https://doi.org/10.1142/9789811278952_0004

Chapter 4

Operational Experience from LHC Run 1 & 2 and Consolidation in View of Run 3 and the HL-LHC

M. Lamont

CERN, DG Department, Genève 23, CH-1211, Switzerland

By the end of Run 2 in December 2018, the LHC had seen seven full years of operation and a wealth of knowledge and experience has been built up. The key operational procedures and tools are well established. The understanding of beam dynamics is profound and utilized online by well-honed measurement and correction techniques. Key beam-related systems have been thoroughly optimised and functionality sufficiently enhanced to deal with most of the challenges encountered. Availability has been optimised significantly across all systems. This collected experience will form the initial operational basis for Run 3 and subsequent HL-LHC operation.

A brief review of Run 1 and Run 2 is given below, firstly to outline the progress made, and secondly to highlight the issues encountered and surmounted along the way. A synthesis of operational features of the machine and the lessons learnt is then presented. The chapter concludes with brief look at consolidation activities in view of the need to sustain high availability and safe operation given the considerable challenges of the HL-LHC operational regime and the time-frame over which it will operate.

1. Overview of Run 1

Following the recovery from the September 2008 incident, Run 1 saw initial commissioning at reduced energy and the inevitable problems of bootstrapping the operations of a 27 km superconducting collider. Nonetheless, having bedded in the core operational and machine protection systems, healthy levels of performance were achieved. A brief overview of 2010 – 2013 operations follows, which aims to highlight the main issues addressed.

1.1. *2010*

Essentially, 2010 was devoted to commissioning and establishing confidence in operational procedures and the machine protection system. At this stage the operational basics were sorted out while climbing a steep learning curve.

Ramp commissioning to 3.5 TeV was smooth and led to very public first collisions at 3.5 TeV unsqueezed on the 30th March 2010 (see Figure 1). Squeeze commissioning subsequently reduced the β^* to 2.0 m in all the four main experiments. After the squeeze was commissioned, there was a period of Stable Beams interleaved with continued system commissioning.

The decision was then taken to operate with bunches of nominal intensity. Consequently, there was a halting push through the introduction of nominal bunch intensity and further operational debugging up to a total stored beam energy of around 1 to 3 MJ. This led, eventually, to a period of steady running that was used to fully verify machine protection and operational procedures.

To increase the number of bunches, the move to 150 ns bunch trains was made and the crossing angles across the interaction regions were deployed. A phased increase in total intensity was then performed. Each step-up in intensity was followed by operational and machine protection validation and

Fig. 1. Tense times in the control room on 30th March 2010 on the occasion of first high energy colliding beams in the LHC.

a few day running period to check system performance. The 2010 proton run finished with beams of 368 bunches of around 1.2×10^{11} protons per bunch, and a peak luminosity of 2.1×10^{32} cm^{-2}s^{-1}. The operational year ended with a successful four week lead-lead ion run.

1.2. *2011*

The beam energy remained at 3.5 TeV in 2011 and the year saw combined exploitation and the exploration of performance limits. Following a ramp-up to around 200 bunches (75 ns bunch spacing) taking about 2 weeks, there was a scrubbing run of 10 days which included 50 ns injection commissioning. After an encouraging performance, the decision was made, to operate with 50 ns bunch spacing, and a staged ramp-up in the number of bunches then took place up to a maximum of 1380 bunches.

Having raised the number of bunches to 1380, performance was further increased by reducing the emittances of the beams delivered by the injectors and by gently increasing the bunch intensity. The result was a peak luminosity of 2.4×10^{33} cm^{-2}s^{-1} and some healthy delivery rates, topping at 90 pb^{-1} in 24 hours.

A reduction in β^* in ATLAS and CMS from 1.5 m to 1 m delivered the next step up in peak luminosity. This step was made possible by careful measurements of the available aperture in the interaction regions concerned. These measurements revealed excellent aperture consistent with a very good alignment and close to design mechanical tolerances. The reduction in β^* and further gentle increases in bunch intensity produced a peak luminosity of 3.8×10^{33} cm^{-2}s^{-1}, well beyond expectations at the start of the year.

1.3. *2012 and 2013*

2012 was a production year at an increased beam energy of 4 TeV. The choice was made to continue to exploit 50 ns bunch spacing and run with a total number of bunches of around 1380. Based on the experience of 2011, the decision was taken to operate with tight collimator settings, which allowed a more aggressive squeeze to a β^* of 0.6 m. Peak luminosity got up close to its peak pretty quickly. This was followed by determined and long running attempts to improve peak performance. This was successful to a certain extent, revealing some interesting issues at high bunch and total beam intensity, but had little effect on integrated rates. Beam instabilities, although never debilitating, were a recurring problem and there were phases when they cut into operational efficiency.

