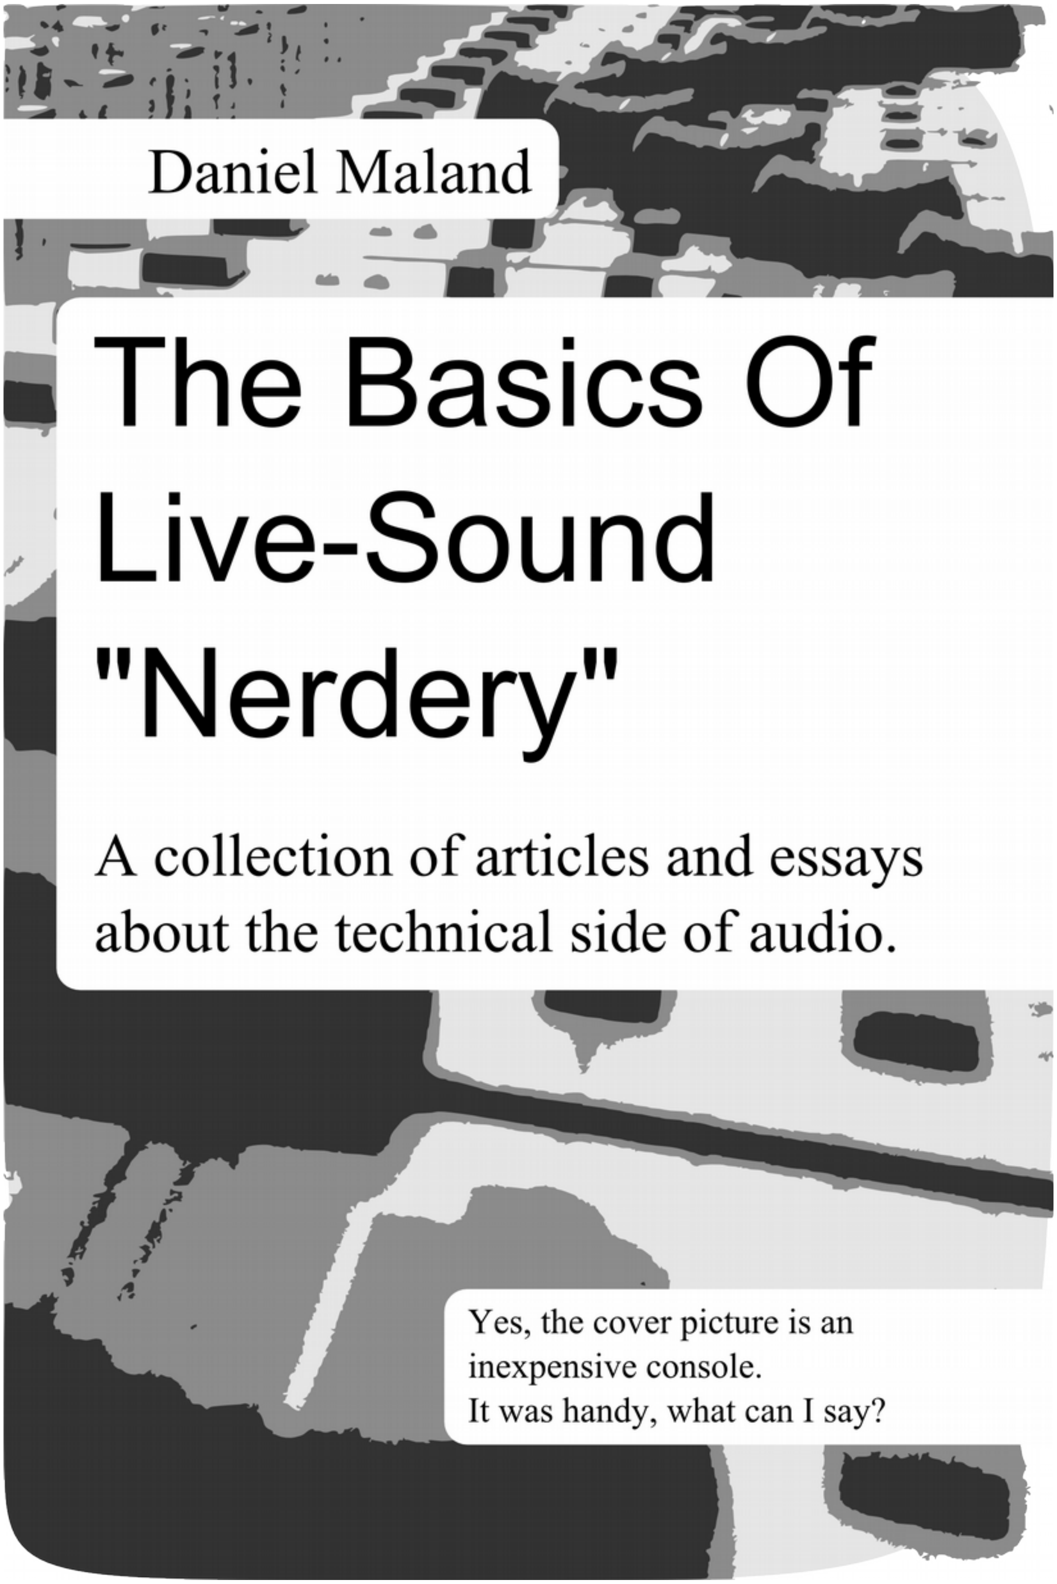

Daniel Maland

The Basics Of Live-Sound "Nerdery"

A collection of articles and essays about the technical side of audio.

Yes, the cover picture is an inexpensive console.
It was handy, what can I say?

The Basics Of Live-Sound “Nerdery”

A collection of articles and essays about the technical side of audio

by

Daniel “Danny” Maland

About This Book

When I set about putting this thing together, a nagging thought was in the back of my mind:

“Do we need another book about live audio, or any kind of audio?”

After all, there is no shortage of publications that deal with sonic physics, studio sound, concert sound, bad sound, good sound...all the kinds of sound that you can imagine. What's the point of throwing anything else into that pool? It's all been talked about, and all the basic facts are discoverable without too much trouble.

After having my nagging thought for a while, I was suddenly blessed with a moment of clarity. (Moments of clarity are rare for me.) I think I was lying in bed. It suddenly came to me that there are a lot of different textbooks dealing with math.

So...what?

Well, the thing is that if any subject is really well documented, really well understood, and factually unassailable, it's mathematics. It's not as though there's a range of opinions about how addition works, or even competing theories on the mechanics of, say, derivative and integral calculus. (Or maybe I'm wrong about that? I'm not a math guy. Anyway...) For some reason, though, we haven't settled on one math textbook that's the math textbook. Different math texts are formatted in ways that different people like. Different math texts get the information across using different analogies, or more of them, or less of them, or have better references in the back, or...

I think you get the picture. This is all just another re-statement of topics that other people have covered. Maybe this treatment is the one that will resonate with you. (Resonate – that's a sound/ music term. See what I did there?) If you've been through other books about topics in live/ concert audio that didn't quite satisfy you, maybe this one will have the missing piece. On the other hand, if this turns out to be the book that you don't really get fired up about, I hope that it somehow helps to point you in the direction of the book that does get you excited.

Important Notes

A good deal of the material in this book got its start as forum posts and blog entries, some of them originally written for AVNetwork.com and Systems Contractor News. I am, of course, very grateful to them for publishing my ramblings in that form. I'm also very grateful to Ms. Kirsten Nelson (the editor of SCN) for giving me the chance to submit my material to them. I'm also very grateful to Rev. Mark Peach for introducing Ms. Nelson and me. (There's a bit of a story.)

For the purposes of getting this material collected and more cohesive, I've made additions, deletions, and clarifications.

You may notice that a lot of the graphics have a sort of "rounded off" or "hand drawn" look to them. This is because most of the graphics were originally meant for online display, and getting them up to a decent resolution required tracing them in a vector program. In cases where clarity of detail was required, I've redone part (or all) of the graphic, and then run it through the same vector trace to keep the style more consistent. There were a few instances where a vector trace was simply inappropriate, and so you'll just have to tolerate the stylistic differences there.

As a final point, this material is about the basics of "nerdery," as opposed to just the basics. This book assumes that you know the essentials of math, like simple algebra, and that you also know how to set up a PA system and get audio through it. "Nerdery" is that next step down the road, where a new level of fascination begins to kick in. This book may still be useful and/ or inspiring to you if you haven't gotten a handle on the bare fundamentals yet, but there may be some things that go over your head if you're not conversant in those concepts.

Part 1: Science And Math

Introduction To Part 1

Over the years, I've seen and heard quite a bit of grumbling from more seasoned audio professionals that the “new guys” don't have an adequate grounding in the academic underpinnings of the audio field. The perception is that there are a lot of people who know how to work equipment, but very few people who have much insight into why basic equipment works the way it does. The inevitable “these kids today” moan-fest ensues (and accelerates) quickly after the current generation is accused of ignorance, which ends up creating an opportunity for the older hands to feel superior about themselves.

Of course, while the moan-and-groan is going on, the “ignorant” young folks are probably out making all the money and connections – background knowledge or no background knowledge.

Anyway.

I am in partial agreement with the folks who want to see a little more scientific rigor in the lives of technical professionals. I think that more knowledge is better than less, and I think that quite a few people are “out there in the water” without a really strong understanding of what makes audio tick at a basic level. There are plenty of audio myths and opinionations that I've run across which betray a lack of understanding about what's really going on with sound and sound gear.

However!

My suspicion is that this is not actually a “these kids today” problem at its core. Rather, I think this is an issue that runs across the generations. My gut feeling (I have no hard data to back this up, I admit) is that the majority of people out there are rarely encouraged to ask: “Why?”

I believe that this has been the case for a very, very long time, and is not necessarily the product of a pernicious societal downturn. Mostly, I think that as equipment and professions become more commoditized, less background knowledge is required to hurdle the barriers to entry. “These kids today” are probably just as teachable and mentally flexible (if not more so) than anyone else, it's just that they haven't had to use their full capacities to get into the business. They haven't been encouraged to ask “Why,” and in the absence of actually needing to ask, they just get on with using the equipment at their disposal.

This may seem like just a question of semantics, until I ask this question of myself and my elders:

> *"Did you understand the math and science behind audio before you got interested in the field, or did you start using the equipment – however basic it was at the time – and then develop an understanding of what was going on?"*

I strongly believe that the honest answer, for most people, is that they started playing with the toys first. The play jump-started the engine of curiosity, and the learning of fundamentals began from there. Yes, there are examples of people with very deep math, science, and engineering backgrounds taking that knowledge into the audio field, but I don't know that such an experience is or was REALLY in the majority. (I could be wrong, though! I often am.)

With all this on the table, this series of topics is going to be an attempt to present some educational concepts to the folks out there who may not have been exposed to it before. The idea is that the whole business may start the "curiosity juices" flowing in some folks, by giving them some background material to chew on that goes beyond the latest and greatest product manual. For the experienced hands out there, none of this will be new – but I'm betting that there are a fair number of folks that could be pointed this way for whom there might be some very juicy tidbits for the taking.

Now then – since we're so very rarely encouraged to ask "Why," what do you say that we start this exercise by asking "Why worry about this kind of material?"

Here are my answers to that query, not necessarily in a particular order:

First, just being a "button pusher" is not enough. Yes, it works for a while, and yes, you can engage in a lot of discovery and learning by experimenting with the various toys available to you. However, just being a button pusher without delving into why the buttons are being pressed (not just what they do) becomes very limiting after a short while. Also, knowing to what extent the knobs and buttons can and can not solve problems can lead to a much greater confidence in one's work, not to mention making one able to propose better solutions to a particular problem.

Second, understanding "why" is the gateway to "how." You may not know how to work a particular piece of gear, or know how to solve a particular problem in terms of "procedure." However, if you understand the concepts that drive gear, or underlie a conundrum, you can sit down and think through creating your own procedure.

Third, more understanding means better navigation of marketing materials. People are always pitching the latest gizmo and gadget, printing spec sheets that claim "this does cool things," and assuring you that their product is THE solution to your needs. If you know what's feasible and what isn't, and you have some sort of idea about why things work, you will be much more able to figure out what manufacturer claims make sense – and where the caveats might lie. Better product choices mean more effective solutions, and that translates into more satisfaction all around.

Fourth, knowing the underlying concepts helps you to be innovative. It's true that you may happen upon a great solution to a problem without knowing all the nuts and bolts. However, if you want a consistent chance at coming up with ground breaking ideas and new ways of using existing tools, you have much better prospects if you know the concepts behind it all. (I couldn't imagine someone like Dave Rat coming up with all the cool deployment ideas and products that he has if he hadn't become educated about how sound and audio technologies really work.)

Fifth, the more you know, the more problems you can troubleshoot. If something breaks, or a proposed solution ultimately fails, you're much more likely to find a fix if you're able to comprehend possible root causes.

Sixth, you can apply general concepts outside of audio. Having a curiosity about how things work, and even more, developing your ability to find out about how things work, isn't just handy in one place. You might get an insight that goes across technical departments. You might also discover that you have a knack for seeing the patterns that govern the financial side of the business of audio – especially if you discover a "bridge" topic that ties things together. (For instance, I have this sneaky suspicion that pretty much all of audio is logarithmic, including equipment package capabilities and costs. I might talk more about this later – I haven't decided yet. I could be totally crackers...)

Finally, knowing the basics provides a foundation for more advanced knowledge. Nobody goes to college before they've been through first grade, right? As you satisfy a simple bit of curiosity, you get a foothold to propel yourself up to understanding another layer of complexity. As I get older, I get more and more convinced that our most convoluted solutions are just a cascade of simpler solutions and refinements that got progressively stacked up on top of one another. (It's the whole "standing on the shoulders of giants" thing.)

Finally, before we truly embark, let me present a standard disclaimer: Everything that I set before you should be read with the idea that “this is how I've come to understand it.” I have tried to ensure that the information I present is sane and can be verified, but I'm not perfect and everything should be subject to scrutiny and checking-over.

Sound Ideas

Audio folks spend vast amounts of time dealing with actual sound in the actual world. That is to say, sonic events that are happening in the air, as opposed to signals that ultimately correspond to sonic events. This being the case, an understanding of what sounds really are, and how they behave, can be a very handy thing to have in the ol' mental toolkit.

Sound is a fairly basic sort of animal, when you get down to it. Any sound pressure wave that you encounter is air molecules being pushed together (compression) and expanding apart (rarefaction) at regular intervals. The pressure wave attempts to travel outward from its source equally in all directions, creating a spherical pattern of radiation. The caveat with this, though, is that truly spherical sound radiation is rarely observed. The source of the sound has to be small (relative to you) and there can be no obstructions for the pressure wave to encounter.

To get into this a bit, the ideal definition of “small” is that of a geometric point, which occupies no space whatsoever. While you can get far enough away from a sound source for it to seem similar to a geometric point, it's rather more difficult to get rid of all obstructions. A small loudspeaker that's several hundred feet away from you, sitting in an empty field, is still sitting on a giant obstruction. That obstruction is the ground.

You can use the following picture to help you visualize all this. The top frame shows a sound source in a 2-dimensional world, with a pressure wave radiating out equally in all directions. The middle panel shows radiation in the horizontal and vertical planes, to help you imagine how a sphere would be created if you had an infinite number of intermediate planes. The last panel illustrates the “source sitting in a field” idea, with the ground obstructing the downward radiation of the pressure wave.

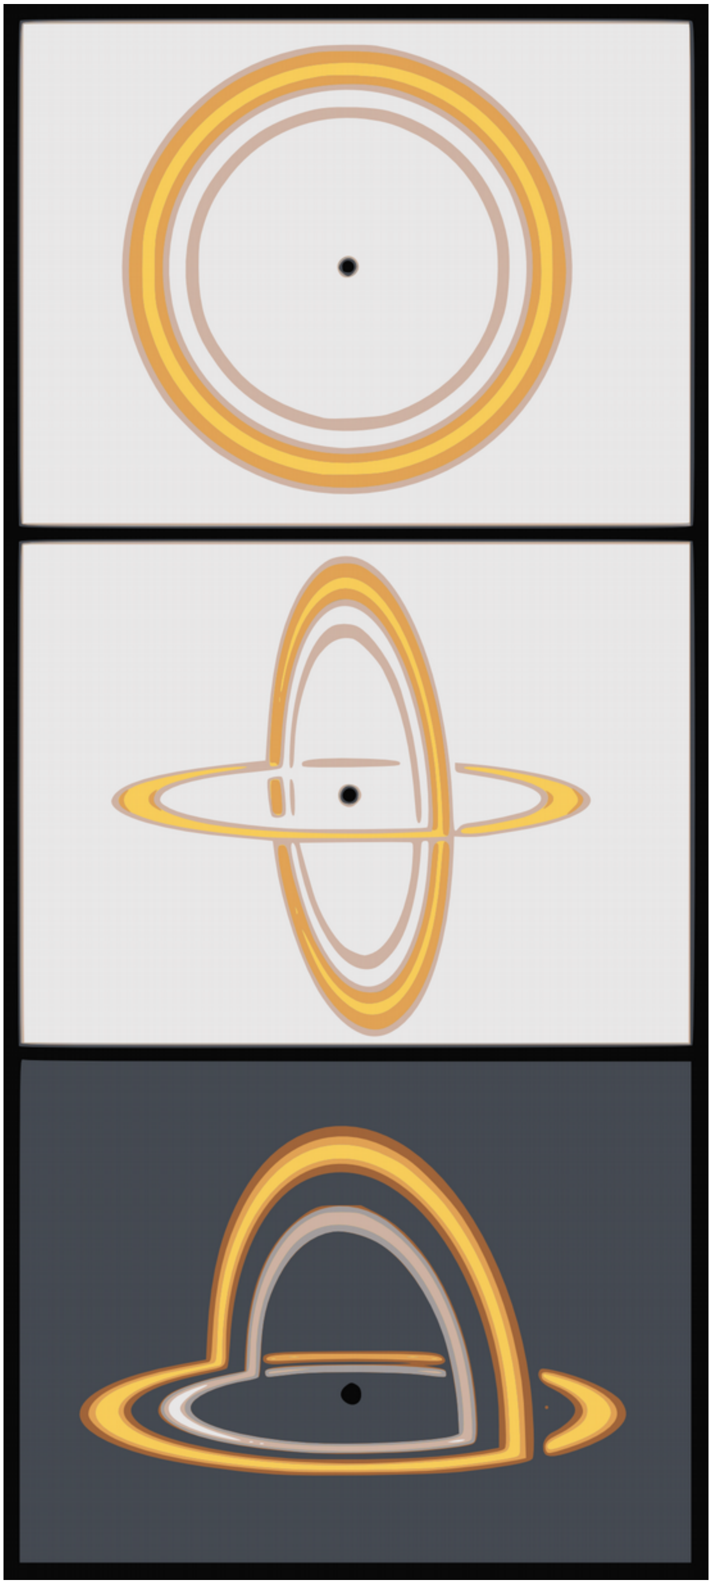

Quite a lot of people represent sound as traveling like “ripples on a pond.” While this is partially true, especially when trying to get an idea of a sound's overall radiation pattern, the analogy omits an interesting detail. Water ripples are what you would call a “transverse” wave. If you look up the meaning of “transverse,” what you get is that the word is an adjective meaning that something is across something else. Transverse waves move between two states (oscillate) in a direction perpendicular to their vector of travel. If a transverse wave is propagating (moving outward from an origin) in a horizontal direction, then the “swing” of the oscillation is in the vertical direction. By contrast, sound pressure waves in air are a longitudinal wave. The oscillation of a longitudinal wave is parallel to the propagation vector. If a sound pressure wave is heading north from Salt Lake City, its oscillation will be to the north and south.

This next picture is a visual example of the difference between transverse and longitudinal waves. The top panel depicts two cycles of a transverse wave, and the bottom panel is one cycle of a longitudinal wave.

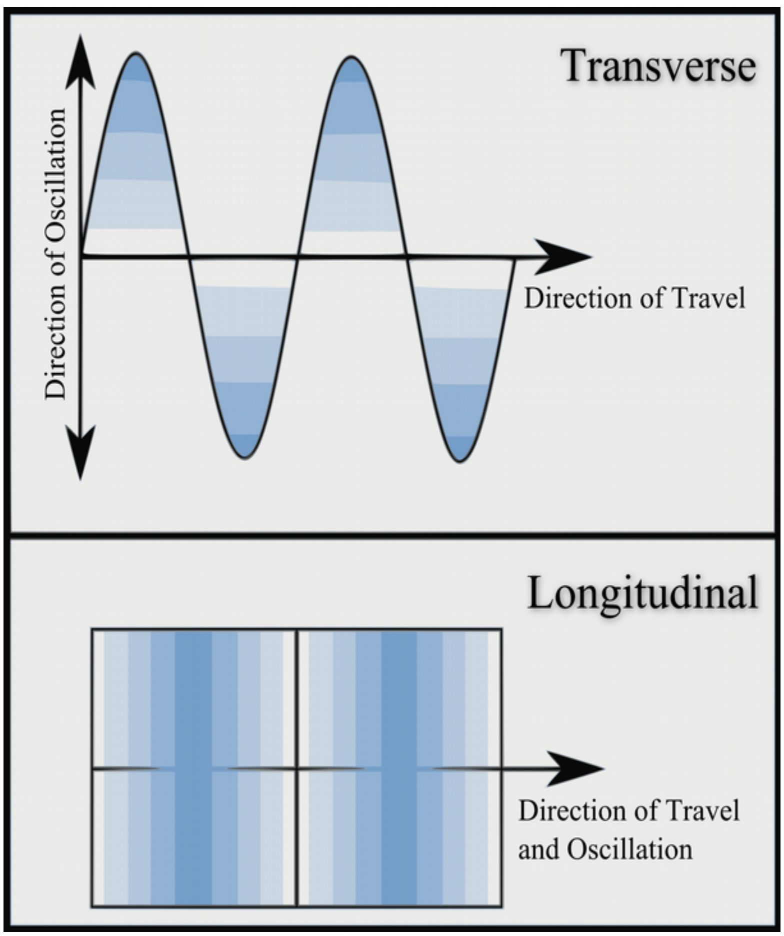

Even with the reality of sound pressure waves being longitudinal when in air, later diagrams will depict sound pressure waves as though they were transverse. The reason for this is convenience. It's rather easier to create complex diagrams of transverse waves than their longitudinal counterparts, mostly because (in my opinion), the separation of visual information into two axes makes for clearer pictures.

Riding Waves

It's time to start talking about how sound pressure waves can be characterized. These characterizations lie at the core of pretty much everything we can make use of in practical reality. I'll talk about that as we go. First, however, I'm going to hit you with another version of the diagram from every physics textbook in history – and yes, I'm going to represent sound as a transverse wave. (You know, clarity, and all that.)

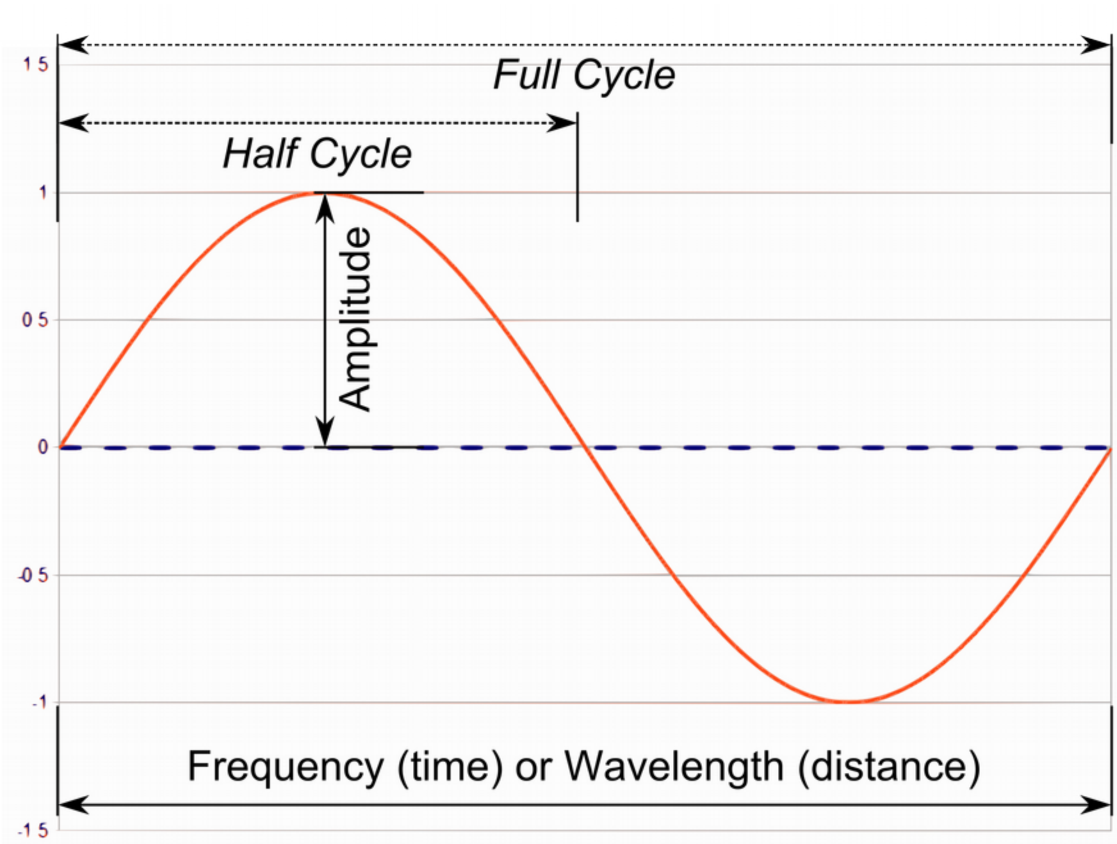

A special mention needs to be made of why the "Amplitude" line looks like it does. The line only indicates the distance from "0" to the wave's positive peak, because that is how most audio technicians express wave measurements. That is, we tend to talk about signal quantities in absolute values – we ignore the "sign." For audio techs, this signal has a peak amplitude of 1 volt. It would be more correct to say that the peak amplitude is plus/minus (±) 1 volt, but that's a bit of a fuss to say all the time.

For audio technicians, amplitude pretty much comes down to one thing: How much force is behind this pressure wave? In air, more force results in more pressure, and thus, larger dB SPL (decibels Sound Pressure Level) measurements. In electrical signals, higher amplitude means more voltage (electromotive force), and by extension, more current and more power into a given device. Voltage measurements for audio are usually in dBu and dBV. I'll go into what all these dB thingamajigs are later.

Frequency is how often we get one full cycle of positive and negative amplitude. A sine wave like the picture above “swings” symmetrically above and below an arbitrary “0” reference point. Ideally, for electrical signals, “0” does actually indicate 0 volts. However, it is entirely possible to have the electrical reference point offset, commonly as a result of DC (direct current) getting into a signal line. Frequency is measured in cycles-per-second, but that's quite a fuss to say or type. For convenience we use the unit Hz (hertz) as a shorthand. Higher frequencies correspond to higher “notes” or pitches. The typical bandwidth of human hearing is 20 Hz (a very low “E”) to 20,000 Hz (a very high “D#”).

Wavelength is the physical distance traveled by a wave as it goes through one full cycle of positive and negative amplitude (while propagating away from its origin). Wavelength is inversely related to frequency. As frequency goes up, wavelength goes down. Wavelength is directly related to the velocity of the propagating wave. As waves travel more quickly, their wavelength increases. These three factors of wavelength, frequency, and velocity are interrelated in a way that gives us this handy mathematical formula:

$$\lambda = v / f$$

The funny lookin' thing on the left is “lambda” (wavelength), whereas velocity (v) and frequency (f) sit over on the right.

“Wait a minute!” you're exclaiming. “Why does this relationship exist the way it does?”

Here's an example. Let's say that you're in an enclosed space, with a lot of downtime and a bucket of tennis balls on your hands. Let's say you toss one of the balls so that it travels 5 feet in 1 second. Along the way, it only bounces once. Here's how that looks when put into the equation:

$$\lambda = 5\ \mathit{feet\ per\ second} / 1\ \mathit{cycle\ per\ second}$$

$$\lambda = 5\ \mathit{feet\ per\ cycle}$$

What if you toss another ball, changing only the speed of its travel away from you? If it's still bouncing at the same frequency, but the travel speed has doubled, then the equation changes to this:

$$\lambda = 10\ \textit{feet per second} / 1\ \textit{cycle per second}$$

$$\lambda = 10\ \textit{feet per cycle}$$

The wave stretches out horizontally, because it's going farther between bounces. If you want to have the ball travel at the faster velocity, but get your old wavelength back, you have to double the bounce frequency, like this:

$$\lambda = 10\ \textit{feet per second} / 2\ \textit{cycles per second}$$

$$\lambda = 5\ \textit{feet per cycle}$$

By making the ball bounce more often, you've counteracted the “wave stretching” effect of the ball moving outward at higher speed.

The natural question at this point is probably this: “Okay, so how fast does a sound pressure wave travel?” The answer to that question is one of my most often given replies. I think it's also my most annoying reply. “It depends.” Sound travels at different velocities in different atmospheric conditions. Find a place with 0% humidity, at sea level, that's at 68 degrees Fahrenheit, and sound will travel at 1126 feet/second. If the temperature rises to 95 degrees, those sound pressure waves will rocket by at 1155 feet/second.

For most purposes, just using 1126 feet/second is likely to be plenty close. Even rounding to an even 1000 feet/second is close enough for ballpark estimates.

As a bit of an aside, I'd like to relate a couple of tidbits about why wavelength and the speed of sound are so important. The design of horn-loaded loudspeakers is dependent on wavelength. A horn that's short relative to a frequency's wavelength, for instance, won't be much help.   General acoustic problems, like nasty standing waves and room treatment needs, also center on the wavelength/ frequency relationship.

In terms of the speed of sound, there's one particular example that was related to me that brings home the importance of sound velocity and changing temperature: An outdoor venue (I don't know which one) had a loudspeaker time alignment system that utilized temperature probes. Based on the sensed air temperature, the alignment delays could adjust themselves to an appropriate setting. This was important because, as the sun would set, half the venue would fall into shadow...and sounds from the “cold” half of the PA would start to arrive “late” when compared to the “warm” half. (You'll see why this is important shortly, when we talk about phase.)

An Aside About Radians

I am most certainly used to expressing angles in degrees. However, radians are worth talking about as an interesting aside. Radians are found by dividing the length of an arc by the radius of that arc. If we're talking about a full circle (360°), then the arc length is simply the circumference. Circumference, of course, can be expressed generally by this formula:

$$C = \pi d$$

That "d," or diameter can easily be expressed in terms of the circle's radius – just double the radius and you get the diameter. This is important because getting radians requires us to work with things in terms of the radius. So, if we rewrite our circumference formula in terms of a radius, and then put together the ratio that I talked about earlier (arc length over arc radius), we get this:

$$\frac{2\pi r}{r}$$

The two radius terms cancel out, and we're left with 2π. So, a complete circle will always encompass 2π radians. Like I said, I'm used to degrees. The rest of this discussion on phase is going to use degrees. Still, understanding how radians work is nifty.

Just A Phase

So – what is phase, anyway? Phase is the measure of the time arrival difference between two pressure waves or signals of the same frequency, as experienced by a particular observer. It is usually expressed in degrees, although some folks might use radians. (In some computing and programming contexts, radians are the native, internal unit.) Phase has a special symbol associated with it: Ø

You will sometimes see the Ø symbol used on switches that invert polarity. This usage is technically incorrect, but it has also been engrained over long years. (We'll talk more about phase versus polarity later.)

