The final version of record is available at http://dx.doi.org/10.1109/TASC.2017.2772805

A Statistical Analysis of Electrical Faults in the LHC Superconducting Magnets and Circuits

L. Bottura, J.Ph. Tock, B. Auchmann, M. Bednarek, D. Calcoen, Z. Charifoulline, G.J. Coelingh, G. D'Angelo, P. Hagen, S. Le Naour M. Mentink, A. Milanese, M. Modena, V. Montabonnet, A. Perin, M. Pojer, F. Rodriguez-Mateos, I. Romera, C. Scheuerlein, R. Schmidt, A. Siemko, M. Solfaroli, J. Steckert, , E. Todesco, A. Verweij, G. Willering, D. Wollmann, M. Zerlauth

***Abstract*—The Large Hadron Collider (LHC) at CERN has been operating and generating physics experimental data since September 2008, and following its first long shut down, it has entered a second, 4-year long physics run. It is to date the largest superconducting installation ever built, counting over 9 000 magnets along its 27 km long circumference. A significant operational experience has been accumulated, including the occurrence and consequences of electrical faults at the level of the superconducting magnets, as well as their protection and instrumentation circuits. The purpose of this paper is to provide a first overview of the most common electrical faults and their frequency of occurrence in the first years of operation, and to perform a statistical analysis that can provide reference values for future productions of similar dimensions and nature.**

***Index Terms*— Superconducting Magnets, Statistics, LHC, Electrical fault**

I. Introduction

The Large Hadron Collider (LHC) at CERN is the largest particle accelerator presently under operation. Since its initial commissioning in 2008, it has produced an integrated luminosity of over 90 fb^{-1} at energies up to 13 TeV in the center of mass of the collisions. This represents a significant experience acquired in operating 9 398 superconducting (SC) magnets powered through 1572 independent circuits rated from 60 A to 13 kA and located in all the sections of the LHC: arcs, dispersion suppressors (DS) and long straight sections (LSS) [1]. Table I lists the 9 398 SC magnets of the LHC machine, their locations and the corresponding number of powering circuits. Corrector magnets account for more than 80% of the total number of magnets and more than 90% of the powering circuits.

The number of magnets, the size of the LHC and the 10-years' timespan since its commissioning has generated a unique set of data on the electrical faults of large systems involving SC magnets. Indeed, in 2016, the LHC operated for 213 days with 153 dedicated to physics data generation. Overall, unavailability due to faults accounted for about one quarter of this time. Amongst them, magnet circuits were only at the origin of 75 hours of fault, i.e. about 8% of the total downtime [2], [3]. The global LHC downtime was respectively 31% and 26% for 2015 and 2016 respectively. So far, SC magnets were not amongst the four first contributors to LHC unavailability.

The authors are with CERN, European Organization for Nuclear Research, 1211 Geneva 23, Switzerland. Corresponding author : e-mail: Jean-Philippe.Tock@cern.ch, phone: +41 22 767 2970

Template version 7.2a, 04 August 2016

Table I
LHC Superconducting Magnets And Circuits Inventory

Quantity	Function	Location	Circuits
1232	Main Bending dipoles	Arc, DS	8
392	Main Lattice quadrupoles	Arc, DS	16
18	Separation/Recombination dipoles	LSS	16
142	Insertion quadrupoles[1]	DS, LSS	86
1044	Dipole correctors	Arc, DS, LSS	1044
512	Quadrupole correctors	Arc, LSS	200
3232	Sextupole correctors	Arc, DS, LSS	112
1584	Octupole correctors	Arc, DS, LSS	64
1232	Decapole correctors	Arc, DS	16
8	Dodecapole correctors	LSS	8
2	Undulators	LSS	2
9398	**TOTAL**		**1572**

.[1]Including triplets quadrupole magnets.