It was a very long operational years and included the extension of the proton-proton run until December resulting in the shift of a four week proton-lead run to 2013. Integrated rates were healthy at around the 1 fb^{-1} per week level and this allowed a total for the year of about 23 fb^{-1} to be delivered to both ATLAS and CMS, who had, on the back of the data delivered in 2011 and the first half of 2012, announced the discovery of the Higgs boson on the 4th July 2012.

Fig. 2. Lyn Evans accepting the plaudits in CERN's main auditorium on 4th July 2012 following the announcement of the Higgs boson discovery by ATLAS and CMS.

1.4. *Long Shutdown 1 (LS1)*

The primary aim of LS1 (2013 to 2014) was the consolidation of the superconducting splices in the magnet interconnects following the incident of 2008. The successful completion of this work allowed, in principle, the current in the main dipole and quadrupole circuits to be increased to the nominal value for 7 TeV operation. The subsequent main dipole magnet training campaign confirmed systematic de-training and the need for a very long training programme to get to 7 TeV, and the decision was taken to operate the machine at a beam energy of 6.5 TeV during Run 2. Besides splice consolidation, a significant amount of maintenance and other consolidation work was performed on all accelerator systems.

2. Overview of Run 2

Important milestones were reached by the LHC during Run 2 and these included the demonstration of reliable operation with 6.5 TeV beams and exploitation with 25 ns bunch spacing and over 2500 bunches. The design luminosity of 1×10^{34} cm^{-2}s^{-1} was passed and a peak of 2.1×10^{33} cm^{-2}s^{-1} reached. Around 160 fb^{-1} was delivered to ATLAS and CMS, along with 6.7 fb^{-1} to LHCb and 33 pb^{-1} to ALICE.

2.1. *2015*

The principle aims were to re-commission the machine without beam following the major consolidation and upgrades that took place during LS1, and, from a beam perspective, to safely establish operations at 6.5 TeV with 25 ns bunch spacing. The beam configuration targeted was close to nominal i.e. 25 ns bunch spacing with around 2800 bunches of near nominal bunch intensity (1.15×10^{11} protons per bunch). A relatively relaxed β^* of 80 cm in ATLAS and CMS was chosen to provide some aperture margin in the Inner Triplets and thereby less rigorous demands on the collimator settings were required to protect said aperture.

Recommissioning at 6.5 TeV with a bunch spacing of 25 ns was anticipated to be more of a challenge than previous operations at 4 TeV with 50 ns beams. The increased energy implies lower quench margins and thus lower tolerance to beam loss. The hardware (beam dumps, power converters, magnets) is pushed closer to maximum with potential knock-on effects to availability. 25 ns beam was anticipated to have significantly higher electron-cloud than that experienced with 50 ns. It also implies higher total beam current and also higher intensity per injection.

UFOs (“Unidentified Falling Objects”) are micrometer sized dust particles that lead to fast, localized beam losses when they interact with the beam. The phenomenon had already appeared during Run 1 and they were expected to become more of an issue at higher energy. All of these factors came into play in 2015, making for a challenging year.

Two scrubbing runs delivered good beam conditions for around 1500 bunches per beam after a concerted campaign to re-condition the beam vacuum. However, electron cloud, as anticipated, was still significant at the end of the scrubbing campaign.

The initial 50 ns and 25 ns intensity ramp-up phase was tough, having to contend with a number of issues, including magnet circuit earth faults, UFOs, an unidentified aperture restriction in a main dipole, and radiation affecting specific electronic components in the tunnel. Combined, these problems made operations difficult during this phase but nonetheless the LHC was still able to operate with up to 460 bunches and to deliver some luminosity to the experiments albeit with poor efficiency.

The second phase of the ramp-up following a technical stop at the start of September was dominated by the electron cloud generated heat load and the subsequent challenge for cryogenics, which had to wrestle with transients and operation close to their cooling power limits. The ramp-up in number of bunches was consequently slow but steady, culminating in the final figure for the year of 2244 bunches per beam.

The overall machine availability was respectable with around 32% of the scheduled time spent in Stable Beams during the final period of proton-proton physics from September to November. By the end of the 2015 proton run, 2244 bunches per beam were giving peak luminosities of 5.5×10^{33} $cm^{-2}s^{-1}$ in the high luminosity experiments with a total delivered integrated luminosity of around 4 fb^{-1} delivered to both ATLAS and CMS. Levelled luminosity of 3×10^{32} $cm^{-2}s^{-1}$ in LHCb and 5×10^{30} $cm^{-2}s^{-1}$ in ALICE was provided throughout the run.

