LHC AVAILABILITY AND PERFORMANCE IN 2012

A.L. Macpherson, CERN, Geneva, Switzerland

Abstract

The LHC performance and overall machine availability for the 2012 proton-proton run are discussed, as well as the factors that contributed to another excellent LHC run

INTRODUCTION

Following on from an excellent year in 2011 in which ~5.5 fb^{-1} of proton-proton collisions at 3.5 TeV were delivered to both ATLAS and CMS, the 2012 run was intended to further extend performance reach. For the 2012 proton-proton run, beam energy was raised to 4 TeV, the beta* squeeze was set to 60 cm, and the target average bunch intensity set at $1.6x10^{11}$ protons. Also, to ensure expedient luminosity delivery it was decided to continue with 50ns bunch spacing, and push back the 25ns scrubbing program to late in the 2012 run.

With this configuration, a target delivered proton-proton luminosity (based on the 2011 luminosity production) was set at 15 fb^{-1}. This was seen as an ambitious goal, given that initially a proton-Lead run, four machine development periods, and four technical stops were also scheduled for 2012. Due to the strong request by the experiments for luminosity delivery, a machine schedule was revised was put in place which allocated the entire 2012 run to proton-proton physics, and moved the proton-Lead run to early 2013. The revised 2012 schedule is as shown in Figure 1. With this schedule the revised target for delivered luminosity was set at 22 fb^{-1} for both ATLAS and CMS.

In actuality the 2012 LHC run exceeded expectations, with a final delivered luminosity of over 23 fb^{-1} for both ATLAS and CMS, and the mid-year announcement of the discovery of a Higgs-like particle based on the combined 2010-2012 data sets [1]. Indeed, this excellent result, along with a proton-Lead pilot run, a high-beta physics program [2], a 25ns scrubbing run [3] and pilot 25ns physics fill, and a vigorous machine development program [4], meant that 2012 exceeded all expectations in terms of machine performance.

The LHC Run for 2012 can be summarised as follows:

- Hardware Commissioning: 35 days
- Beam commissioning: 21 days
- Machine Operation: 257 days
- Physics Operation: 228 days
- p-p Luminosity Production running: 201 days

Within the run, the following were also included:

- 3 Technical stops
- 4 Machine development periods
- 2 Floating Machine development periods
- A 25ns scrubbing run

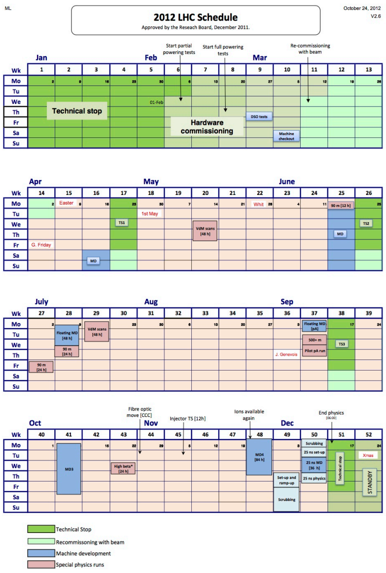

Figure 1: The final LHC machine schedule for 2012.

In terms of fill numbers, the 2012 proton-proton physics run extended from fill 2465 to fill 3457.

LHC AVAILABILITY

After short periods of hardware commissioning and beam commissioning, physics operation started on the 4th of May and continued through till the 17th of December. LHC machine availability for the 2012 proton-proton run is defined by the run period after commissioning, but excluding technical stop and machine development periods, and is shown in Figure 2. Over 36% of the run was spent in physics (stable beams) operation for a total of 73.2 days of ~1757 hours of physics, compared to 32% in 2011.

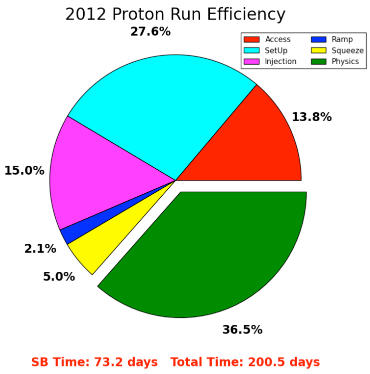

Figure 2: LHC machine availability for 2012.

If we compare the availability figures to those of the luminosity production running in 2011 [5], the difference in percentages for the 6 machine phases can be extracted, and is shown in Figure 3.