The way that degree assignments work for a wave is illustrated in the following picture. The beginning of the wave cycle is 0°, the "positive" peak is 90°, the zero-crossing is 180°, the "negative" peak is 270°, and the end of the cycle is 360°.

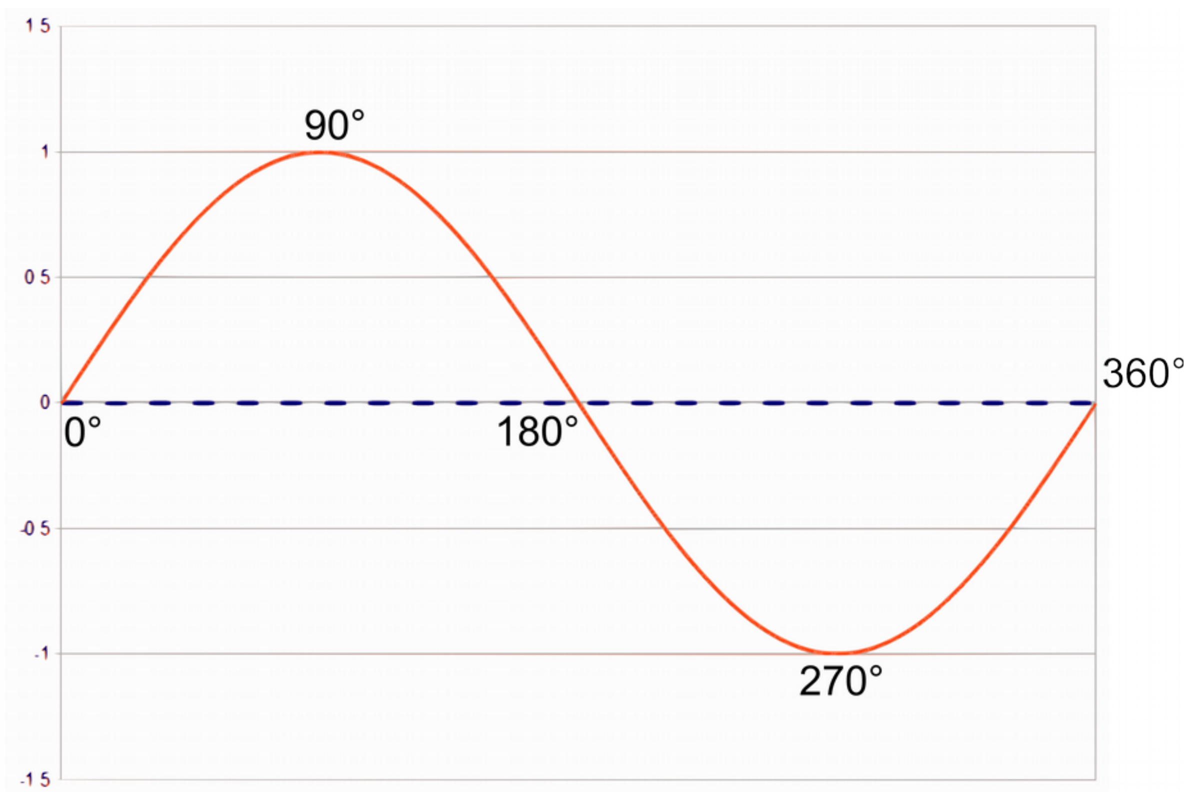

This being laid out, let's say that we have two 1kHz sine waves. Everything about them is the same, except that one of them arrives at our measurement point a half-millisecond behind the other. Since a 1kHz sine wave makes 1000 cycles every second, 1ms is a single cycle. Half a millisecond, then, is one half of a cycle of a 1kHz wave. The half-cycle is where the zero-crossing occurs, so the “late” wave can be said to be 180° out of phase. The following graphic illustrates this situation. The “late” wave (green) begins its positive-going half-cycle where its counterpart is beginning its negative-going half-cycle.

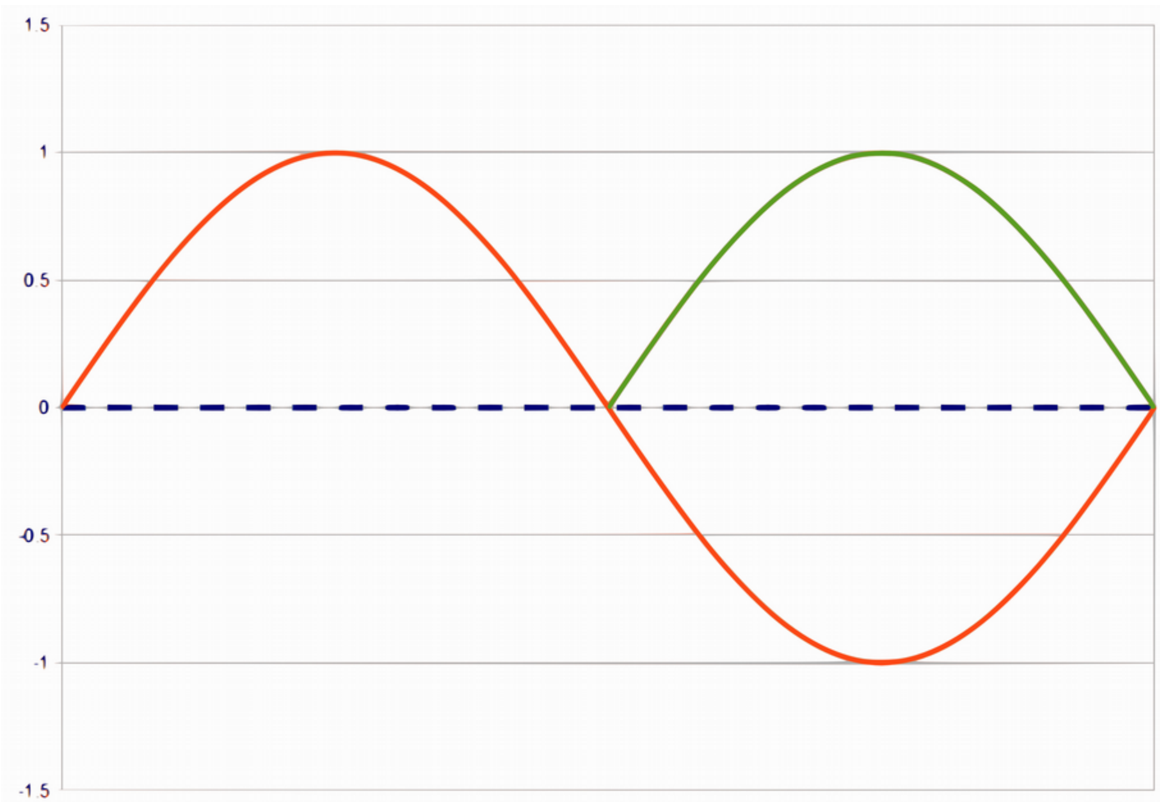

Where things really become interesting is when our two waves are combined. For instance, let's say that this “half millisecond late” condition occurs at the diaphragm of a microphone. The microphone, unlike a human, has no brain to interpret what it experiences – and further, it is only a single transducer. We humans benefit from having two, spatially separated transducers attached to our brains, so we can make differential comparisons. A microphone, on the other hand, “reports” only the total pressure occurring at the diaphragm, so...

The microphone hears silence when the two waves interact, because the sum of the two pressure waves is 0 (in terms of a differential from whatever starting point we have).

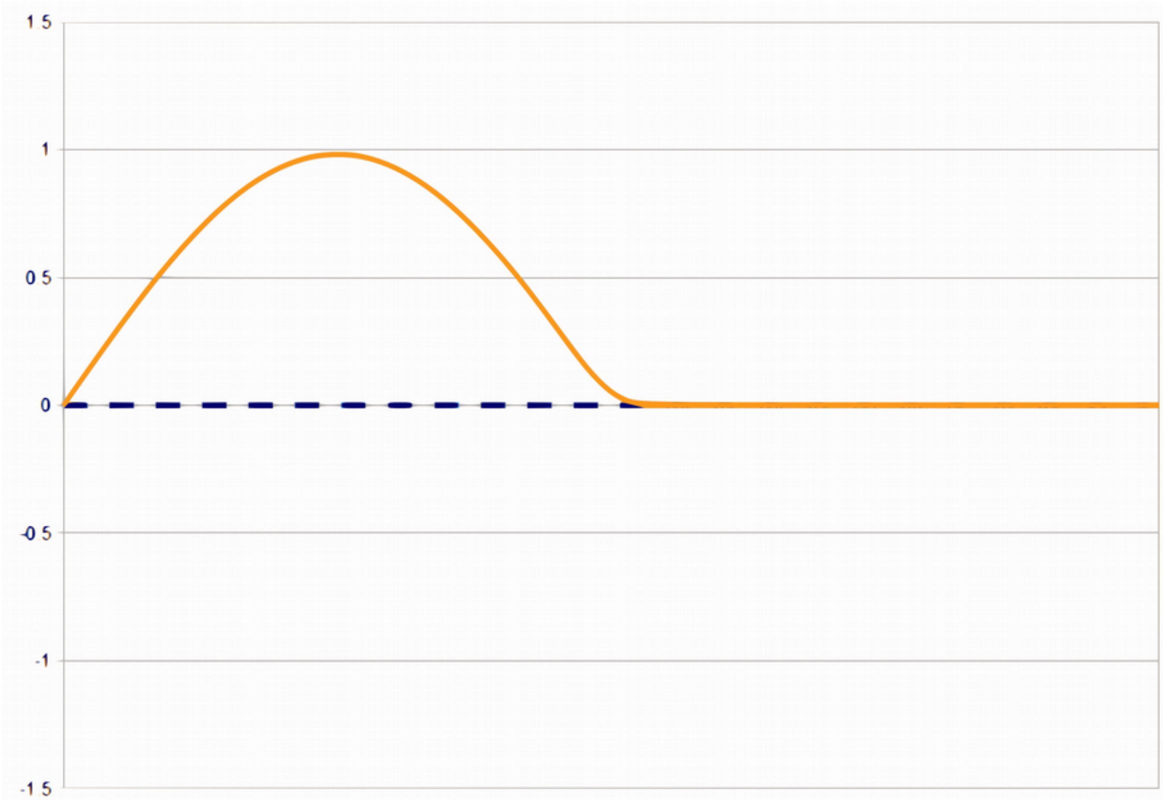

In other words, yes, sound pressure waves and signals do interfere. That interference can be constructive or destructive. The following picture shows two different summations. The red line is the summation of two waves precisely in phase, and the purple line is the summation of two waves that are 90° out of phase. Notice how the vertical scale has changed: The wave summations take them above the “1” and “-1” levels, because this is constructive interference. The strange “ripple” in the purple trace occurs because the wave that starts late destructively interferes at first.

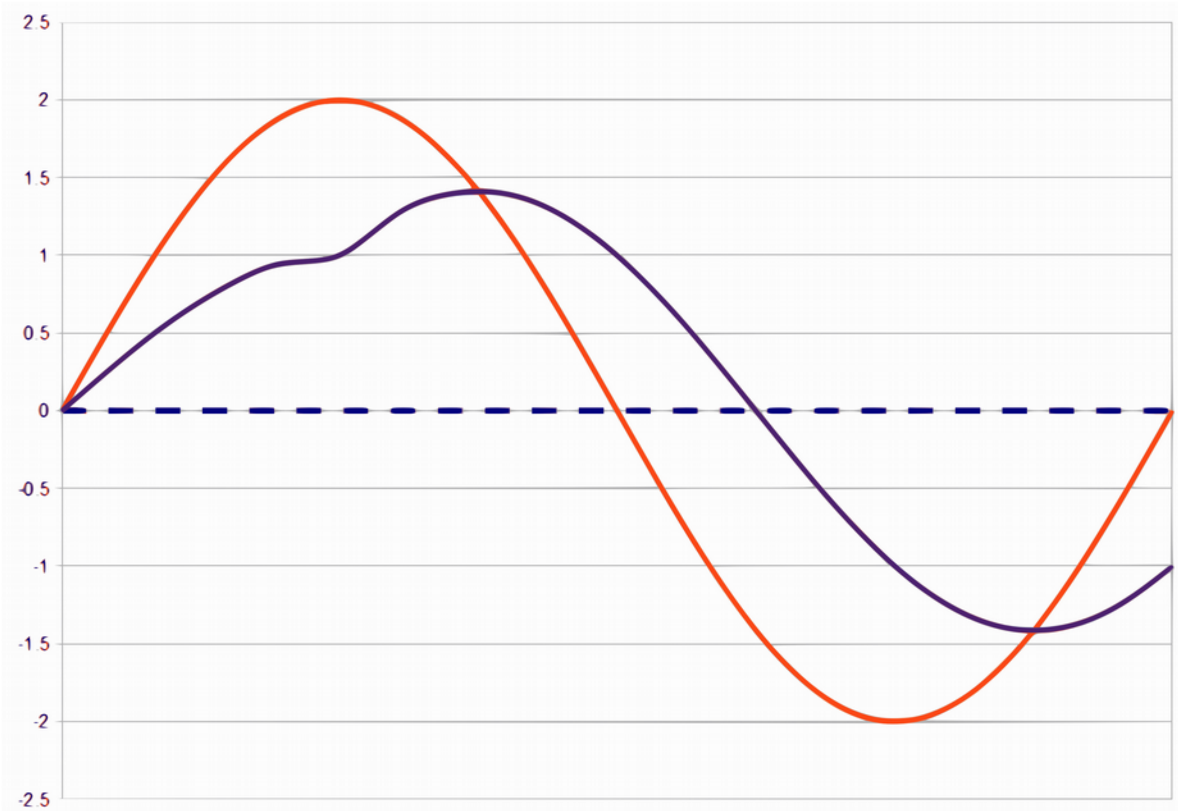

A few paragraphs ago, I set up the “180° out” example by being very specific. The two waves involved had a frequency of 1kHz, and the wave arriving “late” was 0.5ms behind the first wave. The reason I had to be so specific was because of this property of phase: For any given time arrival differential, different frequencies will be more or less out of phase. For instance, if the time differential was still 0.5ms, but the wave frequency was 2kHz, the phase angle would be 360°. If a 1kHz wave cycles once per millisecond, then a 2kHz wave has to cycle twice per millisecond. That being the case, at half a millisecond, the 2kHz wave has finished a full cycle (360°). By extension, a 3kHz wave completes one and a half cycles at 0.5ms, and so is 540° out of phase – which is effectively 180°. (Unless you need to be exact about the time differential, you can “start over” when going beyond 360°.)

In the vein of “interesting applications,” this “wrapping phase” phenomenon is what causes the effect known as “comb filtering.” This name comes from the characteristic look of the effect when viewed as an output trace from a FFT (Fast Fourier Transform) analyzer – like Visual Analyzer 2011 by SillanumSoft.

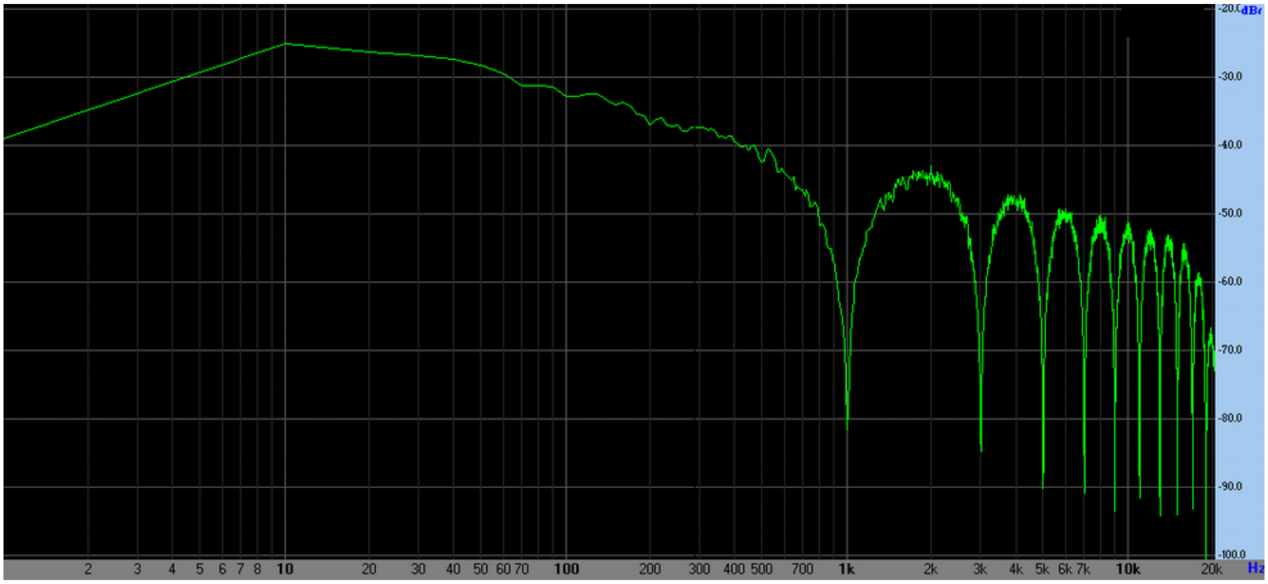

Comb filtering is only apparent with broadband sounds, because the first complete null does not occur at any lower frequency than that which is 180° out of phase for a given time arrival differential. After that, a complete null will appear at every frequency where the time differential corresponds to a wave being effectively 180° out of phase. Interestingly, this re-occurance works out to be the first null frequency, plus the frequency of an octave above the first null. A first null at 1kHz will have additional nulls at 3kHz (1kHz + 2kHz), 5kHz, 7kHz, and so on. A first null at 2kHz will have additional nulls at 6kHz, 10kHz, 14kHz, and so on.

In the vein of bedrock concepts, knowing what phase is and is not is important – especially as it deals with a misapplication of terminology that is still quite rampant in the world of audio: “Phase Reverse” or “Phase Flip”

It is not my intention that you should become a pedantic lecturer to all those around you, but please be aware:

There is no such thing as swapping, reversing, or flipping *phase*.

There is such a thing as swapping, reversing, or flipping *polarity.*

The phase of a frequency is dependent on a specific time arrival difference for that frequency. That being the case, it is not possible to electronically configure a device in such a way that all frequencies passable by the device are shifted 180° relative to themselves. However, what is possible is to re-configure a transducer or signal path such that, in relationship to other devices, positive pressure corresponds to negative signal or vice-versa. In a correctly implemented system, this produces no time differential at all between two signals. There is an “appearance” of all frequencies being 180° out of phase, to be sure, but that is only an appearance. Again, there is no time differential involved, and so phase is unaffected.

(This, by the way, is why it is impossible to actually fix a phase problem by way of polarity inversion on selected signals. You may, of course, find that polarity flipping makes a phase problem less apparent, but to really solve the issue you have to do something that affects time.)

More Phases

Phase effects happen all the time to pro audio types, often without our intention. For instance, if I generate a sonic event via a loudspeaker in a room that is even a tiny bit reflective, I will experience phase effects. The sonic event traveling directly to me from the loudspeaker will arrive at my position at some time, and a reflection of that same sonic event will arrive at my position at some later time. “Later" may be a matter of microseconds, but it's still later.

An unfortunate mentality that can sometimes take hold in audio people is the idea that phase effects are inherently undesirable. To the neophyte, the phrase "out of phase" seems like a bad thing. It sounds even worse when it's coupled with "destructive interference." The newly minted practitioner of sonic wizardry, when inclined to think this way, is robbed of basic understanding and a useful tool for their bag of tricks. Phase is an awfully handy thing – in fact, destructive interference as a result of phase can be downright nifty. Why?

First, consider the fact that you, as a human, are very likely to own a model of the most sophisticated and swift real-time audio processing system created to date. That processing system is your brain, which has two pretty good (if still limited) audio capture and transduction devices attached via a high-speed network. Those ears of yours are constantly talking to your brain via the auditory nerve, and because you probably have two of them, some pretty neat differential analysis can be done.

Because it has two transducers to work with, the brain can deduce quite a bit. The brain will notice the raw difference in SPL between the two ears, plus whatever diffraction and absorption effects are caused by having the head in the way. Just as important, however, is the time arrival difference between the left and right ear. We don't hear comb filtering of every sound, all the time, because we don't simply sum the inputs at our two ears. Instead, the two inputs are compared, and the brain can use all of this "differential" information to figure out where a sound is coming from.

So...what about that whole “destructive interference can be nifty” thing? Probably the most recently popular example of phase as both hurtful and helpful is the area of subwoofer deployment. Let's say that you have an area to cover where you could place stacks of speakers about 30 feet apart, or roughly 10 meters. (My acoustical prediction software only works in metric.) If you have four subwoofers, you could stack two each on either side of the deployment area, except...

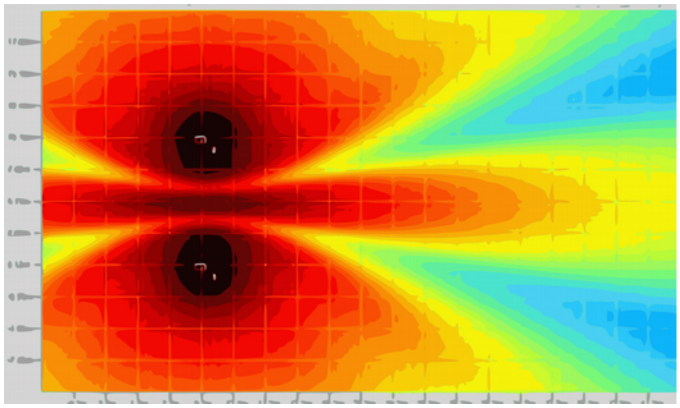

Where the subs are arriving at the same time relative to the frequency cycle times involved, you get a lot of level. When you're close enough to one stack or another of subs that their SPL at your reference point is much higher than that of the other stack, you're also in good shape. However, if you're where the SPL levels of both stacks are similar, and one stack arrives substantially later than the other, you're not in good coverage.

A quick fix is to get all those subs clustered in a central location so that no matter where you are, their combined SPL and relative arrival times are very close. The coverage is much more even, but a lot of acoustical energy is getting thrown behind the subs, which may not be what you want.

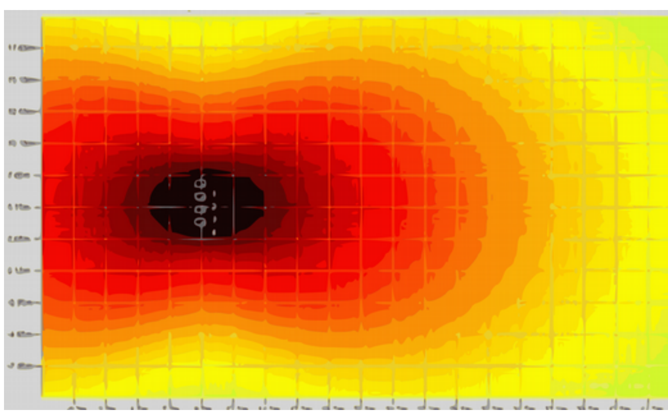

If you have the processing available, and the space required, you can create an end-fired subwoofer array. Because of the relative simplicity of the array, it tends to work best at a narrow range of frequencies, but the problems with this are mitigated by subwoofers not usually being asked to reproduce a very large passband anyway. The trick is to use physical spacing and digital delay (not to be confused with echo) to create a situation where the subs cancel their outputs behind the array, but sum their outputs in front.

The key with delay is to remember that applying it has the effect of pushing the delayed device away from your frame of reference. If you're standing at the first subwoofer, looking down the line, delaying a box pushes it away from the first subwoofer in time. However, a person at the end of the line looking at the first subwoofer perceives that the delayed device is being pushed towards the first subwoofer instead. If we decide, for example, to use a target frequency of 60 Hz, then we can calculate that a quarter wave is about 1.43 meters, or 4.17 ms. I'm using a quarter wave because, if we start with two boxes, spaced a quarter wave apart, and then delay the second subwoofer one quarter wave, it appears to be a half wave away from the first box when standing behind the array, but appears to be precisely aligned with the first box when in front of the array.

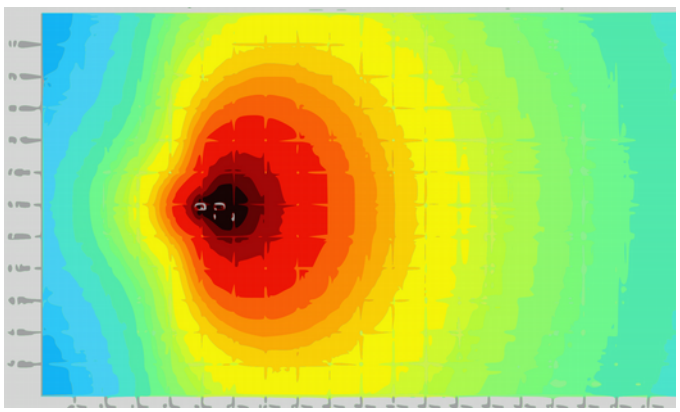

With just two subs in the array, we've already lost a lot of rearward spill, while maintaining smooth coverage to the front. Now, if we add a third sub, we can put it two-quarters of a wave away from the first box, and then delay it by two-quarters of a wave. This does put the number three box in phase with the rearmost sub, but it also combines nicely with the number two sub. The result is a touch more spill directly to the rear, but a more focused area of that "back spill" and stronger output to the front.

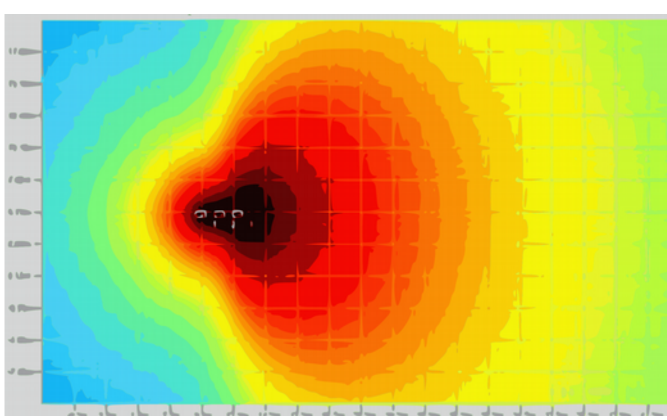

Now, let's add our final sub. It's three-quarters of a wave away from box one, so three-quarters of a wave of delay causes it to be a half-wave out of phase from the reference point of box one. Again, adding another sub has allowed us more forward coverage, while slightly tightening up our spill to the rear sides. The tradeoff is a longer “tail” directly behind the array.

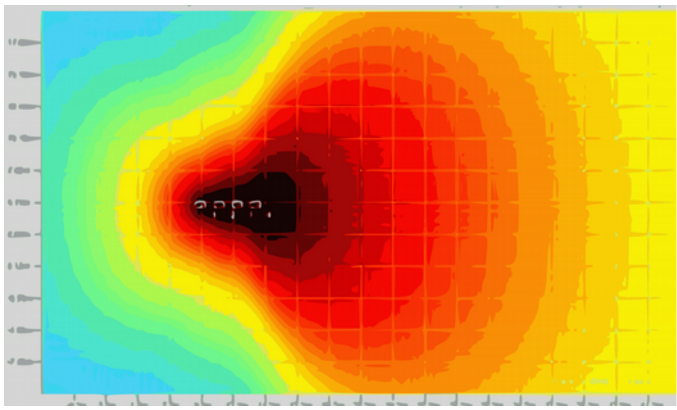

Although deploying an array like this takes up a lot of room (and power), and may not result in exactly the predicted performance (after all, doing this indoors is going to introduce a lot of factors not accounted for in this simple prediction), it's a great example of using phase as a tool to solve problems.

The Mysterious Inverse Square

The experience of sound's intensity being lessened as we put more distance between ourselves and the source of the sound pressure wave is a fairly intuitive thing. If you've ever been right up against the stacks at a concert and found that life was a little too loud, you probably realized that an easy way to remedy the problem was to walk away from the mountain of loudspeakers. The basic mechanics of why that works go like this:

Sound pressure waves radiating from a point source into unobstructed air propagate in a spherical pattern. Effectively, the energy of the sound pressure wave is spread over the surface of that spherical pattern. An observer in the path of the pressure wave is exposed to some section of that sphere, and thus, some portion of the energy spread across that imaginary surface. An observer that is close to the source of the pressure wave gets a far more concentrated dose of energy than someone standing at a distance. This picture offers a visual example of this, using a blue “sound balloon” and a square, gray “observer.” The left side of the picture is an observer that is relatively close to a sound source, and the right side is an observer that is farther away.

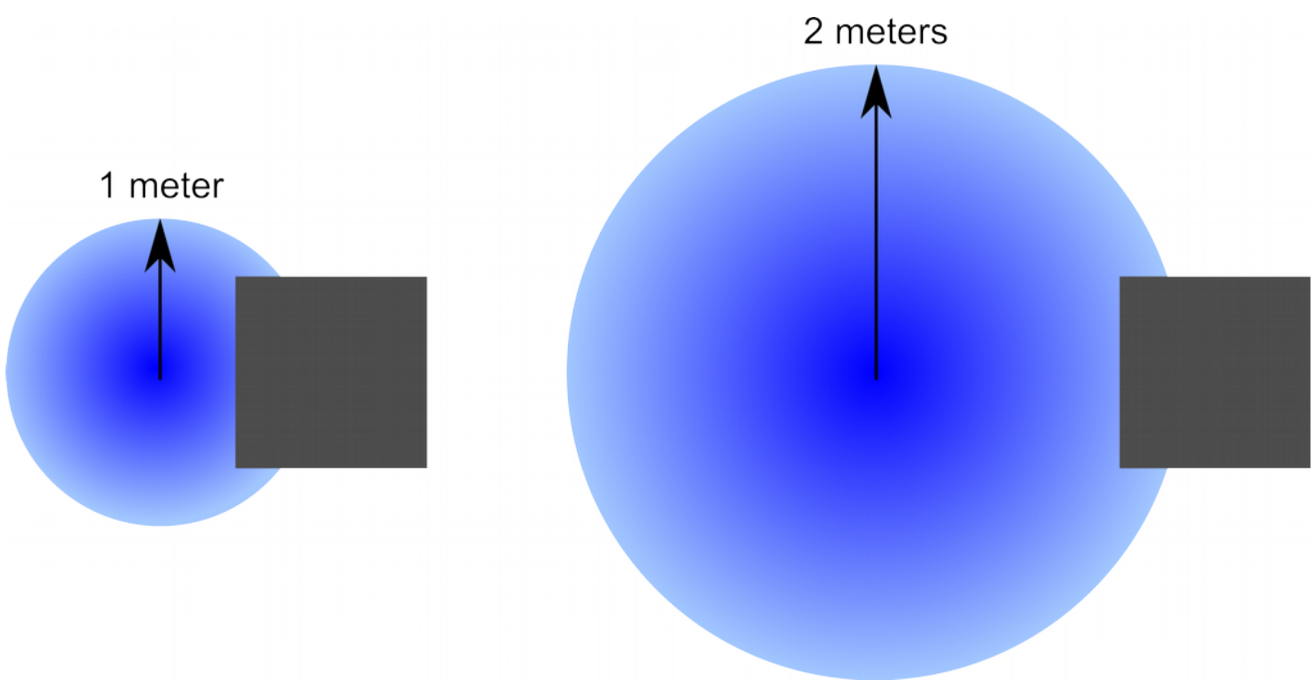

Yes, it's entirely true that we are very unlikely to experience this “idealized” model of sound propagation. Real life just presents too many boundaries and obstacles for a perfect sphere to be our observation of propagating sound pressure waves. However, that natural tendency, shall we say, for sound to “want” to expand in a sphere yields a tidy mathematical model of sonic intensity versus distance. That model is the inverse-square law, which states that, under “ideal” conditions, the intensity of a sound pressure wave experienced by an observer is quartered when the distance to the source is doubled. In mathematical terms, this is how that looks as a generalization:

$$I = 1/d^2$$

If we use a distance of “1,” we get this:

$$I = 1/1^2$$

$$I = 1$$

If we double that distance to “2,” we get this:

$$I = 1/2^2$$

$$I = 1/4$$

That's all fine and good – but why is it like that? Well, it all comes down to that spherical radiation pattern. Physics defines “intensity” as power divided by area:

$$I = P/A$$

In the case of sound emanating from a point source into empty space, the surface in question is a sphere. This is the mathematical generalization for a sphere's surface area:

$$4\pi r^2$$

With that in mind, we can write out a definition of intensity as it pertains to us:

$$I = \frac{P}{4\pi r^2}$$

The bit to take notice of is that “radius squared” at the end. For our purposes, “radius” and “distance” are interchangeable – the observer is, by necessity, sitting at a point on the end of some arbitrary radius. If that radius has a value of “1,” then:

$$I \approx \frac{P}{12.57}$$

If we double the radius to “2,” then we get this:

$$I \approx \frac{P}{50.27}$$

Which is ¼ of the value if the radius was “1.” Notice that the only thing we've changed is the actual number of the answer – the proportionality is the same. The inverse square law, then, is an audio-centric shorthand for the more generalized physics equation. (It's not just for audio. Everything I've been taught indicates that inverse square works in the lighting world as well.)

