MIMO ROOM IMPULSE RESPONSE ANALYSIS AND SYNTHESIS THROUGH A MODAL-DOMAIN APPROACH: APPLICATION TO A REVERBERATION CHAMBER

Giulia Fratoni[1], Michele Ducceschi[1]

[1] *Dept. of Industrial Engineering, University of Bologna, Italy*

This paper is a revised version of the authors' submitted manuscript that was peer-reviewed by at least two anonymous reviewers.

Abstract

A computational pipeline for modal analysis and synthesis of room impulse responses (RIRs) is presented, with application to a reverberation chamber. Modal frequencies and damping coefficients are estimated from measured RIRs using a sub-band least-squares complex frequency-domain (LSCF) algorithm. Synthesis is performed as a direct modal sum in the time domain. The method extends previous single-channel work to a multiple-input multiple-output (MIMO) setting comprising 36 source–receiver pairs. Poles extracted independently for each measurement are consolidated into a common frequency basis via a voting procedure across spatial positions. Residual amplitudes are then obtained for each measurement by banded least-squares fitting. The resulting parametric model yields per-band T_{60} statistics with inter-measurement uncertainty. Reconstructed RIRs are validated against measurements in both frequency and time domains.

Keywords: *synthetic room impulse response, modal analysis, reverberation chamber, least-squares complex frequency-domain, MIMO*

1. Introduction

Reverberation chambers are specialised acoustic laboratories featuring hard reflective boundaries designed to create a diffuse sound field [1], [2], [3]. These laboratories, in which sound waves ideally propagate with equal probability at all angles of incidence, are essential for measuring the sound absorption coefficients of materials in accordance with ISO 354 [4]. Standard methods rely on decay-rate analysis in third-octave bands and are known to overestimate the absorbing properties of specimens, particularly at low-mid frequencies [5]. A modal description of sound absorption — one that resolves individual mode frequencies and damping rather than assuming a spatially uniform decay — would address several of these shortcomings.

This work takes a first step in that direction by presenting a MIMO modal identification pipeline compatible with the multi-source, multi-receiver configurations prescribed by ISO 354 [4]. The present work extends the modal-domain RIR reconstruction framework introduced in [6], which operated on single source–receiver pairs, to a full MIMO configuration involving 36 source–receiver combinations per measurement condition. As in the earlier work, modal parameters are extracted using a frequency-domain, sub-band implementation of the LSCF method, and the RIR is reconstructed via direct time-domain modal synthesis. Poles extracted independently from each measurement are consolidated into a common basis through a frequency-bin voting procedure: a mode enters the basis only if it is identified in a sufficient fraction of the source–receiver pairs. Residue amplitudes are then refitted per measurement using least-squares on the shared pole set. The resulting representation yields per-band T_{60} statistics across the full spatial sample. Since the modal frequencies are set by the room geometry and the damping coefficients by the boundary conditions, comparing the common basis between the empty room and the room with the specimen would give a per-mode measure of the additional absorption introduced. This paves the way for more reliable estimation of absorbing properties, which remain key boundary conditions in any room acoustics modelling techniques [7], [8], [9].

2. Theoretical background

A room impulse response (RIR) can be expressed as a superposition of decaying sinusoids, each associated with one acoustic mode of the enclosure. In the frequency domain, the corresponding transfer function $H(j\omega)$ admits a partial-fraction expansion over conju-

Corresponding author: *michele.ducceschi@unibo.it.*

gate pole pairs:

$$H(j\omega) = \sum_{p=1}^{P} \left[\frac{r_p}{j\omega - \lambda_p} + \frac{r_p^*}{j\omega - \lambda_p^*} \right], \qquad (1)$$

where P is the number of retained modes, $\lambda_p = -c_p + j\sqrt{\omega_p^2 - c_p^2}$ is the complex pole of the p-th mode (real poles are dropped), $\omega_p = 2\pi f_p$ is its natural angular frequency, c_p its damping coefficient (in rad/s), and $r_p = a_p + jb_p$ the associated complex residue. The reverberation time of mode p follows from the damping coefficient as $T_p = 3\ln 10/c_p$.

