

The LEOPARD Project

© John Meyer, 2015

1 Introduction

We developed the LEO family of speakers to produce the most ideal linear speakers possible: that is, to reduce distortion to the theoretical limit. LEOPARD is a small line array element in the LEO family that wasn't targeted to a specific market. Instead, we gave it to our engineers as a challenge to see how well it could be designed.

What do we mean by an ideal loudspeaker system? An ideal loudspeaker would reproduce sound which a listener would find indistinguishable from that coming directly from a sound source.

In short, LEOPARD is more powerful and has less distortion than a comparably sized loudspeaker. Before discussing how this was achieved, we must start with how we perceive and measure sound. When describing LEOPARD, we also discuss loudspeaker theory and history so that this achievement can be seen in context.

2 Background: Measuring Tones & Notes

What do we know about sound to guide us toward creating the perfect loudspeaker?

We perceive sounds because our eardrums (tympanic membranes) move back and forth in sync with the minute pressure variations created by all sound sources, including musical instruments and voices.

The vibration creates a wave in the air that travels away from the source at a speed of about 1130 feet/second at sea level and room temperature. All sound frequencies travel at the same speed in air.¹ Airborne sound waves also behave like electromagnetic (radio or light) waves, which means that fundamental laws of physics can be used to predict its interaction with a solid object, such as a wall.

¹ Note that this is quite different from the case of surface water waves, which are nonlinear and behave differently than pressure waves in air. For instance, waves in water with a large spacing between the crests travel slower than waves with a shorter spacing.




Figure 1. Ideal frequency spectrum of the note A (440 Hz) played on a piano after reaching steady state.

The notes of musical instruments are a combination of steady state pure tones, which exist in a “time invariant space.” Everything else—onset, offset, frequency shifts, gain changes, limiters, compressors, etc.—is in the “time variable space.”

What do the terms “notes” and “pure tones” mean? Let’s look at the tone A440. A tuning fork that is marked A440, when struck, produces a compression wave where the air pressure goes back and forth 440 times every second. We call this a pure tone or a sine wave frequency of 440 Hz.

A note, however, may contain a series of tones. For example, a 440 Hz note on a piano might be made up of a 440 Hz tone (the fundamental) plus some amount of 880 and 1320 Hz tones (the harmonics). In actuality, piano notes can have more than 10 harmonics. Figure 1 shows the amplitude versus frequency spectrum of a theoretical 440 Hz piano note that has reached steady state.

Steady state sound sources do not change their frequency (pitch) or amplitude (loudness) over time, or at least not during the recorded time record. Therefore we need to let a note or tone settle to a steady state condition before measuring it.

If single notes on a piano are recorded with a pressure-sensitive microphone, the notes can be clearly seen in the time record because each note fades away fairly quickly (Figure 2). If we want to find the frequency of a note, we have to first isolate it, as with the 300 ms time record shown in the lower trace.




Figure 2. Time record of a series of single notes played on a piano (above), with an expanded view of a single note (below).

The Fourier transform of the isolated note's time record shows a spectrum that doesn't look much like our ideal piano note (Figure 3). The problem is we have captured the initial strike of the hammer hitting the string along with the resulting




Figure 3. The intensity spectrum of the piano note shown in the bottom of Figure 2. Note differences from the ideal spectrum of Figure 1.




Figure 4. Scott's phonautograph, introduced in 1860 to measure the sound waveform.

decay in amplitude. As you can see, the waveform changes constantly except over very short time periods, and there's no obvious place where the note reaches steady state. This blindness to temporal changes is one of the disadvantages of spectrum analysis, and worth remembering whenever you're looking at a spectrum analyzer. By contrast, our ears are very good at hearing temporal changes. Looking for the note that is buried in the time record is still useful, though, and will be covered in another paper about source independent measurement (SIM); however, many measurement systems and processors routinely throw out the 'stuff between the notes' as noise.

Helmholtz showed how to create notes using multiple single tone sources in the 1850s. For example, we could begin to simulate a piano note by choosing tuning forks with the correct frequencies. Let's start with three separate sound sources, spaced a few feet apart, where each produces a single steady pure tone at 440, 880 and 1320 Hz. Now let's take a measurement of the rapid pressure variations in the air at a point a few meters away. Helmholtz used Scott's phonautograph (patented in 1857), which scratched a pressure wave onto a turning cylinder to visualize the time record of the note (Figure 4). This isn't very different from modern recording equipment or the human ear, which are also pressure sensitive.

Helmholtz found that the pure tones from each source add together in the air to create a combination wave shape (Figure 5). Suppose that you happen to choose a measurement position where the three tones (sine waves) all cross zero at the same point in time (as on the left hand side of the plot). We say that the tones are in sync, or in phase. At another measurement position, the resultant wave will have a different shape (Figure 6). This is because the




Figure 5. A 10 ms time record of three pure tones, 440, 880, and 1320 Hz, in phase. The bottom trace is what you would see on the phonautograph.

distance from each sound source has changed, affecting the phase relationship between the tones. In this case, there is no moment where all tones cross zero at the same time. We now say that the sine waves are *not* in phase, or that the harmonics have phase shift, or phase offset.

