Proceedings of the Seventh International Conference on
Railway Technology: Research, Development and Maintenance
Edited by: J. Pombo
Civil-Comp Conferences, Volume 15, Paper 19.4
Civil-Comp Press, Edinburgh, United Kingdom, 2026
ISSN: 2753-3239, doi: 10.4203/ccc.15.19.4

Open-Source Rail Track Modelling Toolbox: Using Phi-Filter and Multiple Modal Bases for Accelerating Frequency Response Calculation

M. Ammann, M. Meguenni and J. Cugnoni

HES-SO University of Applied Sciences and Arts Western Switzerland, Yverdon-les-bains, Switzerland

Abstract

Switzerland has reduced railway noise exposure for 80% of affected residents since 2000, but challenges persist for the remaining 20%. To support this goal, the Rail Track Modelling Toolbox, developed with the Swiss Federal Office for the Environment (FOEN) and Swiss Federal Railways (SBB), provides open-source 3D finite element models to optimize vibration and noise mitigation. However, high computational costs (hours to days per simulation) limit scalability. This study addresses this challenge by introducing two key innovations: a phi-filter that prioritizes eigenvector components alongside modal mass to improve high-frequency accuracy and multiple modal bases to adapt to frequency-dependent material properties, reducing memory usage and simulation time. Results demonstrate that combining these methods achieves near-reference precision with 25% of the original modal basis, cutting memory requirements from 90 GB to 11 GB and simulation time from 4 days to 5 hours. These advancements enable high-fidelity simulations for soft track systems (e.g., polyurethane rail pads) up to 5000 Hz, facilitating cost-effective noise reduction strategies.

Keywords: railway noise mitigation, finite element modelling, modal basis reduction, viscoelastic materials, computational efficiency, open-source.

1 Introduction

While Switzerland has protected 80% of residents affected by railway noise since 2000, ongoing efforts aim to address the remaining 20% through innovative measures and research. To achieve this, various innovative components have been developed to mitigate noise radiation, dampen ground vibrations, and lower track maintenance costs by protecting the ballast. To address this challenge, an open-source toolbox has been developed [1-3] during multiple projects to support infrastructure managers and industries in designing and prototyping solutions to mitigate railway noise emissions. Those projects, mandated by the Swiss Federal Office for the Environment (FOEN) and carried out in partnership with Swiss Federal Railways (SBB), leverages advanced 3D finite element models to estimate noise generation mechanisms and propagation. The toolbox provides high-fidelity simulations that enable stakeholders to evaluate the impact of novel designs and optimize noise-reduction strategies.

However, the computational cost of such detailed simulations cannot be neglected, with individual runs taking hours or even days to complete. The Multi-Sleeper Model, the main model of this toolbox, performs frequency-domain simulations, with frequency dependent viscoelastic materials of large-scale rail tracks using the open source finite element (FE) solver Code_Aster (www.code_aster.org). This FE-based 3D vibro-acoustic model includes the rails, the rail pads, the sleepers, the under-sleeper pads (USP), the clamps and the ballast (modelled by discrete springs). The acoustic calculation is built on acoustic monopole sources superposition or by exporting the velocity fields to BEM (ToPNoise model).

Four main techniques efficiently reduce simulation time:

1. Dynamic substructuring: it consists in creating a macroelement instanced several times to build a larger model. The goal is to run the harmonic analysis in generalized coordinates using a light but rich modal basis of the macroelement, composed of eigenmodes and static interface modes with the Craig-Bampton method [4].
2. Eigenmodes pre-calculation and reuse: the eigenmodes of the macroelement are pre-computed once in a separate simulation (phase 1). The main harmonic simulation (phase 2) reuses those modes to rebuild an updated modal basis of the macroelement quickly (assumption of constant modal basis for small change in material properties).
3. Modal basis reduction by selecting only the modes with the highest modal mass, typically 95% of the total mass is kept.
4. Parallelization: during the main harmonic simulation, multiple jobs are run simultaneously, each job simulating a different part of the frequency spectrum defined by the user.

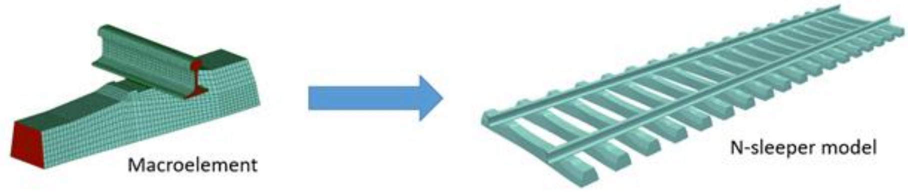

Figure 1: From macro-element to full track model.

