Lies, Damn Lies, and Specifications

Mike Rivers

It's often said that there are three kinds of lies - lies, damn lies, and specifications. This month we'll look at some typical specifications for audio gear and consider their significance in terms of a unit's sound as well as their likely effect on your pocketbook.

Specifications or Characteristics?

A frequently asked question (and really what prompted this article) is "Does anyone have all the specs on the Ultrafirtzalizer UF-2416?" Generally what they really want to know, at least from that initial query, are the functional characteristics rather than detailed specifications. A functional characteristic describes what something does. A specification describes how well it does it in terms that can be quantified by measurement. Characteristics can be vague ("adds tube warmth") but can also be very specific (24 input channels). Specifications, on the other hand, are real numbers (THD less than 0.01%), often with a tolerance attached (Frequency response flat, 20 Hz - 20 kHz, +/- 2 dB). The tolerance allows for manufacturing differences and a certain amount of creativity in marketing. Study the characteristics when you want to learn how many bands of equalization or auxiliary sends a console provides, then read the specifications (frequency response, gain, distortion, noise) to see how well it's likely to perform.

How Important are Specifications?

It depends. Compare spec sheets among similar products and most of the specifications that they have in common will be nearly identical. This is part truth and part marketing. In the realm of pro audio, we aren't always able to quantitatively measure all of the differences that we can hear can when comparing similar units in actual use. Take three mic preamps which claim nearly identical frequency and phase response, noise level, gain range, maximum output level, and total harmonic distortion, all specifications which can be verified in a laboratory (no damn lies here). Yet upon critical listening, one sounds more "open", another sounds more "warm", and you might describe the third as sounding "harsh in the upper midrange". And another listener may disagree. There's obviously something going on that's not told in the spec sheet.

Beginners in the audio field tend to take a spec sheet with a lot more reverence than may be justified, often because it's the only place to find any technical information at all about the unit. In the spec sheet, you'll usually find some important characteristics that will tell you if you're looking in the right ballpark, but rarely will it tell you anything that's likely to predict how the device actually sounds.   Here's some guidance as to what some of the more common specifications mean, and how to put them into perspective.

Frequency Response

Frequency response is the most often quoted specification in the audio field. It describes how the output relates to the input over the range of frequencies within the audio range (20 Hz and 20 kHz) and sometimes beyond. It's usually expressed as a bandwidth and a tolerance within that range, for example: "25 Hz to 20 kHz, +/- 2 dB." A graph of amplitude vs. frequency is also a common way of showing frequency response. A straight line represents no variation in amplitude, hence the term "flat response". The frequency axis is almost always a logarithmic scale, with the equal distance along the axis representing 20 Hz to 200 Hz and 2 kHz to 20 kHz . This tends to go along with the way we hear and makes the graph prettier.

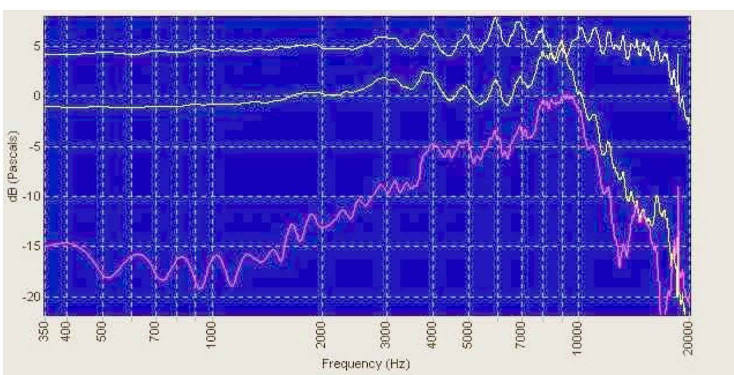

Frequency response graphs aren't very interesting if they're flat. This means that all the frequencies that go in comes out in the same relative amplitudes - pretty much what we expect from an amplifier, mixer, signal processor (when it's not processing) or recorder. This is the measured frequency response of a cardioid microphone . The top curve is the frequency response measured on axis, the middle curve is the frequency response measured 90 degrees off axis, and the lower curve is the frequency response from the rear (180 degrees off axis).

