Audio Engineering Society

Convention Paper

Presented at the 115th Convention
2003 October 10–13 New York, NY, USA

Design of Low-Order Filters for Radiation Synthesis

Peter Kassakian, David Wessel*

Center for New Music and Audio Technologies, U.C. Berkeley, CA

Correspondence should be addressed to Peter Kassakian or David Wessel (`[kassak, wessel]@cnmat.berkeley.edu`)

ABSTRACT

A fundamental goal of sound synthesis is to reproduce, and to control, as many facets of the sound as possible. By numerically solving a carefully constructed optimization problem, we are able to design low-order filters for use with a dodecahedral loudspeaker array to synthesize low order spherical harmonics over specified frequency ranges. The method, a variant of least-squares, is general, allowing for the inclusion of side constraints, arbitrary array geometry, and incorporation of measured loudspeaker characteristics. We compare the predicted loudspeaker array performance with high-resolution measurements of the physical system.

1. INTRODUCTION

In an effort to provide electronic musicians the same possibility of spatial complexity and radiatory interest as their acoustical peers, we study a dodecahedral loudspeaker array and a design method for determining optimal filters for radiation control. The

*The authors would like to thank Meyer Sound Laboratories, Inc. for continued collaboration and research support.

spatial interaction between an acoustic instrument and a small room is an important and exciting feature of a live performance – a feature that can be lost if the instruments are amplified via mono or stereo loudspeakers, which is usually the case with electronic instruments [1] [2]. The desire to bring precise spatial control to the world of electronic instruments initially motivated this research, though the filter design method, developed here, is applicable in other contexts, such as sound reinforcement using line arrays.

The classical engineering approach to radiation control, applicable for the most part in our framework, often neglects to include certain constraints on the problem that are necessary for systems where musical content is to be reproduced [3]. For example, it is important that the transfer functions in the field be reasonable in terms of phase linearity and spectral flatness. In [4], the authors discuss the desire for spectral flatness in the context of a broadband microphone array used to enhance near-field reception and suppress far-field interference. They handle the desire by incorporating constraints that the response be unity at a certain distance and at a several discrete frequencies. In this report, we focus on separating the spatial response from the response across frequency, choosing to equalize the frequency response with separately designed filters. This decomposition yields a more efficiently implementable system. These requests on the flatness and phase of the frequency responses, and the desire for easy implementation, can conflict with the desire for particular radiation patterns. In this study, we address these (often implied) constraints and attempt to formulate them explicitly in a numerically solvable optimization problem. We then perform high-resolution measurements of the physical system and verify the resultant performance.

In the following section, we introduce the specific physical system, along with our fundamental problem. We discuss the filter design problem from the point of view of optimization and propose an iterative method of design, related to least-squares. In Section 3, we discuss the proposed algorithm in detail and relate it to intuition. Based on the system's theoretical performance, we explain how to design filters for good spatial performance only in the frequency ranges where good performance is possible, surrendering in frequency ranges where the spatial performance is doomed.

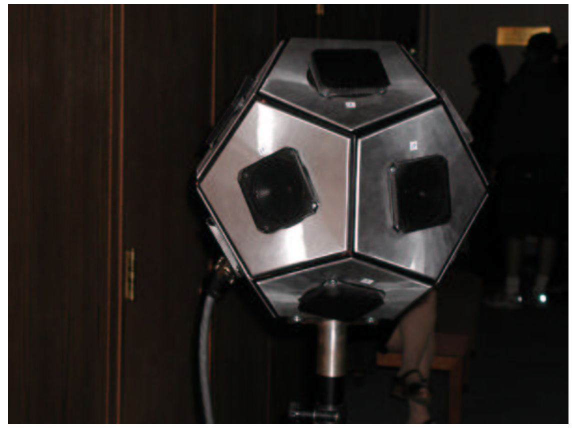

Fig. 1: **Dodecahedral Loudspeaker Array.**

Finally, we explain the implementation and measurement strategies, and present one degree measurements of response of the physical system, taken in an anechoic chamber. Here, we compare the achieved dipole pattern in one direction with an electronically steered version, aimed 45° off-axis. We conclude with potential applications of this easily controllable system and offer insights into the nature of its performance and limitations.

