Status of the LHCb Experiment

Report to October 2010 RRB
by the LHCb Collaboration

1. Introduction

LHCb is successfully taking data with an operational efficiency higher than 99%. Much progress has been achieved in all areas of the experiment. Key performance parameters, such as tracking efficiency, detector alignment, impact parameter, proper lifetime and invariant mass resolution, are gradually reaching their Monte Carlo expectation values. Beauty and charm hadrons have been reconstructed in many different decay modes, in general with an efficiency consistent with the expectations from simulation, and with a good signal-to-background ratio. This gives confidence that LHCb will soon be able to deliver flavour physics results of the highest scientific interest. Indeed, first measurements have been already reported at conferences and published in scientific journals.

In order to maximize the integrated luminosity and physics reach for the 2010 physics run, the experiment is running at $\beta^* = 3.5$ m, which is significantly less than the LHCb design-value of 10 m. Consequently the average number of interactions per beam crossing reaches up to 2, to be compared with the LHCb design value of only 0.4. In order to cope with these extreme conditions a lot of effort has been invested in improving data processing and understanding the analysis of data with multiple interactions per beam crossing.

LHCb wishes to collect as much data as possible in 2010, since an integrated luminosity of only 50 pb^{-1} would allow the experiment to approach the existing world-best precision in the measurements of B$_s\rightarrow\mu\mu$ decay, the CP-violating phase in B$_s$ oscillations, as well as in CP-violation in the charm sector, all topics which have high sensitivity to New Physics contributions.

2. Detector Subsystems

2.1 Vertex Locator (VELO)

All VELO subsystems have performed well throughout the physics run. A single hit resolution of 4 μm has been achieved at optimal track angle, and a primary vertex resolution of 14 μm in X and Y, and 76 μm in Z. The system is aligned at an accuracy of a few μm. The system currently has 0.7% of inactive strips compared with design, a number which may further reduce with calibration improvements. The VELO closing after declaring stable beams and opening before the beam dump has been optimised throughout the operations for detector safety and efficiency; the closing now routinely takes around five minutes.

Changes: A change of the firmware in the data processing boards was required to cope with the higher data rates delivered by the LHC. This was successfully performed in mid-August.

Concerns: The insulation of the pipes of the CO_2 cooling system is defective, leading to ice formation at the joints of the pieces of insulation. The insulation of the full system will be replaced during the winter shutdown.

Further improvements to the firmware of the data processing boards will be required to fully optimise the system performance.

Plans: Given the special nature of the VELO system with its required high accuracy, proximity to the LHC beam, operation in vacuum, and motion system, the system has been operated with 24 hour VELO shift coverage, and two continuous on-call experts this year. While noting that the operation of this complex system will always require significant person effort, preparations for streamlining of operation in 2011 is underway. A review of the performance of relevant areas will occur during the Technical Stop in winter 2010/11. This will include a review of the closing procedure, with a representative from the LHC beam protection community present.

The replacement of the VELO mechanics and detectors are either nearing completion or completed and under test. Work has started on obtaining the required infrastructure for assembly and test of the replacement VELO modules in the detector halves, utilising elements from the assembly and spares of the current detector. The assembly of the system is expected to start early next year.

2.2 Silicon Tracker (ST)

Both the Inner Tracker (IT) and Trigger Tracker (TT) run smoothly and without problems. Day-to-day operation in hands of the central shift crew is supported by an on-call piquet. To ensure safe and efficient data taking, detailed monitoring of the detector status and detector quality is provided. If problems develop alarms are generated and appropriate automatic actions are taken. More than 99.4 % of channels in the Trigger Tracker and 98.4 % of channels in the Inner Tracker are operational.

The non-working channels have several causes. First, three (one) modules in IT (TT) have problems located inside the detector. For TT these problems will be fixed on the spot whilst in the case of IT the service boxes will be dismounted and the effected modules fixed during the 2010/11 Technical stop. The remaining faults are dominated by failing VCSEL diodes located on the digitizer boards in the detector service boxes. For TT these problems are typically fixed during short detector stops whilst for IT it is planned to exchange the VCSELs during the Christmas break.

Alignment and calibration of the detector are progressing well. The S/N of the detector is good and within 20% of the expectation obtained from test-beam studies. Cluster resolutions of 55 microns are found for both detectors, IT and TT. These values are consistent with the expected binary resolution of the detector. Further work is needed to understand discrepancies between the alignment obtained with magnet off and magnet on data.

