Case Study: Congregation Rodeph Shalom

A Tale of Options

Hadi Sumoro – www.HXAudioLab.com

The project featured in this article is a collaboration between HX Audio Lab and Wavelink, LLC in Philadelphia, PA. Raymond Stokes from Wavelink contacted HX Audio Lab to perform a system optimization at the Congregation Rodeph Shalom synagogue in Philadelphia, PA. The goal is to enhance the sound reproduction quality after a loudspeaker upgrade. As a side note, HX Audio Lab has obtained permission to use the name and images of the Rodeph Shalom Synagogue for this article.

Please note that the article will mention several product brands. As independent consultants, we are not affiliated with any specific organization. The products and software we choose are merely tools that help us achieve our objectives. This article offers insights into our thought process from the perspective of an acoustic and electro-acoustic consultant. While intended for educational and informational purposes, it does not delve into any specific audio topic in depth. This article highlights some unusual topics and presents a fraction of our findings on spatial averages, which a few audio practitioners have deemed controversial. Please note that HX Audio Lab has independently developed the processes highlighted in this article, drawing on years of experience in system optimization. At the time this article is written, some findings are inconclusive, some are part of our routine work, and some are conducted as second opinions from a different perspective.

As a technical note, we use AFMG Systune and EASERA for measurements and post-processing. Additionally, we utilize HX Audio Lab's Filter Hose to generate FIR filters, which helps us develop equalization curves. This article includes several screenshots that showcase the programs mentioned in this paragraph. Additionally, it's recommended to read our article: A Tale of "Muddiness" prior to reading this article.

About Congregation Rodeph Shalom, Philadelphia, PA

Founded in 1795, Congregation Rodeph Shalom is the oldest Ashkenazic synagogue in the Western Hemisphere and a vibrant center of Reform Jewish life in Philadelphia. For more than two centuries, Rodeph Shalom has served as a spiritual, cultural, and civic anchor—honoring Jewish tradition while embracing innovation and inclusion.

Our mission is to create a spiritual and communal home that inspires connection, learning, and action. Guided by enduring Jewish values and compelled to moral action, we seek to build a just, compassionate, and inclusive Philadelphia.

Rodeph Shalom offers a full range of worship, education, and community engagement programs for individuals and families of all backgrounds. Our services include weekly Shabbat and holiday observances, lifelong learning opportunities, and social justice initiatives addressing food insecurity, housing, and equity across the city. Through Breaking Bread on Broad—our secular, community-based food program—RS serves hundreds of local families each week with dignity and care.

As a congregation rooted in tradition and open to all, Rodeph Shalom continues to be a place where people come together to learn, pray, celebrate, and repair the world. In 2028, we will mark a historic milestone—the 100th anniversary of our landmark sanctuary—a testament to the enduring spirit and vision of our community.

The Room

The main synagogue is a large historical room. The Renkus Heinz steerable array is the primary loudspeaker, featuring two presets: one to cover the first floor only and another to split the sound beam to cover both the first and second floors. Figure 1 shows the room.

Figure 1 - The Synagogue

The coverage and initial sound reproduction quality are considered good after the steering algorithm is implemented. Over the years, we have worked with various steerable array brands, including Tannoy, Renkus-Heinz, EAW, and others. In today's technology, most steering algorithms are pretty well developed, resulting in excellent direct sound coverage, especially front-to-back. Note that a steerable array typically has a vertical arrangement of drivers. While this can result in great front-to-back coverage, horizontal coverage still largely depends on the loudspeaker design and the steering algorithm's side effects. Since the room is acoustically symmetrical (left and right), we mainly worked with ten microphone locations in the left audience area (on the first floor), as shown in Figure 2.

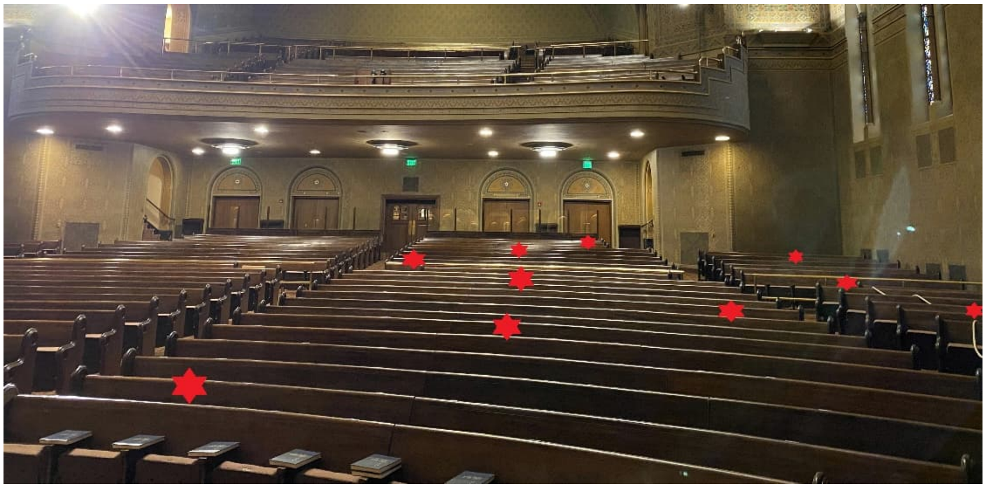

Figure 2 - Ten microphone locations on the first floor

Subjectively, for a large room, the reverberation time is not too long, but the overall quality is slightly dark or muddy. The choice of a steerable array is undoubtedly the right one, as it minimizes the sound beams that can excite unwanted room reflections.