2. SYSTEM AND PROBLEM

2.1. Loudspeaker System

The system under investigation is that of the dodecahedral loudspeaker array, shown in Figure 1. It consists of twelve independent four-inch drivers, each with its own enclosure. We assume the system is driven with a monophonic signal, the spatial variation being attributable only to differing filters, the differing orientations and locations of the speakers, and the inherent directivity and frequency response of the speakers.

To perform the filtering, we deploy a Macintosh running Max/MSP [5] and a twelve-channel audio interface device, designed in-house [6]. This feeds twelve 150 Watt amplifiers that drive the speakers. The focus of this paper is on the design of the filters that will be applied to the twelve speaker signals.

Anechoic measurements of the response of an individual loudspeaker (mounted in the array) were performed at one-degree increments, on two orthogonal rings. These measurements allow us to avoid making idealizations about the directivity of the loudspeakers. Because of the pentagonal symmetry of the dodecahedron, the orthogonal rings are equivalent to rings separated by 18°. Together they represent 10 longitudinal rings of data taken every 18°. Figure 2 shows the magnitude of these responses as functions of frequency and angle. Their similarity leads us to make the simplifying assumption that the response of a single driver mounted in the array does not vary with longitudinal slice angle through the driver's center. In other words, we approximate the pentagonal symmetry as being circular.

Equipped with this data, we seek to find filters that, when applied to the loudspeaker signals, result in far-field spatial responses that are close to some target.

2.2. Problem Formulation

With no filtering, let $\vec{g}(\omega, x) \in \mathbb{C}^{M\times 1}$ be the vector of linear transfer functions of the $M = 12$ loudspeakers (input to field point) as functions both of frequency ω and spatial position x. If $H \in \Re^{N\times M}$ is the matrix of FIR (Finite Impulse Response) filter taps of order N for each of the speakers, and $\vec{p}(\omega)$ is the vector $[1\, e^{-j\omega} \cdots e^{-j\omega(N-1)}]^T$, then the response $y \in \mathbb{C}$ of the filtered system is given as

$$y(\omega, x) = \vec{p}^T(\omega)\, H\, \vec{g}(\omega, x) \qquad (1)$$

The complex response y is linear in the real filter tap matrix H. This means we can solve the following problem (the least-squares problem)[7]:

$$\min_H \int_{\omega\in\Omega} \int_{x\in X} |y(\omega, x) - d(\omega, x)|^2 \, \mathrm{d}\omega \, \mathrm{d}x \qquad (2)$$

where $d(\omega, x)$ is the desired response as a function of frequency and field position in space. Using measurement data to determine $\vec{g}(\omega, x)$ at discrete points in frequency and space, we can approximate the continuous problem above by the discrete version:

$$\min_H \sum_{\omega_i\in\Omega} \sum_{x_j\in X} |y(\omega_i, x_j) - d(\omega_i, x_j)|^2 \qquad (3)$$