The poles are extracted using the LSCF algorithm [10], a frequency-domain estimator that fits a right-matrix-fraction polynomial model to the measured transfer function. Let $z = e^{j\omega T}$, with T a time constant; the model reads

$$H(z) \approx \frac{B(z)}{A(z)} = \frac{\sum_{m=0}^{N} b_m z^m}{\sum_{m=0}^{N} a_m z^m}, \qquad (2)$$

where N is the polynomial order. The denominator coefficients $\boldsymbol{\alpha} = [a_0, \ldots, a_N]^\top$ are obtained from a linear least-squares problem on the measured frequency response, whose normal equations are solved via a Schur-complement reduction with amplitude-invariant Tikhonov regularisation [6]. The poles of $A(z)$ are then recovered as eigenvalues of a companion matrix and mapped to continuous-time via $\lambda = -\ln(z)/T$; only upper-half-plane ($\mathrm{Im}(\lambda) > 0$) and left-half-plane ($\mathrm{Re}(\lambda) < 0$, physically damped) poles are retained.

To separate physical poles from spurious mathematical artefacts, the polynomial order is swept over a prescribed range $[N_{\min} : \Delta N : N_{\max}]$ and a stabilisation diagram is constructed: a pole is accepted as stable when, between two consecutive orders, both its relative frequency shift and its relative damping shift remain below specified tolerances ϵ_f and ϵ_c. An additional constraint rejects any pole whose T_{60} exceeds a ceiling derived from the measured reverberation time of the room. For a detailed derivation of the LSCF formulation and its application to single-input single-output RIRs the reader is referred to [6].

Once the poles $\{\lambda_p\}$ are known, the residues $\{r_p\}$ are determined by a second, independent least-squares solve. From eq. (1), substituting $r_p = a_p + jb_p$ leads to a linear system in the unknowns $[a_1, \ldots, a_P, b_1, \ldots, b_P]^\top$, which is solved per frequency band with Tikhonov regularisation. The time-domain RIR is then synthesised directly via the causal modal sum

$$h_{\mathrm{syn}}(n) = 2\,\mathrm{Re}\left\{ \sum_{p=1}^{P} r_p \exp(\lambda_p n / f_s) \right\}, \qquad (3)$$

where f_s is the sampling rate. Alternatively, the impulse response may be reconstructed using a finite-difference recursion with an exact mapping of the continuous poles, as detailed in [6].

3. Method

The method relies on two main steps: in the first, RIRs are measured in the reverb chamber following standardised procedures; in the second, the room modes are extracted via a custom MIMO implementation of the LSCF method.

3.1. Lab measurements

The authors conducted a measurement campaign in the reverberation chamber of the Department of Industrial Engineering at the University of Bologna (see Fig. 1) [5]. With the exception of the access door,

Figure 1: View of the reverberation chamber under study (top). Sources (S) and receivers (R) layout of reverberation time measurements (bottom).

all boundary surfaces are made of reinforced concrete to approximate an infinite acoustic impedance condition relative to the propagation medium, i.e., the air ($Z \approx 415$ Rayl). The walls are intentionally non-orthogonal to limit strong modal behaviour associated with the room's natural resonant frequencies (standing waves) [11], [12], [13]. Furthermore, while conventional designs often employ suspended diffusers to enhance scattering, the facility under study utilizes a contoured, pre-fabricated ceiling to achieve the necessary degree of spatial uniformity. To prevent high-frequency acoustic absorption caused by surface porosity, a specific coating was applied to the room's walls and ceiling. For the same reason, the room is kept consistently free of dust and particles, with the floor cleaned before each use. Located in the building's basement, the room experiences negligible external background

Table 1: Main geometrical data of the reverberation chamber, sources and receivers layout, and thermo-hygrometric conditions during the measurements campaign (see Figure 1).

