CHALLENGES IN BEAM COUPLING IMPEDANCE FOR FCC-ee

C. Zannini*, C. Antuono, E. de la Fuente Garcia, D. Gibellieri[1], E. Macchia, I. Mases Sole,
G. Rumolo, L. Sito, CERN, Geneva, Switzerland
A. Ghribi [1], GANIL, Caen, France
M. Zobov, INFN - LNF, Frascati, Rome, Italy
M. Migliorati, University of Rome La Sapienza and INFN, Rome, Italy
[1]also at University of Caen Normandy, Caen, France

Abstract

A comprehensive impedance model is required to ensure beam stability and optimize performance in the FCC-ee main rings. The model integrates contributions from a wide range of components, accounting for both resistive-wall and geometric effects. In this paper, we discuss the main challenges introduced by the peculiar FCC-ee parameter regime. A first difficulty arises from the combination of large beam-pipe dimensions and very short bunch lengths, which drives wakefield simulations into an extremely demanding computational regime, where very fine spatial resolution is necessary to accurately capture the beam–environment interaction. In addition, the beam-pipe cut-off lies within the frequency range excited by the FCC-ee beam. As a consequence, several higher-order modes may propagate over long distances, leading to non-local impedance effects and possible cross-talk between different accelerator elements. This means that the impedance environment cannot be treated as purely local, but requires a distributed description and an assessment of how propagating power is transported and potentially absorbed within the machine.

INTRODUCTION

A reliable impedance model is essential to ensure beam stability in FCC-ee. Compared to previous machines, FCC-ee operates in a regime that uniquely combines ultra-short bunch lengths with relatively large beam-pipe apertures. This configuration pushes a significant fraction of the beam-induced electromagnetic spectrum beyond the cut-off frequency of the vacuum chamber. Consequently, a non-negligible portion of the excited fields is no longer evanescent but can propagate over long distances along the machine. This feature is not marginal but intrinsic to the FCC-ee operating regime. The beam spectrum significantly overlaps with the beam-pipe cut-off frequency, as illustrated in Fig. 1. In contrast to the LHC, where most of the spectral content remains below cut-off, FCC-ee exhibits a substantial fraction of power above cut-off. As a result, higher-order modes can propagate along the vacuum chamber, making propagation effects an inherent characteristic of the FCC-ee impedance environment.

This propagation fundamentally alters the physical picture of impedance modeling. In conventional accelerators, impedance is typically treated as a local property, attributed to individual components whose electromagnetic interaction with the beam decays rapidly with distance. In FCC-ee, however, the presence of propagating modes introduces a non-local coupling between elements: a given structure can influence, and be influenced by, components located far upstream or downstream. As a result, the machine behaves less like a collection of independent devices and more like a distributed electromagnetic system, where interference, reflections, and modal conversion effects play a significant role.

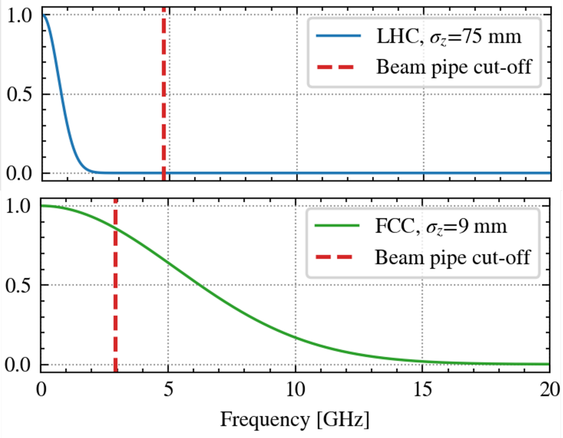

Figure 1: Comparison between beam spectrum and beam pipe cut-off frequency for FCC-ee and LHC. FCC-ee exhibits a substantial fraction of spectral content above beam pipe cut-off.

Such a regime brings new challenges for both simulation and interpretation. Numerical results become intrinsically sensitive to boundary conditions, domain truncation, and the treatment of absorbing boundaries, as these can artificially suppress or enhance propagating content. The very definition of impedance becomes less straightforward, as separating localized contributions from global, propagating effects is no longer trivial. This raises important questions about how to construct a meaningful and robust impedance budget for the machine. Moreover, the presence of propagating power is not only a conceptual challenge for modeling, but also a practical concern for operation. Electromagnetic energy transported over long distances can couple into beam diagnostics, which may therefore respond not only to the local passage of the beam but also to fields generated elsewhere in the ring. This non-local sensitivity complicates the interpretation of measurements and can potentially lead to malfunctions. Altogether, FCC-ee operates in a regime

---

* carlo.zannini@cern.ch

where impedance must be understood as a global property arising from the interplay between localized sources and long-range field propagation. Developing a reliable model therefore requires not only accurate descriptions of the beam coupling impedance of individual components, but also a careful treatment of wave propagation above cut-off, including appropriate boundary conditions.

METHODOLOGY FOR IMPEDANCE EVALUATION

In the presence of propagating modes, wakefield simulations become highly sensitive to the longitudinal boundary conditions. As previously discussed, the combination of a large beam pipe diameter and a short bunch length places FCC-ee in a regime that extends well above the beam pipe cut-off frequency.