Helmholtz observed that phase offset between the tones didn't change how the note sounded. Both notes we created above sound the same. This is where the notion that we are phase deaf first arose. Helmholtz was careful not to make this claim, however, and pointed out that he wasn't able to study nonharmonic sounds at that time. As we know now, this is a special case because the notes are steady state and very simple. More generally, we are not phase deaf for




Figure 6. A 10ms time record of three pure tones, 440, 880, and 1320 Hz, out of phase, showing a different waveform shape than Figure 5.

signals that change with time.

Helmholtz's research on musical notes developed from earlier work by Fourier; this approach is now used extensively and is called Fourier or spectrum analysis. It uses a mathematical function called the Fourier transform to convert a steady state sound wave in the time domain to the frequency domain, where each frequency has a magnitude and phase. Remember that the Fourier transform assumes that the sound source never changes; it's been running forever and will run forever. This is important because spatial relationships in the steady state sound sources are lost, as we will show later on. What we learned here is that, in order to accurately reproduce a steady state note, we need a loudspeaker that is capable of reproducing all the frequencies at the correct loudness.

But that's not all that's required from good loudspeakers. They also need to reproduce the phase accurately in order to reproduce more complex sounds. Helmholtz's experiments show that phase is in many ways equivalent to timing differences between frequencies. Any sound that changes over time needs to be reproduced with a good phase response.

Our three sound sources producing a 440 Hz note is one such example. Keeping the source producing the fundamental tone close to our measurement microphone, let's move the 880 Hz source 500 feet away and the 1320 Hz source twice as far. Now, in order to have the same loudness at the microphone, we need to increase the loudness of the distant sources. After letting the sources run for a while (into steady state), we see the same pressure waveform as before, when all the sources were close. In fact, there is no way to tell in the time record that a source is 500 or 1,000 feet away as long as it produces a continuous tone. However, if we turn off all the sound sources at the same time, the harmonics can be heard for up to a second after the fundamental frequency fades away.

Reproduction of a drum or firecracker provides another example of why phase matters. The pressure waveform in this case looks like a sharp spike that Fourier analysis shows to contain all frequencies that align at a single point in time. The waveform will be changed if any frequency is delayed and arrives later than the others, such as when a loudspeaker reproduces the sound. This kind of phase delay can modify the tonal and transient response of the drum. Note that it's very easy to hear the effect of phase offset in this case because this is not a steady state condition—pressure changes very rapidly over time.

3 LEOPARD

Now let's turn to how loudspeakers produce sound, starting first with frequency response and then moving to discuss power, distortion, and directionality. In each section, we'll discuss how that aspect of LEOPARD was designed.

3.1 Frequency Response

What is meant by a speaker system's frequency response, or transfer function? A "transfer function" describes the nature of the system under test, not how it's measured. For instance, the output of a device with a transfer function of 2 will be double its input.

The transfer function frequency and phase response of a speaker system describes the level and timing of all the tones the system will reproduce relative to the input signal. The speaker system must reproduce all tones from at least 32 to 16,000 Hz to ensure that it can play all forms of music. Since there are an infinite number of frequencies between 32 and 16,000 Hz, we need to define a practical limit of resolution. With 64 tones/octave, we get 576 tones for these nine octaves, which should cover all forms of music.

We could send one tone at a time to measure the loudspeaker. Each tone needs to be played for at least one second to ensure the steady state response is reached. The whole test with 576 tones takes about 10 minutes this way. The power or amplitude response tells us the tones the system can reproduce, but doesn't say anything about their phase, or when they will be reproduced. To obtain the transfer function phase response, a dual channel analyzer is required; it compares the sent and received tones after first subtracting the propagation delay from the speaker to microphone.

The frequency response has been calculated since the 1970s with the Fast Fourier Transform (FFT), an efficient computer algorithm that allows practical implementation of the Fourier transform. Results can be obtained in seconds by driving the system with a complex signal, such as a mix of all 576 tones. The same response can be acquired from either music or random noise with our SIM (Source Independent Measurement) analyzer (which won an R&D 100 award in 1992).

Such measurements can verify that the speaker reproduces all frequencies at the same level. We must also monitor the phase response to ensure that time varying signals like percussion are reproduced accurately, which will be the case if the phase response is zero at all frequencies. The question then becomes, "How important is the phase response?" If you want to reproduce a signal with sharp transients, it is very important.

