Nuclear Instruments and Methods in Physics Research A 848 (2017) 19–30

Contents lists available at ScienceDirect

Nuclear Instruments and Methods in Physics Research A

journal homepage: www.elsevier.com/locate/nima

Reaching record-low β^* at the CERN Large Hadron Collider using a novel scheme of collimator settings and optics

R. Bruce[a,*], C. Bracco[a], R. De Maria[a], M. Giovannozzi[a], A. Mereghetti[a], D. Mirarchi[a], S. Redaelli[a], E. Quaranta[a,b], B. Salvachua[a]

[a] CERN, Geneva, Switzerland
[b] Politecnico di Milano, Milan, Italy

ARTICLE INFO

Keywords:
Circular colliders
Synchrotrons
Collimation
Beam optics
Large Hadron Collider

ABSTRACT

The Large Hadron Collider (LHC) at CERN is built to collide intense proton beams with an unprecedented energy of 7 TeV. The design stored energy per beam of 362 MJ makes the LHC beams highly destructive, so that any beam losses risk to cause quenches of superconducting magnets or damage to accelerator components. Collimators are installed to protect the machine and they define a minimum normalized aperture, below which no other element is allowed. This imposes a limit on the achievable luminosity, since when squeezing β^* (the β-function at the collision point) to smaller values for increased luminosity, the β-function in the final focusing system increases. This leads to a smaller normalized aperture that risks to go below the allowed collimation aperture. In the first run of the LHC, this was the main limitation on β^*, which was constrained to values above the design specification. In this article, we show through theoretical and experimental studies how tighter collimator openings and a new optics with specific phase-advance constraints allows a β^* as small as 40 cm, a factor 2 smaller than β^*=80 cm used in 2015 and significantly below the design value β^*=55 cm, in spite of a lower beam energy. The proposed configuration with β^*=40 cm has been successfully put into operation and has been used throughout 2016 as the LHC baseline. The decrease in β^* compared to 2015 has been an essential contribution to reaching and surpassing, in 2016, the LHC design luminosity for the first time, and to accumulating a record-high integrated luminosity of around 40 fb^{-1} in one year, in spite of using less bunches than in the design.

1. Introduction

The Large Hadron Collider (LHC) [1,2] is a 27 km synchrotron, designed to collide two counter-rotating beams (called B1 and B2) of 7 TeV protons, or heavy ions of equivalent magnetic rigidity, at four interaction points (IPs) inside experimental detectors. These are installed in four of the eight straight insertion regions (IRs). Two general-purpose experiments require high luminosity, ATLAS [3] and CMS [4], which are installed in IR1 and IR5 respectively. The schematic layout of the LHC is shown in Fig. 1 and the key design parameters for proton operation are shown in Table 1 together with what has been achieved so far in operation.

The LHC beams are controlled by superconducting magnets. Many of them have a working temperature of 1.9 K but risk to quench even due to a temperature rise as small as 1.4 K [6]. At the same time, the LHC beams have an unprecedented stored energy, targeting 362 MJ in the design case (see Table 1), so a local loss of even a tiny fraction of the beam could heat the magnets enough to induce a quench [6]. Therefore, a multi-stage collimation system has been installed in order to intercept any regular or accidental beam losses [1,7–11].

Most LHC collimators are made of two movable jaws, where the active part is between 60 cm and 1 m long. Robust collimators close to the beam are made of carbon fiber composite (CFC), while others that are further out are made of tungsten, which has a much higher stopping power but is also more fragile. The collimators are ordered in a strict hierarchy in terms of the half-gap in units of the local betatron beam size σ, calculated as

$$\sigma = \sqrt{\beta \epsilon_n / \gamma_{\mathrm{rel}}}, \tag{1}$$

where we use the convention to take the nominal β-functions[1] and a normalized emittance ϵ_n=3.5 μm unless stated otherwise. The betatron collimation system, installed in IR7, uses primary collimators (TCP) closest to the beam. They are followed by secondary collimators

* Corresponding author.
E-mail address: roderik.bruce@cern.ch (R. Bruce).

[1] The achieved optics corrections at top energy are in general very good and the error on the measured optics is often comparable to the deviation from the nominal [12–14].

http://dx.doi.org/10.1016/j.nima.2016.12.039
Received 20 December 2016; Accepted 21 December 2016
Available online 23 December 2016
0168-9002/© 2016 The Author(s). Published by Elsevier B.V.
This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/BY-NC-ND/4.0/).

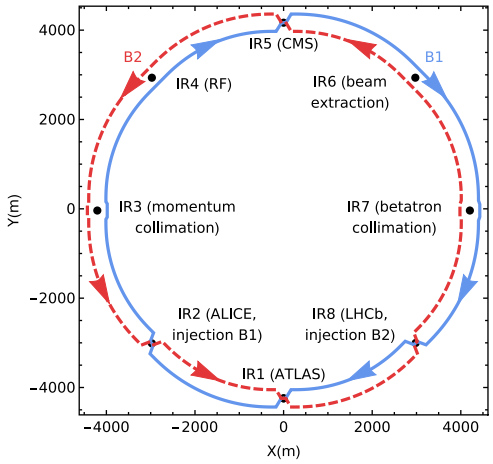

Fig. 1. The schematic layout of the LHC (the separation of the two rings is exaggerated). The two beams are brought into collision at the four experiments ATLAS, ALICE, CMS and LHCb. Adapted from Ref. [5].

initially conservative approach was taken with rather open collimator settings [21]. Later on, a statistical approach was developed in order to reduce the margins, while still keeping the risk of exposing sensitive elements very low [22]. Using this approach, the collimator retractions were calculated based on a probability distribution of drifts of the orbit, optics and other machine imperfections, which were estimated from data in previous running periods. This allowed to reduce β^* significantly in 2011 [23] and further in 2012 to 60 cm [24], i.e. almost down to the nominal value of 55 cm (see Table 1) in spite of the lower energy. This relied also on a tighter TCP cut than in 2011, smaller errors on the available aperture than pessimistically assumed during the LHC design [25], and a smaller beam–beam separation [26].

The same principles for positioning collimators were used initially in Run 2, when the LHC was restarted in 2015 after a 2-year shutdown [27]. The main changes in Run 2 were an increase of the beam energy from 4 TeV to 6.5 TeV and a shorter 25 ns bunch spacing (see Table 1). The TCPs were kept at the same setting in mm as in 2012, which lead to a larger cut in σ due to the higher energy. In addition, 2 σ extra safety margin was taken for the first run at the higher energy of 6.5 TeV, compared to 4 TeV in Run 1 [28]. Because of this, and the larger crossing angle required to compensate for a stronger long-range beam–beam effect with the shorter bunch spacing [29], the LHC was operated in 2015 with β^*=80 cm.

Table 1
Typical proton running conditions in the LHC during operation so far, shown together with the design parameters. The values of luminosity, crossing angle, beam–beam separation, geometric reduction factor, and number of colliding bunches, refer to the high-luminosity experiments in IR1 and IR5 only.

Machine parameter	Design	2010	2011	2012	2015	2016
Beam energy (TeV)	7.0	3.5	3.5	4.0	6.5	6.5
Protons/bunch (average at start of collisions) (10^{11} p)	1.15	1.0	1.3	1.5	1.0	1.1
Max. number of bunches	2808	368	1380	1380	2244	2200
Max. stored energy per beam (MJ)	362	23	112	143	277	272
Bunch spacing (ns)	25	150	50	50	25	25
Transv. normalized emittance ϵ_n, typical value in collision (μm)	3.75	2.6	2.4	2.4	3.5	2.2
Half crossing angle (μrad)	143	100	120	145	145	140/185
Primary collimator (TCP) cut (σ), for ϵ_n=3.5 μm	6.0	5.7	5.7	4.3	5.5	5.5
Secondary collimator (TCSG) cut (σ), for ϵ_n=3.5 μm	7.0	8.5	8.5	6.3	8.0	7.5
Tightest dump protection (TCSP/TCDQ) cut (σ), for ϵ_n=3.5 μm	7.5	9.3	9.3	7.1	9.1	8.3
Tertiary collimator cut (σ), for ϵ_n=3.5 μm	8.3	15.0	11.8	9.0	13.7	9.0
Smallest allowed magnet aperture (σ), for ϵ_n=3.5 μm	8.4	17.5	14.1	10.5	15.5	10.0
β^* (m)	0.55	3.5	1.0–1.5	0.6	0.8	0.4
Max. peak luminosity (10^{34} cm^{-2} s^{-1})	1.0	0.021	0.35	0.77	0.47	1.45
Total integrated luminosity, average over ATLAS and CMS (fb^{-1})		0.048	5.5	22.8	4.2	39.3

(TCSGs), dump protection (TCSP and TCDQ) in IR6, and tungsten tertiary collimators (TCT) in front of the experiments. They are set to protect the aperture bottlenecks of the ring, which during physics operation are in the inner triplet magnets in the final focusing system of the experimental IRs [15–18]. It should be noted that there are also active absorbers (TCLA) as well as a hierarchy for momentum cleaning in IR3, which we do not discuss in detail here. This hierarchy must be respected at all times, in order to guarantee sufficient protection. Because of drifts in orbit and optics over time, margins in σ are needed between collimator families to ensure that the hierarchy is respected.

All magnets must have larger normalized apertures,[2] in units of σ, than all the collimation stages. If β^* is decreased to gain luminosity, the β-function in the triplets increases, so that their normalized aperture becomes smaller and approaches the opening of the TCTs. In this situation, the whole collimation hierarchy has to be moved in, or the margins between the collimation stages decreased, in order to maintain guaranteed protection. It is thus clear that how much the collimator settings can be tightened, without jeopardizing machine safety, determines how small a value of β^* can be accommodated.

In the first run of the LHC (2010–2013, called Run 1) [19,20], an

The year 2015 can be considered a commissioning year, when operation at 6.5 TeV and 25 ns was established, and rather relaxed machine parameters were used. A total integrated luminosity of 4.2 fb^{-1} was collected. In 2016, however, the luminosity had to be increased significantly in order to meet the goals of producing more than 25 fb^{-1} over the year, and more than 100 fb^{-1}of data in Run 2 (up to the end of 2018). One way to do this, which is independent on encountered limitations on the maximum beam intensity such as electron cloud [30], is to tighten the collimator settings and decrease β^*.

This article shows theoretical studies to make such a low-β^* configuration possible, as well as experimental validations. While several factors contribute to the reach in β^*, we focus here on the one which has given the largest gain in β^* for the LHC in 2016. This is a novel method, where we match the machine optics with special conditions on the betatron phase advance $\Delta\mu$. This allows to significantly reduce collimation margins without jeopardizing the machine protection. This method has been essential for introducing β^*=40 cm as the 2016 baseline, a factor 2 below β^*=80 cm used in 2015. Using β^*=40 cm in 2016 has given a major contribution to achieving and exceeding the nominal LHC luminosity for the first time ever, and to increasing the peak luminosity by about a factor 3 compared to 2015.