Despite the unprecedented number of SC magnets in a single installation, the statistics are low, with the LHC being at the beginning of its operational life. For the LHC, the material and human investment is major, i.e. an effort that can be repeated only a few times per century. Hence, and in spite of the above-mentioned difficulties, we make a first attempt to use the experience gained during LHC commissioning and operation in the last decade, suitably screened to focus only on relevant electrical faults, to provide estimates for the expected electrical lifetime and failure rate of the LHC SC magnets. The intention is to support long-term operation of the LHC, but also to provide a digested benchmark for future projects of similar scope, such as the Future Circular Collider (FCC) study at CERN [4].

In the paper, we start recalling the most common electrical faults observed, defining categories that can be used for further interpretation. These elements are then taken as a basis for a failure rate analysis based on Weibull statistics, to compute the probable failure rate and expected lifetime of the magnets. This is finally used to estimate the probability that LHC operation may be interrupted by a magnet fault requiring a lengthy, unplanned magnet exchange.

II. A Catalog of Faults

The origins of faults of SC magnet systems were reported and discussed in previous work. In particular, [5]-[7] provide a good summary and basis for statistics. In essence, the causes of highest incidence are insulation (electrical) faults, accounting for about one third of all recorded system failures. It is interesting to note that under-performing superconductor, possibly the most difficult technology in the magnet, is not the main issue (less than 20 % of system failures).

We focus in this analysis on electrical faults, which can be originated by various causes that we broadly divide in the following families:

- Electrical stress associated with operation (magnet ramps), or quench and ensuing fast discharge. We expect the main magnets to experience O(10^4) powering cycles, and O(10) natural quenches during operation. An additional number of quench discharges will be experienced when including commissioning and diagnostics;
- Mechanical loading and fatigue on coil, structure, busses. These are associated with magnet powering, where the number of cycles during the machine lifetime is O(10^4) per magnet, as well as thermal cycles, which are expected to be O(10) for the whole LHC;
- Radiation and associated degradation of mechanical and electrical strength. Dose is in in the range of O(10) MGy, as expected on the magnets in the triplet region P1 and P5, and in the collimator regions P3 and P7.

Of the above categories, the LHC has already experienced a significant mechanical and electrical stress, having reached powering conditions within 10% of nominal, but the radiation dose is still much below the projected value for the whole lifetime. In fact, locations with high radiation dose have been identified and are monitored, with plans for a protection, preventive maintenance and upgrade that should avoid any visible degradation. For this reason, we neglect radiation effects in our analysis. The fact that the average number of faults located in the high radiation zones is the same as in the other zones supports this approach.

A total of 48 electrical faults have been observed since the LHC commissioning in 2006. We have excluded from the statistics faults such as magnet quenches during commissioning, including training quenches, propagated quenches and quenches during operation at constant current (8 in the main bending dipole and lattice quadrupoles) [9]. We also neglect non-conform values of the resistance of the magnet splices at the level of the connection between poles and apertures, a non-conformity that was discovered in a part of the dipole production after installation [8]. As long as the above behavior and values do not change during the lifetime of the machine, we consider them as built-in features of the magnet itself, rather than a developing fault. [9]-[11].

For our analysis, we have further divided the faults in the following broad categories:

- Dielectric strength, i.e. a failure to withstand the high voltage test performed regularly to qualify the circuits for operation or detected by the on-line continuous monitoring system. This category covers non-conform ground insulation, and insulation between different circuits or between turns of a magnet. We further define a sub-category related to dielectric strength issues at the level of the dipole protection cold diode, discussed later;
- Increase of the splice resistance and/or development of a splice resistance above the specification value;
- Open circuit, a lack of continuity preventing a circuit to be powered;
- Quench heater strips failures, of any type (open circuit, dielectric strength) [11];
- Any other fault, including instrumentation and powering issues.

