

ON THE BENCH

Decibels and The Highs & Lows of Dynamic Range

Are you on the level? Or does your knowledge of technical specs level out pretty darn quickly? This issue we elevate your know-how to a new level.

Text: Rob Squire

► Once upon a time radio stations all across Australia employed chief technical officers. Among the many and varied jobs these engineers performed was to send an annual report to the GPO (Postmaster's General Office). Believe it or not, the original Broadcasting Act was administered by the GPO, a fact which harked back to the days of telegraphy, which the GPO handled. These Annual Reports detailed the frequency response, distortion and signal-to-noise performance of each station's audio system. Every piece of equipment had to be tested and calibrated, and forms detailing the technical performance of each piece of equipment lodged at the GPO. The system as a whole (including cabling) was then tested and specified. To add to this very unlikely picture, now and then men in lab coats would arrive unannounced and wheel in their own test equipment and verify these reported specs.

It's a different scene today, of course, not just in radio stations but also in recording studios, large and small, where even an understanding of equipment specifications and what they mean often runs thin (let alone regular testing and reporting to the authorities!). How times have changed!

Nowadays, from time to time, I get called to studios where the assistant engineer has decided the console needs recalibrating to 'work better' with the DAW, and where recording involves a constant wrestle with gain controls to avoid digital 'overs'. A simple question to a client about what reference level they'd like their VUs calibrated to quickly turns into a philosophical discussion that seems to involve everything but the answer to the question.

STANDARD PRACTICE?

There's no denying that a certain lack of technical understanding of equipment specification and calibration exists within the audio industry today. So this issue I thought I'd have a crack at outlining some of the common terms of technical specification and see if the waters are clearer or muddier by the time we get to the end of this article.

dB - DECIBELS

We love seeing dB written on equipment, and any gear bearing these two little letters is always somehow taken more seriously. It might be fine for guitar amps to have controls labelled '1' to '10' but our *professional* studio gear needs more

detail than that, right?

So what is a dB? Unlike apples or microphones, you can't pick one up. A dB is not a physical object, nor is it even a measurement of a physical property. While we can always say that 'cable X is 10 metres long' – and all around the planet we agree on what that means – we can't similarly assert that 'signal X is 10dB loud' (or long, or large, or whatever). Simply put, we can't come to any agreement on what that means. The fundamental thing about a dB is that it specifies the *relationship* (or ratio) of one measurement to another. It is not an *absolute measurement* in itself.


Since dB expresses the ratio of two measurements *to one another* it makes sense (and is quite valid) to label a gain control or an EQ boost/cut control with a dB scale. In this situation the dB reading is showing the ratio of the signal coming out the circuit associated with this control, to the signal going in.

Where we tend to come unstuck is when we talk about actual signal levels themselves. To use a dB scale as an absolute measurement of a signal level requires a known 'reference level' to which all further measurements are related or referenced. For audio signals, this reference level will be a specified voltage. With this standard reference voltage established we can now measure other audio signals and express them in decibels such that we universally agree on the measurement.

THE DECIBEL REFERENCE

A standard which has existed for a long time, that would be more than familiar to those old school radio station technicians, is dBm. Here the suffix 'm' has been added to define the standard voltage to which all measurements in dBm are referenced. For a measurement in dBm the 'reference voltage' is simply defined as the voltage required to deliver 1milliwatt (1/1000 of a Watt) of power into a 600Ω load. For a sine wave this is a voltage of 0.774 volts. From this we can say that 0dBm equals 0.774V into a 600Ω load, +20dBm equals 7.74V and -60dBm equals 0.774mV.

These days (as mentioned previously in my article on impedance in Issue 53) the standard audio impedance of 600Ω – which had been an industry standard for decades – is largely neglected today. Thus, the dB standard largely used now is the dBu, which is defined by the



reference level of 0.774V. Thus 0dBu equals 0.774V. Note that this is the same voltage given as the reference voltage for dBm, however, in the case of dBu there is no requirement to know the impedances of the equipment involved.

