



Partitioning And Transmuter Research Initiative in a Collaborative Innovation Action

PATRICIA

Grant Agreement Number 945077

Research and Innovation Action

Activity: NFRP-2019-2020

Topic: NFRP-2019-2020-07 Safety Research and Innovation for Partitioning and/or Transmutation

Start date: 01/09/2020 – Duration: 48 months

DELIVERABLE

D10.1 Accelerator reliability studies with emphasis on the fault tolerance capabilities

Authors: Lukas Felsberger, Andrea Apollonio, Jan Uythoven (CERN)

Ulrich Dorda, Dirk Vandeplassche, Jarno van de Walle (SCK CEN)



This project has received funding from the Euratom Research and Training Programme 2019-2020 under grant agreement No 945077.

DOCUMENT CONTROL SHEET

DOCUMENT INFORMATION

Document title	Accelerator reliability studies with emphasis on the fault tolerance capabilities
Authors, (organisation)	Lukas Felsberger, Andrea Apollonio, Jan Uythoven (CERN), Ulrich Dorda, Dirk Vandeplasse, Jarno van de Walle (SCK CEN)
Document type	Deliverable
Document ID	D10.1
Work package n°	WP10 (Domain: ADS System safety)
Work package title	WP10 Accelerator and beam line safety and reliability
Lead beneficiary	CERN
Dissemination level	Public
Date of issue	12/12/2022
Archive ID reference COO	SCK CEN/52402567

DOCUMENT SUMMARY

The MYRRHA Accelerator-Driven-System (ADS) poses very stringent reliability requirements on the accelerator. In particular the requirement of no more than ten beam interruptions longer than three seconds in a period of ninety days is beyond the capabilities of existing, comparable accelerators. To study the feasibility of an accelerator meeting these requirements several reliability studies had been carried out previously. These were initial investigations of possible fault-tolerance strategies by Burgazzi and Pierini, a reliability model of MYRRHA based on reliability data from SNS in the FP7-MAX project, and a further development of this model using experimental reliability data from CERN LINAC4 in the H2020-MYRTE project.

In the PATRICIA 10.1 project a more complete and refined reliability model of the MYRRHA accelerator was developed. Emphasis was put on a detailed understanding of fault-tolerance capabilities, which have previously only been studied conceptually, and missing key input data and associated future studies have been identified.

The objectives have been achieved by an iterative methodology. A reliability analysis of existing accelerators allowed to derive a top-down requirements model, which was used to derive reliability requirements for all sub-systems. These served as reference for detailed simulation studies of the sub-systems to optimize the fault-tolerance schemes as well as to highlight (potentially missing) key input data and methods or experiments to obtain them.

The reliability analysis of existing, comparable accelerators (ORNL SNS, CERN LINAC2/4) revealed that the foreseen fault-tolerance capabilities can cover the biggest sources of failures. However, it also showed that this might not be sufficient to reach the availability goals because about 99% of faults have to be avoided in comparison to existing, comparable accelerators. Hence additional fault-tolerance capabilities and other means of reliability improvement, such as conservative and robust design, careful supplier selection, optimized maintenance, etc., have to be considered. The analysis of existing accelerators also revealed that there is little data and understanding on root-causes of beam trips of a few seconds since they are less relevant for most applications.

Detailed simulation studies allowed a more precise quantification of the fault-tolerance schemes. These are (1) modular-redundant RF power amplifiers, (2) redundant low-energy injectors, and (3)

serial redundancy of RF cavities in the mid- to high-energy section of the accelerator. Results showed that for all the fault-tolerance schemes to be effective they need to have a highly-reliable switching mechanism to go from a reference configuration to a fault-compensating configuration, allow to repair faults as quickly as possible and ideally during beam production, and minimize singular points of failure and common-mode failures.

(1) Modular-redundant RF power amplifiers promise the biggest reliability improvement over existing accelerators at moderate additional cost. The reliability modelling and analysis of the CERN-SPS 200 MHz SS amplifier, 1.6 MW, built by Thales-Gérac revealed that RF amplifiers in line with the MYRRHA reliability goals already exist, even when repairs can only be carried out when no beam is being produced.

(2) Hot-spare redundant injectors imply a relatively costly duplication of the low-energy accelerating section. Their effectiveness strongly depend on fast repairs of the faulty injector and of the reproducibility of the produced beam when switching from the failing to the hot-spare injector. The former seems feasible when compared to time-to-repair data from existing accelerators and because repair is possible when the other injector delivers beam. The latter has yet to be experimentally validated with the MINERVA set-up before a concluding statement on the effectiveness of redundant injector's fault-tolerance capabilities can be made.

(3) Serial redundancy of RF cavities in the mid- to high-energy section of the accelerator requires a significant RF power margin (according to latest beam-dynamics models around 45% across the entire main LINAC and 50-70% in its MINERVA sub-section up to 100 MeV) and individually powered RF cavities. Both considerably increase the cost. The effectiveness of the measure strongly depends on the number of cavities used for compensation and the reliability and speed of switching to a new accelerator configuration after a fault. As most repairs can only be carried out when no beam is being produced, the time to repair a fault is governed by the time between failures of MYRRHA. This leads to a lower effectiveness of the fault-tolerance scheme than estimated in previous studies, which assumed immediate repair (during beam production) for the majority of faults. The switching to a new accelerator configuration has theoretically been shown to be feasible in the required time, but its reliability can only be practically determined with the 100 MeV MINERVA set-up. This will allow a concluding estimate of the effectiveness of the serial redundancy scheme.

With regards to future steps, failure data from MINERVA will be essential to assess how MYRRHA reliability goals can be met, especially data on very short beam trips of a few seconds. Automatic fault tracking solutions have been prototyped at CERN and at SCK-CEN to take care of tracking of short faults as they cannot be tracked manually and are recommended to be implemented in the experimental set-ups of MINERVA and possibly even before at comparable accelerator facilities to obtain a better understanding of root-causes of short beam trips.

An exhaustive list of the required additional key-inputs, associated studies and methods to arrive at conclusive estimations of the MYRRHA accelerator reliability is provided. The most important among these are:

- Reliability test campaigns and fault-tracking with the MINERVA set-up to estimate whether sub-system failure rates are in line with the provided requirements models.
- Beam-reproducibility and beam intensity ramp-up studies using the injector test-stand and/or the MINERVA set-up in view of the redundant injector set-up.
- Experimental testing of the switching to new superconducting cavity configurations in under three seconds for serial redundancy.
- Systematic identification, quantification and mitigation of singular points of failures across systems, such as RF tuning system blockades, controls problems, electrical network glitches, etc.

DOCUMENT HISTORY		
Version	Status	Date
v1	Final	01/12/2022

DOCUMENT APPROVAL		
The author, WP Leader and Coordinator acknowledge and accept delivery of the work completed for this deliverable.		
Date	Authors	Organisation
01/12/2022	Lukas Felsberger, Andrea Apollonio, Jan Uythoven Ulrich Dorda, Dirk Vandeplasseche, Jarno van de Walle	CERN SCK CEN
Date	WP Leader	Organisation
01/12/2022	Jan Uythoven	CERN
Date	Coordinator	Organisation
07/12/2022	Paul Schuurmans 	SCK CEN

DISTRIBUTION LIST		
Project Officer Renata Bachorczyk-Nagy	EC	Copy on PATRICIA SharePoint
PATRICIA Beneficiaries	PATRICIA Consortium	

Table of contents

Contents

1	Introduction.....	8
1.1	MYRRHA	8
1.2	Availability Requirements.....	8
1.3	Concepts and Definitions	8
1.4	MYRRHA Accelerator	10
1.4.1	Proposed Fault-Tolerance Methods.....	11
1.5	MINERVA test-set up	11
1.6	Previous Availability Studies.....	11
2	Research Objectives	14
2.1	Global Objectives (Beyond this Study)	14
2.2	Objectives of this Study	14
2.2.1	Derived Objectives	14
3	Methodology	15
4	Quantitative Availability Evaluation of Existing Comparable Accelerators.....	16
4.1	Top-Down Comparison.....	16
4.1.1	Long beam trips.....	17
4.1.2	Short beam trips.....	19
5	Top-Down Availability and Requirements Model	20
5.1	MYRRHA	20
5.1.1	Default model.....	20
5.1.2	Without fault-tolerance model	21
5.2	MINERVA	21
5.2.1	Default model.....	21
5.2.2	Without fault-tolerance model	21
6	Fault Tracking and its Automation	26
6.1	CERN Accelerator Fault Tracker.....	26
6.2	Automated Fault Tracking	26
6.2.1	LINAC4 Prototype	26
6.2.2	LLN Injector Teest-Stand Prototype	27
6.2.3	MINERVA/MYRRHA Outlook	27
7	Quantitative Availability (Simulation) Studies.....	27
7.1	AvailSim4	27
7.2	SPS SSPA	28
7.2.1	System Description and Study Objectives.....	29

7.2.2	Modelling Approach and Parameters	30
7.2.3	Simulation Results	33
7.2.4	Reliability Impact of Using Cavity-Combiners	34
7.2.5	Conclusions & Discussion	37
7.3	MYRRHA Injector Switchover Analysis	38
7.3.1	System Description.....	38
7.3.2	Modeling Approach and Parameters	40
7.3.3	Simulation Results	42
7.3.4	Conclusions & Discussion	44
7.4	MYRRHA Main LINAC Dynamic Compensation Analysis	45
7.4.1	System Description.....	45
7.4.2	Modeling Approach and Parameters	48
7.4.3	Results	52
7.4.4	Conclusions and Discussions	57
7.5	MYRRHA Main LINAC RF Powering analysis	59
7.5.1	System Description.....	59
7.5.2	Modeling Approach and Parameters	61
7.5.3	Results	63
7.5.4	Comparison with SPS SSPAs	65
7.5.5	Conclusions and Discussion.....	66
7.6	Quadrupole circuits	66
8	Implications for MYRRHA Availability.....	69
8.1	Recommendations based on Current Results	69
8.2	Recommended Studies based on MINERVA	70
8.2.1	Estimation of system failure rates.....	71
8.2.2	General estimation of fault-tolerance effectiveness, singular failures and other unforeseen events.....	73
8.2.3	Redundant Injectors	73
8.2.4	Main LINAC Dynamic Compensation.....	74
8.2.5	SSPAs	75
8.2.6	Glitch data	75
8.2.7	Quadrupole switching matrix	75
8.2.8	Synthesis.....	76
8.2.9	Extrapolating from MINERVA to MYRRHA	77
8.3	Knowledge-Transfer Activities.....	77
9	Summary, Conclusions & Outlook.....	79
10	Bibliography.....	81

11	Annexes	82
11.1	System definitions for top down models	82
11.2	Fault tracking tables for LLN Injector Set-Up	83
11.3	Asymmetric Injector Model.....	85
11.4	Glitch robustness.....	85

1 Introduction

This chapter introduces the MYRRHA accelerator, its availability requirements, reliability concepts and definitions used throughout this document, and previous related availability studies.

1.1 MYRRHA

MYRRHA (Multi-purpose hYbrid Research Reactor for High-tech Applications) is an Accelerator Driven System (ADS) consisting of a proton accelerator and a nuclear reactor core cooled by liquid lead-bismuth-eutectic (LBE), which also serves as spallation target. One of the main benefits of using an accelerator-driven nuclear reactor is that it allows so-called transmutation of nuclear waste, in which long-lived minor actinides are broken down in shorter lived waste.

The particle accelerator is required to drive the nuclear reaction with a continuous wave (cw) beam of 4 mA beam current and 600 MeV beam energy.

1.2 Availability Requirements

The particle beam should ideally be delivered continuously to the reactor as any interruptions, so called beam trips, of more than a few seconds lead to thermal stress of the reactor materials. The simulations show that if the beam is absent for longer than three seconds, the reactor needs to shut-down and enter a restart procedure called reactor turnaround (RT), which is estimated to take 12-24 hours. This has a considerable impact on the availability of the entire ADS. Hence, both duration and frequency of beam trips have to be limited. Moreover, the total number of thermal cycles of a nuclear reactor has an impact on its lifetime and must be reduced. The currently set requirements on beam trips are shown in Table 1. The requirement of no more than 10 trips longer than 3 seconds in a fuel cycle (FC) of 90 days poses the most stringent limit, which is currently not met by comparable accelerators and will be the requirement referred to throughout this document. Studies are being launched at the reactor side to see how the gradual ramp-up of the beam intensity following a beam stop could possibly relax the present limit of 3 seconds.

Table 1: MYRRHA accelerator requirements on beam trip frequency and duration.

# of beam trips > 3 sec	Max. 10 per 90 days
# of beam trips > 0.1 sec	Max. 100 per day
# of beam trips <0.1 sec	Unlimited

Previous reliability studies resulted in a range of measures to address the current beam trip requirements. These measures will be introduced after the necessary reliability concepts and definitions are discussed.

1.3 Concepts and Definitions

Reliability $R(t)$ is defined as a statistical concept expressing the probability p that the lifetime T of a system exceeds a certain time t :

$$R(t) = p(T > t).$$

This is especially useful to express the lifetime for systems, which are not repaired. A popular choice for this function is the exponential distribution,

$$R(t) = \exp\left(-\frac{t}{m}\right),$$

with m being the mean time to failure. The exponential distribution is characterized by a constant-in-time failure rate $\lambda = \frac{1}{m}$ (Footnote: the failure rate is defined as $\lambda = \frac{\frac{dR(t)}{dt}}{R(t)}$), which is characteristic for many systems in the mid of their lifetime before they enter an end-of-life period with an increased failure rate. The exponential distribution is being assumed for all systems in the following, unless explicitly defined otherwise.

Since particle accelerators and their sub-systems are repairable, the concept of the **availability** A will be used frequently in this document. It expresses the fraction of time a system is in operational state:

$$A = \frac{\text{uptime}}{\text{uptime} + \text{downtime}}.$$

To characterize the average time it takes to repair a system one uses the MTTR,

$$MTTR = \frac{\text{downtime}}{\text{number of faults}},$$

and to characterize the average time a system is in a working state between repairs one uses the MTBF,

$$MTBF = \frac{\text{uptime}}{\text{number of faults}}.$$

This allows to express the availability as:

$$A = \frac{MTBF}{MTBF + MTTR}.$$

Note that the failure rate of a system can be expressed as:

$$\lambda = \frac{\text{number of faults}}{\text{uptime}} = \frac{1}{MTBF},$$

And in the following it is generally assumed that the time between failures follows an exponential distribution,

$$P(T > t) = \exp\left(-\frac{t}{MTBF}\right).$$

The parameters of these statistical expressions are usually estimated from experimental failure or lifetime data of systems. Since these data are limited, the estimation of these parameters involves uncertainty. To address this, both a **most likely** and a **pessimistic estimate of the parameters** is being used whenever possible as a simple means of quantifying uncertainty and expressing results. This is best explained with an example:

Assume a particle accelerator has been run with a set of times between failure of {2, 6, 4, 8} days. From this the most likely and a pessimistic MTBF should be estimated.

For the most likely estimate the so-called maximum-likelihood (ML) method is used, which chooses the parameters of a statistical distribution so that the probability of the observed data is maximized. The ML estimate for the MTBF parameter of the exponential distribution is simply the mean of the observed data,

$$\widehat{MTBF} = \frac{2+6+4+8}{4} = \frac{\text{uptime}}{\text{number of failures}} = 5 \text{ days}.$$

For the pessimistic estimate, one determines a lower or upper bound to a parameter with a certain level of confidence, $1 - \alpha$, to express how certain one is that the actual value of the parameter is above

or below the respective bound. E.g. for the example above one can estimate with 90% confidence (i.e. $\alpha = 0.1$) that the actual MTBF value is above 2.5 days. Hence, in a conservative estimation one should use the pessimistic MTBF estimate of 2.5 days instead of the most likely estimate of 5 days.

There are different methods to estimate these confidence bounds. Throughout this study, the X^2 (chi-square) distribution is used to estimate the one-sided lower bound by:

$$\frac{2 \cdot \text{uptime}}{X^2(\alpha, 2 \cdot \text{number of failures} + 2)}.$$

For the above example, this can be evaluated as

$$\frac{2 \cdot (2 + 6 + 4 + 8)}{X^2(0.1, 2 \cdot 4 + 2)} = \frac{40}{X^2(0.1, 10)} = \frac{40}{15.987} \approx 2.5,$$

where the X^2 distribution is evaluated with probability of 0.1 and degrees of freedom of 10.

Another concept which will be referred to is **fault-tolerance**. This means that a system can continue to operate despite some of its sub-systems failing. To achieve fault-tolerance one usually needs:

1. Redundancy, so that a fault of one sub-system can be compensated by other sub-systems.
2. Fault isolation, so that a fault cannot propagate to other sub-systems.
3. Fault diagnostics and notification, so that a problem can be identified and is noticed.
4. (On-line) repair, so that a fault is repaired and a system can return into its original and most-resilient operational state.

By employing these principles, systems can theoretically continuously operate despite their sub-systems failing. In practice, this will also lead to an increase of system cost and complexity, which may introduce additional reliability challenges and requires careful implementation of these measures.

1.4 MYRRHA Accelerator

A schematic overview of the MYRRHA accelerator and reactor is shown in Figure 1. The particle beam is produced at the left end in the sources of one of the two redundant injectors housed inside the injector building. The particle beams are accelerated up to 6.6 MeV in the RF-Quadrupole (RFQ), and so-called CH-cavities. At this energy they are transferred from either one of the injectors into the so-called main LINAC. It consists of 163 superconducting RF cavities, housed inside cryostats, and accelerates the beam up to 600 MeV. The main LINAC is split into three accelerating sections. The first section with 62 single-spoke cavities in 32 cryostats, the second section with 18 double-spoke cavities in 9 cryostats and finally the third section with 74 elliptic cavities in 18 cryostats. Between the single and double spoke sections is a single rebuncher cavity. After every cryostat there is a quadrupole magnet doublet to focus the beam.

After the main LINAC the particle beam goes through a beam-delivery-system (BDS) into the MYRRHA reactor. It mainly consists of a series of bending and focusing magnets that direct the beam to the reactor, where a beam window connects the accelerator with the reactor. At the end of the injectors and after the main LINACs are beam dumps where the beam can be dumped in a safe manner.




Figure 1: Schematic overview of MYRRHA ADS.

1.4.1 Proposed Fault-Tolerance Methods

Fault-tolerance will be applied throughout the accelerator infrastructure. This study is focussing on three strategies in particular. These are (1) modular-redundant RF power amplifiers, (2) redundant injectors, and (3) serial redundancy of RF cavities in the main LINAC. They are treated in more detail in Sections 7.5, 7.3, and 7.4, respectively.

1.5 MINERVA test-set up

MYRRHA is being implemented in phases. One of these phases is called MINERVA, which will represent a 100 MeV proton accelerator and will later be extended to form the entire MYRRHA accelerator. MINERVA will only have a single injector line and the superconducting main LINAC will only consist of 62 cavities as shown in Figure 2 below.




Figure 2: Schematic overview of MYRRHA with the parts which will be realised in the MINERVA phase highlighted in the purple dashed line.

MINERVA will allow to carry out various experiments and tests, which are required to demonstrate the feasibility of MYRRHA. The reliability testing foreseen for MINERVA is treated in more detail in Sec. 8.2.