Room's features	
Room's volume	277 m^3
Floor surface	81 m^2
Room's height	3.4 m
Sources and receivers	
Sources (Dodecahedrons)	N=2
Sources (Subwoofers)	N=4
Receivers	N=6
Thermo-hygrometric conditions	
Temperature	20.6 °C
Relative Humidity	42.4 %
Pressure	1010.8 hPa

noise, with no significant effect on the measurements. Measurements were performed using:

- 2 custom dodecahedrons, calibrated according to ISO 3741 [14];
- 4 custom subwoofers;
- 6 1/2 inch monaural free-field microphones;
- RME Fireface soundcard;
- LD Systems amplifier.

Impulse responses were obtained using exponential sine sweep signals (1024 k samples at 48 kHz) emitted with Dirac (B&K, version 7). The same commercial software was employed to derive room criteria of interest: in this case, the reverberation time, T_{20}. The temperature was 20.6 °C, and the relative humidity was 42.4%. The combination of 36 source–receiver pairs ensures the statistical reliability of the obtained values. With a measured $T_{20} = 7.5$ s at mid-frequencies (500–1000 Hz), the Schroeder frequency is calculated as follows:

$$f_{Sch} \approx 2000\sqrt{\frac{T}{V}} = 330\text{Hz} \tag{4}$$

where V is the room's volume (see Table 1).

3.2. Computational pipeline

The computational pipeline operates in three sequential stages: (i) single-RIR modal analysis, (ii) cross-measurement pole consolidation on a common basis, and (iii) amplitude re-fitting and statistical post-processing. A schematic of the pipeline is presented in Fig. 2.

Stage 1: Per-RIR modal analysis.

Each of the 36 measured RIRs is processed independently. The raw impulse response is onset-aligned and normalised. The aligned RIR is then transformed to the frequency domain through a zero-padded FFT. The analysis range (20–12 000 Hz) is covered by narrow, linearly-spaced processing bands of fixed bandwidth (25 Hz). The choice of a fixed, narrow bandwidth keeps the polynomial order tractable across the entire spectrum. A two-regime strategy is adopted, splitting the frequency axis at the Schroeder frequency, f_{Sch}. Below f_{Sch}, where individual modes are resolvable, LSCF is applied: the polynomial order is sized from the Weyl modal density estimate ($\mathrm{d}n/\mathrm{d}f = 4\pi V f^2/c^3$) and swept to produce a stabilisation diagram, from which stable poles are identified (relative frequency and damping tolerances set to $\epsilon_f = 0.05$ and $\epsilon_c = 0.10$). Above f_{Sch}, high modal overlap prevents reliable per-mode resolution; spectral peaks are instead detected directly on $|H(f)|$ and assigned a uniform damping derived from the Schroeder T_{20} of the corresponding band. Peaks are detected with the following parameters: minimum peak prominence equal to 0.5 dB and minimum inter-peak spacing equal to 0.5 Hz.

For each processing band, the per-pole damping values are averaged to a single band representative value before the residue amplitudes are computed via the banded least-squares solve described in Section 2. An iterative gain-outlier rejection step (based on the median absolute deviation of the residue magnitudes) removes poles with anomalous amplitudes and re-solves the system, preventing a small number of poorly conditioned poles from biasing the reconstruction. The resulting pole set $\{f_p, c_p, a_p, b_p\}$ for each RIR is stored together with per-band statistics: mode count, modal density, mean T_{60} and its standard deviation, and inter-mode spacing.

Stage 2: Common pole basis.