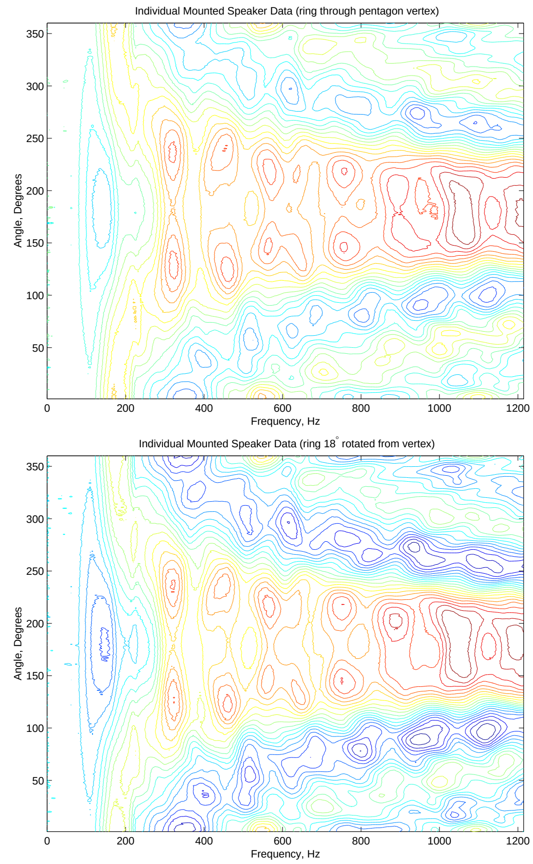

Fig. 2: **Individual Loudspeaker Data.** Magnitude of the response as a function of frequency (Hz), and angle (°). TOP: the circle of measurement is through the center of the driver, and through one vertex of the pentagonal baffle surrounding the speaker. BOTTOM: The circle of measurement also passes through the center of the driver, but passes through no pentagon vertex. It is a circle that is aligned 18° from the line segment joining the center of the driver with a vertex of the pentagonal baffle.

The solution to this problem can now be found by means of standard optimization code, and in fact can be found analytically by solving the associated normal equations [8] [9] [10].

The optimal filters minimize the summed squared error between the desired and realizable complex responses. Implicit in the formulation is that phase error and magnitude error are treated equally. The importance of these errors may not be equal, however. For example, at a given frequency, the absolute phase in the far-field is of no importance. The relevant characteristic relating to the phase of the responses is its linearity across frequency. The authors of [11] discuss this idea in abstraction and point out other subtleties of the far-field pattern synthesis problem. For us, in addition, depending on the spatial pattern, the relative and absolute magnitude of the responses across frequency may be of little import, as long as the normalized patterns are spatially accurate. For the purposes of low-impact implementation, it is advantageous to equalize the average (magnitude) frequency response or frequency response in a particular direction after the fact, recognizing that the application of a common filter to all twelve loudspeaker signals does not change the normalized spatial responses.

The desire for phase linearity, and the relaxing of the constraint on relative magnitudes across frequency motivates the formulation of a revised problem. In the next section we present a simple method for approximately solving such a problem, and discuss the algorithm in detail.

3. ITERATIVE LEAST-SQUARES

3.1. Phase Concerns

Unlike the standard least-squares problem in (3), which is a smooth, unconstrained convex problem, attempting to ignore phase error in the objective function leads to a non-convex problem. As mentioned in the previous section, the absolute phase in the far-field is ill-defined since the specific far-field radius r, is arbitrary; the magnitude of the far-field polar pattern decreases as $1/r$, and the phase is linear as a function of r [12]. We would like to find filters that optimize a *phase normalized* pattern, not the absolute pattern as explicitly specified by $d(\omega, x)$ in (3).

The objective function that minimizes the summed magnitude error squared is given by:

$$e = \sum_{\omega_i \in \Omega} \sum_{x_j \in X} \big| |y(\omega_i, x_j)| - |d(\omega_i, x_j)| \big|^2 \qquad (4)$$

In contrast to the above, which depends only on the magnitude of the desired response, $d_0(\omega, x) \equiv |d(\omega, x)| \in \Re$, formulation (3) requires a specific desired complex function $d(\omega, x) \in \mathbb{C}$. This suggests a method for approximately minimizing (4). By analyzing the phase of the resultant optimal responses generated by (3), we can modify $d(\omega, x)$ to better reflect the expected phases of the responses as a function of frequency. We then can solve (3) again, with a more appropriate target function.