1.6 Previous Availability Studies

An investigation of the availability and reliability of a LINAC conceived for an ADS based on a high-level reliability block diagram approach was presented in (Burgazzi & Pierini, 2007). The paper investigated different LINAC architectures, including one without any fault-tolerance strategies, one with theoretically infinite fault-tolerance for the injectors and the main LINAC, and one with limited fault-tolerance on the injectors and the main LINAC. The latter can be considered closest to the actual MYRRHA implementation and is shown in Figure 3.




Figure 3: Reliability-block diagram of the model used by Burgazzi and Pierini to model the MYRRHA accelerator and its fault-tolerance schemes. (Burgazzi & Pierini, 2007)

As one can read from the block-diagram (from top to bottom), it consists of two injectors (each consisting of source, RFQ and normal-conducting drift tube LINAC) with ideal redundancy, a range of infrastructures and services (cryogenics plant, controls systems, cooling), the main LINAC with spoke and elliptic sections each split into in-tunnel elements with high MTBF repaired during reactor turnaround and out-of-tunnel elements with low MTBF repaired immediately, and finally a beam delivery system.

Notably, the model assumes that repairs on the main LINAC out-of-tunnel elements can be carried out while beam is being produced, which will likely not be possible in practice. Moreover, switching from a nominal to a fault-compensating configuration is assumed to work every time, which is optimistic. Neither cryostats nor quadrupoles failures are modeled. Lastly, the model assumes that any distribution of faulty cavities in the main LINAC can be compensated as long as the number of faulty cavities is below a certain threshold. This might not be possible when faulty cavities are next to each other. With these optimistic assumptions it is estimated that a reduction of failures by a factor of 17.5 (4 failures per fuel cycle instead of 90) is possible by the proposed fault-tolerance methods. The authors point out that their study should be seen as “pioneering effort” to identify paths towards a possible accelerator for ADS and mention that their modeling suffers from lack of accurate input data.

In the FP7-MAX4.4 project (Pitigoi & Ramos, 2014), a more detailed model of the MYRRHA was developed. The lack of accurate failure data was addressed as far as possible by using reliability data from the operation of the Spallation-Neutron-Source (SNS) at the Oak Ridge National Laboratory

(ORNL), whose accelerator is comparable to the MYRRHA accelerator. The model improves upon the previous study by having a certain success-rate of switching from a nominal to a fault-compensating configuration, both cryostats and quadrupoles have been added, and singular points of failure of the fault-tolerance schemes are considered. However, it still assumes immediate access to repair for most systems. Nevertheless, with these more realistic assumptions it is estimated in an optimized model that a reduction of failures by a factor of 5 (6.4 failures per fuel cycle instead of 32.5) is possible by the proposed fault-tolerance methods. The proposals for the optimized model were to improve reliability of badly performing SNS components, extend coverage of fault-tolerance strategies, and apply scrutiny when avoiding singular points of failure. It was also concluded that even with the fault-tolerance strategies the reliability goals can only be achieved when all subs-systems are having a high reliability.

Finally, during the H2020-MYRTE D2.9 project (Apollonio, JM, Pitigoi, Fernandez Ramos, & Uythoven, 2015), the accelerator availability model was further refined by augmenting it with experimental input from CERN's Linac4, which is also comparable to the MYRRHA accelerator even though it lacks the superconducting cavities. The focus on this project was on a model for MINERVA and hence the modeling of the fault-tolerance strategies was not further studied as this is not part of MINERVA, except for providing a test set-up for them. This report also emphasised the importance of fault tracking and registration of short faults.

2 Research Objectives

The PATRICIA D10.1 study aims to build upon the previous reliability studies with specific objectives, which are outlined below. All studies are embedded in the MYRRHA reliability studies with more global objectives, which are detailed in the following.

2.1 Global Objectives (Beyond this Study)

The main objectives of the MYRRHA reliability studies are to identify (1) if and to what extent the MYRRHA reliability goals can be met and (2) to identify the most cost-effective set of measures to meet the reliability goals.

2.2 Objectives of this Study

The objectives of this study are a direct consequence of the global objectives and the achievements in previous studies outlined above. A more complete and refined reliability model of the MYRRHA accelerator shall be developed with an emphasis on a detailed understanding of fault-tolerance capabilities, which have previously only been studied conceptually.

Concretely, the following should be covered:

- Global accelerator models for MINERVA and MYRRHA.
- Detailed availability models for three proposed fault-tolerance strategies: modular-redundant RF power amplifiers, parallel redundant injectors, and serial redundancy of RF cavities in the main LINAC.

Each of these studies contributes to address the global objectives by providing methods to estimate the feasibility of meeting the MYRRHA reliability goals.

2.2.1 Derived Objectives

To address the outlined goals of this study, additional objectives were identified:

- Development of tools and features for the advanced availability simulations.
- Hand-over of the availability simulation models to SCK so that they can be continuously updated when new data arrives.
- Teaching of SCK colleagues on usage and development of the availability simulation tools and models.
- An updated high-level availability comparison with existing accelerators to derive top-down fault budget models for MYRRHA and MINERVA
- Investigation of automated fault tracking methods to keep track of short beam-trips relevant for MYRRHA operation

3 Methodology

To address the outlined objectives, a multi-stage process had been chosen as methodology, which is described in the following. The remaining document follows the order of this multi-stage process, hence the description of it also serves as overview of the coming Sections.

A first estimate of the expected failure rates of the different accelerator sub-systems and the biggest expected reliability challenges was obtained by performing a quantitative availability evaluation of existing, comparable accelerators, for which reliability data was available (Sec. 4). It was immediately noticed that failure data is often not available for faults of few seconds in other accelerator facilities since they contribute little to downtime and are very cumbersome to collect manually. This called for investigating automatic fault tracking solutions, which were prototyped and tested (Sec. 6).

Based on the lessons learned from this initial data collection and analysis process, top-down availability and requirements models for MYRRHA and MINERVA in different fault-compensation configurations have been developed (Sec. 5). They reflect the estimated reliability of the MYRRHA and MINERVA sub-systems and can be used to derive reliability requirements for sub-systems to be in line with the MYRRHA availability goals.

The developed top-down requirement models also serve as reference input to the more detailed availability models of the fault-tolerance methods (Sec. 6), which address the core objectives of this project. This modeling is based on the AvailSim4 tool (Sec. 7.1) for particle accelerator availability simulations, which has been developed at CERN and is now publicly available¹. An availability model has been developed for each of the three introduced fault-tolerance strategies (Secs. 7.2 to 7.5). For each modeling, the core goal was to identify the key parameters to optimize the fault-tolerance scheme and those parameters which still need to be estimated to allow a conclusive estimate of the effectiveness of the schemes and subsequently the feasibility of meeting the MYRRHA availability goals.

Based on the results of these investigations, a recommended course of action in view of the MYRRHA availability goals was derived (Sec 0). This is structured in recommendations based on the current state of knowledge, short-term recommendations on studies and estimations to be carried out to improve the current state of knowledge and give more confident recommendation, and longer-term recommendations and a detailed discussion how MINERVA is most effectively used to address with confidence how the MYRRHA reliability goals can most effectively be reached.

¹ <https://pypi.org/project/availsim4/>

4 Quantitative Availability Evaluation of Existing Comparable Accelerators

To obtain an updated and more complete overview of reliability challenges in existing, comparable accelerators, a quantitative analysis of existing availability data was carried out in a top-down manner.

4.1 Top-Down Comparison

For the Top-Down availability analysis three accelerators were selected on the basis of similarity to the MYRRHA accelerator and accessibility to reliability data. A summary of their (main) characteristics is shown in Table 2.

Table 2: Summary of relevant characteristics of the MYRRHA accelerator in comparison to existing, comparable accelerators, whose reliability data is used to obtain a first estimate of the expected reliability challenges for the MYRRHA accelerator.

Parameter	MYRRHA	ORNL SNS (Oak Ridge National Laboratory Spallation Neutron Source)	CERN LINAC2	CERN LINAC4 (H-)
Accelerated particle	p	p	p	H-
Beam energy	600 MeV	1 GeV	50MeV	160MeV
Beam current	4 mA avg.	38 mA peak, 1.4 mA avg.	180mA peak (150mA before source upgrade)	65-80 mA peak (25 mA operational), 0.032 mA avg.
Beam pulse duration	continuous wave (cw)	1 ms	250 μ s	400 μ s
Commissioning date	na	2006	1978	2016
Superconducting RF cavities	yes	yes	no	no
Availability data source	na	2011/2012 (MAX4.4 based on SNS logbook data)	2018 (CERN accelerator fault tracker)	2021 (CERN accelerator fault tracker)
Use	operation as spallation source (protons on spallation target)	operation as spallation source (protons on spallation target)	operation as proton injector	operation as proton injector
Architectural fault-tolerance	Yes	No	No	No

As one can see, SNS is closest to MYRRHA in beam energy and the only comparable accelerator in this list which is using superconducting RF cavities as well. All of the listed comparable accelerators use pulsed beams, whereas MYRRHA uses so-called continuous wave beams producing an (almost) continuous particle beam. For the interpretation of the reliability data it is also important to note that MYRRHA will be operated far from particle accelerator performance limits, whereas other machines are operated closer to their limits with the goal of setting new records in terms of energy, intensity or similar. Operation close to performance limits usually triggers additional failure modes, which are not seen when operating with sufficient margins. These could not be filtered out from the data.

In the following, the analysis of the available data will be split in data on long and short beam trips, with the former being longer than two minutes and the latter shorter or equal to two minutes. This separation is done since the data for longer faults is of significantly better quality, since longer faults usually require a human intervention to be repaired, contribute more to downtime, and are better documented. Short beam trips are often self-correcting problems or remotely resettable and are often not followed-up in detail. Moreover, without automated fault tracking solutions in place, short beam trips may even go unnoticed unless they significantly impact a user of the beam.

4.1.1 Long beam trips

With regards to long beam trips, the data availability for SNS, LINAC2 and LINAC4 are very good. Figure 4 shows the distribution of faults by system type for the indicated operational periods. The blue bars show the failure count and the orange line the cumulative percentage (summing from the left). The absolute fault count numbers for LINAC2 are significantly lower. This is probably due to the smaller size of the accelerator and because when the data was recorded it had been a mature and optimized machine (in fact 2018 was the last year of operation before de-commissioning). The data for LINAC4 and SNS were taken during the very first years of operation, where many systems still required optimization. Due to this different operational situation it is best to focus on the relative frequency of faults by system type.

For all accelerators, the RF systems are failing most frequently. It must be noted, that all three accelerators make use of tube-based RF power amplifier technology, which contributes to the majority of RF faults. At MYRRHA this shall be addressed by using fault tolerant solid state power amplifiers. The SNS data has a separate category for faults of the RF cavities and cryogenics, whereas for LINAC2/4 the cavities are included in the RF category and cryogenics do not exist because of the lack of superconducting elements.

The second most problematic system type across the three accelerators are electrical systems (SNS) and power converters (LINAC2/4). Here, LINAC2/4 have a separate category for electrical network problems (usually thunderstorm related glitches), whereas it appears to be included in the electrical systems category for SNS. Electrical systems are abundant in particle accelerators and they often contribute a significant amount of failures, which need attention. ‘External’ faults, such as those triggered by electrical network glitches, can be avoided by using sufficient capacitive elements to make electrical systems robust to such disturbances. ‘Internal’ faults, e.g. from electronic component wear, need to be addressed by employing redundancy, de-rating, careful supplier and component selection, reliability testing, or adequate maintenance schemes.

The other fault categories vary between accelerators. Generally, these are controls and beam instrumentation (i.e. diagnostics) systems, vacuum systems, and cryogenics plants among others.

A best-effort analysis of faults longer than 20 minutes at LINAC4 by root-cause confirms the findings based on the data by system type and adds further detail. Indeed, most faults were due to the tube-based RF amplifier technology. It also points towards a significant fraction of faults being due to controls- and software problems and not just because of (hardware) component failures. Lastly, more hardware failures were affecting mechanical systems than electronic systems. A prominent example are air fans, which are redundant but due to the lack of monitoring and timely replacement eventually lead to operational problems.




Figure 4: Fault count by system for existing, comparable accelerators. SNS data from Pitigoi et al. Reliability model of SNS linac No. NEA-NSC-DOC--2015-7. 2015. LINAC2/4 (L2/4) data from the CERN Accelerator Fault Tracker database.

4.1.2 Short beam trips

Short beam trip data is usually less detailed since short trips of a few seconds contribute less to downtime for most accelerator facilities. The data was available for SNS from the MAX 4.4 report (Pitigoi & Ramos, 2014). The data for CERN LINAC4 was recorded by implementing an automated fault tracking prototype (see Sec. 6 for details).

Figure 5 shows the requirement on short beam trips (longer than three seconds) for MYRRHA next to observed beam trip counts per 90 days at SNS and LINAC4 for beam trip durations of three and six seconds. The observed trip frequency is at least two orders of magnitude above the MYRRHA requirements. This means that at least 99% of faults of existing, comparable accelerators need to be prevented in MYRRHA, which may pose a significant challenge.




Figure 5: Number of short beam trips per 90 days for existing, comparable accelerators in comparison to the MYRRHA requirement. Data sources: SNS data: F. Bouly, Proton Linacs as ADS drivers - EUCARD2 Workshop : Status of ADS R&D, CERN. LINAC4 data: analysis based on logged beam current data; prototype explained in Appendix 6.

The origin of beam trips of a few seconds is not well understood. They are different from hardware problems requiring a repair or an intervention. Experience from accelerator operation at CERN suggests that they increase when accelerators are pushed to their performance limits, which should not be the case for MYRRHA. Moreover, experience suggests that they are frequently related to settings issues and further operational optimization could reduce them. However, often the root-causes are not analyzed and addressed, as these short faults are less significant for availability. E.g., at LINAC4 despite these many ‘faults’ per 90 days, the availability was 97% (logged in accelerator fault tracker – AFT, Oct. 21).

Since many of these short faults often do not receive a lot of attention at accelerator facilities, it is hard to judge how difficult it would be to prevent them. Hence, tracking and diagnostics of such short trips is vital in order to address them and should be an integral part of the MINERVA and MYRRHA controls system. Possibly short beam trips should even be studied in existing, comparable accelerators to increase the data and understanding of them. The importance of this is reflected in the fact that a fault-tracking prototype has been developed as part of this project for both LINAC4 and the Louvain-La-Neuve injector test stand, which is described in more detail in Section 6.

5 Top-Down Availability and Requirements Model

Based on the fault count distribution by system of comparable existing accelerators (Sec. 4), top-down models for MYRRHA and MINERVA have been derived in a best-effort estimation by scaling the fault distribution by system (Figure 4) to the MYRRHA requirements and distributing it across its sub-systems. The different proposed fault-tolerance methods, technological choices, and differences of MYRRHA with respect to the other comparable accelerators have been taken into account in several iterative reviews with the experts participating in the work package.

The top-down models generally show the most likely system fault count configuration based on current knowledge, which of course is just one of many possible configurations in line with reaching the MYRRHA availability requirements. That means that if all systems are designed and built to meet the requirements outlined in the top-down model, the MYRRHA reliability goals will be reached. However, since this is not always possible, this also means that once the reliability of a sub-system is known to be better or worse than assumed in the model, the system affected and also the other sub-system requirements can be updated accordingly.

5.1 MYRRHA

5.1.1 Default model

The default model represents the MYRRHA accelerator with all foreseen fault-tolerance strategies. Figure 6 shows the MYRRHA requirements model as a tree of all accelerator systems and sub-systems. At the root is the MYRRHA accelerator with an allowed fault count (i.e. count of beam trips longer than three seconds) of 10 per 90 days, denoted by '(10)'. E.g., the switching magnet of the injectors with (0.05) should have an MTBF of $90/0.05 = 1800$ days or nearly 5 years.

Note that this represents the limit of tolerable faults to be in line with the MYRRHA requirements. It needs to be considered as an upper bound and requirements should be derived with a sufficient margin to ensure conformity with MYRRHA reliability requirements. This applies to all listed fault counts.

The first level lists boxes for the injectors, the main LINAC, the beam delivery system (BDS) to the reactor, general infrastructures and services, and for possible operational and maintenance related faults. Short beam trips of unknown origin are not included as a separate category because it is believed that with the necessary diagnostics they can be addressed and assigned to the respective sub-system that causes them.

When taking a look at the injectors, they are composed of a box representing the switching magnet, which deflects the beam from the active injector to the main LINAC, and a box called Injector Switching representing the fault-tolerance method, which connects to the two injectors. The injector switching box has a fault count of '(0.95)' and a fault-tolerance effectiveness rating of '<91%>', whereas each of the injectors has a fault count of '(5.5)'. This should indicate that of the 11 faults occurring in total in the two injectors, 91% are compensated by the fault-tolerance scheme and the remaining 9% contribute the 0.95 faults, which cannot be compensated. Expressed as formula, it is

$$\begin{aligned} & \text{number of faults impacting beam} \\ &= \text{number of faults of systems covered by fault tolerance} \\ & \cdot (1 - \text{fault tolerance effectiveness}). \end{aligned}$$

For the injector redundancy this evaluates to

$$0.95 = (5.5 + 5.5) * (1 - 0.9136).$$

The effectiveness of '91%' is based on a best effort estimation using an availability simulation and expert reviews, which is described in more detail in Sec. 7.3.

Each of the injectors has a MTTR of '{8}' hours, which is required to achieve the estimated effectiveness of the fault-tolerance scheme. This does not mean that all of the injector sub-systems must have an MTTR of 8 hours but that the averaged MTTR over all failures in the injector should not exceed 8 hours.

Moving on to the Main LINAC, one can see that for the Dynamic Compensation box at the very left the same notation as for the injectors applies. This time, instead of covering two parallel-redundant injectors, both SRF (superconducting radio frequency) systems and SRF cryo-modules are covered, except for the re-buncher cavity at 100 MeV, which is in a separate box to the right. For the SRF and SRF cryo-modules, the system count is stated in square brackets; e.g. there are 58 cryostats housing 152 RF cavities.

Next to it is the fault-tolerance concept for the quadrupole circuits, denoted as 'Quad – Matrix/Compensation'. This fault-tolerance is not treated in detail in this project, as it still needs to be defined in more detail. It is shortly discussed in the Appendix 7.6.

The remaining boxes do not contain explicitly represented fault-tolerances and should be straightforward to read. A full list of the systems with system boundary definitions is given in the Appendix 11.1.

5.1.2 Without fault-tolerance model

Figure 7 shows the entire MYRRHA accelerator without any of the foreseen fault-tolerances. It shows the system reliability requirements that would have to be met to be in line with the reliability goals without employing fault-tolerance strategies. The numbers were derived by keeping the same distribution of faults across sub-systems as in the reference model with fault-tolerance but by scaling them to fit the requirements without fault-tolerance. This can only be achieved when the systems are exceptionally reliable. For example, the RF powering faults with an assumed failure rate of 1 failure per 90 days over 153 systems would require an individual RF amplifier to have a failure rate of $3\text{e-}6$ failures per hour, which is extremely low for a relatively complex system.