The First Floor

Each microphone location on the first floor is shown in Figure 2, denoted with red stars. The measurements for each microphone location are shown in Figure 3. Notably, two locations yield significantly different frequency responses with a wide dip around 1kHz. Figure 4 shows only eight locations, excluding the two worst ones.

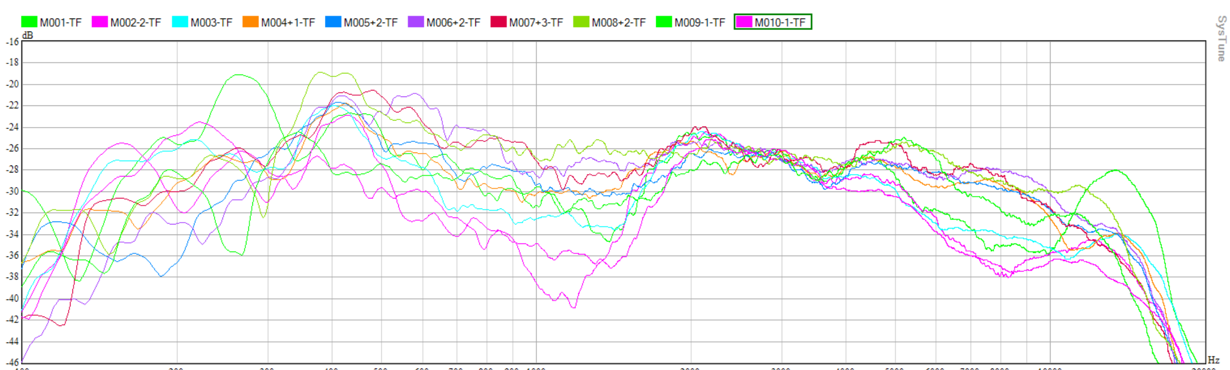

Figure 3 - Measurement result on each microphone location on the first floor

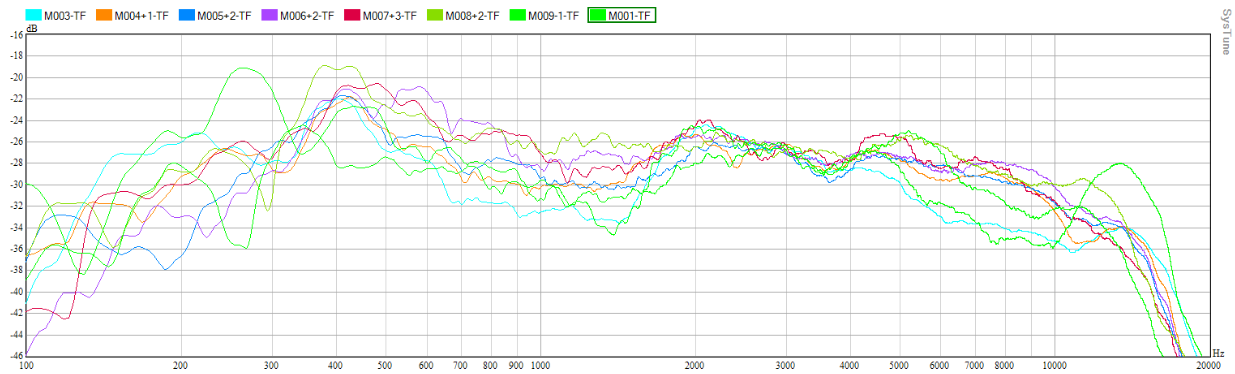

Figure 4 - Measurement result on each microphone location on the first floor, minus the two worst locations

The measurements in Figures 3 and 4 display the direct sound field responses of the (audience) left loudspeaker array, including the direct sound and reflections up to approximately 50ms after the direct sound arrival.

Rather than discussing topics that are usually addressed, let's delve into areas that aren't often explored—some of the nuts and bolts of the spatial average curve. The spatial average curve typically refers to a power average of the frequency responses obtained from different locations. Power average means the magnitude-only (logarithmic) average. However, audio practitioners rarely delve into how the mathematical average is performed, not in a mathematical sense, but in a practical one. Let us compare the spatial average of Figures 3 and 4.

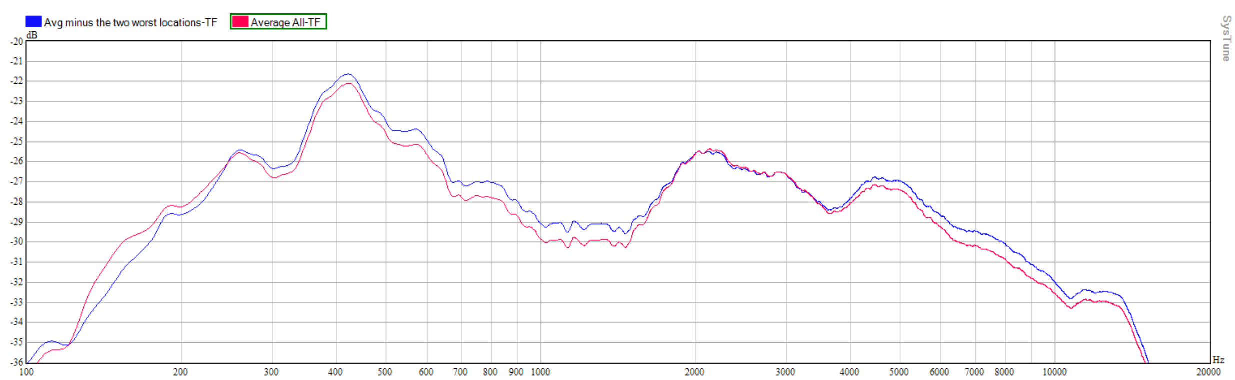

Figure 5 - Spatial average difference with (red curve) and without (blue curve) the two worst locations

At a glance, Figure 5 does not show significant differences. However, we have found that a 1dB difference in the spatial average over a wide frequency range is indeed noticeable. The 'small' difference in the averages makes sense, as the two worst locations (out of 10) do not significantly impact the average. However, let us introduce a can of worms. If the reader looks closely at Figures 3 and 4, the curves are somewhat 'normalized' at near 3000Hz. Without this normalization, Figure 3 would resemble Figure 6 below.