This doesn't come from a manufacturer's spec sheet, but it's really useful information. If, for instance, the spec sheet claims 20 dB of off axis rejection, you can see that this is true, but only with a signal coming into the rear of the mic, and only below 2 kHz. Lie, or damn lie? You decide.

There are, however, limits beyond which the equipment is expected to perform, or where performance is intentionally restricted. A power amplifier's frequency response may drop off sharply below 20 Hz to protect a loudspeaker from low frequency noise below the range of hearing. A console may have its high frequency response restricted so that it doesn't pick up radio stations.

There are good arguments for extending frequency response of audio gear beyond 20 kHz, but how far is still being debated, mostly in audiophile circles. Noted designer Rupert Neve has long contended that near flat response extending beyond 50 kHz is important for transparent sound, and he designs his equipment accordingly. Higher isn't always better, but knowing that a piece of audio equipment can pass 100 kHz tells you something about the designer's philosophy.

How flat should frequency response be? Nothing's perfect, but we can learn a few things about the equipment's performance by studying irregularities in its frequency response. Most electronics exhibits very flat frequency response within its specified limits. Where you see irregularities of any significance is in transducer-based equipment where electrical energy is changed from electrical to mechanical energy or the other way around. Microphones and loudspeakers are the prime example, but analog tape recorders (which convert to magnetic energy and back) are likewise guilty.

A microphone's response may be specified as 50 Hz to 18 kHz, +/- 3 dB, but the manner in which the frequency response varies over that 6 dB range is important in determining the characteristic sound of the mic. Typically all microphones drop off at the high and low ends, so the "-3 dB" points might well be the frequencies where, relative to some mid-band point (1 kHz is typical), the output drops by 3 dB.

But what about the "+" part of the specification? The microphone might have a smoothly rising response up to some frequency, above which it falls off smoothly, or it might have a well defined 3 dB hump at some frequency. This is often by design, to improve vocal presence or to get more "whump" out of a bass drum. The bass response might be rolled off to reduce proximity effect in a microphone that's intended for working close. This is the stuff of which different sounding (though similarly specified) mics are made.

The Shure SM-58 (this frequency response plot is from the mic's data sheet) gets its well known character from significant peaks in the vocal "presence" range, while the AKG D-12 is similarly famous for a peak in the kick drum "thump" range. Neither of these mics are "flat" but both are useful because of their particular frequency response characteristics. A mic with a smoothly rising frequency response up to 12 kHz or so gives an "airiness" to vocals. And a truly flat mic such as those made by Earthworks or DPA is extremely accurate and rather unforgiving of the source or

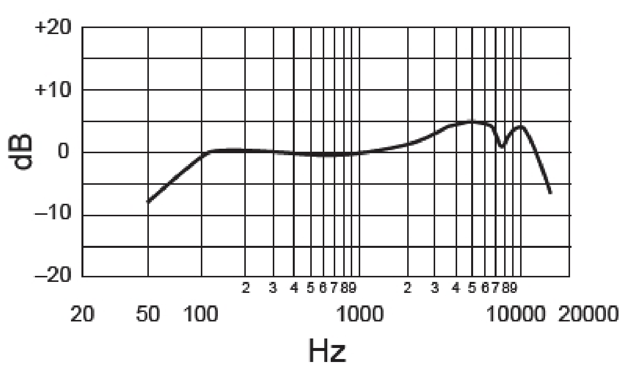

the room acoustics, nor will it add any coloration to the sound. What you hear (or rather, what the mic hears) is what you get.

