Availability of Conceptual Machines - How and Why?

A Case Study of the Future Circular Electron-Positron Collider

of

Hannah Alida Dostmann

from Mannheim

Supervisors: Martin Dazer, Jack Heron

Date of oral examination: 22 July 2025

Institut für Maschinenelemente, University of Stuttgart

2025

Acknowledgements

I would like to thank my supervisor, Jack Heron, for his continuous support and insightful guidance throughout this project, as well as the entire TE-MPE group for their collaboration and encouragement. I also extend my gratitude to all system experts at CERN who provided valuable input, and to the Institute of Machine Components at the University of Stuttgart, in particular Martin Dazer, for the support and guidance. Working on this thesis has been both a professional and personal learning experience, and I am grateful for the opportunity to contribute to a project of this scope and significance.

Abstract (english)

In large-scale, high-criticality technical systems, availability is a key performance driver and can offer significant advantages when analysed at an early stage. However, assessing availability during the conceptual design phase remains a major challenge, as traditional methods rely on mature system architectures and detailed component data—resources typically unavailable at early stages. This limitation hinders strategic design decisions and increases the risk of costly late-stage redesigns.

This thesis presents a generalised framework for availability forecasting in conceptual systems, where design details are incomplete. A Monte Carlo-based simulation environment is developed that estimates key performance indicators and decomposes them into subsystem contributions. This supports performance optimisation and helps to translate top-level goals into reliability requirements.

The methodology is demonstrated using the Future Circular Collider (FCC-ee) as a case study, illustrating its benefits and value in a real-world context. Beyond this application, the framework is designed to be domain-independent and adaptable to other complex systems, offering structured support for early-stage design and reliability-focused planning.

Abstract (german)

In hochkomplexen technischen Systemen ist die Verfügbarkeit ein zentraler Leistungsfaktor und kann bei frühzeitiger Analyse zu besonders großen Vorteilen führen. Die Bewertung der Verfügbarkeit in der konzeptionellen Entwurfsphase stellt jedoch eine besondere Herausforderung dar, da herkömmliche Methoden auf ausgereiften Systemarchitekturen und detaillierten Komponenteninformationen basieren – Ressourcen, die in frühen Entwicklungsstadien typischerweise nicht vorliegen. Diese Einschränkung erschwert strategische Designentscheidungen und erhöht das Risiko kostspieliger Neugestaltungen in späteren Phasen.

Diese Arbeit stellt ein generalisiertes Rahmenwerk zur Verfügbarkeitsprognose in konzeptionellen Systemen mit unvollständigem Design vor. Eine Monte-Carlo-gestützte Simulationumgebung wurde entwickelt, die zentrale Leistungskennzahlen schätzt und diese in Beiträge einzelner Subsysteme zerlegen kann. Dadurch kann das Design auf Leistung optimiert und Zuverlässigkeitsanforderungen aus übergeordneten Zielvorgaben abgeleitet werden.

Die Methodik ist am Beispiel des Future Circular Collider (FCC-ee) demonstriert und zeigt dabei ihre Vorteile in einem realen Anwendungskontext. Darüber hinaus ist das Rahmenwerk domänenunabhängig konzipiert und flexibel auf andere komplexe Systeme übertragbar, wodurch es ein strukturiertes Werkzeug zur Unterstützung von Entwurfsentscheidungen in frühen Phasen sowie zur planvollen, zuverlässigkeitsorientierten Systementwicklung bietet.

Table of Contents

List of Abbreviations

AFT	Accelerator Fault Tracking
CLIC	Compact Linear Collider
DC	Direct Current
DEMATEL	Decision-Making Trial and Evaluation Laboratory
DES	Discrete Event Simulation
FCC	Future Circular Collider
FMEA	Failure Modes and Effects Analysis
FOO	Feasibility of Objectives
GS	General Services
HE-Linac	high-energy linear accelerator
HV	high voltage
IP	interaction point
LCB	lower cold box
LCLS	Linac Coherent Light Source
LHC	Large Hadron Collider
LINAC4	Linear Accelerator 4
LLRF	low level radio frequency
LV	low voltage
MTBF	Mean Time Between Failure
MTTR	Mean Time To Repair
MV	medium voltage
MYRRHA	Multi-purpose hybrid Research Reactor for High-tech Applications

PMU	Phasor Measurement Units
PPBI	pre-polarised bunch injection
pre-TDR	pre-Technical Design Report
PS	Proton Synchrotron
PSB	roton Synchrotron Booster
QPS	quench protection system
QUI	interconnection unit
R&D	research and development
RDP	resonant depolarisation
RF	radio frequency
SLAC	Stanford Linear Accelerator Center
SPS	Super Proton Synchrotron
SS	Straight Sections
TIDVG	Target Internal Dump Vertical Graphite
UCB	upper cold box
UPFC	Unified Power Flow Controllers
WCS	warm compressor station

List of Symbols

Symbol	**Einheit**	**Beschreibung**
A		Overall availability
A_S		Subsystem availability
E		Efficiency
L		Nominal Instantaneous Luminosity
L_{int}		Integrated Luminosity
$MTBF$		Mean Time Between Failure
$MTTR$		Mean Time To Repair
n		Number of subsystems
n_f		Number of faults
N_{IP}		Number of interaction point
R		Overall reliability
R_S		Subsystem reliability
T		Annual Scheduled Physics Time
t		Time
τ_r		Intervention duration
τ_{up}		Up time duration

List of Figures

List of Tables

1 Introduction

In high-criticality technical systems, assessing availability at the early stages of development can deliver exceptional advantage. Availability analysis helps identify potential weaknesses and enables design adjustments, preventing costly and time-consuming issues during operation. Such adjustments typically involve measures to enhance reliability, such as the integration of redundancy, as well as optimisations related to maintenance and repair processes or operational efficiency [1]. The importance of doing this at an early development stage is well justified by Figure 1.1. During early development phases, the freedom to shape the system is greatest and strategic directions can still be easily adjusted. As the project progresses, some aspects become fixed and unchangeable; if changes are possible later on, they tend to be increasingly costly.

A well-known large-scale example that illustrates the consequences of insufficient early availability planning is the Channel Tunnel. In its initial years of operation, the system suffered from frequent disruptions, including train breakdowns, ventilation problems, and even fire incidents, which led to significant interruptions [2], [3]. These issues were largely attributed to the complexity of the interconnected subsystems and the lack of an integrated reliability and availability strategy during the design phase [4]. A thorough availability analysis could have revealed critical subsystems within the tunnel and failure scenarios, enabling the implementation of preventive measures and reducing the risk of such operational setbacks.

The Future Circular Collider (FCC), CERN's leading proposal for a next-generation energy-frontier particle accelerator, presents an even greater chal-

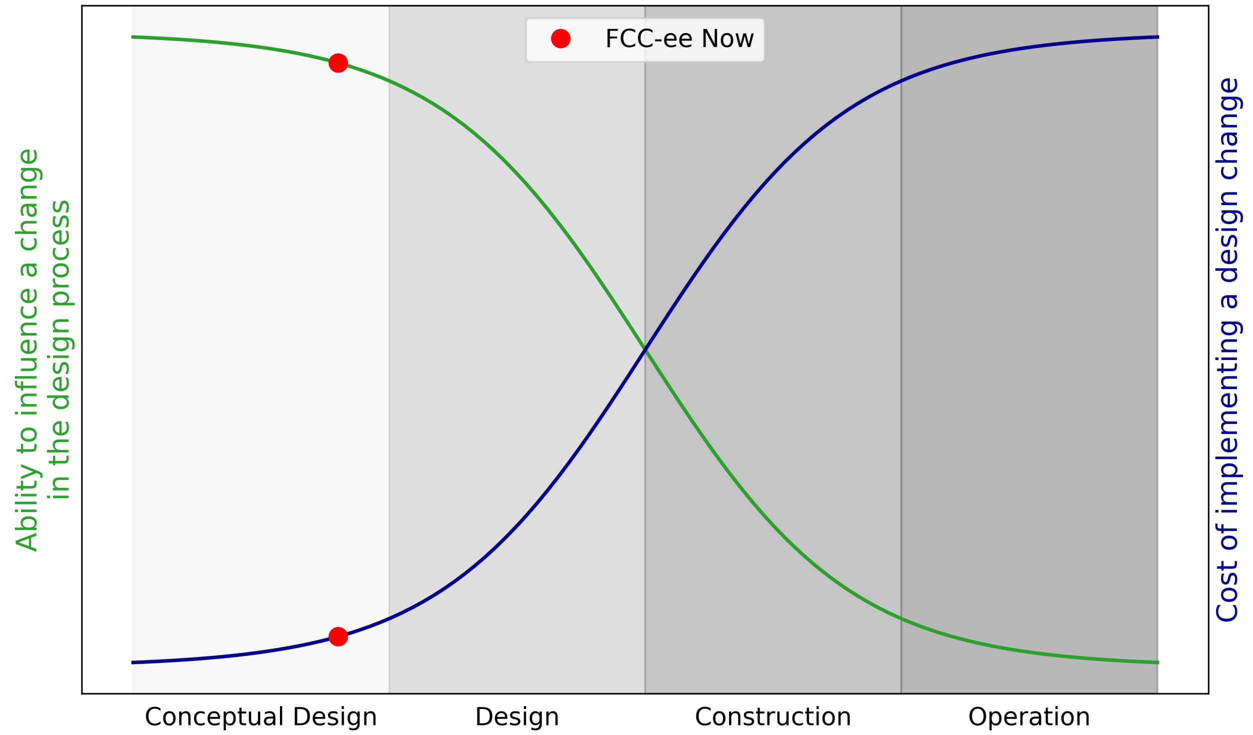

Figure 1.1: Graph illustrating that the ability to influence design changes is highest during the early project lifecycle (Conceptual Design), when implementation costs are low. In later phases (Construction, Operation), design flexibility decreases while the cost of change increases.

lenge in terms of complexity, scale, and therefore availability. It is planned to operate in two stages: first as a lepton collider (FCC-ee), followed by a hadron collider (FCC-hh). With a planned circumference of 91 km, the FCC will surpass all previous accelerators, not only in size but also in technical and organisational demands. This makes availability a critical risk factor for key performance parameters such as luminosity — the number of particle interactions — and operational cost.

The number of interactions per time (nominal instantaneous luminosity) varies with the energy level the accelerator is running at. To achieve its physics targets, the FCC-ee must reach a minimum of total interactions (integrated luminosity), which in turn depends on meeting a minimum avail-

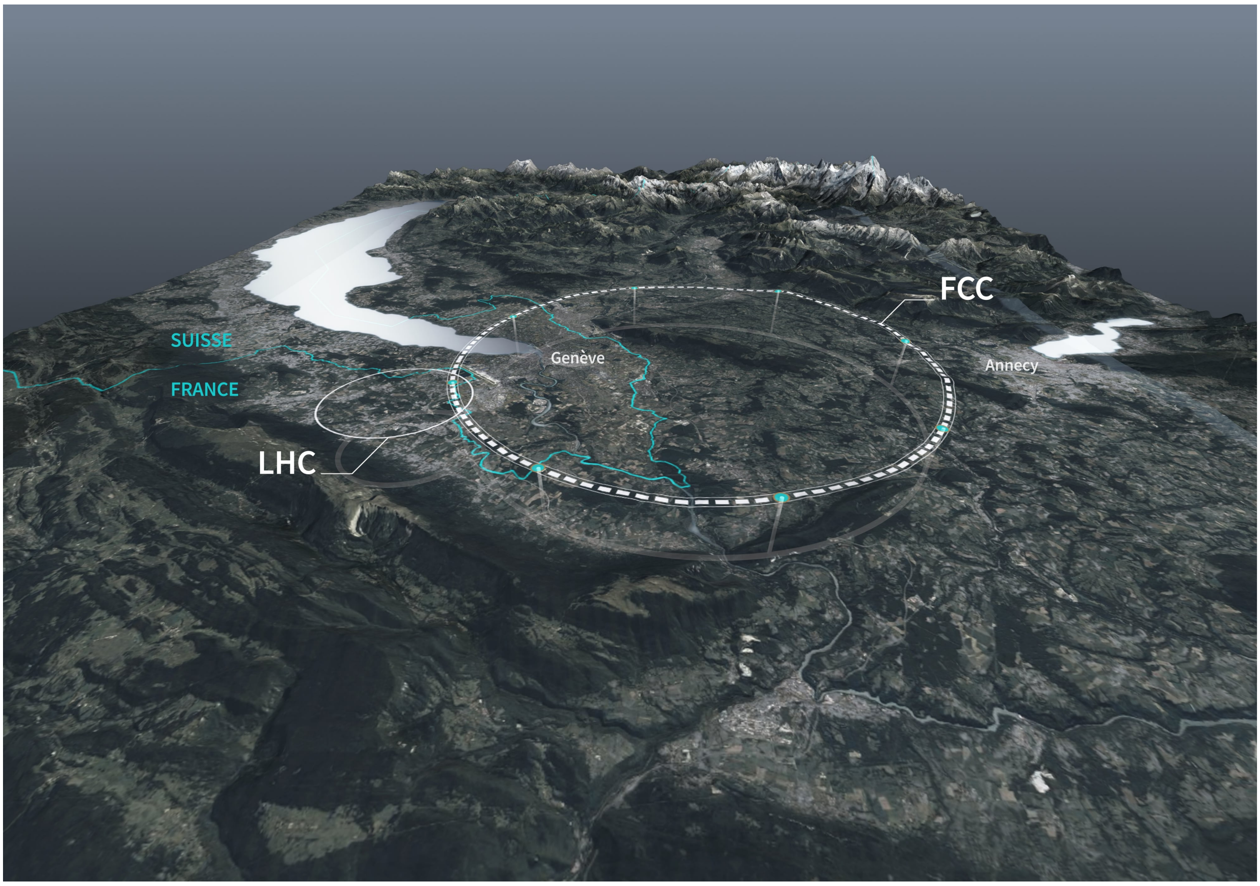

Figure 1.2: Image comparing the FCC with a circumference of 91 km to the Large Hadron Collider LHC with 27 km. The FCC will be approximately three times larger. [7]

ability threshold. Availability is the proportion of scheduled physics days where the machine can deliver beam. Previous studies have defined this requirement to be 80 % [5]. For comparison, the LHC — currently the world's largest particle accelerator, spanning only about one third of the FCC's size (see Figure 1.2) — achieved an availability of 72 % between 2015 and 2024 [6]. This highlights the importance of a systematic and early availability analysis for the FCC-ee, to assess feasibility and to guide and optimise the machine's design from from this early stage in the design process onward.

However, conducting availability analyses at such an early stage poses significant challenges. Traditional availability assessment methods, such as Petri Nets, Markov Models, Failure Modes and Effects Analysis (FMEA) etc.

[1], are typically applied during the design phase, when system architectures are already well defined. In contrast, the FCC-ee has only just reached the pre-Technical Design Report (pre-TDR) phase, meaning that many technical specifications are still under discussion. As a result, the necessary level of detail required by conventional methods does not exist yet. The challenge is further compounded by the system's exceptional complexity and the vast number of components spread across multiple subsystems. Such characteristics necessitate a structured methodology to estimate availability under conditions of limited design maturity and high uncertainty. This leads to the central question of this thesis: "How can availability be analysed for a system that has not yet been fully designed — and why is such an analysis crucial at this early stage?"

1.1 Research Contribution

This thesis addresses the availability challenge of conceptual machines by presenting an enhanced Monte Carlo simulation environment that deconstructs down time contributions into its main constituent systems. The specific contributions of the thesis are as follows:

1. A generalized framework for availability forecasting of conceptual systems that extrapolates fault data of existing operational systems is provided and showcased using the example of the FCC-ee.

2. A formulation to anticipate availability in systems with neither rigorous designs nor comparable failure data is given and applied to the FCC-ee beamstrahlung dump system.

3. Using the above formulations in a Monte Carlo simulation various key performance indicators for the baseline design of the FCC-ee are predicted. This has identified significant shortcomings in various key

performance parameters that are well below design targets and highlights the benefits of an availability analysis at such an early stage of development.

4. Various solutions to the above shortcomings are proposed as research and development (R&D) opportunities, whose impact on accelerator performance is simulated and quantified.

1.2 Publications

The research contributions of this thesis have also contributed to the following publications:

- M. Benedikt, F. Zimmermann, B. Auchmann, et al., Future Circular Collider Feasibility Study: Volume 2 – Accelerators, technical infrastructure and safety. CERN, Geneva, Switzerland, Rep. CERN-FCC-ACC-2025-0004, 2025.
- J. W. Heron, H. A. Dostmann, L. Felsberger, D. Wollmann and J. Uythoven, "Availability Assurance in the Future Circular Electron-Positron Collider (FCC-ee)," in *Proc. 16th International Particle Accelerator Conference (IPAC'25)*, 2025.
- H. A. Dostmann, J. W. Heron, L. Felsberger, D. Wollmann and J. Uythoven, "Qualitative Availability Forecasting in the Beamstrahlung Dump for the Future Circular Electron-Positron Collider (FCC-ee)", CERN Document Server, pending publication.
- J. W. Heron, H. A. Dostmann, L. Felsberger, D. Wollmann and J. Uythoven, "Availability, Efficiency and Luminosity: Rising to the Challenge in the Future Circular Electron-Positron Collider (FCC-ee)", Physical Review Accelerators and Beams (PRAB) journal, (submitted, pending review).

and were presented at the following conferences:

- FCC Week 2025, Vienna, Austria, May 19–23, 2025. (presented by H. A. Dostmann)

- IPAC'25 - the 16th International Particle Accelerator Conference, Taipei, Taiwan, June 1-6, 2025. (presented by proxy)

The rest of this thesis is laid out as follows. Chapter 2 places the research question in the context of the current state of research and identifies gaps in availability analysis for systems during early development stages to justify the contributions of this thesis. It also provides the necessary background on the FCC-ee, which serves as the case study for this work. Chapter 3 introduces the simulation environment and presents a generalised framework that extrapolates fault data from existing equivalent systems. To address cases where no equivalent system can be identified, Chapter 4 proposes a formulation for estimating availability based on system complexity. This enables availability estimation even if the reference system has significant design differences.

Simulation results are presented in Chapter 5, demonstrating the value of early availability analysis by revealing potential shortfalls in the baseline design of the FCC-ee. Chapter 6 explores R&D opportunities for design optimisation, illustrating how the developed simulation framework can inform design decisions beyond initial availability predictions. Finally, Chapter 7 summarises the main contributions of this thesis and draws conclusions, while Chapter **??** outlines directions for future work.

2 Literature Review

This chapter presents a structured review of the scientific and technical literature relevant to availability in the context of large-scale technical infrastructures. It outlines established methodologies, highlights current challenges, and identifies critical gaps in the assessment of availability during early design phases. Furthermore, essential background on the FCC-ee is provided to contextualise the subsequent case study. This will justify the contributions of this thesis.

2.1 Repairable Systems

Throughout its operational lifetime, a technical system is rarely in continuous, uninterrupted operation. Downtime occurs either due to unexpected failures or planned maintenance activities [8]. Figure 2.1 illustrates a typical state diagram, where the system is assigned a value of 1 during operation, and 0 during failure or shutdown periods.

Many machines are not static but dynamic systems that operate under varying conditions and in multiple operational modes. For example, a system may have a low-load mode, in which certain subsystems are either operated at reduced capacity or temporarily deactivated, and a nominal operation mode, where the machine runs at full design load, delivering maximum output. Within these operational states, key factors such as subsystem reliability or system performance may vary.

Even within non-operational states (i.e., when the system is not operating), further differentiation is often necessary. For example, one may distin-

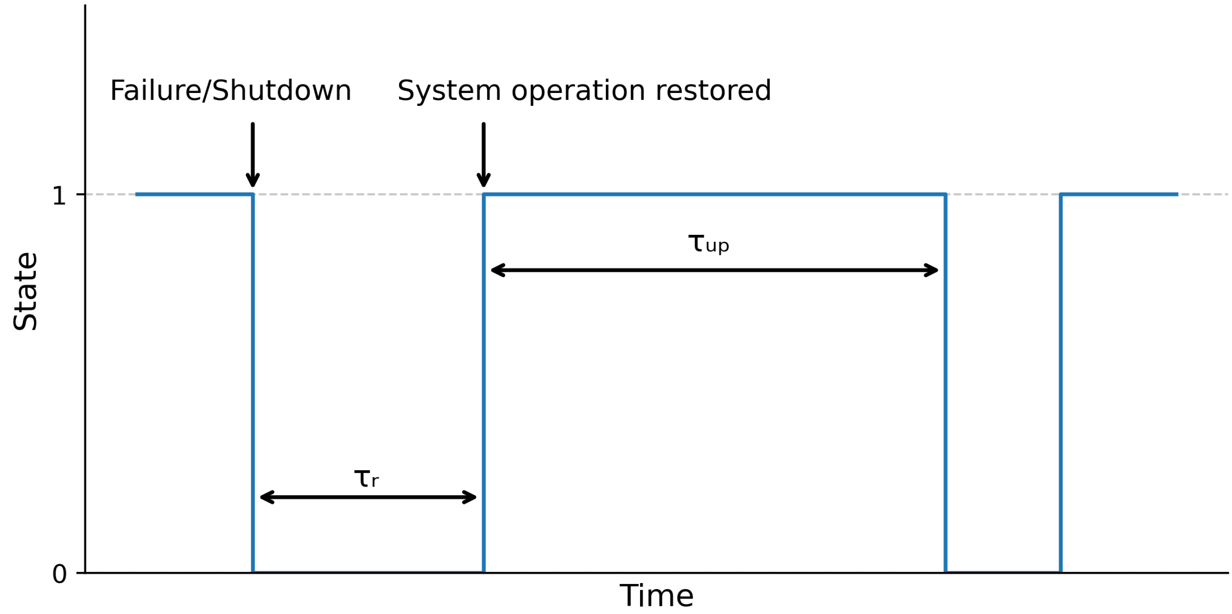

Figure 2.1: State diagram with two states: "up" and "down"

guish between: on-site interventions that require personnel access, remote interventions that can be resolved without physical presence, and planned technical stops, such as scheduled maintenance windows. Each of these operational or non-operational states can be assigned a unique number in such a diagram.

The duration of an intervention, denoted by τ_r, includes not only the repair work itself but the entire period from fault detection (or shutdown) to the restoration of system operation. This covers activities such as diagnostics, travel time, waiting for spare parts, and any other delays.

2.1.1 Availability and Related Terms

Availability describes the proportion of observed time a system remains in an operational state. Assuming the system is considered "as good as new" after each repair and that the observed time span is much greater than both the MTBF and the *MTTR*, the average availability can be expressed as [8]:

$$A = \frac{MTBF}{MTBF + MTTR} \tag{2.1}$$

Here, MTTR denotes the mean downtime of all failures n_f:

$$MTTR = \frac{\sum \tau_r}{n_f} \tag{2.2}$$

The MTBF represents the average up time between the end of a fault and the occurrence of the next und laesst sich auch als Kehrwert der fault rate beschreiben.

$$MTBF = \frac{\sum \tau_{up}}{n_f} \tag{2.3}$$

If such a system consists of several independent subsystems where the failure of any individual component leads to a total operational loss, the system can be considered as a series configuration. In this case, the overall availability A is given by the product of the availabilities A_s of each contributing subsystem:

$$A = \prod_{S=1}^{n} A_S \tag{2.4}$$

The reliability of such a series configuration can be described using the same relationship:

$$R = \prod_{S=1}^{n} R_S \tag{2.5}$$

For instance, if n identical components are connected in series and each follows an exponential failure distribution, the overall system reliability becomes:

$$R(t) = \prod_{S=1}^{n} e^{-\frac{t}{MTBF}} = e^{-\frac{nt}{MTBF}} \tag{2.6}$$

This corresponds to an system-wide MTBF of:

$$MTBF_{sys} = \frac{MTBF}{n} \tag{2.7}$$

2.1.2 Redundancy

Redundancy refers to the integration of additional components that are capable of taking over the functionality of one or several failed units, ensuring that the system remains fully or at least partially operational despite individual failures. It can be differentiated between [9]:

- Active redundancy: All redundant components operate simultaneously, sharing the load. If one component fails, the remaining ones continue without interruption (e.g., parallel power supplies).
- Passive (standby) redundancy: The redundant component is inactive under normal conditions and is only activated when a failure occurs (e.g., backup generators).

2.2 Availability and Efficiency in Particle Accelerators

At CERN, the observed time for calculating availability refers only to the time intended for physics operations, excluding periods such as extended shutdowns, annual commissioning, machine development, and technical stops. This time is divided into the following three proportions (see Figure 2.2):

1. Down Time: Periods where the accelerator is not operating due to unplanned failures or technical issues.

2. Operation: The accelerator is operating but not yet delivering luminosity, as it is preparing and tuning the beam for stable physics conditions. For example during set up periods, beam fill, polarisation or adjustments where no useful particle collisions can be observed.

3. Stable Beams: Attained when the beam is circulating at nominal energy with stable transverse and longitudinal harmonics. Only in this state can useful particle collisions be observed at the experiments

Availability is defined as the proportion of time spent in Operation or Stable Beams relative to the total scheduled physics time. In contrast, operational efficiency refers exclusively to the fraction of time spent in the Stable Beams state.

2.3 Availability in the CERN Complex

Figure 2.3 shows the average contribution of each subsystem to the overall unavailability of the LHC during proton and ion runs from 2015 to 2024. Down time resulting from “child faults” — failures in downstream systems triggered by a “parent system” — is attributed to the responsible parent system.

The Injector Complex accounts for the largest share of downtime. As shown in Figure 2.4 this includes four accelerators: the Linear Accelerator 4 (LINAC4), the roton Synchrotron Booster (PSB), the Proton Synchrotron (PS), and the Super Proton Synchrotron (SPS).

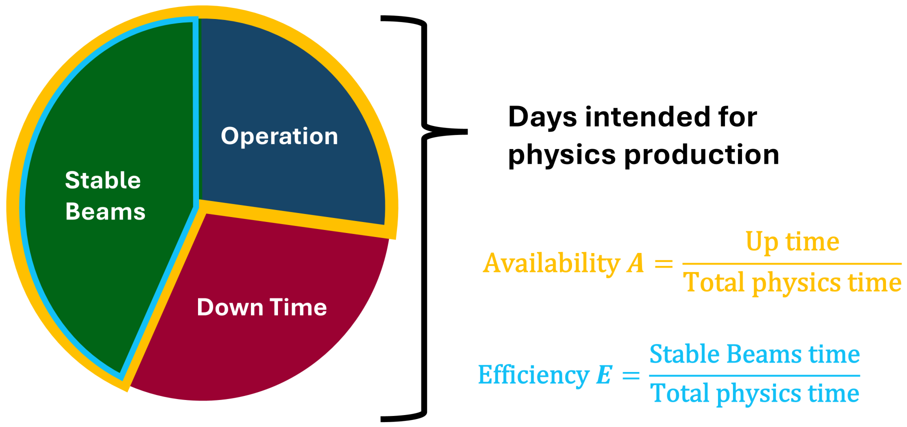

Figure 2.2: Availability and Efficiency at CERN.

Among the subsystems of the LHC itself, the cryogenics system contributes the most to down time. This system is frequently affected by quenches, after which the affected sector must be re-cooled to its operational temperature of approximately 2 K, a process that can take several days [10].

The next major contributor is the quench protection system (QPS), which, in the event of a fault, often triggers cryogenics child faults and leads to extended downtimes. This is followed by the power converters and the electrical network.

A large proportion of electrical network down time can be attributed to child system faults triggered by brief glitches. This is true for nearly all failures caused by external perturbations (see the pie chart in Figure 2.5), as well as for a significant share of faults from internal equipment issues. The bar in Figure 2.5 chart breaks down the downtime by subsystem, with the cryogenics system being the most dominant contributor.

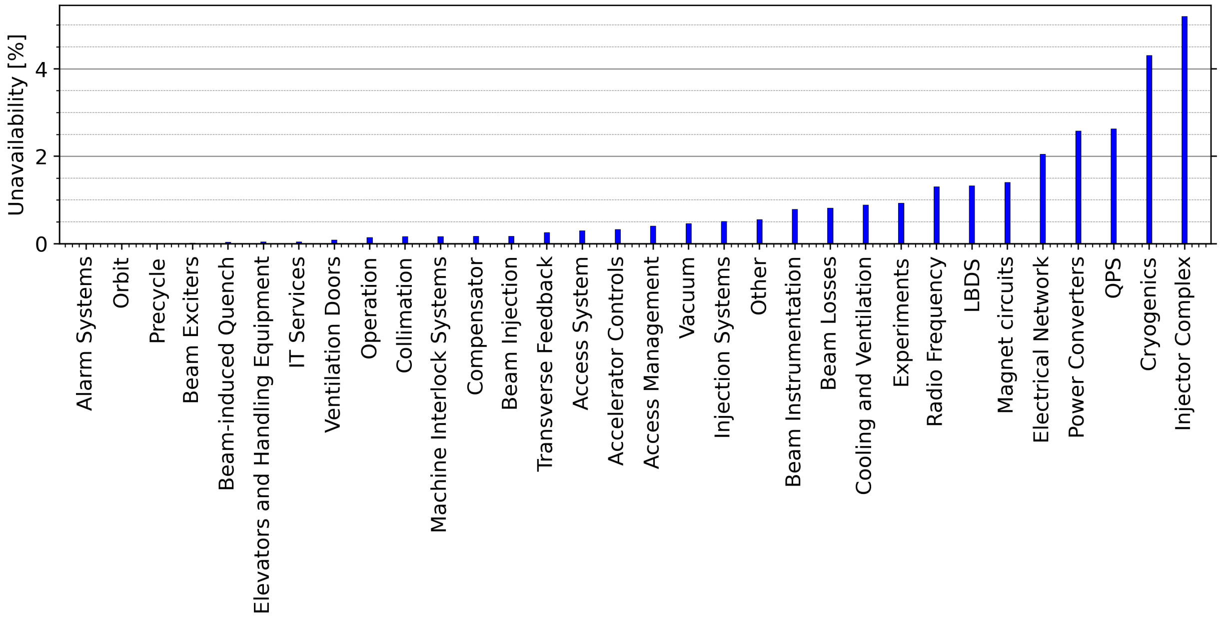

Figure 2.3: Unavailability in LHC during the proton and ion runs from start of stable beams to end of physics, excluding machine developments and technical stops (2015-2024) based on data from [6].