5.2 MINERVA

5.2.1 Default model

Figure 8 shows the requirements for MINERVA to be in line with the requirements for the MYRRHA reference model. The numbers are derived by scaling the fault counts of the MYRRHA reference model by the reduced system counts for MINERVA and, where applicable, corrections by system complexity were made; e.g. for the controls system. The RF powering systems are an example for the scaling by system count. MYRRHA has 152 systems and 2.5 failures and MINERVA has 60 systems and 1 failure, which is $2.5 \times 60 / 152$.

5.2.2 Without fault-tolerance model

The model in Figure 9 shows the requirements for MINERVA to be in line with the MYRRHA requirements model of 10 faults over 90 days without any fault-tolerances. Again, this was obtained by simply scaling the fault counts of the MYRRHA model without fault-tolerances by the reduced system counts for MINERVA and, where applicable, minor corrections by system complexity were made; e.g. for the controls system. The RF powering systems are again an example for the scaling by system count. MYRRHA has 152 systems and 1 failures and MINERVA has 60 systems and 0.4 failure, which is $1 \times 60 / 152$.




Figure 6: MYRRHA Top-Down model with fault-tolerance strategies. Numbers in brackets are fault count per 90 days (worst case to reach requirements). Numbers in square brackets are the system count. Numbers in arrow-brackets indicate the effectiveness of the respective fault-tolerance mechanism (i.e. what percentage of faults in the systems below the fault-tolerance method will be visible to operation).




Figure 7: MYRRHA Top-Down model without fault-tolerance strategies. Numbers in brackets are fault count per 90 days (worst case to reach requirements). Numbers in square brackets are the system count.




Figure 8: MINERVA Top-Down model in line with MYRRHA requirements when MYRRHA is implementing fault-tolerance strategies. Numbers in brackets are fault count per 90 days (worst case to reach requirements). Numbers in square brackets are the system count. Numbers in arrow-brackets indicate the effectiveness of the respective fault-tolerance mechanism (i.e. what percentage of faults in the systems below the fault-tolerance method will be visible to operation).




Figure 9: MINERVA Top-Down model in line with MYRRHA requirements when MYRRHA is not implementing fault-tolerance strategies. Numbers in brackets are fault count per 90 days (worst case to reach requirements). Numbers in square brackets are the system count. Numbers in arrow-brackets indicate the effectiveness of the respective fault-tolerance mechanism (i.e. what percentage of faults in the systems below the fault-tolerance method will be visible to operation).

6 Fault Tracking and its Automation

As a prerequisite to reach the MINERVA and MYRRHA reliability requirements one needs to be able to both measure the frequency and understand the origin of faults. Only then will it be possible to avoid these failures. This Chapter will introduce existing fault tracking solutions at CERN as an example of establishing a common fault database in an organization and then explore possibilities of automatizing fault tracking, as will be necessary for tracking of short faults of seconds or even less.

6.1 CERN Accelerator Fault Tracker

At CERN the Accelerator Fault Tracker (Roderick, Burdzanowski, Anido, Pade, & Wilk, 2017) is an established tool to track, visualize, analyze and store faults in a central place for the accelerator chain.

The registration of faults is done by the accelerator operators. They classify the faults according to which system has caused the downtime. For each class of systems, a group of engineers and specialists is requested to review and approve the fault classification and details. This ensures that the relevant experts are aware of reliability issues and that the data is reviewed and agreed upon.

The fault tracking tool has been in use for LHC since 2015 and for the injector chain since 2017. It is interfaced with a web-GUI, the operations-web-logbooks, the accelerator logging systems, and a database. Since 2020, it is also interfaced with automatic fault tracking clients, which are gradually rolled out and improved across the different accelerators. Their main aim is to reduce the fault-recording workload for operators, especially for short and frequent faults.

6.2 Automated Fault Tracking

Very short or very frequent faults cannot be logged manually. Since they are of particular importance for MYRRHA, it is necessary to implement an automated system to track such beam trips. In the following sections, two prototypes for such fault tracking are discussed and their lessons learned summarized.

6.2.1 LINAC4 Prototype

The CERN LINAC4 accelerator has a beam interlock systems (BIS) that monitors all accelerator equipments that may lead to damage of the machine when they fail. This covers a significant fraction of accelerator systems. As soon as an equipment enters an unsafe state, the BIS will be notified and stops the beam. Since the data of this is recorded back in time, it represents a good starting point to automatically identify short beam trips and to identify the equipment that has triggered it.

To do so, a prototype was implemented in Python. It successfully managed to reconstruct beam trips and the triggering equipment. Since certain beam outages are not due to BIS-connected equipment the beam current and the beam settings were monitored in addition for a more complete picture.

The prototype allows to generate:

- Fault downtime statistics – see e.g. Figure 5
- Fault timelines for in-depth analysis of events and identification of root causes.

The lessons learned were the following:

- The monitoring of beam availability and beam trips is relatively straightforward.
- Identifying root causes automatically is challenging. However, an attribution of responsible systems is possible and detailed timelines of system behavior help system experts in determining root-causes of problems.

- A BIS provides the infrastructure that ensures highly consistent data on accelerator equipment failures. The fact that it is independent from the controls system ensures that in the case of controls problems the data is still not corrupted. However, depending on the implementation of the BIS not all operationally relevant beam trips may be recorded.

6.2.2 LLN Injector Test-Stand Prototype

Following the experience with the LINAC4 fault tracking prototype, a similar approach was tested for the LLN injector test-stand. The controls system was chosen as the infrastructure providing the relevant data. A hybrid approach was chosen in which data can be manually entered by the operators via a graphical user interface or automatically by the controls system. The system uses the fault-tracking tables from the Appendix 11.2 to generate templates for entering the data automatically.

The prototype could not be tested in a realistic scenario. Hence, a conclusion on how fault tracking via the controls system can be compared to using the BIS could not be made.

6.2.3 MINERVA/MYRRHA Outlook

For MINERVA and MYRRHA an automatic fault tracking system will be essential to keep track of short beam trips. The LINAC4 prototype has indicated that a combination of monitoring the equipment via the BIS and the beam production via beam currents can provide a consistent and relatively complete data-set on short beam trips. For MINERVA it is foreseen to embed the fault tracking in the controls systems. It should be taken care that the implementation produces a consistent data-set which provides a complete overview of equipment misbehaviour. Moreover, it is recommended to adopt the approach of the fault tracking at CERN where system experts are automatically informed when a fault is assigned to their system and requested to review and follow-up the incident. Lastly, it will be important that the data collected by the fault tracking systems will be used to compare the actual reliability of the systems with the top-down availability requirements models. For this the classification of the failure data should follow the structure of the top-down models.

7 Quantitative Availability (Simulation) Studies

The top-down models of the previous section allow a compact overview of the expected fault contribution of all accelerator sub-systems. However, to quantify the effectiveness of the fault-tolerance schemes and identify their key dependencies, more detailed models are required. These often cannot be evaluated analytically and are implemented as Monte Carlo simulations in AvailSim4, which is introduced in Sec. 7.1. After that, Sec. 7.2 details the modeling of a recent RF power amplifier installation at CERN, which is similar to the foreseen MYRRHA RF amplifiers, and was used as a reference case to validate the modeling. The validated modeling approach was then applied to the MYRRHA injectors in Sec. 7.3, the Main LINAC dynamic compensation in Sec. 7.4, and to the main LINAC RF amplifiers in Sec. 7.5. Each of these subsections will include a short introduction of the system under study, the modeling approach and parameters, the simulation results and their interpretation. Lastly, a short discussion of a fault-tolerance scheme for the quadrupole magnets is introduced and discussed in Sec. 7.6. Since this investigation was less detailed, it does not follow the usual structure of presentation.

7.1 AvailSim4

AvailSim4 is a Monte-Carlo based availability simulation software specifically made for the particle accelerator domain. Development has been taken over by CERN in 2020. During this time a range of MYRRHA specific features developed, which are e.g. needed to model the complex logic for the dynamic compensation of RF circuits in the main LINAC. It includes numerous features to model complex system behaviour including different phases of operation with associated failure modes, blind

failures that can only be detected by dedicated inspections, different repair policies depending if systems can be accessed during beam, etc.

The software has been made open source and is accessible on the pypi package repository with a link to the source code and a detailed user guide.²

7.2 SPS SSPA

As part of the LHC Injector Upgrade program at CERN a complete overhaul of the Super Proton Synchrotron Radio-Frequency (RF) system took place. New cavities have been installed for which the solid-state technology was chosen to deliver a combined RF power of 2 MW from 2560 RF amplifiers.

This project has been studied as part of the PATRICIA project to benefit from the experience of implementing solid-state technology at CERN since the same type of amplifiers are foreseen to be implemented at MINERVA and MYRRHA. The findings have been published in (Felsberger, et al., 2021).

This strategy of choosing solid-state amplifiers promises high availability as the system continues operation when some of the amplifiers fail. A similar approach is chosen for the RF powering of the MYRRHA cavities. Solid-state amplifiers have crucial advantages over traditional tube-based amplifiers, which include:

- No need for HV
- Easy operation and maintenance, because modules are small and light in comparison to tube-based solutions which may require a crane for handling
- Modular redundancy: no interruption if some modules fail

Figure 10 gives an overview of existing RF power amplifier technologies and their typical power ranges. It shows that SSPAs have significantly lower power so they require power combiners.




Figure 10: Overview of existing RF power amplifier technologies and their typical power ranges. (Jacob, Mercier, Langlois, & Gautier, 2011). The MINERVA SRF power range was added at a frequency of 352.2 MHz.

² <https://pypi.org/project/availsim4/>

The study on the SPS SSPAs quantifies the operational availability that can be achieved with this new type of installation. A model based on lifetime estimations of the RF modules was verified using data from early operational experience. This led to a validated real-world example as well as a validated modeling methodology, which was used for the assessment and modeling of the MYRRHA RF amplifiers and the fault-tolerance models in general.

7.2.1 System Description and Study Objectives

Two new RF powering systems with a peak power of 2 MW each have been installed in the CERN Super Proton Synchrotron (SPS). The 200 MHz solid-state RF powering system, illustrated in Figure 11 has no RF circulators and uses a mix of cavity-combiners and 3 dB-combiners to provide the power from 2560 solid-state amplifiers, distributed over 2x16 towers that are shown in in Figure 12, to the cavities.

At the tower level, the outputs of up to 80 RF modules, each containing a power supply and two solid-state amplifiers, are combined in an 80:1 cavity-combiner. The outputs of the 16 towers are connected to a four stage 3 dB-combiner and transmitted via a coaxial feeder-line to the cavity.

Only 72 of the up-to 80 RF modules are required for SPS operation. This power margin ensures high operational availability of the system. Similar redundant systems using solid-state technology have been implemented at SOLEIL (Marchand, Lopes, Ribeiro, & Ruan, 2011) and ESRF (Jacob, Mercier, Langlois, & Gautier, 2011). Exceptionally high availability of the systems was reported (Marchand, et al., 2013), (Jacob, Farvacque, Mercier, Langlois, & Gautier, 2013).




Figure 11: Architectural overview of the SPS RF powering.



Figure 12: The 2 x 16 towers of the new SPS RF powering installation.

7.2.2 Modelling Approach and Parameters

The goal of the modelling was to quantify the operational availability of the RF powering system for the SPS and use it to identify the best operational strategy considering current and future operation.

A quantitative reliability model was generated based on expert interviews, engineering documentation, testing data, operational data of comparable systems, and manufacturer data. In line with the Design Review Based on Failure Mode method (Schmidt, Riedel, & Kangas, 2011), a detailed system analysis was carried out with particular attention for novel parts of the system. The resulting reliability model was combined with an operational model and potential future operational scenarios and different configurations of the system were simulated in AvailSim4. The results of these simulations serve as evidence for decision making of the actual operational scenario and the modelling can be updated whenever new operational data are available.

Table 3 shows the obtained reliability model parameters in terms of estimated lower bound of MTBF at 90% confidence (pessimistic estimate), MTBF point estimate (most likely value), the type of MTBF distribution, MTTR point estimate, the type of MTTR distribution, and the fault logic for each sub-system.

Table 3: Estimated input parameters for the simulation.

Sub-System	MTBF 90% [h]	MTBF (point estimate) [h]	MTBF Distribution	MTTR [h]	MTTR Distribution	Fault Logic
RF module	[173360, 346720, 693440]		Exp.	1	Fixed	16 x (7200 n)
Driver	158114	420768	Exp.	1	Fixed	16 x
Controls	31500	72300	Exp.	2	Exp.	1 x
Cooling, Combiners, Loads, Feeder	62000	144000	Exp.	24	Exp.	1 x
SPS shadow events	243	173	Exp.	5	Fixed	Providing repair opportunities
Maintenance Stop	1680	1680	Fixed	24	Fixed	

The overall MTBF of an RF module was estimated by the manufacturer as 346`720 hours (Charlot, 2018). To reflect uncertainties of this estimate, potential future aging problems, and inherent pessimistic assumptions of the reliability estimate, a sensitivity analysis on this parameter was carried out (by using [1773360, 246720, 693340] hours for both the lower bound and point estimate model). A second sensitivity analysis was carried out by varying the number of installed RF modules from 72 (no redundancy) to 75 (up to three faulty modules tolerated).

The MTBF of the Driver SSPAs was identified in a reliability test of 64 units over 18 months. The Controls, Cooling, Combiners, Loads, and Feeder lines (referred to as 'Other' in the Results section) are iterations of previous systems that have proven highly reliable at CERN. Hence, the parameter estimates are determined from the demonstrated reliability of comparable systems in other installations with similarity-based scaling laws employing parts-count complexity approximation and confidence-limit estimation methods (Section 8.3.2.5.2 of (DoD, 1998)).

The SPS shadow events (SPS failures leading to interruption of operation for more than five hours) and regular Maintenance Stops (MS) provide opportunities for repairs. The MTBF of shadow events has been determined from SPS operation data of 2018. The time spent in shadow events or MS is not counted as operational time. Generally, faults covered by the redundancy (RF module faults) are being exchanged during shadow events or MS. All faults other than the RF module failures lead to an interruption of operations are repaired immediately.

The described behaviour is summarized in the state diagram in Figure 13. The simulation starts in the OP state. If there is a failure of an RF module, a transition to the Degraded OP state happens. It will remain there, even if additional RF module failures happen. There are two cases in which the Degraded OP state is left. Firstly, it transitions back to OP if there is an external failure allowing to replace the RF module in its 'shadow' or a scheduled maintenance stop and, secondly, it transitions to the Down state if there is a critical failure which cannot be compensated. This is either when there are too many faulty RF modules or when a failure of a non-redundant system occurs. Then the simulation enters the Down state until the repairs of the faulty systems are completed.





Figure 13: State diagram illustrating the simulated behaviour of the system.

A Monte-Carlo simulation of the mentioned system configurations is carried out using AvailSim4. To accumulate sufficient statistics, the system has been simulated on a computing cluster for an


equivalent of 3000 years of operation. The range of evaluated system configurations is provided in the Results section.

7.2.3 Simulation Results


This section presents the simulation results for the studied system configurations.




(a) 90% confidence lower bound.




(b) Less pessimistic point estimate.




(c) No shadow interventions.




(d) Redundant Drivers.




(e) 90% confidence lower bound.



(f) Less pessimistic point estimate.



(g) No shadow interventions.



(h) Redundant Drivers.

Figure 14: Top Row: System failure rate. Bottom Row: System unavailability. The results are shown as function of the number of installed modules (x-axis) and the MTBF of the RF-modules (three bars per Redundancy level). The colors indicate the attribution of failure rate and unavailability to sub-systems.

The simulation results of the reference configuration using pessimistic parameters (90% confidence lower bounds) are shown in Figure 14.a and Figure 14.e, respectively. The results of the same configuration using less pessimistic point-estimates are shown in Figure 14.b and Figure 14.f, respectively.

The failure rates and un-availability are higher for the more pessimistic estimation method. Moreover, the contribution stemming from the RF modules is decreasing once redundancy is employed. Already two additional RF modules (720074) almost completely screen RF module failures from operations for all studied MTBF values.

Subsequent configurations are all based on the point estimates. Figure 14.c and Figure 14.g show a configuration in which faulty RF modules covered by the redundancy are only being exchanged during MS and not during shadow events. In comparison to the reference configuration, the fault-tolerance effect is noticeably decreased.

Finally, Figure 14.d and Figure 14.h show a modified system with redundant drivers by adding a combiner and splitter between the drivers and the towers and by employing the same repair policy as is used in the reference configuration for RF modules. This results in an almost complete screening of driver faults from operations.

Overall, the results underline the effectiveness of the chosen redundancy strategy. RF amplifier failures can be almost entirely screened from operation with 74 installed RF modules (which is well below the 80 modules that can be installed), if no common-mode failures occur. Comprehensive testing and qualification carried out during the development, manufacturing, and installation of the cavity combiner system gives high confidence that common-mode failures will not become an issue during operation. In addition, extensive monitoring is in place to identify any traces of such behavior.

Making drivers redundant would lead to a significant availability improvement and can be recommended. Similarly, performing repairs during shadow events improves the system availability. The expected number of RF module replacements is in the order of three per month, for which sufficient spares have to be maintained.

7.2.4 Reliability Impact of Using Cavity-Combiners

Owing to the fact that the output power of the implemented cavity-combiner solution scales linearly with the number of intact RF modules, it can almost completely screen failures from operations with $n=75$ installed modules. In other words, at a ratio of installed to required power of 75 to 72, almost all faults are screened.

Had 3 dB-combiners been used instead of cavity-combiners, the same powering ratio would have lead to a redundancy of 630064 (note that $n=64$ modules instead of 75 since 3 dB-combiners are 2^n combiners, which can only take powers of two as input numbers) since the output power would scale non-linearly as a function of the intact modules m :

$$P_{\text{required}} \leq \left(\frac{m^2}{n^2} \right) P_{\text{installed}} \Leftrightarrow$$

$$m \geq \sqrt{\frac{n^2 P_{\text{required}}}{P_{\text{installed}}}} = \sqrt{64^2 \frac{72}{75}} = 62.7.$$

Assuming the same RF module MTBF of 346720 h and repair of faulty modules during MS, the cavity combiner solution achieves an MTBF of 836000 h, which is an order of magnitude higher than the 3 dB combiner variant at an MTBF of 35500 h (See Equ. 1 on page 90 of (AFMC, 1993) for calculation

details.). This illustrates that cavity-combiners are the preferred solution to increase the reliability and reduce the cost of an RF powering solution. Presently for MINERVA and MYRHHA 3dB combiners are foreseen due to the small power per amplifiers and the resulting low number of RF transistors per amplifier with commonly used RF transistor output powers, for which use of cavity combiners proves difficult when single transistors fail.

The difference between cavity and 3dB combiners can be explained by considering the details of combining power in 3 dB and in cavity combiners when one input has failed. The 3dB combiner has two inputs and two outputs. When the two inputs are phase corrected, the power at the two outputs can be calculated as:

$$\Sigma = \frac{P_1 + P_2}{2} + \sqrt{P_1 P_2}, \Delta = \frac{P_1 + P_2}{2} - \sqrt{P_1 P_2}.$$

Using the same amplitude P at both inputs, the power gets combined into one output:

$$\Sigma = \frac{P + P}{2} + \sqrt{P^2} = 2P, \Delta = \frac{P + P}{2} - \sqrt{P^2} = 0.$$

However, if one of the input fails, the power gets split across the outputs, which leads to significant unwanted losses:

$$\Sigma = \frac{P + 0}{2} + \sqrt{P \cdot 0} = \frac{P}{2}, \Delta = \frac{P + 0}{2} - \sqrt{P \cdot 0} = \frac{P}{2}.$$

Figure 15 shows an example of a 64:1 3 dB combiner. Figure 16 shows a schematic 3 dB combiner with eight working inputs on the left image, which are combined without losing power (in a realistic scenario there would be minor losses). On the right image it shows what happens if one of the inputs fails. Imbalanced inputs lead to significant losses in the combiner, so that of the 7P input power, only 6.125P is left at the output. This means that amplifiers with 3 dB combiners need to be equipped with bigger margins to continue operation with faulty inputs.



