PROCEEDINGS OF SCIENCE

PoS(FPCP2015)048

The LHCb upgrade

U. Marconi[*]
INFN Bologna (Italy)
E-mail: umberto.marconi@bo.infn.it

Profiting from a data sample of over 3 fb^{-1}, collected during the LHC Run 1, exploiting proton-proton collisions at the centre-of-mass energies of 7 TeV and 8 TeV, LHCb has successfully performed a large number of world-class precision measurements in heavy flavour physics. However, many of the LHCb measurements will remain limited by statistics, even though adding an integrated luminosity of $5-6$ fb^{-1} at 14 TeV, expected in the LHC Run 2. The main obstacle preventing LHCb to run the detector at higher luminosities, with enhanced trigger efficiencies, is due to the current 1 MHz readout system limitation. LHCb will therefore undergo a major upgrade in the Long Shutdown 2 ($2018-2019$) aimed at collecting an order of magnitude more data by 2028. The upgrade consists of a new full readout operating at the LHC bunch crossing rate of 40 MHz, with a data acquisition system with the ultimate flexibility of a software trigger. The detector will be run at instantaneous luminosity of 2×10^{33} cm^{-2}s^{-1}, five times higher than presently. In order to cope with the higher occupancies and radiation doses several sub-detector upgrades are also underway.

Flavor Physics & CP Violation 2015,
May 25-29, 2015
Nagoya, Japan

---

[*] Speaker.

1. Introduction

The understanding of flavour dynamics is considered one of the most important aims of particle physics. The last 15 years have witnessed the triumph of the Kobayashi-Maskawa (KM) mechanism, formalised in the quark mixing matrix, which describes the flavour changing transitions of quarks in the SM. This has been achieved owing to several experiments, notably including those operating at the so called B-factories (BaBar and Belle), at the Tevatron (CDF and D0), and now at the LHC (ATLAS, CMS and LHCb). LHCb has produced a plethora of results, on a broad range of flavour observables, and ATLAS and CMS have given significant contributions to the beauty sector, mainly using final states containing muon pairs. The measurements performed confirm that the level of CP violation in the SM is not large enough to explain the matter-antimatter asymmetry of the universe. This is indeed considered as a compelling evidence that other sources of CP violation BSM must exist.

LHCb can profit from a number of powerful key features. The $b\bar{b}$ production cross-section at the collision energy of $\sqrt{s} = 14$ TeV is expected to be around 500 μb. Running at the instantaneous luminosity of 4×10^{32} cm^{-2}s^{-1} it implies a rate of about 200 kHz of b-events at the LHCb interaction point. The production mechanism gives access to all the b-flavoured hadrons. Moreover the cross-sections for $c\bar{c}$ production is approximately 20 times larger than of $b\bar{b}$. The hadronic states containing quarks of the $b\bar{b}$ and $c\bar{c}$ pairs appear in the same hemisphere and with large boost. This allows a very good proper time resolution (for resolving, for instance, the fast B_s oscillations), and it makes possible the flavour tagging with a single arm spectrometer.

The LHCb detector acceptance spans the polar angles between 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 range $2 < \eta < 5$. Although the acceptance corresponds to only 4% of the solid angle, it includes about 40% of the $b\bar{b}$ pair production cross-section. The full detector coverage, in the complete pseudo-rapidity range, includes tracking, vertex reconstruction, particle identification and calorimetry [1].

The current detector has been efficiently operated at four times the design foreseen pileup. In addition, in order to satisfy an extended physics programme the physics output rate was increased from 2 kHz in 2010 to 5 kHz in 2012. In these conditions the LHCb sub-detectors have performed equivalently or even better than design. The VELO vertex detector has achieved an impact parameter resolution of about 20 μm for tracks with high transverse momenta (anyhow below 10 GeV/c in the LHCb acceptance). The momentum resolution ranges from 0.4% at 5 GeV/c, up to 0.6% at 100 GeV/c. The kaon identification is about 95% with only 5% misidentification of pions. Muons are identified with 97% efficiency with only $1-3\%$ of mis-identification. The energy resolution of the electromagnetic calorimeter has a sampling term less than 10% with a 1% constant term.

