Audio Engineering Society

Convention Paper 7994

Presented at the 128th Convention
2010 May 22–25 London, UK

Characterising studio monitor loudspeakers for auralization

Ben Supper[1]

[1] Focusrite Audio Engineering Ltd, High Wycombe, UK.
ben@focusrite.com

ABSTRACT

A method is presented for obtaining the frequency and phase response, directivity pattern, and some of the non-linear distortion characteristics of studio monitor loudspeakers. Using a specially-designed test signal, the impulse response and directivity pattern are measured in a small recording room. A near-field measurement is also taken. An algorithm is presented for combining the near- and far-field responses in order to compute out the early reflections of the room. Doppler distortion can be calculated using recorded and measured properties of the loudspeaker. The result is a set of loudspeaker impulse and directional responses that are detailed enough for convincing auralization.

1. INTRODUCTION

This research supports a DSP algorithm that aims to simulate loudspeaker listening realistically over headphones (a technique known as auralization). There are four components in this system: a model of loudspeaker sound radiation, a room reverberation simulator, a model of directional head-related cues, and an algorithm for imposing the simulation on an audio signal in real time. Many loudspeaker and room models are included, and these can be interchanged. This paper concentrates on the acquisition of loudspeaker data.

It is clear that there are many subtleties in loudspeaker systems that affect the way that they radiate sound, both on-axis, and into the reverberant field. This makes realistic loudspeaker modeling by theoretical analysis impractically difficult, so we base our simulation on recorded impulse responses. There are some technical challenges to overcome before any degree of fidelity is obtained. These will be covered in the following sections.

Section 2 outlines the main reasons why different loudspeakers sound different on- and off-axis: even those with a fairly flat spectrum over their operating bandwidth. With these in mind, Section 3 covers the stimulus generation and measurement methods for recording loudspeakers. Section 4 presents the post-processing techniques for retrieving corrected impulse responses, and some actual recorded data.

2. INDIVIDUAL PROPERTIES OF LOUDSPEAKERS

There is a wealth of literature discussing the phenomenon of different loudspeakers sounding different, and most introductions to loudspeaker reproduction deal with this concept (for example, [1]). Assuming the drive units in a system to be well-engineered and ordinarily linear, what follows is a brief outline.

2.1. Factors affecting on-axis response

2.1.1. Case size

The physical volume occupied by the enclosure is the most significant variable determining the low frequency response of a loudspeaker. The smaller the cabinet, the greater the stiffness imparted to the mechanical system. This raises the resonant frequency of the bass drive unit, reducing its ability to radiate low frequencies. Also, larger-volume cabinets generally accommodate larger woofers which can move a greater volume of air. Although an active filter can increase the power delivered to the bass driver at low frequencies, and hence extend the low-frequency response, there is a limit to what can be achieved in practice. This is because the required amount of cone displacement increases by a factor of four for every octave of bass response required.

2.1.2. Bass reflex ports, passive radiators, and transmission lines

A reflex port is a Helmholtz resonator added to the cabinet, generally tuned to a frequency about an octave below the resonant frequency of the woofer, and damped using a soft, porous substance such as foam or mineral wool. This extends the lower range of the loudspeaker. The advantage of greater bass extension is counterbalanced by the higher-order low frequency roll-off caused by the port (24dB/octave rather than 12dB/octave). This affects the low-frequency phase response of the system, and also produces a ring in the loudspeaker's impulse response that is imposed on transient signals.

Passive radiators and transmission lines are other purely mechanical ways of affecting the loading on the bass drive unit, and hence increasing bass extension. Their effect is similar to reflex porting, but less pronounced, generally adding 6dB/octave to the low-frequency roll-off.

2.1.3. Crossover design

The electronic crossover that distributes incoming audio to different drive units in a system influences the on-axis response of the loudspeaker. A carefully-designed crossover will be flat on-axis, but passive crossovers often create phase shifts between the drive units meaning that there is a pronounced dip or peak in the response at the crossover frequency.

Furthermore, even constant-voltage crossovers will produce transient ringing at the crossover frequency, which is more severe for sharper filters.

2.1.4. Case diffraction

Sound waves diffract around an obstacle whose size is smaller than its wavelength: the obstacle is essentially ignored. Above a certain threshold, the obstacle's edges present a sudden change in acoustic impedance to the pressure wave, and secondary waves propagate from the interface. The result at the listening point is a comb filter response. Loudspeakers with sharp edges suffer from the effects of diffraction more than loudspeakers with rounded edges, for which the change in acoustic impedance is more gradual. Cube-shaped loudspeakers, where all edges are the same distance from a single drive unit, are most severely affected.