MEANWHILE BACK IN THE REAL WORLD

So how does all this talk about decibels and reference voltages affect your studio? Now that we know what a decibel is, let's take a look at those familiar specs we're always quoted and find out exactly what they mean. If you have any combination of audio equipment hooked together to form a 'studio', an understanding of these terms is important.

These days, studios large and small consist of a number of items of gear: analogue, digital, new and old. Understanding how they all come together to form your 'studio' – how each piece of equipment interacts with all the others – is vital for a healthy audio signal path and maximum signal-to-noise ratios. It's all about optimising the operating level of each piece of equipment, minimising noise and maximising headroom to achieve the best possible performance for your studio as a whole.

NOISE

A low noise level – or noise floor – is always desirable in your studio setup; so what's the best we can expect? Electronic components such as transistors, valves and resistors all generate noise, so there's no escaping it entirely. There are different types of noise, but under most situations thermal (or Johnson) noise is the most significant contributor. Thermal noise is generated by the agitation of electrons in any electrical conductor, resistor etc and is often called 'white noise', which has an equal

energy density in equal width frequency bands. The thermal noise in a 150Ω resistor across a bandwidth of 20Hz to 20kHz is 222nV or -130dBu . Thus, any microphone or line output with a typical output impedance of 150Ω will be generating noise at a level of -130dBu regardless of anything to do with its design or the components used.

GAIN

Gain, as we know, is the control we use to crank a signal up or down. Let's use the example of our ideal microphone, with its 150Ω output impedance, and apply 60dB of gain to get the microphone's output level up to a line level at around 0dBu. This brings the -130dBu noise floor up to -70dBu . This figure theoretically assumes that the electronics in the microphone or the microphone preamp contribute *nothing extra* to the noise performance. In real equipment, of course, the active electronics, transistors or valves *will* inevitably contribute to the noise of the device above the simple thermal noise of the circuit resistances. This is often expressed as the self-noise figure (NF) for microphones and equivalent input noise (EIN) for mic preamps.

A typical microphone may have a self noise figure of 15dB, meaning that the actual noise from the microphone will be (in *relative* terms) 15dB higher than the theoretical minimum of -130dBu , i.e., -115dB . So now with our 60dB of gain the noise floor will be -55dB . Typical noise figures (NF) for microphones range from around 5dB (a very low noise mic) to 25dB.


On top of this, the microphone preamp will also contribute noise. Today it's relatively easy to design and build low-noise preamps and a preamp with an EIN of -128dBu with a 150Ω microphone source connected is not uncommon. The figure of EIN = -128dBu tells us that the preamp contributes 2dB of noise above the noise generated by the 150Ω source resistance. So this brings our noise floor up to -53dB , at a minimum.

SIGNAL TO NOISE RATIOS


'Optimised' signal-to-noise ratios in a recording setup are achieved when the gain in the system is kept as low as possible while still achieving the desired recording level, and attenuators or faders are kept at their maximum. A console fader (real or virtual) doesn't change the signal-to-noise ratio, it only turns the signal up or down to the same extent as the noise. Thus the ratio of audio signal to noise remains the same. A gain control, however, increases the noise output of a preamp as the gain is increased.

HEADROOM

At the other end of the scale we have the maximum available output level of equipment. Primarily this is determined by the power supply voltages that each device employs. Many IC-based audio devices run power rails of +15 and -15 volts. In such a device, the maximum audio voltage that would typically be produced is 22dBu. One trick to get a higher maximum output level from a device running on $\pm 15\text{V}$ power rails is to provide a balanced output where one output amplifier is 'pushing' the signal and another amplifier is inverting this signal and 'pulling' it. This setup gives a further 6dB of maximum output level: 28dBu. This level is only preserved when coupled to a balanced input, if the output is unbalanced (where the 'cold pin' is connected to ground) then we lose this extra 6dB of headroom. An exception to this is if the output is coupled through a transformer. In this case we maintain the level regardless of the balanced or unbalanced connection. Also the use of 'step up' transformers allow the maximum output level to be increased beyond what may be expected based on the power supply voltages that the unit runs on. An example of this is the Neve 1073 module. While the electronics run off a single 24V power supply, by employing a 1:2 step up ratio on its output transformer, a maximum output level of +26dBu is achieved.



