CHALMERS
UNIVERSITY OF TECHNOLOGY

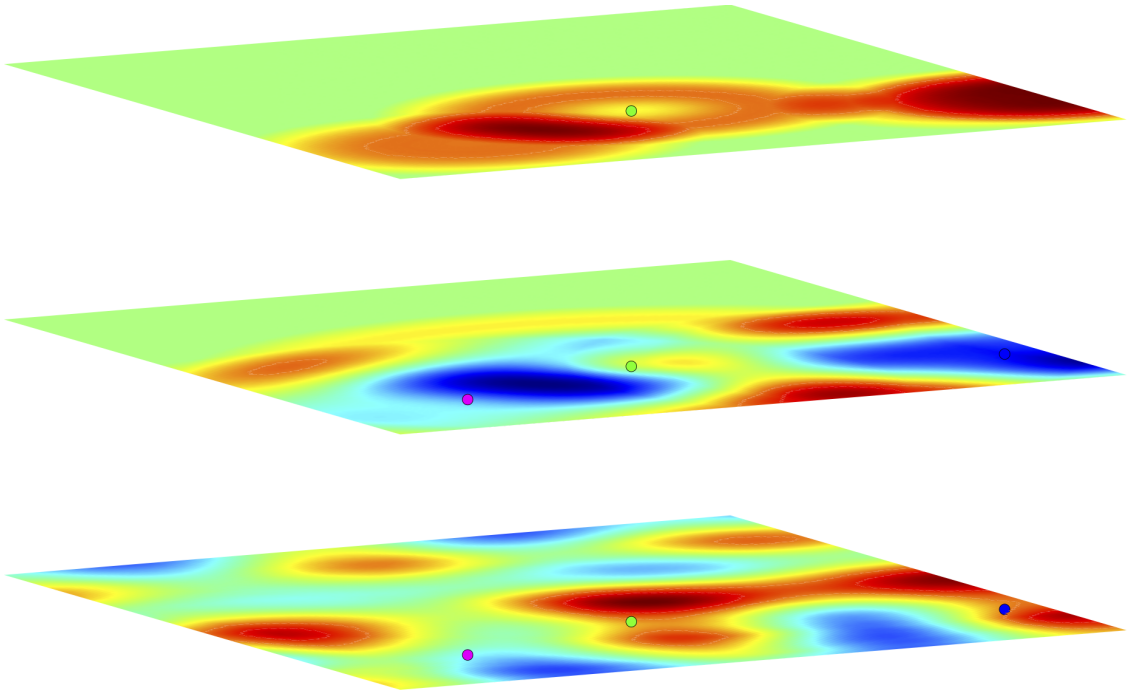

Optimization of Low-Frequency Sound Reproduction in Enclosed Spaces

Master's thesis in Sound and Vibration

JAN FISCHER

DEPARTMENT OF APPLIED ACOUSTICS
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025
www.chalmers.se

Master's thesis 2025

Optimization of Low-Frequency Sound Reproduction in Enclosed Spaces

JAN FISCHER

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Architecture and Civil Engineering
Division of Applied Acoustics
Chalmers University of Technology
Gothenburg, Sweden 2025

Optimization of Low-Frequency Sound Reproduction in Enclosed Spaces
JAN FISCHER

Supervisor: Jens Ahrens, Department of Applied Acoustics
Examiner: Jens Ahrens, Department of Applied Acoustics

Master's Thesis 2025
Department of Architecture and Civil Engineering
Division of Applied Acoustics
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 2210

Cover: Snapshot of the wavefield in a room from three loudspeakers.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2025

Optimization of Low-Frequency Sound Reproduction in Enclosed Spaces
JAN FISCHER
Department of Architecture and Civil Engineering
Chalmers University of Technology

Abstract

Achieving accurate low-frequency reproduction in enclosed spaces is notoriously difficult, as standing waves and modal resonances introduce significant peaks, dips, and extended decay times in the bass range. This thesis presents a multi-loudspeaker corrective strategy that employs adaptive filtering to mitigate these issues, focusing on the 20–100 Hz region where modal effects are most pronounced. By iteratively refining digital filters using a Least Mean Squares (LMS) algorithm, each loudspeaker compensates for problematic resonances, thereby improving the overall frequency response and reducing spatial variability.

The methodology is validated through both virtual and real-world evaluations. In simulations based on modal summation, the proposed approach demonstrates substantial reductions in modal peaks and ringing, establishing an idealized benchmark. Subsequent experiments in an actual listening environment confirm that LMS-based filters effectively smooth out low-frequency irregularities caused by the primary loudspeaker, although the integration of multiple sources highlights challenges related to phase alignment, overlapping modes, and precise calibration requirements. Despite these complexities, the results underscore the viability of multi-loudspeaker adaptive filtering for localized modal compensation, pointing to future avenues such as fully joint multi-channel optimization and machine learning–based filter tuning. Overall, this work advances the development of adaptive room correction systems by illustrating both the benefits and constraints of a self-optimizing, low-frequency–focused design.

room acoustics, low-frequency correction, multi-loudspeaker, itterative filtering, LMS algorithm, room modes analysis, sound reproduction

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

A	Absorption Area
ADC	Analog-to-Digital Converter
BEM	Boundary Element Method
C50	Clarity Index (common in speech/music clarity studies)
DAC	Digital-to-Analog Converter
DSP	Digital Signal Processing
EDT	Early Decay Time
FEM	Finite Element Method
FFT	Fast Fourier Transform
f_S	Schröder Frequency
IR	Impulse Response
LMS	Least Mean Squares (adaptive filter algorithm)
MIMO	Multi-Input Multi-Output
MSE	Mean Squared Error
RG	Room Gain
RIR	Room Impulse Response
RT	Reverberation Time
SNR	Signal-to-Noise Ratio
SPL	Sound Pressure Level
STFT	Short-Time Fourier Transform
STI	Speech Transmission Index
T60	Reverberation Time (time to decay 60 dB)

Nomenclature

Below is the nomenclature of indices, sets, parameters, and variables that have been used throughout this thesis.

Indices

n_x, n_y, n_z	Mode indices along the x, y, and z axes (e.g., $n_x = 0, 1, 2, \ldots$).
t	Discrete time step index (used for signals $x[t]$, $y[t]$, etc.).

Sets

$\mathcal{M}$	Set of acoustic modes (e.g., $\{(n_x, n_y, n_z)\}$).
$\mathcal{S}$	Set of loudspeakers.
$\mathcal{R}$	Set of microphones.

Parameters

A	Total absorption area (m^2).
α	Absorption coefficient (dimensionless).
c_0	Speed of sound in air (typically $343\,\mathrm{m\,s^{-1}}$).
d	Damping constant for modal summation or other low-frequency models.
df	Frequency increment (Hz) based on FFT size and sampling rate.
f_S	Schröder frequency (Hz).
f_1, f_2	Start and end frequencies of a swept-sine signal (Hz).
m	Adaptive filter length (number of LMS coefficients).

M	Total number of modes in a summation approach.
μ	Step size (LMS algorithm convergence parameter).
$NFFT$	FFT size (block length) in frequency-domain processing.
T_{60}	Reverberation time (s), time for sound to decay by 60 dB.
V	Volume of the room (m^3).
Δf	Frequency step (Hz) in FFT-based analyses.

Variables

f_{n_x,n_y,n_z}	Modal frequency for indices (n_x, n_y, n_z) (Hz).
$\frac{\Delta M}{\Delta f}$	Modal density (modes per Hz).
ω	Angular frequency ($\mathrm{rad\,s^{-1}}$), $\omega = 2\pi f$.
$H(\omega)$	Frequency response or transfer function at angular frequency ω.
$\psi_i(r)$	Eigenfunction corresponding to the i-th room mode, evaluated at position r.
K_i	Normalization factor for eigenmodes, often $\int \psi_i(r)\, dV$.
$x(n)$	Input signal at discrete time n.
$d(n)$	Desired (reference) signal at discrete time n.
$y(n)$	Filter (or system) output at discrete time n.
$e(n)$	Error signal at discrete time n, where $e(n) = d(n) - y(n)$.
$w(n)$	Vector of LMS filter coefficients at time n.
$w_1(n), w_2(n), w_3(n)$	Adaptive filter coefficients for loudspeakers S_1, S_2, and S_3.
τ	Time delay (in seconds or samples) found via cross-correlation or alignment.
$\mathrm{RIR}(t)$	Room impulse response in the time domain.
$\mathrm{IR}(t)$	General impulse response in the time domain (e.g., of a filter).
fs	Sampling frequency (Hz).
t_{impulse}	Time axis for plotting or analyzing an impulse response.

Contents

1

Introduction

High-fidelity sound reproduction in enclosed spaces remains a persistent challenge due to the pronounced effects of room acoustics, particularly at low frequencies. In these ranges, standing waves and modal resonances cause significant peaks and dips in the frequency response, leading to uneven sound fields and degraded audio fidelity [1]. Addressing such issues is paramount in professional and consumer contexts alike, from recording studios and concert halls to home theaters.

This thesis proposes a self-optimizing algorithm for low-frequency room correction that capitalizes on multiple loudspeakers. By analyzing room impulse responses and iteratively compensating for modal resonances, the method produces a more uniform low-frequency response across a defined listening area [2].

Room impulse responses are obtained via a swept sine technique, chosen for its high temporal and spectral resolution [3]. The correction filters employ the Least Mean Squares (LMS) algorithm, which adaptively minimizes the discrepancy between desired and measured outputs [4]. By iterating over frequency-dependent delays and amplitude distortions, the LMS-based approach flattens the low-frequency spectrum while preserving overall signal integrity.

Focusing on the sub-400 Hz region delivers two key advantages. First, it targets the domain where room modes are especially disruptive [5]. Second, it lowers computational overhead by avoiding full-range processing. In addition, leveraging multiple loudspeakers not only broadens the correction zone, it also harnesses the potential of multi-point (or MIMO) configurations that distribute the corrective effort among different sources [6].

Conventional single-channel equalization frequently suffers from narrow sweet spots and insufficient spatial coverage. In contrast, the multi-loudspeaker strategy demonstrated herein achieves more robust correction for multiple listener positions, minimizing modal decay time differences and diminishing seat-to-seat variations at low frequencies. Objective metrics—such as response flatness and reduced modal ringing—are accompanied by subjective listening assessments to underscore the system's real-world benefits. Empirical results confirm that the proposed approach surpasses traditional single-channel methods, particularly in multi-listener setups.

Additionally, this thesis situates the new algorithm within the broader context of existing room correction techniques, highlighting where conventional solutions fall

short and how multi-loudspeaker and adaptive strategies can overcome these limitations. By integrating theoretical modeling with practical experimentation, the work offers insight into the design and deployment of real-time acoustic correction systems for the modern audio environment.

Overall, the self-optimizing low-frequency correction algorithm presented here marks a significant advancement in both room acoustics and audio engineering. By systematically managing modal behavior through adaptive, multi-loudspeaker filtering, it paves the way for scalable solutions that address pressing acoustic challenges while maintaining flexibility for evolving listening spaces.

2

Theory

Understanding and addressing the challenges posed by low-frequency room modes requires a solid grasp of acoustical theory and mathematical modeling. This chapter provides an overview of the physical mechanisms that generate low-frequency resonances, the mathematical framework for predicting modal frequencies, and the adaptive filtering methods used to correct room-induced anomalies. It emphasizes how these theoretical foundations guide the custom MATLAB scripts for measuring impulse responses, aligning signals, and designing filters to mitigate problematic modes.

2.1 Fundamentals of Room Acoustics

Room acoustics investigates how sound behaves in enclosed spaces, accounting for both time-dependent (transient) and steady-state (modal) phenomena. Low frequencies are especially prone to forming standing waves (*room modes*), which cause peaks and dips in the frequency response. High and mid frequencies, by contrast, predominantly face reflections, scattering, and diffusion challenges. This section covers core acoustic concepts—reverberation, absorption, and the behavior of reflections—while briefly noting the significance of low-frequency modes, which are detailed further in Section 2.3.

Reverberation Time (RT)

Reverberation time (T_{60}) is the duration required for the sound in a room to decay by 60 dB once the source stops [7]. Sabine's approximate formula is:

$$T_{60} = 0.161 \cdot \frac{V}{A}, \tag{2.1}$$

where: T_{60} = Reverberation time in seconds
V = Volume of the room in cubic meters (m^3)
A = Total equivalent absorption area (m^2)

Concert halls often favor longer RTs (about 2 s), enriching musical performances, whereas lecture spaces aim for shorter RTs (about 0.6 s) to preserve speech intelligibility [8].

Schröder Integration for Reverberation Measurements
An influential method for measuring reverberation time in enclosed spaces was introduced by M. R. Schröder [9]. Rather than analyzing the raw sound-pressure decay directly, this approach relies on what is often termed *Schröder integration*. The foundation is the impulse response $h(t)$ of the room—obtained by any suitable measurement technique (e.g., pistol shot, sine sweep, or maximum-length sequence). Once $h(t)$ is captured, its squared magnitude $|h(t)|^2$ represents the instantaneous acoustic energy in the room at time t. Schroeder's insight was to integrate this squared impulse response in a *reverse* time direction to form a smoothly decaying energy curve:

$$E(t) = \int_t^\infty |h(\tau)|^2 \, d\tau, \tag{2.2}$$

where: E(t) = total remaining acoustic energy in the room at time t in $\mathrm{Pa\,s^2}$
h(τ) = impulse response in Pa

where $E(t)$ is the . Plotting $E(t)$ in decibels versus t yields a near-monotonic decay free from the random fluctuations that often arise if one tries to measure the decay of band-limited noise bursts directly. In practice, one typically applies a logarithmic scale to $E(t)$ (e.g., $10\log_{10}[E(t)]$) and performs a linear regression over a specified dB range (e.g., $-5\,\mathrm{dB}$ to $-25\,\mathrm{dB}$ or $-5\,\mathrm{dB}$ to $-35\,\mathrm{dB}$) to extract the reverberation time T_{20} or T_{30}. This integration technique simplifies the detection of multi-slope decays and facilitates more precise measurement of early decay rates, making it especially useful for diagnosing room-acoustic anomalies such as insufficient diffusion or rapid "initial" decays. Modern measurement systems commonly implement Schröder integration internally, providing a stable and repeatable estimate of parameters like T_{60} from a single measured impulse response.

Absorption Coefficients

Materials dissipate incident acoustic energy through mechanical or thermal losses, described by the absorption coefficient α:

$$\alpha = \frac{P_{\mathrm{abs}}}{P_{\mathrm{in}}}, \tag{2.3}$$

where: α = Absorption coefficient (dimensionless)
P_{abs} = Absorbed sound power (W)
P_{in} = Incident sound power (W)

Standards like DIN EN ISO 354-2003 prescribe measurement methods. Summing surface areas weighted by α gives the total absorption A:

$$A = \sum_{n=1}^{N} \alpha_n \cdot S_n. \tag{2.4}$$

Porous materials such as fiberglass efficiently absorb mid and high frequencies [10], while other treatments (e.g., resonant absorbers) can target lower frequencies.