A chart of the various faults is reported in Fig. 1. It is evident that the dominant fault mode is a reduction or loss of dielectric strength, accounting for 25 (16+9) faults, summing up the two first categories i.e. diodes insulation faults are also considered as a loss of dielectric strength at the level of the dipole protection diode box. As they are quite numerous and accounting for one third of all the dielectric strength faults, a specific category was created. We expand on this failure mode below, in a dedicated section. It is also clear that quench heaters are a delicate part of the system, with a significant number of faults registered during the machine lifetime. This has been traced to specific manufacturing issues, i.e. (a) insufficient quality of the soldered connection of the quench heater strips to the powering leads, and (b) a folding deformation of the quench heater strips taking place during collaring. Finally, circuits are monitored continuously during operation. This diagnostics generates warning in case the resistance changes as a function of time, which then triggers close examination. Increased splice resistance could point to a fatigue phenomenon. So far, all high resistance circuits are confined to low-current correctors, and the resistances have been found to be stable with time.

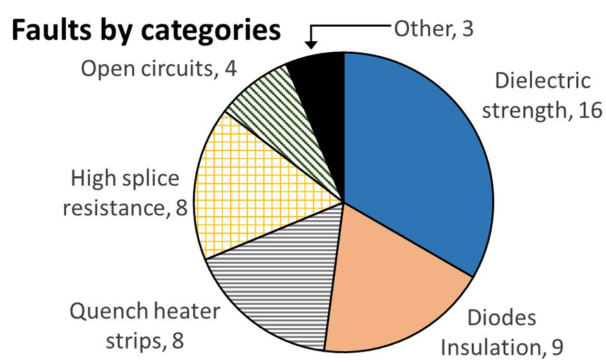

Fig. 1. Distribution of registered electrical faults by the category defined in the text.

A. Dipole Cold Diodes

Since the installation of the SC magnets in the LHC tunnel, 9 short-circuits to ground were registered in the containers of the dipole protection cold diode. They are primarily caused by metal debris falling on live bus-bars of the cold diode bypass circuits. They are present in the cold masses since their manufacturing and are transported by helium flows. Such helium flows take place during cool-down and warm-up phases but also in case of main magnet quenches. Seven cases were detected at room temperature and easily fixed by opening the helium enclosure and cleaning the diode containers. The two other cases took place during training campaigns, at 1.9 K.

They were solved thanks to a capacitive discharge, using a so-called "earth fault burner". [13]. Out of 255 quench events at high current, 2 have provoked a subsequent short to ground. An analysis of this fault mode led to the conclusion that the risk associated with the extensive training campaign of the LHC dipoles required to reach the nominal beam energy of 7 TeV [9],[14] is not acceptable. Specifically, training at high current leads to secondary quenches, increasing the helium flows, and so enhancing the severity of this issue. Methods to reinforce the electrical insulation in this critical location are being developed at CERN. This systematic intervention on the 1232 main dipole cold masses will require the opening of the corresponding interconnections, including the helium enclosure to obtain access to the diode container. The electrical insulation of the cold diode containers will be reinforced to reduce the risk of a short to the minimum. After this large intervention, we are confident that this fault will not occur anymore; being cured for the rest of the LHC operational lifetime.

B. *Operational experience*

For the main magnets, the highest requirements on their robustness was applied. When possible, as in the case of quench heater circuits for example, redundancy was implemented. In case of failure of one quench heater, it is possible to reconfigure the protection scheme to allow pursuing operation up to the next planned long technical stops. Only in one case, when a potential interturn short in a cryodipole was suspected, the replacement of this cryomagnet took place in a long technical stop though it was only planned a few month in advance. The impact on the schedule was minimized, especially thanks to the stock of available spare magnets.

For the corrector magnet circuits, there is some redundancy in the number of available circuits. As the field quality of the main magnets is better than initially specified, the available margin is even higher. Condemning one circuit has a minimum impact on operation. Thanks to the implementation of spare busbars in the corrector magnets powering circuits, it is often possible to apply a solution that minimizes the number of unavailable magnets, thus restoring almost the full correction capabilities. This fix can only be applied with a machine at room temperature so it is carried out during a planned long technical stop.

