INTERLOCKING OF LATTICE SKEW QUADRUPOLE CIRCUITS IN THE LHC

C. Lannoy[*], C. Hernalsteens, S. Kostoglou, J. Uythoven, D. Wollmann
European Organization for Nuclear Research, Geneva, Switzerland

Abstract

The Large Hadron Collider operates with 24 power converter circuits supplying the lattice skew quadrupole magnets for coupling correction. These circuits are currently not interlocked via the Powering Interlock Controller. Recent operational experience shows that a trip of one of these circuits can cause beam losses that can lead to beam dumps via beam loss monitors. This study employs Xsuite beam tracking simulations to assess the impact of individual skew quadrupole circuit failures on beam losses under realistic LHC Run 3 operating conditions. The simulations identify the circuits that were subsequently connected to the beam interlocking system during the 2025 Year-End Technical Stop.

INTRODUCTION

The Large Hadron Collider (LHC) uses 24 lattice skew quadrupole circuits (RQS), 12 per beam, to correct the linear coupling between the horizontal and vertical planes [1]. The circuits are split into two families: eight 4-magnet circuits (4 per beam, named “RQS.A[1–8][1–8].{B1/B2}”), with magnets (MQS) distributed across an arc between two insertion regions and equipped with Energy Extraction (EE) systems, and sixteen 2-magnet circuits (8 per beam, named “RQS.{R/L}[1–8].{B1/B2}”), with magnets on one side of an insertion region, without EE. For identical orbit offsets and currents, the 4-magnet circuits produce a larger integrated kick upon a sudden current decay and are therefore expected to have a larger impact on the beam. Figure 1 shows the position of the Beam 1 RQS magnets around the ring together with the measured orbit (horizontal and vertical) from the Beam Position Monitors (BPM) and the theoretical closed orbit defined by the optics. The orbit at the RQS magnet locations is highlighted.

The RQS circuits are currently not connected to the Powering Interlock Controller (PIC). A trip therefore does not trigger a pre-emptive beam dump request via the Beam Interlock System (BIS): the beam is only extracted reactively if the beam losses exceed the Beam Loss Monitor (BLM) dump thresholds. During LHC Run 3, 36 RQS trip events were recorded [2], of which 17 occurred without beam, 10 were simultaneous with other circuit trips (e.g. electrical grid perturbations), 5 occurred at top energy with visible losses but with the beam surviving, and 4 triggered a beam dump via BLMs due to high losses. All four BLM-triggered dumps took place in 2025 at top energy, caused by circuits RQS.A81B1, RQS.A67B1 (twice) and RQS.L7B2. The latter is a 2-magnet circuit, which shows that the magnet count alone is not sufficient to predict criticality. A previous analysis of RQS circuit trips before 2017 is given in Ref. [3].

This paper summarises Xsuite [4] beam tracking simulations used to assess the impact of individual RQS circuit failures on beam losses under realistic Run 3 conditions, and derives a prioritisation for PIC interlocking of a subset of circuits, deployed during the 2025 Year-End Technical Stop (YETS).

FAILURE SIMULATIONS

Single RQS circuit failures are simulated with Xsuite using the Run 3 LHC lattice at 6.8 TeV in the 2025 pp end-of-levelling configuration ($\beta^* = 18$ cm in IR1/IR5, crossing angle 160 µrad) with the nominal collimator settings. The failure is modelled by ramping one RQS circuit at a time from 0 A to 200 A, the approximate maximum operational current reached during Run 3. The MQS magnets are designed for a nominal current of 550 A, corresponding to a skew-quadrupolar gradient of 120 T/m [1]. The operational currents during Run 3 therefore remain well below the design limit.

A single particle initialised on the unperturbed closed orbit is tracked turn by turn. At each turn, the particle coordinates at the primary collimator (TCP) in IR7 are converted to normalised coordinates and actions J_x, J_y using the unperturbed Twiss parameters. The actions quantify the orbit shift induced by the failure at the TCP location. The transverse beam distribution in each plane ($u = x$ or y) is modelled as a static, conservative double Gaussian,

$$\rho_u(u) = 0.65\,\mathcal{G}(0,1) + 0.35\,\mathcal{G}(0,2), \tag{1}$$

with $\mathcal{G}(\mu,\sigma)$ a normalised Gaussian. The TCP is treated as a black absorber, meaning that all particles of the static distribution whose action is large enough for their amplitude to exceed the TCP half-gap are considered lost instantaneously. The loss fractions in x and y are computed independently and added, and the cumulative maximum over turns is taken to enforce monotonically increasing losses. The resulting total is converted to MJ via the nominal stored beam energy. This procedure only concentrates on the fast turn-by-turn orbit excursion driven by the dipolar feed-down, which is the fastest failure mode and relevant for machine protection. Long-term effects such as emittance growth or luminosity performance are outside the scope of this study.

