

MACHINE AVAILABILITY AT THE LARGE HADRON COLLIDER

M. Pojer, R. Schmidt, M. Solfaroli Camillocci, S. Wagner, CERN, Geneva, Switzerland

Abstract

One of the most important parameters for a particle accelerator is its uptime, the period of time when it is functioning and available for use. In its second year of operation, the Large Hadron Collider (LHC) has experienced high machine availability, which is one of the ingredients of its brilliant performance. Some of the reasons for the observed MTBF are presented. The approach of periodic maintenance stops is also discussed. Some considerations on the ideal length of a physics fill are drawn.

INTRODUCTION

Machine availability, the fraction of uptime with respect to the aggregate of the up- and down-time, plays a fundamental role for a particle accelerator. The performance must be maintained for the time needed to accumulate the required statistics.

In its second year of operation, the CERN Large Hadron Collider (LHC) has known an excellent productivity, with the peak luminosity constantly increased and brought to 40% of nominal. It is the result of an accurate and well monitored strategy of performance increase, based on the optimization of the operation and achieved by pushing the machine parameters towards their limit. The integrated luminosity exceeded the 2011 target by a factor 5, as a consequence of the high machine availability during the year. The chosen approach to maintainability and the capacity of quickly identifying and fixing all issues are key elements of this availability.

2011 PERFORMANCE

Machine Operation

The performance of the LHC was progressively pushed in the first half of 2011 by increasing the number of bunches and bunch intensity [1]. After summer, the maximum luminosity was further increased by reducing β^* to 1 m. The integrated luminosity at the end of the year was around 5.6 fb^{-1} for the high luminosity experiments, well above expectations.

Machine Availability and Stable Beams

The commissioning of the machine in 2011 was carried out between Jan 27 and Mar 13 (commissioning of the technical system first and then commissioning with beam for 4 weeks). Taking this last date as the start of operation, until Dec 6 (date of the last beam, prior to the shutdown of the machine for the end of the year) 269 days were devoted to luminosity production.

Figure 1 shows the distribution of the different phases of the machine along the 269 days of operation. About one quarter of the time (corresponding to a total of 63

days) was invested in collisions (stable beams, i.e. when the experimental detectors are switched on and collect data). Another quarter was spent in preparation of the collisions, while a bit more was used to setup the machine (preparation before a physics fill, machine protection validations and specific studies). From this statistics we conclude that the machine is operational for more than 82% of the time, with only 17.7% of the time lost for access. It is important to note that this does not correspond to an availability of 82% since not all faults require an access. This time is counted in the setup period. If a fault happens in stable beams, the time to recover the conditions, i.e. to go back in stable beams should be counted as unavailability of the machine, which is not the case.

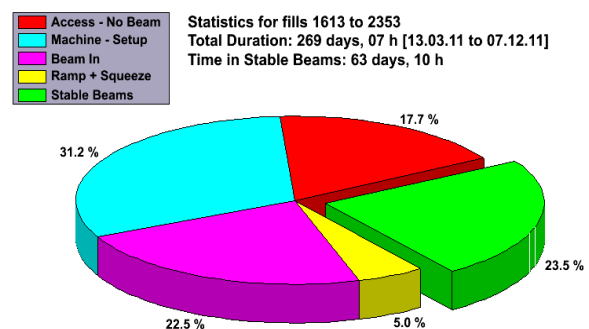


Figure 1: Time distribution during 2011 run.

Another option is to estimate the availability as the ratio between the minimum time required by all fills that produced luminosity and the total time dedicated to physics. In other words, let's consider 400 fills terminated with stable beams in 2011 and assume a turn-around time of 3 h (i.e. the time between two consecutive stable beams, which was a minimum of 2h07min last year). If we multiply 400 by 3 (= total overhead) and add the time spent in stable beams (63 days), we get 113 days of effectively used machine time. Over 269 days, this translates in 42% machine availability.

It is worth noticing that stable beams that terminated prematurely are reducing machine availability and should not have the same weight as the others in the above calculation: we should normalize with respect to the produced luminosity. One approach in this direction considers the efficiency as the ratio between total integrated luminosity and maximum luminosity that could be obtained in case of absence of faults. Further reduced machine availability would result from it.

One parameter that is frequently used for the efficiency of a collider is the Hübner factor, i.e. the ratio of delivered luminosity to the amount that could be collected by running continuously at the peak luminosity. In its first year of luminosity production, LHC operation resulted in

a Hübner factor of about 0.22-0.23, a very good result for a new machine with extreme complexity. For comparison, the Tevatron reached, after decades, of operation a value of 0.3 in 2011. It should be noted that the luminosity lifetime at the LHC is higher than the one of the Tevatron, contributing to a relatively large Hübner factor during the early days in LHC.