2.1.5. Reflections from the back of an enclosure

The travelling wave from the tweeter is often reflected from the back of a loudspeaker enclosure. This reflection hits the rear surface of the tweeter a few hundred microseconds later, radiating an echo. Such an echo comb filters the tweeter's response.

All of these features are readily observed in typical loudspeaker frequency responses. Figures 1 and 2 demonstrate some of these aberrations.

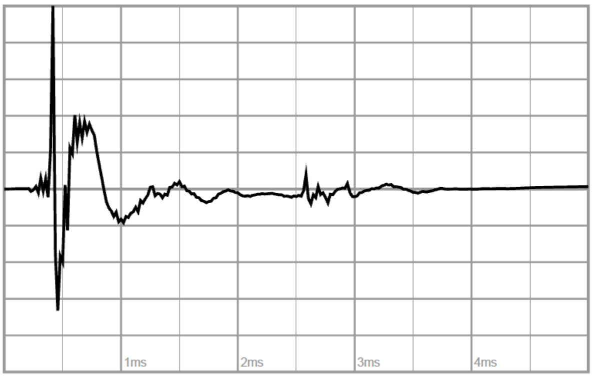

Figure 1. On-axis impulse response of a passive, two-way, sealed-box loudspeaker (Yamaha NS-10M) showing crossover ringing (the initial decaying sinusoid, before the 1.5ms mark), diffraction effects (the rippling around and after 1.5ms) and a back-of-cabinet reflection. Low-frequency effects are too small in amplitude, and too slow in time, to be seen at this scale.

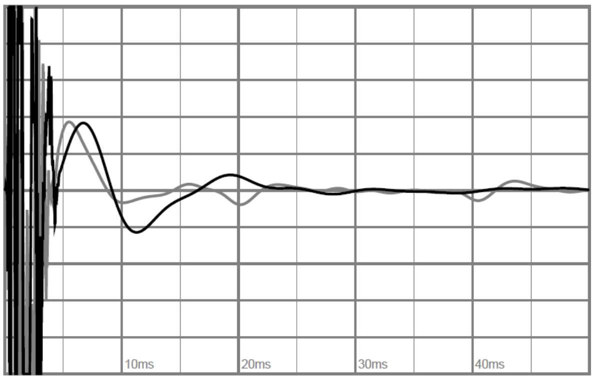

Figure 2. Extended impulse response for a sealed-box loudspeaker (Yamaha NS-10M: grey) and a ported loudspeaker (KRK RP6: black). The vertical scale is one hundred times that of Figure 1.

2.2. Factors affecting off-axis response

The loudspeaker's off-axis properties will now be considered. Generally, it is the energy radiated off-axis that reflects from the walls of the listening room, and is responsible for colouring reverberant energy.

2.2.1. Directivity versus frequency

Cone-shaped loudspeakers approximate a piston. Because these have a physical width, they cannot radiate omnidirectionally at all frequencies. Apparent phase cancellation caused by time-of-arrival differences between the nearest and farthest edges of the cone will cause local nulls in some listening positions. When the wavelength becomes high enough, this effect becomes very pronounced.

Above a certain frequency, standing waves will form on the loudspeaker cone and it becomes a phased array, focusing sound forwards at the expense of directional response. However, well-designed crossovers and stiff drive units prevent this happening in the active frequency range of the loudspeaker.

2.2.2. Panel resonance

A loudspeaker enclosure serves the function of containing the pressure wave created behind the drive units. Were it not to exist, the forward displacement of a loudspeaker cone would simply circulate air between the front of the drive unit and the back, and little energy would be radiated forward. In practice, no enclosure is perfectly rigid, and the loudspeaker's panels will resonate at different frequencies. The result is a non-ideal directional response, particularly at low frequencies. In extreme situations, this could also influence the on-axis response.

2.2.3. Acoustic shadowing

When the wavelength being radiated is less than or approximately equal to the size of the loudspeaker enclosure, the enclosure casts an acoustic shadow that increases the high frequency directionality of the loudspeaker.