[2] We define the normalized aperture as the smallest available space between the beam center and the vacuum chamber, divided by the local transverse beam size in the corresponding plane.

Other contributions to the luminosity gain, not treated in this article, come from a smaller transverse emittance, a higher bunch charge, a smaller normalized beam–beam separation, and a shorter bunch length. These results were achieved although it was not possible to inject the full number of bunches in the LHC (see Table 1), due to restrictions on a beam dump in the injector complex.

2. Studies of losses from beam dump failures in the 2015 machine

Table 1 shows that in 2015, the smallest allowed aperture was 10 σ larger than the TCP cut, and there was 6.4 σ margin between the three last hierarchy stages (TCDQ, TCT and aperture). In order to reduce β^* through tighter collimator settings, it is very important to find ways of further reducing the large margins outside the TCDQ aperture. They were put in place to protect sensitive equipment from miskicked particles during beam dump failures [22,31–33]. During a standard beam dump, 15 horizontal extraction kickers in IR6, called MKDs, fire during a 3 μs abort gap in the filling pattern when no beam is passing.[3] When the next bunch arrives, the MKDs are already at their full strength so that the bunch is kicked into the extraction channel.

However, rare failure modes exist in which one (single-module prefire, SMPF) or all MKDs (asynchronous beam dump, AD) would trigger outside the abort gap. The passing bunches would then be kicked by intermediate MKD strengths, which could deflect them directly onto the machine aperture. During standard physics operation, the sensitive elements with the tightest apertures are the TCTs and the triplets behind them. The TCTs are made of tungsten and hence not robust against high-intensity impacts [34,35]. Damage to accelerator components in such a scenario would be extremely costly, in particular for the triplet, both in terms of equipment and downtime, and must be avoided. The role of the robust TCDQ absorber is to intercept such miskicked beam and the margin between TCDQ and TCTs has to be large enough that the TCTs or triplets can never be hit even if imperfections change the effective cuts in σ.

It has so far been assumed that the dump protection should always be at a smaller normalized opening in σ than the TCTs, independently on the optics, to account for the case in which the TCDQ/TCSP and the horizontal TCTs are at the same $\Delta\mu$=90° from the MKDs [22]. However, this is sometimes a pessimistic assumption, and dropping it could bring a gain. Fig. 2 shows schematically the oscillating orbit caused by a misfiring MKD and that if TCT_1 is placed at a fractional phase advance $\Delta\mu_{TCT}$ from the MKDs close to 0° or 180°, it can be moved in much closer to the beam (in units of σ) than TCT_2 with $\Delta\mu_{TCT}$ close to 90° or 270°, without increasing the risk of damaging impacts.

To understand the potential gain, we simulate first the expected losses during an SMPF, which is more critical than an AD, accounting for the actual $\Delta\mu_{TCT}$, using the 2015 operational configuration. This is done using SixTrack [36–38], where an ensemble of macroparticles is tracked through the magnetic lattice of the machine. It includes also particle–matter interaction inside collimators leading to possible outscattering. Further details on the SixTrack simulation chain are given in Ref. [11], where the results are also shown to give a good agreement with LHC beam loss data.

In SixTrack, we include the misfiring MKDs. The MKD kicks are taken from measured data of the worst mode of SMPF observed, called type 2, which occurred without beam [39,40]. This is the failure mode which has the slowest rise of the total kick, summed over all MKDs, and hence implies the largest amount of protons kicked at amplitudes where they risk to hit sensitive machine elements such as the TCTs. Pessimistically, the most upstream MKD with the largest β-function triggers first, and the others follow after delays of about 1–2 μs, depending on MKD. During the rise of the field, each passing bunch

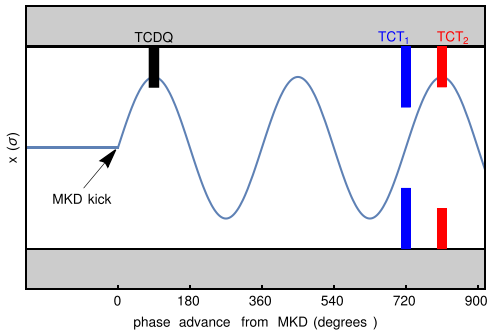

Fig. 2. Schematic illustration of the horizontal orbit in normalized phase space resulting from a misfiring MKD during an asynchronous beam dump, as it performs betatron oscillations along the machine, shown together with the one-sided TCDQ at 90° phase advance from the MKD, and two different TCT positions at $\Delta\mu_{TCT}$=0° and $\Delta\mu_{TCT}$=90°.

encounters a different kick. We therefore simulate several consecutive Gaussian bunches spaced by 25 ns, each encountering different MKD strengths. On the turn after the misfire, it is assumed that all MKDs have reached the full field and remaining particles are extracted.[4] Only a limited number of bunches contribute to the TCT losses, since a bunch kicked at a too small amplitude passes inside the TCT aperture, while one with a large kick is fully intercepted by the TCDQ. Therefore, only about 20 bunches are simulated around this window and the losses from these bunches are summed and normalized to 1.5×10^{11} protons per bunch.[5] The results show that only about 8 bunches give significant contributions. The simulations include a maximum orbit shift of 1.2 mm away from the TCDQ, at the limit of the present interlocks to stay pessimistic, but a perfect optics and orbit otherwise.

Fig. 3 shows an example of the simulated losses around the ring in the 2015 machine (see Table 1). The main loss location is at the TCDQ, with minor losses appearing also at other elements. The highest TCT loss is in IR5 in B2, and the triplets do not intercept any losses due to the TCT protection. However, at tighter TCT settings this changes, as illustrated in Fig. 4. It shows the simulated losses on horizontal TCTs in IR1 and IR5, when everything is kept constant at the 2015 settings except for the TCTs, which are simultaneously closed in steps from their nominal 13.7 σ setting. The losses in Fig. 4 have been split into primary losses, defined as protons that have not previously interacted with any collimator, and secondary losses, which are protons outscattered from upstream collimators. Furthermore, we show results both for perfect collimators, and when imperfections are included (random tilt angles, gap errors, center errors, and surface flatness errors), using the collimator imperfection models and parameters described in Ref. [11]. In each case, 20 random seeds were used, limited by available computing resources, and the result shown is their average. The standard deviation of the TCT losses over the seeds is typically around 30–50% for the tightest TCT openings, while it is sometimes over 100% at larger openings, where imperfections have a stronger influence.

The results are analyzed in view of the 2015 values of $\Delta\mu_{TCT}$ shown in Table 2. On the IR1 TCT in B1, which has $\Delta\mu_{TCT}$ close to 60°, the primary losses increase steeply when it is moved in, while the losses fall off to zero closely outside the TCDQ cut. This TCT is hit also by a significant number of secondary protons, which have scattered in TCSGs in IR7. These impacts increase also at smaller TCT openings, where the IR7 cleaning is less efficient. This is particularly pronounced, since the IR1 TCT is at a $\Delta\mu$ of about 130° from the IR7 TCSG that receives by far the highest number of hits during an SMPF. The secondary losses drop off at larger openings, when the TCT gets further

[3] As a comparison, the revolution period of the LHC is 89μs.

[4] The flat-top part of the MKD pulse length extends over at least 90μs, which is more than one turn and enough to extract all remaining particles [1].

[5] This was chosen to be slightly more pessimistic than the foreseen maximum bunch intensity in Run 2 [41].

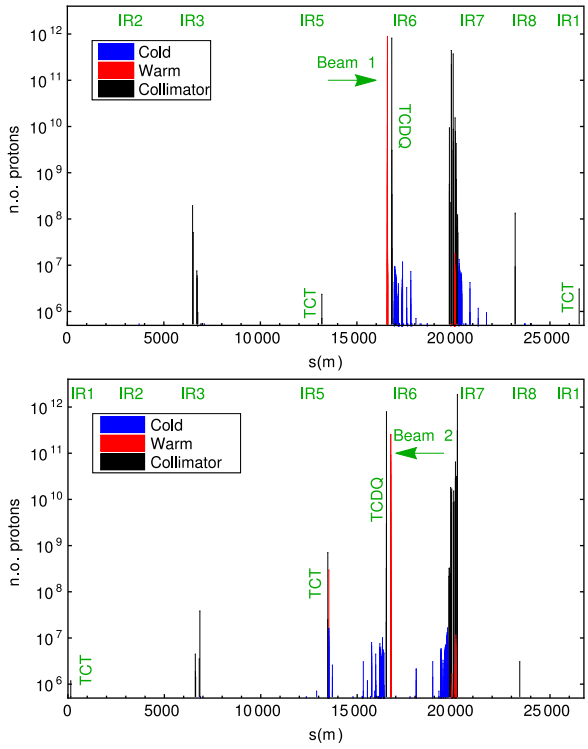

Fig. 3. Loss distribution around the LHC ring for B1 (top) and B2 (bottom) during a SMPF as simulated with SixTrack for the 2015 standard physics configuration (see Table 1). Only 20 bunches, with MKD kicks in the range where parts of the bunches could end up on TCTs, are simulated and summed. If all bunches over the MKD sweep would be included, higher losses would be expected on the TCDQ but not on the TCTs.

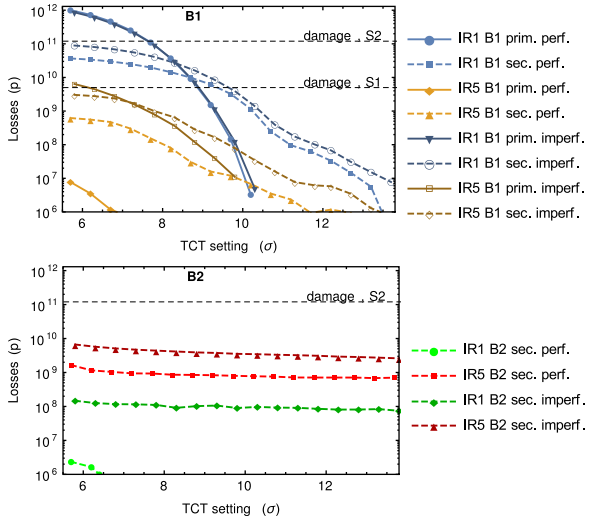

Fig. 4. Losses during an SMPF event, simulated with SixTrack, on the horizontal TCTs in IR1 and IR5 for B1 (top) and B2 (bottom), as a function of the TCT setting. The results are split into primary ("prim.", solid lines) and secondary ("sec.", dashed lines) impacts, and shown both for a perfect machine ("perf.", lighter colors) and for the average over 20 seeds with collimator imperfections ("imperf.", darker colors) as in Ref. [11]. The 2015 physics configuration at 6.5 TeV and β^*=80 cm was used. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

out relative to the IR7 cleaning hierarchy. The B1 TCT in IR5 intercepts much fewer particles than the one in IR1, although $\Delta\mu_{TCT}$ is not so different, since it is partly shadowed by the IR1 TCT. With imperfections, the shadowing can be partly canceled, leading to a large increase in losses.