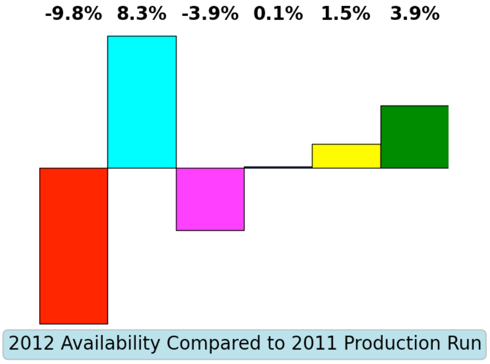

Figure 3: Comparison of 2012 LHC machine availability to that of 2011. What is shown is the difference between the 2012 and 2011 percentages for each of the 6 machine phases defined in Figure 2.

Figure 3 shows less of the available run time was spent in access in 2012 than 2011, but that this was almost completely countered by the relative increase in the beam setup phase (machine closed but no beam injected). In terms of improvements, Figure 3 shows a decrease in the percentage of time spent in injection, and an increase in the stable beams percentage. This tends to suggest that in 2012, we improved the injection procedure, and this improvement translated into more available stable beams time. It should also be noted that in 2012, the beta* squeeze was split into a two step procedure in order for LHCb to transition from a horizontal crossing angle at injection to a vertical one at physics settings (in order to improve operational conditions under polarity flips of their external crossing angle coming from the LHCb dipole).

To quantify the availability improvements it is worth comparing the 2012 Hubner factor with that of 2011. 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). The Hubner factor in 2012 is H_{2012} = 0.175 which assumes a physics duration of 200.5 days, a peak luminosity of 7695 (μb.s)$^{-1}$ and a delivered luminosity of 23.269 fb^{-1}. The equivalent 2011 value is H_{2011} = 0.156, and so implies a clear indication of improvement of machine performance and physics availability.

LHC PERFORMANCE - LUMINOSITY

For the proton-proton 2012 run, the default filling scheme was with 1374 bunches per beam and 50ns spacing between bunches, which gave 1368 colliding bunches in ATLAS and CMS, 1262 in LHCb, and no colliding bunches in ALICE. Due to detector constraints, ALICE data taking was done with collisions generated by main bunch-satellite bunch collisions, which gave a reduced rate compatible with the ALICE detector.

With the schedule and availability as outlined above, the machine was able to deliver luminosities of 23.27 fb^{-1} for both ATLAS and CMS (see Figure 4) and over 2.1 fb^{-1} to LHCb.

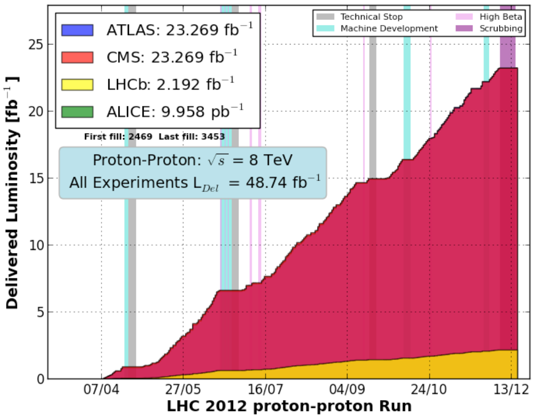

Figure 4: Delivered luminosities for the 2012 LHC proton-proton run. ATLAS and CMS delivered luminosities are almost identical, hence almost indistinguishable, while ALICE delivered luminosity is not visible due to the absolute scale.

In terms of delivered luminosity, the performance of the machine is best put in context when the target estimates are considered. Figure 5 shows both the target luminosity delivery for the run, and the actual luminosity delivery for CMS, and it can be seen that without the run extension, the machine was on target to reach the 16 fb^{-1} expected. With the run extension, the machine exceeded expectations and delivered ~1 fb^{-1} above target by the end of the run.

Such impressive performance was based not only on machine availability, but also on careful attention to the optimisation of operational parameters. This optimisation was done throughout the year, and a summary plot of the machine tuning over the year is given in Figure 6.