2.2.4. Phase alignment of drive units

Phase differences caused by the signals from the different drive units of a loudspeaker arriving at slightly different times will affect the resulting frequency response when a listening position is closer to one drive unit than another. However, when loudspeakers are aligned correctly, this phenomenon does not generally affect on-axis listening; neither does it affect the total sound power radiated into a room. It is also unusual for this phenomenon to influence sound in the reverberant field, except in the case of strong, isolated early reflections. In our simulation, individual reflections are not modified for drive unit phase alignment discrepancies.

Many of these effects can be seen in the directional plots in Figure 3.

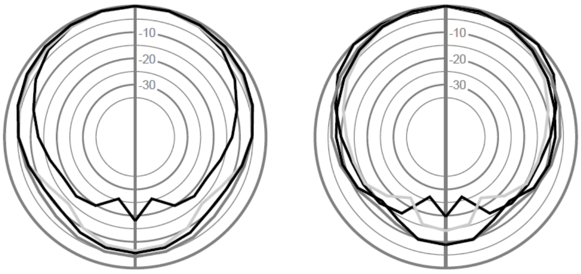

Figure 3. Third-octave smoothed directivity plots, measured for a small loudspeaker (Rogers LS3/5a: left) and a larger loudspeaker (Quested S8R: right). Directivity is shown at 200 Hz (black), 400 Hz (dark grey), 800 Hz (light grey) and 3.2 kHz (black).

In practice, the directional response of a loudspeaker will have a cone filter component and a low-pass component that is related to the size of the drive unit and the cabinet. Much of the energy that is radiated into the room is affected by this phenomenon.

For example, a typical softly-furnished living room may have dimensions of around 6×4×2.5 m, and a reverberation time of about 0.6 seconds. In this room, the critical distance, when there is more reverberant energy reaching a listener than direct sound energy, is only about one metre. A penchant for shiny floors and hard furnishings can easily reduce this distance to 75 cm. Listeners will tend to want a greater distance than this between themselves and their loudspeakers, so off-axis (or power) response becomes important. It is therefore essential to characterise the polar plot of the loudspeaker under test, and to deal with this information carefully.

Nevertheless, we do not need to be too fastidious when considering off-axis response. The law of the first wave-front, where the human auditory system suppresses early reflections, assists in mitigating the perceived problems of poor off-axis response. Furthermore, humans integrate loudness over critical bands of approximately a third of an octave, so any frequency distortion of a higher resolution than this does not really need to be considered.

It is for this reason that the loudspeakers' directional responses are recorded only in one plane. In reality, the responses in different planes are affected by the relative phasing of the drive units around the crossover frequency, and the physical dimensions of the loudspeaker cabinets. However, the average sound power output is not affected by drive unit phasing, and it is disregarded here. For this reason, we measured our loudspeakers in the horizontal plane, with the drive units vertically-aligned.

There are some situations where this technique is not valid. For example, certain loudspeakers are designed to radiate via circular reflectors, and are therefore diffuse and omnidirectional in the horizontal plane. In this case, the loudspeaker's response must be characterised in the vertical plane so that the reverberant field is excited properly. Therefore, these loudspeakers are best measured in the vertical plane (or tipped upon their sides).

2.3. Stereo matching

Cheaper loudspeakers do not use high-tolerance components in their crossovers. The quality control of materials used in the manufacture of such drive units may also suffer. Phase and frequency response discrepancies between the left and the right channels will result. These harm the stereo imaging ability of the loudspeaker system at mid-frequencies. For an authentic simulation of budget loudspeakers, it is therefore necessary to record information from all the loudspeakers in a stereo system to obtain these discrepancies.

However, we can safely assume that the directional properties of the loudspeakers will be similar, even if their on-axis responses differ, as these characteristics are overwhelmingly dominated by the spatial dimensions of the case and the drive units. Therefore, only the on-axis response of the individual loudspeakers needs to be changed.

2.4. Non-linear distortion

As well as the distortions detailed above, which can be simulated by convolution, there are a number of non-linear distortions associated with loudspeakers.

Power compression occurs over a period of time, as a drive unit's voice coil warms up, increasing its resistance. As this happens, the same driving voltage no longer moves the voice coil as far. This reduces the output sound pressure. The effect of power compression can become quite pronounced over a period of time. However, as power compression is effectively a memory effect (requiring knowledge of long-past signals), it is fairly onerous to include. It would also be cumbersome for a user to have to 'rest' our simulation in the same way that a loudspeaker might need to be rested to counteract power compression. Therefore, we have not simulated it in the auralization algorithm.