CERN SPS 64 to 1 combiner @ 200 MHz

Figure 15: CERN SPS 64 to 1 200 MHz combiner. (Montesinos, 2017)




Figure 16: Schematic 3 dB combiner with eight inputs. Left: All inputs work and are combined without losing significant power. Right: One input failed, which leads to significant losses.

Cavity combiners, as shown in Figure 17, do not suffer from this limitation. If one input fails (Figure 18) this does not lead to significant losses inside the combiner. Hence, amplifiers with CC combiners can be equipped with smaller margins to run with faulty inputs.



CERN 144:1 cavity combiner 144 kW CW @ 200 MHz
calculations by ESRF within CRISP, construction by CERN

Figure 17: CERN 144:1 cavity combiner. (Montesinos, 2017)



Figure 18: Schematic cavity combiner with eight inputs. Left: All inputs work and are combined without losing power. Right: One input failed, which leads to no additional losses.

This shows that different power scaling results in different system reliability estimates if the power margins are kept the same for CC and 3dB combiners. This is illustrated in Figure 19, where the resulting MTBF of the redundant scheme (y axis) is plotted as a function of the power margin (x axis). One can see that an amplifier with 3dB combiner needs a 720080 powering margin to match the reliability of a CC with 720075 margin, which would represent a significant cost saving for the CC. Lastly, it has to be pointed out that CCs can only be realised when there is a sufficient amount of inputs (i.e. at least several tens of inputs). Hence, for low power applications where only a few SSPAs are combined, it cannot be made use of this advantage. This is the case for most of the MYRRHA amplifiers.




Figure 19: The different power scaling between 3dB and cavity combiners (CC) results in different system MTBF (y-axis) at the same power margins (x-axis).

7.2.5 Conclusions & Discussion

An availability study of a solid-state RF powering solution using cavity-combiners was presented. The results show that the employed redundancy effectively screens failures of RF powering modules from operation even at moderate margins of installed versus required RF power. Even margins as low as 4% would be sufficient to screen faults effectively as is shown in Figure 14: Top Row: System failure rate. Bottom Row: System unavailability. The results are shown as function of the number of installed modules (x-axis) and the MTBF of the RF-modules (three bars per Redundancy level). The colors indicate the attribution of failure rate and unavailability to sub-systems. Figure 14 b, where the contribution of RF module failure vanishes as soon as 72oo75 redundancy is used.

A similar 3 dB combiner based solution would achieve an order of magnitude lower MTBF at the same RF power margin, as is shown in Figure 19.

The generated model can be adapted to optimize similar RF powering solutions requiring high availability, e.g., for accelerator driven systems such as MYRRHA. The chosen modeling methodology could be validated for a realised fault-tolerant system and will be reused for modeling of MYRRHA fault-tolerance schemes.

The operation of the installation is monitored and it has been able to continuously operate during the first three months of operation (no faults detected). This represents a demonstrated failure rate of 2 per 10⁴ hours ($= 1/(2 \text{ installation} \cdot 92 \text{ days} \cdot 24 \text{ hours})$), as shown in Figure 20. It is expected that the system continues to operate without failures for several months to come to get closer to the predicted reliability. Should any unexpected failure modes occur, detailed investigations will be launched.




Figure 20: The red dashed line shows the demonstrated failure rate after three months of failure-free operation in comparison to the predicted simulation results.

The relevant parameters and characteristics of the SPS fault-tolerance scheme are summarized below:

Cost/margin: <5%

Benefit: in comparison to tube-based RF amplifiers, which is a key fault contributor in linear accelerators, failures are essentially removed

Key factors: SSPA technology; large overall power to single amplifier power ratio, cavity combiner, frequent access for amplifier replacement

Risks/unknowns: low, as system is operational. Low possibility that degradation may reveal unknown weaknesses.

Alternative: The proposed solution already presents a highly optimized scheme.

7.3 MYRRHA Injector Switchover Analysis

7.3.1 System Description

The injectors are around 30 meters long and produce a beam of 17 MeV protons. They consist of a range of specialized accelerator systems, such as the source, LEBT line (Low-Energy-Beam-Transfer line including beam diagnostics, solenoid magnets, chopper, collimator), RFQ (Radio Frequency Quadrupole), MEBT line (Medium-Energy-Beam-Transfer line including beam diagnostics, corrector magnets, beam-rebunching cavities), and normal-conducting CH cavities. General infrastructure and service systems, such as controls, water cooling, and pressurized air, are provided for the injectors.

The proposed fault-tolerance scheme is to duplicate the entire injector and use a fast switching magnet to select the active one. Most likely a hot-spare redundancy will be implemented, meaning that most systems in the non-active spare injector will remain powered, and a probe-beam is continuously produced and either sent to the injector dump or dumped in earlier sections (e.g. in LEBT). The injector dump continues is bent at 45 degree from the injector beam line. The hot spare configuration implies that the failure rate of the spare injector will be practically the same as for the active injector, which has to be considered for the modeling.

A schematic overview of the injector layout is shown in Figure 21 with the beam marked as a grey line, currently being injected from injector one. In case the beam stops being delivered from injector one due to e.g. a fault, a switching to operation from injector two has to be carried out. This involves the following steps:

1. Rapid detection of beam instability or an equipment fault and stopping of beam production in the active injector. Repair has to be initiated as soon as possible.
2. The fast switching magnet (the green magnet where the beam from both injectors leads to) has to switch polarity to receive beam from the other injector
3. The other injector, which has been in hot-spare standby previously, has to send a test beam and then gradually ramp up the beam intensity to nominal conditions.

Once the repair of the previously active injector has succeeded, the process can be repeated when a failure in the currently active injector occurs.




Figure 21: Schematic illustration of the MYRRHA accelerator and its fault-tolerances with the beam currently being delivered from injector 1.

All these steps have to be carried out in under three seconds to not interrupt the reactor operation.

At the level of the redundant injector set up, the list of possible failure scenarios includes:

1. The non-active injector is not ready to take up operations; e.g. because of a recent failure in hot-spare standby mode (most systems are powered even when no beam is produced) or because the repair of a failure from a previous active mode takes very long.
2. The switching mechanism does not get triggered; e.g. because a beam instability originating from the active injector is only noticeable further down in the main LINAC and cannot be identified as an injector problem.
3. The switching execution does not succeed; e.g. because the spare injector cannot reproduce the beam of the active injector and this prevents ramping up to nominal beam intensity in under three seconds.
4. External, singular and common mode failures. E.g. a failure of the fast switching magnet.³

Points 2 to 4 will require special attention and are best addressed when more detailed information about the systems are available. Point 1 can be assessed by availability simulation to identify requirements on the MTBF and MTTR of each injector, which will be discussed in detail in the following.

³ An analysis of an asymmetric injector layout which can continue operation from one injector without the fast switching magnet being active is presented in Appendix 11.3.

7.3.2 Modeling Approach and Parameters




Figure 22: Overview of the modeled systems within the MYRRHA top-down model.

Figure 22 shows an overview of the elements that are discussed in this section. To study the redundant injectors set-up and identify the injector MTBF and MTTR requirements, it is sufficient to model the systems at the top injector level without considering the individual components of the injectors. Moreover, all reasons for the switching not to succeed can be grouped in a single success rate of the switching process. Lastly, external, singular and common mode failures are additive and do not need to be modeled explicitly to study the fault-tolerance scheme. However, they are modeled as failures of external systems, which trigger external reactor turnarounds and allow for more time to carry out repairs of the explicitly modeled systems. This is summarized in Table 4 including all the chosen parameters to carry out a multi-variate sensitivity analysis of the fault-tolerance scheme to identify critical dependencies.

Table 4: Chosen input parameters for the simulation. TTR stands for time-to-repair and the statement of $TTR < T$ is different from $MTTR < T$ in the sense that in the former all repair times are shorter than T and in the latter the mean of all repair times is shorter than T . Here, this means that the repair should not delay a reactor turnaround.

System	Logic	Failure behavior	Repair behavior	comment
Injectors I1, I2	(1oo2)	MTBF=[32,...,1024] hours, exponentially distributed	MTTR = [2,8,32] hours, exponentially distributed	MTTR of LINAC2/4 is around 1 hour; MTBF of LINAC2 (50 MeV) around 100h
Switching mechanism SW (excluding switching magnet)	When required	$P_{fail} = [0., 0.05, 0.1]$, probability to fail per use	$TTR < \text{reactor turnaround time}$	Assuming $P_{fail} = 0.05$ in top down model
Other systems triggering reactor turnaround Ext.	Allowing more time for repairs	MTBF = 30 days, exponentially distributed	$TTR < \text{reactor turnaround time}$	Note that fewer failures here means a worse situation for the fault-tolerance scheme due to less time for repairs.

The chosen parameters cover wide ranges of possible operational scenarios. They have been chosen to cover the expected actual situation based on experience in other facilities and discussions within the working group. Note that it was chosen to model the hot-spare injector with the same failure rate as the active injector, which represents a quite realistic but possibly slightly pessimistic assumption.

The state diagram in Figure 23 gives an overview of the states that the simulation transitions from and to. The simulation starts operation with injection from injector 1 (I1 OP). If there is a fault in the active injector; the switching mechanism works; and the other injector is ready to continue operation, then the transition to operation from injector two (I2 OP) happens. If either the other injector two is not ready or the switching mechanism does not work, the reactor turnaround (RT) is triggered. After the reactor turnaround set time (12 hours) is over, operation continues from any of the two injectors that is ready to take up operation. If no injector is ready, the Long Repair state is entered until either of the two injectors has their repairs finished and is ready for operation. If any external system fails, the external reactor turnaround (Ext. RT) is triggered, which allows for more time for repairs. If the set time (12 hours) is over, operation continues from one of the two injectors. If the simulation has spent 90 days in either I1 OP or I2 OP the simulation is terminated.




Figure 23: State diagram illustrating the simulated behaviour of the redundant injectors system.

The relevant metrics are how often the reactor turnaround is triggered within 90 days of operation and the effectiveness of the fault-tolerance, which is measured as the fraction of times the reactor turnaround is entered of all the times that the active injector fails. These metrics are explored for the range of parameters shown in Table 4 using 100 MC evaluations per evaluation.

7.3.3 Simulation Results

Figure 24 shows a scenario in which the switching mechanism always succeeds. Hence, reactor turnarounds are only triggered when the active injector fails and the spare injector is not ready to take up operations due to recent failure or longer repair time.

The y axis shows the number of reactor turnarounds as a function of the injector MTBF on the x axis. For each injector MTBF, three injector MTTRs are chosen and represented in different bars. As reference, a non-redundant situation with only one injector is plotted as blue crosses. The requirement to be in line with the top-down model (Figure 6) of 0.95 failures per 90 days ($=4.4\text{e-}4$ /h) is shown as red bar.




Figure 24: Estimated failure rates as a function of injector MTBF, MTTR with perfect switching between injectors.

Figure 25 and Figure 26 show the situation for a switching mechanism failing 5 and 10 percent of times, respectively. One observes that with higher MTBF and lower MTTR the failure rate generally decreases. Moreover, it increases with the switch failure probability, which is intuitive. Upon closer inspection, one can see that the blue bars with injector MTTR of 2 hours are more dependent on the switch failure probability than the green bars with MTTR of 32 hours. The interpretation of this is that for small MTTR, the reactor turnaround occurrence is triggered more often by the switching than the injector unavailability. Hence, the impact of the switching failure probability is more visible.




Figure 25: Estimated failure rates as a function of injector MTBF, MTTR with 5% failure probability when switching between injectors.




Figure 26: Estimated failure rates as a function of injector MTBF, MTTR with 10% failure probability when switching between injectors.

In comparison to the non-redundant solution, even with a switch failure probability of ten percent the redundant solution achieves a ten times lower failure rate as long as the MTTR is a few hours. Hence, this represents an effective fault-tolerance method.

The MYRRHA reliability requirement is satisfied for a range of parameter combinations, with the injector MTBF being at least 256 hours, the MTTR being less than 8 hours and the switching failure rate less than 5 percent. These numbers represent the absolute limit, which is tolerable, especially when considering that common mode failures have not been considered yet.

The initial assumption was that the expected switching success rate will can reach 95% and that the common mode failures can be minimized. At the same time, considering feasible MTBF and MTTR values as stated in the previous paragraph, it was concluded to assign a fault-tolerance effectiveness of 91% in the top-down model.

7.3.4 Conclusions & Discussion

The proposed injector fault-tolerance scheme is relatively expensive due to the doubling of the infrastructure but represents an effective means of fault-tolerance. However, the successful and reliable switching between injectors still needs to be tested, especially in terms of beam reproducibility for successful injection into the main LINAC. This can be tested with the MINERVA setup. An overview of the key facts of the redundancy scheme are given in the following:

- Cost/margin: >100%
- Benefit: reducing injector faults by a factor ten as can be seen in Figure 26 by comparing the non-redundant solutions with the redundant solutions
- Key factors: reliable switching between injectors, combination of high injector MTBF and low MTTR
- Risks/unknowns: Reproducibility of beam (especially when switching infrequently between injectors; switch by default on a regular basis?)
- Alternatives: identifying sources of reliability problems during MINERVA and applying more targeted improvements instead of duplicating the entire injector.
- Improvements: When the second injector is being designed and built one can profit from the experience with the first injector and correct its weaknesses while keeping proven designs. Care has to be taken to avoid common mode failures by duplicating the design.

7.4 MYRRHA Main LINAC Dynamic Compensation Analysis

7.4.1 System Description

An overview of the main LINAC section within the MYRRHA facility is shown in Figure 27. It mostly consists of a chain of superconducting RF cavities, which are housed in cryostats. There are three different sections. The first section with 62 single-spoke cavities in 32 cryostats, the second section with 18 double spoke cavities in 9 cryostats and finally the third section with 72 elliptic cavities in 18 cryostats. Between the single and double spoke sections is a single rebuncher cavity.




Figure 27: The entire MYRRHA ADS with the main LINAC highlighted by the orange box.

Each cryostat has a valve box attached, which is the link to the cryo-plants via a cryogenics distribution system. The RF tuning mechanism, consisting of an electro-mechanical and a piezo actuator, is integrated in the cryostat.

Each RF cavity is individually powered and controlled. This means that for each cavity there is an RF power amplifier, circulator, feeder line, power coupler to the cavity, and LLRF (Low Level RF controls), forming an independent RF circuit

This individual powering and controls allows a serial redundancy scheme in which the failure of one RF circuit does not impact other circuits and can be compensated by ramping up the accelerating fields of other (neighbouring) cavities. In case of a cryo module failure, all the cavities within the cryo module will be detuned and have to be compensated by even more neighbouring cavities. This compensation scheme involves a sequence of steps, as shown by each row in Figure 28 and listed below, which need to be carried out successfully in under three seconds:

1. A cavity is not operational (row 1)
2. Within three seconds during a so called dynamic compensation (DC) procedure (row 2):
 1. Beam stops
 2. Fault detected & localized
 3. Identify compensating configuration & detune faulty cavity by moving the tuning mechanism 'out'
 4. Neighboring cavities set-point changed to pre-set compensation configuration & RF fields ramped up
 5. Beam restart and iterative tuning corrections
 6. Beam intensity ramp-up
3. Repair of faulty RF circuit initiated as soon as possible and later completed (row 3)
4. Repaired cavity brought back 'online' (row 4)

If these steps cannot be carried out successfully, the reactor has to enter a reactor turnaround (RT) procedure.




Figure 28: Illustration of the sequence of steps involved in the dynamic compensation scheme.

After completion of steps 3 and 4 the fault-tolerance scheme leaves the more vulnerable compensating mode. Hence, they should be done as early as possible. However, regarding step 3 it seems that the majority of faults are not accessible for immediate repair but only when there is no beam, which is the case during reactor turnarounds or at the end of the fuel cycle. Moreover, even if step 3 could be carried out immediately, carrying out step 4 during a dynamic compensation (DC) poses additional challenges since the procedure of step 2 has to be repeated with re-tuning the repaired RF cavity into normal operation, which is expected to be difficult as the repaired cavity will have been completely detuned. Hence, the foreseen scheme will be to wait for repairs until a reactor turnaround or the end of a fuel cycle. This means that more faults will accumulate before a repair can be carried out.

Notable scenarios in which the fault-tolerance scheme fails include the following:

1. Too many failures accumulate and cannot be compensated anymore.
2. The switching from normal to compensating configuration does not succeed in under three seconds.
3. Controls or LLRF systems are giving wrong instructions to the systems under control.
4. The tuning mechanism gets stuck in a way that does not allow to detune a faulty RF circuit.
5. External, singular and common mode failures.

Points 3 to 5 will require special attention and are best addressed when more detailed information about the system designs is available.

Point 1 can be assessed by availability simulation to identify both requirements on the reliability of the RF and cryogenics systems as well as the flexibility of the compensation scheme, taking into account an assumed behaviour regarding points 2 – 5. This will be addressed in the following subsections.

Regarding the flexibility of the compensation scheme, once in a compensating mode, beam dynamics simulations have confirmed that the accelerator can compensate a range of fault configurations (Bouly, 2022), which will be discussed in more detail in the modeling subsection. It generally shows that the longitudinal beam matching when compensating cavities in the low-energy section will be one of the main challenges, due to beam halo creation. At SNS it has been demonstrated that the accelerator can

be run in compensating mode after a certain transition time from normal to compensating mode, which is largely exceeding the three seconds required for MYRRHA by far .

This leads over to point 2 of the list: the transition from normal to compensating mode in under three seconds. This is posing several challenges. It requires a very rapid movement of the tuning mechanism, which has been estimated to be feasible, in principle. The tuning will first be set to a pre-defined compensation setting, which has to be identified 'off-line' during extended testing in the commissioning period of the accelerator. On top of that pre-set configuration some fine-tuning will be necessary, which adds an iterative component requiring very fast 'on-line' optimization during operation. Good off-line optimization during commissioning tests will be essential to keep the on-line optimization as short as possible to stay under the required three seconds rule.

For both the off-line and on-line optimization the number of compensating cavities requiring re-phasing is critical as more of them means more possible configuration to test during off-line optimization and more tuning iterations during on-line optimization. Hence, fewer compensating cavities would promise a faster and more reliable transition from normal to compensating mode.

Distributing the compensation across fewer cavities implies a higher accelerating gradient per compensating cavity and this directly translates to a higher power margin of the RF powering of the cavities. The power margin is strongly impacting the cost of the system as it needs more up-front investment cost to install a higher power amplifier and also a higher operating cost as they generally operate at lower efficiency when ran further from their nominal power in the nominal conditions. Hence, from a cost-point of view it would be best to distribute the compensation across as many cavities as possible. Lastly, when compensating with more cavities the difference in the set-points between normal operation and compensating operation is smaller and this could be easier to establish.