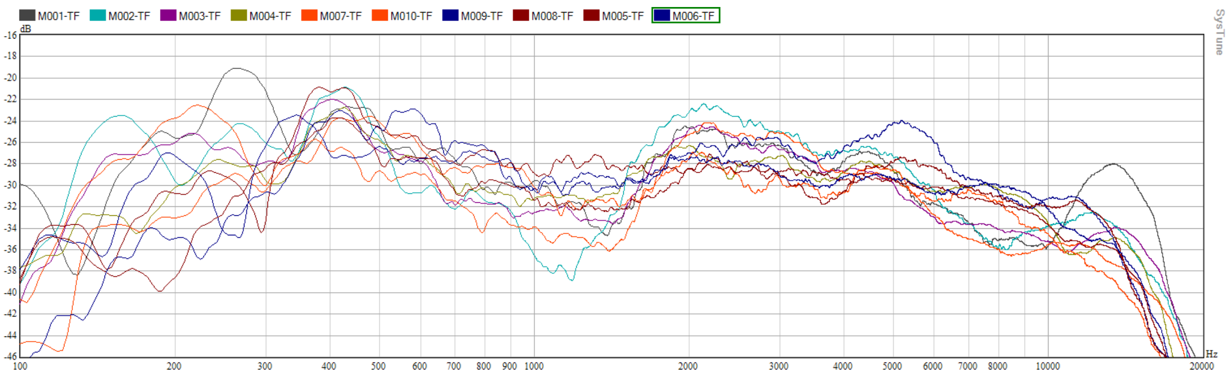

Figure 6 - Measurement result on each microphone location on the first floor without normalization (raw)

The spatial average of Figures 3 and 6 is illustrated in Figure 7 below. Relative to the red curve (spatial average for everything, normalized at 3000Hz), if we created an EQ curve based on these findings, Figure 5 will result in a slight cut at 400 – 1600Hz, and above 4000Hz, and Figure 7 will result in a slight cut at 1600 - 4000Hz, and below 300 Hz.

We genuinely care about where the 1dB adjustment is made. For this reason, we always use and periodically update the microphone's compensation curve. After many years of experience, we have come to understand that normalization is a necessary step before pressing the average button. The details of this are saved for future publication, as there is more to it than meets the eye. Additionally, this is ongoing research, and some conclusions remain unclear. For this project, we used Figure 5 to support our EQ development work.

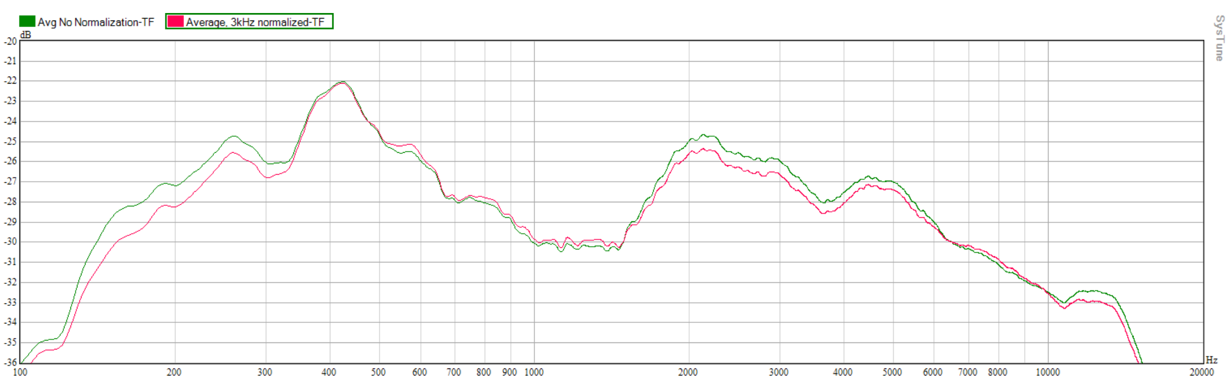

Figure 7 - Spatial Average curves of Figures 3 (red) and 6 (green).

Due to the 'small' difference (1 dB), we understand that many may dismiss these findings as insignificant. As everyone has their own methods for developing an optimization EQ curve, we study individual responses and various approaches to obtaining spatial average curves. Spatial average curves are numbers that do not represent any specific audience locations, but they're not entirely random, either. We consider it a 'trend'. When a spatial average is used in conjunction with individual responses, we obtain more comprehensive data for studying the interaction between the loudspeaker and the room. Here are several cases where we recognize that additional post-processing may be necessary before clicking the average button.