The load is applied at the quarter span on the top of the rail (one fourth of the distance between two consecutive sleepers), red points on the figure 2. The model is composed of 400 sleepers, 800 macro-elements, and the load is applied on the 201st sleeper on only one rail. The force is applied vertically or laterally depending on the case.

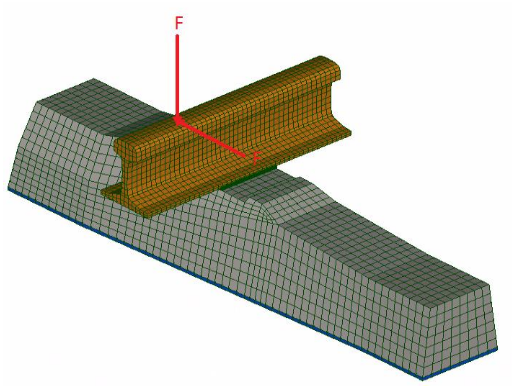

Figure 2: Macro-element with loading point.

Despite those technics, if the rail track support become very soft (i.e. soft rail pads and soft under-sleeper pads) and the frequency range goes higher than 2000Hz, the convergence of the dynamic sub-structuring method is poor due to low track decay rate (TDR) and very heterogeneous vibration shapes. Thus, this approach requires either a full modal basis leading to a prohibitive amount of modes and memory to converge or leads to a low precision at high frequency >2000 Hz if the modal basis is filtered with the modal mass.

2 Methods

2.1 Limitations of Modal Mass Filtering

The modal mass-based filter appears insufficient, particularly for high-frequency responses. To address this, we explored alternative filtering methods. A posteriori techniques are suboptimal because they require running the resource-intensive simulation before reducing its complexity by modal basis reduction methods. Therefore, an a priori technique was necessary.

2.2 Phi-Filter: Eigenvector-Based Mode Selection

To address the limitations of the existing filtering approach and identify a more effective mode selection strategy, we began by considering the general expression of the Frequency Response Function (FRF) in the modal domain [5]:

$$H_{ij}(\omega) = \sum \frac{\phi_i(r) \cdot \phi_j(r)}{m_r(\omega_r^2 - \omega^2 + 2i\zeta_r\omega_r\omega)} \tag{1}$$

with,

$\phi_i(r)$, $\phi_j(r)$: Components of the mode shape r at DOFs i and j
m_r : Modal mass associated with mode r
ω_r : Natural frequency of mode r
ζ_r: Damping ratio of mode r
ω : Excitation frequency (rad/s)

Since our objective is to evaluate the response over the entire frequency range, we focused on the terms independent of the excitation frequency. Consequently, the most relevant quantities for mode selection are the modal mass and the eigenvectors.

While the original filtering method—based solely on modal mass—is a reasonable approach, we investigated whether incorporating information from the eigenvectors, could enhance mode selection and, thus, improve the accuracy of the computed response. To achieve this, we retained the modal mass filter and introduced an additional filter based on the eigenvector components. This new filter, referred to as the *phi-filter*, is constructed as follows:

We focused on the point mobility of the system at a location corresponding to a quarter-span of the rail (i.e., one quarter of the distance between two sleepers). Based on several convergence studies, the point mobility at quarter span is a good indicator of the general convergence of the model as most structural modes have a significant contribution at that point. Thus, the FRF is evaluated at that particular single point *i*, and for point mobility *i =j*. The phi-filter computes the value of $\frac{\phi^2}{m}$ or each mode at the considered point. Two lists of modes are then created: one sorted by $\frac{\phi^2}{m}$ in descending order, and another sorted by modal mass in descending order.

To combine the filters, the user defines the proportion of modes to be selected from each list (e.g., 75% from the phi-filter list and 25% from the modal mass list). For example, if 100 modes are requested, 75 modes are selected from the top of the phi-filter list, and the remaining 25 modes are selected from the modal mass list, ensuring no duplicates. If duplicates are encountered, the next mode in the modal mass list is added until the desired number of modes is reached

2.3 Multiple Modal Bases for Frequency-Dependent Materials

In parallel, to further reduce memory usage and simulation time, we explored a second solution: using multiple modal bases to cover a wide frequency range. Since material parameters are frequency-dependent, the modal basis is computed based on the material properties at a given reference frequency, typically 1000 Hz. This frequency was initially chosen because it provided accurate results up to 1500 Hz, which was the upper limit of our previous harmonic simulations.

