

LHC OPERATIONAL EFFICIENCY IN 2010

W. Venturini Delsolaro, CERN, Geneva, Switzerland

Abstract

An overlook on the beam and machine statistics in the 2010 run is given. We report on the machine availability and efficiency for physics and give a breakdown of the downtime according to the various technical systems. We revise the frequency and duration of the technical stops with respect to their impact on the machine availability. Finally the tools presently available for the collection of this kind of data are reviewed and needs for 2011 are defined.

INTRODUCTION

LHC beam operation in 2010 was mainly driven by commissioning activities, although a significant collision data set was eventually delivered to the experiments.

The time period considered in this paper spans from the 1st of March to November 30 (6600 hours). Machine statistics were collected by surveying the e-logbook and cross checking with minutes of various meetings and with logged data for the beam intensity. The machine states considered here are beam setup, stable beams, setup without beam (the joined time of these three being defined as the machine availability), technical stop and fault (machine not available due to some system fault). Time spent in supplementary hardware commissioning was included in technical stops.

MACHINE STATISTICS

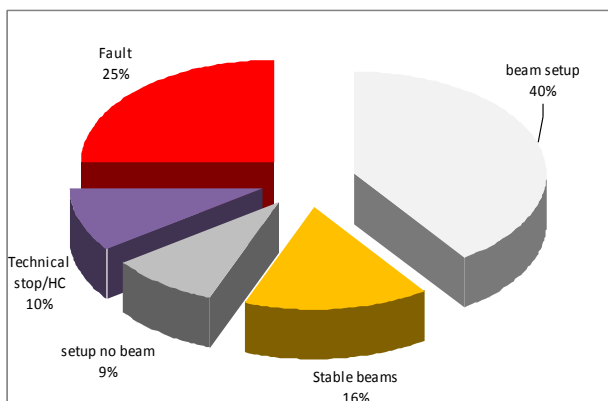


Fig. 1 Global 2010 machine statistics

Setup without beam

At the LHC, even with 100% availability, there would always be a physiological time without beam: the magnetic machine must be brought back to injection energy at the end of physics fills, and it needs to be pre-

cycled whenever the magnetic history deviates from the established standard, as for example after an access. In addition, the injectors have to prepare the required beam, which must be steered down the transfer lines and injected into the collider; a delicate operation in itself, which cannot always be carried out parasitically.

The time spent setting up the machine without beam was 9% of total. Cycling the machine as a consequence of faults was considered as downtime (machine not available) and attributed to the faulty system.

Beam setup

Under this category fall both the physiological phases with beam which are preliminary to collision data taking by the experiments (injection, adjustments at injection and at high energy, ramp, squeeze, steering of collisions); and all the machine commissioning and development with beam. These activities represent the highest fraction of total time (40%). Because of the way statistics were collected, this bin contains as well a good deal of inefficiencies, i.e. time when the beam was present but some problem was being handled (wrong settings, interpretation of doubtful measurements or unexpected events, struggles with the software, hesitation, panic, etc.), both during physics runs and during commissioning activities.

Stable beams

The time spent in stable beams was 16% of the total over the year. This rather low average was due to prevalence of the above mentioned commissioning periods. Figure 2 shows the evolution of this value along the run. In periods entirely devoted to physics we managed to have up to 29% of the time in stable beams. Noteworthy are the dips in June and in September, when the efforts were focused on commissioning the machine for higher intensities. Both were followed by an upward trend, which did not seem to reach saturation.

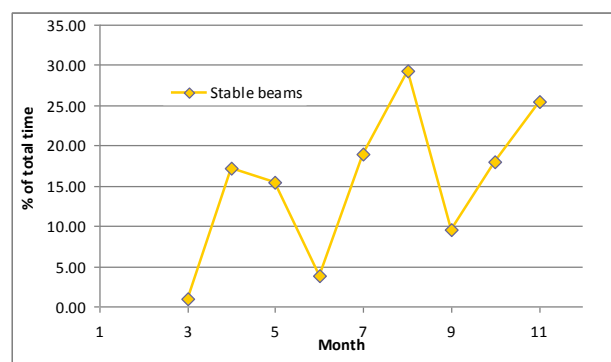


Fig. 2 Stable beams fraction along the run

Technical stops

There were six scheduled technical stops, with very little adjustments of the actual dates with respect to plans. The average duration of stops was 4 days, and the average spacing was 39 days. The main activities driving the frequency and the duration of technical stops were the maintenance of the cryogenics systems (de icing, replacement of malfunctioning valves), of the QPS (repair of quench heater power supplies, replacement of defective cards, etc), and replacement of power converter modules. In some cases hardware upgrades were carried out, for example to allow the QPS coping with higher current ramp rates in the main magnets. On these occasions some hardware commissioning had to follow the technical stop.

The time devoted to technical stops was 10% of total.

It is well known that maintenance activities, besides beneficial effects, can introduce new problems (as the say goes: as long as it works, do not touch it!). Trying to assess if the frequency of technical stops had been appropriate, I have considered the three days preceding and the three days following each technical stop and looked at the number of faults occurred in these two periods. Preventive maintenance reduces the number of faults after the technical stops, but on the other hand new problems appear. At the start of the run, the net result was that the number of faults after technical stop was (much) higher than before! However, in the course of the year, this phenomenon went decreasing and eventually it disappeared.




Fig. 3 Increment of faults after technical stop

Figure 3 shows the degradation (difference in downtime due to faults before and after technical stop), along the run.