As illustrated in Fig. 2, when only sub-cut-off modes are excited, the computed impedance is largely insensitive to the choice of boundary conditions. In contrast, the presence of propagating modes introduces a strong dependence on the boundary conditions. Figure 2 highlights the fundamentally different physical nature of modes below and above the beam-pipe cut-off. For frequencies below cut-off, the fields excited in the beam pipe are evanescent and cannot transport power away from the cavity region. As a result, these modes remain confined within the cavity volume, and their resonant frequencies and shunt impedances are essentially independent of the external boundary conditions. This explains the excellent agreement between the electric and perfectly matched layers (PML) solutions observed in the low-frequency region. Above the cut-off frequency, the situation changes significantly. The beam pipe supports propagating modes, allowing electromagnetic energy to travel along it [1]. In this regime, the choice of boundary condition becomes critical. Electric (perfectly conducting) boundaries enforce total reflection of outgoing waves, artificially transforming the structure into a closed resonator and producing narrow, high-Q resonances. In contrast, open boundaries, typically implemented using PML, approximate an open domain by absorbing outgoing power. The resulting resonances are therefore broader, slightly shifted in frequency, and reduced in amplitude. While the open-boundary approach provides a more realistic description, it also complicates the interpretation of the results. A significant fraction of the beam-induced power is transported away from the device, potentially contributing to non-local effects. In this regime, cross-talk with neighboring elements should be considered for a more accurate impedance evaluation.

Propagating Modes and Non-local Effects

Above cut-off, part of the excited field propagates along the beam pipe and can interact with distant components. This leads to non-local impedance contributions, cross-talk between accelerator elements, ambiguity in defining a purely local device impedance. In this regime, the impedance of a device cannot be interpreted independently of its environment, and a distributed description would in principle be required. However, modeling an entire 90 km machine with full electromagnetic detail is not computationally feasible.

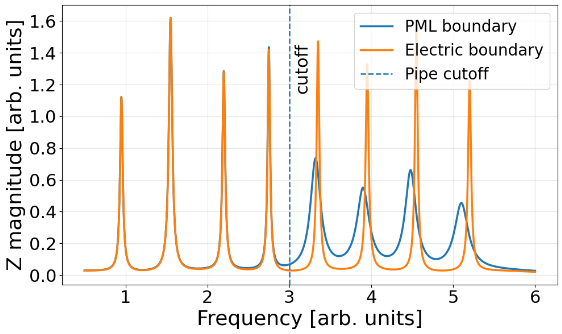

Figure 2: Comparison of impedance computed with electric (perfectly conducting) and open (PML) boundary conditions. Significant differences arise in the presence of propagating modes.

Several strategies can therefore be considered:

- *Artificial localization:* enforcing reflective boundaries at the entrance and exit of each device. This approach suppresses propagation but leads to a significant overestimation of impedance and is therefore not suitable.
- *Use of absorbers or couplers:* introducing dedicated elements to damp propagating modes and reduce cross-talk. While effective, this approach increases design complexity and is not always practical, although it remains relevant for specific cases.
- *Pragmatic approach:* allowing propagation while still analyzing individual elements (or small assemblies). In this case, results must be carefully interpreted by separating locally dissipated power from propagating power.

The latter approach is discussed in this work. The key point is to disentangle the power deposited locally in the device from the power carried away by propagating modes. The latter can be considered as distributed along the machine and must be accounted for in global power-loss estimates.

Power Decomposition To clarify the interpretation of the simulation results, it is useful to distinguish between total and propagating power. The total beam-induced power is obtained from the impedance using the standard formalism [2, 3], while the locally deposited power is computed from material losses within the structure. The propagating power is then determined as the difference between the two.

Figure 3 shows an illustrative example of an FCC-ee collimator [4] with a taper angle of 15°. As observed in the top plot, below the beam-pipe cut-off frequency the propagating power is negligible, and the small fluctuations are within the uncertainty of the evaluation methodology. Above cut-off, however, a significant fraction of the power is found to propagate. As a validation of the method, simulations with electric boundary conditions are used as a reference (see

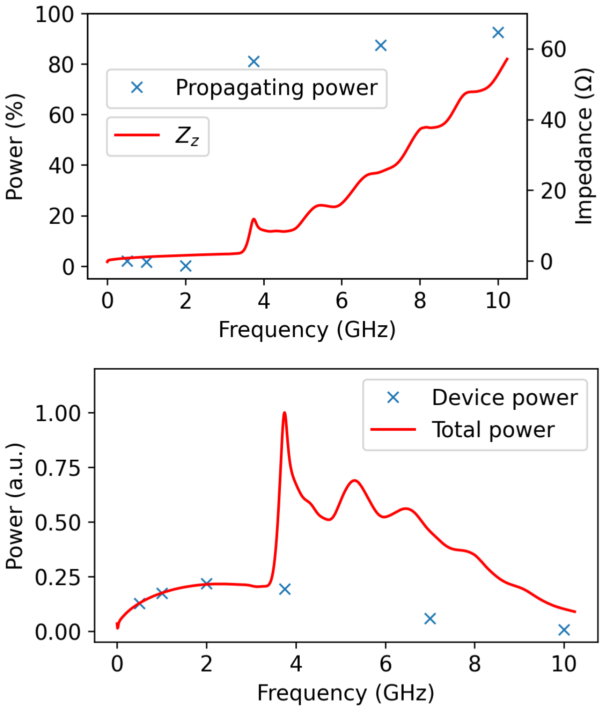

Figure 3: Illustrative example of an FCC-ee collimator with a taper angle of 15°. Top: real part of the impedance together with the fraction of propagating power. Bottom: total beam-induced power and locally deposited power, whose difference defines the propagating contribution.