To balance these aspects it had been proposed to use four RF circuits to compensate a faulty one. This implies a relatively low number of cavities requiring re-phasing during a dynamic compensation and reasonable RF power margins. The margins have been evaluated by (Bouly, 2022) and are between 40% and 70% of the required RF power for nominal operation. A detailed overview of the dynamic compensation margins is shown in Figure 29. Averaged over the entire MYRRHA main LINAC, the required margin for dynamic compensation is approximately 45%. This neither includes margins for tuning and losses (around 30%) nor for fault compensation inside the RF amplifiers (10-20%; see Sec. 7.5 for details), which have to be added in addition. Since margins are leading to significant costs, the goal of the modeling is also to assess whether the chosen scheme of four compensating RF circuits is required or if the margins can be reduced by using more compensators. In the model this will be assumed to be a fixed parameter even though in practice this could be implemented dynamically depending on the actual margins present in the installations, which may only be precisely known by measurement.




Figure 29: Estimation of required RF power margins for dynamic compensation when four compensating cavities are being used. Estimated by beam-dynamics simulations (Bouly, 2022).

7.4.2 Modeling Approach and Parameters

Figure 30 shows an overview of the systems to be modeled. These are the RF systems and the cryo-modules covered by the dynamic compensation fault-tolerance scheme.




Figure 30: Overview of the modeled systems within the MYRRHA top-down model.

The modelling described in this section will aim to find the best trade-offs between RF margins, compensation flexibility, system failure rates, the MYRRHA reliability goals and the fault-tolerance effectiveness. To do so, the model should:

- Allow to study sensitivity with respect to:
 - MTBFs of sub-systems
 - Repair accessibility of sub-systems
 - Different compensation scheme parameters
 - Edge cases at transition from one accelerating section to other
- Not require a cluster computing facility for exploratory simulations
- Be understandable & maintainable in the longer term

In this regard, it is based on a dynamic compensation model including the following aspects:

- RF systems and Cryostat failures
- Imperfect switching (the results in the following will often assume perfect switching)
- Dynamic compensation schemes explicitly modeled for each RF circuit
- Reactor turnaround & fuel cycle procedures and associated repair policies/accessibility

As for the redundant injector model, to study and optimise the effectiveness of the dynamic compensation scheme it is not necessary to model the following aspects:

- Cryostats & RF circuits have single failure mode, instead of modelling all their sub-systems
- Other systems are only modelled to trigger external reactor turnarounds, giving opportunities to repair cryostats and RF systems
- The different accelerating stages are not combined in one model. Instead, there are different models for the different sections to avoid a mix of effects when carrying out sensitivity analyses.

More concretely, Figure 31 shows the modeled components. These are RF elements, each representing an entire RF circuit, and Cryo modules, each representing a cryostat and its valve box, housing two cavities in the spoke section or four cavities in the elliptic section. Both have a single failure mode attached. They form a chain of repeating elements making up the main LINAC.




Figure 31: Illustration of the modeled components.


To explicitly model the compensation scheme, two parameters are used:

- Number of compensating cavities per detuned cavity – N_c : E.g., if $N_c=4$, this means that if a single spoke cryo-module stops working, the two cavities therein need to be compensated by two times four cavities, making eight compensating cavities in total. This parameter has a


strong influence on the required RF powering margin for RF circuits since the power that is missing in the detuned cavity needs to be distributed across the N_c compensating cavities. In theory this scales with $1/N_c$. So, the more compensating cavities per detuned cavity, the fewer additional power needs to be provided by the compensating cavities, resulting in a smaller RF power margin. In practice, beam dynamics simulations indicate that there are limits to reducing the margin by distributing the compensation over more and more cavities since the direct neighbours of the detuned cavity will often have to bear more of the compensation load.

- Maximal compensation distance to detuned cavity – C_d : E.g., if $C_d=4$, this means that the compensating cavities can be located up to four cavities away from the detuned cavity. For the case described earlier of a faulty single-spoke cryo-module which requires eight compensating cavities, a C_d of four is necessary so that four compensating cavities can be placed to each side of the faulty cryo-module. This is also shown in the fourth row ($N_c=4$, $C_d=4$) of Figure 32, which also shows a few other examples of possible compensation configurations and one which cannot be compensated (row three) given the chosen parameters for N_c and C_d .


$N_c = 4$, $C_d = 2$:




$N_c = 4$, $C_d = 3$:



$N_c = 4$, $C_d = 3$:



$N_c = 4$, $C_d = 4$:



$N_c = 4$, $C_d = 4$:




Figure 32: Examples of fault configurations which can or cannot be compensated by a compensation scheme with the respective parameters for N_c and C_d .

These are the basic building blocks of the model. Figure 33 shows the state diagram, describing the system behaviour. The simulation starts in the normal operation state (OP). Should a (in principle) compensatable fault of an RF circuit or a cryo-module happen, the dynamic compensation state (DC) is entered for three seconds. Once there, it is evaluated whether there is a possible compensating configuration, given the chosen parameters for N_c and C_d . If this is not the case, the simulation enters the reactor turnaround state (RT), which means that beam production is stopped for 12 hours. If there is a compensating configuration, it is evaluated whether the switching succeeds using a pre-set failure probability. If it does not succeed, the system enters RT for 12 hours again. If it succeeds, it enters the

compensating operation state (degraded OP), in which the faulty RF circuit is compensated and beam keeps being delivered to the reactor. Whenever a (in principle) compensatable fault happens, the steps involving the dynamic compensation are repeated.

Once a non-compensatable fault is encountered and the system transitions to a RT, RF repairs can be initiated. After the set time of a RT, the system returns to normal OP if all repairs have been finished, to degraded OP if not all repairs could be finished but sufficient to run in compensating mode, or to a Long Repair state if not sufficient repairs could be finished to return to any operational mode. E.g., cryo-module failures are only repaired after the 90 days fuel cycle unless operation cannot resume without the repair, in which case a Long Repair is scheduled to complete the normally lengthy cryogenics repairs. Note that the latter two options have been omitted from Figure 33 for a better overview.

Should any other MYRRHA system fail, the simulation will enter an externally triggered reactor turnaround (ext. RT), which allows for repairs of the explicitly modeled systems, but is not counted towards the performance metrics of the fault-tolerance scheme. If the simulation has spent 90 days in any of the operational modes, the simulation is terminated.




Figure 33: State diagram illustrating the simulated behaviour of the system. Abbreviations: RT – reactor turnaround, DC – dynamic compensation, OP – operations, Ext. RT – Reactor turnaround triggered by systems other than main LINAC SRF & Cryo-module.

Note that the described state diagram represents the reference configuration of the system. In the analysis alternative configurations will be tested, such as having repairs during beam production, which would slightly alter the state diagram.

To fully define the reference configuration, its complete set of parameters is listed below. It models the single-spoke section with the parameters chosen so that they are in line with the MYRRHA top-down requirements models. It consists of

- 30 cryo modules (i.e. Cryostat without tuning mechanism + valve box)

- MTBF = 256 kh
- Repair after fuel cycle
- 60 RF circuits (i.e. RF amplifier, feeder, coupler, cavity, tuner, LLRF)
 - MTBF = 32 kh
 - Repair when reactor turnaround
- ,Rest of MYRRHA'
 - MTBF = 45 days (twice per fuel cycle)
 - Note that an optimistic MTBF is actually a pessimistic assumption as it leads to fewer repair opportunities.
- Singular (non-compensatable) points of failure
 - Not used in the reference model, but in a 'pessimistic' model derived in the results section
- Number of compensating cavities per detuned cavity – $N_c = 4$
- Maximal compensation distance to detuned cavity – $C_d = 4$
 - Note that this allows compensation of faults at beginning/end of the LINAC
- 100% success rate of dynamic compensation execution **if** compensating configuration exists

The following simulation parameters have been chosen:

- Simulation terminated after 90 days of delivered beam (time spent in any operational state)
- 2000 Monte Carlo iterations per data point

The reported performance metrics are the following:

- The number of faults – S = (Sum of RF & Cryo-module faults not covered by dynamic compensation). Note that they will be multiplied by 152/60 to scale up to the full length of MYRRHA. This is a deliberate choice to make it comparable with top-down assumption. These faults are defined as those triggering a reactor turnaround (RT) and do not include faults of external MYRRHA systems triggering Ext. RTs.
- The non-coverage, which is the fraction of RF & Cryo-module faults that are not covered by dynamic compensation - non-coverage = $S / (\text{number of dynamic compensations triggered})$. This is a measure for efficiency of dynamic compensation; i.e. a lower percentage is a better dynamic compensation performance.

The results shown in the following subsection use the reference model with no singular points of failure and a 100 % switching efficiency. However, these two factors are expected to be present. Hence, a second, more pessimistic scenario will be derived, in which these factors are present.

7.4.3 Results

As pointed out, there is a relatively big range of parameters influencing the effectiveness of the fault-tolerance scheme. These include the following, among others:

- The switching success rate from a normal configuration to a compensating configuration, or from one compensating configuration to another one
- The singular and common mode failures, which cannot be compensated by the scheme
- The MTBFs of the RF and Cryo-modules
- The 'flexibility' of the scheme to compensate failures, expressed by the compensation distance C_d and the number of compensators N_c
- The accessibility to repair systems

Since studying all these factors simultaneously would be overly complicated to interpret, each factor will be studied individually by varying it with respect to the reference model. This will allow to identify upper and lower bounds to these parameters, which can then be used when studying combinations of factors.

7.4.3.1 *The switching success rate and singular points of failure*

The first investigated parameter is the success rate of switching between different accelerator fault-tolerance configurations in under three seconds. In a first attempt, this can be done without simulation. In line with the reference configuration (as shown in the MYRRHA top-down model in Figure 6), the number of RF and cryo-module faults within a fuel cycle of 90 days, which should be compensated, is ten. This means that for each of this ten faults a dynamic compensation is initiated.

In the most pessimistic case, if the switching success rate was 0 %, all ten faults would trigger a reactor turnaround, which is clearly above the requirement of not triggering more than 3.3 RTs due to the RF systems or the cryo-modules per fuel cycle. In the most optimistic case, i.e. no singular points of failure and compensating configurations are always available, the switching success rate would have to be at least 67 % to conform to the requirements ($3.3RT's = 10 \text{ RF and cryo faults} * (1 - 0.66)$ where 66% is the fraction of compensated faults = switching success rate), which represents a lower acceptable bound.

The same argument can be made for singular points of failure, e.g. a stuck cold-tuning-system. Here the most optimistic case is that the switching success rate is 100 % and that there are always compensating configurations available. Hence, all of the failures except the singular ones can be compensated. In this case, up to 3.3 singular (or common mode) failures can be tolerated, which represents an upper bound.

In practice, one can expect that both imperfect switching and singular points of failure may occur, even though the aim is to remove them as far as possible. As it is very difficult to estimate which category will dominate, realistically achievable values are assumed.

Simulating the reference model with perfect switching and without singular failures shows that on average 0.56 triggered RTs are due to non-availability of compensating configurations (because statistically too many faults accumulate in certain regions in which not sufficient RF circuits are available for compensation – blue line in Figure 35 for $Cd = 4$). Adjusting the switching success rate to 85 %, which seems feasible, the amount of RTs from non-availability of compensating configurations and imperfect switching sum up to 2.08 triggered RTs on average (orange line in Figure 35 for $Cd = 4$), which is a bit less than two thirds of the tolerated RTs. The remaining third is assumed to be ‘consumed’ by singular points of failure. It was estimated that having 1.65 singular failures (e.g. CTS stuck or LLRF misbehaving) on top of the imperfect switching and the non-availability of compensating configurations leads to having 3.18 triggered RTs on average overall (grey line in Figure 35 for $Cd = 4$), which is very close to the requirement of 3.3 triggered RTs. Note that the singular failures cannot simply be added to the 2.08 triggered RTs since they provide additional repair opportunities which reduce the count of triggered RTs due to non-availability of compensating configurations.

Hence, the values of an 85 % switching success rate and 1.65 singular failures per fuel cycle represent a realistic set of requirements, which will be used as ‘pessimistic’ model in the following. They also represent the case used for the top-down model with and overall fault-tolerance effectiveness of approximately “<66%>”, as shown in Figure 6.

7.4.3.2 *RF and Cryo-Module Failure Rates*

As argued above, for ten RF and cryo-module failures the reliability requirements can be fulfilled if the switching success rate is at least 85% and the 1.65 of the ten faults are singular and non-compensatable. This puts very stringent requirements on the system failure rates (MTBF = 256 kh for cryo-modules and 32 kh for RF circuits - where kh stands for thousand hours, see 7.3.2) and one may wonder if this could be relaxed in case the switching success rate is improved and the fraction of singular points of failure reduced.

The results of this investigation are shown in Figure 34, where the sensitivity of the number of RTs triggered (left y axis) with respect to the total number of (non-singular) RF and cryo-module failures (x axis) is tested, whilst assuming the idealized scenario of 100% switching success rate and no singular points of failure. In such a conditions up to 30 faults could be compensated whilst triggering not more than 3.3 RTs. It is important to note that the ‘theoretical coverage’ (right y axis), the fraction of RF and cryo-module faults not triggering a RT, is going down with increasing system failure rates since more faults will accumulate before an ‘externally’ triggered RT allows for repair of the systems.⁴ Note that the ratio of RF and cryo-module faults is kept constant across the range of the x-axis.




Figure 34: Impact of RF & Cryo module failure rates (x-axis) on effectiveness of fault-tolerance scheme expressed in number of triggered RF (left y-axis) and coverage (right y-axis).

7.4.3.3 Compensation Distance C_d – The flexibility of the scheme

The compensation distance determines how far a compensating RF circuits can be located from the faulty one. Hence, the larger the distance, the more flexible the scheme and the more faults within a local area can be accumulated. Note that this is because the number of compensating cavities per faulty one is kept the same. This effect is studied by evaluating the sensitivity of the number of RTs triggered (y-axis) with respect to the compensation distance (x-axis), as shown in Figure 35. Throughout, the number of compensators is kept at four, as is the case for the reference configuration.

One can see that the reference configuration with perfect switching and without singular points of failure triggers the fewest RTs. There is a strong dependence with C_d . Firstly, going from no dynamic compensation (‘without redundancy’) to an inflexible compensation scheme with $C_d = 2$, already reduces the triggered RTs from 10 to 2. A strong further reduction to 0.74 RTs is observed when the reference configuration with $C_d = 4$ is approached. Going to larger C_d can reduce the number of RTs further, but is likely not possible due to beam dynamics limitations. Moreover, were it possible, it may lead to a less reliable switching to compensating configurations because of increased numbers of cavities to re-phase if the compensators are located far from the detuned cavity. At the same time, the reduced change in parameters may lead to a more reliable switching.

⁴ Expressed as formula: $x\text{-axis-value} * (1 - \text{right-y-axis-value}) = \text{left-y-axis-value}$




Figure 35: Impact of the compensation distance on number of triggered RTs.

When imperfect switching (85% success rate) is considered, the number of triggered RTs already increases significantly. Finally, for the ‘pessimistic’ case of imperfect switching with 85% success rate and 1.65 singular points of failure per fuel cycle, the number of RTs are further increased to the point that $C_d = 4$ only just meets the requirements.

Overall, it can be concluded that the expected foreseen compensation distance of $C_d = 4$ allows for sufficient compensation flexibility reaching a significant coverage of faults (9.25 out of 10 faults compensated). The limitations of the scheme are more likely due to imperfect switching and/or singular points of failure.

7.4.3.4 The number of compensators N_c – The flexibility of the scheme

The number of compensators has a strong impact on the required RF power margin, as it can be smaller the more RF circuits are used to compensate a detuned one. However, the number of RF circuits that can be chosen as compensators is limited by the compensation distance C_d due to beam dynamics limitations. Hence, the more compensators are required within this distance, the sooner a situation will arrive in which insufficient compensating RF circuits are available once failures accumulate. This effect is studied here by varying the number of compensating RF circuits, while keeping the compensation distance fixed at four as in the reference model. The sensitivity of the number of RTs triggered (left y-axis) with respect to the number compensators (x-axis) is shown in Figure 36. In addition, the theoretical impact on the RF powering margin is shown (right y-axis), where 100% corresponds to the margin required for four compensators (left end of the grey dashed line in the figure). It is important to note that this is based on the assumption that the compensation can be distributed equally among the compensating cavities. However, in practice there are limits to this and the reduction of the margin is likely smaller than indicated.




Figure 36: Impact of number of compensators per faulty RF circuit on number of RTs triggered.

In the investigated range of N_c (4 to 8), the power margin can theoretically be halved. However, the price for that is an increase of RTs triggered due to a reduced coverage of the fault-tolerance scheme. Concretely, the number of triggered RTs goes from 0.75 to 3.3 for the idealised reference scenario and from 3.3 to 4.7 for the ‘pessimistic’ scenario.

The top-down model reliability requirement could only be met for $N_c > 4$ if the actual system failure rates, the compensation switching success rate and/or the number of singular failures are better than in the assumed ‘pessimistic’ model. However, using more compensating cavities implies a more complex transition to new compensating configurations as more cavities need to be re-phased and fine-tuned. This likely leads to an even lower switching success rate.

Overall it can be concluded that increasing the number of compensators leads to a relevant reduction of the RF powering margin but the fault-tolerance scheme is losing effectiveness. Hence, it should only be implemented if the actual systems perform significantly better than assumed in the top-down requirements models and reliable switching can be guaranteed for more compensating cavities (i.e. if in reality the orange curve in Figure 36 is approaching the blue curve).

7.4.3.5 The accessibility to repair systems - Immediate repair of RF failures

Although the scenario of being able to repair RF failures as soon as a failure happened is likely not possible to realise, it is still interesting to study it to identify the theoretical potential of the dynamic compensation scheme. To do so, the results with immediate RF repairs are compared with the reference models, in which RF repairs are only carried out during RTs, in Figure 37.




Figure 37: Difference in number of RTs triggered between repair of RF components during beam or during RTs only.

For the idealized case with perfect switching and without singular failures, there is a significant difference in the number of RTs triggered up to $Cd=4$ (lower two curves). For the more 'pessimistic' scenario with imperfect switching and singular failures, the difference is less noticeable and practically irrelevant (upper two curves).

The interpretation of this is that for the 'pessimistic' scenario the frequency of RTs is high enough to be able to timely repair faulty RF components. However, for the idealized scenario RTs are rare and being able to access RF components for repair during beam, brings a significant reduction of the time until a component is repaired. Since the 'pessimistic' scenario is expected to be closer to the real situation, the conclusion is that unless the dynamic compensation scheme is already very effective, repairs during RTs are sufficient and sooner (or immediate) access to repair of RF components does not bring significant advantage.

7.4.3.6 Integrating dynamic compensation and power amplifier modular-redundancy margins

The current models assume that an RF circuit is either faulty or operational and able to be in either normal or compensating mode. This means that even if one or several RF pallets inside the RF amplifier are faulty, it needs to be able to deliver sufficient power for the compensating mode. A result of that is that the power margin for modular redundancy has to be counted on top of the margin for dynamic compensation. If one assumes an RF circuit only being able to go into compensation mode when all RF pallets are functional, the margin for modular redundancy can be integrated into the dynamic compensation margin⁵.

