PERFORMANCE–COST DESIGN TRADE-OFFS FOR AVAILABILITY AND INTEGRATED LUMINOSITY IN THE FCC-ee

J. W. Heron*, H. Dostmann, L. Felsberger, J. Uythoven, D. Wollmann, 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, it is ambitious in size, complexity and technical objectives. Availability is a significant challenge. Rising to this ambition requires a coordinated availability-driven design strategy. Three objectives are identified: (1) R&D opportunities must be evaluated and compared across holistic metrics to enable early and informed design decisions; (2) Precise and balanced targets for availability must be defined ready for detailed system design; (3) Viable solutions to improve performance must be optimised against cost constraints to deliver efficient as well as performant solutions. Towards these objectives, this paper presents early results from *Ramilab*, an integration-level performance-cost optimisation tool designed to evaluate, compare and optimise accelerator designs through their holistic effect on availability and integrated luminosity.

INTRODUCTION

At 91 km circumference, the Future Circular Collider (FCC) would be the largest particle accelerator 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, including four years at the Z pole, two years at the WW pair production threshold, three years at the ZH production peak, and five years for top/anti-top $t\bar{t}$ production.

The machine must be operational for a minimum amount of time at each energy in order to meet its physics objectives. There are 185 days 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 minimum 80 % [1]. For comparison, the 27 km Large Hadron Collider (LHC) was available for 72 % of the time during the physics production scheduled between 2015-2025 [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.

Ramilab [3] is an enhanced Monte Carlo simulation tool designed to study availability, operational efficiency and integrated luminosity in the FCC-ee. This forecasts overall accelerator performance by extrapolating the reliability and maintainability of its top-level systems from state-of-the-art representations in existing working machines. Figure 1, from the authors' previous publication [4], shows a significant shortfall in achieved integrated luminosity at all energies, but especially at lower energies (Z and WW) where the polarisation phase significantly reduces the available time for stable beam physics operation. This insight is timely and valuable at the current stage of accelerator development, and the obstacles to hitting design targets are considerable.

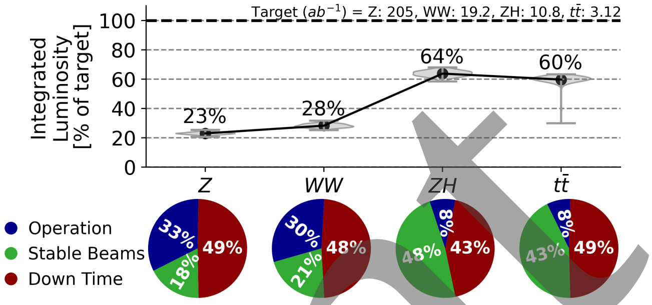

Figure 1: FCC-ee predicted integrated luminosity as percent of target in each energy mode, extrapolating from the availability of systems in current working accelerators. The pie charts show time spent in operation cycle phases [4].

Expanding on the FCC-ee availability challenge from [4], this paper identifies three principal objectives for supporting the technical design stage: (1) Evaluating R&D opportunities; (2) Allocating system-level design targets; (3) Cost-optimising the design process. Each begins with an objective description of the intended support, and follows with an example from recent studies.

R&D OPPORTUNITY ASSESSMENT

Technical Design Support

The FCC-ee design landscape encompasses a wide range of alternative technical solutions, across numerous subsystems and disciplines, that deviate from the baseline accelerator design. The *Ramilab* framework enables quantitative comparison of these alternative design proposals in terms of holistic machine performance-i.e. availability, operational efficiency, and integrated luminosity. This can deliver valuable insights for informing engineering design decisions in advance of committing to significant R&D investment.

Example: Alternative Design Opportunity

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.

---

* jack.heron@cern.ch

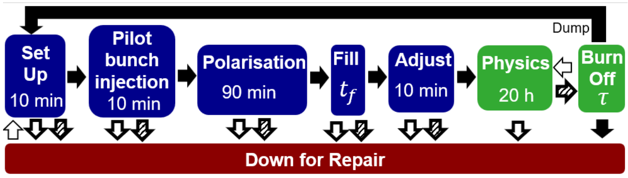

(a) Baseline design: Polarisation phase with wigglers

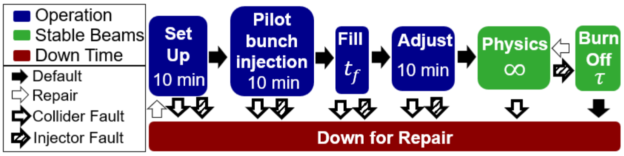

(b) Alternative design: Polarised injections.

Figure 2: FCC-ee operation cycle in Z and WW modes.

Baseline Design: Polarisation phase with wigglers
Pilot bunches are polarised in the collider ring at the start of every fill using wiggler magnets, a process that takes 90 minutes. The wigglers are then turned off before injection of the main colliding bunches. Once used for measurement, a pilot bunch is depolarised and cannot be used again until it naturally self-polarises, estimated at ~20 h [5]. This gives the operation cycle shown in Fig. 2a.