If the phase variation with spatial position x_j is expected to be small, (for example if the set of field points X lie on a sphere of constant radius surrounding the array and cover same-sign lobes), we can approximate (4) by:

$$e \approx \sum_{\omega_i \in \Omega} \sum_{x_j \in X} \left| e^{-j\rho_i} y(\omega_i, x_j) - d_0(\omega_i, x_j) \right|^2 \qquad (5)$$

where, for any particular x_{j0}, ρ_i can be chosen to make the approximation exact, i.e., such that $|y(\omega_i, x_{j0})| = e^{-j\rho_i} y(\omega_i, x_{j0})$. Formulation (5) is, for fixed ρ_i, the least-squares problem.

Hence it's possible to move toward solutions that minimize (4) by solving first a least-squares problem (with ρ_i fixed), and then a "parameter approximation problem" that finds the best ρ_i, given a set of filters. We can view this simply as adjusting the phase of the complex target function $d(\omega, x)$ to better reflect reality as we converge on an approximate solution to (4).

By iterating several times, we can find solutions that perform better with respect to objective function (4) than a single solution to (3) does. This formulation does not explicitly address the desire for linear phase across frequency; it merely eliminates the artificial imposition of *zero* phase across frequency or, more precisely, the artificial imposition of a rigid phase specification across frequency. As we shall see, for short filters, or scalar gains, approximate linear phase is a natural side effect.

3.2. Magnitude Concerns

Because of the desire to implement the system efficiently, we opt to separate the magnitude response equalization across frequency from our spatial response synthesis problem. Applying a common filter to all of the twelve signals does not change the normalized spatial pattern in the far-field; it only alters the response across frequency. It is for this reason that we are motivated to generalize the algorithm, accounting for interfrequency magnitude differences.

We assign a magnitude parameter m_i for every frequency ω_i, and adjust the target response accordingly, as described above in the context of phase normalization. Hence, in conjunction with the phase normalization, we iteratively generate appropriate target functions across frequency: $d(\omega_i, x) = m_i e^{j\rho_i} d_0(\omega_i, x) = q_i d_0(\omega_i, x)$. A summary of the algorithm:

Algorithm

Let $q_i = 1$
Repeat until q converges:
A. Solve:

$$\min_{H} \sum_{\omega_i \in \Omega} \sum_{x_j \in X} \left| y(\omega_i, x_j) - q_i d_0(\omega_i, x_j) \right|^2 \quad (6)$$

B. Let:

$$q_i = |y(\omega_i, x_{j0})| e^{-j\Phi(y(\omega_i, x_{j0}))} \quad (7)$$

where,

$$\Phi(y) = \begin{cases} \mathrm{Arg}(y) & : \ \mathrm{Re}(y) \geq 0 \\ \mathrm{Arg}(y) + \pi & : \ \mathrm{Re}(y) < 0 \end{cases} \quad (8)$$

We have used this method to find optimal single-tap gains associated with each of the loudspeakers to achieve a dipole pattern across a range of frequency. The ability to specify the sets of frequencies Ω and field points X allows for the assignment of more or less weight in certain frequency ranges or in certain directions.

4. RESULTS

4.1. Specific Design and Measurement

The method described above is applicable in the design of arbitrary length FIR filters. We have found that, for a desired pattern of a dipole, single-tap filters (gains) perform very well, and only minimal extra performance is obtained by increasing the filter order. An added benefit of designing a set of twelve scalar gains is that the filters are spectrally flat and have linear phase responses. This does not guarantee flatness and linear phase in the far-field, since the response there is given as (1). It does, due to the reasonable response of the loudspeakers, and the simplicity of the array geometry, result in a response that is clearly adequate from a perceptual standpoint, and that possesses near linear phase in directions of geometric symmetry. The method provides us with a simple system that is easily equalizable.

