ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-TUPAB325

DATA-DRIVEN RISK MATRICES FOR CERN'S ACCELERATORS

T. Cartier-Michaud*, A. Apollonio, G. Blarasin, B. Todd, J. Uythoven
CERN, Geneva, Switzerland

Abstract

A risk matrix is a common tool used in risk assessment, defining risk levels with respect to the severity and probability of the occurrence of an undesired event. Risk levels can then be used for different purposes, e.g. defining subsystem reliability or personnel safety requirements. Over the history of the Large Hadron Collider (LHC), several risk matrices have been defined to guide system design. Initially, these were focused on machine protection systems, more recently these have also been used to prioritise consolidation activities. A new data-driven development of risk matrices for CERN's accelerators is presented in this paper, based on data collected in the CERN Accelerator Fault Tracker (AFT). The data driven approach improves the granularity of the assessment, and limits uncertainty in the risk estimation, as it is based on operational experience. In this paper the authors introduce the mathematical framework, based on operational failure data, and present the resulting risk matrix for LHC.

INTRODUCTION

CERN's accelerators have been successfully operated for many years, both in terms of beam performance and machine availability, thanks to the experience developed in system design, operation and maintenance. Traditionally, accelerator systems were designed based on the know-how from experts, physicists and engineers, and on the lessons learned from previous machines. The increased damage potential of the LHC and its size and complexity changed this paradigm. New requirements for machine protection were therefore derived from reliability engineering best practices and tools. The LHC machine protection was designed according to European standards for safety-critical electronic systems, based on the concept of risk matrices and Safety Integrity Levels (SIL) [1,2]. At that time, the LHC failure modes were identified by experts and assigned estimated probabilities/frequencies of occurrence and severity/consequences [3]. Risk matrices adapted to the LHC were first presented in [4].

After several years of operation, we are now in the position to feedback operational experience into the definitions of these risk matrices, to more accurately evaluate the reliability requirements for future system designs, in particular for the different upgrades of the LHC.

In the first section of this paper, the concept, parameters and limitations of traditional risk matrices are recalled. The data-driven risk curves are introduced in the second section. Using risk curves based on the data gathered in the Accelerator Fault Tracker (AFT) [5], data-driven risk matrices tailored to CERN's accelerators are presented in the third section.

LIMITATIONS OF RISK MATRICES

Risk matrices are used as a tool for risk management and decision making, allowing to map the failure modes of a system or machine in a 2D table, discretized with respect to the likelihood of a failure mode (the y-axis in this paper) and its consequences (x-axis, see Fig. 1). Both axes could be expressed in any relevant unit, the likelihood could be a frequency or a probability, the consequence could use several reference quantities, as recovery time or cost impact. Both dimensions can be scaled in a quantitative way, e.g. 'one failure per month' or '10 % probability of occurrence', or in a qualitative way, e.g. 'rare/frequent/certain', 'low/high probability'. The progression of both scales is generally logarithmic to better cover wide ranges of frequencies and consequences.

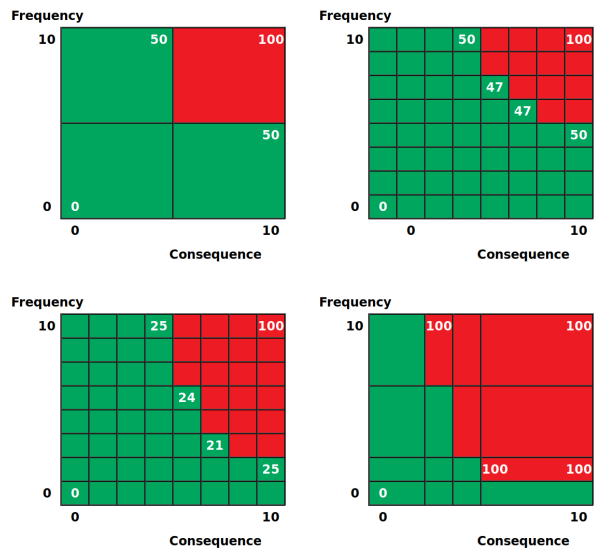

Figure 1: Four risk matrix sketches. The red color defines unacceptable risks, the green defines acceptable risks. Top left, low resolution risk matrix with the threshold $T = 50$. Top right, high definition risk matrix. Bottom left, definition of the risk as $R(F,C) = F \times C^2/10$ and $T = 25$. Bottom right, high definition and custom definition of the risk.