The target average bunch intensity at injection was set at 1.6x10^{11} protons per bunch, and Figure 6 shows that this translated into an achieved bunch charge of ~1.5x10^{11} protons at declaration of stable beams, and that only a moderate increase over the year was possible. Similarly, transverse emittance stayed relatively constant over the year despite the mid-year move from the Q26 to Q20 SPS optics[6] and other optics corrections to the injectors, which significantly reduced the transverse emittance at injection. However these improvements in injector optics can be seen in Figure 6 in terms of increased peak luminosity and transverse beam brightness, but also in the growth of the longitudinal bunch length, indicating that in terms of bunch charge, the machine was running close to its limit.

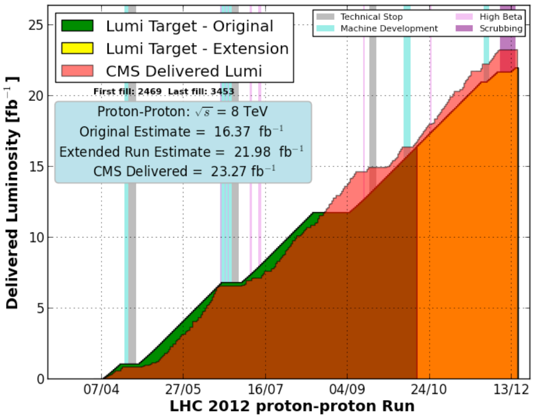

Figure 5: Comparison of target and delivered luminosities for the 2012 LHC proton-proton run. By the completion of the original run period (green) the actual and the target CMS delivered luminosities are almost identical. During the extended run period (yellow) delivered luminosity exceeds expectations.

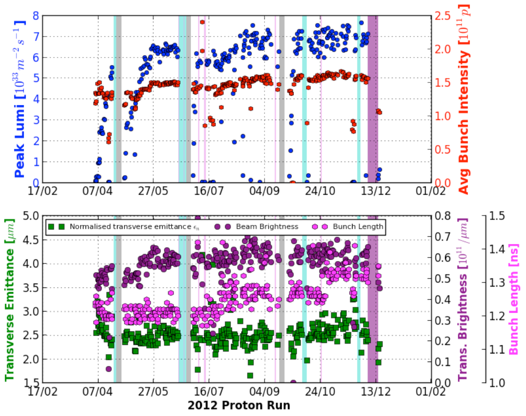

Figure 6: Beam parameter evolution during the 2012 LHC proton-proton run.

The mid-year improvement in injected beam quality, combined with enhanced satellite production in the PS enabled a significant reduction in bad background conditions in ALICE, and an increase in the rate of main bunch - satellite collisions. This had a dramatic effect on the both the ALICE data taking efficiency and their delivered luminosity, as shown in Figure 7.

For LHCb, extensive use of luminosity levelling by separation was made to ensure data taking with controlled trigger rates. Data taking efficiency was further enhanced by the two-step squeeze process that rotated the crossing angle so that it was orthogonal to the external crossing angle from the LHCb dipole. This change had the advantage that the machine operation was more transparent to the regular LHCb dipole polarity flips, and so helped improve machine turnaround. Figure 8 shows the LHCb delivered luminosity and the dipole polarity flips performed during the year. The final ratio of luminosity taken with positive and negative LHCb dipole polarities is 49.2%:50.8%, and so meets the LHCb requirements of balanced data sets, needed to reduce systematic errors in their physics analyses.

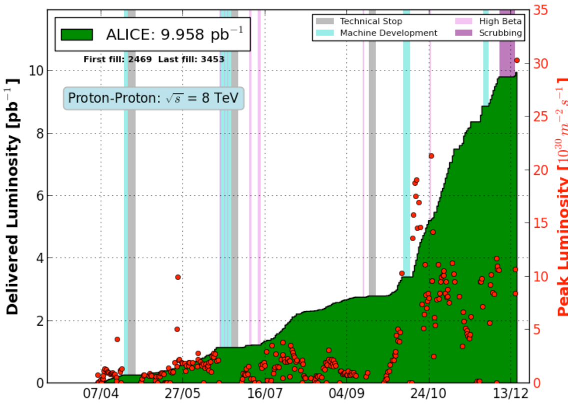

Figure 7: Delivered and peak luminosities for ALICE over the course of the 2012 proton-proton run. A clear improvement is seen after technical stop number 3, due to improvements in the beam quality from the injectors, and enhancement of the satellite population in the PS.