III. Statistical Analysis

The concept of reliability, and specific indicators such as the Failure Rate (FR), or the Mean Time To Failures (MTTF), are a part of well-established engineering practice for components produced in series and operated over sufficiently long time to accumulate relevant statistics. As mentioned earlier, this is not the case for a large part of the magnetic system of the LHC, and especially for the SC magnets. These are in several cases first-of-a-kind productions, barely beyond the prototyping stage (e.g. inner triplet). In any case, the accumulated statistics on failure rate and failure consequences is not sufficiently high (namely for main magnets) to allow for an extrapolation using standard methodology. In addition, if we try to put it in the framework of a Failure Mode and Effect Analysis (FMEA), the LHC may still be in the "early-operation failure" regime of the bathtub curve describing failure rate vs. time.

In spite of these *caveats*, we decided to use the present statistics on electrical faults to attempt an extrapolation of the rate of occurrence of such events for the duration of the LHC lifetime, as this is a large distributed superconducting electrical machine. We consider that the basic unit for the analysis is an instrumented *cryomagnet*. This, in the LHC jargon, is the cryostated cold mass containing an assembly of single SC magnets with different optical functions. The single magnets in the cryomagnet are instrumented, possibly complemented by quench heaters and protection diodes, and connected to bus-bars that lead to the terminals at the end of the cryomagnet. The cryomagnet is the unit built and qualified by tests installed in the accelerator as a single object and interconnected to form the powering circuit. A total of 1748 cryomagnets are installed in the LHC.

The LHC cryomagnets, though in great variety and types, share a great deal of common technology. In our analysis we assume that they are identical from the point of view of the probability to fail, so that we can cumulate the inventory of faults described earlier, and assign them to all types of cryomagnets. It should be understood that one such failure does not necessarily entail a complete loss of performance. However, one such event will eventually call for a maintenance operation at the level of the cold mass, or a magnet exchange, during a technical stop of sufficient length, or shutdown period. As an example from the previous discussion on fault consequences, a number of corrector magnets have been condemned due to electrical faults, without disruption of the machine performance. In the meantime, investigations are pursued to prepare an intervention on the affected cryomagnets, to restore the whole or a part of the circuit by intervening at the next opportunity. This invariably entails long times for the warm-up, radiation cooling and recommissioning that impact the amount of integrated luminosity delivered by the LHC, irrespective of the magnet affected. In practice, this kind of interventions are regrouped and carried out in the frame of planned long shutdowns. This is also why we do not distinguish among the fault types for this first analysis.

Finally, we have summed faults attributed to the cryomagnets themselves with the ones that are linked to the cold part of the powering circuit. Both types represent approximately the same number of faults.

To estimate the FR and MTTF we consider each fault as a *failure*, and generate a cumulated plot of failures versus time, normalized to the total number of cryomagnets. The result is shown in Fig. 2.

The data is then modeled with a two-parameter Weibull distribution, as customary in failure analysis:

$$F = 1 - e^{-(t/\lambda)^{k}}, \tag{1}$$

where F is the distribution, t is the time from the beginning of commissioning tests in years, λ is the characteristic time of the failure mode, and k is a parameter indicating whether the failure rate is constant ($k = 1$), increases ($k > 1$, ageing) or decreases ($k < 1$, early-operation failure) in time. The distribution has a Mean Time To Failure (MTTF) of:

$$MTTF = \lambda\, \Gamma\left(1 + \frac{1}{k}\right) \quad (2)$$

and predicts a failure rate given by:

$$FR = \frac{k}{\lambda}\left(\frac{t}{\lambda}\right)^{k-1} \quad (3).$$

The result of the fit is also reported in Fig. 2, as well as the parameters obtained from the fit, i.e. λ = 1764 ± 976 (yrs) and k = 0.71 ± 0.05 (-). The uncertainties on the model coefficients, computed by applying error propagation formulae, correspond to a 95% confidence level. A first comment on these values is that the characteristic time of the failures corresponds to a MTTF (Eq. (2)) in the range of 2000 years. Given the 1748 cryomagnets installed, and ignoring the shape of the failure distribution function, we hence expect that at least one cryomagnet would develop an electrical fault per year. The second observation is that the parameter k is smaller than one, even when considering the uncertainty of the fit coefficient. This seems to indicate that the system is in a regime of *early-operation failure*, and we expect the present failure rate to be above the MTTF, but to reduce over the coming years.