However, as the frequency range of interest extends beyond 1500 Hz, the precision of the simulation decreases, particularly above 3000–3500 Hz. To address this, we introduced the use of multiple modal bases, each computed at a different reference frequency, to cover the full frequency spectrum of interest.

To implement this feature in the finite element model, the input parameter is no longer a single frequency for defining material properties. Instead, it is a list of center frequencies for estimating the material properties of the N modal bases. Additionally, a list of cut-off frequencies is provided to select the appropriate modal basis for each frequency band of the harmonic simulation.

The final implementation in the multisleeper model supports two modal bases, which was found sufficient for frequency range up to 5kHz. This approach minimizes additional computation time, as calculating a modal basis is computationally expensive, and reducing the number of bases is therefore advantageous. The tradeoff solved by this multi modal basis method is that the computational cost of calculating a second modal basis on the macro element is offset by a much more efficient modal basis requiring a lower number of modes for the full harmonic simulation.

3 Results

3.1 Phi-filter

To evaluate the effectiveness of the new phi-filter, a series of simulations were conducted on a macro-element containing a total of 1893 modes up to 15kHz. The model is composed of 400 sleepers, 800 macro-element. The macro element has EVA standard Swiss rail pad, rail 60E2, sleeper B91 with USP. The load is applied vertically and horizontally (two simulations), to extract the point accelerance at the quarter span of the 201st sleeper. The following mode counts were tested: 300, 600, 900, 1200, and 1893 modes. For each subset, both the modal mass filter and the phi-filter were applied with varying selection ratios:

1. **100% phi-filter**: Only the phi-filter is used.
2. **75% phi-filter**: 75% of the selected modes come from the phi-filter, and 25% from the modal mass filter.
3. **50% phi-filter**: Equal distribution between both filters.
4. **25% phi-filter**: 25% from the phi-filter, and 75% from the modal mass filter.
5. **0% phi-filter**: Only the modal mass filter is used.

These simulations were performed for the vertical loading case, with a subset of tests also conducted for the lateral loading case (25%, 50%, and 75% phi-filter, for 300, 600, 900, and 1893 modes).