Changes: None

Concerns: In recent months several VCSEL diodes located on the digitizer boards in the detector service boxes have failed. Significant discrepancies remain between global alignment parameters obtained with the magnet on and off.

Plans: Further understanding of the VCSEL failures. Setup of test-stands to allow repair and testing of IT and TT modules during the 2010/2011 Technical stop. Further work to understand the alignment of the detector.

2.3 Outer Tracker (OT)

The OT detector is routinely taking data since the beginning of LHC operation in 2010. The detector and readout electronics operate reliably and the amount of dead channels is below 1%. Four out of the 435 VCSEL transmitters on the on-detector serializer boards died; all were successfully replaced and all Front-Electronics spares are repaired and available at CERN.

Calibration parameters, such as noisy channels, drift time-space T(r) relation, time offset and space alignment, have been determined, transferred to the appropriate databases, and found to be stable during detector operation in 2010. Correspondingly, Monte Carlo simulation has been updated (mainly T(r) relation and single-cell efficiency) in order to accurately describe the detector performance. The spatial resolution determined with data (about 270 micrometers for tracks with momentum higher than 10 GeV/c) is worse than the one predicted by Monte Carlo simulation (about 200 micrometers), but at present this is not limiting the physics performance. The discrepancy will further be reduced with refinements in the space alignment and possibly in the dead material description.

In 2010 the detector was operated with 1.5% of molecular oxygen added to the nominal Ar/CO_2 gas mixture, in order to reduce the effect of aging; further irradiation in situ during short accesses confirmed the beneficial effects already observed in laboratory tests. A large analysis effort has been made to produce tools for a timely determination of the onset of aging effects, based on the usage of tracking to "map" the stability of the detector amplification in horizontal and vertical bins (aging effects will be highly dependent on the large variations in beam intensity on the detector surface). Since effects on tracking performance will only start to be visible for very large gain deterioration, we complemented the method by performing threshold (or HV) scans. These procedures, taking approximately 1 hour of beam time, have been demonstrated to be sensitive to gain variations of the order of 10%, and are thus a powerful technique for a timely detection of aging effects. At the same time, a possibility of curing existing gain losses through a HV training procedure has also been demonstrated in situ and a remotely-controlled procedure is ready to be tested during 2010/2011 Technical stop.

Changes: None.

Concerns: Uncertainty in the long term behaviour of the gain loss remains the main concern, most notably the differences in the behaviour of different modules.

Plans: Since the analysis software as well as the data taking procedure is ready and tested, we intend to request periodic (about once every 2 months) threshold/HV scans in order to

measure 2D gain maps and monitor aging effects. These will be complemented by scans with a radioactive source performed in situ during short accesses.

2.4 RICH

Both RICH1 and RICH2 detectors are extremely stable and taking data efficiently. Beam collision data have been used to align the system; the angular resolution from the gaseous radiator is approaching the design value. The angular resolution with the Aerogel is also improving steadily, and work is continuing to better understand the calibration pattern.

The evolution of each HPD in terms of its ion feedback (IFB) is continuously monitored, thus being a reliable indicator of the HPD lifetime. During the 2010/2011 Technical stop those tubes with IFB already above 5% , or expected to exceed 5% in 2011 or 2012, will be replaced, in both RICH1 and RICH2. This will make the running in 2011 stable. All the repaired tubes that have been installed since the first intervention in March 2009 continue to behave satisfactorily. In RICH1, three VCSEL diodes are dead, and consequently are not transmitting HPD data. In RICH2 two VCSELs are dead. Their replacement (VCSELs or the entire L0 boards on which they are mounted) will be performed in December 2010.

Still a few residual discharges have been observed in the HV power supplies. New improved units with one output cable and a splitter made in-house have been thoroughly tested in the laboratory. They show a much better stability. One such unit has already been installed on the detector and is operating optimally. More of these improved units will be installed during the next Christmas break to take out even the few that show residual discharges.

The pressure and temperature of the gaseous radiators inside the RICH1 and RICH2 vessels are monitored continuously and the gas purity is excellent and stable (> 95% and >99% respectively).

The performance of the online and data quality monitoring are continuously improving, and the RICH reconstruction software is providing excellent PID for many of the physics analyses ongoing in LHCb.

Changes: None.