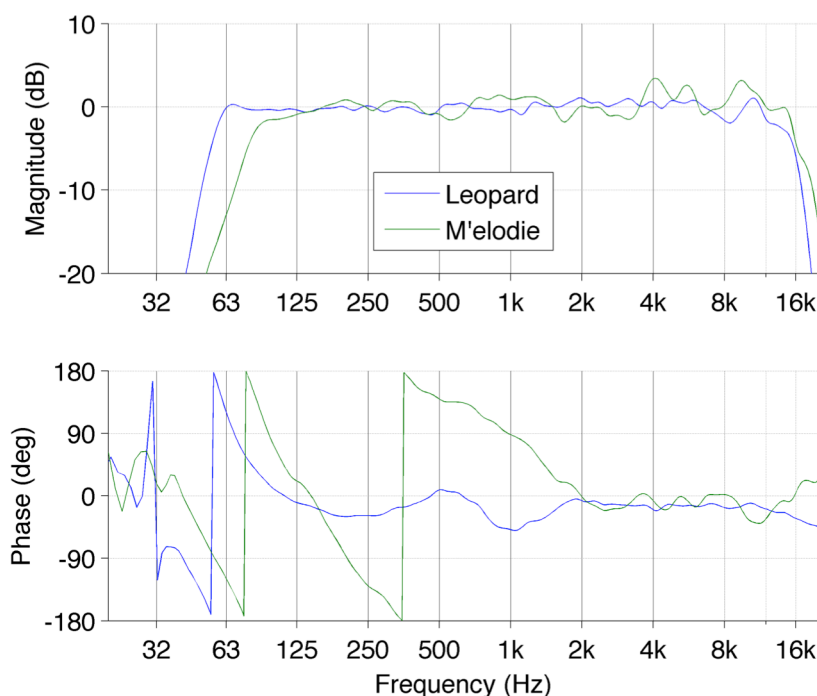


Figure 7. Frequency response of M'elodie and LEOPARD at 2 m; the magnitude response is shown in the upper panel, and the phase response in the lower. All responses have 1/6 octave smoothing. Note that this response was used to design LEOPARD; its released response is different.

Cinema loudspeakers in the 1930s had large amounts of phase distortion, because it was then widely believed that our hearing was phase insensitive and that the phase offset shouldn't affect sound quality. Nonetheless, moviegoers complained about the poor sonic image of tap dancing.

Now let's look at the frequency response of a real system. The M'elodie loudspeaker set a benchmark in sound quality when it was introduced. It has a very flat magnitude response—all tones are equally loud—and the phase response is flat and zero above 2 kHz (Figure 7). In comparison, LEOPARD has an equally flat magnitude response, but its phase response is effectively zero from 125 Hz and above, which is a remarkable improvement.

The low frequency phase response is very difficult to make zero. This is a fact of physics related to how small a low driver is compared to the size of the wavelengths that it reproduces. A 100 Hz tone has a wavelength of 10 feet, ten times larger than a 12" driver. One solution would be to make a 10-foot driver, but that has many obvious drawbacks. A reasonable compromise is to make the phase zero above a certain frequency, typically around 200 Hz.

Be careful when comparing the phase response between two loudspeakers: the comparison is only valid when the magnitude response is the same. This is because any change in magnitude, such as produced by equalization, produces a change in the phase response.

Also, note that the frequency response shown in Figure 7 is what the engineers used when they were designing a single LEOPARD loudspeaker. Since LEOPARD is meant to be used as part of a line array, never alone, the final frequency response is one that gives the ENTIRE line array a flat response. Making loudspeaker arrays sound consistent is a complex process and will be covered in another paper.

The loudspeaker's job is to reproduce not only all tones, but also all the information between the notes, which includes transients and non-steady state events where phase is crucial. Music can only be reproduced accurately by speaker systems with a flat magnitude and flat phase response. The next aspect to consider is distortion, which—if it isn't controlled—can make a speaker with even a perfect frequency response sound bad.

3.2 Power

But before we describe distortion, we need to first determine how loud a loudspeaker must be, because distortion increases with level. This is a crucial part of designing any loudspeaker, especially for a high-power model like LEOPARD.

Power, or loudness, is related to the sound pressure level (SPL), which is a root-mean-square (RMS) measurement averaged over time that can be computed by almost any sound pressure meter. Because power is very difficult to measure directly, most loudspeakers are rated instead by the maximum SPL they can produce at a certain distance.

One important point about power is that it's independent of the signal being generated. With a sine wave, all the power goes into a single frequency. With pink noise, the power is spread out over time and the audible frequency range. The total power can be the same in both cases.

The signal type also determines the required headroom, or the difference between peak and RMS levels. The peak level is an instantaneous measure; it's

3 dB greater than RMS for a sine wave and 12 dB greater for pink noise.² A loudspeaker must reach the peak levels without clipping (which produces large amounts of distortion) to accurately reproduce a signal. Even if a speaker can track a sine wave, there's no guarantee that it can reproduce pink noise at the same level, because the latter requires more headroom.³

LEOPARD was designed to be twice as powerful as M'elodie. Therefore, it needs to reach 106 dB SPL at 2 m with 12 dB of headroom to reproduce pink noise. LEOPARD's two LF drivers must each produce 100 dB SPL.

The prototype LEOPARD 9" LF driver can produce a measured output of 106 dB SPL for a 90 Hz sine wave at 2 m. In comparison, a generic 8" LF driver produces 102 dB SPL under the same conditions.