⁵ Not vice versa as the dynamic compensation margin is larger than the margin for modular redundancy, as is estimated in Section 7.5.3.

This can be simulated by distributing the compensating cavities only on the fully functional RF circuits within the compensation distance. Instead of simulating all individual RF pallets, it was estimated that the MTBF for at least one RF pallet failing is 6250 hours⁶, which is relatively frequent.

In addition, it has to be considered that each time there is an RF pallet failure in an RF circuit which is already in compensating mode, a dynamic compensation needs to be triggered to move the compensating function to another RF circuit which has no faulty RF pallet. This increases the numbers of dynamic compensations that need to be executed and can cause additional failures if the switching is not highly reliable.

Figure 38 shows the results of these simulations when using four neighbours for compensation ($n_c = 4$). For the situation with perfect switching and no singular points of failure, one can see that the difference between using and not using degraded RF circuits is relatively moderate and the requirements could be reached with the foreseen scheme of $cd = 4$. However, once imperfect switching is taken into account, the situation is dramatically different, since the solution of not using degraded amplifiers will lead to switching much more often (almost 60 times instead of 10 times). Hence, such a scheme can only be recommended if the switching practically never fails or if the RF pallets are so reliable that they rarely fail. In that case it would lead to saving the additional RF power margin for modular redundancy, which is estimated in Section 7.5.3 to be 10-24%.




Figure 38: Effect of avoiding to use degraded RF circuits (i.e. those with faulty RF modules) for compensation.

⁶ This was estimated for the single spoke section by assuming a pessimistic RF module failure rate of $1e-5$ per hour and 16 modules per cavity on average. This results in an overall failure rate of $16e-5$ per hour or 6250 hours.

7.4.4 Conclusions and Discussions

The modeling and simulation of dynamic RF compensation show that the proposed fault-tolerance scheme with four compensating circuits per faulty cavity seems to represent a good trade-off between fault coverage, RF powering margins and number of cavities to rephase when switching to compensating configurations. A further reduction of the power margins by using more compensating RF circuits is possible when experience with MINERVA shows that the expected system failure rates are lower, the singular points of failure well managed, and the switching success rate is higher than assumed in the 'pessimistic' scenario used here. For this it will be essential to test the scheme in depth at MINERVA. An overview of the main characteristics of the fault-tolerance scheme is given below:

Cost/margin: ~45% in RF power, individual powering and control of cavities

Benefit: if switching is reliable and singular points of failure minimal, can reduce faults up to a factor of ten

Key factors: reliable switching, system MTBFs, singular and common mode failures, access to repair, number of compensating cavities, flexibility to distribute compensating cavities

Risks/unknowns: reliable switching, system MTBFs, singular and common mode failures

Alternatives: identifying sources of reliability problems during MINERVA and applying more targeted improvements instead of this rather complex scheme

Possible Improvements: if switching is very reliable, number of compensators could be increased to reduce margins and cost

7.5 MYRRHA Main LINAC RF Powering analysis

7.5.1 System Description

The RF powering provides the high-power RF fields to the cavities to accelerate the beam. Traditionally, tube-based technologies were used for power amplification but there has been a gradual shift towards solid-state amplifiers, which allow for redundancy and fault-tolerance, as described in detail for the SPS SSPAs in Sec. 7.2. This is important for linear accelerators, where a significant fraction of failures may stem from tube-based RF amplifiers (see Sec. 4 for data on LINAC4 and SNS).

The RF amplifiers of MYRRHA will be using solid-state technology and employ fault-tolerance through redundancy in a similar manner as is presently applied for the SPS RF power amplifiers at CERN. Figure 39 shows an overview of the elements of the RF circuits in the MYRRHA main LINAC.




Figure 39: Schematic overview of an RF circuit.

One can see that the RF power amplifier represents only a small part of the entire circuit. However, it includes the majority of high-power active elements, which need particular attention from a reliability point of view. Hence, the analysis will focus on the RF power amplifier in the following. The remaining systems will still require deeper analysis, which is best carried out when more detailed information about them will be available.

Shifting the focus back on the RF amplifier, Figure 40 shows a prototypical architecture of a power amplifier for MYRRHA and a 3D rendering of the cabinet. It consists of several PSUs (power supply units), which provide the electrical power to the SSMs (Solid-state modules). The solid-state modules consist of several (2-16 based on a survey of suppliers) RF pallets, which contain (usually 1 KW) solid-state transistors for generating the RF fields. The outputs of the pallets within the SSMs are entering 3dB combiners. The outputs of the SSMs are connected to progressive 3dB combiners. As one can see in Figure 40, there are actually two amplifiers housed in one so-called cabinet. This will be common for the MYRRHA main LINAC powering since the required power output of a single amplifier is relatively small. This may pose some restrictions on carrying out repairs when one amplifier is faulty and the other one in the cabinet is operational. Finally, the cabinet also houses the controls electronics and interfaces for water cooling circuits, for which no redundancy is foreseen.




Figure 40: Left: Prototypical architecture of a power amplifier for MYRRHA. Right: 3D rendering of a power amplifier cabinet.

The amplifiers will be dimensioned so that failures of RF pallets and PSUs can be compensated. The continuous operation with few faulty SSPAs has been demonstrated several times (see Sec. 7.2.1 for references) and is considered highly reliable.

Failed pallets will be replaced by swapping entire SSMs when there is no beam during RTs or after the end of the fuel cycle. Hot-swapping (i.e. exchange during beam) would be possible in principle, but it is expensive to develop, not a mature application, and it would require that an exchange of faulty pallets can be done at the pallet level instead of exchanging an SSM, as removing a whole SSM would reduce the output power too much to continue operation. Failed PSUs can be individually exchanged and implementation of hot-swapping has been commonly done.

The RF pallet and PSU fault-tolerance strategies address the systems which are expected to be most prone to failure. There remain a range of notable possibilities for how the RF powering could still fail:

- Accumulation of too many failures, which cannot be compensated anymore.

- External, singular or common mode failures. In particular failures of systems not covered by fault-tolerance; e.g. water cooling, controls, combiner, etc.

The latter point requires a detailed analysis of the considered systems, which should be carried out in an integrated manner during system design and subsequent development stages. The former point can already be assessed by an availability analysis to minimize the required redundancy margins whilst satisfying the reliability targets for MYRRHA. This analysis is described in below.

It is to be noted that the study deals with the RF pallets only and does not include PSUs as the dimensioning of the redundancy is more intricate for the pallets. For the PSUs the dimensioning of the redundancy can be carried out in a similar manner, the difficulties arising from the 3 dB combiners can be ignored.

7.5.2 Modeling Approach and Parameters

To identify the necessary redundancy margin for the RF pallets, the reliability target, the estimated failure rate of the pallets, and the output power requirements need to be known.

The reliability target is defined by the top down models from Sec. 5. For the MYRRHA reference model the requirement for failures from the RF powering is 2.5 failures per 90 days per 152 amplifiers, which translates to a failure rate of 7.6×10^{-6} per hour. The RF powering includes the redundant pallets and PSUs as well as non-redundant controls, cooling, combiner and other systems. Since the failure rates of all these systems are not known yet, a pragmatic approach is to assume that the systems covered by fault-tolerance should not contribute more than one tenth of the overall RF powering failure rate. This translates to a requirement of 7.6×10^{-7} per hour for the pallet redundancy assembly.

Regarding the estimated failure rate of the pallets, a survey among pallet suppliers has returned five failure rate estimations between 1×10^{-5} and 1×10^{-6} per hour including both reliability predictions and measured field-reliability. Generally it is observed that predicted reliabilities were significantly lower than measured reliability of operational devices. The standard procedure of estimating the most likely failure rate and a statistical confidence bound cannot be applied here since the sample sizes used for estimating the field reliability as well as the reliability prediction details are not known. Instead a pessimistic model with an RF pallet failure rate of 1×10^{-5} per hour, a likely model with a failure rate of 4×10^{-6} per hour and an optimistic model with a failure rate of 1×10^{-6} per hour will be assumed.

The output power requirements were identified in beam dynamics simulations for both nominal operation without detuned cavities and compensating operation with four compensating cavities per faulty one (Bouly, 2022). They are shown in Figure 29.

With these three factors it is possible to estimate the required redundancy margin to meet the estimated reliability requirement for a given pallet failure rate and required output power. The steps will be explained using the example of a cavity with a nominal power requirement of 9.2 KW, which is characteristic for cavities towards the end of the single-spoke section up to 100 MeV. The same approach applies throughout the LINAC by adapting the power requirement.

For the case of 9.2 KW nominal power, first the margins for tuning and losses of 30% and the margin for dynamic compensation of 55% (in this specific section as estimated by beam dynamics simulations – see Figure 29) need to be added. This results in a power requirement of 18.54 KW (12 KW for tuning margin only).

The most demanding requirement is that the RF pallets need to be able to deliver these 18.54 KW despite some pallets having failures. The number of tolerated faults is related to the resulting failure rate of the redundancy scheme and is expressed by the formula:

$$\lambda_{\text{redundancy}} = \frac{(m + q)! \lambda^{q+1}}{(m - 1)! \mu^q},$$

with $\lambda_{\text{redundancy}}$ being the resulting failure rate of the redundant assembly, λ being the pallet failure rate, n being the number of installed pallets, $q = n - m$ being the number of tolerated faulty pallets within the assembly, m being the minimum number of pallets to deliver the required power, and μ being the mean corrective repair rate for a pallet; i.e. assuming interventions every 30 days on average the mean corrective repair rate is 1/15 days.

The parameters m and q need to be chosen so that the failure rate of the redundancy is in line with the requirements. Intuitively, the more failures are tolerated, the lower the resulting failure rate of the redundancy assembly. Assuming the pessimistic pallet failure rate of 1e-5 failures per hour and a repair opportunity every 30 days on average, a 22oo24 redundancy has a resulting failure rate of 1.6e-6 failures per hour. A 22oo25 redundancy, which tolerates three failures, has 1.4e-7 failure per hour, which is almost an order of magnitude less. Hence, a change of redundancy from 22oo24 to 22oo25 implies an order of magnitude gain in availability. The 22oo24 redundancy does not comply with the reliability requirement of 7.6e-7 failures per hour, but the 22oo25 redundancy does.

To choose m the power that is being delivered when some pallets are faulty needs to be estimated. Since 3 dB combiners are used, the scaling is not linear with the number of active pallets but follows a quadratic law:

$$P_{\text{required}} \leq \left(\frac{m^2}{n^2} \right) P_{\text{installed}} \Leftrightarrow m \geq \sqrt{\frac{n^2 P_{\text{required}}}{P_{\text{installed}}}}.$$

In the former example case the values are: $P_{\text{required}} = 18.54 \text{ KW}$, $P_{\text{installed}} = 25 \text{ KW}$, $m = 22$ and $n = m + q = 25$. Hence,

$$18.54 \text{ KW} \leq \left(\frac{22^2}{25^2} \right) 25 \text{ KW} = 19.36 \text{ KW}$$

Hence, the chosen 22oo25 redundancy would be adequate to power cavities at the end of the single spoke section. A 22oo24 redundancy would also deliver sufficient power (20.17 KW) but with only two tolerated faults, the reliability requirement would not be met if the pallet has a failure rate of 1e-5 failures per hour.

For the likely scenario of a pallet failure rate of 4e-6 per hour, a 21oo23 redundancy would both fit the output power requirement with delivering 19.17 KW from 21 active pallets and the reliability requirement with a resulting failure rate of 0.9e-7 per hour.

For the optimistic scenario of a pallet failure rate of 1e-6 per hour, even a 20oo21 redundancy would both fit the output power requirement with delivering 19.05 KW from 20 active pallets and the reliability requirement with a resulting failure rate of 1.51e-7 per hour. One can see that between the pessimistic and the optimistic scenario, the redundancy margin can be reduced from 35% (25 KW installed when 18.54 KW required) to 13% (22 KW installed when 18.54 KW required).

One can see that the optimal number of pallets can be determined for each cavity. However, it has to be considered that the number of RF pallets per SSM is fixed and depends on the chosen supplier. E.g., if the suppliers only offer six pallets per SSM the number of pallets per cavity have to be multiples of six. For the 22oo25 redundancy, this means that actually 30 pallets will have to be installed as 24 would not be sufficient. This is referred to as grouping effect and it further inflates the margins.

7.5.3 Results

In this section, the procedure of determining the required number of pallets per cavity to conform to the reliability and input power requirements is carried out for the entire MYRRHA main LINAC for the pessimistic, likely, and optimistic pallet failure rate assumption. This will be carried out for the MYRRHA reference configuration of the accelerator with implemented dynamic compensation (four compensators per faulty cavity; Top-Down model as in Figure 6) as the power requirements from beam dynamics simulations are available for it (Bouly, 2022). Note that the grouping effect will not be taken into account as it depends on the chosen supplier. Hence the minimum number of required pallets will be stated and then depending on the number of pallets per SSM for the chosen supplier, the minimum number has to be 'rounded up' to the next corresponding multiple of the number of pallets per SSM.

Figure 41 shows the resulting number of pallets (expressed in power; y axis) for each cavity (x axis) for the pessimistic scenario with pallet failure rate of $1e-5$ per hour (light green dots), the likely scenario with pallet failure rate of $4e-6$ per hour (grey dots), and the optimistic scenario with pallet failure rate of $1e-6$ per hour (yellow dots). One can see that in comparison to the dynamic compensation margin (darker green line) and the tuning losses margin (dark blue line) the margins for pallet redundancy are relatively small. Only in the pessimistic scenario are they comparable to the other margins.




Figure 41: Required RF power margins across the main LINAC.

Whereas the tuning losses require a margin of around 30% across the LINAC and the dynamic compensation on average around 45%, the redundancy requires average margins of 10% for the optimistic scenario, 16% for the likely scenario, and 24% for the pessimistic scenario. Taking into

account the grouping effect, these margins may grow slightly. It has to be emphasized that these redundancy margin are likely leading to the biggest single reliability improvement in comparison to the comparable accelerators studied in Sec. 4 since the tube-based RF powering failures are essentially removed.

To study the potential of hot-swapping, the calculations were repeated for the likely scenario with a mean corrective time of three hours for hot-swapping faulty RF pallets. The resulting installed power requirements is shown in Figure 42. The dark green line shows the estimated pallet counts (again expressed in power) for hot-swapping. The other lines are the same as in the previous plot. The average margin across the LINAC is reduced from 16% to only 8% when implementing hot-swapping. This would result in a significant cost saving. However, this has to be balanced against the additional cost of developing reliable hot-swapping and producing SSMs with only a single RF pallet, which can be hot-swapped.




Figure 42: Required RF power margins across the main LINAC when hot-swapping of RF pallets is used.

Another relevant scenario is when dynamic compensation does not have to be implemented because the failure rates of the systems can be kept sufficiently low. In this case, the required failure rate for each RF powering reduces to $2.27\text{e-}7$ per hour. Repeating the calculations for the likely scenario with a pallet failure rate of $4\text{e-}6$ per hour that fulfills this requirement gives a result as shown in Figure 43.




Figure 43: Required RF power margins across the main LINAC when no dynamic compensation is used.

The orange line shows the scenario without dynamic compensation and the grey line the previous reference situation with dynamic compensation. Notably, the power requirements without dynamic compensation is much lower. This is because the increase in pallet redundancy margin is much smaller than the decrease due to the lack of dynamic compensation margin. Concretely, the pallet redundancy margin goes from 16% to 23% when moving from a required failure rate of $7e-7$ to $2.27e-7$ per hour. This mere difference of 7% is outweighed by the missing dynamic compensation margin of 45%.

Based on this one could conclude that dynamic compensation should not be implemented. However, it is unlikely that the even more stringent required failure rates for all accelerator systems can be met if no dynamic compensation is implemented.

7.5.4 Comparison with SPS SSPAs

When comparing the margins to those of the SPS SSPAs analysed in Sec. 7.2, one notices that the margins are bigger. This is due to two factors:

- The use of 3 dB combiners for MYRRHA, which have significant losses when some of the input amplifiers are faulty, instead of cavity combiners, which do not suffer from this.
- The bigger ratio of pallet power to cavity power requirement leads to a disadvantageous quantization effect. This means that a power requirement change from 9.9 kW to 10.1 kW leads to a change from 10 to 11 pallets required. Or in other words a 2% input power requirement change leads to a 10% jump in powering requirement. This effect will be further exacerbated by the grouping effect. E.g. for the mentioned example the jump would be from 10 to 15 kW installed power if there are 5 pallets per SSM.

7.5.5 Conclusions and Discussion

The proposed strategy of using modular and redundant solid-state RF amplifiers instead of tube-based technologies can be recommended from a reliability perspective. Tube-based RF amplifiers historically contribute the majority of failures in linear accelerators. These failures can be avoided by using solid-state amplifiers with a redundancy margin of 15-32%. To determine a more precise margin it will be necessary to narrow the range of the expected failure rate of an RF pallet, which should be possible with experimental data from MINERVA.

The key characteristics of this fault-tolerance scheme are listed below:

Cost/margin: 10-25%

Benefit: Single biggest reduction of failures in comparison to existing linear accelerators with tube-amplifiers. As for SPS SSPAs, it has been shown that failures are essentially removed from operation.

Key factors: SSPA technology, 3db combiner, less frequent access for amplifier replacement make margins bigger than for SPS SSPAs, cavity power to single pallet power ratio

Risks/unknowns: low, as comparable systems are operational.

Alternatives: Hot-swapping, cavity combiner.

7.6 Quadrupole circuits

In the main LINAC there will be 59 quadrupole doublet circuits, which require less than 1 failure per 90 days overall. For each doublet circuit this means 0.017 faults per 90 days or an MTBF of 130000 hours. One should keep in mind that this is the worst case leading to 10 faults per 90 days and that the reliability target should be set higher.

The first question is whether existing, comparable power converters reach MTBFs that are compatible with this requirement. To answer this, typical field-MTBFs of CERN power converters were investigated and are shown in Figure 44. Power converters with currents between 60 A and 220 A and voltages between 20 V and 500 V were selected to try to have comparable systems and filter out custom designs for CERN. It was not possible to filter exclusively for types of converters used for powering quadrupoles. The MTBF is calculated by counting the number of non-external (e.g. operator putting the wrong settings) power converter faults during the duration of the operational periods. However, the quality of failure data is varying and the operational periods might not be tracked to full detail as the data was collected by and for the equipment group and not the accelerator fault tracker tool, where this is tracked more accurately. Nevertheless it should provide a good MTBF estimation.




Figure 44: CERN magnet power converter observed MTBFs as a function of the quantity per converter type operated by CERN.

One can see that the requirement of 130000 hours is only satisfied by two out of the 31 plotted converters. Moreover, this only represents faults from the converter part and the controls, cabling, connectors and magnets have not been included yet. Hence, it becomes evident that a standard-solution will not be sufficient.

As solely relying on careful selection of systems with high MTBF is risky, it is necessary to foresee a possible use of a fault-tolerance scheme in case the equipment does not perform as expected. Different possibilities exist:

- Use of a dynamic compensation scheme similarly as it is done for the RF cavities.
- Making the power converters redundant by
 - Modular redundancy as for the RF amplifiers
 - A switching matrix with spare converters that get re-connected to the electrical circuit whose original converter failed.
- A combination of the above.