1. For >80% usage, rooms are mainly occupied at the half-front only. Hence, the EQ curve can be slightly weighted towards the front audience. In this case, individual responses can be gain-adjusted by 1-3dB before creating a spatial average curve to accommodate VIP areas, less occupied areas, etc.
2. A room that receives loud sound at the front but has a drastic SPL (sound pressure level) drop in the back area may require specific normalization or exclusion to examine how to create the spatial average. Without additional post-processing, the result may be overly weighted toward the front audience areas due to the high SPL.
3. Coverage differences (or problems). This is typically found when loudspeakers are not correctly specified/oriented, or for a variety of other reasons. Normalization and/or weighting may be beneficial to create a meaningful starting point for EQ development.
4. Impulse Response Windowing used for individual measurement will affect the individual curve and certainly affect the spatial average. The choice between a regular window and a multi-window, and how much you want to ‘zoom in’ on direct sound versus room reflections, are all essential considerations. This part is another substantial topic for other publications.

More thoughts can be incorporated into the development of the equalization curve than can be covered in a single article. This article presents only partial findings and thoughts on how we post-process data, with a specific focus on this room.

Implementing an EQ for the First-Floor Coverage

Figure 8 displays the individual microphone responses after EQ application, and Figure 9 presents the spatial average curves. The spatial average curves were calculated using normalization, with and without windowing.

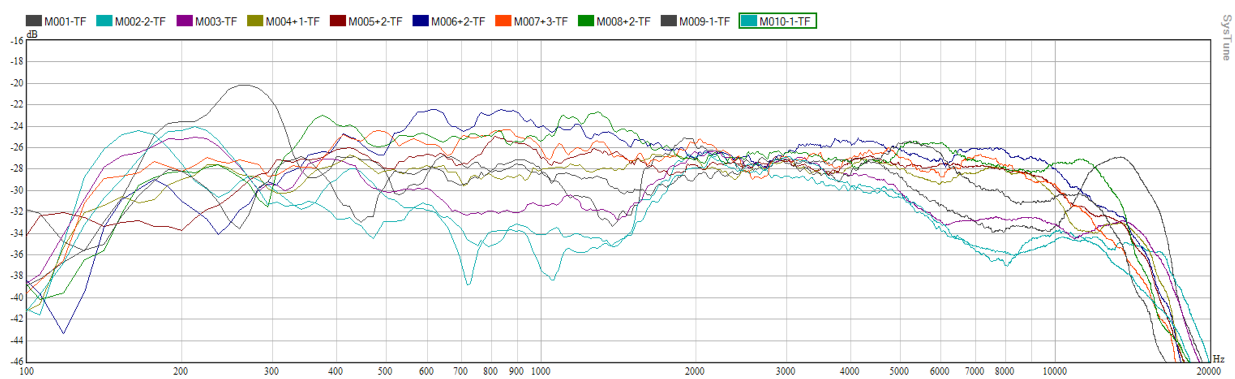

Figure 8 - Individual measurement with the first-stage EQ applied

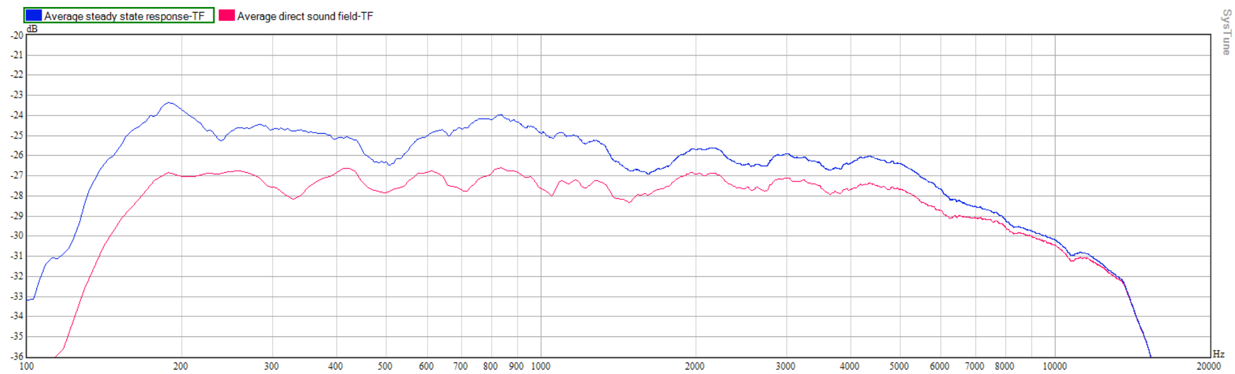

Figure 9 - Spatial Average (with normalization and includes all microphone locations) with first-stage EQ applied, direct sound field (red), and steady state responses (blue)

The first-stage EQ filter is developed in two processes. First, we identify negative coloration problems. This can be deduced initially from several figures above.

Figures 3 and 4 – we can notice an overall bump at approximately 400Hz, a wide dip at approximately 1000Hz, and variations at high frequencies, including a bit of barkiness at 2kHz.

Figures 5 or 7 – we can notice a bump at approximately 400Hz, a wide dip at approximately 1000Hz, and a wide bump around 2000Hz. Due to the zoomed-in graph, the high frequency seems to roll off steeply. It is a 4-dB drop from 5000 to 10000 Hz.

As audio practitioners, we recognize that the world of audio is a blend of 50% subjectivity and 50% objectivity. Therefore, the goal is to transform what currently sounds unsatisfactory (subjective) into a better-sounding system (subjective), through a process of technical work (objective). This line of thinking is illustrated in Figure 10. We emphasized the scientific part for the majority of the time, leaving the subjective tweak as the last 10-20% part.