A dipole response is unrealizable at high frequencies due to the increasing directivity of the loudspeakers and the decreasing ratio of wavelength to array dimension [12] [1]. Hence we let the set of optimization frequencies Ω include only frequencies where we expect to achieve a good fit. We know, since we are designing single-tap filters, that the frequency coloration due to the filters will be nearly flat and nearly linear phase across the entire audible spectrum.

Response measurements were performed in an anechoic chamber at one degree increments on the half-circle indicated in Figure 3. Figure 4 shows the predicted versus measured responses for a collection of frequencies in the range where the dipole pattern is achievable with reasonable accuracy. Also included in Figure 4 (bottom), is the predicted pattern for the simple system consisting of two speakers mounted opposite each other in the array, driven in opposite phase. The added speakers in the array produce a better dipole fit, and also allow electronic steering in any direction.

4.2. Pattern Control

Due to the cubic symmetry of the dodecahedron and the orthogonality of the three fundamental dipole modes, the dipole is easily reoriented by synthesizing three independent patterns with appropriately chosen scalar weights [1]. The bottom right plot in Figure 4 shows the results of such an electronic rotation. Hence we see that the spatial accuracy is good enough to allow us to use an analytic method

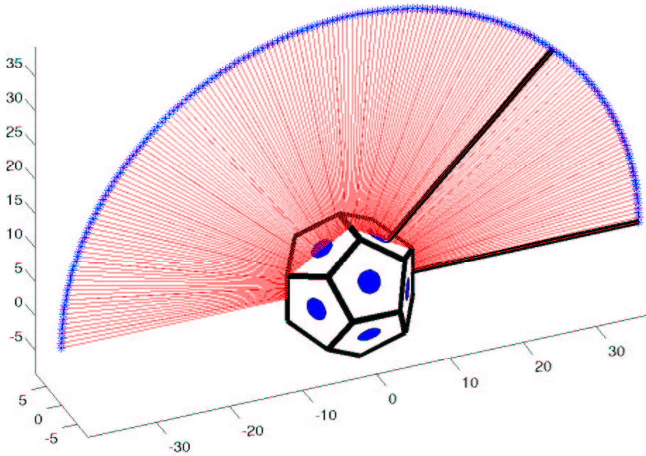

Fig. 3: **Measurement Orientation.** The dodecahedral array was positioned on a servo-controlled arm and responses were measured in the directions shown in the figure. The gains applied to the loudspeaker signals were chosen to produce two dipole patterns: one with its lobe direction at 0°, and the other rotated electronically by 45°.

for rotation (as opposed to generating distinct optimization problems for different rotation angles).

4.3. Pattern Deterioration

We conclude this section with two plots of measurements over a wider frequency range, showing the pattern deterioration as the frequency increases. The left plot of Figure 5 shows responses of the system equipped with unity gains (the optimal gains to create a monopole). The right plot shows the response of the system with the 45° rotated dipole gains. We see that pattern fit becomes less attainable as the frequency exceeds 400 Hertz for these two low-order spherical harmonics.

5. CONCLUSION

We have outlined a simple method for designing FIR filters for use with a dodecahedral loudspeaker array. The design method relies on a convenient algorithm that, in effect, trades response across frequency for better response across spatial angle. The algorithm can be thought of as an enhancement to the standard least-squares minimization in the complex domain. Though no guarantees are provided with respect to the optimality of solutions, we have found that the method provides us with a simple way to discover well-performing filters for our physical system. The filters perform better than standard least squares solutions, with respect to a specific frequency response independent objective function.

The utility of casting the problem as a specific optimization problem is multifaceted. Importantly, as contrasted with ad-hoc methods, we know in what sense our solutions are near optimal. We also have a framework of reasoning with which we can think about the problem. In addition, we are able to extend the reasoning to other, more constrained or less intuitive problems. For example, the generality of the outlined method is sufficient to allow for inclusion of constraints such as hard limits on magnitude frequency response for subsets of loudspeakers. This might be useful if the array consisted of loudspeakers of differing sizes.