2.2. *2016 – 2018*

2016 started with four weeks of relatively smooth commissioning with beam with the machine fully validated for $\beta^* = 40$ cm. The first part of the operating period was hit by a number of serious problems in both the LHC and the injectors – in particular a leak from a cooling circuit to the beam vacuum in the SPS beam dump which limited the beam intensity to the LHC. However, after recovery from the main LHC problems, things progressed well. The number of bunches was increased to 2040 per beam – the maximum with the SPS limit of 72 bunches per injection. A bunch population of 1.1×10^{11}

gave a peak luminosity of $\approx 8 \times 10^{33}$ cm^{-2}s^{-1}. Design luminosity was reached on the 26th June thanks to the reduced β^* and lower transverse beam sizes from the injectors, following significant effort to optimise beam brightness via: continuous optimisation; the change of the PS Booster's working point; and the deployment of the batch compression, merging and splitting (BCMS) scheme in the PS.[2] An increase in the peak luminosity of around +20% and a new record of 1.2×10^{34} cm^{-2}s^{-1} was obtained as a result.

The smaller emittances allowed the reduction of the crossing angle from 370 μrad to 280 μrad and a concomitant increase in the geometrical reduction factor from around 0.59 to 0.70. Performance was also helped by the use of a reduced bunch length in Stable Beams. Thus, despite the limit in the number of bunches and a limit in bunch intensity from injection kicker vacuum issues, the peak performance of 40 – 50% over nominal was obtained.

2016 was also blessed by unprecedented machine availability: the machine was available for operation 72% of the time scheduled for physics. Overall Stable Beam efficiency was of order 49% (to be compared to 36% in 2012, and 30% for the short production period in 2015).

2017 saw a further reduction in beam size at the interaction point (β^* = 30 cm), which, together with small beams from the injectors, gave a peak luminosity of 2.1×10^{34} cm^{-2}s^{-1}. Despite the effects of an accidental ingress of air into the beam vacuum during the winter technical stop, referred to as "16L2" after the location of the contamination, around 50 fb^{-1} was delivered to ATLAS and CMS.

2018 essentially followed the set-up of 2017 with a squeeze with ATS optics[3] to 30 cm in ATLAS and CMS. Soon after the intensity ramp up the debilitating effects of 16L2 returned, limiting the maximum bunch intensity to approximately 1.2×10^{11} protons per bunch.

Despite the limitation from 16L2, the peak luminosity was systematically close to the 2×10^{34} cm^{-2}s^{-1} and somewhat more integrated luminosity was possible thanks to the levelling strategy pursued:

- continuous crossing angle reduction ("anti-levelling") in Stable Beams, from an initial 160 μrad smoothly to 130 μrad as a function of the beam current;
- β^* levelling: for the first time the LHC was operated with a dynamically changed optics in Stable Beams, with the β^* in ATLAS and CMS being reduced from 30 cm to 27 cm to 25 cm while colliding.

3. Performance

3.1. *Run 1*

One of the main features of operations in Run 1 was the use of the high bunch intensity with 50 ns bunch spacing offered by the injectors. The injector complex has succeeded in delivering beam with significantly higher bunch intensities with lower emittances than nominal. This is particularly significant for the 50 ns beam. Happily the LHC was capable of absorbing these brighter beams, notably from a beam-beam perspective. The clear cost was increased pile-up for the high luminosity experiments, which they successfully dealt with.

The corresponding values for the main luminosity related parameters at the peak performance of the LHC through the years are shown in Table 1. The design report values are shown for comparison. Remembering that the beam size is naturally larger at lower energy, it can be seen that the LHC has achieved 77% of design luminosity at 4 sevenths of the design energy with a β^* of 0.6 m (cf. design value of 0.55 m) with half nominal number of bunches.

Table 1. Run 1: Proton performance related parameter overview.

Parameter	2010	2011	2012	Design value
Energy [TeV]	3.5	3.5	4	7
Bunch spacing [ns]	150	75/50	50	25
Number of bunches	368	1380	1380	2808
Bunch population (10^{11})	1.2×10^{11}	1.45×10^{11}	1.7×10^{11}	1.15×10^{11}
β^* in IP 1 and 5 [m]	3.5	1.0	0.6	0.55
Normalised emittance (μm)	$\approx$2.0	$\approx$2.4	$\approx$2.5	3.75
Peak luminosity [cm^{-2}s^{-1}]	2.1×10^{32}	3.7×10^{33}	7.7×10^{33}	1×10^{34}
Pileup	4	17	37	19
Stored beam energy [MJ]	$\approx$28	$\approx$110	$\approx$140	362

3.2. *Run 2*

Following a conservative and indeed difficult 2015, peak luminosity in ATLAS and CMS was resolutely pushed throughout the run, principally by:

- a staged reduction of the β^* down to 30 cm at the start of Stable Beams;

- operational use of luminosity levelling via separation, crossing angle reduction and change of β^* – all during Stable Beams;
- provision of high-brightness beams from the injectors (BCMS).

This resulted in a peak luminosity of over twice design and was in fact limited there by the cryogenic cooling capacity of the inner triplets.

Table 2. Run 2: Proton performance related parameter overview.