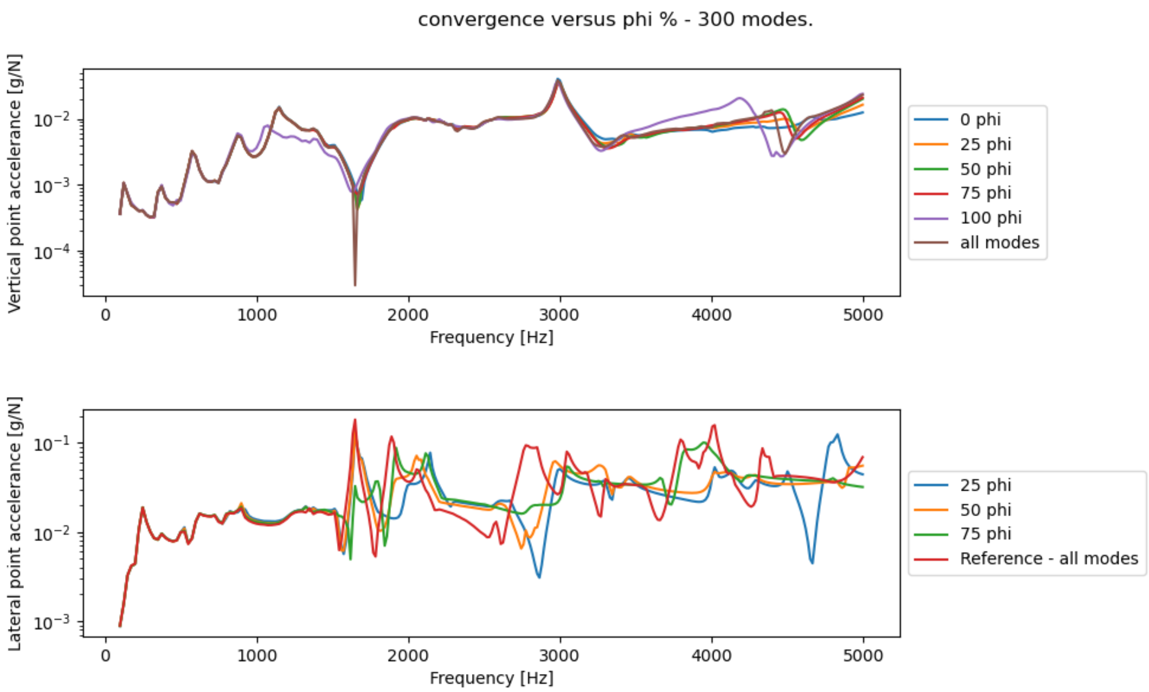

Figure 3: Comparison of the different ratio of modes selected with the Phi-filter for a given number of modes, for a modal basis of 300 modes.

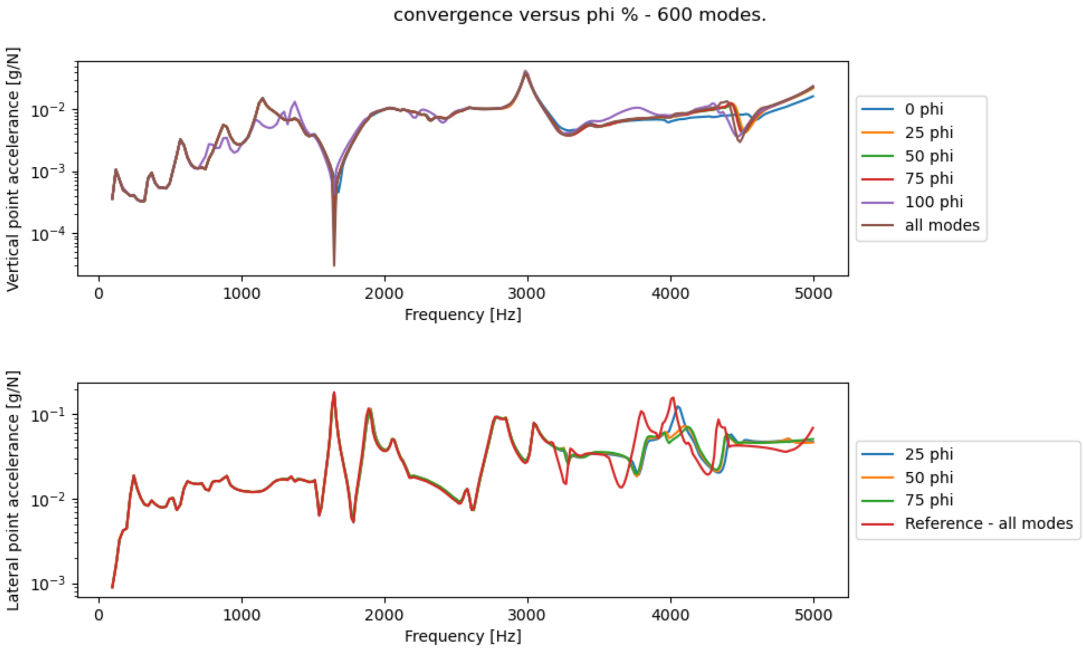

Figure 4: Comparison of the different ratio of modes selected with the Phi-filter for a modal basis of 600 modes.

