Design of analog audio filters for individual speaker models in a distributed system

William Saranpää
wi4034sa-s@student.lu.se
Simon Alveteg
si2781al-s@student.lu.se

Axis Communications AB

Supervisor: Iman Ghotbi and Alex Lam

Examiner: Pietro Andreani

June 4, 2025

Printed in Sweden
Tryckeriet i E-huset, Lund

Abstract

This thesis investigates the performance and feasibility of passive and active analog frequency compensation filters for audio signal processing in a distributed speaker system. One passive and one active circuit topology were investigated more thoroughly and then the active filter was realized on a PCB and evaluated with electrical, thermal and acoustic measurements.

The work also resulted in two methods of designing frequency compensating or equalizing filters based on active or passive filters respectively. The design was partly automated with the use of optimizing algorithms that integrated with a LTSpice model of different circuits. The active filters were based on a constant-Q graphic equalizer topology, and performed better than the passive bridged-T based topologies in relation to flattening the speaker's frequency response. The passive filters had less power dissipation and produced less noise than their active counterparts, at the expense of lower selectivity as well as a lower output volume due to the added insertion loss. The design of the active filters was also quicker and gave more reliable results.

The filters were evaluated on two different speakers with different frequency responses in order to ensure the generality of the method. The passive filters could reduce the average deviation from 3.52 dB to 1.68 dB for speaker 1 and from 2.64 dB to 1.58 dB for speaker 2, while the active filter with 16 filter sections could reduce it to 1.57 dB for speaker 1 and 0.82 dB for speaker 2. Further reduction is most likely possible, as manual adjustment of just five filter sections brought the average deviation down from 1.75 dB to 1.63 dB for speaker 1.

Analog filters are a viable option for applications where the power budget, cost or limitations in signal availability or routing pose challenges. The method created simplifies the design process of equalizing filters and reduces the engineering time needed.

Acknowledgements

We would like to thank our supervisors, Alex Lam and Iman Ghotbi, for their valuable insight and guidance during this master thesis. We also want to thank Felix Brandel, Christoffer Axebring and everyone else at NB Audio for their help and encouragement.

Our sincere thanks go to Axis Communications for providing access to their facilities and for the opportunity to contribute to such an engaging and rewarding project. We also want to thank all the friends we made along the way.

Popular Science Summary

Sannolikheten att det finns en högtalare nära dig är hög. De finns där för att spela musik i en affär, för att meddela att brandlarmet har gått i byggnaden du befinner dig i eller för att berätta att tåget du ska åka med är inställt. Alla högtalare låter på lite olika sätt, de har sina egna skavanker som gör att en viss frekvens kanske låter lite mer än andra. I detta projekt har vi funderat på detta och undersökt hur man kan kompensera olika högtalares egenheter med analoga filter för att få ett platt frekvenssvar, det vill säga att alla frekvenser låter lika högt.

Det är nu för tiden vanligt att högtalare frekvenskompenseras digitalt med något datorchip. Detta kan bli problematiskt för vissa sorters högtalarsystem där det finns olika typer av begränsningar som exempelvis strömförbrukning eller kostnad. Därför undersöker detta arbete en lite äldre teknik, nämligen att frekvenskompensera med analoga filter.

Analoga filter kan delas upp i passiva eller aktiva filter. Passiva filter har oftast bara komponenter som motstånd, kondensatorer, och spolar medan aktiva filter även kan innehålla förstärkare. Det visar sig att båda sorterna har för- och nackdelar, men om man vill platta till frekvenssvaret och hålla ljudvolymen lika hög, är aktiva filter det bättre valet.

Detta arbete har kommit fram till att det finns situationer där de äldre analoga filtrena är ett bra val för att filtrera signaler till högtalare. De är ofta billigare, det går fortare att utveckla dem samt att de kan vara mer energieffektiva. En annan fördel är att varje högtalare kan byggas på ett enklare sätt, man behöver inte programmera varje högtalare i rummet för att allt ska låta bra. Digitala filter presterar bättre i frekvenskompensering jämfört med analoga filter, men ibland behövs inte all prestanda utan bara tillräcklig. Därför bör den lite äldre tekniken inte förbises i alla situationer, utan ingå i utvecklingsfasen då det, beroende på situationen, kan vara det bättre alternativet.

Abbreviations

BPF	Bandpass Filter
dB	Decibels
HIFI	High Fidelity
SPL	Sound Pressure Level
dBSPL	Sound Pressure Level in decibels
DSP	Digital Signal Processor
ESR	Equivalent Series Resistance
ESL	Equivalent Series inductance
MFB	Multiple Feedback Filter
PCB	Printed Circuit Board
SNR	Signal to Noise Ratio
VCVS	Voltage Controlled Voltage Source filter

Table of Contents

List of Figures

List of Tables

Chapter 1
Introduction

In audio engineering and electrical engineering in general, a trend has been to move from the analog domain into the digital due to digital circuits providing more flexibility and processing power. Digital signal processors, or DSP, are programmable and small, but can be limited by processing power and power consumption, amongst other aspects. There are applications where more traditional analog circuitry is used for audio signal processing, either due to personal preference or due to digital circuits being unfeasible in the design. Analog circuits can have several benefits, including but not limited to, smaller size, smaller power consumption and reduced engineering effort.

Filters are an integral part of electronic systems and serve several functions. They can mitigate noise coming from or going into a device, they are central parts in signal processing and can be used to combat non-linearities in devices, reject interference as well as many more applications. Analog filters have a rich history and have been used for most of the 20th century in a vast variety of systems. Therefore, they are interesting to revisit and uncover overlooked solutions in order to see if some of their benefits might outweigh the drawbacks for certain applications in audio signal processing.

In this project we will investigate the use of passive and active analog audio filters in speaker systems in order to mitigate speaker non-linearities. Different types of analog filters will be investigated and evaluated with respect to their performance and feasibility. This will occur on two different speakers in order to ensure that the method is general enough to be applicable for many different speakers and projects. The work will result in some guidelines and tools for designing analog audio filters for different speakers as well as a filter prototype that will be manufactured and evaluated.

The thesis aims to raise awareness of analog filters and their design for audio signal processing as well as investigate whether they could be viable for certain applications where the digital alternative might not be practical. This is investigated in the context of distributed speaker systems where size and power consumption is limited.

Chapter 2

Background

2.1 Acoustics

Sound is the brain's perception of periodic variations in atmospheric pressure, known as sound-pressure waves. These variations over time are often represented graphically as waveforms. These pressure waves can be measured in pascal, but it is more common to use the unit dBSPL, which is the sound pressure level in decibels [8]. The sound pressure in dBSPL can be expressed as

$$L_p = 20 \log_{10} \left(\frac{p}{p_0} \right) \tag{2.1}$$

where p is the sound pressure in Pa and p_0 is equal to 20 μPa, which is approximately the smallest sound pressure a human ear can perceive [8]. Hence, the smallest sound pressure that humans can hear is 0 dBSPL and the pain threshold is around 120 dBSPL.

The range of human hearing is from around 20 to 20 000 Hz, commonly referred to as the audio spectrum [9]. The audio spectrum can be divided into the low-, mid- and high-frequency band, as can be seen in figure 2.1. As evidenced by the equal-loudness contours of figure 2.2, the ear is non-linear and is extra sensitive to frequencies in the upper mid-range, which need a lower volume than frequencies in the low- and high-bands in order to be perceived as equally loud. Each curve in the figure indicates the sound pressure levels that are required in order for our ear to perceive the frequencies as being equal in loudness level to a 1 kHz reference level. The difference is larger at lower loudness levels.

The human voice produces frequencies in most of the audible range, but not all of them are needed for us to understand what is being said. The 'telephone bandwidth' used in early telecommunications only ranged from 300 to 3400 Hz, and was adequate at the time, albeit at a noticeably lower quality than face-to-face communication [10].

2.2 Speakers

A speaker is a non-linear device, meaning that the behavior of the speaker is hard to predict and can distort based on frequency and or input signal amplitude. The goal of loudspeaker design is therefore to produce a speaker that is as linear as possible [11]. The physical design of the speaker and its enclosure will greatly

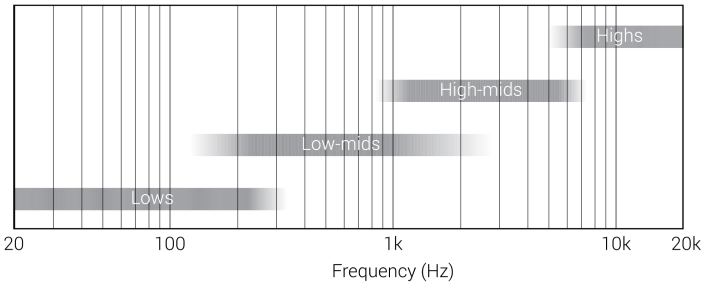

Figure 2.1: The different frequency bands of the audio spectrum. Lows and highs are commonly referred to as bass and treble respectively [1].

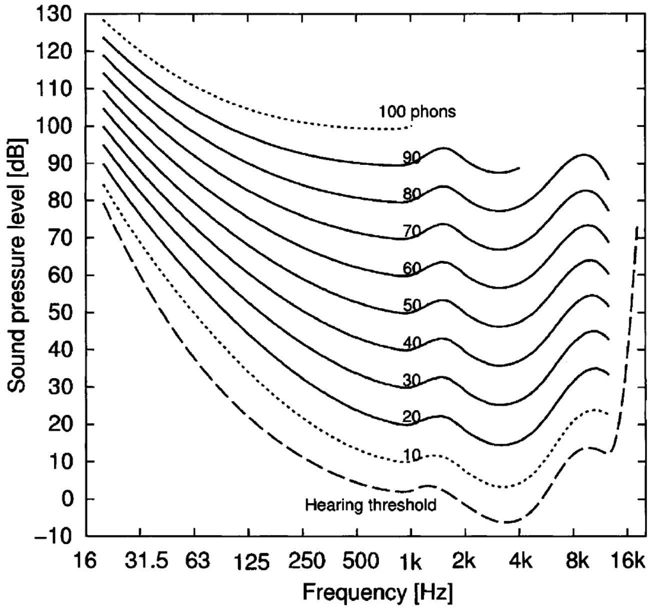

Figure 2.2: Equal-loudness-level contours from the ISO standard. A phon is a unit of perceived loudness, defined as the SPL of a 1 kHz tone [2].

affect its sound reproduction and frequency response. The perfect speaker does not exist, and the requirements placed on the speaker vary greatly with the context in which it is to be used. For example, if its only purpose is to give announcements in a public space, the quality most likely does not matter as much as it does for a speaker used as part of an audio engineer's monitoring environment.

A common measure of the quality of a speaker is its harmonic distortion, which refers to unwanted tones that are harmonically related to the original signals [12], i.e. their frequencies are integer multiples of the original tone. These harmonic distortions are dependent on the physical design of the loudspeaker and its enclosure [11]. The total effective value of all harmonic components, normalized with respect to the fundamental component, is the Total Harmonic Distortion, THD, which can be expressed as the ratio

$$\mathrm{THD} = \frac{\sqrt{\sum_{i=2}^{\infty} V_i^2}}{V_1} \tag{2.2}$$

where V_i is the rms voltage or SPL of the ith harmonic frequency, with $i = 1$ being the fundamental frequency.

Even if the system has low THD, noise and interference can enter the system from sources such as the power supply, electromagnetic induction or noisy circuit components and degrade the audio quality. As such, it is often useful to extend THD to include noise components in the measurement, resulting in the THD plus Noise metric, THD + N. This provides a broader estimate of unwanted signal content, and is also easier to measure, making it useful in comparing different loudspeakers. THD + N can be very different depending on the speakers, but is often in the range of -60 to -20 dB, where the lower it is, the better the loudspeaker.

The frequency response curve of the speaker is another common measure of quality, which represents how the device will affect different frequencies in the audio spectrum [9]. The output of the device is measured with a microphone as a reference tone is swept across the frequency spectrum at equal level, giving an accurate measure of any volume changes with respect to frequency that the device exhibits. An ideal speaker would have a completely flat frequency response, and therefore not affect the balance of the input signal at all. In reality this is not achievable, and for that reason it can be useful to tune the frequency response of a specific speaker model by applying different filtering techniques on the signal beforehand.

Speakers can be organized into different types of speaker systems. The spatially distributed system shown in figure 2.3a sends the same full-range audio signal to each speaker to cover a larger area, which is common in public address or background music applications. In contrast, the crossover network of figure 2.3b divides the audio signal into separate frequency bands using crossover filters, directing bass, midrange and treble content to speakers optimized for each range. This type of network is more common in HIFI, high fidelity, audio systems. These approaches are not mutually exclusive, and a speaker with a built-in crossover filter can easily be part of a spatially distributed system.

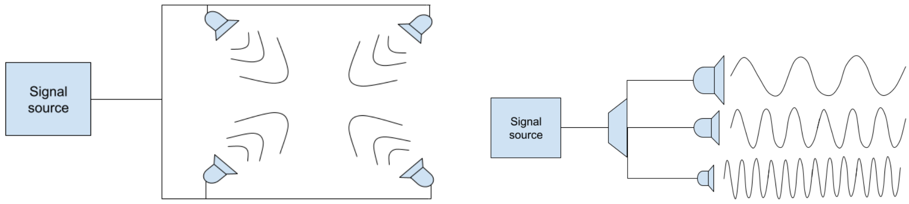

(a) Spatially distributed network. **(b)** Crossover network.

Figure 2.3: Two different kinds of speaker networks. The spatially distributed network sends the same audio signal to all speakers, whereas different speakers are responsible for different parts of the frequency spectrum in the crossover network.

When using a spatially distributed speaker system with a centralized signal source, maintaining consistent sound quality can be challenging. Differences in speaker characteristics, or even entirely different speaker models can make it difficult to apply effective centralized signal processing since the processing needed might differ from speaker to speaker. One solution is to process the signal locally at each speaker.

In modern sound systems that include microprocessors, the signal processing is typically performed in the digital domain [13]. Digital signal processors can be very flexible, efficient, and programmable. They enable operations that are impossible to implement in the analog domain, such as adaptive filtering, encryption and error correction. One example is the SigmaDSP ADAU1761, a low-power DSP that draws between 6-60 mA during operation and features configurable pins for various applications, including analog or digital audio inputs and outputs [14]. Because it is mostly programmable it can be fairly easily altered to match new requirements. For some systems however, limited processing power can constrain the ability to perform real-time digital filtering, and in certain cases it might not be practical or cost-effective to include a DSP for each speaker.

2.3 Analog filters

Analog filters are concerned with analog signals, i.e. varying continuous voltages and currents [4]. This contrasts with digital filters, which operate on data encoded in a binary format. This is a large field and analog filters can be further classified depending on the components used in their construction. Including, but not limited to, crystal, mechanical, microwave or switched-capacitor filters. Some of the more common are passive RLC filters, where the components are resistors, capacitors and inductors, and active RC filters, which consists of resistors, capacitors and active elements, such as transistors or amplifiers.

An analog filter can be characterized by its transfer function [4], which is essentially the ratio of the output signal to the input signal in the s-domain. It can also be described as the Laplace transform of the system's impulse response. It is in general a rational function, which means it is a fraction of two polynomials of s. A generic transfer function can therefore be written on the following form

$$H(s) = \frac{\sum_{i=0}^{M} a_i s^i}{\sum_{j=0}^{N} b_j s^j} = \frac{\prod_{i=0}^{M}(s - z_i)}{\prod_{j=0}^{N}(s - p_j)} \tag{2.3}$$

where the polynomials can be expressed as a product of the poles, p_j, or zeros, z_i, or as a general polynomial.

The transfer function can tell much about the behavior of the filter [3]. The order of the filter corresponds to the number of poles, i.e. N in equation 2.3, and each pole corresponds to a reactive element, i.e. a capacitor or inductor, in the circuit realization. By evaluating the transfer function on the imaginary axis $s = j\omega$, the frequency response of the filter can be determined [4].

The phase of the frequency response is a measure of how much the filter shifts the phase of each frequency component of the input signal. This phase shift corresponds to a time delay experienced by sinusoidal components as they pass through the system. Group delay provides a practical way to assess these time-domain effects. It is defined as the slope of the phase response with respect to frequency and indicates the extent to which different frequency components are delayed relative to one another [11]. A flat group delay means all frequencies arrive at the same time, while a peak in group delay suggests some frequencies have a larger time-delay than others, which can alter the shape of the waveform.