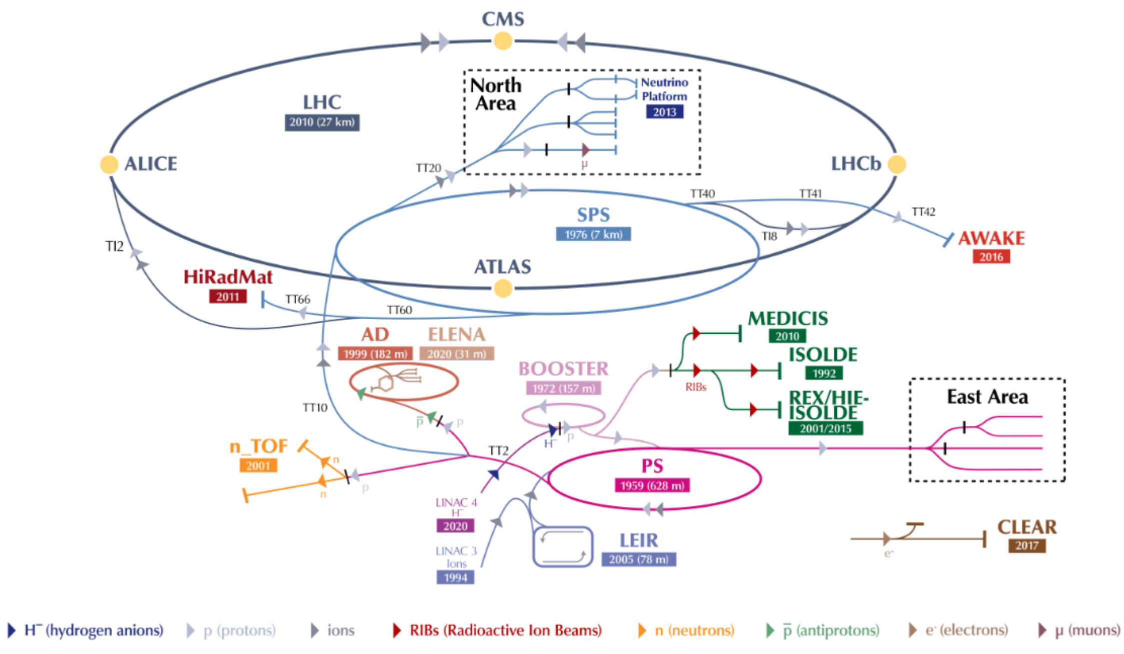

Figure 2.4: The CERN accelerator complex [11]

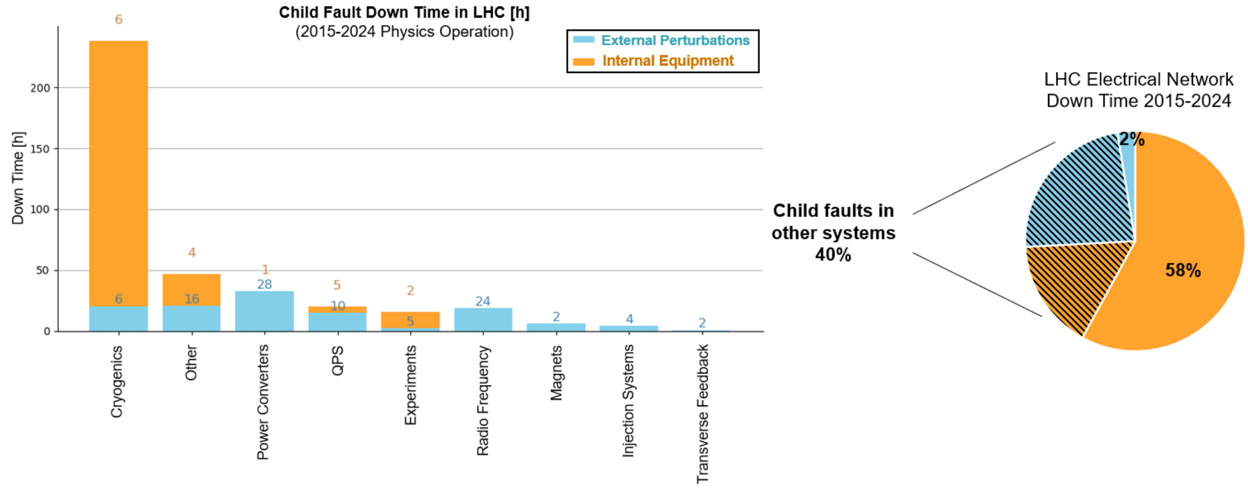

Figure 2.5: Child faults of the LHC electrical network

Despite the LHC having only 16 RF cavities, the RF system still contributes significantly to overall downtime. In many accelerators, this system represents one of the main challenges for availability [6], and consequently, several studies already exist that explore R&D opportunities to improve its performance. One such example is presented in [12], where the availability of the new RF system of the SPS is studied. As part of its upgrade, two new solid-state RF powering systems have been installed. These systems are based on a highly modular architecture distributed across 32 towers. Each tower combines the output of up to 80 RF modules. Only 72 of these are required for nominal operation.

The main objective of the study was to develop a quantitative availability model of this RF system to identify the best operational scenario. Following a Digital Twin paradigm, a detailed reliability model was created, incorporating data from expert knowledge, component testing, and operational experience with similar systems.

The results of the simulations demonstrate that even with just three spare Solid State Power Amplifiers (i.e., 75 installed vs. 72 required), nearly all RF amplifier failures are fully absorbed without impacting system operation—provided no common-mode failures occur. The simulations also show that performing repairs during planned maintenance or "shadow events" – periods of downtime caused by unrelated system faults – significantly improves system availability. Furthermore, introducing redundancy at the driver level could yield additional improvements.

Summing up the contributions from all systems, the LHC was operational only 72% of the scheduled time. For accelerators that are even larger and more complex, such as the FCC-ee, this number could be even lower if availability and reliability are not prioritized from the very beginning of the design process.

2.4 Availability in Future Machines

Early availability analysis of particle accelerators has proven to be important in recent years, leading to the development of several methods and approaches. In [13], a method is presented for breaking down the overall machine availability target into specific availability requirements for its subsystems based on complexity criteria. This provides initial targets for the subsystems, which can later be used as reference points in subsequent availability analyses. The approach is especially valuable when detailed subsystem designs are not yet available or when new technologies are being developed and no failure data exists for accurate reliability assessment.

Six factors are introduced, that influence subsystem complexity: repair time, criticality, intricacy, state of the art, performance time, and environmental conditions. These factors are scored on defined scales by system experts. Subsystems with higher scores are considered more complex and are thus allocated less strict availability requirements. To transform these complexity scores into availability requirements, the paper examines four allocation techniques: the Feasibility of Objectives (FOO) technique, factors product, factors sum, and the Bracha technique. Additionally, the authors propose combining these techniques with the Decision-Making Trial and Evaluation

Laboratory (DEMATEL) procedure, which considers interdependencies between subsystems.

The proposed methodology is demonstrated using two case studies. First, it is applied to the Compact Linear Collider (CLIC) – a proposed future electron-positron collider – to identify which systems require the most attention during design to meet the overall machine availability target of 80%.

In the second case study, the methodology is applied retrospectively to the LHC to compare the complexity-based availability estimations with the actual performance of the subsystems. For most subsystems, the expected and observed availabilities align, confirming the validity of the approach.

Initial efforts to develop availability models have also been undertaken. The process presented in [14] approaches with a functional decomposition of the machine into subsystems or even individual components. For each of these, system experts must be consulted using a standardised failure catalogue template. Based on this information, availability models are developed.

This methodology was applied to the CLIC power converter system. In contrast to a klystron-based powering scheme, a second Drive Beam is proposed for powering. Concerns were raised regarding the reliability and availability due to the Drive Beam magnet. To address this, a new powering strategy has been proposed: instead of individual converters, a single large power converter supplies current to a group of 10 to 60 magnets. The converters are implemented with redundant modules to allow failure-tolerant operation and planned maintenance.

The objective was to assess how much the proposed powering strategy improves availability and reliability compared to individual powering. In addition, the study revealed that further significant gains can be achieved by increasing the reliability of the power converter failover mechanisms.

In [15], the study was extended to include availability modelling of both the main beam RF powering and the main linac magnet powering, each in

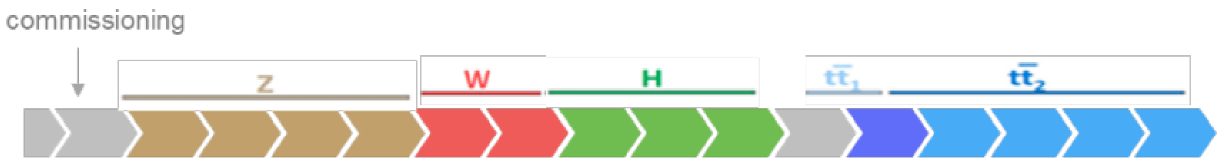

Figure 2.6: Timeline of FCC-ee Operation. Each Arrow represents one year. [18]

both klystron- and drive-beam-based configurations. The drive beam design demonstrated significantly better availability for a single powering sector.

Studies of machine availability to evaluate integrated luminosity performance and to guide both component and accelerator design are also planned for other future accelerators, such as the muon collider [16].

2.5 FCC

The FCC is a visionary and ambitious project led by CERN, intended to introduce a new era in particle physics. It is planned to operate in two stages. The FCC-ee – its first stage – foresees the construction of a 91 km underground circular tunnel, which will initially host a high-luminosity electron–positron collider. The machine is designed to run in four energy modes ranging from 91 GeV to 365 GeV to cover the e^+e^- centre-of-mass energies of the four heaviest elementary particles known today – the Z-Pole (Z mode), WW threshold (W mode), Higgs-strahlung peak (ZH mode), and top-quark pair production threshold ($t\bar{t}$ mode) [17]. Figure 2.6 shows the planned timeline of FCC-ee operations, where each arrow represents one year. A technical stop is foreseen between ZH and $t\bar{t}$ to allow for an upgrade of certain systems.

2.5.1 Status

According to the 2025 Feasibility Study [17]–[19], the FCC-ee is both scientifically justified and technically feasible. With the feasibility study now

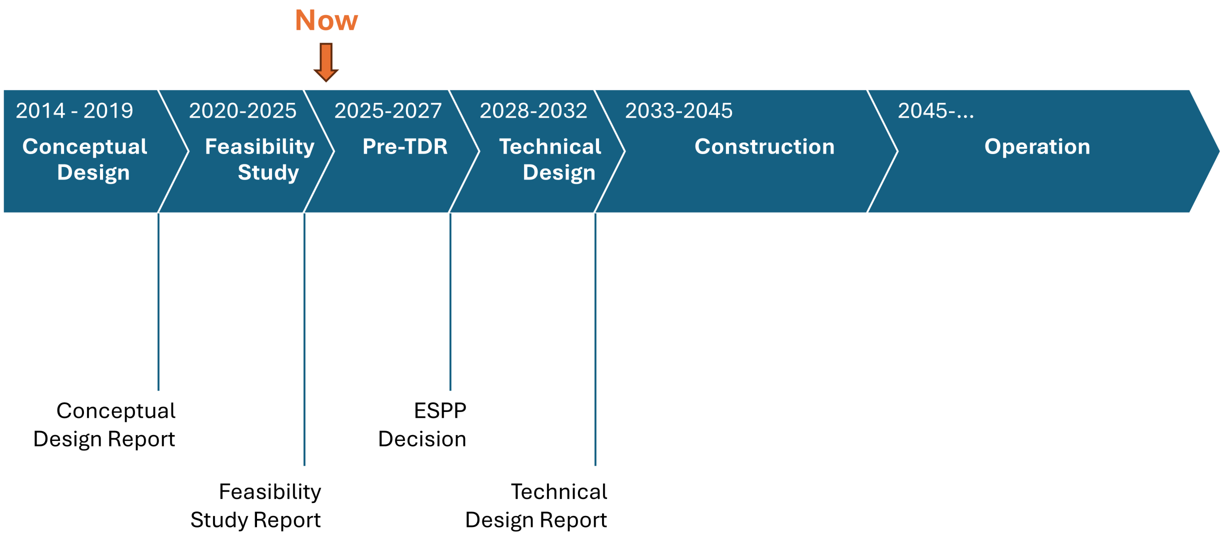

Figure 2.7: Project lifecycle of the FCC project, which is currently in the pre-TDR phase.

complete, the next step is a formal decision by CERN and its international partners, expected around 2028, regarding project approval and funding. If approved, construction could begin in the 2030s, targeting first beams around 2045 (see Figure 2.7. The FCC-ee will serve as a precision tool to study fundamental particles and forces: in its second stage, the infrastructure will be upgraded to host the high-energy proton–proton collider (FCC), with planned operation starting in the 2070s and continuing beyond the end of this century [20].

2.5.2 Top-Up Injection

Some accelerators, such as the LHC, are injected with particle bunches at energies well below the final collision energy. These bunches are then ramped up to the required collision energy within the main accelerator—a process that takes approximately 20 minutes in the LHC, during which no luminosity is produced. In contrast, due to the short beam lifetime during collisions—well below one hour—the FCC-ee requires continuous top-up injection. This means that particle bunches are injected directly at full operational energy from a dedicated full-energy booster ring. This concept is already implemented in most circular lepton colliders **Furukawa_topup**, [21] and eliminates the need for energy ramping, ensuring stable, high-luminosity operation. Among the four operational modes, the Z mode is particularly demanding, requiring injection every three seconds of up to 1120 bunches [18].

2.5.3 Operation cycle

In the Z and W modes, ultra-precise measurements of beam energy are required, which requires the pilot bunches to be polarised at the start of each fill, a process that takes roughly 90 minutes [18]. The required precision in ZH and $t\bar{t}$ modes is lower, eliminating the need for polarisation and avoiding the associated lifetime limitations [22]. Consequently, the electroweak (Z, W) and high-energy $(ZH, t\bar{t})$ modes require distinct operation cycles, shown in Figure 2.8. In [23], the operation cycle of the FCC-ee was presented for the first time and has been updated according to latest developments in this thesis. Phases in this cycle are as follows:

1. *Set Up:* Preparation of the equipment for injection.
2. *Pilot Bunch Injection:* Injection of pilot bunches and adjustment of equipment and optics.
3. *Polarisation:* Polarisation of pilot bunches for resonant depolarisation (RDP) measurement (Z and W modes only).
4. *Fill:* Injection of main colliding bunches. Fill time t_f in each energy mode is shown in Table 2.1.

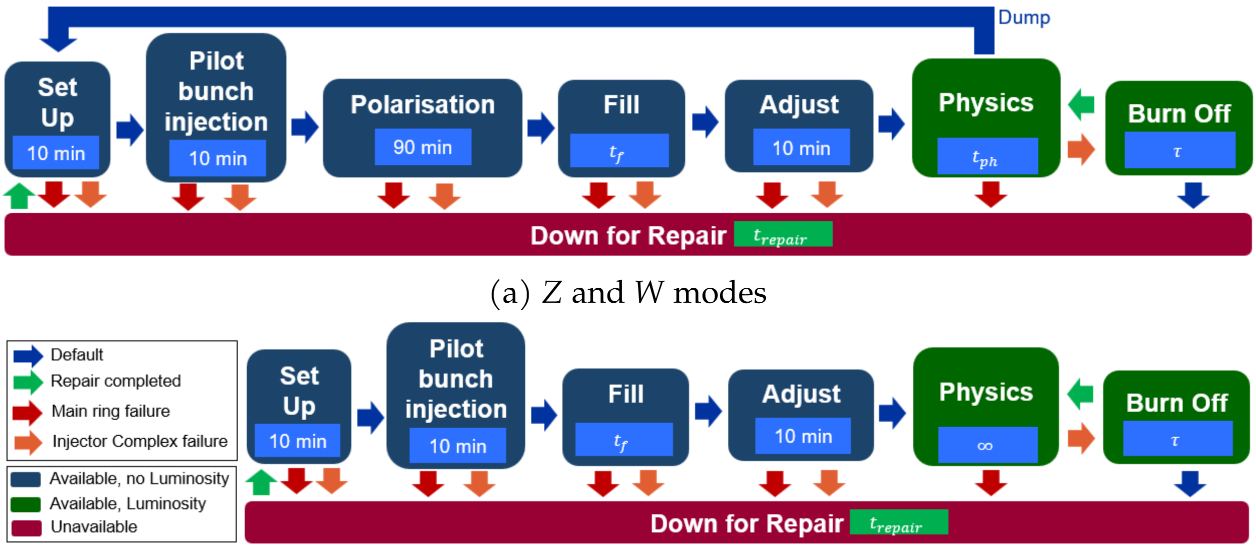

(a) Z and W modes

(b) ZH and $t\bar{t}$ modes

Figure 2.8: Baseline Operation Cycles in the FCC-ee.

5. *Adjust:* Final adjustments to equipment and beam.

6. *Physics:* Beginning of collisions. With top-up injection, flat luminosity can be maintained. In Z and W modes, an upper limit on the duration of physics $t_{ph} = 20$ h is applied corresponding to the Touschek and gas scattering lifetime of the pilot bunches [22].

7. *Burn Off:* Should the injector complex fail, the beam does not need to be dumped immediately. It can remain circulating in the main colliding rings for a characteristic lifetime τ (see Table Table 2.1), while the luminosity diminishes. If the injector complex remains non-operational, the beam is dumped after time .

8. *Down for Repair:* Repair in case of equipment failure. The accelerator is stopped.

Table 2.1: Operation cycle phase durations.

	Z	W	ZH	$t\bar{t}$
Fill time, t_f [min]	7.7	2.5	1.52	1.45
Burn off lifetime, τ [min]	15	12	12	11
Max. physics duration, t_{ph} [h]	20	20	∞	∞

2.5.4 Luminosity

One of the most important parameters in particle physics experiments is the number of useful interactions (events), also known as integrated luminosity [24]. The higher the number of observed events, the more likely rare events, that advance research, will be observed.

The luminosity achieved annually for FCC-ee at each mode of operation are derived from three parameters [5]:

- Nominal Instantaneous Luminosity L: represents the interaction rate analogous to the number of particle interactions observed each second. The instantaneous luminosities for each interaction point (IP) can be found in Table 2.2. To account for machine commissioning and beam tuning, the first two years of Z and the first year of $t\bar{t}$ operation are assumed to reach, on average, 50 % and 65 % of this value, respectively.

- Annual Scheduled Physics Time T: It is assumed that 185 days per year are scheduled for physics. This is obtained by subtracting from one year (365 days): 17 weeks of extended shutdowns (120 days), 30 days of annual commissioning, 20 days for machine development and 10 days for technical stops.

- Operational Efficiency E: In the Z and W modes, a lower efficiency is assumed compared to ZH and $t\bar{t}$, due to the limited pilot bunch lifetime and longer turnaround times, linked to the need to polarise

Table 2.2: Design instantaneous luminosities for each IP, the targeted integrated luminosity and time spent in each energy mode [18].

	Z	W	ZH	$t\bar{t}_1$ / $t\bar{t}_2$
Inst. Luminosity/IP [10^{34} cm^{-2}s^{-1}]	140	20	7.5	1.8 / 1.4
Targeted Integrated Luminosity [ab^{-1}]	205	19.2	12	0.42 / 2.7
Run time [years]	4	2	3	1 / 4

pilot bunches. Assuming an availability of 100 %, efficiency could reach about 90 % in Z and W, and close to 100 % in ZH and $t\bar{t}$ modes.

From the above factors, the achieved annual luminosity at each IP can be calculated as follows:

$$L_{int} = ETL \tag{2.8}$$

The physics goals of FCC-ee require integrated luminosities as listed in Table 2.2 for the four operation modes, summed over four interaction points. Assuming top-up operation without losses due to external influences, this could be reached with an efficiency of about 75 % for all energy modes.

2.5.5 Cool Down Time

When charged particles are accelerated along curved paths, synchrotron radiation is emitted. In the FCC-ee, especially at high beam energies like $t\bar{t}$, this radiation can activate nearby equipment. To protect personnel, zones with elevated radiation levels are inaccessible. If access is needed, such as for corrective maintenance, these areas must first undergo complete air exchange to lower radiation concentrations to safe levels [18]. The required duration of ventilation – or cool down time – varies significantly depending on the area: some zones can reach acceptable conditions within 1.5 hours, whereas more critical sections, like collimation straight sections without a bypass tunnel, may require up to 24 hours before entry is permitted.

2.5.6 Beamstrahlung Dump

In the interaction regions of the FCC-ee, several radiation-producing processes can lead to particle showers that may endanger both sensitive components and tunnel infrastructure, such as electronics and cabling. Beamstrahlung is a particularly intense radiation source in these areas, occurring when charged particle bunches are deflected by the electromagnetic field of the opposing beam during collisions [25]. This process is similar to synchrotron radiation but is localized at the IPs. Because of dense bunches and tight focusing, each beam emits several hundred kilowatts of beamstrahlung power. Consequently, two intense photon beams are emitted along the outgoing beam lines [26], requiring dedicated dumps for their disposal.

2.5.7 Layout

The layout of the FCC-ee is illustrated in Figure 2.9. It will consist of three rings: one for the booster and two for the colliding beams, enabling continuous top-up injection and efficient operation. Additionally, the FCC-ee features an injector complex consisting of low-energy injectors and a high-energy linear accelerator (HE-Linac). The tunnel is divided into Straight Sections (SS) connected by arcs and will be accessible through eight access points, four of which will be equipped with experimental detectors and the remaining four will contain technical accelerator equipment [18].

Electrical Network

The electrical network of the FCC-ee is connected to the French national grid at three access points. PA shares a transmission line with the existing

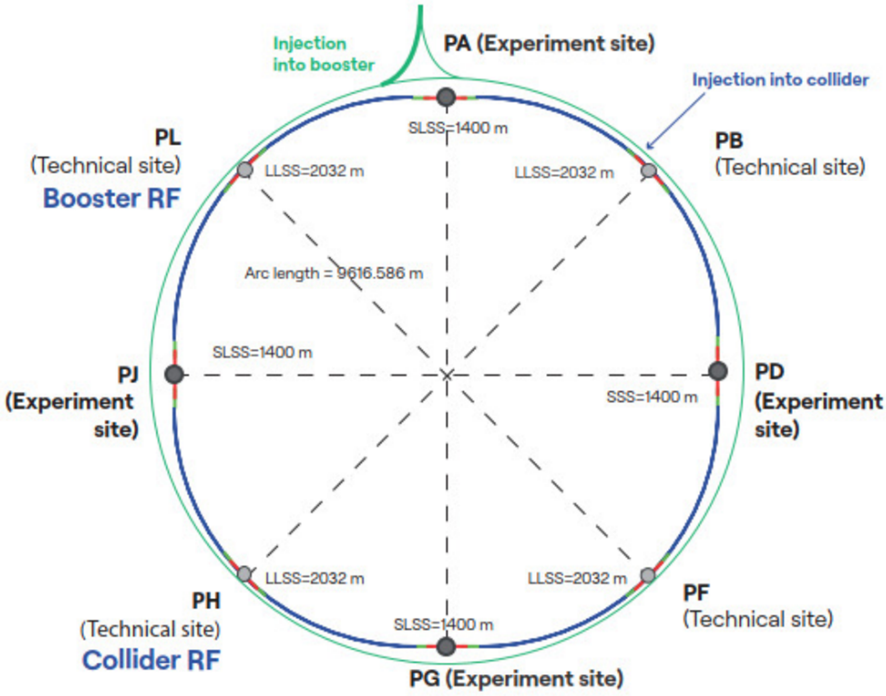

Figure 2.9: Layout of the FCC-ee/

LHC infrastructure, while the other two connection points, PD and PH, are linked via an additional transmission line originating from the same grid connection point in Génissiat. High-voltage substations at these sites -— three at PA, three at PD, and one at PH -— receive power at 400 kV and step it down to 63 kV for the internal HV network, which distributes power to all access points around the ring. Power distribution is organized as illustrated in Figure 2.10: PA supplies PB, PJ, and PL; PD supplies PF and PG; and PH operates independently because the RF systems located there constitute the largest power demand in the machine and require a dedicated, parallel power line. In the event of any HV transport line failure or failures of the supply at PA or PD, the network can activate a redundant connection to reroute power from one of the remaining operational substations, ensuring continuity of supply. A complete shutdown of the machine, without any possibility of network reconfiguration, will only occur in the event of a sup-

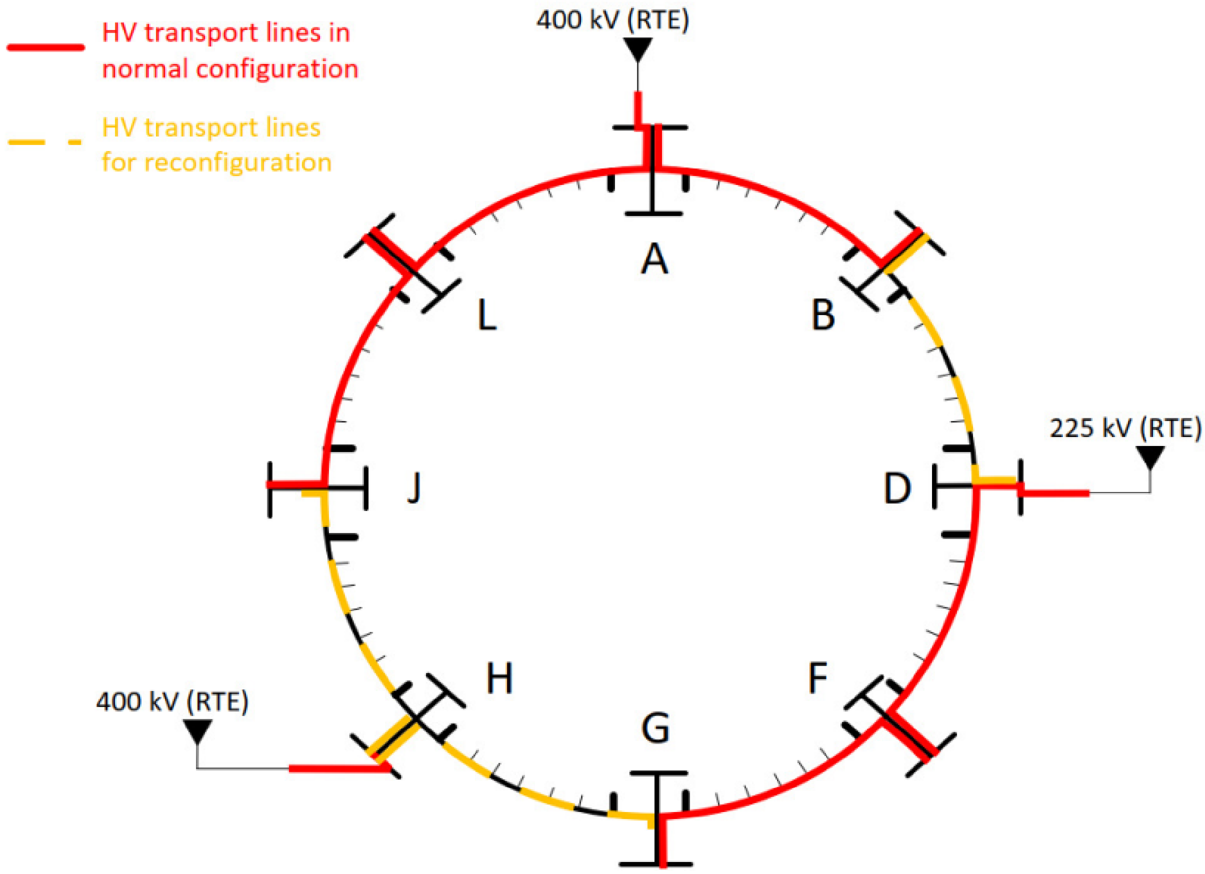

Figure 2.10: Layout of FCC-ee external electrical network [18].

ply failure at PH, where the RF loads exceed the compensating capacity of the other power sources.

The internal network layout at each access point is illustrated in Figure 2.11. The HV transmission lines feed two 63/20 kV transformers, which supply two separate medium voltage (MV) networks: one dedicated to General Services (GS), and the other to Machine systems. These two networks distribute power to 20/0.4 kV transformers, which in turn feed low voltage (LV) switchboards responsible for local distribution to end-use equipment. In addition to the main and MV supplies, a secured network is planned at each access point. This network includes backup power sources capable of supplying critical loads within seconds of a main power outage. The primary backup consists of an emergency power station installed at the surface near

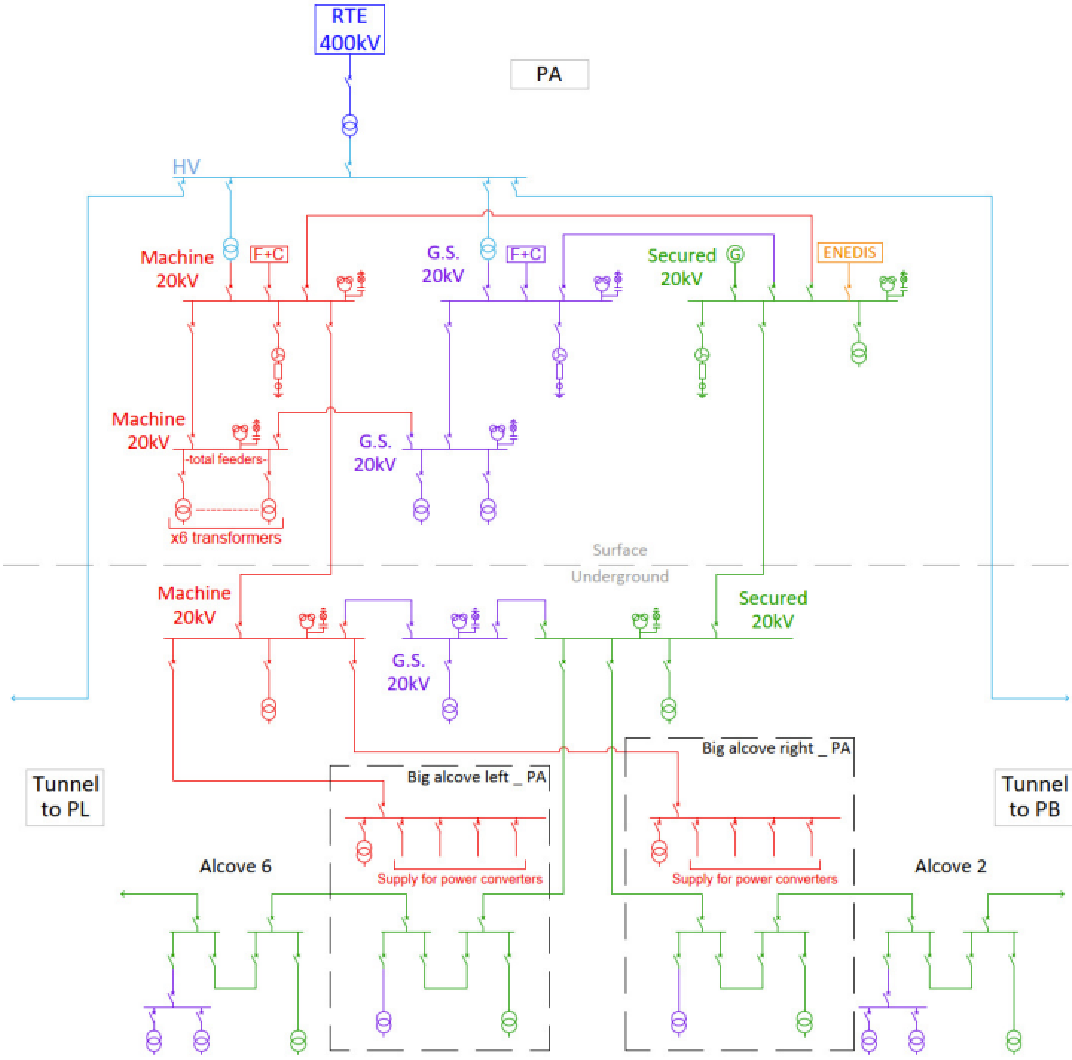

Figure 2.11: Layout of FCC-ee internal network [18].

each substation, with diesel generators as the baseline technology. Furthermore, a secondary backup will be ensured by a connection to the public MV utility network, established during the construction phase and maintained throughout operation [18].

Radio Frequency System

The RF system provides the accelerating voltage required to increase the energy of charged particle bunches and compensating for synchrotron radiation losses by generating high-frequency electromagnetic fields within RF cavities.

The booster consists exclusively of 800 MHz cavities operating at 2 K, while the collider is primarily equipped with coated elliptical 400 MHz cavities operating at 4.5 K. An exception occurs during $t\bar{t}$ operation, where additional 800 MHz cavities are deployed in the collider to enhance the total accelerating gradient.

The cavities for the collider will be concentrated in a single straight section at Point PH, while the booster's cavities will be installed at Point PL. This centralisation, rather than distributing RF systems around the ring, is primarily driven by the requirement for ultra-precise centre-of-mass energy calibration at the Z mode. In addition, this simplifies the integration of supporting systems, including cryogenics, electrical distribution, and maintenance operations. The required number of cavities per energy mode is given in Table 2.3.

A time-domain model was developed to study beam–cavity interactions in the event of RF system failures [18] in the collider. For the Z mode, the failure of a single RF cavity can be detected by the low level radio frequency (LLRF) system within one turn and compensated using strong direct RF feedback from the remaining cavities. However, due to increased impedance, coupled-bunch instabilities can occur. These can be mitigated using a dedicated longitudinal damping system. For the W mode, the system is more robust thanks to lower beam current, higher energy, and strong synchrotron radiation damping. Even in the worst-case scenario involving six simultaneous RF cavity failures, voltage and power transients remain moderate and do not lead to instabilities. In the booster, the transients required due to energy ramping complicate beam stability calculations, and it is presently unclear whether the beam could be preserved in the event of any lost cavities in Z or W mode. The higher-energy modes ZH and $t\bar{t}$ are even less critical and,

Table 2.3: Cooling capacity eq. @ 4.5 K and number of RF cavities [18].