Harmonic, intermodulation, and other non-linear distortion is caused most commonly by four further mechanisms:

1. Loose or poorly-damped speaker components resonating (or bass ports chuffing);

2. The voice coil or suspension mechanism being pushed beyond its linear operation at high sound pressure levels;

3. The travelling acoustic wave tilting forwards owing to non-isothermal compression;

4. Doppler shift of high-frequency audio caused by lower-frequency audio radiating from the same cone.

Of these types of distortion, the first is a pathology of poorly-assembled or worn-out loudspeakers. For our purposes, it is more desirable to assume the ideal than to include the effects of a worn-out loudspeaker in auralization. The second is dependent on knowing the absolute sound pressure level at the loudspeaker, does not scale easily with volume, and does not generally occur at reasonable listening levels. For these reasons, its effects are also neglected. The third type is significant only in systems where very high local sound pressures are generated, such as in horn-loaded loudspeakers and large public address systems. It is neglected in this simulation because studio monitor loudspeakers generally do not operate under such conditions.

Doppler shift interests us. Its strength depends on absolute sound pressure level, but its effect can be derived from the following formula, linking cone displacement to physical properties of the drive unit and sound wave:

$$s = \frac{\rho\, p}{f^2\, A}$$

where s is the maximum displacement of the cone in metres, ρ is the density of air (approximately 1.2 kgm^{-3} at room temperature), p is the sound pressure in Pascal, f is the frequency of the input signal, and A is the moving surface area of the cone. For ported loudspeakers, or those with passive radiators, this formula will not hold, as the radiating area at low frequencies will be greater than the cone area. However, it allows for a good first approximation.

The effect of Doppler distortion on the bass drive unit is orders of magnitude more significant than for any other drive unit in a system. Audio that passes through the bass unit is effectively re-sampled, with the time delay at any time proportional to the absolute cone displacement. This can be applied to the input audio fairly efficiently using a bicubic resampling algorithm.

So, for example, a six-inch cone in a sealed box radiating 60Hz at 90 dB SPL has an excursion of (1.2 × 0.63) / (60² × 0.018) = ±1.2 cm. This is equivalent to ±1.6 samples of delay at 48 kHz.

Inspecting the relationship between input signal and displacement, we can derive a formula for transforming the input audio signal into instantaneous displacement:

1. Integrate the input audio twice, first by taking the running cumulative sum of the input samples, and then doing this again. This creates a signal whose value at any frequency is changed in proportion to $1 / f^2$.
2. Use a simple high-pass filter tuned to a very low frequency (typically below 1 Hz) to remove the dc that results from the first process.
3. Multiply the resulting signal by an appropriate scale factor to calculate the number of samples of deviation required to resample the audio data.

The scale factor that converts the doubly-integrated input signal to displacement in samples is found using the following formula:

$$k_d = \frac{\pi^2\, \rho\, p_0}{f_s\, A\, c}$$

where p_0 is the desired sound pressure (in Pascal) at unity input level, c is the speed of sound at room temperature (approximately 345 ms^{-1}) and f_s is the audio sampling frequency.

The great advantage of knowing this formula is that the effects of Doppler distortion can be reproduced over headphones based upon only the dimensions of the bass drive unit, the crossover frequency, and the intended peak sound pressure level. It can also be arbitrarily adjusted.

Using the equation given in Beers and Belar [2] , the distortion of a 1kHz sine tone played simultaneously from the same drive unit as the 60Hz, 90dB sine tone calculated above would be 1.30 × 1000 Hz × 0.012 m = 15%. This is above the 2% threshold of audibility for a pure tone that is stated in their paper, and also above that anecdotally stated by Klipsch [3].

However, two sine tones is a very critical signal for Doppler distortion, and one which is rarely encountered in normal listening circumstances. The audibility of distortion is highly dependent of the kind, as well as the amount, of distortion introduced. Informal listening tests revealed that the effect of Doppler distortion was unnoticeable on simulated systems with six-inch drive units, when the low-frequency roll-off of these systems was considered. This was true even when its influence was artificially doubled. We therefore left it out of the final run-time simulation.

2.5. Summary of section 2

It is important to measure three properties of a set of loudspeakers for realistic auralization:

1. The impulse response of the loudspeaker, so that both the frequency and phase response are accounted for. This allows bass ports, crossovers, and passive radiators to be simulated properly.