When two similar waveforms with different phases are summed directly, e.g. electronically, frequency distortion can occur due to constructive and destructive interference [1]. In extreme cases, such as phase inversion, destructive interference will cause the signals to cancel out completely [9]. However, if the signals are played through different speakers, such as in a stereo setup, phase differences do not sum directly at the ears. Instead, the brain interprets the resulting differences in the time it takes the sound to reach each ear, and other spatial cues, to localize the sound source. In such cases, small phase shifts typically do not cause noticeable artifacts, and their perceptual impact is minimal compared to situations where signals are summed electronically.

The magnitude of the frequency response is also an important way to characterize a filter, and is also known as the amplitude response. This is illustrated in figure 2.4 where the four common filter types, low pass, high pass, band pass and band stop or notch filters are shown [3]. These show the passbands and stopbands of the different filter types, as well as the cut-off frequency ω_c or the bandwidth between ω_1 and ω_2. The attenuation is normally large in the stopband and the signal can be assumed to be largely unaffected in the passband.

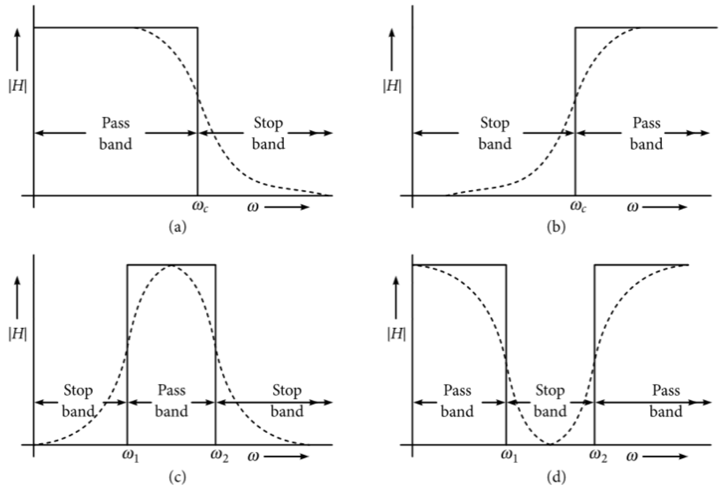

Figure 2.4: The graphs show the bode amplitude plot of ideal filters and dashed lines show a practical filter for (a) low pass, (b) high pass, (c) band pass, (d) band stop or notch [3].

The ideal response shown in figure 2.4 for the different types of filters are not attainable in the real world. One main difference from the ideal is how fast the passband transitions into the stopband, or vice versa, which is called roll-off [3]. Roll-off is generally proportional to the order of the filter, where higher order results in a higher roll-off. However, the rate of roll-off is also dependent on the design of the filter. For example, a Chebyshev filter is designed to roll off quickly, at the expense of introducing ripples in either the stop- or passband. In contrast, a Butterworth filter provides a maximally flat passband, but has a more gradual roll-off of $20n$ dB/dec, where n is the order of the filter.

For a filter with a resonant peak, such as a band pass filter seen in figure 2.4c, the amount of roll-off as you move away from the resonant frequency is referred to as the selectivity of the filter [15]. A filter with high selectivity strongly emphasizes the resonant frequency while attenuating nearby frequencies, resulting in a sharp resonance curve. Conversely, a filter with low selectivity exhibits a broader response, allowing a wider range of frequencies around the resonance to pass with similar amplitude.

The sharpness of the resonance is commonly described using the quality factor, Q, which is defined as the ratio of stored energy to dissipated energy per cycle [15]. It can also be expressed as

$$\mathrm{Q} = \frac{f_0}{BW_{3dB}} \tag{2.4}$$

where f_0 is the resonant frequency and BW_{3dB} is the bandwidth of the resonant peak, measured between the points where the amplitude has decreased by 3 dB. Higher Q systems are more selective, having narrower bandwidths and greater frequency discrimination.

An important time-domain measure is the step response, from which parameters such as overshoot, settling time, delay time, and rise time are derived. These parameters, illustrated in figure 2.5, provide insight into the transient behavior of the filter [4]. The amount of ringing, damped oscillations in the output signals, exhibited by the step response, together with the rise time, is the time required for the output to remain within a specified tolerance band around the steady-state value, also known as the settling time [16]. If the ringing never settles, and instead increases over time, the system would be classified as unstable [17]. The larger the bandwidth of the filter, the shorter the settling time becomes [13].

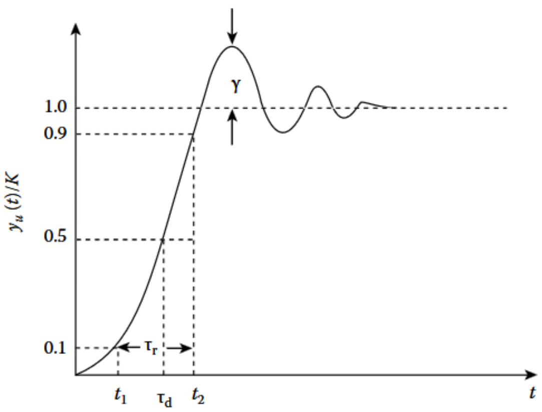

Figure 2.5: A normalized step response of a generic filter. In the figure the delay time, τ_d, rise time, τ_r, and overshoot, γ, are illustrated [4].

It can be shown that in order to realize an arbitrary filter's frequency response, a designer only needs a circuit that realizes a number of biquadratic transfer functions [4]. This can be expressed as

$$H_{BQ}(s) = \frac{a(s+z_1)(s+z_2)}{(s+p_1)(s+p_2)} \tag{2.5}$$

where the zeros $z_1 = z_2^*$ and poles $p_1 = p_2^*$ are complex conjugates [4]. There has been a lot of work on biquad theory and realization. A topology that realizes the biquadratic transfer function can be used to implement all common filter shapes in a predictable and stable way.

2.3.1 Deviations from Ideality

Ideal inductors and capacitors are purely reactive, but in reality they have parasitic components that can greatly affect their performance in a circuit [16]. The losses and reactive impurities present in capacitors and inductors cannot be neglected in the design of analog filters, especially not when a very sharp response curve or narrowband filter is desired [18]. The losses tend to decrease the rate of attenuation roll-off, increase the attenuation within the pass-band, and in certain cases, prohibit the realization of narrow band-pass filters.

Electrically, capacitors exhibit some equivalent series resistance, ESR, and series inductance, ESL [16]. The ESR is frequency dependent, generally reaching a minimum around 100 kHz to 1 MHz, being smaller for larger capacitance values. The ESL together with the capacitance forms a series LC circuit, giving the overall impedance a minimum at the resonant frequency, where the reactances cancel each other out.

Similarly, inductors have imperfections such as internal resistance and capacitance, that make them deviate from ideality, resulting in a resonant frequency at which the inductor's distributed capacitance resonates with the inductance, cancelling each other out [15]. In practice inductors are generally avoided, if possible, since they tend to be more bulky and expensive, are susceptible to magnetic pickup of interference, and depart further from the ideal than capacitors by having significant ESR [13]. This is especially true at frequencies under 1000 Hz, where the parasitic resistance of the inductor becomes so large that a good resonant circuit is difficult to realize [18], not to mention the cost and physical space a large inductor requires.

Another deviation from ideality that electrical components experience is that their real values differ from their nominal values [3]. This is partly due to the fabrication process tolerance, which is expressed as the deviation in percent from the nominal value [15], and partly due to the aging of the circuit, and environmental effects such as temperature and humidity, causing the components to drift from their initial values [19].

The amount of introduced error by differing component values depends on the sensitivity of the circuit [3]. The concept of sensitivity is an important criterion for comparing circuit configurations and evaluating how well they meet design requirements [19]. Understanding how sensitive a circuit is to variations in specific component values can inform design trade-offs. For example, if the filter's performance is relatively insensitive to the value of a particular capacitor, a designer may opt for a less precise, lower-cost component without significantly affecting the circuit performance [4].

2.3.2 Noise

Noise is a factor in all filter circuits and every component produces some form of noise which in turn affects the SNR of the filter [4]. It is measured in $V/\sqrt{Hz}$ or $A/\sqrt{Hz}$ if it is current-based. All resistors exhibit thermal noise, which can be expressed as

$$E_t = \sqrt{4k_B T R \Delta f} \tag{2.6}$$

where E_t is the thermal voltage noise, k_B is Boltzmann's constant, T is the temperature, R the resistance of the component and Δf is the noise bandwidth. The noise bandwidth is the bandwidth of the filter the component is used in. The thermal noise exists in capacitors and inductors as well, since no component is ideal, and they all have some type of resistance.

In capacitors, since the noise bandwidth of a RC circuit is $4RC$, the thermal voltage noise can be expressed as

$$E_t = \sqrt{\frac{k_B T}{C}}. \tag{2.7}$$

where C is the capacitance. Similarly, for inductors, the thermal current noise can be expressed as

$$I_t = \sqrt{\frac{k_B T}{L}} \tag{2.8}$$

where I_t is the thermal current noise and L is the inductance.

Active circuits can have other types of noise sources [4]. In addition to thermal noise produced by the passive components, active components add to the noise, which is often modelled by using current and voltage noise sources at the input. These are, in general, larger than the thermal noise from passive components and are often reported in their respective data sheets. The noise of an active component is also affected by the gain of the circuit due to noise gain, and therefore it is important to consider the noise performance when designing since the ideal gain of an active stage might add too much noise to the output.

2.3.3 Circuit Loading

Connecting a load to a circuit can significantly affect its output, depending on the proportion of the load impedance to the circuit's output impedance [13]. In circuits that can be modeled using a Thevenin equivalent, a voltage divider is formed between the output impedance and the load impedance. As a result, the voltage seen across the load depends directly on their ratio. To minimize loading effects, the output impedance should be made as low as possible, while the load should present a high input impedance to avoid drawing excessive current and altering the intended behavior. Conversely, for circuits better represented by a Norton equivalent, the goal is to maintain a high output impedance, ensuring that most of the output current is directed into the load rather than being dissipated internally.

To realize higher-order filters one can cascade several lower-order sections, but as more sections are added circuit loading will cause the interactions between them to increase [3]. This interaction causes the component sensitivities to go up, meaning that any variation in component values will affect the performance of all cascaded sections [20]. One solution is to make each successive filter section have much higher impedance than the preceding one, but the issue can be avoided completely by using active circuits, like transistors, operational amplifiers, and buffers, to isolate the sections from each other [13].

2.4 Passive filters

Passive filters are a category of analog filters that rely solely on passive components, most commonly resistors, capacitors, and inductors [3]. They are typically responsive to frequencies between around 100 Hz and 300 MHz, with the lower limit resulting from the large capacitor and inductor values that would be required [15]. They are linear, inherently stable, and do not introduce much noise, but due to the absence of active components they can not provide any gain, resulting in an output power that is always less than the input power [21]. This power reduction is called insertion loss, as the signal power is reduced when the filter is inserted into the signal path [9].

Passive filters can be classified either by the components they consist of (e.g. RLC), or by the configuration of said components, also known as the circuit topology [18]. Among the main considerations when designing passive filters are how these topologies affect the filter's performance in terms of insertion loss and impedance matching at the input and output terminals [22].

Kimball [5] presents and analyzes several general circuit topologies, six of which are shown in figure 2.6. They are all made to depend on two impedances, Z_1 and Z_2, which are the inverse to each other with respect to the line impedance R_0, such that

$$Z_1 Z_2 = R_0^2. \tag{2.9}$$

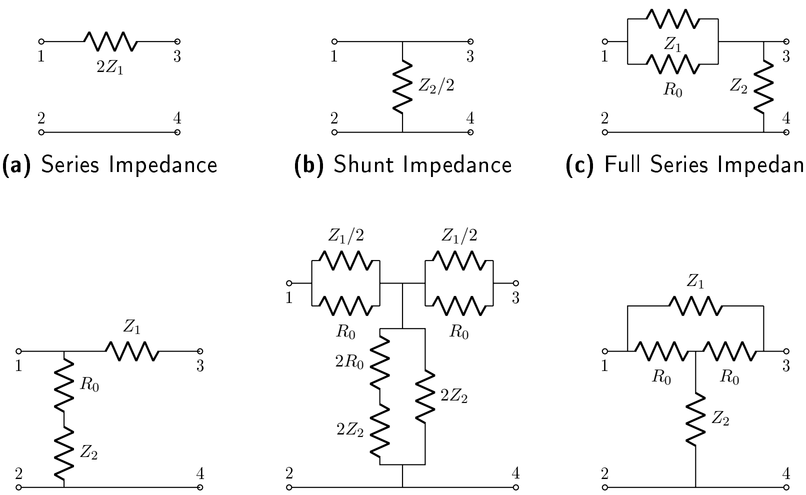

(a) Series Impedance **(b)** Shunt Impedance **(c)** Full Series Impedance

(d) Full Shunt Impedance **(e)** T **(f)** Bridged-T

Figure 2.6: General circuit topologies, as presented by [5].

The same insertion loss characteristic can be achieved using all six topologies, but the impedance match at the input and output terminals differ [5]. For the series and shunt impedance types, the input and output resistances vary with frequency. In contrast, the full series and full shunt configurations maintain a constant input resistance, while the T and Bridged-T topologies provide constant and matched input and output resistances.

The bridged-T section uses fewer components than the regular T section [23], and two possible implementations are shown in figure 2.7. One of the main advantages of this topology is that it exhibits constant-loss, in the sense that varying the attenuation does not change the insertion loss introduced by the filter. The bridged-T network is also constant resistance [22], meaning that the input resistance does not change with frequency if the network is correctly terminated at one end by its image impedance, R_0, regardless of the termination at the other end [24]. A key benefit of such constant-resistance behavior is that multiple sections can be cascaded with their insertion losses adding directly.

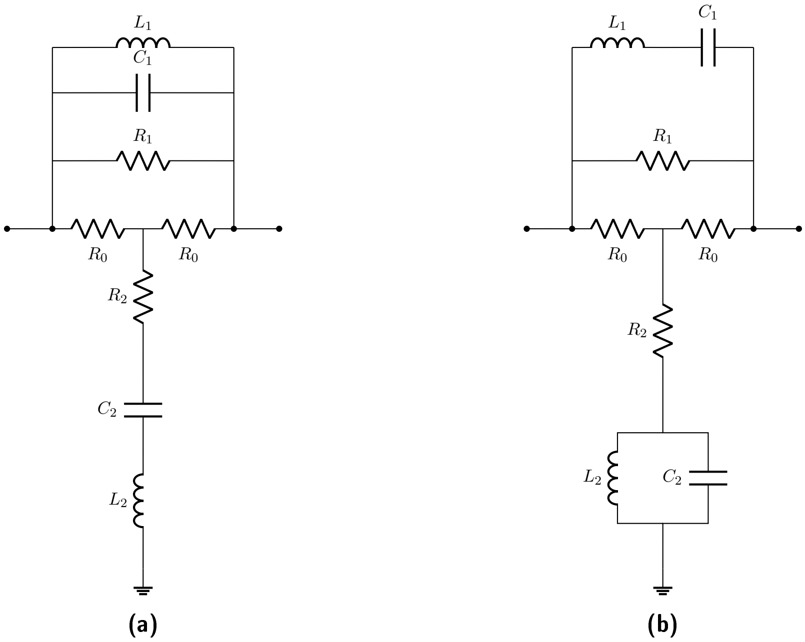

Figure 2.7: Bridged T filters implementing the notch (a) and narrow gain (b) filter functions [5].

Achieving high selectivity in passive filter design often necessitates cascading several filter sections containing both inductors and capacitors. The more sections are added, the harder it becomes to avoid an increased insertion loss [15]. As discussed in section 2.3.1, inductors can be undesirable due to their size, cost, and susceptibility to parasitics. Consequently, realizing highly selective filters without inductors is challenging. In such cases, active RC filter designs are typically employed as a practical alternative [13].

2.5 Active filters

With active filters, any RLC filter characteristic can be synthesized without the use of inductors [13]. Active filters gained traction as operational amplifiers started becoming cheaper and more reliable [3], and they are especially useful in integrated circuits. One benefit of operational amplifiers is that they have, in the ideal case, infinite input impedance and zero output impedance [25]. This could be used to buffer passive circuits, i.e. add a voltage buffer between circuit sections, which would reduce the effect of circuit loading and insertion loss. Another benefit is that they can provide both gain and attenuation to a circuit.

As discussed in the previous section, the possibility to avoid inductors is a big advantage of active circuits. Inductors can be replaced with gyrators that are an electrically equivalent active circuits [3]. Figure 2.8 shows a gyrator that only requires one operational amplifier, with the limitation that it can not replace floating inductors, since it needs to be grounded at one end. There are ways to realize a true floating inductor gyrator circuits with the drawback of needing at least four operational amplifiers, and a multitude of passive components, per inductor [3]. The drawbacks of using gyrators closely mirror those of active circuits in general, including power consumption, added noise, susceptibility to interference, and potential instability due to feedback [21].

