

LHC AVAILABILITY AND PERFORMANCE IN 2011

A.L. Macpherson, CERN, Geneva, Switzerland

Abstract

The LHC performance and its overall machine availability are discussed, as well as the factors that contributed to an excellent LHC performance in 2011.

INTRODUCTION

In 2011, the LHC had an excellent year, with delivered luminosity for proton-proton running at 3.5TeV per beam well in excess of the target of 1 fb^{-1} set prior to the 2011 run. Indeed, with both record totals for delivered luminosities in both proton-proton and lead-lead running, several special optics runs and a vigorous machine development program, 2011 exceeded all expectations in terms of machine performance.

The LHC Run for 2011 can be summarised as follows

- The 2011 LHC Run went from 19/02/2011 (end of Hardware commissioning) until 07/12/2011. A total 264 fills reached Stable Beams.
- For the LHC proton-proton run at 3.5TeV there were 219 fills that reached Stable Beams. (First stable Beams fill: 1613, last fill: 2267).
- For the LHC ions run there were 39 fills that reached Stable Beams. (First Stable Beams fill 2290: last fill 2351.)

LHC AVAILABILITY

In 2011 the beam energy for the proton-proton run was set at 3.5TeV per beam and 50ns bunch spacing. After a short period of beam scrubbing with 25ns beam, to suppress the electron cloud effects on the beam vacuum, the total beam intensity was ramped up by a steady increase in number of bunches, and by 92 days into the run the machine was operating at 1380 bunches per beam, with an average bunch intensity of 1.1×10^{11} protons and a normalised transverse beam emittance of $\sim 2.2 \mu\text{m}$. For the remainder of the proton-proton run the machine operated with 1380 bunches, and this we label as the luminosity production phase of the proton-proton run.

During the luminosity production phase, beam and optics improvements were made that permitted the machine to deliver $\sim 72\%$ of the total luminosity. These improvements were:

- Reduction of the normalised transverse emittance of the beam from $\sim 2.5 \mu\text{m}$ to $\sim 2 \mu\text{m}$ at the start of the luminosity production phase. Throughout the run, beam 2 consistently exhibited a degraded emittance with respect to beam 1, particularly in the vertical plane, but as the blowup was not prohibitive and as the source not understood, no corrective intervention could not be taken.

- The average bunch intensity was adiabatically creased by $\sim 27\%$ to a final value of $\sim 1.45 \times 10^{11}$ protons.
- Mid way through the luminosity production phase, the beam squeeze optics were changed so that the β^* at ATLAS and CMS was reduced from 1.5m to 1.0m (ALICE and LHCb remained at a β^* value of 10m and 3m respectively)

These changes are summarised in Figure 1 and 2


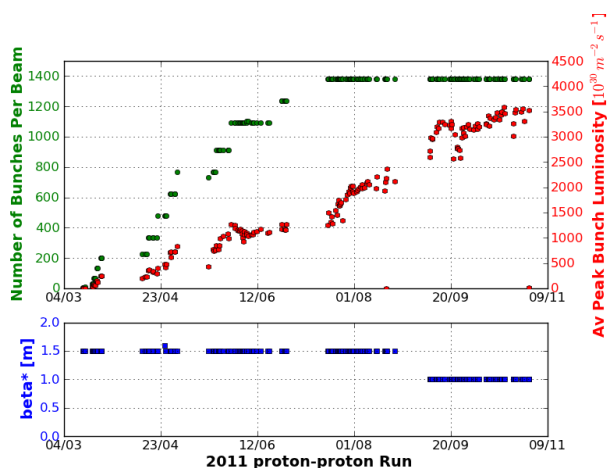


Figure 1: The evolution the average bunch intensity, number of injected bunches, and average normalised emittance (as determined from ATLAS and CMS peak luminosities at the declaration of Stable Beams) over the



duration of the 2011 proton run.

Figure 2: The evolution of the average bunch Luminosity over the 2011 proton Run. The lower plot shows the evolution of the optics squeeze parameter, β^* , over the 2011 proton run

DELIVERED LUMINOSITY

With these improvements, careful attention to the beam, and continuous optimisation of the machine, the achieved delivered luminosity was pushed well beyond expectations, both for the proton-proton and for the lead-lead run. While ATLAS and CMS pushed for the maximum delivered luminosity in the proton-proton run, LHCb and ALICE operated with luminosity levelling in order to optimise the data taking with respect to their detector capabilities. In addition, LHCb ran with a reduced squeeze of $\beta^* = 3\text{m}$ and ALICE was un-squeezed with a β^* of 10m. On switching to Ions, ATLAS, ALICE, and CMS all ran with a β^* of 1m while LHCb elected not to take data.