The 36 pole sets from the empty-room condition are consolidated into a single *common pole basis*. A frequency-domain voting scheme is employed: the analysis range is binned into intervals of width $\Delta f = 0.5$ Hz, and a histogram counts the number of IRs that contribute at least one pole to each bin. A bin enters the common basis only if at least 40% of the IRs vote for it. The representative frequency and damping of the basis pole are taken as the median of the contributing per-IR values. This majority-vote criterion retains modes that are consistently excited across the spatial positions while discarding weakly or spuriously identified poles.

Stage 3: Amplitude re-fit and synthesis.

With a shared set of pole frequencies and condition-specific damping coefficients now established, the residue amplitudes are re-computed for every IR. The poles (f_p, c_p) are held fixed and only the real and imaginary parts of the residues (a_p, b_p) are solved via the banded least-squares formulation of Eq. (1), applied over narrow linear bands (50 Hz width, with 10% overlap at the edges to mitigate boundary artefacts). The selected band size of 50 Hz ensures that a manageable number of peaks are processed within each slice. At low frequencies, it is wide enough to capture

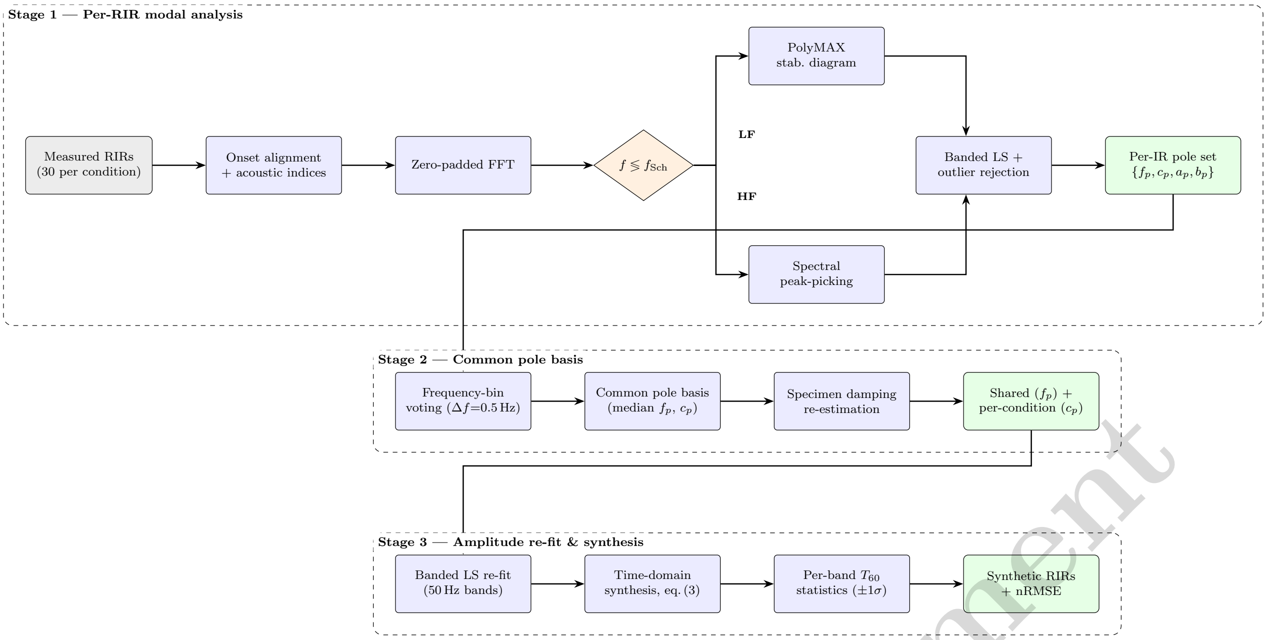

Figure 2: Block diagram of the three-stage modal analysis pipeline. Stage 1 is applied independently to each of the 36 measured RIRs; Stage 2 consolidates the pole sets into a shared frequency basis; Stage 3 re-fits amplitudes on the common basis and produces per-band statistics.

dominant individual modes, whilst at high frequencies, it prevents the matrix size from exploding due to an excessive number of overlapping peaks. This places all 36 IRs of each condition on a common parametric footing, enabling direct amplitude comparison across source–receiver pairs.