Given F a frequency and C a consequence, the risk R is defined as $R(F,C) = F \times C$. The acceptability of different risks is defined using a threshold T to split the space $(F \times C)$ in two areas: "acceptable" and "unacceptable". The failure modes of a system are then placed onto the risk matrix, and critical ones can be identified according to these definitions and mitigation actions put in place to reduce either the frequency and/or the consequences of critical failure modes.

* thomas.cartier-michaud@cern.ch

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-TUPAB325

The main limitation for the use of risk matrices is in the uncertainty of the estimates for likelihood and consequences of a failure mode, as these have to be defined very early during the design and potentially for a new system, for which previous experience might not be available. Placing failure modes in a risk matrix defined a-priori, with very limited accuracy in the definition of the axes' discretization (see Fig. 1 top left), can lead to significant over- or underestimation of the risk. With a refined discretization in both dimensions (see Fig. 1 top right), the classification of each cell of the matrix in either "acceptable" or "unacceptable" offers the possibility to use alternative cost functions $R(F, C) = F \times C^2/10$. With this different function, a higher weight on the consequence is put compared to the frequency in order to penalise the faults with important consequences (see Fig. 1 bottom left).

As an example, in accelerator-driven systems such as MYRRHA, an accelerator is coupled with a sub-critical nuclear reactor. For these type of accelerators, faults of a duration longer than few seconds have to be avoided because of mechanical stresses induced on the reactor in case of a beam stop and the heavy restart procedures [6]. For such a system, a configuration with many faults with a duration shorter than a second is preferable to a configuration leading to a few faults longer than a few seconds, as shown in Fig. 1 bottom right. This example of risk matrix uses the minimum number of discretization points in both axes to match the desired frontier of acceptability. It can be convenient to still keep additional discretization points, such as Fig. 1 top right and bottom left, because a risk matrix is often reused for different applications by simply changing the threshold T.

The choice of a particular discretization and the frontier of acceptability, i.e. function $R(F, C)$ and threshold T, should be optimized for each application. This optimisation is performed in this paper using a data driven approach on systems already in operation, the LHC and CERN's injector complex, with the aim of limiting the loss of information when translating the data used as an input into the resulting risk matrix as described in the next section.

DATA-DRIVEN RISK CURVES

The experience based failure data from the different accelerators at CERN is stored in the Accelerator Fault Tracker (AFT) [5]. The AFT data have been collected and reviewed by a team of experts from the beginning of 2015 for the LHC and from 2017 onwards for the injectors. The proposed metrics used for the CERN's risk matrices are based on failure frequency and recovery time. The latter is chosen to measure the time in which the accelerators are not available to produce physics or provide beam to the downstream machines. Taking the data from AFT without any additional filtering leads to risk matrices which directly reflect the past operation. The availability of the LHC and its injectors has been remarkable in the last years [7–9]. With the aim of improving future operation and providing guidelines for new system designs, a selection of faults coming from AFT could be considered today as "unacceptable", and thus are excluded from the data used in the process of generating risk matrices and risk curves.

In order to avoid the discretization problem described in the previous section, we introduce the use of risk curves. A continuous risk curve is defined across all frequencies and recovery times based on the existing failure data. This is performed by approximating existing data with a curve. Using a parameter α, one defines the interval $I_{\alpha,d_i} = [d_i/\alpha, d_i \times \alpha]$ for a failure of duration d_i. Given a fault $d_1 = 25$ min, one could say this fault is of the same order than a fault $d_2 n = 60$ min with an $\alpha = 3$ margin because $25 \in [60/3, 60 \times 3]$ but it is not true if $\alpha = 2$ as $25 \notin [60/2, 60 \times 2]$. I_{α,d_i} is then used to define the function evaluating the frequency of faults of a duration d with an α margin: $F_\alpha(d) = \sum_{i=1}^{N} \mathbb{1}_{I_{\alpha,d_i}}(d)/D$, with $\mathbb{1}_{I_{\alpha,d_i}}$ the indicator function which equals 1 if $d \in I_{\alpha,d_i}$ and 0 otherwise, N the number of faults used by the data driven approach, d_i being the duration of the i-th fault and D the period of time used to acquire the list of faults d_i.