The evolution of the delivered luminosity for the proton and Ions Run is shown in Figure 3 and 4. These plots are based on preliminary off-line luminosity data published by the Experiments.

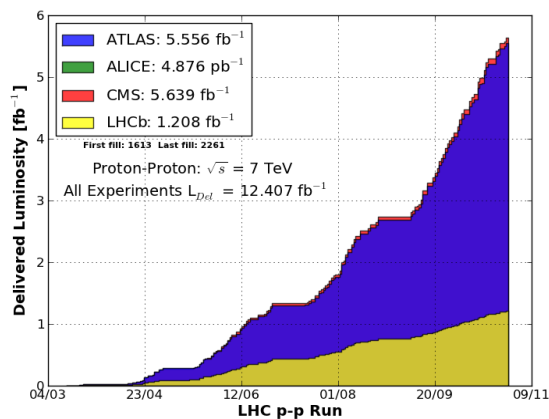


Figure 3: The evolution of the delivered luminosity over the 2011 proton-proton Run

Figure 4: The evolution of the delivered luminosity per experiment over the 2011 lead-lead Run. For the Ions running LHCb was not taking data.

MACHINE AVAILABILITY


In terms of LHC machine availability, the 2011 run was analysed in terms of six phases of operation that define the operational cycle. These phases are:

- **NB: No Beam.** The machine is prepared for access and there is no possibility for beam. This condition is defined by the state of the EIS Beam safety elements at LHC Points 3 and 7. Having both elements are off is a necessary condition for access into the LHC.
- **SU: Machine setup.** This is the time where the machine is closed but there is no beam in the machine.
- **INJ: Injection:** There is at least one circulating beam. The mode is defined by at least one of the Beam presence flags to be true.
- **RE: Ramp Energy.** This is the period of time that the beam is accelerated from injection to top energy, and is defined by the Beam mode. It also requires at least one of the Beam presence flags to be true.
- **SQ: Flat Top, Squeeze, and Adjust.** This phase is for the period of time at flat top energy after the ramp, for the β^* optics squeeze, and for the adjust period when the beams are put into collision (but before Stable beams is declared). Again, at least one of the Beam Presence flags has to be true is required.
- **SB: Stable Beams.** This is the period when the beams are in collision and the LHC experiments take data. This is defined by the Stable Beams beam mode but also requires both Beam presence flags to be true.

Note that the Setup time includes the time to setup the machine prior to injection, but also the ramp down and any downtime when the LHC is not prepared for access.

With these definitions, the machine availability can be assessed for the different operational periods. As an example, Figure 5 shows the availability for the Luminosity production period of the 2011 proton Run.

2011 Proton Run: Luminosity Production



SB Time: 26.6 days Total Time: 81.4 days




Figure 5: Machine performance during the 2011 luminosity production phase of the proton-proton Run.

Run	Days	NB %	SU %	INJ %	RE %	SQ %	SB %
2011	299.3	25.7	30.5	17.4	1.7	4.3	20.5
2011 no TS	277.9	23.3	29.5	18.7	1.9	4.7	22.0
p-p	156.6	22.0	20.4	19.2	2.2	3.8	33.8
P-P LP	81.4	23.6	19.3	18.9	2.0	3.5	32.6
Pb-Pb	24.1	25.0	20.8	13.6	2.2	5.5	32.9
MD	33.2	22.9	32.3	36.8	1.2	6.0	0.8

Run	Days	NB %	SU %	INJ %	RE %	SQ %	SB %
High β	4.2	6.2	43.7	10.3	3.2	35.4	1.1




Table 1: Machine Availability for the different Run epochs. Note 'p-p LP' refers to the luminosity production phase of the proton-proton Run. Also, data for all rows except the first do not count the periods allocated to technical stops (TS).

In terms of the machine performance over the different phases of the 2011 run, the availability for these different phases is summarised in Table 1, but it can be easily seen that the percentage of time spent in Stable Beams, both for proton and Ions running, is at the level of 33%. In terms of performance this can be evaluated in terms of the Huber factor H, which is given by

$$L_I = 0.00864 D L_P H$$

L_I is the integrated luminosity in fb^{-1} , L_P is the peak luminosity in $\mu\text{b/s}$, and D is the number of running days for physics assuming no unplanned stoppages.