3.3 Single-Tone Distortion

A perfectly linear driver is defined as a loudspeaker piston whose motion precisely tracks its input. Any deviation from linearity is called distortion, and is never desirable. But even a completely linear piston will react with the air to produce distortion. This is the theoretical limit of distortion for any airborne sound.

In general, air distortion arises because pressure can increase indefinitely, but can't decrease below a vacuum (zero); that is, linear motion doesn't result in linear pressure changes.

More specifically, consider the case where the speaker is a piston driving a closed space, as shown in Figure 8. When the piston goes forward halfway, the volume decreases by 50% and the pressure in the space doubles. When the piston goes back the same amount, the volume increases to 150% and the pressure decreases to 67%, not 50%. This phenomenon was described and calculated by Olson in the 1940s.

² Engineers prefer to think about and measure the RMS level because it relates directly to power. Peak levels don't, which makes it much harder to compare performance across different signals.

³ While pink noise requires more headroom from a system, it will mask distortion. As a rule of thumb, it's impossible to detect distortion of less than 10% using pink noise.




Figure 8. Demonstration of air distortion. A piston drives a closed space; equal displacements forward and backwards produce the same change in volume ($\frac{1}{2}$ times here), but not the same change in pressure.

Now let's consider harmonic distortion, which we have some control over. We don't want the system to alter tones, for example, by adding its own harmonics. This arises because the driver cannot move too far forward or back without damaging itself (often called clipping distortion).

Some harmonics are more audible than others. For instance, almost all musical notes include some 2nd harmonic (twice the fundamental frequency). The 2nd harmonic is hard to detect, and is only noticeable when there's more than 3% present. Higher harmonics are easier to detect: 2% for 3rd, 1% for 5th or 7th. However, frequencies that are not harmonically related to the original note are detectable at as little as 0.1%.

Harmonic distortion is straightforward to measure: send a pure tone to the test system and look for any harmonics. The total harmonic distortion (THD) is the total power of all harmonics relative to the power of the fundamental. The tone can be made loud enough to generate 10% THD, which is considered the maximum distortion before losing control of the system. At this point, the LEOPARD driver is 4.5 dB louder than the generic or M'elodie systems. Be careful with these numbers, though, because this test doesn't indicate what happens below the 10% distortion threshold. For a good driver, distortion would drop quickly as the level is reduced; for a poor driver, distortion would drop slowly.

Let's look at harmonic distortion in the custom-built LEOPARD driver. Figure 9 shows the spectrum of a complete LEOPARD system (high and low drivers, amplifier, and crossover) producing a 500 Hz tone at 106 dB SPL at 2 m. Note that the measured 2nd harmonic distortion equals the lower limit of air distortion: 0.2%. A loudspeaker in air cannot produce 106 dB SPL with any less




Figure 9. Spectrum response of a complete LEOPARD system driven with a 500 Hz tone to 106 dB SPL and measured at 2 meters. Note that the measured 2nd harmonic distortion is right at the 0.2% limit of air distortion.

distortion. When a loudspeaker system produces 30% third harmonic distortion, suppressing the harmonic is very difficult. Adding a filter upstream of the speaker at the frequency of the third harmonic won't have any effect, because the harmonic is generated downstream at the loudspeaker. I encountered this problem once at a STYX concert in the 1980s, and it was a very frustrating experience. The only way to suppress the harmonic is to reduce the level of the fundamental with a filter or by reducing the system loudness. Linear filters only work properly on linear sound systems. A corollary is that it's easy to add harmonics to a linear system.

3.4 Multi-tone Distortion

Distortion resulting from interactions between tones is harder to measure than the single-tone variety, so we'll need to use special test signals for each different kind of multi-tone distortion.

Let's start with frequency modulation (FM) first, another fundamental limitation for loudspeakers. Even a perfectly linear loudspeaker will produce FM distortion, also termed Doppler distortion or flutter, whenever two or more tones are reproduced simultaneously. The FM distortion occurs because the velocities of the two tones add when both are moving in the same direction (and subtract when they are moving in opposite directions).

If we send 80 Hz and 440 Hz tones to a loudspeaker, the 80 Hz signal will frequency modulate the 440 Hz carrier frequency. If the piston of the low

frequency speaker moves with a peak excursion of 0.5 cm, then the peak velocity is $2\pi \times (\text{modulating frequency}) \times (\text{peak excursion})$, or

$$2\pi \times 80 \text{ Hz} \times 0.5 \text{ cm} = 251 \text{ cm/s.}$$

The amount of frequency shift is given by the (carrier frequency) \times (peak velocity) / (speed of sound). For this example, FM distortion is

$$(440 \text{ Hz}) \times (251 \text{ cm/s}) / (34400 \text{ cm/s}) = 3.2 \text{ Hz,}$$

which is considered below the threshold of audibility.

The 440 Hz tone shifts 3.2 Hz in frequency due to the motion of the low driver. This flutter is a rapid frequency change back and forth (80 times a second in this case). It is not a tone, it is a changing frequency much like vibrato. In fact, FM distortion is the same for all carrier frequency levels. In this case, excursion at 80 Hz alone determines the flutter frequency.