Concerns: The stock of repaired HPDs is limited and will not allow the replacement of the least degraded tubes in RICH2. This will however have only marginal effect on the performance, as the central regions with high occupancy will be equipped with good or refurbished tubes.

Plans: Collecting more collision data will help in understanding and further improving of the detector performance.

2.5 Calorimeters

During the start of 2010 data taking the inter-cell time alignment work was pursued and time alignment with a precision below 1ns for ECAL and HCAL and 1.5ns for PS and SPD was achieved. With this precision the reconstruction of π^0 and η was performed easily from their

photon pair decay; the reconstructed mass and width were matching well with the expectations given by the initial inter-cell calibration.

For the early 2010 data the emphasis had been put on the absolute calibration of each calorimeter cell using the following procedure:

- ECAL, HCAL, PS cell inter-calibration using the smoothing of the energy flow in each detector that leads to an inter-calibration accuracy of respectively 4%, 5% and 8% for the 3 detector elements.
- Inter-cell calibration of pre-shower cell using charged particle minimum ionization, giving a 5% inter-cell calibration accuracy.
- Fine calibration using π^0 from two separated photons, with 80 Million events giving an absolute calibration for the electromagnetic calorimeter below 2%.

The SPD threshold calibration has been adjusted using the charged particle response to a threshold scan.

With the quality of calibration that had been reached in May 2010, π^0 and η mesons as well as D decays involving π^0 in the final state and $J/\psi \rightarrow e^+e^-$ decays have been reconstructed with a better resolution than expected.

The gain stability is permanently monitored by the LED system to a 1% level during data taking. Since April 2010 periodical calibration runs with a radioactive source confirm the stability behaviour of the Hadronic Calorimeter.

Using online data, the monitoring has improved considerably towards an automatic alarm processing. The system surveys pedestal, noise, and gain stability for the PMTs. More than 99.9% of the calorimeter detector channels are operational.

Changes: None

Concerns: The work done on calibration with low luminosity has to be adapted for the high luminosity we are facing since August. The importance of the pile up effect has to be understood, and work to monitor the pile up on an event-to-event basis is in progress.

Plans: Pursue calibration task for photons and electrons, to reach a 1% resolution for high energy.

2.6 Muon system

The Muon Detector has been fully commissioned using the first data delivered by the LHC. The operating voltage of the chambers has been optimised in order to maximize the efficiency and to minimize the aging effects. A full time and threshold calibration has been performed in this new operating condition in order to have the best possible working conditions for the LHC run. The behaviour of the detector is very stable since the beginning of the LHC physics run. The amount of dead channels is of the order of 0.1% and the noise is negligible. The efficiency measured with data is above 99% in all the detector stations. With increasing beam intensity some high voltage instability has been observed, leading to

repeated trips of a few chamber gaps. In total, about 0.5% of the gas gaps are affected by this phenomenon. Due to the fourfold redundancy of the LHCb Muon chambers, this HV instability does not affect the overall detector efficiency; this anomalous behaviour is in any case actively monitored and actions are being taken to solve the problem by reconditioning the chambers.

The Muon control system has been improved resulting in a better monitoring of the hardware configuration. Online and offline data quality monitoring tools have been refined allowing a quick feedback on problems affecting the readout electronics and the chamber performance.

Changes: None.

Concerns: HV instability in a small number of chambers in high beam intensity conditions.

Plans: Replace and/or recondition a few chambers with HV problems; improve recovery procedures after HV trips during the 2010/2011 Technical stop.

2.7 Online System

The LHCb Online system was operated very successfully during LHC operation. Overall efficiencies were around 90% including VELO closing and high-voltage ramping procedures. The High-Voltage control is now centralized and operated regularly by the shift crew.

Problems in the readout network have been resolved and no more errors have been observed. The purchasing procedure of the last tranche of the HLT farm is completed and the purchase is well under way. We expect the farm to be gradually upgraded in the coming few weeks, as CPUs are delivered. Total CPU power of the farm will be increased by approximately 400%.

Changes: None.

Concerns: Long-term manpower coverage.

Plans: Consolidation of the system and preparation of work plan for the 2010/2011 Technical stop. Improve efficiency wherever possible.

3. Trigger

The LHCb trigger is fully operational and now running with thresholds close to our nominal conditions due to the high multiplicity events delivered by the LHC. As only 20% of the event filter farm (EFF) is available the main limitation is the CPU used in the High Level Trigger (HLT). Strategies have been developed to minimize this CPU usage, in particular the use of global event cuts to keep the event size and the occupancy acceptable.