Parameter	2015	2016	2017	2018
Energy (TeV)	6.5	6.5	6.5	6.5
No. of bunches	2244	2220	2556 – 1868	2556
No. of bunches per train	144	96	144 – 128	144
Bunch population (10^{11})	1.2	1.25	1.25	1.1
β^* [cm] in IP 1 and 5 [cm]	80		40	$40 \to 30 \to 27 \to 25$
Normalised emittance [μm]	2.6 – 3.5	1.8 – 2	1.8 – 2.2	1.8 – 2.2
Peak Luminosity [$\text{cm}^{-2}\text{s}^{-1}$]	0.6×10^{34}	1.5×10^{34}	2.0×10^{34}	2.1×10^{34}
Half Crossing Angle (μrad)	185	$185 \to 140$	$150 \to 120$	$160 \to 130$

This peak performance was accompanied by impressive availability and a low level of premature dumps following a concerted program of measures outlined in more detail below (5).

The resultant integrated luminosity delivered to ATLAS and CMS is shown in Figure 3.

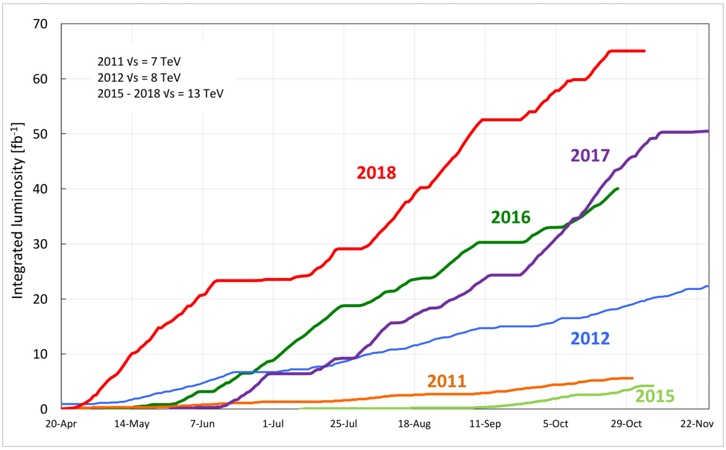

Fig. 3. Average integrated luminosity delivered to ATLAS and CMS during Run 1 and 2.

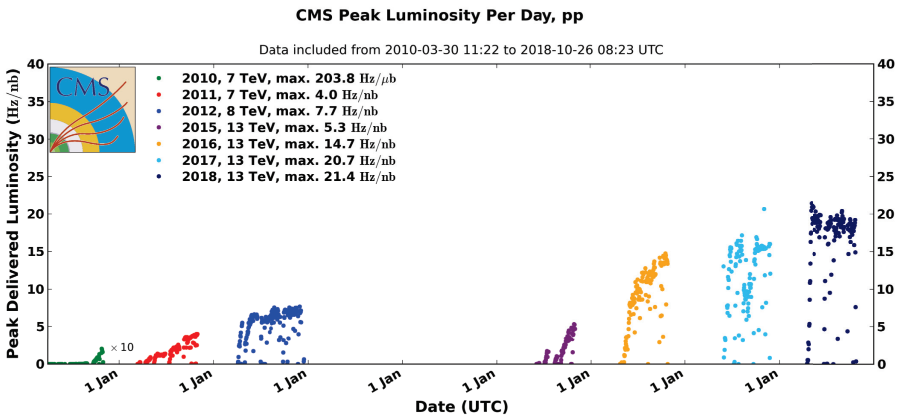

Fig. 4. CMS peak luminosity by day 2010 – 2018.

CMS's peak luminosity by day is shown in Figure 4. This illustrates nicely the results of all the measures outlined above.

An interesting snapshot of the LHC's overall performance during Run 1 and Run 2 is given by ATLAS's collection of performance records as of the end of 2018 – see Figure 5.