Equipped with easily realizable filters that produce controllable patterns, we have a system that can be used not only for amplification but for musical composition, novel room measurement techniques, and reverberation enhancement. Future theoretical work includes a continued investigation into non-iterative algorithms for array pattern synthesis for musical application. Another area of interest is the problem of simultaneous control of separate patterns, perhaps in different frequency ranges, without the use of cross-over networks.

6. ACKNOWLEDGEMENT

The authors wish to thank Meyer Sound Laboratories for use of their anechoic chamber, their collaboration with respect to finding solutions to the technical problems described in this paper, and their continued financial support of the work. Support for this project was also provided by a UC Discovery Grant in Digital Media from the University of California's Industry-University Cooperative Research Program (IUCRP). We also wish to specifically thank John Meyer, Perrin Meyer, Roger Schwenke, Jon Arneson, Laurent El-Ghaoui, Ervin Hafter, Adrian Freed, Matt Wright, Rimas Avizienis, Richard Andrews, and Brian Vogel.

7. REFERENCES

[1] O. Warusfel, P. Derogis, R. Causse, "Radiation synthesis with digitally controlled loudspeakers" presented at the AES 103rd convention, New York, (1997 September 26–29).

[2] Gabriel Weinreich, "Directional Tone Color", Journal of the Acoustical Society of America, Vol. 101, No. 4, pp. 2338–2346, (1997 April).

[3] Aharon Ben-Tal, Arkadi Nemirovski, *Lectures on Modern Convex Optimization*, MPS-SIAM Series on Optimization, Philadelphia, PA, pp. 210–219, (2001).

[4] H. Lebret, S. Boyd, "Antenna array pattern synthesis via convex optimization" IEEE Trans. on Signal Processing, Vol. 45, No. 3, pp. 526–532, (1997 March).

[5] See http://www.cycling74.com for details.

[6] R. Avizienis, A. Freed, T. Suzuki, and D. Wessel, "Scalable Connectivity Processor for Computer Music Performance Systems" Proceedings of the International Computer Music Conference, Berlin, Germany, (2000).

[7] E. W. Cheney, *Approximation Theory*, Second Edition, AMS Chelsea Publishing, Providence, RI, 101–105, (1982, Reprinted 1998).

[8] Gene H. Golub, Charles F. Van Loan, *Matrix Computations*, Third Edition, The Johns Hopkins University Press, Baltimore, MD, pp. 236–247, (1996).

[9] Stephen Boyd, Lieven Vandenberghe, *Convex Optimization*, Draft. pp. 213–222, (30 May, 2001).

[10] Lloyd N. Trefethen, David Bau III, *Numerical Linear Algebra*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, pp. 77–85, (1997).

[11] Simon Doclo, Marc Moonen, "Design of far-field and near-field broadband beamformers using eigenfilters" Accepted for publication in Signal Processing, (2003).

[12] Leo L. Beranek, *Acoustics*, Published by the Acoustical Society of America, Woodbury, NY, pp. 93–100, (1993).

