

I'm Doug Fearn and this is My Take On Music Recording

When we think of distortion, we are usually talking about harmonic distortion. But there are several other types of distortion, and understanding all of them will be helpful to you. In this discussion, I define "distortion" as any deviation from the original sound.

This episode is a bit technical, but I'll try to keep it as simple as possible. The electrical engineers out there may object to my over-simplifications, but I do it to make this more accessible to everyone.

We will start with harmonic distortion. In equipment specifications, this is usually shown as THD, for total harmonic distortion, or THD+N, which is the total harmonic distortion plus noise. THD doesn't really tell you much about how the piece of equipment sounds.

To understand harmonic distortion, we first have to know what harmonics are. These are sometimes called "overtones" or "partials" in music. In engineering terms, a harmonic is an integral multiple of a fundamental frequency. OK, what's a fundamental frequency? It is the pure tone equivalent of any musical note. For example, A440 is a standard of musical pitch. It means that A4 has a frequency of 440Hz. All musical notes can be defined by their frequency. And that frequency simply means how many times per second a plucked string vibrates, or a column of air in a wind instrument vibrates.

The second harmonic of A4 is 880Hz, twice the fundamental frequency of 440. Each even-numbered harmonic is an octave above the previous one and is therefore perfectly in tune with the fundamental.

But the third harmonic is not an octave above the second harmonic. It is three times the fundamental frequency of 440Hz, or 1320Hz. When you combine the 440Hz fundamental with the 1320Hz third harmonic, it does not sound good because 1320Hz is not a musical note in the Western scale. The closest actual note is an E, but the third harmonic is flat and does not blend in.

That's the case for all the odd harmonics. The fifth harmonic is 2200Hz, but the nearest note is C#7, but the harmonic is once again flat.

The seventh harmonic is 3080Hz. The nearest note is G7 at 3135.964. This harmonic is very flat.

And so on, up the scale, as far as your hearing goes, and the distortion products extend.

Of course, scale temperament enters into this, but that is beyond the scope of this discussion.

So far, we are just talking about pure tones. But actual instruments are far richer sounding than pure tones. In the evolution of musical instruments, designers worked to produce instruments that had a pleasant array of overtones. Among other things, this meant reducing the odd-order harmonics in the instrument as much as possible.

Each type of instrument has different proportion of overtones, which gives it a distinctive sound. That's why it is easy to tell if the same note is played on a piano or a guitar, for example. The two instruments have very different harmonic proportions. And individual instruments of the same type can sound different, too, because each may have a slightly different mixture of overtones.

And the note and all its harmonics are dynamic in nature, with different overtones fading away at different rates as the note decays. It is the richness of the harmonics that give us beautiful-sounding instruments.

So, what does all this have to do with distortion?

If our equipment adds odd-order harmonics to the sound, the instrument is going to sound discordant and unpleasant. And that's harmonic distortion.

Also, if we cannot reproduce the exact proportions of the fundamental note and all its overtones, our recording is not going to sound like the source instrument.

Emphasizing the even-order harmonics out of proportion is also harmonic distortion. But that distortion is much less objectional, because the harmonics are perfectly in tune with the fundamental. In fact, adding even harmonics makes the sound fuller, louder, and more pleasing to our ears, within reason.

This is sometimes imprecisely called "saturation." Even-order harmonic distortion is a big part of the "saturated" sound, but it is not the entire defining aspect of what we call saturation.

All electronic circuits and devices introduce distortion. As I explained in Episode 22 on microphone preamplifiers, well-designed vacuum tube circuits tend to have mostly even harmonics in their distortion. Solid-state distortion is often mostly odd-order harmonics. That's why to many of us, tube gear sounds musical and pleasant, while solid-state can sound discordant and harsh. Take a listen to episode 22 to get a more complete explanation, especially why the dynamic characteristics of distortion are so important.

In both solid-state and tube circuits, another source of odd-order harmonic distortion can occur in push-pull circuits, which by their nature, cancel out most even-order harmonics. And an asymmetrical problem in push-pull circuits, called "crossover distortion" will also add odd harmonics to the sound. Note that the term "crossover" in this context is entirely different from the crossover network in a loudspeaker system. To minimize crossover distortion, many push-pull circuits have a balance control to adjust the symmetry of the waveform and minimize the odd harmonics.