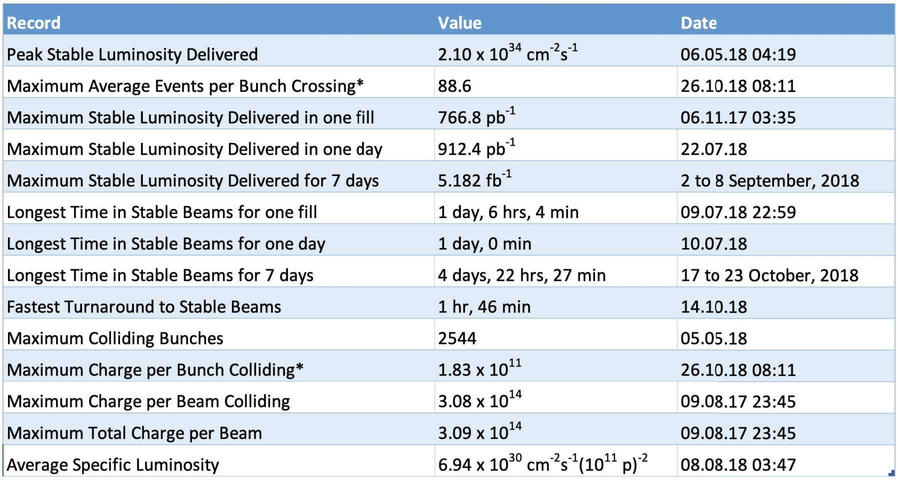

Record	Value	Date
Peak Stable Luminosity Delivered	2.10 x 10^{34} $cm^{-2}s^{-1}$	06.05.18 04:19
Maximum Average Events per Bunch Crossing*	88.6	26.10.18 08:11
Maximum Stable Luminosity Delivered in one fill	766.8 pb^{-1}	06.11.17 03:35
Maximum Stable Luminosity Delivered in one day	912.4 pb^{-1}	22.07.18
Maximum Stable Luminosity Delivered for 7 days	5.182 fb^{-1}	2 to 8 September, 2018
Longest Time in Stable Beams for one fill	1 day, 6 hrs, 4 min	09.07.18 22:59
Longest Time in Stable Beams for one day	1 day, 0 min	10.07.18
Longest Time in Stable Beams for 7 days	4 days, 22 hrs, 27 min	17 to 23 October, 2018
Fastest Turnaround to Stable Beams	1 hr, 46 min	14.10.18
Maximum Colliding Bunches	2544	05.05.18
Maximum Charge per Bunch Colliding*	1.83 x 10^{11}	26.10.18 08:11
Maximum Charge per Beam Colliding	3.08 x 10^{14}	09.08.17 23:45
Maximum Total Charge per Beam	3.09 x 10^{14}	09.08.17 23:45
Average Specific Luminosity	6.94 x 10^{30} $cm^{-2}s^{-1}(10^{11}\ p)^{-2}$	08.08.18 03:47

Fig. 5. LHC performance records at the end of 2018 as noted by ATLAS. * indicates a record achieved during machine development.

3.3. *Other users*

Throughout Run 1 and Run 2, the operational flexibility of the LHC has allowed the pursuit of a rich variety of physics programmes ranging through lead-lead, lead-proton, xenon-xenon, and an interesting, and sometimes demanding, forward physics programme.

The time limited ion programme inevitably represents a challenge for LHC operations.[4] The team has to commission new configurations and provide stable physics operation within time frame of one month and meet demanding requirements from the experiments which include multiple changes of beam conditions (intensity ramp-up, solenoid reversal, beam reversal, low/high/levelled luminosity, special beam energies, Van der Meer scans). Nonetheless, heavy-ion operation of LHC has surpassed initial expectations, both quantitatively (3.5 times design luminosity after about 10 weeks of Pb-Pb operation since 2010) and qualitatively (asymmetric p-Pb collisions, unforeseen in the design, have yielded almost 6 times their nominal luminosity and a rich harvest of unexpected physics results). The fact that it has been possible to rapidly recommission the LHC in multiple new configurations efficiently is testament

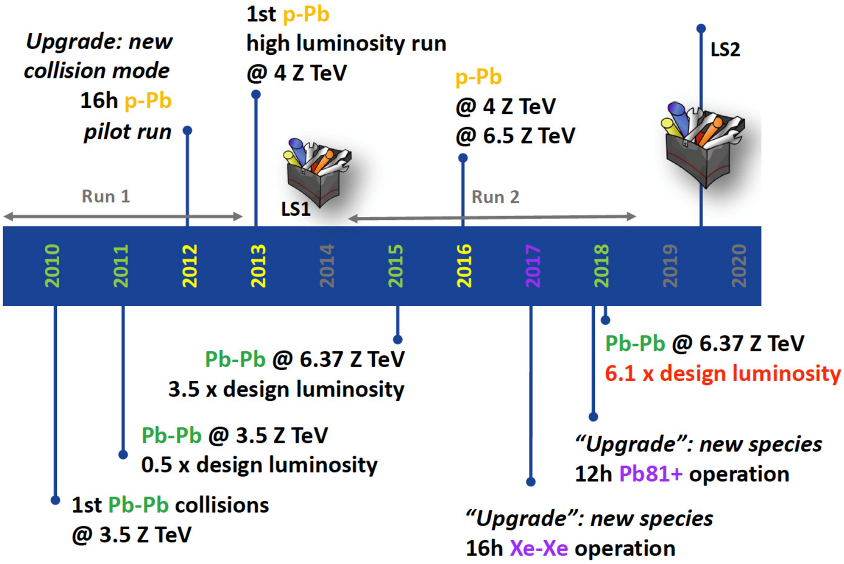

Fig. 6. Timeline of the heavy-ion runs during Run 1 and Run 2. Figure courtesy John Jowett and Michaela Schaumann.[5]

to the understanding and level of control that has been established; the salient points are summarized below.