Faults

The total downtime time due to faults (including the time needed to bring back the machine after the repair) was 25% of total. In many cases we had coupled faults, for example a QPS board would not come back after a trip due to a power converter fault, or a loss of cryogenics conditions or an electrical perturbation.

Faults statistics

The distribution of downtime according to the technical system is shown in figure 4.




Fig. 4 histogram of LHC faults

Data are raw: no attempt was made to normalize the downtime to the complexity of the systems. Therefore it is no surprise that a hugely complex system such as the QPS it at the top of the score. Since the integral of the histogram equals 25%, numbers can be multiplied by four to get the fraction of downtime for a given system.

The faults statistics of such complex systems show that, although improvements are still possible, their reliability is already remarkable. As an example, Table 1 gives some details of the QPS “internal” statistics.

Equipment type	Faults	Quantity	Availability [%]	MTBF [hours]
Quench heater power supplies	26	6076	99.998	1145760
Quench detection systems	19	10438	99.999	3362135
DAQ caused by radiation (SEU)	12	1624	99.997	828240
DAQ other causes than radiation	8	2532	99.999	1936980
DAQ all faults combined	20	2532	99.997	774792
EE600	6	202	99.988	206040
EE13 kA	5	32	99.939	39168

Table 1 detailed QPS statistics (courtesy R. Denz)

Although less frequent, faults in the cryogenics systems, in particular cold compressor stops, have a big impact on the machine availability because of the long recovery times.

Power converters have the third position. Again, this is expected due to the large number of elements.

Electrical perturbations from the supply network are the fourth source of downtime; the immunity of the LHC to this kind of events is somewhat lower than that of the injectors. The cryogenics systems, present only in the LHC, were sensitive to electrical perturbations at the start of the run, but the cryogenics team managed to increase their immunity in the course of the year.

The injectors contributed to the downtime due to faults for a little more than 8%. This is not the downtime of the injectors, but the injector faults seen from the LHC, i.e. the cumulated time when the LHC was requiring beam and the injectors could not deliver it due to some internal fault.

These five systems alone account for 70% the downtime. The remaining 30% is shared among 23 other categories. It should be noted that “small” systems may have low MTBF without becoming “visible” in the statistics. Also, systems which give “small”, i.e. easily recovered faults would create a “dust” of sub threshold incidents, which escape completely the present approach as they would not appear in the logs. Examples of this are small software bugs and many controls issues.

OPERATIONAL EFFICIENCY

For operational efficiency it is meant here the ability to use the available machine time in order to produce maximum integrated luminosity. It is the efficiency of the operations team running the machine in the control room. Once a refilling policy is given, there is a theoretical maximum fraction of the total time in which the machine can run in stable beams mode. In other words, operational efficiency is defined with respect to the minimum turnaround time.

It is useful to consider as well other definitions of efficiency: since after all the goal of the LHC is to produce integrated luminosity, then ultimately its efficiency is the fraction of runtime which is spent in stable beams. This is rather the efficiency of the collider, which considers downtime due to faults, technical stops, but also machine commissioning and development time as inefficiencies. Such a crude definition is certainly ungenerous, but not depleted of sense, from certain points of view. Other possible definitions would exclude some combination of machine commissioning, machine development time, and technical stops from the runtime. The so called Hubner factor was used at the time of LEP to relate the integrated luminosity to the peak luminosity and the scheduled time for physics [1]. In this case, operational inefficiency and hardware faults occurring during the scheduled physics time, but also the optimization of refilling, contribute to the final result.

As reported above, the LHC was producing luminosity in stable beams mode for 16% of the runtime in 2010.

Normalizing to the available time (i.e. not considering faults and technical stops) the resulting operational efficiency (for physics) would be 24% over the year. However that is not very meaningful as it includes commissioning and machine development in the operational inefficiencies.

During the last two weeks of August, when the only aim of the operation crews was to produce luminosity, the operational efficiency was 50%. This must be compared to the theoretical maximum, i.e. with minimum turnaround: in the period considered the operational efficiency could have been 83%, which indicates the margin for improvement from the side of operations. The analysis of operational inefficiencies is the subject of another contribution [2]

SUMMARY

The overall availability of the LHC during the first operation year was a remarkable 65% [Fig. 1], steadily increasing along the run.

The dominant activity was beam commissioning in a quest for higher intensities (first single bunch, then total), which eventually paid off with doubling of the luminosity goal for 2010 and delivery of $\sim 50 \text{ pb}^{-1}$ to the experiments.

Downtime due to faults amounted to 25% of the total; the top 5 systems were QPS, Cryogenics, EPC, EL-UPS, and the injectors. The hardware teams are working on identified weak points, although the reliability of the equipment is already very high.

Recovery from technical stops was initially troublesome, with a clearly visible degradation of the machine availability due to new faults after the stops. The detrimental effect of technical stops was steadily decreasing and disappeared at the end of the year.

Operational efficiency reached 50% (60% of the theoretical maximum) when running the machine in physics mode.

Finally, a word on tools: statistics are extremely important as they provide the input to understand and improve the exploitation of the LHC. Digging out the information from the logbook at the end of the year is time consuming and error prone. Automatic tools for data collection are missing, and needed for 2011.

ACKNOWLEDGEMENT

I like to thank R. Denz, S. Claudet, V. Montabonnet and M. Zerlauth for their help

REFERENCES

- [1] J. Wenninger, these proceedings
- [2] S. Redaelli, these proceedings