At this point, you may be asking if this is actually useful in real life. After all, a truly spherical radiation pattern is something that most pro-audio folks will never run across. Although we'll need to go over decibels and how they work in order to get “real numbers,” the inverse-square law is an approximation that I've used many times. It's a great “worst case scenario” rule-of-thumb for predicting how much output a PA system might have at a certain distance, and I'm a big proponent of using conservative estimates when predicting how much coverage you'll get from an audio setup. Further, this idea of power radiating over an area can awaken the question of whether or not we can focus power into a smaller area, for purposes of greater intensity. All of us have probably done it to light, by way of a magnifying glass – and yes, you can do it with sound pressure waves as well. Have you ever noticed that some manufacturers refer to their loudspeaker waveguides as acoustical lenses?

Sonic Bumpercars

Just above, we got into talking through the basic "inverse square" behavior of sound pressure waves. The short-form recap of that is: Without obstructions or intentional intervention, sound pressure waves propagate in a sphere. The intensity of sound that an observer experiences is inversely proportional to their distance from the sound source. If the observer doubles their distance from the source, the apparent sound pressure wave intensity drops to one-quarter of the previous intensity.

Righto. That's great, and all, but...

If you set up a loudspeaker in a room, dig out your SPL (Sound Pressure Level) meter, find yourself a tape measure, and then start taking readings against, say, broadband noise, you will probably not observe inverse-square intensity falloff. What you are likely to get are readings that are higher than you would expect. Why?

A really big factor (amongst other, potentially major factors) is that whole bit about being inside a room. In addition to the section of the sound pressure wave that propagates to you and strikes you directly, you are also in the way of sound pressure waves that are propagating to you indirectly. Unlike an “idealized” model of pressure wave propagation, all that energy that begins by radiating away from you is not “lost.” A pretty significant portion of it does end up reaching you – it's just that it's a bit late.

The following picture is a visual illustration of this. The top part of the graphic shows a loudspeaker producing sound pressure waves that strike an observer. With no boundaries present, any part of a sound pressure wave traveling away from the listener never gets a chance to interact with that person. The contrasting part is the (much more muddled) lower piece of the diagram. Inside an enclosed space, the portions of the sound pressure waves not traveling directly to the observer are, at least partially, redirected into the space. Because of this, that “indirect” energy does have a chance to eventually meet with the observer.

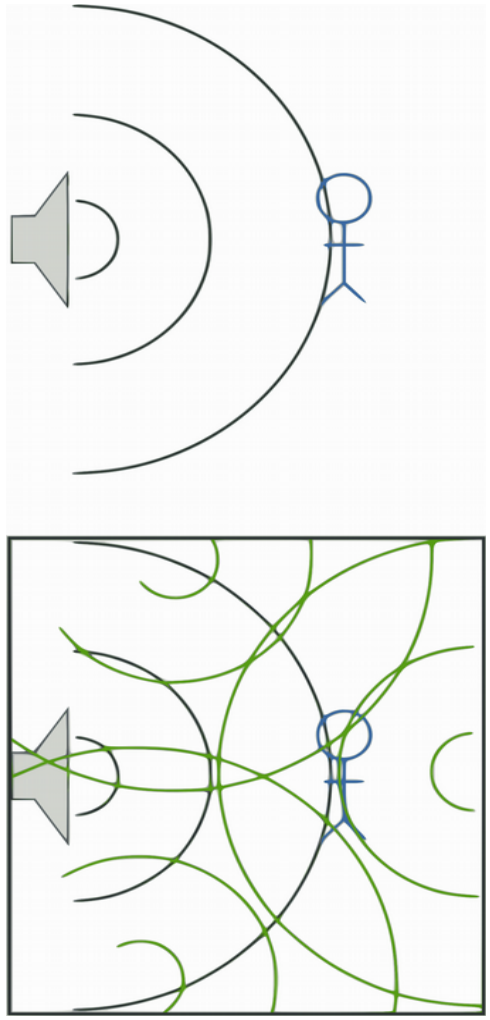

The reason you are likely to get readings higher than expected is because that indirect energy is added to the direct energy that the simple inverse-square model expects.

Now, if you want to be brutally frank, the bottom part of the picture is highly simplified. In real life, there's a lot more to sound bouncing around a room (and everything in the room) than what's depicted above. The good news is that everything basically comes down to three factors: Reflection, absorption, and diffraction. Sound can be made to refract, but acoustic refraction in the practical reality of most audio people is due to temperature gradients – not interactions with solid objects.

Simple acoustical reflection occurs when a sound pressure wave encounters a smooth, flat, rigid object with a surface area that is large in comparison to the sound's wavelength. The object must also have sufficient mass to prevent the pressure wave from simply passing through it. With all of this in place, you get a behavior like that shown below. The pressure waves strike the object at an angle, and then change direction to continue on at an angle opposite that of their arrival.

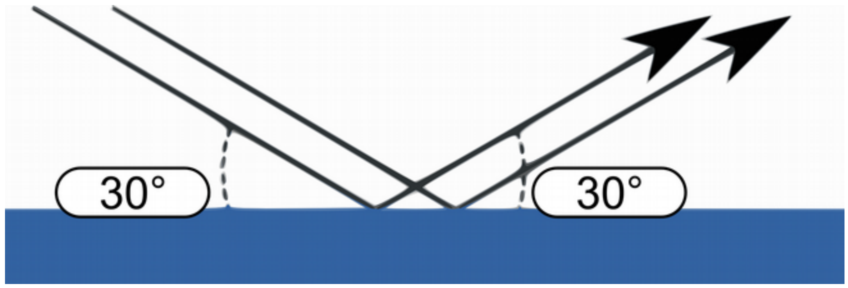

If you've ever been in a room dominated by simple acoustical reflection, you would probably describe the sound of the room as very “live,” and probably “springy.” This is because the sounds reaching you in an indirect fashion are a large number of still-somewhat-distinguishable-from-one-another echoes, and in many cases, those echoes may be able to bounce back and forth between surfaces repeatedly (flutter echo). One could argue that the “liveness” of a space is directly proportional to the sound pressure level of indirect reflections versus that of the direct sound. You could also say that the “springiness” of a room is directly proportional to the sound pressure level of the echoes and amount of time that the flutter echo can be sustained.

If you introduce a rigid, sufficiently massive, and sufficiently sized object to a space, but give it a varied surface instead of a smooth one, what you're likely to get is diffuse reflection. Diffuse reflection occurs when there is sufficient surface variation of an object to cause a sound pressure wave to apparently reflect in multiple directions at once. This isn't some sort of magic, of course – what you really have is just a lot of simple reflections. It's really a question of scale. If all of those echoes happen very closely together in time from a human perspective, then we can't tell them apart as “separate” acoustical experiences. Instead, they form a smooth “wash” of echoes that we otherwise call reverberation. Here is a visual example of diffuse reflection:

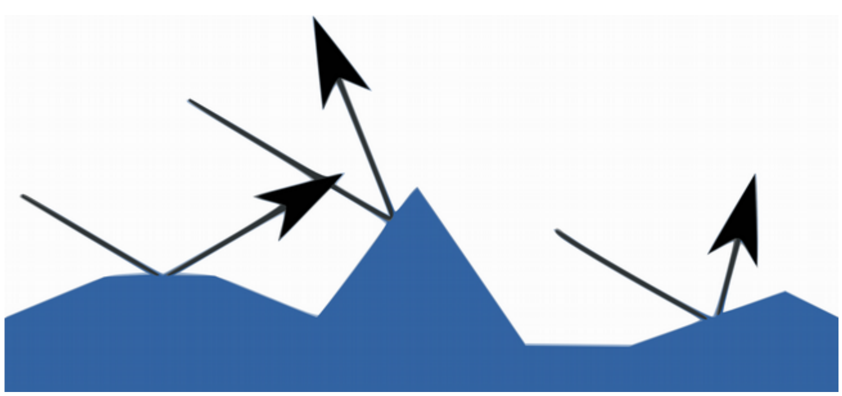

Aesthetically, humans have a tendency to prefer diffuse reverberation to simple "slap" echo – when it comes to room design, that is – and you may find that reverberation is more pleasant when its high frequency content is "rolled off" to some extent.

Absorption occurs when a sound pressure wave strikes an object and a large portion of the associated energy is converted into heat. Like reflectors, absorbers need to have a large surface area compared to a sound's wavelength, but in contrast, absorbers usually need to be a soft, porous material. (It is possible to create absorbers from non-porous materials, but examples of those absorbers are quite rare in terms of room acoustics.) If all other factors are unchanged, absorbers with greater mass will outperform absorbers of lesser mass.

Rooms with a great deal of absorption tend to be referred to as "dead." A sufficiently large amount of absorption in a room can create what is called an anechoic chamber, where a human's perception of reflected sound is negligible. Because any sound pressure wave you encounter in such a room is direct sound, you are very likely to observe sound pressure level decay that is consistent with the inverse square law. Humans make fairly decent absorbers, which is why so much is made of being ready for the difference in the sound of a room when it is empty, versus when it is filled with people.

The final topic for this chunk of audio math and science is diffraction. Diffraction is the phenomenon of sound pressure waves bending around obstacles and through openings. Diffraction occurs when the obstacle or opening in question is small relative to a sound's wavelength. In the case of a small obstacle in front of a sound pressure wave, you will still be able to hear the sound. The wave simply reforms itself on the other side of the obstacle. In the case of sound encountering a small opening, a more curious effect occurs. From an observer's reference point, the small opening acts as though it is the original emitter of the sound pressure wave. This picture shows the "small object" condition in its top portion, whereas the bottom portion depicts a small opening.

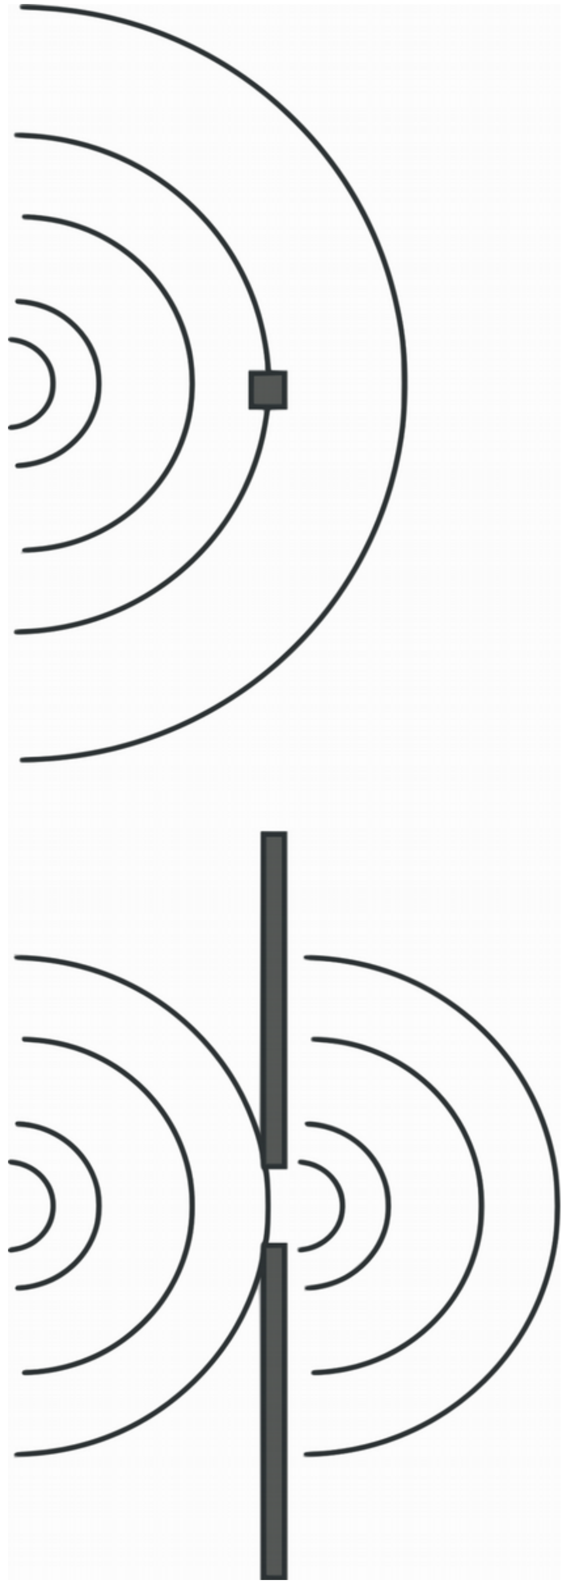

Easy As Fallin' Off A Log

When we're talking about sound, the quantitative descriptions of frequency and wavelength are pretty straightforward and familiar. Hertz is nothing more exotic than wave cycles per second, and wavelength is any measure of linear distance that you might prefer. Counting is something we do every day, and simple distance is similarly common. When it comes to pressure in air, though, most of us don't have the same sort of intuitive experience. I'm going to guess that you've never heard an exchange like this:

"Dude, the guitar player's amp was creating observed pressure waves of 20 Pa RMS (Pascals Root Mean Square), and we were 100 feet away!"

"Seriously, Dude? He had to have it so loud that you were seeing 0.15 mmHg (millimeters of mercury) at that distance? Wow."

Air pressure units just aren't something that we're "calibrated" in. Even voltage is still a little unwieldy:

"It's okay, we've got level to spare. All those inputs together are peaking at 6.15 V (volts), and the processor can handle a little over 12.27 V. No problem."

Beyond us not having an intuitive "feel" for what the measurements mean, there's also this whole problem of the range that the standard measurements cover. For instance, my "loud guitar player" example used 20 Pa RMS, but an undamaged human hearing system is capable of detecting sound pressure waves of only 0.00002 Pa RMS. (One atmosphere is 101, 325 Pa, and that is considered the upper limit of an undistorted sound pressure wave in air. That is, the rarefaction, or decompression cycle of the wave just reaches vacuum. Suddenly, that guitar rig isn't quite so impressive...)

So, what's an audio human to do? We need some sort of abstract measuring stick, one which we can relate to fairly easily, and one that's also fairly compact. To generate such a measuring stick, we can use a handy mathematical maneuver called the logarithm, or "log" for short.

Logarithms are the inverse function of exponents. This picture shows you how they work. A given base, raised by some exponent, gives you a number. If you take the logarithm of that number using the original base, you get the exponent involved.

$$b^x = n$$

$$\log_b n = x$$

For example:

$$10^2 = 100$$

$$\log_{10} 100 = 2$$

This next picture shows you a comparison of exponential and logarithmic functions. The top graph is an exponential curve, and the bottom graph is a logarithmic curve.

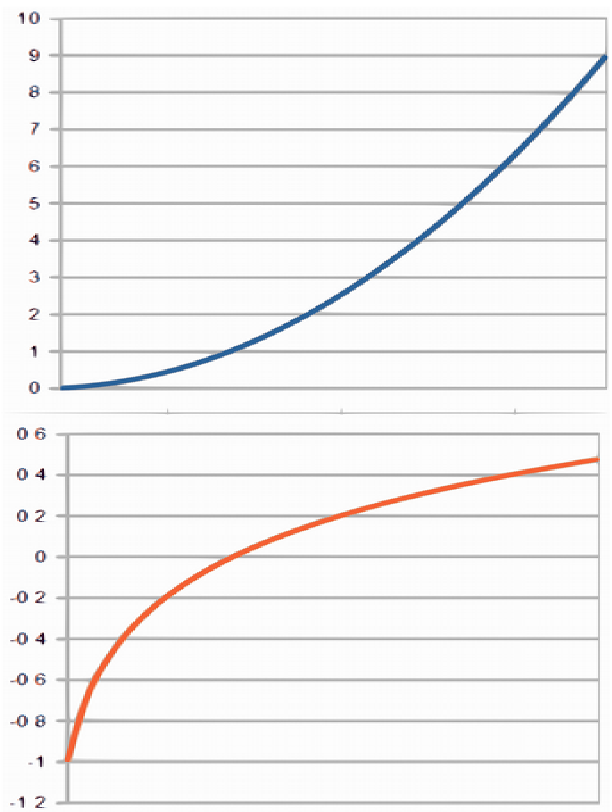

So, that's neat, but how does it help us make a compact measuring stick?

If you look at the exponential curve, you can see that it “accelerates” as the x-axis (horizontal) values rise. The farther you travel to the right on the x-axis, the more the corresponding space required for the graph on the y-axis (vertical) increases. The logarithmic curve, being the inverse case, “decelerates” as the x-axis progresses. This being so, the corresponding y-axis space for the graph becomes more and more compact for each unit increase in x.

Yes, I realize that the last paragraph was incredibly abstract. The bottom line is that logarithms, and the use of logarithmic scales, let us compact and view very wide ranges of measurement values in convenient and comprehensible ways. This doesn't only apply to pressure or voltage, by the way. Frequency response charts benefit greatly from using logarithmic scales, because human hearing has a range of about 19,980 "whole number" frequencies. Try to represent that in a linear fashion, and you get something like the picture below. A very important set of frequencies gets smashed to one side: The very-critical-for-human-hearing range of 1 kHz to 4 kHz, plus everything below that. (This is also where most of the fundamental frequencies of pleasant musical notes reside, like "Concert A" at 440 Hz.) The red highlight accents that range.

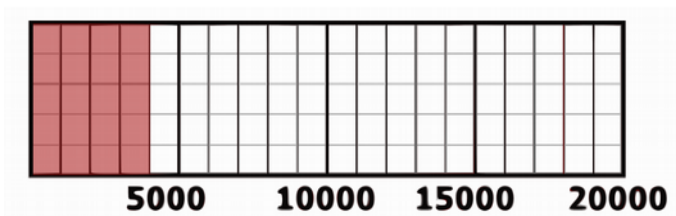

So, what if we used a logarithmic scale? One where each unit of horizontal space in the graph represented one octave? (An octave is a doubling of frequency, and is a fundamental "unit" in both music and sound.) If we choose this kind of representation, we get a picture like this:

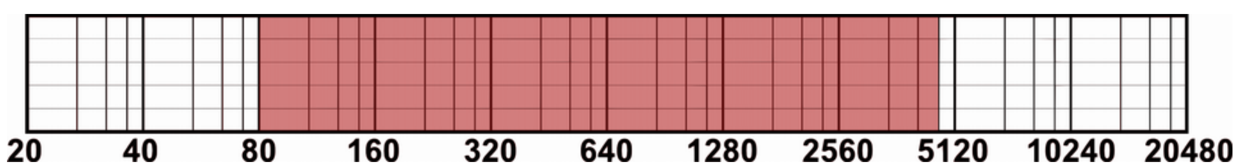

The most critical range is nicely spread out, making the whole thing much easier to read. The highest couple of octaves are compressed horizontally, which means that much less space is wasted on representing 10,000 whole number frequencies at the top of the scale.

Ringing Decibels

What in tarnation is a decibel anyway?

If you're familiar with SI units, you know that “deci” refers to something being one-tenth of the basic unit. For instance, a decimeter is 1/10 of a meter. In the same way, a decibel is 1/10 of a “bel,” which was a unit name dreamt up by – wait for it – Bell Telephone Laboratories. The decibel is a convenient shorthand unit for describing the ratio of measured power to a reference power level. Those two elements are essential to understanding decibels. They are not an absolute measurement in themselves, and the basic, expected ratio is in terms of power. (If the ratio isn't in terms of power, some conversion has to be done at some point.) The convenience of the decibel comes by way of using a logarithmic scale to effectively compress large ratios into a manageable numerical value. For the decibel, the chosen logarithmic base is 10 (also called the “common log”). What this results in is that 10 times more power gives you 10 more decibels.

In audio, the chosen behavior for the decibel is provided by this mathematical generalization:

$$10\left(\log_{10}\frac{P_{measured}}{P_{reference}}\right)$$

If we decide to use a reference power of one watt, we can get a handful of useful examples:

$$10\left(\log_{10}\frac{1\,w}{1\,w}\right)=0\,dBw$$

$$10\left(\log_{10}\frac{2\,w}{1\,w}\right)=3.01\,dBw$$

$$10(\log_{10}\frac{10\,w}{1\,w})=10\,dBw$$

$$10(\log_{10}\frac{100\,w}{1\,w})=20\,dBw$$

Notice that the “output unit” is dBw. This denotes “decibels referenced to 1 watt.” Decibels are technically meaningless without a reference point, and so you should always try to give some indication as to what that reference is. There are certainly situations where the context of the measurement provides all necessary clues, but it's good to push ambiguity as far away as one can.

Of particular note in the above examples is the 2:1 ratio giving 3 dB greater output. Depending on who you ask, humans subjectively perceive “twice as loud” when they get 6 – 10 dB more out of an audio source. The problem is that “twice as loud” is objectively just twice the power, or 3 dB. If you have 500 watts flowing continuously through an audio system, and then want 3 dB more output, you're going to have to find a way to get 1000 watts continuous to flow through that rig. Worse, if people really want to feel like that rig is twice as loud, you will probably have to find a way to get the equivalent output of 2000 – 5000 watts. (I say “equivalent” because simply increasing the power output into a given set of drivers is not always a smart idea. In fact, just buying bigger amplifiers and not getting any more drivers involved is a prime way to cook an audio system.)

Now – what happens if measured power is less than the reference power? If that's the case, we get a negative number. If a 2:1 ratio gives us +3 dB, then a 1:2 ratio gives us -3 dB, like so:

$$10(\log_{10}\frac{0.5\,w}{1\,w})=-3.01\,dBw$$

Now, let's take a moment to delve into one of the sticky bits, which is a piece of the puzzle that has thrown both green and seasoned audio folks for a loop. That sticky bit is taking the time to be aware of precisely what any particular decibel measurement is referencing. If you don't know what the reference point is, you are wide open to all manner of misunderstandings. For instance, I can clearly remember the case of an experienced but confused audio technician being very upset about his new, digital mixing console. The console had meters that would not read beyond 0 dB. The tech in question, being seasoned in the world of analog audio, “knew” that a good, solid level was 0 dB. Having set his levels to be consistently reaching about 0 dB on his meters, he was quite angry at how terribly distorted his signals sounded. What had to be pointed out to him was that his new console used a meter reference where 0 dB was 0 dBFS – decibels referenced to Full Scale. In other words, when you reached 0 you had also reached clipping. What he had wrongly assumed was that the scale was the same as his analog consoles, where the 0 dB reference was up to 24 dB below the point of a signal clipping.

I myself get a bit lazy when using decibels to refer to sound pressure level. I can have a meaningful conversation with most folks by simply saying things like, “Yeah, those guys were pretty mellow. We came out of the gate at only about 100.” The reason, though, that the conversation is meaningful is that we all have experience in what that sloppy reference to “100” actually means. To be really correct, though, I would have to be much more specific: “Yeah, those guys were pretty mellow. We came out of the gate at only about 100 decibels SPL (Sound Pressure Level), C-weighted, averaged over 1 second.” That sentence is rather a fuss to pronounce, but it would be the only way to give a complete stranger a correct picture of what 100 dB meant in our context.

So – how about a handy list of some common decibel reference points? No problem. Please note that the reference points listed are the needed value for a 0 dB measurement, unless otherwise indicated. Also, be aware that none of these references utilize power directly, and so they don't abide by the rule of “double the reference to get +3 dB.” (There will be more explanation about this later):

dB SPL - Sound Pressure Level, referenced to 20 micro Pascals rms. Usually coupled with some description of the weighting curve used when measuring, such as A, C, or Z, and should also be coupled with some description of whether the measurement used a slow or fast average. A-weighted measurements mostly take into account the frequencies that humans are most sensitive to. C-weighting includes much more low-frequency information than A-weighting, and Z-weighting takes into account all audio frequencies equally...insofar as the measuring device can detect them.

dBu – Voltage referenced to 0.775 volts RMS (Root Mean Square).

+4 dBu - 1.23 Volts RMS, the "pro audio" standard reference voltage for VU meters - see below.

dBVU – Voltage referenced to +4 dBu, unless somebody has decided to calibrate the meter to something else. It is quite common to have the VU meters on professional, analog, multitrack tape machines calibrated to something other than +4 dBu. VU metering is good for measuring average levels, but is essentially unhelpful for measuring large, fast transients. These large, abrupt increases and decreases in level simply occur too quickly for VU meter ballistics (the rate at which the meter can rise and fall) to show accurately.

dBV - Voltage referenced to 1 volt RMS.

-10 dBV - The "consumer" audio voltage reference of 0.316 Volts RMS. A very common mistake is to say that "consumer level" is 14 dB below "professional," but this is incorrect because the scales are different. "Professional" level (+4 dBu) is +11.8 dBV.

dBFS – Signal referenced to "full scale." This is commonly found in digital systems, as it is a very handy indicator of exactly how close you are to clipping a digital signal path or converter. Many modern systems maintain an internal digital signal path with far more dynamic range than any digital to analog output stage can handle. As such, a meter placed before the final bit-depth reduction point can display positive dBFS numbers when clipping occurs. The output signal is still clipped, of course – it's just that you have an indication of the reduction in gain required to avoid that clipping.

Of special note are the decibel numbers used for faders. On a fader, the 0 dB reference point is not necessarily tied to any specific voltage. Rather, 0 dB (or "unity") indicates that the net effect of the fader plus the upstream electronics is to neither attenuate or boost the signal flowing out of the fader. That output signal may be very small, or it may be in hard clipping. The fader markings only indicate the net effect on that signal's level, so you need to rely on other metering to determine how strong that signal actually is.

As an aside, the reason for saying "net effect" is because a bog-standard fader at its labeled 0 dB point is actually attenuating the signal flowing through it. Conventional wisdom states that a fader at "0" is neither attenuating nor boosting, but this isn't actually true for a simple, analog fader. As a simple, passive device, it is incapable of increasing the voltage of an input signal – it can't boost anything unless it's paired with active electronics. This being the case, the upstream, active circuitry is applying some fixed amount of gain, and when the fader is at "0" it is actually counteracting that gain.

Volts, Power, And Decibels: A Love Triangle

In the last chapter, life was just about to get useful. I was about to give you that list of decibel references, and then I tacked on this whole bit about how the references didn't use power...and thus behaved differently. The problem here is that decibels are meant as an expression of power ratios, but in pro audio we also apply the decibel to voltage ratios. When we do that, our decibel conversion math changes from this (power):

$$10\left(\log_{10}\frac{P_{measured}}{P_{reference}}\right)$$

to this (voltage):

$$20\left(\log_{10}\frac{V_{measured}}{V_{reference}}\right)$$

That's all fine and good, but why the change in the math? It all comes down to how voltage translates into power. To get there from here, the freeway entrance ramp is Ohm's Law.

This next picture depicts a very simple electrical circuit. There's a source of 1 Volt, and the current produced by that 1 Volt potential is flowing across a 1 Ohm resistor. It's true that the diagram is depicting a direct current, or DC circuit, and it's also true that audio folks actually use AC, or alternating current, but this will get us "plenty close enough."

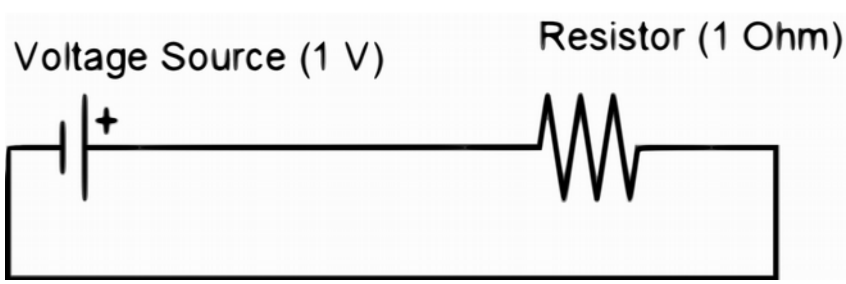

In this circuit, the relationship between voltage (which may also be called electromotive force, or ε, or electrical potential), current (which is Amperes or I), and resistance (which is Ohms, or R, or Ω) works out to being the exact representation of Ohm's Law:

$$I=\frac{V}{R}$$

In other words, 1 Volt across 1 Ohm equals 1 Ampere of current. If we increase the voltage and keep the resistance constant, more current will flow through the circuit. If we increase the resistance and keep the voltage constant, then the current flow will decrease.

You can think of this in terms of water. The voltage source is like a pump, the resistor is like a turbine or water wheel, and the current is how much water is flowing through the system. If we keep the pump as it is, but put in a turbine that requires more force to turn at a certain rate, then we would intuitively expect less total water flow through the system. If we then got ahold of a stronger pump, we could increase the water pressure and restore the total amount of flow.

So, how does power figure into all this? Well, Joule's Law lets us get a mathematical equation that relates power (watts or P) with two items that we're familiar with from Ohm's Law, namely voltage and current.

$$P=IV$$

That being the case, our simple circuit from Figure 3 gives us 1 watt of power. We have 1 Volt flowing across 1 Ohm, which gives us 1 Ampere, and 1 Ampere multiplied by 1 Volt gives us 1 watt. That's simple enough, but what happens to the power when we double that voltage? Take a look here:

$$2\,A=\frac{2\,V}{1\,\Omega}$$

$$4\,w=2\,A*2\,V$$

By doubling the voltage, we've gotten four times the power. In terms of decibels and power, that works out to this:

$$10(\log_{10}\frac{4w}{1w})=6.02\,dBw$$

To get an equivalent decibel answer when working directly with voltage, we would need to multiply by 2. Thus, “decibel math” with voltages requires us to multiply the logarithm result by 20 (2*10) instead of 10.

$$20(\log_{10}\frac{2V}{1V})=6.02\,dBV$$

This may seem strange, but remember that what we're thinking of are sine waves, which have a symmetrical “swing” away from a neutral point. In audio, we most commonly talk about the magnitude of that swing in absolute terms – we disregard whether the swing is positive or negative. If we say that the voltage swing has doubled, we mean that it has twice the magnitude in both the positive and negative directions. If that's the case, then power increases with the square of voltage...as evidenced if we rewrite our power calculations to use voltage and resistance only. It works like this:

$$P=IV$$

$$I=\frac{V}{R}$$

$$P=\frac{V}{R}*V$$

$$P=\frac{V^2}{R}$$

RMS: It's Better Than Average

In a couple of spots, you may have caught me referring to voltages with “RMS” tacked on to the end of the numbers. The reason to qualify a voltage number as “RMS” is so that we can reduce a constantly changing voltage into a simple, single number that represents an average of the voltage in a circuit.