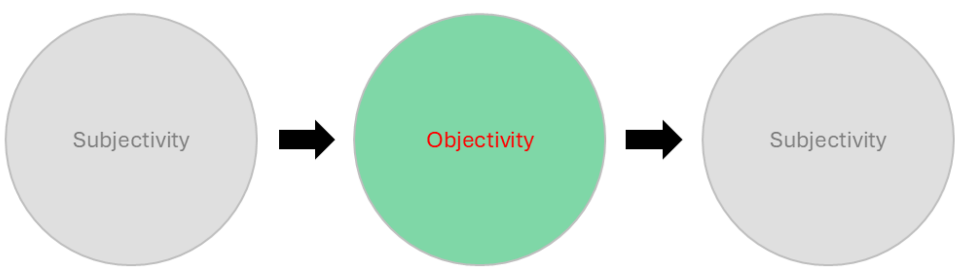

Figure 10 - The subjective-to-subjective process

The first iteration of the FIR filter is shown in Figure 11.

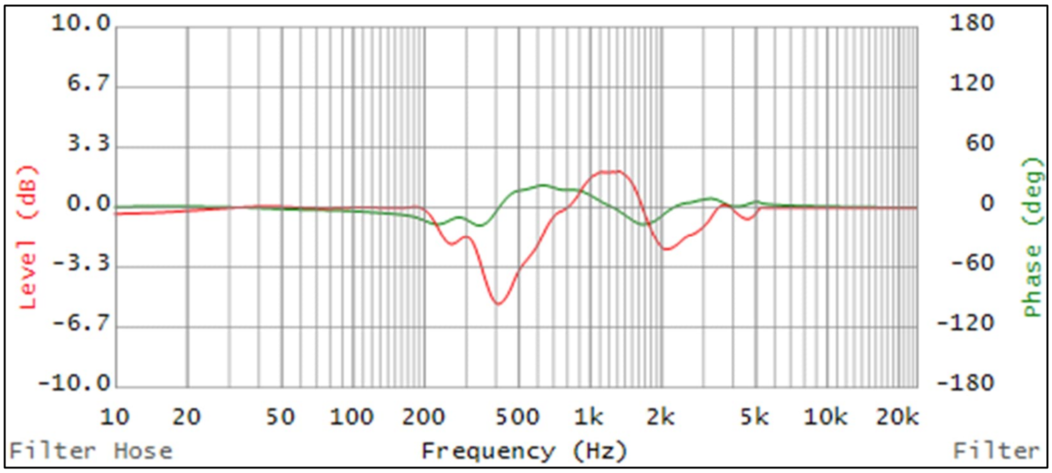

Figure 11 - The first FIR filter iteration

Following a quick listening session, we created the second iteration as shown in Figure 12. Note the 2 dB bump at 10kHz and around 180Hz. The second FIR filter iteration becomes the first-stage EQ. Subjectively, the sound reproduction quality becomes fuller, flatter, and more 'in front of the face' sound. However, we also observed low- to mid-frequency resonances being excited.

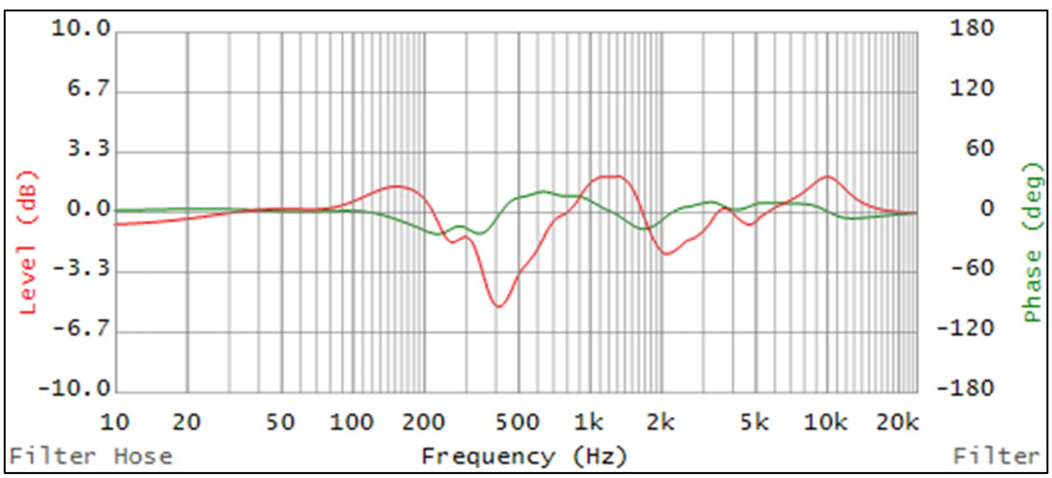

Figure 12 - The second FIR filter iteration

Next, the second-stage EQ is created to minimize negative coloration that may not be analyzable from the direct sound field responses. This is where Figure 9 can help. The red curve shows the spatial average from the direct sound field, and the blue curve shows the spatial average of the steady state response (i.e., no windowing is used). While working in the spatial average world, we can notice the curve differences as small as 1 dB. Two areas that caught our attention are bumps at approximately 850Hz and 190Hz. It turns out that the slight boost we added around 180Hz could introduce negative coloration (accentuating unwanted room resonance sound), so we undo that filter. Adding a filter to adjust the area around 850 Hz further cleans up the sound reproduction quality.

The reader may ask: “How about the two worst locations? Is there anything we can do about it?” We could redo the steering algorithm. However, it was outside of our scope of work at that time. Before further judging these two worst locations based on the direct sound field responses, we should consider the bigger picture. Subjectively, these two areas do not sound as bad as how they “look”. Figure 13 shows the direct sound field responses at the two worst locations. These curves were used in the EQ making process, but they are not entirely what we heard. Figure 14 shows the steady-state responses of these worst locations. It turns out the steady-state responses aren’t as bad-looking as the direct field response. “Bad looking” means larger up-down variation in this context.