A linear loudspeaker will also reproduce a “beat frequency” from two closely spaced tones. For example, 440 and 450 tones will create a 10 Hz beat frequency. This 10 Hz sound is not a true tone because the speaker doesn’t create this frequency; its perception is created in our ear. In this example, a 10 Hz tuning fork would not vibrate along with it.⁴

Now let’s look at a kind of distortion that we can reduce by good design: intermodulation distortion (IM), which is created when two frequencies interact nonlinearly with each other. The signature of IM is in-phase steady state tones flanking the upper tone. If we send two frequencies, 80 and 440 Hz, to the non-linear loudspeaker, we will see frequencies on a spectrum analyzer at 80, 440, $440-80=360$, and $440+80=520$ Hz. These IM products, as they’re called, have harmonic distortion of their own, which then in turn create more IM products with all the other frequencies present. It has been known since the 1950s that IM distortion creates a “veiled” sound quality.

IM and other kinds of distortion can be shown by combining several sine tones into a more complex waveform shape. A good test signal is produced by adding tones of 80, 100, 125, 225, and 440 Hz in phase (Figure 10). Their sum has an RMS amplitude 9.1 dB below the peak level. These frequencies are chosen

⁴ To make the 10 Hz into a real tone, you would either need to increase the SPL to the point where the air is very nonlinear, or use a very nonlinear driver. This is called heterodyning.




Figure 10. Time record of the 5-tone signal for distortion measurements, containing 80, 100, 125, 225, and 440 Hz tones.

because any distortion created will have discrete spectral lines in the frequency domain.

Figure 11 shows three sound systems driven with our 5-tone test signal: LEOPARD, M'elodie, and a generic system. The complex 5-tone signal is played at the loudest sound pressure level that doesn't engage the protection limiters on M'elodie, which is 100 dB SPL RMS (109 dB peak) at 2 m. Notice that the

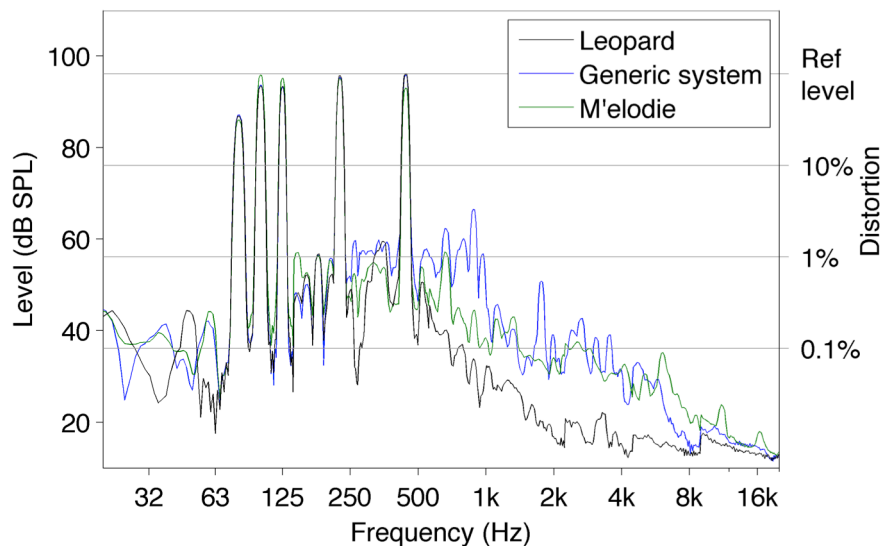


Figure 11. The spectrum response of three loudspeakers to the 5-tone test signal from Figure 10. The tone is played at 100 dB SPL RMS and measured at 2 m. The higher the baseline between the tones, the greater the distortion.

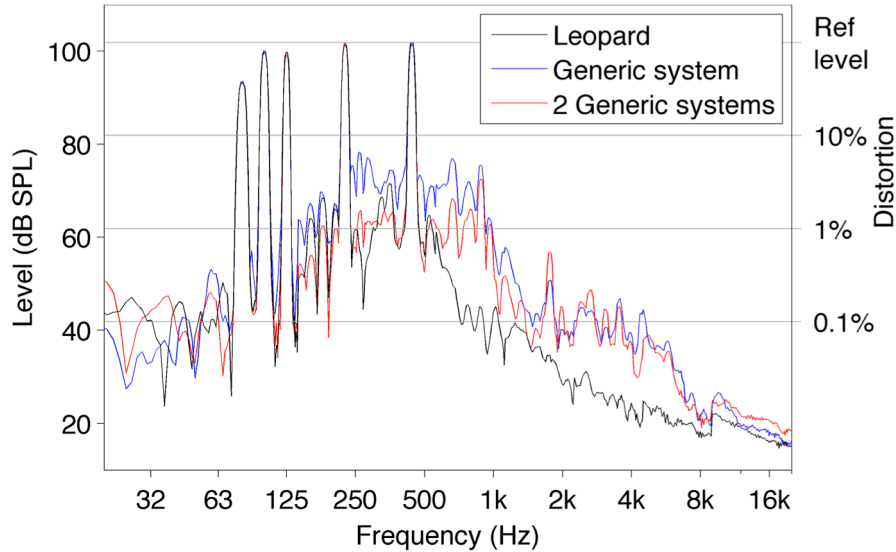


Figure 12. Spectrum response to the 5-tone test signal played at 106 dB SPL (115 dB peak) at 2 m. Each of the two generic systems is driven at -6 dB so the resulting level stays at 106 dB SPL.

generic system has the greatest distortion from 500-1000 Hz, but matches M'elodie's distortion from 1-4 kHz. LEOPARD has much less distortion than either.