Another interesting characteristic of our hearing is that our brain interprets a range of even-order harmonics as defining a lower musical note that does not even exist in the recording. Or sometimes isn't even produced by the instrument. That's why you can reproduce a decent-sounding bass guitar on a small speaker that does not actually reproduce the fundamental frequency of the bass note. Your brain "assumes" that it hears the fundamental.

Both even- and odd-harmonic distortion will measure the same way on test equipment, which is why you have to use your ears, rather than read the specs, to determine how a piece of gear sounds.

To characterize the true nature of harmonic distortion, we need to look at the frequency spectrum. That will show the fundamental and all the harmonics in a graphical way. We can see each harmonic, and easily determine whether it is an even or odd harmonic, and its level relative to the fundamental and all the other harmonics.

Our hearing is exquisitely sensitive to low levels of discordant distortion. When I am developing a new audio circuit, I not only measure the distortion with sophisticated lab equipment, but I also listen to the pure tones through good monitor speakers. Quite often I hear distortion increase with a very small increase in level, while the test equipment shows no change at all. I suspect that we can hear deficiencies in audio that we do not have to tools to measure. And that's OK, because the goal is good sounding audio to our ears, not to test equipment.

But sometimes unpleasant distortion is exactly what we want. If the goal of the recording is to put people on edge and annoy them, then the odd-harmonic distortion is quite good at doing that. It's a valid way to achieve a sonic goal.

It's always been interesting to me that the recording technology over the last century has often gravitated toward equipment that emphasizes the even harmonics. The phonograph record, optical audio film tracks, and tape machines all tend to have high amounts of distortion – perhaps 10 percent or more. Yet, the sound can be quite pleasing.

Optical film audio tracks, for example, sound to me like they have about 20 percent distortion, but this is almost all even harmonic distortion and not nearly as objectionable as you might imagine. The distortion makes the dialog, music, and sound effects of those old films really loud (along with some intrinsic compression in that format).

Tape machines usually are spec'd at 3% distortion, but it's mostly even harmonics and it is part of the sound of tape. Overloading, or saturating, the tape increases the even harmonic distortion even more.

Guitar amps are lousy audio amplifiers, with all sorts of deficiencies, including high harmonic distortion. Tube guitar amps, like other vacuum tube audio gear, have predominantly even-order harmonic distortion. The distortion level can be very high – so high that it is difficult to measure with test equipment. But that's the sound of the electric guitar. The guitar/amp combination has evolved over the years to utilize the distortion as part of characteristic sound of the electric guitar.

In almost all cases, the amount of distortion is dynamic in nature. As the sound gets louder, the distortion goes up, emphasizing the dynamics of the performance.

In addition to the guitar amplifier, there are pedals which are designed to increase the harmonic distortion of the guitar. This results in a much louder sound, especially when the even-order harmonics are intensified.

Some guitar pedals take advantage of “clipping.” Clipping occurs when the audio is so loud that the electronic device can no longer even approximate the original waveform, and the peaks of each audio cycle are “clipped” at the power supply voltage. At its extreme, this clipping converts the original waveform to a square wave, which, by definition is the fundamental and all the odd-order harmonics. That’s a very unpleasant sound, but it can be useful in creating an extremely loud and irritating guitar sound. Clipping also occurs in solid-state audio devices, as I explain in the episode on mic preamps.

There are other kinds of distortion in addition to harmonic distortion. Remember, distortion is broadly defined as any change in the characteristic of the original sound.

One of those is what I will call amplitude distortion. This is where the balance over the frequency spectrum changes. In other words, we do not have flat response.

For example, if our electronic device rolls off at the low end or at the high end, we no longer have an accurate reproduction of the original sound. Same thing if there is a peak or dip in response anywhere in the audio range.

That sounds like equalization, doesn’t it? And, yes, technically speaking, equalization is a type of distortion. But it’s one that we use all the time for creative effect, or to make instruments fit together in the mix. In this case, that “distortion” is something we want.