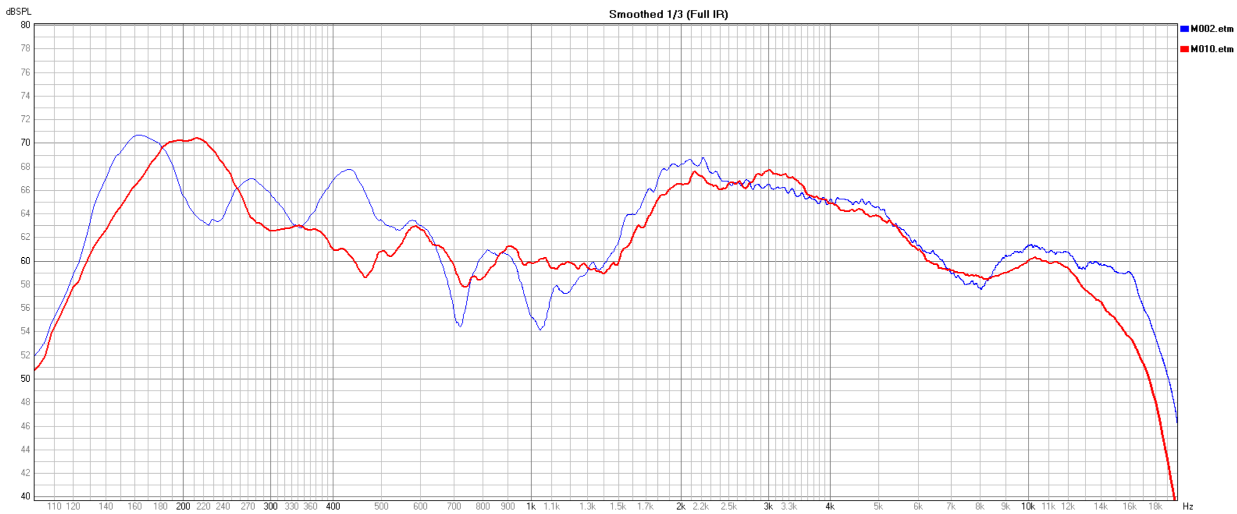

Figure 13 - The two worst locations, direct sound field responses

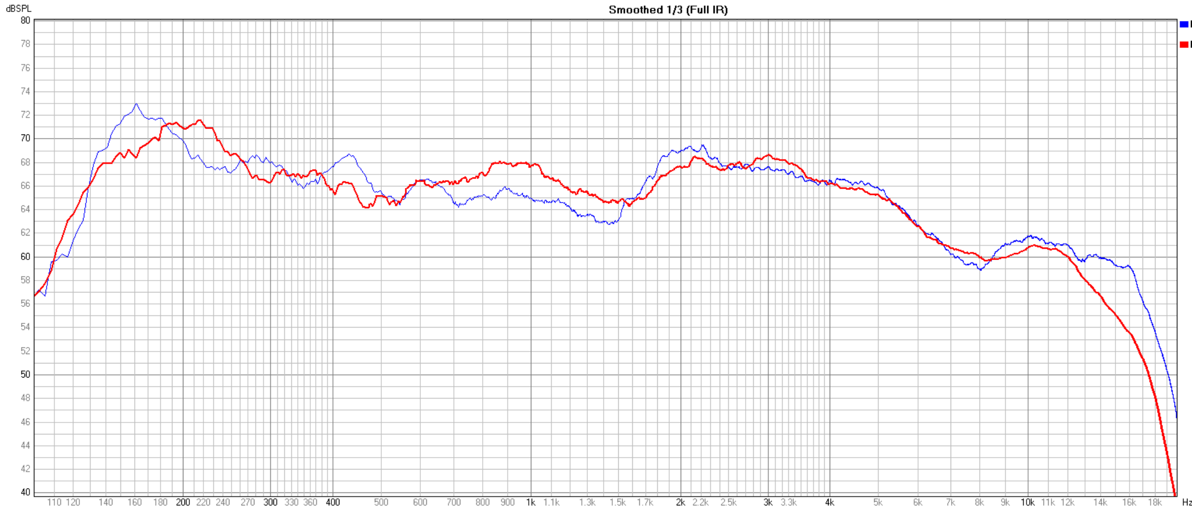

Figure 14 - The two worst locations, steady state responses

This also serves as a reminder that what we hear is not always what we see in the frequency domain curves, and we may work with something that does not directly represent what we hear. Objectivity nowadays is often represented by numbers, graphs, and digital zeros and ones. It's the practitioner's ability to translate and connect them into reality.

The Second Floor

There's more! The loudspeaker array contains the second preset, which splits the beam to cover both the first and second floors. We also developed an EQ to improve the system's reproduction for this preset. A side note: the measurement below includes the subwoofer.

The ten microphone locations on the first floor are reused, and eight additional microphone positions are selected to represent the second-floor audience area. Figure 15 shows the direct sound field responses of the first-floor microphone locations, and Figure 16 shows the steady state responses. Both were taken using the EQ curve developed for the first-floor-only preset.