4. Overview of LHC Operational Characteristics

The performance described above is built on the back of some excellent system performance and some fundamental operational characteristics of the LHC. Very good understanding of the beam physics and a good level of operational control was established and the following features related to beam-based operation may be elucidated.

- The linear optics is well measured and is remarkably close to the machine model. The bare beta-beating is acceptable and has been corrected to excellent. The availability of multi-turn orbit measurements and impressive analysis tools should be noted.
- There is excellent single beam lifetime and on the whole the LHC enjoys very good vacuum conditions.
- Head-on beam-beam is not a limitation although long-range beam-beam has to be taken seriously with enough separation at the long-range encounters guaranteed by sufficiently large crossing angles. The tolerance to high head-on beam-beam tune shifts can be partially attributed to: well-corrected lattice errors, via both an excellent magnet model and a superb optics measurement and correction programme; low external noise, and other perturbations. A full analysis may be found at Ref. [1].
- Better than nominal beam intensity and beam emittance is delivered by the injectors and it has proved possible to collide nominal bunch currents with smaller that nominal emittances with no serious problems.
- Collective effects have been seen with high bunch intensities and with nominal bunch intensities in the presence of electron cloud. Single and coupled bunch instabilities have been suppressed using a range of tools (high chromaticity, Landau damping octupoles and transverse feedback).
- There is better than expected aperture due to good alignment and respect of mechanical tolerances.

- There is excellent field quality, coupled with good correction of nonlinearities. The magnetic machine is well understood and the modelling of all magnet types has delivered an excellent field description at all energies. This model includes persistent current effects which have been fully corrected throughout the cycle.
- A strict pre-cycling regime means the magnetic machine is remarkably reproducible. This is reflected in the optics, orbit, collimator set-up, tune and chromaticity. Importantly orbit stability (or the ability to consistently correct back to a reference) means that collimator set-up remains good for a year's run.
- There is low tune modulation, low power converter ripple, and low RF noise. Power converters are delivering remarkably stable and accurate currents ranging from single digits to several thousand amps. Tracking between power converters in the ramp and squeeze is exceptional and the whole system is complemented by a very good front-end control system.
- Efficient, stable, operating procedures and supporting software are in place.

5. Operational Cycle and Availability

The nominal operation cycle provides the framework driving luminosity production. Given the high stored beam energy, the nominal cycle must be fully mastered for effective, safe operation. As of Run 2, the operational cycle was well established for 50 and 25 ns and bunch population exceeding nominal.

The turnaround time is defined as the time taken to go from the dump of a physics fill at top energy back into colliding beams following a refill. Following converted effort over the years and numerous operational improvements, by 2018 the minimum turnaround time had been reduced to around 110 minutes.

Availability is defined as the overall percentage of the scheduled machine time left to execute the planned physics program after removing the total time dedicated to fault resolution. Faults cover an enormous range from a simple front-end computer reboot to the loss of a cold compressor of the cryogenics system with a corresponding loss of time to operations from 10 minutes to potentially days. Availability has, in general, been excellent considering the size, complexity, and operating principles of the LHC. The percentage of scheduled proton-proton physics time spent delivering collisions to the experiments ("Stable Beams") was around 36% in 2012. Following a prolonged campaign of consolidation and targeted system improvements, the corresponding number in 2017 and 2018 was around 50%. As of the end of Run 2, there is good overall system performance and availability based on solid foundations and vigorous follow-up of problems. This is the result of a sustained, targeted effort across the board by all teams, backed by effective fault tracking. Beam related issues such as radiation to electronics, UFOs, beam induced heating have all been relentlessly addressed.

Operations also depends heavily on the superb performance of machine protection and associated systems. These include the beam interlock system, the beam dump system, the beam loss monitors, and the collimation system. There is rigorous machine protection follow-up, qualification, and monitoring; all non-conformities are carefully examined. The importance of this to the success of the LHC so far cannot be over stressed and due credit must be given to the teams involved for ensuring the safety of the machine during beam based operation over the two runs.

Remarkable operational flexibility has been demonstrated, and allowed the team to handle, for example, the slower than expected electron cloud conditioning, and the effects of the accidental air ingress in Sector 12 – the now infamous 16L2.

6. Issues

There have inevitably been a number of challenges during the exploitation of the LHC. Initially, single event effects (SEEs) caused by beam induced radiation to tunnel electronics was a serious cause of inefficiency. However, this problem had been foreseen and its impact was considerably reduced following sustained program of mitigation measures. There were several shielding campaigns prior to the 2011 run including relocation "on the fly" and equipment upgrades. The 2011/12 Christmas stop saw some "early" relocation and additional shielding and further equipment upgrades. Further improvement followed an extensive campaign of relocation, shielding, and hardware upgrades during LS1.