The thing is that the voltage coming from the amplifier, being an audio signal, is time variant. We could talk about the voltage at some particular instant in time if we wanted to, but that wouldn't really help us much, especially if we chose the particular instant where the voltage was at zero. (In order to go from positive voltage to negative, and then back again, you have to make “zero crossings.”) We could take an arithmetic average of the voltage, but if we're talking about a nicely symmetrical sine wave, that simple mean of the voltage values is zero. What we need is an RMS (Root Mean Square) average. RMS uses the square-root of the mean of the squares of the values involved.

Let's say that we've done a little bit of (very coarse) measuring on a sine wave:

Measured Values	**Squared Values**
0	0
1	1
3	9
1	1
0	0
-1	1
-3	9
-1	1
0	0

What's ever so handy is that multiplying a negative number by a negative number gives us a positive result, and so the negative-going half of the sine wave no longer poses an arithmetic problem. We add up our new values and divide by the number of values (9) to get a simple mean of about 2.44. The square root of 2.44 gives us an RMS figure of 1.56.

I should note that this example is only to show the basics of how RMS works. Sine waves, being continuous, aren't really represented well by a handful of numerical values. In reality, the RMS value of a sine wave with a peak amplitude of “3” is much higher than 1.56 – our answer was skewed downwards because of all those “0” values I included. For a pure, undistorted sine wave, it's much more accurate (and much easier) to use this formula to find the RMS value. The “a” is the peak amplitude of the wave:

$$RMS_{sine} = \frac{a}{\sqrt{2}}$$

For example, using the sine wave from above:

$$2.12 = \frac{3}{\sqrt{2}}$$

For sine waves, the RMS voltage quite handily translates into what the voltage would be if the circuit were DC. In this case, the RMS value is represented by the greenish line that intersects the red-colored sine curve:

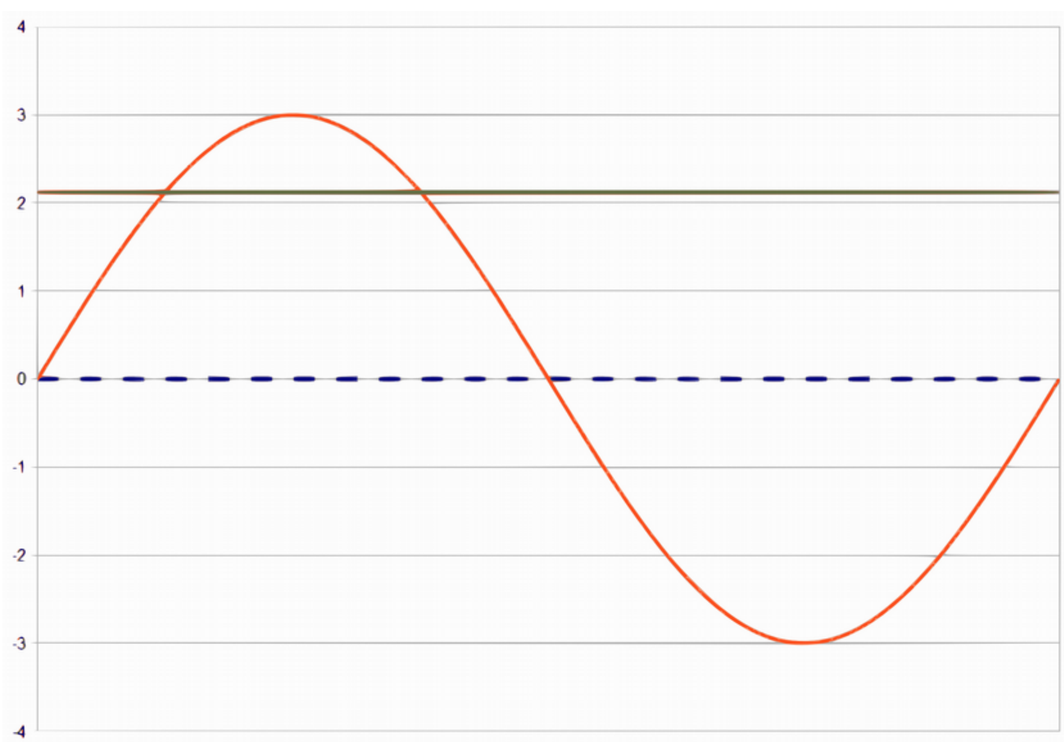

I should also mention that amplifiers are usually rated in RMS, or "continuous" power. Some folks rail against this as being meaningless for reasons of purity in the discipline of physics. I don't particularly see it as a problem, as long as you realize that what it is really meant by "RMS power" is "the average power derived from an RMS voltage provided by this amplifier into a given load."

When an amplifier manufacturer publishes this number, it most certainly is useful for getting a basic idea about amplifier performance – but you also have to be aware that it doesn't give you the whole picture. Actual music isn't much like test signals, especially test signals that are pure sine tones. Musical peaks can go far above the average level of the program material, and so an undistorted amplifier without dynamic range compression happening pre-input may be delivering nowhere near it's stated "RMS" or "continuous" power rating. (Some people use this as an excuse to connect dangerously over-powered amps to their loudspeakers. It's all fun and games until there's an accident, or someone decides to run the system too hard...)

Power, And The Loudspeakers It Pushes Around

This topic is a big thing for audio folks, because pretty much everything we do comes down (eventually) to our output transducers. That is to say, loudspeakers. If we didn't have some way of converting electrical signals into air pressure, audio dudes and dudettes wouldn't have much of a job – at least, not as we recognize the job today.

Loudspeakers are an enormous topic in and of themselves, but the main thing to tackle right now is the idea of how loudspeakers and power come together.

The first thing to get at is the idea of impedance. Impedance, like resistance, is opposition to current flow. Impedance is more complex than resistance, however, and is meant to describe opposition to flow in circuits where voltage varies over time. Circuits that utilize alternating current are not only subject to "straight line" DC (Direct Current) resistance, but also feel the effects of shifting magnetic fields and capacitance. (As an aside, capacitance is the ability to store electrical energy by means of electrical fields, as opposed to, say, chemical means.) For audio folks, the big thing to realize is that impedance changes with the frequency of voltage being put across a circuit.

A visualization of this can be seen below. The picture is a (very) roughly generalized representation of a loudspeaker driver's opposition to current flow when graphed against frequency. (Please note that this plot assumes a "bog standard" loudspeaker driver that uses a voice coil and a magnet.)

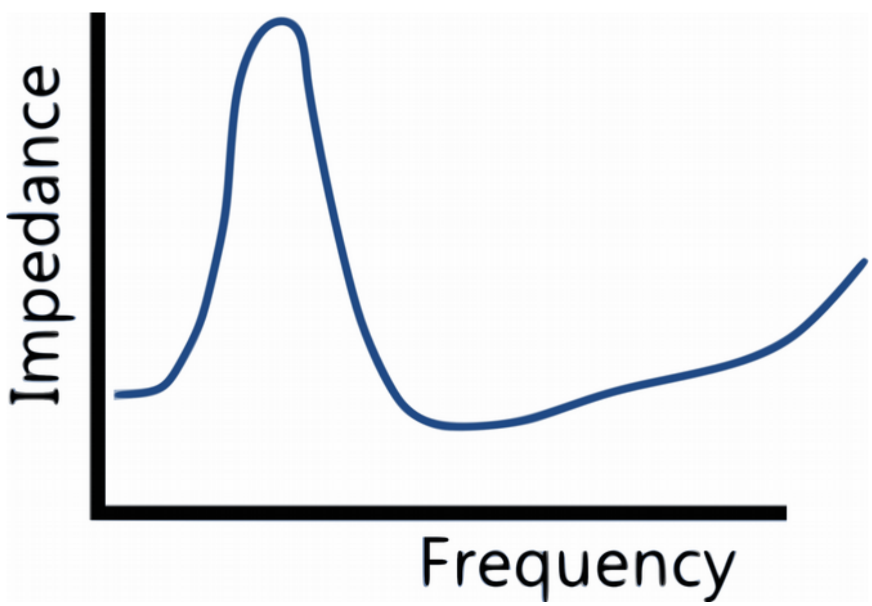

There's this marked impedance peak down at the resonant frequency of the driver, because that's where the driver vibrates the most readily. You would think that this would cause an impedance drop instead, but the driver being able to cycle the most freely means that it can also move the most through the magnetic field provided by the magnet. This being the case, the driver can “push back” to the amplifier with the most “force” at its resonant frequency.

If you think this is wild and wooly, it gets even more interesting when you stick that loudspeaker in a ported box, or mate it to a horn, or add a passive crossover, or...you get the idea.

Now, the reason that manufacturers rate drivers and loudspeaker systems with a single number for “nominal impedance” is that it provides a “pretty close enough” shorthand for what a loudspeaker's impedance looks like through the frequency range that the manufacturer expects to be in play. An audio human ought to be aware that impedance doesn't come down to a constant number, but there's no reason to be overly concerned about the curve – unless you're designing drivers or loudspeaker enclosures, of course.

So, having it in our heads that loudspeakers oppose current flow as a kind of exotic resistor, we can begin to understand what happens when we start connecting loudspeakers and the amplifiers that love them. Here's a picture of a simple setup:

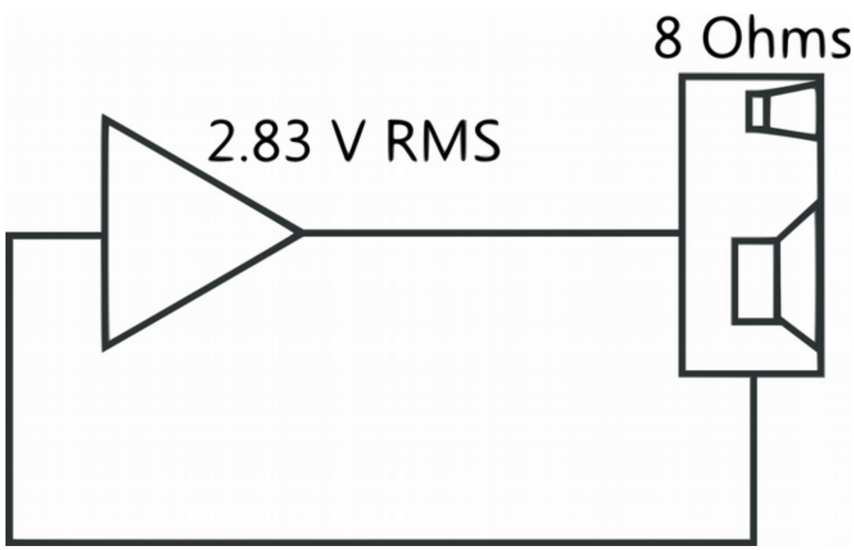

We've got one amplifier, and one loudspeaker system. We're going to take it on faith that the loudspeaker system's nominal impedance really is what the manufacturer says it is, 8Ω. (The Omega is the symbol for Ohms.) Our power amplifier is giving us nice, steady output of 2.83 V RMS.

Just a little while ago, we were talking about calculating power based on voltage and resistance. If we go back to that, and then substitute the numbers from the simple amp-and-speaker system that was pictured earlier, we get this:

$$P=\frac{V^2}{R}$$

$$1\,w=\frac{2.83^2\,V}{8\,\Omega}$$

So, 2.83 V RMS into an 8Ω load is just about 1 watt of continuous output. (It's actually 1.0011125 watts, but referenced to an even 1 watt, the decibel difference is 0.004 dB. Not worth losing sleep over.) Incidentally, this is why manufacturers use 2.83 V RMS so often when talking about loudspeaker sensitivity – put that into an 8Ω loudspeaker, measure the SPL from 1 meter away, and there's your 1watt/ 1 meter sensitivity figure. (You can also measure from 10 meters and then mathematically derive the 1 meter SPL...)

No Free Lunches Will Be Served

I used a simple setup to help illustrate voltage and power into loudspeakers, but you may have started to wonder about what happens in more complex scenarios. To crack that topic, we need to get out the age-old metaphor of electricity as water.

Let's say that we have a water pump, some pipes, and a couple of water wheels or turbines. The pump is like an amplifier, the pipes are like interconnection cables, and the turbines are like speakers. Just like an electrical system, there is pressure (voltage), a flow amount (current), and opposition to that flow (impedance/ resistance).

The picture below shows an arrangement of these items where two turbines are connected in series; the second turbine is fed from the output of the first. In this case, the pump creates a certain water pressure at the input to the first turbine. The first turbine resists the flow of water somewhat, lowering the pressure available to the second turbine, which also resists the water flow to some degree. The total opposition to flow that the pump experiences is simply the addition of the resistance of the two turbines.

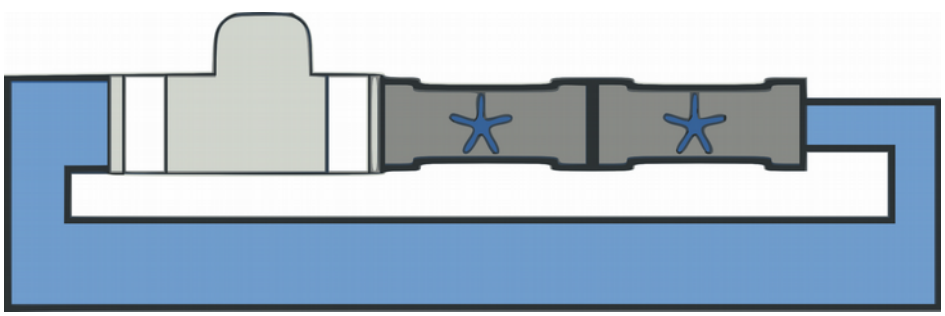

This kind of connection scheme is possible in the world of audio, but it's rather rare in the case of loudspeakers. For amps and speakers, it is far more common to see a parallel connection, like this:

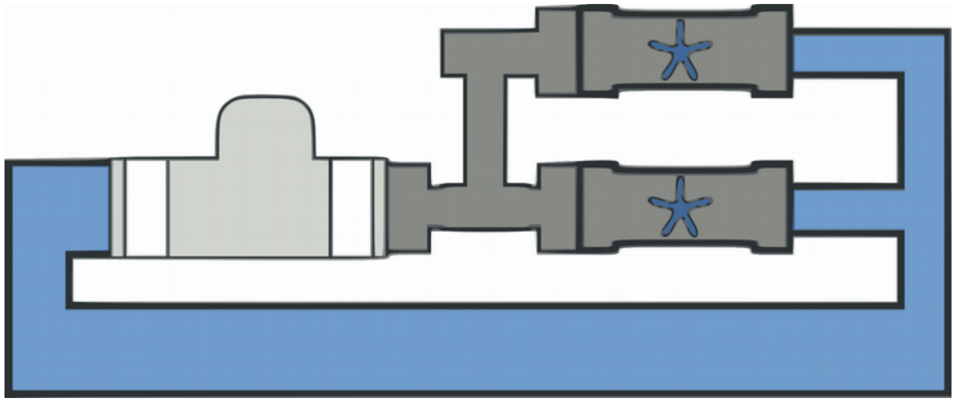

In a parallel connection, both turbine inputs are driven from the output of the pump. Each turbine experiences the same pressure, and the flow opposition experienced by the pump is the reciprocal of the sum of the reciprocals of the turbine resistances. In electrical math, that looks like this:

$$R_{total}=\frac{1}{\frac{1}{R_1}+\frac{1}{R_2}+\frac{1}{R_3}\ldots+\frac{1}{R_n}}$$

But...why?

Assuming that the pump can provide an adequate amount of flow for either system at a pressure of, say, “1”, two situations can develop. In the first picture, the pump has to create a total system pressure of “1” where the flow of water has only a single path. The flow is restricted once, and then restricted again. Each turbine added to the path makes the system harder for the pump to push against, because the pump has to push against one turbine, and then the next, and then the next. Assuming that the system isn't leaking, and the pump isn't letting water go the wrong way, then each turbine has to see the same amount of flow. That is, if one liter of water is going down the pipe every second, then each turbine has to pass one liter of water every second – no ifs, ands, or buts. The pump creates a total system pressure necessary for one liter of water to go through through every second, and the pressure across each turbine will be enough to keep the flow constant. That pressure will not necessarily be the same for each turbine, and each turbine only gets part of the total system pressure. This makes sense, because if the first turbine is very hard to push, and the second turns very readily, we would expect the second turbine to see less pressure than the first – though still enough to keep the flow through the system constant.

If that can be said to be fair enough, then the happenings for the second picture go as follows. (Remember that the system isn't leaking, and the pump is such that water can't flow the wrong way through.) Just as in the previous example, the pump works to get the total system pressure up to “1”. Because that pressure can flow directly to each individual turbine, each turbine sees the same pressure. The pump experiences the opposition to flow from both turbines, but it pushes across both turbines simultaneously. It doesn't have to get across turbine one before it can push against turbine two. While there is a total system flow (as you would expect), a very resistant turbine outputs a smaller amount of flow into the total, whereas a very free turbine contributes a larger flow output. If we keep adding branches with turbines to the pump connection, the pump effectively sees a larger and larger “pipe” connected to its output. As the pipe gets bigger, the flow opposition drops – but it also gets more and more difficult to pressurize. (You can try this with a kitchen sink. Take out anything that reduces flow through the drain, and you can run the faucet at maximum pressure without filling the sink. Put a food catch in the drain hole, and the faucet can now “pressurize” the sink to the point that it starts to fill.)

So...why is there no free lunch?

Some folks get all excited when they realize that they get more total power out of a power amplifier as the total system impedance is reduced. A power amplifier, as we experience them, is a device that is meant to create a requested voltage across its output terminals into whatever is connected. Effectively, it's an electrical pump. Ask it for 50 V RMS, and it will do its darndest to give you exactly that as a total system voltage. Assuming that you're within the design limits of the amplifier, this total system pressure is created whether the amperage (flow) is small or large.

If we take that amplifier, and connect a single 8Ω loudspeaker to the terminals, we get the following:

$$V = IR$$

$$50\,V = 6.25\,I * 8\,\Omega$$

$$P = \frac{V^2}{R}$$

$$312.5\,w = \frac{50^2\,V}{8\,\Omega}$$

If we connect another 8Ω speaker in parallel to the first, the total system impedance looks like this (the reciprocal of the sum of the reciprocals of the impedances) :

$$R_{total} = \frac{1}{\frac{1}{R_1} + \frac{1}{R_2} + \frac{1}{R_3} \ldots + \frac{1}{R_n}}$$

$$4\,\Omega = \frac{1}{\frac{1}{8\,\Omega} + \frac{1}{8\,\Omega}}$$

Now, this theoretical power amplifier I've conjured up can definitely put 50 Volts into 4Ω, so we get this:

$$625\,w = \frac{50^2\,V}{4\,\Omega}$$

“Cool! More power!” Well, yes...except power amplifiers won't quite give you twice the power for half the impedance, if you want the distortion created to be constant. (I'm assuming that this amplifier isn't running at “full tilt boogie.”) Further, there's this whole issue of the current involved:

$$50\,V = 12.5\,I * 4\,\Omega$$

The current that the amplifier has to provide has doubled. The amplifier is working harder, generating more waste heat, and is probably creating more distortion artifacts (whether they're audible or not is another thing).

This is why amplifiers are given a minimum impedance rating by the manufacturers. At some point, the poor amplifier just can't fill an “enormous pipe” with the necessary pressure – not without melting itself down, at any rate. Drop the impedance to 2Ω, and the amplifier has to generate a flow of 25 Amperes to keep up a “pressure” of 50 V RMS. There are certainly a good number of power amplifiers “out in the wild” that can do 25 Amperes, and more, but you can't just blindly assume that the amplifier you have on hand is one of them.

Also, please remember that what we've been calculating are totals for the whole system. Even in the best-case scenario, each 8Ω enclosure connected to our power amplifier is still getting 312.5 Watts. The amplifier has not become some sort of magical creature that puts more power into each box – rather, it has supplied more power to the system as a whole, and each component in the system still has to share that power as per the demands of physics.

The Area Under The Curve

The compact form of this topic is that “power is effectively the area under the voltage curve.” What this naturally brings up is integral calculus, because mathematical integration is the finding of the area under a curve.

Now, don't worry. We're not going to get into actually doing integrals on various waves, mostly because I myself can't do them “by hand” in a meaningful way. I don't have the background. However, we can talk about the implications of the concept, and how it works in a general sense.

The first thing to define is what “under” a curve means. For our purposes, “under the curve” means “between the curve and the x-axis.” This is important to understand, because one's first instinct is to define “under” as the area from the graph, all the way down to the bottom of the page. However, this wouldn't give you much in the way of meaningful answers – you'd get “negative infinity” all the time, if you think about it. No, what we're concerned about is how far away our graph/ function/ voltage/ whatever is from “0,” or any other reference point that we might choose. This picture shows the area under the curve (shaded red for the -x range, and blue for the +x range) for the function y=x. We're only looking at x from -1 to 1, for the sake of simplicity:

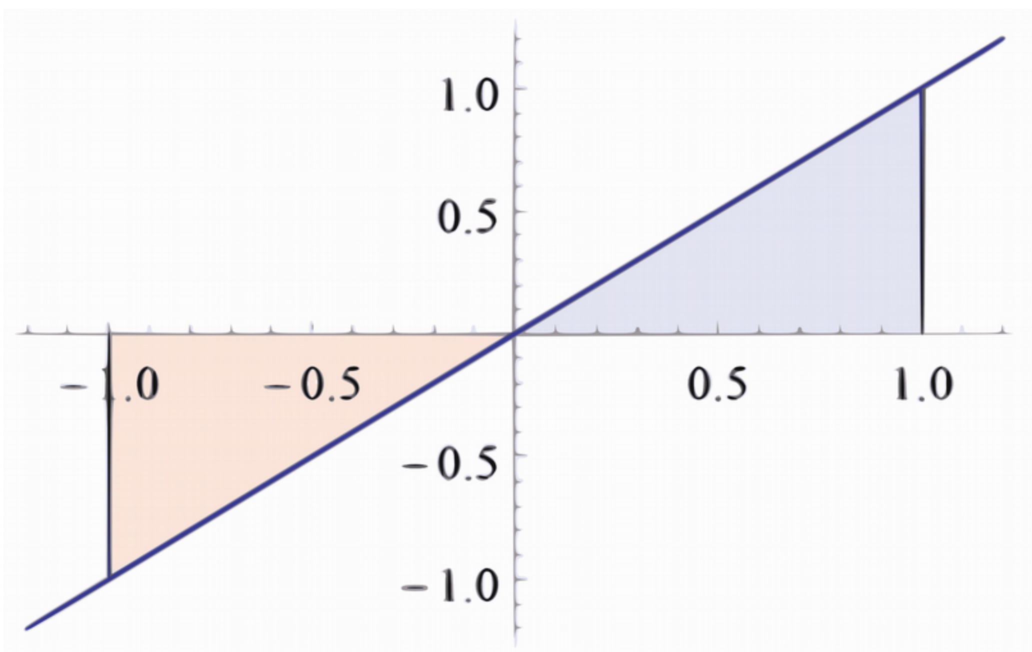

Now lets look at another picture. This is a sine wave that peaks at 2 volts. We're looking at it from 0 to 0.001 seconds (or 1 millisecond), and we can see that one full cycle of the wave has occurred. In other words, it's a 1 kHz sine wave. I've highlighted the area under the curve:

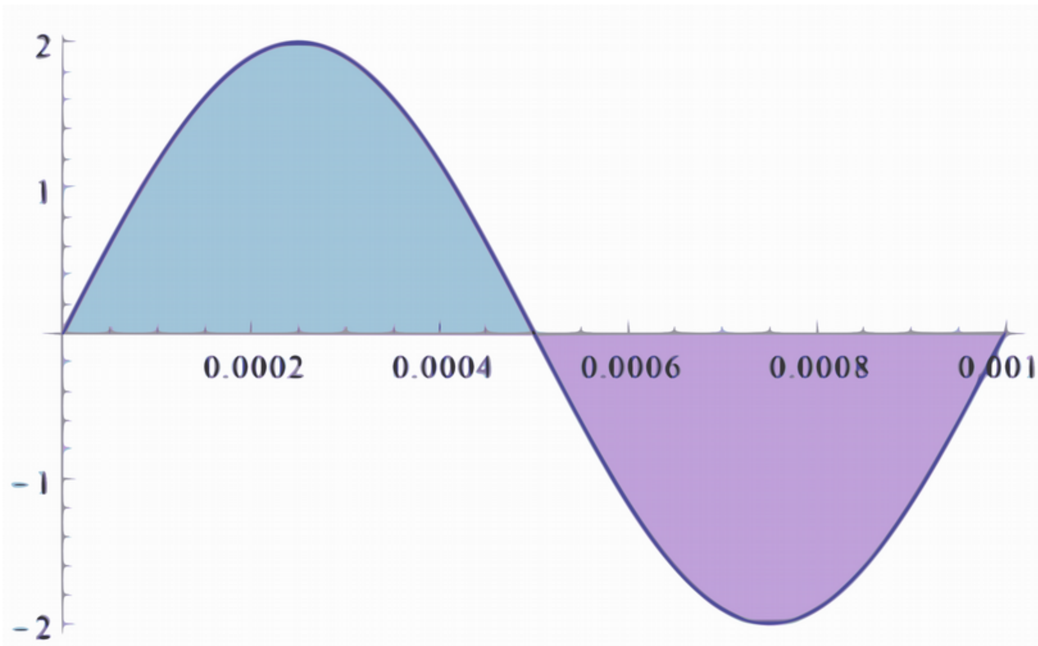

This helps us to visualize, say, the power being delivered to a loudspeaker by an amplifier. You can see that, if the voltage delivered is time variant, so to is the power derived. At each “zero crossing” (three being visible), the voltage delivered to the loudspeaker is zero, and so there is no power being delivered to the loudspeaker. However, that is only for those instantaneous points in time. At every other point in the wave cycle, some amount of voltage is generating some amount of power. In the same vein, maximum power is generated at the peaks of the wave – but only for an instant. The majority of the time, the loudspeaker is subjected to a voltage that is both non-zero and non-peak. This is why pro audio types are so keen on “continuous power,” or the not-correctly-named-but-accepted-anyway “RMS power.” Continuous/ RMS power ratings give us a meaningful way to predict what we might observe from a power amplifier delivering voltage to a loudspeaker.

Now, let's examine the curious case of a power amplifier being driven into VERY hard clipping or limiting. In the case of clipping, we get a whole bunch of harmonic distortion products being added to our desired signal. In the case of limiting, especially if the limiter is well designed, large amounts of harmonic distortion are avoided while the audibility of the limiter's activity is also minimized. In either case, the output signal becomes more and more like a square wave. This picture shows us our sine wave with a square wave overlaid:

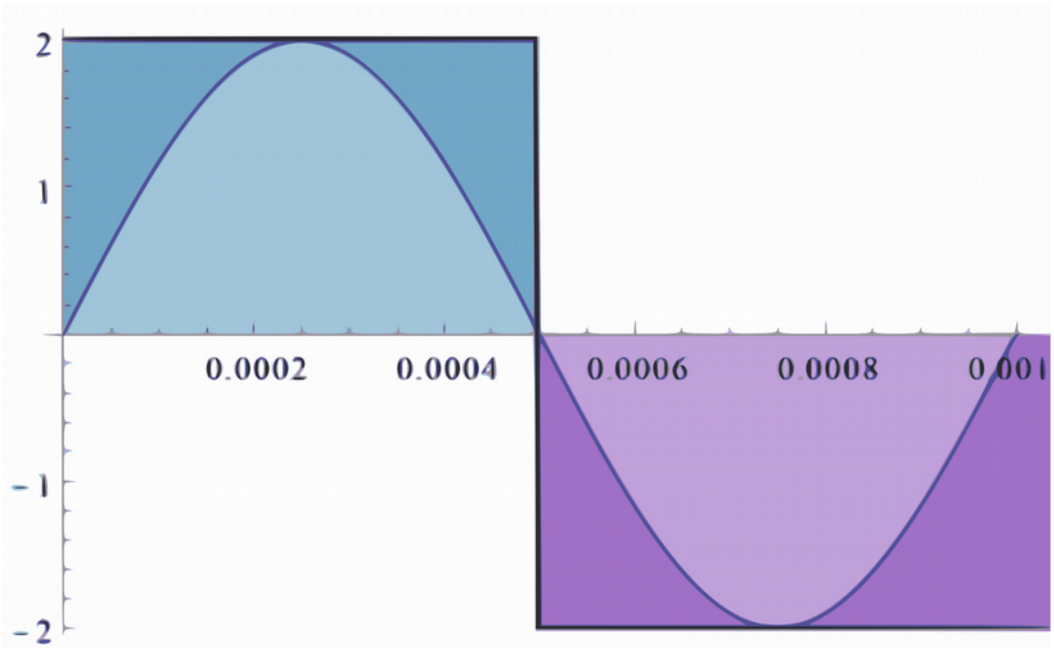

What you should immediately notice is just how much more area the square wave has under its curve. As the power amplifier is driven harder and harder, the continuous power delivered to the loudspeaker more and more closely approaches the peak power. As opposed to a condition where most of the voltage is non-zero and non-peak, the "mathematically ideal" square wave spends all of its time at peak voltage, instantaneously switching from positive to negative-going amplitude.

Side note: Mathematically ideal square waves aren't possible with power amplifiers and speakers. We'd need infinitely fast and powerful power supplies, as well as a way to make the loudspeaker cone move instantly from positive to negative displacement. We'd also have to remove the modern protection systems from our power amplifiers that keep us out of really hard clipping. Still, we can get "plenty close enough" to be dangerous...

The major point in all of this is to be aware of the assumptions being made between the rated continuous power output of an amplifier, and the rated continuous power of a loudspeaker. From the standpoint of preventing thermal failure – that is, not overheating or "cooking" your loudspeaker components, the thing to realize is that the continuous power rating of an amplifier is not actually "maximum power deliverable." Instead, it is "the DC (direct current) equivalent power deliverable with distortion characteristics deemed to be tolerable.” The amplifier is actually capable of delivering a good deal more than its rated continuous power, if one is willing to accept distortion and/ or limiting as a side effect.

This is one of the roots (there are other important ones, of course) of the myth of “underpowered amps cause people to blow up loudspeakers.” People want more SPL than a system can safely deliver, so they drive their amplifiers “into the clip lights”, and then cause thermal failure of their loudspeaker components. They then wrongly conclude that clipping is “dangerous” without respect to the power being delivered, and advocate for larger power amplifiers into the same loudspeakers. The reality is that the amplifier connected to their loudspeakers became too powerful as the RMS voltage approached the peak voltage. They kept asking for more area under the curve, and the amplifier obliged – to a point, of course.

Sir Isaac And Your Transducers

Quite a bit of the material presented so far is “popular” for the analytically-minded audio person. The reason for this is because a lot of the math and science translates directly into usable numbers. Decibels and watts, for instance, are near-instantly usable, quantitative descriptions of sonic events, or the signals corresponding to sonic events.

By my (admittedly limited) observations, what is far less popular in the mathematical and scientific arsenal of pro-audio folks is classical mechanics. I believe that this is largely, though not exclusively, due to how the numerical output of the relevant equations is not directly useful for pro-audio types. I myself cannot claim to have any desire to know exactly how fast a subwoofer driver accelerates, nor can I claim to have any use for knowing the number of Newtons required to accelerate that driver at a certain rate. However, understanding how aspects of classical mechanics, especially Newton's second law, relate to audio devices can be quite eye-opening.