Healthy waveforms with a good balance of recording level and headroom.



Unhealthy waveforms recorded too 'hot' into the DAW, causing distortion in the analogue path and digital clipping.

DYNAMIC RANGE

A system's dynamic range is the ratio of signal levels between the noise floor and the maximum output level. Take our mic preamp with an EIN of -128dBu and a maximum output level of $+28\text{dBu}$. With 40dB of gain dialled up, the noise floor will be $(-128\text{dBu} + 40\text{dB})$ or -88dBu . But, of course, with a real microphone plugged in that has its own 15dB noise floor, it will now be 15dB worse, at -73dB . This setup has a dynamic range of 101dB , which coincidentally, is pretty close to the dynamic range of our hearing, from our perception of the faintest whisper to the point of sound becoming painful.

DIGITAL SYSTEMS

Relating all this information about voltage and headroom to the digital domain is, unfortunately, where it all goes haywire! A digital recording system essentially has two parts: the computer/software and the analogue-to-digital/digital-to-analogue converters. Within the computer software, absolute levels are meaningless with regard to the outside world. The only reference of importance is the full-scale digital code – when all digital bits are full and a larger signal simply cannot be represented by the system. This reference point is called 0dBFS and unlike an analogue system where a 0dBu signal has headroom above it, within the digital system 0dBFS has no headroom above it, only grief and hideous distortion. This is the starting point for much confusion when relating signals in the analogue domain to signals in the digital domain.

The relationship between our analogue 0dBu and our digital 0dBFS is arbitrary, and dictated solely by the converters and the options they provide (or don't provide) to set up that relationship. Since the introduction of digital recorders, various standards have been floated to define this relationship. Two that are widely accepted are $-18\text{dBFS} = +4\text{dBu}$ and $-14\text{dBFS} = +4\text{dBu}$. The reasoning behind these figures is that $+4\text{dBu}$ is well entrenched as the traditional 'operating point' – or reference level – for analogue equipment. In the system where the converters are calibrated for $-18\text{dBFS} = +4\text{dBu}$, we have a maximum analogue level that the converters could be expected to deal with of $+22\text{dBu}$ ($+4\text{dBu} + 18\text{dB} = 22\text{dBu}$). Note how this ties in nicely with the discussion above on headroom and devices running on typical power supplies of $\pm 15\text{V}$. For the system where $-14\text{dBFS} = +4\text{dBu}$, the maximum analogue level the converters could be expected to deal with would be $+18\text{dBu}$. By setting up these relationships between analogue and digital levels you optimise the operation of your equipment, avoiding digital overs while also running the analogue equipment in its 'sweet spot', where the average signal is around the 4dBu level. Many studios use the -18dBFS reference for recording, where audio is less processed and the dynamic range is wider. Meanwhile the -14dBFS reference is often used in mixdown and mastering where the audio is less dynamic and more controlled.

BIG WAVES ARE TROUBLE

One aspect of digital recording level setting is the tendency of engineers to use the waveform window as a level metering tool. This is fraught with problems and leads to recordings that are generally too hot, inclined to clip on unexpected peaks while tracking, and problems with setting levels in mixdown. The dynamic range between peak and average levels of unprocessed audio can be very wide – for percussive sounds: often in the order of 20dB or more; and even with vocals 12dB is not unusual. Take a 12dB range and display it in a waveform window and the average parts of the waveform will be one quarter the size of the peaks, filling only a quarter of the height of the window. What's worse, this assumes you have set the recording level so that your predicted peaks hit 0dBFS (which we're trying to avoid), and there's no headroom left. Give yourself a 6dB margin on those peaks to avoid digital overs and now the waveform is on average filling *one eighth* of the window... not much, eh!?

So always remember, conservative levels in a waveform window will have the average height of the waveform filling only a small fraction of the maximum height available. Big full waveforms are an accident waiting to happen.

Hopefully, the next time you read the specifications data for an audio unit (you do read this stuff, don't you?) you'll have a better understanding of the 'numbers' and be better able to work towards setting the optimum operating levels of your analogue and digital equipment. ■