The Level-0 trigger is fully operational and has been run with constantly increasing thresholds as the luminosity increased. Because of the CPU limitations of the EFF we have not yet been able to saturate the L0 bandwidth. The multiplicity in the SPD detector is now

used as a filter to reject too busy events. This considerably reduces the CPU consumption in the HLT while affecting the physics output very little.

The HLT has been commissioned with all sub-detectors except RICH participating in the decision. Thanks to the good performance of the VELO and tracking detectors it was found that we can operate with offline quality tracking, which was not anticipated before. The HLT has been re-optimized for speed while keeping a high efficiency. Several trigger configurations have been developed and tested offline before being run in the EFF. For the present data taking period we have four scenarios ready that have been tested in the EFF and are used in turn depending on the instantaneous luminosity. The fast change from one configuration to the next is now used routinely without any loss of data.

The main difference between the configurations is the efficiency on charm decay reconstruction that decreases rapidly with increasing thresholds. Our core B physics potential is very little affected.

In parallel to the physics triggers we run rate-limited minimum bias and beam-gas triggers, which allow a determination of the luminosity and of trigger efficiencies. The total output rate varies between 0.3 and 2 kHz depending on luminosity.

After installation of the remaining CPUs during the 2010/2011 technical stop we will be able to run at lower thresholds during the start of next year. The HLT algorithms will be re-optimised based on the experience acquired with the 2010 data. We anticipate that we will then be able to profit from the full L0 bandwidth.

Changes: Retuning of HLT due to high multiplicity.

Concerns: CPU limitations at high luminosity.

Plans: Installation of remaining CPUs in EFF. Commissioning of full bandwidth in L0.

4. Computing

Between March 30[th] and August 2010, LHCb collected 95.5 TB of raw data with high efficiency. In accordance with the computing model, the RAW data were distributed to the 6 Tier-1s where reconstruction and stripping has all been carried out.

The running conditions have been challenging. The reconstruction application had to be adapted to deal with an average of 1.5 visible collisions per crossing instead of 0.4. This has led to larger event sizes. The very high occupancy in detectors required special measures to be taken in the pattern recognition and particle identification algorithms. Due to the fast evolution of the machine luminosity and the "learning period" with real data we have performed several reconstruction and stripping cycles, more than foreseen in the steady state. The retention rate for stripping is still a factor of 2 above the canonical 10% primarily due to the decision to retain additional events for charm and cross section measurements with the early data, but we are working to reduce this in the longer term.

User analysis jobs have also all been run at the Tier-1 sites. The split of CPU time used in the last 6 months is approximately 50% production and 50% users. Monte Carlo simulation continues mainly at Tier-2 sites. During a particularly intense period in June (prior to the summer conferences) there were 25,000 jobs running simultaneously.

All in all, conditions have resulted in longer job execution times for the high occupancy events, and some pressure on disk space, but both have been managed within the existing CPU and disk resources.

Changes: none

Concerns: We continue to have long-term manpower shortages, resulting in delays to developments and to critical dependency upon individuals.

Plans: Adapt computing model to new running conditions with increased pile-up.

We plan to reprocess the full 2010 dataset starting immediately after the end of 2010 data taking. A corresponding Monte Carlo production will be launched. These datasets will be used for analyses to be presented at the 2011 winter conferences. All previous reprocessing will be removed from disk, to make room for the 2011 data.

We plan to introduce the use of micro-DSTs as output format of the stripping for channels requiring a high retention rate. This, combined with tighter selection cuts, should allow us to achieve the 10% stripping retention rate foreseen in the computing model (currently this rate is 20%).

We expect to pilot the use of T2 centres for data reconstruction.

We hope to develop additional monitoring tools to aid those performing computing operations shifts.

5. Reconstruction performance and physics studies

Good progress has been achieved on tracker alignment using data with reconstructed muons from J/ψ decays. The impact parameter resolution for high p_t tracks is now close to the Monte Carlo expectation. A small discrepancy between data and simulation is still observed at low p_t which seems to be related to an underestimation of the material in the Monte Carlo. Studies are underway with data to map out the amount and distribution of material by reconstructing hadronic secondary interactions in the VELO. The most recent alignment has also significantly improved the invariant mass resolution, in particular for D and B mesons and for the Y resonances, which are 50% narrower than was the case in early summer, and are now very close in width to Monte Carlo expectation. A small offset remains in the absolute mass scale, which could be explained by a ~0.5 per mille bias in the magnetic field calibration. However, this imperfection has no consequences for the physics topics under study, and can be corrected in the end-of-year reprocessing.