4. Results

The present section reports the validation of the MIMO modal identification pipeline applied to the reverberation chamber under study.

4.1. Statistical analysis of global decay

The pipeline produces per-band statistics computed across the 36 IRs. Figure 3 presents a comprehensive statistical evaluation of the reverberation time derived from the modal-domain model compared to classical acoustic indicators. For each ISO third-octave band, the following quantities are reported: mean $T_{60} \pm 1\sigma$ derived from the synthesized RIRs and mean $T_{20} \pm 1\sigma$ derived from measured RIRs using ITA-Toolbox [15]. The thin lines represent the reverberation time calculated for each of the 36 individual source–receiver pairs, highlighting the spatial non-uniformity of the sound field, particularly in the low-frequency (LF) regime where individual modes dominate. Both series include $\pm 1\sigma$ standard deviation bands to quantify spatial uncertainty. In the LF region (before the vertical green line at 330 Hz), the higher variance in the lines and the wider σ bands reflect the sensitivity of the reverberation time to specific modal shapes and source–receiver placement in the non-diffuse regime.

Overall, these results indicate that the discrepancies observed below the Schroeder frequency are likely due to the different treatment of the late decay. The presence of background noise in measured RIRs affects the estimation of the decay slope, particularly in the final portion of the impulse response, leading to slightly longer reverberation time estimates. In contrast, the synthetic RIRs exhibit an uninterrupted exponential decay because they are generated from a modal model without measurement noise. As a result, the late decay continues to decrease without flattening, producing a steeper decay profile and, therefore, lower reverberation times than the measured ones. Moreover, excluding poles with anomalous residue amplitudes may further reduce the estimated reverberation time by removing outlier peaks.

4.2. Spectral and temporal accuracy

Figure 4 provides a high-resolution view of the reconstruction accuracy for a representative source–receiver pair (S1-R1). It provides a dual-domain validation of the synthetic RIR against the measured target data. The upper panel displays the magnitude of the transfer function $|H(f)|$ in decibels in the frequency domain (20–140 Hz). The synthetic response (red) is reconstructed using 5,174 poles over the full bandwidth, and the zoomed view confirms that individual modal peaks at this spatial coordinate are tracked with high precision against the target (black). The lower panel presents a zoomed-in view of the impulse response, $h(t)$, over the initial 20-35 ms of decay. The overlay of the target (black) and synthetic (red) signals re-