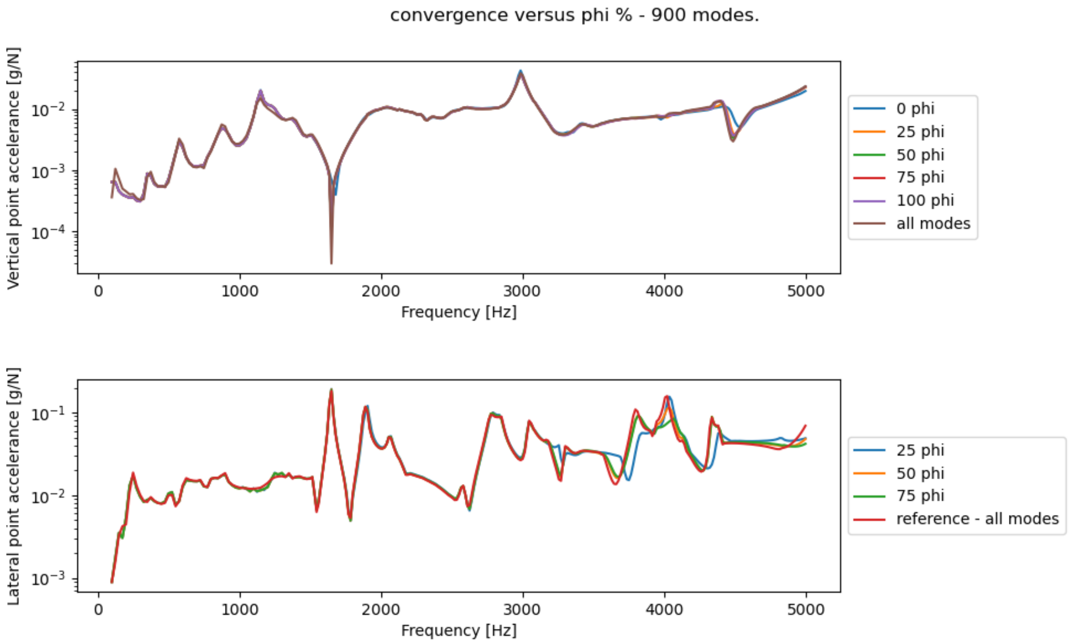

Figure 5: Comparison of the different ratio of modes selected with the Phi-filter for 900 modes.

The first conclusion drawn from these simulations is that the Phi-filter significantly improves convergence for the vertical loading case. For any number of modes, the selection based solely on the mass criterion consistently gives the poorest results.

Introducing just 25% of modes selected with the Phi-filter already increases the accuracy, even for a modal basis of 300 modes (the smallest case tested).
The second conclusion for the vertical loading case is that, when using only the mass filter, a minimum of 900 modes is required to obtain a reliable solution. In contrast, with 25% of modes selected via the Phi-filter, 600 modes are sufficient, and even 300 modes capture the correct trend, although the curve is not yet converged to the reference.

The third conclusion is that increasing the proportion of modes selected by the Phi-filter systematically improves convergence, up to a certain point. The case with 75% of modes chosen using the Phi-filter provides the best agreement with the reference solution. Interestingly, the case with 100% of the modes selected with the Phi-filter shows poor convergence when using fewer than 900 modes but becomes very accurate beyond this threshold. This behaviour can be explained by the fact that, at low frequencies, the mass contribution dominates — hence the mass filter performs better in this range. Consequently, keeping a small percentage of modes selected with the mass filter is beneficial for the vertical loading case.

When considering the lateral loading case, the conclusions must be slightly adjusted. In this configuration, the best compromise is obtained with a 50%-50% ratio between Phi-filter and mass filter selections. Convergence only becomes satisfactory from 900 modes onwards, indicating that a larger number of modes is required for the lateral case compared to the vertical one. This also suggests that the modal mass plays a more significant role in the lateral response, which explains why a balanced combination of Phi- and mass-based selection yields better results in this situation.

The full modal basis is composed of 1893 modes, then 900 modes represent 48% of the total. Having precise result with only half of the total modal basis is a already a significant gain and allow to achieve much faster harmonic simulation and save memory (more than quadratic scaling in computation cost vs number of modes for harmonic simulation).

3.2 Multiple modal basis

In addition to assessing the phi-filter, we evaluated the benefits of using multiple modal bases to improve accuracy across a wide frequency range. This approach is motivated by the frequency-dependent viscoelastic properties of the rail pads and under-sleeper pads (USPs, when present). During the creation of a modal basis, these properties are evaluated at a fixed reference frequency, typically 1000 Hz. While this approximation yields precise results up to approximately 2000 Hz, accuracy significantly decreases beyond this range, particularly above 3500 Hz. Simply increasing the number of modes without filtering is not a viable solution, as it leads to a prohibitive increase in memory and computational requirements.

To evaluate this approach, we selected a very soft track configuration, featuring soft polyurethane rail pads and sleepers equipped with USPs. The full modal basis computed at 1000 Hz consists of 3142 modes up to 15kHz, but harmonic simulations could not be successfully executed due to insufficient memory (requiring over 128 GB). Consequently, we gradually reduced the number of modes to identify the threshold at which the point accelerance begins to deviate from the highest resolution solution. The following mode counts were tested: 2000, 1571, 900, 786, and 526 modes.