The performance of the particle identification is continuously monitored using clean signals of reconstructed π^0, η, K_S, Λ, ϕ and D* decays. Improvements in the RICH alignment and knowledge of the sensor positions have resulted in Cherenkov angle resolutions very close to Monte Carlo expectations. The muon identification has a performance in data which is essentially identical to that of the simulation.

To take full advantage of the rapidly evolving LHC luminosity, effort has been invested in understanding how the trigger, reconstruction and analysis can perform efficiently with pile-up events, a scenario beyond the design specifications of the experiment. This work has been

largely successful – for example, a fake rate is achieved in the muon identification which is rather independent of track multiplicity. The CPU performance of the reconstruction and stripping algorithms has been optimised to cope with high multiplicity events without compromising the performance.

LHCb is now producing important and wide ranging physics results. The experiment has performed the first measurements of the open beauty and charm cross-sections at LHC energies. These results have validated the theoretical predictions and confirmed one of the assumptions that underpin the expected physics reach of the experiment. Similarly important production studies are being conducted with charmonia.

The forward geometry of the spectrometer gives many observables in electroweak and minimum bias physics an interest that is highly complementary to those of the central detectors, and studies are advanced in many of these topics, with several results already having been presented.

The data sample now accumulated is sufficient to observe significant signal peaks in a large fraction of the key decay modes of the LHCb flavour-physics programme. For example, hadronic B-decay modes with branching ratios of 10^{-6} are now observed. The event yields are similar to those expected from simulation, showing that the trigger and tracking are performing as expected, and the background level is satisfactory in all cases. An integrated luminosity of ~50 pb^{-1} will allow for several significant flavour-physics measurements to be performed already with the 2010 dataset.

6. Preparation for LHCb upgrade

The running conditions of recent months have had the unexpected benefit of allowing us to progress in our upgrade studies using real data. Since the current bunch intensities are larger than originally foreseen, we can use the present data to study the effects of multiple interactions per crossing, which will be one of the key challenges for the upgraded detector running in the 10^{33} cm^{-2} s^{-1} regime. Understanding the detector performance with actual data is of great benefit in deciding the optimum upgrade strategy and the critical pieces of hardware that need to be upgraded.

We have already decided that having a flexible trigger strategy will allow us unprecedented access not only to new B and D decay physics that it will be necessary to pursue, as a natural second stage programme to our first set of 'critical flavour measurements', but also will allow us to investigate other physics in the very forward direction. This may include exotic particle searches, such as hidden valley Higgs decays, or exotic QCD phenomena, such as rapidity gaps as postulated by Bjorken. Thus the key element of the Upgrade is to replace all the detector electronics so that the system can be readout at the 40 MHz beam crossing rate and the trigger decision can be executed in software alone. The key element in this upgraded trigger is the VErtex LOcator because it provides track impact parameters and vertices. The Upgraded VELO will be based on the Timepix2 chip under development at CERN.

Preparations have started also from the organizational and resource sides. In order to speed up data processing and exploit running at significantly higher than nominal luminosity, the

Swiss groups are requesting a contribution at the level of 0.6 MCHF to increase the CPU power of the computing farm.

A Letter of Intent is being prepared for submission.

7. Cost and funding issues

The status of the accounts healthy and there is no cash flow problem foreseen. This has been the first year of continuous data taking. Nevertheless, based on extrapolations of the spending profile until end of August, we expect that the M&O Category A budget forecast for this year is roughly correct.

8. Collaboration Issues

Ulrich Straumann (University of Zurich) has been re-elected for the second two years term of the Chair of the Collaboration Board. His second term ends in December 2012.

Pierluigi Campana (INFN, Frascati) has been elected as Spokesperson. He will start his three years mandate in summer 2011.

Guy Wilkinson (University of Oxford) has been re-elected for a second one year term as Physics Coordinator. His second term ends in December 2011

Rolf Lindner (CERN) has been appointed as the Technical Coordinator for a term of three years starting from July 1, 2010.