A loudspeaker’s actual measured frequency response curve is typically considerably uglier than that of a microphone, though to read a spec sheet you’d think speakers were all pretty flat. A speaker’s frequency response may indeed be quite flat when measured in an anechoic chamber, but this means little in a normal listening environment. The low end of a speaker’s frequency response will tell you if it will reproduce the pedal tones of an organ, and the high end response will tell you if you’ll be able to hear that 10 dB “air band” boost you cranked in. But what happens to frequencies in the middle can only be determined by critical listening. Two speakers with similar specs can (and do) color the sound in very different ways.

Frequency response of a loudspeaker or microphone varies according to position - where you’re listening or playing with respect to its axis. If only one spec is given for frequency response, it’s always the on-axis response and, unless it’s a very good omnidirectional mic, you can expect it to be less than flat anywhere off axis. When choosing a microphone, it’s important to know its off-axis (polar) response characteristics, either to take advantage of them or know where you’ll run into problems. If, at 45 degrees off axis, a microphone drops several dB of highs, it’s the wrong mic for recording a wide sound source, but it might be just the ticket for rejecting leakage from another instrument off to the side of the one you’re miking. The off-axis frequency response of a loudspeaker can help you to determine how well the PA will cover an audience, or how wide your monitoring “sweet spot” will be.

Often a spec sheet will show off-axis frequency response as a polar plot, with selected frequencies represented by a family of blob-shaped curves (roughly heart-shaped in the case of a cardioid microphone) surrounding a central point (the mic or speaker). The polar response of a Shure SM58 (taken from their data sheet) is shown here. Note that while the response from the rear is always several dB lower than that of the front, what happens as you go around the mic varies with frequency.

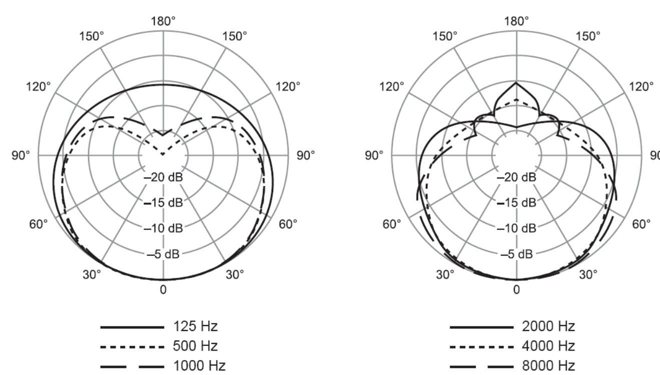

The more familiar frequency-vs.-amplitude graph may have two or more lines plotted on the same scale to show frequency response at one or more off-axis angles. Typically mic spec sheets (if they address the issue at all) will show response on axis and 180 degrees off-axis response. Be aware, though, that a supercardioid mic has two directions for minimum pickup, about 120 degrees clockwise and counterclockwise from the front. This is often important since the rear of a mic often ends up being pointed at some other sound source - another instrument or perhaps a PA monitor speaker.

A good analog tape recorder is rarely flatter than +/- 2 dB over the audible range, but at a tape speed of 30 IPS, reasonably flat frequency response may extend well beyond 25 kHz. Digital recorders, A/D, and D/A converters have frequency response flat within a few tenths of a dB over the audible range, but dropping like a rock just below half the sampling frequency (the Nyquist frequency, from the Nyquist/Shannon Sampling Theorem). For a 44.1 kHz sample rate, generally frequency response isn't given above 20 kHz. In fact, when measuring the frequency response and distortion, it's common to restrict the bandwidth of the measuring equipment to the Nyqist frequency to avoid measuring a stray clock signal that makes its way to the output. But an honest spec sheet will state how much clock leakage appears at the output. When measuring an analog tape recorder, the measurement bandwidth is also limited to avoid including the bias signal in the measurements.