	Z	W	ZH	$t\bar{t}_1$ / $t\bar{t}_2$
# RF cavities (Collider)	264	264	264	672
# RF cavities (Booster)	112	112	112	448
Cool. Capacity [kW] (Collider)	2 x 39.4	2 x 39.4	2 x 39.4	2 x 82
Cool. Capacity [kW] (Booster)	1 x 13.3	1 x 13.3	1 x 13.3	2 x 22

according to current studies, should be able to preserve the beam with 90% of the RF cavities operational in both the Booster and Collider [27].

Cryogenics

The cryogenics system of the FCC-ee is an infrastructure system designed to provide the low-temperature environment required for the operation of superconducting RF cavities and the magnets at the experiment points PA, PD, PG, and PJ.

Since the RF concept involves installing and operating all cavities needed for the ZH mode from the outset, all corresponding cryomodules will likewise be installed at once. This approach reduces the cryogenic system to two main operational stages (with cooling capacities equivalent at 4.5 K as in Table 2.3):

- A unified stage for Z, W, and ZH running, with cryoplants dimensioned for the ZH-mode heat load
- A second stage for $t\bar{t}$ operation, following the same cryogenic scheme.

To optimise system integration, the RF systems of the collider and booster are distributed across two locations: Point PH houses the collider RF, while Point PL hosts the booster RF. This separation also reflects differing cryogenic demands, as the booster must ramp particle energy, while the collider

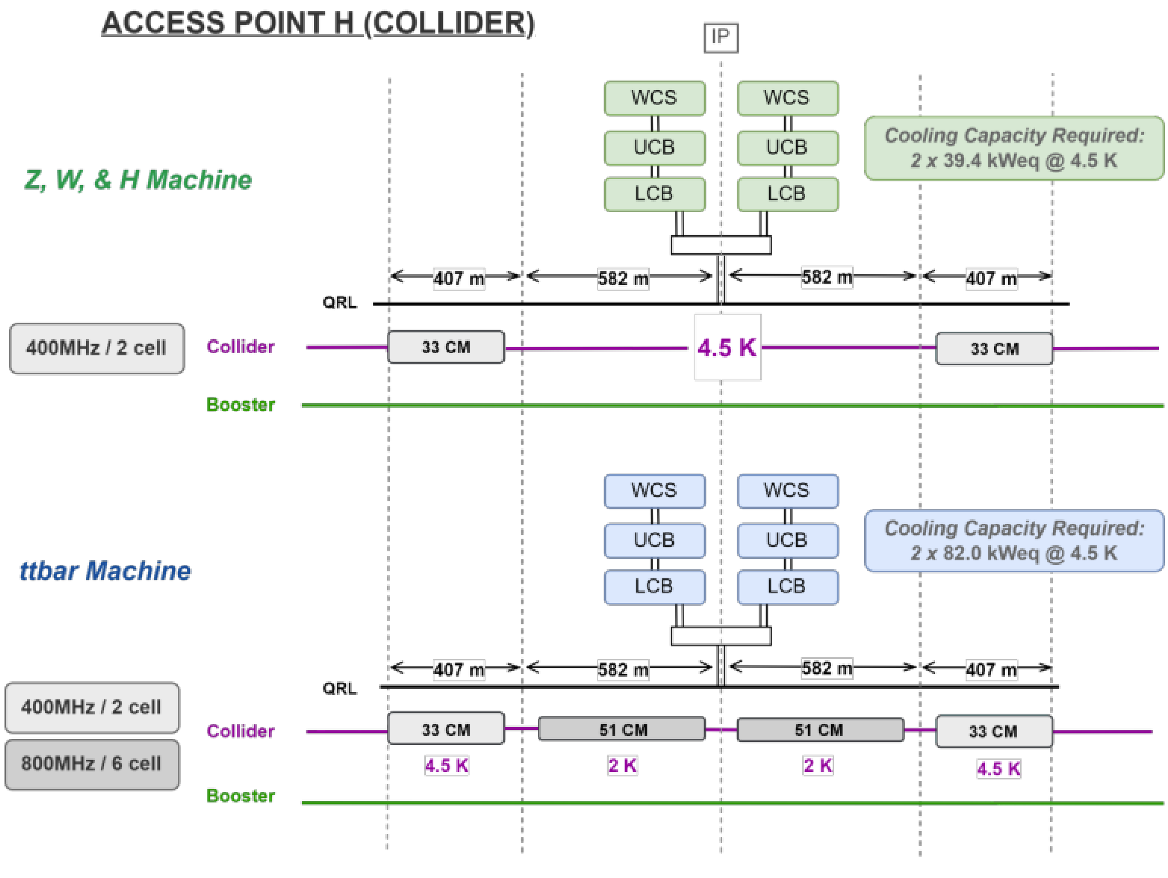

Figure 2.12: Layout of collider cryogenics system [18].

maintains constant energy. Consequently, the static and dynamic heat loads differ significantly.

Figures 2.12 and 2.13 illustrate the cryoplant architectures at points PH and PL, respectively. Configurations with dual cryoplants – existing for all modes in the collider and for $t\bar{t}$ in the booster – enable redundancy between the two cryoplants and facilitate maintenance. In these cases a interconnection unit (QUI) is included to link the refrigeration units to the distribution lines. Other key components include the warm compressor station (WCS) and the upper cold box (UCB) on the surface, as well as the lower cold box (LCB) in the tunnel.

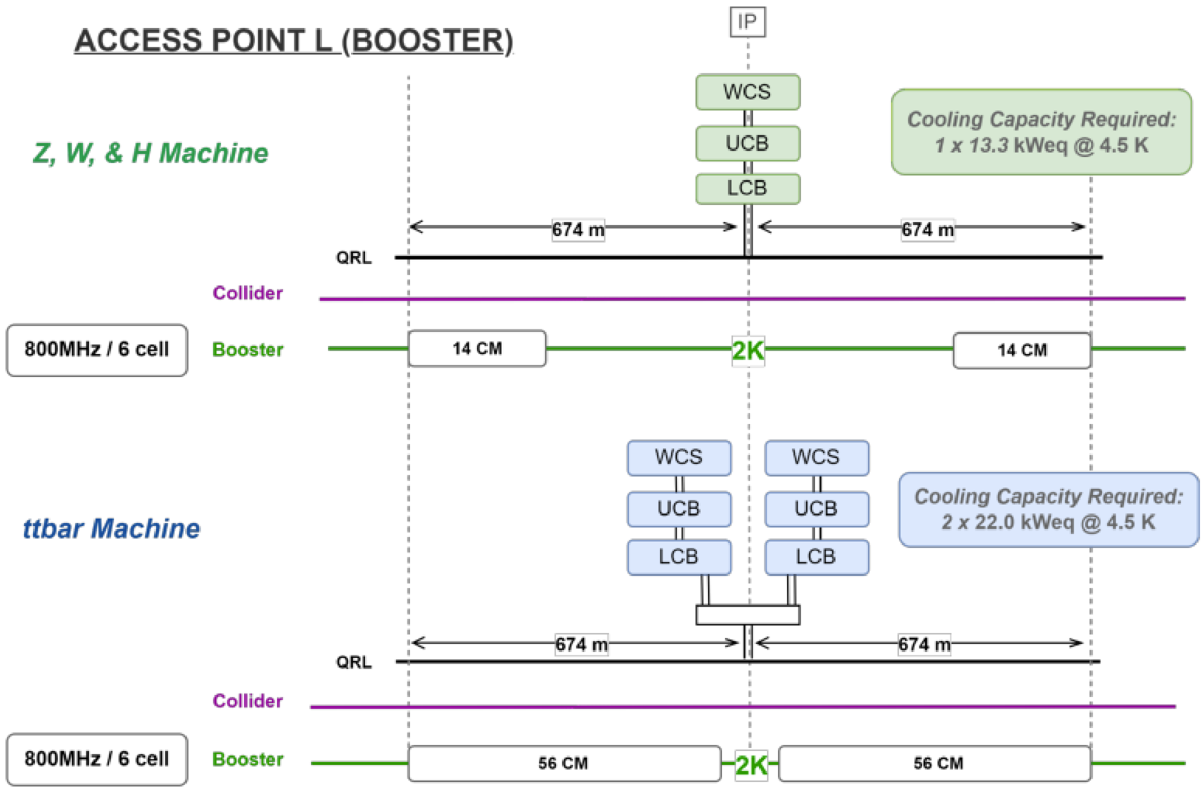

Figure 2.13: Layout of booster cryogenics system [18].

2.5.8 Previous Studies on FCC Availability

The importance of analysing the availability of the FCC-ee was recognised early on. As early as 2016, the first availability studies were carried out for the FCC-hh [28]. In 2023, Heron et al. shifted the focus to the FCC-ee. In [29], he defined availability targets for different FCC-ee subsystems, referring to an overall availability target of 80 %. Instead of setting a uniform availability goal for all systems, the study proposes scaling the requirements according to the complexity of assuring availability, similar as in the previously presented study [13]. Less complex systems, where availability is easier to ensure, are assigned less stringent targets, allowing greater margins for highly complex and critical systems. This serves as a reference point for evaluating the severity of the FCC-ee availability challenge and the risk each subsystem poses to overall availability, although the complexity assessment was based on an earlier development stage and current evaluations may yield a different distribution of target values. Such values can also be used as input for further studies to translate changes in the availability of one system to the overall availability of the FCC-ee.

This was done in [23], where a detailed availability model of the RF system was simulated in conjunction with the target availability values of the remaining systems. The RF system emerges as the most problematic subsystem, particularly in Z and W mode due to the lack of redundancy. Several strategies to improve this performance are discussed.

- Ferroelectric fast reactive tuning: Beam dumps could be prevented by detuning tripped cavities quickly enough to avoid damage. E.g., using a technique as presented in [30].
- Improving RF circuit reliability: Enhancing hardware reliability using strategies as studied in [12] for the SPS.
- Drive Time: The large scale of the FCC-ee means repairs may involve long delays. Stationing personnel around the ring or using fast transport (e.g., helicopters) could reduce downtime with minimal R&D.
- Robot maintenance: Automating repairs could eliminate delays due to drive times and allow some fixes to occur during beam operation.
- Fault prediction: By using sensors, simulations, and machine learning, faults might be predicted and addressed before failure. This could dramatically reduce the frequency and impact of long, human-repair faults.

Simulations show that, without additional measures, the RF system would require a seven-fold and fifteen-fold increase in mean time between failures for the Z and W modes, respectively. Robot maintenance and drive time optimisation could help reduce this gap by up to 30%, which would correspond to a fault prediction rate of approximately 20–30%.

The influence of redundancy measures, such as ferroelectric fast reactive tuning, are further investigated in [31]. For the Z mode, a minimum of 4% redundancy in RF cavities is shown to be necessary to meet the luminosity goal, with diminishing returns beyond that point.

The repair of redundant components was simulated using three distinct repair schedules:

- Optimal: This optimistic schedule assumes that all repairs—both for blocking and non-blocking failures—begin immediately after a failure occurs. Once the blocking failure is resolved, any ongoing non-blocking repairs are interrupted and deferred to the down time.
- Blind: Representing a more pessimistic scenario, this schedule also initiates all repairs immediately after a failure. However, the accelerator is only restarted once all failures, including non-blocking ones, have been fully resolved.
- Realistic: In this more practical scenario, a one-hour attempt is made to remotely fix the blocking failure from the control room following a failure event. If unsuccessful, maintenance personnel are dispatched to repair both blocking and non-blocking faults. The accelerator is restarted only after the completion of all repairs.

The optimal and blind schedules deviate from the realistic scenario by approximately $\pm$5–10%, depending on the energy mode.

In W mode, integrated luminosity falls well short of the physics objectives, even with RF redundancy. This highlights the critical importance of not just subsystem reliability, but also of optimising operational efficiency and fault recovery strategies. To mitigate these limitations, two alternatives to the operation cycle described in Section 2.5.3 are explored for Z and W mode.

- Indefinite physics: If the lifetime of pilot bunches in the main ring can be extended beyond the natural polarisation time, they could be topped up immediately after each measurement and re-polarise naturally. This would eliminate the need to dump the beam after 20 hours for re-injecting polarised pilot bunches. As a result, physics in the Z and W modes could continue uninterrupted until a system failure triggers a beam abort. While beneficial at low fault rates, this scheme offers limited improvement under realistic conditions with frequent interruptions.

- pre-polarised bunch injection (PPBI): If pilot bunches can be polarised prior to injection and their polarisation preserved through the injector complex and the booster, the 90-minute polarisation phase currently included in each Z and W operation cycle could be eliminated entirely. This would enable continuous physics operation in these modes, aligning their operation cycle with that of the ZH and $t\bar{t}$ modes. As a result, the integrated luminosity increases by 15–40% across Z and W modes, with the advantage growing at higher fault rates.

2.6 AvailSim4

AvailSim4 is a simulation tool developed at CERN for predicting the reliability and availability of complex systems. Its primary application areas include accelerators within the CERN complex as well as those operated by partner institutions and projects, such as Multi-purpose hybrid Research Reactor for High-tech Applications (MYRRHA). At the core of AvailSim4 is a Discrete Event Simulation (DES) engine that models system behaviour over time by simulating component failures and repairs within a specified operational time.

The DES framework offers high flexibility and control over the simulation process, as it allows system behaviour to be modelled in fine detail and driven through individual discrete events. The simulations are Monte Carlo based, meaning that DES is performed repetitively, each time sampling from user-defined probability distributions. While this approach is well-suited for capturing the stochastic nature of real-world failures and recoveries, it comes with the drawback of slow statistical convergence, particularly for rare events, which may require millions of iterations to produce reliable estimates.

The input to an AvailSim4 simulation is structured into three files: one describing the modelled system, one defining the simulation configuration, and, optionally, one for sensitivity analysis. The first file, which describes the model, includes the following sheets:

1. ARCHITECTURE Sheet: This sheet defines the fault tree structure of the system. The nodes of the tree represent components, with the leaves corresponding to "basic" components -— actual hardware elements or groups of components to which failure modes are assigned. Compound components serve as logical groupings of basic components and form the internal nodes of the tree. These are used to represent active redundancies, system-level symmetries, or to define the level of aggregation for statistical reporting. Each compound component has child components (which can be basic or compound). AvailSim4 allows for shared children, enabling multiple parent components to depend on the same subcomponents, as well as for the modelling of ??s -— scenarios in which the failure of one component may trigger the replacement or repair of others.

2. FAILURE MODES and FAILURE MODE ASSIGNMENTS Sheets: This sheet specify the failure modes including a failure and repair distribution and assigns them to the basic components. Additionally, failure modes can be linked to specific operational phases, which must be reached before repair can begin. This feature becomes particularly important when defining repair schedules.

3. PHASES Sheet: This sheet defines all possible operational states (or phases) of the system, as described in Section 2.1.

4. PHASES JUMP Sheet: This sheet defines all state transitions, allowing the system phase to change based on any component and any status.

The file describing the simulation parameters defines the numerical parameters, such as the seed, simulation duration, and number of iterations. With the third file, the framework offers support for a basic sensitivity analysis, enabling to examine how changes in input parameters affect simulation outcomes. Selected parameters can be varied by specifying a list of alternative values, which will automatically replace the original entries from the system input file during the simulation runs.

2.7 Summary

This chapter reviewed the role of availability in the context of particle accelerators, demonstrating that it has long been a critical topic in their design and operation. Historical data from the LHC highlights significant availability limitations, particularly within the cryogenics system – due to their long recovery times – as well as within the power converters and the electrical network, which are often the root causes of secondary (child) faults. RF systems have proven to be particularly prone to failure in other accelerators, and corresponding redundancy strategies have already been explored in the literature.

Availability studies have also begun for the FCC-ee, offering initial insights into operational vulnerabilities -— especially in the RF system and the *W* operating mode – and outlining early R&D opportunities. However, existing FCC-ee studies so far rely on availability target values, an RF system model, and in some cases outdated assumptions. To identify more potentially critical systems, these targets must be replaced by detailed, system-specific availability models embedded in a generalized modelling framework grounded in historical fault data. This is done in chapter 3. For subsystems lacking sufficient historical data, a dedicated methodology is developed in chapter 4.

In addition, this chapter provided extensive technical context about the FCC-ee, enabling the application of the proposed methodological framework to this machine as a case study (chapter 5) and supporting the identification of further R&D opportunities (chapter 6).

3 Simulation Environment

This chapter introduces the simulation environment including a generalized framework for availability forecasting of conceptual systems, thereby addressing the research gap on how to model availability for systems in an early stage of development. The approach relies on extrapolating historical fault data from currently operating facilities and is demonstrated through its application to the FCC-ee and selected subsystems.

3.1 Phases/States

As described in Section 2.1, complex machines typically feature multiple phases or states, which must be defined with their transitions and outputs. In the FCC-ee, the phases are described by the operation cycle (see Section 2.5.3). Where only during Physics or Burn Off Luminosity is produced. By tracking the time spent in these operation phases, achieved integrated luminosity is calculated [31]

$$L_{int} = \begin{cases} N_{IP}Lt, & t \in \text{physics} \\ N_{IP}L\tau(1 - e^{-\frac{t}{\tau}}), & t \in \text{burn off} \\ 0, & \text{otherwise} \end{cases} \tag{3.1}$$

Where $N_{IP} = 4$ is the number of interaction points and L is the nominal instantaneous luminosity as per Table 2.2. In the first two years of Z operation, and first year at $t\bar{t}$, reduced luminosity is expected at 50 % and 65 %, respectively to account for machine commissioning and beam tuning.

Table 3.1: Approach and cool-down times relevant to human repair faults at each access point.

	PA	PB	PD	PF	PG	PH	PJ	PL	Arc
Approach Time [h]	0.00	0.75	0.55	0.80	0.63	0.60	0.42	0.28	–
Cool-Down Time [h]	1.54	1.70	1.54	24.0	1.54	0.72	1.54	0.72	1.70

3.2 Repair Schedule

Particularly when redundant and non-redundant systems are combined, the schedule with which repairs are implemented has significant impact on availability [31]. The simulation should reflect a realistic repair strategy, ideally informed by experience from existing facilities. Otherwise, availability estimates can be bounded by evaluating both an optimistic and a pessimistic repair scenario.

The realistic repair schedule for the FCC-ee was initially defined in [23] based on operational experience from the LHC (see Fig. 3.1) and updated to the current findings. Following a dump, remote diagnostics and repair is attempted from the control room for up to one hour. If the fault cannot be resolved within this time, technicians are dispatched to repair the failed component responsible for the dump, as well as any previously failed redundant components. Access to the tunnel is only granted once sufficient ventilation is achieved, and the accelerator is restarted only after all repairs have been completed. An overview of drive and cool down times by location is provided in Table 3.1.

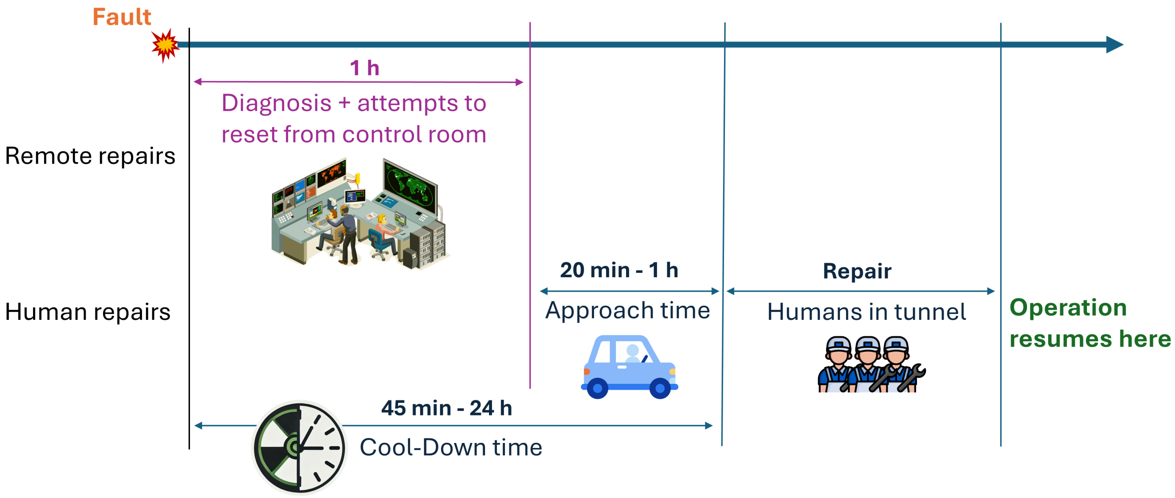

Figure 3.1: Realistic repair schedule of the FCC-ee.

3.3 Generalised Framework for Systems with Representative Fault Data

Once all subsystems that could have a significant contribution to top-level system unavailability have been identified, an availability model can be developed for each of them according to the following steps.

1. An equivalent or similar state-of-the-art system currently in operation at an accelerator and providing fault data is identified. In the case of the FCC-ee, the corresponding system of the LHC was used for comparison in nearly all cases. The data base Accelerator Fault Tracking (AFT) [6] provides well documented fault data for LHC.

2. Fault data is filtered to include only blocking faults during physics operation from 2015-2024 as described in section 2.2. Down time that occurs in the shadow of other faults is excluded, and the down time of

child faults is attributed to their respective parent fault — where a parent fault is the primary cause of failure and child faults are secondary effects or resulting issues triggered by the initial fault. This filtering assumes, for the time being, that the system has the same level of redundancy and the same effect on downstream systems as the reference system.

3. Where applicable, the representative system is broken down into subcomponents/fault mechanisms to allow for more fine-grained analysis.

4. Fault mechanisms are distinguished into:

   - Remote Repairs: can be executed from the control room, for instance by resetting or bypassing the faulty component.
   - Human Repairs: require on-site intervention by technical personnel. Beyond the repair duration itself, the time needed for a technician to reach the affected area must be taken into account.

   Ideally, the representative failure data includes information on whether a human or remote intervention is required. If such information is not available, a simplifying assumption can be applied, whereby failures exceeding a certain duration are classified as requiring human intervention. Alternatively, the data could be separated using techniques such as those described in [32], which enable the decomposition of mixed distributions. Using LHC fault data, failures lasting longer than one hour were categorised as human repairs. This threshold reflects operational experience at the LHC, where remote recoveries were usually completed within minutes, whereas human interventions — including diagnostics, travel time, and repairs — almost never took less than one hour.

5. Distribution Functions are fitted to the grouped failure data

- Failure Distribution: Since the FCC-ee's systems will be regularly maintained and the available failure data does not allow for conclusions about early-life, random, or ageing-related failures, all failure mechanisms are modelled using an exponential distribution. This neglects ageing effects and reduces computing time.
- Repair Distribution: Since many repairs are either very short or fall within a medium range, the more flexible Weibull distribution was selected to more accurately represent the data. In cases where fewer than three comparative values were available, the exponential distribution was used instead.

6. The distribution functions are then scaled or shifted from the reference system to the study system:
   - Failure Distribution: A suitable quantity that can be used for scaling must be identified. This could be, for example, the number of components within the subsystem that can fail individually in relation to the reference system. With twice the number of components, the expected MTBF is assumed to be halved (see Equation (2.7)). Since the exponential distribution is fully characterised by the MTBF, no further steps are required. If the number of components is not yet precisely known, an approximation was made assuming a similar design specification to the reference system.
   - Repair Distribution: The Repair Distribution for remote repair faults is left unscaled. The additional time required to access the failed component, compared to the reference system, is added as an offset – as the third parameter of the three-parameter Weibull distribution, known as the location (or threshold) parameter. It should be noted that the downtime of the reference system may already include approach time. For the case study the drive time from PA to the respective location was taken, as the LHC is only

located around the CERN site at PA. If the subsystem is located inside the tunnel, the cool down time (see Section 2.5.5) to fully ventilate tunnel was also considered.

7. If applicable, a formulation for redundancy is defined.

3.3.1 Illustration based on Selected FCC-ee Subsystems

This section provides specifics on some systems that were modelled as a part of this thesis; details on the modelling of all remaining systems can be found in [18]. All steps require close collaboration with the relevant system experts. They were consulted to identify the critical components to use for scaling, their quantity in both the reference systems and FCC, and any system-specific characteristics such as redundancy schemes or typical fault types.

Radio Frequency

The RF model, based on the 16 superconducting 400 MHz cavities of the LHC was already developed for previous studies [23], [31] and has been updated in this work to reflect the latest FCC-ee design specifications. This includes the number of total and redundant cavities as summarised in Table 3.2.

- Collider: Redundancy in the collider is defined as described in Section 2.5.7. In Z and W, the loss of more than one or six cavities, respectively, must dump the beam. In ZH and $t\bar{t}$ the beam can be maintained with only 90% of the RF cavities operational
- Booster: Pending further study, no redundancy was assumed in the lower energy (higher beam current) modes for the booster. The higher energy (lower beam current) modes are expected to preserve the beam

Table 3.2: Number of cavities and implemented redundancy in the RF system across all operation modes. Higher-energy modes feature increased redundancy, while lower-energy modes operate with little or no redundancy.

	Z	W	ZH	$t\bar{t}_1$ / $t\bar{t}_2$
# Cavities (Collider)	264	264	264	672
# Cavities (Booster)	112	112	112	448
Redundancy (Collider) [%]	1.0	4.5	10.0	10.0
Redundancy (Booster) [%]	0.0	0.0	10.0	10.0

not only in the collider, but also in the booster with only 90% of the RF cavities operational.

As the system incorporates active redundancy, the failure distribution can be scaled to an individual cavity, while the total number of cavities and the corresponding redundancy logic are defined within the system architecture.

Electrical Network

The electrical network model is structured into three components: external perturbations, the dedicated RF supply branch at point PH, and internal equipment. All three components are based on failure data from the LHC. The AFT database provides the LHC data already categorised according to the breakdown presented in Figure 3.2. The blue bars show down time of the electrical network only. In contrast, the green bars indicate the total down time attributed to the system, including additional down time caused by child faults, as described in Step 2 of the generalised framework for systems with representative fault data.

- **External Perturbations**: External Perturbations are voltage deviations originating in the transmission network outside the hardware designed

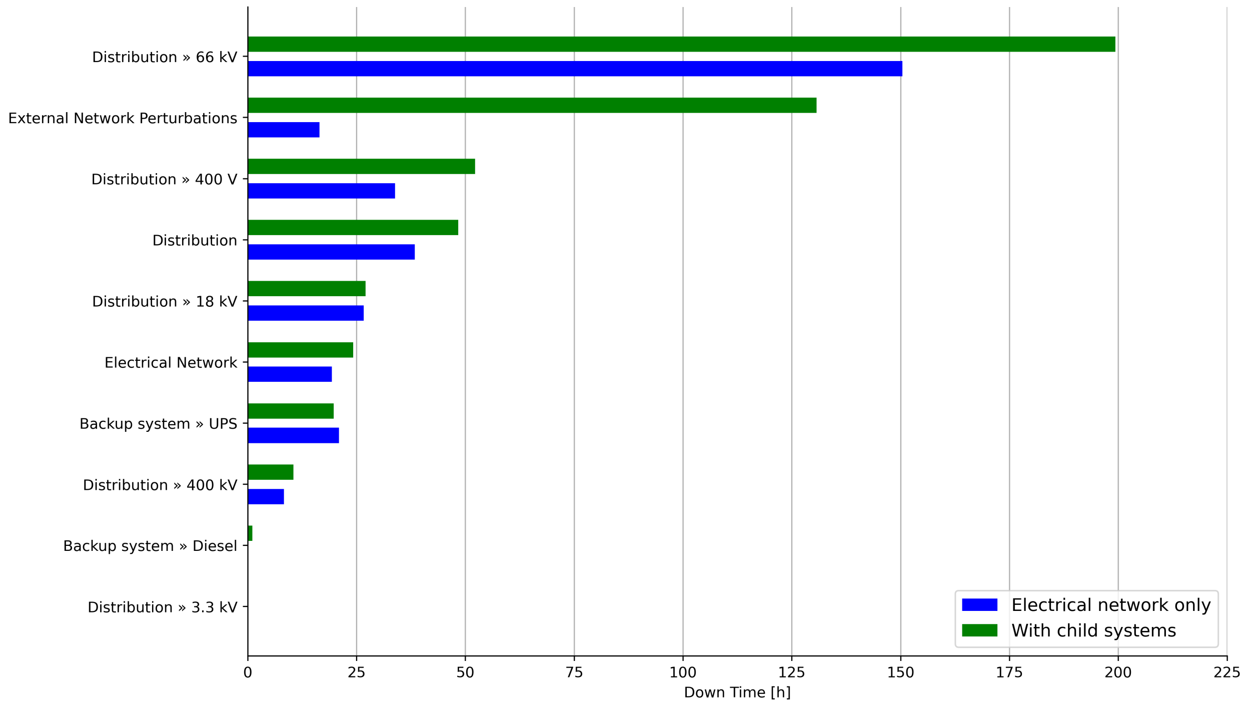

Figure 3.2: LHC downtime caused by electrical network failures during physics operation (2015–2024), highlighting the relevance of child faults within the network and providing a basis for detailed modelling of the FCC-ee electrical network.

and operated by CERN. These events are typically very short in duration (on the order of a few milliseconds) but can trigger child faults in sensitive equipment throughout the accelerator complex. If only the direct down time of external perturbations is considered, as shown by the blue bars in Figure 3.2, this category appears among the least significant. However, when all resulting child faults are included, it becomes the second-largest source of downtime in the LHC.

As described in Section 2.5.7, the FCC will be supplied with power at three IPs from the external transmission network. At PA, the same substation with the same connection to the primary transmission line as used for the LHC will be employed (see Figure 2.5.7, Line 1). Consequently, a similar number of perturbations is expected at this location. In contrast, PH and PD will connect to a second primary transmission

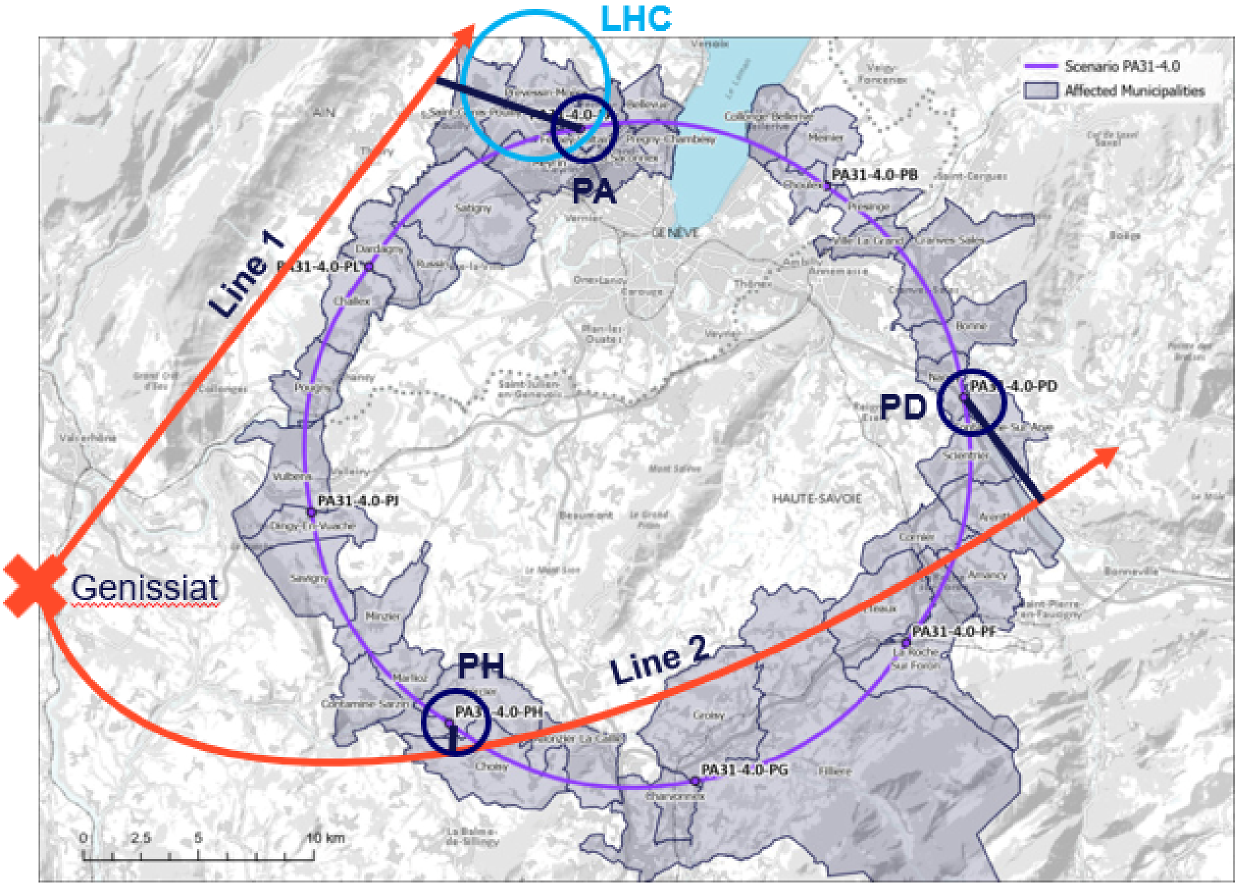

Figure 3.3: External Power Supply of the FCC [33].

line, which also originates from Génissiat. Some external perturbations may scale with the number of main transmission lines; others, since both lines originate from the same source, could affect both lines simultaneously and should thus not be scaled. Additionally, certain perturbations may be specific to individual connections, which would justify scaling by a factor of three. Due to the lack of detailed failure mode categorisation in the available data, the MTBF for external perturbations was conservatively scaled by a factor of two, reflecting the presence of two primary transmission lines.