The present LHCb trigger operates in two levels. The first level consists of a low-latency fully-pipelined hardware trigger, acting within 4 μs. It uses simple events signatures, as transverse energy and momentum of electrons, photons, hadrons and muons, to reduce the readout rate to about 1 MHz. The LHCb visible event rate of about 12 MHz is reduced to 1 MHz by thresholds of about 1 GeV/c for muons and $3-4$ GeV/c for electrons, photons and hadrons. The available bandwidth to the HLT is shared between 150 kHz of electron and photon triggers, 450 kHz of hadron triggers and 400 kHz of muon triggers. The the High Level Trigger (HLT) has access to the full detector information and performs exclusive and inclusive event selections by means of order of 50,000 copies of the trigger algorithm running on 2,500 CPU nodes. Operating at $\sqrt{s} = 8$ TeV, with an output rate of 5 kHz, the combined trigger efficiency is of 90% for selecting B decays to muons, 30% for selecting B decays to hadrons, and 10% for charm decays (efficiencies are referred here to events that will be useful to the offline analysis).

LHCb has remarkably pioneered several operational developments during Run 1, aiming at maximising the physics yield and reducing the systematics. A real-time control system has been put in place that allows LHCb to run at a constant and optimised luminosity throughout the fills. It implies that a single trigger configuration and pre-scales settings can be used permanently. It leads to stable detector performance over time and easier calibrations, which in itself reduces systematics.

Fraction of the events accepted at the first level trigger are temporarily stored into the local disks of the HLT farm. This way the HLT processing can be partially deferred, to take place afterwards during the inter-fill time. With this technique the HLT CPU effective capacity has been increased by more than 20%. The physics results of Run 1 relies on these features, that will be exploited also in the future. In particular, in order to optimally exploit the farm capability, doubling the effective CPU capacity of the online farm, in RUN 2 the HLT will act in two steps: HLT-1 will select and store events at 100 kHz on local disks; then HLT-2 will perform the deferred final event selection at 10 kHz.

In addition to the routine operation LHCb has taken many occasions to explore the detector response at increasing luminosity and higher pileup, up to the foreseen upgrade luminosity. The results show that the S/B is independent of the pileup up to $5-6$ primary vertices, since of the separation between them is of the order of centimetres, while typical resolutions of primary and secondary vertices are about 60 μm and 200 μm respectively.

With its upgrade LHCb aims at reaching experimental sensitivities comparable to the uncertainties affecting the theoretical predictions [2]. Experience with the detector operation and the analysis of the data sets from Run 1 indicates that systematics errors can be managed and they won't spoil the final precisions. Table 1 shows the expected statistical precisions achievable both in 2018 and with the upgraded detector, for a number of representative physics modes, to be compared with the current theoretical uncertainties [3].

2. Trigger and detector upgrade

The LHCb upgrade is scheduled for 2018 and consists of a complete redesign of the readout and the trigger systems. Figure 1 shows how the trigger efficiencies vary as a function of the instantaneous luminosity for some reference channels. While decay modes with muons would be handled efficiently with the present detector also at higher luminosities, the hadronic modes saturate, due to the first level trigger bandwidth limitation. An efficient trigger requires the removal of the first level trigger and the adoption of a prompt event reconstruction, based on the the whole detector information. It can be achieved by removing the hardware trigger bottleneck and reading out all the sub-detectors at 40 MHz, performing solely a software triggering on a dedicated CPU farm.

Table 1: Expected statistical sensitivities after LHC Run 2 and after 50 fb^{-1} with the LHCb Upgrade as compared to current theoretical uncertainties for a list of key observables .