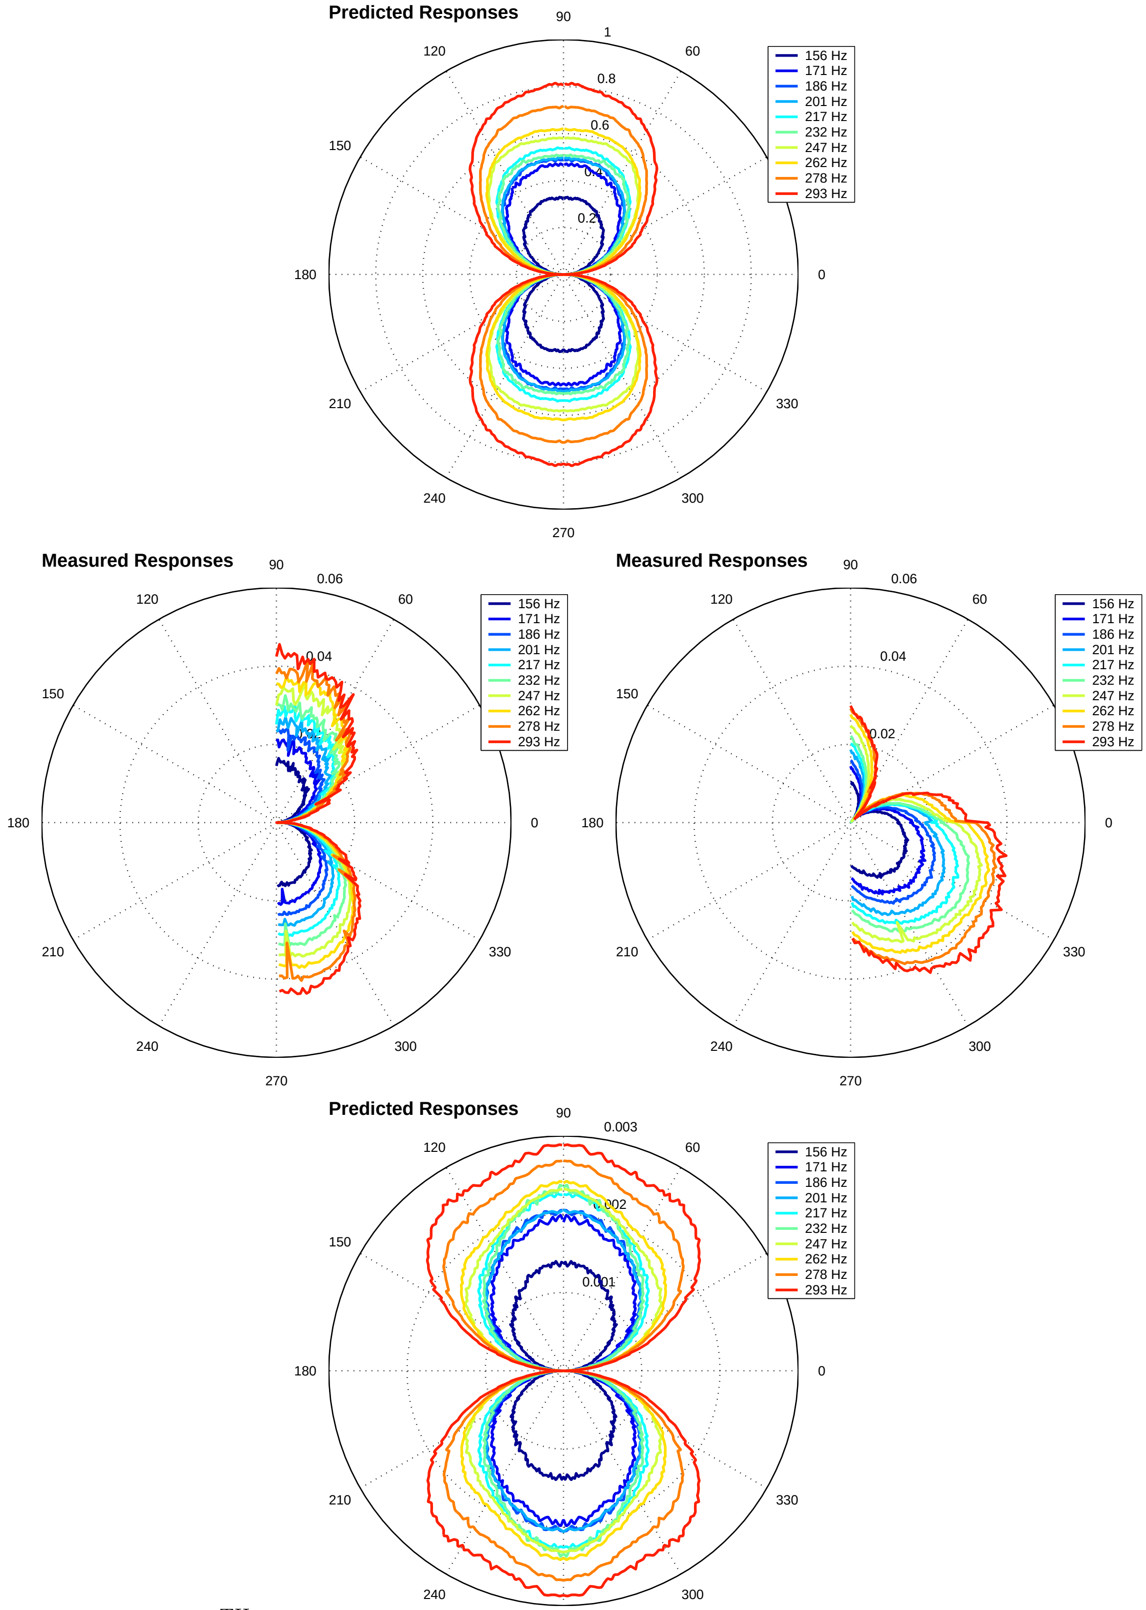

Fig. 4: **Predicted and Measured Responses.** TOP: The predicted pattern, calculated during the optimization. MIDDLE: Measured responses for the same frequencies. RIGHT: Measured pattern. The electronic rotation method works well for frequencies where the normalized dipole is achievable. BOTTOM: Predicted