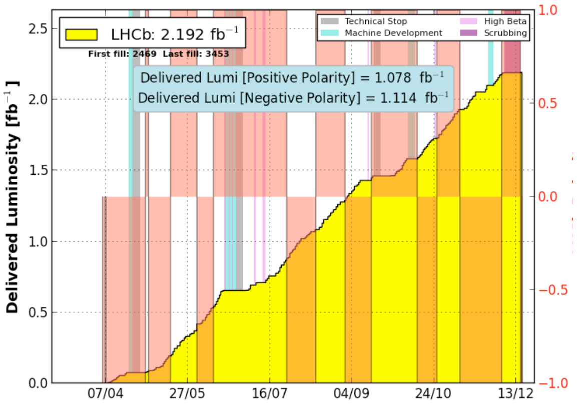

Figure 8: LHCb delivered luminosity and LHCb dipole polarity for the 2012 proton-proton run.

Finally, in terms of luminosity delivery, the weekly performance is given in Figure 9, and it shows that during the course of the year, the recovery from technical stops improved, although the statistical significance of this improvement may not be strong.

LHC PERFORMANCE - DOWNTIME AND SYSTEM FAULTS

Delivered Luminosity is not the only measure of machine performance; machine downtime and fault statistics are also key indicators that show not only performance but also the possible onset of operational issues and equipment failure modes. To examine these availability factors, both postmortem data from all beam dumps above 450 GeV and operations fault tracking data from the e-logbook and the Technical Infrastructure Major Event tracking has been used to extract performance estimates. Unfortunately this data does not form a complete set, and it has already been identified that the fault recording and tracking mechanisms need to be revised before the LHC restart in 2015 [7].

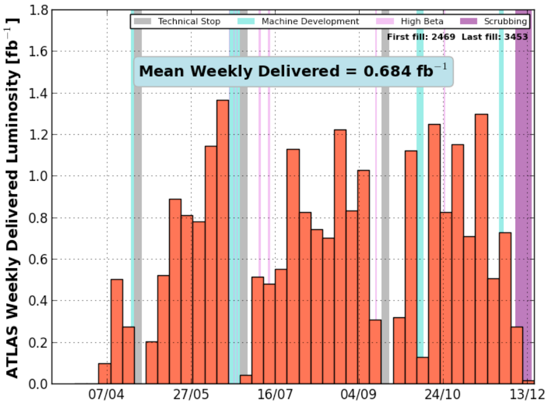

Figure 9: Weekly delivered luminosities during the 2012 LHC proton-proton run.

As a first estimate, independent of fault type, the time to recover from a beam dump until the start of ramp can be considered, and is shown here in Figure 10. Two setup time distributions are shown; the raw distribution as extracted from machine operation, and the fault corrected setup time, which for any given fill, is this setup time minus any declared fault time. This choice of setup time definition was chosen as an indicator, as it allows inclusion of both for recovery from faults in the previous fill, and the inclusion of unrelated delays from injection.

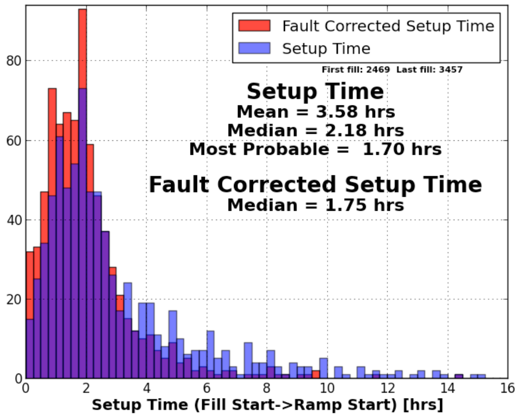

Figure 10: Machine setup time (including a 45 minute ramp down sequence). The fault corrected setup time is based on the faults logged in the LHC logbook. evolution during the 2012 LHC proton-proton run.

While the most probable setup time remains the same for both distributions, for the fault corrected setup time there is a reduction in the median setup time and a clear reduction of the tail of the distribution. Both distributions appear to follow a log-normal shape; log-normal distributions can be thought of as the multiplicative product of many independent random variables each of which is positive. The median fault corrected setup time is 105 minutes, which includes the ramp down procedure which has a minimum duration of 45 minutes. This discrepancy between actual average setup time and minimum ramp down duration, suggest that either the machine setup procedure is well away from being in the shadow of the ramp down, or that not all faults and delays have been fully accounted for.