Alternative Design: Polarised injections
If pilot bunches can be polarised in a dedicated damping ring prior to injection, the 90 minute polarisation phase could be eliminated entirely. This delivers Fig. 2b.

The holistic machine performance gained with polarised injections is shown comparing the 2025 clusters of Fig. 3. It more than doubles the achieved integrated luminosity because, in removing the polarisation phase at the start of every fill, the duration of stable beams is extended accordingly. This is the single most compelling opportunity to improve physics performance identified to date.

AVAILABILITY DESIGN TARGETS

Technical Design Support

System design teams must work to reliability and maintainability targets throughout the technical design stage. How to divide the overall machine unavailability budget (currently 20%) between the systems that compose it remains an open question. The *Ramilab* framework permits a data-driven decomposition of the global unavailability budget across systems, thereby providing a rigorous basis for setting design targets. This is an ongoing study, exemplified below.

Example - Reliability Extrapolation

An initial estimate of the scale of the task at hand is displayed through a parameter scan on system reliability. Two simplifying assumptions are adopted:

1) Fewest Systems: Reliability improvements are applied to the minimum number of top-level systems required to achieve a target global dump rate. This approximates optimal reliability consolidation, where only systems that constrain global performance are improved. Consolidation is first concentrated on the highest-failing system (i.e. the Radio Frequency (RF) [4]). Once its dump rate matches that of the next most critical system, both are improved in parallel. This continues, incorporating additional systems as required, until the target overall reliability is reached.

2) Even Scaling: Within each top-level system, consolidation is modelled by uniformly scaling the Mean Time Between Failures (MTBF) of all underlying subsystems and components. Although coarse, this assumption is analogous to system-level redundancy and enables each system's reliability improvement to be characterised by a single multiplier.

The resulting trend on Z mode integrated luminosity with and without polarised injections is shown in Fig. 3. The leftmost 2025 clusters correspond to present-day extrapolated performance. Progression along the horizontal axis reflects increasing machine-wide reliability, requiring improvements across an expanding set of top-level systems.

To reach the target 100 % of design luminosity with polarised injections (cluster II), 16 systems require consolidation: the RF must increase its MTBF by a factor of 27, the low-energy injectors by a factor 20, power converters by a factor 13, etc. Without polarised injections (cluster III), 18 systems must improve, and the RF requires a 61-fold increase in MTBF.

These results demonstrate the scale of the availability challenge. The FCC-ee demands a fundamentally reliability-driven design approach in order to succeed. Nevertheless, these estimates should be interpreted in context: this performance is extrapolated from smaller machines that were not designed under FCC-ee-level reliability constraints. Significant gains can therefore be expected through targeted design improvements and robust quality frameworks. In many cases, achieving high reliability may be less a technological limitation than a question of design prioritisation. Furthermore, reliability improvements alone are unlikely to be the most efficient route to achieving luminosity targets and architectural or operational design changes, such as polarised injections, need to be considered to significantly reduce reliability requirements.

DESIGN OPTIMISATION

Technical Design Support

The challenging FCC-ee requirements for reliability and maintainability introduce additional complexity for design teams, who must concurrently satisfy other performance metrics such as lifecycle cost, sustainability, safety, etc. *Ramilab*'s capability extends to systematic optimisation of technical solutions under capital and operational cost constraints. It therefore supports decision-making not only at the conceptual stage, but throughout the design lifecycle, helping to ensure that resources allocated to availability are deployed with maximal effectiveness.

