in Fig. 4, at the end of 2010, LHCb took data at 80% of the design luminosity with only 344 colliding bunches instead of 2622. This is equivalent to six times the design in terms of the number of average visible interactions for bunch crossing.

B. 2011 Run

Due to a relatively fast LHC re-commissioning in 2011, the ultimate luminosity of 2010 was reached after only one month. The nominal number of bunches for the LHC Run I with 50ns (~1300) was reached at about 1/3 into the 2011 running year. As a result the second year of LHCb operation was already largely characterized by stable data taking at high efficiency with the opportunity to explore the LHCb performance in the luminosity domain while running with a significantly lower event pileup than in 2010. Fig. 4 shows the stepwise increase in the instantaneous luminosity and the extended stable periods as defined by the need to monitor the sub-detector behavior in conditions well above the design luminosity.

The trigger in 2011 was tuned to a luminosity of 3.5×10^{32} $cm^{-2}s^{-1}$ for the bigger part of the year and the event storage rate was increased to 3 kHz in order to accommodate 1 kHz for the continued charm physics program.

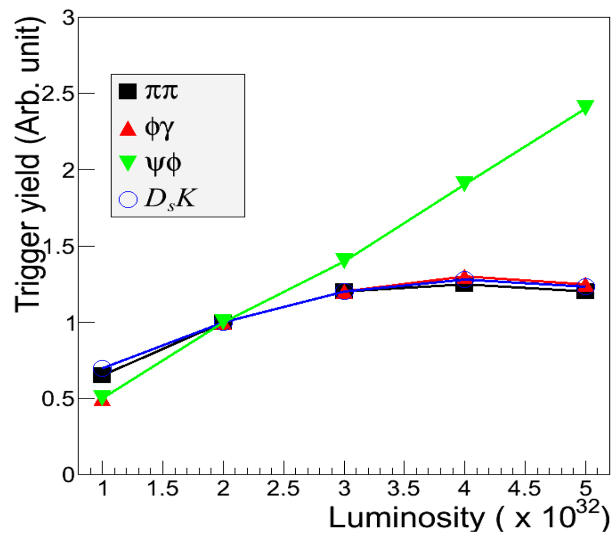

Fig. 5: Trigger yield as normalized to the yield at the nominal design luminosity of 2×10^{32} $cm^{-2}s^{-1}$ for B decays with muons and hadrons in the final state.

Fig. 5 shows the trigger yield for final states with muons and with hadrons as a function of instantaneous luminosity. The yield is normalized to the yield expected at the nominal design conditions of 2×10^{32} $cm^{-2}s^{-1}$. The figure illustrates the effect of the event pileup which ranges from an average pileup per visible crossing of one up to 2.4 in the plot. As seen in Fig. 5 the trigger yield for hadronic channels saturates at high luminosity. The reason is that in order to respect the 1 MHz readout limitation after the first level trigger, a stronger cut on transverse energy is needed with event pileup in order to suppress the increasing rate of triggers due to overlapping clusters faking particles with larger transverse energy. The optimal luminosity was found to be around 4×10^{32} $cm^{-2}s^{-1}$.

In total over one fb^{-1} was collected by the end of 2011 at an operational efficiency of 91%. The offline data quality evaluation qualified more than 99% of the data good for all physics analyses.

C. 2012 Run

A very fast LHC re-commissioning allowed reaching a luminosity of 4×10^{32} $cm^{-2}s^{-1}$ with the nominal number of bunches for 50 ns operation after only a month. Yet another data taking mechanism of significant importance was put in operation at the beginning of 2012 which allowed LHCb to profit from 20-30% more CPU power in the event filter farm. Instead of increasing the size of the farm, the mechanism defers a fraction of the High-Level Trigger processing to the inter-fill time of several hours between LHC collision periods when virtually no computing power is needed for the detector.

With the extended LHCb physics program in mind together with the upcoming 2-year shutdown of the LHC accelerator 2013-2014, the output event rate to storage was further increased to 5.5 kHz in order to collected additional statistics.

As a result, the LHCb data taking in 2012 has been characterized by a stable trigger configuration and very stable data taking at the optimized instantaneous luminosity of $4\text{x}10^{32}$ $cm^{-2}s^{-1}$ and with an operational efficiency of 95%.