Reflections, Diffusion, and Scattering

When sound waves encounter walls, ceilings, or other surfaces, they are reflected. Excessive reflections can blur transients, create flutter echoes, and distort spatial cues [6]. At low frequencies, repeated reflections between parallel boundaries form *room modes*, described in Section 2.3. At mid and high frequencies, reflection paths become more directional, and surfaces can be engineered to scatter or diffuse sound:

- *Diffusion* redistributes reflected energy uniformly in multiple directions. Quadratic residue diffusers (QRDs) spread sound over a broad bandwidth while preserving overall energy [10].
- *Scattering* refers to the redirection of sound waves by irregular surfaces. Coarse or uneven textures help break up reflections that would otherwise travel on parallel paths, reducing coloration and flutter echoes.

These techniques manage higher-frequency reflections without overdamping the room.

2.2 Fundamental DSP Concepts: Fourier Transform and Sampling

Digital Signal Processing (DSP) underpins the computational aspects of room acoustics analysis and filter design. Two core ideas are:

Fourier Transform

Fourier analysis reveals the frequency content of signals by transforming time-domain data into the frequency domain. The continuous-time Fourier transform of a function $f(x)$ is given by:

$$F(\omega) = \int_{-\infty}^{\infty} f(x)\, e^{-i\,\omega\,x}\, dx \tag{2.5}$$

and its inverse is:

$$f(x) = \frac{1}{2\,\pi} \int_{-\infty}^{\infty} F(\omega)\, e^{i\,\omega\,x}\, d\omega \tag{2.6}$$

Since real signals are often discretized, the Discrete Fourier Transform (DFT) is used:

$$X_k = \sum_{n=0}^{N-1} x_n\, e^{-i\,2\pi\frac{k}{N}n}, \quad x_n = \frac{1}{N} \sum_{k=0}^{N-1} X_k\, e^{i\,2\pi\frac{k}{N}n} \tag{2.7}$$

Operations like convolution in the time domain become simple multiplications in the frequency domain.

Sampling and Alias Avoidance

Digital Signal Processing (DSP) relies on converting analog signals into a stream of discrete-time samples. The *Nyquist–Shannon sampling theorem* dictates that the

sampling rate f_{sampling} must be at least twice the highest frequency content $\left(f_{\text{max}}\right)$ in the signal:

$$f_{\text{sampling}} \geq 2\, f_{\text{max}}. \tag{2.8}$$

If this condition is not met, *aliasing* occurs: high-frequency components fold back into lower-frequency bands, producing distortions in the recorded signal. To prevent aliasing, hardware or software-based *anti-aliasing filters* remove frequency components above the Nyquist limit ($f_{\text{sampling}}/2$) before sampling.

Practical Considerations in Audio Sampling:

- *Sampling Rate Choices.* Common rates include 44.1 kHz (CD audio) and 48 kHz (broadcast). Higher rates like 96 kHz can provide extra headroom for anti-aliasing filters or address ultrasonic content, though they increase data size.
- *Bit Depth and Quantization.* After sampling, signals are quantized to a finite number of bits (e.g., 16, 24). Higher bit depth decreases quantization noise and increases dynamic range, allowing more precise representation of low-level signals.
- *DAC Reconstruction.* On playback, Digital-to-Analog Converters (DACs) apply interpolation and smoothing filters to recreate the continuous waveform from the discrete samples, ideally matching the original analog signal within the system's resolution.

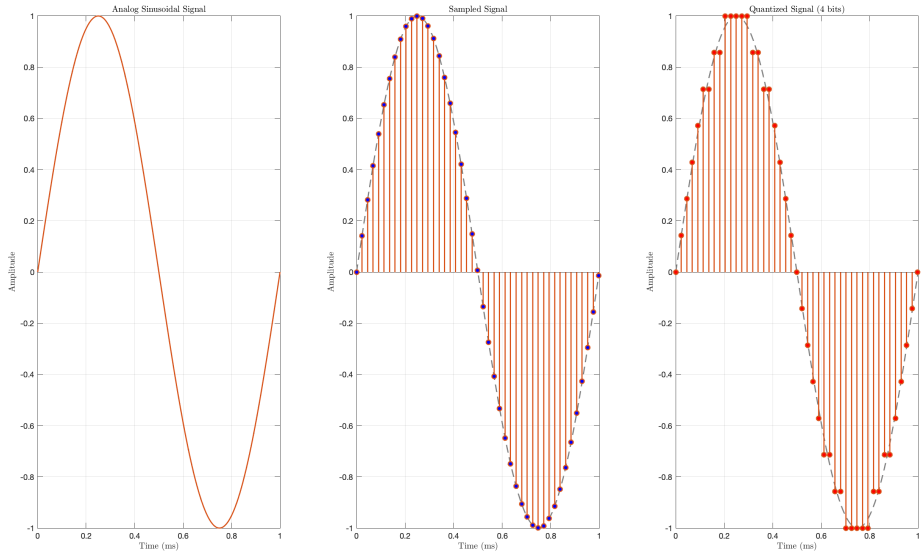

Figure 2.1: Analog sine wave (black dashed), sampled points (blue markers), and quantized samples (red markers).

Figure 2.1 illustrates an analog sinusoidal signal being sampled (blue markers) and subsequently quantized (red markers). In practice, maintaining a sampling rate com-

fortably above $2\,f_{\max}$ ensures minimal aliasing in the recorded or processed audio. Furthermore, steep anti-aliasing filters near the Nyquist frequency ($\frac{f_{\text{sampling}}}{2}$) must be carefully designed to avoid passband ripple or phase distortion while sufficiently attenuating higher frequencies.

2.3 Prediction of Low-Frequency Room Modes

At low frequencies, the wavelength of sound is comparable to typical room dimensions, giving rise to standing waves or *room modes*. These modes cause substantial peaks and dips in the frequency response, often leading to uneven bass.

2.3.1 Formation of Standing Waves

In a rectangular room of dimensions (L_x, L_y, L_z), standing waves form when reflected waves interfere constructively and destructively. The modal frequencies are given by:

$$f_{n_x,n_y,n_z} = \frac{c_0}{2}\sqrt{\left(\frac{n_x}{L_x}\right)^2 + \left(\frac{n_y}{L_y}\right)^2 + \left(\frac{n_z}{L_z}\right)^2}, \tag{2.9}$$

where: f_{n_x,n_y,n_z} = Frequencies of the room modes (Hz)
c_0 = Speed of sound ($\approx 343\,\mathrm{m\,s^{-1}}$)
n_x, n_y, n_z = Mode integers $\{0, 1, 2, \dots\}$
L_x, L_y, L_z = Room dimensions (m)

In practice, modes are enumerated up to the *Schröder frequency*:

$$f_S = 2000\sqrt{\frac{T_{60}}{V}}, \tag{2.10}$$

beyond which modes overlap too densely for purely modal analysis [1].

2.3.2 Modal Summation and Damping

Room responses can be computed by summing discrete mode contributions. In a so-called "shoebox" room, the frequency response $H(\omega)$ is:

$$H(\omega) = -\frac{4\pi c^2}{V}\sum_i \frac{\psi_i(r_s)\,\psi_i(r_r)}{(\omega^2 - \omega_i^2 - j\,\delta_i\,\omega_i)\,K_i}, \tag{2.11}$$

where: c = Speed of sound ($\mathrm{m\,s^{-1}}$)
V = Volume ($\mathrm{m^3}$)
ψ_i = Eigenfunction at source/receiver
δ_i = Damping factor (dimensionless)
K_i = Normalization constant

where δ_i accounts for boundary and air losses. Inverse FFT of $H(\omega)$ yields an impulse response for the predicted low-frequency behavior. In order to calculate

the Transfer function one needs to calculate the eigenfunctions for the sender and receiver postion $r = (x, y, z)$. This is done using the following formula:

$$\psi_{n_x,n_y,n_z}(r) = cos\left(\frac{n_x \pi x}{L_x}\right) cos\left(\frac{n_y \pi y}{L_y}\right) cos\left(\frac{n_z \pi z}{L_z}\right) \tag{2.12}$$

2.3.3 Non-Rectangular Rooms

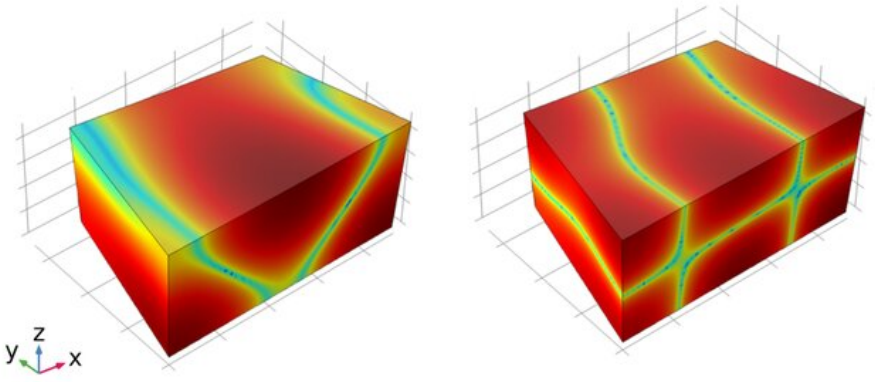

Figure 2.2: Representative room modes in a non-rectangular space [11].

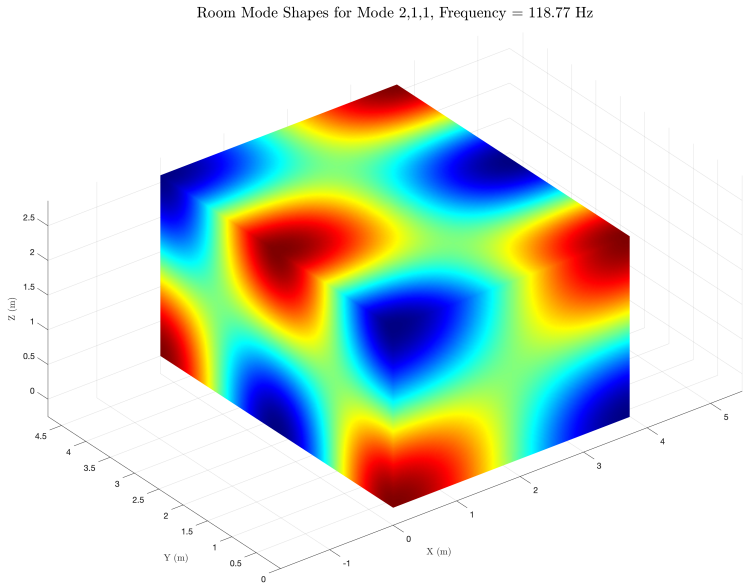

Figure 2.3: Room mode shape in a rectangular room, $\{2, 1, 1\}$, $f = 118.77\,\text{Hz}$.

While modal analysis in rectangular or "shoebox" spaces can be handled with closed-form equations (as illustrated in Figure 2.3), many real-world rooms deviate significantly from these idealized geometries. Walls may be angled or curved, ceilings may slope, and architectural features such as alcoves, columns, or irregular floor plans introduce additional complexity. Consequently, the underlying mode shapes become

more complicated than simple sinusoidal patterns, as shown in Figure 2.2 [11].

In rectangular rooms, each dimension leads to distinct standing-wave patterns that can be described by closed-form solutions, enabling faster calculations and simpler analyses for tasks such as loudspeaker placement or preliminary acoustic design. By contrast, non-rectangular rooms feature walls arranged at oblique angles or constructed from varying materials, which distribute resonances more unevenly and may occasionally reduce the severity of sharp peaks. However, the same geometric complexity can introduce intricate interference zones and localized "acoustic lens" effects, making it harder to predict standing waves through traditional formulas. Numerical simulations in these irregular spaces demand fine spatial discretization to resolve local sound-pressure variations accurately, increasing memory usage and processing time. Higher frequencies—where smaller wavelengths must be modeled—can exacerbate this computational burden, requiring meticulous mesh design and powerful hardware. Furthermore, while finite element or boundary element methods can yield exceptionally detailed insights into local pressure and velocity fields, they also require carefully specified boundary conditions. Real-world factors, such as partly absorptive surfaces, air losses, or small openings, further complicate the setup and validation of these numerical models. Despite these challenges, advanced modeling approaches offer a higher degree of accuracy and realism, particularly when dealing with unconventionally shaped spaces whose mode structures deviate substantially from the neat sinusoidal patterns of rectangular geometries.

2.4 Room Modes and Loudspeaker Positioning

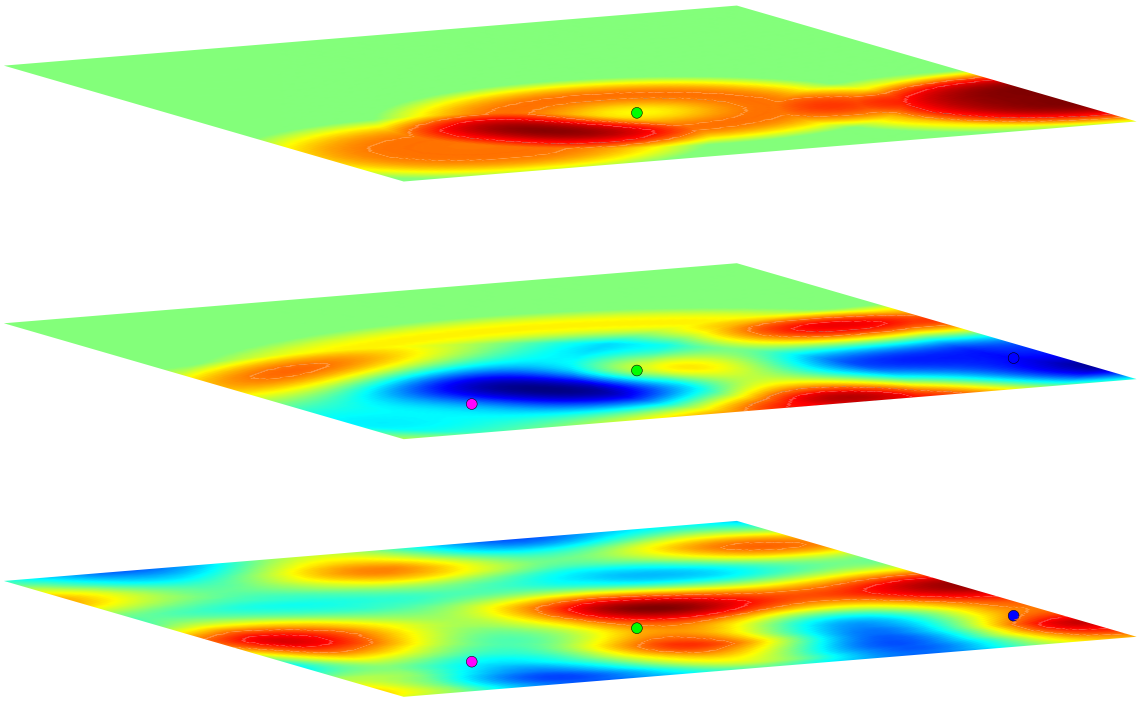

Figure 2.4: Wavefield simulation illustrating the influence of loudspeaker placement on room modes.