The use of a dynamic compensation scheme would be novel and require extensive development and testing. However, in terms of hardware the current quadrupole design with independent powering seems to be in principle ready for it. It should be studied whether the beam optics also permit such a scheme. The main advantage of this scheme is that it can also cover faults in the electrical circuit which are outside the power converter; e.g. magnet cooling failures.

Redundant power converters have been successfully realised in other facilities or are foreseen to be implemented. Hence, the implementation risk seems to be lower than for dynamic compensation schemes. Obviously it can only cover failures of the power converters themselves. Hence, these solutions can only be chosen if it is evident that doublet circuit failures outside the power converters are below one fault in 90 days.

The choice between modular redundancy and the switching matrix is partially governed by the rated power of the converters because at the order of tens of Amperes it will not be cost-effective to split a converter in multiple sub-modules. This seems to be the case for the MYRRHA quadrupole circuits.

Instead it is more effective to realize redundancy through a switching matrix, which can flexibly connect m quadrupole circuits with $n=m+q$ converters. The idea is that in normal operation m quadrupole circuits are connected to m converters. As soon as converter i fails, the switching matrix

will re-connect the magnetic circuit i to one of the q spare converters. This is effectively an m on n redundancy. The choice of m has to be optimized taking into account the expected converter MTBF and the choice of n will depend on physical limitations regarding the cabling length and the number of converters connected to a single switching matrix. As for the other fault-tolerance schemes, the switching has to function reliably and in under three seconds.

Based on this discussion it is recommended to foresee the possibility to install switching matrices and spare converters. In addition to that, dynamic compensation for quadrupole circuits should be tested with MINERVA so that it could be deployed in case it is required due to an higher than expected failure rate of magnets and connectors, which would not be covered by a switching matrix.

8 Implications for MYRRHA Availability

Based on the results of the investigations in previous sections, a recommended course of action in view of the MYRRHA availability goals is presented in this section. It is structured in recommendations, based on the current state of knowledge in Sec. 8.1 and longer-term recommendations in Sec. 8.2.

8.1 Recommendations based on Current Results

Despite the remaining uncertainties due to limited information about MYRRHA systems and their expected reliability, some recommendations can already be made.

In terms of the general feasibility of the MYRRHA reliability goals it must be said that given the available data of comparable accelerators, the MYRRHA reliability goals seem extremely ambitious even when the proposed fault-tolerance strategies can be implemented and perform well. This is best explained by the fact that at least 99% of faults of the comparable LINAC4 and SNS accelerators need to be removed to be in line with the MYRRHA reliability goals. Hence, besides focusing on improving the reliability of the MYRRHA accelerator, the discussion on the reliability requirements should be continued and possibilities to relax them should be studied. In particular it should be studied how the beam should ideally be ramped up after a beam trip to reduce the thermal stress inside the reactor. This may allow beam trips longer than the currently estimated three seconds.

Regarding the choice of reliability improvement methods to reach the MYRRHA reliability goals, the following can be stated:

- The foreseen use of modular-redundant RF power amplifiers can clearly be recommended. They are expected to yield the single most significant reliability improvement over existing, comparable linear accelerators with tube-amplifiers at a moderate additional cost.⁷ The same applies to the modular-redundant PSUs inside the RF power amplifiers.
- The proposed duplication of the low-beam-energy injectors can be a very effective fault-tolerance strategy if it can be ensured that the switching of beam production between injectors is reliable. However, doubling the injector means a doubling of the costs. If the sources of its faults are concentrated on certain sub-systems it may be more cost-effective to re-design the affected sub-systems to reach the desired reliability. Hence, it can only be recommended once it becomes clear that doubling the entire injector will be the most cost-effective way to reach the target failure rates. This will probably only be possible to determine with experimental data from MINERVA.
- The proposed dynamic compensation scheme, representing a serial redundancy of RF circuits in the main LINAC section, can also be a very effective fault-tolerance strategy if the switching to degraded configurations with detuned faulty RF circuits works reliably and singular failures are reduced. It also requires significant additional investments especially in terms of RF power margin, although not a doubling of costs as for the redundant injectors. The alternative of not using dynamic compensation and improving the reliability of the RF circuits is possible, in principle, but extremely challenging due to the stringent reliability requirements. Hence, it is recommended to continue to study and optimise the scheme despite the complexity and additional costs. A conclusive recommendation on its implementation will be possible with MINERVA.
- Regarding other fault-tolerances, which are not explicitly treated in previous sections, but shortly discussed in the appendices, the following can be stated based on experience with similar systems in other accelerator infrastructures:
 - A fault-tolerance should be foreseen for the electrical power converters of the main LINAC quadrupole doublets. The most promising approach is to use a switching matrix,

⁷ This ignores operational cost, which may increase to an often cited lower efficiency of solid-state based amplifiers. The financial impact of this should be investigated.

which allows to reconnect spare converters to a magnetic circuit whose original converter failed. This cannot cover failures of the magnet or the electrical connections of it, but covers faults of power converters which are expected to be much more frequent. A serial redundancy of magnetic quadrupole circuits (as is done in the dynamic compensation for the RF circuits) would additionally cover failures of magnets and their circuitry. It would be good to test the feasibility of this with MINERVA as it should not require additional hardware, but mostly modifications of the controls systems.

- To make the accelerator operation robust against electrical network glitches, the required power converter standards should be identified based on glitch data of the local electrical network in Mol and applied to all power converters with an additional margin to account for a possible increase of extreme weather conditions in the future decades.
- Besides the mentioned fault-tolerances, high reliability has to be aimed for across all accelerator systems to reach the ambitious reliability goals.

To improve the precision of recommendations, following actions are recommended, which can be carried out immediately:

- The top-down requirements models can be circulated among the people in charge of the various accelerator sub-systems to make them aware of these requirements and ask them for feedback on the feasibility. Where necessary, the top-down models should be updated.
- Similarly, communication channels between system experts and the AvailSim4 reliability model owners should be established to ensure that the models can be updated and system experts can be informed about possible reliability impacts of certain changes of system specifications.
- For the comparison with similar linear accelerators LINAC4 and SNS were chosen. Although they are similar in design, the reliability was not as much a top-priority for these accelerators as is the case for MYRRHA. Instead they excel in terms of certain performance measures. As a result, a considerable amount of their faults may come from the fact that they are pushing performance limits. It is difficult to filter out these types of faults from the available data. Hence, the actual fraction of faults that are not expected to occur in MYRRHA due to its more conservative design is unknown. Therefore, it is recommended to also study the reliability of possibly less comparable accelerators by spec-sheet, which focus on high reliability. E.g. this might be the case for light sources.
- As soon as the designs of certain key systems are more concrete, a more detailed reliability analysis should be carried out. E.g., for the Cold-Tuning-System or a system component analysis using standardised fault data for the main LINAC LLRF systems.
- As far as possible, the reliability of similar systems in other accelerator facilities should be experimentally measured and lessons learned should be applied to MYRRHA systems.
- The development of automatic fault-tracking solutions to track very short beam trips reliably should be continued and ideally improved in real-world scenarios, such as test-stands or other accelerator facilities.

8.2 Recommended Studies based on MINERVA

MINERVA was conceived as a first step in the phased implementation of MYRRHA and to demonstrate the feasibility of the MYRRHA accelerator and its reliability goals. As such, MINERVA has to demonstrate mainly two things from a reliability point of view:

- The overall accelerator failure rate and its sub-systems failure rates, as a basis to estimate the expected reliability of MYRRHA.
- The feasibility and maturity of the fault-tolerance strategies and whether they will allow to reach the MYRRHA availability goals

It is important to note that the data of the former point will only be available a few years after the first beam was produced in MINERVA. This is because valid reliability statistics are only available after a good fraction of teething problems has been overcome and sufficient beam time has been accumulated. The data for the latter point will only be available after the systems to test the fault-tolerance strategies have been deployed and dedicated test-time has been allocated.

In the following the collection of these data will be discussed in more detail for each fault-tolerance scheme. Thereafter, the extrapolation from MINERVA reliability data to a MYRRHA reliability estimate is addressed.

8.2.1 Estimation of system failure rates

MINERVA has to demonstrate that the system failure rates are small enough to reach the MYRRHA reliability goals. For this the MINERVA top-down requirements models were introduced in Sec. 0, which give the requirement for MINERVA sub-systems to be compatible with MYRRHA.

One question is how long MINERVA needs to be run so that sufficient run-time has been accumulated to demonstrate the required failure rate. This generally depends on the target failure rate, the number of failures observed during the observation period, and the number of systems under observation. Generally, the accumulated run-time for a system, i.e. the time the systems are under observation times the number of systems, has to be larger than the target failure rate which needs to be demonstrated because of failures that occur during the test and because one wants a certain degree of confidence in the estimate.

These examples should illustrate that the actual time it will take to demonstrate certain reliability targets may vary significantly between systems. However, once MINERVA has reached a certain maturity, it should be possible within months of operation to get a good idea of the failure rate of MINERVA and its main sub-systems. It will be more difficult for (sub-sub-) systems which have very low failure rates and a small system count; e.g. a specific power supply of the injector source. However, for the estimation of the MYRRHA failure rate it is sufficient to know the failure rates of the main systems without detailed knowledge of all sub-system failure rates. An exception to this are certain key elements of the fault-tolerance schemes, which are discussed later in this section.

Table 5 gives an overview of the accumulated run-times for different MINERVA sub-systems after one month or one year of operation. Looking at the infrastructure and services, their target failure rate in the MINERVA requirements model in Figure 8 is 2.62 per 90 days. If MINERVA was run for one month with two infrastructure and services faults, its maximum-likelihood estimate of the failure rate is two per 90 days, which would be in line with the requirements. However, the 60% confidence estimate would only be 3.1 failures per 90 days since very little runtime has been accumulated. If MINERVA was run for a year with 8 failures, the ML estimate would be 2 failures per 90 days again but the 60% confidence estimate would be 2.3 failures, which is much closer to the ML estimate since more experience had been accumulated. Note that 60% confidence is still very low though, so when the observed failure rate is very close to the target failure rates, the required test time may be longer than a year.

For the 62 main LINAC SRF circuits the target failure rate is 4 faults per 90 days or 0.065 failures per RF circuit. After a fuel cycle of 90 days with one failure for all circuits, the ML estimate is 0.016 failures per RF circuit per 90 days and even the 90% confidence interval would be 0.063 failures. This is a case where there are many systems under test and the observed failure rate is significantly lower than the target failure rate. In this case, less test time is needed to demonstrate with high confidence that the reliability targets can be met.

These examples should illustrate that the actual time it will take to demonstrate certain reliability targets may vary significantly between systems. However, once MINERVA has reached a certain maturity, it should be possible within months of operation to get a good idea of the failure rate of MINERVA and its main sub-systems. It will be more difficult for (sub-sub-) systems which have very low failure rates and a small system count; e.g. a specific power supply of the injector source. However, for the estimation of the MYRRHA failure rate it is sufficient to know the failure rates of the main systems without detailed knowledge of all sub-system failure rates. An exception to this are certain key elements of the fault-tolerance schemes, which are discussed later in this section.

Table 5: Overview of the accumulated run-times for different MINERVA sub-systems after one month or one year of operation.

MINERVA System	Accumulated after 1 month of beam [h]	Accumulated after 1 year of beam [h]	MINERVA target failure rate [h]
Infrastructure & Services (x 1)	720	8640	820
Main LINAC SRF circuit (x 62)	44640	535680	33480
Main LINAC Cryo module (x 31)	22320	267840	267840
Main LINAC Quadrupole doublet (x 32)	23040	276480	26784
Dynamic compensation (x 1)	Not hours	Not hours	Not hours
Injector (x 1)	720	8640	393

An important aspect to consider is that MINERVA, as all new accelerators, will take a considerable time to transition from a very high failure rate after first beams to a state of maturity with a low failure rate, which will be representative for future operations. This can be addressed by the use of reliability growth tracking (i.e. plotting the failure rate over time and extrapolating into the future). It may help to estimate the failure rates even before all teething problems have been overcome.

8.2.2 General estimation of fault-tolerance effectiveness, singular failures and other unforeseen events

Generally, for the estimation of the effectiveness of fault-compensation schemes the following factors need to be considered:

- Observed vs. required failure rate (i.e. whether there is a need for fault-tolerance)
- Covered systems by the fault-tolerance
- Covered systems failure rate
- Covered systems failure modes
- Demonstrated/estimated fraction of faults covered by fault-tolerance (of covered systems)
- Cost/risk of implementation
- Estimated improvement (i.e. how much closer to reliability targets) after implementation

Special attention to fault tracking will be necessary so that teething problems, problems not relevant for MYRRHA, those that will be covered fault-tolerances and especially those that will not be covered by fault-tolerances, are appropriately marked. The developed fault tracking templates (Appendix 11.2) have respective fields where this can be recorded in sufficient detail. This will also allow for a more accurate estimation of the effectiveness of future fault-tolerance strategies and the expected MYRRHA failure rate. In the following, the estimation of fault-tolerance effectiveness are discussed in more detail for each scheme.

8.2.3 Redundant Injectors

Besides the sub-system failure rates, the main factors determining the failure rate of the redundant injector fault-tolerance are the MTTR of an injector and the reliable switching from one injector to the spare injector.

The estimation of the MTTR can be made using expert interviews and needs to be complemented by experimental data with MINERVA.

Determining the reliability of switching between injectors will be challenging since MINERVA will only be equipped with one injector. The main concern of the switching is whether the beam can be reproduced by the spare injector so that it can be sent almost instantly (but certainly in under three seconds) to the main LINAC, especially when the spare injector has not delivered beam for extended periods of time. To test this with a single injector it may be necessary to switch off the injector for a longer period of time and check upon its restart whether the beam can still be sent to the main LINAC without applying lengthy iterative corrections. This would mean that once an injector is set-up for beam delivery its beam parameters will not drift even when it is not used for longer periods of time.

Another possible failure scenario is that the beam produced by the injector causes problems further down in the main LINAC without it being observable by the beam diagnostics of the injectors. Such a failure mode would have to be identified and correctly attributed to a faulty injector, otherwise the switching to the spare injector would not be initiated, which renders the fault-tolerance useless. Any signature of such behaviour has to be looked out for and addressed by sufficient beam diagnostics.

If the experience with MINERVA shows that the injector has a failure rate of less than the requirements model for non-redundant injectors (i.e. two faults per 90 days with the current models) or if it is evident that the re-design of certain sub-systems reduces the failure rate below the requirement, then a duplication of the injectors may not be necessary.

8.2.4 Main LINAC Dynamic Compensation

For the dynamic compensation it will be important to test that the reference configuration with four compensating cavities works reliably. That means that the switching between configurations in under three seconds works consistently, that as few as possible singular failure modes exist, and that the faulty configurations that are considered compensatable by beam dynamics simulations are actually compensatable in practice.

To accelerate this it will be necessary to be able to 'inject' or simulate certain failure conditions and verify the correct reaction of the system. This may reveal certain failure modes that cannot be covered, which would reduce the coverage of the fault-tolerance. Detailed recordings of this are necessary to update the dynamic compensation AvailSim4 model, which allows to estimate the reliability impact.

If (or once) the dynamic compensation works reliably in the reference configuration, it will be interesting to test its limits. There may be large accumulations of faults in certain areas, which are not considered compensatable by beam dynamics simulations, or the usage of more than four compensators, which is expected to pose challenges for the reliable switching in under three seconds. This would possibly allow for a scheme with more compensators to minimize the required RF power margins or a dynamic scheme depending on the available power margin.

As for the redundant injectors, if the experience with MINERVA shows that the main LINAC has a failure rate of less than the requirements model without dynamic compensation (i.e. 2.74 faults per 90 days) or if it is evident that the re-design of certain sub-systems reduces the failure rate below the requirement, then the implementation of dynamic compensation may not be necessary. The same applies if dynamic compensation cannot be successfully implemented or its additional complexity leads to a significantly increased failure rate of other sub-systems.

8.2.5 SSPAs

The key parameter for the dimensioning of the margin of the RF pallet redundancy is the failure rate of a pallet. This can be estimated from MINERVA operation as exemplified in Sec. 8.2.1 above. However, the question will be if the same failure rate can be assumed for the final 600 MeV extension of the accelerator. This is more likely if the same pallets from the same supplier will be used than if a new supplier is chosen. In case of choosing a new supplier it has to be seen whether running extensive tests to estimate the pallet failure rate or simply overdimensioning the margins is more cost-effective. Ideally, extensive testing and reliability runs are part of any contract, even if the same supplier is chosen for MINERVA and MYRRHA.

Besides the pallet margins it will be important to verify that the hot-swapping of the PSUs works reliably and that failures of RF pallets do not lead to a beam trip of more than three seconds. This is expected to be the case as it has been successfully implemented.

8.2.6 Glitch data

The collection of data on glitch duration and magnitude of the electrical network should be continued throughout MINERVA operation to see if any extreme weather trends become visible over the years on the electrical network and to follow-up if the power converters can withstand the observed glitches according to their specifications. A more detailed discussion is provided in the Appendix 11.4

8.2.7 Quadrupole switching matrix

For the fault-tolerance of the quadrupoles it will be important to test whether the switching can be carried out reliably without interrupting beam for more than three seconds. Moreover, collection of failure rate and repair time data will allow for possible modifications of the number of spare converters and the number of switching matrices.

8.2.8 Synthesis

Table 6 summarizes the key quantities that should be determined during MINERVA for each of the main fault-tolerance schemes to estimate their effectiveness and usefulness.

Table 6: Summary of the key quantities that should be determined during MINERVA for each of the main fault-tolerance schemes to estimate their effectiveness and usefulness.

Strategy	Redundant injectors	Modular-redundant RF amplifiers	Dynamic compensation	Quad-switching-matrix
Cost/margin relative to non-fault-tolerant solution	100% (doubling of injector)	25% power margin	45% power margin	~20%
Possible benefit relative to non-fault tolerant solution	Order of magnitude	1-2 orders of magnitude	Order of magnitude if switching is reliable and singular points of failure reduced	Order of magnitude if switching is reliable and singular points of failure reduced
Covered systems	Injector	RF powering	Main LINAC RF circuits and cryomodules	Main LINAC Quadrupole power converters
Estimated fraction of accelerator faults aimed at	~10%	~30%	~25%	~20%
Key factors for effectiveness	Reproducibility of beam by other injector Successful switching and ramping under three seconds Injector failure rate Time to repair	Power margin RF Module MTBF Time to repair (accessibility of repair)	Switching success rate and covered failure configurations Power margin (number of compensating cavities) Time to repair (accessibility of repair) Observed failure modes, failure rates and their coverage	Switching success rate Singular points of failure Power converter MTBF How large are switching clusters and how big is the margin? (defines flexibility of compensation) Time to repair (accessibility) Observed failure modes and failure rates
Risk/Maturity	Test beam reproducibility – Needs to be tested with MINERVA	Mature technology Demonstrated for CERN SPS SSPAs	Needs extensive test campaigns with triggered faults Implemented with switching time of several minutes at SNS	Foreseen for implementation in ESRF EBL with more stringent switching time requirements Implemented at SNS with less stringent switching time requirements

8.2.9 Extrapolating from MINERVA to MYRRHA

Once the described data from MINERVA has been collected, it can be scaled to obtain the expected failure rate of MYRRHA. This requires a set of rules for extrapolation, which will be referred to as scaling laws. E.g. if MYRRHA uses three times more vacuum pumps (of similar type in similar conditions) than MINERVA, the failure rate of all vacuum pumps will be three times higher.