IV. Operational Highlights and Performance

Fig. 6 shows the LHCb online system [3][4] with the main performance numbers. The first three years of operation of the LHCb experiment saw the emergence of several new operational improvements and developments which allowed taking the LHCb detector beyond its design performance and extend the physics program. Three of the more interesting developments are described below in some more detail, together with the detector performance.

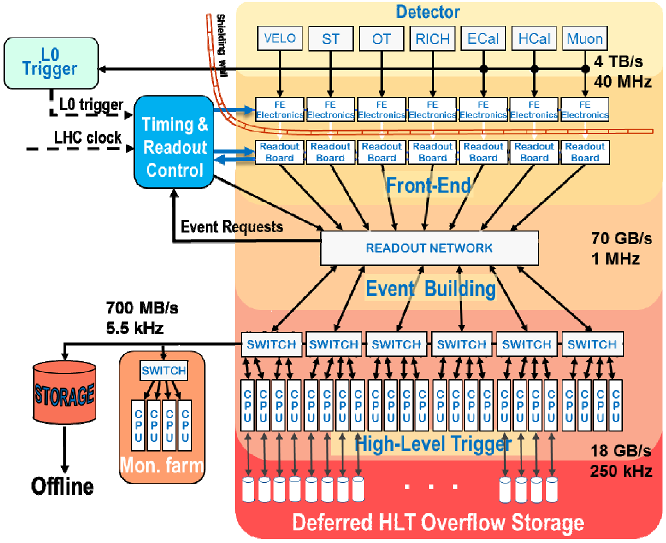

Fig. 6: Overview of the LHCb online system together with the current main performance parameters.

A. *Trigger Architecture*

Fig. 7 shows the two-level LHCb trigger architecture [5]. The first level trigger (L0) is a low-latency trigger implemented in hardware. It selects events containing muon, electron, photon or hadron candidates with relatively high transverse momenta characteristic of the decays of the heavy hadrons. At a visible bunch crossing rate of 12 MHz at $4\text{x}10^{32}$ $cm^{-2}s^{-1}$, thresholds of O(1) GeV/c for muons and O(3-4) GeV/c for electrons, photons and hadrons allow limiting the accept rate to 1 MHz while maintaining a very good efficiency for B decays and D decays. In this configuration the 1 MHz bandwidth is shared between 150 kHz of electron and photon triggers, 450 kHz of hadron triggers and 400 kHz of muon triggers. In addition, a few kHz of random no-bias triggers on beam-beam crossings, random luminosity triggers on all type of beam crossings, and detector pulsed calibration triggers allow LHCb to study trigger efficiencies, determine the integrated luminosity of data sets offline, and monitor the performance of the sub-detectors.

The software High-Level Trigger also aims at being largely inclusive and is performed in two stages. The first stage is aimed at being fast by selecting events with tracks with non-zero impact parameter. A matching of the VELO tracks with the main tracking stations and the muon stations is performed for selected tracks in order to also apply transverse momentum requirements, track fit quality cuts and a simple muon identification. The HLT1 stage achieves a rate reduction of a factor of 20.

At $4\text{x}10^{32}$ $cm^{-2}s^{-1}$, HLT2 operates at a rate of 50 kHz and performs a global event reconstruction specially tuned for the real-time environment together with full muon and electron identification. Triggers on B hadrons aim at being inclusive by triggering on partially reconstructed decays, while the physics processes which require unbiased lifetime are selected by fully reconstructed decays together with mass requirements. The latter technique is also applied to prompt D hadron decays.

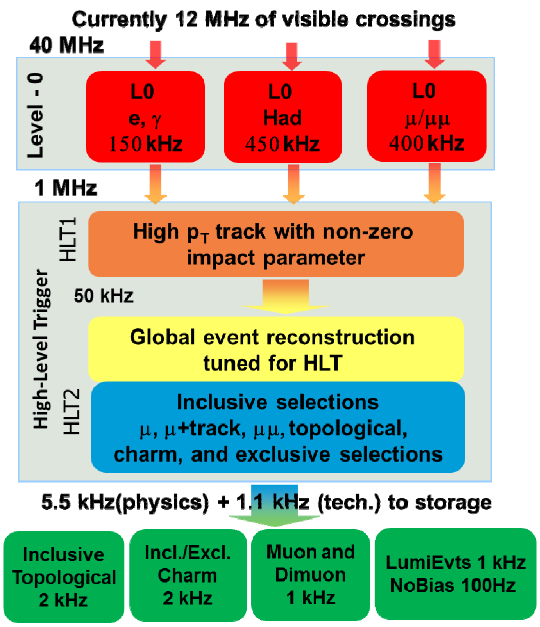

Fig. 7: LHCb trigger architecture as it has been configured and tuned during 2012.

