Audio Engineering Society
Convention Paper
Presented at the 111th Convention
2001 September 21–24 New York, NY, USA

Loudspeaker Transfer Function Averaging and Interpolation

David W. Gunness
Eastern Acoustic Works, Inc.
1 Main St., Whitinsville, MA 01588 USA
David.Gunness@eaw.com

ABSTRACT

Transfer functions of acoustical systems usually include significant phase lag due to propagation delay. When this delay varies from one transfer function to another, basic mathematical operations such as averaging and interpolation produce unusable results. A calculation method is presented which produces much better results, using well-known mathematical operations. Applications of the technique include loudspeaker complex directional response characterization, complex averaging, and DSP filter design for loudspeaker steering.

0 INTRODUCTION

The phase response of a loudspeaker typically consists of a nearly minimum-phase characteristic plus excess phase lag due to the propagation time from the source to the microphone. Put another way, loudspeakers exhibit a relatively short impulse response that does not begin at zero time.

This results in rapidly wrapping phase plots at high frequencies. Phase wrap can be minimized during the measurement process by adjusting the offset delay of the measurement system; but if the objective of a series of measurements is to characterize variations of phase response, the relative differences in arrival time must be preserved. Inevitably, some of the measurements will exhibit significant wrapping.

Phase wrapping is not necessarily detrimental, in and of itself. However, it can lead to some misleading results when complex transfer functions are subjected to mathematical operations such as interpolation, smoothing, and averaging.

There are several tasks that call for acoustical transfer functions to be interpolated, smoothed or averaged. These tasks include complex directional response characterization, real-time measurement and tuning of loudspeaker systems, conversion of data for use in loudspeaker modeling programs, and the application of complex smoothing for the purpose of frequency-scaled time windowing. The algorithm presented in this paper was developed as part of a project to measure and represent the three-dimensional complex response of loudspeakers. That project is the topic of a paper by William Hoy and Charles McGregor [1].

0.1 Complex Polar Response

A source's directionality is directly related to the physical shape and size of the wave front it emits. A point source has no directionality. If an acoustical source has size, it is not possible to remain equidistant from it while sweeping a measurement microphone through a circular arc to measure its polar response (except in the case of a spherical source). Consequently, the propagation time from real sources will vary with direction. In fact, the particular way in which propagation time varies with direction is the most critical piece of information provided by complex polars. This information is required for accurate interference calculations, by which the performance of arrays of loudspeakers can be predicted.

There has been much spirited debate within the Audio Engineering Society in the last decade concerning the angular resolution required for complex polar data. Various conclusions have been reached, ranging from 10 degrees to 0.5 degrees per data point. Even without consulting transcripts of these debates, one can reasonably assume that proponents of the finer resolutions have no faith in the effectiveness of interpolation. Such mistrust is not unfounded, as the following example illustrates:

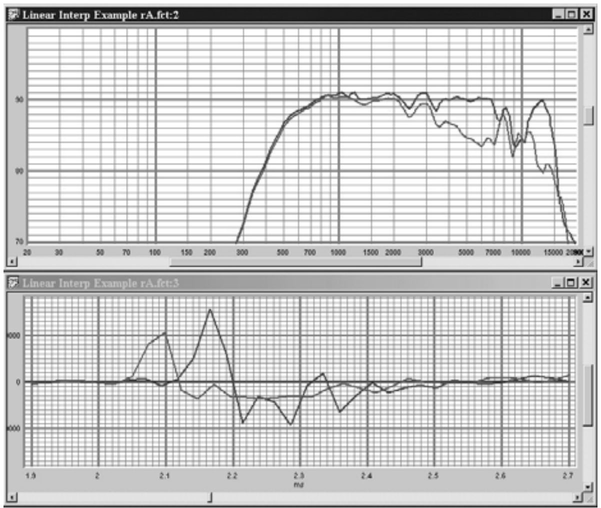

Figure 1: Magnitude and Impulse Response at 15 & 20 degrees Off Axis

Figure 1 presents two response measurements of a highly directional high frequency horn, taken at 15 degrees and 20 degrees off axis. This particular device and these off-axis angles were selected because they represent a practical worst case, as far as the angular resolution required. The impulse response chart shows clearly that the propagation time varied by about 0.07 ms between the two measurements.