An equalizer's function is to intentionally create a non-flat frequency response, but in a useful way. The action of an equalizer is often presente as a frequency response curve, or family of curves, showing the boost/cut range of each control. An equalizer that specifies a range of +/- 15 dB means that it can provide 15 dB of boost or cut at its operating frequency. It's useful to know if this boost is a narrow or wide peak, or has an adjustable shape. Bandwidth of an equalizer's boost or cut is sometimes expressed in terms of octaves - a 1/10 octave notch, or a two octave gentle curve. It may also be expressed as a number called Q. A high Q (generally greater than 5) describes a sharp peak (or notch if you're cutting rather than boosting), one that affects a narrow range of frequencies, useful for "corrective surgery". A low Q less than 1) is a gentle hump or dip, generally described as "musical".

Distortion

All electronics is non-linear to a degree - you never get out exactly what you put in. Distortion is the term that we use to describe the difference between input and output, and it comes in many flavors. The most common distortion specification, because it's easiest to measure and has the most appealing (lowest) number, is Total Harmonic Distortion, or THD. If a spec sheet says "distortion" without saying what kind, you can be pretty sure it's THD.

THD is measured by putting a sine wave (single frequency) test signal into the unit and measuring everything that comes out after filtering out the test frequency. What's left behind are harmonics of the test signal that result from the non-linearity of the electronics, plus whatever noise and hum is present.

THD is generally expressed as a percentage of the amplitude of the test signal (before it's filtered out), or as a number of dB lower than the test signal's level.

A couple of damn lies associated with THD (or more accurately, sins of omission) are that the test frequency or its amplitude isn't always specified. Since hum and noise are independent of the signal (an exception is digital quantization noise, but that's worth another article), using a high amplitude test signal can make a noisy unit look cleaner on the spec sheet than it may be in your studio. But the truth is that nearly all THD measurements are made as close as possible to the maximum operating level.

Most power amplifiers exhibit what's known as "crossover distortion", which becomes more significant at very low levels. This has nothing to do with a loudspeaker crossover network, but rather, it's a "glitch" right around zero voltage when the signal swings between positive and negative. Power amplifiers are typically measured at sufficiently high power so that any crossover distortion is insignificant compared to other contributors to the THD. Crossover distortion, while buried in the THD spec, might make an amplifier sound "grainy" at low levels. At high power levels, THD varies considerably with frequency. An "honest" power amplifier spec sheet presents a graph of THD vs. frequency at both high and low levels.

These days, THD of solid state small signal equipment tends to be in the range of a few hundredths of a percent, with power amplifiers getting up to maybe a tenth of a percent when pushed near their limit. Back in the bad old days, distortion less than 1% was considered inaudible, but we've learned to hear better than that, and now we can hear differences between units with very low distortion.

On the other hand, we've learned that we like some forms of distortion - tube equipment is a good example. Some highly regarded new and vintage tube gear measures a few tenths of a percent THD (if there's a "Drive" control, you can usually crank this up to 10% or so) and we love how they sound. Some signal processors make constructive use of harmonic distortion (primarily second and a little third harmonic) on the order of .5 to 5 percent, which we often perceive as an increase in clarity, apparent loudness, and that mysterious "analog warmth." We also love the sound of gear with THD that's well under .005%. So THD is a specification that, in absolute terms, isn't one to live or die for.

Intermodulation distortion (IM or IMD) is a little nastier. While harmonic distortion adds frequencies which are harmonically (and mostly musically) related to the input signal, IM adds frequencies that are not related. IM is a measure of how two frequencies present at the same time interact. We hear IM as a buzziness or mushiness.