The High-Level Trigger tasks run as ~30 000 independent copies of the executable on ~1500 CPU nodes. Each task receives the full events and performs the selection with an average processing time per event of 35-40 ms (2012). The output rate is about 5.5 kHz with the bandwidth sharing shown in Fig. 7 together with an additional ~kHz of no-bias and luminosity triggers. At an event size of O(55) kB for full events, the output data rate is about 300 MB/s. For a full running year, this amounts to 10^{10} events at a total data size of about 700 TB.

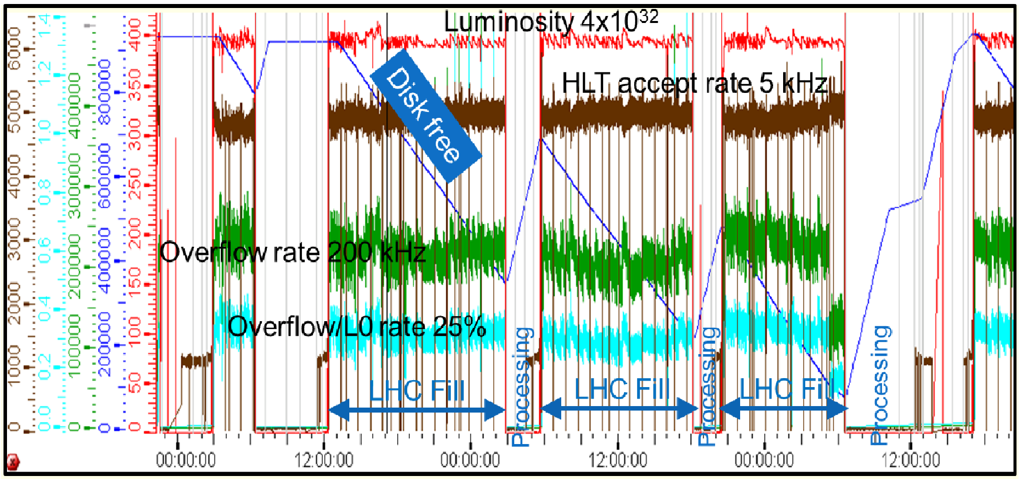

Figure 8: Example of the deferred HLT scheme. The trend shows three consecutive long fills of >10h with relatively short inter-fill time. During the collision period at $4x10^{32}$ $cm^{-2}s^{-1}$, about 25% of the L0 rate is deferred to local disk on the event filter farm nodes. The deferred events are subsequently processed by the HLT in the inter-fill time. The blue trend represents the free disk space. According to the tuning of the system, occasionally the complete disk space totaling a petabyte is only recovered in a longer inter-fill time.

In this scheme, the LHCb trigger performance is extremely high with a trigger efficiency for muonic B decays of ~90% and for hadronic B decays of ~30%. In addition, sufficient bandwidth is reserved to maintain an efficiency of 10-20% for charm decays at the highest luminosity throughout 2012.

B. *Deferred HLT*

While the event filter farm with 30 000 HLT tasks allows an average processing time per event of 30ms at the input event rate of 1 MHz, the 35-40ms processing time per events in 2012 is based on a scheme whereby a fraction of the HLT processing is deferred to the LHC inter-fill gaps. As illustrated in Fig. 6, 1000 of the 1500 event filter farm nodes are each equipped with a terabyte of local disk. Instead of producing dead-time by reducing the L0 trigger rate, the local disk space allows overflowing a fraction of the L0 input rate during the physics fill. The overflow mechanism is completely dynamic and is handled by the event builder task in each node which in order of priority:

1. first serves the HLT tasks on the node,
2. secondly overflows the excess rate of events to the local storage,
3. and only in rare case that the disk fills up, stops requesting events from the central readout controller.

The last case protects against event loss and allows proper dead-time accounting in case the system is limited temporarily to run at a lower rate due to the disks being full.

In the LHC inter-fill time the system is automatically configured into a read-back mode to process the locally stored events at a time while the farm would otherwise be idling. In reality a fraction of the farm is not configured in this mode allowing the LHCb detectors to be operated during the inter-fill for tests and calibrations.