If we apply simple linear interpolation to estimate the complex transfer function at the midpoint between the two (their average) we get an impulse response which lies between the impulse responses of the two measurements at every point, as shown in Figure 2. The Fourier transform of that impulse response (the frequency response), on the other hand, does not lie between the original frequency responses. In intuitive terms, adding the two transfer functions and dividing by two shows the interference of the two functions, rather than estimating the midpoint of a transition from one to the other.

The interpolated impulse response consists of ½-amplitude arrivals at the times of the original impulses. However, what is required is a single impulse, similar in shape to the original impulses, but occurring at a time halfway between the two original impulses. Furthermore, the interpolated magnitude response should lie between the two original magnitude responses.

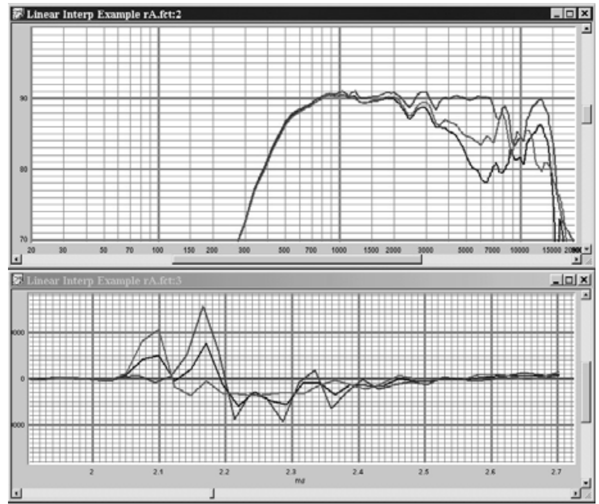

Figure 2: Linear Interpolation of Complex Transfer Function

1 AN ALGORITHM FOR INTERPOLATING TRANSFER FUNCTIONS

Conceptually, the interpolation of loudspeaker transfer functions can be divided into two separate problems: interpolating the propagation delay, and interpolating the characteristic complex frequency response (which is to say, the shape of the impulse). If we can eliminate the difference in arrival time, the phase response of the two transfer functions will become similar enough that interpolation is effective.

1.1 Using the Cross-correlation Function to Eliminate Excess Delay

The cross-correlation of two impulse response functions gives an indication of their similarity for various offset times. If the shapes of the impulse responses of the two functions are similar enough, then the cross-correlation function will have a well-defined peak that indicates the relative arrival time, or “excess delay”.

Cross-correlation is defined as the convolution of one sequence with a reverse-order second sequence. In the present case, the sequences are impulse responses. The cross-correlation can be efficiently calculated in the frequency domain by multiplying the Fourier transform of the first impulse response by the complex conjugate of the Fourier transform of the second response: then taking the inverse Fourier transform (IFT) of the result. In pseudo-code, this looks like:

$$\text{Cross-Correlation} = \text{IFT}(\text{FFT}(\text{Impulse}_1)\cdot\text{Conj}(\text{FFT}(\text{Impulse}_2)\,)\,) \quad (1)$$

For the particular case of interpolating complex polar response, adjacent transfer functions should be extremely similar. So, the cross-correlation function would be expected to have a well-defined peak. Figure 3 shows the cross correlation of the two transfer functions illustrated in Figure 1, indicating an excess delay of .075 ms.

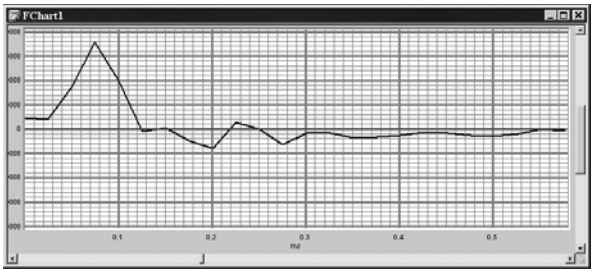

Figure 3: Cross-correlation of 15-degree and 20-degree Transfer Functions

The propagation time difference should be backed out of the ”late” transfer function. This is accomplished by synthesizing a transfer function representing the excess delay. This transfer function will have unity magnitude and a phase response corresponding to the delay. The complex value at each frequency is defined by,

$$F(\omega, t_e) = \cos(-\omega t_e) + j\sin(-\omega t_e), \qquad (2)$$

where ω is 2πf, and t_e is the excess delay.