IM is measured by applying a test signal consisting of two or more different tones mixed together to the equipment. Non-linearity will result in generating new frequencies that are the sum and difference of the two test frequencies, often a family of new frequencies including sums and differences of the harmonics as well as the fundamental frequencies. IM measures frequencies that are present at the output other than the ones used for the test. Loudspeakers are the most notorious IM sources because a certain amount of mixing is inevitable in the mechanical motion of the speaker cones, but good electronics should have IM on the order of a few hundredths of a percent. There area few standard test methods for IM (if you're reading a good spec sheet it will say which method was used, though you may have to look it up to see what the actual test was) which produce a single number, or it may

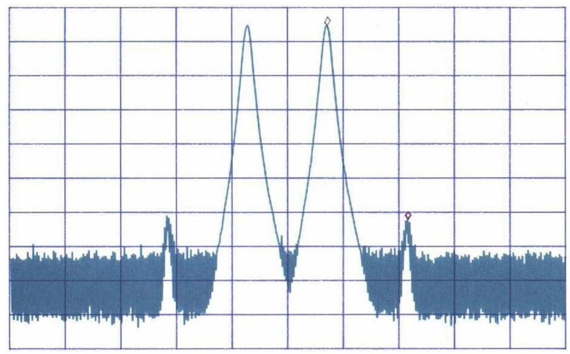

be presented as a frequency spectrum plot. The graph above shows the two test frequencies (the large peaks) with two IMD products represented by the smaller peaks.

Transient Response

Accurate reproduction of transients is very important to music recording, yet its one of the least quantifiable characteristics. For good transient performance, a device not only must respond quickly to a very fast change at its input, but it must also not contribute anything of its own when responding. This means that when the stimulus goes away, so should the device's output. A microphone with extended high frequency response has the potential for providing good transient response, but you don't want its diaphragm to continue ringing like a drum head after the initial sound dies away. Transient response is best evaluated by listening. Slow transient response sounds dull.

Slew rate is a common specification for transient behavior of electronics. It's expressed in volts per microsecond (V/usec). This tells how fast the output rises, given a fast (in theory, going from zero to full level instantaneously) rising input. A slew rate of 40 V/usec means that it would take one microsecond for the output to change by 40 volts. Since studio equipment typically operates from a ±15 volt internal power supply, its output can never really swing 40 V. In reality, it might go to 5 V in 1/8 of a microsecond. Power amplifiers, on the other hand, can swing relatively high output voltages. You can never have too high a slew rate, but expect it to be higher for small signal devices (which don't have to swing many volts) than for power amplifiers.

Noise

Nowhere is specsmanship so skillfully applied as when describing a unit's self-noise. In order to learn from a specification sheet just how much noise you're likely to hear in unit's output, you need to know the test conditions, which are often not specified. Noise floor, signal-to-noise ratio, equivalent input noise, and just plain noise all describe what comes out of the box with nothing going in.

Noise is random by definition, but output noise usually includes a certain amount of hum at the power line frequency and its harmonics. Typically, noise is measured over a limited bandwidth so as not to include what that you can't hear.

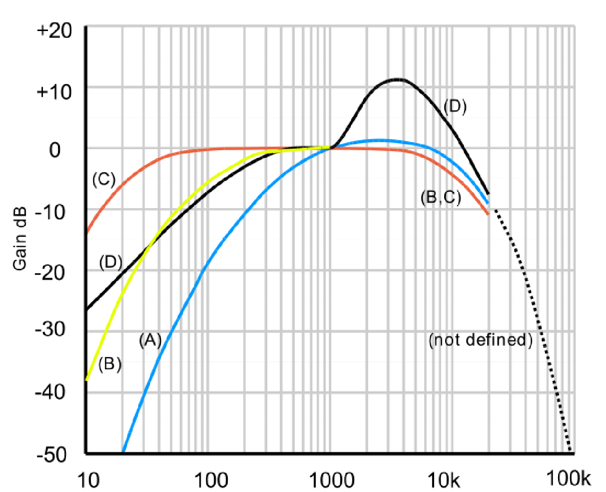

A-weighting (blue), B (yellow), C (red), and D-weighting (blk)

Sometimes the measurement bandwidth is stated as part of the noise spec, sometimes not. Neither can you always assume that an unspecified bandwidth is 20 Hz to 20 kHz. Sometimes a "weighted" noise specification is offered, where the noise is measured after filtering out frequencies where the ear is less sensitive. The intent of such weighting is to emulate the way we actually hear. The result is a more impressive number, since not all of the noise that's present contributes to the measured value. Note that with A-weighting, the fundamental power line frequency is attenuated by about 30 dB before measuring the noise output, and high frequency hiss is attenuated as well. C-weighting is about as close a you'll get to a flat noise measurement, D-weighting is used for measuring aircraft noise, and nobody uses B-weighting any more.