Table 2
Fractional phase advance between the first extraction kicker and the TCTs in the high-luminosity insertions, calculated from the 2015 LHC optics with β^*=80 cm.

IR	B1	B2
IR1	61°	188°
IR5	52°	193°

The B2 TCT in IR5 has $\Delta\mu_{TCT}$ close to 180° and intercepts only secondary protons (scattered out of the TSCP and TCDQ in IR6), which are approximately independent of the setting. We note that if $\Delta\mu_{TCT}$=180° from the MKDs, $\Delta\mu$ from the TCDQ to the TCT must be 90°. Therefore, any angular kick given in the scattering is transformed into a spatial amplitude at the TCT, explaining the losses. Furthermore, the curve is flatter than for the secondary losses on the IR1 TCT in B1, since no other collimators are in between IR6 and IR5 in B2 (see Fig. 1), so the IR5 TCT is expected to intercept secondary particles up to very large amplitudes. Secondary losses from IR7 are not seen due to the ring geometry, since they are extracted on the second turn in IR6. The B2 TCT in IR1, with a similar $\Delta\mu_{TCT}$, intercepts only negligible losses, since most secondary protons are absorbed upstream in IR5.

In Fig. 4, we show also the estimated damage limit, expressed as the number of protons needed to cause an onset of permanent plastic deformation, for two different impact scenarios from thermo-mechanical simulations of LHC TCTs [34,35]. This damage limit depends strongly on the distribution of the impacts. In scenario 1 [34] (S1), the TCT is subject to primary impact by a full bunch hitting it directly. The standard deviation of the transverse impact distribution is 0.5 mm. In scenario 2 [35] (S2), the TCT is hit by secondary protons, which are much more diluted, with a standard deviation of 6 mm. Therefore, the stress on the material is much lower and many more protons are needed to achieve equivalent damage. The damage limit for S2, expressed as the smallest number of impacting protons causing damage, is about a factor 20 higher than for S1.

In our studies, the primary impacts have an impact distribution very close to S1, so we assume that this damage limit is approximately applicable. The secondary impacts in B2 are much more spread out and resemble the distribution in S2. Therefore, we show only the damage limit from S2 in Fig. 4 for B2. The secondary impacts in B1 are less spread out than in S2, since the TCLAs in IR7 block particles at large amplitudes, but still more spread out than in S1. We assume therefore that the damage limit for secondary impacts in B1 should be somewhere in between S1 and S2.

We point out that these damage limits have a bit of uncertainty, since the scenarios are not exactly identical. A new thermo-mechanical study would be needed in each case for a detailed estimate. Nevertheless, they are good enough to give an approximate estimate. It should be noted also that onset of plastic deformation of the material is a rather pessimistic damage limit. If the deformation occurs only in a small volume, the collimator can still be used in operation. The onset of ejection of tungsten fragments from the jaw, representing a more severe damage which could cause significant LHC downtime and should be avoided, requires about a factor 4 more protons. This pessimistic use of the damage criterion, together with the normalization to a rather high bunch population of 1.5×10^{11} protons, compensates for the uncertainties coming from the differences in proton impact distributions.

It can be seen in Fig. 4 that the B1 TCT in IR1 risks to be damaged by primary losses from an SMPF around 9 σ and the losses increase rapidly below. Nevertheless, the margin from the nominal 13.7 σ setting is largely sufficient to cover optics variations and orbit shifts evaluated with the methods in Ref. [22]. At small TCT openings, the secondary losses from IR7 on the IR1 B1 TCT are also potentially harmful. The simulated B2 IR5 losses are instead, with imperfections, still about a factor 30 below the S2 damage limit applied in this case, as seen in Fig. 4, and the worst seed is still more than a factor 10 below.

In view of the clear correlation between TCT loss profile and $\Delta\mu_{TCT}$, the operational TCT settings might be significantly tightened without risk if all TCTs had $\Delta\mu_{TCT}$ such that they do not intercept primary losses, as the IR5 TCT in B2. Depending on the implementation, this could also decrease the secondary losses in IR1 from IR7, since the phase advance between the most impacted TCSG in IR7 could be moved, as well as the phase between the TCSG and the IR1 TCT. With a $\Delta\mu_{TCT}$ close to 0° or 180°, the TCT setting and allowed aperture could therefore be possibly decoupled from constraints given by the TCDQ setting and beam dump failures. This could allow a significant gain in the collimation hierarchy margins, but the magnitude of the gain depends on other constraints on the collimator settings.

3. Other limitations on the collimation hierarchy

The closest TCT setting and the allowed aperture are not only determined by the risk of losses during beam dump failures, but also by beam cleaning during regular operation. Even with a perfect $\Delta\mu_{TCT}$, the TCT setting cannot be tightened further than these other constraints allow. Therefore, in this section, we quantify the inner limit on the TCT in the absence of dump failures.

Since the collimators are aligned once in the beginning of the year and are then deterministically driven back to be centered around the same central orbit in every fill, margins are needed between the different stages of the cleaning hierarchy in order to compensate for any drifts in orbit and optics over time. In order to assess these margins in an empirical way, experimental studies have been performed in the LHC where the integrity of the hierarchy was checked over the year 2015 and with different collimator settings [42]. This was done using the so-called loss maps, where a low-intensity beam was excited using a white-noise excitation of the transverse damper to produce controlled losses [43]. The resulting loss distribution is observed using the beam loss monitors (BLMs) around the ring [44,45].

The BLM signals within IR7 during loss maps at two different settings of the TCSGs, several months after the collimator alignment, are shown in Fig. 5. In the lower image the retraction TCP-TCSG is 2 σ, i.e. 0.5 σ less than in regular 2015 operation. The losses are highest close to the TCPs, at the beginning of the cleaning insertion, and then decrease steadily downstream, indicating a correct cleaning hierarchy. The upper image shows losses with a 1 σ TCP-TCSG retraction. In this case, the highest loss occurs in the middle of the cleaning insertion, on a secondary collimator, indicating an incorrect cleaning hierarchy that would not be acceptable to use in operation.

Loss maps have been repeated on several occasions in 2015, and the settings with 2 σ retraction have always shown a correct hierarchy. If the orbit fluctuations are not worse than in 2015, it is thus compatible with the long-term stability of the cleaning hierarchy, without the need of intermediate alignments, to have as 2016 baseline a 2 σ TCP-TCSG retraction. Similarly, a small 0.3 σ reduction of the 2015 retraction between the dump protection in IR6 and the TCSGs in IR7 is put in place, which is compatible with that the TCDQ and TCSP should not act as secondary collimators.

Another very important limitation on how close to the beam various collimators can operate is their impedance. In fact, the LHC machine impedance is dominated by the contribution from CFC collimators over a wide range of frequencies [46]. Experimental and theoretical studies on this have been carried out in 2015 [47,48]. It is beyond the scope of this article to describe in detail these studies, but it is concluded that the tighter collimator settings in IR7 and IR6, described above, do not increase the impedance beyond acceptable levels [48]. These settings, which are summarized in Table 1, are therefore confirmed as the 2016 baseline.

With the IR7 and IR6 settings fixed, we study the constraints on the TCT setting and the triplet aperture. The TCTs must be well outside the TCSGs in IR7 in σ, as they would otherwise intercept secondary halo, from which outscattered showers risk to increase the power load on the triplet and cause an intolerably high experimental background. With these constraints, we tentatively put an ultimate inner limit on the TCT setting at the same level as we propose for the TCDQ and the TCSP, in order to stay clear of secondary halo, meaning about 8.3 σ if the TCSGs are at 7.5 σ. However, it should be noted that this limit might be lowered further in the future if the TCSG setting is reduced, either by a smaller margin TCP-TCSG or by tighter TCPs.

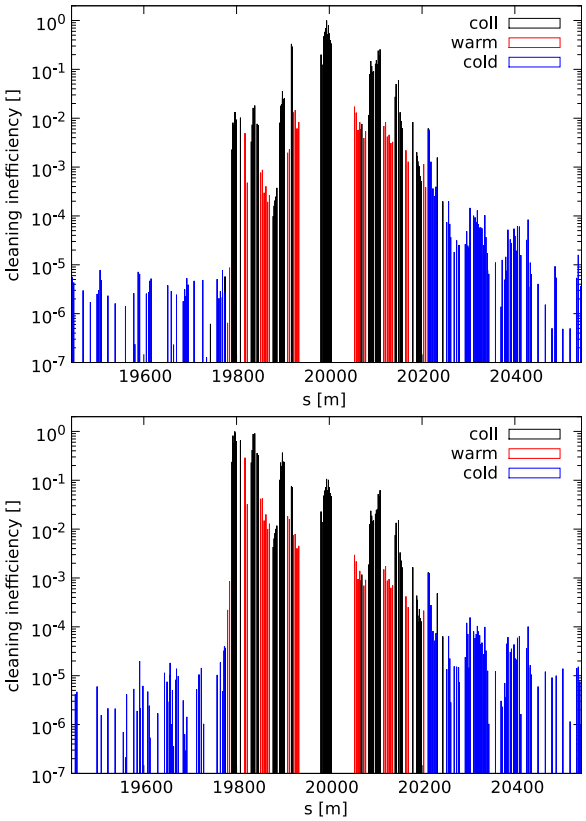

Fig. 5. Measured BLM signals in IR7, normalized to the highest loss, with TCSGs at 6.5 σ (top) and at 7.5 σ (bottom), recorded during special tests in 2015 [42]. In the top plot, the highest loss occurs at a secondary collimator in the middle of the insertion, which indicates a broken collimation hierarchy.

Even if the TCTs are outside the TCSGs, the losses on the TCTs and the induced experimental background from tertiary halo increase steeply when the setting is reduced, as shown in dedicated experimental and theoretical studies [49,50]. The results confirm that a TCT setting at 9 σ, i.e. 1.5 σ outside the TCSGs in IR7, gives sufficient margin, but this is not necessarily the ultimate limit, due to the uncertainties in extrapolating from the special conditions in the beam tests.