Optimizing the system for the fill time distribution of LHC (average about 6h) and the inter-fill time distribution (average about 8h), and requiring that >95% of the physics are processed before the next fill starts, allows over-committing the farm by >25%. This is effectively equivalent to increasing the CPU power by hardware by the same amount.

This scheme has been in operation throughout 2012 starting with a 10% HLT deferral in the first month and 25% throughout the rest of the year, that is, 250 kHz. **Error! Reference source not found.** shows a case of three long physics fills with short inter-fill gaps during which the processing of the deferred HLT events was not completed until two fills later.

C. *Luminosity Control*

Most of the LHCb precision measurements have potential systematic effects associated with changing running conditions and the progressive trigger configurations which were initially considered in order to follow the luminosity decay in each LHC fill. In addition, the sensitivity to changing running conditions is further aggravated by event pileup. The prospect of running at higher luminosity than the design was to a large extent made possible with the idea of controlling the optimal luminosity directly from the LHCb online system by adjusting the transversal beam overlap at the LHCb interaction point. This method was proposed and tested first in July 2010 and was put in operation at the beginning of 2011 [2].

By running with focused beams and begin the fill with a large beam separation, and then progressively increase the overlap, this method allows ramping the luminosity to the desired value at the start of the LHC fill in a controlled way, and allows maintaining the luminosity virtually constant throughout entire fills, and even months of running. This made it possible to maintain the same carefully optimized trigger configuration over months of running. Thus, calibration and ageing effects can be carefully monitored and managed on a continuous basis. On the immediate term, the system also dynamically takes into account environmental effects and temporary technical constraints.

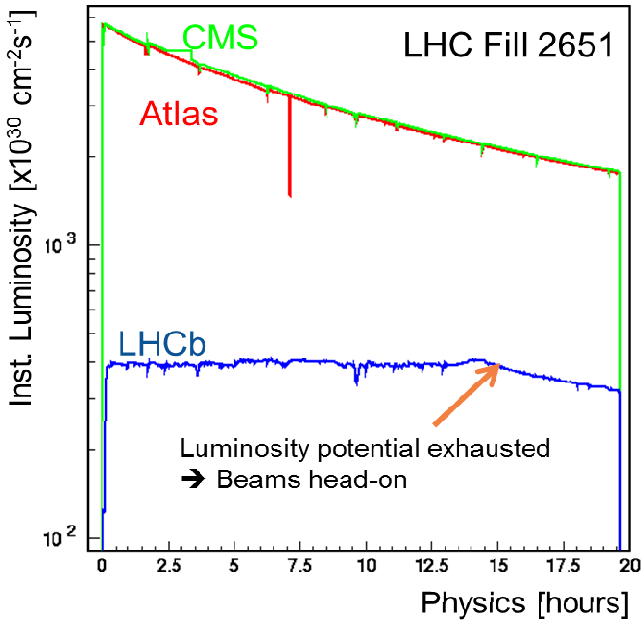

Fig. 9: The plot illustrates the LHCb luminosity control which initially ramps the luminosity up to the desired value and subsequently maintains the luminosity stable in a range of about 5% by adjusting the transversal beam overlap. In this particularly long fill the point at which the beams end up head-on is visible.

The luminosity control consists in practice of an automatic slow real-time feedback system controlled by software in the readout control system of LHCb which communicates directly with an application in LHC, which in turn drives a set of corrector magnets in the LHCb intersection region. The LHCb control manager monitors the instantaneous luminosity and a number of environmental and technical data taking parameters, and controls the luminosity in an iterative procedure together with the LHC control application.

Fig. 9 shows the evolution of the instantaneous luminosity in a particularly long fill. The beam focussing is five times stronger in Atlas and CMS than in LHCb. After 14 hours the beams are head-on at the LHCb interaction point and the typical luminosity decay follows. The system has been in routine operation during all of 2011 and 2012 and has allowed producing data at an extreme operational stability.

D. Detector Performance

In the beyond-design conditions described in Section 3, the LHCb sub-detectors have demonstrated extremely good performance. Performance numbers are either equivalent to design or better than design. The VELO has achieved an impact parameter resolution of 20 µm for tracks with high transverse momentum, and a proper time resolution of 45 fs for $B\rightarrow J/\psi\,\phi$ and $B\rightarrow D\pi$ decays. The momentum resolution ranges from 0.4% at 5 GeV/c up to 0.6% at 100 GeV/c. A mass resolution of 15 MeV/c^2 has been achieved for $J/\Psi \rightarrow \mu\mu$ decays and of 8 MeV/c^2 for $B\rightarrow J/\Psi\phi$ decays. The kaon identification is 95% with only 5% mis-identication of pions. Muons are identified with 97% efficiency with only 1-3% of mis-identifications. The energy resolution of the electromagnetic calorimeter has a sampling term of 10% with a 1% constant term.