The late transfer function is divided by the excess delay transfer function. The resulting transfer function is no longer late, but is approximately synchronous with the early transfer function. In fact, its phase response would be expected to be as similar as possible to that of the early transfer function.

Figure 4 shows the two example impulse responses, as well as the phase response (unwrapped for readability) after removing the excess delay.

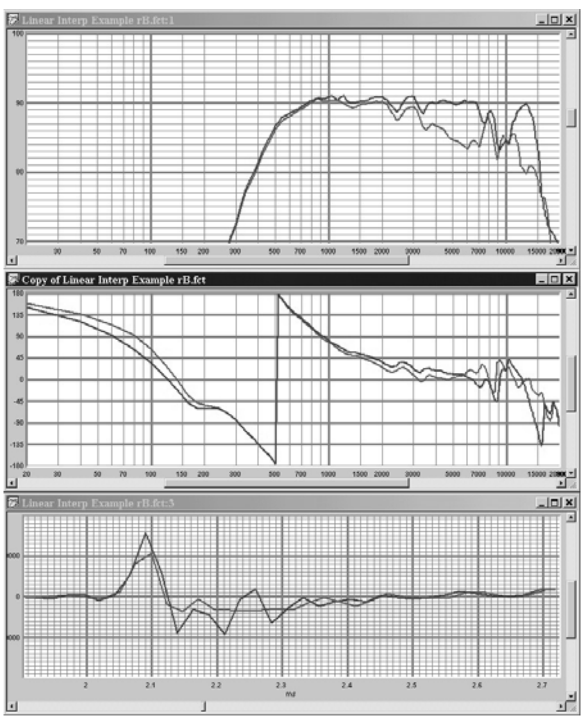

Figure 4: Transfer Functions with Excess Delay Removed

We are now ready to interpolate the impulse response.

1.2 Geometric Interpolation

The requirements for an effective interpolation method are that:

a) The interpolated magnitude response should vary logarithmically from one magnitude response to the other. Effectively, we would like to linearly interpolate decibel values.
b) The phase response should interpolate correctly regardless of wrapping. Correct results will be evidenced by an impulse shape that morphs smoothly from one impulse shape to the other.

These requirements are met if we adopt an interpolation method that corresponds to the geometric mean, rather than the arithmetic mean. For purposes of illustration, let’s compare the linear mean of two quantities (interpolation of the midpoint) to the geometric mean:

$$Linear\_Mean = \frac{(Value_1 + Value_2)}{2} \qquad (3)$$

$$Geometric\_Mean = \sqrt{Value_1 * Value_2} \qquad (4)$$

Consider two complex quantities which differ by nearly 180 degrees in phase: 1 + 0j, and -0.9 - 0.1j. The linear mean of these two values is (1 – 0.9 – 0.1j) / 2, or 0.05 –0.05j. The magnitude of this result is much smaller than the magnitude of either of the original values, so it obviously fails to satisfy requirement (a). The geometric mean is $\sqrt{-.9-0.1j}$, or .05 - .95j. This result falls nicely between the original values in both magnitude and phase. Working a few more examples will demonstrate that the geometric mean satisfies both requirements, as long as the phase difference is no greater than π.

In order to exploit the desirable behavior of the geometric mean, we must expand the averaging equation to an interpolation equation:

$$Geometric\_Interpolation = Value_1 * \left(\frac{Value_2}{Value_1}\right)^{ratio}, \qquad (5)$$

where ratio is the fraction of interpolation (if the ratio is 0.0, the function returns $Value_1$; if the ratio is 1.0, the function returns $Value_2$). This expression should not be reduced to $Value_1^{(1-ratio)} \cdot Value_2^{ratio}$, because the reduced expression only provides valid results for phase differences of less than π/2 and absolute phase of +/- π/2; whereas, equation (5) is valid for any absolute phase value, and phase differences of +/- π.

When geometric interpolation is applied to the example transfer functions, the result is as shown in Figure 5.