- **RF Branch**: The RF branch requires a dedicated substation, as the RF system represents the largest power consumer in the machine. This branch is modelled separately, based on failure data from the LHC substation at PA—labelled "Distribution - 400 kV" in Figure 3.2. The MTBF was scaled according to the installed power, resulting in 0.35 times the failure rate compared to the LHC.

- **Internal Equipment**: The remaining internal equipment was categorised into the HV, MV, LV networks, the backup system, and uncategorised faults. Due to the near-identical voltage levels in the LHC and FCC-ee (see section 2.5.7), LHC fault data can be directly mapped onto the relevant FCC-ee subcategories. The specific LHC data used as a basis for each failure mode is summarised in Table 3.3. These failure modes were further subdivided by the power insertion points PA, PD, and PH, and scaled according to the number of locations each substation is required to supply. This results in an overall scaling factor of 8.

  For HV-related failure modes at PA and PD, the redundancy scheme described in Section 2.5.7 was included. This represents a form of passive redundancy, which cannot be directly implemented in the fault tree, as reconfiguration requires manual or automated switching after a failure. According to system experts, this reconfiguration takes approximately one hour. Thus, while the redundancy does not prevent down time entirely, it can significantly reduce it -— provided no child faults occur. If child faults are triggered, the accelerator remains non-operational for their full repair duration, despite the reconfiguration. To account for this behaviour, the affected failure modes were further divided into two categories: those with and without child faults. For the subset without child faults, the model assumes a MTTR of one hour.

Table 3.3: Summary of the LHC fault data used to model each failure mode of the FCC-ee electrical network, including the corresponding scaling factors (SF) applied to adjust the fault rates.

FCC-ee Subsystem	Failure Mode	LHC Fault Data	SF
External Network	Ext. Perturbations	Ext. Perturbations	2
RF Branch	RF Branch	400 kV	0.35
Internal Equipment	High Voltage (HV)	66 & 400 kV	8
	Medium Voltage (MV)	18 kV	8
	Low Voltage (LV)	0.4 & 3.3 kV	8
	Backup System	Backup System	8
	Uncategorised	"Distribution"	8

Cryogenics

Failure modes of the cryogenics system were maintained in the same categories as for the LHC, although some LHC-specific faults were excluded. The LHC employs superconducting magnets distributed throughout the entire ring, all of which require continuous helium cooling. This distributed layout results in a complex helium flow system, which has occasionally led to failures due to helium level oscillations in magnets or distribution feed boxes. In contrast, the FCC-ee requires cryogenic cooling only for its RF cavities, which are concentrated at two points: PH and PL. This localised configuration significantly reduces the complexity of helium distribution. As a result, failures due to helium oscillations are no longer expected in the FCC-ee and were therefore not considered in the model. Almost all remaining failure modes can be assigned to the categories Production (failures in equipment on the surface and within the cryogenic plants), Tunnel (failures in equipment located in the tunnel), and Operation (e.g., non-optimised refrigeration tuning and user-related issues).

In the absence of a detailed design specifying the number of components, the failure rate at each point was assumed to scale with the reference 4.5 K power level, relative to the eight LHC cryoplants, each rated at 18 kW. This scaling remains constant for the lower-energy modes (Z, W, ZH), which require two collider cryoplants with 39.4 kW each and one booster plant at 9 kW. Prior to $t\bar{t}$ operation, a long shutdown is foreseen during which the collider cryoplants will be upgraded to 82 kW and the booster plant to 22 kW, with an additional booster plant of equal power being installed. This results in scaling factors ranging from 0.1 (in the booster during Z, W, and ZH operation) to 1.1 (in the collider during $t\bar{t}$ mode), relative to the LHC baseline.

Access System

The access system manages the operation of tunnel entry doors. Down time of other systems can be extended if a door fails to open or if its status is incorrectly registered, e.g., being reported as open when in fact it is closed, potentially preventing beam operation due to safety interlocks. As both the LHC and the FCC-ee feature eight access points, a comparable number of access doors is expected. Consequently, no scaling of the associated failure rates was applied.

Cooling & Ventilation

The fault rate was scaled based on the number of critical circuits per access point, which is comparable between the LHC (160 circuits) and the FCC-ee (168 circuits). This approach does not account for the significantly larger volume covered by these circuits in the FCC-ee.

Elevators & Handling Equipment

Each access point in the LHC is equipped with a single elevator shaft, meaning that elevator failure directly prevents tunnel access for maintenance. In contrast, the FCC-ee features two elevator shafts per access point, introducing active redundancy that ensures accessibility unless both elevators fail simultaneously.

Lower Energy Injectors

In the injector complex, the electron and positron source, linac, and damping ring are modelled based on unscaled fault data collected from the SuperKEKB electron/positron source, injector linac, and damping ring over the period 1998–2023, as detailed in [7]. These systems feature a high level of redundancy, resulting in particularly short MTTR.

Power Converters

To model failures in the powering systems of various magnet families, fault data was drawn from power converters of comparable types in the LHC. Only failures that resulted in operational down time in the LHC were considered, implying that a similar level of redundancy is assumed for the FCC-ee power converters and magnets. The large number of sextupoles and corrector magnets in the FCC-ee optics design leads to an increased overall failure rate in these categories, with MTBFs on the order of 3 to 5 days.

Smoke Detection System

The smoke detection system in the FCC-ee must identify fire incidents within the accelerator tunnel. Two main failure types are relevant: false positives, where a fire is mistakenly indicated, resulting in unnecessary beam interruptions; and false negatives, where a real fire goes undetected, potentially leading to equipment damage. To ensure safety, the system is intentionally biased towards false positives. Any alarm—regardless of its validity—requires an on-site inspection before beam operation can resume.

Since the LHC tunnel is primarily equipped with superconducting components, its alarm system is designed with a focus on detecting oxygen deficiency due to helium leaks, making it a less suitable reference for the FCC-ee. Instead, the FCC-ee system is modelled based on the aspirating smoke detection system used in the SPS, which offers better comparability. Fault rates were scaled proportionally to tunnel length, assuming an equal detector density per kilometre (7 km in SPS vs. 91 km in FCC-ee).

Vacuum

The vacuum system of the LHC has proven to be highly reliable, with only a small number of failures recorded during the observed period. This low failure rate allowed for a detailed analysis of each individual event, enabling precise identification of the affected components — including pumps, gauges, valves, and controllers. For the FCC-ee, these failure modes were extrapolated by scaling according to the expected number of corresponding components in both the collider and booster vacuum systems.

3.4 Summary

This chapter has presented a comprehensive methodological framework for forecasting the availability of complex technical systems based on representative failure data from existing facilities. As illustrated in Figure 3.4, first the machine is broken down into all subsystems that could significantly contribute to unavailability. For each of these subsystems, a comparable system currently in operation is identified. By applying appropriate scaling techniques, failure and repair distributions are derived from the failure data

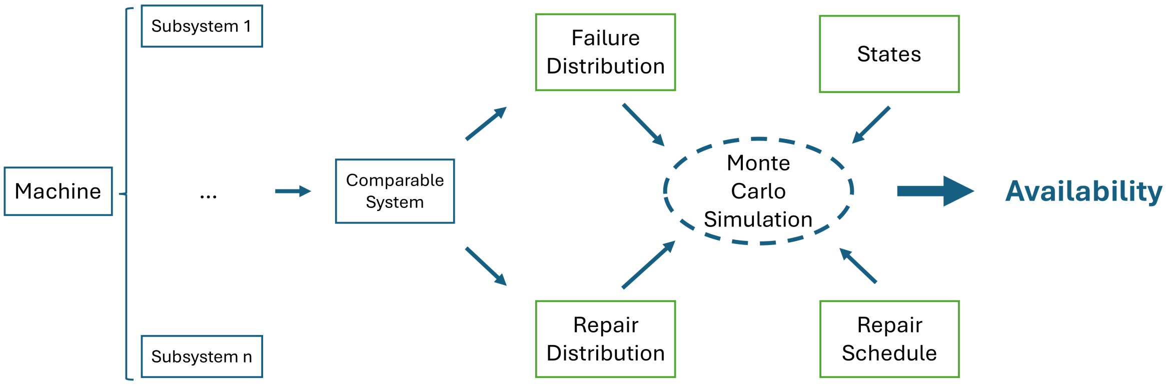

Figure 3.4: Schematic illustration of the methodology, showing all required simulation inputs to estimate availability and, consequently, luminosity.

of these reference systems. After defining the state transitions and a repair schedule, the machines availability is simulated.

The applicability and flexibility of the model were demonstrated through selected subsystems of the FCC-ee, such as the electrical network, cryogenics, and others. For the Beamstrahlung Dump System of the FCC-ee, however, finding a suitable reference system proved challenging, as its design is unique and has no direct equivalent in existing accelerators. To complete the model, an additional methodology will be developed in the following chapter. Together, these chapters provide a solid foundation for further studies, including the simulation of the impact of design changes or R&D opportunities aimed at improving the FCC-ee's performance, as they will be proposed in chapter 6. Moreover, this framework can be applied to future studies involving other technical systems.

4 Systems without Representative Fault Data

If Step 1 of the in Section 3.3 presented methodology is not feasible because no comparable state-of-the-art reference system exists or no failure data is available, but the system is still deemed high risk to machine availability, a provisional availability estimate can be assigned. This placeholder serves until a more detailed reliability analysis becomes possible.

To address such cases, this chapter develops a methodology, which allows any system to be used as a reference, even if it differs significantly in design from the system under study. Similar to the approach proposed in [29], the methodology relies on scaling by complexity, where complexity refers to the difficulty of ensuring availability. The methodology is then applied to the beamstrahlung dump of the FCC-ee.

4.1 Methodology

"Unknown" systems without representative fault data require a new quantity called "complexity" for availability estimation. Such a system is typically in its pre-design stage and therefore exists only conceptually, in the form of requirements and a technological outline. It has not yet been designed in sufficient detail for more traditional methods of reliability and availability estimation to be applicable. The process for estimating the availability of an unknown system involves first finding a "known" system – i.e. one that has already been designed and for which the availability is known – that is most comparable to the unknown system. By assigning a quantified value to their complexities, the availability of the unknown system can be assessed in relation to the complexity and availability of the known one. This holds for all equally competently designed systems.

Complexity represents the difficulty of assuring system availability, and is proportional to the expected down time. It can be assessed by giving numbers to six qualitative scores, each of which independently influence the difficulty of assuring availability in the design stage [13]:

1. **Repair Time**: The repair time represents the average time to restore operation after a failure. Next to the repair itself it also includes time needed for identification, diagnostics and recovery to the nominal operational state. A longer repair time leads to greater complexity in the context of availability.

2. **Criticality**: This score represents the capacity of a system to continue working under partial failure or degraded operation. For some subsystems, the overall system must be immediately shut down in the event of a fault. In others, operation can continue for a certain duration or at reduced efficiency. A higher criticality increases complexity.

3. **Intricacy**: The intricacy refers to the interlocks within the system that can independently block the machine. The higher this number, the greater the overall complexity.

4. **Technical Maturity**: The technical maturity of a subsystem considers the design maturity or level of development of a given technology. This allows for a distinction between systems that exist only theoretically and those that have been in operation for many years. This is a significant factor for availablity, as systems generally get more reliable as their underpinning technology grows more mature.

5. **Performance Time**: The performance time represents the fraction of the total operation time that the subsystem is requested to be available. A system required all year will cause more unavailability and is therefore allocated with a higher complexity than a system required only one day.

6. **Environment**: The category environment can refer to all environmental factors that accelerate the aging of materials, such as particularly high humidity or radiation. High radiation levels for example can destroy cables or other components and complicate availability assurance

The scores B (ranging from 1 to 10) rely on heuristic and relative scoring metrics. This makes the approach flexible to significant design differences between known and unknown systems. But the more conceptually similar a known system is, the more simple the analysis becomes. Various solutions to prioritize and combine these scores are presented and discussed in previous assessments on the CLIC [13] and FCC-hh [34]. The factors product is chosen, as it is not affected by the scores' absolute magnitude. This applies weights w_S for each system according to

$$w_S = \prod_{i=1}^{6} B_{Si} \tag{4.1}$$

With these weights, the complexity is calculated.

$$C_S = \frac{w_S}{\sum_{S=1}^{6} w_S} \quad , \quad \sum C_S = 1 \tag{4.2}$$

As down time scales with complexity, the availability of a subsystem can be described as a function of its complexity and the overall availability [29]:

$$A_S = A^C \tag{4.3}$$

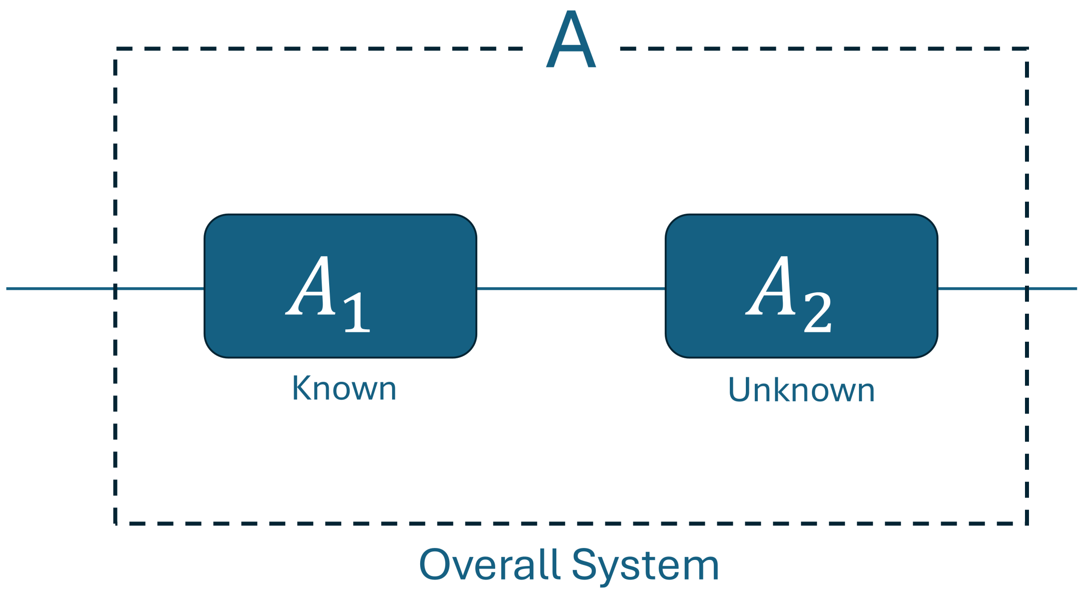

Figure 4.1: System connection of a known and unknown system.

Even if the known and unknown systems originate from two different machines, they are now considered as virtually connected in series for the purposes of evaluation (see Figure 4.1). Disconnecting this virtual series arrangement will not affect the availability of the individual systems. Accordingly, Equation (2.4) holds.

Using Equations (2.4) and (4.3), the following relationship between the two systems can be derived:

$$A_u = A_k^{\frac{C_u}{C_k}} \tag{4.4}$$

With adequately defined numbers for C_k and C_u, as well as the availability of the known system A_k, the availability the unknown system can be calculated.

4.2 Application to Beamstrahlung Dump

Typically, dump systems in particle accelerators use graphite as an absorbing material due to its robustness, capability for high-temperature operation, and relatively good radiation hardness. Depending on the energy mode, the beamstrahlung dump system of the FCC-ee will have to deal with nominal power levels between 80 and 370 kW, with photon energies reaching from 100 MeV up to several GeV. In case of vertical offsets between beams at the interaction point, the power is expected to reach up to 500-600 kW, resulting in temperatures exceeding 4000 °C. For such high thermal loads the thermal conductivity of graphite is insufficient. Consequently, a completely new concept is introduced for the beamstrahlung dump, wherein liquid lead was identified as suitable for the requirements of the FCC-ee's beamstrahlung dump in the case of quasi-continuous wave energy deposition [35], [36]. This would be the first implementation of a liquid dump system in a particle accelerator. In terms of availability analysis, this means that there is no comparable system whose failure data can be used for modelling, and an alternative method must be applied, using a system with significant design differences.

4.2.1 Complexity Evaluation

In collaboration with relevant experts within the design team of the FCC-ee beamstrahlung dump system [37], [38], the Target Internal Dump Vertical Graphite (TIDVG) of the SPS was chosen as the most comparable "known" system for the purposes of this qualitative analysis. The TIDVG consists of several graphite blocks and was designed for 300 kW deposition in the most demanding scenario [39]. The experts were asked to rate the beamstrahlung dump of the FCC-ee and the TIDVG of the SPS relative to each other according to the previously defined categories. The scores are provided in Table 4.1.

As the known and unknown systems are deliberately conceptually similar, many scores are the same for both systems. Since both dump systems must be available during active beams due to safety reasons, and any detected fault will block operation, Criticality and Performance Time are both rated at 10. Additionally, both subsystems operate under extreme environmental conditions and are exposed to radiation levels in the MGy range, leading to an identical rating in the Environment category. Repair Times are also expected to be similar, with an estimated MTTR of approximately four hours in both cases.

However, the systems differ in two key aspects. While the TIDVG is present only once in the SPS, the FCC-ee requires two beamstrahlung dumps per IP (one for each beam). With four IPs, this sums to eight units, each of which can individually block the beam. Furthermore, the beamstrahlung dump is innovative and still in a highly theoretical development stage, whereas the TIDVG has been in use since 2021 and operates with conventional materials. As a result, the ratio of Technical Maturity was rated as 1.5:8. The resulting complexities can be found in Table 4.1.

4.2.2 Availability Estimation

If the two systems are assumed to be in a virtual series configuration, as shown in Figure 4.1, the availability of the beamstrahlung dump can be calculated from (4.4). Using fault data from proton and ion runs during stable beam periods (2015–2024), excluding machine development and technical stops [6], the TIDVG achieved an availability of 99.7%. Consequently, the predicted availability of the beamstrahlung dump system results in 99.4%. This equates to the beamstrahlung dump being responsible for roughly 20 times more downtime than the TIDVG. Extrapolating to the 185 days of physics that are scheduled for the FCC-ee, this would result in about 26 hours of downtime per year.

Table 4.1: Scores from the six complexity categories averaged from two relevant system experts in the FCC-ee beamstrahlung dump design space. The resulting complexity and availability of each system is given below.

Category	**TIDVG** ("known" system)	**Beamstrahlung Dump** ("unknown" system)
Repair Time	8	8
Criticality	10	10
Intricacy	1	8
State of Art	1.5	8
Performance Time	10	10
Environment	10	10
Complexity	0.045	0.955
Availability	99.97%	99.4%

4.3 Summary

This chapter presented a methodology to predict availability of a system in the conceptual stage of its design process, in the absence of sufficient detail for an estimation as proposed in section 3.3.

The proposed methodology is general to any engineering machine or system with arbitrary complexity and technical maturity. The resulting estimated availability of a conceptual (or "unknown") system can be used as part of reliability analysis in a larger engineering hierarchy. It can also be used to assess the level of risk that availability poses to the conceptual system's performance. If the risk is deemed high, this can trigger more detailed reliability-focused treatment in the design stage.

The methodology was showcased with the beamstrahlung dump system of the FCC-ee, a pioneering liquid-based dump concept. The results indicate that the beamstrahlung dump will contribute significantly more downtime than other dump systems. This serves as input to complete the FCC-ee availability model developed in this thesis, enabling the simulation of the machine's overall performance, with results presented in the following chapter.

5 Results

FCC-ee availability is simulated using the discrete-event Monte Carlo-based simulation platform AvailSim4. Each energy mode was simulated separately for its full operation term, with performance averaged over 100 iterations. The results are presented in this chapter. This will help identify shortfalls and critical systems or components within the current baseline design. These insights can inform future design optimisations and demonstrate the value of early availability analysis.

5.1 Overall Physics Performance

The graph in Figure 5.1 shows the achieved integrated luminosity normalised to the target value for each energy mode. The projected integrated luminosity remains well below the target values in all energy modes, with the shortfall being most pronounced in the electroweak sector (Z and W). The pie charts below illustrate the proportions of down time, operation, and stable beams. In Z and W, the time spent in stable beams ranges from 17–22%, whereas in the higher-energy modes (ZH and $t\bar{t}$), it increases to 43–49%. For comparison, the design requirements specified in Table 2.2 require a stable beam time (efficiency) of 75%.

Efficiency is particularly low in the Z and W modes due to the 90-minute polarisation phase at the beginning of each fill, which significantly delays recovery after equipment failures. As shown in Table 5.1, the MTBF in these modes is comparable to the time required to restore stable beams. As a result, the FCC-ee reaches the physics stage in only about 30% of runs before

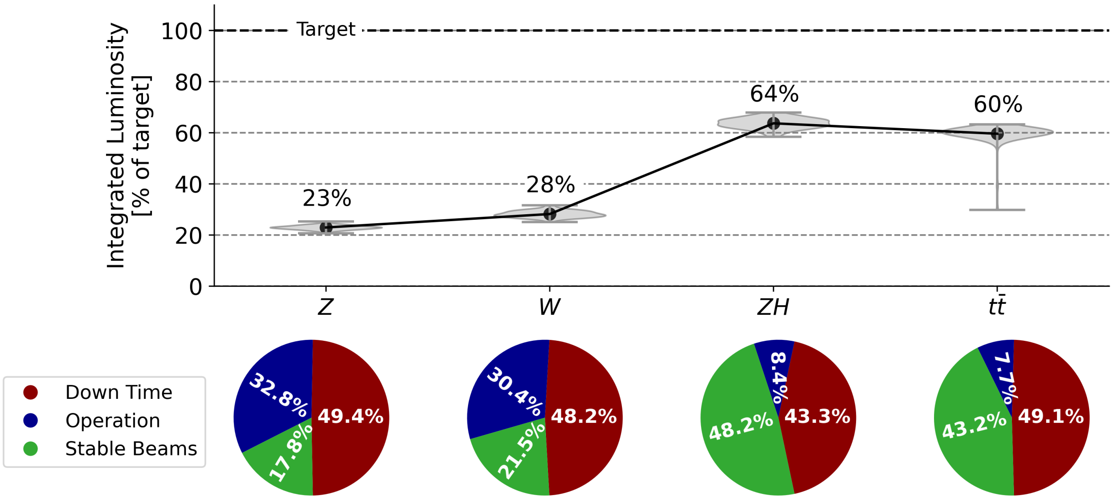

Figure 5.1: Achieved integrated luminosity under the FCC-ee baseline design, normalised to the target values for each energy mode. The pie charts illustrate the time spent in various operation cycle phases. Luminosity and efficiency are particularly low in the Z and W modes.

a dump occurs. In contrast, for the ZH and $t\bar{t}$ modes, this rises to approximately 80%. The combination of short MTBF and long recovery time leads to a substantial shortfall in integrated luminosity at lower energy modes. To achieve the FCC-ee's luminosity goals, significant improvements in reliability, availability, and operational efficiency will be necessary.

	Z	W	ZH	$t\bar{t}$
MTBF [h]	1.9	2.1	3.2	3.1
MTTR [h]	1.8	2.0	2.4	3.0
Recovery Time (to reach stable beams) [h]	2.1	2.0	0.5	0.5

Table 5.1: MTTR and MTBF of the overall machine for each energy mode. In the lower energy modes, the MTBF is similar to the time required to recover stable beams after a failure.

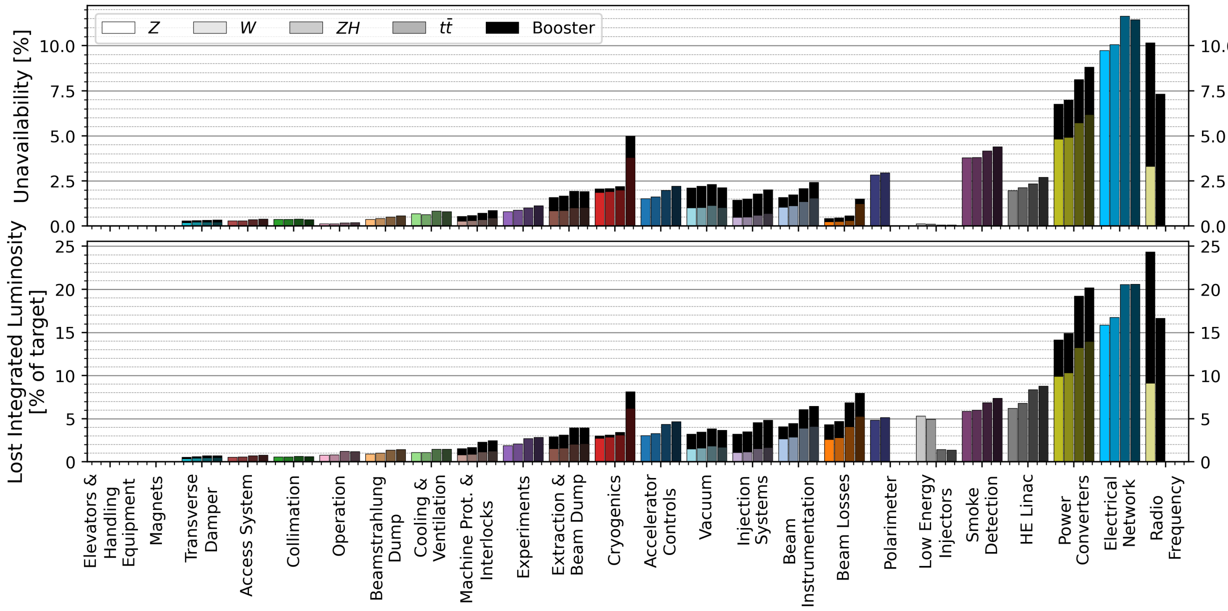

Figure 5.2: Unavailability and lost luminosity contribution from each main constituent system in the FCC-ee. Lost luminosity is normalised to the integrated luminosity target in each energy mode, shown in Table 2.2.

5.2 Contribution of the Subsystems

Figure 5.2 illustrates the contribution of each system to the FCC-ee's unavailability -— defined as downtime relative to the total time scheduled for operation -— as well as its impact on lost luminosity, normalised to the integrated luminosity targets in each energy mode (see Table 2.2). If a system exists in both the collider and the booster, the proportion corresponding to the booster is shown in black. The systems are ordered by their lost luminosity in the Z mode. The primary contributors to lost luminosity are the RF systems, the electrical network, the power converters and the Injector Complex.

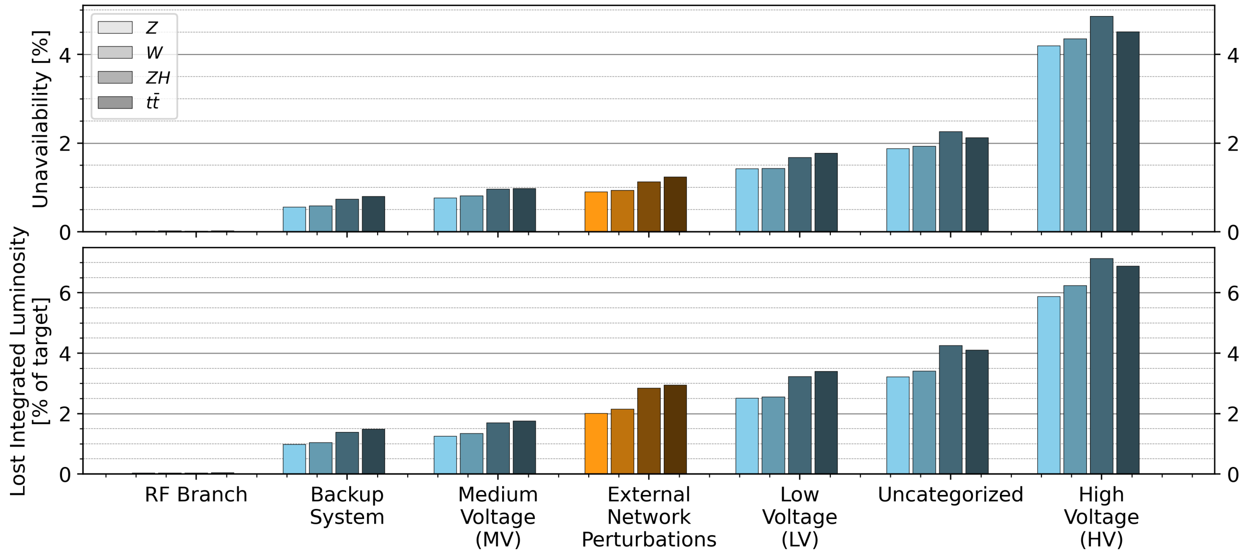

Figure 5.3: Unavailability and lost luminosity contribution from the electrical network in the FCC-ee. Lost luminosity is normalised to the integrated luminosity target in each energy mode, shown in Table 2.2.

Radio Frequencies

The RF system has highest unavailability in Z mode due to the strict constraint of allowing only a single-cavity redundancy in the collider. In contrast, the W mode benefits from the ability to tolerate the loss of up to six cavities without beam loss, effectively reducing the collider's unavailability contribution from the RF system to zero. This aligns with the findings in [31], where a redundancy greater than 4% showed little to no further improvement. The booster, lacking any redundancy, contributes significantly to unavailability in both low-energy modes. In ZH and $t\bar{t}$ modes, the 10% cavity redundancy is sufficient to entirely eliminate RF-related downtime, as failed cavities can be repaired during longer interruptions caused by other subsystems.

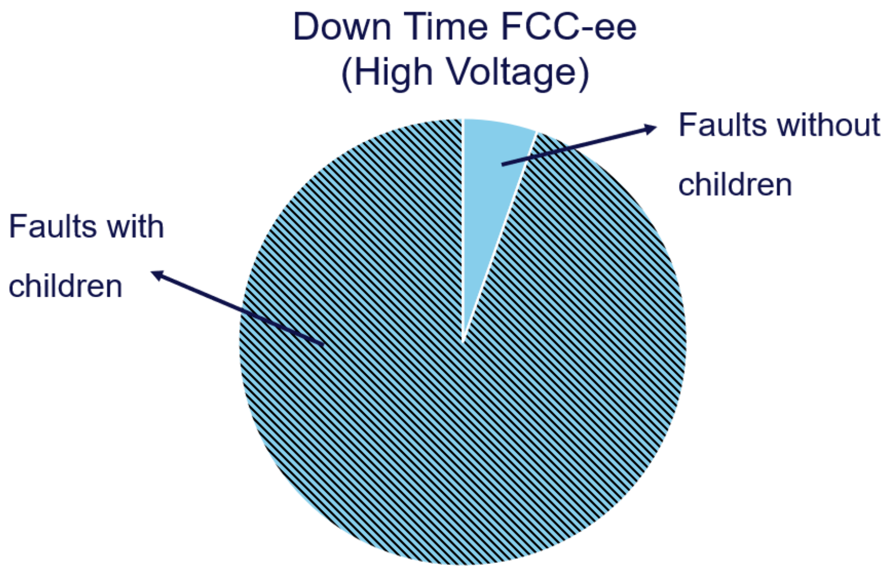

Figure 5.4: Down time of faults in the HV network with and without children.