Fig. 4). In this case, all power is artificially confined within the structure, and the total power coincides with the locally deposited power. This provides a useful consistency check for the power evaluation procedure.

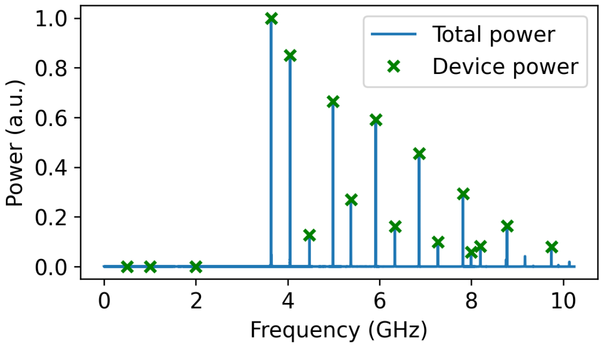

Figure 4: Illustrative example of an FCC-ee collimator with a taper angle of 15°. Validation of the power evaluation method using electric boundary conditions. The total power is shown in blue, while the device (locally deposited) power is indicated by green crosses.

Example of Application

The method has been applied to the representative case of FCC-ee collimators. Figure 5 shows that design parameters, such as the taper angle, influence not only the total beam-induced power but also its partition. For larger taper angles, a significant discrepancy between total and device power is observed over a broad frequency range, indicating enhanced propagation of electromagnetic energy. In contrast, smaller taper angles reduce this difference, suggesting a mitigation of propagating power.

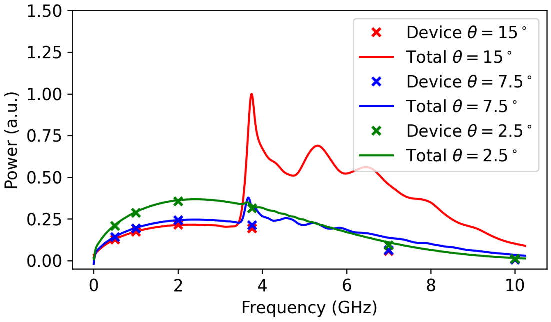

Figure 5: Comparison of total (lines) and locally deposited (device) power (markers) for different taper angles of the FCC-ee collimator under study.

This highlights an additional design criterion: beyond minimizing beam-induced power and related instabilities, geometries can also be optimized to limit the fraction of power that propagates away from the device, thereby reducing non-local effects. Such effects are expected to complicate both the modeling of beam coupling impedance and the overall operability of the machine.

CONCLUSION

FCC-ee operates in a regime where a significant fraction of the beam-induced fields lies above the beam-pipe cutoff frequency. This results in the excitation of propagating modes, giving rise to non-local effects and a strong sensitivity of the impedance to boundary conditions.

In this context, a meaningful interpretation of impedance requires distinguishing between locally deposited and propagating power. A dedicated methodology to perform this separation has been established, and consistency checks have been carried out to validate the approach.

The analysis presented in this work shows that geometrical details can strongly influence not only the total impedance but also the fraction of power that propagates along the machine. This introduces an additional design criterion: beyond minimizing beam-induced power and related instabilities, accelerator components should also be optimized to limit the propagating power.

A comprehensive approach combining accurate electromagnetic simulations, appropriate boundary conditions, and careful interpretation of power deposition is therefore required to ensure reliable predictions and robust machine performance.

REFERENCES

[1] R. E. Collin, *Foundations for Microwave Engineering*. New York: McGraw-Hill, 1992.

[2] C. Zannini, G. Iadarola, and G. Rumolo, “Power Loss Calculation in Separated and Common Beam Chambers of the LHC”, in *Proc. IPAC’14*, Dresden, Germany, Jun. 2014, pp. 1711–1713. doi:10.18429/JACoW-IPAC2014-TUPRI061

[3] E: de la Fuente, L. Sito, F. Giordano, G. Rumolo, B. Salvant, and C. Zannini, “A Python Package to Compute Beam-Induced Heating in Particle Accelerators and Applications”, no. 68, pp. 611–614, Mar. 2024. doi:10.18429/JACoW-HB2023-THBP52

[4] D. Gibellieri *et al.*, “Impact of the geometric impedance of collimators on beam stability in fcc-ee”, in *Proc. IPAC’25*, Taipei, Taiwan, pp. 2137–2140, Nov. 2025. doi:10.18429/JACoW-IPAC2025-WEPM074