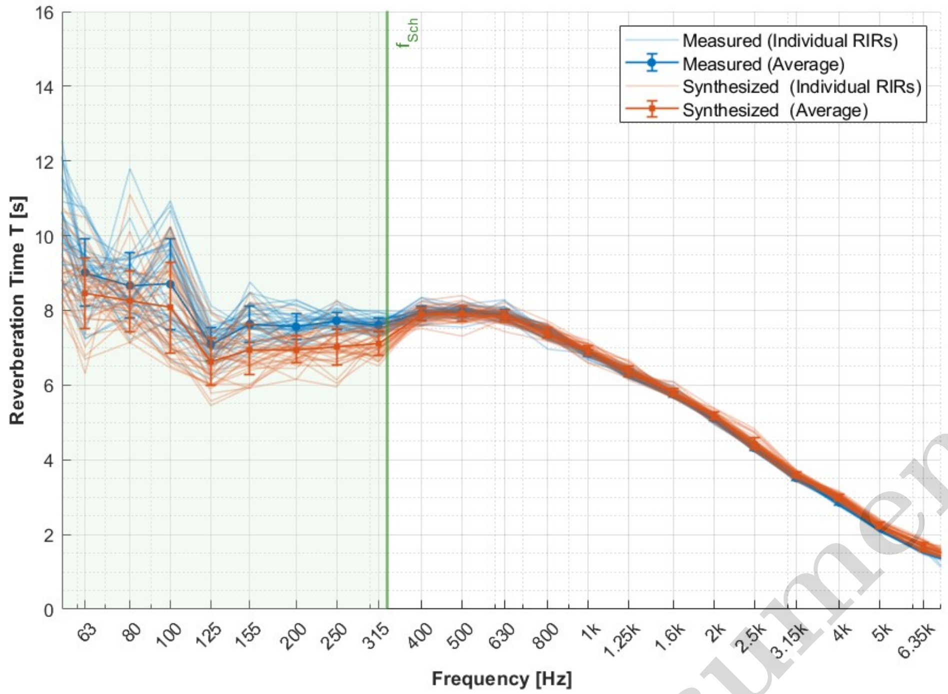

Figure 3: Mean reverberation time values derived from the synthesized (orange) and measured RIRs (blue). Standard deviation (error bars) and the results for each source-receiver pair (thin lines) are also provided.

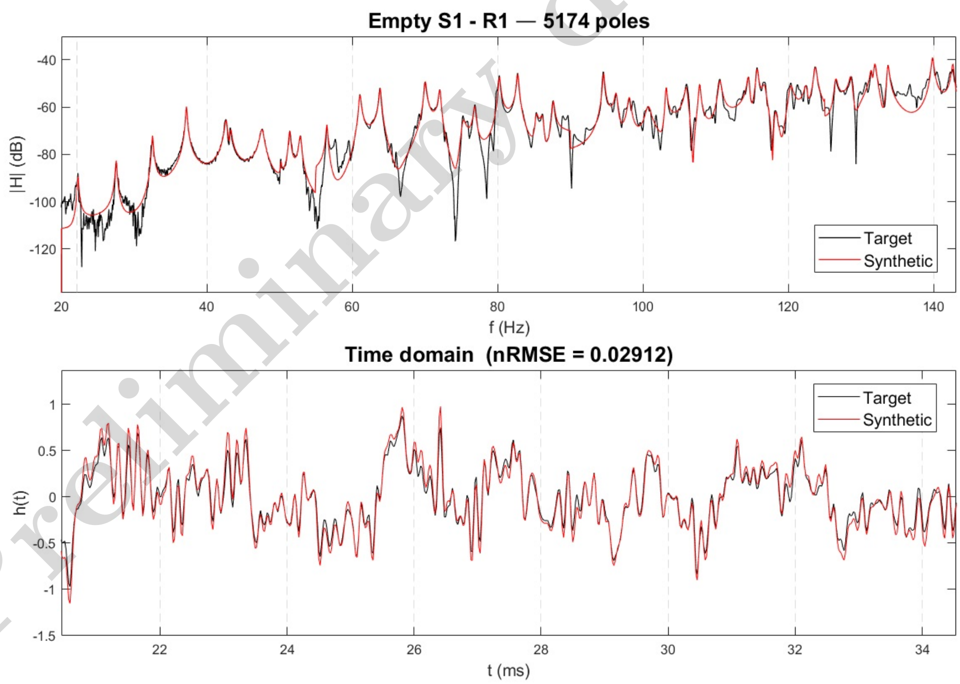

Figure 4: Comparison between target (black) and identified (red) responses in low-range frequency domain (top) and in time domain up to 35 ms (bottom) for the S1-R1 pair, here taken as example (see reverberation chamber's plan in Fig. 1).

veals a high degree of temporal correlation, thereby preserving the early reflection pattern. The accuracy of the synthesis is quantified by a normalised Root Mean Square Error (nRMSE) of 0.029. This low error value indicates that the modal-domain representation—based on consolidated poles and re-estimated damping coefficients—faithfully reproduces the energy decay and phase characteristics of the physical space. Although the poles were derived from a common basis, determined by 36 independent measurements, the specific residue amplitudes for the S1-R1 pair were uniquely determined to match the local pressure field.