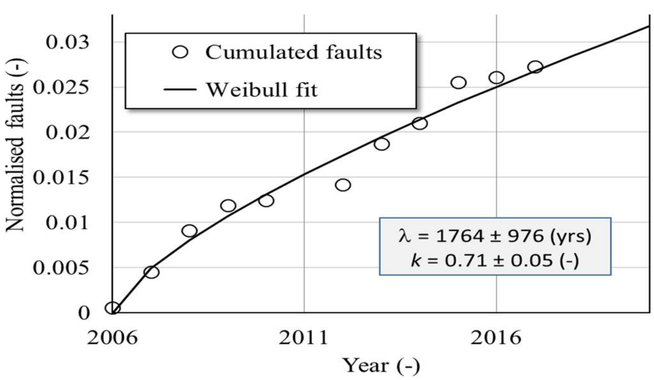

Fig. 2. Plot of normalised cumulated faults of electrical nature in the LHC SC magnets vs. time of detection of the failure.

This can be evaluated more formally using the failure rate FR defined in Eq. (3), and plotted in Fig. 3 where we compare it to the binned distribution of registered faults, with bin equal to one year. The registered faults show large variability, which is due to the fact that the machine goes through cycles of normal operation, maintenance and commissioning, discussed below. The average value of the yearly FR in the period considered is ≈ 4 cryomagnets/year, reducing in time. A projection to the expected end of life of the LHC, after 35 years of operation, gives a FR of ≈ 2 cryomagnets/year, still significant.

The typical cycle of operation and maintenance for the LHC is 5 years, with 3 years of operation, and 2 years of maintenance, including the associated re-commissioning. Long shutdowns took place in 2008-2009 and in 2013-2014. They were obviously preceded by a warm-up to room temperature and followed by re-cool-downs and re-commissioning phases. It is interesting to note in Fig. 3 that the number of faults registered during stable operating periods (roughly 2010-2012 and from 2016 onwards) are lower than during other phases. This confirms that electrical faults are triggered by stresses generated by thermo-mechanical cycles like the global warm-ups and cool-downs but also quenches occurring during training and commissioning campaigns. Most faults are detected during the commissioning or the interventions taking place during long shutdowns. Though it cannot be excluded that the faults developed during commissioning, it is very likely that they were present before but not detected. Therefore, the yearly binning is probably not the most adequate one and a further refinement by phases can bring additional information. This can be the subject of a deeper analysis with more solid statistics.

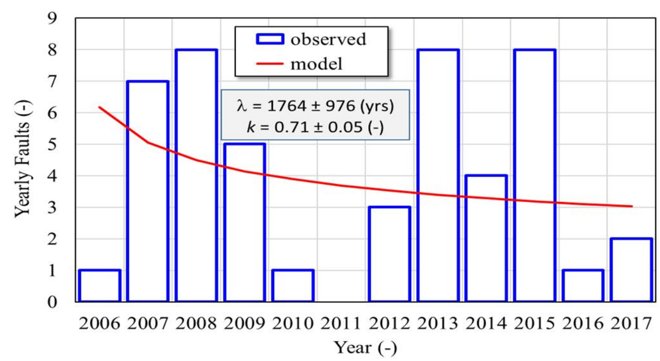

Fig. 3. Observed and modeled failure rate, binned on yearly basis. The model is obtained from the Weibull distribution fit described in the text.