Electrical Network

The electrical network is forecasted to be the highest contributor to luminosity losses in the higher energy modes. The breakdown in Figure 5.3 shows that, despite the implemented redundancy, the HV network accounts for the majority of these losses. This is primarily due to child faults, which significantly reduce the effectiveness of the redundancy. As illustrated in the pie chart (Figure 5.4), the proportion of down time caused by failures without child faults -— i.e., those that could benefit from the reconfiguration -— is minimal. To fully benefit from the redundancy in the HV network, child faults must be prevented. Following the HV network, the next largest contributors to losses are uncategorised faults and the LV network. Losses due to external perturbations rank lower in the FCC-ee than in the LHC due to the lower scaling factor; nevertheless, they remain significant.

In the electrical network of the LHC, child faults are particularly relevant. Without additional mitigation measures, this can also be expected in the FCC-ee. However, it is important to note that the most affected downstream

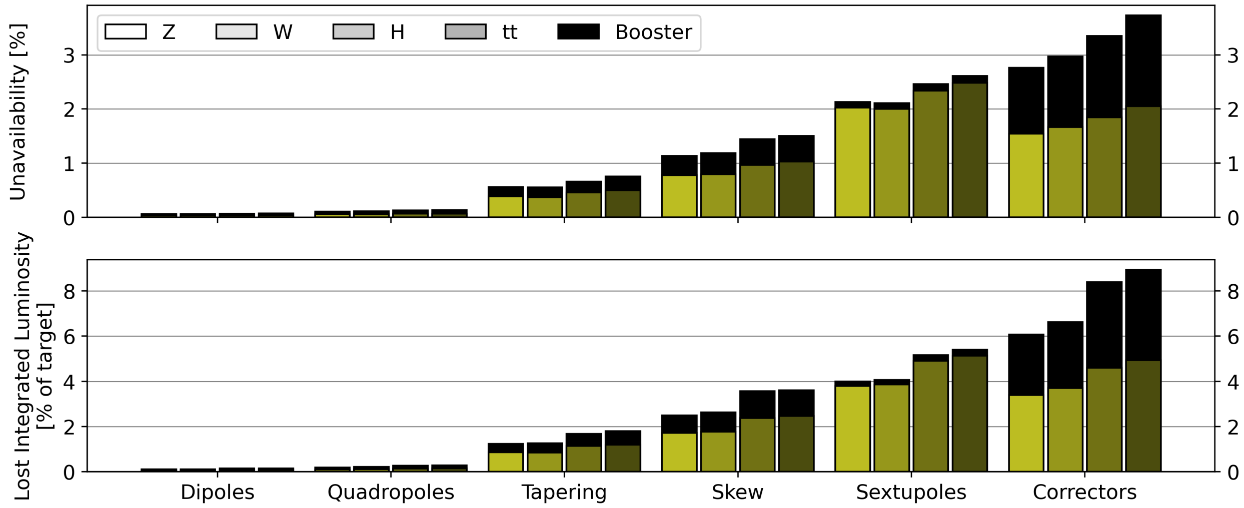

Figure 5.5: Unavailability and lost luminosity contribution from each power converter group in the FCC-ee. Lost luminosity is normalised to the integrated luminosity target in each energy mode, shown in Table 2.2.

systems may differ. The current model assumes the same distribution of affected systems as observed in the LHC (see Figure 2.5). In contrast, the FCC-ee will, for example, have significantly fewer cryoplants, which could result in a reduced susceptibility of this subsystem. On the other hand, the number of RF cavities and power converters will increase, potentially making them more vulnerable to child faults.

Power Converters

Power converters contribute significantly to unavailability across all energy modes. Figure 5.5 breaks down their impact by powered magnet type. The exceptionally high number of converters for sextupole and corrector magnets results in a low MTBF, making it the dominant factor behind the substantial luminosity shortfalls attributed to this system.

Injector Complex

The HE-Linac was modelled based on data from the Linac Coherent Light Source (LCLS) of the Stanford Linear Accelerator Center (SLAC). Together with the low energy injectors it forms the injector complex, which represents the fourth-largest contributor to luminosity shortfalls.

The low-energy injectors, while contributing almost no direct downtime, still cause significant luminosity losses in the lower energy modes. This is due to their relatively high fault rate, despite a very short mean time to repair (MTTR) of approximately three minutes. The fast recovery is made possible by the high level of redundancy in these systems. However, each failure still triggers a turnaround penalty associated with restoring the collider's operational phase. This penalty can reach up to 1.5 hours in the lower energy modes, while it is limited to around 30 minutes at higher energies (see Table 2.5.3).

Furthermore, down times typically remain well below the pilot bunch lifetime of 11–15 minutes (see Table 2.1). During this period, the collider can preserve the beam in the burn-off phase and return to stable beams, provided the injector complex is restored in time. This mechanism helps explain the performance gap between the lower and higher energy modes: since the ZH and $t\bar{t}$ modes spend significantly more time in stable beams, they are more likely to benefit from this resilience to short faults.

Cryogenics

In the $t\bar{t}$ mode, the cryogenics system also has a notable contribution, as it will be upgraded alongside the RF cavities. Figure 5.6 breaks down its impact by fault category. The majority of losses attributed to this system are due to failures of the category Production.

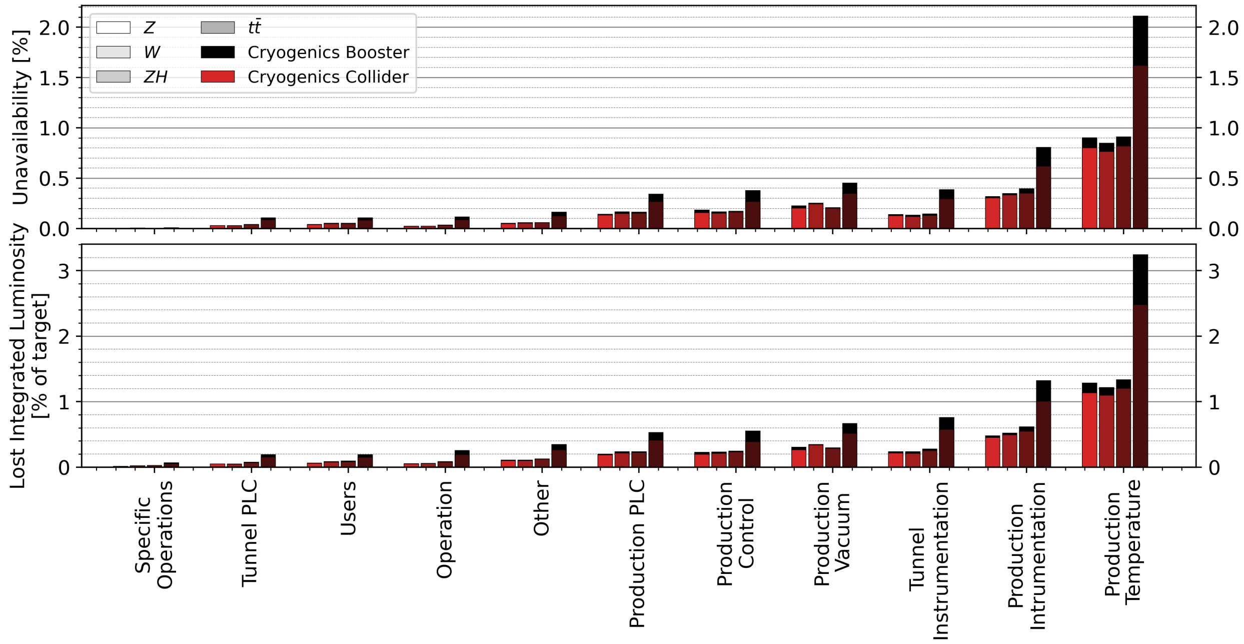

Figure 5.6: Unavailability and lost luminosity contribution from cryogenics in the FCC-ee. Lost luminosity is normalised to the integrated luminosity target in each energy mode, shown in Table 2.2.

5.3 Operational Costs

In addition to predicting performance indicators such as availability and luminosity, the simulation framework also supports intervention planning and the estimation of operational costs. For example, it allows estimating the number of human interventions required per week. Figure 5.9 breaks this number down by access point and Figure 5.7 shows a breakdown by system. In total, approximately 20 on-site interventions per week are required in the $t\bar{t}$ operation mode, corresponding to roughly three per day. The RF system accounts for the largest share, with about 10 human interventions per week. Redundancy may reduce or even eliminate system-induced machine failures, but it does not protect against failures in the cavities themselves, which still require intervention in the shadow of down time caused by other systems. During a human intervention, multiple components within the same system may need to be repaired. This is particularly relevant for the

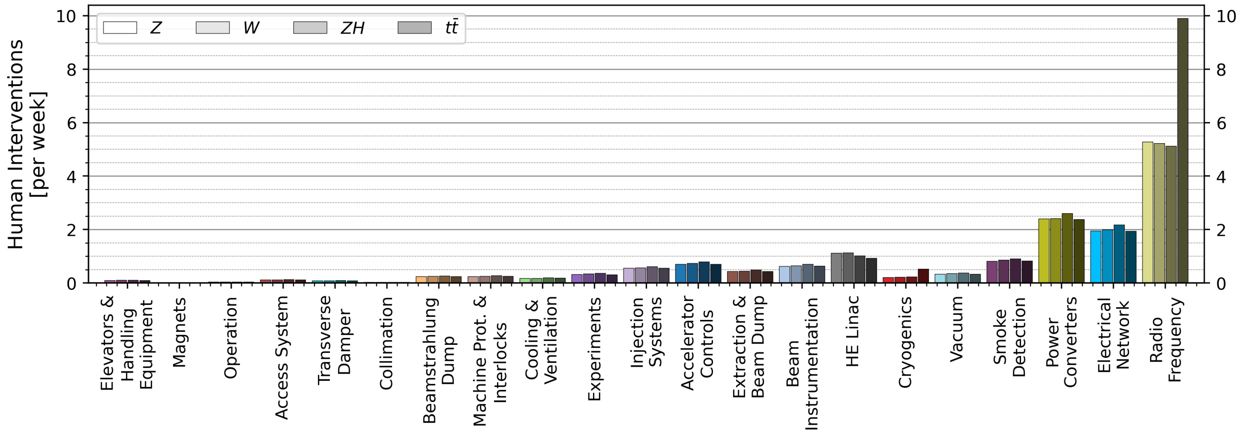

Figure 5.7: Weekly maintenance interventions by personnel per system

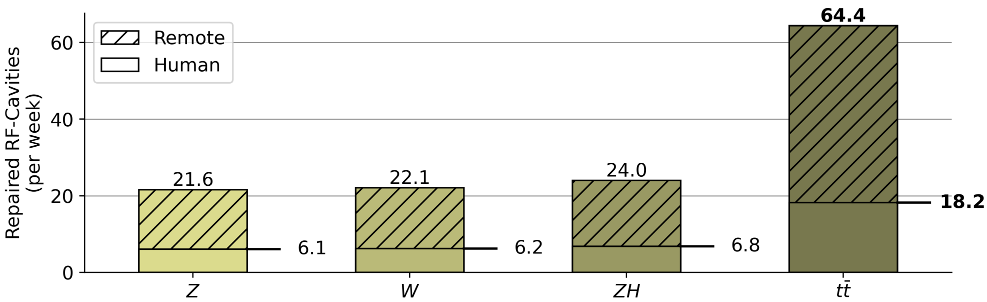

Figure 5.8: Weekly number of failed ?? cavities. In the $t\bar{t}$ mode, 18 cavities require on-site repair by personnel each week, adding up to a total of 64 failures per week when including remotely recoverable faults.

RF system with its redundant cavities. Figure 5.8 shows the number of cavities repaired on-site (blank) and remotely (hatched) per week. During the 10 human interventions required by the RF system each week, an average of 18 cavities must be repaired on-site. Moreover, the number of failing remote cavities is significant, bringing the total count in $t\bar{t}$ mode to as many as 64 per week. These results clearly indicate that a highly responsive and capable maintenance service will be necessary, and that the current design assumptions may lead to high operational costs.

As the collider cavities are located at access point PH and the booster cavities at PL, these access points experience the highest intervention rates.

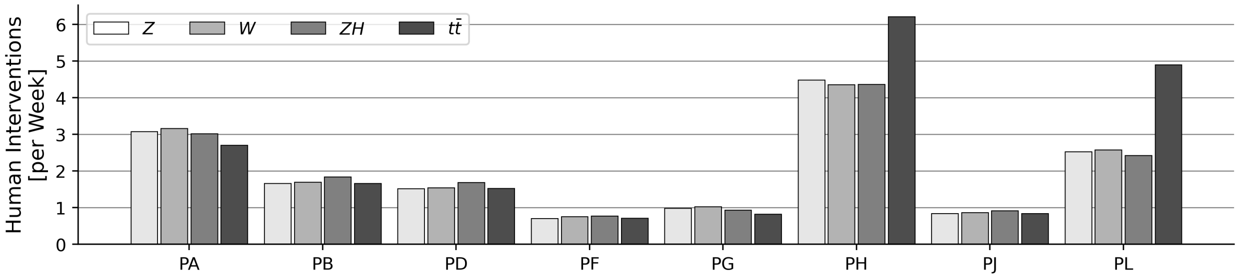

Figure 5.9: Weekly maintenance interventions by personnel per location

The high intervention rate at access point PA is due to the large number of access points the electrical network must supply from this location. PA currently provides the fastest access, as it is located on the main CERN campus, where most personnel are based. However, based on these results, it would also be advisable to strategically allocate resources and staff near PH and PL. Additionally, measures aimed at reducing cool down times, for example, would have the greatest impact if implemented at these access points.

The current figures only include interventions due to beam-blocking failures and the redundant components explicitly modelled in the fault tree, such as those in the RF system. The total number of interventions is expected to increase when systems that are already redundant in the reference facilities are also taken into account. Moreover, the figures depend on the chosen repair schedule.

5.4 Summary

The simulation results clearly demonstrate that the FCC-ee, in its current baseline configuration, falls significantly short of its integrated luminosity targets across all energy modes, indicating a pressing need for machine-wide reliability improvements. This is especially critical for the Z and W modes, where the long polarisation phase further emphasises the importance of solutions to enhance overall operational efficiency.

The subsystem breakdown identifies the most significant contributors to luminosity losses:

- The RF system dominates unavailability in Z mode due to tight redundancy constraints, though higher redundancy in ZH and $t\bar{t}$ modes mitigates its impact.
- The electrical network becomes the leading cause of losses in high-energy modes, with redundancy compromised by complex failure cascades (child faults).
- The power converters consistently affect all modes due to their sheer number and low MTBF, particularly in sextupole and corrector circuits.
- The injector complex, while causing minimal direct down time, imposes significant turnaround penalties -— and thus luminosity losses -— particularly in the low-energy modes, due to its high failure rate.
- Finally, cryogenics begins to play a noticeable role in $t\bar{t}$ mode, especially due to faults of the production category.

Moreover, these results highlight the necessity of a highly responsive and capable maintenance service. The current design assumptions suggest potentially high operational costs. The RF system, in particular, emerges as a key driver of cost and maintenance demands, accounting for approximately half of all weekly on-site interventions and a large share of cavity failures. The concentration of RF-related repairs at specific access points (PH and PL) further suggests opportunities for improvements in access logistics and maintenance scheduling.

To address these shortcomings, the next chapter will focus on identifying R&D opportunities. Since not only the most critical systems have been identified but also the underlying causes of their significant contributions, targeted measures can be developed.

6 R&D Opportunities

This chapter outlines R&D opportunities aimed at enhancing system reliability, availability, and operational efficiency in response to the shortfalls identified in the previous chapter. These measures can be modelled and their impact on overall performance quantitatively assessed, illustrating how the developed simulation environment can inform design decisions beyond initial availability predictions. While some of the proposed solutions will be specific to the FCC-ee, they may also provide transferable insights for the optimisation of systems.

6.1 Radio Frequency

As illustrated in the W, ZH, and $t\bar{t}$ mode, the unavailability and related luminosity losses of the RF system can be fully mitigated by implementing only 4–5% redundancy. However, this effect depends on the failure rate of other subsystems: if the human failure rates of these decrease, fewer opportunities will exist to repair RF components in their shadow, potentially necessitating greater redundancy to maintain the same performance. In the Z and W collider modes, introducing additional redundancy is technically challenging due to strong beam loading effects. Nevertheless, this option should be investigated further, as the performance improvements could be substantial.

It is also important to recognise that the current simulation assumes failures only affect single cavities and ideal behaviour, where seamless compensation between cavity circuits is possible. In practice, shared failure modes may lead to system-level faults that cannot be compensated instantly by backup cavities, reconditioning or component replacement may be necessary. A close collaboration with the RF system designers will be essential in the near future to systematically decompose the RF system into subsystems, identify failure-prone components, and characterise their impact. Extending the modelling beyond the RF cavities to include higher system levels would allow for a more comprehensive analysis and open up opportunities for infrastructure optimisation. For example, following the presentation of results to system experts, it was noted that the klystron gallery remains accessible during collider operation and could therefore serve as a strategic location for installing high-failure auxiliary equipment. If such components can be made redundant, modular, and replaceable without down time (hot-swappable), the system's overall reliability could be significantly enhanced.

6.2 Power Converters

In the current optics layout, groups of four sextupole magnets are powered together via a single power converter, while corrector magnets are each connected to individual units. As a result, the failure of just one converter can simultaneously disable four sextupoles. To address the operational risk this poses, close collaboration with the optics team is necessary to explore strategies for beam stability under such fault scenarios. A shift towards a system-level redundancy concept could offer substantial gains in reliability. The addition of redundant power converters is less a question of technical feasibility and more one of spatial constraints and cost considerations.

It is important to highlight that the present simulation reflects the redundancy levels currently applied in the LHC. For this approach to be effective in the FCC-ee context, sextupole and corrector converter systems would need to be enhanced beyond LHC standards. Furthermore, reducing the number of corrector magnet families — particularly for sextupoles — would

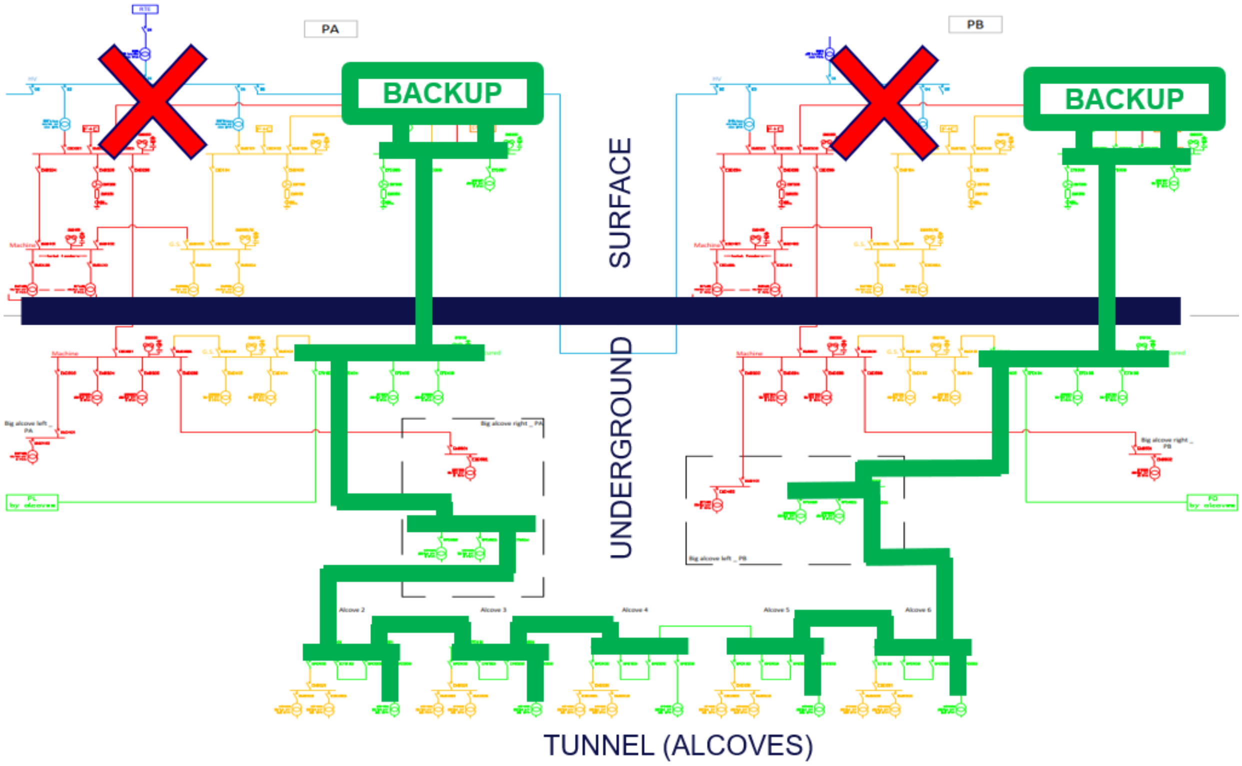

Figure 6.1: Layout of the medium voltage network at PA and PB. [40]

reduce the number of required converters and thus lower the overall fault rate.

6.3 Electrical Network

The electrical network contributes significantly to lost luminosity, despite the presence of a one-hour changeover redundancy between the high voltage network. To make this redundancy more effective, child faults must be reduced. If this can be achieved, it would also be efficient to add further redundancy within the internal network. For example, as illustrated in Figure 6.1, a backup MV network at each access point is planned and has been accepted by safety team under specific conditions [40].

Avoiding Beam Dumps from Short Glitches

Further gains would be made by avoiding beam dumps due to the short glitches. These include several internal network failures and nearly all faults caused by external perturbations. In these cases, damping the voltage deviations to make them invisible for child systems or reducing the child systems' sensitivity would prevent a system failure. Currently, the simulation assumes that the FCC-ee will exhibit a similar resilience to electrical glitches as observed in the LHC. Since the FCC-ee is expected to incorporate a significantly larger number of electrically sensitive components, such measures become even more important.

Achieving the required resilience demands a deeper understanding of the nature and dynamics of these perturbations. Preliminary studies are already underway [41]: Phasor Measurement Units (PMU) have been installed in CERN's electrical network to collect real time data on perturbations. This enables a more precise analysis of power flows within the grid and helps identify correlations between network disturbances and their impact on connected downstream loads. Once these perturbations are better understood, a twofold approach will be required:

- Child systems: The insights gained will inform the definition of detailed electrical resilience requirements for the downstream systems. These requirements must then be communicated to relevant system experts and incorporated as a priority in the ongoing design process.
- Electrical network: The electrical network itself should also be made as "user-friendly" as possible. This requires solutions to dampen these voltage deviations, such as the use of Unified Power Flow Controllers (UPFC) or Direct Current (DC) networks. These technologies can be further explored and validated using a digital twin of the CERN electrical grid, developed using the real time data obtained from the PMU [41].

6.4 Cryogenics

Particularly in the $t\bar{t}$ mode, the cryogenics system may become a critical concern due to the necessary upgrade required for the RF cavities. As a potential solution, the use of cryoplants based on an alternative process -— Turbo-Brayton cycles -— has been proposed. These have been specifically developed with a strong focus on reliability and are expected to be significantly more reliable than conventional cryoplants [42]. Their potential to improve the overall system reliability, will be investigated and modelled in future studies.

6.5 Injector Complex & Booster

Should a failure occur within the booster or injector complex, the main collider rings are capable of sustaining stable beams for approximately 10 to 15 minutes, depending on the energy mode (see Section 2.5.3). If top-up injection is restored within this time, beam dumps can be avoided and physics can resume seamlessly. This significantly reduces the FCC-ee's sensitivity to short-duration fault types in the booster and injector complex and leads to an inherent advantage for performance. To realise this potential, two complementary strategies can be pursued:

1. Extending the collider beam lifetime without top-up injection would provide more time for injector complex or booster systems to recover before a beam abort becomes necessary. Achieving this requires optimisation of the beam optics under various possible failure mechanisms.

2. In parallel with efforts to extend the beam lifetime, the repair duration of remote failures in the booster or injector complex should be minimised. If the recovery time remains below the bunch lifetime

across all energy modes (currently less than 11 minutes), the natural resilience of the system can be fully exploited in all modes. Physics operation may then resume before a beam dump becomes necessary. One promising approach is the automatic reset of failed components, which is applicable to a wide range of fault types, particularly in the RF system. Alternatively, equipping the system with redundant or backup modules that can be rapidly activated to replace faulty components presents a compelling measure.

While some R&D opportunities become less effective as the fault rate decreases, this measure becomes more effective with decreasing fault rate. Figure 6.2 shows the distribution of the durations of remote failures in the booster or injector complex compared to a dashed line representing the lowest burn off lifetime (Z mode) (see Table 2.1). Systems such as power converters, transverse dampers, beam losses, low-energy injectors, and beam instrumentation already exhibit failure durations that, in most cases, fall below the limited beam lifetime. These systems can therefore already benefit from the resilience during stable beams operation. In contrast, system designers of the extraction & beam dump and injection systems should aim to develop solutions that reduce the duration of diagnostics and enable faster remote repair.

6.6 Operation Cycle

Two opportunities already exist within the operation cycle [31] and were explained in Section 2.5.8:

1. **Indefinite Physics:** The potential of this approach is illustrated in Figure 6.3, which shows the distribution of stable beam durations across energy modes. Stable beams rarely last longer than 20 hours due to equipment-related failures. Therefore, extending the maximum

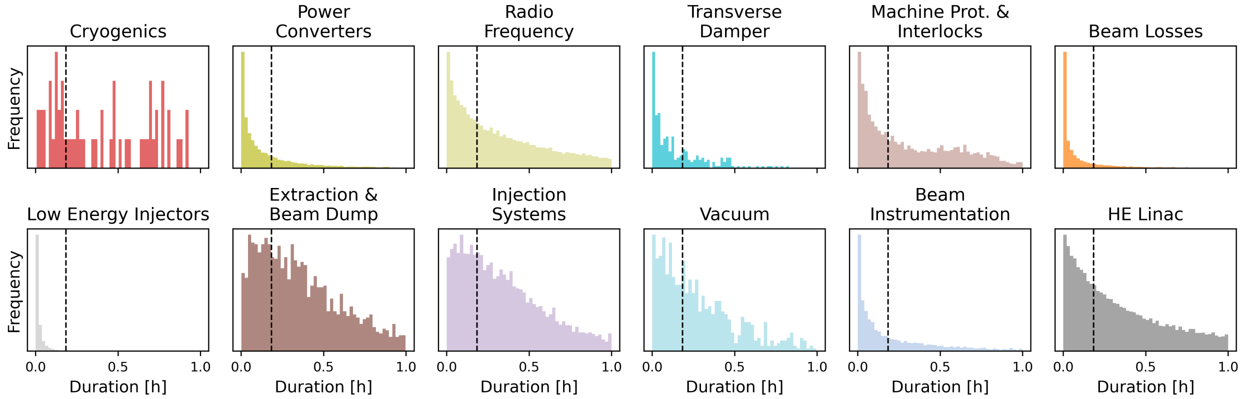

Figure 6.2: „Duration of remote interventions in the booster and injector complex compared to the shortest burn-off lifetime (Z mode) (see Table 2.1). The more the distribution lies to the left of the limit, the more frequently the system benefits from resilience to short faults during stable beams.

physics duration offers minimal benefit unless it is accompanied by a general improvement in machine reliability.

2. **Polarised Bunch Injections:** While this measure has negligible impact on overall availability, it significantly enhances operational efficiency. The benefit is illustrated in Figure 6.4, which shows that in both Z and W modes, the integrated luminosity more than doubles. Through the removal of the polarisation phase at the beginning of each fill, the time spent in and the frequency of reaching stable beam increases significantly. This represents a potentially transformative R&D opportunity to support physics goals under the demanding reliability constraints of the FCC-ee.

   However, it should be noted that this approach interacts with R&D opportunities aimed at reducing failure rates. A lower number of beam dumps leads to fewer repetitions of the operation cycle, and therefore reduces the cumulative benefit of eliminating the 1.5-hour polarisation period.

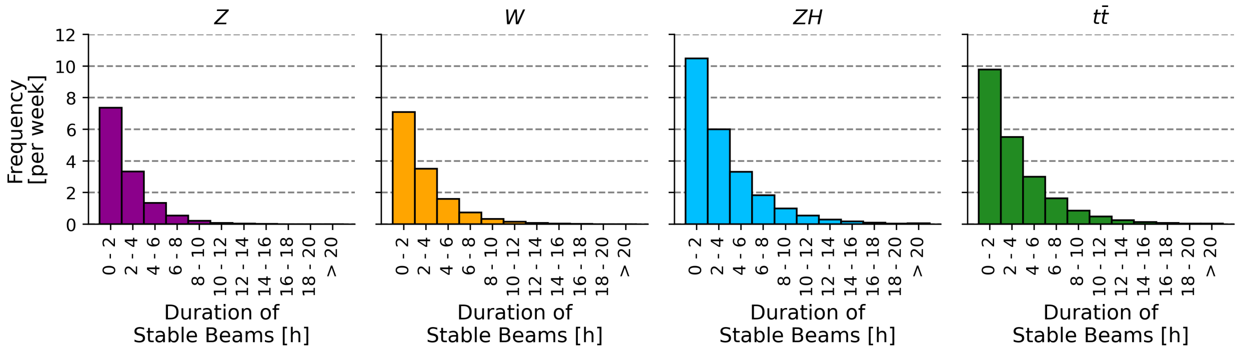

Figure 6.3: Stable beam duration distributions by energy mode. The 20-hour limit defined in the operation cycle for Z and W mode is rarely exceeded, so removing this limit would have negligible impact.

6.7 System Reliability Requirements

For the FCC-ee to meet its performance targets, a machine-wide improvement in reliability is required. However, translating performance shortfalls into quantitative requirements for all systems is complicated by many open design decisions. To nevertheless provide an estimate of what such requirements might look like, a study was conducted in which a maximum acceptable dump rate was defined. The dump rate only accounts for failures that resulted in a complete machine shutdown. Systems exceeding this dump rate threshold had their input MTBF scaled according to the ratio between the actual and the allowed dump rate. Systems with a dump rate below the threshold remained unchanged.

The required scaling factors to reach 100% of the target in Z mode solely by improving system reliability are shown in the bar chart (i) in Figure 6.5. The figure also presents the scaling factors necessary to: achieve similar luminosity gains as with polarised injections only (iii), and to reach the target with polarised injections enabled and reliability improvement (ii). The resulting integrated luminosities for these scaling factors are illustrated for both the baseline operation cycle and polarised injections, along with an approximated curve plotted over the machines overall dump rate.

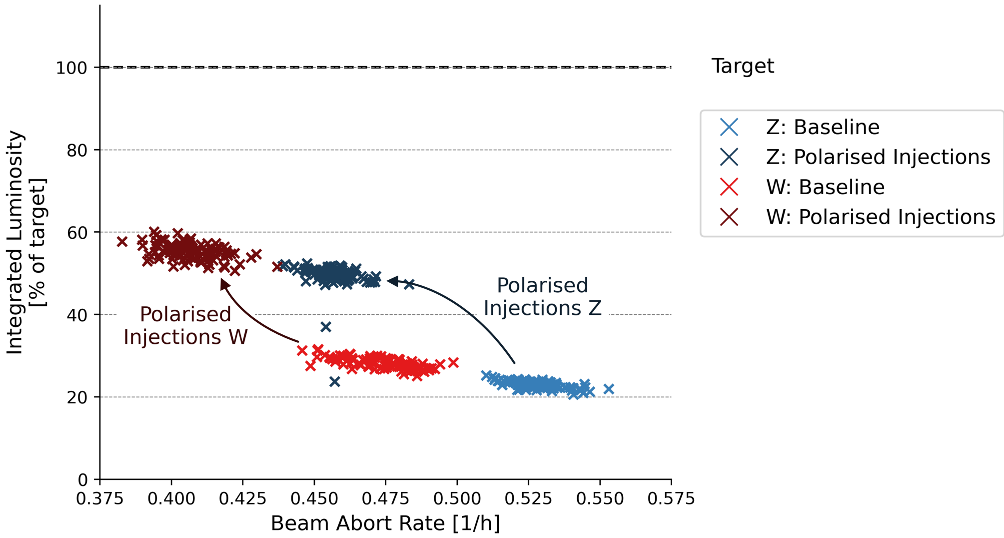

Figure 6.4: Effect of Polarised Injections on integrated luminosity. The achieved luminosity more than doubles.