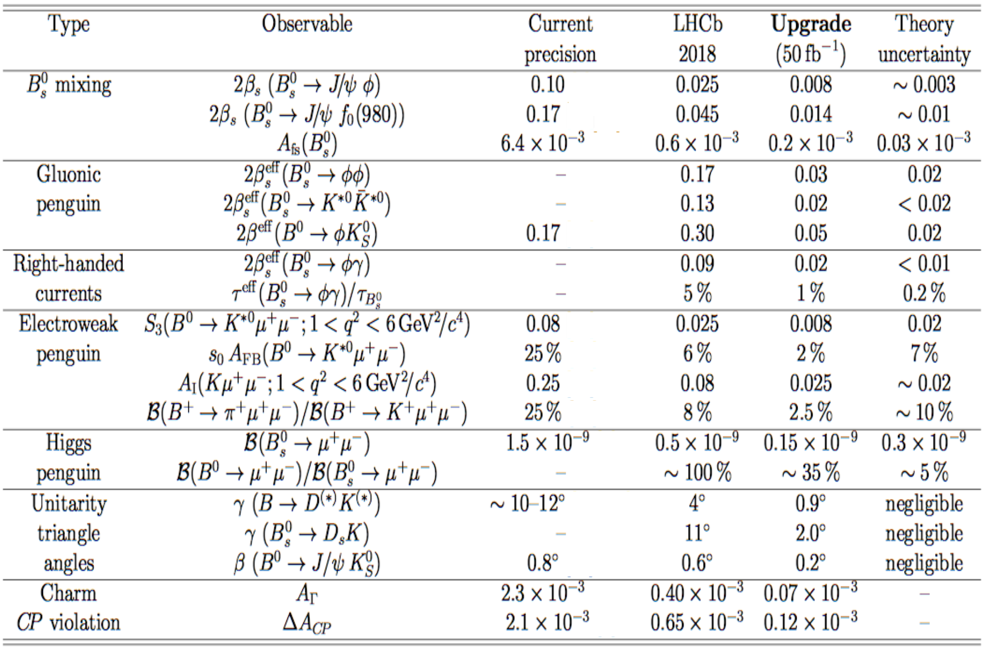

Type	Observable	Current precision	LHCb 2018	**Upgrade** (50 fb^{-1})	Theory uncertainty
B_s^0 mixing	$2\beta_s\ (B_s^0 \to J/\psi\ \phi)$	0.10	0.025	0.008	~ 0.003
	$2\beta_s\ (B_s^0 \to J/\psi\ f_0(980))$	0.17	0.045	0.014	~ 0.01
	$A_{\rm fs}(B_s^0)$	6.4×10^{-3}	0.6×10^{-3}	0.2×10^{-3}	0.03×10^{-3}
Gluonic penguin	$2\beta_s^{\rm eff}(B_s^0 \to \phi\phi)$	–	0.17	0.03	0.02
	$2\beta_s^{\rm eff}(B_s^0 \to K^{*0}\bar{K}^{*0})$	–	0.13	0.02	< 0.02
	$2\beta^{\rm eff}(B^0 \to \phi K_S^0)$	0.17	0.30	0.05	0.02
Right-handed currents	$2\beta_s^{\rm eff}(B_s^0 \to \phi\gamma)$	–	0.09	0.02	< 0.01
	$\tau^{\rm eff}(B_s^0 \to \phi\gamma)/\tau_{B_s^0}$	–	5 %	1 %	0.2 %
Electroweak penguin	$S_3(B^0 \to K^{*0}\mu^+\mu^-; 1 < q^2 < 6\,{\rm GeV}^2/c^4)$	0.08	0.025	0.008	0.02
	$s_0\ A_{\rm FB}(B^0 \to K^{*0}\mu^+\mu^-)$	25 %	6 %	2 %	7%
	$A_{\rm I}(K\mu^+\mu^-; 1 < q^2 < 6\,{\rm GeV}^2/c^4)$	0.25	0.08	0.025	~ 0.02
	$\mathcal{B}(B^+ \to \pi^+\mu^+\mu^-)/\mathcal{B}(B^+ \to K^+\mu^+\mu^-)$	25 %	8 %	2.5 %	~ 10 %
Higgs penguin	$\mathcal{B}(B_s^0 \to \mu^+\mu^-)$	1.5×10^{-9}	0.5×10^{-9}	0.15×10^{-9}	0.3×10^{-9}
	$\mathcal{B}(B^0 \to \mu^+\mu^-)/\mathcal{B}(B_s^0 \to \mu^+\mu^-)$	–	~ 100 %	~ 35 %	~ 5 %
Unitarity triangle angles	$\gamma\ (B \to D^{(*)}K^{(*)})$	~ 10–$12°$	4°	0.9°	negligible
	$\gamma\ (B_s^0 \to D_sK)$	–	11°	2.0°	negligible
	$\beta\ (B^0 \to J/\psi\ K_S^0)$	0.8°	0.6°	0.2°	negligible
Charm *CP* violation	A_Γ	2.3×10^{-3}	0.40×10^{-3}	0.07×10^{-3}	–
	ΔA_{CP}	2.1×10^{-3}	0.65×10^{-3}	0.12×10^{-3}	–