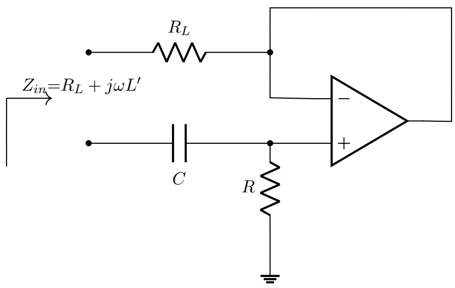

Figure 2.8: A realization of a gyrator circuit [6]. Where $L' = R_L RC$ is the simulated inductance value of the gyrator.

There are a multitude of different active filter topologies which are different in many ways. The Sallen-Key topology, also known as a voltage controlled voltage source, VCVS, is particularly valued for its simplicity [3]. This topology is a biquad, which means that it can be used to realize any arbitrary transfer function. The topology can be used to construct any of the filters seen in figure 2.4, but a bandpass filter implementation of VCVS can be seen in figure 2.9b. The gain of the filter is decided by the ratio of R_a and R_b and care needs to be taken in order to ensure stability, since too high gain will lead to oscillations.

Another popular configuration is the multiple feedback filter, MFB [3]. This topology is shown as a bandpass filter in figure 2.9a, and is also a biquad circuit. It needs two fewer resistances, and exhibits lower sensitivities compared to the VCVS bandpass filter [26].

If sensitivity or stability is a big concern, the state variable filter is a possible route [3]. This type of filter can realize a transfer function of order n by cascading n integrators and a summing amplifier. One practical implementation is called the Tow-Thomas biquad, which uses two integrators and a summing amplifier. This filter requires only six passive elements and has really low sensitivity and can achieve very high Q-values while still being stable. However, state variable filters use several operational amplifiers to achieve its performance, which can become costly.

With most active filters, the Q factor affects the noise levels [27]. The higher the Q factor the higher the gain and also the noise gain of the circuit. This leads to higher output noise level at the frequencies close to the center frequency.

The output swing of an operational amplifier is limited by the supply voltage, if a signal is boosted beyond the supply it will lead to clipping of the signal [3], which is illustrated in figure 2.10. Clipping would introduce distortions in the system due to higher harmonics caused by the sharp transitions. Another issue that clipping

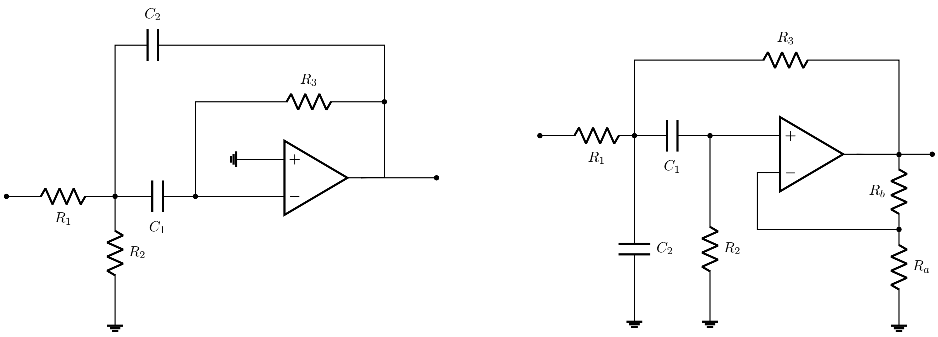

(a) MFB **(b)** VCVS

Figure 2.9: Two different bandpass filters topologies using a multiple feedback topology, (a), and a voltage controlled voltage source topology, (b),

introduces is loss of information, which in e.g. audio systems would mean that the sound produced would sound different or in the worst case be unintelligible.

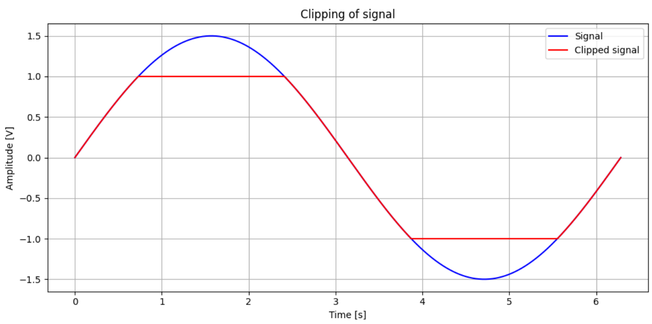

Figure 2.10: A graph illustrating the clipping of a signal, which would lead to distortion and loss of information in the signal.

2.6 Equalizers

As noted by [28] analog filters or tone control filters have found extensive use as a means to shape the audio signal before it reaches the speakers. Rather than completely rejecting frequencies in the stop bands, tone control filters attempt to either attenuate or boost them by a small amount [29]. The first tone control circuits were passive analog circuits and could either be fixed or utilize potentiometers to vary the amount of cut or boost to a certain degree.

An equalizer can be seen as a more advanced form of tone-control. It is a four-terminal network whose attenuation loss varies with frequency in some desirable manner [22]. Sound equalization is an important part in audio signal processing, and has many use-cases [30]. One use case is illustrated in figure 2.11 where equalization is used to produce a flatter speaker response.

While frequency distortion is easily perceived due to the ear's sensitivity to intensity, phase distortion is much less distinguishable [4]. Phase distortion occurs when the phase response of the transfer function is non-linear, which is common in higher-order filters. However, since phase distortion is generally less perceptible, equalization efforts tend to focus on shaping the amplitude response rather than correcting phase variations. Phase distortion effects can be mitigated by modulating the phase in the digital domain, so that the group delay is the same for all frequencies, before the signal is sent through the analog systems [3].

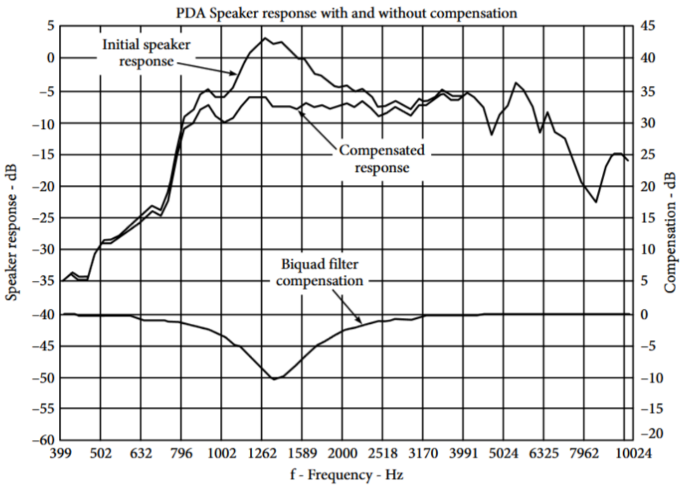

Figure 2.11: The speaker response before and after compensating with a filter [3].

Equalizer circuits can be divided into a multitude of different, at times overlapping, categories [23, 24]. These include, but are not limited to, fixed, variable, graphic, and parametric equalizers.

Fixed equalizers have no variable elements and are designed to combat known non-linearities or problematic frequencies, such as 50 Hz noise from the power grid [3]. They are commonly used as permanent parts of recording and reproducing equipments [22]. In contrast, variable equalizers expose various knobs that the operator can turn in order to vary different parameters of the filter [23].

Graphic equalizers and parametric equalizers are two types of variable equalizers. Parametric equalizers provide the operator with a set of bands, where the gain or attenuation, bandwidth and center frequency can be varied as needed. On a graphic equalizer on the other hand, the bands are fixed in frequency and bandwidth, only allowing the gain of each band to be varied. Since these circuits are designed to be operated and adjusted with human input, they need to be rather complex in order for the behavior to be smooth and predictable.

2.6.1 Graphic Equalizers

There are several factors or design considerations when building a graphic equalizer [23]. It can be either passive or active, either constant Q or proportional Q and a designer needs to choose the width, and number of frequency bands at which it operates. As the name might imply, decreasing the gain of a proportional Q equalizer also decreases the Q value proportionally, yielding a lower selectivity compared to a constant Q graphic equalizer, which can achieve the same sharpness of the notch regardless of the amount of attenuation or gain.

The number of bands in a graphic equalizer determines its resolution, and depends on the desired selectivity and the frequency range of interest [30]. If there is one filter at each octave frequency, then the equalizer would need 10 bands to cover the audio spectrum. If the bands instead cover just half or a third of an octave, it would require 20 and 30 bands respectively. There is also room for variation within graphic equalizer design since all parts of the frequency range might not be of the same importance, meaning that more interesting regions could have narrower bands for more control.

One implementation of a constant Q graphic equalizer can be seen in figures 2.12 and 2.13, which achieves constant Q by separating the volume section from the filter section [31, 32, 7]. The basis of the equalizer are the two inverting summing amplifiers, into which a number of inverting band-pass filters are connected with their inputs at the output of the first summer and the outputs connected to adjustable taps of potentiometers individual to the filters. The ends of the potentiometers are connected to the input of either summer, contributing two signal components to the signal passing from the input to the output.

The output of the first summing amplifier is fed into the filters and the second summing amplifier. The filter output that gets fed back to the first amplifier, U1, will therefore be phase inverted relative to the input, meaning that the filtered signal will be cut away from the input signal through destructive interference, leading to an attenuation of the signal at the frequencies selected by the filter. Similarly, at the input of the second summing amplifier, U2, the output of the first amplifier will be compared to the signal coming from the filters, both of which will be inverted relative to the input, which will lead to the signal getting a boost in amplitude at the frequencies selected by the filter through constructive interference.

The maximum amount of boost or cut at each filter frequency is controlled by the third amplifier in figure 2.12, U3, as it decides the amplitude of the filter signal relative to the input signal by

$$V_{\text{sig}} = -\frac{R_4}{R_3} \cdot V_{n3} \tag{2.10}$$

where V_{n3} is the voltage at the negative input of U3.

The actual amount of boost or cut, on the other hand, is decided by the ratio between the feedback resistors, R2 and R6, of the summing amplifiers and the resistance through the potentiometer and R1 of the filter. By adjusting the value of the potentiometer you can control how much gain each of the summing amplifiers will apply to the filter signal.

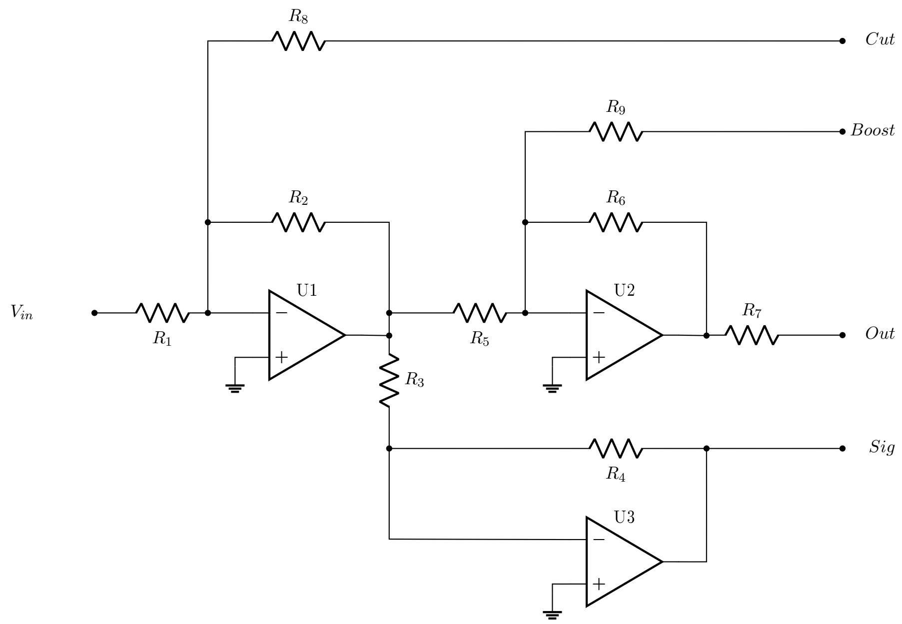

Figure 2.12: The summing section of the constant Q graphic equalizer, with U1 and U2 being the two inverting summing amplifiers, and U3 controlling the gain of the filter-signal [7].

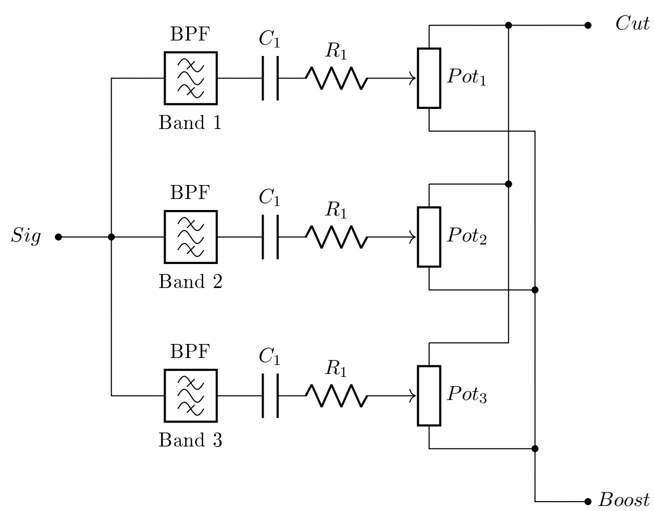

Figure 2.13: The filter section of a constant Q graphic equalizer [7], where a plurality of bandpass filters are connected in parallel, the output of which goes through the potentiometers that control how much of the signal is sent to which of the summing amplifiers in fig 2.12.

2.7 Optimization

The conventional circuit synthesis methods require a transfer function to be found, or given, which simultaneously satisfies all specifications and is realizable for the chosen circuit configuration [33]. Many important practical problems can not be solved using this method, and optimization is likely to be needed when the specifications are difficult to approximate with a transfer function of reasonable complexity. Moreover, [4] states that computer optimization is necessary if completely arbitrary gain characteristics are desired.

Trying to design filters in order to compensate for non-linear and complex behavior such as speaker response can be difficult. This is due to the large amount of variables that can be changed and the non-traditional shapes the filter will have to adapt. However, there are tools available that can help through different optimization algorithms. Measurements of the speaker response can be used to define the optimal filter goal function and then tune the parameters of an analog circuit until it is reached.

A tool that uses an iterative optimization approach to designing analog filters is LTspice Opt [34] which integrates with the open source simulator LTSpice. The program takes in a circuit topology in the form of an LTSpice file and a target frequency response and then varies the values of specified parameters in order to minimize the distance between the goal function and the circuit behavior. The tool does two rounds of optimization, first a particle swarm optimization in order to avoid local minima, then a least square optimization in order to get to the minima faster.

2.7.1 Least square

Least square optimization is a supervised algorithm that fits certain parameters to a certain curve in order to minimize the sum of squares of the difference between the model and the goal, often denoted as the cost function [35]. It is used when working with a non-linear or complex model and updates the parameters according to different methods, most commonly Gauss-Newton or gradient descent. The Gauss-Newton method assumes that the objective function, the sum of squares, behaves as a quadratic function when close to the optimum and utilizes this in order to quickly approach the optimal parameters. Gradient descent on the other hand alters the parameters in the direction of the steepest gradient in order to approach a minimum of the cost function. These work well together, since gradient descent works better far from a minimum, while Gauss-Newton converges much quicker closer to a minimum.

A positive aspect with this optimization algorithm is that a model of the problem is included. Therefore, the gradient descents updates values with respect to their effect on the model of the problem. This can help convergence speed since the gradient knows how much or little to alter a variable in order to take a step in the right direction.

The main issue with gradient descent based algorithms is that they might get stuck in a local minimum, without noticing the global minima of the cost function [35]. Therefore, it is important to explore several parts of the possible solution space in order to arrive at the global minima.

2.7.2 Particle swarm

Particle swarm is a blind optimization algorithm that does not rely on gradients in order to approximate an optimal solution [36]. Instead, the algorithm is inspired by the social behavior of flocks of birds or schools of fish in order to approach an optimum.

This optimization algorithm first began as a simulation of groups of people or other animals who explore a space in order to avoid predators and find food, or to adjust ones own beliefs according to the majority [36]. Large populations of animals have been observed to collectively find some optimum in a large space based on the movements of their peers, which is what particle swarm optimization seeks to achieve. The solution space is populated with different particles, and they start to move around to explore the space. The particle with the best current solution is known, and each particle is then guided towards the current best solution. This occurs at the same time as each particle's direction and velocity is stochastically modified as well, in order to avoid converging to a local minimum too quickly.

There are many parameters that could be introduced, but common ones are cognitive, social and inertia parameters. The cognitive parameter takes a particles current position into account, the social parameter takes the group's best performing location into account and the inertia parameter decides how hard or easy it is to change direction and converge.