Figure 2 summarises the results for Beam 1 with the accumulated loss, the peak orbit change at the TCP, and the sum of the initial orbit offset $\Sigma|x| + \Sigma|y|$ at the magnets of the failing circuit prior to the failure. The circuits showing

---

[*] christophe.lannoy@cern.ch

Figure 1: Layout of the Beam 1 RQS magnets around the LHC ring (top) and comparison between the theoretical (Xsuite) and measured (BPMs) horizontal (centre) and vertical (bottom) closed orbit at their longitudinal positions.

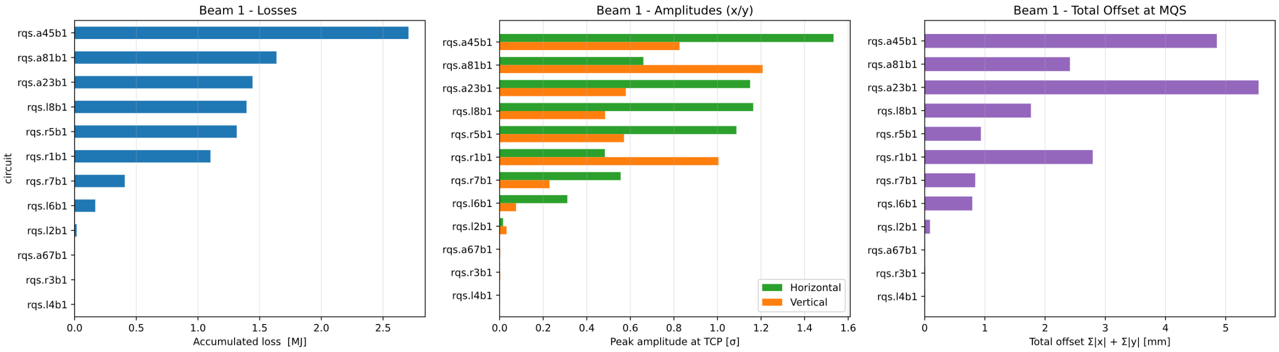

Figure 2: Xsuite simulation of single RQS circuit failures for Beam 1: accumulated losses at the primary collimators (left), peak horizontal and vertical oscillation amplitude at the TCP normalised to the beam σ (centre) and total orbit offset $\Sigma|x| + \Sigma|y|$ summed over the magnets of the circuit (right).

the largest losses (RQS.A45B1, RQS.A81B1, RQS.A23B1) are 4-magnet circuits, but several 2-magnet circuits (e.g. RQS.L8B1, RQS.R5B1, RQS.R1B1) also produce losses comparable to those of the 4-magnet families. A direct comparison of the three panels shows that the losses are generally larger where the total initial orbit offset at the magnets of the failing circuit is largest, and that the losses correlate more strongly with this initial offset than with the magnet count. A non-zero offset produces a dipolar feed-down from the skew-quadrupolar field, which kicks the beam into the collimators. Conversely, circuits such as RQS.A67B1, RQS.R3B1 and RQS.L4B1 show essentially no losses in the simulation because the closed orbit crosses the magnetic centre of their magnets in both the x and y plane.These per-circuit offsets are specific to the optics studied here; different optics, crossing-angle or dispersion-correction settings would shift the orbit at the RQS magnets and change which circuits appear critical.

IMPACT OF ORBIT OFFSET UNCERTAINTY

The previous conclusion, that losses are driven by the initial orbit offset at the failing circuit, is only useful if the offset

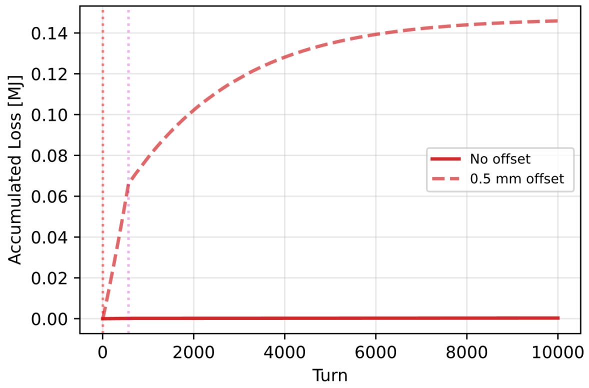

Figure 3: Time-dependent simulation of the RQS.A67B1 trip with a theoretically centred orbit (solid) and with an artificial 0.5 mm initial offset in the magnets (dashed). A sub-millimetre offset, which is within the accuracy of the BPM-to-magnetic-centre alignment, is sufficient to turn a no-loss case into a significant loss event.

itself is predictable. In practice, the comparison of measured (BPMs) and theoretical orbits in Fig. 1 shows an RMS difference in the order of 0.5 mm for the vertical plane. This can be explained by BPM alignment tolerances and residual orbit imperfections after correction. The observation that RQS.A67B1 caused two dumps in 2025, while showing essentially no losses in a simulation using the theoretical orbit, is a direct manifestation of this uncertainty.