More generally, scaling laws can factor in system count, complexity, power/load/stress, technological maturity among other factors. However, accurate and empirically verified scaling laws exist for very few systems. The general exception is to scale by system count for (nearly) identical systems in (nearly) identical usage scenarios. This can be applied for the extrapolation from MINERVA to MYRRHA. Table 7 gives an overview of the proposed scaling law for the main systems.

Table 7: Proposed scaling laws for the extrapolation from MINERVA to MYRRHA.

Systems	Proposed scaling	Comment
RF systems	By number of cavities (x2.5)	Elliptic cavities more or less problematic than spoke?
Cryo modules and valve boxes	By number of cryostats (x2.5)	Cryostats for elliptic more or less problematic than spoke?
Quad doublet circuits	By number of circuits (x2.5)	
One Injector	Possible scaling if weak points redesigned	
Controls & timing system	unknown	Complexity expected to increase with introduction of fault-tolerance schemes
Cryo plant	~x 2-3 (will have a second plant)	Options of reliability improvement unclear
Beam vacuum pumps	By number of quad circuits (x2.5)	Assuming they are put next to the quads
Electrical network	Possible scaling if grid stability improving/worsening	
Cooling and Ventilation	Similar to cryo	Could scale with system count

Note that for the cryo plant there it is foreseen to build a second one for MYRRHA. The above scaling assumes that the first and second one will be both necessary to operate MYRRHA and they will not be redundant. If they are made redundant, the scaling above does not apply.

The scaling laws were already applied to the top-down models introduced in Section 5. So they should provide a good first reference for this extrapolation. The scaling laws should be refined once more details about the systems become available.

8.3 Knowledge-Transfer Activities

As described in the recommended course of action, the MYRRHA reliability activities will be most effective when the tools and models developed in this project are used throughout the phased implementation of MYRRHA. Since this was identified early in the project, the availability tools and models have been made available to the project participants.

For the open-source AvailSim4 tool, weekly practice sessions have been held from March to November 2022. In this sessions, the colleagues from SCK-CEN had hands-on practice with the tool and worked

on the availability simulation models for the fault-tolerance schemes introduced in Sec. 7.3 and 7.4. Commonly used gitlab repositories are available for the AvailSim4 tool and the simulation models.⁸

⁸ The tools and models are available on <https://gitlab.cern.ch/availsim4> .

9 Summary, Conclusions & Outlook

Building upon previous reliability studies, a more complete and refined reliability model of the MYRRHA accelerator was developed with an emphasis on a detailed understanding of fault-tolerance capabilities, which have previously only been studied conceptually.

Concretely, the following was covered:

- Global accelerator models for MINERVA and MYRRHA.
- Detailed availability models for three proposed fault-tolerance strategies: modular-redundant RF power amplifiers, redundant injectors, and serial redundancy of RF cavities in the main LINAC

To achieve the mentioned objectives, supplementary tasks were identified and addressed. An updated high-level availability comparison with existing accelerators has been done to derive top-down fault budget models for MYRRHA and MINERVA. Tools and features for the advanced availability simulations were developed. The availability simulation models were made available to SCK so that they can be continuously updated when new data arrives. To be able to do so, there has been a continuous transfer of know-how to SCK colleagues for using and developing the availability simulation tools and models. Automated fault tracking methods have been prototyped to keep track of short beam-trips relevant for MYRRHA operation.

An in-depth discussion was provided of how the developed models and tools shall be used in the future and in view of data from MINERVA. This allows to address the global MYRRHA reliability studies objectives in the future, which are to identify (1) if and to what extent the MYRRHA reliability goals can be met and (2) how to best define which measures to take to meet the reliability goals.

Based on the work performed, the following conclusions can be made:

- In terms of the general feasibility of the MYRRHA reliability goals it can be concluded that based the available data of comparable accelerators, the MYRRHA reliability goals seem extremely ambitious even when the proposed fault-tolerance strategies can be implemented and perform well. This is best explained by the fact that at least 99% of faults of the comparable LINAC4 and SNS accelerators need to be removed to be in line with the MYRRHA reliability goals. Hence, besides focusing on improving the reliability of the MYRRHA accelerator, the discussion on the reliability requirements should be continued and possibilities to relax them should be studied.
- The foreseen use of modular-redundant RF power amplifiers can clearly be recommended. They are expected to yield the single most significant reliability improvement over existing, comparable linear accelerators with tube-amplifiers at a moderate additional cost. The same applies to the modular-redundant PSUs inside the RF power amplifiers. Cavity combiners can achieve better reliability improvement than 3dB combiners when the same RF power margins are used. However, they require a sufficient number of RF transistors as input to be effective and fault-tolerant, which is likely not the case for MYRRHA amplifiers.
- The proposed duplication of the low-beam-energy injectors and the dynamic compensation scheme for RF cavities can be a very effective fault-tolerance strategy if it can be ensured that the switching to compensating configurations in case of a failure is working reliably. However, implementation of the schemes represent a significant cost. If the sources of faults are concentrated on certain sub-systems it may be more effective to re-design or duplicate only the affected sub-systems to reach the desired reliability. Still the proposed scheme can be recommended once it becomes clear that it will be the most cost-effective way to reach the target failure rates. This will only be possible following the operational experience with MINERVA.

- MINERVA will provide sufficient data to estimate the expected failure rates of the main accelerator systems within several months of stable operation. Dedicated test-programmes are required to assess key factors influencing the effectiveness of fault-tolerance methods, which are listed in Table 6. Both failure rates and fault-tolerance effectiveness will allow to judge whether the MYRRHA accelerator is expected to meet its availability goals.

To gain confidence in the designs for MYRRHA as well as the models to predict the accelerator reliability, following actions are recommended, which can be carried out immediately:

- The top-down requirements models and the AvailSim4 simulation models should be continuously used and further developed at SCK. People responsible for the various accelerator sub-systems should be made aware of the reliability requirements, which should be part of the technical specifications of systems. Communication channels should be set up between systems experts and reliability model experts to ensure that the models are used to assess any changes the systems undergo or to incorporate more precise information about them.
- As soon as the designs of certain key systems are more concrete, a more detailed reliability analysis should be carried out. E.g., for the Cold-Tuning-System or the main LINAC LLRF systems.
- As far as possible, the reliability of similar systems in other accelerator facilities should be analysed and lessons learned should be applied to MYRRHA systems.
- The development of automatic fault-tracking solutions to track very short beam trips reliably should be continued and ideally improved in real-world scenarios, such as test-stands, other accelerator facilities and MINERVA in particular.

The performed studies and deduced recommendations represent a significant contribution towards reaching the MYRRHA reliability goals. Nevertheless, reaching the MYRRHA availability goals will require continued reliability efforts and attention across all accelerator systems, all the organisational aspects, and all implementation stages leading up to the finalization of MYRRHA.

10 Bibliography

- AFMC. (1993). *Rome Laboratory Reliability Engineer's toolkit*. Griffiss AFB, NY: Systems Reliability Division Rome Laboratory, Air Force Material Command (AFMC).
- Apollonio, A., JM, A., Pitigoi, A. E., Fernandez Ramos, P., & Uythoven, J. (2015). *H2020-MYRTE D2.9: LINAC4 reliability analysis for MYRRHA*.
- Bouly, F. (2022, February). *Multiple failure compensation*.
[https://extranet.sckcen.be/sites/patriciapartnerarea/WP10/Task%2010_1%20Accelerator%20Reliability/Weekly%20Meetings/2022_05_17_47th_weekly_meeting_materials/MYRRHA__power_localcompensation_2022\(1\).xlsx](https://extranet.sckcen.be/sites/patriciapartnerarea/WP10/Task%2010_1%20Accelerator%20Reliability/Weekly%20Meetings/2022_05_17_47th_weekly_meeting_materials/MYRRHA__power_localcompensation_2022(1).xlsx).
- Burgazzi, L., & Pierini, P. (2007). Reliability studies of a high-power proton accelerator for accelerator-driven system applications for nuclear waste transmutation. *Reliability Engineering & System Safety*, 449-463.
- Charlot, J. (2018). *Etude MTBF système amplificateur de puissance pour le CERN*. Thales SIX GTS France SAS - 4, avenue des Louvresses - 92230 Gennevilliers.
- DoD. (1998). *MIL-HDBK-338, DoD*. Department of Defense, USA.
- Felsberger, L., Apollonio, A., Cartier-Michaud, T., Montesinos, E., Oliveira, J. C., & Uythoven, J. (2021). Availability Modeling of the Solid-State Power Amplifiers for the CERN SPS RF Upgrade. *JACoW IPAC 2021*, (pp. 2308-2311).
- Jacob, J., Farvacque, L., Mercier, J., Langlois, M., & Gautier, G. (2013). Commissioning of first 352.2 MHz--150 kW solid state amplifiers at the ESRF and status of R&D. *IPAC2013*, (pp. 2708-2710). Shanghai.
- Jacob, J., Mercier, J. M., Langlois, M., & Gautier, G. (2011). 352.2 MHz--150 kW solid state amplifiers at the ESRF. *IPAC*, (pp. 71-73). San Sebastian.
- Marchand, P., Baete, J., Cuoq, R., Dias, H., Diop, M., Labelle, J., . . . Petit, S. (2013). Operational Experience with the SOLEIL Superconducting RF System. *SRF2013*. Paris.
- Marchand, P., Lopes, R., Ribeiro, F., & Ruan, T. (2011). Development of high RF power solid state amplifiers at SOLEIL. *IPAC2011*, (pp. 376-378). San Sebastian.
- Montesinos, E. (2017). *RF power & LLRF*. Barcelona: LLRF2017.
- Pitigoi, A., & Ramos, P. (2014). *MAX Deliverable D4.4: Reliability model of the MYRRHA Linac*.
- Roderick, C., Burdzanowski, L., Anido, D. M., Pade, S., & Wilk, P. (2017). Accelerator fault tracking at CERN. *Int. Conf. on Accelerator and Large Experimental Control Systems*, (pp. 397-400). Barcelona.
- Schmidt, R., Riedel, G. J., & Kangas, K. (2011). Risk assessment using design review based on failure mode. *RAMS*, (pp. 1-6).

11 Annexes

11.1 System definitions for top down models

The system (boundary) definitions for all the systems listed in the top-down requirements models from Section 5 are given below.

- MYRRHA accelerator: The entire 600 MeV accelerator including all infrastructures and services up to the beam window of the reactor.
- Injectors: The injector(s) including the injector dumps, the beam lines leading to the main LINAC and the injector switching magnet.
 - Switching magnet: The injector switching magnet circuit including its power converter, electrical connections.
 - Injector switching: Not a system but the fault-tolerance scheme for the injectors.
 - Injector ½: The physical injector
 - Source: The proton source assembly
 - LEBT: The Low-Energy-Beam-Transfer line including all its beam-line elements
 - RFQ: The Radio-Frequency-Quadrupole
 - MEBT: The Medium-Energy-Beam-Transfer line including all its beam-line elements
 - CH: The normal-conducting CH RF circuits
 - Vacuum: The injector vacuum pumps, valves and gauges at the beamline
- Main LINAC: The main LINAC starts after the injector switching magnet and ends at the beam delivery system (BDS), which starts after the 600 MeV dipole.
 - Dynamic compensation: Not a system but the fault-tolerance scheme for the main LINAC RF circuits and cryo-modules.
 - SRF: The main LINAC superconducting RF circuits
 - Powering: The RF amplifiers providing high-power RF fields.
 - Feeder, circulator, coupler: The systems delivering the high-power RF fields to the cavity.
 - Cavity & tuning: The RF cavity and its mechanical tuning system.
 - LLRF: The dedicated RF controls system.
 - SRF Cryo-modules: The cryostats and their valve boxes housing the superconducting RF cavities
 - Valve box: The interface between cryostat and cryogenics distribution system.
 - Cryostat: The system providing cryogenic environment for the superconducting RF cavities. May include a possibly active insulation pumping.
 - Non-redundant RF&Cryo: Essentially the foreseen re-buncher RF circuit and its cryo-module at 100 MeV.
 - Quad – Matrix/Compensation: Not a system but the fault-tolerance scheme for the main LINAC quadrupole doublet circuits.
 - Quad-Doublets: The quadrupole doublet circuits.
 - Converters: The power converters and their controls.
 - Cabling and connectors: The circuitry between converter and magnet.
 - Doublet magnet: The normal-conducting quadrupole magnet doublet.
 - Vacuum: The main LINAC vacuum pumps, valves and gauges at the beamline.
 - BI: The main LINAC beam instrumentation and diagnostics systems.

- Non-redundant kickers and magnets: The electrical circuits and magnets that divert the beam at 100 MeV and 600 MeV.
- BDS to reactor: The BDS starts after the 600 MeV dipole and ends at the beam window into the reactor as well as the 600 MeV beam dump. It includes a fast deflecting magnet to choose the destination of the 600 MeV beam.
- Infrastructure and Services: This groups all systems that are providing 'services' to both the injectors, the main LINAC and the BDS.
 - Controls and Timing: The controls and timing system which represent the interface between all accelerator system's controls and the operating personnel. It also includes the fault-tracking, machine protection, personnel protection and similar systems. It does not include controls dedicated to accelerator systems such as the LLRF (RF controls).
 - Electrical Network: The systems interfacing between the electrical grid and all the electrical power consuming beam-line elements.
 - Cooling plant and ventilation: The systems providing and distributing cooled water and the systems ensuring correct ambient temperatures for operation of the accelerator.
 - Cryo plant and distribution: The systems generating and distributing cryogenic cooling.
 - Pressurized air: The systems generating and distributing pressurized air.
 - IT services: The systems interfacing (mostly) with the controls and timing systems to provide the necessary computing and storage infrastructure.
- Operation and Maintenance: This is not a system but should represent errors due to operation and maintenance mistakes.

Generally, the boundaries between beam-line elements (cavities, magnets,...) and infrastructure and services (vacuum, cooling, ...) are put so that a failure of a beam line element leads to a failure of itself only, whereas a failure of an infrastructure or service leads to an outage of all beam-line elements connected to it. Hence, a valve box is part of a beam-line element (SRF Cryo-module in this case) whereas the cryo-plant is an infrastructure and service. For the distribution systems between the cryo-plant and the valve boxes it may not always be clear where the boundary should be drawn. E.g. a distribution valve may impact only one cryo-module if it is stuck in the closed position but it may impact the entire accelerator if it starts leaking. In this case, the failure modes impacting the entire accelerator would have to be counted towards the infrastructure and services and the failure mode impacting the cryo-module only towards the cryo-module. Clearly this is an edge case which should be the exception.

11.2 Fault tracking tables for LLN Injector Set-Up

A system breakdown table detailing all systems has been generated for the LLN injector fault tracking. A few rows are shown below. The entire table is on the sharepoint.

SBS ID	SBS Title	Sub-Section	Device instance	System	Sub-System	Component	Sub-Component
		<u>Ion Source</u>					
MYR-1459	SRC-SRCA-A		Ion Source Assembly				
				vacuum pumps			
					primary pump		
					turbo pump		
					angle valve		
						01	
						02	

						03	
					Pirani gauge		
						02	
						01	
						03	
					Cold cathode vacuum gauge - Penning		

To log sufficient detail of the operation of test-stands or MINERVA, a fault tracking template has been compiled. It should allow to know all necessary details about the faults to correct them and to understand their impact on the reliability of MYRRHA. Its fields are shown in Table 8.

Top category	Sub category	Field	Example/unit
Failure Reporting		Report ID:	[1,2,3,...]
		Sub-Section:	[use]
		Device instance:	[see system breakdown table]
		Component:	[see system breakdown table]
		Sub-Component	[see system breakdown table]
		Date and Time (of Failure)	yyyy-mm-dd hh:mm:ss
		Failure Description:	[free text]
	Expected future failure behaviour	MINERVA 100MeV (Phase1)	[yes/no]
		Minerva 600MeV (Phase 2)	[yes/no]
		MYRRHA+REACTOR (Phase 3)	[yes/no]
		comment:	[free text explaining why (not) failure expected to happen in MINERVA/MYRRHA operation]
Repair reporting		Repair action:	[free text, list replaced components (if any)]
		Duration OP was interrupted because of failure	[m]
	Expected future repair duration:	MINERVA 100MeV (Phase1)	[m]
		Minerva 600MeV (Phase 2)	[m]
		MYRRHA+REACTOR (Phase 3)	[m]
		comment:	[free text explaining expected differences for MINERVA/MYRRHA in terms of repair (duration)]
Failure Analysis		Previous occurrences (report ID)	[IDs of previous identical failures]
		Cause(s)	[free text]
		Parent Fault (report ID)	[ID of event causing this failure]
		Comments	[free text]
Corrective Action		Description	[free text]
		Status	New/in progress/closed

		Resolution status	[done, won't fix, duplicate, incomplet, cannot reproduce, invalid]

Table 8: Fault tracking template to capture details of operational problems.

It is important to point out that the faults should be classified in the same way as the top-down requirement models, so that it can be assessed whether the system failure rates are in line with the requirements.

In addition, a template for tracking operational times was generated, which is not shown here. Together with the fault tracking template it allows to estimate the necessary reliability statistics as it allows to infer the duration each systems was operational.

11.3 Asymmetric Injector Model

The current layout foresees a symmetric arrangement of the injectors, where both injectors need an injector dump magnet to deflect the beam from being dumped in the injector dumps to the switching magnet to steer the beam into the main LINAC. In case of the failure of the injector switching magnet, beam could not be delivered anymore even if both injectors still work.

The use of an asymmetric arrangement, in which one injector only needs one magnet to deflect the beam away from the dump and directly into the main LINAC and another injector would have both an injector dump magnet and the injector switching magnet, would have the advantage that the failure of the injector switching magnet would not impact beam production from one injector.

A quick availability comparison between the symmetric and the asymmetric model revealed that the benefit of the asymmetric model is very minor since the failure rate of the injector switching magnet is orders of magnitude smaller than the failure rate of the injectors. Hence, there is no recommendation to implement the asymmetric model from an availability perspective.

11.4 Glitch robustness

Electrical network glitches are a common reason for power converter trips leading to the interruption of particle accelerator operation. Hence, it was quantified how often this is expected to happen given the electrical network situation in Mol.

Electrical network glitches have been recorded at Mol from 2017 onwards. Figure 45 shows the recorded voltage dips. The y-axis shows the magnitude of the dip and the x axis shows the duration of the dip. Points above the red line indicate dips that would be tolerated by a power converter which conforms to the IEC 61000-4-11 standard. 5 out of 20 voltage dips would have led to an interruption of accelerator operations if all converters would be built to the standard. This would mean 0.56 triggered reactor turnarounds per fuel cycle, which would be in line with the top-down requirements models.




Figure 45: Voltage dips recorded at Mol site.

However, one needs to consider that most of the voltage dips probably accumulate during summer months when thunderstorms are frequent and that extreme weather conditions might intensify the frequency of voltage dips. Hence, it may pay off to apply even more stringent standards against voltage dips for power converters.