This type of algorithm is good for exploring the state of possible solutions due to the population of the solution space and the stochastic behavior of the particles [36]. Therefore, it is a popular choice to use for non-linear or complex systems. However, there are some drawbacks of particle swarm such as comparatively slow convergence speed and the fact that it is not guaranteed to find the optimum.

Since it is a blind optimization algorithm, the knowledge of the design space, such as circuit design, can help somewhat, but in a limited fashion. It is important to decide upon good bounds and number of particles. Depending on the problem to solve, the social and or inertia parameters could be tuned in order to avoid converging to local minima too quickly and explore the solution space more thoroughly. However, this would significantly increase the execution time of the algorithm.

Chapter 3

Measurements and Simulations

In order to design frequency compensating filters, it is important to thoroughly understand and measure the different speakers that are considered. The simulation setup should also be realistic and good models are important in order to ensure accurate results when building a product.

3.1 Measurement Setup

Several different types of measurements on the speakers or filter have been conducted in this thesis. There have been electrical and thermal tests on the filter. The filter and the speakers have also been analyzed with frequency response measurements as well as subjective listening tests.

Electrical measurements on the circuits were taken using a Keysight 34465A digital multimeter and a WaveSurfer 4024 HD oscilloscope. The multimeter measured the current drawn from the DC supply voltage in order to measure the power consumption of the filter at different frequencies. The oscilloscope was used in order to measure the frequency response.

Thermal measurements were taken with a UTi260B professional thermal imager. The measurements were taken when the speaker was driven at 1 W or 2 W for different stimuli. The stimuli were white noise or single tone sinus waves at different frequencies. The input signal stayed on for 2 minutes or more if the temperature was not stable before taking a measurement and before testing with a new signal the PCB was cooled with a freezing spray.

3.1.1 Frequency response measurements

Measurements on the speakers were conducted in an anechoic chamber which has a noise floor of 8.3 dBSPL. The chamber has been certified by IAC acoustics and uses both software and hardware from Audio Precision. The microphone was a GRAS 46AE 1/2" CCP free-field standard microphone with a sensibility of 50 mV/Pa and a dynamic range of between 17 dB to 138 dB.

The measurement setup is depicted in figure 3.1. The speaker was placed one meter from the microphone in order to ensure far field measurements of the frequencies of interest.

The frequency response, the harmonic distortion and noise are of interest and these are captured by sending a logarithmic sweep of the audible range and listening with the microphone. These metrics are captured by the microphone and

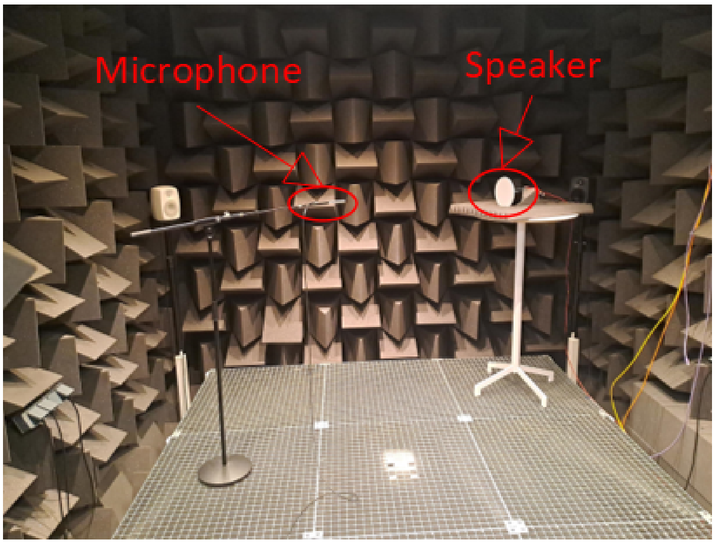

Figure 3.1: A picture showing the measurement setup while measuring the frequency response of the different speakers.

audio precision calculates the frequency response and THD+N. Most of the measurements are done in order to drive the speaker at 1 W and the source is calibrated to deliver the same power over the entire frequency range.

The input signal can be varied as well in order to drive the speaker at different power levels and this is used to look at the linearity of the speaker's frequency response as well as to analyze the power consumption of the filter.

3.1.2 Listening tests

In order to assess whether the frequency compensation is successful or not, a subjective listening test was conducted. This test included several participants that would rate whether the speaker with or without filter sounded best, in some aspects, without knowing which audio was filtered or unfiltered. The test was conducted using the same speaker, but switching between a filtered and unfiltered audio signal.

In order to have the same volume for the filtered and unfiltered audio, a portable SPL-meter was used while playing white noise for both signal paths. The sound was measured at a distance of 1 m and was 78 dBSPL for both signals.

When designing subjective listening tests, it is important to have a reproducible test which can give clear results. According to a reference manual from the international telecommunication union [37] it is important to have good assessors that can distinguish certain audio artifacts in order to get a good results. Our population of assessors consisted of some audio professionals as well as master thesis student's. Since we can not ensure that every participant could accurately hear different audible artifacts, we decided to simplify the listening test somewhat.

Another important factor to a listening test is to use mock stimuli in order to assess the participants [37]. This was done by playing the same audio twice and checking if they could accurately identify that the audio was identical.

The questionnaire used in the listening tests can be seen in appendix B. The results from these questionnaires were later compiled into graphs.

In the listening test, seven audio files were used where the first three clips were of speech and the last four clips were different genres of music. The audio clips included a similar spread of bass, mid and treble frequencies. During the test the different audio were referred to as speaker 1 or speaker 2. Which alternative, speaker 1 or 2, that had the filter was altered. One time for the speech clips and one time for the music clips, speaker 1 and 2 were the same in order to assess the participants' ability to accurately hear audio differences.

3.2 Simulation Setup

Various filters were first simulated and evaluated before moving on to realization. LTspice was used to determine their behavior and characterize their figures of merit and drawbacks. However, a large part of filter design is how one decides the component values. Since the speaker frequency response is not suitable to be compensated with a traditional filter shape, an optimizer is a good choice for tuning the filter response. We used LTspice opt [34] in order to optimize the component values to best match the speaker response.

The speaker response from the measurements were used as a basis for the optimal filter that the optimizer used. The ideal behavior of the filter can be described by

$$|H(j\omega)| = L - |FR(j\omega)| - IL \tag{3.1}$$

where the magnitude of the ideal transfer function is decided by the desired output flatness of the speaker, L dBSPL, the measured frequency response, $|FR(j\omega)|$ dBSPL, and eventually some accepted insertion loss, IL dB. The IL factor could be zero if active filters are used. However, since passive filters do inherently add attenuation, a good matching passive filter needs to take the insertion loss into account. The value of L is based on the output level of the frequency response where the response is the flattest. This is done because we want to only address the peaks of the frequency response and not shift large areas up or down.

The optimization tool LTSpice Opt that was used needs a goal function and a control file as inputs [34]. The goal function is the response that the circuit should emulate, this can be an ideal goal function or something that is based on another type of circuit. We used a goal function which is based on equation 3.1. The control file has information about which values the optimizer can vary and within which bounds. The control file also contains information about the error weights for that run. The error weight tells the optimizer which part of the frequency spectrum is more important and an error weight of zero indicates that we do not care about those frequencies. Information about the tolerances of the components is also contained in the control file in E-series format, which indicates which tolerance the components have. We settled on using E96 which corresponds to a tolerance of 1%.

The optimization is run in two phases, one particle swarm optimization and one least square optimization [34]. The particle swarm optimization utilizes the python library Pyswarms [38] and the parameters used are $c1 = 0.5$, $c2 = 0.3$, $w = 0.9$, the number of particles are 30 and the lower and upper bounds come from the control file. The parameter $c1$ is the cognitive parameter which determines how much the particles current position should matter. The parameter $c2$ determines social behavior and decides how much the groups best position should affect the particle. The final parameter w describes a particle's inertia, i.e. how hard it is to change a specific particles direction and velocity.

The least square optimization is run using an algorithm written by the author of LTSpice Opt [34]. The cost function used is a linear one based on the squared distance to the goal function.

Chapter 4

Method, Results and Evaluation

In this section we will explore the use of analog filters in order to reduce speaker non-linearities by flattening their frequency response. The different types of filters will be evaluated on two different speakers in order to compare the performance as well as to ensure the method is versatile enough to work for many different types of audio systems. Both passive and active analog filters will be explored and evaluated.

4.1 Requirements of the filter

The speakers show frequency distortion, as can be seen in figure 4.1 and 4.2. The two speakers have different output sound pressure levels and deviations in the frequency response. The goal of the filter is to flatten the response, to ensure that most frequencies have a predictable and similar output sound pressure. Other enhancements to the sound, such as boosting bass and treble frequencies, can then be done centrally, and a similar response can be expected in different types of speakers in the same system.

The frequency response for speaker 1, see figure 4.1, shows that it has a higher output sound pressure for the same input power compared to speaker 2, see figure 4.2. It can also be noted that the average deviation for speaker 1 is larger than for speaker 2, 3.52 dB and 2.64 dB variation respectively. Furthermore, the highest rate of change occurs in speaker 1 around 2 kHz, which would require a very high selectivity in order to compensate, whereas the variations in speaker 2 do not require as high selectivity to compensate.

According to equation 3.1, the ideal filters for speaker 1 and 2 without added attenuation can be seen in figure 4.3a and 4.3b respectively. These frequency responses were the goal functions for the active filter implementations and the passive implementations had different levels of added insertion loss.

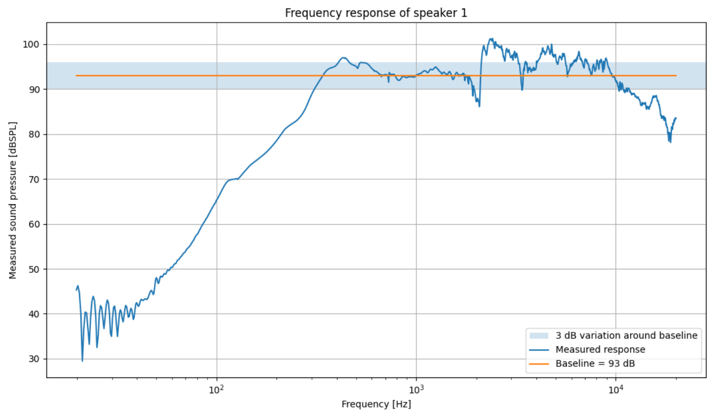

Figure 4.1: The speaker response for speaker 1. The average level of output sound pressure level was 93.24 dBSPL and the average deviation from the baseline, 93 dB, was 3.52 dB within the frequency range 400 Hz - 16 kHz. The speaker was driven at 1.125 W.

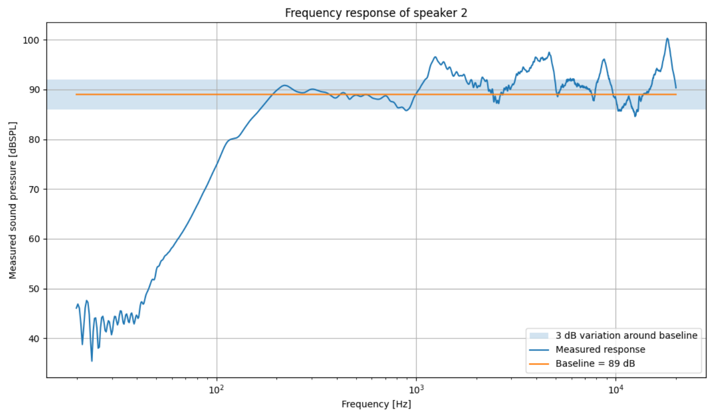

Figure 4.2: The speaker response for speaker 2. The average level of output sound pressure level was 90.02 dBSPL and the average deviation from the baseline, 89 dB, was 2.64 dB within the frequency range 400 Hz - 16 kHz. The speaker was driven at 1.125 W.

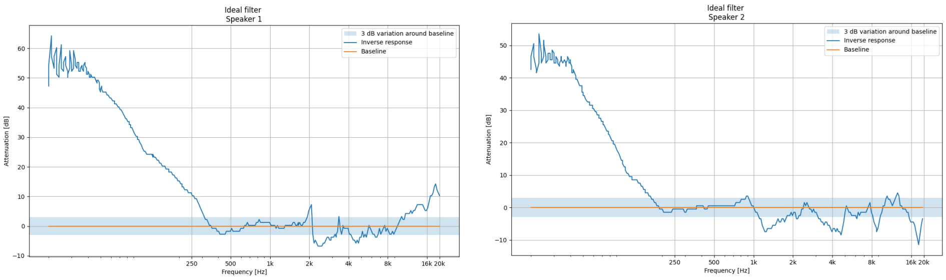

(a) Speaker 1 **(b)** Speaker 2

Figure 4.3: The inverse speaker response showing the ideal filter for speaker 1, (a), and 2, (b).

The frequency response of the speakers for different input levels can be seen in figures 4.4a and 4.4b for speaker 1 and 2 respectively. It can be seen that the response is fairly linear, and the overall shape is not dependent on the input signal power. The varying signal Vin is the input signal to the speaker driver.

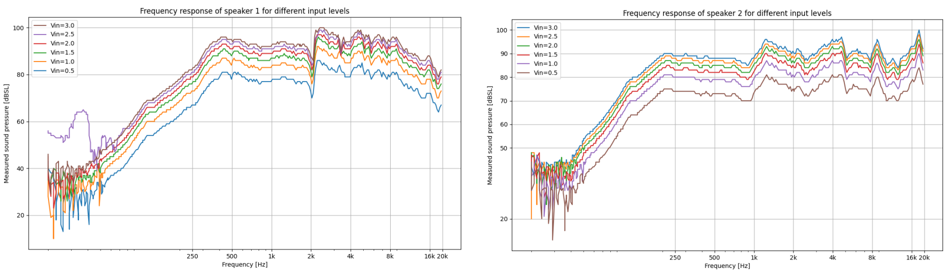

(a) Speaker 1 **(b)** Speaker 2

Figure 4.4: The speaker response for speaker 1, (a), and speaker 2, (b) at different input signal voltages.

Taking the average output sound pressure level for the different input voltages seen in figure 4.4 and converting it from dBSPL to Pascal, the linear behavior of the speakers within this input signal range is evident. This can be seen in figure 4.5. Thus, an equalization of the signal seems to be possible due to the linearity of the transfer function of the speaker.

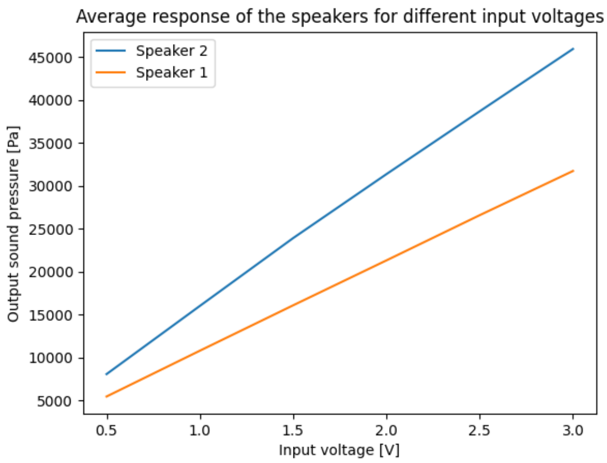

Figure 4.5: The average output sound pressure for different input signals for speaker 1 and 2.

Another important characterization of the two speakers are their nonlinearities and noise performance, which can be measured by THD+N, that can be seen in figure 4.6. It can be seen that speaker 1 exhibits worse THD+N performance since it is more noisy, and around 20 dB higher for frequencies near 2kHz compared to speaker 2.

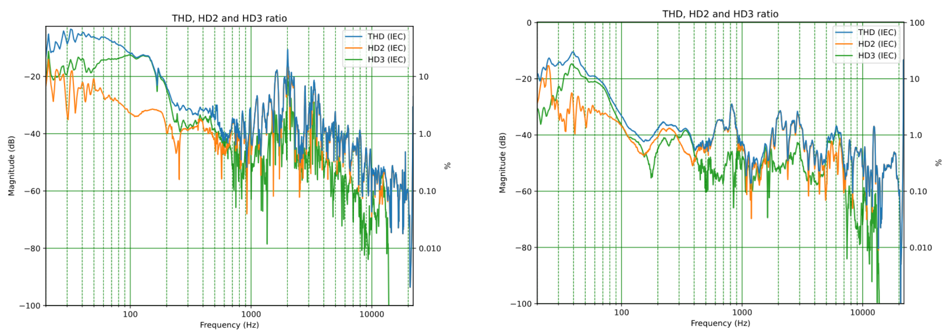

(a) Speaker 1 **(b)** Speaker 2

Figure 4.6: The total harmonic distortion and noise, THD+N, as well as the harmonic distortion for the second and third fundamental harmonics for speaker 1, (a), and 2, (b).