Figure 6: Turn around time for fills during the 2011 Proton Run. The Luminosity production phase is defined as fills with 1380 bunches per beam injected.

The Hubner factor is the ratio of actual delivered luminosity to the amount you could collect by running continuously at the peak luminosity, and the expected value was $H=0.2$, (as achieved at LEP). For both the Luminosity production phase of the proton run, and for the Ions run, Hubner factors above 0.2 were achieved, with values of 0.22 and 0.24 respectively.

To achieve this sort of performance, attention was paid to an optimisation (whenever possible) of the turnaround. In terms of the proton Run the turnaround is summarised in Figure 6, which gives a most probably turnaround time for the 2011 proton run of 5hrs 13.8 min. It is noted that average turn around times are influenced by the long tail of Figure 6, which is due primarily to equipment faults and cryogenics recovery times.

By comparison, the duration of stable beams times is shown in Figure 7, where the average Stable Beams duration for the 2011 proton run was 5hrs 45.6 minutes.

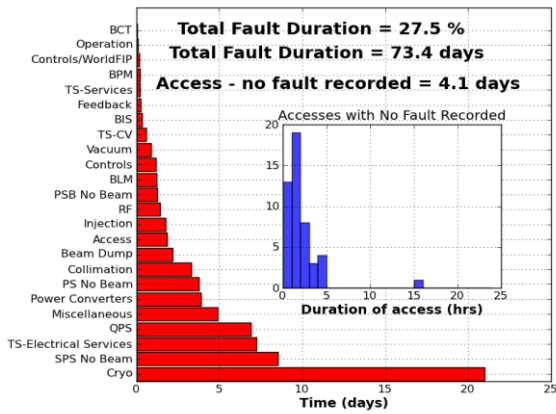


Figure 7: Distribution of stable beams durations over the 2011 proton Run. The average duration is 5.76 hrs.

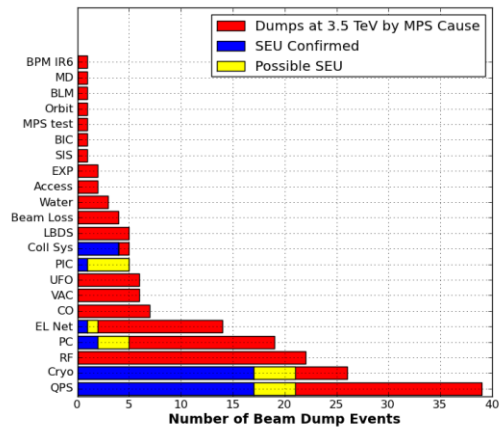
FAULTS AND DOWNTIME

Despite the impressive machine performance there are several issues associated to machine turnaround time and down time that need to be addressed.

LHC downtime was accrued due to system faults or failures, with faults defined as any incident, hardware or software failure, which prevented normal operation. Typically the shift crews register faults via the e-logbook interface, and this allows for fault statistics by system as shown in Figure 8. From this it appears that cryogenics is the clear leader in downtime, with a global down time of ~21 days. However Figure 9, which is a detailed accounting of the cryogenics downtime shows a downtime of 25.9 days.

This highlights an issue with the fault tracking in the e-logbook, as it appears that not all faults were recorded by the shift crews, and some faults went undocumented or were hidden in the shadow of others. This can be seen by the inset of Figure 8, which shows the time spent in access during the run when there was no fault recorded in the e-logbook. These unaccounted faults were typically faults in the QPS system, and this is consistent with Figure 10, where QPS dominates the MPS dump cause for all beam dumps at 3.5TeV. Clearly, for efficient and effective fault tracking and for downtime counting, fault-tracking procedures for 2012 have to be revised.

Figure 8: Breakdown of LHC downtime by system as registered in the e-logbook fault tracking. The inset shows



the distribution of downtime events where no fault was recorded in the e-logbook, but for which the machine was in a state compatible for access.