6.1. *UFOs*

UFOs (Unidentified Falling Objects) are microscopic particles of the order of 10 microns across. These fall from the top of the vacuum chamber or beam screen, become ionised by collisions with circulating protons and then are repelled by the positively charged beam. While interacting with the circulating protons they generate localised beam loss which may be sufficient to dump the beam or, in the limit, cause a quench. They have now been very well studied and simulated. There were occasional dumps in 2012 following adjustment of BLM thresholds at the appropriate time-scales (the beam loss spike caused by a UFO is typically of order 1 ms). With the increase in energy to 6.5 TeV and the move to 25 ns the UFOs become harder (energy) and more frequent (25 ns). Indeed, during the first half of 2015 they were a serious issue but happily there was conditioning and the UFO rate fell to acceptable levels as the year progressed. It should also be noted that it was fortunate that UFO rates have conditioned down, accompanied, as elsewhere, by excellent diagnostics, well thought through mitigation actions and understanding through simulation.

6.2. *Beam induced heating*

Beam induced heating has been an issue and essentially all cases have been local and, in some way, due to non-conformities either in design or installation. The guilty parties have been clearly identified. Design problems have affected the injection protection devices and the mirror assemblies of the synchrotron radiation telescopes. Installation problem have occurred in a low number of vacuum assemblies. These singularities have all been addressed and the issue is not expected to be problem in the long term.

6.3. *Beam instabilities*

Beam instabilities were an interesting problem that dogged operations through 2012. It should be noted that this problem paralleled a gentle push in bunch intensity with the peak going into stable beams reaching around 1.7×10^{11} protons per bunch i.e. ultimate bunch intensity. In 2015 operations with 25 ns bunch spacing and lower bunch population meant that intrinsically instabilities should have been less of an issue. However, high electron cloud proved to be a driver and defence mechanisms were deployed in the form of high chromaticity, high octupole field strength and the transverse damper system.

6.4. *Electron Cloud*

Electron cloud is the result of an avalanche-like process in which electrons from gas ionisation or photo-emission are accelerated in the electromagnetic field of the beam and hit the beam chamber walls with energies of few hundreds of eV, producing more electrons. The electron impact on the chamber wall causes gas desorption as well as heat load for the cryogenic system in the cold regions. High electron densities in the beam chamber can lead to beam oscillations and blow-up of the particle bunches due to the electromagnetic interaction between electrons and protons. Electron bombardment of a surface has been proven to reduce drastically the secondary electron yield (SEY) of a material. In a process known as scrubbing, deliberate invocation of high electron cloud with beam thus provides a means to reduce or suppress subsequent electron cloud build-up.

Although electron cloud was not an issue with 50 ns beam, 25 ns operation proved to be a challenge in 2015, and extensive scrubbing – both dedicated at low energy and while delivering collisions to the experiments – was required. Conditioning thereafter has been slow and the heat load from electron cloud to the cryogenics system remained a limitation in 2018.

7. Conclusions

After seven full years of operation, in the beam parameter regime concerned, the extended LHC team has managed to develop an impressive mastery of the LHC and the delivery of the requisite beam from the injectors. A concise summary of the salient observations is attempted below.

- Good peak luminosity via exploitation of all available parameters (β^*, bunch population, bunch length, crossing angle, transverse emittance).
- Stunning availability following sustained effort from hardware groups accompanied by effective fault tracking.
- Few premature dumps allowing long fills: the UFO rate conditioned down and radiation to electronics effects have been largely mitigated, again after a sustained and successful campaigns.

- Excellent and improved system performance across the board, for example, the new developments of the transverse damper system; collimator alignment software; improved injection kicker performance via hardware modifications.
- The magnets, circuits and associated systems are behaving well at 6.5 TeV.
- Good beam lifetime through injection, ramp, and squeeze with tight control of tune and closed orbit, reflecting that operationally things are very well under control.
- Excellent luminosity lifetime in general with only moderate emittance blow-up in Stable Beams and minimal non-luminosity beam loss after the first hour or so.
- Well established and tuned magnet model, good compensation of persistent current decay and snapback, which couple with a strict magnet cycling give excellent magnetic reproducibility.
- The optics of the machine has been measured and corrected to a impressive level, both linear and higher orders, and a superb level of understanding has been established.
- Aperture is fine and compatible with the collimation hierarchy.
- The collimation system has consistently demonstrated excellent performance and impressive robustness.
- A reliable and well designed machine protection system coupled with a disciplined regime has assured safe exploitation.

2016 was really the first year when it all came together: injectors; operational efficiency; system performance; understanding and control; and availability. In 2017, and 2018, the LHC was able to build on this to move into a true exploitation regime, accompanied, as always, by continued efforts to improve integrated luminosity delivery.

The LHC has moved haltingly from commissioning to exploitation, and is now enjoying the benefits of the decades long international design, construction, and installation effort – it's clear that the foundations and fundamentals are good. It's present performance is worthy reflection of this effort and the huge amount of experience and understanding gained and fed-forward over the last years. Remarkably, not only can a 27 km superconducting collider work, it can work well!