There are several aspects to take into account while designing a filter for a speaker. Firstly, it should achieve the desired results, which in this case is reducing frequency distortions. Secondly, the method should be expandable and general, since each individual speaker can exhibit different distortions. The frequency response will also be altered when introducing a new enclosure, other electronics or mounting the speaker in different locations. Therefore, it would be ideal to be able to have a design which is easily adapted to different iterations in the design process. Due to the poor bass performance shown for both speakers, see figure 4.1 and 4.2, as well as the fact that high treble frequencies are hard to hear, it was decided that the frequencies of interests are 400 Hz - 16 kHz. The criteria used when designing the filters are summarized below.

- Keep the output sound level high.
- Produce as flat a response as possible.
- Keep the power consumption low.
- Minimize the size and cost.
- Have a simple design procedure.
- The filter design should be easily modified and expandable.

When both passive and active filters have been investigated, one type will be realized and tested as a proof of concept. In order to do this, the design will be modified to comply with existing hardware that will be used for testing.

4.2 Passive filter design

The first step of the passive filter design was to choose a suitable filter topology. The Bridged-T notch filter, see figure 2.7a, was selected for a number of reasons. The design is expandable due to its constant resistance property reducing the circuit loading effect when cascading multiple sections [5], and the low attenuation in the passband is helpful in order to minimize the attenuation of the overall signal.

4.2.1 Methodology

The bridged T topology was implemented in LTSpice with varying amount of filter sections cascaded. These files were then used as a basis for the optimization script. We investigated 1 to 8 filter sections at different levels of acceptable insertion loss. A filter section is a bridged T filter and if there are several filter sections they are connected in cascade.

The optimizer is allowed to vary every value in the bridged T section except R_0, since R_0 decides the input and output resistance, see figure 2.7a. The resistors could vary between 0 Ω and 10 MΩ, the capacitors between 0 F and 10 mF and the inductors between 0 H and 1 H. The upper bounds are unrealistic, but if components receive a very large or very small value, it can be replaced with an open or closed circuit, reducing the component count. The optimizer then repeats the process for different amounts of acceptable insertion loss, between 5 and 30 dB.

Several passes of optimizations are conducted on the same circuit due to the large number of variables. The different optimization sessions focus on different frequency bands. The error weights are selected so that the frequencies that are not within the current frequency band are disregarded during optimization. The strategy is summarized in table 4.1 where the columns show which iteration of optimization it is and what frequencies are considered, and each row represent a certain filter that has one to eight total filter sections. The entries represent how many sections are varied for each iteration for the different filters, i.e. some filter sections are not considered for some passes of the optimizer in order to shrink the solution space somewhat.

Total filter sections	Pass 1 0.4-2 kHz	Pass 2 0.4-8 kHz	Pass 3 0.4-16 kHz	Pass 4 0.4-16 kHz
1	1	1	1	1
2	2	2	2	2
3	3	3	3	3
4	1	2	3	4
5	1	3	4	5
6	2	3	5	6
7	2	4	5	7
8	2	4	6	8

Table 4.1: Table showing how many filter sections within a certain filter are being varied for each pass of the optimizer and which frequency bands are considered.

The performance of the passive filters is evaluated by measuring how well they fit the ideal filter as well as how much insertion loss they add. The fit is measured by using regression fit, R^2, and the average deviation from the ideal filter in dB.

4.2.2 Simulation Results

The simulations and optimization were performed for both speakers. An example of a performance summary for 20 dB insertion loss can be seen in figure 4.7. It can be seen that for 20 dB insertion loss, the best performing filters had more filter sections, 7 sections for both speaker 1 and 2. There are also outliers, for example in speaker 1 with 8 filter sections, the R^2 score is -0.6 which suggests a worse fit than no filter. A similar result can be seen for speaker 2 with 6 filter sections.

The performance of the passive speakers after optimization for the different acceptable insertion loss are compiled in appendix C. It can be seen that the results are not showing a clear trend and there are large variations between different optimizations. However, a general trend is that for lower insertion loss, the best performing filter has fewer sections than for higher insertion loss.

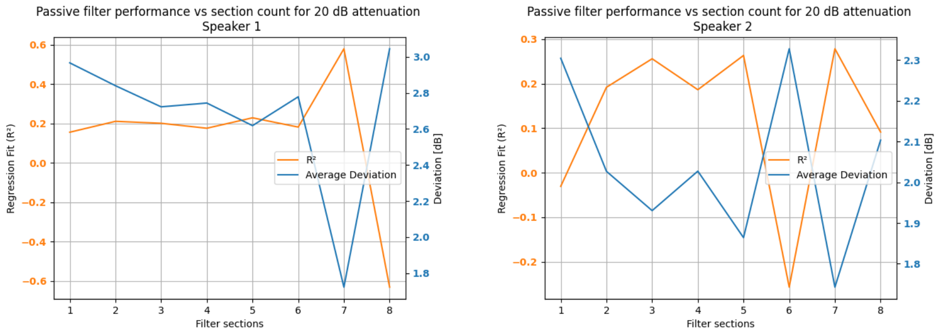

(a) Speaker 1

(b) Speaker 2

Figure 4.7: Performance of different amount of filter sections for 20 dB insertion loss for speaker 1 (a) and 2 (b).

The best performing filter for speaker 1 can be seen in figure 4.8a and has 8 filter sections with 15 dB acceptable insertion loss. It has a R^2 score of 0.62 and an average deviation of 1.68 dB. For speaker 2, the best performing filter can be seen in figure 4.8b and has 5 filter sections with 10 dB insertion loss. It has a R^2 score of 0.50 and an average deviation of 1.58 dB. In both figures the top graph shows the ideal filter for their respective speaker in orange and the filter's transfer function in blue. The bottom graph shows the original frequency response of the speaker in orange and the compensated response where the filter's frequency response is added to the speaker's is shown in blue. There is also a shaded region which indicates the 3 dB acceptable deviation from the baseline.

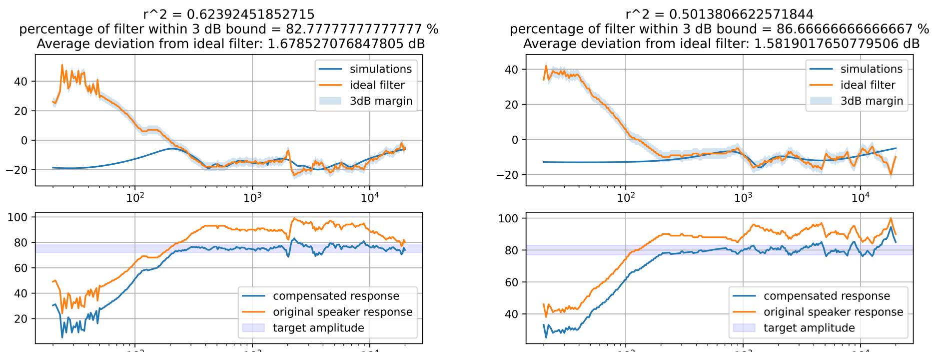

(a) Speaker 1

(b) Speaker 2

Figure 4.8: The best performing filters according to R^2 value for both speakers. Speaker 1 (a) has 8 filter sections and 15 dB insertion loss. Speaker 2 (b) has 5 filter sections and 10 dB insertion loss.

The values for the different components are, naturally, different for each iteration. But, in general, the values for the capacitances range between 100 nF to

100 μF, the values for the inductors range between 10 μ H and 10 mH and the resistors range between 10 Ω to 10 MΩ. There are some outliers in certain circuits as well, such as capacitors around 1 mF and inductors around 100 mH.

The power dissipation of the passive circuits depend mostly on the amount of insertion loss, as can be seen in table 4.2. The data is based on simulations in LTSpice where the mean power dissipation of the 8 different filters over the frequency range 20 Hz to 20 kHz for the different acceptable insertion loss modes was measured. The table shows that output power decreases whilst power dissipation increases with more insertion loss in the ideal filter goal function. Except for certain outliers such as 20 dB insertion loss for speaker 1 or 10 dB insertion loss for speaker 2.

Acceptable insertion loss	**Power dissipation Speaker 1 [mW]**	**Output power Speaker 1 [mW]**	**Power dissipation Speaker 2 [mW]**	**Output power Speaker 2 [mW]**
5 dB	24.68	56.65	22.73	59.95
10 dB	29.26	29.99	30.20	23.01
15 dB	29.36	16.53	25.58	8.02
20 dB	27.00	4.74	25.46	5.55
25 dB	40.77	7.23	31.58	2.91
30 dB	34.28	2.31	33.26	3.03

Table 4.2: Table showing the average power dissipation and output power for all 8 filters based on different acceptable insertion loss for both speakers.

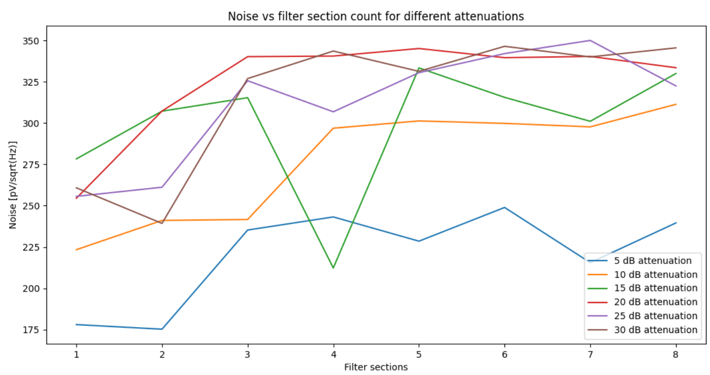

Figure 4.9: The average noise produced by the passive filters against the section count for different levels of acceptable insertion loss.

The noise introduced by the different passive filters were simulated and is shown in figure 4.9. It shows that the noise increases with more insertion loss, and somewhat with added filter sections. However, the noise is quite negligible in general.

4.2.3 Evaluation

Using passive components to compensate the non-linear frequency response of the speakers has some promise. The design is highly expandable due to the bridged T topology experiencing low circuit loading effects, so it is possible to cascade many filter sections. The best performing passive filters, see figure 4.8, managed to reduce the deviation from ideality by roughly 1 or 2 dB, from 3.52 dB to 1.67 dB and from 2.64 dB to 1.58 dB for speaker 1 and 2 respectively. Thus, passive filters can be used to flatten the frequency response somewhat.

However, the passive filters necessarily attenuate the signal and for example the filter in figure 4.8a has 15 dB insertion loss. This affects the sound level that the speaker can produce and this high of a reduction in volume is not ideal. Therefore, it would be better to try and keep insertion loss low. Additionally, the selectivity is fairly low, limiting which peaks that are possible to compensate.

While it is possible to design a passive filter with high selectivity, the Q value is often inherently tied to the attenuation the filter introduces. In our case we would ideally only want a few decibels of attenuation while maintaining a high Q. This is not possible with passive filters. There are topologies that could provide a very sharp notch, but the attenuation would be orders of magnitude larger than desired.

The power dissipation compiled in table 4.2 shows that, in general, the dissipation is fairly small, between 20-40 mW. This is reasonable since only passive components were used in the design, and they do not dissipate much compared to active circuits that actively consume power. The power dissipation increases while output power decreases for larger insertion loss in the goal function, which is reasonable since the input voltage was constant for all simulations. There are some outliers, but that seems reasonable due to the varying performance of the optimizer.

The noise for the different filters is also quite low and in the order of 100 $pV/\sqrt{Hz}$. This is due to thermal noise from passive components in general being quite low. The noise increases somewhat with both number of filter sections and insertion loss, which is reasonable since more passive components increases the noise. The larger the overall resistance of the circuit is responsible for the increase in insertion loss, and more resistance is also related to more noise.

The results generally show that the better the filter matches the frequency response of the ideal filter, the more insertion loss it introduces. However, the results do not show a very clear trend and some filters perform worse than no filter, as can be seen in appendix C. This large variation can be explained by the optimizer. It seemed to have a hard time finding the optimal value of all the components. This could be due to the use of particle swarm optimization, since this method is not guaranteed to find an optimal solution. The solution space might be too large, or the convergence might be too quick, and the solution ends up in a non-ideal local minimum. The least squares optimization is also only applied after the particle swarm, so it does not explore the whole space but rather the local minimum that the particle swarm round ended up in.

The optimizer was also very slow when trying to optimize the passive filters, which makes the design process a bit more difficult due to the long wait times. This is probably due to the large amount of variables, especially in the larger filters, since each filter section has six variable component values and the fact that several iterations of optimization was used for each filter.

Another factor that slowed down the optimizer was the large bounds of the different components, especially capacitors and inductors, since these upper bounds would result in unrealistic components. The optimized values for the passive filters lead to inductor values in the range of μH to 100 mH; capacitor values were in the range of 100 nF to 1 mF. Inductors of 100 mH size are fairly large and expensive and usually have a non-negligible parasitic resistance. The largest capacitors would also be very bulky and exhibit non-ideal characteristics. However, these outlier components could likely be removed and replaced with an open or closed circuit and then the optimizer could do another pass in order to tweak the other values with respect to the new circuit, which is why the large upper bounds were included in the method. It would be interesting to investigate the removal of unrealistic components and how they affect the overall performance as well as power consumption, noise and price.

The optimization problem could also be mitigated by providing a better starting point for the optimizer by doing some manual filter design before and letting the optimizer do some tweaking instead. This could shrink the solution space and guide the optimizer slightly, at the expense of a more difficult design process. Another way to approach this problem is to further investigate the optimizer parameters, as explained in section 3.2. The parameters could be tweaked to be better equipped at handling an optimization problem with this many variables.

The decrease in volume that is caused by the overall insertion loss of the signal can be mediated by adding a preamplifier before the filter to compensate for the insertion loss, but then the filter solution would no longer be fully passive. The extra source of noise, nonlinearity as well as the need for external power supply would negate some of the benefits with passive filters.

4.3 Active filter design

Active filters do not suffer from the same limitations that passive filters do, and unlike passive filters they can be constructed to have zero insertion loss. As such the overall volume of the signal does not need to be attenuated in order to flatten out the response. The same method of cascading filter sections can be used with active filters as well, with the added benefit that the active elements can act as buffers between sections, eliminating the loading effect. The problem that a very selective, high Q, filter usually has very high attenuation or gain still remains as with the passive circuits.

One way to solve this is to use the topology of a constant-Q graphic equalizer, as shown in figures 2.12 and 2.13. The gain-section is completely separate from the filter-section, meaning that each filter only needs to be concerned with its center frequency and bandwidth, while the gain or attenuation is controlled separately. This allows for sharp cuts and boosts while keeping the attenuation or gain low.

Multiple Feedback bandpass filters were selected for the filter section. This is due to their good performance and low sensitivity, as discussed in section 2.5. Another reason was that MFB has fewer passive components compared to e.g. VCVS.

4.3.1 Methodology

The main design problem to solve when it comes to a constant-Q graphic equalizer, is to decide on the number of filters, their center frequencies, and the amount of boost or cut each filter section should have.

In order to simplify this process a simple algorithm was written in python, which analyzes the frequency response of the speaker and, given a target amplitude baseline, generates a list of 16 frequencies sorted by importance. The algorithm repeatedly added the frequency with the highest deviation from the baseline to a list, and removed all frequencies within a 1/3 of an octave band of that frequency from consideration before selecting the next frequency. As such, no two frequencies on the list could be closer than 1/3 of an octave apart, which is the bandwidth of a bandpass filter with a Q-value of 4. Frequencies with a deviation below 2 dB were deemed insignificant and filtered out.

Based on the selected frequencies, 16 different schematics were designed, each containing a progressively increasing number of filter sections — from 1 to 16. These were then run through the optimizer in order to set the values of the potentiometers. This allowed for evaluation of how the number of filters impacts the overall performance and accuracy of the equalization, which served as a basis for selecting the number of filter stages.

Before realizing the circuit, some manual adjustments were made of the chosen filter, and the power consumption and transient behavior of the adjusted circuit were simulated. Each filter section originally consisted of a multiple feedback bandpass filter with a Q of four. By adjusting the Q-values, amount of attenuation or gain, and in some cases slightly altering the center frequencies of individual filter sections, the performance of the filter could be increased further.

4.3.2 Simulation results

The result of running the frequency selection algorithm on the frequency responses of our two speakers can be seen in figure 4.10, and the resulting frequencies are listed in table 4.3 in order of importance.