Three distinct modal bases were computed to cover the frequency range of interest:

1. The first basis was calculated at 1000 Hz and used for harmonic frequencies between 100 and 1750 Hz.
2. The second basis was calculated at 2500 Hz and applied for harmonic frequencies between 1750 and 3500 Hz**.**
3. The third basis was calculated at 4000 Hz and used for harmonic frequencies between 3500 and 5000 Hz

For each modal basis, modes were selected using the combined mass- and phi-filter described in Section 3.1, with 50% of the modes chosen via the phi-filter and 50% via the mass filter, until the desired number of modes was reached.

The following plots show only three cases because from 2000 to 786 modes there is almost no variation is visible on the point accelerance curves.

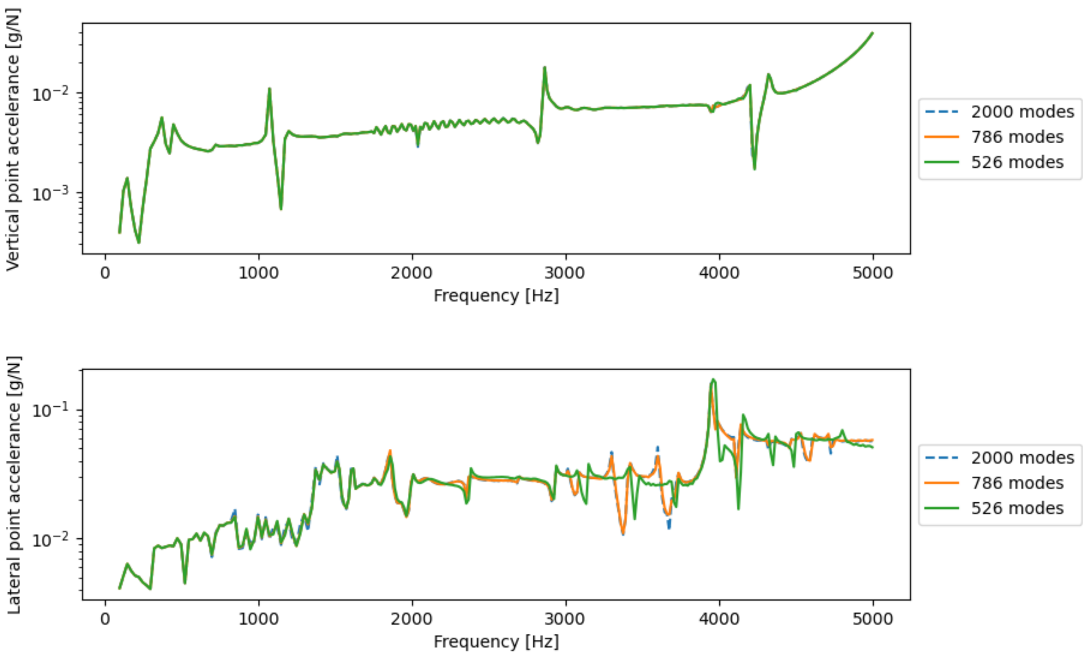

Figure 6: Convergence of the point accelerance versus the number of modes with three modal basis.

The results indicate that, for the lateral loading case, the point accelerance begins to deviate from the reference solution when fewer than 786 modes are used. In contrast, for the vertical loading case, the precision remains high even when reducing the modal basis to 526 modes. This discrepancy highlights the differing sensitivity of vertical and lateral responses to the number of modes retained.

Notably, the plots reveal that no significant variation in point accelerance is observed when reducing the modal basis from 2000 to 786 modes. This suggests that the multiple modal bases approach effectively preserves accuracy while substantially reducing computational demands.

This evaluation demonstrates the performance benefits of using multiple modal bases to achieve more efficient modal bases. By combining the phi- and mass-filters, we previously showed that the modal basis size could be halved without sacrificing precision. With the addition of two supplementary modal bases, we further reduced the required number of modes. In this study, the modal basis was reduced from 3142 to 786 modes—representing only 25% of the original size—while maintaining high precision in both vertical and lateral loading cases.

This reduction translates into significant computational savings, including lower memory requirements, enabling simulations on systems with more limited resources, and faster computation times, as fewer modes need to be processed for each harmonic simulation. The overall computational tradeoff of computing several modal basis is easily offset by the computational cost reduction of the harmonic phase which is globally the bottleneck of this method.