2. The directional response of a loudspeaker, so that the correct spectrum of energy is imparted to the reverberant field. This needs to be measured precisely, although fine frequency detail is not important.

3. The on-axis response of every loudspeaker in the system, so that any phase and frequency discrepancies that affect stereo listening can accurately be characterised.

We found that the non-linear aspects of loudspeaker reproduction could safely be neglected as either inaudible or undesirable within the intended context of our simulation.

3. MEASUREMENT TECHNIQUE

3.1. Generating the stimulus

To capture frequency and phase characteristics, the impulse response of the loudspeaker is required. A raw impulse is an impractical measurement stimulus, as it has very little energy in any frequency band, is hard to pass through the elements of a replay chain with accuracy, and is vulnerable to noise in the recording process. Practical measurement methods derive the impulse using other means. A noise-based stimulus, such as a maximum-length sequence, is tonally complex and of a high general output level. As a loudspeaker stimulus, it is vulnerable to non-linear distortion that cannot be cancelled by subsequent processing.

An approximation to our measurement stimulus is obtained by starting with an impulse response, passing it through a FFT, slicing the resulting frequency-domain rectangle into octave bands, and then passing each slice through an inverse FFT. A 16384-point FFT can thus be divided into thirteen non-overlapping octave-wide segments: the first accounting for frequency bin 1, the second for bins 2 and 3; the third for bins 4 to 7, and so on until the thirteenth segment, which accounts for bins 4096-8191. This series, arranged into order of ascending frequency, would take 4.44 seconds to play at f_s = 48 kHz.

We could directly synthesise these functions by adding sinusoids together of the appropriate levels. However, if instead of sinusoids, we use sin x / x functions, each slice becomes an audio signal defined by the following function:

$$F_S(n) = \sum_{b=S}^{2S-1} \frac{\sin(2\,\pi\,b\,n/16384)}{2\,\pi\,b\,n/16384}$$

where $F_S(n)$ is the audio signal for slice S (where S is a power of 2 up to 4096). $n = 0$ is the centre of the slice, where l'Hôpital's rule yields $\sin 0 / 0 = 1$. Unlike the FFT-based waveform, this function does not have periodicity, but the slices still sum to an impulse. The first three of these slices are shown in Figure 4.

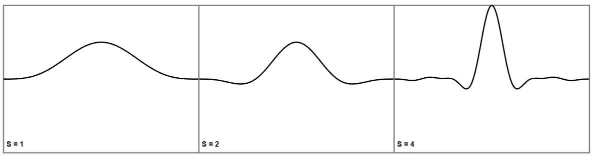

Figure 4. The first three slices, $F_S(n)$, for S = 1, 2, 4.

The usefulness of this series as a loudspeaker test signal arises from a number of its properties: it is band-limited to about an octave at any time, reducing Doppler distortion. It is spread out in time, and has a reasonable amplitude that decays symmetrically either side of a central peak. It is simple to replay, record, and recover the signal. Finally, the slices sum to form an impulse, so that retrieving an impulse response from the recovered signal is trivial. However, two further refinements are necessary before the signal is ready to use.

The first refinement is to equalise each slice. Taking octave slices means that the high-frequency segments contain more FFT bins, and hence more energy, than the low-frequency segments. One way of equalizing these is to scale each sin x / x component to $1 / f$, so that the sum across different bands is approximately constant. The resultant sum can then be differentiated to recover an impulse response. The method we chose is to boost each slice after it has been generated according to a $1 / f$ law, so that the scaling applied to each successive slice is half that applied to the previous slice. The reciprocal of this scaling is applied when the signal is recovered.

Our second refinement is to realise that there are still sharp transitions in the function at the boundaries of each slice. The derivative of the signal is therefore not smooth. These discontinuities could be windowed out towards the edges of each slice, but a more elegant approach is to extend each slice function into adjacent slices, so that its response is allowed to decay to a lower amplitude. In this implementation, the functions are continued in a circular manner so that presenting the train of slices more than once in immediate succession (so that measurement noise can be averaged out) continues the pattern. Mathematically, the slices will still sum to form an impulse, as the incoherent tails of each slice will cancel exactly.

The first few equalised and overlapped slices are shown graphically in Figure 5.