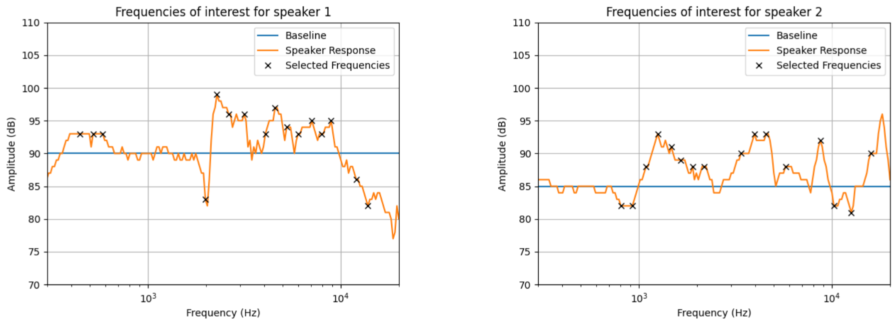

(a) Speaker 1 **(b)** Speaker 2

Figure 4.10: Frequencies of interest for speaker 1 and 2.

Section	Speaker 1 (Hz)	Speaker 2 (Hz)
1	2279	1250
2	13819	3969
3	1984	4559
4	4559	8706
5	2618	1469
6	3150	3376
7	7071	15874
8	8909	1649
9	5237	12599
10	12030	806
11	442	925
12	519	1088
13	583	1895
14	4061	2176
15	6015	5744
16	7937	10234

Table 4.3: Frequencies of interest for each section and speaker.

The corresponding performance for each configuration, ranging from one to sixteen filter sections, is presented in figure 4.11. As expected, performance improves with more filter sections, but the gains in performance diminish beyond a certain point since each additional filter section will add to the cost, size and power consumption of the circuit. The full 16-section filter, shown in figure 4.12 serves as a reference for the best achievable correction.

A five-section configuration was chosen for speaker 1 as a compromise between performance and complexity. By manually adjusting the five section equalizer, the performance was further improved, as seen in the comparison in figure 4.13. The

Q of section 1 and 5 was increased to 8 to provide narrower corrections, and the potentiometer values were adjusted slightly for most bands. This tuning step led to a noticeably closer match to the target response.

The step response of the filter with 5 sections can be seen in figure 4.14. The step had a rise time of 1 ns and the response had a rise time of 400 ns. The ringing died out after roughly 4 ms and the overshoot was roughly 0.11 V. This suggests that the filter is stable.

The result from the power consumption simulation showed a constant consumption of 115 mW independently of both the frequency and amplitude of the input signal.

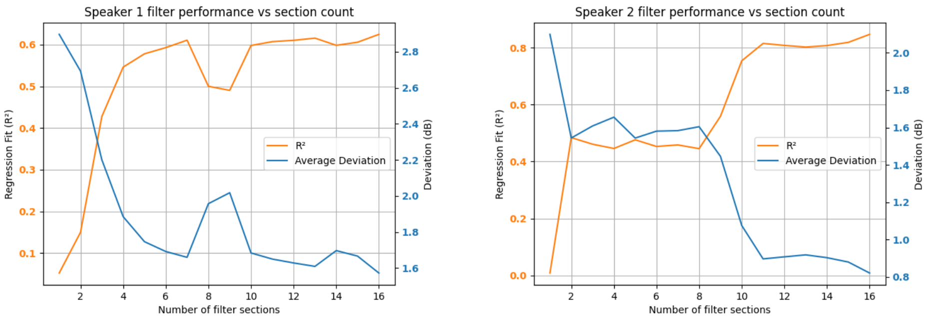

(a) Speaker 1 **(b)** Speaker 2

Figure 4.11: Filter performance for different numbers of filter sections, for speaker 1 and 2.

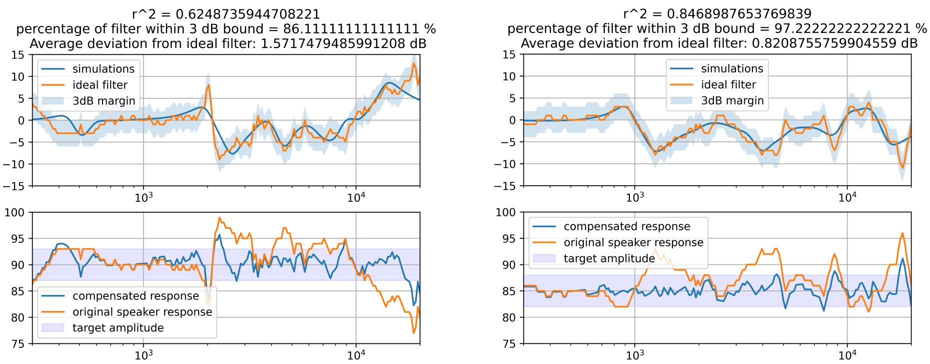

(a) Speaker 1 **(b)** Speaker 2

Figure 4.12: Frequency response with 16 filter sections, with no manual adjustments.

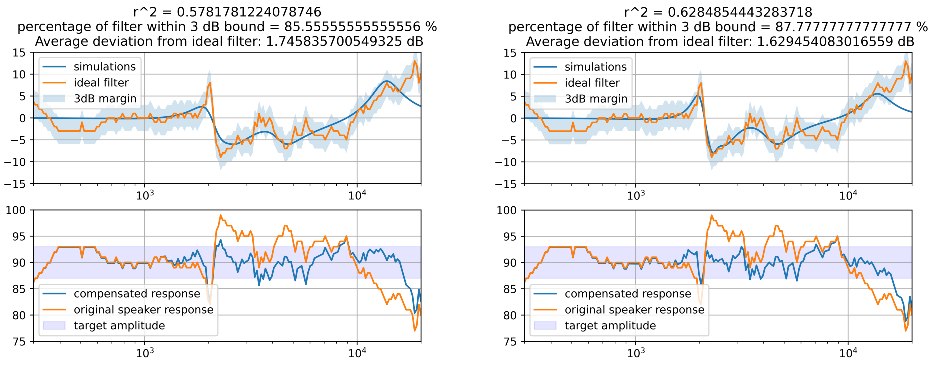

(a) Before

(b) After

Figure 4.13: Frequency response of speaker 1 before and after manual adjustments.

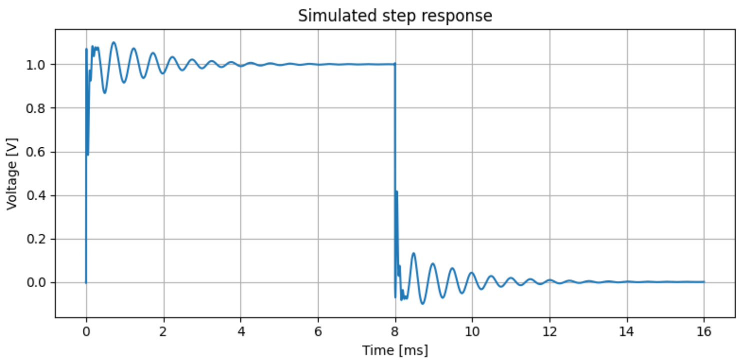

Figure 4.14: Step response of the active filter with 5 filter sections. The step had a rise time of 1 ns and the filter had a rise time of 400 ns and the ringing stops after roughly 4 ms.

4.3.3 Evaluation

From these results one can conclude that the constant-Q graphic equalizer structure is a great fit for this particular problem, where a varying amount of attenuation or boost needs to be applied to different parts of the frequency spectrum. The design is very flexible, essentially allowing any number of filter sections, each with varying Q-values, center frequencies and gain or attenuation. Importantly, the constant-Q property ensures that the selectivity of each section remains independent of how much cut or boost is applied — a key advantage over other active filter topologies.

The flatness of the frequency response with the filter applied is of course dependent on the original frequency response of the speaker. In our case, speaker 1 has much sharper transitions and more variation in amplitude overall, requiring sharper and more selective filter sections than for speaker 2. This is clearly shown in figure 4.12, where the resulting response for speaker 2 is visibly flatter than that for speaker 1, even though all filters use the default Q-value of four. If similar manual adjustments, as presented in figure 4.13, were to be performed for both speakers with 16 filter sections, the performance could most likely be improved further.

From the step response in figure 4.14 it is clear that the filter is stable, although there is some ringing present. The rise time of 400 ns, primarily determined by the opamp's slew rate and the circuit's capacitances, is much faster than required for audio frequencies and therefore poses no issue in this application. The power consumption was also stable at 115 mW regardless of frequency or input signal amplitude, which makes the performance more predictable.

The use of the frequency selection algorithm significantly reduces the workload on the optimizer in the active filter design compared to the passive approach. Since the key center frequencies have already been selected and the BPF have been designed, the optimizer only needs to worry about choosing the amount of cut or boost of each filter by changing the values of the potentiometers, dramatically shrinking the solution space and resulting in a quicker convergence.

Significant effort can be spent on selecting the optimal center frequencies for the filter sections, a process which we attempted to simplify using a python algorithm. While the algorithm successfully identifies problematic frequencies, its simplicity leads to some suboptimal results. As shown in figure 4.11b the performance plateaus between three and eight filter sections, suggesting that the corresponding frequency bands may not be the most critical. This indicates that the algorithm's ranking does not always reflect the true impact of each frequency on overall correction. This could be due to the fact that the algorithm does not take the surrounding frequencies into account when selecting a peak. Thus, some peaks might be detected on a sharp edge and a filter placed with a bandwidth wide enough to negatively affect other frequencies.

The ideal algorithm would select frequencies in an order that yields the largest increase in performance with the first filter sections, which is not what we see in neither figure 4.11a nor 4.11b. By reordering the filters for the filter section evaluation of speaker 2, by placing sections 9 to 11 earlier in the sequence, performance at lower section counts was noticeably improved, as demonstrated in figure 4.15, proving that the algorithm used is not ideal.

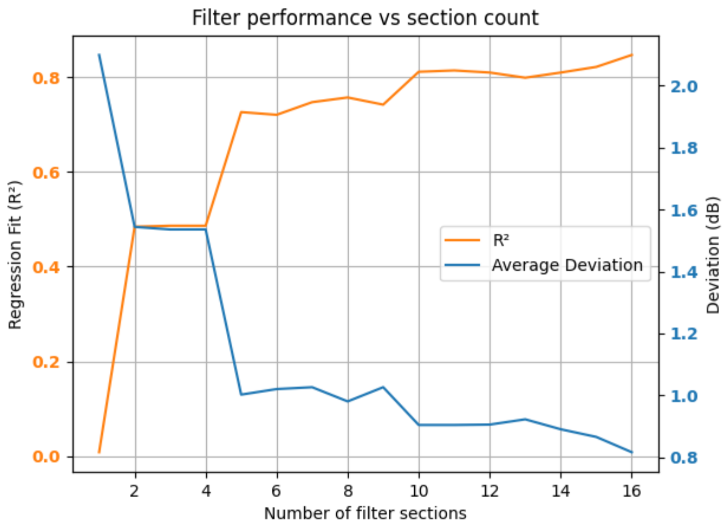

Figure 4.15: Filter performance for different numbers of filter sections for speaker 2, after reordering the frequencies.

There might be several ways to improve the frequency selection algorithm which would be interesting to investigate. It would be interesting to make the frequency selection more iterative by choosing a frequency, simulating the behavior and possibly choosing another Q-value and then choosing the next frequency based on the compensated frequency response. This might help the frequency selection order the different peaks more accurately based on importance since some peaks might be able to be compensated by a nearby peak. Another possible solution might be to take the rate of change in the frequency response into account by weighting the deviation of the peak with the derivative or second derivative of the peak or nearby points. There are other ideas which could be interesting to investigate such as e.g. machine learning.

4.4 Filter Realization

The best performing active circuits were notably better than the passive circuits in flattening the frequency response of the speakers. Thus, we decided to realize the active topology described in section 4.3 with five filter sections. Some modifications to the design had to be made in order to integrate the filter with the existing system used for testing of the prototype. The input signal was reduced by 6 dB to comply with the signal requirements and avoid clipping. A bypass switch was also added to the design in order to simplify the testing process.

A prototype PCB was designed using KiCad, the blueprint of which can be seen in figure 4.16, and the final components used are listed in table 4.4. Connectors were included to provide power and balanced input and output to the board. The passive components were chosen in the 0402 package size, and the RC4558 dual op-amp was selected for its low noise and solid performance in the audio frequency range. The components and connectors were soldered to the board by the authors. The dimensions of the board was 23 by 41 mm, and the final PCB can be seen in figure 4.17.

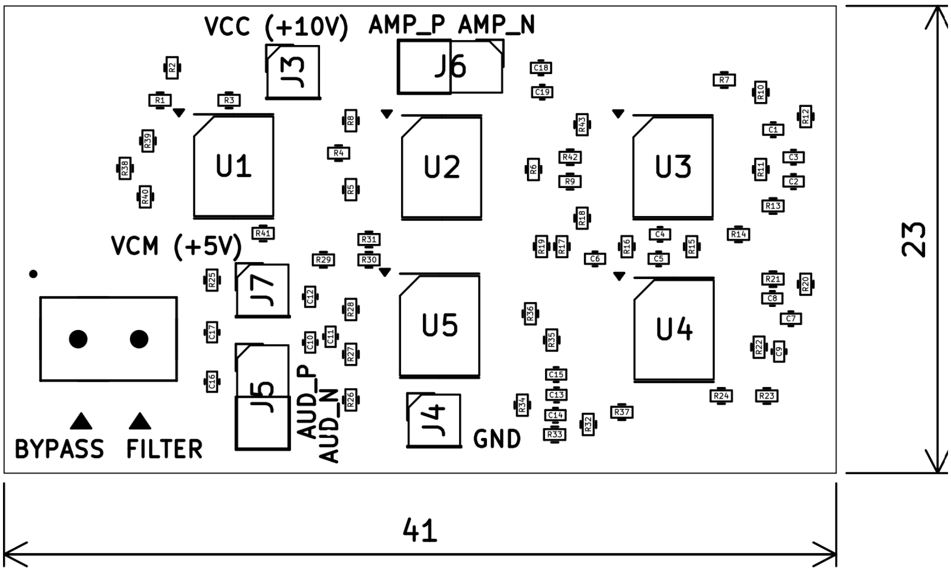

Figure 4.16: The blueprint of the designed PCB where U-sites are for opamps, J-sites are holes for pins and a site for a bypass switch. The side lengths are expressed in millimeters.

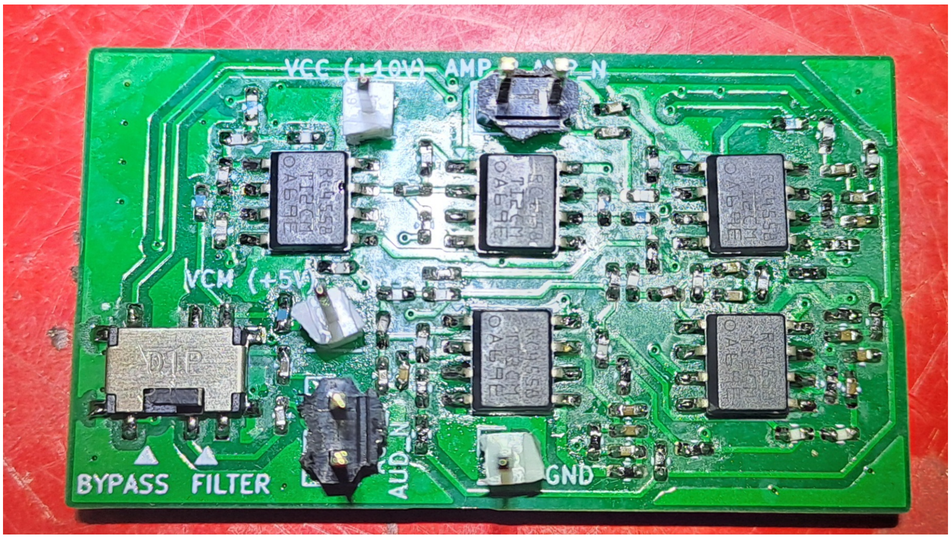

Figure 4.17: A photo of the filter PCB.