To understand better the time delays associated with machine setup and overall turnaround, it is instructive to look at the recorded system faults, both in terms of occurrence and fault resolution time. From the data available in the e-logbook, the system faults histogram of Figure 11 can be produced. In this figure a fault is defined as any incident, hardware fault, or software failure which prevented normal operation.

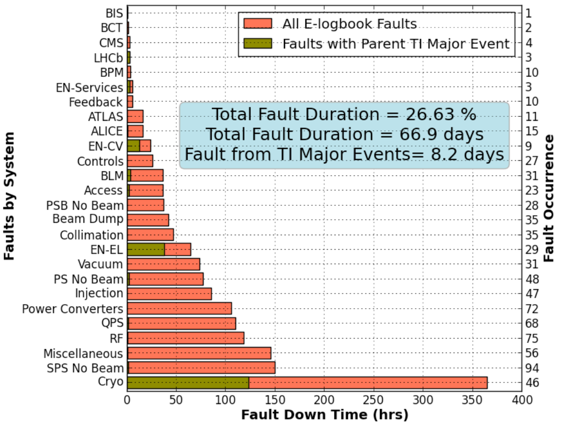

Figure 11: Breakdown of total fault time by system during the 2012 LHC proton-proton run.

From Figure 11, cryogenics is the clear leader in downtime, with a global down time of ~15 days. By comparison, detailed cryogenics availability data gives the total cryogenics availability (technical stops excluded) as 13.6 days (see Figure 12).

As noted, in terms of downtime, cryogenics dominates, but this is to be expected due to the reset procedure of the cold compressors and the thermal inertia the cryogenics system. However it is worth noting that in Figure 11, it can be seen that ~1/3 of the down time was associated to external events (as recorded in TI Major events) that triggered trips of the cryogenic sectors. A typical example of such and external event is an electrical network perturbation. Further, it is extremely encouraging to note cryogenics availability rose from 87.1% in 2011 to 94.4% in 2012. This reflects the consolidation within cryogenics, and the mitigation of communication faults which perturbed cryogenics in 2011.

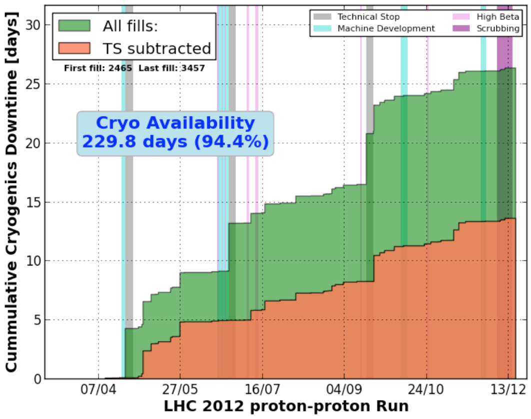

Figure 12: Cryogenics downtime during the 2012 LHC proton-proton run. The integrated downtime excluding the Technical stop periods is shown in red.

For machine performance in terms of system faults, it is also instructive to look at the beam dump statistics of fills above injection energy. Data for such statistics is taken from the LHC postmortem data base, and a breakdown by beam dump cause is given in Figure 13. As in 2011, the QPS is the most proficient, and was also the leader in terms of beam dumps triggered by radiation induced single event upsets (SEUs). This is as to be expected, as the planned QPS upgrade is not foreseen to be completed until after the 2013-2014 long shutdown. However, what is encouraging is that both the total number of cryogenics induced beam dumps and the number SEU triggered cryogenics induced beam dumps significantly decreased in 2012, and is attributed to the aforementioned consolidations and mitigations.

For an overall comparison the percentage of SEU induced beam dumps dropped from 17.5% in 2011 to 9.5% in 2012, and implies a significant improvement in performance and a validation of the R2E consolidation activities [8].

In terms of recovery time after the beam dump, Figure 14 gives the breakdown by system of the cumulative sum of recovery times from beam dump back to injection. While this distribution may be susceptible to error due to individual fills with multiple systems failing, the distribution shows, like in 2011, that the QPS system, due to both the high occurrence of faults, and the cost in terms of system recovery (LHC access or full power reset of circuits that then require precyling), is the leading system in terms of cumulative post beam dump recovery time. Indeed the top five systems are the same in 2011 and 2012, and apart from cryogenics, the other four usually involve LHC access to address the fault that triggered the beam dump. Naturally, this incurs extra downtime due to resolving of the fault, and the process of LHC machine access (especially if it has to be coordinated with the radiation piquet outside normal working hours).