To quantify the sensitivity to the initial orbit offset, the failure of RQS.A67B1 was re-simulated twice, once with the theoretical orbit (no initial offset at the magnets of the circuit) and once with an artificial 0.5 mm offset added to all magnets of the circuit. In both cases the simulation uses a time-dependent current waveform reproducing the actual converter fault behaviour, with an initial fast drop due to quenchback — a sudden increase of the magnet resistance triggered by the rapid current change induced by the EE system — followed by a slower exponential decay.

The result is shown in Fig. 3. Using the initial theoretical centred orbit, both the orbit shift at the TCP and the accumulated losses remain at zero. With an additional 0.5 mm offset, the orbit shift at the TCP becomes non-negligible and the accumulated losses reach ~0.14 MJ within a few thousand turns, sufficient to trigger a BLM dump in operation. A sub-millimetre offset is therefore enough to turn a circuit that appears safe in simulations into one that causes a beam dump. This explains the discrepancy between the simulation results of Fig. 2 and the operational behaviour of RQS.A67B1, and more generally prevents the use of the initial orbit offset as a stable selection criterion between LHC runs.

INTERLOCKING STRATEGY

Three parameters determine the criticality of an RQS failure: the number of magnets in series, the initial orbit offset in the magnets of the circuit, and the operational current. Only the first is a *stable* criterion, in the sense that it depends on the hardware configuration and not on the operational settings that may change every year.

Based on these considerations, the following interlocking strategy is proposed:

- **Baseline (stable criterion):** interlock all eight 4-magnet circuits, i.e. RQS.A23B1, RQS.A45B1, RQS.A67B1, RQS.A81B1 on Beam 1 and RQS.A12B2, RQS.A34B2, RQS.A56B2, RQS.A78B2 on Beam 2.
- **Addition for 2026 (operational history):** interlock the 2-magnet circuit RQS.L7B2, that was responsible for one of the BLM-triggered dumps in 2025, under the assumption that initial orbit and operational current remain similar to those of 2025.

This brings the total number of RQS circuits to be interlocked to 9 out of 24. The selected circuits are configured as *auxiliary* circuits in the PIC, so that the associated beam dump request via the BIS is *maskable* at low intensity. Operational flexibility is therefore preserved for commissioning and machine development. The deployment is purely a PIC reconfiguration that can be performed remotely with no impact on the hardware commissioning schedule, and was implemented during the 2025–2026 YETS.

A retrospective application of this strategy to Run 3 shows the expected trade-off between increased machine protection and availability. The four reactive BLM-triggered dumps of 2025 would have been replaced by pre-emptive PIC-triggered dumps, but two additional pre-emptive dumps would also have occurred for trips where the beam would have survived without interlocking. Early operational experience in 2026 already provides a first validation of the strategy. A QPS trip of RQS.A78B2 at −140 A with beam triggered a pre-emptive PIC dump. This circuit carries one of the largest initial orbit offsets among all RQS circuits (almost 2 mm in each of its four magnets), and would likely have caused significant BLM losses without the interlock.

CONCLUSION AND OUTLOOK

Xsuite tracking simulations show that the impact of an RQS circuit failure on the LHC beam is dominated by the initial orbit offset at the magnets of the failing circuit, and that a sub-millimetre offset is sufficient to turn a trip that appears safe in simulation into a BLM-dump event. Since the orbit and the operational currents are not stable criteria from year to year, the prioritisation for the 2026 run relies on the only hardware-driven criterion (number of magnets in series), complemented by the operational history of 2025. Nine RQS circuits have been interlocked as auxiliary, maskable PIC channels during the 2025–2026 YETS, and early 2026 operation confirmed the relevance of the strategy. The interlocking will be reassessed after the high luminosity upgrade of the LHC, when the new optics and orbit configuration will be consolidated and may justify a different selection of 2-magnet circuits. The 4-magnet circuits are expected to remain interlocked.

REFERENCES

[1] O. S. Brüning *et al.*, *LHC Design Report*. Geneva, Switzerland: CERN, May 2004. doi:10.5170/CERN-2004-003-V-1

[2] C. Lannoy, “Interlocking proposal for skew quadrupoles (RQS)”, presented at the 265th Machine Protection Panel meeting, CERN, https://indico.cern.ch/event/1635712/, Jan. 2026,

[3] A. Antoine, “Configuration change of PIC to interlock all RQS circuits”, presented at the 155th Machine Protection Panel meeting, CERN, `https://indico.cern.ch/event/682863/`, Nov. 2017,

[4] G. Iadarola *et al.*, “Xsuite: An integrated beam physics simulation framework”, in *Proc. HB'23*, Geneva, Switzerland, pp. 73–80, Mar. 2024. doi:10.18429/JACoW-HB2023-TUA2I1