Finally, the margin between TCT and triplet has to be assessed also for the case where none of these elements risks catastrophic damage during asynchronous dumps. This margin can then be slightly relaxed and we assume that it should not be violated during 90% of the time in physics, as a trade-off between availability and performance. If too high losses would occur during the other 10% of the time, the beam would be safely dumped by the BLMs. Using the methods in Ref. [22], this gives a retraction of 0.9 σ between TCTs and triplets if we assume the same statistical distribution of measured orbit shifts as in 2015. With a 9 σ TCT setting, a triplet aperture of 9.9 σ can thus be accepted. Dedicated aperture measurements with beam have shown that the triplet aperture at β^*=40 cm is acceptable [51].

4. Reach in β^* with improved phase advance

Assuming that the inner limits on the TCT setting and triplet

aperture in the absence of constraints from asynchronous dumps are 9 σ and 9.9 σ respectively, we study the influence of the phase advance between MKDs and TCTs on the allowed settings and hence on β^*. For this purpose, we study losses on TCTs during a type 2 SMPF as a function of the phase advance using the phase-space integration (PSI) method in the horizontal plane. This simple and fast method, introduced in Ref. [52] and refined in Ref. [22], consists of a numeric integration of the beam distribution over the part of the normalized phase space which is not intercepted by the TCDQ but later on the same turn hits a TCT. If R_i is the phase space region where particles are intercepted by a certain element i, and the phase space region not intercepted is the complement R_i^c, we can write the fraction of a bunch intercepted by the TCT as

$$f_{\mathrm{TCT}} = \iiint_{R_{\mathrm{TCT}} \cap R^c_{\mathrm{TCDQ}}} \rho(X_1, P_1, \delta)\mathrm{d}X_1\mathrm{d}P_1\mathrm{d}\delta, \tag{2}$$

where ρ is an assumed Gaussian beam distribution, (X_1, P_1) are the initial normalized phase space coordinates at the first MKD, and δ is the fractional momentum offset.

To determine the integration regions R, we propagate (X,P) with linear transfer matrices. Assuming that MKD j gives a kick θ_j to the passing bunch, the region R_i outside aperture i is given in linear approximation by the inequality

$$\left| C_{1i}X_1 + S_{1i}P_1 + D_i\delta + \sum_{j=1}^{15} (S_{ji}\theta_j) \right| \geq A_i, \tag{3}$$

which simply states that the sum of the initial amplitude and the total kick given by the MKDs, propagated to i, is larger than A_i (the normalized aperture of element i). Here (C_{ji}, S_{ji}) are the matrix elements from j to i, which in normalized phase space are simply $(\cos\Delta\mu_{ji}, \sin\Delta\mu_{ji})$. Furthermore, D is the linear dispersion normalized by σ, which we assume to have its nominal value, since the contribution from δ is small compared to the considered oscillation amplitudes.

Compared to SixTrack, the PSI method is much faster and we do not need to match a full LHC optics for each studied scenario, since (C_{ji}, S_{ji}) depend only on the phase advance $\Delta\mu_{ji}$. However, all collimators and aperture cuts are by definition considered as perfect absorbers, so the secondary impacts cannot be studied. The method has been shown to agree well with SixTrack for the primary impacts [22].

As for SixTrack, the kicks θ_j are based on measured data from the type 2 SMPF, and f_{TCT} in Eq. (2) is calculated for several consecutive bunches receiving different MKD kicks. To estimate the total impacts, we sum f_{TCT} over all contributing bunches and normalize again the result to 1.5×10^{11} protons per bunch. For simplicity, we include only the cuts of the TCDQ and one TCT.

Fig. 6 shows losses on the TCT, calculated using Eq. (2), where the TCT setting and $\Delta\mu_{\mathrm{TCT}}$ are varied. The phases between the MKDs, as well as the 2016 tighter TCDQ setting of 8.3 σ, were kept constant. There is, as expected, a dramatic improvement in losses when $\Delta\mu_{\mathrm{TCT}}$

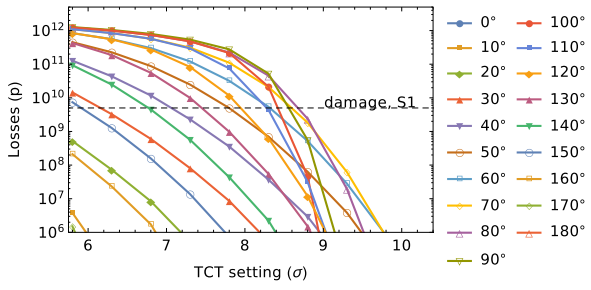

Fig. 6. Losses, calculated using the PSI method, on a TCT during an SMPF event as a function of the TCT setting, for various values of the fractional phase advance between the first MKD and the TCT, and assuming a TCDQ setting of 8.3 σ, corresponding to the 2016 baseline without imperfections. The simulations are normalized to 1.5×10^{11} protons per bunch.

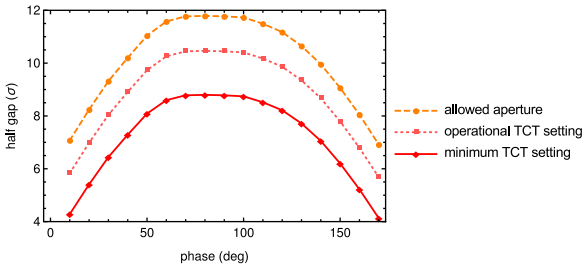

Fig. 7. The minimum allowed effective TCT opening for beam dump failures, the minimum allowed operational TCT setting, and the minimum allowed aperture, as a function of the phase advance $\Delta\mu_{\mathrm{TCT}}$ from the dump kickers. The aperture is calculated from the operational TCT setting, which in turn is derived from minimum TCT setting, adding the operational drifts of 2015 as in Ref. [22]. The minimum TCT setting is the smallest opening where the losses in Fig. 6 are still a factor 2 below the plastic damage limit from S1 [34].

approaches 0° or 180°. At the intersection between the S1 damage limit[6] and the TCT losses for different $\Delta\mu_{\mathrm{TCT}}$, it can be seen that every 10° that can be gained in $\Delta\mu_{\mathrm{TCT}}$ contributes by about 1 σ, if $\Delta\mu_{\mathrm{TCT}}$ is far from 90°, but the gain becomes smaller when $\Delta\mu_{\mathrm{TCT}}$ approaches 90°. It should be noted that for any a there is a small difference in losses between 90°+a and 90°−a, which is introduced by the fact that the $\Delta\mu$ between the first MKD and the TCDQ is not exactly 90°, but 96°, as in most LHC optics.

For each $\Delta\mu_{\mathrm{TCT}}$ we calculate now the TCT setting, where the interpolated losses are half of the estimated S1 damage level, as shown in Fig. 7. We assume this as the minimum allowed setting for dump failures, after any drifts in optics and orbit have been subtracted. This accounts for the uncertainties on the S1 limit which assures that we stay on the pessimistic side.

The minimum operational TCT setting can be obtained if we add to the minimum setting the operational margins needed to account for orbit and optics drifts. We do this according to the method in Ref. [22], but starting from the values in Fig. 7 instead of the TCDQ setting. Our calculation can be considered a generalization of this method, accounting for arbitrary $\Delta\mu_{\mathrm{TCT}}$. Using the same method, we calculate also the margin from TCT to the aperture, and show the results in Fig. 7.

To derive the reach in β^*, we first calculate analytically, as in Ref. [22], the needed aperture for each new β^* by scaling the β-function and orbit at the triplet bottleneck from the most pessimistic aperture measurement with beam (in this case 10 σ aperture at 6.5 TeV with β^*=40 cm and a crossing angle of 185 μrad [51]). The new crossing angle is calculated by using a fixed normalized beam–beam separation d_{BB}. The resulting aperture, shown for a few different values of d_{BB} and emittances, is shown in Fig. 8.

At any given $\Delta\mu_{\mathrm{TCT}}$ and d_{BB}, we can now determine the minimum acceptable β^* by finding the β^* where the required aperture equals the protected aperture at that $\Delta\mu_{\mathrm{TCT}}$. For the start of operation in 2016, we use as input d=10 σ for 3.75 μm emittance, which is assumed as limit to preserve an acceptable dynamic aperture [53]. This is also an improvement compared to 2015, where 11 σ was used [29], and the limit is based on experimental studies from 2015 [54,55]. This smaller d_{BB} contributes also to an improved β^*-reach.

The resulting minimum β^* as a function of $\Delta\mu_{\mathrm{TCT}}$ is shown in Fig. 9. For each $\Delta\mu_{\mathrm{TCT}}$, the maximum of the allowed aperture from dump failures constraints and cleaning has been used. For the 2016 baseline of a 9.9 σ aperture, the curve is flat for $\Delta\mu_{\mathrm{TCT}}$ below 30° and above 150°, where cleaning is limiting, while the dump failure constraints are limiting over the rest of the range. At the worst possible phase, a β^* just below the nominal 55 cm is possible. This includes already the gain from the tighter settings in IR7 and IR6, and d=10 σ. If, in addition, a

[6] Since we only consider focused primary losses in the PSI method, we compare to the more pessimistic S1 damage limit.

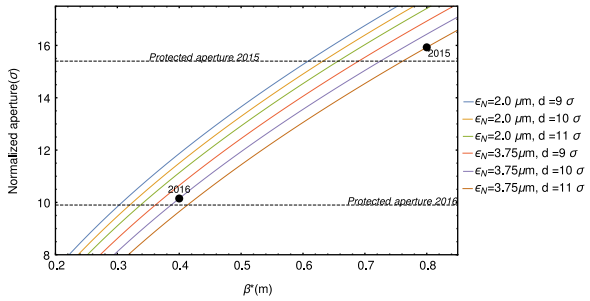

Fig. 8. The estimated aperture as a function of β^*, for given normalized beam–beam separations d_{BB} and normalized emittances ϵ_n. The protected apertures in the 2015 and the 2016 configurations are shown as horizontal lines, and the dots indicate the operational configurations in the two years.

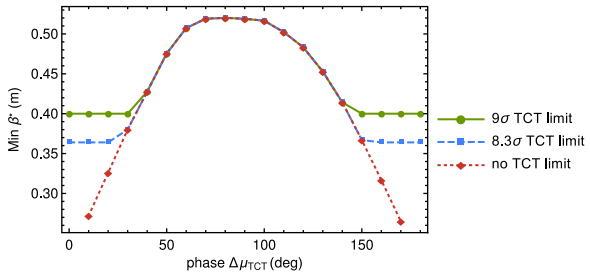

Fig. 9. The minimum allowed β^* as a function of the phase advance $\Delta\mu_{\mathrm{TCT}}$ between the dump kickers and TCTs, calculated as the maximum over the different limits from cleaning (flat part of curve) and beam dump failures. The calculations assume an 8.3 σ TCDQ setting, a 10 σ beam–beam separation [53], and that secondary losses from dump failures are not limiting, and that the measured aperture in 2015 does not degrade.