Figure 2.4 presents a series of simulated snapshots of the low-frequency wavefield inside the room. These visualizations derive from the modal summation script, which computes the pressure distribution by summing individual standing waves up to a specified frequency limit. The color gradients represent regions of varying sound pressure: red and yellow indicate higher pressure (antinodes), whereas green and blue reflect lower pressure (moving toward nodal lines).

Several key insights emerge:

1. Influence of Corners and Boundaries: As expected in small-room acoustics, corners and walls host regions of constructive interference (e.g., bright red spots). This occurs because each boundary reflection reinforces specific modal frequencies, often creating strong pressure maxima at or near corners. Consequently, a loudspeaker placed in a corner is more likely to excite room modes vigorously, causing pronounced peaks in the transfer function.

2. Nodal Lines and Interior Zones: Within the central area of the room, we observe lines and zones of lower pressure—blue and green bands indicating nodal or near-nodal locations. A loudspeaker positioned exactly on such a node would minimally excite that particular mode, explaining why certain speaker placements yield deep dips in measured or simulated responses. This pattern underscores the

importance of carefully balancing loudspeaker location to avoid unintentionally reinforcing or nullifying key frequencies.

3. Interplay of Multiple Modes: The snapshots often reveal overlapping lobes of pressure maxima and minima. In reality, multiple modes combine to form a complex interference pattern. Even a slight shift in loudspeaker position can move it from a low-pressure zone to a high-pressure zone for one mode, while simultaneously altering its coupling with other modes. This explains why different loudspeakers in the same room (as shown in later figures) may display significantly different frequency responses.

4. Comparison with Transfer Function Results: Earlier transfer-function plots (Figures 4.1 and 4.2) demonstrate that a speaker located in a high-pressure zone at a specific frequency experiences pronounced peaks, while one positioned near a pressure null shows a dip. Figure 2.4 visualizes these effects, linking them directly to the underlying spatial patterns of wave interference. By combining wavefield snapshots with transfer-function measurements, one gains a more intuitive grasp of why certain frequencies are amplified or attenuated at specific positions.

5. Practical Implications for Loudspeaker Placement: Placing a loudspeaker away from corners or strong antinodes generally reduces the severity of modal peaks. Conversely, positioning a subwoofer in a corner can be advantageous if one aims to excite all modes deliberately and then apply global equalization (e.g., via adaptive filters). However, corner placement can also exacerbate unwanted resonances if not carefully managed. These visualizations serve as a guide for systematically identifying "hot spots" and more neutral zones in the room.

Overall, Figure 2.4 demonstrates that spatial pressure distributions are highly frequency-dependent, and that loudspeaker placement profoundly influences how modes are excited. While the simulation provides an idealized snapshot, real rooms may introduce additional complexities—such as partial absorption, furnishing irregularities, and non-rectangular geometries—that further color the actual wavefield. Nonetheless, the patterns shown here form a foundation for understanding why certain positions yield boomy bass, while other placements exhibit relatively smooth responses.

2.5 Adaptive Filtering Using the LMS Algorithm

The most straightforward method for equalizing a loudspeaker's response in a room involves measuring the frequency response, calculating the inverse of the result, and then inverse transforming it back into the time domain to obtain a perfectly fitting FIR filter. This straightforward approach is not without its limitations. First, if there is a deep null existing in the original frequency response, the filter will result in a very large gain. It is important to note that the validity of this filter is limited to specific combination of sender and receiver position.
In order to sidestep this issue the *Least Mean Squares* (LMS) algorithm offers an iterative method to refine filter coefficients $w(n)$. At each sample n:

$$e(n) = d(n) - y(n), \quad w(n+1) = w(n) + \mu\, x(n)\, e(n), \tag{2.13}$$

where: $e(n)$ = Error signal
$d(n)$ = Desired or reference signal
$y(n)$ = Filter output
$x(n)$ = Input signal
μ = Step size (convergence parameter)

By adjusting μ, one can balance convergence speed against stability [12]. To achieve low-frequency correction, the desired signal $d(n)$ should be a time-aligned sweep that ideally produces a flat response in the room. Repeated iterations reduce modal peaks without over-boosting nulls.

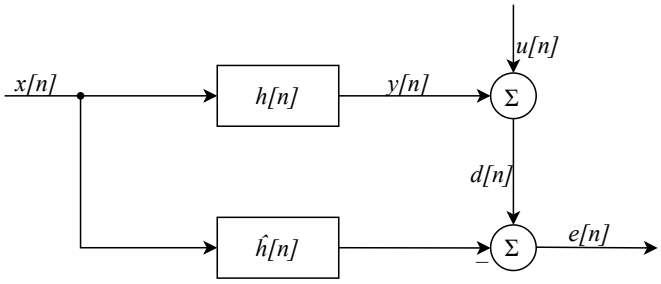

Figure 2.5: LMS filter signal path for room-acoustics correction [12].

The above graph shows the standard deployment of an LMS algorithm to identify an unknown linear time-invariant (LTI) system, in this case $h[n]$. If we assume now that this is a room response, the input signal $x[n]$ would be the measurement signal, either a swept sine or a noise signal, which is then played back on a loudspeaker, interacting with the room and being recorded by the microphone. This signal is then the output of the room, or the unknown LTI System or $y[n]$. In the process of recording and DAC/ADC conversion will introduce some added noise $u[n]$. This is then our desired signal $d[n]$.
The other strand of our signal chain is then taking the input signal $x[n]$, which is filtered by our adaptive filter $\hat{h}[n]$ and finally subtracted from the desired signal $d[n]$ resulting in the error signal $e[n]$. The LMS aims to update the filter coefficients of $\hat{h}[n+1]$ using the formulas 2.13 ti minimize the error signal.

If we modify the standard system identification scenario by changing the desired signal $d[n]$ to be the input signal $x[n]$, the iterative filter will attempt to minimize the difference between the room response and the perfectly flat sweep. To ensure the efficacy of this process, it is necessary to incorporate a delay into the input signal. This will allow the filter to modify signals in both the present and the future, while avoiding attempts to rectify events that have already occurred. This is known as causality, and the necessary delay depends on the LTI system, the filter length, and the sampling frequency.

2.6 Signal Alignment and Summation

Cross-Correlation to Find Delays

During measurements, each microphone and speaker can incur different path lengths and latency. A cross-correlation method locates the sample offset τ that maximizes:

$$\mathrm{corr}(\tau) = \sum_{n=1}^{N} x_1(n)\, x_2(n+\tau). \tag{2.14}$$

MATLAB routines then apply zero-padding or trimming to align signals precisely, ensuring coherent summation without artificial interference.

Summation and Multi-Channel Measurement

Once time-aligned, multi-channel data can be:

- Summed or averaged for frequency-response checks,
- Fed into an LMS routine for sequential or simultaneous speaker correction.

Correct alignment prevents destructive or constructive interference artifacts that might mask or exaggerate the true room response.

2.7 Other Room Correction Approaches and Multichannel Solutions

A variety of techniques exist to enhance sound quality and mitigate low-frequency issues in enclosed spaces, ranging from purely passive means to advanced digital signal processing. Traditional approaches often begin with modifying the physical properties of the room itself. By installing absorbers or resonators such as Helmholtz resonators or membrane-based traps—one can specifically target problematic low-frequency resonances [10].

Although effective, these devices typically require considerable volume to address bass frequencies, sometimes making them impractical in smaller or multi-purpose rooms. Diffusion and scattering elements can also be added to break up reflections at higher frequencies, improving clarity without over-damping the space. However, none of these passive methods can adapt dynamically to changing conditions, and

their success in controlling deep modes remains limited by size constraints and placement logistics.

Beyond passive treatment, active acoustic control systems use additional loudspeakers or actuators to counter unwanted sound fields in real time. One notable example, the double bass array, places extra subwoofers on a room's rear wall and drives them with delayed and inverted signals to cancel low-frequency standing waves [13]. This approach often succeeds in smoothing the bass response in a defined listening area, but it demands careful timing and calibration, and any misconfiguration can reintroduce problems elsewhere. Active absorption is another innovative concept, using compact electroacoustic absorbers that sense and counteract resonances [14]. Although these devices can offer high efficiency with a smaller footprint than bulky passive traps, they still add complexity and require ongoing sensor-actuator feedback. Also add considerable cost into a playback system.

Inverse filtering is a more direct DSP-based strategy wherein the measured room transfer function is inverted to flatten the overall response [15]. While straightforward in theory, this method risks instability if the inversion encounters deep nulls, and excessive boosts at certain frequencies can cause significant pre-ringing artifacts or amplify noise. Techniques like regularization partially alleviate these pitfalls, but single-point or single-channel inverse filtering often yields inconsistent performance across multiple listening positions. Recent approaches improve upon basic inverse filtering through clustered response analysis [16], wherein multiple room impulse responses from various locations are measured and grouped by similarity. Filters are then designed to address shared resonances across different seats, thereby improving the uniformity of the sound field while balancing trade-offs among multiple positions.

Robust mixed-phase multipoint correction further refines these ideas by optimizing both the amplitude and temporal aspects of the response over an extended listening area [17]. This technique limits pre-ringing artifacts and overly aggressive gain boosts by imposing constraints on filter phase response and magnitude, resulting in a more natural-sounding correction. Some systems also adopt multi-domain metrics, such as steady-state deviation and modal decay time, to evaluate correction effectiveness comprehensively [18]. The net result can be a perceptually convincing improvement in bass clarity and decay, albeit at the cost of increased design complexity.

Although each of these methods can be applied independently, recent research highlights the advantages of *multichannel* or *MIMO (multi-input multi-output)* solutions, which jointly optimize all loudspeakers and measure performance across multiple listening points. Rather than designing each speaker's filter in isolation, a MIMO scheme models how every speaker excites the room at each location. Correction filters are then derived to minimize errors simultaneously at multiple seats, distributing corrective action among different sources. In effect, lower seat-to-seat variation and more moderate filter gains become possible, reducing the likelihood of large boosts or unwanted pre-ringing. Commercial implementations (e.g., Dirac Live®)

illustrate how carefully calibrated multichannel corrections can smooth the overall low-frequency response while offering more consistent performance for several listening positions. However, MIMO strategies inevitably demand additional hardware (multiple DSP channels and carefully matched amplifiers), extensive measurement data, and higher computational overhead for filter design. They can also be "over-constrained" if too many points must be equalized perfectly, risking a compromise in any single seat's response.

In practice, no single method definitively resolves every acoustic issue. Passive modifications remain essential for broad-spectrum control and to avoid relying solely on high filter gains, active control can further suppress standing waves at the lowest frequencies, and advanced inverse filtering or multipoint clustering can refine local improvements while minimizing side effects. Ultimately, *multichannel MIMO approaches* represent a powerful extension of these techniques, enabling more uniform bass reproduction across multiple locations and mitigating many of the limitations inherent in single-channel or single-point correction systems. Yet as with any comprehensive approach, the gains in flexibility and coverage come at the price of greater design complexity, measurement rigor, and computational resources. In modern listening environments—ranging from home theaters to professional mixing studios—combining various strategies into a cohesive correction framework often yields the best overall balance of performance, adaptability, and practical feasibility.

3

Methods

This chapter describes the measurement, signal-processing, and adaptive filtering procedures used to mitigate low-frequency room modes in a small-room environment. A series of custom MATLAB scripts were developed to handle swept-sine playback, multi-microphone recording, time alignment, transfer-function computation, and iterative LMS filter design. The following sections detail how these scripts interconnect, the rationale behind each step, and how these methods were implemented in practice.

3.1 Overview of Measurement Framework

Measurements took place in a room whose dimensions are approximately:

- Length (L_x) $=$ 3.725 m,
- Width (L_y) $=$ 4.777 m,
- Height (L_z) $=$ 2.604 m.

Room modes in the 45–180 Hz region are especially prominent under these conditions due to the small room dimensions and low-frequency characteristics of sound. To explore multi-loudspeaker correction, three Genelec 8020B monitors (labeled S1, S2, S3) were deployed. Additionally, three microphones—two AKG C414 and one AKG C314—were strategically positioned at specific listening-area coordinates to capture the room's acoustic response from multiple perspectives. The choice of sound sources and microphones was driven by their availability. It is important to note that the microphones have an adaptable directivity pattern and are set to the omnidirectional setting. The size of the condenser capsule may pose a challenge when the measured frequency region of interest is higher. However, at these large wavelengths, the microphones do not produce a problematic influence on the sound field.

The experimental setup is shown in Figure 3.1, where the placement of loudspeakers and microphones within the room is depicted. This layout was chosen to ensure adequate coverage of the room's modal region while capturing both axial and oblique mode interactions.

Figure 3.1: Experimental setup showing the room layout and microphones used for low-frequency modal analysis.

The room was equipped with heavy curtains along the walls to minimize high-frequency reflections, as visible in the image. However, the curtains have minimal impact on the low-frequency range under consideration. The placement of microphones and loudspeakers was carefully measured. This configuration allowed for consistent and repeatable measurements across the experimental procedure.

Audio Equipment

An Antelope Audio Discrete 4 Pro audio interface handled digital-to-analog (`DA`) and analog-to-digital (`AD`) conversions. An RME QuadMic preamplifier boosted the microphone signals. All measurements operated at a 9600 Hz sampling rate to balance computational load and frequency resolution below 200 Hz. The measurement scripts stored all signals in FLAC format to preserve quality without excessive file size.

Coordinate Choices

The loudspeakers were placed as follows:

- **S1:** (1.836, 0.950, 0.996) m,
- **S2:** (0.886, 0.143, 0.100) m,
- **S3:** (3.525, 0.200, 1.280) m.

This arrangement deliberately excites various axial, tangential, and oblique modes to assess the impact of different loudspeaker placements on room responses.
The microphones (R1, R2, R3) were located near:

- **R1:** (1.900, 3.200, 1.030) m,
- **R2:** (1.571, 3.542, 0.777) m,
- **R3:** (2.135, 3.536, 1.309) m.