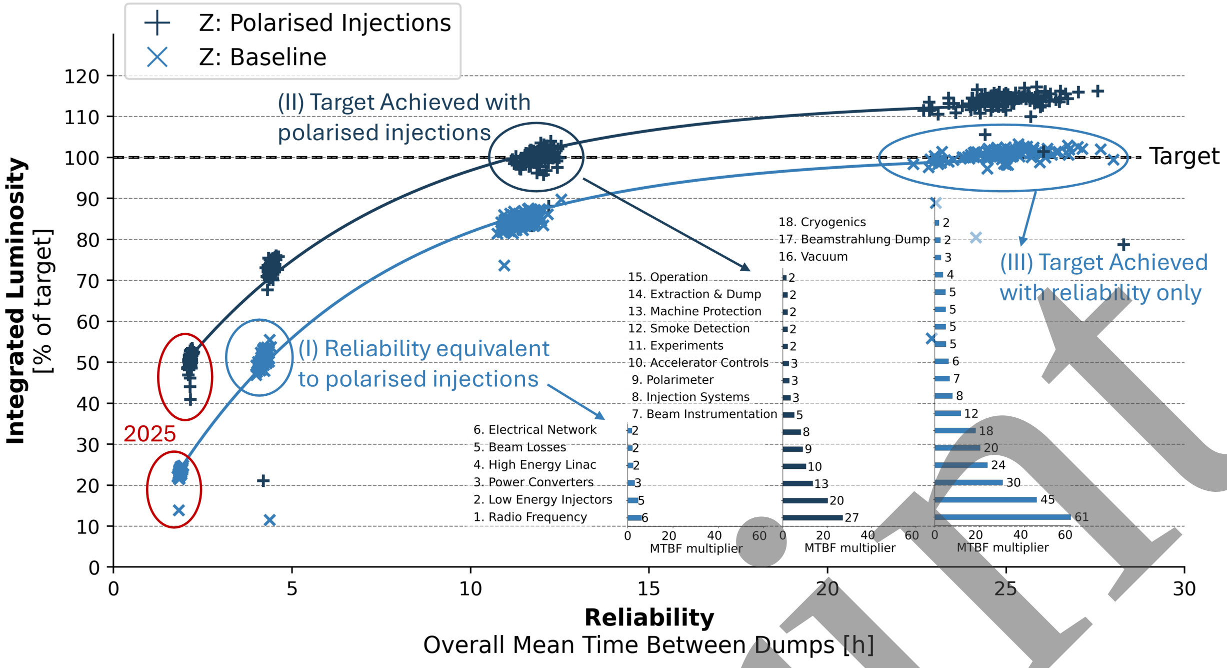

Figure 3: Reliability improvement required to achieve Z mode integrated luminosity targets, (II) with and (III) without polarised injections.

Table 1: R&D Opportunity Merit Table

R&D Opportunity	**Luminosity Gain*** **(ab^{-1})**	**CapEx Gain*** **(MCHF)**	**OpEx Gain***† **(MCHF)**	**Merit** $\left(\frac{\Delta \textbf{Performance}}{\Delta \textbf{Cost}}\right)$
Polarised Injections	ΔL	Δc_1	Δc_2	$\frac{\Delta L}{\Delta c_1 + \Delta c_2}$
⋮	⋮	⋮	⋮	⋮

* Difference over the baseline configuration equivalent value.
† Operational cost per year multiplied by the 15 years of FCC-ee operation.

Example - Performance/Cost Ranking

The *Ramilab* framework enables direct, quantitative comparison of R&D opportunities by mapping their impact on machine performance and cost, by capturing their effect on integrated luminosity through changes in reliability, operational efficiency, or both. In parallel, associated capital (CapEx) and operational (OpEx) cost drivers can be parameterised. This provides a common evaluation space in which heterogeneous design options—spanning hardware upgrades, redundancy strategies, and operational innovations—can be assessed on equal footing.

Table 1 illustrates the power of this approach. For each R&D opportunity, performance gain relative to the baseline configuration and corresponding cost variations are combined to produce a figure of merit, e.g. the ratio of performance gain to total cost increase. This enables rapid ranking of competing options according to their efficiency at converting resources into physics output. Importantly, the framework has scope for more refined analyses beyond this simple metric, including multi-objective optimisation and sensitivity studies with respect to cost assumptions and operational scenarios.

In this way, *Ramilab* provides a structured basis for prioritising R&D investments. Above qualitative judgment or subsystem-level optimisation, design choices can be guided by quantified impact on global machine performance.

CONCLUSION

Achieving the FCC-ee physics objectives places unprecedented demands on machine availability and requires a step-change in how reliability, operational efficiency, and cost are integrated into the accelerator design process. Closing the gap between predicted and target integrated luminosity requires a coordinated, availability-driven design strategy. The *Ramilab* framework provides decisive capability: it establishes a system-level, performance–cost modelling approach that directly links design decisions to global machine outcomes. This enables rigorous comparison of R&D opportunities, data-driven allocation of reliability and maintainability design targets, and optimisation of technical solutions under realistic constraints. It can therefore expose the true leverage of design proposals ranging from subsystem improvements to architectural changes. These capabilities are actionable at both the current stage of FCC-ee development and throughout the technical design lifecycle. *Ramilab* is not merely a supporting tool, but a necessary framework for prioritising resources, managing design trade-offs, and ensuring delivery of the FCC-ee ambition.

REFERENCES

[1] M. Benedikt *et al.*, "Future Circular Collider Feasibility Study Report Volume 2: Accelerators, technical infrastructure and safety", CERN, Geneva, Rep. CERN-FCC-ACC-2025-0004, 2025. doi:10.17181/CERN.EBAY.7W4X

[2] L. Felsberger *et al.*, "Report Accelerator Complex Availability 2024", Rep. CERN-ACC-NOTE-2025-0003, 2025.

[3] Ramilab: FCC-ee reliability, availability, maintainability, and integrity study workspace, 2026, `gitlab.cern.ch/jheron/ramilab`

[4] J. Heron, H. Dostmann, L. Felsberger, J. Uythoven, and D. Wollmann, "Availability assurance in the future circular electron-positron collider (FCC-ee)", in *Proc. IPAC'25*, Taipei, Taiwan, May 2025, pp. 3018–3021. doi:10.18429/JACoW-IPAC2025-THPS027

[5] A. Blondel *et al.*, "Polarization and Centre-of-mass Energy Calibration at FCC-ee", `http://arxiv.org/abs/1909.12245`, Sep. 2019, doi:10.48550/arXiv.1909.12245,