$\Delta\mu_{\mathrm{TCT}}\leq 30°$ could be guaranteed, β^*=40 cm is within reach.

The values of $\Delta\mu_{\mathrm{TCT}}$ in Fig. 9 are the limit that should be respected also in an imperfect machine. Studies of phase beating resulting from random optics imperfections, compatible with the measured β-beat [12–14], show that $\Delta\mu_{\mathrm{TCT}}$ could vary by a few degrees. For most random seeds with larger variations in $\Delta\mu_{\mathrm{TCT}}$, the β-beat is also so large that the margins using the previous model, relying on a β-beat below 10% [22], would be violated.

The phase could also drift due to off-momentum beating. We consider the worst case of the largest possible fractional momentum error of 2×10^{-4} that can still be tolerated in the LHC without triggering a beam dump. At larger energy offsets, the orbit excursions would trigger interlocks on the beam position. Optics calculations using MAD-X [56] shows that this could give rise to a phase beating of 7° between MKDs and TCTs.

In total, this gives a possible phase beating of about 10°. Therefore, for the 2016 baseline with TCDQ at 8.3 σ and a TCT setting not smaller than 9 σ, we demand that $\Delta\mu_{\mathrm{TCT}}$ should be within a 20° range from 0° or 180° for all TCTs. Provided that this constraint if fulfilled, β^*=40 cm is possible.

5. LHC optics with improved phase advance

Nominal LHC optics [1,57], used in Run 1 and Run 2 (2015) for physics operation, undergoes a transition once top energy is reached, where β^* is decreased from its relaxed injection value to the smaller value used for physics, at three of the four LHC experiments. This transition, called squeeze, is implemented by varying the current of the quadrupoles nearby the IP that have individual power converters (Q1–Q13 left and right). Only the Q1–Q3 quadrupoles, closest to the IP, affect both beams at the same time. The rest of the quadrupoles have separate vacuum chambers and coils for the two beams. They can be powered independently, provided the ratio of the two currents does not exceed 50%. During the squeeze, the Twiss parameters at the boundaries of the experimental IRs (β-functions, horizontal dispersion and phase advance) are kept constant to match the arcs, while $\Delta\mu_{\mathrm{TCT}}$ as a function of β^* changes since it is not explicitly constrained.

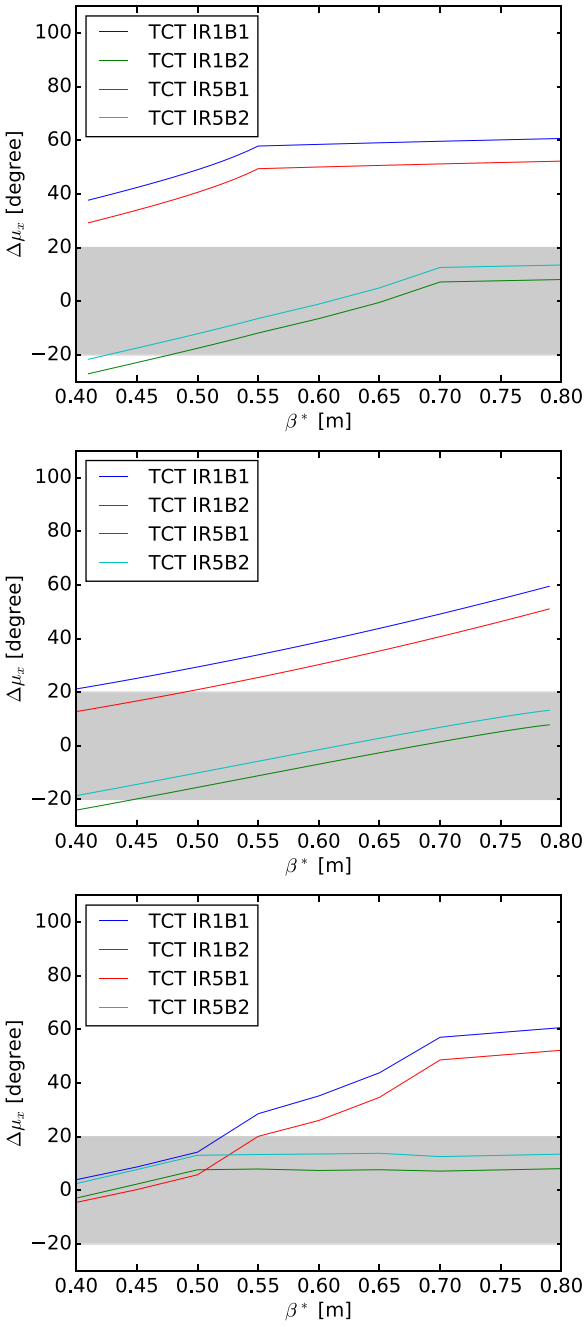

Fig. 10. Phase advance $\Delta\mu_{\mathrm{TCT}}$, modulo 180°, as a function of β^* for the nominal 2015 optics (top), the IR Option (middle), and the Arc Option (bottom) for B1 and B2 in IR1 and IR5. The Arc Option is well within the target from β^*=55 to 40 cm and represents the safest option.

Fig. 10 shows on the top $\Delta\mu_{\mathrm{TCT}}$ as a function of β^* during the squeeze for the standard LHC optics. The allowed range of phases, according to the 2016 collimation constraints discussed above, is shown as a gray band. Even though $\Delta\mu_{\mathrm{TCT}}$ improves at smaller β^*, it is not within the permitted interval at β^*=40 cm. Therefore, the standard optics has to be modified to meet these requirements. It is also important to note that, since $\Delta\mu$ between IP1 and IP5 is sufficiently close to a multiple of 180° (about 5–8°), it is possible to optimize the phase for both IRs by adjusting $\Delta\mu_{\mathrm{TCT}}$ to the IR5 TCTs only. Apart from the change of $\Delta\mu_{\mathrm{TCT}}$, it is desirable to minimize changes with respect to

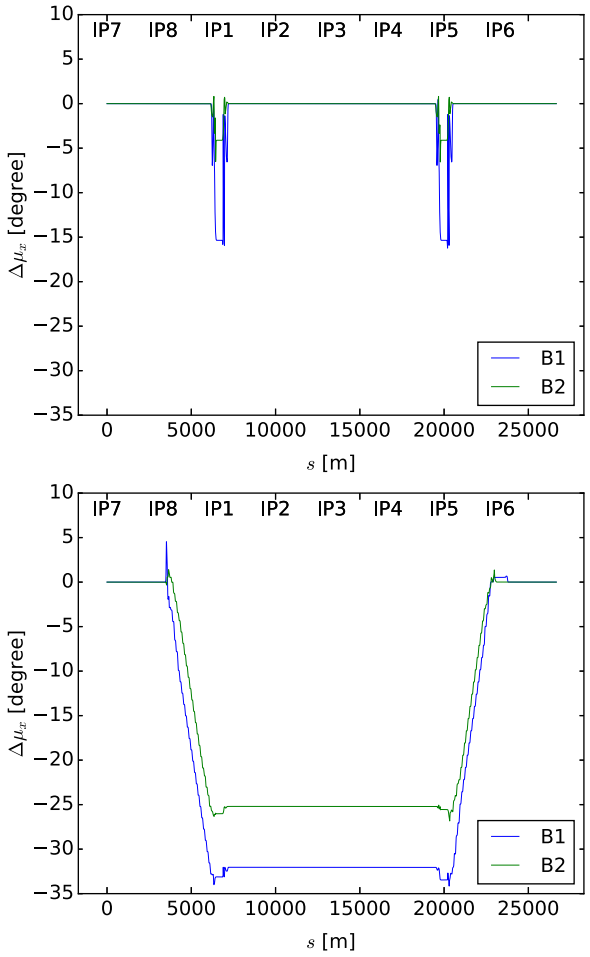

Fig. 11. Comparison with the 2015 optics in terms of phase change $\Delta\mu_x$ for the IR Option (top) and the Arc Option (bottom) at β^*=40 cm. For better readability, the coordinate system has been shifted to have s=0 in IR7.

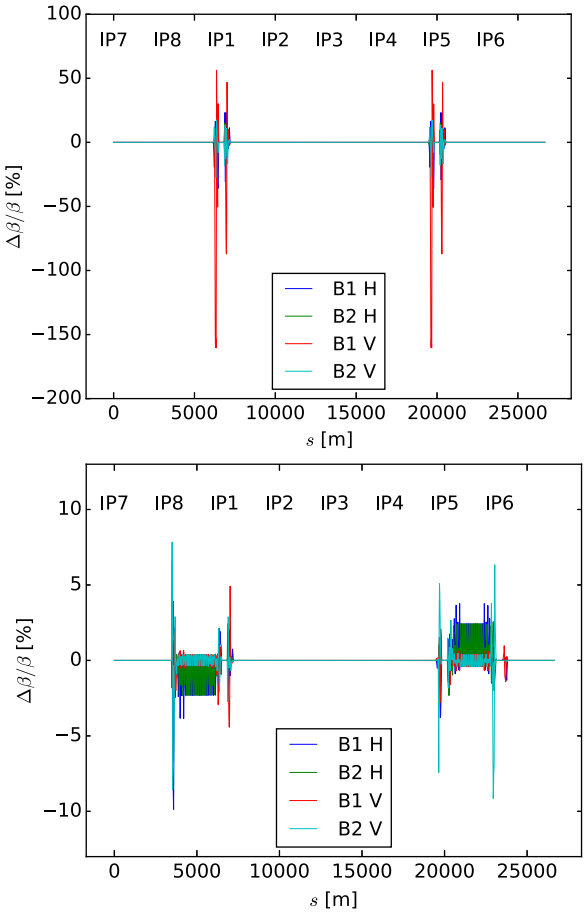

Fig. 12. Comparison with the 2015 optics in terms of $\Delta\beta/\beta$ for the IR Option (top) and the Arc Option (bottom) at β^*=40 cm. For better readability, the coordinate system has been shifted to have s=0 in IR7.

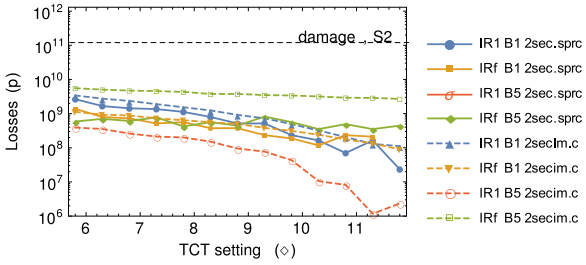

Fig. 13. Losses during an SMPF event, simulated with SixTrack, on the horizontal TCTs in IR1 and IR5 in both beams, as a function of the TCT setting, normalized to 1.5×10^{11} protons per bunch. Results are shown for a perfect machine ("perf.", solid lines) and for the average over 20 seeds with collimator imperfections ("imp.", dashed lines) as in Ref. [11]. Only secondary impacts are obtained. The new β^*=40 cm optics for 2016 was used, as well as the 2016 collimator settings.

the 2015 optics, in order to re-use the optics corrections and experience gathered so far.