The consequence of the 40 MHz readout requirement is that all the sub-detectors readout electronics must be rebuilt. Sub-detectors upgrades are also needed in order to maintain the same performance as now, while coping with significantly higher pileup and occupancy.

The main detector replacement concerns the tracking system [4], involving the VErtex LOcator (VELO) and the tracking stations. They will be completely replaced. The VELO will be based on pixel technology to provide significantly higher granularity than now. The plan is to use planar silicon sensors 55 × 55 μm^2, incorporating 256 × 256 pixels. In order to improve the impact parameter resolution the sensor distance to the beam will be reduced by 2 mm. The inner radius of the RF foil will be about 3.5 mm from the beam instead of 5.5 mm, with reduced thickness. The upstream TT tracker (renamed UT) will be rebuilt using silicon strips as presently, but with higher segmentation and improved sensor overlap. It will be placed closer to the beam pipe to improve the small-angle coverage. The VELO-UT system shall allow fast tracks momentum measurements, essential feature to perform fast forward tracking and discard fake tracks. The downstream tracker (T1,T2 and T3) will be entirely replaced with large surface of scintillating fibre trackers. The fibre trackers will use 250 μm diameter fibres, read out with Silicon Photo-Multipliers (SiPMs). The expected spatial resolution is about 60 − 100 μm.

Due to the high occupancy the scintillating pad detector (SPD) and the calorimeter pre-shower (used for e/γ separation) will be also removed. This is expected to improve the resolution of the electromagnetic calorimeter and ease the calibration. The e/γ separation will be done in the

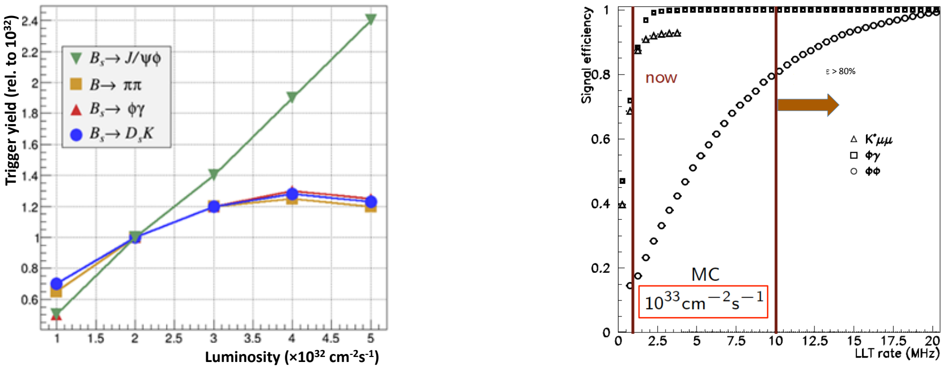

(a) Event yield as a function of luminosity for different decay modes with the 1 MHz bandwidth limitation.

(b) Signal efficiency as a function of the first level trigger bandwidth for different decay modes.