Other times we might use eq to fix a problem in the original recording. That might help, but be careful that the cure isn’t worse than the problem.

Personally, I find I use less and less eq on the sessions I do. Most of the time, I just use a D.W. Fearn VT-5 equalizer on the mix bus, and rarely equalize individual tracks. This helps me keep the sound more solid-sounding. It works for the kinds of projects I do, but it may not be practical for you.

Messing with the frequency response has other implications, too. And that leads us to another form of distortion, called phase shift distortion.

An equalizer will always introduce a phase shift. It’s the nature of modifying the frequency response, and the benefits usually outweigh the drawbacks of equalization.

Equalizers have been designed that have close to zero phase shift. When I first heard about these maybe 30 years ago, I thought it was a brilliant idea. But when I got my hands on one of the zero-phase shift equalizers, I quickly realized that I didn’t like that sound at all.

The inherent phase shift in frequency response manipulation with a well-designed equalizer should sound natural to us, because our acoustic environment often modifies the frequency response. Think of the highs attenuated by sound-absorbing materials.

When I was designing the VT-5 equalizer, I was very aware of how changing component values in the passive eq part of the circuit affected the phase response. Careful design resulted in an equalizer that affects the phase as little as physics will allow. Mimicking nature certainly made the VT-5 sound better.

All audio devices have a limited range of flat response. This isn't a major problem in modern gear, which usually has very little variation in response within the human hearing range of 20Hz to 20kHz. But the frequency response may roll off abruptly just past those frequencies, which can cause problems. This is especially true of low sample rate PCM digital recording, where very sharp cutoff filters need to limit the response above the Nyquist frequency, which is defined as half the sample rate.

For a 44.1kHz sample rate, all frequencies above 22.05kHz must be attenuated a lot to prevent aliasing artifacts. To extend the response as high as possible, the filter has to act like a "brick wall" to frequencies above 22kHz. Such a filter can have serious implications for the sound of the converter. Modern designers are really good at making the filters as inobtrusive as possible. But it is still a distortion.

The higher the sample rate, the less stringent the requirements for the filter, and the less distortion will be introduced.

You probably recognize the term "phase" from the "phase reverse" button on your hardware or software. Technically, this control reverses the polarity of a signal, swapping the positive peaks and the negative peaks of the audio waveform. In terms of phase, it's a 180-degree phase shift.

This polarity reversal switch is a useful tool in the recording process. Although well-designed audio gear should not have an inherent phase reversal, sometimes it does. This is true in some very old gear, or it could be the result of a construction error. Fortunately, it's easy to correct with your "phase" button.

Improper cable construction or wiring can also inadvertently cause a polarity reversal.

There are other times when a polarity reversal is useful, like when you are using two mics that pick up the same sound. Mics should be standardized in polarity, but some are not. If those out-of-polarity mics are combined in mono, the audio will partially cancel out. In fact, the audio may almost disappear.

Mic'ing a snare drum on the top and bottom is a classic situation where reversing the polarity of one of the mics often provides a better sound. Or at least a different sound.

Then there is "absolute polarity." Take a kick drum as an example. When hit, the drum produces an increase in pressure at the microphone. When played back, your speakers should also produce an increase in pressure. But what if the speaker cone moved in the opposite direction? The air reaching your ears would be "sucked in" rather than "pushed out." Does this matter?

Fortunately, our hearing is largely insensitive to the absolute polarity of a sound. That's a good thing, because in the journey from the microphone to the listener, there are literally thousands of places where the polarity could be reversed. The chances that the pressure wave from a sound is reproduced by the listener's speaker the same way it was generated is probably no better than 50%. That's OK, though, since our hearing seems to adapt to the absolute polarity either way.

Some people can detect a difference in absolute polarity, and one way sounds better than the other. But most of us don't notice that.

One exception is if you are monitoring your own voice in headphones. For many of us, we like the sound in the headphones with one polarity or the other. For me, it is always a polarity reversal that sounds better. When I did the final testing before shipment on each of our preamps, I could use test equipment to determine if the polarity was correct. But I found that with a mic and headphones, I could determine that instantly, just by how my voice sounded. You might want to experiment with this when recording a vocalist. They may find one polarity sounds better. Many do not detect a difference.