By clustering the microphones at slightly different heights and lateral offsets, the method captures seat-to-seat bass variation typical of small control or listening rooms.
Figure 3.2 shows an example configuration of the experimental setup with the loudspeakers and microphones placed according to the above coordinates. It is important to note that this setup is only one of several configurations tested during the study and does not necessarily represent the specific measurements presented later in the results. Other arrangements were explored to assess the variability in room responses under different conditions.

Figure 3.2: Example configuration of the loudspeakers used in the experimental setup. It is about the Genelec 8020 Loudspeakers. The Neumann Array and the subwoofers where not part of the tested setup. This arrangement represents one of the tested configurations and may not reflect the specific measurements reported later.

3.2 Measurement procedure and script description

3.2.1 Signal Generation and Data Acquisition

A core part of the procedure relied on generating linear swept-sine signals in the 45–180 Hz band. Each sweep was 12 s in duration, with 2 s of appended silence for separation. The MATLAB script then handled:

- **Signal Creation:** A function call (e.g., `chirp`) produced the sweeps at a defined start frequency $f_1 = 45\,\text{Hz}$ and end frequency $f_2 = 180\,\text{Hz}$, using a specific initial phase (e.g., $-90°$) to mitigate amplitude spikes.
- **Playback and Recording:** The custom routine `playAndRecordAudio` streamed the swept-sine file through each loudspeaker channel. Meanwhile, all microphone inputs were simultaneously recorded, preserving multi-channel data in real time.
- **FLAC Storage:** Each set of raw microphone signals was saved as a FLAC file with an appropriate naming scheme, ensuring lossless capture and easy re-loading for analysis.

In some measurements, only one loudspeaker was active; in others, multiple speakers were driven to study possible multi-subwoofer approaches.

3.2.2 Time Alignment of Signals

Due to different path lengths and minor latency differences, the raw signals contain misalignments in both speaker-to-mic and mic-to-mic data. The scripts addressed these lags with cross-correlation and zero-padding:

1. **Short Alignment Sweeps:** Each speaker played a brief chirp that was recorded by all microphones. This step was repeated for each loudspeaker, helping estimate relative speaker delays.
2. **Cross-Correlation (`calculateDelayFromRecordedAudio`):** For each channel pair, the code computes the lag τ at which the correlation peak occurs. That shift best aligns the signals.
3. **Zero-Padding (`applyDelays`):** Once τ is found, the script inserts zeros or truncates samples to align each track. Fine adjustments can be done with `fineTuneAlignmentZeroPadding`.
4. **Validation:** The function `calculateCoherence` and an `MSE` check (`calculateAlignmentMSE`) confirm that the signals are nearly sample-accurate in time.

This alignment ensures that subsequent computations, especially for iterative LMS filtering, capture coherent phase relationships rather than artificially shifting signals.

3.2.3 Computation of Room Transfer Functions

Once aligned, the script sums or averages each microphone's response as needed to form an effective transfer function from each loudspeaker. Typically:

- **FFT Analysis:** The function $\widehat{H}(f)$ is derived by comparing the recorded sweep's spectrum to the known input sweep.

- **Windowing or Deconvolution:** In some cases, the code uses short-time or direct deconvolution with the original chirp, yielding an impulse response that can be time-gated or truncated.
- **Storage and Plotting:** Magnitude responses are stored for quick lookups, while plots show main resonances that coincide with the room's predicted modes around 60–120 Hz.

This frequency-domain analysis identifies the peaks and dips that the LMS filter will later attempt to flatten.

3.2.4 Adaptive LMS Filter Design

For each loudspeaker channel S_i, the method constructs a desired signal $d(n)$ by shifting the original sweep to match measured delays. The recorded signal $x(n)$ is the room's actual output. The LMS algorithm updates the filter coefficients $w(n)$ to minimize:

$$e(n) \;=\; d(n) - y(n), \qquad w(n+1) \;=\; w(n) + \mu\, x(n)\, e(n).$$

Implementation details:

- **Filter Length:** $m = 4096$ samples was chosen, enough to cover up to 180 Hz with sufficient FFT resolution at 9600 Hz sampling.
- **Step Size (μ):** An empirical scaling factor of 24 ensured stable convergence over about 40 iterations.
- **Coefficient Constraints:** To avoid huge boosts at deep nulls, maximum amplitude limits could be enforced on the filter coefficients (though generally optional if the step size is chosen conservatively).

3.2.5 Sequential Multi-Channel Procedure

The loudspeakers were equalized one after another in a *cascaded* fashion. Although this is not a true MIMO (joint) optimization, it allows each secondary loudspeaker to "see" the correction already applied to the previous ones and thus captures a first-order share of the loudspeaker interaction.

1. **Step 1 – Design of the main filter $\hat{w}_1$.**

   (a) Loudspeaker 1 (**S1**) is driven with a swept-sine in the 45–180 Hz band while all three microphones record.
   (b) The three microphone tracks are time-aligned (cross-correlation, zero–padding) and vector-averaged to form one reference room-response signal $x_1[n]$.
   (c) The same sweep, shifted by $M/3$ samples, is used as the desired signal $d[n]$ so that the LMS produces a linear- or mixed-phase FIR whose main lobe sits roughly in the middle of the M-tap window.[1]
   (d) Running the LMS on $\langle d[n], x_1[n] \rangle$ yields $\hat{w}_1$, which is henceforth convolved with every S1 playback signal.

[1] A fractional delay of $M/3 \ldots M/2$ samples proved empirically to give a good compromise between causality margin and filter pre-ring.

2. **Step 2 – Design of the companion filter $\hat{w}_2$.**

   (a) With $\hat{w}_1$ now active in real time, Loudspeaker 2 (**S2**) is measured in the *presence* of the corrected S1 (both loudspeakers are playing).
   (b) The three microphone signals are again aligned and vector-averaged, producing $x_2[n]$.
   (c) The same delayed sweep acts as $d[n]$; an LMS run on $\langle d[n], x_2[n] \rangle$ delivers $\hat{w}_2$.

3. **Step 3 – Design of the tertiary filter $\hat{w}_3$.**
   The procedure of Step 2 is repeated for Loudspeaker 3 (**S3**), this time with $\hat{w}_1$ *and* $\hat{w}_2$ active, producing the final filter $\hat{w}_3$.

After all three stages, the system owns a set of FIR filters $\{\hat{w}_1, \hat{w}_2, \hat{w}_3\}$ whose primary objective is to linearise the frequency response of S1 but which inherently account for the first-order cross-coupling introduced by S2 and S3. A verification sweep played through the fully filtered array confirms whether any residual peaks remain or whether additional (joint) optimisation would be advantageous.

3.2.6 Music Playback and Verification

After the three FIR filters $\{\hat{w}_1, \hat{w}_2, \hat{w}_3\}$ were finalised, their practical benefit was assessed with programme material.

1. **Low–/High-band split.** A `crossoverFilter` routine separated each stereo track at $f_c \approx 180\,\text{Hz}$. The low-band signal was convolved with the respective $\hat{w}_i[n]$, while the high band remained unaltered so that mid-/high-frequency timbre was not affected.
2. **Subjective check.** Listeners compared unfiltered and filtered playback, focusing on *(i)* seat-to-seat level consistency, *(ii)* perceived "tightness" of kick drum and bass guitar transients, and *(iii)* reduction of lingering "boom" on sustained notes.
3. **Objective re-measurement.** A fresh swept-sine was played with all filters active. The three microphones were time-aligned and vector-averaged as before. Any remaining peaks (or new artefacts) were inspected against the pre-correction baselines.
4. **Optional spot checks.** Short broadband bursts and narrow-band noise were also recorded to ensure that the FIRs did not introduce audible pre-ring or excessive gain outside the design band.

Overall, the filtered playback exhibited a noticeably more uniform low-frequency presentation, while verification sweeps confirmed that no major new resonances or ringing artefacts had been introduced.

3.3 Overview of Experimental Setup and Signal Processing Pipeline

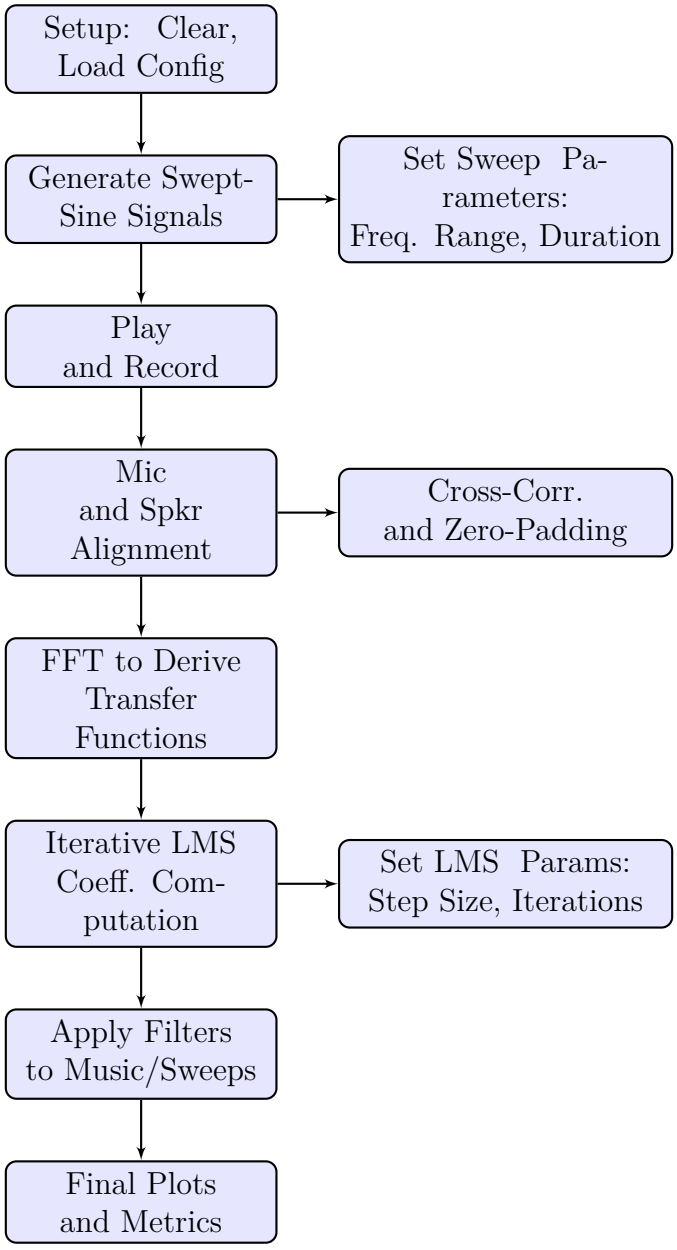

Figure 3.3: Expanded flowchart for measurement, alignment, and LMS filtering, referencing key scripts such as `playAndRecordAudio`, `calculateDelayFromRecordedAudio`, `fineTuneAlignmentZeroPadding`, and `dsp.LMSFilter`.

Setup: Clear, Load Config. The script initializes the environment by clearing variables and loading pre-defined configurations, such as sampling rates and device settings. This step ensures consistency and prepares the system for measurements.

Generate Swept-Sine Signals. Swept-sine signals are created within a specific frequency range (e.g., 45–180 Hz) to excite room modes. The MATLAB script `playAndRecordAudio` handles these signals, generating them based on parameters like duration, frequency range, and phase alignment. Side step `sweepParams` configures these settings to avoid artifacts or clipping.

Play and Record. The loudspeakers play the sweeps, and microphones simultaneously record the room responses. The script interfaces with audio hardware (e.g., RME devices) to handle real-time playback and capture. Multiple loudspeakers can be tested sequentially or simultaneously to observe interactions.

Mic and Spkr Alignment. Recorded signals often include delays due to path differences. The `calculateDelayFromRecordedAudio` script uses cross-correlation to estimate these delays. The `fineTuneAlignmentZeroPadding` function adjusts signal alignment by zero-padding or shifting samples.

FFT to Derive Transfer Functions. Once signals are aligned, the script performs FFT to convert recorded impulse responses into the frequency domain, producing transfer functions that reveal modal peaks and nulls. These results guide subsequent LMS filter design.

Iterative LMS Coeff. Computation. The LMS algorithm iteratively adjusts filter coefficients to minimize the error between a desired flat response and the measured transfer function. Parameters such as step size and iteration count are tuned via `dsp.LMSFilter` or equivalent custom routines. These settings are configured in `lmsParams`.

Apply Filters to Music/Sweeps. The computed filters are applied to test signals (music tracks or sweeps) to validate their effectiveness. The script convolving the filter with audio data ensures proper frequency-domain shaping.

Final Plots and Metrics. Results are visualized through plots of frequency responses, impulse responses, and error metrics. These plots confirm the filter's performance and highlight residual issues, such as persistent peaks or excessive gain.

3.4 Calibration Assumptions for Microphones and Loudspeakers

No loudspeaker- or microphone-equalisation file was applied in this work. Instead, we rely on three practical observations that hold well in the 45 Hz–180 Hz band of interest:

1. **Equipment Linearity.** Modern studio monitors and 1-inch condenser capsules exhibit less than ±2 dB deviation in the first audio octave. At 100 Hz the acoustic wavelength ($\lambda \approx 3.4\,\mathrm{m}$) is orders of magnitude larger than the diaphragm diameter ($\approx 25\,\mathrm{mm}$), so the capsule behaves as an acoustically small, pressure-mode sensor; diaphragm resonances lie far above the band of interest.
2. **Modal Dominance vs. Schröder Frequency.** Using Sabine's estimate[2] the room's Schröder frequency is roughly $f_S \approx 190\,\mathrm{Hz}$. We therefore limited all processing to $\leq 180\,\mathrm{Hz}$, i.e. the lower side of the transition band where discrete modes are still clearly identifiable. Above this point modal density increases quickly and the finite-length FIR filters began to introduce ringing artefacts during early pilot runs.
3. **Empirical LMS Tuning.** The 45 Hz to 180 Hz window was confirmed empirically as the range in which the adaptive filter converges reliably with the chosen settings (initial step size, μ; FIR length; 40 iterations; 12 s sweep length; −12 dBFS playback level and magnitude normalisation). Extending the upper bound forced either a smaller step size (slower convergence) or produced overshoot and pre-ringing in the final filter.

Additional Safeguards

- *Dynamic-range management:* All sweeps were rendered 12 dB below clip, producing SNR > 75 dB at the microphones while keeping loudspeaker excursion in its linear region.
- *Coefficient clipping:* The iterative LMS update included a soft limit of ±12 dB in the magnitude response to prevent exaggerated boosts at very deep nulls.