The Approach to a Reliable Operation

From a design point of view, the availability of a process or a machine is defined as:

$$A_d = MTBF / (MTBF + MTTR) \quad (1)$$

A_d is a function of the mean time between failures (MTBF, which is the predicted time between inherent failures of a system during operation) and the mean time to repair (MTTR).

Very large MTBF values compared to MTTR result in high availability. This is correlated to the reliability of the machine (i.e., the probability that the machine will perform as designed). If reliability decreases (i.e., MTBF becomes smaller), better maintainability (i.e., shorter MTTR) is needed to achieve the same availability. Contrarily, when the reliability increases, then maintainability is less important to achieve the same availability. A trade-off is required to achieve the objectives.

This is why preventive maintenance is performed to improve reliability (longer time between failures) and maintainability (shorter recovery periods). It is common to introduce the definition of operational availability as

$$A_o = MTBM / (MTBM + MDT) \quad (2)$$

In this case the MTBF is replaced by a mean time between maintenance, which considers all corrective and preventive actions, and the mean down time includes all time associated with the system being down for both corrective and preventive maintenance.

Preventive maintenance to increase the reliability and therefore the availability of the machine is also applied for LHC. Several stops for preventive maintenance during the year (one technical stop each two months approximately) have been chosen to avoid premature wear out and long stops due to unexpected failures of vital systems, like cryogenics, cooling and ventilation or the powering systems.

One could object that this kind of approach is self-killing, in the sense that the down-time is artificially increased; but a careful planning of all interventions can finally reduce the mean time for maintenance. In fact, the total time for maintenance is the sum of the time to prepare the intervention, the time for interventions and the start-up time. Planning allows to reduce to the minimum the time to prepare, thus improving the process. Ideally one would like to reduce to a minimum as well the restart time. Unfortunately, the number of interventions

(hardware and software) during a technical stop is high and some re-commissioning of the machine is required, in particular for machine protection systems. Sometimes the restart is difficult. Experience has shown [2] a learning curve in 2011. Still, more can be done, in particular a tighter control on the maintenance activities and in increasing of the awareness among the teams intervening on the machine, since sometime interventions have side effects that are only observed when beam operation resumes.

MAJOR SYSTEM UNAVAILABILITIES

Preventive maintenance interventions are not sufficient to ensure a full availability in between technical stops. The system reliability is playing a role. In addition, starting from early 2011, the effect of radiation to electronics started reduce the MTBF of equipment, e.g. by single event upsets in the electronics leading to a beam dump.

Looking at the statistics of dump causes, it turns out that among five systems with the highest recurrence of faults, three of them are those that more seriously were affected by single event upsets (SEU): the quench protection system (QPS), the cryogenics and the power converters, which are accounting for 50% of the total dumps [3]. The other two systems, RF and electrical distribution, suffered from operation with high intensity and from the high sensitivity to network perturbations.

Cryogenics

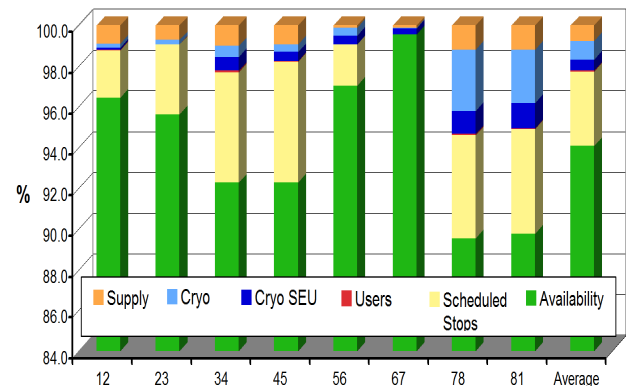


Figure 2: Availability of the cryogenics in the 8 sectors.

In terms of machine availability, the cryogenics system dominated the LHC downtime in 2011, with a total of 21 days when the temperature conditions of the superconducting circuits could not be guaranteed. For such a complex system the availability of the cryogenics was still very high, with a minimum of 94%, if we don't consider the scheduled stops (Figure 2).

SEU

With the increase of the number of bunches and their intensity, i.e., with the increase of the luminosity (between 2 and 10 times from the nominal performance), electronic equipment sensitive to radiation and installed in

areas exposed to ‘high’ doses, started to manifest spurious trips [4]: the cryogenics controls, the QPS detectors, the controls of the power converters and of the collimators; all systems installed in areas with high-energy hadron fluence larger than few 10^6 cm^{-2} were affected by this phenomenon. In Figure 3, the number of SEUs observed during the year is shown, with about 150 observed and confirmed events [5].