So – what is this “Newton” business? Well, Sir Isaac Newton was the sort of chap who was interested in how and why things moved like they did. He was interested enough to be a developer of the methods of calculus, which is the branch of mathematics that specifically deals with things that are in motion. He also came up with three “laws” of motion, the second of the laws being very important for audio. When you do the kind of work that Newton did, you are very likely to have a scientific unit named after you, and so we got the Newton unit. One Newton is the force required to accelerate a one kilogram object by one meter per second, per second.

For pro-audio folk, Newton's second law is the most immediately important, even if they don't realize it. The second law states that an object accelerates in proportion to the force applied to it, and acceleration is inversely proportional to the object's mass. Put less elegantly, the harder you push on something, the faster it gets moving – and heavier things have to be pushed harder to get them moving at the same rate as lighter things.

The way this concept boils down is usually presented as “Force equals Mass times Acceleration”:

$$F = ma$$

As it turns out, this relationship is very important in the audio world – actually, it's critical. The reason for it being critical is because it deals with the make-or-break components that form the entry and exit points of our signal chains. These components are the transducers.

Transducers are devices which convert one form of energy into a different, yet corresponding form of energy. Microphones, for instance, convert sound pressure waves into electrical signals that represent those waves. Loudspeakers convert electrical signals into sound pressure waves. These two activities, done poorly, can completely negate otherwise excellent devices sitting in between them.

Since the last few articles have been talking about loudspeakers, let's use a couple of those as an object example. A fictional company, Sonny's Subwoofers makes all kinds of low-frequency loudspeaker drivers. Their basic model is the 15” diameter “Kaboom,” which has a deluxe version in the 18” diameter “Kaboom+.” The Kaboom+ has a diaphragm mass that's 1.5 times greater than that of the Kaboom. Not only is it bigger, but it has more additives to increase its performance – there's more material in it.

If we invent our own units, we get these “synthetic” equations:

$$1\, amplifier\ force = 1\ Kaboom\, mass * 1\, displacement\, unit / second^2$$

$$1.5\, amplifier\ force = 1.5\ Kaboom\, mass * 1\, displacement\, unit / second^2$$

To move the significantly heavier cone of the Kaboom+ at the same rate of acceleration as the lighter Kaboom, the amplifier needs to have the necessary power to let us achieve 1.5 times the physical force on the driver.

There's something else to consider. Newton's first law says that an object in motion will stay in motion unless a force acts upon it. What this means is that that Kaboom+ is not only more difficult to start, it's also more difficult to stop. If the Kaboom+ is not appropriately damped mechanically (designed so that it stops moving as quickly as possible when signal is removed), and the amplifier has trouble damping the driver's movement electrically, the Kaboom+ will have more problems with “ringing” than the basic Kaboom. This ringing causes the sounds produced by the loudspeaker to “smear,” that is, to last longer than they should, and can also cause the loudspeaker to produce entirely unwanted tones.

Important: The above is not to imply that a larger diameter driver is necessarily inferior to a smaller diameter driver. Bad designs come in all sizes. A well-built driver can have a large mass and sound excellent in its intended application. Indeed, we need sufficiently-sized drivers to reproduce low frequency material.

So, what about microphones? If we have two microphones picking up the same airborne sonic event, then (all else being equal), a pickup with low mass will be more sensitive than a pickup with higher mass. Microphone elements with lower mass tend to be superior in terms of their transient response (their ability to react to large changes in level over short time periods), as they have greater acceleration for a given input force. However, better transient response does not necessarily make a microphone better for every application. It is easier to make a rugged microphone if you allow the diaphragm to have a greater mass, and damping that mass to keep it from ringing may not be so difficult as to be prohibitive. It's true that the most accurate microphones available are low-mass devices, but a higher-mass and lower accuracy device may be a better choice for a particular use.

Microphonic Pickups

Now that I've talked your ears off with all of these conceptual bits, let's see how some of them integrate when we deal with a couple of microphones. Microphone one will be a dynamic moving-coil type, and microphone two will be a condenser.

Microphone one works by way of electromagnetic induction. It has a diaphragm that's connected to a wire coil. The wire coil is suspended in a magnetic field. When the diaphragm moves, the coil also moves through the magnetic field. This generates an electrical current which is analogous to the sonic event at the diaphragm.

Microphone two uses the principles of capacitance to operate. It also uses a diaphragm, but the diaphragm does not have a coil attached. Instead, the diaphragm is some distance away from a non-moving backplate. The whole assembly has a bias voltage applied so that the diaphragm and backplate become charged plates with an insulator between them – the insulator being air. At rest, the assemblage has a certain capacitance, or ability to hold a charge. When exposed to sound pressure waves, the diaphragm moves in relationship to the backplate, and the system capacitance changes as a result. If the diaphragm moves toward the backplate the capacitance drops, which means that the system must release charge. The reverse is true if the diaphragm moves away from the backplate. This release and taking up of charge naturally results in current flow which is analogous to the sonic event at the diaphragm.

Here are some simplified cutaways of these microphone designs:

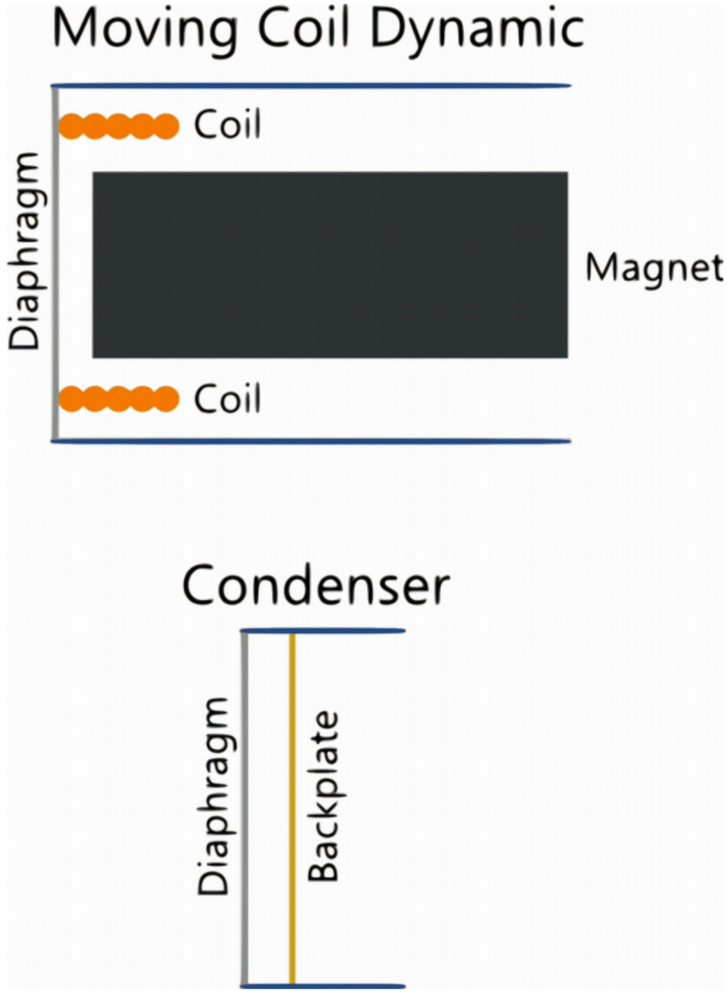

Our microphones are set up side by side in a perfectly anechoic chamber; acoustical reflections simply don't occur. Each microphone's element is precisely 2 meters from a sound source. The sound source is producing a 1 kHz tone at 100 dB SPL, averaged over 1 second, measured at 1 meter from the source.

Question: What is the SPL at the microphone elements?

Since we don't have to account for reflected sound, we can simply assume that the basic inverse square law is in effect. If the distance to the source is doubled, the observed intensity should only be one quarter of that experienced at the closer distance. One quarter intensity is 6 dB down from the original intensity – we use the “10 log” formula because we are concerned with power (intensity being power divided by area).

$$-6\, dB\, SPL = 10\left(\log_{10}\frac{1}{4}\right)$$

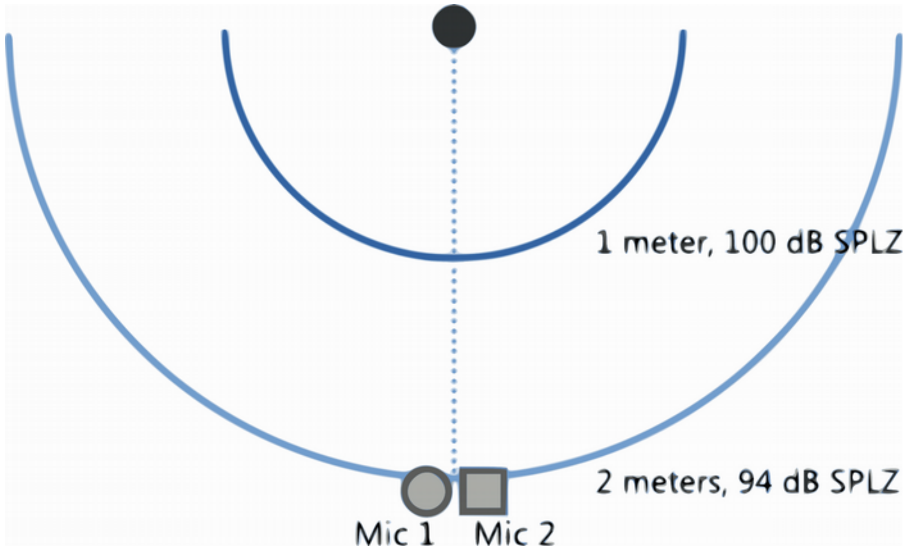

So, the SPL at the microphones is 94 dB SPL, unweighted, averaged over 1 second. (I wanted this number specifically. The reason is that microphone sensitivity is most often expressed in terms of a decibel output relative to 94 dB SPL at 1 kHz.)

Microphone one has a sensitivity of -55 dBV at 1 Pascal/ 94 dB SPL. Microphone two has a sensitivity of -38 dBV at 1 Pascal/ 94 dB SPL.

Question: What might account for the higher sensitivity of microphone two?

Remember that microphone two is a condenser microphone. The mass of its moving part, which is just a diaphragm, can be much smaller than the mass of the dynamic mic's moving part, which is a diaphragm and a wire coil. For the same amount of force on each mic's diaphragm, the diaphragm of mic two will have substantially greater acceleration.

Question: What is the voltage output of each mic in this situation?

Since the sensitivity measurements given to us are in dBV, we know that the reference point involved is 1 volt. We're working in volts and not power, so the multiplier we'll need to use is 20, instead of 10. For the first mic:

$$-55\,dBV = 20\left(\log_{10}\frac{v}{1\,V}\right)$$

$$\log_{10} v = \frac{-55}{20}$$

$$v = 10^{-2.75} = 0.00178\,V$$

...and here's the second microphone:

$$-38\,dBV = 20\left(\log_{10}\frac{v}{1\,V}\right)$$

$$\log_{10} v = \frac{-38}{20}$$

$$v = 10^{-1.9} = 0.0126\,V$$

This illustrates an adage that was handed down to me when I was in school: “The mic pre is the highest gain stage in the chain.” Whether or not you're into all the nuances of different microphone preamps, the fact remains that the sometimes 50+ dB they have to add to a signal to get it up to “nominal” level in a professional system is a lot of work to do. That work has to be done cleanly and accurately (or somewhat inaccurately in a pleasing sort of way). Even large power amplifiers don't have to apply that kind of voltage gain.

Going Loud

The last section was about microphones, but now we'll be transducing in the opposite direction; we'll be turning electricity into sound pressure waves. To do that, we'll be using a device that looks remarkably like a dynamic moving-coil microphone. The difference is that this device is meant to handle much bigger excursions of its diaphragm, is built at a much larger scale, and (of course) is meant to produce sounds of its own.

As an aside, yes, you can wire up a garden-variety loudspeaker to a microphone input and get sound into a system with it. The frequency response won't be much to write home about, but you can do it. As a general-purpose input transducer, this sort of "kludge" doesn't make the grade. However, it can work handily as a single-purpose or special-effect trick. Indeed, on the market today is a mostly-meant-for-low-frequencies input transducer that is a loudspeaker mounted in a drum shell.

Here are the essential innards of a loudspeaker:

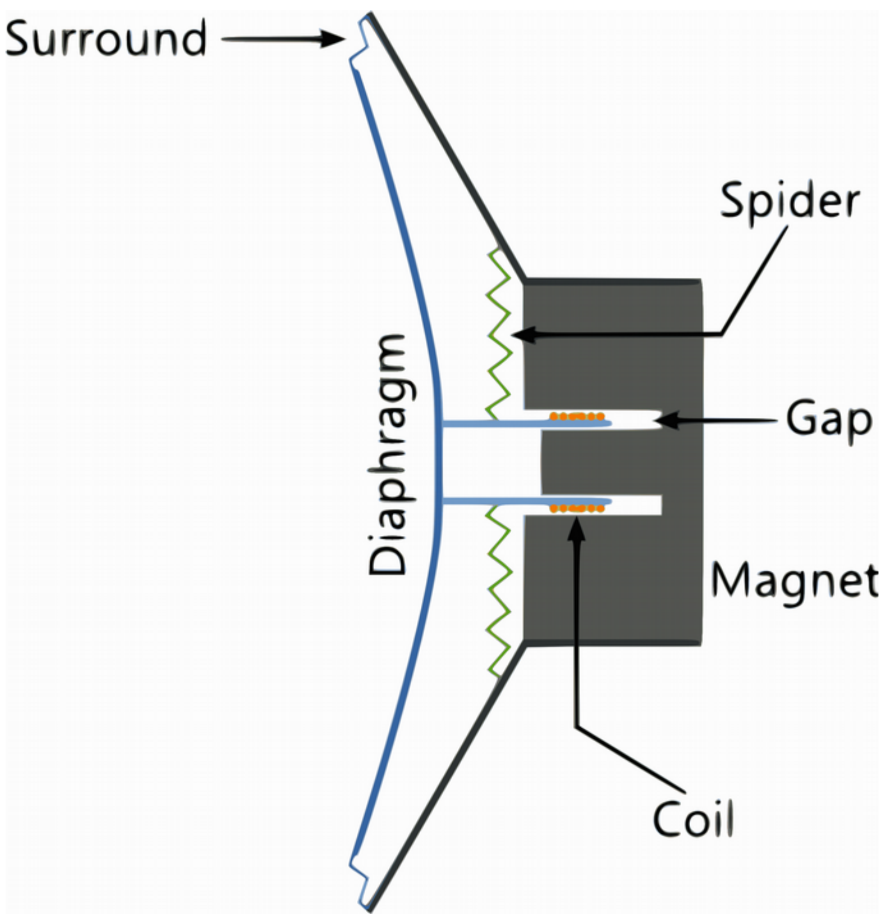

A loudspeaker's diaphragm is the "piston" used to impart a push to air molecules. The piston has to be suspended in some sort of frame to work properly, which is where the surround and spider come in. The suspension of a loudspeaker also has the important task of ensuring a consistent "return to zero" (or, "return to start") when the assemblage stops moving. Much like a dynamic microphone, the loudspeaker's coil is suspended in a magnetic field. The difference in operation is that a relatively large electrical input is presented to the coil, which turns the magnet/ coil pairing into an electric motor.

Let's assume that our loudspeaker has been loaded into an appropriate enclosure by an equipment manufacturer. The manufacturer claims that this particular driver can handle 500 watts of continuous input. They call it "RMS power" on the spec sheet, in defiance of the folks who will complain that this is the wrong terminology. (It IS the wrong terminology, but as mentioned previously in the series, you will be okay if you know what that term is really meant to symbolize.) The manufacturer also has a frequency response chart that shows the unweighted SPL produced by the assemblage when driven with 1 watt of input – and measured 1 meter away. At another spot in the spec sheet, they've reduced that measurement to a shorthand "sensitivity" number of 99 dB/1w/1m. That number is an average that they arrived at in some way that may or may not be clear to us.

This picture shows the frequency response chart from the spec sheet. The passband we're interested in for our application is highlighted in blue:

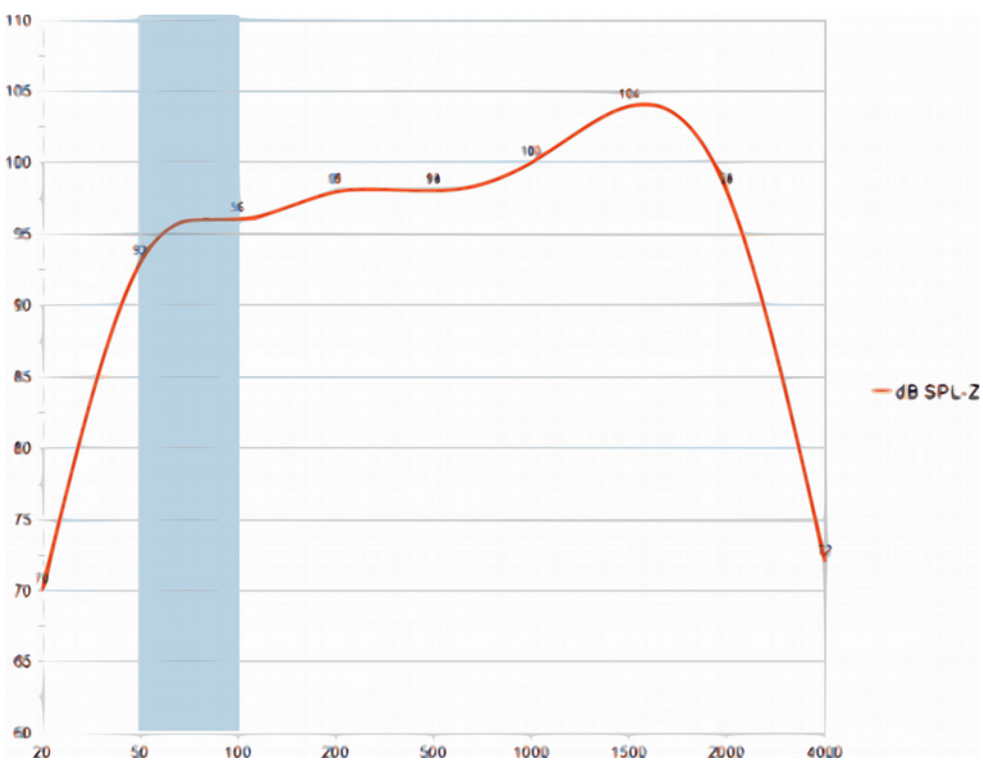

Question: What "worst case" unweighted SPL measurement would we expect to see if the loudspeaker was driven with the maximum recommended continuous power, and measured at 1 meter?

To answer this question, we first have to determine where, in our intended passband for this device, we get the least SPL. Taking a look at the chart, we can see that the lower limit of our passband, 50 Hz, is where the loudspeaker produces roughly 93 dB SPLZ at 1 watt/ 1 meter.

To figure out the theoretical SPL at 500 watts, we need to get out our logarithmic math. The good news is that this calculation is very easy. We're dealing directly in power, and our reference power is 1 watt, so there is effectively no calculation required for the division.

$$26.99\,dB = 10\left(\log_{10}\frac{500\,w}{1\,w}\right)$$

$$119.99\,dB\,SPLZ = 26.99\,dB + 93\,dB\,SPLZ$$

In theory, the device can produce a continuous level of 119.99 dB SPL-Z when driven at full power, measured at one meter. Again, this is assuming the worst case in our passband, as the loudspeaker system in question is noticeably less sensitive at 50 Hz than at 100 Hz.

Question: Why might the actual SPL measurement be higher than we expect? Why might it be lower?

It is quite possible that our theoretical numbers will not match up with what we actually experience from this device. For instance, it is possible that the manufacturer measured their device in an anechoic chamber, and it is unlikely that we are in an anechoic environment. Depending on the manufacturer's test setup, their anechoic measurement may have effectively been "full-space," where all sonic energy not directly traveling to the measurement point is lost to the measurement.

As you might remember from the beginning of this series, such an environment for actually listening to a loudspeaker's output is very rare. If we're indoors, it's far more likely that we're in a situation that is somewhere at or between "half-space" and "quarter-space." Half-space is called that because it assumes that sound is being forced to radiate into half of a sphere, instead of a full sphere. Quarter-space and eighth-space follow on from this. In theory, each halving of the radiation space allows the observed SPL to increase by 6 dB, although this is dependent on the walls involved being large enough to reflect the wavelengths being emitted, and also on the walls being perfect reflectors.

The main reason for us to see less than the maximum output we expect is that of power compression. Power compression is caused by voice coil heating. In metals like copper, which is the most common material used for voice coils, resistance to current flow increases as temperature increases. If we run our loudspeaker with the maximum continuous input power that it can withstand, the voice coil will inevitably heat up. As it does so, it becomes more difficult for the amplifier to drive the loudspeaker assembly. As a result, we do not appear to get the full “theoretical” benefit from the power delivered. (This is not strictly true, of course – our problem is that our theoretical numbers don't take power compression into consideration.)

Another part of power compression is that a driver undergoing large excursions may very well have larger-than-intended portions of its voice coil leaving the gap in the magnet. As a result, the influence of the magnetic field on the coil is reduced. This mostly results in distortion products, but it can also contribute to voice coil heating. Thermal management for loudspeaker voice coils loses efficacy when the voice coil has traveled beyond where a designer expects it to be.

Question: What is the maximum distance from the loudspeaker for 85 dB SPL unweighted, assuming full continuous power is applied, and assuming a full-space environment?

To answer this question, we first start with our “maximum output at 1 meter” determination. In this case, our starting number is 119.99 dB SPL-Z. The difference between 85 and 119.99 dB SPL-Z is 34.99 dB. Theoretically, each time we double our distance to the loudspeaker, we lose 6 dB of measured SPL. Ultimately, we have to figure out how many doublings of distance it takes to get to a 34.99 dB loss.

Each time we double the distance, we can represent it as a power of 2. For instance, 2 raised to the first power is 2 meters, or double the distance from 1 meter. The double of that distance is 4 meters, or 2 raised to the power of 2. The double of that is 8 meters, or 2 raised to the power of 3. That 6 dB loss, then, becomes a number we can use to find the total “doublings of distance” involved. What this turns into is dead-simple algebra.

$$34.99\, dB = 6\mathrm{dB} * x$$

$$5.832 = \frac{34.99\,dB}{6\,dB}$$

So, 2 raised to the power of 5.8312 should give us our distance for 85 dB SPL-Z continuous:

$$57\,meters = 2\,meters^{5.832}$$

In theory, we should be able to go out 57 meters and still have 85 dB SPL-Z to work with. Of course, we do have to remember that we won't be running the loudspeaker at full throttle all the time, but at least we have some sort of idea about what is "in the ballpark" of expectation for the moments when we are running at our absolute maximum.

Part 2: Anecdotes and Observations

Introduction To Part 2

Quantitative descriptions of how sound works are very helpful, but they don't contain the whole story of the experience of live sound. You can have all the science completely memorized, and still be tripped up when you get a real sound system and real musicians in a real room.

Why?

Knowing all the theory doesn't necessarily create a good practitioner. A good practitioner is produced by knowing what theories apply at different times, how to apply those theories, and how to synthesize new and greater understanding of how theories come together in real situations. Some people are better at this than others, but I believe that most people can achieve equivalent levels of professionalism if they get the sufficient amount of experience. This amount is different for different people.

This part of the book, then, is a collection of insights and experiences that have shaped my understanding of live audio as a practical (and sometimes impractical) scientific art. Or...maybe it's an artsy science?

I guess it depends on the day.

There's still a good deal of “sciencey” and “analytical” thought in this part of the book, but it's far more on the end of application and overall results than getting some sort of numerical representation of an audio event.

It's interesting that my justification for including this part of the book is so much less involved than my justification for the math and science portion. This probably has to do with the tendency for “stories about the experience” to naturally be received as instructive. Very few people ask “when are we going to use this in real life?” when the example is about real life.

The hope for this part of the book is that you can compare and contrast your own experiences with what you find here. If there's an experience presented that you've never had, hopefully you can file it away as a useful tidbit for the day that you do have that experience.

Some Basic Principles Of Live Sound, Stated Briefly

1. The most basic principle of mixing a live show, to me, is being able to hear everything clearly without excessive SPL (sound pressure level). Sometimes this isn't possible to achieve, at which point the goal changes to prioritizing the relative importance of different sources, and focusing one's effort on making the most important sources clearly audible.

2. Great live sound starts with great sources. A source that sounds bad without reinforcement can only be helped so much by the application of technology – and the more of your technology you use to fix one problem, the less you have left over for everything else. After good sources, an engineer that makes good choices is critical, followed by having enough power and coverage for what one wants to accomplish. A step beyond that is to have truly superior transducers on both ends of the signal chain. However, I must emphasize that the very best equipment can be made to sound truly awful if the humans involved are unskilled or make poor decisions.

3. Live Sound vs. Studio Sound is an enormous discussion, but from my experience, one might boil it down to studio work having many more safety nets in terms of ability to go back and get things as close to perfection as possible. Live audio, on the other hand, is much more immediate, and mistakes at the time of presentation to an audience can not be fixed (though you can react to them so that they no longer apply for the duration of the show).

Why Is It So !@#$%^& Loud?

I guarantee you have been to a show that was too loud. Maybe it sounded good but was just too intense, or maybe it was really intense and also sounded terrible. You may have ended up asking yourself, “How can this gig sound so terrible with all of the money and time these people have invested in it?” In answer, let me share a few reasons why shows can get to be too loud. These reasons aren't necessarily in a particular order.

1: Mix-Person Preferences...or Deafness

Of all the possible reasons for a loud show, this one is the most obvious. Some folks mixing FOH (Front Of House) are talented, but are unfortunately possessed of the notion that “overwhelmingly loud is powerful.” There are audio humans out in the field who, when presented with any sound rig of any size, will try to run that rig as hard as it can be run. Some of these people can hear when a live-sound system is under strain, some can't, and some don't care.

Interestingly, one of the traditional sonic cues of “loud” is the sound of a system creating distortion. For some people, “loudness” isn't perceived until they start hearing that distorted edge, and so they will push a rig until they start hearing it. Modern, high-performance audio rigs can get very, very loud before this distortion begins to happen in the way that some folks are used to.

Even if they don't purposefully run the system all the way to the clip lights or limiting indicators, some audio techs just think that “this genre of music doesn't sound correct until it's a certain loudness.” In other cases, a sound-human may not even be going for “loud,” but may just be unconsciously compensating for accumulated hearing-loss. This can get especially nasty when that compensation takes place in a specific frequency range, like the “hurt zone” around 4 kHz.

2: It's Monitor World's Fault!

If a band is using traditional monitor wedges, as opposed to IEMs (In Ear Monitors), those monitor wedges are most definitely contributing to the volume in the room. Even though they aren't pointing at the audience, they radiate some sound in the audience's direction. If there's something acoustically reflective in their direct path, whatever sound was pointed at the musicians may end up (effectively) pointed at the crowd anyway.

Now, take the above paragraph and combine it with bands tending to want their monitors run at higher levels. Musicians, being into music (logical, right?), often like to have that music hit them at higher volume than an average person might prefer. Also, to really hear what *you're* playing, your monitor has to get ahead of all the other sounds around you, which can include both the crowd and everybody else playing on stage. Pretty soon, those monitor speakers are making a lot of noise.

With monitor-world cranking away, and that sound being audible to the audience, a problem starts to occur. The spill from the monitors imposes a natural lower limit on how quiet a show can be. What's worse, that spill usually sounds pretty muddy and incoherent – the high frequencies are probably much lower in level than you would typically want, and all that sound has been bouncing around everywhere anyway. To get the show's sound back into balance, the FOH engineer has to make the mix for the audience loud enough to compensate for what's audible from monitor land. The audio tech out front can choose not to run the main PA at all, of course, but this usually sounds so poor that it really isn't an option. As a result, the FOH PA must add something to the volume in the room, even if it's only a little. If the monitors are already very, very loud, then you can guess what happens next.

3: Oh, Those Funny Musicians

My monitor-world example above is actually just a subset of a whole issue called “stage volume.” Stage volume is the total loudness of everything happening in the area the performers occupy. Monitor wedges are part of that, as are drum kits, amplifiers, pianos, horns, acoustic guitars, and absolutely anything else that makes any kind of sonic contribution.

In exactly the same way as with monitor wash, the loudest person on stage imposes the minimum volume. If the drummer is drowning out everyone else, then the only way to get a balanced mix is to be louder than the drummer. There is no practical way to electronically reduce the level of the person banging away on the kit, and if they can't be convinced to tone down what they're doing...you're pretty much stuck.

Of course, drummers aren't the only loud people out there. Guitar players with amplifiers can make epic amounts of noise. There are plenty of loudspeakers used for guitar amplifiers that can make around or over 100 dB SPL (Sound Pressure Level) with just one measly watt of input. If the guitar player's amp can provide a still-not-that-much-power of 30 watts, that speaker can now hit you with 114 db SPL, depending on where you are. That's a lot of level in most clubs and bars. Oh, and let's add something on to that. A good amount of the most highly prized guitar amps get their signature sounds from distorting the tubes that provide the output power for the amp, meaning that the guitar player has to push them into their maximum level to get them to “sound right.”

Again, if any musician is drowning everybody else, the dude or dudette mixing for the audience can either compete, and balance the mix, or not compete and have an unbalanced mix. The correct choice is not always obvious, and with different people in the audience wanting different things, there may not even be a correct choice.

In The Zone

Some bands can walk up on stage, start playing, and sound amazing with no help. A good audio tech can get their overall sound to better fit a specific venue, but they would still be brilliant fun with nothing but a vocal mic or two. Some bands don't fit this mold. If there's no competent audio tech at hand, or very minimal PA, they just sound like a blast of notes that are hard to tell apart. (Sometimes they sound like nothing more than a blast of notes, even with a PA and a competent operator.) Why does this happen?

There are a number of potential reasons, but one of the biggest culprits I've run into over the years is the "people aren't in their zone" problem. This is not "The Zone" that people get into where playing is no longer a conscious thought. The zones I'm talking about are tonal ranges, or frequency areas where an instrument or voice is strongest.

Quite a few instruments have audible content across vast swaths of the audible frequency range. A drum kit, for example, can have enormous output across every frequency that you can hear. An electric guitar can have a lot of character defined from 100 Hz to 500 Hz, be very strong from 500 Hz to 2 kHz, and then still have a lot to say in higher frequency ranges. Still, there is a tendency for different instruments to naturally have a range where most of their sonic information lies, with other ranges being more about the particular instrument's "flavor."

Really excellent bands aren't just good at playing their instruments and knowing their songs. They also have a very good grasp of what zones their instruments and voices occupy, and know how to make room for each other. Although they may not use exactly this kind of language when getting their sound, they may decide that they really want the deep bass of 40 – 100 Hz to be provided by the kick drum. That being the case, the bass player dials up a tone that emphasizes 100 – 300 Hz on the low end. The guitar player who usually sets up on stage-left opts for sounds (and musical parts) that emphasize 400 Hz to 1.2 kHz, whereas the guitar player on the other side of the stage picks tones and notes that lie more in the range of 1.2 kHz – 3.6 kHz. (If that guitar player is good, those mid-high frequencies will lend clarity without being brash or harsh.) The lyrics are very important to the band, and the critical range for those is from about 1kHz to 5 kHz. As a result, vocal passages require the guitar players to back off a bit, and the drummer takes it easy with his cymbals - as cymbals can obliterate 2 kHz and up.