To reach the full luminosity target in *Z* mode under the baseline cycle, improvements would be required in 18 systems. For example, the RF system would need to improve its reliability by a factor of 61, followed by a factor of 45 for the low-energy injectors and 30 for the power converters. Although these improvement factors may initially appear overly ambitious, the previously discussed measures -— particularly when combined -— have the potential to these requirements. For instance, in *W* mode, a redundancy level of just 4% in the FCC system was sufficient to improve its reliability by more than a factor of 61. The use of polarised injection alone more than halves the required system-level reliability improvements (ii). Furthermore, some systems, such as the low-energy injectors, automatically improve as the overall MTBF increases, since they benefit more frequently from resilience to short failures when the time spent in Stable Beams becomes longer.

Nevertheless, Figure 6.5 peroves, that a machine wide improvement in reliability is essential, higlighting the need for a highly-prioritised reliability and quality support framework in the technical design stage.

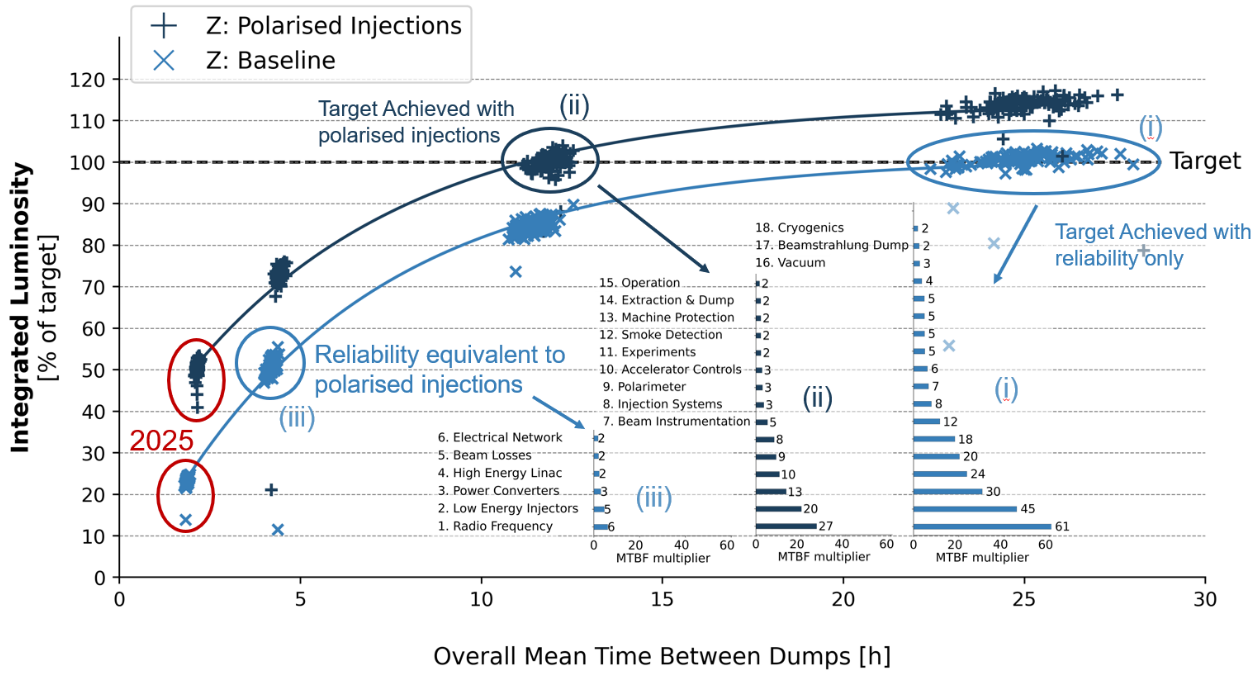

Figure 6.5: Reliability improvement requirements in order to achieve integrated luminosity in the Z mode baseline design with and without polarised injections.

6.8 Summary

This chapter has identified and analysed numerous R&D opportunities aimed at enhancing the FCC-ee's performance. While the initial focus of the FCC-ee is on achieving luminosity goals rather than minimising operational costs, measures that reduce failure rates will naturally have a positive effect on operational expenses. Key subsystem-level strategies were discussed, including redundancy enhancements in the RF and power converter systems. An approach to increase resilience against electrical network glitches was presented, alongside two opportunities to exploit resilience against short faults in the booster and injector complex. Moreover, operational cycle innovations such as polarised bunch injections continue to show strong potential for significantly enhancing integrated luminosity. This illustrated how the developed simulation framework can inform design decisions beyond initial availability predictions. In collaboration with system experts, the design can then be iteratively refined and tested with the simulation environment to optimise machine performance.

Finally, translating the performance shortfalls into quantitative system-level reliability targets reveals ambitious but achievable requirements—especially when multiple mitigation measures are combined. Achieving these targets, however, will require a highly prioritised reliability and quality support framework during the technical design stage. The proposed measures can also serve as guidelines for future projects beyond the FCC-ee.

7 Conclusion

This thesis has addressed the challenge of early availability analysis in the design of high-criticality technical systems. The ability to influence design decisions and implement cost-effective reliability measures is highest during the early phases of development. Historical examples, such as the LHC, have shown that delayed prioritisation of availability can result in long-term operational inefficiencies and high maintenance costs. The FCC-ee, with three times the size and even more ambitious luminosity targets, presents an even greater challenge in terms of availability. Therefore, early-stage availability forecasting is not merely beneficial, but essential. However, traditional availability modelling methods are inadequate at this conceptual stage, necessitating the development of new methodologies capable of influencing the design process before extensive technical details are committed. This led to the central research question of this thesis: How can availability be analysed for a system that has not yet been fully designed and why is such an analysis crucial at this early stage?

7.1 Summary of the Contributions

Chapter 2 established the scientific and technical context for the FCC-ee and reviewed previous availability studies in the FCC-ee and other particle accelerators. It identified gaps in current methodologies for estimating availability during the early design phases. This justified the need for a novel, generalised, and data-driven modelling framework.

Chapter 3 introduced a simulation framework used to estimate availability by extrapolating failure data from existing, operational state-of-the-art systems. This framework enables early forecasting of both overall and subsystem-level performance. The model also supports testing the impact of design choices, operation cycle strategies, and maintenance approaches — thereby offering a versatile foundation for guiding system-level decisions.

Chapter 4 addressed subsystems for which no comparable reference systems exist. A complementary methodology was developed to estimate availability based on system complexity. This method was applied to the beamstrahlung dump system to obtain a first availability placeholder and to complete the FCC-ee baseline availability model. In doing so, methodology supports systematic risk identification and the prioritisation of critical subsystems

Chapter 5 applied the simulation framework to the FCC-ee, revealing that the current baseline design falls significantly short of its luminosity targets across all energy modes. A breakdown of the contributions by subsystems highlighted the most critical ones, such as the RF system, electrical network, and power converters. Furthermore, concerns regarding operational costs were identified. These findings quantify and prioritise the need for an availability-focused design between the accelerator's top-level systems. They also indicate that measures to increase overall machine efficiency — especially in lower-energy modes — are necessary. This includes optimisation of the operation cycle and maintenance strategies.

Chapter 6 proposed R&D opportunities that could improve availability, operational efficiency and integrated luminosity in the machine. For example, injecting already polarised bunches could more than double the achieved luminosity. Combined with the additional R&D opportunities proposed, the performance targets are expected to be achievable.

The key contributions of the thesis are as follows:

1. A generalized framework for availability forecasting of conceptual systems that extrapolates fault data of existing operational systems was provided. This enables early-stage availability analysis -— and consequently, availability assurance -— for the FCC-ee as well as for other large-scale future engineering projects.

2. A formulation to anticipate availability in systems with neither rigorous designs nor comparable failure data was given. It can stand on its own or support a larger engineering hierarchy as in this thesis. Applied to the Beamstrahlung dump system, it completed the FCC-ee availability forecast.

3. Using the above formulations in a Monte Carlo simulation various shortfalls and most critical systems in the FCC-ee could be identified, enabling targeted and well-founded R&D opportunities to improve operational performance.

4. Various solutions to the above shortcomings were proposed as R&D opportunities, whose impact on accelerator performance was simulated and quantified. Pre-polarised injections, in particular, showed promising improvements and, in combination with systems designed for reliability, could enable the accelerator to meet its performance targets. This illustrates how the developed simulation framework can effectively support informed design decisions at an early stage.

7.2 Impact and Importance of Early Availability Forecasting

In summary, this work has demonstrated that early and structured availability analysis is both feasible and essential for conceptual systems like the FCC-ee. The developed methodologies allow for actionable insights well before detailed technical designs are fixed, enabling well-placed investments to positively impact system performance and long-term operational sustainability. Without such an analysis, the FCC-ee could face prohibitively high operational costs and risk failing to meet its ambitious performance goals. As each year of development progresses, flexibility decreases and the cost of change increases — reinforcing the need to integrate availability considerations from the very beginning of the design process.

The FCC-ee will continue to benefit from this availability model throughout its entire design process. To maximise this benefit, clear and achievable reliability and availability targets must be defined and communicated to the system design teams. As designs mature, progress must be systematically tracked and updated regarding availability, efficiency, integrated luminosity, and cost factors. In parallel, systems — particularly the critical ones — should be modelled in more detail to identify high-failure components and quick-win opportunities. This will facilitate a transition toward modular, redundant, and hot-swappable system architectures, and requires continued close collaboration with the relevant design teams. For example, RF fault data will be analysed in collaboration with system experts to identify and categorise underlying failure mechanisms. This will enable the development of a multi-level availability model and support the simulation of R&D strategies as outlined in Section 6.1.

7.3 Paradigm Shift in Accelerator Design

This thesis also demonstrates that the FCC-ee systems cannot simply follow the same paradigm as previous accelerators; a paradigm shift in design methodology using for reliability designed systems is needed. To enable this shift, the technical design phase of the FCC-ee must be guided by a systematic and highly prioritised reliability and quality assurance framework which should endure through construction, commissioning and operation to ensure a consistent and integrated approach across all systems — encompassing validation and testing, tools and training, diagnostics, health monitoring and fault prediction strategies as well as fault tracking and analytics solutions. This will support the process from design through to installation and operation of the FCC-ee, ensuring that resources committed to lifecycle reliability are effectively and efficiently spent.

7.4 Outlook

The general framework developed in this thesis is designed to be extensible and can be enhanced in the future to incorporate additional layers of detail. By refining the model – for example, through the separate treatment of parent and child faults or by explicitly modelling redundancy already present in the reference system, its predictive capability regarding key performance indicators can be significantly improved. This, in turn, would strengthen its capacity to evaluate and prioritise R&D opportunities. This study also highlighted the critical importance of high-quality failure data from existing machines. Without structured and detailed fault documentation, availability models remain highly uncertain. Future projects would greatly benefit from systematic fault logging standards across all subsystems of any complex machine.

7.5 Publications

The research contributions of this thesis have also contributed to the following publications:

- M. Benedikt, F. Zimmermann, B. Auchmann, et al., Future Circular Collider Feasibility Study: Volume 2 – Accelerators, technical infrastructure and safety. CERN, Geneva, Switzerland, Rep. CERN-FCC-ACC-2025-0004, 2025.

- J. W. Heron, H. A. Dostmann, L. Felsberger, D. Wollmann and J. Uythoven, "Availability Assurance in the Future Circular Electron-Positron Collider (FCC-ee)," in *Proc. 16th International Particle Accelerator Conference* (*IPAC'25*), 2025.

- H. A. Dostmann, J. W. Heron, L. Felsberger, D. Wollmann and J. Uythoven, "Qualitative Availability Forecasting in the Beamstrahlung Dump for the Future Circular Electron-Positron Collider (FCC-ee)", CERN Document Server, pending publication.

- J. W. Heron, H. A. Dostmann, L. Felsberger, D. Wollmann and J. Uythoven, "Availability, Efficiency and Luminosity: Rising to the Challenge in the Future Circular Electron-Positron Collider (FCC-ee)", Physical Review Accelerators and Beams (PRAB) journal, (submitted, pending review).

and was presented at the following conferences:

- FCC Week 2025, Vienna, Austria, May 19–23, 2025. (presented by H. A. Dostmann)

- IPAC'25 - the 16th International Particle Accelerator Conference, Taipei, Taiwan, June 1-6, 2025. (presented by proxy)

Bibliography

[1] "DIN EN 60300-3-1:2004-10 – Zuverlässigkeitsmanagement – Teil 3-1: Anwendung von Zuverlässigkeits-, Verfügbarkeits-, Instandhaltbarkeits- und Sicherheitsmanagement (RAMS) – Leitfaden zur Anwendung bei Systemen," Deutsches Institut für Normung e.V. (DIN), Tech. Rep., 2004, Berlin: Beuth Verlag.

[2] M. A. Soomro and X. Zhang, "An analytical review on transportation public private partnerships failures," *International Journal of Sustainable Construction Engineering & Technology*, 2011, Accessed: 2025-05-28, [Online]. Available: https://www.researchgate.net/publication/267364741_AN_ANALYTICAL_REVIEW_ON_TRANSPORTATION_PUBLIC_PRIVATE_PARTNERSHIPS_FAILURES.

[3] R. Carvel, *Fire dynamics during the channel tunnel fires*, https://www.academia.edu/73046156/Fire_Dynamics_During_the_Channel_Tunnel_Fires, [Online]. Accessed: 2025-05-28, 2010.

[4] N. Srivatsav, J. Joys, Z. Wang, K. Safty, and R. Kapoor, *The chunnel project: Management retrospective*, https://www.academia.edu/7035250/The_Chunnel_Project_Management_Retrospective, [Online]. Accessed: 2025-05-28, n.d.

[5] F. Zimmermann, A. Apollonio, M. Benedikt, O. Brunner, Y. Funakoshi, C. Milardi, S. Myers, A. Niemi, K. Oide, Q. Qin, J. Seeman, and J. Wenninger, "FCC-ee Operation Model, Availability & Performance," WEPAB03, 2019, DOI: 10.18429/JACoW-eeFACT2018-WEPAB03.

[6] CERN, *Accelerator Fault Tracking*, https://aft.cern.ch/.

[7] CERN. “Feasibility study.” Accessed: 2025-06-10. CERN. (2025), [Online]. Available: https://fcc.web.cern.ch/overview.

[8] B. Bertsche and M. Dazer, *Zuverlässigkeit im Fahrzeug- und Maschinenbau: Ermittlung von Bauteil- und System-Zuverlässigkeiten*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2022, ISBN: 978-3-662-65024-0, DOI: 10.1007/978-3-662-65024-0_13.

[9] A. Birolini, *Reliability Engineering: Theory and Practice*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2017, ISBN: 978-3-662-54208-8, DOI: 10.1007/978-3-662-54209-5.

[10] L. Felsberger, P. Alexaki, A. Asko, N. Charitonidis, J. Heron, B. Mikulec, B. Rae, J. Uythoven, and D. Wollmann, “Report Accelerator Complex Availability 2024,” Tech. Rep., 2025, CERN Document Server, [Online]. Available: https://cds.cern.ch/record/2925562.

[11] CERN. “Cern’s accelerator complex.” Accessed: 2025-06-10. CERN. (2025), [Online]. Available: https://www.home.cern/science/accelerators/accelerator-complex.

[12] L. Felsberger, A. Apollonio, T. Cartier-Michaud, E. Montesinos, J. C. Oliveira, and J. Uythoven, “Availability Modeling of the Solid-State Power Amplifiers for the CERN SPS RF Upgrade,” *12th International Particle Accelerator Conference*, vol. 2021, pp. 2308–2311, 2021, Conference location: Campinas, SP, Brazil, DOI: 10.18429/JACoW-IPAC2021-TUPAB345.

[13] O. Rey Orozco, A. Apollonio, M. Jonker, and J. Uythoven, “Availability Allocation to Particle Accelerators Subsystems by Complexity Criteria,” *9th International Particle Accelerator Conference*, WEPAF075, 2018, Conference location: Vancouver, Canada, DOI: 10.18429/JACoW-IPAC2018-WEPAF075.

[14] O. Rey Orozko, A. Apollonio, and M. Jonker, "Methodology Applied for Dependability Studies on the Compact Linear Collider," TUPIK100, 2017, DOI: 10.18429/JACoW-IPAC2017-TUPIK100.

[15] O. R. Orozco, *Availability studies for clic*, Presentation at the CLIC availability studies meeting, Slides available at: https://indico.cern.ch/event/739111/, Jun. 2018.

[16] C. Accettura *et al.*, "The Muon Collider," The International Muon Collider Collaboration, Tech. Rep., 2025, [Online]. Available: https://edms.cern.ch/ui/#!master/navigator/document?D:101765627:101765627:subDocs.

[17] M. Benedikt *et al.*, "Future Circular Collider Feasibility Study Report Volume 1: Physics and Experiments," CERN, Geneva, Tech. Rep., 2025, DOI: 10.17181/CERN.9DKX.TDH9, arXiv: 2505.00272.

[18] M. Benedikt *et al.*, "Future Circular Collider Feasibility Study Report Volume 2: Accelerators, technical infrastructure and safety," CERN, Geneva, Tech. Rep., 2025, DOI: 10.17181/CERN.EBAY.7W4X, arXiv: 2505.00274.

[19] M. Benedikt *et al.*, "Future Circular Collider Feasibility Study Report Volume 3: Civil Engineering, Implementation and Sustainability," CERN, Geneva, Tech. Rep., 2025, DOI: 10.17181/CERN.I26X.V4VF, arXiv: 2505.00273.

[20] *Future circular collider*, https://home.cern/science/accelerators/future-circular-collider, Accessed: 2025-06-03. CERN, CERN, 2025.

[21] Institute of High Energy Physics, *Top-up operation mode successful on BSRF*, https://english.ihep.cas.cn/se/fs/bb/naa/201703/t20170324_289547.html, Accessed: 2025-06-24, 2025.

[22] A. Blondel *et al.*, *Polarization and centre-of-mass energy calibration at FCC-ee*, 2019, arXiv: 1909.12245 [physics.acc-ph], [Online]. Available: https://arxiv.org/abs/1909.12245.

[23] J. Heron, L. Felsberger, C. Hernalsteens, C. Wiesner, D. Wollmann, J. Uythoven, and F. Rodriguez Mateos, "Machine Protection and Availability in the FCC-ee. FCC-ee: Future Circular Lepton Collider," 2023, CERN Document Server, [Online]. Available: https://cds.cern.ch/record/2880188.

[24] W. Herr and B. Muratori, "Concept of luminosity," 2006, DOI: 10.5170/CERN-2006-002.361.

[25] K. Yokoya and P. Chen, "Beam-beam phenomena in linear colliders," *Lecture Notes in Physics*, May 1995, DOI: 10.1007/3-540-55250-2_37.

[26] M. Boscolo and A. Ciarma, "Characterization of the beamstrahlung radiation at the future high-energy circular collider," *Phys. Rev. Accel. Beams*, vol. 26, p. 111 002, 11 Nov. 2023, DOI: 10.1103/PhysRevAccelBeams.26.111002.

[27] I. Karpov, A. Butterworth, and O. Brunner, *Private Communication*, September 2024.

[28] A. Apollonio, "Fcc availability studies," in *Proc. of FCC Week 2016*, 2016.

[29] J. Heron, L. Felsberger, D. Wollmann, J. Uythoven, and F. Rodriguez Mateos, "Availability Targets Scaled According to Assurance Complexity in the FCC-ee.," 2023, [Online]. Available: https://cds.cern.ch/record/2880189.

[30] A. Kanareykin, "Ferroelectric fast reactive tuner for srf cavities - material properties and its applications," in *Proc. of FCC Week 2023*, 2023.

[31] J. Heron, M. Blaszkiewicz, L. Felsberger, J. Uythoven, and D. Wollmann, "Availability and Luminosity in the Future Circular Electron-Positron Collider (FCC-ee)," *JACoW IPAC*, vol. 2024, WEBN3, 2024, DOI: 10.18429/JACoW-IPAC2024-WEBN3.

[32] B. Klein, "Numerische analyse von gemischten ausfallverteilungen in der zuverlässigkeitstechnik," Available at https://elib.uni-stuttgart.de/bitstreams/c7712d1f-bd95-4ba5-8ae0-ca4c93286b67/download, PhD thesis, University of Stuttgart, Stuttgart, DE, Jul. 2013.

[33] CERN. "Future circular collider." Accessed: 2025-06-10. CERN. (2025), [Online]. Available: https://home.cern/science/accelerators/future-circular-collider.

[34] A. Apollonio, "Conclusions on availability studies for the future circular collider," in *Proc. of FCC Week 2018*, 2018.

[35] M. Calviani *et al.*, "Beamstrahlung dump concepts, design, considerations & r&d roadmap," in *Proc. of FCC Week 2024*, 2024.

[36] A. Frasca, A. Ciarma, A. Lechner, A. Perillo Marcone, B. Humann, C. Järmyr Eriksson, C. Welsch, D. Calzolari, G. Lerner, M. Boscolo, M. Calviani, N. Kumar, and R. Kersevan, "Energy deposition and radiation level studies for the FCC-ee experimental insertions," *JACoW IPAC*, vol. 2024, TUPC66, 2024, DOI: 10.18429/JACoW-IPAC2024-TUPC66.

[37] A. Lechner, *Private Communication*, September 2024.

[38] A. P. Marcone, *Private Communication*, September 2024.

[39] A. Francia *et al.*, "Design and Early Operation of a new-generation Internal Beam Dump for CERN's Super Proton Synchrotron," *Phys. Rev. Accel. Beams*, vol. 27, p. 043 001, 4 Apr. 2024, DOI: 10.1103/PhysRevAccelBeams.27.043001.

[40] C. Marcel, "Fcc electrical grid and infrastructure update," in *Proc. of FCC Week 2025*, 2025.

[41] M. C. Moratalla and M. Parodi, "New technologies for electrical transmission and distribution in fcc," in *Proc. of FCC Week 2025*, 2025.

[42] J. J. Breedlove, M. V. Zagarola, G. F. Nellis, J. A. Gibbon, F. X. Dolan, and W. L. Swift, "Life and reliability characteristics of turbo-brayton coolers," in *Cryocoolers 11*, R. G. Ross, Ed. Boston, MA: Springer US, 2002, pp. 489–497, ISBN: 978-0-306-47112-4, DOI: 10.1007/0-306-47112-4_62.

Appendix A Publication 1

https://cds.cern.ch/record/2928793/files/2505.00274.pdf

Appendix B Publication 2

AVAILABILITY ASSURANCE IN THE FUTURE CIRCULAR ELECTRON-POSITRON COLLIDER (FCC-ee)

J. W. Heron*, H. A. Dostmann, L. Felsberger, D. Wollmann, J. Uythoven,
European Organisation for Nuclear Research (CERN), Geneva, Switzerland

Abstract

The Future Circular Electron-Positron Collider (FCC-ee) is CERN's leading proposal for the next generation of energy-frontier particle accelerators. At 91 km long, it is ambitious in size, complexity and technical objectives. Availability is a main challenge. This paper presents results from a Monte Carlo simulation that extrapolates reliability and maintainability from systems in current working accelerators to the FCC-ee. Significant integrated luminosity shortfall appears in all energy modes due to low availability and operational efficiency. The primary contributors are highlighted, exposing several compelling R&D opportunities.

INTRODUCTION

The Future Circular Collider (FCC) would be the largest synchrotron ever built. Two collision stages are planned: First, leptons (FCC-ee) from ~2045; then hadrons (FCC-hh) from ~2070. The FCC-ee will operate at four energy modes: four years at the Z pole; two years the WW production threshold; three years at the ZH production peak; and five years at the top/anti-top $t\bar{t}$ production threshold.

To meet physics objectives, the machine must be operational for a minimum amount of time in each energy mode. 185 days are scheduled for physics operation each year. Availability is the proportion of these days where the machine can deliver beam, as opposed to being under repair. The global availability target for the whole machine is 80 % [1, pp.79]. For comparison, the 27 km Large Hadron Collider (LHC) averaged 72 % availability in physics production 2015-2024 [2]. Additional challenges in the FCC-ee, like size, complexity and ambitious technical objectives, make availability one of the main challenges to its deliverables.

This paper presents a Monte Carlo simulation to study availability, operational efficiency and integrated luminosity performance in the FCC-ee, by extrapolating the reliability and maintainability of its top-level systems from state-of-the-art representations in existing working accelerators.

SIMULATION MODEL

FCC-ee availability is modelled in an enhanced Monte Carlo simulation environment using *AvailSim4* [3], a discrete-event Monte Carlo tool developed in-house at CERN for analysing availability and reliability in complex engineering hierarchies. This deconstructs availability into the accelerator's top-level systems, capturing probabilistic failure and repair durations and various performance features.

Energy Calibration

A principle task for the FCC-ee is ultra-precise measurement of electroweak (Z and WW) observables, for which accurately determined collision energy is key. This involves beam energy calibration every 10-15 minutes using non-colliding polarised pilot bunches, which circulate simultaneously with the colliding bunches. Beam energy is measured by Resonant Depolarisation (RDP), where a kicker magnet's excitation frequency is adjusted until polarisation vanishes.

Pilot bunches are polarised in the collider at the start of every fill using wiggler magnets for 90 min. Wigglers are then turned off to inject colliding bunches. Pilot bunches have a combined Touschek and beam-gas lifetime less than 20 h [4], after which the beam must be dumped and re-filled.

In ZH and $t\bar{t}$ modes, the energy spread makes RDP impossible. Measurement is instead achieved by observing collisions at the Interaction Point (IP). This is less accurate, but removes the need for polarisation at the start of every fill. Further, with top-up injection, physics can continue until beam dump due to machine fault or scheduled end of fill.

Operation Cycle

FCC-ee operation cycles are shown in Figure 1:

1. *Set Up:* Equipment is prepared for injection.
2. *Pilot Bunch Injection:* Pilots are injected and equipment/optics is adjusted.
3. *Polarisation:* Z and WW modes only. Wigglers are turned on to polarise pilot bunches ready for RDP.
4. *Fill:* Wigglers are turned off and the main colliding bunches are injected. Duration t_f = 2-7 min.
5. *Adjust:* Final adjustments to equipment and beam.
6. *Physics:* Collisions begin. With top-up injection, flat luminosity can be maintained.
7. *Burn Off:* If the injector complex fails, luminosity decays with lifetime τ = 10-15 min, after which the beam must be dumped as it quickly becomes unstable.
8. *Down for Repair:* Down time on equipment failure.

Luminosity

Luminosity is produced only in Physics or Burn Off phases. With time t spent in these phases, integrated luminosity is

$$L_{int} = \begin{cases} N_{IP} L t, & t \in \text{physics} \\ N_{IP} L \tau (1 - e^{-\frac{t}{\tau}}), & t \in \text{burn off} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

Where $N_{IP} = 4$ is the number of IPs and L is nominal instantaneous luminosity ($1.25\text{-}141 \times 10^{34}$ cm^{2}s^{-1} depending on energy mode). In the first two and one years of Z and $t\bar{t}$ operation, respectively, a 50 % and 65 % reduced luminosity is applied for machine commissioning and beam tuning.

* jack.heron@cern.ch

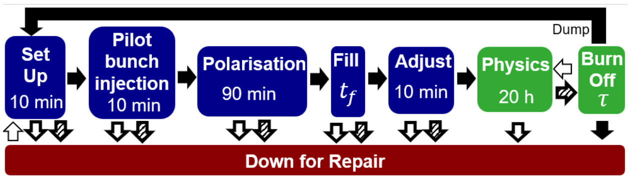

(a) Z and WW modes

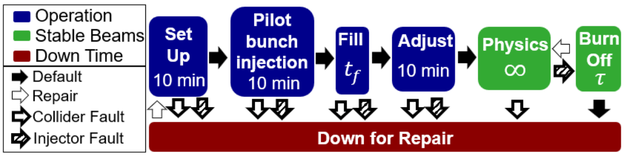

(b) ZH and $t\bar{t}$ modes

Figure 1: FCC-ee operation cycles.

Contributing Systems

In collaboration with relevant system experts, each main constituent system in the FCC-ee is approximated for availability. Fault data from a representative ("rep") system in a current working accelerator is extrapolated to the equivalent system in the FCC-ee using a generalised framework.

1. Where possible, the rep system is deconstructed into subcomponents to achieve better granularity.
2. Fault mechanisms are categorized according to:
   - Remote Repairs can be completed from the control room, e.g. by re-setting or by-passing the failed component.
   - Human Repairs require on-site intervention by personnel and an approach time to the site of the fault (20-60 min.)
   - Surface Equipment can be accessed for repair immediately when a technician arrives.
   - Tunnel Equipment cannot be accessed before a 1-24 h radiation cool-down time, depending on location.
3. Mean Time Between Failures (MTBF) and Mean Time to Repair (MTTR) is extracted. If possible, a bespoke redundancy is formulated. MTTR includes down time from any child faults in other systems that result directly from the root cause.
4. MTBF and MTTR are extrapolated to the FCC-ee:
   - MTBF scales inversely with the relative number of components between rep system and FCC-ee.
   - MTTR is left unchanged. An approach time and/or cool down time is added to human repairs.

Failure data came from the Accelerator Fault Tracking tool at CERN [5], unless otherwise stated. Some system specifics:

Radio Frequency (RF): The LHC's 16 superconducting cavities were extrapolated to the FCC-ee's 376-1120. Redundancy varies with energy mode: In Z and WW, the beam is dumped if more than one and six cavities trip, respectively. There is no redundancy in the booster. In ZH and $t\bar{t}$, 10% of cavities are redundant in both collider and booster.

Electrical Network: The LHC has one connection to the external grid. The FCC-ee will have three (PA, PD and PH) connected to two transmission lines. External network perturbations are applied at twice the rate of LHC. Internal equipment faults are scaled with the number of substations at each connection. A redundant link between PA and PD can restore power in ~1 hour.

Power Converters: Comparable LHC converters are chosen, where an equivalent redundancy is assumed.

Injector Complex: The High Energy (HE) Linac is modelled from Linac Coherent Light Source (LCLS) at SLAC [6]. The low energy injectors from the SuperKEK-B e^-/e^+ source, injector linac and damping ring [7].

Smoke Detection: Faults in the Super Proton Synchrotron (SPS) long-range aspirating smoke detectors are scaled from the 7 km tunnel to the FCC-ee's 91 km.

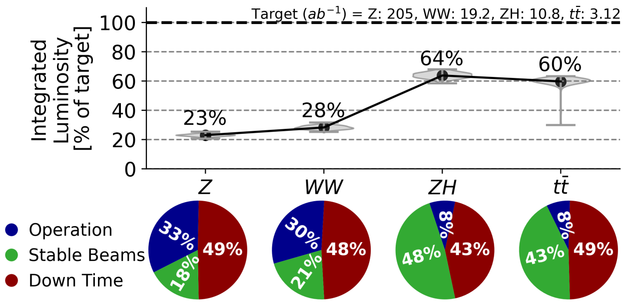

Figure 2: FCC-ee: Predicted integrated luminosity as percent of target in each energy mode, extrapolating from the availability of systems in current working accelerators. The pie charts show time spent in operation cycle phases.

Repair Timeline

A repair timeline similar to the LHC was assumed: Following a dump, operators try for one hour to diagnose and restore operation remotely. If operation cannot be restored, only then are personnel sourced for human repairs, which will all finish including redundancies.

RESULTS

Each energy mode was simulated for its full run (four years in Z, two years in WW, etc.) over 100 iterations. Predicted integrated luminosity and availability in each energy mode is shown in Figure 2. Integrated luminosity is significantly below target in all energy modes. The pie charts show percent of time spent in operation cycle phases. Stable beams composes 18-21 % of scheduled physics in Z and WW modes due to the polarisation phase, which strongly penalises the recovery time following beam dump.