pattern for two opposite loudspeakers driven in opposite phase.

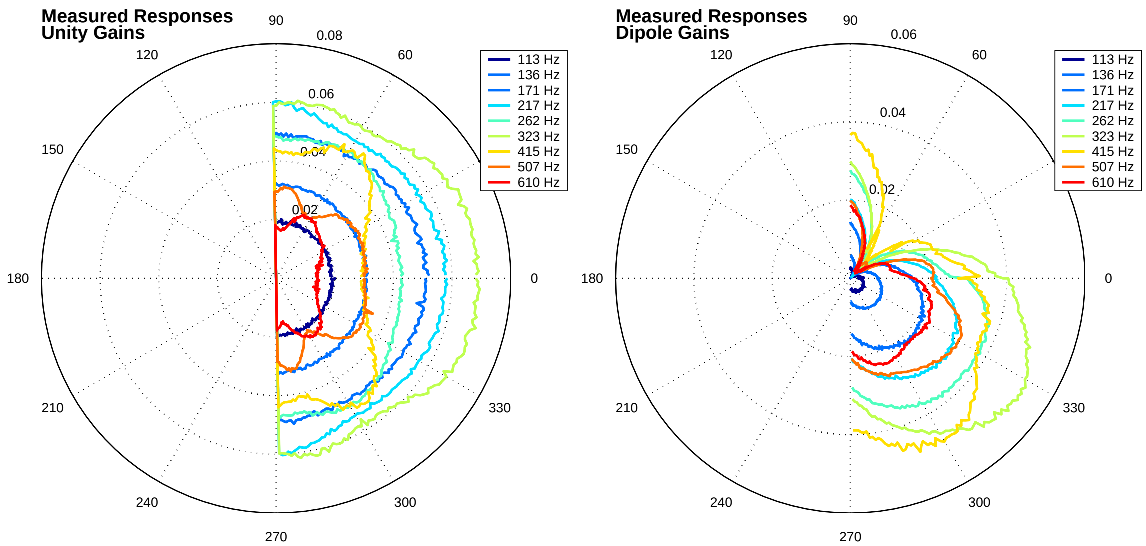

Fig. 5: **Responses at Higher Frequencies.** The left plot shows the response of the system outfitted with unity gains. Above 400 Hertz, the fit begins to deteriorate. The right plot shows the rotated dipole for the same frequencies.