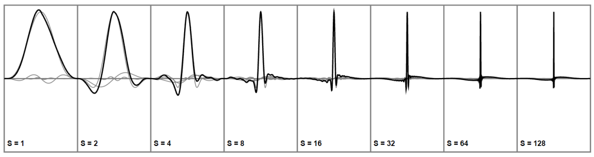

Figure 5. The first eight slices of the stimulus waveform. These slices are equalised, overlapped, and summed to create the test stimulus. The light grey traces are individual slice responses.

In our program, a 100ms 1kHz burst precedes the start of the measurement stimulus, which is then presented ten times. The 1kHz burst acts as a locator, and allows the analyser to correct automatically for any delays in the replay, acoustic propagation, and recording chain. This is especially necessary when recording the response behind the loudspeaker, when the first wall reflections usually exceed the direct sound in amplitude, and automatic time alignment would thus be very difficult.

For all the studio monitor loudspeakers we measured, peak replay level was set using a warble tone stimulus at 85dB SPL. This was made to correspond to -20 dB FS on our recording system to allow plenty of headroom. An Earthworks M30 omnidirectional reference microphone was used to measure the readings. This was oriented at 90 degrees to the loudspeaker in order to obtain the microphone’s diffuse field frequency response.

4. MEASURING LOUDSPEAKERS IN A SMALL ROOM

4.1. Measuring technique

The loudspeaker measurements were made in a small, acoustically damped room, a little smaller than 4×4×2.2 metres. This does not approximate an anechoic chamber: for example, it has strong floor and ceiling reflections. There are already established methods for dealing with these measurement conditions, which are reviewed by Vanderkooy and Lipshitz [5].

The path length of a single floor reflection for a loudspeaker mounted a height h above the ground and a distance d from the microphone is:

$$l = \sqrt{d^2 + 4h^2}$$

The maximum difference, in dB, that this extra path makes to the level of the received audio signal occurs if it arrives out-of-phase with the direct sound:

$$L = -20 \log \frac{d}{l}$$

At d = 1 m, and h = 1.1 m, the interfering reflection will create a 4.6dB dip in the frequency response. This formula assumes a perfectly reflecting floor. Nevertheless, it is clear that data recorded in a real reverberant environment is heavily compromised by early reflections. To reduce the maximum ripple caused by floor reflections to 0.5dB, it is necessary to record at no more than 12 cm. In practice, this distance can be increased by a few centimeters, as the floor is not a perfect reflector.

Our main readings are taken at 1.5 m from the loudspeaker with the microphone aligned with the tweeter, so that the single-reflection comb response would cause 7dB dips in the frequency response. At high frequencies, the ringing and resonances of the loudspeaker tend to die down quickly, so that the impulse response can be safely truncated by the time that the first reflection arrives, 3.3ms after the direct sound.

To obtain the low-frequency response, we combine this data with a near-field recording, taken 15 cm from the loudspeaker, with the microphone aligned with the woofer. There are problems with using such near-field responses, caused by the different relative distances of parts of the cone's surface and the cabinet edge at close quarters when compared with more distant measurements. Near-field responses have a shallower low-frequency roll-off and a less pronounced corner frequency than anechoic responses taken at a distance. However, the alternative methods for computing out reflections require a priori knowledge of the characteristics of the loudspeaker. We do not have this privilege when dealing with arbitrary loudspeaker systems, for which we possess no driver, cabinet, or crossover data.

Vanderkooy and Rousseau consider near-field responses and their shortcomings explicitly [4], and reject them for their purposes. However, their method is intended for critical measurement of loudspeaker system performance, and the theoretical low-frequency performance of their loudspeaker is already known. In our application, a room response will be imposed for auralization. The addition of room modes and the relative insensitivity of the human auditory system to very low frequencies allows the near-field discrepancies to be regarded as less serious, and in practice, the apparent quality of the resulting simulation is quite satisfactory.

Because the low-frequency response of a loudspeaker is generally devoid of significant directional anomalies, the same near-field response was used to correct measurements taken at every location around the loudspeaker. Thirteen directional far-field measurements were recorded for each type of loudspeaker, in fifteen-degree intervals over a semicircle in the horizontal plane.

4.2. Data processing

After recovering the impulses responses, the next stage is to combine near-field and far-field measurements. The flowchart for the entire process is shown in Figure 6.

The first step of this process is to gather data about the relative time and level alignment of the two measurements, so that the two responses can be combined properly. Near- and far-field responses are first windowed using a raised cosine function to remove all content after 3.3ms, when the first floor reflection reaches the microphone. The resulting data is then band-limited in the frequency domain to 500–800Hz. This treated data will be referred to as the *alignment responses*.