Two options have been developed to meet the desired constraints. The first one acts on the internal phase advance within IR1 and IR5 (called "IR Option"). The differences are local to IR1 and IR5, however the relative local optics changes are rather large. In addition the duration of the squeeze would have been longer with respect to 2015. The second option (called "Arc Option") acts on the phase advance of the arcs IR5–IR6 and IR8–IR1, while absorbing the residual β-beating with the corresponding insertions. This option integrates large phase advance changes by small local modifications of the β functions and the variations of quadrupole strengths are smaller than for the IR Option. $\Delta\mu_{\mathrm{TCT}}$ for both options is shown in Fig. 10.

The two options can be directly compared in terms of how they reach the phase target, and the extent of the optics change, in particular the differences in phase and β-functions (see Figs. 11 and 12). Clearly the Arc Option better fulfills the requirements on $\Delta\mu_{\mathrm{TCT}}$ at the cost of a large fraction of the machine being involved in the modification. For this reason the Arc Option has been retained.

To verify that this solution is indeed a usable optics for LHC operation, several checks must be performed. First, the gain in losses on TCTs during beam dump failures was assessed in simulations. The resulting losses on all TCTs in IR1 and IR5 at β^*=40 cm, as simulated with SixTrack using the 2016 collimator settings, are shown in Fig. 13 as a function of the TCT setting. As in Fig. 4, an SMPF was assumed with the first MKD firing, and results are shown both for the perfect machine and with imperfections. Comparing the results with Fig. 4, a very significant gain in losses is seen on the critical TCTs. No primary impacts are observed on any of the TCTs, and the secondary impacts on the previously most critical TCT in IR1, B1, are also much lower, since the phase advance between the IR7 TCSG emitting secondary particles and the TCT is also changed. Considering the damage S2 limit, since these secondary losses are spread out over several millimeters, the simulated number of impacts on the most critical TCT is still more than an order of magnitude below damage. Based on these results, we conclude that the possible impacts during beam dump failures are, as expected, not limiting the TCT setting with the new optics.

Further SixTrack simulations have been performed, showing that the regular beam cleaning performance is equivalent in the new and old optics versions and that the dynamic aperture of the machine does not

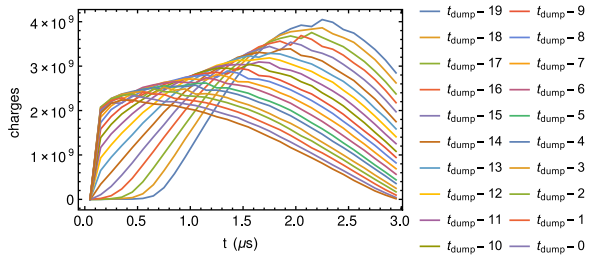

Fig. 14. Example of longitudinal beam profiles along the abort gap, as measured by the synchrotron light monitor, during the 20 s preceding the asynchronous dump test.

degrade with the new optics. The new optics can thus be used operationally in the LHC.

6. Experimental validation

The β^*=40 cm LHC configuration with tighter collimator settings and the new optics should be validated also experimentally, using a safe low-intensity beam, before high-intensity beams are allowed for physics operation. In particular, the losses on TCTs during beam dump failures must be investigated to show that the improved $\Delta\mu_{TCT}$ works and that the machine is safe also with the tighter TCT settings. Therefore, special asynchronous dump tests have been carried out, which are anyway a part of the standard commissioning to verify the safety of the LHC.

In these tests, the RF system is switched off with a single low-intensity bunch in a bucket next to the abort gap. As the beam debunches, the abort gap is populated, and a standard beam dump is triggered. This means that the beam in the abort gap is kicked at intermediate amplitudes by the MKDs, as would be the case during a real asynchronous beam dump. It should be noted that these tests resemble an AD, while the more critical SMPF event cannot easily be tested experimentally with beam.

The abort population is measured online with a synchrotron light monitor [58,59] as shown in Fig. 14. Here the longitudinal distribution of beam in the abort gap is evolving over 20 s. It can be seen how the bunch enters the abort gap from the right and moves to the left while it spreads out.

The losses around the LHC ring are measured with the BLMs during the asynchronous dump test. Examples of observed loss distribution are shown in Fig. 15. Losses on both beams are superimposed, since they were dumped simultaneously. As expected, the main loss location is at the dump protection devices in IR6, but losses above the noise level are observed at several other elements, and in particular at the TCTs.

With the new β^*=40 cm optics, asynchronous beam dump tests were carried out with special settings, where all collimators except the TCTs and the TCDQ and TCSP in IR6 were kept open at parking position. This situation is significantly more pessimistic than what is deployed during standard operation, where additional protection is provided by IR7. Furthermore, this simplified setup makes it easier to compare the data with simulations. Tests were performed at three different TCT settings, and always with a 1.2 mm orbit bump away from the TCDQ, in order to reflect the most pessimistic situation. These measurements can be compared with the 2015 results at β^*=80 cm, where the full collimation system was kept at standard settings. These tests were done with two different TCT settings. Fig. 15 shows examples of loss distributions in both configurations.

To get an empiric estimate on the risk for TCT damage during an asynchronous dump at high intensity and to compare with simulations, we use the measured TCT losses $N_{loss,test}$ during the low-intensity test to calculate the potential TCT losses $N_{loss,full}$ with a full physics beam as

$$N_{loss,full} = N_{loss,test} \times \frac{N_{AG,full}}{N_{AG,test}} = \frac{S_{BLM}}{F_{Gy/p}} \times \frac{N_{AG,full}}{N_{AG,test}}. \quad (4)$$

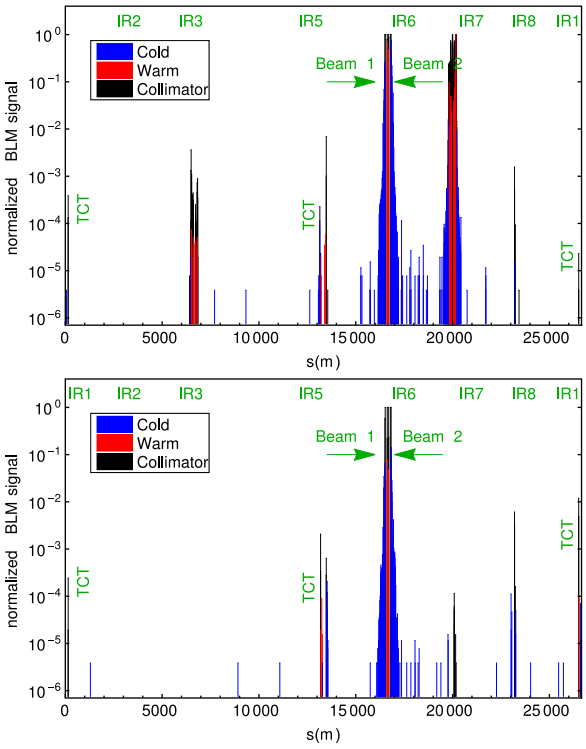

Fig. 15. The measured loss distribution around the ring at the time of the asynchronous dump test, using a 40 μs integration time on the BLMs, for the 2015 configuration at β^*=80 cm and TCTs at 13.7 σ (top) and for β^*=40 cm with TCTs at 9 σ (bottom), where in the latter case the IR3 and IR7 collimators have been moved out to parking position. All signals have been normalized to the highest one. The loss profile at the peaks appears flat since some BLMs reached saturation.

Here $N_{AG,test}$ and $N_{AG,full}$ are the number of miskicked particles in the abort gap, going to the TCT, with a full physics beam and during the test respectively. We calculate $N_{loss,test}$ from the measured TCT BLM signal S_{BLM} in Gy, integrated over 40 μs[7] around the dump, using a conversion factor $F_{Gy/p}$ between impacting protons on the TCT and BLM signal (in Gy). The BLMs are placed downstream of the TCTs and intercept secondary shower particles, so $F_{Gy/p}$ depends on the local geometry and the materials that the induced showers have to pass to reach the BLM, as well as on the distribution of the impacts on the TCT. We derive $F_{Gy/p}$ experimentally from a dedicated test, where a horizontal TCT was moved in so that it was the tightest aperture in σ of the ring. In this configuration, a low-intensity beam was excited to create losses, which were all intercepted by the TCT. By relating the BLM signal to the loss in intensity, $F_{Gy/p} = 2.1 \times 10^{-11}$ Gy/p was obtained.

In Eq. (4), we must account for that with a full physics beam, we expect an equal bunch passing every 25 ns, while in the test, the abort gap is not homogeneously populated, as can be seen in Figs. 14 and 16. In Fig. 16, the horizontal axis has been scaled by the expected normalized MKD kick in σ as a function of time in the abort gap, assuming the ideal MKD kicker waveform and that all kickers fire simultaneously. The gray band indicates the range of the 8 bunches that in simulations cause significant TCT impacts, with amplitudes in the range 6–12 σ. As can be seen, this is only a small fraction of the abort gap, at the beginning of the kicker rise.[8] We therefore use for $N_{AG,test}$ the integrated population only in this gray window of Fig. 16. For $N_{AG,full}$ we simply assume 8 bunches of 1.5×10^{11} protons each, to stay on the pessimistic side for LHC Run II. Since the density is rather

[7] This is the shortest possible integration time for the BLM electronics.

[8] Beam kicked at higher amplitudes is intercepted by the TCDQ or extracted, and beam kicked at lower amplitudes does not reach the TCTs.

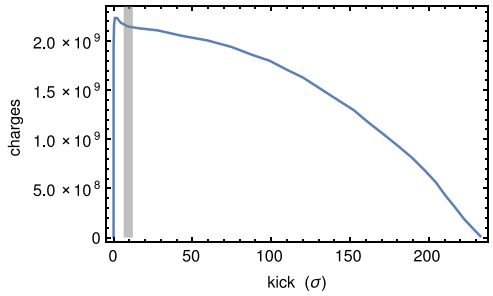

Fig. 16. Measured longitudinal beam profile from Fig. 14 along the abort gap, at the time when the dump was triggered, given as a function of total kick in σ. The kick has been summed over all MKDs, and assuming that all kickers fire at the same time. The gray band indicates the region in σ where particles could end up at the TCTs.

flat over a few σ around the chosen window in Fig. 16, the results do not depend much on its exact limits, as long as we always use the same time interval for both $N_{AG,test}$ and $N_{AG,full}$.