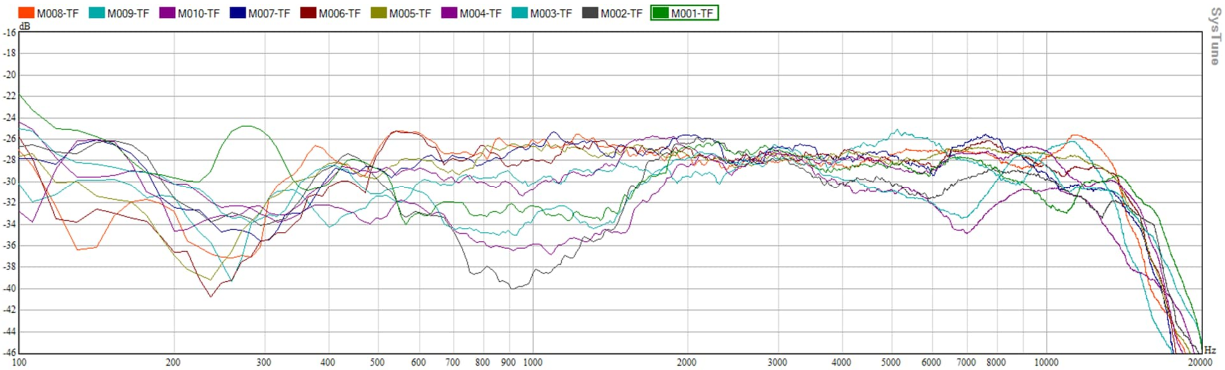

Figure 15 - Direct sound field responses of the first-floor mic locations with the first-floor-only EQ

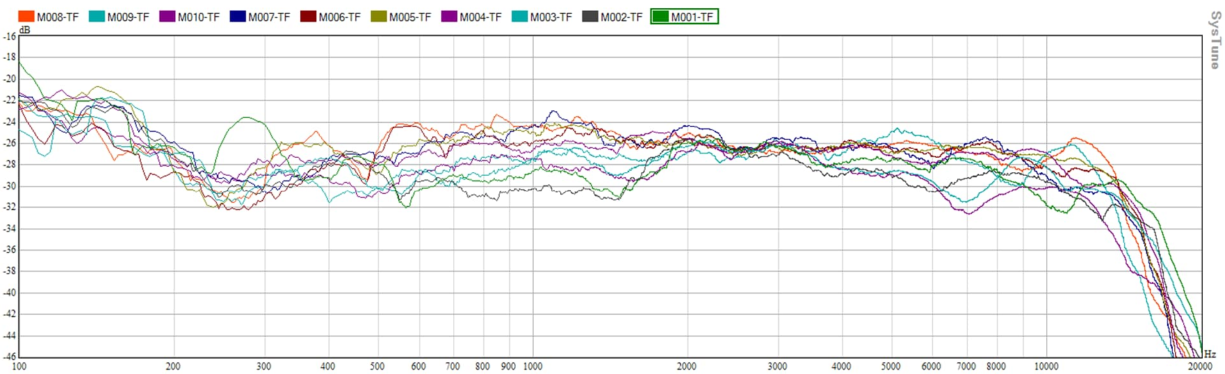

Figure 16 -Steady state responses of the first-floor mic locations with the first-floor-only EQ

Figure 15 resembles Figure 8. This indicates that the steering algorithm for both floors provides good coverage on the first floor, comparable to the first-floor-only preset. This is certainly good news! Subjectively, we didn't notice much difference between the first-floor-only preset and the two-floor preset with the same EQ when listening on the first floor.

Let us observe the microphone locations on the second floor. Figure 17 shows the second-floor direct sound field responses, and Figure 18 shows the steady state responses.

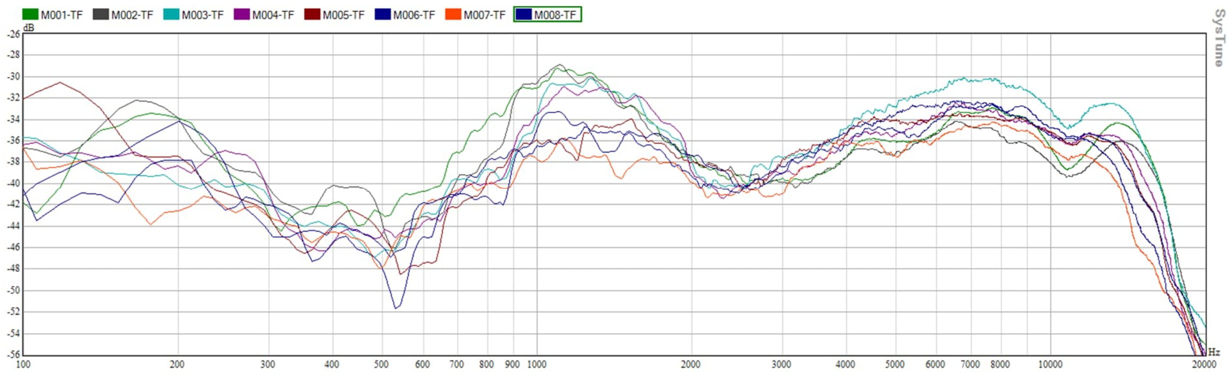

Figure 17 - Direct sound field responses of the second-floor mic locations with the first-floor-only EQ