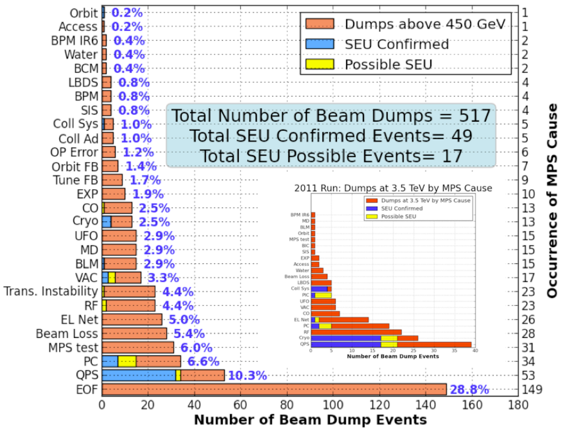

Figure 13: Beam dumps above injection energy by cause. Blue and yellow bars are stacked histograms representing SEU confirmed and SEU possible triggers, whereas the total number of beam dumps (red) is not stacked.

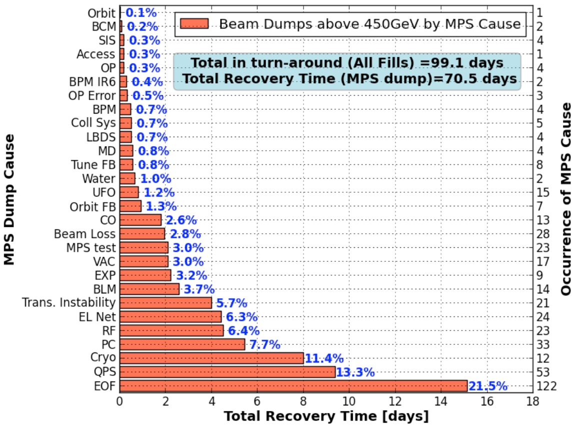

Figure 14: Cumulative recovery time by machine protection dump cause.

SUMMARY

The 2012 the LHC again completed a very successful proton-proton physics run, delivering in excess of 23 fb^{-1} to both ATLAS and CMS, which gave sufficient events for the discovery of a Higgs-like particle. As a measure of performance, the delivered luminosity was again beyond target, while the fraction of time in physics in 2012 improved by ~4% compared to 2011. This is reflected by the improvement of the machine Hubner factor from $H_{2011} = 0.156$ to $H_{2012} = 0.175$.

As this was a luminosity production run, beam quality, beam optics, and the operational cycle in both the injectors and the LHC itself were addressed, and this allowed for optimised luminosity delivery, which is perhaps best typified by the improvements made for ALICE experiment (Figure 7).

Unfortunately the fault and downtime tracking system is still not ideal, but the picture that emerges in 2012 is similar to 2011 in terms of the systems that have the biggest contribution to LHC downtime. These systems (particularly QPS and Power Converters) are to undergo substantial upgrades in the upcoming long shutdown, which should help reduce the LHC downtime.

Yet it is very encouraging to observe both the number of single event upset related beam dumps in the cryogenics system and the total cryogenics downtime were drastically reduced. This is due to R2E mitigations and system consolidation, and offers some assurance that the overall machine availability can be further improved.

REFERENCES

[1] Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC, ATLAS Collaboration. Physics Letters B, Volume 716, Issue 1, 17 September 2012. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC, CMS Collaboration Physics Letters B, Volume 716, Issue 1, 17 September 2012.

[2] Commissioning and operation at beta* = 1000 m in the LHC, H. Burkhardt et al., IPAC 2013.

[3] Electron Cloud and Scrubbing, G. Iadarola and G. Rumolo, these proceedings.

[4] LHC Machine Development Program, https://espace.cern.ch/lhc-md/default.aspx

[5] LHC Availability and Performance in 2011, A. Macpherson. LHC Beam Operation workshop - Evian 2011, 12-14 Deember 2011. CERN Reports Number: CERN-ATS-2012-083.

[6] Q20 in SPS, H. Bartosik, LHC Injector Upgrade 2013, 12 April 2013 https://indico.cern.ch/conferenceDisplay.py?confId=238152

[7] A Look Back on 2012 LHC Availability, B.Todd et al., these proceedings.

[8] R2E Experience and Outlook, G. Spiezia these proceedings.