Not-so-great bands don't do this as well. Their instruments may sound really good individually, but put them in a group and they all manage to be present and absent at the same time. They don't give each other space. The bass player and drummer might both want to be driving the deep notes, and so you can either hear kick drum or bass guitar, but not both at the same time. The drummer might be totally in love with his cymbals, and smash them constantly, which eliminates vocal intelligibility with a wall of hiss. The guitar players both like to have a "scooped" sound, but they dial out so much of the actual soul of the guitar that all they have left is a bit of where the bass player already is – and a lot of the vocal crispness and cymbal wash range. The rhythm guitar player hammers on their instrument at all times, which means that the lead-guitar's solo lines disappear. Further, the singer has to be incredibly strong, because nobody pulls back when it's time for the lyrics to come across.

The above examples aren't the only configurations that work or don't work, but you get the picture.

The not-so-great example can be fixed in the mix at times, but it requires the audio tech to get the PA loud enough to obscure the offending sounds with corrected ones. In the end, it's much better for the band to already "sound like a band" without desperately needing the PA. In cases like that, the audio tech and sound system actually have room to work, whereas fixing the most obnoxious parts of a band that sounds like a roaring storm of test-noise may make it impossible for the mix-person and audio rig to do much of anything else.

So, really, it's much better if people are mindful of what zone they need to be in at any given time.

An Irreverent Look at Why "PA In A Box" Systems Don't Get Along With Rock Bands

I'm not opposed to inexpensive equipment. I have some inexpensive equipment that's done exactly what I've needed it to do, and has stayed in working order for years. However, inexpensive PA gear often struggles to keep up with the rock bands who buy that gear. This can be mystifying, because that inexpensive "PA In A Box" may be rated at, say, four times the wattage of the individual guitar amps in the band. It seems like a good investment, but what's missing is an understanding of the whole picture.

First, the guitar players' amplifiers only have to do one thing. Guitar. It doesn't have to necessarily sound clean (in fact, that's usually exactly the opposite of what one wants, right?), and feedback (unless it's really nasty) is actually a welcome addition. Along with the frequency range of the sound being produced and less pronounced dynamics than some other instruments, this adds up to some pretty large objective (and subjective) SPL.

Let's say that a small PA has been acquired to only do one thing, namely vocals. Essentially, it's a "singer amp" instead of a "guitar amp." Let's also say that the power amp section is rated to produce a continuous 200 watts with reasonably imperceptible distortion when mated to the supplied speakers.

Basically, you're in serious trouble.

See, you've got that 50 watt guitar amp over on stage left. You've also got another similar guitar amp on stage right. Just one of them can produce 115 dB SPL at full tilt, and two of them together can make 118 dB SPL or more. Both are potentially being wielded by players who have bought into the idea that the audience is there to appreciate the nuances of how their amps and cabs react at high volume. ("Dude, if I turn it down too far, it starts to sound like solid-state and not Led Zeppelin!") The audience doesn't really care, of course, especially since the acoustics of the venue are smearing whatever nuance may actually be present into a giant wall of midrange blast. It's *epic*, dude. Legendary, even.

Now, to compete with the raging wall of guitars, the drummer starts smacking the kit. He's not amplified, but the acoustic output is pretty gi-freaking-inormous. He's even a little louder than the guitar amps, which is disconcerting, because anything you could hear in the guitar amp area is periodically annihilated by his hammering on the snare (the one in the catalog with the description: "Guaranteed to slice through the densest of mixes! Louder than the forces of evil after lunch has been canceled!"). He's also merciless with the cymbals mounted on the kit, of which he has seven or more.

Sounds loud, right? Ah, but we have one instrumentalist left. The bass player.

To keep up with the guitar players, the bass rig (probably several hundred watts) is dialed into the "space shuttle launch" range. Both onstage and in the house, the bass is not just audible - its downright tactile. Everything not bolted down is rattling across the floor. He's so loud that the kick drum is completely inaudible. There are birds experiencing mild incontinence several thousand feet above the venue.

So...back to the PA

Here's the problem. You've got a hurricane of SPL going on from just the backline. If your eardrums weren't rattling from overexcursion, you'd actually be able to hear your spleen pulsing with the rhythm. It's *loud*.

Now, that 200 watt PA might actually produce some respectable level. On paper, it might be able to get to 113 or 118 dB SPLC. Here's the thing, though. Unlike the guitars, feedback is both much more possible, and pretty much unacceptable. Similarly, audible distortion is to be filed in the "bad things" category. In reality, that little PA just can *not* be run at full-tilt-boogie during every millisecond of the show. (No PA can, actually – well, not within reason, anyway.)

There's also the little problem of the fact that pretty much the entire audible spectrum, is occupied at all times. Especially the vocal "meat" range, which is located smack dab in the region of the earth shattering guitars (and the occasional snare hit), and the "intelligibility" range, which lives precisely in the area of the brain shattering cymbals and guitar "presence" zone.

Now, we're going to add that the singer simply cannot vocalize like a guitar. That is, the poor vocalist simply cannot sing at the top of his lungs all the time. Plus, the singer may have bad mic technique, and even be a bit quiet.

The bottom line?

The vocals are not audible, but the PA is rated as having more output than a guitar amp.

The PA is unlikely to have flexible enough EQ to really get its gain before feedback maximized, and it can't be run into the red zone because it sounds like one of those "Death Metal" pedals. The vocals have nowhere to "sit" in the mix, because all the frequency real estate has been used, and so the only option is to push them over the SPL mountain - which isn't going to happen – because the PA can only theoretically make as much SPL as just one guitar amp is making in real life.

A real question might be if there is ever a way to make a small PA adequate for a rock band.

There absolutely IS. However, it does not involve technological fixes. Here's the secret.

Make the band quieter.

If you get the band's level to match what the PA is capable of giving to you, the whole thing can be pulled off without spending thousands on an enormous PA. Otherwise...

Regarding A Show That Went Badly

So, the opening band for a show brings in an upright bass.

An upright bass is potentially one of the best tests of whether or not you have your crap together. Spoiler: My test score for the night was epic fail, dude.

Yours truly decides to just slap a mic on that sucker and go, so I grab one of the mics that was originally slated to be grabbing some crowd noise for the live recording that night. (I was told that the opener was going to be using only the equipment provided by the headliner, but that wasn't actually the case.)

Now, I'm not going to discuss mic makes and models, because I'm going for general applicability here. I will say that the mic was inexpensive, and chosen for use because of its low profile, integrated stand mount, and the fact that it "basically sounds like a vocal mic" (because you can point it at something and it will sound reasonably like that something). Also, it was handy, whereas the equipment room was difficult to get to with all the gear on stage.

Maybe you know where this is going? The above is called laziness, or complacency in more sophisticated circles.

Hindsight tells me that what I needed was a mic with good "selectivity" and feedback rejection, and that I should have probably mic'ed more at the bridge area than trying to jam that stupid mic right up to the bass's soundhole. Of course, that's hindsight. What I did was exactly the opposite of that, combined with a stubborn belief that EQ was going to fix my problems.

Should I have known better? Yep.

Did I know better? Yep.

Did I still manage to become totally inflexible and idiotic in the heat of the moment? *Uh-huh*.

To cut to the chase, we could not get anywhere near the needed amount of bass in the monitors. There was an enormous amount of drum bleed making it into the mic (both directly, and I suspect, indirectly from the reflections off the bass body), and there was one very specific angle where the interactions of the mic, bass, and monitor set off some of the most sudden, earsplitting, and out of control high-frequency feedback I've encountered before or since.

The feedback happened three our four times. The vocalist had her hands over her ears and a very distraught look on her face. Tip: As an engineer, if you ever create a situation where anyone puts their hands over their ears, you have failed. *Failed*. Insert another quarter. (That's a reference to arcade games, by the way. In ancient times, you put metal coins into a box to play a video game contained in that box.)

The band and I eventually give up trying to get sufficient monitor level, and hunkered down for the set. With the first feedback problem fixed, my new problem was the low-frequency feedback I was getting (constantly) while trying to just make the upright audible in the house. The band's set was basically ruined. They were really nice about it, though.

So, let me tell you what was cemented in my mind after that experience:

1) Engineering occurs about 90% on stage and 10% at the console. If you get your setup wrong, you are screwed. *Screwed*. All the knob twisting in the world *can not save you.*

2) Corollary to #1: If you get behind the console and the EQ can't fix your problem quickly and easily, the EQ will never fix your problem at all. If it's taking you more than a few seconds to "dial out a problem," the problem needs to be fixed at the point where it is. The problem is on stage.

3) Use the correct mic. NOT the convenient mic. The two minutes you think you've just saved may cost you an entire set in the end. Effort during the gig is saved by preparation before the gig; never by cutting corners. Not ever.

4) Effort can neither be created or destroyed, it can only be transferred from one part of the execution of a production to another. If you haven't spent most of your effort in an aspect of preparation (like mic choice and placement), then you should watch out.

5) The instant you get complacent about "knowing the room and the rig," you are moments from getting flattened. Depend on it.

The worst part about the whole thing is that I actually knew all this going in, but failed to act on it effectively. The results were flat out disastrous. It does not matter that the headliner's set went extremely well, because an entire band's prep was completely blown by one person. Me.

Be careful!

Measuring Doesn't Always Help

As time has gone on, it has gotten easier and cheaper to get ahold of a computer, audio measurement software, and microphones that are actually usable for measurement duties. I hold this to be a good thing, because I'm all for people being able to explore and do greater things than they could at previous times. The whole situation is not without its problems, however. Not knowing the limits of a measurement system, or trusting in it blindly (or perhaps “deafly”) can cause problems.

For instance, some measurement systems are entirely blind to phase. They take one measurement without respect to anything else, analyze the measurement to get a frequency response, and that's all you get. Even with a system that does honest-to-goodness transfer functions (where a measurement is compared to a reference), people may not know how to even generate a valid phase trace...and if they could, they might not know how to read it anyway. What can happen in either case is that an audio tech starts applying gobs of positive or negative gain to different EQ bands, desperately trying to fix frequency response peaks and dips that can't be fixed with EQ anyway. The root problem is some kind of acoustic summation or cancellation, and while you can alleviate the problem with corrective EQ, you can't actually solve it entirely. Not knowing this, though, people struggle valiantly to make that response curve look like they think it should.

What's even worse is when people completely ignore what their ears are telling them. I've been guilty of this myself. Instead of using the measurement system to get us quantitatively “lined up” with the qualitative experience of our ears, we get engrossed in the pure numbers and forget to listen. The measurement system gets treated as a cure all, but in the end the system either sounds good or it doesn't - and no amount of making something look good on a display matters if the end result sounds bad or "wrong."

My case in point for this would be a band I used to work for. They used one of those AutoEQ systems to get their system "flat." I have no idea of the circumstances behind how they ran the AutoEQ, but I'll be danged if the result was not one of the thinnest, most nasal, over hyped high end messes I've ever encountered. They were constantly cooking high-frequency drivers because of the enormous positive gain the unit was applying in that range.

Pressing bypass on the main EQ was an enormous improvement, and I think my only baseline tweak was to put a 3 dB cut at 400 Hz into the rig by way of a parametric.

The apparent result was that, when they ran their own sound (on days I was unavailable), they no longer had as many feedback problems, and felt a lot more confident about running their own mix in general. Their only real mistake was that they hadn't trusted their own hearing to tell them that something was very "not right" about what the very sophisticated measurement and correction system had done to the PA. The problem was that the one mistake had big consequences.

An Opinion About Digital Mixing Consoles

For me, the small digital mixing console is no less than miraculous. So much so that going back to the "traditional" way makes me desire my own workflow that much more.

Case in point: I was mixing the farewell show for a favorite band of mine, and I was using a house system with a...something or other. It was analog. My normal console was digital. Just to start, the house rig had been through the process of being rewired and retweaked all day long, and the console was anything but zeroed. Having to get everything back to a zero state by hand is time consuming and frustrating when, on my own rig, I could have reset the mic pres, scrolled down to my "Zero Split Monitor" preset, and hit recall. Done. Five seconds. Ready to go.

Further, while contending with an unzeroed console, I had none of my usual goodies. With the analog desk at this venue, I was stuck with an EQ comprising four bands of what the manufacturer thought would fit most needs...and it was okay. However, just tweaking one channel of vocals was making me miss *my* EQ's, which were configured exactly the way that I wanted them (*all* fully parametric filters, by default), with the Q's that I wanted, and sweepability from 20 Hz – 20 kHz on all bands. For whatever reason, the manufacturer configuration of the EQ sections I had on hand just didn't feel like it could quite do enough to get exactly what I wanted – and it bugged the crap out of me. Also, on the digital, I would have had dynamics processors on every channel, right at my fingertips. No patching, no cabling, and I could have chosen where I wanted the compressor to be in the signal path. The analog console had none of that.

Of course, this is only a single case on a single "other" console, but I think it illustrates my point – the point being functionality and flexibility. Years ago, back in school, when I saw a pair of digital consoles running a whole studio, I knew it was the way I wanted to go. Similarly to the functionality of the SSL G series (a large-frame analog console) in the studio next door, the digitals were packing flexible EQ and dynamics processing right into the frame of the console - and that immediately set my mouth to watering. Obviously, the sonics of the digitals weren't necessarily the same as the SSL. Still, who wants to lug an SSL to a club gig?

I can definitely understand why some people don't like the digital workflow, even when steps have been taken to get more knobs on the control surface, but I like what digital does for me. Digital consoles aren't always the fastest for all tasks, but the amount of control you can have seems to “pay the difference” for me.

The Problem Of “Bleed”

Some folks mistakenly believe that the problem of unwanted sounds leaking into a mic can be fixed by using a microphone with lower sensitivity. The assumption is that the mic will somehow filter out sounds that fall below a certain SPL. This is simply not true.

Mics with high sensitivity just put out more electricity for a given sound pressure. Mics can't distinguish between sound pressure that you want, and sound pressure you don't want. A creature with a brain is required for that particular feat. (Mics can be more or less "selective" by having tighter or wider polar patterns, but if your unwanted sound is within the polar pattern, the mic itself is not going to reject it.) If a sound source on stage is so loud that it overpowers a different, desired source at the mic, then the undesired source is too loud.

In small venues, bleed can be a very big problem with vocal mics. Especially when running at high gain, the drums and other backline can end up being greatly amplified by your vocal transducers. The undesired sound is arriving at the vocal mics within the sensitive part of the mic pattern, no two ways about it. The proportionality of the desired source (say, the vocalist) to the undesired sources is what you're going to get through the PA. If the sound pressure produced by the vocalist doesn't compete greatly with the sound of the backline arriving at the mic element, then cranking the mic up isn't going to get you the big separation you might expect.

The above is also why gates aren't a panacea for bleed problems. A gate *can* help with bleed, but only until it opens. If its open (to allow the singer to be heard), anything bleeding into that mic is coming through (although the singer is, hopefully, loud enough to help mask the bleed). Again, a gate isn't intelligent. Even if you can do all kinds of magic with the gate key input, filtering out as much spurious "triggering" signal as possible, the gate is still going to pass everything arriving at its inputs when it does open - whether it's "signal" or "bleed" is not something the gate can know. In a similar vein, gates can only help with feedback if they are closing off a signal before feedback can build up to the threshold. If the feedback is in the range allowed by any filtering on the key input, and is high enough, the gate is just going to stay open while the signal rings.

There are ways to get extremely high levels with incredible clarity and no feedback. Most of these ways center around excellent onstage sources, high performance FOH loudspeakers, monitor loudspeakers, and excellent mics. These ways (especially the technological ones) tend to be very expensive.

It's a lot cheaper, and easier on your ears to fix bleed and feedback problems by turning things down until the problems go away. Fighting battles by turning things up is hard. Fighting them by turning things down is easy – assuming, of course, that getting all the people involved on board with making less noise is easy. It often isn't, unfortunately.

One Size Doesn't Fit All

My Dad is the product of a rural upbringing. Even though he “moved to the city,” he still has that small-town tendency of expressing wisdom in pithy sayings. Here's one of this favorites:

> *“The Law of Unintended Consequences is that there are always unintended consequences.”*

When my church found a new space to worship in, we ran into “The Law” in short order.

The old building we rented had its audio needs handled by a long-term-but-still-temporary install that greatly prioritized music. The FOH (Front Of House) rig was arranged in “small rock and roll show” fashion. There were two stacks up against the proscenium wall, and that was it. No delay speakers, no dedicated system for spoken-word, nothing. In a room that small, though, it all worked just fine. Rather suddenly, though, an opportunity came along for us to utilize a much better space. The new lease puts us into a flat-out gorgeous chapel. A chapel with an honest-to-goodness permanent audio install.

Right here is where the unintended consequences come in. In this case, the unintended consequence was that reducing difficulty for one set of users may increase difficulty for others.

The company handling the chapel install probably had only one set of users to work with at the time. Those users had a set of clearly defined needs, and those users were mostly non-technical. That being the case, some parts of the system interface were hidden from the user, and some parts of the system interface were highly simplified. The main mixer was four input channels and a master – no channel EQ, no routing.

The above was perfect for the normal chapel occupants, but it just didn't work for us. I needed much more mixing flexibility for a start. I had to bring in monitor wedges for the band anyway, so adding a mixing console and a snake wasn't a huge problem – it was more work, though.

The bigger issue was how to interface my console with the install. If I had just run a line to the accessible XLR input, I could have reduced the gain on the installed mixer to allow for line-level. However, if I did that, I knew that one day I would forget to reset that gain control. if the next user of the chapel doesn't know what to look for, there was a very good chance that someone would receive an unpleasant phone call.

Ultimately, I left the install mixer alone, and stuck a spare direct box in between my console and the system input. Inelegant in some ways, but it did work.

Beware The Little Critter With Sharp Teeth

Just a few paragraphs ago, I told you that my dad grew up in a rural setting. A number of his "wise sayings" can originally be attributed to my grandpa – and grandpa had quite a library. I don't remember this being part of the official collection, but I can certainly imagine grandpa saying it:

"Even a little critter can have teeth as sharp as anything."

Ultimately, what this means is that some little issue can end up having an enormous effect later on. You might even say that it relates handily to "The Law of Unintended Consequences." In the end, you can restate this nugget of know-how regarding critters like this:

There are details even beneath the details that we're all used to noticing, and decisions that affect those details may have surprising outcomes.

Here are a few examples from my own experience:

Critter 1: The North-American, Black Bodied DI Box

At first glance, these DI boxes seemed like a perfect purchase. They were inexpensive, had some nice extras, and most things plugged into them sounded about like what you would want them to sound like.

The issue?

The PAD enable switches were on push buttons, and these buttons were near the input jacks. If the DI was placed on stage such that the jacks were easily accessible to performers, the PAD switches were also easily accessible to an errant toe. On more than one occasion during the unbridled hilarity of a rock show (and also during the much more bridled hilarity of a church service), the aforementioned toe and PAD switch came in contact. Source drops 20 dB, source is now inaudible against the rest of the mix, performer is concerned, performer doesn't know to check the switches, I can't get to the stage quickly (or at all)...frustration ensues.

I have now learned to purchase DI boxes with recessed, sliding switches.

Critter 2: The Common, Switched Microphone

Switched microphones may seem amateurish, but they are useful in some contexts. Especially when a tech tends to work alone, having that switch can allow for “performer POV” checking of high-gain monitors without fear that the system will “run away” on you. My favorite switched mics have all been durable, inexpensive, and have competed well enough with industry standard microphones to be very usable.

The issue?

Rather like the DI boxes I talked about before, the switches were far too easy to toggle. They were, in fact, a little too “smoothly” engineered. So, once again, the performer brushes the switch, performer's voice disappears, I have no input at console, performer can't feel that they've moved the switch, performer starts jiggling the cable, I try to get to the stage as quickly as possible...frustration ensues.

The beauty of this situation was that there was a fix. A bit of minor surgery allowed me to remove the “finger pad” from the switch. This made the switch recessed, and caused the switch to require slightly more force to move. You have to very intentionally jam your finger down into where the switch resides to make contact, and you have to actually want to disengage the switch to make it do so.

I have now learned to purchase switched mics that either “lock” the switches, or allow for easily making the switch difficult to accidentally come in contact with.

Critter 3: The Desert-Dwelling, 24 Channel Snake

Years ago, I purchased a snake that had more channels than I needed, and had the added nicety of locking features on the female XLR jacks. This was welcome, in that it made it difficult for a mic line to get yanked out of the snake head accidentally.

The issue?

The lock release levers weren't really engineered very well. With repeated use, they tended to get loosened inside the XLR connector, to the point where they would actually lose mechanical contact with the locking mechanism. This was a real problem when you had to get the cable disconnected again. Luckily, I was in a position to be able to take the stagebox apart when necessary – but that got old after two or three go-arounds.

This was another situation where a bit of surgery fixed the problem. I got so fed-up one day that I opened the stagebox, removed and discarded all the lock release levers, wadded up bits of masking tape, and jammed all the locks into the “open” position. The female XLR jacks now work perfectly with just friction to hold the cables connected to them.

I have now learned to purchase snakes that only use non-locking female XLR jacks, or to look specifically for high-quality lock release mechanisms.

Each example that I've presented is an easily overlooked, “little critter” sort of issue. They're things that can be hard to screen for if you're not prepared to look for them. However, each issue ended up having a pretty significant “bite” later on.

You Can't Fix Everything With Electronics, Part 1

At some point, every live-sound practitioner is likely to end up in a situation that can be described in terms similar to this:

> *"We can bring in all the fancy boxes and toys we want, but I'm tellin' ya, this rig is still going to sound like World War II being waged inside a giant box made of tin and concrete."*

A more refined and generalized statement of the above is that the proper deployment and operation of electronic devices can mitigate, but not eliminate undesirable sonic characteristics arising from the acoustical properties of a given space.

So - why talk about this?

In my opinion, there is a strong temptation present in tool-centric disciplines (mechanical work, branches of engineering, audio, etc.) to focus very intently on solutions brought about by the creation and use of more sophisticated and/ or refined versions of the tools. Perfectly intelligent, rational, and professional people can become somewhat (or completely) deaf to the idea of fixing a problem's root cause, and become overly enamored with what their discipline's tools can do to "patch" the problem. I believe that this can happen more easily when the tools are particularly impressive (and inspiring of passionate responses) in and of themselves. I also believe that achievements by outstanding practitioners of a discipline can help to cause this deafness.

Now, please note that I think admiring good engineering and good products for what they are is entirely appropriate! Gear is awesome. I could shop for it all day long. I could also spend all day admiring the accomplishments of Thomas Danley, David Gunness, James Lansing...well, you get the idea. My point, though, is that adoration of audio processing and reproduction equipment can lead to seeing those tools as immediate, accessible methods of fixing problems that might be better handled by other means. (If I really love hammers, and it's easy to reach one, I might just start banging away at that danged screw that's sticking up over there...)

Here's an example from my House of Worship experience:

My church's old worship space ("The Basement") had some rather problematic acoustical qualities. The RT60 (reverberation time) of the space was plenty long for the size of the room, being a bit over 1.5 seconds, and that reverberation was dominated by frequencies between 50 and 200 Hz. Every transient in that room was smeared like warm peanut butter. I called the room “the midrange cannon” or “the horn-loaded stage,” whereas a local music promoter came up with the analogy that “this whole room is a kick drum, and the entrance door is the hole.”

In that room, a whole ton of equalization helped - but when push came to shove, solving an intelligibility problem required a big sound pressure level differential between the “masking” sound and the sound you wanted to be clear and understandable. The building didn’t belong to us, so that made extensive, permanent acoustical fixes out of reach. That being the case, along with me being a “lover of the tools,” I pined for more and better PA loudspeakers. (“More power!” as Tim the Toolman would say.) I would EQ and re-EQ...and you know? We got along pretty well.

Here’s the thing, though.

My tendency to focus on electronic tools as an immediate and accessible way to “patch” the room’s problems made me overlook something which would later become completely obvious to me. For a reasonable investment (even to a very small congregation), we could have installed non-permanent acoustic treatment. Not nearly enough to make the space perfect, but certainly enough to put a major dent in the problem. Enough, probably, to make a much more expensive PA upgrade unnecessary...but I never proposed such a solution. To use a farming analogy, I was so focused on my horse and plow that I completely overlooked how a couple of days of getting rocks out of the field could have helped.

You Can't Fix Everything With Electronics, Part 2

It can certainly be faster and easier to use an equipment-oriented solution to an acoustical problem, but that "mitigation of undesirable issues through electronics" is only that. Mitigation. The root cause of an issue isn't being addressed - rather, the symptoms are being soothed.

It's all fine and good to say, but there's a deeper question involved. Why would the last few sentences be true? Ultimately, the undoing of our tools is that their adjustability to acoustical conditions will fall short in at least one of two areas: Frequency content with respect to time, and frequency content with respect to an observer's location. Whatever other factors are in play, time invariance and spatial invariance are the key issues surrounding the limitations of electrical and electro-acoustic tools. In a human, this inability to vary behavior is seen as being not very smart, but in our gear, it's pretty much the normal state of affairs.

Let's first consider a very common electro-acoustical device, namely, the loudspeaker. Let's also consider a common acoustical problem, that of a room that is highly reverberant. What's quite obvious is that the speaker can not de-reverberate the room – in fact, as an acoustical source it causes reverberation to occur. What may not be so immediately obvious is that the loudspeaker's output is spatially and temporally invariant.

"Wait a second!" I can hear you exclaim. "The speaker's output is most definitely variant with respect to location! It has different measured output if you're at a different angle, and its SPL decays as you get farther from the box."

What you're saying is true, but I would respectfully put it to you that what you're describing is different observed acoustical responses based on different frames of reference. If you move your frame of reference, the characteristic coverage pattern of the loudspeaker has not changed. If it's nominal coverage pattern is 90°, it remains 90° whether you're standing in that coverage or not. The coverage of the loudspeaker is spatially and temporally invariant, in that the characteristic directivity of the loudspeaker for any given input signal remains fixed for all observation locations and path length delays. Further, the acoustical effects of the loudspeaker itself (including resonances, distortion products, and other issues) do not change based on observation location and time, although certain observation locations and times may reduce your ability to detect them.

Okay, that’s all very “sciencey,” but why does it mean that the loudspeaker can’t fix our reverberant room? The first issue is that spatial invariance. With a loudspeaker we have to choose one given coverage pattern and live with it. A wide coverage pattern (especially with a bit of distance between the loudspeaker and the observers) lets us hit all of our listeners with a small number of loudspeakers, but it also hits a lot of the reflective surfaces that cause reverberation. Worse, those people sitting in the back are probably quite a ways beyond the critical distance (which is where the SPL of the sound coming directly from the source and the SPL of the reverberation are equal). Here's a visual example:

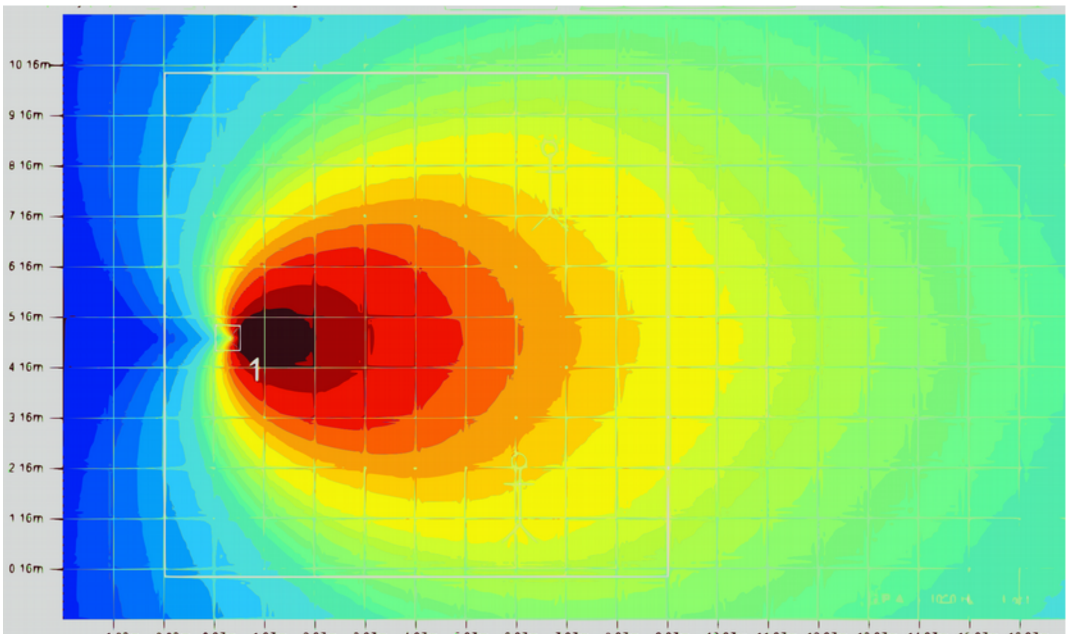

Since we can’t vary the speaker’s coverage based on listener location, and we’re unable to resort to any practical trickery that would cause the speaker to produce a signal that reduces or eliminates the space’s reverb time, we’re basically stuck with having to choose a spatially and temporally static coverage pattern, speaker placement, and overall output that reduces the effects of the reverberation on our listeners. We can probably get pretty good results by using higher directivity speakers producing smaller outputs at shorter distances - but the number of loudspeakers needed goes up as those factors (coverage, output, and distance) are reduced. You also start to experience more pronounced phase effects as more sources reproducing the same material interact.

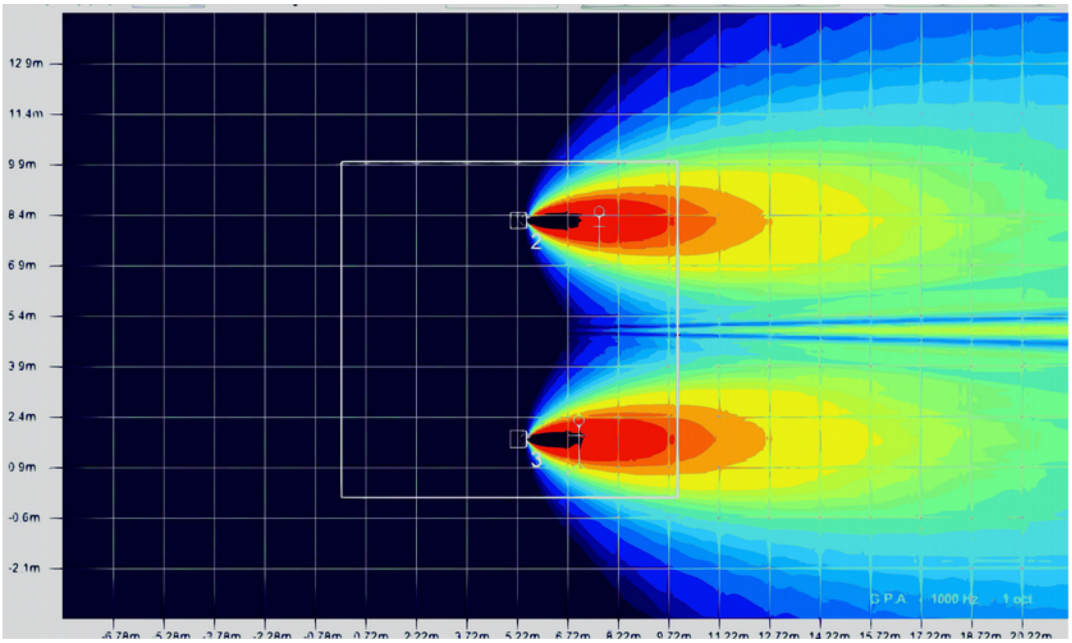

Now, if we really want to go "whole hog," we can just give everybody headphones. Those would effectively eliminate the room from the equation. Of course, we'd have to buy a lot of headphones.