Although these measures cannot replace a full IEC calibration chain, they keep the combined loudspeaker–microphone transfer sufficiently flat—and, crucially, predictable—within the modal band. Future work could introduce reference microphones and loudspeaker ground-plane measurements to extend the usable frequency range or to quantify any residual system colouration.

3.5 Potential Error Sources and Mitigation

Despite the controlled laboratory setting, several factors can bias low-frequency measurements. Table 3.1 summarises the main risks and the counter-measures adopted in this study.

[2] $f_S \approx 2000\sqrt{T_{60}/V}$ [**?**]. With $L_x = 3.725\,\mathrm{m}$, $L_y = 4.777\,\mathrm{m}$, $L_z = 2.604\,\mathrm{m} \Rightarrow V \approx 46.4\,\mathrm{m}^3$ and a measured mid-band $T_{60} \approx 0.45\,\mathrm{s}$, we obtain $f_S \approx 190\,\mathrm{Hz}$.

Table 3.1: Major error mechanisms and mitigation strategies.

Source of error	**Mitigation**
Environmental noise	Measurements taken during off-hours; sweep level chosen for SNR > 50 dB.
Positioning tolerances	Laser-measured coordinates; three repeat sweeps averaged.
Room variability	Furniture kept fixed; total measurement time $<$ 30 min.
Equipment non-linearity	Processing band-limited to 45 Hz to 180 Hz; playback 6 dB headroom.
Numerical artefacts	Double-precision processing; scripts validated on synthetic test cases.

3.6 Summary of the Method Chain

Combining band-limited swept-sine measurements, sub-sample signal alignment, and iterative LMS equalisation, the procedure systematically reduces dominant room modes:

1. *Multi-position acquisition:* each loudspeaker is captured with three spatially separated microphones.
2. *Phase-coherent stacking:* cross-correlation aligns all traces to within one sample, after which a complex vector average represents the listening zone.
3. *Constrained LMS design:* a 4096-tap FIR is adapted over 40 iterations while soft-limiting excessive gain.
4. *Programme-material check:* filters are auditioned on music and re-validated with verification sweeps.

The next chapter presents the measured improvements—both spectral and temporal— and discusses the residual challenges that motivate a future move toward fully-joint MIMO optimisation.

4

Results

This chapter presents both simulation-based outcomes and real-world measurements, illustrating how the interplay between loudspeaker positioning, room modes, and adaptive filtering affects low-frequency performance. The MATLAB scripts referenced in earlier chapters were used extensively here—particularly those computing modal responses (`RoomSimMultiCh`) and those implementing LMS-based filter design—to compare predicted behaviors with actual measurement data and to assess the effectiveness of the multi-loudspeaker correction approach.

4.1 Simulation Insights

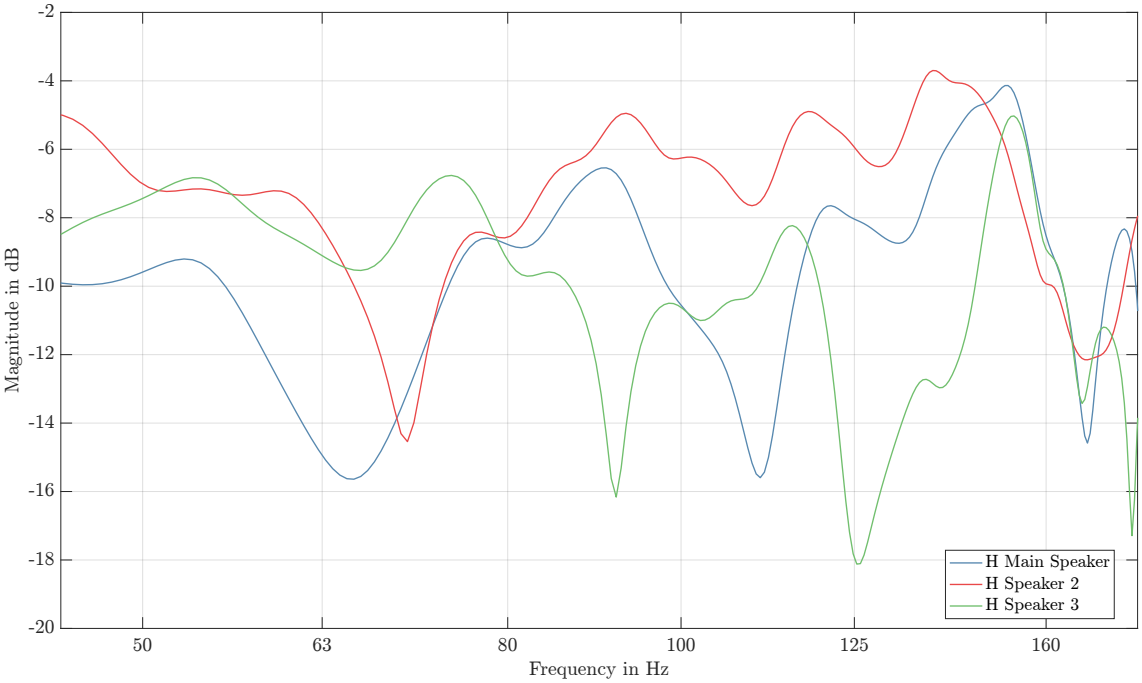

Figure 4.1: Simulated room transfer functions for the three loudspeakers.

Figure 4.1 illustrates three distinct simulated transfer functions, each representing a different loudspeaker position within the same room. These results stem from summing the contributions of low-frequency room modes up to a specified `n_max` in the MATLAB scripts, effectively capturing how each loudspeaker interacts with the standing waves.

One immediately noticeable feature is the strong modal activity concentrated near 70 Hz. Speaker 3 exhibits a pronounced peak at this frequency, while Speaker 2 shows a corresponding dip. This contrast underscores how even slight shifts in loudspeaker location can move a source from an antinodal region (where a mode is strongly excited, causing a peak) to a nodal or near-nodal position (yielding a marked dip). Such behavior is characteristic of small-room acoustics, where the wavelength of low-frequency sound is comparable to the room dimensions, causing drastic differences in pressure build-up from one spot to another.

Another noteworthy aspect involves the variation in overall level between approximately 50 Hz and 80 Hz. Speaker 1's curve maintains a more moderate response in this range, suggesting that its placement may be less coincident with primary mode axes or corners. Conversely, Speaker 3's heightened peak at around 70 Hz points to stronger modal coupling at that location. This interplay of peaks and dips provides valuable insight into how each mode interacts with the speaker's position, reinforcing the importance of speaker placement in low-frequency optimization.

Beyond 80 Hz, the responses begin to show multiple fluctuations—smaller peaks and dips that reflect higher-order modes. These higher-frequency modes are more densely packed and may overlap, giving rise to complex interference patterns. Although each loudspeaker experiences these modes differently, one can observe that all three plots share certain frequency regions (e.g., near 125 Hz and around 160 Hz) where modal contributions cause significant dips or humps.

The amplitude variations in each curve also highlight the sensitivity of a room's pressure distribution at low frequencies. Small changes in boundary conditions—such as the presence of furniture, wall characteristics, or even minor repositioning—can further alter these transfer functions. In practical scenarios, if Speaker 2 and Speaker 3 were physically swapped, one could expect the red and green curves to effectively exchange their features, reaffirming how source location dictates the excitation of specific modes.

Finally, these simulations emphasize the importance of considering multiple loudspeakers in a single room. While one speaker may produce a peak at a problematic frequency, another might present a dip, potentially offering opportunities for complementary placement or multi-source equalization strategies. In home theater or studio applications, this knowledge can guide both physical placement (to avoid severe nulls or boomy resonances) and electronic correction methods such as adaptive filtering.

Overall, the figure underscores the significant impact of speaker positioning on low-frequency room responses. In real-world measurements, one would likely find similar variation among the different loudspeakers, validating these simulation results and encouraging strategies like subwoofer crawling, multi-subwoofer solutions, or careful placement to minimize the most prominent peaks and dips.

4.2 Measurements in the Actual Room

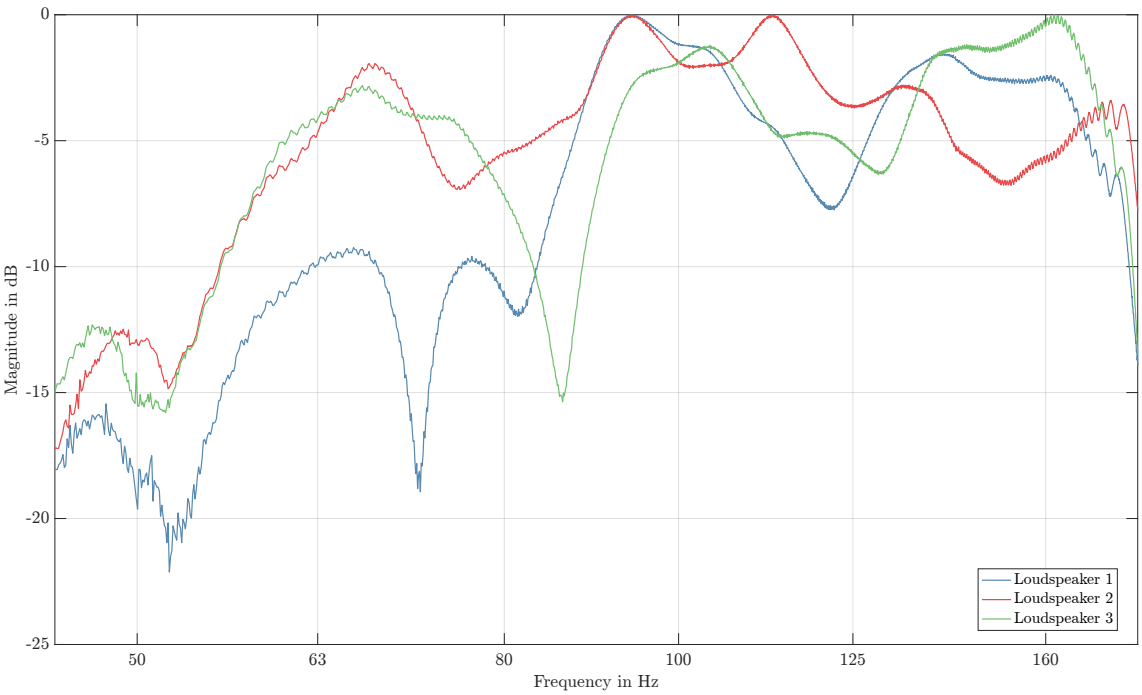

Figure 4.2: Measured room transfer functions for the three loudspeakers.

Figure 4.2 shows the measured transfer functions for three different loudspeakers placed in the actual room under investigation. Each curve was obtained through the multi-microphone setup and swept-sine procedure detailed in Chapter 3, capturing how real boundary conditions, speaker roll-off characteristics, and partial damping affect the low-frequency region.

A notable discrepancy from the simulations is observed below roughly 60–63 Hz. All three curves exhibit a more pronounced drop-off in level than predicted by the mode-based simulations. This roll-off largely reflects the inherent low-frequency limitations of the physical loudspeakers, along with added boundary absorption and small leakage paths that were not modeled precisely in the simulated environment. In practice, most consumer and even studio-grade loudspeakers begin to lose output capability in the very low-frequency region, an effect that the simulations did not fully replicate.

Despite this overall reduction in energy below 63 Hz, the familiar modal "signature" still appears between approximately 50 Hz and 120 Hz. Each loudspeaker shows a series of peaks and dips that align reasonably well with the mode distribution indicated by the simulations. However, the measured amplitudes of these peaks and notches tend to be lower in magnitude than predicted. This suggests that real-world factors like wall construction materials, furnishings, and additional damping elements (e.g., carpeting, acoustic panels) have partially mitigated some of the strongest modes.

Comparing each speaker:

- **Loudspeaker 1 (blue curve)** exhibits a deep dip around 55–60 Hz and again near 80 Hz, suggesting placement near or at a pressure node for certain modes. Above ∼120 Hz, the response flattens out slightly before rolling off near 160 Hz.

- **Loudspeaker 2 (green curve)** follows a less steep slope below 50 Hz than the other two, but still shows a drop-off that is steeper than in the simulations. Between 60 Hz and 110 Hz, it has a notable dip around 70–75 Hz, akin to what was observed in the simulations, though not as pronounced.

- **Loudspeaker 3 (red curve)** displays a comparatively smoother descent below 60 Hz but then reveals stronger peaks around 80 Hz and 110 Hz. The larger peak near 80 Hz was also seen in the simulation, though it appears at a slightly different frequency here, reflecting the real room's boundary conditions.

These variations among speakers illustrate how each position in the room excites different modes with distinct strengths. Although the dominant modes roughly match those anticipated by the simulation, the real room introduces additional complexities—non-uniform wall absorption, minor phase mismatches, and slight deviations in loudspeaker placement—that lead to smaller or shifted peaks and dips. The general shape of the measured responses, however, confirms the primary modal issues lie in the 50–120 Hz band, exactly where the simulation predicted trouble spots.

In essence, the measured data validate the core conclusions drawn from the simulations: that strong modal resonances dominate the low-frequency region, and that small positional changes can significantly alter the transfer function. Nevertheless, the measured responses also highlight real-world damping and speaker limitations not fully captured by idealized modeling. Such comparisons underscore the importance of both simulation and measurement when implementing room acoustic treatments and low-frequency correction filters. The next chapters leverage these measurements to refine the adaptive filtering strategies and evaluate their effectiveness in mitigating the most problematic room resonances.

4.3 Evaluating Adaptive Filtering

4.3.1 Individual Speaker Corrections

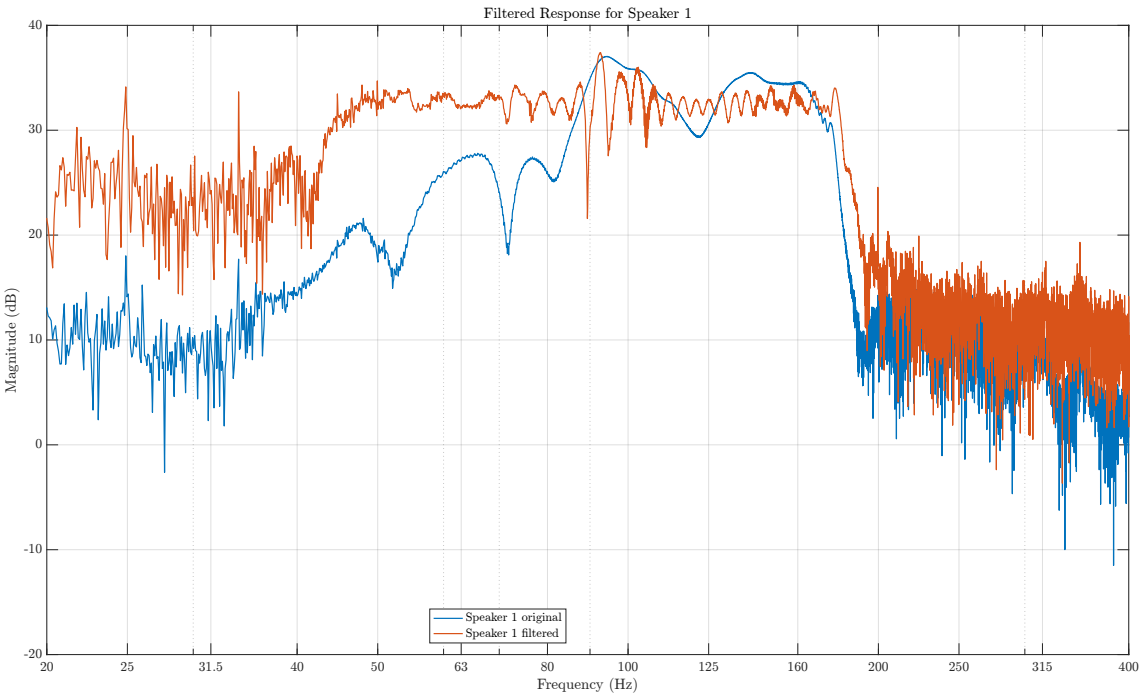

Figure 4.3: Original vs. LMS-filtered transfer function of Loudspeaker 1.