3.3 Final Implementation and Comparative Performance

The final implementation of our method allows users to define two modal bases with customizable parameters. Specifically, users can specify:

- The two frequencies at which the material properties are evaluated for each modal basis.
- The transition frequency that splits the harmonic frequency range, determining which modal basis is used for a given frequency interval.
- The number of modes retained for the harmonic simulation.

To assess the performance of these modifications, we conducted three simulation cases:

1. A reference simulation using all 1893 modes.
2. A simulation using 95% of the modal mass, which retained only 73 modes (approximately 4% of the total).
3. A simulation leveraging the new features, including:

   a. Two modal bases, computed at 1000 Hz and 3500 Hz.
   b. A transition frequency set at 3000 Hz.
   c. A 50% mass-filter and 50% phi-filter ratio.
   d. A total of 460 modes (approximately 25% of the full modal basis).

For all those cases the macro-element is composed with rail 60E2, EVA Swiss standard rail pad, B91 sleeper with USP. The model has 400 sleepers and is loaded at the quarter span of the 201st sleeper.
The following table summarizes the memory requirements**,** number of modes, and simulation times for each case:

SIMULATION CASE	MEMORY REQUIRED	NUMBER OF MODES	SIMULATION TIME
1	90 Gb	1893	4 days and 5 hours
2	3 Gb	73	4 hours
3	11 Gb	460	5 hours

Table 1: Summary of the simulation cases performances.

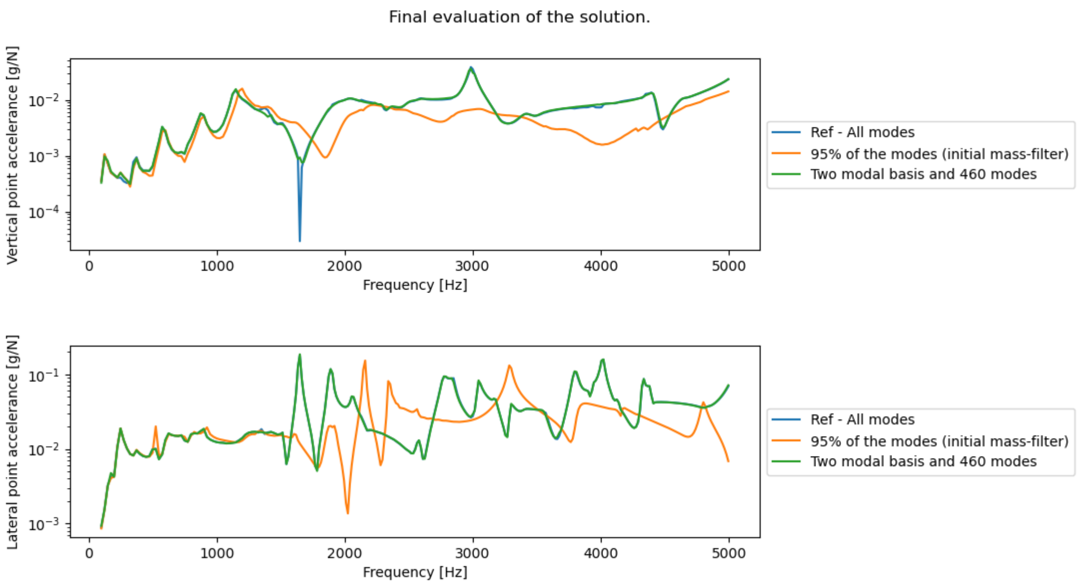

Figure 7: Evaluation of the solution: reference case compared to the original mass filter and the new implementation.

The first observation concerns Simulation Case 2, where retaining 95% of the modal mass results in only 73 modes being selected. This phenomenon arises because the modal mass is concentrated in a small number of modes. The remaining modes contribute very little to the total mass, leading to a threshold effect: increasing the mass percentage slightly (e.g., from 95% to 96%) can suddenly include a significantly larger number of modes, potentially reaching the full modal basis. Consequently, this approach lacks precision, particularly for frequencies above 1500 Hz where many modes have a low modal mass.
In contrast, the new solution (Simulation Case 3), which combines two modal bases and the phi-filter, delivers highly precise results. The Frequency Response Function (FRF) obtained with this method almost perfectly overlaps with the reference solution, despite using only 25% of the total modes. This reduction in the number of modes leads to a significant decrease in memory usage, dropping from 90 Gb to just 11 Gb.
The most striking improvement, however, is in simulation time. The reference simulation (Case 1) required 4 days and 5 hours to complete, whereas the new method (Case 3) achieved nearly identical results in just 5 hours. This represents a reduction in computation time by a factor of 20, making the approach highly efficient for practical applications.