"Ah ha!" you say. "You haven't talked about EQ. We'll use EQ to make things sound nice!"

Well, yes, you can do a lot with EQ. EQ is my favorite audio tool by far. However, I would again put it to you (respectfully) that EQ can't fully fix the issues you're facing. If your acoustical problem is really nasty, even a whole mountain of EQ can't stop the issue from ruining your day - it may even bring attention to the problem!

Why?

First off, EQ (even more than loudspeakers) is THE poster child for spatial invariance. In fact, the highly talented, well respected, greatly experienced, and just flat-out cool Bob McCarthy has this quote attributed to him:

> *"Any electrical filter on an acoustical device is a spatially invariant solution to a spatially variant problem. As grownups we have to decide what part of the space to tune the solution to and let the others go."*

In other words, any "tuning" that you do with an equalizer can only be 100% "correct" at one point in a room. Don't get me wrong! The tuning might be 99% correct at every other possible listening position, but you still have to accept that the filtering you’ve applied simply can’t change based on a listener’s physical positioning. True, you can always add more discrete channels of reproduction, each with discrete filters and discrete loudspeakers so as to get solutions that are more correct at more locations - but there are always practical limits to how much gear you can throw at a problem.

Here’s a graphical depiction of what I’m talking about. The three different traces show the measured response of a loudspeaker in a room, under three different conditions. The red trace is with no EQ, and a relatively close observation point. (I’m pretty sure that big notch at around 1k was being caused by a number of electrical and acoustical factors, because even a huge amount of corrective EQ couldn’t fully fix it. Hey – there’s an object example, eh?) The green trace is with corrective EQ for the original observation point, and the gold trace is that same EQ being applied with an observation point much farther from the loudspeaker.

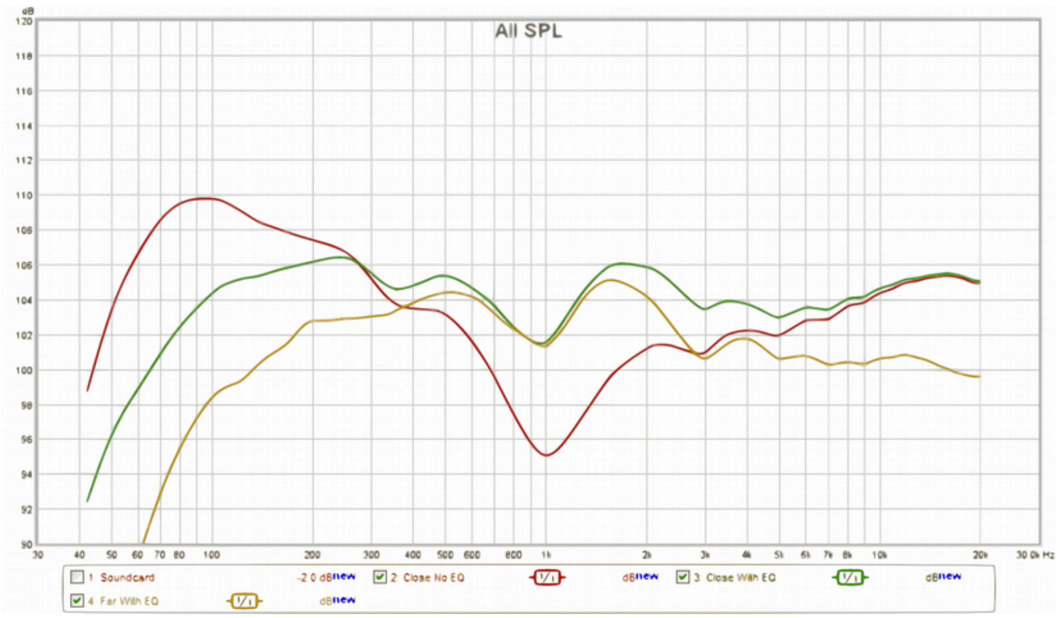

The problem should become apparent fairly quickly. If we use EQ to get the frequency response “just so” for one observation position, it doesn’t guarantee that any other observation position will get the same results. If we extrapolated the situation above to a more severe situation, we might have to be very mindful that getting the frequency response to be really balanced for the folks in the back might submerge the closer listeners in a sea of low frequency material, while slicing their heads off with high frequency content. Again – if we have enough money, time, and space in the room we could have different amplifiers and speakers with discrete EQ for folks in different areas.

The final major point to make in all of this is that EQ is also time invariant. For simple filters, this is easy to understand. A filter with a given frequency, bandwidth, and gain is set and left. After being configured, it is always operational with those settings, at all times. One might choose to mention that there are such things as dynamic equalizers, which alter filter characteristics in real time, based on inputs or feedback methods. However, I think a strong argument can be made that even though a dynamic EQ changes a filter over time, it does not actually change a filter with respect to time experienced by an outside observer. The filter is changed with respect to some detected signal level, and this probably involves detection of how long a signal level is greater or lesser than a given threshold, but the change over time is incidental. The time that a signal spends at any particular point other than the threshold is essentially ignored. Further, signal being detected is usually an electrical input to the device, and not a measurement of total acoustic response in the actual room. (Please notice that I said "usually." There are some devices that can do dynamic EQ by measuring the "in room output" in some way. However, the spatial invariance problem still exists in that you can only have so many measurement devices, and you ultimately have to pick one average that you can live with.)

The next picture is a graphical depiction of why time invariance matters for acoustical problems. It shows the RT60 (reverb time) and waterfall plots for the gold-colored trace from before. What is apparent rather immediately is that all the frequencies present do not decay at the same rate. The lower frequencies tend to "hang around" a bit. The problem, then, is that the corrective EQ I'm using does not change over time. If those resonant frequencies in the room reverberation are causing me a problem (such as the masking of sounds), then I have to simply reduce their output level. I have no way of reducing their reverberation time without modifying the room. Even a dynamic EQ can't really reduce reverberation time - it can only reduce a frequency that is being judged to be getting beyond a certain level.

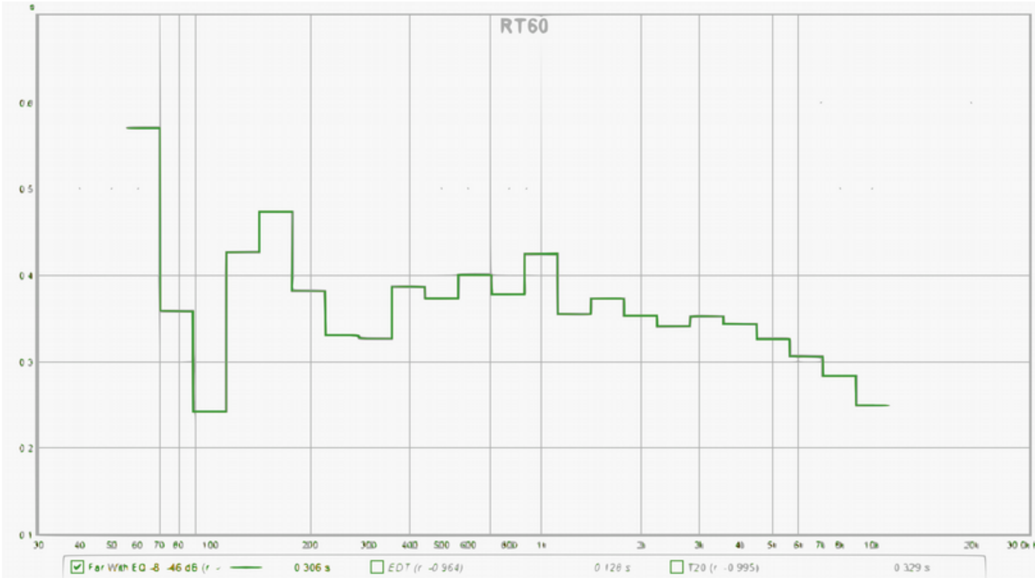

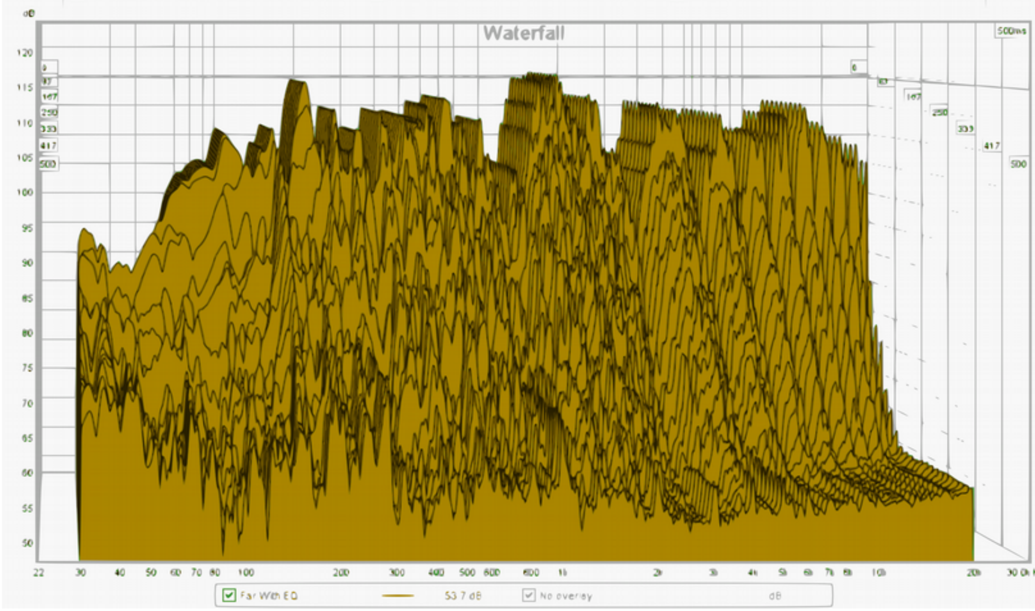

Now, here’s why this situation can actually make EQ draw attention to an acoustical problem. In a highly reverberant space, especially where, say, low frequencies dominate the reverb, you may end up in a situation where the only way to tame the problem is to severely cut the problem areas. What you can end up with is a PA system that very obviously has a huge “hole” in its low frequency response. You ultimately are forced to make sounds that have almost no low frequency content in them to avoid a low frequency reverberation problem. The reverberation can’t fill in content that isn’t there, of course, so you can end up with an overall effect of “gorgeous room, but why does the sound system remind me of my old clock radio?” Again, the room reverberation doesn't add low frequency information to an acoustical event in the room. It merely reinforces and sustains the low frequency information that is already present.

Humans Are Smart, Signal Processors Are Dumb

This isn't about, say, how one EQ implementation is better or worse than another. Rather, this is about the basic, inherent limitations of audio processors in general. I'm going to use equalizers as my object example, but pretty much any audio processing tool can fit into what I'm getting at. To generalize:

Electronic devices are usually incapable of evaluating their own contribution to the sonic character of a space. The few devices that are capable of evaluating that contribution are usually subject to marked limitations in their data acquisition and interpretation - that is, as compared to a human operator.

Well, there's a mouthful. Let's step through that a bit at a time. As a quick note, when I'm talking about "total acoustic response," what I'm referring to is the output of a device as it is being heard or measured after being converted into airborne sound pressure waves in a physical location.

A basic EQ does not have any method for measuring or interpreting the signals being input, nor the results of the output. They simply act upon a signal as they receive it. If that device's output - and the resulting total acoustic response in a space - is "correct" for an observer, then that's wonderful. However, if something changes in such a way as to render the total acoustic response "incorrect" to an observer, the device is helpless without the intervention of an operator. Further, the device has no way of "knowing" whether what it is doing happens to be "correct" or "incorrect" to an arbitrary observer.

Now - what about devices that can make some interpretation about what they're doing to a signal? For instance, we could consider an equalization device that can measure the total acoustic response involving its own output in the room, and make adjustments such that the response matches a set of parameters. (This is usually called something like Auto EQ or an EQ Wizard.) Firstly, most of these devices can only measure one set of data points at a time. They will usually accept only one measurement microphone input, and may or may not be able to compensate for any imperfections in the measurement microphone itself. Second, the device has no way of knowing why a given measurement is the way it is. A large notch in a measured frequency response might be because of an acoustical phase cancellation, but the device has no way of knowing if that is the case or not. It merely "hears" a large difference in sound pressure level from one frequency range to the next. Third, useful measurements rely on predictable input signals that can be meaningfully compared to an output. It's all fine and good to have a transfer function created by comparing the signal running through the outputs of a mixing console to the total acoustic response in the room - but what happens if that total acoustic response includes a sound that is not being produced by the PA system? The measurement (and any automated response to it) is invalid, whether wildly or by a small degree.

A perfect object example of this is the AutoEQ I talked about in “Measuring Doesn't Always Help.” The AutoEQ generated a curve based on a measurement that was taken in who-knows-what room, with the measurement mic in who-knows-what position relative to the PA. The AutoEQ got a correct solution for that single, precise situation, but the result sounded poor in the situations I experienced. The automatic EQ simply had no way of knowing that what it was doing was wrong in changing situations. All it knew was that the operators wanted a “flat” curve, and it had done everything it could to give them that...all based on what was happening at the measurement mic on a single occasion of measurement.

You can see here what I mean by limited data acquisition and interpretation. The device had only one measurement to work from, and it could only interpret whether or not the resulting frequency response conformed to a target curve. The humans involved were the only “devices” in the process that could tell that the results were inappropriate in most situations.

The Worst Kind Of Problem

What is the worst kind of problem that can arise during the installation or operation of an audio system?

I certainly can recognize that a number of answers might be considered to be completely valid, but I'm going to opine that the most correct answer is this one:

The worst kind of problem that can afflict an audio system is an intermittent, show-stopping problem.

The key word up there is "intermittent." Why?

Let's first consider a relatively minor issue – say, an audible-but-not-loud (maybe 40 dB or so below the continuous level of "nominal" program material) buzz or hum stemming from a ground loop. If that issue is "steady-state," then there are two things working in favor of a person working on the system. The first favorable condition is that the system, having predictable behavior, can be meaningfully and scientifically troubleshot. The original system configuration and steady-state problem act as a sort of control, and changes made to the system during troubleshooting are the variables. Consistent system behavior means that potential causes for the ground-loop can be investigated and eliminated until the actual cause is found and remedied. The second favorable condition is that a not-overbearing steady-state issue has a real chance of eventually being "filtered" by the brain of a person who is listening to the system. It becomes like any other kind of continuous noise or distraction, in that is taken as a "given," recognized as not being part of the program material, and then ignored after a few minutes. This is rather like the hiss, crackles, and pops that people would hear while playing a vinyl record, or like the "swish" and garble that you might come across in a low quality digital encoding.

Of course, our ground-loop problem isn't made acceptable by the ability of the listener to ignore it. It's just that it's actually tolerable until it can be fixed.

Now - what if that buzz comes and goes? First, troubleshooting becomes very hard. The system now has unpredictable behavior, and so it becomes difficult to ascertain whether any given change to the system has actually had an effect on the problem or not. The system's original configuration and problem can no longer be considered "controls" in the troubleshooting experiment. Instead, the problem enters into the list of variables...and this can make it impossible to actually isolate in more severe cases. The second annoyance is that the intermittent buzz is now much harder for a listener to ignore. When it is gone for a while, and then unpredictably appears, the listener's brain is very likely to latch on to it as a disturbance to the program material. The listener's experience is jarred by a signal that, by being intermittent, indicates that something is going wrong and they need to pay attention to that something going wrong.

Obviously, a ground-loop buzz or hum isn't a fatal error. However, extrapolating this up to a "show-stopper bug" is just a matter of upping the stakes.

Consider a show being mixed on a digital console. The stage crew has just finished changing over the stage from "Bobby's Boomers" to "Camille's Crooners," so the console operator recalls the scene written for the band during sound check. The console loads the scene...and something goes horrifically wrong.

The console's control surface locks up. The audio engine stops with a bang. The operator asks the stage crew to shut off the power to the PA system's amplifiers, and reboots the console. Being a savvy sort, the operator then tries to reproduce the problem. The operator again attempts to load the scene for the new band, and again, the console throws a fit.

This is troublesome, of course, but not unpredictable. After a quick conference with the band and crew, everybody decides to just wing it. A basic mix is thrown together, and the show goes on. The operator doesn't load the offending scene, and the console happily works as it should.

Contrast this situation with one where the console also stops with a bang upon a scene recall, but then recalls the scene without a hitch upon reboot. Then, 30 minutes later, the console throws a fit during a fader move. Reboot. The console is then happy for two hours, after which it locks up when the operator selects the EQ page for channel 14.

This problem is also a show-stopper, but it's far worse than the first scenario. Being that the issue is intermittent and unpredictable, there's no way to nail it down. Could the power supply be failing? Could there be a problem between firmware version 24.18.a.12.99 and this one particular configuration? Could an option card be working loose? Could it be... Could it be...

Further, every second with the console is a nail-biter. When is it going to go down again? Will it reboot each time? How many more nasty looks will the event coordinator throw at the mix position?

In the end, a problem's severity is only one aspect of how bad the problem really is. A severe problem that occurs at random is far worse than a severe problem that occurs under predictable conditions. I can imagine Grandpa Maland right now, leaning on his fence, and complaining about a troublesome horse: "I wish that son of a gun would just go lame and stay that way. Then I might be able to figure out what's wrong with 'im."

Waxing A Wreck

A story Dad once told me was about how he had become enamored of an outwardly appealing but not-running-so-hot motorcar. (I can’t remember the make or the model.) He had it in the driveway, and was lovingly washing and shining his acquisition. Grandpa Maland walked up, and imparted his assessment of the situation:

“Son, you can’t polish a turd.”

See, Dad was dealing with cosmetics, but the weak link in the car wasn’t the body, or even the frame. It was the engine - Dad should have been working on the bit that actually made the car go somewhere reliably.

The way I apply this to the audio world is that time, money, and effort are easily apportioned to “deceptively sexy” parts of a project, when they would really be better spent on the nuts and bolts that will have the biggest and longest-lasting impact on the whole. Specifically, I’m talking about how folks can get roped into endless tweaking of things that are already strong links in a signal chain. In so doing, they rob themselves of the much greater benefits afforded by investing in the weakest links. In effect, they “wax a wreck.”

Okay...so, what are the weakest links?

My experience in audio has taught me that the weakest inherent links in a signal chain are at the extreme ends. At the very beginning and very end of a signal chain is a transducer of some sort. A transducer is a bridge between a sonic event and the electrical signaling that represents that event. It acts as a converter between one form of energy and a different, yet corresponding form of energy.

Please note that, in this case, the terminology of “weakest” is not an implication of unsatisfactory performance. What it is describing is a proportionality between event integrity of event reproduction across devices in a signal chain. Maintaining signal integrity across a transducer is significantly more difficult than across a device which does not have to convert from one form of energy to another.

Audio transducers are microphones and loudspeakers. The microphone is the entry point to the signal chain, and the loudspeaker is the exit. Now that we’ve got a definition of what the weakest links are, what’s all this to-do about being disproportionately concerned about the stronger links? Here’s an example:

Every so often, somebody wanting to improve a live sound reinforcement system will start asking questions of others to get a feel for what they should do. What's interesting is how often the asking party will fixate upon something in the middle of the signal chain. You'll get questions about whether or not switching to a significantly more expensive mic preamp will result in a more pleasing overall sound, or a worry about whether purchasing console "X" will cause the system to sound less "warm and punchy" than console "Y." There are discussions about whether or not amplifier "X" is better on subwoofers than amplifier "Y."

Now, let's be clear.

I'm not saying that different preamps, consoles, and amplifiers are incapable of producing measurable and human-detectable differences in the sound of a signal chain. I believe that they can, even though I myself am not the type to spend a lot of time worrying about them. (I'm not even really convinced about amplifiers though. I think you've either got enough power or your don't, and that's about it. I could be wrong. I often am.) What I do find very interesting is how often the above worries have been presented by people who have not yet spent a proportionate amount of time and money on improving the extreme ends of their signal chains. Doubling or quadrupling their current investment in microphone or loudspeaker quality would get them noticeable gains, yet they've bought into a line of thinking that claims the same quality gains from a double or quadruple investment in devices which perform no transduction at all.

Why would they buy into that thinking?

My best guess is that, in the audio world, we've traditionally seen the "big boys" spending a lot of time talking about things in the middle of the signal chain. Preamps, consoles, compressors, etc. In my opinion, what often goes unsaid in those conversations is that the guys at the very top of the game are already using top-notch equipment at the ends of their signal chains. Further, I think that a good number of folks look at what is "reachable" about a more highly placed situation. That is, a giant concert-style PA with effectively unlimited power reserves is totally outside the realm of possibility, but this or that processing device is at least somewhat accessible.

I should also note that we may be seeing a bit of a conversational trend reversal as "digital" becomes more and more commonplace. With so much processing going "inboard," there's more time to discuss mics and loudspeakers again. Sometimes.

It depends.

Anyway – what I’m going for with all this is a note of caution. I’m not holier than thou in any way, as I’ve certainly been guilty of putting lots of money into things sitting “mid-chain” while my endpoints were lonely and neglected. Having experienced those mistakes, though, I think I can legitimately suggest the asking of these two questions:

1) Is this addition to the system just a nifty gizmo, or does it have an impact on the quality of the whole that seems generous compared to the time and money involved?

2) Is this really a weak link, or am I just thinking or being told that?

If you're careful about the answers, you'll tend to put your money into the things that matter.

Expectation Management

I made it all the way through high school without actually using my assigned locker. I spent four years hauling every book and supply needed for each day in backpacks and handle bags, sometimes two at a time. There was one particular year where I broke the strap on a rather nice, black leather bag that I liked.

Dad's reaction was to say, “You know, that bag just wasn't designed to carry that amount of weight.”

My teenage know-it-all reply was, “But it all fits in there! They should design the bag so that it's able to carry that volume of lead, if I want.” (It's amazing how much more we know as absolute certainty when we're younger, wouldn't you agree?)

The problem here was something called “Expectation Management.” In my mind, the bag was capable of providing enough interior volume for the books I wanted to put inside it, and I felt that it was thus reasonable that it should handle the mass of those books as well. The real issue was that I was misreading the application of the bag. It was actually designed to hold clothing, which is usually far less dense than textbooks.

Problems with expectation management happen all the time with audio equipment.

A notorious expectation management problem is the one described by the myth of “underpowering.” The assertion is that Jane or Johnny User toasted their loudspeaker drivers because they drove an amplifier into clipping (or hard limiting, as is more likely in these modern times) and that they would not have burned the voice coils if they had been provided an amplifier capable of “cleanly” swinging greater voltage at its outputs.

Not entirely true. While hard clipping can cause high-frequency distortion artifacts, and those artifacts can be sent through a passive crossover to cook HF drivers, that just isn't the whole picture.

As discussed earlier in this book, every power amplifier that you're likely to come across has been rated for a certain amount of power at a certain amount of distortion. If you're willing to live with greater distortion (or greater limiting artifacts), you can drive the amp beyond that rated power. The peaks of your signal will get smashed flatter than pancakes, of course, but this does not prevent the average power of the output signal from approaching that same maximum peak.

(Incidentally, this is essentially how the “loudness war” in recorded music works. Inter-sample clipping notwithstanding, you can't master a CD with a level above 0 dBFS. You can, however, slap a limiter across the signal, and then drive the signal into the limiter to get a higher average level.)

What's killing Jane or Johnny's loudspeakers isn't that the amplifier is too small. The really deadly thing is that Jane or Johnny expects to be able to get more Sound Pressure Level (SPL) out of their loudspeakers than is actually possible. If they were willing to drive the “little” amp beyond its linear performance limits to get a desired SPL, then it is reasonable to assume that they will end up driving a more powerful amplifier to the same level of continuous output as the smaller unit. They will then likely cook their drivers just as fast, because the thermal stress applied to the loudspeakers will be just as bad. (They will get less distortion or limiting artifacts from the amp, of course.) They may even kill their drivers faster, as the peak voltage deliverable from a sufficiently powerful amplifier may allow Jane or Johnny to quickly exceed the mechanical limits of their loudspeakers.

Bottom line?

Jane or Johnny is in much the same position as I was with my bag of books. They expect that their sound system should be able to do something that it actually can't, and are damaging it as a result. Expectation management, then, is an inexpensive and (if genuinely received by Jane or Johnny) effective way of fixing their problem. If they can have their expectations of the system brought into line with what the system can actually do, then they are far less likely to push their system into the damage zone. Yes, some sort of system protection to guard against accidents and misuse is a good idea, and so too is a system redesign if the whole shebang just doesn't do what it ought to do. Still, having an appropriate mental picture of what can and cannot be accomplished with a given system is key.

Classy Watts

These days, it seems somewhat less inevitable that a live-sound person (or any sound person) will hear something to do with “Class A watts.” I'm grateful for this, because there is no such thing as “Class A watts.” There's also no such thing as a particular manufacturer's watts, or watts as referenced to a certain model of device.

The phenomenon of distinguishing watts by amplifier topology, amplifier manufacturer, or amplifier model is an incorrect description of a legitimate observation: Different amplifiers deliver different amounts of power into a given loudspeaker. Like I said, this is a legitimate observation – in fact, it's obvious. Not all amplifiers are equal. The problem is that the explanation for why the different amplifiers are unequal is simply wrong.

A watt is an SI unit. It's one joule (a unit of energy) per second. For audio types, it's much easier to talk about this in terms of electricity. Watts are the square of the voltage, which is then divided by the resistance. This definition is not changed by amplifier type, model, or manufacturer. It's a scientific fact. You either have a certain number of watts flowing in a system, or you don't.

So, what could fuel this incorrect description of varying performance?

One possible reason is a misplaced belief that all control labels mean the same thing. This is far more common with instrument amplifiers than power amplifiers, but it still bears some discussion. A guitar player might conclude that “tube watts are louder than solid-state watts” after switching two amp heads (that we'll assume are genuinely capable of the same wattage output) between the same cabinet. The guitar player sets all the controls to “5” on both amps, and the tube-powered model is observed to create a greater acoustical output. So – tube watts are more powerful, right?

Wrong.

First of all, for gain and volume controls, that setting of “5” may not be the same across both amps. Sure, all the potentiometers are halfway through their travel, but there's no guarantee that being halfway through the travel of the master volume on the tube amp corresponds to the same percentage of possible output voltage as half the travel of the master volume on the solid-state amplifier. The tube amp could have 8 dB of volume left in reserve, and the solid-state amp might have 10 dB in reserve at that point. Then, there's the whole problem of the tone controls on the amps probably having very different results with all the knobs set to “5.” The scales are arbitrary. They're not calibrated in dB, and so you don't really know if you have multiple independent variables or not.

Moving outside of the example above, what about the issue of “my Class A combo is only 30 watts, and it gets louder than his 50 watt Class AB?” Again, there's the whole issue of the controls not necessarily being to the same scale. Beyond that, though, is the question of whether the loudspeakers in both amplifiers are the same sensitivity. It's entirely possible that the Class A combo has a loudspeaker that produces 3 or more dB SPL at 30 watts than the loudspeaker mated to the Class AB combo produces at 50 watts.

What about pro-audio, though? Most professional power amplifiers don't have tone controls. They have input attenuators, and that's often the extent of what you can do at the front panel. (Unless you have one of the new, super-nifty amps with integrated DSP for crossover, EQ, bass enhancement, etc.) The input attenuators are, by and large, calibrated in dB. There's a lot less that can be different...or so you'd think.

To start, you have to be careful that you're really comparing the same wattage rating. Some folks read a number on a spec sheet and assume it's a number that somehow applies equally to all other manufacturers. They then make proclamations that one company makes louder watts than another. This isn't true, however. One company might be claiming higher output under a very specific set of conditions which don't match up to another company's conditions for that kind of output. For instance, two amps might both be capable of 1 kW into 4 Ohms, but one might only be able to sustain that 1 kW for a second or two, whereas the other amp might be able to do that all day long. Another example is if one company only publishes numbers corresponding to 0.1% THD (Total Harmonic Distortion), and the other company allows for up to 1% THD. If you're willing to tolerate more distortion, you can get more power out of a unit.

Then, there's the whole issue of voltage gain. If two amplifiers have different voltage gains and you don't take that into account, you might wrongly come to the conclusion that “this amplifier's watts are louder than that amplifier's watts.” Both amps might be rated for the same maximum output, but if you don't run both of them right up to that maximum then one might be louder than the other. Finding an amplifier with a voltage gain of 32 dB, and then finding a not-otherwise-incredibly-different amplifier with a voltage gain of 37 dB isn't difficult to do these days. If we're running a rig such that the 32 dB voltage gain amplifier is putting about 1 watt into a loudspeaker, then swapping for the 37 dB amp without changing anything else puts more than 3 watts into the same load. The amp with higher voltage gain doesn't have watts that are magically more powerful – it's just delivering more of them.

The Curious Case Of Divergent Solutions

Most of the time, live-sound technicians get to work in situations where solutions converge, or are at least convergent enough to handle easily. We get ourselves a mix, and the mix is pretty much correct throughout the entire venue area that we're concerned about. The “mix solution” at one arbitrary point is essentially the same as the mix solution at another arbitrary point – the solutions converge. There are situations out there, however, where this isn't the case. When these conundrums arise, it's helpful to know how they've come to be, because this knowledge can help greatly in either fixing or choosing to tolerate an issue. (There are venues out there where the audience has to travel behind the FOH PA to, say, get to the restroom. There's obviously a divergent solution in play there, but it doesn't really get considered because it simply gets accepted as part of the room.)

In the world of live-sound, divergent solutions come about when a spatially variant acoustical condition invalidates a sound reinforcement decision. The variant acoustical condition may be purely acoustical, or it may be electro-acoustic in nature. Nevertheless, the bottom line is that your mix becomes inappropriate at some physical point where you would like it to remain appropriate.

A very simple case of divergent solutions is that of critical distance. Critical distance is a point or set of points away from a sound source where the reverberant sound is equal in intensity to the direct sound. At one distance your mix may sound fine, but go too far and the sonic equivalent of gaussian blur has made it impossible to figure out where one thing starts and the other ends. You can sometimes fix this by using more PA (like delay stacks) to get more direct sound into the audience, but any addition of volume from the PA will also increase the total volume of the reverberant field. Too much SPL from the delays or “spot coverage” PA elements may cause the reverberant field to overwhelm the main PA even more quickly.

As an aside, the rule of thumb that I currently subscribe to is that you can alleviate critical distance problems with better PA coverage only if the total SPL from all PA and other acoustical elements remains fixed. If you start with, say, 100 dB SPLZ (measured at some point in the room), and then add a whole bunch of delays and spot coverage to fix a lack of intelligibility, then you have to reduce the levels of all the PA elements and stage volume until you get back to 100 dB SPLZ. Otherwise, you've probably just created more problems for yourself – and even if you do keep the volume down, you still might have problems if you didn't aim your extra coverage carefully.

A divergent solution strongly related to the above is the problem of audience members just plain being out of coverage. This problem can manifest if you don't have enough PA, or you deploy an adequate PA incorrectly. Some audience members end up being so far “off axis” from the sound reinforcement system that the mix becomes unintelligible. The PA effectively stops competing adequately with the stage volume, or becomes overwhelmed by the sound of the reverberant field, or both. If you're sitting in the effective throw of the PA, the mix is fine, but if you're not, the mix is terrible. The solutions have diverged, and you either need more PA to point at different locations or you need to spread out the coverage of the PA you have.