Noise is generally expressed in dB below the unit's nominal output (which you should be able to determine elsewhere in the spec sheet). At what gain setting? You may have to guess.

Equivalent Input Noise (EIN) is becoming fairly common on the spec sheet for a mic preamp. This is a derived figure. equal to the noise measured at some gain setting, minus the gain. If a mic preamp puts out -85 dBV of noise when set for 40 dB of gain, (pretty good performance, by the way) the EIN is -125 dBV. (-85 -40) Specsmanship warning! When it comes to quiet, -125 looks better than -85 even though they both represent the same amount of noise. EIN is actually a fairly useful figure to a designer because it tells him about the noise performance of a particular component. It's less significant to the system engineer (that's you), however if you know the gain at which EIN was measured and under what conditions, you can estimate the actual noise output level. EIN is usually measured at maximum gain because that's usually where the number comes out best, and with a 150 Ω resistor connected to the input to simulate a microphone (which is actually a good non-cheating way to do it).

Signal-to-noise ratio (S/N) is the ratio of the noise output to the maximum signal output. If a unit's maximum output is +22 dBu and the S/N is specified as 75 dB, expect to see about -53 dBu (1.7 millivolts) of noise at the output with no input signal present.

Environmental Specs

Power requirements (including estimated battery life if you're looking at a battery-powered unit), operating temperature range, dimensions, weight, and mounting configuration are all important, as are what accessories are supplied and what you might need. How about controls? What can you do with the knobs and front panel display, or must you connect a computer or monitor and keyboard in order to do certain things? These are usually not specifications in the traditional sense, though you may find some of them on the spec sheet. For others, you may need to look in the user's manual or in other places in the manufacturer's literature or web site.

Digital Specs

This is the part you're all waiting for, but it's where I'm copping out. We all know that digital reproduction is perfect, so why should we need any specs? About this bridge in Arizona you'd like to buy . . . We're still learning how to measure digital equipment and what those measurements mean in terms of sound. There is no spec that will tell you whether the last two bits on a twenty bit D-to-A converter contain anything other than noise. (I wrote that in 1997. Today's 24-bit converters usually can be relied upon to be pretty good to at least 20 bits) There are specifications on clock jitter that look like such small numbers you'd think they were insignificant. And technology is moving faster than the industry and professional organizations can standardize measurement techniques. If it sounds good, it probably is good. If it sounds bad, it probably isn't good. Don't sweat the specs, at least not this year.

One specification for an A/D or D/A converter or a device that contains one, like a mic preamp with digital output, is the relationship between input or output level (depending which way you're going) and bits. For an output device, how many volts or dBu come out when converting from a full scale (0 dBFS) sine wave? Does 0 dBFS = +20 dBu, +12 dBu, or maybe +24 dBu? There's no standard for this, so it's important to know when you're connecting that output to a speaker system. Also, understand that the gain specification for a preamp with digital output, unless there's also an analog output, is usually meaningless. What's important is what voltage it takes (usually at the maximum gain setting) at the input to get full scale digital output.

The Bottom Line

Specifications can give you a clue as to a unit's design or intended application, and in a general sense, what level of performance you can expect. They don't, however, tell you very much about how something will sound. Look over the spec sheets, but don't use them as your sole guide to a decision. Ask questions of other users, look over the controls and connectors, and if at all possible, listen and play with it. Buyer beware, but be informed, too.