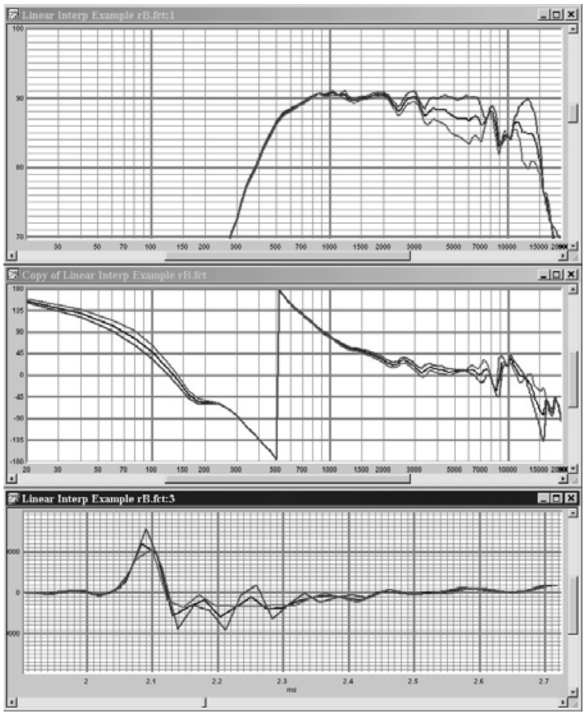

Figure 5: Aligned and Interpolated Transfer Functions

1.3 Interpolating the Propagation Delay

Interpolating the propagation delay is a simple matter. We simply work the method described in section 1.1, in reverse. For the delay we use the excess delay multiplied by the interpolation ratio. For the example, we multiply the interpolated result by the transfer function equating to ½ of .075 ms, or .0375 ms.. The interpolated response is shown in Figure 6, overlaid with the actual measured response it estimates. In this example, the estimated and measured results are nearly identical.

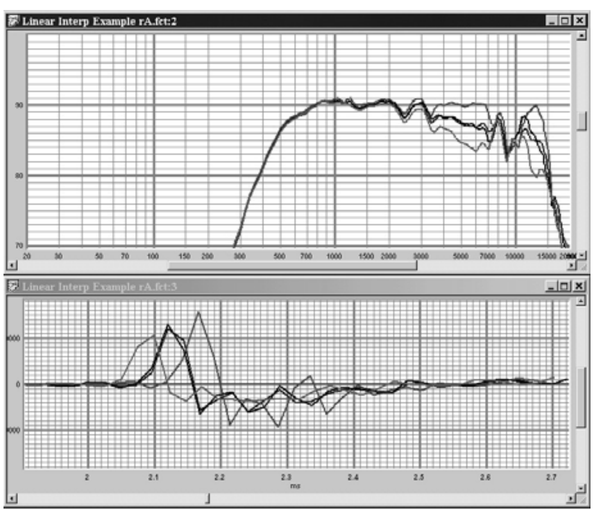

Figure 6: Interpolated vs Measured Response

1.4 Additional Techniques

It has been found that for marginally matched transfer functions, identification of the excess delay is accomplished more reliably if the envelope of the cross-correlation function is used, rather than the cross-correlation itself. The envelope may be calculated by treating the cross-correlation as an impulse response, and calculating its corresponding analytic impulse [2]. The envelope of the cross-correlation is the magnitude of its analytic impulse.

The cross-correlation result could be thought of as being weighted by the response magnitudes of different frequencies. The frequencies at which the response magnitude is relatively large have a greater effect on the cross-correlation result than do the frequencies at which the response magnitude is relatively small. For most loudspeaker applications this is a desirable effect, because it naturally weights more heavily the frequency range over which the loudspeaker operates effectively. However, in some applications it may be desirable to equalize the levels of the various frequencies before calculating the cross-correlation. In the field of computer imaging, the function so obtained is referred to as the “phase correlation” [3], because only the phase of the input functions affects the result - the magnitude response is effectively ignored.

1.5 Limitations of the Algorithm

The method described in this paper is an algorithm, not a transform. It falls far short of the level of robustness expected of a mathematical transform. It only results in good estimates if:

a) The two transfer functions to be interpolated are similar to each other.
b) Each of the transfer functions can be characterized as having a particular arrival time (the impulse response is sufficiently focused).

A particular condition that violates the second requirement is when an impulse response contains two discreet arrivals, nearly equal in level. Generally, the polar response of a single loudspeaker satisfies both constraints. But, if a loudspeaker consists of multiple sources, they should be measured separately and mathematically summed to produce the net acoustical response. A complete treatment of this process is given in Hoy/McGregor [1].