Recall our goal that LEOPARD is to be 6 dB louder than M'elodie. We could reach this by using the same power amplifiers as M'elodie, but drive them 6 dB louder much like the generic system. Unfortunately, doing this would also produce more distortion.

Instead we designed the LEOPARD drivers specifically for this higher output level. This can be seen in Figure 12, where we increased the level to 106 dB SPL (115 dB peak) at 2 m, again without engaging the protection limiters.

At this louder level, the generic system has 10 times more distortion and is clearly in a very non-linear state. A rule of thumb is that distortion drops proportionally with level; therefore, we can expect about a 10-fold reduction for a 20 dB drop in SPL. For instance, a mic specified to have 1% distortion at 140 dB SPL has an estimated distortion of 0.1% at 120 dB. This is only true if the system is well-behaved. Interestingly, when we drop the level of the generic system by 6 dB, the distortion drops as expected below 500 Hz but not above. This means that two generic systems together, each reduced by 6 dB, would still produce more HF distortion than one LEOPARD system.

This 5-tone test reveals non-linearities not detectable with a single tone. Going back to Helmholtz's experiment, the multi-tone waveform will change if the phase

of any one component changes. We could, for example, adjust the phases to make a test signal with less peak level but the same RMS level. A lower peak means that the driver will produce less distortion in a non-linear system (Perrin Meyer and David Wessel). This is one way to test for system distortion.

We can't cover all the nonlinear mechanisms of drivers, but the improvements, not surprisingly, come at a higher cost: larger magnets and dual voice coils, each driven by a separate amplifier.

Before leaving the low driver, we need to determine the upper frequency limit. This is ultimately constrained by FM distortion (FMD), which increases with the frequency range (lower to upper) produced by the driver. For example, a LEOPARD driver with full excursion at 50 Hz also producing a tone at 2 kHz has FMD of 18 Hz. Lowering the upper tone to 600 Hz reduces the FMD to 5.4 Hz.

We must cross over at or below 600 Hz to keep the FMD < 6 Hz. SPL can be increased by either enlarging the piston, increasing its excursion, or both. Since we are already at the maximum excursion for an acceptable level of FMD, our only choice is to increase cone area. One way to do so is to use two LF drivers.

4 Directionality

4.1 Low Driver Polar Response

Now that distortion has been covered, we next consider the directionality of loudspeakers, or their polar patterns. This is crucial for building loudspeaker arrays that sound consistent over their coverage area. We begin with two LEOPARD low frequency drivers mounted next to each other.

In Figure 13, we are outside looking down at two sound sources mounted on a baffle wall and pointed into empty space. No other walls are present.

When we map this sound pressure field onto a polar chart, we can plot a line connecting equal SPL points, which is called an isocontour (Figure 14). This is similar to the isobar on weather maps. The deep side lobes where the response drops to zero are more noticeable here than on Figure 13. These lobes are caused by destructive interference where the sine waves from each loudspeaker are 180° out of phase and cancel each other completely. The sine waves are in phase and add together in constructive interference in the center of the lobe. This lobe/null structure means that this tone disappears at certain angles, which is the worst case possible for a loudspeaker.




Figure 13. Sound pressure field for two drivers in a baffle wall (shown on the bottom), with the drivers spaced 2.5 times the wavelength of the pure tone driving them. In this case, the frequency is 1.4 kHz, and the spacing is 2 feet.

The side lobes will finally disappear when the driver spacing is reduced to half a wavelength (Figure 15). There are two ways to view a polar plot like this, which can be seen by paying attention to the -3, -6 and -10 dB points marked with dots.




Figure 14. Polar plot showing the equal SPL response of two sound sources, spaced as in Figure 13. The sound pressure is shown at a fixed distance as a function of the angle, and is normalized to the level at 0° (Olson).