Components	**Qty**	**Value / Part Number**
C1, C6, C9, C12, C15	5	10 μF
C2, C3, C7, C8, C13, C14	6	27 nF
C4, C5	2	4.7 nF
C10, C11	2	12 nF
C16, C17, C18, C19	4	100 nF
R1, R3, R4, R5, R6, R10, R17, R23, R29, R30, R35, R42, R43	13	10 kΩ
R2, R9	2	10Ω
R7, R24	2	100 kΩ
R8	1	20 kΩ
R11	1	41.2 kΩ
R12	1	20.5 kΩ
R13	1	162Ω
R14	1	9.76 kΩ
R15	1	316Ω
R16	1	19.6 kΩ
R18	1	5 kΩ
R19	1	95.3 kΩ
R20	1	23.7 kΩ
R21	1	187Ω
R22	1	47.5 kΩ
R25, R38, R40	3	120 kΩ
R26	1	11.5 kΩ
R27	1	374Ω
R28	1	23.2 kΩ
R31	1	90.9 kΩ
R32	1	13.7 kΩ
R33	1	191Ω
R34	1	26.7 kΩ
R36	1	15 kΩ
R37	1	84.5 kΩ
R39, R41	2	30 kΩ
J3, J4, J7	3	Conn_01x01_Pin
J5, J6	2	Conn_01x02_Pin
S2	1	MSS6-V-T_R dipswitch
U1, U2, U3, U4, U5	5	RC4558 opamp

Table 4.4: Component list with quantities and values for the PCB design.

The frequency response of the realized filter can be seen in figure 4.18, and the performance is very similar to the simulations. The high frequency response differs slightly, but the main difference is in the low frequency response where the attenuation is much larger for the lowest frequencies.

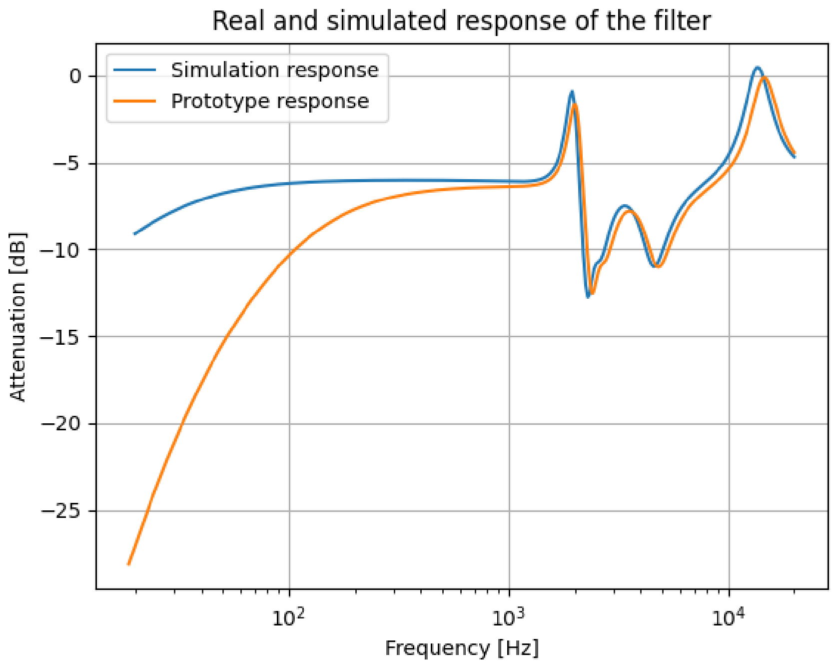

Figure 4.18: The frequency response of the realized prototype compared to the simulated response of the filter.

The filter was added to speaker 1 in order to investigate how well it actually performed in flattening the frequency response. The results can be seen in figure 4.19, where it can be seen that the filter managed to flatten the frequency response from an average deviation of 4 dB to 1.67 dB. Compared to the simulated frequency response of speaker on in figure 4.13b the filtered frequency response is not quite as good, and the measured response had more peaks that exceeded the 3 dB limit. However, in terms of the average deviation, the simulation predicted an average deviation of 1.63 dB whereas the real measurement was 1.67 dB.

The noise was investigated and measurements of THD+N for the speaker with and without the filter connected can be seen in figure 4.20. It can be seen that the noise is not changed significantly overall and remains at the same level. There are two main differences between figure 4.20a and 4.20b, namely at around 2 kHz and around 16 kHz. At 2 kHz in the filtered measurement, see figure 4.20b, the filter introduces a peak in the noise. However, at 16 kHz the filtered signal seems to reduce the noise at higher frequencies.

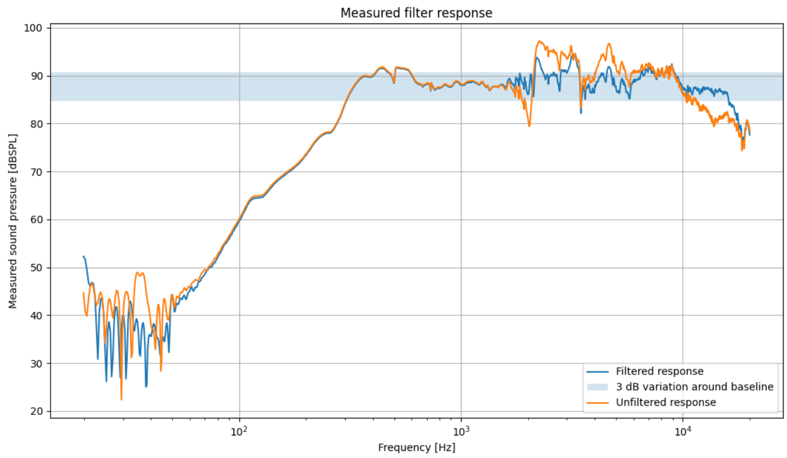

Figure 4.19: The frequency response of speaker 1 with and without the filter. The average deviation was 4.0 dB without the filter and 1.67 dB with the filter.

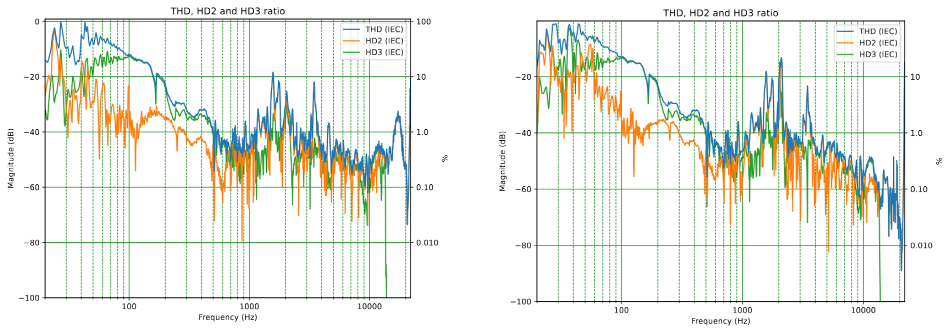

(a) Unfiltered **(b)** Filtered

Figure 4.20: Total harmonic distortion and noise of speaker 1 without the realized filter, (a), and with the filter, (b).

Temperature measurements were also conducted on the realized filter. The results can be seen in table 4.5. The filter exhibited an increase in temperature of approximately 9 to 12 $°C$. The magnitude of the input signal did not significantly affect the temperature, as the results remained similar across different levels. Interestingly, some stimuli resulted in lower temperatures at higher amplitudes. Among the tested signals, white noise caused the highest temperature increase, while the 500 Hz tone caused the least. The 2 kHz tone resulted in the highest temperature among the pure tones. Additionally, the power consumption was measured at 180 mW during operation, regardless of frequency or input signal amplitude.

Stimuli	Max temperature [$°C$]	Temperature increase [$°C$]
No Stimuli	25.3	3.1
White noise (2.8 V)	33.9	11.7
White noise (4 V)	33.0	10.8
500 Hz (2.8 V)	31.6	9.4
500 Hz (4 V)	31.0	8.8
2 kHz (2.8 V)	33.2	11.0
2 kHz (4 V)	32.3	10.1
12.5 kHz (2.8 V)	31.7	9.5
12.5 kHz (4 V)	31.7	9.5
16 kHz (2.8 V)	32.7	10.5
16 kHz (4 V)	32.7	10.5

Table 4.5: Table showing the maximum temperature measured on the filter during different input signal magnitudes. Room temperature was 22.2 $°C$.

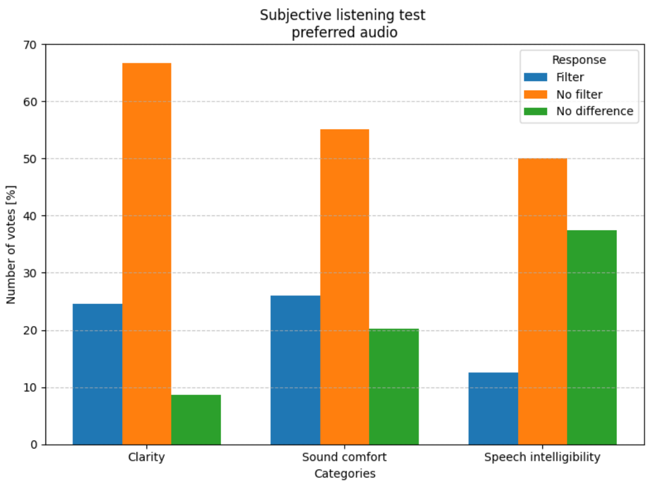

Figure 4.21: The responses from the subjective listening test. The bars show the frequency of answers for each respective audio for the different categories.

The subjective listening test was conducted with 14 participants, as described in section 3.1.2. The results from this test can be seen in figure 4.21. It is clear from this response that the unfiltered audio was preferred in most cases. However, only 54 % of the participants correctly identified when there were no difference between the 2 different audio.

The participants had the ability to write down their thoughts on the different audio tracks. Six participants expressed that there was a noticeable difference in volume between the filtered and unfiltered audio on 3 audio tracks, track 3-5, and that it therefore was hard to fairly judge which sounded better. Two participants seemed to notice higher frequencies in the filtered signal as well.

4.4.1 Evaluation

Judging by figure 4.18, the prototype has a very similar performance compared to the simulations. However, when comparing the frequency response of the speaker in figure 4.19 and the simulated frequency response in figure 4.13b, the simulated compensation is clearly better. This seems reasonable since the model used to predict the frequency compensation does not include a model of the speaker, and assumes a fully linear system. This simplistic model can be used to somewhat accurately predict the effect of filters on speaker performance, but if more accuracy is needed, the speaker could be modelled electrically.

Another possible reason for the discrepancy between the simulations and measurements on the speaker response is that the audio signal is routed through existing hardware before arriving at the filter. The output of the filter is also fed through a power amplifier before reaching the speaker. Neither the amplifier nor the other hardware is included in the simulation model.

The PCB was fairly small at 23 by 41 mm, see figure 4.16, but could be decreased further by using smaller components. If the filter is to be included in a product, it could also be integrated into the existing PCB. Assuming this is the case, the cost of the PCB is not taken into account, leaving only the total cost of the components at roughly 0.4 - 0.5 $ USD, which is fairly low. In relation to a DSP, with a size of 4 by 4 mm and a cost of roughly 10 $ USD [14], the analog solution is cheaper, but larger. However, since a DSP needs more input signals and different supplies than the analog filter it needs more peripheral circuitry to function, which would have to be taken into account for a proper size comparison.

It can be seen from figure 4.20 that the filter did not introduce much noise, since the noise floor seems to be at the same level. However, there is a new peak around 2 kHz that can be seen in the filtered measurements, see figure 4.20b. This seems reasonable since this is the frequency where the corresponding BPF has the highest Q-value, which also means the highest added noise. Another interesting observation is the reduction of noise for the highest frequencies in figure 4.20b. It is unclear exactly why this is, but since the speaker is non-linear, it is hard to predict how it will act and difficult to describe without a good model of the speaker.

The power consumption of 180 mW is comparable to, albeit slightly higher than, the 115 mW predicted by the simulations. This discrepancy could also be explained by the simplicity of the model, where all the passive components were completely ideal.

The temperature measurements of the filter were not ideal and the increase in temperature from room temperature was quite large in all stimuli cases. One interesting aspect is that the power consumption was constant across different frequencies while the temperature increase was not. For example the 2 kHz tone lead to 10 $°C$ increase whilst the 500 Hz tone increased the temperature by 9 $°C$.

The power consumption could increase the temperature by a constant amount for all frequencies, but the frequencies where the temperature increased the most correspond to the frequencies where the filter boosts the signal. Therefore, the extra temperature increase could be explained by the added power dissipation in other passive components due to boosting the signals. The white noise performed the worst leading to an increase in temperature of 11.7 $°C$ at 1 W, which is to be expected since white noise contains all frequencies.

When increasing the input signal amplitude, the temperature was slightly lower compared to the normal input signal of 2.8 V. This is interesting and could be due to several factors. Either the opamps are more suited and optimized to work with that level of voltage swing or there were some measurement error. It could be the case that the temperature had not stabilized after 2 minutes, thus leading to a lower temperature measurement in some cases.

If temperature performance is critical in a certain application, there are ways to mitigate the high temperature increase of the circuit. One large factor is the design of the PCB where different substrates and possibly heat sinks can be incorporated. The choice of components is also a possible area of improvement. The opamp used, RC4558, has fairly high thermal metrics and there are other options which could radiate less heat [39]. The same can be said for the passive components.

The subjective listening tests seem to imply that the unfiltered signal resulted in better audio, as can be seen in figure 4.21. However, there are several factors which makes this result questionable. One major aspect is the difference in volume between the filtered and unfiltered signal. Louder volume is usually associated with better sound quality which could explain the skewed results. If the results from the test is compiled without the three audio tracks where the volume was different, the results still favor the unfiltered signal, however not by as large a margin as before. The filtered signal was preferred by 27.5 % of participants and the unfiltered signal was preferred by 52.2 % of participants, when merging the three different categories together.

Another factor to take into account is the fairly low, 54 %, number of people who correctly identified identical audio signals. According to [37], the recommendation states that participants who can not identify the hidden reference, the same audio, should be excluded. During the test only three participants correctly identified both hidden references and their score was 31 % preferred the filtered signal and 44 % preferred the unfiltered signal.

When taking the tracks with no volume difference and only using the participants who correctly identified the reference, the results instead become 30 % preferred the filtered audio and 50 % preferred the unfiltered audio.

Taking these factors into account, it is hard to fully trust the results of the listening tests. Even though both audios were calibrated to have the same volume, due to technical issues there was a difference in volume for some audio tracks which should have been accounted for. The low accuracy of identifying the hidden reference also suggests that the difference between the two signals was not large. It would be necessary to test all the different audio files in order to determine the correct volume. However, the focus of this thesis was to flatten the frequency response of the speaker, not to make it sound better. It is possible that the unfiltered signal is more appealing.

Chapter 5

Conclusion

There are, naturally, advantages and disadvantages with both passive and active analog filters. In this thesis we have specifically studied analog filters in the context of compensating frequency distortions in loudspeakers, and for this use-case the constant-Q graphic equalizer structure proved to be the most promising.

While passive filters were able to reduce the average deviation from 3.52 dB to 1.68 dB for speaker 1 and from 2.64 dB to 1.58 dB for speaker 2, and the problematic effects of circuit loading could be mitigated by using the constant-resistance structure of the bridged-T, they still come with significant limitations. The combination of low-noise and not needing an external power supply are the two biggest reasons to choose a passive filter solution. However, the inevitable introduction of insertion loss, and the inability to achieve a high selectivity as the amount of attenuation is decreased are major drawbacks that inhibit the use of passive filters in this context. The insertion loss could be removed by using active elements to introduce gain, but then the filter would require an external power supply, and would no longer be low-noise, resulting in a hybrid solution that combines the low selectivity of passive filters with the disadvantages of active filters.

In contrast, the constant-Q graphic equalizer allowed for independent control of gain, bandwidth, and center frequency for each band, enabling more precise compensation of the frequency response. Even the narrowest peaks could be properly compensated, limited mainly by the higher component tolerances needed to realize bandpass filters with higher Q-values. Additionally, insertion loss is avoided due to the added gain of the active elements, at the expense of a higher power consumption and more noise compared to the passive filters.

The number of filter sections is a key design trade-off. More sections allow for finer compensation and a flatter frequency response, but also increased complexity, cost, and power consumption. With 16 filter sections the average deviation could be reduced to 1.57 dB for speaker 1 and 0.82 dB for speaker 2, which is significantly better than what was achieved with the passive filters. These numbers can most likely be reduced further, as manual adjustment of just five filter sections brought the average deviation down from 1.75 dB to 1.63 dB.

The optimization based method used gave mixed results. Especially for the passive filters, where the optimization did not show a clear trend in performance, suggesting that the method is more suited for problems with fewer variables and a smaller solution space, as was the case in the active filters. For the active filter design the optimizer could quickly choose values to match a generic frequency response, given that the frequency bands chosen by the frequency selection algorithm were adequate. The simplicity of this algorithm was the main limiting factor in the design process for active filters.