Fig. 17 shows the comparison between the simulated and the measured TCT losses during an AD as a function of TCT setting, for both the new β^*=40 cm optics with only the TCTs and the IR6 collimators at the beam, and the 2015 β^*=80 cm configuration with the old $\Delta\mu_{TCT}$ in Table 2 and the full collimation system active. We show the results for the most critical TCT (IR1 B1). As in previous simulations, we normalize the results to a full 25 ns physics beam of 1.5×10^{11} protons/bunch and include also collimator imperfections.

Figure 17 shows an excellent qualitative agreement between simulations and measurements. At β^*=80 cm with the old $\Delta\mu_{TCT}$, the measured losses rise sharply with decreasing TCT settings, as expected from the simulations. On the other hand, with the new β^*=40 cm optics with matched $\Delta\mu_{TCT}$, the three measurement points with different TCT settings cause very similar TCT losses, consistent with that these are indeed secondary impacts as foreseen by the simulations.

Quantitatively, the agreement between the average over the imperfection seeds and the measurements is generally within a factor 3, except the point at 10.7 σ at β^*=80 cm, where the discrepancy is about a factor 5. Considering that the losses span many orders of magnitude, the many unknowns, and the complexity of the simulation chain, we consider this a very good agreement. These discrepancies are of similar order of magnitude as reported in Ref. [11] for simulations and measurements of standard collimation losses.

Several uncertainties could by far make up for the discrepancies. The imperfections are not known in detail, and the typical spread between different random seeds gives variations of about 50%. Some minor variations can be expected also between measurements due to small deviations of e.g. orbit, optics, and beam distribution, which was not known in detail (a Gaussian distribution was assumed in the simulations). Furthermore, the BLM response factor could be different in this particular loss scenario compared to when it was measured, due to variations in the TCT impact distribution. Finally, timing imprecision and noise in the measurements of the abort gap population add further uncertainties.

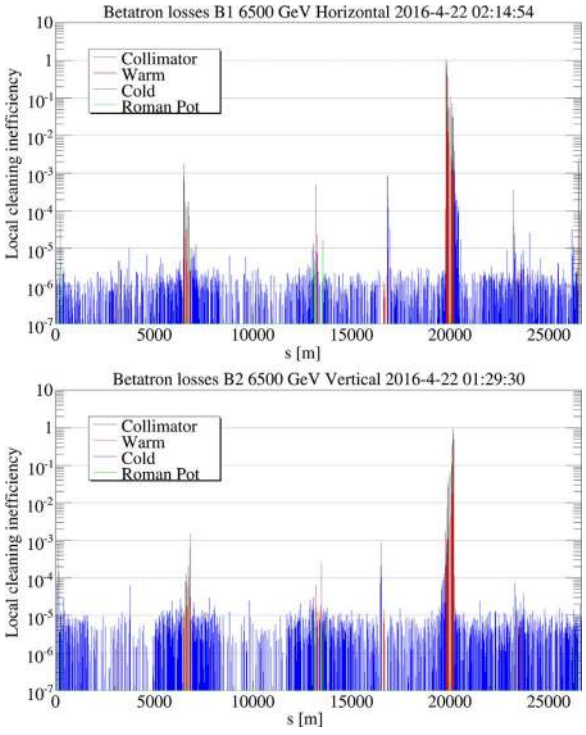

Fig. 18. Measured beam loss distributions around the ring, for losses in the horizontal plane of B1 (top) and the vertical plane of B2 (bottom), using the 2016 collimator settings in Table 1 and the newly developed β^*=40 cm optics with improved phase advance.

Based on these studies, we conclude that the simulations work very well within the expected accuracy, and that the danger of beam dump failures does not limit the TCT setting with the new phase advance at β^*=40 cm. The margin to damage for the secondary impacts in S2 is more than an order of magnitude, in spite of the extremely pessimistic configuration studied, with an orbit bump away from the TCDQ and with IR7 is open.

Other necessary commissioning and validation steps before high intensity beams are allowed with the new collimator settings include validations of the cleaning using loss maps. Examples of the measured loss distribution around the ring are shown in Fig. 18, with the main peaks in IR7. Higher losses than in 2015 are seen at the TCTs, as

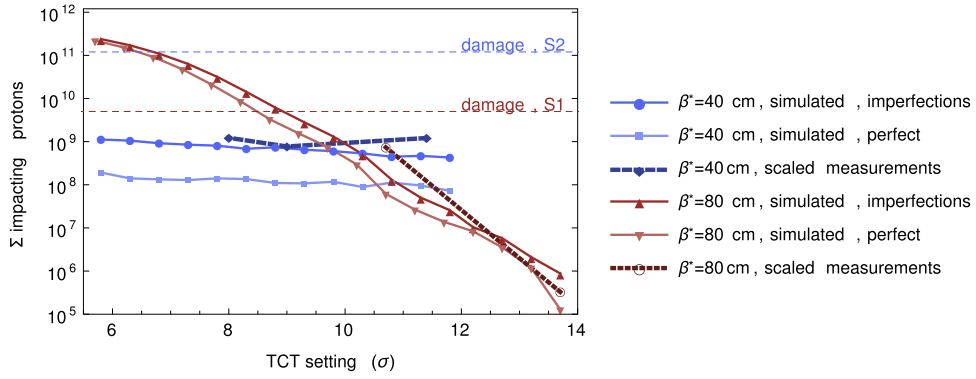

Fig. 17. Losses on the most critical TCT (IR1 B1) during an AD, from SixTrack simulations and from scaled measurements, normalized to 1.5×10^{11} protons per bunch. Results are shown for the new β^*=40 cm 2016 optics with all collimators retracted to parking except TCTs and IR6 dump protection, and for the 2015 β^*=80 cm optics with the full collimation system active. The measured data, taken with a low-intensity beam, were scaled up by the ratio of beam in the abort gap kicked on TCTs. Simulation are presented both for the perfect machine and as the average over 20 seeds of collimator imperfections as in Ref. [11].

expected from the tighter setting. The measured IR7 cleaning hierarchy with the final optics is very similar to the one from the 2015 tests in Fig. 5 and the cleaning efficiency of the collimation system is, with these tighter settings, slightly better than what was achieved in 2015 [60] and in Run 1 [11]. Furthermore, the available aperture was also validated with beam-based measurements as in [18]. The machine configuration relies also on a very good optics commissioning and correction [12–14], which is important for the collimation hierarchy. After the successful beam tests, the β^*=40 cm configuration was validated as baseline for the 2016 physics run in the LHC [61,62].

7. Outlook

While β^*=40 cm has been successfully commissioned and put into operation in 2016, it is likely not the lower limit, and studies for improved reach in β^* have continued during the year. Using beams with smaller-than-nominal emittance from the LHC injector chain, further experiments in late 2016 have shown that the crossing angle at β^*=40 cm could be reduced from 185 μrad to 140 μrad without significant detrimental effects on the beam lifetime [63], which corresponds to d_{BB}=9.3 σ for a 2.5 μm emittance. The smaller crossing angle was successfully introduced in operation in late 2016, but without squeezing β^* further.

Furthermore, work is ongoing to understand if the collimation hierarchy margins can be even further reduced in the future. New tests in 2016 give hope that introducing an angular tilt of the TCSG, which broke the hierarchy in Fig. 5, could compensate for a likely angular misalignment [64], so that the TCSGs could be moved in by at least 0.5 σ while keeping a correct hierarchy. Furthermore, in another study it is investigated whether the TCP cut could be reduced, and the results indicate that a reduction by 0.5 σ might be feasible [65].

The margin between TCDQ and TCT could also be tightened further, since Run I studies have shown that the dominating source of machine-induced experimental background is beam-gas collisions [5], so that the contribution from an increased tertiary halo might not have a significant negative impact. The preliminary outcome of the operational experience with tighter TCT settings and new tests in 2016 [66] support this conclusion and indicate that the TCTs might be closed further by at least 0.5 σ if $\Delta\mu_{\mathrm{TCT}}$ is kept close to zero.

In total, combining a 0.5 σ tighter TCP, a 0.5 σ reduction of the margin between TCP and TCSG, and the TCTs moved 0.5 σ closer to the TCDQ, a further reduction of the allowed protected aperture by 1.5 σ would be possible. The combination of a tighter collimation hierarchy and the smaller beam–beam separation would allow β^* to be squeezed further down to around 30 cm. This could be a possible operational configuration in the future, but it might be that an even smaller β^* could be within reach.

8. Conclusions

We have reviewed the calculations of LHC collimator settings, which should protect the magnets and also the less robust collimators. In the previous LHC run, the potentially damaging losses from miskicked protons during beam dump failures, which could hit the tertiary collimators and the inner triplets in front of the experiments, were driving the limit on the setting on the tertiary collimators and the allowed normalized aperture in the triplets. This in turn constrained the achievable β^* and peak luminosity. We have shown the development of a modified LHC optics, where the fractional betatron phase advances between the dump kickers and the sensitive elements are matched close to 0° or 180°, which alleviates this constraint. Simulations and measurement results both show that, using the new optics, the sensitive elements do not risk to be damaged during dump failures in all realistic running scenarios.

In the absence of potential losses during beam dump failures, other constraints from beam cleaning and experimental backgrounds set limits on the collimator settings. Still, the tertiary collimators can be moved in significantly closer to the beam, and a much smaller normalized aperture and β^* can be allowed. The use of the new optics and collimator settings was a key factor for allowing β^*=40 cm in 2016 LHC operation. Other contributions that helped making this possible came from the use of a smaller retraction between the betatron collimators, based on experimental studies, and a smaller crossing angle. The 2016 LHC baseline of β^*=40 cm is half of what was used in 2015 and about 30% less than the nominal design value. The record-low β^* has been an important factor for reaching and surpassing the LHC design luminosity for the first time in 2016, despite lower-than-design beam energy.

Without the constraint of dump failure losses, there is potential to squeeze β^* to even smaller values below 40 cm in the future. The settings of the main betatron collimators are still more open than foreseen in the design and could potentially be tightened. New tests in 2016 indicate that β^*=30 cm could be within reach, but it is also not sure that this will be the lower limit in the future.

Acknowledgments

We thank the LHC operational crew and the LHC collimation team for assisting during the LHC beam tests. We would like to thank also G. Arduini, N. Biancacci, A. Dabrowski, M. Fraser, B. Goddard, M. Guthoff, M. Huhtinen, M. Lamont, T. Manousos, E. Metral, T. Pieloni, B. Salvant, J. Uythoven, and J. Wenninger for helpful discussions and input.