The change in sound is due to the dual paths of your voice reaching you through bone conduction and through the headphones.

Digital latency will also affect this. That's a subject for an entire episode in itself.

In stereo or multichannel playback, a part that is out of phase in the stereo channels will sound very odd to the listener. For one thing, your ability to localize a sound will be diminished. This is sometimes used in video productions to create a spacey effect. But keep in mind that this will not work if the listener is hearing the audio in mono. The out of phase sound will disappear.

But I want to talk about a subtler form of phase shift distortion.

As I mentioned earlier, an amplitude distortion, in the form of a roll-off, affects the sound of the audio. It is almost always accompanied by a shift in the phase with frequency. Good designs minimize this phase shift so the adverse effects are less obvious.

Here's where the problem comes in. As an example, let's take a recording of a double bass, or stand-up bass. This is a beautiful-sounding instrument that has its lowest fundamental note at C2, or 65.41Hz. That note, and all the notes of the bass, are accompanied by a rich collection of overtones, mostly even harmonics, that give the bass its characteristic interesting sound.

In the real world, the overtones of the bass are mostly in phase with the fundamental note. I say mostly because the overtones may shift in phase depending on the position of the listener or the microphone. Regardless, the alignment of the fundamental and the overtones is what makes the bass sound like it does.

Also, the various overtones decay at different rates.

However, if our electronic circuitry has any phase shift, then the relationship between the fundamental note and its overtones will change. Therefore, the sound of the instrument will change, and this is almost always for the worse.

The bass will lose its luster and may start to sound distant or indistinct. I use the bass as an example, but this applies to all instruments.

To preserve the sound, the overtone structure must be maintained with as little phase shift as possible. This requires careful design of the audio equipment.

This is something I strive for when I am designing a piece of equipment. Through careful design, the phase shift can be kept to a minimum, and the sound preserved.

That way, the instruments, especially instruments in the lowest register, maintain their solidity. It really does make an obvious difference when done well.

But wait. All audio circuits will eventually roll-off in frequency response. Won't that affect the sound?

It certainly will. That's why I design my products to extend the low frequency response as low as possible. For example, our mic preamps have -3dB point of 0.5Hz. Actually, it's nearly impossible to directly measure response that low, so the half-a-Hertz figure is derived from computer modeling of the circuit, confirmed by extrapolation of the response as low as it can be measured.

That extended low-frequency response minimizes the phase shift at low frequencies. I can measure the phase shift pretty accurately, and it is only a few degrees at 20Hz. And that makes the low-end sound of our preamps very solid. Our line-level products like the VT-5 equalizer and the VT-7 compressor have similar extended low-end response.

The same thing applies to the high-end of the audio spectrum. The actual energy of a percussive sound may extend well above the 20kHz generally accepted as the upper range of our hearing. Certainly, low sample-rate antialiasing filters will eliminate those overtones. At higher sample rates, the overtones are less affected.

There can be subtle changes in the sound of the attack of the percussive notes if the high end rolls off. That's also distortion. If we roll off the high frequency response, the percussive sound will sound different, lacking in "punch." Or the attack may become indistinct. That's because we have distorted the waveform of the initial transient of the sound. Extended high frequency response helps preserve the impact of percussive sounds. D.W. Fearn products have a high frequency response that rolls off smoothly. It is down 3dB somewhere between 55kHz and 100kHz, depending on the product. You will preserve the transient information if the recording system has equally good high frequency response. And extended high frequency response requires high sample rates. I think a 96kHz sample rate is about the minimum for clean percussive transients.

Another kind of distortion that you might encounter is called intermodulation distortion. This type of distortion arises when there are two distinct frequencies present. Of course, in music, there are many frequencies produced, but just looking at two tones makes it easier to understand.

With intermodulation distortion, the distortion products are the sum and difference of the two frequencies.

For example, let's take two pure sine waves, one at 261.63Hz, which is C4, and another at 329.63Hz, the E above the C. This is a very common harmony combination, a major third, and it sounds nice.