In conclusion, a constant-Q graphic equalizer can be used to compensate frequency distortions in loudspeakers. The more filter sections added, the flatter the compensated response can become, given that the center frequencies and bandwidths are chosen correctly. It is cheap option that only needs a few square millimeters of space and is not limited by processing power. Depending on the application, active analog filters could definitely be a viable solution, but the drawbacks need to be carefully considered on a case-by-case basis, based on the limitations and constraints of the specific speaker system in question.

5.1 Future work

Several potential directions for future work have been identified based on the findings and limitations of this thesis. Due to time constraints, only a limited number of topologies could be properly investigated. A more exhaustive study of other passive and active filter configurations could reveal alternatives with improved performance, greater stability, or lower component sensitivity. For instance, asymmetric EQ structures may be advantageous in systems with limited headroom, and passive LC filters could help mitigate insertion loss. Other bandpass filter topologies are also worth considering, such as the state variable bandpass filter, which is known for its low sensitivities and high stability.

Further evaluation of the filters could also be of interest. A sensitivity analysis could help determine the acceptable tolerances for different components, potentially reducing overall costs by allowing lower-precision components to be used where appropriate. Additionally, more precise thermal measurements could be useful, as temperature increases in the filter may pose limitations in the system.

The effect of phase distortion in speaker systems can also be investigated further, particularly in relation to group delay and the spatial positioning of loudspeakers relative to the listener. Other psychoacoustic effects could also be taken into account in order to improve the listening experience, for example by adjusting the filter goal function to better align with the frequency sensitivity of human hearing.

Regarding the optimization process in the passive case, future work could involve finding starting conditions and values of the optimization parameters that are better suited for the specific circuit topology at hand. Additionally, the frequency selection algorithm could be made more sophisticated, e.g. by opting for a more iterative approach where one center frequency is selected and compensated for at a time.

Finally, a thorough comparison with digital filter implementations would help contextualize the relevance of analog solutions. While this thesis focused on analog designs, many of the same goals could be achieved in the digital domain. Future work could investigate performance, power consumption, noise, and implementation cost trade-offs between analog and digital solutions for the same loudspeaker system.

References

[1] R. Izhaki, *Mixing Audio, Concepts, Practices, and Tools, 4th Edition*, ch. 12. Routledge, 2024.

[2] Y. Suzuki and H. Takeshima, "Equal-loudness-level contours for pure tones," *The Journal of the Acoustical Society of America*, vol. 116, pp. 918–933, 08 2004.

[3] M. A. Siddiqi, *Continuous Time Active Analog Filters.*, ch. 1, 9, 12. Cambridge University Press, 2021.

[4] W.-K. Chen, *Passive, Active, and Digital Filters*, ch. 1,2,4. Crc Press, 2009.

[5] H. Kimball, "Motion picture sound engineering," 1938.

[6] D. Berndt and S. Roy, "Inductor simulation using a single unity gain amplifier," *IEEE Journal of Solid-State Circuits*, vol. 4, no. 3, pp. 161–162, 1969.

[7] R. Elliot, 2001. https://sound-au.com/project75.htm.

[8] J. G. Švec and S. Granqvist, "Tutorial and guidelines on measurement of sound pressure level in voice and speech," *Journal of Speech, Language, and Hearing Research*, vol. 61, no. 3, pp. 441–461, 2018.

[9] D. M. Huber and E. Caballero, *Modern Recording Techinques, 10th Edition*, ch. 2. Routledge, 2024.

[10] R. V. Cox, S. F. d. C. Neto, C. Lamblin, and M. H. Sherif, "Itu-t coders for wideband, superwideband, and fullband speech communication [series editorial]," *IEEE Communications Magazine*, pp. 106–109, october 2009.

[11] V. Dickason, *Loudspeaker Design Cookbook, 7th*, ch. 5. Audio Amateur, Incorporated; 6th edition, 2005.

[12] Audio Precision. https://www.ap.com/blog/the-big-six-audio-measurements.

[13] P. Horowitz and W. Hill, *The Art of Electronics, 3rd Edition*, ch. 1, 5, 6. Cambridge University Press, 2015.

[14] Analog Devices, 2024. `https://www.analog.com/media/en/technical-documentation/data-sheets/ADAU1761.pdf`.

[15] P. Sherz and S. Monk, *Practical Electronics for Inventors, 4th Edition*, ch. 3. McGraw-Hill Education, 2016.

[16] P. Horowitz and W. Hill, *The Art of Electronics, The X-Chapters*, ch. 1x, 4x. Cambridge University Press, 2020.

[17] K. Warwick, *An introduction to control systems (second edition)*, ch. 4. World Scientific Publishing Co., 1996.

[18] A. I. Zverev, *Handbook of Filter Synthesis*, ch. 1. John Wiley & Sons, Inc, 1967.

[19] M. Ghausi, R. Schaumann, and K. R. Laker, *Design of Analog Filters*, ch. 1,2,4. Prentice-Hall Inc, 1990.

[20] H. Zumbahlen, *Linear Circuit Design Handbook*, ch. 8. Newnes/Elsevier, 2008.

[21] C. van Ryswyk, "Filters for equalization: active or passive?," *journal of the audio engineering society*, october 1976.

[22] V. Nostrand, *Motion Picture Sound Engineering*, ch. 16. D. Van Nostrand Company, inc., 1938.

[23] D. A. Bohn, "Operator adjustable equalizers: An overview," in *Audio engineering society conference: 6th international conference: Sound reinforcement*, Audio Engineering Society, 1988.

[24] B. F. R., "Attenuation equalizers," *journal of the audio engineering society*, vol. 1, pp. 125–136, january 1953.

[25] M. A. Siddiqi, *Continuous Time Active Analog Filters*, ch. 1, 10, 12. CRC press, 2020.

[26] T. Deliyannis, "High-q factor circuit with reduced sensitivity," *Electronics Letters*, vol. 4, no. 26, pp. 577–579, 1968.

[27] S. Micheal, 2006. https://www.ti.com/lit/an/sboa114/sboa114.pdf?ts=1746452131270.

[28] H. H. k., "equalization simplified," *journal of the audio engineering society*, vol. 22, pp. 700–703, november 1974.

[29] E. James, "Simple tone control circuit: Bass and treble, cut and lift," *Wireless World, Feb, pg*, vol. 48, no. 50, pp. 156–7, 1949.

[30] R. Elliot, 2021. https://sound-au.com/articles/eq.htm.

[31] K. J. Gundry, "Adjustable equalizers useable in audio spectrum," U.S. Patent 3921104A, 1973.

[32] D. A. Bohn, "Constant-q graphic equalizers," *Journal of the Audio Engineering Society*, vol. 34, no. 9, pp. 611–626, 1986.

[33] G. C. Temes and J. W. LaPatra, *Introduction to synthesis and design*, ch. 1,11. McGraw-Hill, 1977.

[34] B. Adams, "Python ltspice opt," 2023. https://github.com/radams2000/Python_LTspice_opt/tree/main.

[35] H. P. Gavin, “The levenberg-marquardt algorithm for nonlinear least squares curve-fitting problems,” *Department of Civil and Environmental Engineering Duke University August*, vol. 3, 2019.

[36] J. Kennedy and R. Eberhart, “Particle swarm optimization,” in *Proceedings of ICNN'95-international conference on neural networks*, vol. 4, pp. 1942–1948, ieee, 1995.

[37] B. Series, “Method for the subjective assessment of intermediate quality level of audio systems,” *International Telecommunication Union Radiocommunication Assembly*, vol. 2, 2014.

[38] Lester James V. Miranda et al., “Pyswarms,” 2017. `https://github.com/ljvmiranda921/pyswarms`.

[39] Texas Instruments, 2024. `https://www.ti.com/lit/ds/symlink/rc4558.pdf?ts=1746514641578`.

Appendix A - Listening test questionnaire

Subjective Listening test

Instruktioner: You will listen to 7 audio clips. For each audio clip, compare Speaker 1 and Speaker 2 and check which performed the best within these 3 aspects:

- Clarity - How clear was the sound?
- Speech intelligibility - How easy was it to perceive the words?
- Sound comfort - Which speaker was the most comfortable to listen to?

Check the box of the alternative (Speaker 1 or Speaker 2) that you perceived to perform best for each aspect. Alternatively, you can check "No difference" if you can't distinguish anything between the 2 speakers.

Sound 1

Clarity:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Speech intelligibility:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Sound comfort:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Comments:			

Sound 2

Clarity:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Speech intelligibility:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Sound comfort:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Comments:			

Sound 3

Clarity:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Speech intelligibility:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Sound comfort:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Comments:			

Sound 4

Clarity:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Sound comfort:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Comments:			

Sound 5

Clarity:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Sound comfort:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Comments:			

Sound 6

Clarity:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Sound comfort:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Comments:			

Sound 7

Clarity:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Sound comfort:	Speaker 1 [ ]	Speaker 2 [ ]	No difference [ ]
Comments:			

Appendix B

Appendix B - Listening test protocol

Subjective Listening test protocol

Setup:

- Mount the speaker on the wall
- Supply the filter and connect all the cords
- Plug in the audio amplifier to a computer
- Ensure sound is audible when played from the computer
- Ensure sound level is the same when the sound is filtered and when it's bypassed
- Setup chairs so that they are facing away from the speaker

Introduction:

- Say welcome and thank them for their participation
- Go through the questionnaire and explain what aspects they should listen to and what the options are
- Play a sound and check that the volume is loud enough for everyone

Procedure:
Remember to switch on/off bypass with regard to table below.

Audio file and speaker configuration

Audio file	Speaker 1	Speaker 2
Talk 1	Filtered	Bypassed
Talk 2	Bypassed	Bypassed
Talk 3	Filtered	Bypassed
Ballad 1	Bypassed	Filtered
Jazz 1	Filtered	Bypassed
Pop 1	Bypassed	Filtered
Jazz 2	Filtered	Filtered

Appendix C

Appendix C - Passive filter performance

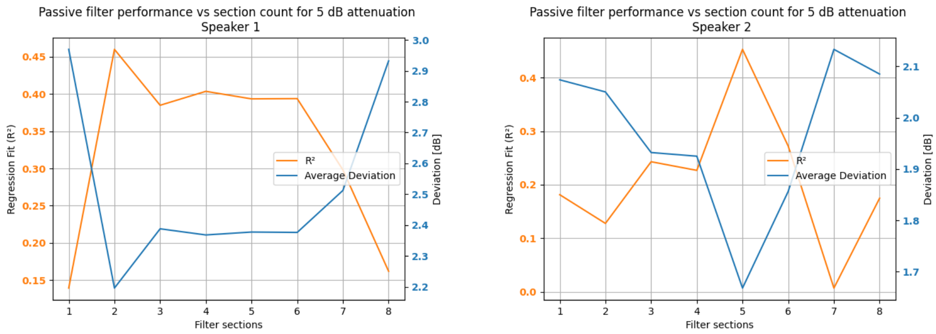

(a) Speaker 1

(b) Speaker 2

Figure C.1: Performance of different amount of filter sections for 5 dB insertion loss for speaker 1 (a) and 2 (b).

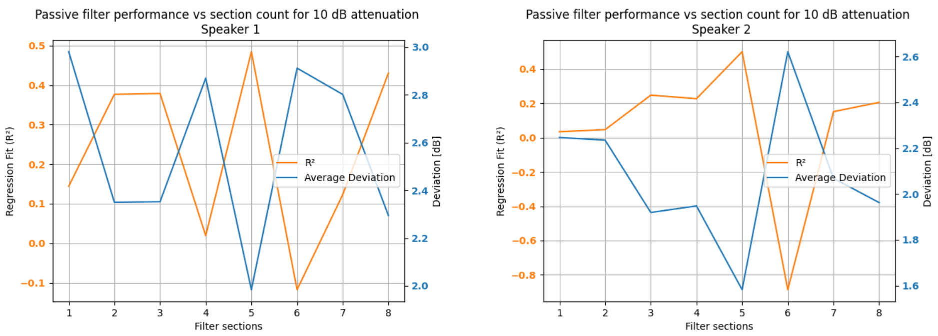

(a) Speaker 1

(b) Speaker 2

Figure C.2: Performance of different amount of filter sections for 10 dB insertion loss for speaker 1 (a) and 2 (b).

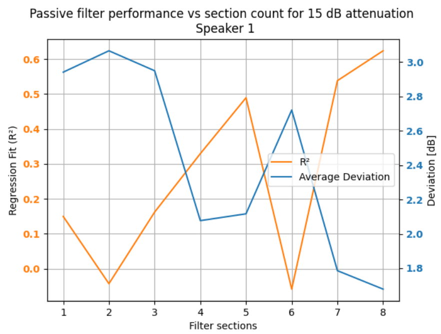

(a) Speaker 1

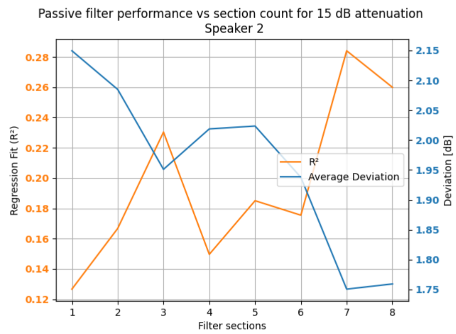

(b) Speaker 2

Figure C.3: Performance of different amount of filter sections for 15 dB insertion loss for speaker 1 (a) and 2 (b).

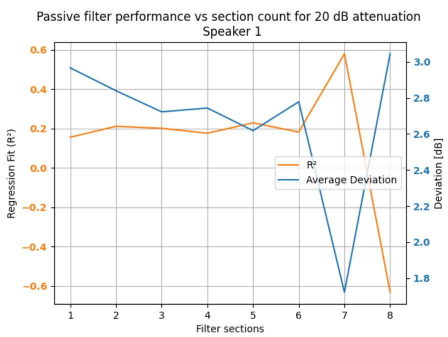

(a) Speaker 1

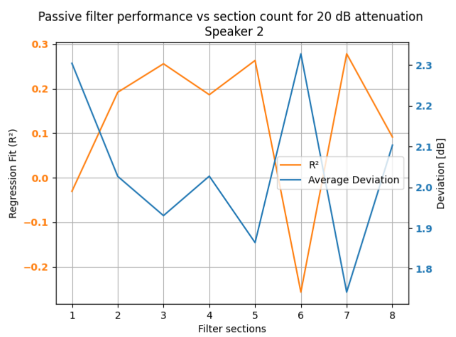

(b) Speaker 2

Figure C.4: Performance of different amount of filter sections for 20 dB insertion loss for speaker 1 (a) and 2 (b).

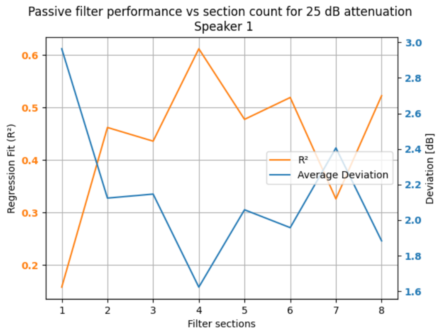

(a) Speaker 1

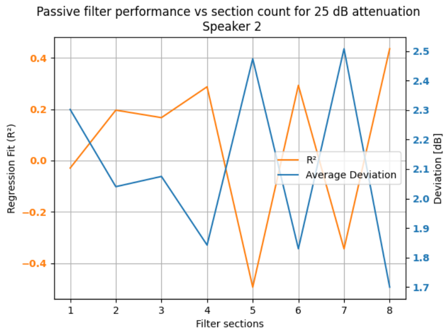

(b) Speaker 2

Figure C.5: Performance of different amount of filter sections for 25 dB insertion loss for speaker 1 (a) and 2 (b).

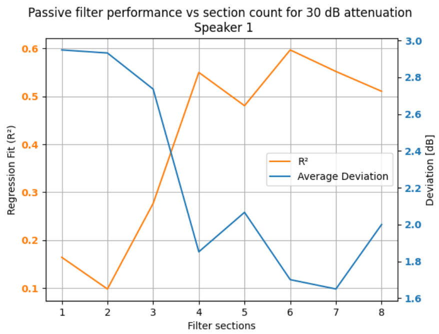

(a) Speaker 1

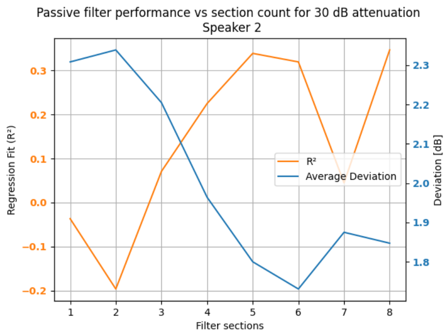

(b) Speaker 2

Figure C.6: Performance of different amount of filter sections for 30 dB insertion loss for speaker 1 (a) and 2 (b).