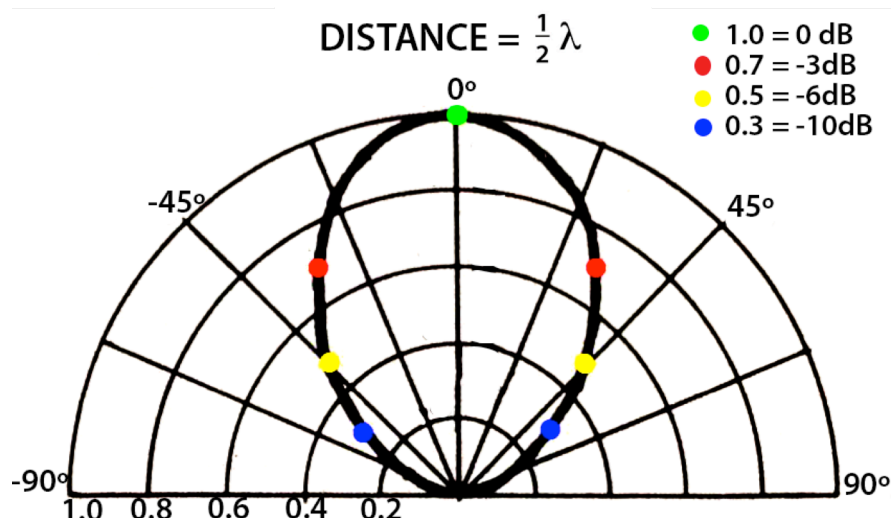


Figure 15. Polar plot showing equal SPL contours for two sound sources spaced a half wavelength apart.

You would measure the same SPL everywhere if you walked along this line. The further you walked toward the side (90°), the closer you'd get to the drivers. Think of the 1.0 or unity point on the polar chart as the distance from the sound source. We'll use 1 meter at 0° as the unity reference. At 45° off center, you'd be 0.5 m from the speaker (since -6 dB is equivalent to a factor of $\frac{1}{2}$).

Alternatively, you would see the SPL drop by the same factor (6 dB) if you walked on the outmost circle. At 45°, you would measure half the SPL (-6 dB), and at 55° the level would be 0.3 (-10 dB). Therefore, if we had 106 dB SPL at 0°, we would see 100 dB SPL at 45°.

As an aside, there's a huge difference between a polar plot with a linear scale and one with a logarithmic scale. Log scale makes the main lobe look wider and the rest of the plot look narrower. Also, a log scale means the polar plot has no direct physical interpretation; one cannot overlay the log plot on a linear space, such as a room, or walk the line and get an equal loudness contour as described above.

The width of this polar pattern is called the beamwidth (BW). In the 1950s, the audio industry defined the beamwidth of a directional source as the total angle between the -6 dB points. In the case of Figure 15, the BW would be 90°. ⁵

⁵ Other disciplines that use waveguides, such as microwave or underwater acoustics, use the half power point of -3 dB.




Figure 16. Mechanical design of LEOPARD system.

Be careful, though: some speaker manufacturers define beamwidth as the angle between the -10 dB points. With that definition, the BW of Figure 15 would be 110° , or 20° wider. We've shown you three ways to look at the same pressure field: the 2-D pressure plot, a polar plot, and beamwidth (discussed more in Figure 19).

So with the two-foot spacing from Figures 13 and 14, the highest frequency without side lobes is 280 Hz. This spacing allows an HF horn to be mounted between the two LF drivers. However, we cannot run both drivers together above the lobe limit of 280 Hz. Our past designs have run both LF drivers together below 280 Hz and only one driver above this frequency. A single 8" driver can be used up to 2 kHz, above which the single driver becomes too directional. This design was used for the M'elodie speaker system and gives a lot of leeway in where to cross over to the high frequency horn-driver.

An alternative approach is to run both drivers up to 500 Hz in order to maintain a high SPL output. To avoid any side lobes, their effective spacing would need to be about half the wavelength of 500 Hz, which is 1 foot. This was accomplished by the use of an innovative design (Figure 16). Ripple in the low frequencies introduced by the HF horn is solved by special electronic filters in the LEOPARD's integrated power amplifier.




Fig 17. Polar response of three different exponential horns to five frequencies (Olson).

4.2 The History of High Driver Polar Response

Now with the low driver fully specified, let's turn to the high driver. Here, horizontal coverage is controlled by a horn (or wave guide). We wanted our horn to have a constant bandwidth from 500 Hz to 16 kHz. This has been very difficult to achieve in the past, but the use of high power computers for modeling has led us to a novel solution for LEOPARD. Before we describe that, we should review the history of horns.

Let's look at some polar charts of exponential horns developed in the 1930s (Figure 17), which are useful for narrow coverage. Unfortunately they have high distortion levels.

The next designs incorporated an array of exponential horns into single units called multi-cellular horns, introduced in the 1940s (Figure 18). This horn array has a width of 32". You can see the polar response is narrowest around 500 Hz.

Beranek introduced a beamwidth chart for multi-cell horns, measuring BW as the angle between the -6 dB points (Figure 19). With this kind of plot, the narrow pattern in the mid-range is very obvious. This is where the width of the horn approaches the wavelength. What is lost with this kind of plot is any lobe structure beyond the -6 dB points.

One last horn of interest is the parabolic horn (also known as a radial horn), which has an exponential flare in the vertical plane with two straight sides for the horizontal walls (Figure 20). A variation is the diffraction horn, where the vertical




Figure 18. Polar response of a multi-cell horn to six frequencies (Olson).

walls are parallel, which has the same horizontal polar pattern. However, diffraction horns produce the highest amount of distortion of all horns. These horns have a narrow coverage angle around the wavelength of the horizontal width of the horn.