The choice of a balanced ratio between the mass-filter and the phi-filter proved to be particularly effective. This combination ensures that:

1. The modal mass contribution is adequately represented, which is critical for low-frequency behaviour.
2. The phi-filter enhances accuracy at higher frequencies, where modal shapes play a more significant role.

This balance allows the method to capture the essential dynamics of the system while maintaining computational efficiency.

These results demonstrate that the new implementation offers substantial advantages:
Reduced memory requirements, enabling simulations on systems with limited resources.
Faster computation times, making it feasible to perform multiple simulations or parametric studies in a reasonable timeframe.
High precision, even with a significantly reduced modal basis, ensuring that the results remain reliable for engineering applications.

4 Conclusions and Contributions

This study introduces two innovative and complementary methods to overcome the computational bottlenecks in high-fidelity railway noise simulations, enabling faster, more efficient, and precise analyses across a wide frequency range.
Key Contributions

1. **Phi-filter**: A novel mode selection strategy that improves high-frequency accuracy by incorporating eigenvector-based criteria. This approach ensures

that modes contributing significantly to the dynamic response are prioritized, enhancing the precision of simulations without increasing computational cost.

2. **Multiple Modal Bases**: A method tailored to frequency-dependent material properties, enabling accurate simulations up to 5000 Hz. By computing separate modal bases at different reference frequencies, this approach adapts to the varying viscoelastic behaviour of materials, ensuring reliability across the entire frequency spectrum.

The proposed methods yield substantial computational and practical benefits:

1. 75% reduction in modal basis size (from 3142 to 786 modes), drastically lowering the complexity of simulations.
2. 88% reduction in memory usage (from over 128 GB to just 11 GB), making high-fidelity simulations accessible on standard workstations.
3. 92% reduction in simulation time (from 4 days to 5 hours), enabling rapid iterative design and optimization cycles.

These advancements are particularly timely in the context of Switzerland's noise reduction targets. By making high-fidelity simulations more accessible, this work empowers infrastructure managers and engineers to optimize track designs iteratively, balancing noise mitigation, cost efficiency, and maintenance requirements. The integration of these methods into an open-source toolbox further democratizes their use, providing a practical solution for real-world railway engineering challenges.

The ability to perform faster, memory-efficient, and accurate simulations not only accelerates the design process but also supports the development of quieter and more sustainable rail infrastructure, aligning with both national and international environmental goals.

Last but not least the modelling toolbox is open-source and free to use for anyone. It can be download on the github web page [1].

Acknowledgements

This work was supported by the Swiss Federal Office for Environnement in the framework of project no 1337000683 "Tool for the Prediction of railway noise and vibrations (ToP-noise)". The authors would like to thank Mehdi Meguenni & Joël Cugnoni from HEIG-VD, Slimane Ouakka & Chakib Mohamed Drias & Bart Van Damme from EMPA, Jakob Oertli and Marc Schmid from SBB for their contributions to this research.

References

[1] Rail Track Modelling Toolbox, https://github.com/jcugnoni-heig/RailTrackModellingToolbox/tree/TopNoise_ongoing

[2] Ammann, M.; Nardin, R. and al. Railways modelling toolbox: open source finite element models for rail pads and track optimisation. Presented at the Proceedings of the fifth international conference on railway technology:

research, development and maintenance, Montpellier, August 22-25, 2022; Pombo, J., Ed.; Civil-comp conferences; https://doi.org/10.4203/ccc.1.20.4

[3] Nardin, R.; Ammann, M. and al. Open source vibro-acoustic finite element modelling of rail tracks components. Presented at the Proceedings of the fifth international conference on railway technology: research, development and maintenance, Montpellier, August 22-25, 2022; Pombo, J., Ed.; Civil-comp conferences; https://doi.org/10.4203/ccc.1.22.10

[4] Roy R. Craig Jr. and Mervyn C. C. Bampton, Coupling of Substructures for Dynamic Analysis. AIAA Journal Vol. 6, No. 7, July 1968

[5] Del Pedro M.; Pahud P. Mécanique vibratoire: Systèmes discrets linéaires. Presses polytechniques et univertsitaires romandes (PPUR). 14.10.1993