The breakdown by system is shown in Figure 3. The RF has highest unavailability in Z due to the strict technical limit on redundancy. In WW, the six redundant cavities eliminates collider unavailability. The booster is significant in both energy modes. In ZH and $t\bar{t}$ the 10% redundancy leads to zero RF down time, as this is sufficient reserve for cavity repair to wait for down time in other systems.

The Electrical Network has high contribution due to child faults in other systems, which cannot be compensated by its redundant link. Power Converters have high contribution due to the large number of sextupoles (powered in quartets) and corrector magnets (powered individually) in the current

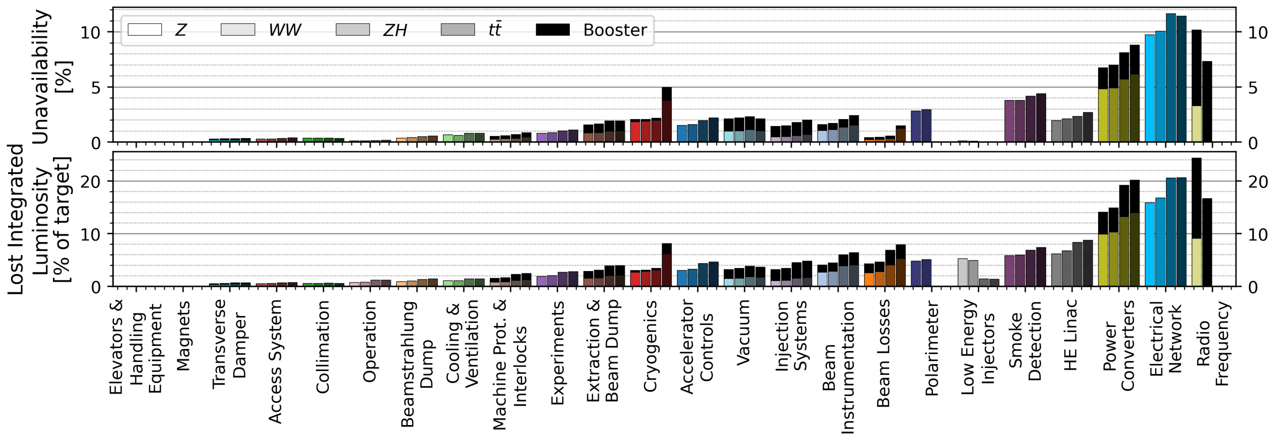

Figure 3: Unavailability and lost luminosity contribution from each system, ordered by Z mode lost luminosity contribution. Lost luminosity is normalised to the target in each energy mode.

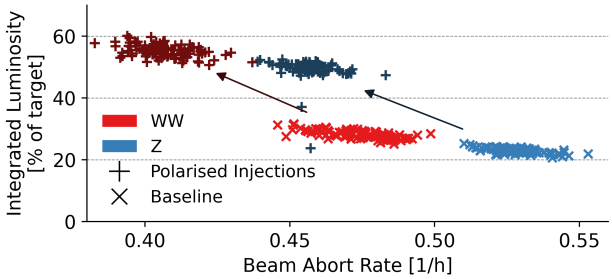

Figure 4: Injection of polarised bunches stands to double achieved integrated luminosity in Z and WW modes.

optics [1, pp.177, 285]. In Smoke Detection, the number of alarms combined with mandatory on-site safety evaluation in the event of detected fire leads to significant down time.

In the Low Energy Injectors, high implemented redundancy allows for short but nevertheless frequent blocking durations. It therefore has low contribution to down time, but high contribution to lost luminosity due to the turnaround time to restore stable beams.

R&D OPPORTUNITIES

To improve operational efficiency in Z and WW modes, a compelling solution is to inject already polarised bunches. This would remove the polarisation phase from the electroweak operation cycle. As shown in Figure 4, this stands to more than double achieved integrated luminosity.

In the RF, additional redundancy should be further pursued due to the high possible gains, particularly in the booster. The klystron gallery will be accessible to personnel during collider operation. If high-failing auxiliary systems can be relocated there and designed to be redundant or hot-swappable, significant gains could be achieved.

The electrical network suffers from many child faults in downstream systems. Two parallel streams must be researched: First, detailed design specifications for resilience to electrical perturbation to be issued to relevant design teams. Second, solutions to dampen these deviations, for example with unified power flow controllers.

For power converters, the optics configuration must consider reducing the number of sextupole and corrector magnets, as well as how to preserve the beam if converters fail. In parallel, power converters must be made more reliable than the LHC.

In the injector complex and booster, a fault during stable beams must last longer than the burn off lifetime (10-15 min.) in order to dump the beam, creating natural resilience to short fault types. That would be better exploited if the burn off lifetime was longer. Measures to bring the recovery time below the beam lifetime would also be powerful.

To achieve physics objectives in the FCC-ee, a machine-wide improvement in reliability is required. Starting with the largest contributing systems, deeper research to deconstruct and isolate problematic fault types must support a redundant, modular and hot-swappable design paradigm. A rigorous reliability/availability support framework must be implemented into the design process for success in the upcoming technical design phase.

CONCLUSION

This paper showcases results from an enhanced Monte Carlo simulation environment to study availability, operational efficiency and integrated luminosity performance in the FCC-ee. This extrapolates the reliability and maintainability of its top-level systems from state-of-the-art representations in existing working accelerators.

Results expose significant shortfall in predicted integrated luminosity at all energy modes, but especially in Z and WW due to low operational efficiency. The most challenging systems for lost luminosity are the RF, electrical network, power converters and injector complex.

Various R&D opportunities exist. The most compelling for Z and WW is the injection of already polarised bunches, which could more than double achieved integrated luminosity. Meanwhile, a machine-wide reliability improvement must be given high priority in the technical design. A rigorous and systematic reliability/availability support framework must persist to inform and guide engineering decisions.

REFERENCES

[1] M. Benedikt *et al.*, "Future Circular Collider Feasibility Study Report Volume 2: Accelerators, technical infrastructure and safety," CERN, Tech. Rep., 2025. doi:10.17181/CERN.EBAY.7W4X

[2] L. Felsberger *et al.*, "Report Accelerator Complex Availability 2024," 2025. https://cds.cern.ch/record/2925562

[3] M. Blaszkiewicz *et al.*, "AvailSim4: Open Source Framework For Availability and Reliability Simulations," en, in *European Safety and Reliability Conference (ESREL); Advances in Reliability, Safety and Security, Part 4*, 2024, pp. 29–38.

[4] A. Blondel *et al.*, "Polarization and Centre-of-mass Energy Calibration at FCC-ee," *arXiv*, 2019, http://arxiv.org/abs/1909.12245.

[5] C. Roderick, L. Burdzanowski, D. Martin Anido, S. Pade, and P. Wilk, "Accelerator fault tracking at CERN," TUPHA013, 2018. doi:10.18429/JACoW-ICALEPCS2017-TUPHA013

[6] W. Allen, A. Brachman, W. Colocho, M. Stanek, and J. Warren, "Availability Performance and Considerations for LCLS X-Ray FEL at SLAC," Tech. Rep., 2011, SLAC-PUB–14422, p. 4. http://inis.iaea.org/search/search.aspx?orig_q=RN:42098221

[7] K. Furukawa *et al.*, "Achievement of 200,000 hours of operation at KEK 7-GeV electron 4-GeV positron injector linac," en, *Journal of Physics: Conference Series*, vol. 2420, no. 1, p. 012 021, 2023. doi:10.1088/1742-6596/2420/1/012021

Appendix C Publication 3

Availability, Efficiency and Luminosity: Rising to the Challenge in the Future Circular Electron-Positron Collider (FCC-ee)

John W. Heron, Hannah A. Dostmann, Lukas Felsberger, Daniel Wollmann, and Jan Uythoven
European Organisation for Nuclear Research (CERN), Geneva, Switzerland
(Dated: June 30, 2025)

The Future Circular Electron-Positron Collider (FCC-ee) is CERN's leading proposal for the next generation of energy-frontier particle accelerators. To reach integrated luminosity goals, it aims to be operational for minimum 80 % over the scheduled 185 physics days each year. For comparison, the Large Hadron Collider (LHC) – the world's largest and most complex particle accelerator, but still a factor three smaller than the FCC – averaged 72 % in 2015-2024. Characteristics of the FCC-ee relating to size, complexity and ambitious technical objectives make availability one of the main challenges to its physics deliverables. This paper presents results from an enhanced Monte Carlo simulation environment that extrapolates reliability and availability performance of all major systems and subsystems in the FCC-ee from current working accelerators. Shortfalls in integrated luminosity are identified in all energy modes in the current baseline design, compounded in the electroweak sector by low operational efficiency. The primary contributors to unavailability and lost luminosity are highlighted. Compelling R&D opportunities to boost availability, efficiency and integrated luminosity are proposed and quantified. The most powerful is the direct injection of polarised bunches, which stands to double integrated luminosity in the lower energy modes. Finally, a parameter scan for reliability of the highest-failing systems shows that dramatic machine-wide improvement in dump rate is required under the current baseline assumptions. To this end, a systematic reliability and availability assurance framework is proposed to support the accelerator design process from conception through to installation and operation.

I. INTRODUCTION

At 91 km circumference, the Future Circular Collider (FCC) would be the largest particle collider ever built. Collisions are planned in two stages: First, leptons (FCC-ee) starting ∼2045; then hadrons (FCC-hh) starting ∼2070. The FCC-ee will operate at four baseline centre-of-mass energies, corresponding to four years at the Z pole, two years the WW pair production threshold, three years at the ZH production peak, and five years at the top/anti-top $t\bar{t}$ production threshold.

The machine must be operational for a minimum amount of time in each energy mode in order to meet its physics objectives. There are 185 days scheduled for physics operation each year, in which luminosity can be produced and collisions can be observed. Availability is the proportion of physics days where the machine can deliver beam, as opposed to being in down time or under repair. The global availability target for the whole machine is minimum 80 % in all energy modes [1]. For comparison, the 27 km Large Hadron Collider (LHC) was available for 72 % of physics production in 2015-2024 [2]. Additional challenges in the FCC-ee, like its size, complexity and ambitious technical objectives, make availability one of the main challenges to its physics deliverables.

This paper presents an enhanced Monte Carlo simulation environment to study availability, operational efficiency and integrated luminosity performance in the FCC-ee, by extrapolating the reliability and maintainability of its top-level systems from state-of-the-art representations in existing working accelerators. Results expose significant shortfall in achieved integrated luminosity at all energy modes, but especially at lower energies (Z and WW) due to low operational efficiency. This insight is both timely and valuable at the current early stage of accelerator development. Further, the simulation allows for three major advantages for the upcoming technical design:

a. Performance Benchmarking: The difference between the availability of current working systems and that required to reach physics goals in the future machine is quantified. The systems and components with highest contribution to down time and lost luminosity reveal tangible areas for improvement that must be addressed. The simulation can be used to benchmark gains in performance as the accelerator design progresses.

b. Opportunity Assessment: The FCC-ee research space includes numerous proposals for alternative designs spanning diverse fields and applications. Each alternative generally requires investment in Research and Development (R&D) prior to implementation. The simulation presented in this paper can quantify and compare these proposals in terms of machine performance, i.e. availability, operational efficiency and integrated luminosity, but also several operational cost factors. This has already exposed and quantified some unexpected effects, interrelations and nuances that deviate from the popular zeitgeist. In this way, the simulation can inform the overall design strategy.

c. Design Optimisation: To meet physics objectives in the FCC-ee, its systems must be designed according to unprecedented reliability and maintainability requirements. This superposes a significant challenge on design teams who must already balance other requirements such as performance, lifecycle cost, sustainability and safety. The simulation presented in this paper is further capable of optimising technical solutions over relevant cost constraints through parameter-scans. It can therefore provide support beyond conceptual stages in the design process, ensuring that resources committed to reliability consolidation are efficiently spent.

This paper presents essential insights gained from the FCC-ee availability study to date. This begins in Section II with a description of the features and assumptions used in simulation to accurately reflect baseline operation of the machine. Section III then presents the forecasted physics performance of the baseline design, identifying the most challenging systems for availability and lost luminosity. Various R&D proposals to promote integrated luminosity of the machine are then modelled and discussed in Section IV, the most compelling of which is the inclusion of a polarisation ring to inject already polarised bunches into the collider. Looking ahead, a parameter scan for component reliability in the most challenging systems is conducted in Section V, which reflects the scale of the availability challenge at hand. Finally, to rise to the challenge, Section VI describes a systematic reliability/availability assurance framework to be implemented through the technical design process. Section VII then concludes the topic.

II. SIMULATION MODEL

FCC-ee availability is modelled in an enhanced Monte Carlo simulation environment using *AvailSim4* [3], a discrete-event Monte Carlo tool developed in-house at CERN for analysing availability and reliability in complex engineering hierarchies. This deconstructs availability into the accelerator's main constituent systems, incorporating probabilistic failures and repair durations, and includes various features of FCC-ee operation relevant to achieved integrated luminosity. This section reviews features and assumptions used in the simulation model.

A. Energy Calibration

A principle goal of the FCC-ee is ultra-precise measurement of electroweak (Z and WW) observables, for which an accurately determined collision energy is key. This involves beam energy calibration every 10-15 minutes using non-colliding polarised pilot bunches, which circulate simultaneously with the main colliding bunches. The energy of these pilot bunches is measured by Resonant Depolarisation (RDP), where the frequency of a kicker magnet is adjusted until the pilot bunch's polarisation vanishes.

Pilot bunches are polarised in the main ring at the start of every fill using wiggler magnets, a process that takes roughly 90 minutes. The wigglers are then turned off before injection of the main colliding bunches. During physics operation, the pilot bunches have a combined Touschek and beam gas lifetime less than 20 h [4]. Once used for energy measurement, a pilot bunch becomes depolarised and can be topped-up at full energy injection. The self-polarisation effect, where the spin of particles gradually aligns with the dipole magnetic field, will fully re-polarise the pilot bunch in ∼200 h in the Z mode. If RDP can be achieved with less than 10% polarisation, this process could be repeated before the pilot bunch lifetime expires.

However, the minimum required degree of polarisation for accurate RDP is not yet precisely know. Further, self-polarisation in the main ring competes with radiative depolarisation, where momentum is transferred from the transverse to the longitudinal plane due to multiple effects such as Touschek scattering, errors and correction optics. As a result, the net rate of polarisation is unclear. Due to this uncertainty, it is not currently known whether pilot bunches will naturally achieve sufficient polarisation for energy calibration before expiring. The baseline simulation pessimistically expects that after 20 h the beam must be dumped to re-fill again in Z and WW modes.

In ZH and $t\bar{t}$ modes, the energy spread makes RDP impossible. Measurement is instead achieved by observing collisions at the Interaction Point (IP). This is significantly less accurate, but removes the need for polarisation at the start of every fill. Further, with top-up injection, physics can continue uninterrupted until a beam dump occurs due to machine fault or schedule end. In these modes, pilot bunches are used only to verify optics before injection of bunches with nominal intensities.

B. Operation Cycle

Energy calibration imposes distinct operation cycles for the electroweak (Z, WW) and high energy (ZH, $t\bar{t}$) modes, shown in Figure 1. Phases in this cycle are as follows:

1. *Set Up:* Main magnets are cycled, Radio Frequency (RF) system is set and other equipment is prepared for injection.

2. *Pilot Bunch Injection:* Pilot bunches are injected and equipment/optics are adjusted.

3. *Polarisation:* Z and WW modes only. Wigglers are turned on to begin polarising the pilot bunches, a process taking approximately 90 minutes.

4. *Fill:* Wigglers are turned off and the main colliding bunches are injected. Fill time t_f in each energy mode is shown in Table I. These times are indicative, and must be studied in detail according to constraints in the injector complex.

5. *Adjust:* Time to bring the beams into collision and make final adjustments to equipment before the detectors can begin to observe useful collisions.

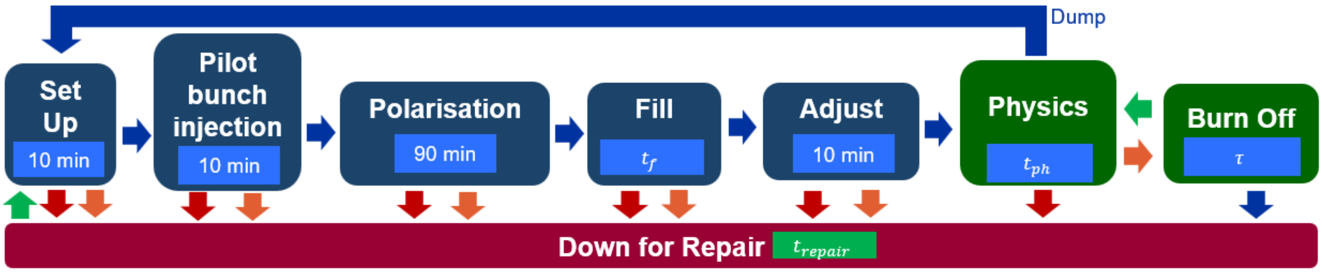

(a) Z and WW modes

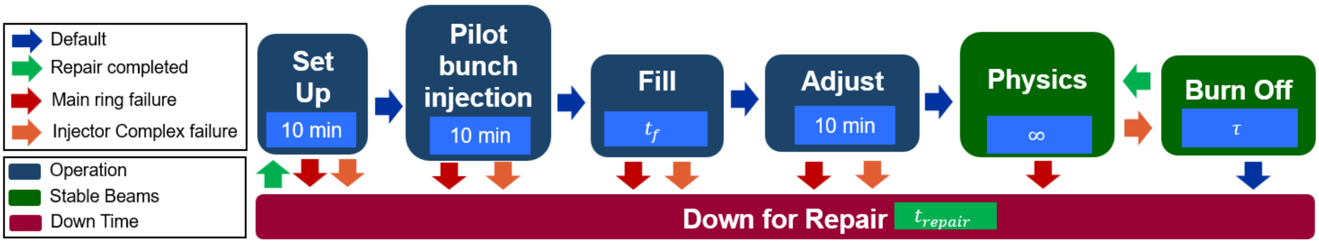

(b) ZH and $t\bar{t}$ modes

FIG. 1: Baseline Operation Cycles in the FCC-ee.

6. *Physics:* Collisions begin at nominal luminosity. With top-up injection, flat luminosity can be maintained. In Z and WW modes, a 20 h upper limit on the duration of physics t_{ph} is applied corresponding to the Touschek and gas scattering lifetime of the pilot bunches.

7. *Burn Off:* If the injector complex fails, the beam can be preserved in the main colliding rings with decaying luminosity corresponding to lifetime τ, Table I, after which, the beams quickly become unstable. If the injector complex cannot be restored before this time then the beam is dumped.

8. *Down for Repair:* On equipment failure, the accelerator is stopped for repair.

C. Luminosity

Luminosity is produced only in Physics or Burn Off phases in the operation cycle. By tracking the time t spent in these operation phases, achieved integrated luminosity is calculated.

$$L_{int} = \begin{cases} N_{IP}Lt, & t \in \text{physics} \\ N_{IP}L\tau(1 - e^{-\frac{t}{\tau}}), & t \in \text{burn off} \\ 0, & \text{otherwise} \end{cases} \qquad (1)$$

Where $N_{IP} = 4$ is the number of interaction points and L is the nominal instantaneous luminosity. In the first two years of Z operation, and first year at $t\bar{t}$, reduced luminosity is expected at 50 % and 65 %, respectively, to account for machine commissioning and beam tuning.

TABLE I: Operation Cycle Times

	Z	WW	ZH	$t\bar{t}$
Fill time, t_f (min)	7.7	2.5	1.52	1.45
Burn off lifetime, τ (min)	15	12	12	11
Physics duration, t_{ph} (h)	20	20	∞	∞

D. Contributing Systems

In collaboration with relevant system experts, each main constituent system in the FCC-ee is approximated for availability. If fault data from a comparable representative system could be found, this is extrapolated to an equivalent system in the FCC-ee using a generalised framework.

1. Generalised Extrapolation Framework

a. A similar or comparable state-of-the-art representative system is identified that exists currently in a working accelerator.

b. Where possible, the representative system is deconstructed into subcomponents and fault mechanisms to achieve better granularity of reliability estimation.

c. Fault mechanisms are categorized according to repair type and location:

- Remote Repairs can be completed from the control room, e.g. by re-setting or by-passing the failed component.

TABLE II: Approach time (the drive time from PA) and cool down time [1] relevant to human repair faults at each access point.

	PA	PB	PD	PF	PG	PH	PJ	PL	arc
Approach (h)	0.0	0.8	0.6	0.8	0.6	0.6	0.4	0.3	
Cool Down (h)	1.5	1.7	1.5	24[a]	1.5	0.7[b]	1.5	0.7[b]	1.7

[a]Collimation straight section without bypass tunnel.
[b]Klystron gallery only. A longer cool down time is required to access the beam tunnel.

- Human Repairs require on-site intervention by personnel. In addition to the time it takes to complete the repair, the approach time to get a person to the site of the fault is significant, shown Table II.
- Surface Equipment can be accessed for repair immediately when a technician arrives.
- Tunnel Equipment cannot be accessed until a suitable cool-down time has passed, shown for each access point in Table II.

d. Fault data is extracted from the representative system to gain Mean Time Between Failures (MTBF) and Mean Time to Repair (MTTR) information for its components. By default, only failures leading to down time are considered, therefore a similar degree of redundancy is assumed. However, if applicable, a bespoke formulation for redundancy of components and subsystems is defined and implemented.

e. MTBF and MTTR are scaled from the representative system to the the FCC-ee system. These are modelled probabilistically using exponential and Weibull distributions, respectively.

- MTBF scales inversely proportionate to the relative number of components. The more components in the FCC-ee system relative to the representative system, the lower the MTBF. If the number of components in FCC-ee is not yet precisely known, an approximation is made assuming a similar design specification to the representative system.
- MTTR for remote repair faults is left unchanged. An approach time and/or cool down time is added to human repair times corresponding to the equipment location, Table II.

2. Top-Level Systems

Fault data is taken from CERN's Accelerator Fault Tracking (AFT) database [5] specific to LHC proton physics 2015-2024, unless otherwise stated.

a. Accelerator Controls: Controls faults are grouped into software and hardware categories. Software faults occur irrespective of the number of instances, so their MTBF was applied without scaling. Hardware faults increase in frequency with the amount of implemented hardware. These were scaled according to the number of front-end computers (607 in LHC, 2400 in FCC-ee).

b. Access System: This controls the doors to enter the accelerator tunnel. Down time may be extended if a door does not open, or if it is incorrectly reported as open after it is closed. Since both LHC and FCC-ee have eight access points, a similar number of doors is assumed, and no scaling was applied.

c. Beam Instrumentation: These faults are grouped into Beam Position Monitor (BPM), Beam Loss Monitor (BLM) and "other" categories and scaled according to the relative number, respectively, in the main collider rings and booster.

d. Beam Losses: Down time due to Unidentified Falling Objects (UFO) and beam instabilities in the main collider rings and booster. UFO events are scaled with the length of beam pipe, assuming similar amounts of dust and beam interaction as in the LHC. Beam instabilities were applied without scaling. The LHC sees relatively long down times due to beam loss effects that cause quenches in nearby superconducting magnets, which is not representative for FCC-ee. Therefore down time due to quench recovery was subtracted to calculate the MTTR. FCC-ee dumps due to beam losses then appear frequently, with short recovery time.

e. Collimation: There are 9 moveable halo collimators and 24 moveable synchrotron radiation collimators in the FCC-ee main collider. Faults are scaled from the LHC's 108 moveable collimators. This category does not include injection and extraction absorbers, which are not yet modelled.

f. Cooling & Ventilation: Fault rate was scaled using the number of critical circuits at each access point, which is similar to the LHC (160 in LHC, 168 in FCC-ee). This does not take into account the overall volume served by these critical circuits, which is significantly larger in the FCC-ee.

g. Cryogenics: The FCC-ee requires only two cryoplants to serve the collider and booster superconducting RF systems at PH and PL, respectively. In the absence of a detailed design to indicate number of components, failure rate at each point is assumed to scale with the reference 4.5 K power level relative to one LHC cryoplant at 144 kW. This is constant for lower energy modes (Z, WW, ZH) at 80 kW in the collider and 9 kW in the booster. During the long shutdown prior to $t\bar{t}$ operation, each cryoplant is upgraded to 168 kW and 53 kW, respectively.

h. Electrical Network: The FCC-ee will draw power from two transmission lines in the French national grid at three points: PA, PD and PH. Faults fall into two categories:

1. External network perturbations are voltage deviations that appear in the transmission network outside of the hardware designed and managed by CERN. These are a high contribution to down time in the LHC, as they can trigger child faults in sensitive hardware across the

complex. These are statistically modelled to appear independently on each transmission line in the FCC-ee. Therefore, failure rate and down time is applied twice: Once to PA and once to PD-PH.

2. Internal faults occur in the distribution network designed and managed by CERN. LHC faults are first categorised into High Voltage (HV) 66-400 kV; Medium Voltage (MV) 3.3-18 kV; and Low Voltage (LV) <3.3 kV equipment. Respective failure rates are then scaled according to the number of substations at each location in PA, PD and PH. The parallel line feeding the collider RF at PH is scaled according to reference power level, approximately 0.34 that of LHC. PA and PD benefit from a redundant HV link between them such that if one fails, power can be drawn from the other with approximately 15 minutes of changeover time. This redundant link can only be used for HV hardware failure, and only if there were no additional child faults downstream. At PH there is no such redundant capacity due to its high power demand.

This formulation assumes that the same proportion of child faults will appear in the event of a voltage deviation as does currently in the LHC.

i. Elevators and Handling Equipment: LHC access points each have one elevator shaft. If the elevator breaks, the tunnel cannot be accessed for repairs. The FCC-ee has two elevator shafts per access point, so access is blocked only if both elevators break simultaneously. The resulting redundancy significantly mitigates this fault type.

j. Experiments: The LHC's experiments can cause down time for multiple reasons, e.g. if they require access to the tunnel for work on the detectors, faults in movable equipment close to the machine aperture, faults in beam loss protection equipment, etc. The four largest experiments (ALICE, CMS, ATLAS and LHCb) account for the majority of this down time. This was applied to the four IPs in FCC-ee without scaling.

k. Extraction and Beam Dump: Describes systems required for beam abort. Extraction systems are the controls and kicker hardware to divert the beam into the dump block. These appear twice in the LHC (one for each circulating beam), as for the FCC, so they are applied without scaling. However, only extraction system faults are applied in the booster as the dump blocks are shared with the collider rings.

l. High Energy (HE) Linac: The injector complex HE linac with energy up to 20 GeV is modelled using fault data from the Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory (SLAC) [6].

m. Injection Systems: Includes kickers, septa and control systems needed to inject and extract the beam under normal operation. The LHC and FCC-ee collider both have two injection systems (one for each beam), so these are applied without scaling. The booster has two injection systems from the injector complex, and two extraction systems towards the collider, so these faults appear at twice the rate.

n. Lower Energy Injectors: In the injector complex, the electron and positron source, linac and damping ring with energy up to 2.86 GeV are modelled from fault data in the SuperKEK-B e^-/e^+ source, injector linac and damping ring between 1998-2023, detailed in [7].

o. Machine Protection & Interlocks: Scaling for Fast Magnet Current Change Monitors (FMCM), Beam Interlock System (BIS), Safe Machine Parameters (SMP), Powering Interlock System (PIC), Warm Magnet Interlock Controller (WIC) categories are scaled from the relative number of hardware instances foreseen in the FCC-ee.

p. Magnets: Failure data from 3 532 normal-conducting magnets around the CERN complex is applied to 25 000 magnets in the FCC-ee, split evenly between collider and booster.

q. Operation: Down time from mistakes made by operators in the control room, e.g. settings input incorrectly, were applied without scaling to the FCC-ee.

r. Power Converters: Failure data was taken from comparable families of power converters to represent powering hardware for various magnet types. As only faults leading to down time in the LHC are considered, a similar degree of redundancy for each FCC-ee magnet group is assumed as exists currently in the LHC. Due to the large number of sextupoles and corrector magnets foreseen in the latest optics configuration, an MTBF of 3-5 days is observed in these categories.

s. Radio Frequency: The LHC has 16 accelerating superconducting RF cavities. The number for collider/booster in FCC-ee varies with energy mode: 264/112 in Z, WW and ZH, and 752/600 in $t\bar{t}$. All cavities are horizontally tested to a margin 10 % above their nominal voltage so, theoretically, nominal beam energy can be preserved if no more than 10 % of the cavities are unavailable. Practically, this is feasible only for the ZH and $t\bar{t}$ modes, where the beam current is relatively low. In Z and WW, the level of beam loading is such that loss of more than one or six cavities, respectively, must dump the beam.

t. Smoke Detection System: The FCC-ee smoke detection system must closely monitor for fire in the tunnel. Alarms may block the beam in the event of a false positive, i.e. indicating a fire when there is none; or a false negative, i.e. failing to detect a fire which then causes damage to other equipment. For safety reasons, the former category is much more common. In the event of any detected fire, a safety assessment must be conducted in person before operation is restored. As the LHC tunnel contains predominantly superconducting equipment, its alarm system is skewed towards detecting oxygen deficiency in the event of a helium leak, and so it is less representative here. The FCC-ee alarm system is better approximated by the long-range aspirating smoke detectors in the Super Proton Synchrotron (SPS). Scaling was applied assuming the same number of alarms per kilometre of tunnel (7 km in SPS, 91 km in FCC-ee).

u. Transverse Damper: The FCC-ee collider will have two transverse damper systems (one for each ring), as for the LHC, so failures are applied without scaling. The booster will have one transverse damper system, so failures occur at half the rate of LHC

v. Vacuum: LHC vacuum faults can be categorised according to the failed component (pump, gauge, valve, controller). These were then scaled according to the relative number of components in the FCC-ee collider and booster vacuum systems.

If no representative fault data could be found, but the system is still deemed a high risk for machine availability, it was given an availability placeholder pending a more rigorous reliability assessment. This applied to two systems: the *Beamstrahlung Dump Absorber* was approximated from the SPS Target Internal Dump Vertical Graphite (TIDVG) using a framework described in [8]; and the *Polarimeter*, which aims for availability of 95 % [9].

E. Repair Timeline

A repair timeline similar to the LHC was simulated: Following a dump, operators attempt for one hour to diagnose and restore operation remotely from the control room. If operation cannot be restored during this time, only then are personnel sourced to begin human repairs. Once humans are in the tunnel they will finish all repairs, including all redundant components.

III. RESULTS

Each energy mode was simulated for its full run (four years in Z, two years in WW, etc.), with performance averaged over 100 iterations. The achieved integrated luminosity and overall availability in each energy mode is shown in Figure 2. Forecasts for achieved integrated luminosity are significantly below target in all energy modes, especially in the electroweak sector (Z and WW).

The pie charts in Figure 2 show the percent of time spent in various phases in the operation cycle. The time spent in stable beams in Z and WW modes is 17-21 %, while in the higher energy modes ZH, $t\bar{t}$ they are 43-49 %. For comparison, the Feasibility Study design requirements quote a time spent in stable beams (efficiency) of 75 % [1, p. 80].

Efficiency is especially low in the Z and WW modes due to the requirement for a 90 minute polarisation phase at the start of every fill. This is a significant burden, as it severely penalises the recovery time following equipment failure. The turnaround penalty is two hours at these lower energies, while only 30 minutes in ZH, $t\bar{t}$. This is a strong factor in the observed luminosity shortfall.

The breakdown by system for unavailability and lost luminosity contribution is shown in Figure 3. The RF system has highest unavailability in Z mode due to the strict limit of one-cavity redundancy. In the WW mode, the ability to preserve the beam in the event of six lost cavities is powerful, reducing the collider unavailability contribution to zero. The booster makes significant contribution in both energy modes as it has no redundancy at all. In ZH and $t\bar{t}$, 10% redundancy leads to zero down time from the RF system, as this is sufficient reserve for failed cavities to wait to be repaired until there is long enough down time in other systems.

The electrical network is a high contributor to down time in all energy modes; partly due to hardware faults, but especially due to the high number of child faults that appear in other systems in the event of electrical perturbation. These child faults impede the potency of the redundant link between PA and PD, as the reconnection can only be activated if there are not other systems needing repair further downstream.