But if we put it through an audio device that has intermodulation distortion, we hear two other frequencies, the sum, which is 591.26Hz, and the difference, which is 68Hz. Neither of those falls on a musical note.

It's an ugly sound, for sure. Intermodulation distortion has no redeeming qualities.

Fortunately, modern equipment has very low intermodulation distortion. If you hear this kind of sound, you can almost be confident that something has failed inside the equipment.

But subtle intermodulation distortion may be present even without a failure. And its causes can be surprising. More on the sources of intermodulation distortion in the next section.

Let's look at how each of these types of distortions come about. First, poor equipment design can cause any of these distortions.

But sometimes you start to hear distortion in equipment that used to sound fine. This can be caused by certain electronic components that have reached the end of their lives. For example, vacuum tubes will eventually wear out and the first clue may be an increase in distortion. You might actually prefer that sound, if the distortion is even harmonics.

In solid-state gear, one source of distortion is electrolytic coupling capacitors in the audio path. Because solid-state audio amplifiers are intrinsically low impedance circuits, the capacitors linking one stage to the next have to be fairly high capacitance, typically 10 to 100mFd. The only way to get that much capacitance in a small space is to use electrolytic capacitors. And those have problems, one of which is a finite life, usually measured in a few years to a couple of decades. Gradually the capacitance starts to drop and the function of transferring audio within the equipment degrades. This is often first noted as a loss of low frequencies.

But even more insidious is the non-linear characteristic of electrolytic capacitors. In other words, they modify the audio going through them, adding harmonic and intermodulation distortion. They can also affect the phase response as they deteriorate, as the frequency response changes. All the distortions we talked can be in one electrolytic coupling capacitor.

Multiply that by dozens of capacitors in the audio path and you can see why consoles and other solid-state gear greatly benefit from periodic replacement of those capacitors.

You rarely hear about "re-capping" vacuum tube gear, because they don't need it. The coupling capacitors are much smaller in capacitance, due to the intrinsic high-impedance nature of vacuum tubes. We don't have to use electrolytic capacitors in the audio path at all. Instead, quality gear uses film capacitors, which have a very long life – maybe hundreds of years.

In both solid-state and tube equipment, the power supplies use electrolytic filter capacitors. They do not affect the audio path, but as they age, they can no longer do a complete job of filtering out the noise component of the power supply DC output. This is generally noticed as an increase in hum at the mains frequency, which will be either 50 or 60Hz, depending on what part of the world you live in.

Poor power supply design can introduce its own problems for the audio device. There are several ways that sub-optimum power supplies can produce distortion in your audio. There's not much you can do about that except to buy well-designed equipment. The distortions caused by the power supply tend to be dynamic in nature, becoming more apparent as the level is increased. This has implications for transient response.

Audio transformers can be a source of harmonic distortion. Transformers can also cause a phase shift. And they might not have flat frequency response.

Top-quality audio transformers require very careful design, expensive metal alloys, and precision manufacturing. That's why in all the D.W. Fearn products, we use custom transformers made for us by Jensen Transformers. There are other top-quality transformer manufacturers out there, but some equipment designers use cheaper transformers, which may introduce distortion.

Another source of distortion is patch bays. Most of the small TRS patch panels and cords have very little actual contact area for the audio circuit, and the jacks and plugs are subject to oxidation over time.

If that oxidation prevents a solid, low-resistance connection, a non-linear junction may result. This could be heard as an increase in distortion, particularly intermodulation distortion. But it could be any of the distortions I have talked about.

The oxidized contacts may form a semi-conductor junction, like a diode. Diodes can be used to generate both harmonic and intermodulation distortion in non-audio devices where such a thing is necessary. Of course, we definitely do not want those distortions in our audio.

This distortion generally increases gradually, and you might not notice it until one day you realize that things just don't sound right anymore.

If you are stuck with TRS patching, then cleaning and polishing the plugs and jacks is vital. In the old days, this was a classic intern project. Every week at my studio, an intern would use brass polish to clean all the plugs on the patch cables. Cleaning the jacks was more difficult, and required a special tool.