Figure 1: Effects of the 1 MHz trigger bandwidth limitation on the trigger efficiency.

software trigger with the full information available. The calorimeter will reduce the gain of the photomultipliers and compensate with an increased electronics gain. The hybrid photo-multipliers (HPD) of the two RICHes, will be replaced with multi-anode photo multipliers. The RICH1 optics will be re-optimised. Its Aerogel radiator will be removed, since it gives too few photons to actually allow reconstructing the rings in the higher multiplicities of the upgrade . The first muon detector layer (M1) located in front of the calorimeter will be removed, cause of the high occupancy. The muon stations, located after the calorimeter, will remain the same [5].

New readout boards, named the PCIe40, will be used for data acquisition. The PCIe40 is a generic hardware component, designed for reading out all the sub-detectors, and to provide the distribution of the timing, commands and control signals [6]. The PCIe40 is based on the PCI Express Generation 3 standard protocol. They will be used to connect the detector front-end through 300 m long optical fibre directly to the motherboards of the event-builder servers. Embedding the PCIe40 readout boards into PC servers will conveniently allow exploiting up-to-date, high-speed, data-centre technology. The event-builder farm will require order of 500 servers to restore event fragments in the complete event at 40 MHz. The expected aggregate data throughput is about 4 TB/s. The logic scheme of the data acquisition system foreseen for the upgrade is shown in Figure 2.

3. Conclusions

LHCb will undergo a major upgrade between 2018 and 2019 to allow operating the experiment at a luminosity of 2×10^{33} cm^{-2}s^{-1}. The detector upgrade, performed to read out the full detector at the bunch crossing rate of 40 MHz, consists of a re-build of the sub-detectors and a complete redesign of the readout system. Events will be processed at the full speed by a versatile software trigger. The full software trigger allows to maximise the trigger efficiencies and to minimise systematic uncertainties in the selection of the interesting flavour decays. The upgraded readout and trigger system, as well as the detector optimisation, will allow the LHCb physics program and running conditions to adapt to any signature, which may come out as the consequence of a changing

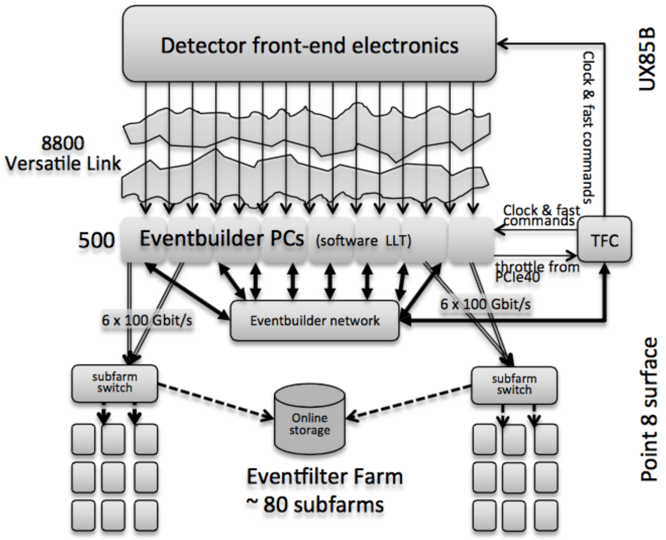

Figure 2: The new data acquisition system architecture for the LHCb upgrade.

physics scene. The plan is to collect a dataset of at least 50 fb^{-1} in less than ten years. It will allow LHCb to reach unprecedented precision in flavour physics.

References

[1] LHCb collaboration, A. A. Alves Jr. et al., The LHCb detector at the LHC, JINST 3, 2008, S08005.

[2] LHCb Collaboration, Letter of Intent for the LHCb Upgrade, CERN-LHCC-2011-001

[3] LHCb Collaboration, Framework TDR for the LHCb upgrade, CERN-LHCC-2012-007

[4] LHCb Collaboration, LHCb Tracker Upgrade Technical Design Report, CERN-LHCC-2014-001

[5] LHCb Collaboration, LHCb PID Upgrade Technical Design Report, CERN-LHCC-2013-022

[6] LHCb Collaboration, LHCb Trigger and Online Upgrade Technical Design Report, CERN-LHCC-2014-016