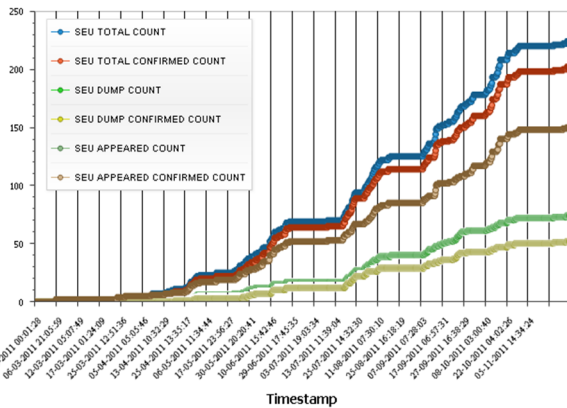


Figure 3: Number of SEUs observed in 2011.

The appearance of many radiation-related dumps led to implement some mitigation measures earlier than initially planned, such as the installation of additional iron shielding in critical areas and the relocation of critical sensitive equipment. The control system for the cryogenics for one of the four cryogenics plants was relocated during the summer period. Many SEUs were avoided thanks to the actions taken and the reliability of cryogenics increased. A massive campaign of relocation of equipment is planned for the first long shut-down at the end of 2012.

MAXIMISING THE TOTAL INTEGRATED LUMINOSITY

Once the availability of the machine is acceptable, one can think about the way to maximize the total integrated luminosity. Assuming constant maximum luminosity, the integrated luminosity can be increased by finding the right compromise between turn-around time (imposed by the time needed to prepare the machine after a dump, plus the time to recover from a possible fault) and fill length.

In fact, even if not many 10% of the fills in stable beams were dumped voluntarily by the operator, all other fills were dumped due to a failure. Voluntarily ending a fill leads to a reduction of the total integrated luminosity, but this might be compensated if the turn-around time is very good. When the beams are degraded, the luminosity becomes very poor, and even a long interval between stable beams could increase the total integrated luminosity if a high peak luminosity is restored. The ideal moment to dump a stable beam is not easy to find.

Consider having a series of physics fills. Let's assume an exponential decay of the luminosity of each fill, with a time constant of 20h. If we want to calculate the optimum

time for the dump, we have to calculate the integrated luminosity along the whole period, which corresponds to calculate the average integrated luminosity of a fill of mean length (mfl), separated from the next one by a mean turn-around time ($mtat$). A different maximum of the integrated luminosity is then found for different $mtat$ at an optimum of fill length (ofl), as shown in Table 1 for few cases. The maximum integrated luminosity can be calculated as well for different $mtat$.

Table 1: Ideal Fill Length as a Function of Turn-around Time (L_0 is the initial luminosity)

$mtat$ [h]	2	3	4	5	6	8	10	12	15
ofl [h]	8.3	10.0	11.4	12.6	13.8	15.7	17.3	18.9	20.5
$maxintL$ [x L_0]	950	873	814	767	727	664	615	574	525

For the LHC minimum turn-around time (2 h), the ideal fill length would be about 8 h. Unfortunately, the average turn-around time along 2011 was more than 12h, which means that a fill should possibly never be dumped voluntarily.

To complete the figure, it is important to notice that the average stable beams duration in 2011 was less than 6 h, with a maximum of stable beams length of some 25 h.

CONCLUSIONS

Good machine availability was obtained in 2011, which was counterbalanced by an avalanche of faults. The corrective/preventive actions taken were beneficial and could increase already in 2011 the reliability of some key systems, but still much has to be done.

ACKNOWLEDGMENT

The authors would like to thank all people involved in the operation of the LHC; in particular, operators, engineers in charge and machine coordinators, for the incredible work done along the year. Many thanks, especially to Alick Macpherson for the work done on the LHC statistics, and to Enrico Bravin for the fruitful discussions during the nights of (un)stable beams.

REFERENCES

- [1] J. Wenninger, "Operation of the LHC at High Luminosity and High Stored Energy", these proceedings.
- [2] M. Solfaroli Camillocci, "Technical Stops: What Were the Issues in 2011, Minimizing Impact and Improving Recovery", proceedings of the LHC Beam Operation workshop, Evian, Dec. 2011.
- [3] M. Zerlauth, "Magnet Powering with Zero Downtime – a Dream?", proceedings of the LHC Performance Workshop, Chamonix, Feb. 2012.
- [4] G. Spiezia, "R2E – Experience and Outlook for 2012", proceedings of the LHC Beam Operation workshop, Evian, Dec. 2011.
- [5] "The LHC Post Mortem system", <http://lhc-postmortem.web.cern.ch/lhc-postmortem/>