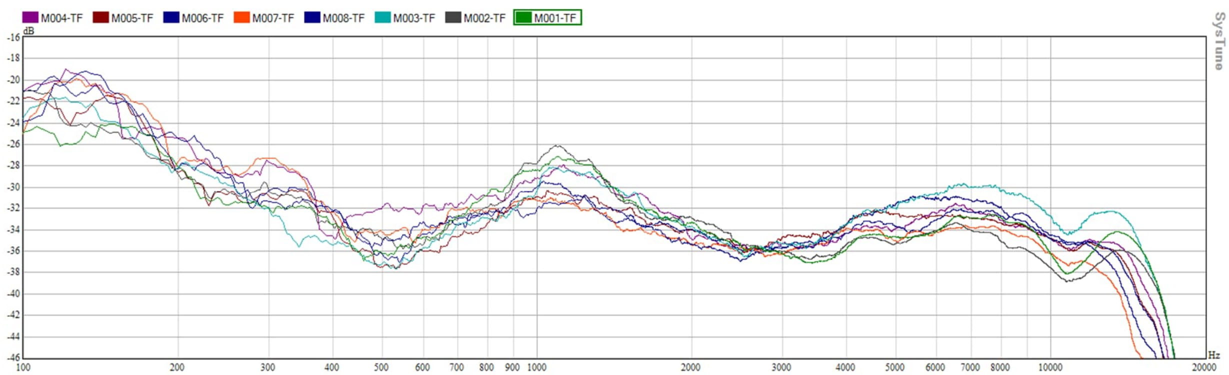

Figure 18 - Steady state responses of the second-floor mic locations with the first-floor-only EQ

Figures 17 and 18 are measured using the EQ for the first-floor-only preset. Readers can compare Figures 17 and 15, as well as Figures 18 and 16. To facilitate the comparison, Figure 19 displays the spatial averages of the steady-state responses, with individual responses normalized near 3kHz.

Figure 19 - Comparison of the first-floor (red) and second-floor (blue) spatial average of the steady state responses

The second-floor curve's look is undoubtedly far from ideal, but it is usable. Subjectively, the second floor sounded a bit 'small', softer, and hollow. There are two significant dips, at 500Hz and 3000Hz.

Since the first floor is already tidy, any changes made to help the second floor will have a more negative impact on the first floor. Our goal is to create a single EQ filter to improve the overall situation. Therefore, rather than making two boosts, our initial thought is illustrated in Figure 20 below, which focuses on reducing variation.

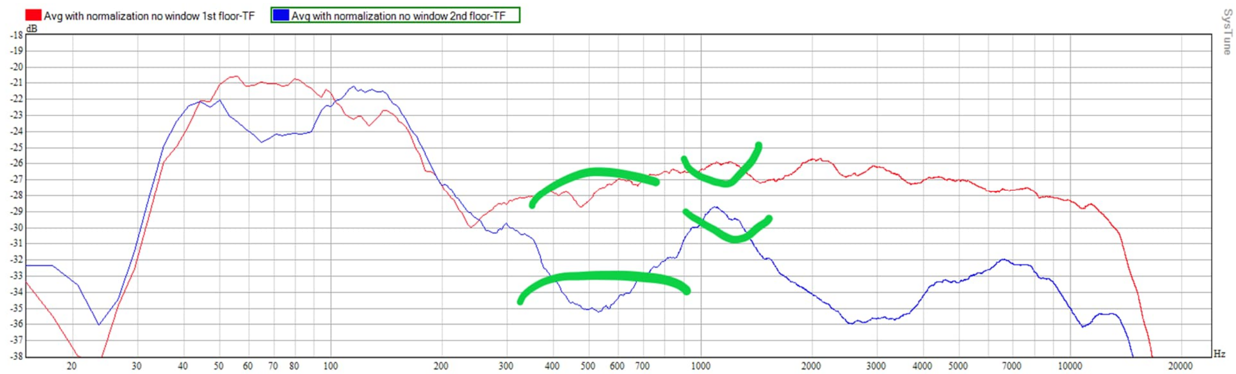

Figure 20 - EQ filter first thought

A slight boost at 500Hz and a slight cut at 1100Hz were first sought. By using two sets of ears (one person on the second floor and the other on the first floor), we can identify a point at which the filters create sufficient positive change on the second floor without significantly altering the first floor. Several listening sessions were also conducted to ensure that negative room reflections/resonances are not triggered. We implemented these additional filters as IIR filters for convenience.

Additionally, we attempted to add more filters at higher frequencies, particularly to minimize the noticeable dip at 3000Hz. Since 3000Hz is a very sensitive region in the human ear, additions and subtractions near this range have substantial effects. From another perspective, reducing the 7000 Hz (instead of boosting the 3000Hz) bump may also make the overall curve flatter. At this point, the process involves more subjectivity, along with objective goals. Sadly, we do not have the final measurement after all EQs are implemented due to the time constraint.

As a side note, the subwoofer output in the measurement above appears to be high. The subwoofer was aligned and set up using the full range preset for the first floor only, which directs all the full range transducers' output to the first floor. Using beam splitting in a two-floor setup reduces the overall output on each floor, as the total output energy is divided between the two floors. Hence, the subwoofer level seems high. This issue was addressed later.

Closing Remarks

As mentioned in the beginning, "*This article offers insights into our thought process from the perspective of an acoustic and electro-acoustic consultant. While intended for educational and informational purposes, it does not delve into any specific audio topic in depth.*" This article goes deeper into our thought process, rather than focusing on exact solutions, precise step-by-step instructions, or typical 'mainstream' methods. It also illustrates our complex thinking process for post-processing measured frequency responses and demonstrates various parallel thinking processes. There is a proverb: "all roads lead to Rome". There are many different methods to achieve the same goal (i.e., a good-sounding system). This article presents a fraction of our thought process and a part of our broader internal research on loudspeaker-room interaction.

www.HXAudioLab.com