Figure 4.3 illustrates the unfiltered (orange curve) and LMS-filtered (blue curve) responses of Loudspeaker 1 over a range from approximately 20 Hz to 400 Hz. One can observe significant modal activity between 50 Hz and 120 Hz in the unfiltered response; strong peaks and dips in this region correlate with the standing waves previously identified in both simulations and earlier measurements. By applying the adaptive LMS routine, the filter coefficients converge iteratively toward a flatter magnitude response

In particular, several peaks between about 60 Hz and 100 Hz are noticeably reduced, bringing the response closer to a reference target (often conceptualized as a "flat" response). The dip around 80 Hz is similarly smoothed, demonstrating the LMS algorithm's ability to adjust phase and amplitude in a way that mitigates destructive interference. While there is an overall improvement, minor artifacts persist, reflecting the practical limitations of finite filter lengths, potential loudspeaker nonlinearities, and the complexity of multi-mode interference. Nevertheless, the measured difference—especially in the critical low-frequency range—shows that adaptive filtering can substantially reduce some of the most problematic modal peaks.

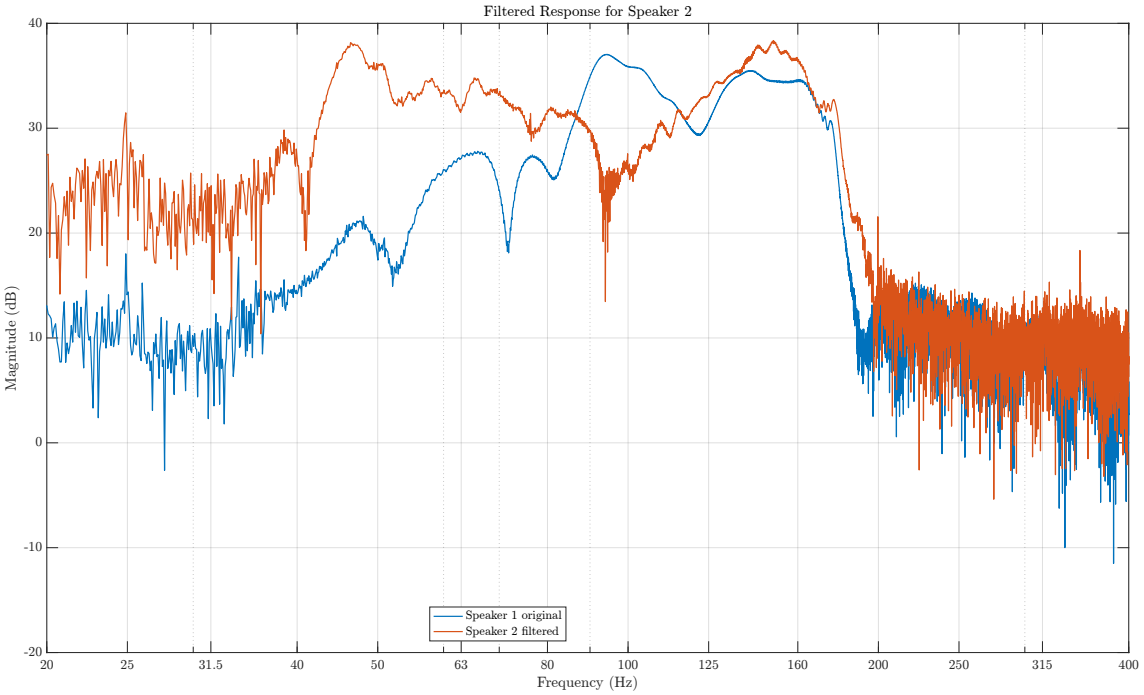

Figure 4.4: Original vs. LMS-filtered transfer function of Loudspeaker 2.

Figure 4.4 shows moderate improvement for Loudspeaker 2. The original response (in orange) reveals prominent peaks around 60 Hz and 100 Hz, alongside a deeper null near 70–75 Hz. These features are consistent with corner-excitation effects observed in earlier simulations. After LMS filtering (blue curve), the most significant peak around 60 Hz is reduced, and the dip near 70 Hz is partially filled in.

However, compared to Loudspeaker 1, the corner placement here introduces more complex room interactions. Certain strong modes remain, likely due to the limited capability of the filter to address deep nulls without risking instability or overcompensation. Additionally, corner placement tends to energize multiple modes simultaneously, making it harder for the LMS algorithm to flatten the entire range evenly. Still, the reduction of the most pronounced resonances suggests a clear benefit from adaptive equalization.

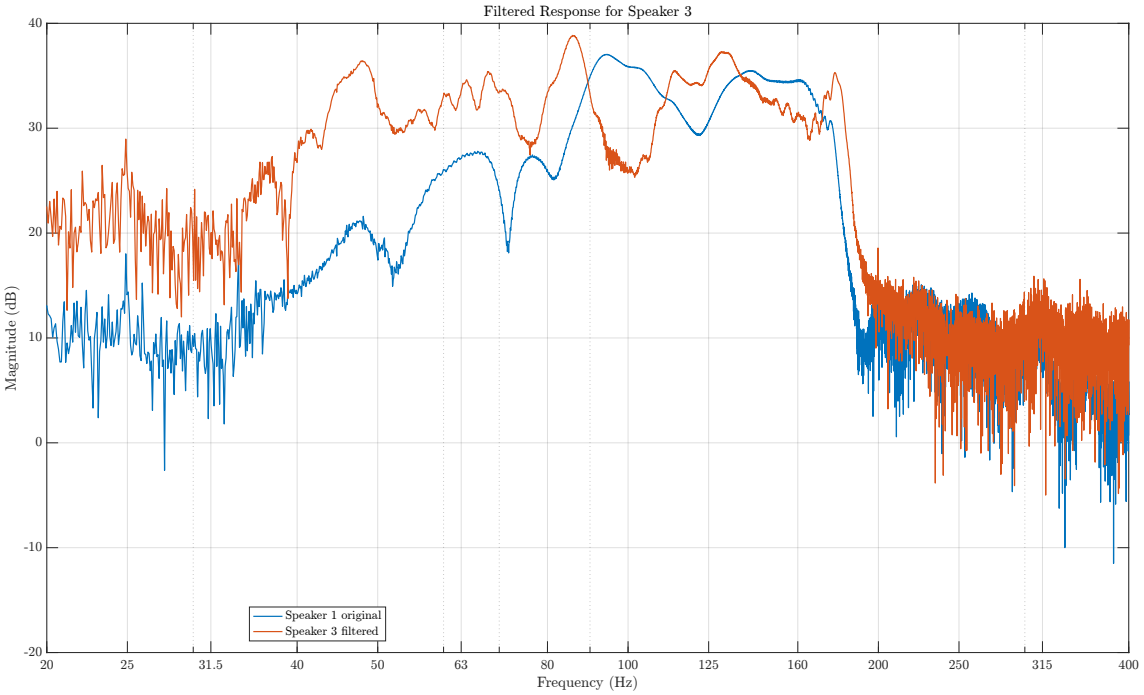

Figure 4.5: Original vs. LMS-filtered transfer function of Loudspeaker 3.

Figure 4.5 displays the original (orange) and filtered (blue) transfer functions for Loudspeaker 3, which had previously shown a pronounced peak around 70–80 Hz in both simulations and measurements. Here, the LMS filter manages to reduce that peak significantly, producing a more even response from roughly 60 Hz to 120 Hz. Smaller peaks remain near 90–100 Hz, highlighting that while adaptive filtering can diminish dominant modes, it may not entirely eliminate every resonance—especially those overlapping in frequency or interacting through complex node/antinode patterns.

In summary, all three figures demonstrate a clear trend: applying the adaptive LMS algorithm moves the transfer functions closer to a flatter profile in the low-frequency domain. The degree of improvement varies based on each loudspeaker's placement relative to strong modal axes and boundary conditions. Moreover, the extent of correction depends on the filter's length, step size (μ), and the loudspeaker's inherent linearity. Yet even with these constraints, the filtered curves show a meaningful reduction of peaks and, in some cases, a partial filling-in of nulls—results that align well with the aims of low-frequency correction in small rooms.

Overall, these measurements confirm that adaptive filtering can be a powerful method to alleviate modal buildup, improve the listening experience, and maintain more uniform bass response across different loudspeaker placements. Subsequent sections explore multi-channel or combined-speaker scenarios, where these individual filters interact and further challenges may arise in balancing multiple sources.

4.3.2 Combined Multichannel Response

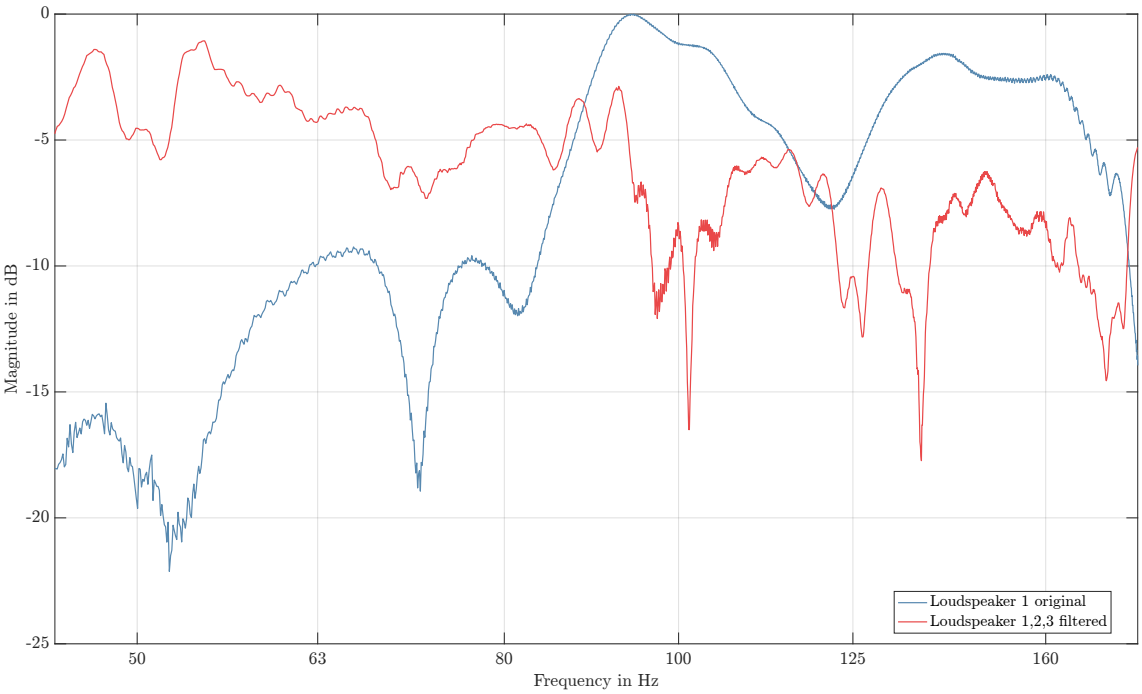

Figure 4.6: Overall transfer function after individual LMS filters are applied to each speaker.

Figure 4.6 presents the resulting transfer function when all three loudspeakers are driven simultaneously, each with its own independently derived LMS filter. In principle, applying a separate filter per speaker aims to correct the primary peaks and dips observed for that particular source in isolation. However, when multiple loudspeakers operate together in the same room, several additional factors come into play:

1. Overlapping Modes and Phase Interactions: Even if each loudspeaker is individually equalized, their combined output may still trigger local constructive or destructive interference at certain frequencies. If two speakers are positioned such that one is exciting a strong mode near $\sim$ 70 Hz while another partially cancels it, the net effect could be a residual dip or peak that was not addressed by the single-speaker filters. The same phenomenon is visible in Figure 4.6, where the red (filtered) trace still exhibits fluctuations around 60–80 Hz and near 100 Hz.

2. Summation in the Modal Region: Low-frequency modes dominate the 50–120 Hz range in this room, as shown in the previous sections. When three loudspeakers emit overlapping frequencies simultaneously, each speaker's corrected response interacts with that of the others, sometimes reinforcing specific modes in corners or along walls. Hence, a slight attenuation from one loudspeaker's filter could be offset by another speaker's energy in the same band, leading to a residual peak or dip in the combined response.

3. Potential for Global Multichannel Correction: Because the LMS filters here were optimized per speaker individually, the algorithm did not account for multi-speaker interplay when determining filter coefficients. A "global" approach—where all channels are corrected in unison—may further smooth the overall room response. Multi-channel LMS or other advanced approaches (e.g., multiple-input multiple-output, or MIMO equalization) could potentially yield more uniform coverage across listening positions, at the cost of increased computational complexity.

4. Practical Significance: In home-theater or studio configurations, multiple loudspeakers often reproduce overlapping frequency content. The curves in Figure 4.6 demonstrate that while individual adaptive filters mitigate key modal peaks for each speaker, the final collective response can still show newly formed or remaining resonances. Users might find an overall improvement in subjective bass clarity and reduced "boominess," yet there remains room for further optimization.

Despite these challenges, the figure confirms that individual LMS filters, applied simultaneously, can reduce some of the most severe modal peaks each speaker would otherwise generate on its own. The result is a partially improved low-frequency response compared to a scenario without any adaptive correction. Moving forward, a dedicated multi-channel optimization strategy—factoring in each speaker's position and contribution to shared modal regions—could provide a more consistently flattened combined response.