5. Conclusion

Initial results demonstrate that modal-based reconstruction provides reliable reverberation time estimates in comparison with the classical Schroeder integration. They demonstrate that the parametric modal model can accurately reconstruct the specific physical behaviour of any given point in the room. As a preliminary study, this research lays the groundwork for evaluating the physical impact of absorbent materials on individual room modes. Future work will explore the reverberation room equipped with a specimen to compare the sound absorption coefficients derived from ISO 354 with the direct variation in modal damping factors.

6. Acknowledgments

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program, Grant agreement No. 950084 NEMUS.

References

[1] W. C. Sabine, *Collected papers on acoustics*. Courier Dover Publications, 2015.

[2] F. Jacobsen, *The diffuse sound field: Statistical considerations concerning the reverberant field in the steady state*. Acoustics Laboratory, Technical University of Denmark, 1979.

[3] M. Nolan, S. A. Verburg, J. Brunskog, and E. Fernandez-Grande, "Experimental characterization of the sound field in a reverberation room," *The Journal of the Acoustical Society of America*, vol. 145, no. 4, pp. 2237–2246, 2019.

[4] *ISO 354:2003 Acoustics — measurement of sound absorption in a reverberation room*, ISO, 2003.

[5] C. Scrosati et al., "Towards more reliable measurements of sound absorption coefficient in reverberation rooms: An inter-laboratory test," *Applied Acoustics*, vol. 165, p. 107 298, 2020.

[6] M. Ducceschi, C. J. Webb, and G. Fratoni, "Efficient synthesis of large room impulse responses in the modal domain," in *2025 Immersive and 3D Audio: from Architecture to Automotive (I3DA)*, IEEE, 2025, pp. 1–8. [Online]. Available: `https://mdphys.org/I3DA_2025.html`.

[7] V. Wittstock, "Determination of measurement uncertainties in building acoustics by interlaboratory tests. Part 2: Sound absorption measured in reverberation rooms," *Acta Acustica united with Acustica*, vol. 104, no. 6, pp. 999–1008, 2018.

[8] G. Fratoni and D. D'Orazio, "Boundary conditions for hybrid simulations in a rectangular environment with sound-absorbing ceiling," *Applied Acoustics*, vol. 240, p. 110 933, 2025.

[9] G. Fratoni, D. D'Orazio, and M. Garai, "Uncertainty of input data for wave-based room acoustic simulations in large non-trivial environments," in *Proceedings of Forum Acusticum 2023*, DOI: 10.61782/fa.2023.1114.

[10] B. Peeters, H. Van der Auweraer, P. Guillaume, and J. Leuridan, "The polymax frequency-domain method: A new standard for modal parameter estimation?" *Shock and vibration*, vol. 11, no. 3-4, pp. 395–409, 2004.

[11] R. H. Bolt, "Normal modes of vibration in room acoustics: Angular distribution theory," *The Journal of the Acoustical Society of America*, vol. 11, no. 1, pp. 74–79, 1939.

[12] D.-Y. Maa, "Distribution of eigentones in a rectangular chamber at low frequency range," *The Journal of the Acoustical Society of America*, vol. 10, no. 3, pp. 235–238, 1939.

[13] P. M. Morse and R. H. Bolt, "Sound waves in rooms," *Reviews of modern physics*, vol. 16, no. 2, p. 69, 1944.

[14] *ISO 3741:2010 Acoustics — determination of sound power levels and sound energy levels of noise sources using sound pressure — precision methods for reverberation test rooms*, ISO, 2010.

[15] M. Berzborn, R. Bomhardt, J. Klein, J.-G. Richter, and M. Vorländer, "The ita-toolbox: An open source matlab toolbox for acoustic measurements and signal processing," in *Fortschritte der Akustik – DAGA 2017*, Kiel, Germany, 2017, pp. 222–225.