Console patch bays often include insert points for outboard processing. When the patch plugs are removed, the audio stays within the console. This is accomplished by using switch contacts within the jack itself. And those can get dirty or oxidized, creating distortion. And cleaning the switch contacts is not easy. Spraying contact cleaner into the jack may help, but eventually the audio quality is going to degrade no matter what you do.

Sometimes this problem is obvious, like when audio stops going through the patch panel. But other times it is subtle. An upgrade of your monitoring system, for example, might suddenly make the distortion problem more obvious.

One way around this is to use XLR connectors for your patch panel. Sure, it takes up about 8 times as much space, but the larger contact area of an XLR plug and jack provides much more contact area, and the design of the connector includes a self-cleaning action every time the connector is used. Some studios use an XLR patch panel for mic lines and preamp inputs, with a conventional TRS patchbay for all other interconnections.

I plan to talk more about patch panels in an upcoming episode on wiring and connections.

Your microphones can also be a source of distortion. It could be from the active electronics in condenser mics, which use the same sort of circuitry I discussed before. Tubes in condenser mics eventually will need replacement, and the capacitors in solid-state mics will have to be replaced at some point.

Also, the capsule itself may cause distortion. In condenser mics, the very thin plastic diaphragm will eventually become contaminated, and this will change the characteristics of the mic. It can create all the distortions I have talked about. Even a perfect capsule will introduce some unavoidable distortion. This is difficult to measure and rarely found in the mic specifications.

Ribbon mics also have some degree of inherent distortion, which to my ear is always less than a condenser mic's distortion. The ribbon might be damaged, from excessive air pressure if the mic is used outdoors, or even by rapidly swinging the mic on a boom. Result: distortion.

Dynamic mics are more rugged, but also subject to many types of inherent distortion. And they often lead a rough life; dropping a dynamic mic could cause misalignment of the inner workings and introduce distortion.

And we can't ignore loudspeakers and other sound reproducers like headphones and earbuds. These are probably the greatest source of distortion of all the elements in the chain from performer to listener. They have all the types of distortion I have talked about, although harmonic distortion dominates their distortion profile. That a loudspeaker sounds as good as it does is amazing to me, since they are incredibly difficult to design well.

So, does any of this information help you in your job? Maybe not. Maybe you just replace a piece of equipment when the distortion becomes apparent. Well, I design the D.W. Fearn products to have a minimum of a 50-year life, and should any distortion arise in one of our units, it is usually pretty easy to repair.

Your knowledge of distortion can be helpful to you in troubleshooting a problem in your studio. Judging by the high distortion levels I hear on many recordings, a little insight into the causes and cures might just help all our projects sound a lot better.

I have to tell you one story from early in my recording career, an experience that taught me a lesson in distortion that I have never forgotten.

I was doing a vocal overdub session with a singer I have worked with occasionally over nearly 50 years now. As we were running down the song, I immediately noticed that his voice was distorted. Where was that coming from? I quickly changed mic inputs on the console, but the distortion remained. I went out to the studio and changed the mic. No change in distortion. After going through every single device, connector, and wire that the audio went through, the distortion was still there.

I went out into the studio, wracking my brain trying to think where this distortion was coming from. What had I missed? While I was standing there, the vocalist started singing.

Only then did I realize that there was no technical reason for the distortion. It was his voice! I listened in amazement and chagrin as I heard the natural distortion in his voice.

Lesson learned – your sound source may be distorted, and there's not much you can do about it.

By the way, over the succeeding decades, his voice has mellowed and matured, and when we record his vocals these days, there is none of that distortion.

And one more thing before I go, let me tell you about a new feature I recently added at dougfearn.com. For those episodes of this podcast that have written scripts, I have posted them in PDF format that you can download if you like. I can't guarantee that they are a perfect word-for-word version of the podcast, because I sometimes do some editing in the process, or add some additional material. And, of course, the interviews are not scripted. And some episodes are just done from notes, so there is no transcript for those.

Let me know if the transcripts are useful to you.

And your comments are always welcome. You can email me at dwfearn@dwfearn.com

This is My Take On Music Recording. I'm Doug Fearn. See you next time.