Power converters also show high contribution due to the high quantities of sextupole and corrector magnets assumed in the current optics configuration. This leads to a group MTBF of 3 and 5 days for converters in these magnet families, respectively.

The injector complex, composed of the HE linac and low energy injectors, also form high contribution to lost luminosity. Both of these accelerators have availability dominated by RF breakdowns. In the low energy injectors, the high implemented redundancy allows for especially short blocking durations in the event of faults. As a result, it has low contribution to down time, but the frequent dump rate leads to high contribution to lost luminosity, especially in the lower energy modes where the turnaround time to restore stable beams is significant.

The smoke detection system features relatively high due to the scaling factor from SPS (7km) to the FCC-ee tunnel (91 km). The number of alarms combined with mandatory on-site safety evaluation in the event of fire detection leads to large amounts of down time.

Beam losses, similar to low energy injectors, have lower contribution to unavailability as the recovery time to restore operation is short. But their contribution to lost luminosity is high as they cause frequent dumps, each of which incurs a turnaround penalty.

IV. R&D OPPORTUNITIES

Significant improvement in the reliability, availability and operational efficiency of the machine is required in order to meet integrated luminosity targets. Several opportunities are identified.

A. Operational Efficiency

Two solutions are proposed within the operation cycle to improve the percentage of operation time spent in stable beams.

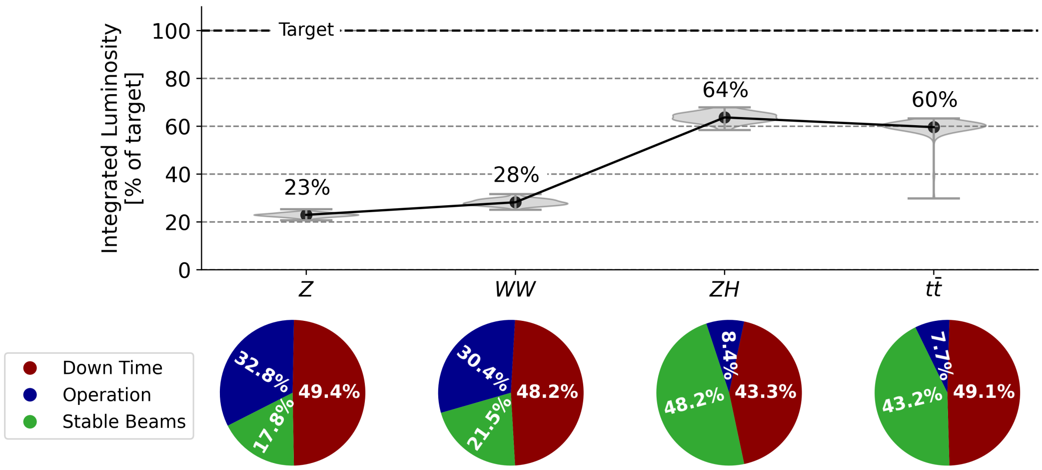

FIG. 2: Achieved integrated luminosity under the FCC-ee baseline design, extrapolating from the reliability and maintainability of state-of-the-art systems in current working accelerators. Units are normalised to the target integrated luminosity at each energy mode. The pie charts show corresponding time spent in various operation cycle phases from Figure 1.

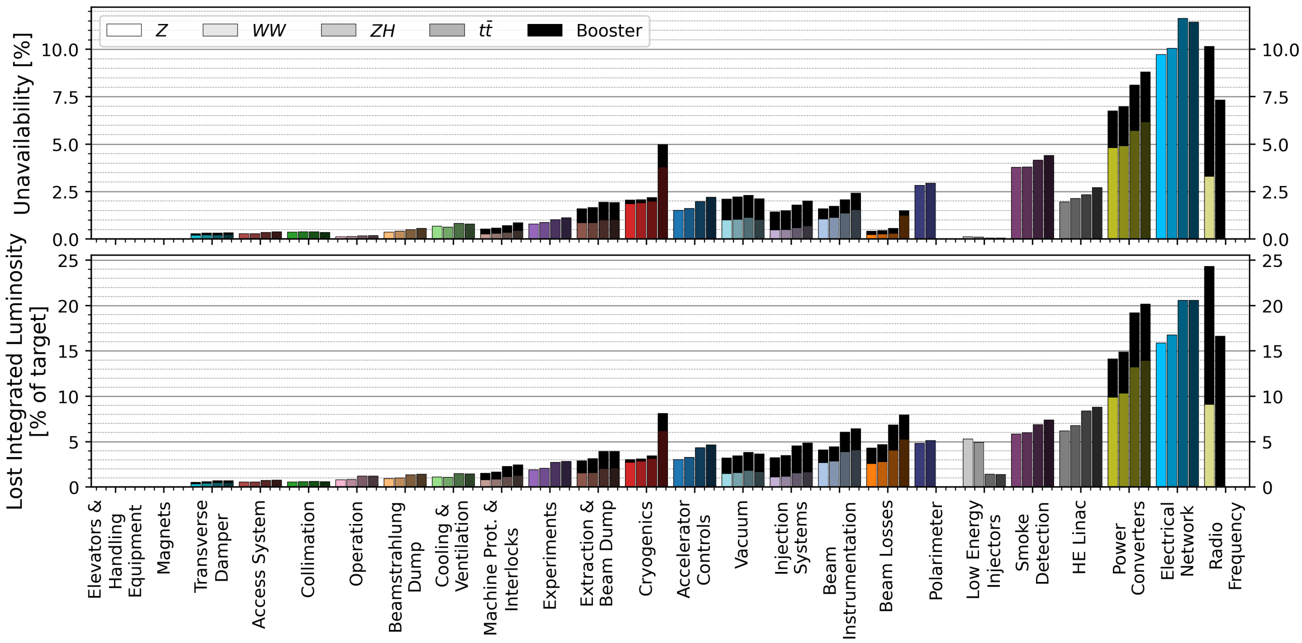

FIG. 3: Unavailability and lost luminosity contribution from each top-level system in the FCC-ee. Lost luminosity is normalised to the target integrated luminosity in each energy mode. Systems are ordered according to Z mode lost luminosity contribution.

1. *Indefinite Physics:*

If the lifetime of pilot bunches in the main ring can be made longer than the natural polarisation time, pilot bunches could be topped up immediately after being used for measurement and allowed to naturally repolarise. This avoids the need to dump after 20 h in order to re-fill with polarised pilot bunches. Physics can then continue uninterrupted in Z and WW modes, until equipment failure dumps the beam.

The utility of this proposal is best illustrated by the distribution of stable beams durations at each energy mode, shown in Figure 4. This shows that stable beams

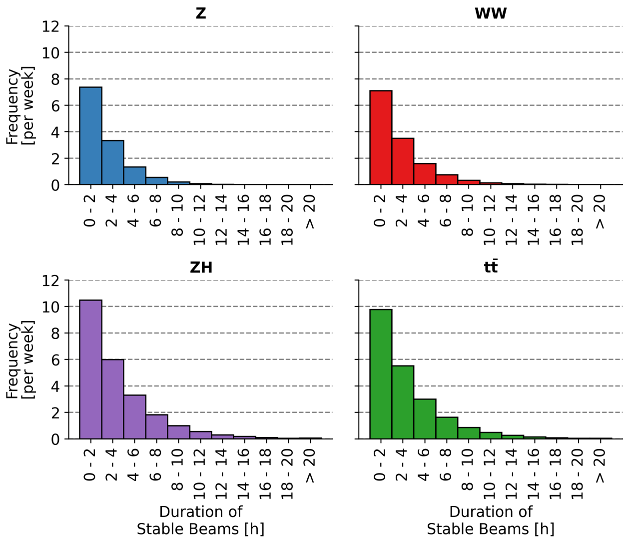

FIG. 4: Distributions for duration of stable beams in each energy mode. In Z and WW, physics cannot endure beyond 20 h. In ZH and $t\bar{t}$, there is no upper limit, however physics can rarely be sustained longer than 20 h due to failures dumping the beam.

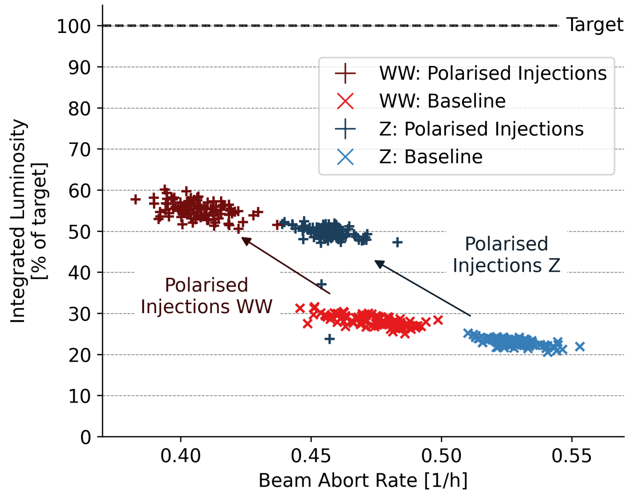

FIG. 5: Effect on achieved integrated luminosity from injecting already-polarised bunches, as an alternative to the baseline operation cycle shown Figure 1a. In both energy modes achieved integrated luminosity is more than doubled. This is the single most compelling opportunity to improve physics performance.

can in fact rarely be sustained longer than 20 h without beam abort due to system failure. Therefore, the gain from extending the upper limit on the duration of physics is negligible, unless it is accompanied by a machine-wide improvement in reliability.

2. Polarised Bunch Injection:

If pilot bunches can be polarised prior to injection, and their polarisation preserved through the injector complex and booster, the 90 minute polarisation phase could be eliminated from the Z and WW operation cycle entirely. It would effectively render the same operation cycle as the higher energy modes (ZH and $t\bar{t}$) while preserving capability for precise energy calibration.

The potential of this scheme is shown in Figure 5. In both Z and WW modes this more than doubles the achieved integrated luminosity. This is because, in removing the 90 minute polarisation phase at the start of every fill, the same duration of stable beams is effectively created. This is a potentially game-changing opportunity to support physics objectives in this challenging reliability environment.

B. Availability

To assure availability in the FCC-ee, a rigorous reliability and availability assurance framework is required to ensure a coherent approach for all systems in the machine. This must be designed and implemented in close collaboration with the relevant hardware teams such that it iterates in parallel with the design process and guides engineering decisions. Several opportunities within high priority systems are already identified.

1. Radio Frequency

The current technical limit for the collider RF cavity redundancy in Z mode is one single cavity. Additional redundancy is challenging due to high beam loading effects, but should be further pursued due to the high possible gains. This is demonstrated by the six redundant cavities in WW mode, which effectively eliminate RF down time.

It must be noted that zero RF down time is only possible as the simulation currently assumes perfect redundancy between cavity circuits. In practice, common mode failures would still occur that require reconditioning or replacement, and cannot be immediately compensated by another redundant cavity. Additional modelling is required to better understand these fault types. The coming years will require close collaboration with the RF design team to deconstruct the RF circuit into subsystems and isolate problematic fault types and high-failing component families.

Analysis of the auxiliary systems around the superconducting cavities would unearth further reliability opportunities. For example, the klystron gallery is accessible to personnel while the main colliding rings are in operation. If more high-failing auxiliary systems can be located there and designed to be redundant, modular and

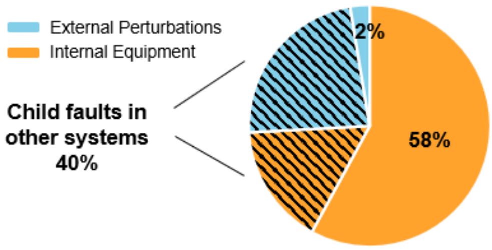

FIG. 6: Down time from the electrical network during LHC physics operation periods 2015-2024. Internal equipment faults can be mitigated through installation of redundant capacity. In parallel, the resilience of electrically connected systems to voltage deviations must be improved, to reduce child faults.

hot-swappable, significant gains could be achieved.

RF redundancy in the booster system should also be further explored, as this is a significant contributor to unavailability in both energy modes. Presently it is unclear whether redundancy could be implemented in the booster RF cavities as the transients due to ramping lead to additional complications. However, the beam current within the booster is significantly lower than in the collider, suggesting beam loading effects should be less problematic. Implementation of redundancy in this system could benefit achieved integrated luminosity significantly.

2. *Electrical Network*

The electrical network suffers from external network perturbations and internal hardware faults, both of which manifest as voltage deviations that may cause child faults and down time in other systems. Proportions of down time are shown for LHC physics operation 2015-2024 in Figure 6. This pie chart illustrates that two parallel chains of R&D are necessary for improvement.

a. Internal Network Redundancy: The internal electrical network forms the highest contribution to down time, and therefore has the most to gain from a well-designed redundant capacity. This could take the form of redundant lines and normally open switches that would allow operators to quickly bypass equipment faults and restore operation without having to wait for the equipment to be repaired. The existing redundant HV link between PA and PD will reduce down time from HV substations significantly at both these points. More could be achieved from redundant configurations in the MV and LV networks as well. However, child faults cannot be compensated in this way, nor can external glitches in the overhead transmission lines. The gains seen from redundancy will inevitably be constrained by the number of child faults due to voltage deviations.

b. Resilience to Voltage Deviations: The current model assumes that child faults will occur at the same rate and for the same duration as they do in the LHC. This is a coarse assumption, and the actual down time will scale with at least the type and number of connected sensitive equipment, e.g. high-precision power converters, RF components, cryogenics, etc., which could be significantly higher (or lower) in the FCC-ee. This risk can be mitigated. First, the precise nature of typical voltage deviations in the grid must be studied and quantified. With proper statistical understanding of the magnitude and phase of these voltage dips, the problem can be tackled from both sides:

On the child side, detailed requirements for resilience to voltage and power fluctuations must be issued to all equipment design teams. Further, these requirements must be prioritised in the design process to ensure delivery. Adequate validation and testing must be enforced for all connected equipment such that no device is installed in the machine that cannot tolerate these minimum fluctuations.

For the electrical side, it must be researched how to dampen voltage deviations to within the above design requirements. This can be achieved, for example, with unified power flow controllers strategically placed in the network to control active and reactive power flow, and compensate voltage glitches.

3. *Power Converters*

The current optics configuration powers sextupoles in groups of four. One failed converter would therefore lead to four magnets failing simultaneously. Collaboration with the optics working group is required to consider how to preserve the beam if combinations of converters and magnets fail. A system-level redundancy approach could significantly improve reliability in this case.

It must be noted that the present simulation assumes the same degree of redundancy in power converters as is currently implemented in the LHC. Therefore, the FCC-ee sextupole and corrector converters must be made more robust than those in the LHC for a solution to take effect. A detailed study into power converter reliability is essential in the coming years.

Reducing the number of corrector families, especially for the sextupole magnets, can drastically reduce the number of power converters required and in turn reduce their overall fault rate.

4. *Injector Complex*

In the event of outage in the booster and injector complex, stable beams can be maintained in the main collider rings for a lifetime of 10-15 minutes depending on energy mode (see Table I). If top-up injection can be restored in this time, normal physics can resume. This significantly reduces the FCC-ee's sensitivity to short duration fault types in the injector complex, and leads to an inherent advantage for availability. The contribution to lost luminosity of subsystems in the injector complex is therefore biased towards those with longer duration fault types.

R&D opportunities exist in exploiting this natural resilience to shorter fault types. If a failed component can be brought back online before the beams in the main collider expire, a lengthy turnaround time may be avoided. To this end there are two parallel approaches:

a. Beam Lifetime: If the collider beam lifetime without top-up injection can be made longer, systems will have better opportunity to recover before a beam abort is necessary. This requires intricate optimisation of the beam optics under various possible failure mechanisms; however, this could greatly alleviate the design task on the collider technical systems.

b. Recovery time: Measures to bring the recovery time of failed systems to below 10 minutes could be extremely powerful for the booster. For example, the automatic reset of recurring RF trips; or the installation of redundant or back-up modules that can quickly be brought online to replace failed components.

V. RELIABILITY REQUIREMENTS

Rising to the FCC availability challenge will inevitably require a machine-wide improvement in reliability. The degree of required improvement depends strongly on many open design decisions relating to instantaneous luminosity, operation paradigm, automation, maintenance processes, etc. Nevertheless, to gain some appreciation of the scale of the task at hand, a parameter scan was conducted for reliability of the highest failing systems.

Two assumptions are made in this analysis:

a. Fewest Systems: Reliability is improved in the smallest possible number of top-level systems to reach an overall dump rate for the whole machine. Reliability consolidation starts with the single highest-failing system (i.e. the RF). Once the RF has the same dump rate as the second highest-failing system (Low Energy Injectors), effort is made to reduce dump rate in both systems. This carries on, including more systems if necessary, until the required overall mean time between dumps is achieved. This simulates the most efficient way to distribute resources for reliability consolidation: No system improves that does not need to improve.

b. Even Scaling: Within the top-level systems, there are many ways in which reliability consolidation could be applied, and the exact solution will vary according to each respective design. For simplicity, this parameter scan scales the MTBF of all subsystems and components beneath a top-level system by the same multiplier. This coarse assumption is analogous to a system-level redundancy approach, and allows for the reliability improvement of each top-level systems to be entirely described by its multiplier.

The effect on achieved integrated luminosity from scaling reliability in this way is illustrated for the Z mode with and without polarised injections in Figure 7. The two leftmost clusters show achieved physics performance extrapolating from systems at the time of writing, which are identical to the Z mode clusters in Figure 5. To move right along the graph's horizontal axis, more and more top-level systems must increase their MTBF in order to raise the overall time between dumps of the whole machine.

For illustration, cluster (I) in Figure 7 shows this reliability approach under the baseline operation cycle where the delivered integrated luminosity is equivalent to that from polarised injections at the 2025 standpoint. This is achieved by improving the reliability of six systems: RF must increase its MTBF six times, low energy injectors must improve five times, power converters three times, etc.

More striking differences are apparent when the requirements to reach 100 % of target integrated luminosity are compared. With polarised injections, shown by cluster (II), the reliability of 16 systems must improve, most significantly the RF must increase its MTBF 27 times. Without polarised injections, shown by cluster (III), 18 systems must improve, and the RF must extend its mean time between dumps 62 times.

These dauntingly high numbers expose some bittersweet truths concerning availability in the FCC. On the one hand, the technical challenge is clear – the FCC-ee cannot simply be a "scaled up" combination of accelerator systems today. A radical rethink of the design process is required. On the other hand, this insight, several years in advance of the technical design phase, is well timed.

There are flickers of hope. The current reliability performance was extrapolated from smaller systems in smaller accelerators that were not necessarily designed with such high priority on reliability as is necessary for FCC-ee. Under a well-implemented reliability and quality framework, large gains can be achieved from small design details if applied in the right way, and in the right place. This is already observed in simulation for the WW mode RF system, where only 6 redundant cavities is sufficient to eliminate RF down time in the collider. In many systems, the constraints on high reliability may not be technological, more, simply, a matter of ensuring that it is adequately prioritised in the design.

Finally, Figure 7 shows that a reliability-focused approach is, in isolation, unlikely the most efficient means of securing luminosity in the machine. Polarised injections, which act mainly on operational efficiency, are capable of reducing reliability requirements to a third of those without it. Other R&D opportunities stemming from alternative operational paradigms, such as automation and robot maintenance, may reduce these requirements even further.

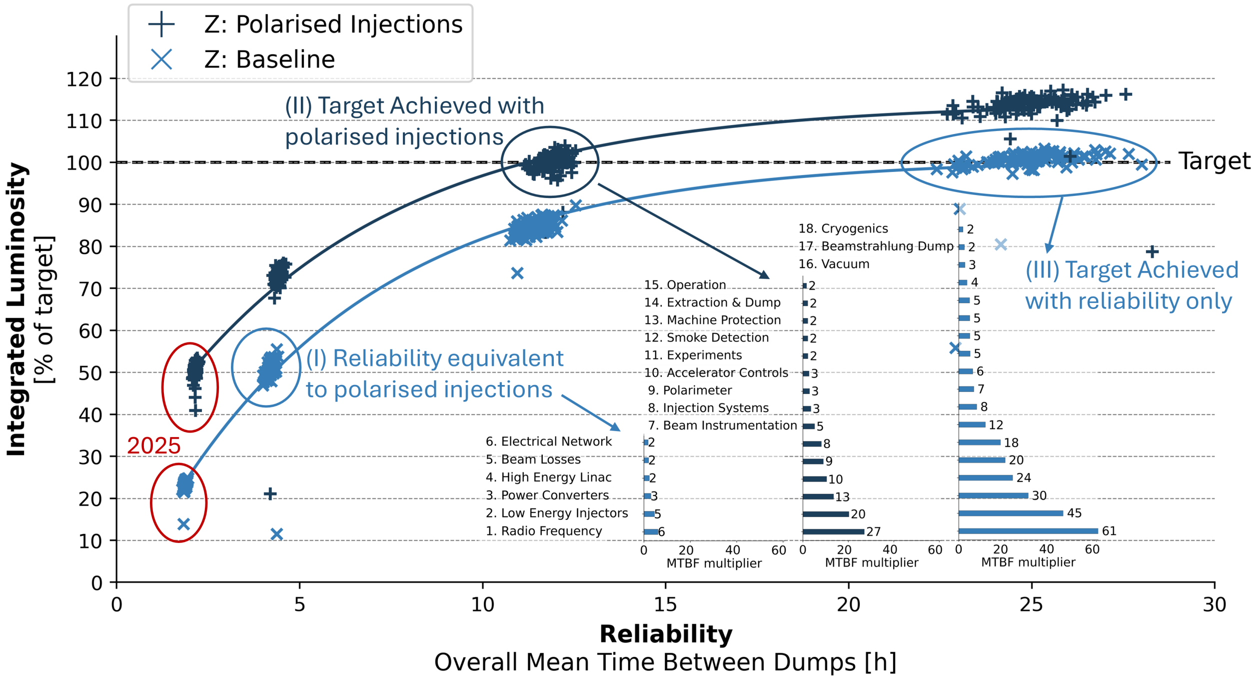

FIG. 7: Reliability improvement requirements in order to achieve integrated luminosity in the Z mode baseline design with and without polarised injections. Capability for polarised injections reduces the required system-level reliability gains to a third of those without it.

VI. RISING TO THE CHALLENGE

Systematic priority on reliability and availability is essential to ensure physics performance in the FCC-ee. To this end, a dedicated reliability/availability assurance framework is proposed to run in parallel with the accelerator design process, supporting development from R&D through to deployment and operation. A timeline is illustrated in Figure 8.

In the first instance, tangible and achievable reliability/availability end-requirements must be issued to system design groups, and milestones outlined for achieving these requirements. As designs develop, progress against these milestones must be tracked, benchmarked and updated, incorporating a holistic view of overall machine performance in terms of availability, efficiency, integrated luminosity and various cost factors.

In parallel, starting with the most problematic systems, equipment must be modelled in greater detail to single out high-failing components and subsystems, as well as low-hanging fruit to promote a redundant, modular and hot-swappable design paradigm. This will require close collaboration with the relevant system design teams. Several opportunities to reduce dump rate in high-failing systems are already proposed in Section IV. With closer scrutiny, more will appear. These must then be evaluated and compared in terms of their direct effect on accelerator performance, as well as optimised against cost constraints.

In the longer term, a rigorous and systematic support framework is required to ensure a coherent approach for the design of all systems in the machine. This framework must be developed and implemented such that it iterates in parallel with the design process and guides engineering decisions from conception through to prototyping and industrial roll-out. In addition to the analysis and optimisation proposed above, this must include validation and testing processes for prototyping and mass-manufacture. Tools, training and expertise must be readily accessible across the complex, ensuring skills and knowledge are where they need to be.

This support framework should endure through construction, commissioning and operation of the machine. Diagnostics, health monitoring and fault prediction strategies may be designed to gain advantage over uncertainty in novel and pioneering systems where the underpinning technology has low maturity. Ready for machine operation, fault tracking and performance analytics must be developed to enable swift and unambiguous root cause evaluation of trends in down time, which can greatly expedite responsive reliability consolidation during technical stops and long shutdowns.

With a well-implemented reliability and availability strategy, the design process can be supported from conception through to installation and operation of the FCC-ee, ensuring that resources committed to lifecycle reliability are effectively and efficiently spent.

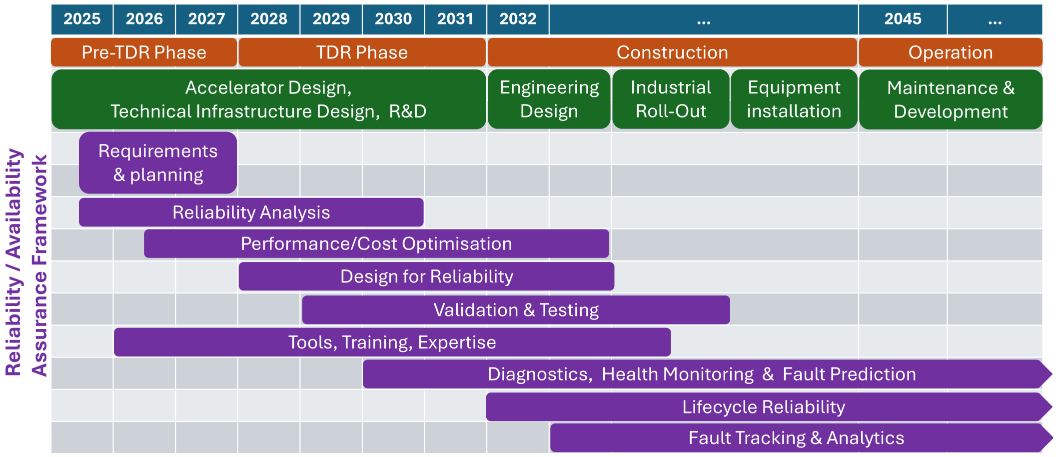

FIG. 8: A rigorous and systematic reliability and availability support framework is essential for success in the FCC. This must persist in parallel with the design process, informing and guiding engineering decisions from conception through to installation and operation.

VII. CONCLUSION

Availability is one of the primary challenges to physics deliverables in the FCC-ee. This paper presents an enhanced Monte Carlo simulation environment to study availability, operational efficiency and integrated luminosity of this future machine, by extrapolating reliability and maintainability of its top-level systems from state-of-the-art representations in existing working accelerators.

Results expose significant shortfall in integrated luminosity at all energy modes in the current baseline design. This is especially pronounced at lower energies (Z and WW) due to the low operational efficiency implied by the 90 minute polarisation phase in the electroweak operation cycle. In these energy modes, the long turnaround time required to restore stable beams in the event of a beam dump severely penalises systems with high fault rate. As a result, a system with high dump rate but quick recovery time, such as beam losses or low energy injectors, can have low contribution to down time but high contribution to lost luminosity.

Two R&D opportunities are explored to resolve this low operational efficiency. *Indefinite Physics* seeks to extend the lifetime of pilot bunches in the collider such that stable beams can be maintained indefinitely. This has negligible effect due to the high dump rate from failed systems. *Polarised Injections* involves direct injection of polarised bunches to avoid the 90 minute polarisation phase at the start of every fill. This is the most compelling solution to improving integrated luminosity to date, with capacity to more than double the achieved integrated luminosity.

From the simulation results, systems are ranked according to their contribution to lost luminosity. These are the RF, Electrical Network, Power Converters, and the Injector Complex. The quantification of these contributions is an essential first step in order to benchmark and improve system reliability in the technical design phase. Several opportunities to reduce dump rate in these highest failing systems are proposed.

Rising to the FCC availability challenge will inevitably require a machine-wide improvement in reliability. For the current baseline design, a parameter scan of the highest-failing systems shows that the dump rate of 18 top-level systems must improve by up to 62 times in order to meet the Z mode integrated luminosity target. With polarised injections, this reduces to 16 systems by up to 27 times. These dauntingly high numbers expose two bittersweet conclusions: First, a radical rethink of the accelerator's operation and maintenance paradigm is essential to expose further game-changing opportunities that reduce these reliability requirements. Second, a rigorous and systematic reliability and availability assurance framework is needed to support the accelerator design process from conception through to installation and operation.

Finally, the structure of a reliability and availability support framework is expanded. In the short term, this includes performance requirements, milestones and benchmarks, as well as more detailed reliability analysis to expose low-hanging fruit in the most problematic systems. In the longer term, a rigorous assurance framework is required to ensure a coherent approach for all systems, including validation and testing, tools and training, diagnostics, health monitoring and fault prediction strategies as well as fault tracking and analytics solutions.

One fact is evidently clear – to rise to the FCC availability challenge, we must act now.

[1] M. Benedikt, W. Bartmann, J.-P. Burnet, C. Carli, A. Chance, P. Craievich, M. Giovannozzi, C. Grojean, J. Gutleber, K. Hanke, A. Henriques, P. Janot, C. Lourenco, M. Mangano, T. Otto, J. H. Poole, S. Rajagopalan, T. Raubenheimer, E. Todesco, L. Ulrici, T. P. Watson, G. Wilkinson, F. Zimmermann, and B. Auchmann, *Future Circular Collider Feasibility Study Report Volume 2: Accelerators, technical infrastructure and safety*, Tech. Rep. (CERN, Geneva, 2025).

[2] L. Felsberger, P. Alexaki, A. Asko, N. Charitonidis, J. Heron, B. Mikulec, B. Rae, J. Uythoven, and D. Wollmann, Report Accelerator Complex Availability 2024, (2025).

[3] M. Blaszkiewicz, J. W. Heron, A. Apollonio, T. Buffet, T. Cartier-Michaud, L. Felsberger, J. Uythoven, and D. Wollmann, enAvailSim4: Open Source Framework For Availability and Reliability Simulations, in en*European Safety and Reliability Conference (ESREL); Advances in Reliability, Safety and Security, Part 4* (Krakow, Poland, 2024) pp. 29–38.

[4] A. Blondel, P. Janot, J. Wenninger, R. Aßmann, S. Aumon, P. Azzurri, D. P. Barber, M. Benedikt, A. V. Bogomyagkov, E. Gianfelice-Wendt, D. E. Kerchen, I. A. Koop, M. Koratzinos, E. Levitchev, T. Lefevre, A. Milanese, N. Muchnoi, S. A. Nikitin, K. Oide, E. Perez, R. Rossmanith, D. C. Sagan, R. Tenchini, T. Tydecks, D. Shatilov, G. Voutsinas, G. Wilkinson, and F. Zimmermann, Polarization and Centre-of-mass Energy Calibration at FCC-ee, arXiv (2019), http://arxiv.org/abs/1909.12245.

[5] CERN, Accelerator Fault Tracking (AFT), online tool, https://aft.cern.ch/.

[6] W. Allen, A. Brachman, W. Colocho, M. Stanek, and J. Warren, *Availability Performance and Considerations for LCLS X-Ray FEL at SLAC*, Tech. Rep. (United States, 2011) sLAC-PUB–14422.

[7] K. Furukawa, M. Akemoto, D. Arakawa, Y. Arakida, Y. Bando, H. Ego, Y. Enomoto, T. Higo, H. Honma, N. Iida, K. Kakihara, T. Kamitani, H. Katagiri, M. Kawamura, S. Matsumoto, T. Matsumoto, H. Matsushita, K. Mikawa, T. Miura, F. Miyahara, H. Nakajima, T. Natsui, Y. Ogawa, S. Ohsawa, Y. Okayasu, T. Oogoe, M. A. Rehman, I. Satake, M. Satoh, Y. Seimiya, T. Shidara, A. Shirakawa, H. Someya, T. Suwada, M. Tanaka, D. Wang, Y. Yano, K. Yokoyama, M. Yoshida, T. Yoshimoto, R. Zhang, and X. Zhou, enAchievement of 200,000 hours of operation at KEK 7-GeV electron 4-GeV positron injector linac, Journal of Physics: Conference Series **2420**, 012021 (2023).

[8] H. A. Dostmann, J. W. Heron, D. Wollmann, and J. Uythoven, Qualitative availability forecasting in the future circular electron-positron collider (FCC-ee) beamstrahlung dump system, Engineering; Accelerators and Storage Rings (2025), https://cds.cern.ch/.

[9] R. Kieffer, ATDC FCC Polarimeter Update: Energy calibration and polarisation overview (2024).