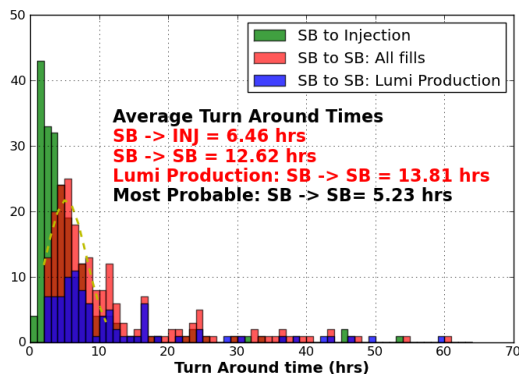
From Figure 10 it can also be seen that downtime due to radiation induced single event upsets (SEUs) was a significant factor for several systems, including QPS, Cryogenics, Power Converters, the electrical network, PIC, and Collimation. Mitigation actions during the 2011 run and the 2011-2012 Christmas stop have been performed in order to reduce sensitivity of equipment to SEUs, and indeed Figure 9 shows the effect of one such mitigation. On August 5th, redundancy improvements were made on PLC logic controllers of the cryogenics, and a significant decrease in the rate of accumulated cryogenics global downtime can be observed.

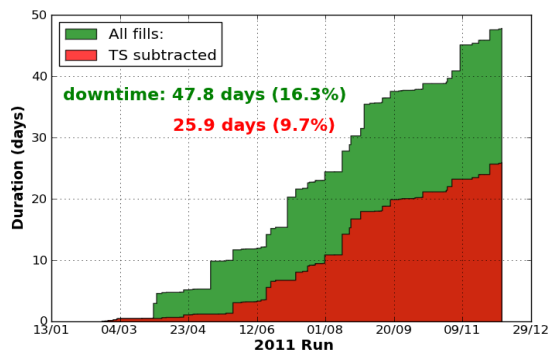
Figure 9: Total cryogenics downtime as a function of the run duration. Downtime is defined as the absence of the global AND of all the Cryo_Maintain signals.

Further mitigation actions have been implemented across the machine, and it is expected that most if not all SEU affected systems will show a marked reduction in downtime (especially QPS and Cryogenics)

Figure 10: Beam dump cause by system for all beam dumps that occurred at 3.5TeV during the 2011 Run. The beam dump cause is as assigned by the Machine Protection Expert Panel analysis that is performed off-line. Note: blue and yellow bars are stacked histograms showing the relative contribution of beam dumps coming from single event upset events produced by radiation.

BEAM DUMP STATISTICS






dumps in 2012, but it is not clear if the percentage of non-programmed beam dumps in the post squeeze phase will move from the 2011 Proton Run value of ~50%, due to the demand for higher beam intensities and increased performance reach of the LHC in 2012.

Over the course of the 2011 run, faults such as discussed above, often triggered beam dumps. Of a total of 482 beam dumps in the 2011 Proton Run, 375 (78%) were non-programmed beam dumps, and of those non-programmed beam dumps, 168 (35%) were in stable beams. The breakdown of the 375 non-programmed beam dumps by beam mode is given in Figure 11, and shows that ~50% of all such beam dumps occurred in the post squeeze phase of operation.

Figure 11: Breakdown by beam mode of non-programmed beam dumps that occurred in the 2011 Proton Run.



If these beam dumps are analysed in terms of the initial trigger and by actual cause (as determined by the Machine Protection Expert review panel), the PIC (Powering Interlock Controller) is the dominant first trigger system (Figure 12), despite not being the system with the fastest reaction time. This is due to the fact that the PIC handles the interlock requests from all the cold magnets, their UPS units, their cryogenic conditions, and their emergency stop inputs, and due to the large scale of the PIC surveillance, it is natural that it is the leading first-trigger system. However, Figure 10 demonstrates that the actual cause of the beam dump is spread over a number of systems, with QPS, Cryogenics, RF, Power Converters, and the Electrical network accounting for three quarters of the non-programmed beam dump triggers.

Figure 12: Breakdown by first trigger of beam dumps at 3.5TeV for all fills in the 2011 Proton Run where a non-programmed dump occurred which was not classified as having a Machine Protection beam dump cause of End-of-Fill.

Mitigations against SEU effects are expected to suppress the QPS, Cryogenics, Power Converter, and Electrical Network triggers in 2012. Yet the number of non-programmed beam dumps will not necessarily decrease, and that the distribution shown in Figure 10 will evolve, due to beam related effects and increased beam intensities as LHC performance is again pushed in 2012.

SUMMARY

The 2011 run has been hugely successful, with luminosity delivery well beyond target, successful implementation of luminosity levelling procedures, and Stable beam availability of ~32 %, which corresponds to a Hubner factor of 0.22 and 0.24 for proton and Ions runs respectively. While fault tracking was not fully complete, successful mitigations against downtime due to radiation induced single event upsets, have already been observed, and further improvements are expected in 2012. This should help reduce the number of non-programmed beam