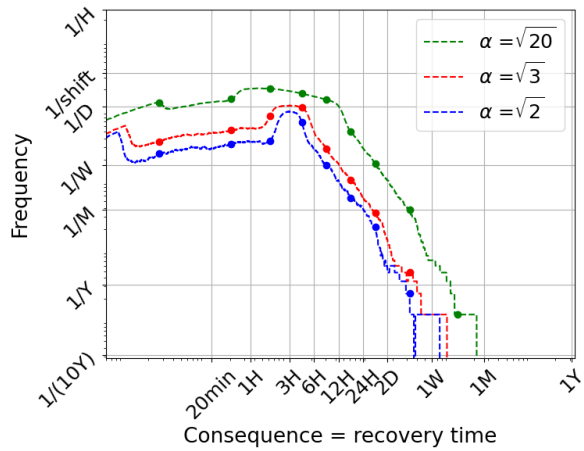

Figure 2: $F_\alpha(d)$ data-driven risk curves based on LHC fault recovery times for three different values of α.

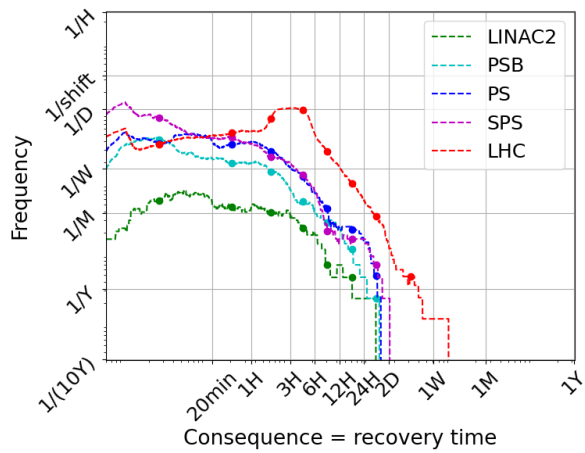

Figure 3: $F_\alpha(d)$ data-driven risk curves based on LINAC2, PSB, PS, SPS and LHC fault recovery times for $\alpha = \sqrt{3}$.

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-TUPAB325

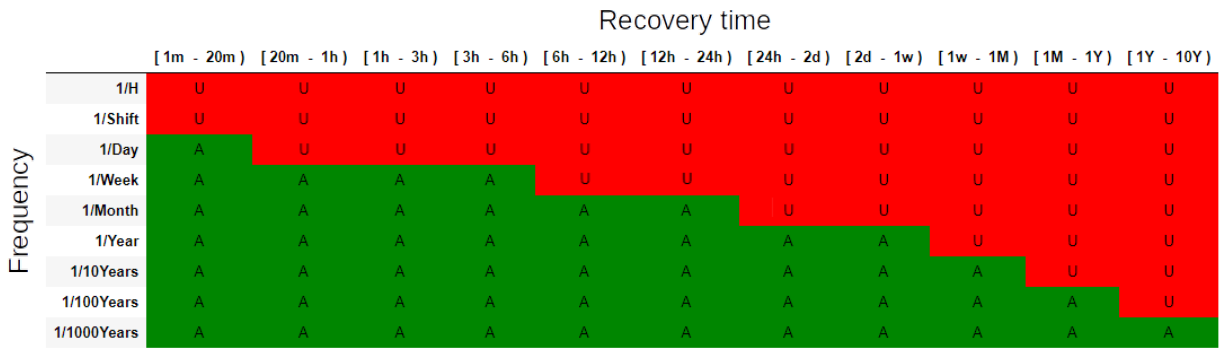

Figure 4: LHC data-driven risk matrix derived from four years of operation. Green cells are considered "acceptable", red cells are "unacceptable".

In Fig. 2, LHC data-driven risk curves $F_\alpha(d)$ are represented for three values of α. For the considered use cases, the values of α which make the more sense are larger than $\sqrt{2}$ and smaller than $\sqrt{20}$. Inferior α would introduce too many intervals $I_{\alpha,d}$, e.i. too many categories of faults, to cover the relevant range of duration considered, from minutes to years. Superior α would mix faults of too different duration to be considered equivalent from an operational point of view. In Fig. 3, curves are represented also for the accelerators in the CERN's injector complex for $\alpha = \sqrt{3}$. Beams are produced in LINAC2, then transferred to the Proton Synchrotron Booster (PSB), which feeds the Proton Synchrotron (PS), feeding itself the Super Proton Synchrotron (SPS), the last accelerator before the LHC. The integral of the risk curve, which is a measure for the unavailability of an accelerator, increases from LINAC2 to the LHC. This is explained by the increasing complexity, size and number of faults of each downstream accelerator as compared to the previous one, closer to the beam source.