Regular calibration and ageing scans with beam are performed to maintain the detector performance. Detector ageing is visible but it is entirely compatible with the ageing expected from the accumulated integrated luminosity. No accelerated ageing or deteriorated detector performance is observed running at the higher luminosity. About 99% of the detector channels of all sub-detectors are operational currently.

Taking together the data recorded at a center-of-mass energy of 7 TeV and at 8 TeV, a total integrated luminosity of over 3 fb^{-1} has been recorded by LHCb at the time of writing. The cumulated operational efficiency across the first three years of operation is close to 93%. The 7% inefficiency contains two currently incompressible sources, the first of which is the closing procedure of the VELO detector amounting to ~1%. The second comes from non-conformities in the implementation of the readout protocol by several sub-detector FE systems, which introduces 2.4% dead-time at 1 MHz. The remaining 3.5% are mainly related to short technical problems with the sub-detector electronics or the central readout system.

All of the LHCb sub-detectors are needed for the LHCb physics analyses and the data must have perfect quality. For this reason the operational strategy is based on aiming at diagnosing and curing immediately any problem which may have a potential impact on data quality already during data taking. In addition, there is an offline data quality evaluation after reconstruction. In total, only 2.5% of the data recorded in the three years of operations has been rejected.

V. Conclusion and Outlook

The impressive start-up and operation of the LHC accelerator has enabled the LHCb experiment to pave the way for heavy flavour physics at an entirely new level of precision in only the first three years of operation. With the demonstration that LHCb could also successfully perform precision measurements with event pileup, the operational strategy evolved very rapidly in 2010 and matured at the beginning of 2011. Largely thanks to an increased readout and trigger processing capacity, and a development to control real-time the luminosity at the LHCb interaction point, LHCb has been able to collect over 3 fb^{-1} in the first three years of operation. Most of the data was recorded at an instantaneous luminosity of $4x10^{32}\ cm^{-2}s^{-1}$, which is equivalent to an average event pileup of visible interactions of 1.7. The strategy opened the door for LHCb to extend the physics program. The additional statistics and well managed systematic effects with the stable trigger and data taking conditions have led to a very large number of world-class measurements and dominance in heavy flavour physics [6], in addition to a reputation of an excellent forward general purpose detector at the LHC.

The Long Shutdown I (2013-2014) will allow LHCb to fully explore the large statistics collected and prepare LHCb for Run II (2015 – 2017) at an energy close to the nominal 14 TeV.

An upgrade of the LHCb experiment [7] is underway for the Long Shutdown II in 2018. With the need for reconstructing the event topology in order to efficiently trigger on leptonic and hadronic decays of beauty and charm hadrons, the 1 MHz readout limit is currently the main bottle neck to run at higher luminosity and with higher trigger efficiencies. The upgrade consists of a complete readout at the LHC bunch crossing rate (40 MHz) with only a software trigger and running at an instantaneous luminosity of up to $2x10^{33}$ $cm^{-2}s^{-1}$. Several sub-detectors upgrades are also underway to cope with the higher occupancies and provide a fast reconstruction.

ACKNOWLEDGMENT

LHCb would like to thank the LHC warmly for excellent performance and excellent collaboration.

REFERENCES

[1] The LHCb Collaboration, "The LHCb Detector at the LHC", *JINST* 3 S08005 (2008).

[2] R. Jacobsson, R. Alemany, F. Follin, "LHCb Luminosity Control and Monitoring" unpublished.

[3] P.R. Barbosa-Marinho et al., "LHCb Online system technical design report: Data acquisition and experiment control", CERN-LHCC-2001-040.

[4] R. Jacobsson et al., "The LHCb readout system and real-time event management", IEEE Trans. Nucl. Sci., vol. 57 (2010) 663.

[5] R. Aaij, "The LHCb Trigger and its Performance", arXiv:1211.3055v1 [hep-ex] (November 2012).

[6] The LHCb Collaboration, "Implications of LHCb measurements and future prospects", LHCb-PUB-2012-009, July 2012.

[7] LHCb Collaboration, "Letter of Intent for the LHCb Upgrade", CERN-LHCC-2011-001, March 2011.