The next 50 years produced many horns that were considered to have constant directivity. While they may have been over-marketed, each was an improvement on previous designs. One difficulty in determining performance was that beamwidth was calculated after smoothing the frequency response. This has the effect of filling in holes and making the polar plot appear smoother than in actuality. Furthermore, many data sheets gave little or no indication of how the data were taken.




Figure 19. Beamwidth of a multi-cell horn constructed from a variable number of cells (Beranek).




Figure 20. Polar response of a 12" radial horn to eight frequencies (Olson).

We decided to publish high-resolution linear polar data in 2000, available on the Internet through the MAPP Online™ program. This program can accurately predict the cancellations and lobes at any distance because it uses high-resolution data without smoothing.

4.3 The LEOPARD HF Horn

Returning to the LEOPARD system, we need a small 12" horn that can fit between the two low drivers and that has a horizontal coverage of 110° from 500 to 16,000 Hz. Our intent is to run both LF drivers up to 600 Hz, which greatly reduces our crossover range options compared to the M'elodie speaker system (which had a crossover between 280 and 2000 Hz). The generic horn is the best attempt to solve the BW narrowing between 500 and 2000 Hz that we have seen from other manufacturers to date. The measured responses of the generic horn and M'elodie are shown in Figure 21, with the horns mounted in their cabinets.

The polar response of the LEOPARD horn measured in free space (Figure 22), shows a remarkably constant 110° beamwidth—much more constant than either of the horns in Figure 21. This design was achieved by extensive computer modeling before any prototypes were built.

4.4 The LEOPARD HF Driver

LEOPARD uses a new compression driver for the horn just described. It's designed to produce 6 dB more output level than the compression driver used in the M'elodie. The 3" diaphragm (piston) has no breakup modes below 18 kHz. This makes for not only lower distortion at higher levels, but an extremely long life at high output levels. The air overload in the horn should be the limiting factor for distortion, provided we make the driver linear.




Figure 21. Measured horizontal beamwidth for two horn designs. Beamwidth is measured at 2 m and computed from the -6 dB point with 1/6 octave smoothing.

The LEOPARD horn flare is a balance of directional control and decreasing air distortion, the latter achieved by a rapid expansion rate near the throat so that SPL reduces rapidly as the wave moves down the horn. The best condition for generating distortion is to run high pressure sound through a tube. The longer and narrower the tube, the greater the distortion, which is why diffraction horns have such high distortion. Second harmonic air distortion increases with frequency for the same SPL output in an exponential horn.

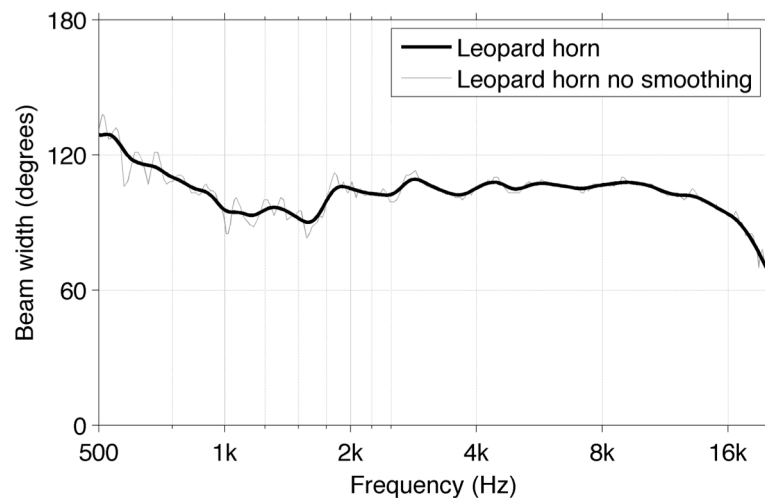


Figure 22. Horizontal beamwidth of the prototype LEOPARD horn, measured as in Figure 21.

The measured LEOPARD horn response matches the calculations very well for 106 dB SPL output at 2 m. Both measured and calculated 2nd harmonic air distortion is 1.2% for a 1 kHz tone. The calculated 2nd harmonic air distortion is 2.5%, while LEOPARD produces just 3% at 2 kHz. The third harmonic is below 0.3%, which is an indicator of very low mechanical “soft clipping” distortion.

This type of HF compression driver has not generally been available on the market before because it lacks any physical protection. One kind of HF driver developed in the 1940s for passive crossover systems, and still in use today, has very stiff, low compliant suspensions to offer protection at frequencies below the horn cutoff. They also react less to horn variation in impedance, which causes peaks and dips in the frequency response. By contrast, the LEOPARD HF driver relies solely on electronic protection.

5 References

Beranek, Leo (1954). *Acoustics*. New York: McGraw-Hill.
Olson, Harry F. (1957). *Acoustical Engineering*, 3rd Ed. Philadelphia: Professional Audio Journals.