Trickier territory is entered when a PA system is overdeployed. This kind of divergent solution occurs when a correct mix solution at one point is a specific blend of the PA and stage volume, and where the PA completely overwhelms the stage volume at other points. This can happen with delay speakers that are improperly volume shaded with respect to the main PA, or if a PA is deployed on both sides of a stage for a band that is only on one side of the stage. In both cases, the apparent location and intensity of sound sources is wildly different at varying points in the audience. In the case of unshaded delays, whatever it is that the PA is reproducing does not exhibit similar SPL decay to the stage volume radiating into the audience area. As a result, whatever is being reproduced electro-acoustically washes out the stage volume, and invalidates the mix.

Similarly to the previous paragraph, if the band is off to one side of the stage with a PA that is deployed on both sides, a mix that blends with the stage volume becomes progressively invalid at various distant points from the mix position. A mix position that mostly hears the side of the stage where the band isn't will tend to produce mixes that have insufficient reinforcement when one is near the side of the stage where the band actually is. On the other hand, mixing where you can mostly hear the band may lead you to make a mix that is too overwhelming for people near the PA on the other side. You may be able to get a decent mix if you're far enough from the stage for there to be only a minimal difference between the blends derived from the stage volume and one side or other of the PA. However, you have to be careful that the mix at points closer to the stage doesn't diverge too far from the chosen solution.

Engineyes

I once was witness to an Internet forum thread where one participant had loaded a digital audio file into an editor and looked at the waveform. The participant in question was amazed at the sound of the file, because it sounded very pleasing...while it had the same *look* as many of the overcompressed (in terms of dynamic range) tracks being heard at the time.

Unfortunately, this visual analysis and observation rubbed another participant the wrong way. The other participant, who had worked on the audio file being looked at, is a major industry player. He decried that the “looker” had not been restricted to making a judgment using only sonic information. The “looker” and all like him were given the derisive nickname “Engineyes,” implying that a recording or live-sound engineer uses only their eers. Ears. Sorry.

I have to say that I forcefully disagree with this nickname, and also the dislike of using tools other than one's own hearing system to analyze audio. In my opinion, the job of a professional is to bring *every* helpful tool and technique to bear in completing a project. Yes, one should use one's ears, but if you have a real time analyzer to help calibrate your ears, you should use that too. If you have meters to help you make sense of what your processors are doing, you should use those in conjunction with what your ears are telling you. If your sense of smell can somehow help you make a show sound better, then for goodness sake, follow your nose!

The notion that we should hobble ourselves – that we should not make use of *every* tool that is even somewhat useful – in order to conform to some purist and dubiously practical ideal of how audio ought to be done, is (to me) unhelpful at best. If using something other than my ears will make the artists on stage happier and the audience more satisfied, why wouldn't I use it?

There's nothing wrong with being an “engineyes” if it makes your shows better.

The Joy Of Analysis

Analysis is, obviously, a great way to get a handle on what's going on with the sound of an audio rig – indoors or out. However, I would also make the point that analysis is a great way to learn about your gear. Sure, your gear has metering and other displays built in, but is the displayed information accurate? If I make this change to the controls, what really happens? Sure, your ears tell you a lot, but they don't relay that information in a quantitative way. Getting out a measuring rig and trying some things can be a fun (and very informative) experiment.

For measuring audio, I do recommend getting your hands on a system that can do dual-FFT/ transfer function measurements. There are some free offerings out there, like Visual Analyzer. (Visual Analyzer is pretty decent, but for some reason it doesn't yet have easy delay compensation for reference signals. This can make it hard to get a meaningful phase trace, and it can also make the magnitude trace more unstable.) Systems capable of performing analysis via transfer functions are helpful because they can compare a known reference signal to a measured signal – the measured signal being the reference plus processing and acoustics.

If I set up a test such that the reference signal passes through the same basic input chain as the measured signal, a dual-FFT system will much more effectively ignore phase and magnitude problems induced by the measurement rig's input system itself. Also, a properly smoothed transfer function trace is much easier to read for small details than just a single channel analysis. The reason for this is that, especially with random-noise measurement signals, a single-channel trace is quite "rough." The roughness comes from the system simply displaying the absolute magnitude of each frequency at each point in time. For noise, each individual frequency's magnitude can vary a fair amount from millisecond to millisecond. With a transfer function trace, the visualization is much more stable, because you're looking at the difference between the reference and the measurement. Both signals can have individual frequencies jumping around in level from millisecond to millisecond, but as long as they jump around in the same way, the trace stays flat.

What I mean by the above is that if the reference signal isn't delay compensated to the measurement, the overall magnitude trace will destabilize to some degree. At some time point, a frequency will have different magnitude when compared to the reference. The effect of this delay can be damped with smoothing features, but it's really best done by figuring out exactly what the delay on the measurement input is, and then applying that delay to the reference. Some systems can do this in a highly automated (or, at least, painless) way, and those systems will usually cost you a good amount of money. Of course, delay compensation is also important to get a valid phase trace (as I alluded to earlier), because phase is all about time differences. You can't say much that's meaningful about the phase of the processing applied to a signal if phase shift from traveling a distance to your measurement mic is already slapped on to your measurement.

But, enough about all that. Let's have some examples.

I have some plugins that I'm interested in using for processing live-audio. (As of this writing, the mixing console that I use for live shows is a modular, digital system that allows me to load certain VST plug-ins for processing.) One of these plugins is a very steep filter. Now, I'm already aware that steep filters tend to "ring" or "overshoot" around their cutoff frequency, but I want to see just how much my the filter actually does this.

In this case, I don't have to worry about phase. The latency of the filter processing (which is essentially negligible for this plugin anyway) is compensated for by the VST host I'll be using. Further, I'll be measuring directly off the internal digital signal from the VST host. Running a 20 kHz tone through the system shows time alignment that's perfect (within the error that the system can actually display).

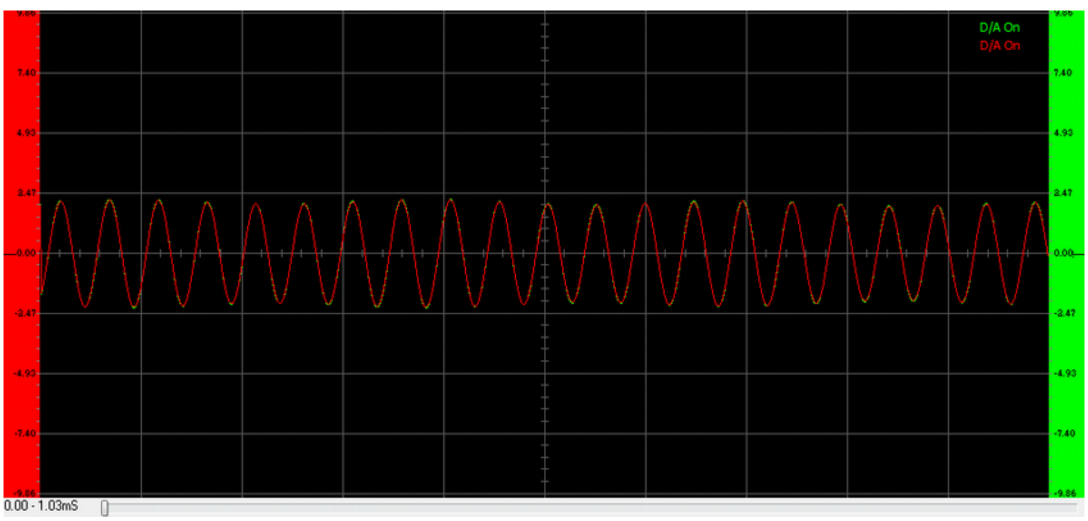

If I switch over to pink noise, and get a transfer function trace, the result (when smoothed to damp any little anomalies) is as flat as a laser beam:

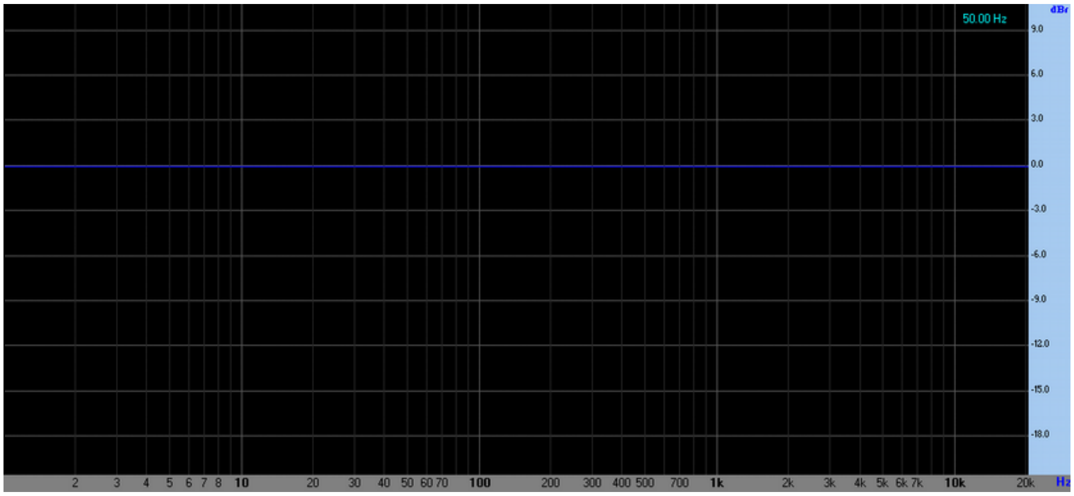

Now, I set up my filter to have a cutoff frequency of 100 Hz. I engage the filter, and...very interesting:

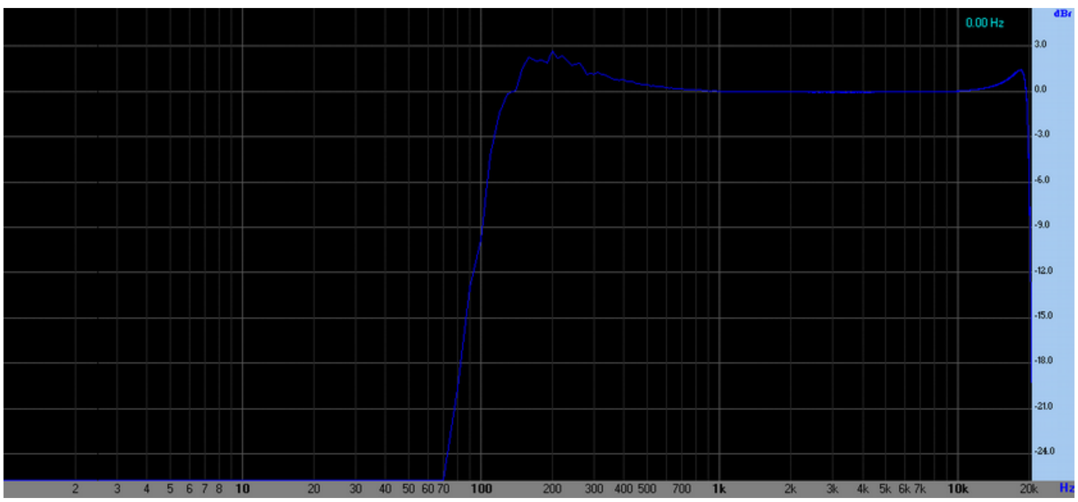

The filter has a pretty wide area where it “rings,” starting at around 700 Hz and (descending through the frequencies) peaking at roughly 200 Hz. The chosen cutoff frequency is about 8 or 9 dB down from flat. A very curious little thing is the added “ring” that's up very high – beyond 10 kHz. What's also curious is that the filter isn't quite as steep as advertised. It's supposed to be 72 dB/ octave, but the difference from 40 Hz to 80 Hz is more like 48 dB/ octave. (Not that I'm complaining very loudly.)

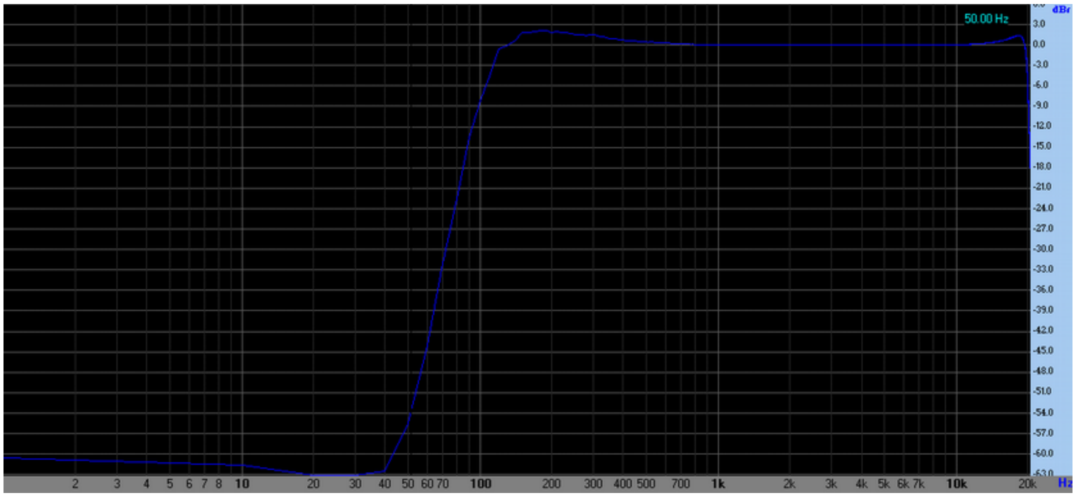

Well, maybe this filter isn't the best choice for the live rig. I wonder what a filter from my favorite EQ plugin looks like...

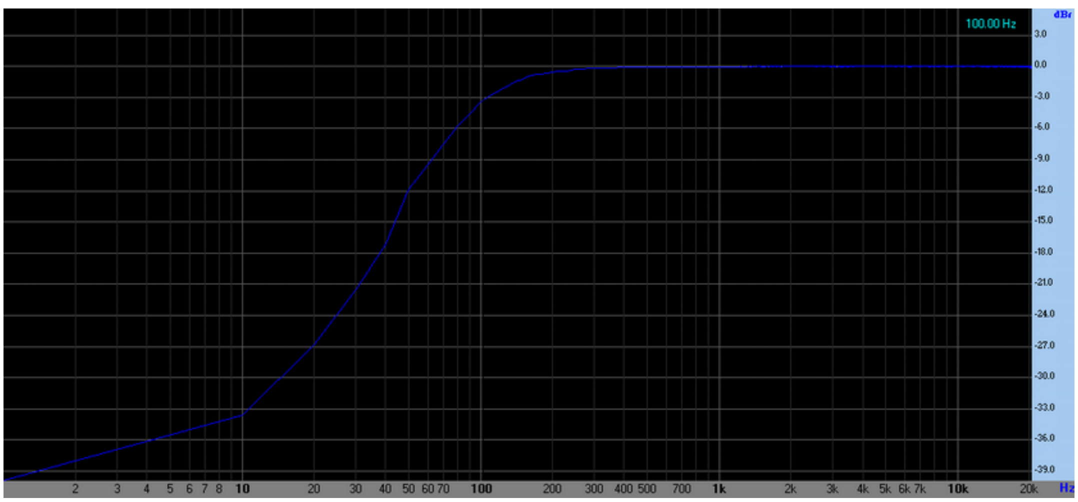

Ah, much nicer. No overshoot, and no weirdness going on in the top octave. The filter slope is, of course, quite gentle in comparison. It's about 10 – 12 dB/ octave. Hmm...what happens if I tweak the filter to be a touch flatter, and then cascade a couple more with the same settings?

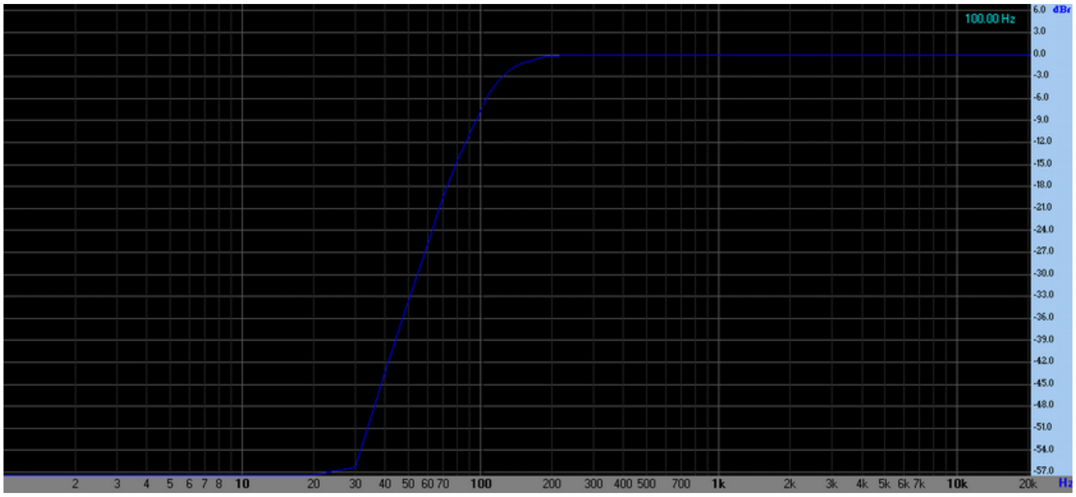

Well, it looks like I've got a 30 dB/ octave filter, with my selected cutoff frequency being about 7 – 8 dB down. The filter has effects between 100 and 200 Hz, but the rest of the trace is pretty much untouched. Neat!

Now, I wonder what a filter from this EQ is really doing when it's set to peaking, -9 dB, 1 octave wide, and 1 kHz?

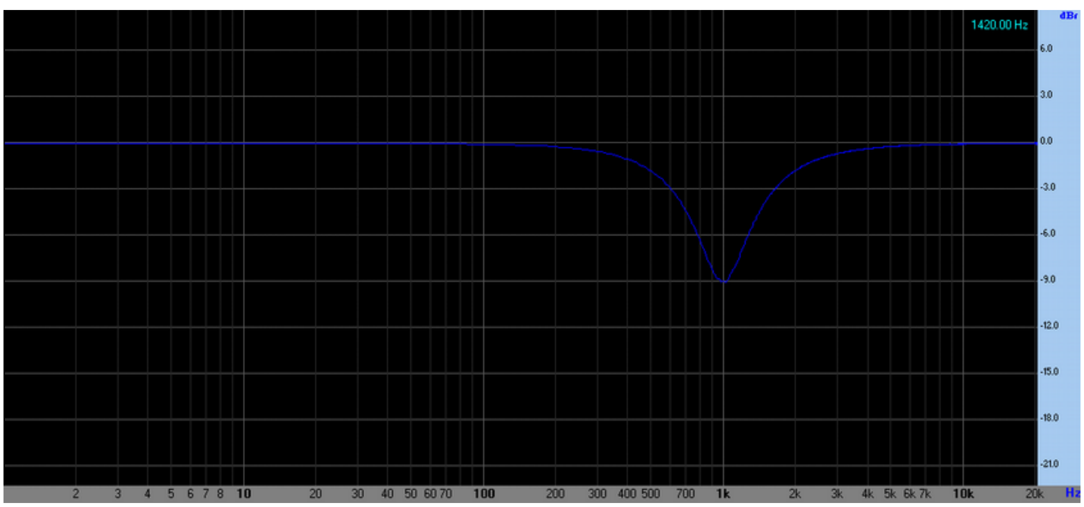

It's pretty much "bang on the setting" of -9 dB at 1 kHz, and not being too destructive below 500 Hz or above 2 kHz.

Now then. I wonder what I can find out about this "knock-around" analog mixer I've got sitting over here. Let's find out what its channel EQ is actually doing when it's set to 9 dB down at 800 Hz, which is an easy setting to get to based on the panel labeling. I'm going to route one side of my noise generator straight back to one of the analysis rig's inputs (as the reference), and the other side is going to pass through the console's input, through the channel EQ, out the console, and then to the other analysis input. I try to get a basic level match between the reference and analysis signals...

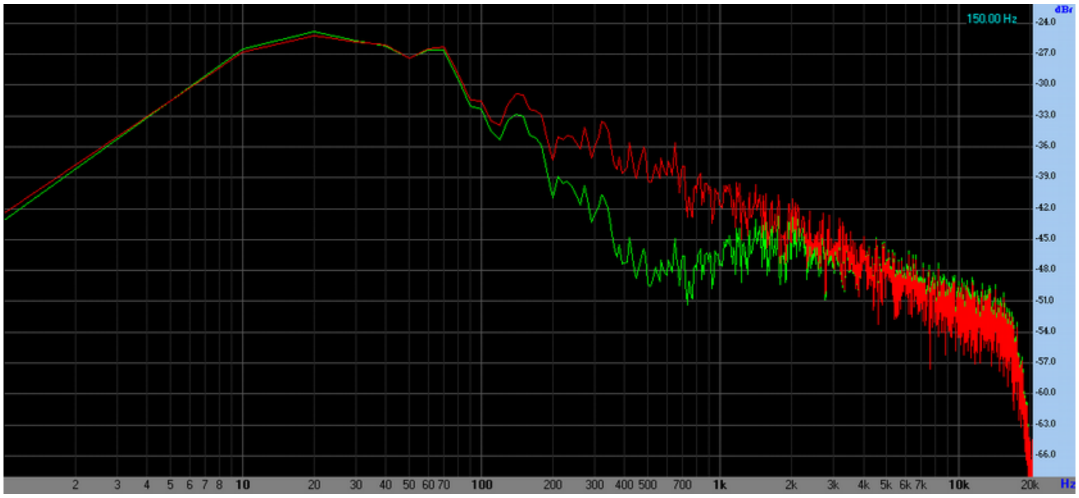

...and then I switch over to the transfer function. After a little more level tweaking to get what I think should be the baseline to 0 dB, here's what shows up:

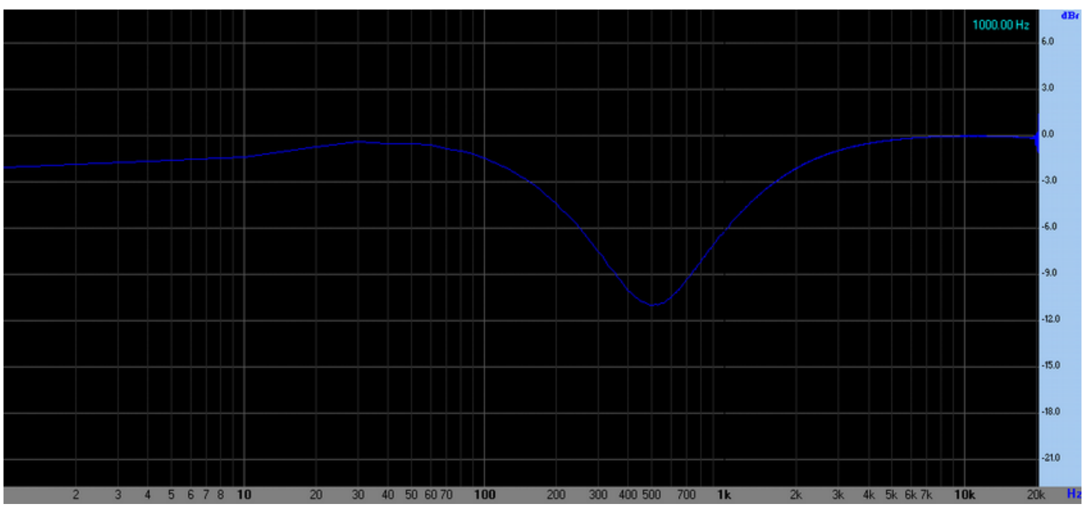

The odd little dip below 30 Hz, and the weirdness near 20 kHz are things that I'm not sure I can comment on meaningfully (my guess is that they're in the realm of experimental error), but other parts of the trace do have some interesting implications. For one, when the EQ frequency selection knob is pointing at the middle “0” in 800 Hz, what you really have selected is 500 Hz. The mark that you think might correspond to -9 dB on the gain knob actually gives you about -11 dB. The filter is a touch wide, but that's about what I would expect from a fixed-bandwidth sweepable EQ like this one.

What about distortion characteristics? Well, with a scope and an analyzer, you can eyeball the point where harmonic distortion starts to take off. Here, we seem to be just fine. The scope shows a nice, smooth sine wave, and the frequency analyzer shows one big spike at 1 kHz, just as it should:

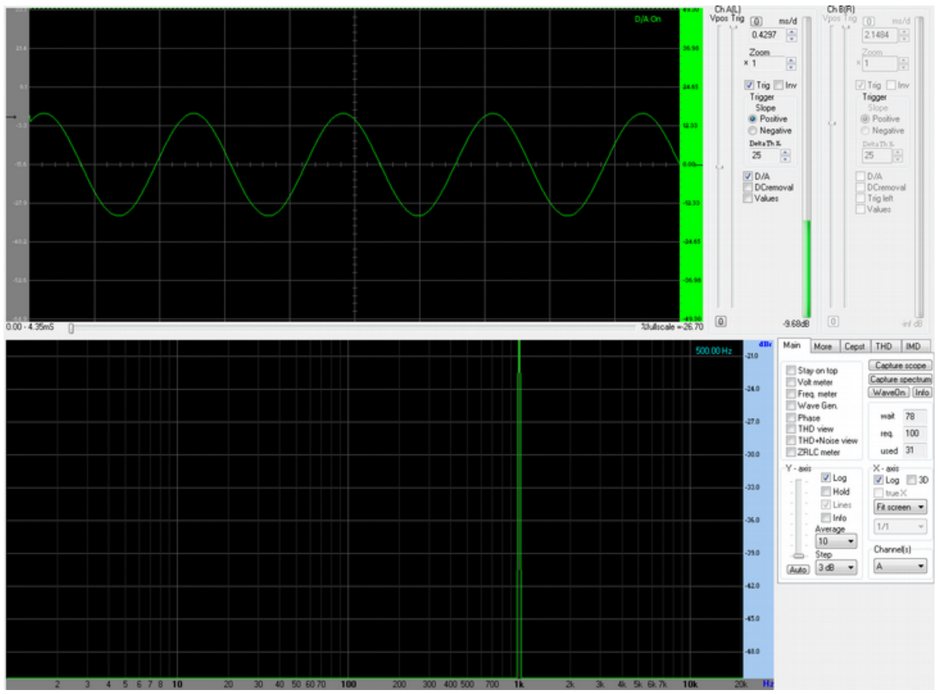

Now, let's push the level up until things don't look so good:

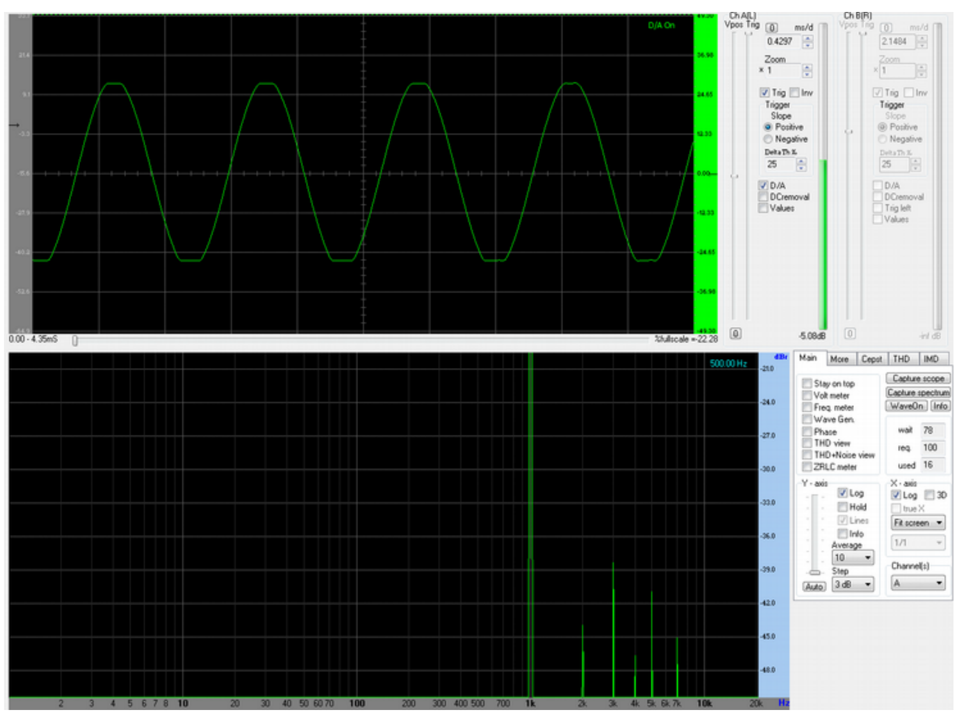

This correlates with the channel clip indicator on the console, in that the clip light is steadily illuminated.

Some analysis systems can do an automated calculation of total harmonic distortion. My rig has this feature, so I'm going to enable it to show something curious about the console. Here's the same situation as above, showing a THD of a bit over 3%:

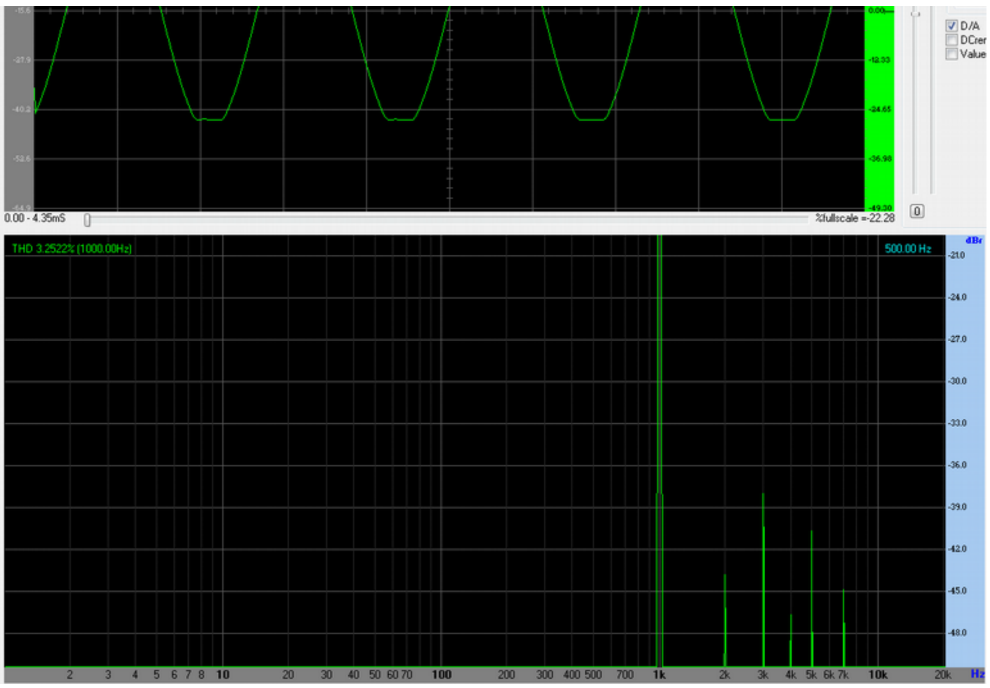

Interestingly, if the gain is backed down just to the point where the sine wave looks nice on the scope, and the distortion components have dropped out of visibility on the magnitude display, the clip light remains brightly illuminated. On this console, the clip light is not completely off until you get down to a voltage corresponding to roughly 0.02% THD.

I could keep going on for pages and pages, but I think you've more than gotten the point by now. Doing some testing can reveal a lot about your gear, and help you to be more knowledgeable about exactly what your gear is doing.