4.4 Extended Analysis of Filtered vs. Unfiltered Responses

In pursuit of a comprehensive understanding of how adaptive filtering impacts low-frequency room behavior, this section examines the normalized unfiltered and filtered signals from multiple perspectives:

1. Changes in the frequency response of a single loudspeaker
2. Spatial variability among three different loudspeakers
3. Evaluations of decay time (T_{20}) in selected low-frequency bands using a Schroeder-based integration approach

By combining these viewpoints, we can assess the filter's efficacy in mitigating individual loudspeaker resonances, enhancing consistency across multiple speakers, and influencing modal decay characteristics.

4.4.1 Spatial Variability in the Frequency Domain

The objective is not merely to linearise one position or one loudspeaker, but to ensure that *every listener position* experiences a similar spectral balance when *all three loudspeakers* operate together. Spatial variability was therefore quantified as the *standard deviation* σ of the magnitude responses measured at the three microphone positions while *S1*, *S2* and *S3* reproduced the same swept-sine *simultaneously*. In other words, we fix the loudspeaker set and ask: *How different is the bass at Seat A, B and C?*

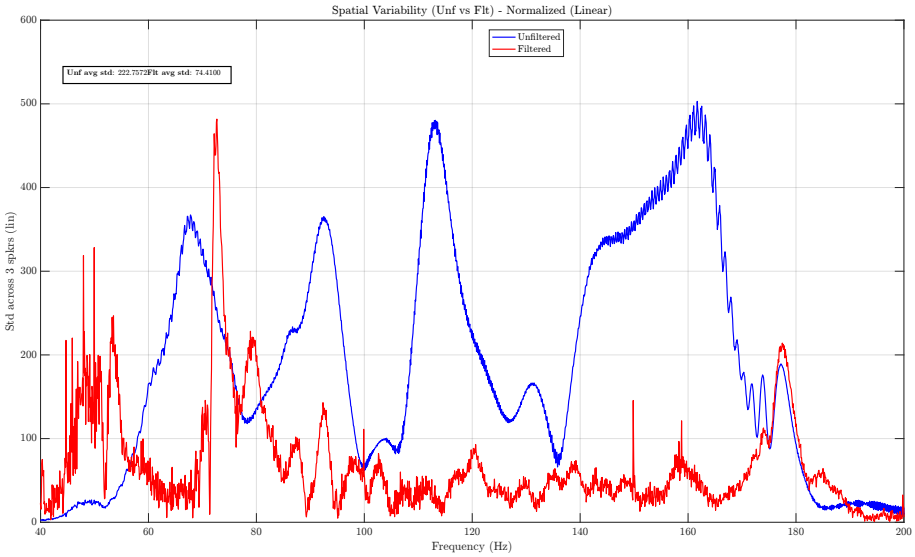

Figure 4.7: Spatial standard deviation (linear magnitude) across the three receiver positions with all loudspeakers active, before and after filtering.

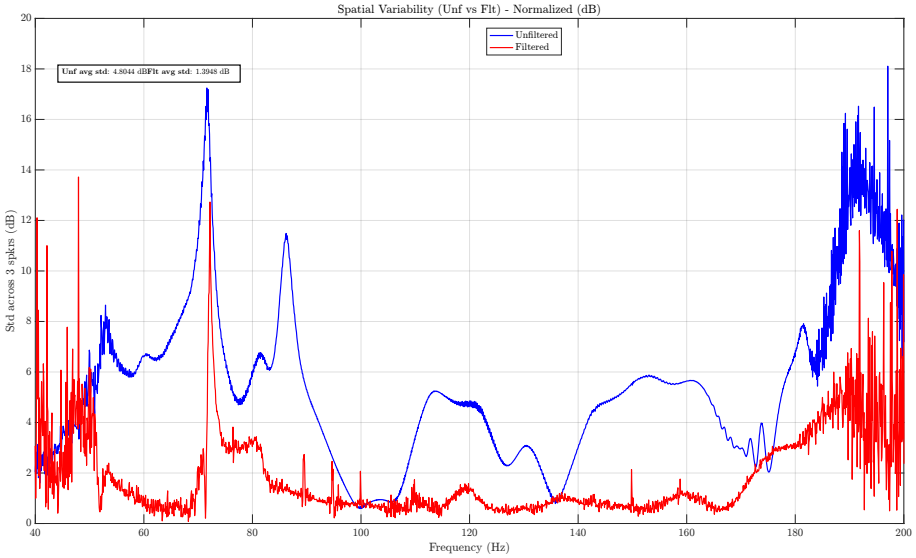

Figure 4.8: Same data as Fig. 4.7 expressed in dB.

Across the 45 Hz to 180 Hz band the adaptive filters tighten the seat-to-seat spread substantially:

- **Linear scale:** σ drops from $\sim$222.8 to 74.4 (–67%).
- **Log scale:** σ_{dB} falls from 4.8 dB to 1.4 dB (–3.4 dB).

Hence the dominant peaks are not only attenuated on average; they are attenuated *more uniformly* across the listening zone, shrinking the difference between the "best" and "worst" seat.

4.4.1.0.1 Residual peak at $\approx$70 Hz. The filtered curve still exhibits a narrow spike in σ at $\approx$70 Hz. This aligns with the $(0, 2, 0)$ axial mode:

$$f_{0,2,0} = \frac{c_0}{2}\frac{2}{L_y} = \frac{343}{2}\frac{2}{4.777} \approx 71.8\,\mathrm{Hz}.$$

The mode has nodal planes at $y = L_y/4 \approx 1.19$ m and $y = 3L_y/4 \approx 3.58$ m. Two microphones (R2, R3) lie within $\approx$4 cm of the upper node ($y = 3.54$ m to 3.58 m), whereas R1 ($y = 3.20$ m) is closer to an antinode. Even a modest residual excitation therefore produces large seat-to-seat differences, yielding the local maximum in σ.
A secondary factor is the loudspeaker geometry: S1 and S3 are asymmetrical with respect to these nodal planes, so their combined radiation pattern after filtering may couple more strongly to this particular mode than before. In short, the $\approx$72 Hz peak is a geometric artefact of microphone placement relative to the modal field, accentuated by unavoidable residual energy in that mode.
These results confirm that adaptive filtering delivers a markedly more uniform bass response, while also highlighting the persistent influence of specific room modes that warrant further multi-channel optimisation.

4.4.2 Decay Time Estimates via Schröder Integration

Here, the script `schroederDecay_60to80Hz.m` automates the steps of the Schröder Integration 2.2 for targeted frequency bands: it applies a bandpass filter (e.g., 60–80 Hz), squares the resulting waveform, and performs a reverse cumulative sum to produce a clean decay profile in decibels. The slope of that integrated decay in a chosen dB interval (e.g., −5 to −25 dB) then yields T_{20}, signifying the time needed for the energy to drop by 20 dB. This approach often reveals subtleties such as multi-slope decays or partial overlaps between neighboring resonances, which can be missed when relying on raw noise-decay measurements.
While steady-state frequency-response analyses reveal which resonances are most prominent, another critical aspect is how quickly the room's energy decays after an excitation—particularly in the low-frequency modal region. Extended decay results in lingering "boomy" bass, making transients less clear and potentially overshadowing other frequency content. To quantify these effects, the script applies bandpass filters around three critical modal bands (60–80 Hz, 83–100 Hz, and 105–125 Hz) and then computes T_{20} decay times using Schröder integration.[1]

Band 60–80 Hz

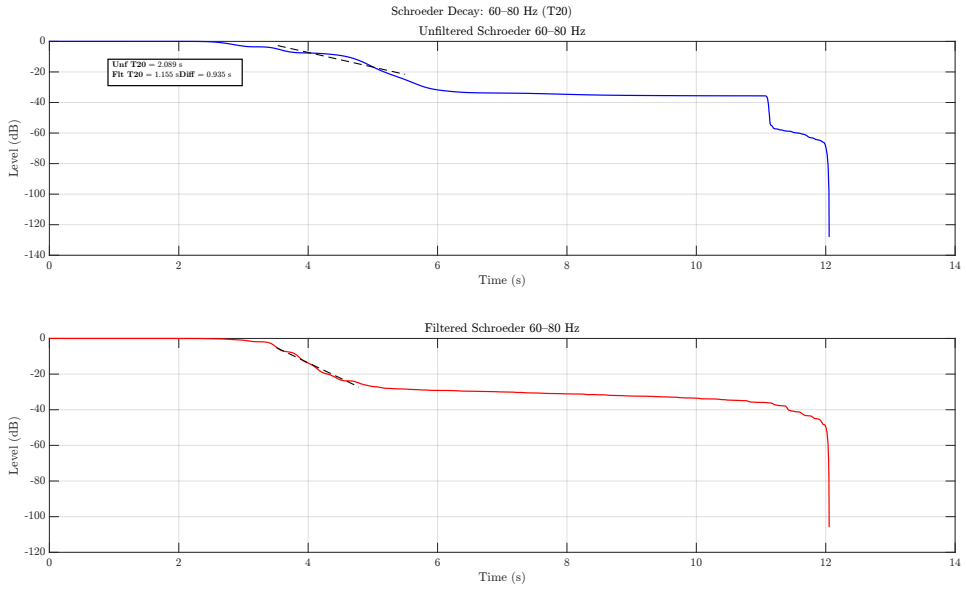

Figure 4.9: Decay time (T_{20}) for the 60–80 Hz band, unfiltered vs. filtered, using Schröder integration.

Figure 4.9 focuses on the 60–80 Hz range, which includes a strong resonance near 70 Hz. In the unfiltered case, the room requires approximately 2.089 s to decay by 20 dB, underscoring a pronounced low-frequency ringing. After applying the LMS-based filter, T_{20} drops to about 1.155 s, indicating a significant improvement of roughly 0.935 s. This reduction usually translates to tighter bass reproduction, as the lingering energy in this sub-bass band is dissipated faster. Musically, one would

[1] See *Room Acoustics* by Heinrich Kuttruff for additional details on the Schröder method.

expect cleaner transients and less masking of higher-frequency detail by extended low-frequency reverberation.

Band 83–100 Hz

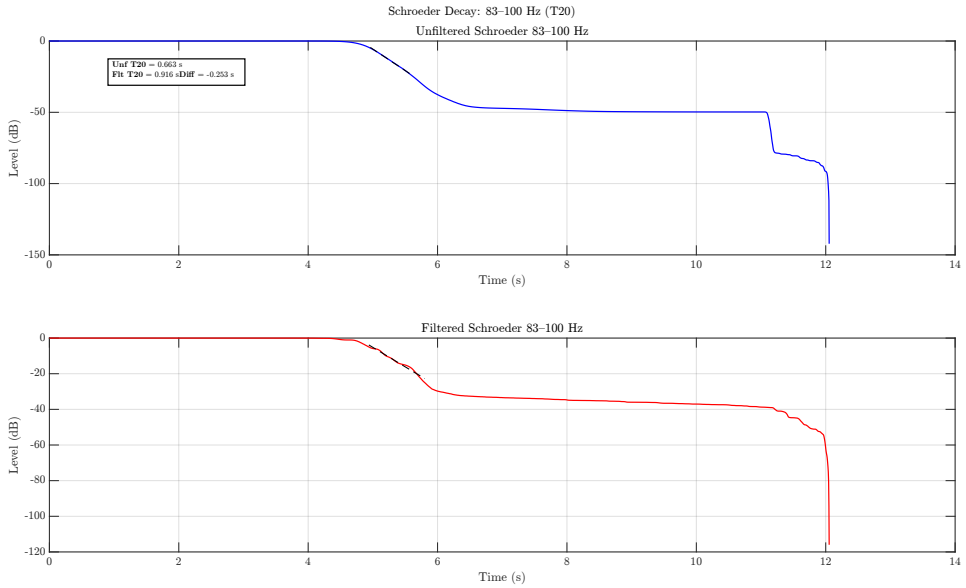

Figure 4.10: Decay time (T_{20}) for the 83–100 Hz band, unfiltered vs. filtered, via Schröder integration.

In the 83–100 Hz band (Figure 4.10), the unfiltered decay time (0.663 s) is surprisingly shorter than the filtered decay (0.916 s), resulting in a -0.253 s difference. This anomaly suggests the filter inadvertently boosts or redistributes energy in this frequency range instead of uniformly damping it. Such behavior can arise when strong modes overlap, and reducing one peak shifts energy into an adjacent band, or when the filter's phase adjustments cause partial constructive interference in nearby frequencies. These results underscore the interdependence of frequency bands: improvement in one area can sometimes come at the cost of increased decay in another.

Band 105–125 Hz

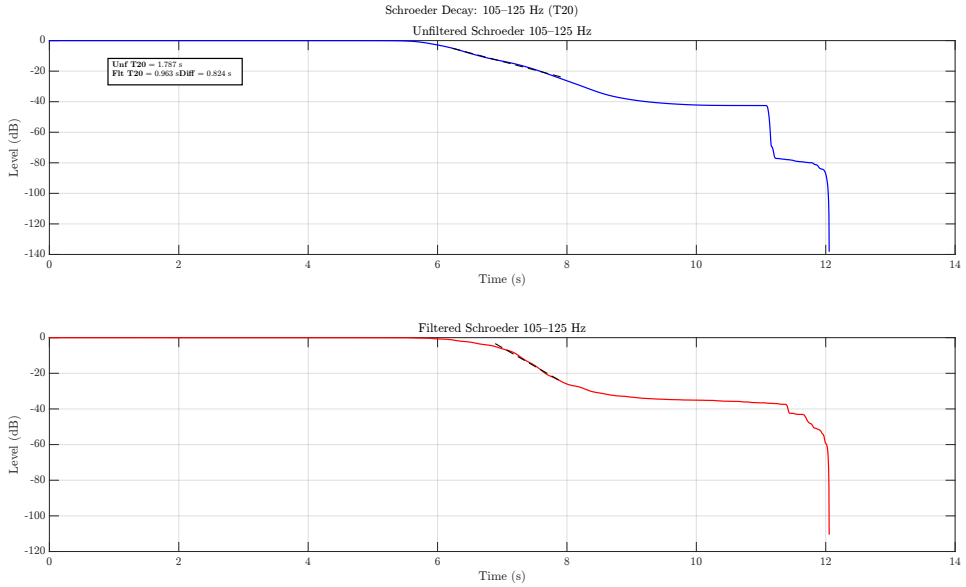

Figure 4.11: Decay time (T_{20}) for the 105–125 Hz band, unfiltered vs. filtered, via Schröder integration.

Conversely, Figure 4.11 shows a more typical outcome. Before filtering, the decay time in the 105–125 Hz range stands at about 1.787 s. After applying the adaptive filter, T_{20} is reduced to around 0.963 s, an improvement of 0.824 s. This reduction aligns with the expectation that controlling a prominent mode often yields a faster decay, making musical bass notes less prone to "hang" and improving overall clarity.