References

[1] O.S. Brüning, et al. (Eds.), LHC Design Report V.1: The LHC Main Ring, CERN-2004-003-V1.
[2] L. Evans, P. Bryant (editors), LHC machine, J. Instrum. 3 (2008) S08001.
[3] ATLAS Collaboration, The ATLAS experiment at the CERN Large Hadron Collider, J. Instrum. 3 (2008) S08003.
[4] CMS Collaboration, The CMS experiment at the CERN LHC, J. Instrum. 3 (2008) S08004.
[5] R. Bruce, et al., Sources of machine-induced background in the ATLAS and CMS detectors at the CERN large hadron collider, Nucl. Instrum. Methods Phys. Res. Sect. A 729 (2013) 825–840. http://dx.doi.org/10.1016/j.nima.2013.08.058.
[6] B. Auchmann, et al., Testing beam-induced quench levels of LHC superconducting magnets, Phys. Rev. ST Accel. Beams 18 (2015) 061002. http://dx.doi.org/10.1103/PhysRevSTAB.18.061002.
[7] R.W. Assmann, Collimators and beam absorbers for cleaning and machine protection, in: Proceedings of the LHC Project Workshop – Chamonix XIV, Chamonix, France, 2005, p. 261.
[8] G. Robert-Demolaize, Design and performance optimization of the LHC collimation system (Ph.D. thesis), Universite Joseph Fourier, Grenoble, 2006
[9] R.W. Assmann, et al., The final collimation system for the LHC, in: Proceedings of the European Particle Accelerator Conference 2006, Edinburgh, Scotland, 2006, p. 986.
[10] C. Bracco, Commissioning scenarios and tests for the LHC collimation system (Ph. D. thesis), EPFL Lausanne, 2008.
[11] R. Bruce, et al., Simulations and measurements of beam loss patterns at the CERN large hadron collider, Phys. Rev. ST Accel. Beams 17 (2014) 081004. http://dx.doi.org/10.1103/PhysRevSTAB.17.081004 URL ⟨http://link.aps.org/doi/10.1103/PhysRevSTAB.17.081004⟩.
[12] R. Tomás, et al., Record low β beating in the LHC, Phys. Rev. ST Accel. Beams 15 (2012) 091001. http://dx.doi.org/10.1103/PhysRevSTAB.15.091001.
[13] A. Langner, et al., LHC Optics Commissioning at β^*=40 cm and 60 cm, CERN-ACC-NOTE-2015-0035.
[14] A. Langner, et al., Optics model, in: Proceedings of the 6th Evian Workshop, Evian, France.
[15] C.A. Pons, et al., IR1 and IR5 Aperture at 3.5 TeV, CERN-ATS-Note-2011-110 MD.
[16] R. Assmann, et al., Aperture determination in the LHC based on an emittance blowup technique with collimator position scan, in: Proceedings of IPAC'11, San Sebastian, Spain, 2011, p. 1810.
[17] S. Redaelli, et al., Aperture measurements in the LHC interaction regions, in: Proceedings of IPAC12, New Orleans, Louisiana, USA, 2012, p. 508.
[18] P. Hermes, et al., Improved aperture measurements at the LHC and results from their application in 2015, in: Proceedings of the International Particle Accelerator Conference 2016, Busan, Korea, 2016, p. 1446.
[19] M. Lamont, Status of the LHC, J. Phys.: Conf. Ser. 455 (1) (2013) 012001.
[20] R. Alemany-Fernandez, et al., Operation and Configuration of the LHC in Run 1, CERN-ACC-NOTE-2013–0041.
[21] D. Wollmann, et al., Multi-turn losses and cleaning, in: Proceedings of the 2010

LHC Beam Operation Workshop, Evian, France.
[22] R. Bruce, R.W. Assmann, S. Redaelli, Calculations of safe collimator settings and β^* at the CERN large hadron collider, Phys. Rev. ST Accel. Beams 18 (2015) 061001. http://dx.doi.org/10.1103/PhysRevSTAB.18.061001 URL ⟨http://link.aps.org/doi/10.1103/PhysRevSTAB.18.061001⟩.
[23] R. Bruce, R. Assmann, How low can we go? Getting below β^*=3.5 m, in: Proceedings of the 2010 LHC Beam Operation Workshop, Evian, France, 2010, p. 133.
[24] R. Bruce, R. Assmann, LHC β^*-reach in 2012, in: Proceedings of the 2011 LHC Beam Operation Workshop, Evian, France.
[25] R. Bruce, et al., Parameters for HL-LHC Aperture Calculations, CERN Report CERN-ACC-2014–0044.
[26] T. Pieloni, et al., Beam stability with colliding beams at 6.5 TeV: in the betatron squeeze and collisions, in: Proceedings of the 2012 LHC beam operation workshop, Evian, France.
[27] R. Bruce, S. Redaelli, Collimation and β^*-reach, in: Proceedings of the 5th Evian Workshop, Evian, France.
[28] R. Bruce, et al., Baseline LHC machine parameters and configuration of the 2015 proton run, in: Proceedings of the LHC Performance Workshop (Chamonix 2014), Chamonix, France.
[29] T. Pieloni, et al., Two beam effects, in: Proceedings of the 5th Evian Workshop, Evian, France.
[30] G. Iadarola, et al., Electron cloud effects, in: Proceedings of the 6th Evian Workshop, Evian, France.
[31] R. Assmann, B. Goddard, E. Vossenberg, and E. Weisse, The Consequences of Abnormal Beam Dump Actions on the LHC Collimation System, LHC Project Note 293, CERN.
[32] R. Schmidt, et al., Protection of the CERN large hadron collider, New J. Phys. 8 (11) (2006) 290.
[33] T. Kramer, LHC beam dump system: analysis of beam commissioning, performance and the consequences of abnormal operation (Ph.D. thesis), TU Graz, 2008.
[34] A. Bertarelli, et al., Updated robustness limits for collimator materials, in: MPP Workshop, Annecy, France, March 2013.
[35] E. Quaranta, et al., Updated simulation studies of damage limit of LHC tertiary collimators, in: Proceedings of the International Particle Accelerator Conference 2015, Richmond, VA, USA, 2015, p. 2053.
[36] F. Schmidt, SixTrack. User's Reference Manual, CERN/SL/94-56-AP.
[37] G. Robert-Demolaize, R. Assmann, S. Redaelli, F. Schmidt, A new version of sixtrack with collimation and aperture interface, in: Proceedings of the Particle Accelerator Conference 2005, Knoxville, 2005, p. 4084.
[38] C. Tambasco, An improved scattering routine for collimation tracking studies at LHC (Master's thesis), Università di Roma, Italy, 2014.
[39] M. Fraser, et al., A beam-based measurement of the LHC beam dump kicker waveform, in: Proceedings of the International Particle Accelerator Conference 2016, Busan, Korea, 2016, p. 3911.
[40] M. Fraser, Private communication.
[41] R. Steerenberg, et al., Beams from the injector in 2016, in: Proceedings of the LHC Performance Workshop (Chamonix 2016), Chamonix, France.
[42] A. Mereghetti, et al., β^*-reach – IR7 Collimation Hierarchy Limit and Impedance, CERN-ACC-NOTE-2016-0007.
[43] W. Hofle, et al., Controlled transverse blow-up of high-energy proton beams for aperture measurements and loss maps, in: Proceedings of IPAC12, New Orleans, Louisiana, USA, 2012, p. 4059.
[44] E. Holzer, et al., Beam Loss Monitoring System for the LHC, IEEE Nucl. Sci. Symp. Conf. Rec. 2 (2005) 1052. http://dx.doi.org/10.1109/NSSMIC.2005.1596433
[45] E.B. Holzer, et al., Development, production and testing of 4500 beam loss monitors, in: Proceedings of the European Particle Accelerator Conference 2008, Genoa, Italy, 2008, p. 1134.
[46] N. Mounet, The LHC transverse coupled-bunch instability (Ph.D. thesis), EPFL Lausanne, 2012.
[47] L. Carver, et al., MD 755: Instability Threshold and Tune Shift Study with Reduced Retraction Between Primary and Secondary Collimators in IR7, CERN-ACC-NOTE-2016-0005.
[48] L. Carver, et al., Instabilities and beam-induced heating in 2015, in: Proceedings of the 6th Evian Workshop, Evian, France.
[49] R. Bruce, et al., Collimation with Tighter TCTs at β^*=40 cm, CERN-ACC-NOTE-2015-0036.
[50] ATLAS Collaboration, Beam Backgrounds in the ATLAS Detector during LHC Loss Map Tests at $\beta^* = 40$ cm and $\beta^* = 80$ cm at E_{beam}=6.5 TeV, ATLAS PUB note, in press.
[51] R. Bruce, et al., IR Aperture Measurement at β^*=40 cm, CERN-ACC-NOTE-2015–0037.
[52] R. Bruce, D. Bocian, S. Gilardoni, J.M. Jowett, Beam losses from ultraperipheral nuclear collisions between Pb ions in the large hadron collider and their alleviation, Phys. Rev. ST Accel. Beams 12 (7) (2009) 071002. http://dx.doi.org/10.1103/PhysRevSTAB.12.071002.
[53] T. Pieloni, et al., Beam–beam effects: long range and head-on, in: Proceedings of the 6th Evian Workshop, Evian, France.
[54] M. Crouch, et al., Long Range Beam–Beam Interaction and the Effect on the Beam and Luminosity Lifetimes, CERN-ACC-NOTE-2016–0019.
[55] M. Crouch, et al., Impact of long range beam–beam effects on intensity and luminosity lifetimes from the 2015 LHC run, in: Proceedings of the International Particle Accelerator Conference 2016, Busan, Korea, 2016, p. 1422.
[56] MAD-X Program, URL ⟨http://cern.ch/mad/⟩
[57] LHC Optics Web, URL ⟨http://lhc-optics.web.cern.ch/lhc-optics/www/⟩
[58] A. Fisher, Expected Performance of the LHC Synchrotron-Light Telescope (BSRT) and Abort-Gap Monitor (BSRA), CERN LHC Performance Note-014.
[59] T. Lefevre, et al., First operation of the abort gap monitors for LHC, in: Proceedings of the International Particle Accelerator Conference 2010, Kyoto, Japan, 2010, p. 2863.
[60] G. Valentino, et al., Performance of the LHC collimation system during 2015, in: Proceedings of the 6th Evian Workshop, Evian, France.
[61] Minutes of the LHC Machine Committee, April 20, 2016.
[62] R. Bruce, et al., LHC run 2: results and challenges, in: Proceedings of the 57th ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams, Malmo, Sweden, 2016.
[63] T. Pieloni, et al., Crossing angle reduction, Presentation in the LHC Machine Committee (LMC), 2016.08.31.
[64] A. Mereghetti, et al., Collimation Hierarchy Limits, 2016, (CERN note, in press).
[65] D. Mirarchi, et al., Collimation with Tighter TCPs, 2016, (CERN note, in press).
[66] R. Bruce, et al., TCT Closure Test in Physics, CERN-ACC-NOTE-2016–0076.