The peak cross-correlation of the alignment responses yields a delay value. This value can be used to synchronise the two impulse responses to the nearest sample.

An appropriate scale factor is also found to match the near-field and far-field levels. A number of methods were tried for finding this. The most reliable scale factor was obtained by calculating the sum-squared signal levels in the interval 500 microseconds either side of the peak sample in the two alignment responses, and dividing one by the other.

Finally, the FFT of the windowed far-field data is combined with the FFT of the time-aligned, level-corrected impulse response of the un-windowed near-field data. As the responses are synchronised to within a sample, this may validly be done at low frequencies by simply weighting and combining the lower frequency bins of the responses. A linear cross-over between 250Hz (beneath which only the near-field measurement is used) and 450Hz (above which only the far-field measurement is used) appeared to be suitable for every loudspeaker we encountered. Figure 7 shows recorded and corrected responses for one of the loudspeakers measured in this experiment.

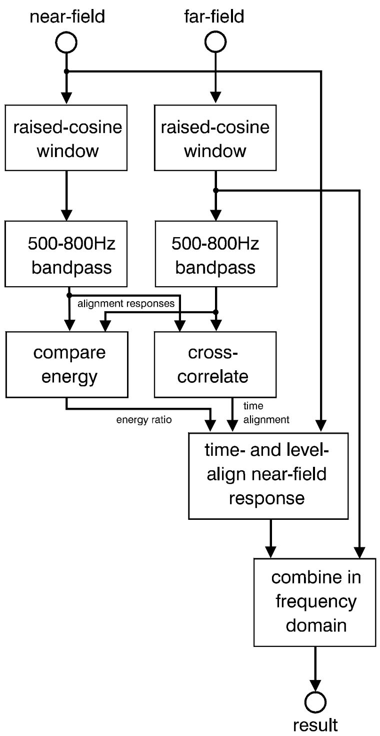

Figure 6. Flowchart for near- and far-field data combination.

5. SUMMARY

This paper presents an automatic method for measuring the impulse response and directivity of loudspeakers in a normal room. This involves generating and processing a test stimulus that avoids exciting non-linear distortion, and using near- and far-field loudspeaker measurements to compute out the room reflections in a manner that is accurate enough for auralization.

Although we present a method for applying Doppler distortion to the simulation, it was decided that this does not make enough difference in practice to be worth including in the run-time simulation.

Subjectively, the resulting auralization works well.

REFERENCES

[1] Watkinson, J. *The Art of Sound Reproduction.* Oxford: Focal Press, 1998.

[2] Beers, G.L., and Belar. H. 'Frequency-Modulation Distortion in Loudspeakers'. *J. Aud. Eng. Soc.*, 29(5), May 1981. pp.320–326.

[3] Klipsch, P. W. 'Modulation Distortion in Loudspeakers'. *J. Aud. Eng. Soc.*, 17(6), April 1969. pp.194–206.

[4] Vanderkooy, J., and Rousseau, M. 'A Study of Low-Frequency Near- and Far-field Loudspeaker Behaviour'. *126th AES Convention*, May 2009. Convention Paper 7747.

[5] Vanderkooy, J., and Lipshitz, S. P. 'Can One Perform Quasi-anechoic Loudspeaker Measurements in Normal Rooms?' *124th AES Convention,* May 2008. Convention paper 7525.