IV. Conclusion

We have presented a first statistical analysis of the time profile of the occurrence of electrical faults during LHC operation. We based the analysis on the definition of a cryomagnet as the unit circuit element in the accelerator, and applying Weibull statistics to the observed electrical faults, once systematic features are removed. Though we recognize that this analysis is approximate, as the statistics is not extensive, the obtained indicators are nonetheless interesting in showing a system that is operating at an approximate decreasing failure rate. The different behaviours depending on the phases, i.e. thermal cycles, commissioning, operation, are also evidenced. The yearly failure rates (≈ 4 in the last 10 years, decreasing to ≈ 2 at the end of LHC lifetime) have been used to establish maintenance plans for the cryomagnets, as well as the size of the stock of spares that will cover long-term operation. We finally believe that this analysis can provide a useful bench-mark for future projects such as the next step in accelerators for High Energy Physics, an FCC [4], or other large-scale SC magnet systems such as ITER [15] and DEMO [16].

Acknowledgment

The analysis presented in this paper would not have been possible without the dedication of the Members of the “LHC Magnet Circuits, Powering and Performance Panel” who identify and track all electrical performance issues.

REFERENCES

[1] O. Bruning et al. (eds) LHC Design report,Volume 1, The LHC Main Ring, CERN-2004-03
[2] J. Wenninger, “Approaching nominal performance at LHC”, Proceedings of IPAC2017, Copenhagen, Denmark
[3] A. Apollonio et al., “LHC accelerator fault tracker – First experience”, Proceedings of IPAC2016, Busan, Korea
[4] M. Benedikt, F. Zimmermann, “Status of the Future Circular Collider Study”, Proceedings of RuPAC2016, St. Petersburg, Russia
[5] D.B. Montgomery, Review of Fusion Magnets System Problems, Proc. 13th IEEE Symp. Fus. Eng., 27, 1989.
[6] Y. Iwasa, Case Studies in Superconducting Magnets, Plenum Press, 1994.
[7] J.H. Schultz, in Engineering Superconductivity, J. Wiley & Sons, 2001.
[8] Z. Charifouline et al., “Resistance of splices in the LHC Main Superconducting Magnet Circuits at 1.9 K”, Magnet Technology Conference 25, Amsterdam 2017
[9] G. Willering et al., “Performance of CERN LHC Main Dipole Magnets on the Test Bench from 2008 to 2016”, *IEEE Trans.Appl.Supercond. 27 (2016) no.4, 4002705*
[10] E. Todesco et al., “Training Behaviour of the Main Dipoles in the Large Hadron Collider”, ASC2016, Denver, 2016, IEEE Transactions on Applied Superconductivity (Volume: 27, Issue: 4, June 2017)
[11] J. Ph. Tock et al., “The Consolidation of the LHC Superconducting Magnets And Circuits: “, Magnet Technology conference MT24, IEEE Trans. Appl. Supercond. 26 (2016) 4002706
[12] J. Ph. Tock et al., “Consolidation of the LHC superconducting circuits: a major step towards 14 TeV collisions”, IPAC’12, New Orleans, 2012, Proceedings of IPAC2012, New Orleans, Louisiana, USA, pp. 3575-3577
[13] M. Bednarek et al., “Design, Assembly and Use of a Device to Eliminate Earth Faults Caused by Metallic Debris in the LHC Main Dipole Circuit”, MT25 Amsterdam 2017
[14] E. Todesco et al. “Training of the Main Dipoles Magnets in the Large Hadron Collider towards 7 TeV Operation”, MT25 Amsterdam 2017
[15] N. Mitchell et al., “The ITER Magnets: Design and Construction Status”, *IEEE Trans. Appl. Supercond.*, vol. 22, no. 3, June 2012, 4200809
[16] G. Federici et al., Overview of the design approach and prioritization of R&D activities towards an EU DEMO, Fus. Eng. Des. 109-111 (2016) 1464–1474