8. Consolidation in view of HL-LHC

The Accelerators & Technology sector strives for a maximum reliability and availability of the whole accelerator complex and the associated experimental areas. Given the age, complexity, and operational lifetime of the complex, the provision of spares and consolidation of the existing equipment and associated technical infrastructure are key issues in ensuring that the needs of the Organisation's diverse physics program are assured. This, of course, includes the flagship LHC programme, which is and will be intimately dependant on the performance of the injectors.

The ongoing consolidation programme consists in the replacement (or renovation) of equipment and related technical infrastructure at the end of the normal lifetime, i.e. when:

- An equipment can no longer be used with sufficient reliability;
- The equipment has been exposed to levels of radiation that compromise its functionality;
- Commercially available spare parts are lacking;
- Technical support is no longer available for components or software;
- The systems no longer meets safety regulations and standards.

The performance and availability of all technical systems is actively monitored with the Accelerator Fault Tracking system. This can give some indication of potential issues and the need for targeted consolidation. However, consolidation of most major systems has to be anticipated before impacting the performance of the machine. For example, the end of life of electronic components, where there are long lead times for product design, prototyping, tendering, and production. Other factors such as maintenance cost, availability of expertise, availability of spares, standardisation, modern functionality/reliability have also to be taken into consideration.

Typical, during an operational year, active consolidation continues as far as possible during technical stops, along with provision of spares, and development and production of components for a major programme of deployment in the long shutdowns. The LHC consolidation program has over 100 consolidation activities ongoing at any one time, and besides long-term activities, the consolidation program also had to respond to a number of punctual demands resulting from issues arising from regular operations.

On the magnet front, spares for the many types of magnets, both superconducting and room temperature must be assured. Of note, as of 2021, is the ongoing production of five spare main quadrupole magnets. A number of corrector magnets have been lost over the years, and remain in the machine, and are either electrically by-passed or simply removed from operations. In the interest of nominal performance, and to avoid the potential loss of a debilitating number of magnets, plans to recover these circuits in the future long shutdowns should be made.

The machine protection group continues targeted revision and renovation of the key elements of the quench protection and energy extraction systems, in particular the tunnel electronics. Vacuum, heavily implicated around the whole machine, continues a rolling program of upgrades which includes mobile pumping stations, bake-out systems, and other elements. In close collaboration with HL-LHC project, the collimation team has developed and partially deployed the next generation of robust, low impedance collimators.

Beam instrumentation is targeting upgrades of its big distributed systems (beam position monitors and beam loss monitors). This will take several years and in 2017 they started on the BLM system, part of the effort being dedicated to the development of radiation hard front-end electronics for the HL-LHC era. In addition, they perform consolidation and upgrades of a number of stand-alone systems (wire-scanners, bunch current transformers, interlocked BPMs); here the goal is performance enhancing consolidation, taking the opportunity of not only replacing equipment, but also leveraging experience and technology to improve system performance to match the needs of the HL-LHC era.

The use of industrial controls is widespread and a number of teams (e.g cryogenics, cooling and ventilation) are renovating and upgrading their systems. The control group continues maintenance of its fundamental infrastructure (field bus installations, repeaters, timing system, control room hardware).

Technical Infrastructure (cooling, ventilation, electrical distribution, lifts, cranes etc.) has a long-term rolling consolidation program with staged replacement and renovation of the enormous amount of site wide systems. Cooling and ventilation continues renovation of HVAC units in surface buildings and industrial control renovation. Heavy engineering will continue to execute its rolling replacement of lifts, overhead cranes and hoists.

In Long Shutdown 2, the LHC beam dump system (LBDS) saw urgent consolidation of the beam dump blocks following issues in Run 2. This experience will be fed forward into the design of new dump blocks which will be produced for the HL-LHC intensities. The LBDS pulse forming networks (PFNs), switches and electronics are the subject of diligent scrutiny and appropriate consolidation as befits their criticality.

References

1. Sondre Vik Furuseth and Xavier Buffat. Modeling of nonlinear effects due to head-on beam-beam interactions. *Physical Review Accelerators and Beams*, 21(8):081002, August 2018.
2. Heiko Damerau, Steven Hancock, Alexandre Lasheen, and Damien Perrelet. RF Manipulations for Special LHC-Type Beams in the CERN PS. In *9th International Particle Accelerator Conference*, 6 2018.
3. S. Fartoukh *et al.*, *Achromatic telescopic squeezing scheme and by-products: From concept to validation*, Phys. Rev. Accel. Beams **24**, 021002.
4. J. Jowett, Colliding Heavy Ions in the LHC, 9th International Particle Accelerator Conference doi:10.18429/JACoW-IPAC2018-TUXGBD2
5. Jowett, John M. and Schaumann, Michaela, Overview of Heavy Ions in LHC Run 2 Proceedings of the 9th LHC Operations Evian Workshop CERN-ACC-2019-059, 15-25 2019.