DATA-DRIVEN RISK MATRICES FOR CERN'S ACCELERATORS

The previous section showed that data-drive risk curves help to better define the discretization of risk matrices with respect to the frequency and consequences of failures, allowing to identify where to place discretization points according to the main variations of the curves. Adding those points reduces the errors in the risk estimation with too coarse discretization intervals. In this section additional practical aspects are discussed for defining discrete risk matrices with the help of risk curves. It is important to include feedback of accelerator experts from a design and operational point of view. It is interesting to notice that the ratio between the highest and lowest recovery time in each interval retained, as shown in Fig. 4, is at most 20 for the first column and at least 2. It is equivalent to use $\alpha = \sqrt{20}$ and $\alpha = \sqrt{2}$ when computing data-driven risk curves as the ratio between the two boundaries of $I_{\alpha,d}$ is α^2. For the LHC, this translates into nine frequency intervals (nine rows) as can be seen in Fig. 4. The frequency of 1/shift (1/8h) is used to reflect what operators potentially encounter during one of the three daily shifts. The observed failures lie in the first seven rows. As the longest stop observed in operation since 2015 has been of the order of one week for the LHC, observed failures are within the first eight columns out of eleven. In the "recovery time" dimension, the discretization reflects the way machines are operated and maintained. Failures of less than 20 minutes can usually be resolved remotely or do not require any action. Failures that are resolved remotely or self correcting require less than 20 minutes. Failures in the range of 3-12 h generally require access in the machine. Failures of 24 hours or more are typically outliers that required dedicated follow-up and recovery procedures.

To be able to place failure modes with very low frequency or very high consequence, an extension of the data-driven risk curves is necessary. This is done analytically, following a simple linear model. Given U_{OE}, the measured unavailability of a machine according to the "Observed Events", f_{RE} is a factor defining an acceptable unavailability budget due to the possible occurrence of "Rare Events" with high-impact, such that $U_{RE} = U_{OE} \times f_{RE}$. The total unavailability according to the risk matrix is then $U = U_{OE} + U_{RE}$. The value of f_{RE} depends on the damage potential of the machine and the operational experience. For the LHC, this factor is fixed to 10% and U_{RE} is distributed over the four last columns of the risk matrix.

CONCLUSIONS AND OUTLOOK

A data-driven approach to build risk matrices for systems based on historical failure data has been presented. The failure data of CERN's accelerators have been collected and reviewed by experts with the AFT. Data-driven risk curves have been introduced as an intermediate step, providing insights on the optimal risk-matrix discretization, the assumed uncertainty margins, and the operational conditions for the different accelerators. The result is a series of data-driven risk matrices for the LHC and CERN's injector complex. Data-driven risk matrices are already in use at CERN to define the acceptable failure frequency for newly designed systems or upgrades of the LHC. However, the approach presented in the paper is generic and applicable to any system or facility for which historical failure data are available.

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-TUPAB325

REFERENCES

[1] B. Todd, A Beam Interlock System for CERN High Energy Accelerators, 2006, http://cds.cern.ch/record/1019495/files/thesis-2007-019.pdf

[2] M. Kwiatkowski, Methods for the Application of Programmable Logic Devices in Electronic Protection Systems for High Energy Particle Accelerators, 2013, https://cds.cern.ch/record/1632194/files/CERN-THESIS-2013-216.pdf

[3] R. Schmidt, Safe LHC parameters generation and transmission, 2006, https://edms.cern.ch/ui/file/810607/0.1/LHC-CI-ES-0004-00-10.pdf

[4] M. Blumenschein, J. Spasic, J. Steckert, and J. Uythoven, "An Approach to Reliability Assessment of Complex Systems at CERN", presented at 65th Annual Reliability Maintainability Symposium, 2019, pp. 1–6, unpublished.

[5] Accelerator Fault Tracker (AFT), aft.cern.ch.

[6] F. Bouly, M. A. Baylac, A. Gatera, and D. Uriot, "Superconducting LINAC Design Upgrade in View of the 100 MeV MYRRHA Phase I", in *Proc. 10th Int. Particle Accelerator Conf. (IPAC'19)*, Melbourne, Australia, May 2019, pp. 837–840. doi:10.18429/JACoW-IPAC2019-MOPTS003

[7] A. Apollonio, B. Todd *et al.*, LHC availability 2016: Proton run, http://cds.cern.ch/record/2237325.

[8] A. Apollonio, B. Todd *et al.*, LHC availability 2017: Proton run, https://cds.cern.ch/record/2294852.

[9] A. Apollonio, B. Todd *et al.*, LHC availability 2018: Proton run, https://cds.cern.ch/record/2650574.