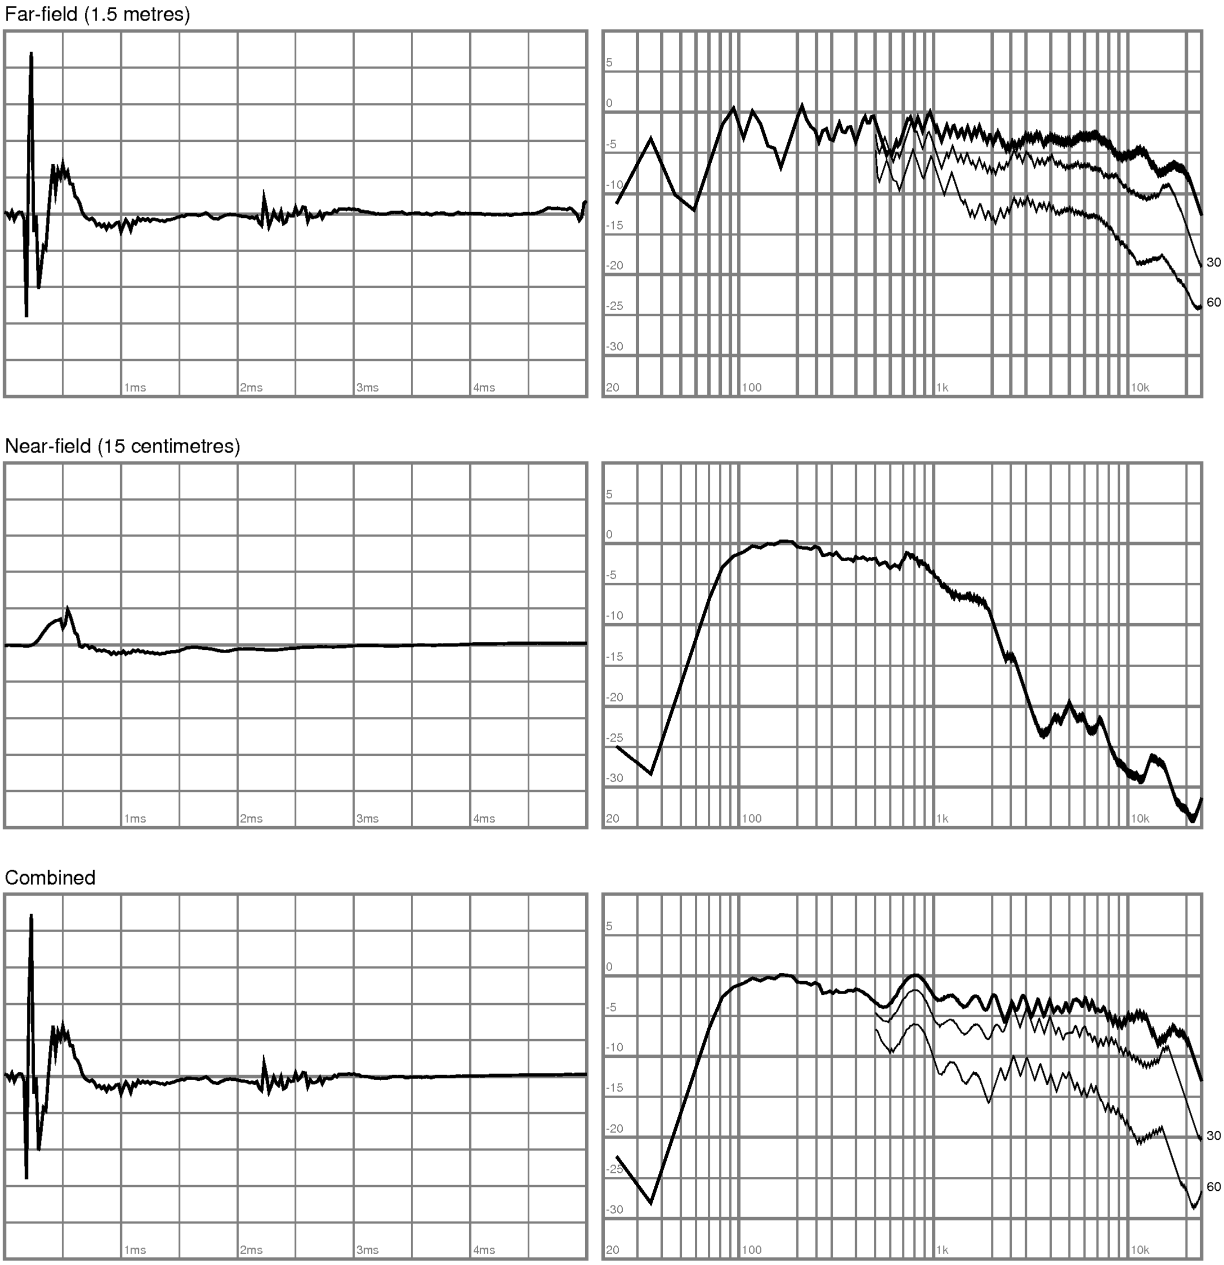

Figure 7. Impulse and frequency responses for the far-field, near-field, and combined responses for a measured studio loudspeaker (Genelec 1031A). Frequency plots are also shown for 30 and 60-degree off-axis responses.