2 OTHER APPLICATIONS

In addition to loudspeaker polar response, there are several other audio-related applications of this technique.

2.1 Measurement Averaging

Averaging is often helpful when applied to acoustical measurements. Most commonly, it is used to increase the signal to noise ratio of a measurement. It can also be used with a source-independent measurement system such as SIA SmaartLive to reduce the inclusion of late reverberant energy in the measurement results. For this to be effective, the complex values of the transfer functions must be averaged. This technique is referred to as “vector averaging”, or “complex averaging” [4].

When acoustical measurements are conducted in-situ, some amount of variability of arrival times will be observed. When making measurements at long distances outdoors (i.e., 100 m or more), arrival time variability of 5 to 20 ms is not unusual. Even in large indoor venues, the variability might be on the order of one or two milliseconds.

The effect of arrival time variability is similar to the interpolation illustrated in Figure 2. Successive measurements interfere, resulting in a misleading reduction in high frequency levels. Consequently, complex averaging has only been practicable indoors, at distances up to 50 meters or so.

Complex averaging can be carried out using the interpolation algorithm described in this paper. Using this technique, the measurement process becomes tolerant of arrival-time variability, and a better figure for the arrival time itself can be determined. The benefits of complex averaging can then be realized outdoors and in large venues.

2.2 DSP Filters for Loudspeaker Steering

When the interpolation algorithm was first conceived, a number of reasonableness tests were performed. One of these tests was to interpolate between some standard high-pass and low-pass filter types. If the magnitude, phase, and impulse response all appeared to be reasonably similar to the input functions, it would support the efficacy of the new method. A surprising result was observed.

When a second-order Butterworth filter and a fourth-order Butterworth filter are interpolated with a ratio of 0.5, the result is an exact match in both magnitude and phase to a third-order Butterworth filter. Similar results are obtained between orders of Bessel filters. The interesting aspect of this observation is that the essential behavior that characterizes a particular filter type survives interpolation. An interpolation between two maximally flat (Butterworth) filters is another maximally flat filter.

By extension, it stands to reason that fractional-order filters can be defined by geometric interpolation; and that these fractional order filters will also share the essential character of the filter type. Experiments have verified that this presumption is true. A 3.5-order Butterworth filter, for example, has an impulse response that rings slightly longer than a 3^{rd}-order filter, and slightly shorter than a 4^{th}-order filter. It has a slope of 21 dB per octave.

While continuously variable low-pass and high-pass filters might prove useful to designers of multi-way loudspeakers, a more interesting value of such a filter set is in loudspeaker steering. It has been found that the smoothest steered directional response is obtained with a programmed-optimizer and high-precision complex polar data [5]. Unfortunately, the order of a low-pass filter was previously an integer-valued function. It was not productive to allow an optimizer to change it, because its effect on the result was discontinuous. Implementing continuously variable filters could lead to better optimization results.

3 CONCLUSIONS

A new transfer function interpolation algorithm has been presented. This algorithm produces better estimates than linear interpolation when the transfer functions being interpolated include pure delay. The new algorithm enables complex polar data to be interpolated with sufficient accuracy for loudspeaker interference calculations; and it allows complex averaging to be employed in outdoor acoustical measurements, despite the presence of changeable wind. Fractional order high-pass and low-pass filters were proposed, based on interpolation of the familiar types.

4 REFERENCES

[1] W.Hoy and C. McGregor, “Loudspeaker Complex Directional Response Characterization”, to be presented at the 111th Convention of the Audio Engineering Society, 2001.

[2] A. Duncan, “The Analytic Impulse”, *Journal of the Audio Engineering Society*, vol. 36, pp. 315-327. (1988 May).

[3] C. Kuglin and D. Hines, ``The Phase Correlation Image Alignment Method," Proc. Int. Conf. Cybernetics and Society, 1975, pp. 163-165.

[4] J. Anderson and C. Dayton, SmaartLive Help, SIA Software Company, Topic “Averaging and Smoothing”.

[5] D. W. Gunness and J.F.S. Speck, “Phased Point Source Technology and the Resultant KF900 Series, “ Eastern Acoustic Works, Inc., 1997)