In summary, the Schroeder-based T_{20} estimates confirm that the adaptive filter successfully shortens decay times in the most problematic low-frequency regions. Nevertheless, the increase in decay time within the 83–100 Hz band reveals the complexity of balancing multiple resonances, especially when they partially overlap. The net effect is still a more controlled acoustic environment, but a multi-band or more advanced approach might be required to avoid inadvertently heightening decay in adjacent frequency ranges.

4.5 Discussion of Results and Implementation Challenges

The measurements and analyses confirm that the adaptive LMS filtering approach can significantly reduce problematic modal peaks for individual loudspeakers, leading to a smoother overall room response—particularly below 125 Hz. As observed in frequency-response plots, decay-time (T_{20}) measurements, and spatial variability metrics, the iterative filtering process yields audibly cleaner bass with shorter modal decay times. Nonetheless, certain frequency bands, such as 83–100 Hz, illustrate how overlapping modes and sequential speaker corrections may inadvertently redistribute

energy, occasionally increasing decay at adjacent frequencies. These findings highlight the complexity of multi-speaker interactions and underline the need for more holistic or fully "global" methods if absolute uniformity is desired.

Overall, the LMS routine demonstrates several encouraging outcomes:

- *Single-Speaker Gains:*
  - *Flatter Frequency Response:* Strong resonances (e.g., those affecting Speaker 1) are substantially reduced post-filtering.
  - *Decay Times:*
    - **60–80 Hz:** Decrease of about 0.935 s in T_{20}, indicating reduced lingering energy near 70 Hz.
    - **83–100 Hz:** Increase of roughly 0.253 s in T_{20}, illustrating that correction in one range can shift energy into another.
    - **105–125 Hz:** Reduction of around 0.824 s in T_{20}, pointing to improved sub-bass control.
- *Multi-Speaker Spatial Variability:*
  - On a linear scale, standard deviation among the three loudspeakers dropped from 222.7572 to 74.4100.
  - In dB, variability decreased from 4.8044 dB to 1.3948 dB, signifying closer alignment in the 45–180 Hz region.

These improvements confirm the system's ability to lower modal peaks, reduce seat-to-seat variation, and shorten the most critical low-frequency decay times. Even so, achieving a fully flat bass response in small and medium-sized rooms remains challenging, given the inherent cross-coupling between modes and speakers.

Beyond these intrinsic room-acoustic factors, real-world implementations face additional hurdles. Repeated measurements revealed that subtle shifts in temperature, furniture placement, and boundary conditions can disrupt the algorithm's iterative updates. The LMS framework assumes a relatively stable environment, so time-varying conditions sometimes led to partial convergence or minor anomalies in the final response. Further complicating matters, most loudspeakers exhibit measurable nonlinearity at very low frequencies or when operating near their physical limits. These nonlinearities skew the measured transfer functions in ways that a linear adaptive filter—like basic LMS—cannot fully capture. Pre-calibration or higher-order modeling may be necessary to maintain consistent performance under demanding conditions.

Interaction among loudspeakers themselves is another key issue, as sequentially filtering each speaker can shift or amplify energy in shared modal ranges. Diminishing one speaker's peak might inadvertently emphasize another's dip or alter decay in an adjacent band. Fully joint or "global" adaptation schemes—where all channels are optimized simultaneously—could more effectively manage these multi-speaker trade-offs.

Finally, practical tuning of the convergence parameter (μ) proved critical. Using too small a step size slowed adaptation, risking incomplete compensation of strong modes, whereas too large a step size caused overshoot and occasional instability. Balancing these extremes demanded extensive empirical tests, adding to setup complexity. Despite these complications, the LMS algorithm demonstrated a remarkable ability to adapt and deliver meaningful low-frequency corrections. Its relative simplicity—compared to computationally heavier methods or machine-learning approaches—makes it attractive for small-scale or home-audio applications, provided one accepts the sequential-tuning trade-offs and carefully manages parameters.

In summary, the results highlight both the effectiveness and the real-world challenges of LMS-based adaptive filtering. Substantial gains in single- and multi-speaker low-frequency performance can be achieved, but future work in areas such as global multi-channel optimization, improved loudspeaker calibration, and robust step-size selection may further enhance consistency, reduce artifacts, and deliver a more uniform listening experience across diverse room conditions.

4.6 Comparison with State-of-the-Art Techniques

To place these findings in context, it is helpful to compare the LMS-based approach with other existing strategies for low-frequency room mode correction:

- **Finite Element Methods (FEM):** Numerical FEM simulations can produce highly accurate predictions of room modes by solving the wave equation with fine spatial resolution. While excellent for design-phase analysis or research, FEM tends to be too computationally demanding. In contrast, the roo, simulation method using modal superposition allows rectangular rooms to be reliable predicted with less resources and gives accurate results for further filter designs.
- **Frequency-Specific Equalization:** Traditional single-channel equalization designs fixed filters for known modal frequencies, typically determined through static measurements. Although effective when the environment remains unchanged, such methods lack adaptive capabilities. The LMS algorithm's advantage lies in dynamically refining filter coefficients as conditions shift—an asset in rooms where small fluctuations could otherwise degrade fixed-filter solutions.
- **Global Optimization Algorithms:** Recent work on multi-speaker optimization uses machine learning, genetic algorithms, or advanced optimization frameworks that search for global solutions to reduce seat-to-seat variance and mode excitation. These methods can outperform simpler adaptive filters but often demand large data sets, high computational power, or extended calibration steps. LMS offers a more direct path for real-time or near-real-time correction, albeit with potential compromises in multi-channel scenarios.
- **Active Dereverberation (ADR):** Active dereverberation methods aim to reduce reflections and reverberant buildup by generating signals that destructively interfere with the room's decaying sound field. While such techniques can be effective at specific locations or frequencies, they often struggle to main-

tain performance over a broader listening area, as precise timing and phase relationships must be sustained. By contrast, an LMS-based correction iteratively addresses the entire frequency spectrum across multiple measurement points, typically yielding more uniform improvements in reverberation control for a wider audience rather than a single "sweet spot."

Strengths of the LMS Algorithm:

- Can adapt in real time to incremental changes in the environment.
- Straightforward to implement on standard audio DSP platforms.
- Reduces the need for detailed prior knowledge of room geometry or speaker characteristics.

Weaknesses of the LMS Algorithm:

- Less effective at coordinating multiple speakers simultaneously, unless extended to a full multi-channel adaptation scheme.
- Susceptible to measurement noise, minor system nonlinearities, and drift in room or equipment parameters.
- Requires careful tuning of convergence parameters (μ) and filter lengths to avert instability or incomplete correction.

Overall, the LMS method represents a robust option for small- to medium-sized rooms where persistent low-frequency resonances hamper audio fidelity. The improvements observed in amplitude response, decay times, and spatial consistency across loudspeakers validate LMS as a viable tool for low-frequency correction, especially when computational resources, measurement equipment, or time are limited. While more advanced multi-channel or global optimization strategies may yield superior uniformity, the results here highlight how a relatively accessible and flexible DSP technique can offer significant acoustic benefits in real listening environments.

5

Conclusion

Small- to medium-sized rooms commonly suffer from strong low-frequency standing waves and modal resonances, causing significant peaks, dips, and ringing that impair tonal balance and clarity. This thesis addressed these issues through an adaptive filtering strategy focused on the most troublesome low-frequency region and implemented via custom MATLAB scripts. By measuring swept-sine responses and iteratively refining filter coefficients with a Least Mean Squares (LMS) algorithm, the system counteracts room-induced phase and amplitude distortions.

5.1 Key Findings

The study delivered the following key insights:

- **Targeted Low-Frequency Correction.** Limiting the algorithm to the 45–180 Hz range effectively balanced computational efficiency and filter performance. This focus allowed the filtering process to directly target the modal region where room-induced anomalies are most prominent, as demonstrated through simulation and real-world measurements. It also allowed for a lower sampling frequency which enables the blocks to be smaller and have the same impact of latency of the system
- **Substantial Gains for Individual Loudspeakers.** Measurements confirmed that iterative error-minimization via the LMS algorithm effectively flattened peaks and dips in single-speaker responses. For most loudspeakers, standard deviation metrics showed a reduction of modal peaks, leading to a smoother and more consistent low-frequency performance.
- **Challenges in Multi-Loudspeaker Scenarios.** While single-speaker corrections proved highly effective, simultaneous multi-speaker operation highlighted residual complexities due to overlapping modes and phase interactions. These challenges resulted in occasional spectral inconsistencies and emphasized the need for more holistic optimization strategies.
- **Practical Workflow Integration.** The proposed methodology seamlessly combined multi-microphone measurements, cross-correlation for time alignment, and adaptive LMS filtering, providing a replicable and modular approach for addressing room-mode challenges in practical settings.

5.2 Contributions

This work provides a systematic framework for low-frequency correction in small rooms, characterized by:

1. **Measurement and Alignment.** Utilizing multi-microphone swept-sine recordings and cross-correlation alignment to ensure accurate timing and coherence between sources.
2. **Iterative Adaptive Filtering.** Applying the LMS algorithm to iteratively refine filter coefficients, focusing on mitigating dominant modes while avoiding instability or excessive computational complexity.
3. **Multi-Loudspeaker Exploration.** Demonstrating the strengths and limitations of sequential filtering in multi-loudspeaker setups, highlighting the interplay between local speaker corrections and global room acoustics.
4. **Modal Analysis Integration.** Incorporating modal predictions via simulation scripts (`RoomSimMultiCh`) to validate measured responses and guide filter design.

These contributions bridge the gap between theoretical acoustics and real-world implementation, offering practical tools for tackling low-frequency challenges in home studios, small venues, and similar environments.

5.3 Recommendations for Future Research

Based on the findings, the following directions are suggested for further exploration:

1. **Holistic Multi-Channel Optimization.** Future work could implement real-time joint optimization strategies for multiple loudspeakers, addressing cumulative phase interactions and ensuring a cohesive sound field across seating positions.
2. **Subwoofer Integration and Hybrid Filtering.** Integrating subwoofers with adaptive crossover designs could extend the system's range and offload deep-bass reproduction from smaller loudspeakers, resulting in cleaner overall performance.
3. **Learning-Based Methods.** Machine learning approaches, such as neural networks or reinforcement learning, could enable more robust adaptive filtering, particularly under non-linear speaker responses and dynamic room conditions.
4. **Enhanced Modal Pre-Analysis.** Pre-computing modal distributions using room geometry and material properties could further streamline filter initialization and placement strategies, reducing setup time while improving effectiveness.
5. **Dynamic Environmental Adaptation.** Incorporating real-time feedback systems to detect and adapt to changing acoustic environments (e.g., moving furniture or temperature shifts) could further enhance the practicality of the system in real-world applications.

5.4 Final Remarks

This thesis demonstrates the potential of adaptive LMS filtering to substantially reduce low-frequency anomalies in single-speaker setups and highlights the additional complexities in multi-speaker scenarios. By combining modal analysis, multi-microphone measurements, and iterative correction techniques, the proposed framework offers a replicable methodology for tackling small-room acoustic challenges. However, the findings also underscore the importance of future innovations in joint optimization, advanced hardware integration, and adaptive algorithms to achieve more consistent and uniform bass reproduction in diverse listening environments.

The methodology and results presented here contribute valuable insights to the field of room acoustics and provide a foundation for future research aimed at further refining adaptive correction techniques for low-frequency modal issues.

Bibliography

[1] Heinrich Kuttruff. *Room acoustics*. CRC Press, Boca Raton London New York, sixth edition edition, 2016.

[2] Todd Welti and Allan Devantier. Low-Frequency Optimization Using Multiple Subwoofers. *Journal of The Audio Engineering Society*, 54:347–364, 2006.

[3] Angelo Farina. Simultaneous Measurement of Impulse Response and Distortion With a Swept-Sine Technique. November 2000.

[4] Bernard Widrow and Samuel D. Stearns. *Adaptive signal processing*. Prentice-Hall signal processing series. Prentice-Hall, Englewood Cliffs, N.J, 1985.

[5] M. R. Schroeder. Frequency-Correlation Functions of Frequency Responses in Rooms. *The Journal of the Acoustical Society of America*, 34(12):1819–1823, December 1962.

[6] F. Alton Everest. *The master handbook of acoustics*. McGraw-Hill, New York, 4th ed edition, 2001.

[7] Wallace Clement Sabine. *Collected Papers on Acoustics*. Harvard University Press, Cambridge, MA, 1922.

[8] Marshall Long, editor. *Architectural acoustics*. Applications of modern acoustics. Elsevier/Academic Press, Amsterdam Boston, 2011.

[9] M. R. Schroeder. New Method of Measuring Reverberation Time. *The Journal of the Acoustical Society of America*, 37(3):409–412, March 1965.

[10] Trevor J Cox and Peter D'Antonio. *Acoustic Absorbers and Diffusers: Theory, Design and Application*. CRC Press, 2009.

[11] Thach Pham Vu, Etienne Rivet, and Hervé Lissek. *Low Frequency Sound Field Reconstruction in Non-rectangular Room: A Numerical Validation*. May 2018.

[12] Simon S. Haykin. *Adaptive filter theory*. Always learning. Pearson, Upper Saddle River, NJ, 5. ed edition, 2014.

[13] Adrian Celestinos and Sofus Nielsen. Low frequency sound field enhancement system for rectangular rooms using multiple low frequency loudspeakers. *Audio Engineering Society - 120th Convention Spring Preprints 2006*, 3:1352–1370, January 2006.

[14] Herve Lissek, Romain Boulandt, and Anne-Sophie Moreau. Practical Active and Semi-Active Strategies for the Control of Room Acoustics in the Low-Frequency Range. In *Internoise 2010 noise and sustainabitlity*, Lisbon, Portugal, 2010.

[15] O. Kirkeby, P.A. Nelson, H. Hamada, and F. Orduna-Bustamante. Fast deconvolution of multichannel systems using regularization. *IEEE Transactions on Speech and Audio Processing*, 6(2):189–194, March 1998.

[16] Sunil Bharitkar and Chris Kyriakakis. System And Method For Automatic Room Acoustic Correction In Multi-Channel Audio Environments. *The Journal of the Acoustical Society of America*, 129:1668, March 2011.

[17] Lars-Johan Brannmark and Mikael Sternad. *Controlling the impulse responses and the spatial variability in digital loudspeaker-room correction.* November 2015.

[18] P Thakkar. MULTI-DOMAIN ANALYSES OF ACOUSTIC ROOM CORRECTION SYSTEMS. In *Reproduced Sound 2022*, Bristol, February 2023. Institute of Acoustics.

DEPARTMENT OF SOME SUBJECT OR TECHNOLOGY
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY