CHAPTER 5

# Decibels For Dummies

## I. Introduction

This chapter summarizes the late 20th century approach to metering and leveling; it can be read as a preface to Chapter 15 in which we take these concepts into the 21st century. In the 20th century, because of their use of recording media with poor signal-to-noise ratios (SNR) engineers were often concerned with the signal peaks and with maintaining quality by maximizing the levels. With the advent of 24-bit recording, the SNR of our media is no longer an issue, but it is still crucially important for us to understand what the decibel scales on our meters are really telling us.

So many of us take our meters for granted—after all, recording is so simple: *all you do is peak to 0 dB and never go over!* But things only appear that simple until you discover one machine that says a recording peaks to -1 dB while another machine shows an OVER level, and yet your workstation tells you it just reaches 0 dB! We need to explore the concepts of the digital OVER, analog and digital headroom, machine meters, gain staging, loudness, signal-to-noise ratio and take a fresh look at the common practices of dubbing and level calibration.

## II. Digital Meters and OVER Indicators

Recorder manufacturers pack a lot in a little box, often compromising on meter design to cut production costs. A few machines even have meters which are driven from analog circuitry—a definite source of inaccuracy. Even manufacturers who drive their meters digitally (by the values of the sample numbers) cut costs by putting large gaps on the meter scale (avoiding expensive illuminated segments). The result is that there may be a -3 point and a 0 dB point, with a large unhelpful no man's land in between. The manufacturer may feel they're doing you a favor by making the meter read 0 if the actual level is between -1 and 0, but even if the meter has a segment at every decibel, when it comes to playback, the machine can't tell the difference between a level of 0 dBFS (*FS = Full Scale)* and an OVER. That's because once signal has been recorded, it cannot exceed full scale again, as illustrated below.

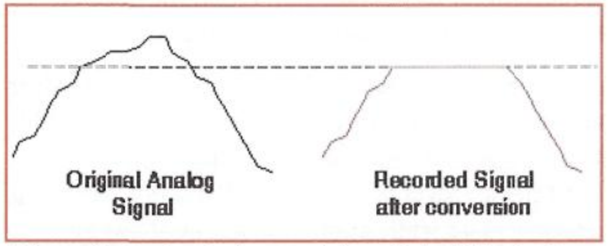


*While an original analog signal can exceed the amplitude of 0 dB, when that recording is reproduced, there will be no level above 0, yielding a distorted square wave. This diagram shows a positive-going signal, but the same is true on the negative-going end.*

One way a signal can go OVER is during recording from an analog source. An early-warning indicator is a level sensor in an A/D converter, driven by the analog portion of the signal, which causes the OVER indicator to illuminate if the analog level is greater than the voltage equivalent to 0 dBFS. If the analog record level is not reduced, then a maximum level of 0 dB will be recorded for the duration of the overload, producing a distorted square wave.

**MYTH:**
*The red light came on while I was recording, but when I played it back, there weren't any overs, so I thought it was OK.**

* Contributed by Lynn Fuston.

After the signal has been recorded, distinguishing between a full scale recording and one that actually went OVER requires more meter intelligence than I've ever seen on a typical machine or DAW. I would question the machine's manufacturer if the OVER indicator lights on playback; it's probably a simple 0 dB detector rather than an OVER indicator. There are more sophisticated, calibrated digital peak meters such as those from Dorrough, DK, Mytek, NTT, Pinguin, RTW, Sony, and others, each with unique features (including custom decay times and meter scales), but all the good meters agree on one thing: the definition of the highest measured digital audio level. A true digital audio meter reads the numeric code of the digital audio, and converts that to an accurate reading.[1]

### The Paradox of the Digital OVER

A well-designed digital audio meter can actually distinguish between 0 dBFS and an OVER. But if the digital levels on the medium cannot exceed 0 dB, how can the meter distinguish an OVER **after** the recording has been made? The answer is that a specialized digital meter determines an OVER by counting the number of samples in a row at 0 dB. The Sony 1630 OVER standard is three contiguous samples, because it's fair to assume that the analog audio level must have exceeded 0 dB somewhere between sample number one and three. Three samples is a conservative standard—most authorities consider distortion lasting only 33 **microseconds** (three samples at 44.1 kHz) to be inaudible. Depending on the nature of the music, distortion lasting as long as one or two **milliseconds** is likely inaudible. Thus, at higher sample rates, where many more samples go by in a short time, a case can be made to count many more contiguous full scale samples before warning the operator. Manufacturers of digital meters often provide a choice of setting the OVER threshold to 4, 5, or 6 contiguous samples, but it's better to err on the conservative side, to let the meter warn you before a problem could occur. If you stick with the 3-sample standard, you'll probably catch audible OVERs. But stand by, I'm about to recommend why you should mix at even lower peak levels!

### Using External A/D Converters or Processors

There is no standard for communicating OVERs on an AES/EBU or S/PDIF line. So if you're using an external A/D converter, the recorder's OVER indicator will probably not function properly, if at all. Some external A/D converters do not have OVER indicators, so in this case, there's no substitute for an accurate external meter; without one I would advise not exceeding -1 dB. I've already received several overloaded tapes which were traced to an external A/D converter that wasn't equipped with an overload indicator.

When making a digital dub through a digital processor you'll find that most do not have accurate metering. Equalizer or filter sections can cause OVERs even when dipping levels! Contrary to popular belief, an OVER can be generated even if a filter is set for attenuation instead of boost, because filters can ring; they also can change the peak level as the frequency balance is skewed. Digital processors can also overload internally in a fashion undetectable by a digital meter. Internal stages may "wrap around" when they overload, without transferring OVERs to the output. In those cases, a digital meter is not a foolproof OVER detector, and there's no substitute for the ear, but a good digital meter will catch most other transgressions. When you hear or detect an overload from a digital processor, try using the processor's digital input attenuator, or simply attenuate its output if you are sure the processor has sufficient internal headroom, explained later in this chapter.

### Oversampled Meters: Even More Sophisticated

Reading the simple numeric code from the digital stream may not be enough to detect OVERs in the converters that reproduce that signal. During the conversion from PCM digital to analog, built-in low-pass filtering causes occasional peaks **between the samples** that are higher than the digital stream's measured level, or even higher than full scale. **Digital designers have known for years that the actual output level of audio from a D/A converter can exceed 0 dBFS** but very few have taken this into account in the design. TC Electronic has performed tests on typical consumer D/A converters,[3] showing that many of them distort severely since their digital filters and analog output stages do not have the headroom to accommodate levels which exceed 0 dBFS! Besides D/As, certain processing elements of the signal chain can distort with intersample peaks, including sample rate converters and digital equalizers as we just explained. 0 dBFS+ peaks may reach as much as +3 dBFS with certain types of signals; what this means is that to make the cleanest recordings and to be perfectly safe, you should never exceed –3 dBFS on a simple (non-oversampling) digital meter! To demonstrate the problem and since this goes against typical *wisdom*, TC have developed an oversampling limiter and special oversampling peak meter in the System 6000.

### Practice Safe Levels

Although there have been no psychoacoustic studies on their adversity, intersample 0 dBFS+ peaks cause some following processing circuits to linger and extend the distortion, which makes post-processing and broadcasting seriously problematic.[4] And some critical listeners report improvements when measured intersample OVERs are eliminated. It makes sense for production engineers to **practice safe levels during recording and mixing** by staying well away from 0 dBFS on a standard peak meter and leaving the decision on whether and how to raise levels to the mastering suite, where we make an educated decision. Mastering engineers, if maximizing levels, should at least use an over-counting meter, plus a digital limiter whose ceiling is set to –0.2 dB (see Chapter 10)[2] but preferably an oversampling limiter and oversampled meter (to prevent downstream problems with DACs and radio processing). Clipping of any type is to be avoided especially if a recording is to undergo further processing, as demonstrated in Appendix 1.[5]

*"You would have to lower the peak level of a 24-bit recording by 48 dB to yield an effective 16-bit recording!"*

**MYTH OF THE MAGIC CLIP REMOVAL:** *Turn it down after clipping and the clip will go away.*

### The Myth of the Magic Clip Removal

If the level is turned down by as little as 0.1 dB, then a recording which may be full of OVERs will no longer measure any overs. But this does not get rid of the clipping or the distortion, it merely prevents it from triggering the meter. Some mastering engineers deliberately severely clip the signal, and then drop the level slightly, so that the meters will not show any OVERs. This practice, known as **SHRED**, produces very fatiguing (and potentially boringly similar) recordings.[6]

### Peak Level Practice for Good 24-bit Recording

Even though 24-bit recording is now the norm, some engineers retain the habit of trying to hit the top of the meters, which is totally unnecessary as illustrated at left. Note that a 16-bit recording fits entirely in the bottom 91 dB of the 24-bit. You would have to lower the peak level of a 24-bit recording by 48

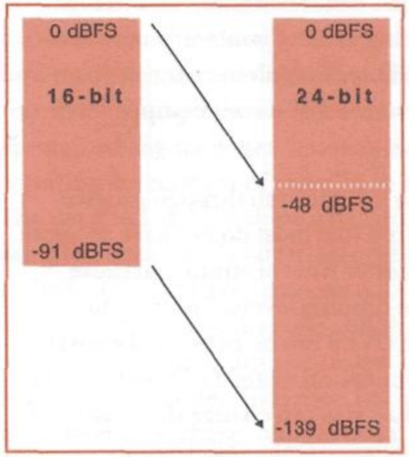


*A 24-bit recording would have to be lowered in level by 48 dB in order to reduce it to the SNR of 16-bit. The noise floors shown are with flat dither.*

dB to yield an effective 16-bit recording! So there is a lot of room at the bottom, and you won't lose any dynamic range if you peak to –3 dBFS or even as low as –10 dBFS; you'll end up with a cleaner recording. Distortion accumulates,[7] and at the mastering studio, a digital recording which is too hot can cause a digital EQ or sample rate converter to overload. A digital mix that peaks to –3 dBFS or lower makes it easier to equalize and otherwise process without needing an extra stage of attenuation in the mastering.

A number of 24-bit A/Ds advertise *additional headroom* by employing a built-in compressor at the top of the scale. As we have seen, there is no audible improvement in SNR by maximizing a 24-bit recording and no SNR advantage to compressing levels with a good 24-bit A/D.

#### How Loud is It?

Contrary to popular belief, the levels on a digital peak meter have (almost) nothing to do with loudness. For example, you're doing a direct to two-track recording (some engineers still work that way!) and you've found the perfect mix. Now, keep your hands off the faders, and let the musicians make a perfect take. During take one, the performance reached -4 dB on the meter; and in take two, it reached 0 dB for a brief moment during a snare drum hit. Does that mean that take two is louder? If you answered "both takes are about the same loudness," you're probably right, **because in general, the ear responds to average levels, not peak levels when judging loudness**. If you raise the master gain of take one by 4 dB so that it, too reaches 0 dBFS peak, it will now sound 4 dB louder than take two, even though they both now measure **the same** on the peak meter.

Do not confuse the peak-reading meters on digital recorders with VU meters. Besides having a different scale, a VU meter has a much slower attack time than a digital peak meter. In Chapter 15 we will discuss loudness in more detail, but we can summarize now by saying that the VU meter responds more closely to the response of the ear. For loudness judgment, if all you have is a peak meter, use your ears. If you have a VU, use it as a guide, not an absolute, because it is still fairly inaccurate.

Did you know that an analog tape and digital recording of the same source sound very different in terms of loudness? Make an analog tape recording and a digital recording of the same music. Dub the analog recording to digital, peaking at the same peak level as the digital recording. The analog dub will sound about 6 dB louder than the all-digital recording, which is quite a difference! This is because the peak-to-average ratio of an analog recording can be as much as 12-14 dB, compared with as much as 20 dB for an uncompressed digital recording. Analog tape's built-in compressor is a means of getting recordings to sound louder (oops, did I just reveal a secret?).[8] That's why pop producers who record digitally may have to compress or limit to compete with the loudness of their analog counterparts.

**MYTH:** *Normalization Makes the Song Levels Correct*

### The Myths of *Normalization*

**The Esthetic Myth:** Digital audio editing programs have a feature called **Normalization**, a semi-automatic method of adjusting levels. The engineer selects all the segments (songs), and the computer grinds away, searching for the highest peak on the album. Then the computer adjusts the level of all the material until the highest peak reaches 0 dBFS. **If all the material is group-normalized at once**, this is not a serious esthetic problem, as long as all the songs have been raised or lowered by the same amount. But it is also possible to select each song and *normalize* it individually, which is part of the esthetic mythology—it's a real no-no. If you're making an album, never normalize individual songs, since the ear responds to average levels, and normalization measures peak levels, the result can totally distort musical values. A compressed ballad will end up louder than a rock piece! In short, **normalization should not be used to regulate song levels in an album.** There's no substitute for the human ear, and currently there is no artificial intelligence that does as well.*

**The Technical Myth:** It's also a myth that normalization improves sound quality of a recording; in fact, it can only degrade it. Technically speaking, normalization only adds one more degrading calculation and resulting quantization distortion. And since the material has already been mixed, it has already been quantized, which predetermines its signal to noise ratio—SNR of the recording cannot be further improved by raising it. Let me repeat: Raising the level of the material will not change its inherent signal to noise ratio but will only add more quantization distortion in an unnecessary step. **If the material is going to be mastered, do not normalize** since the mastering engineer will be performing further processing anyway.[9]

### Judging Loudness the Right Way

Since the ear is the only judge of loudness, is there any objective way to determine how loud your CD will sound? The first key is to use a single D/A converter to reproduce all your digital sources and maintain a fixed setting on your monitor gain. That way you can compare your *CD in the making* against other CDs, in the digital domain. Judge DATs, CDs, workstations, and digital processors through this single converter.

## III. Calibrating Studio Levels: Headroom and Cushion

Protecting your A/D and mix from clipping does no good if your analog console, preamplifiers or processors are *distorting in front of the A/D!* Since mastering engineers usually chain multiple pieces of gear, it's important to understand how to optimize analog levels, distortion and noise when making signal chains in front of your A/D converter. Ostensibly, typical balanced analog gear has a *nominal* level of +4 dBu (reference .775 volts[10], yielding 1.23 volts with sinewave. Unfortunately however, not all analog gear is created equal, and +4 dBu may be a bad choice of reference level. I use the term *nominal* to mean the *average* voltage level that corresponds with 0 VU, typically 20 dB below full scale digital (0 dBFS). We need to examine some

* When a client asks me if I *normalize* I reply that I never use the computer's automatic *normalization* method, but rather songs are leveled by ear. I avoid the term normalization because it has been misused.

easily overlooked factors when deciding on an in-house standard analog (voltage) level.

One factor is the clipping point of consoles and outboard gear. Before the advent of inexpensive 8-buss consoles, most professional consoles' clipping points were +24 dBu or higher. But a frequent compromise in low-priced console design is to use internal circuits that clip earlier, around +20 dBu (7.75 volts). This can be a big impediment to clean audio, especially when cascading amplifiers. To avoid the *solid-state edginess* that plagues a lot of modern equipment, the *minimum* clip level of every amplifier in a system should be 6 dB above the potential peak level of the music. The reason: Many opamps and other solid state circuits exhibit an extreme distortion increase long before they reach the actual clipping point, as they change from class A to class AB operation. This means clipping point should be at least +30 dBu (24.5 volts RMS) if 0 VU is +4 dBu!

### You Can Never Have Enough Headroom!

A lot of solid-state designs start to sound pretty nasty when used near their clip point.[11] All other things being equal, the amplifier with the higher clipping point will sound better. Perhaps that's why tube equipment (with its 300 volt B+ supplies and headroom 30 dB or greater) often has a *good* name and solid state equipment with inadequate power supplies or headroom has a *bad* name. Most of the *robust-sounding* solid-state equipment I know uses very high power (but very expensive) supply rails.

Traditionally, the difference between average level and clip point has been called the *headroom*, but in order to emphasize the need for even more than the traditional amount of headroom, I'll call the space between the peak level of the music and the amplifier clip point a *cushion*. With analog tape, a 0 VU reference of +4 dBu with a clipping point of +20 dBu provided reasonable amplifier headroom, because musical peak-to-average ratios were reduced to the compression point of the tape, which maxes out at around 14 dB over 0 VU. Instead of clipping, analog tape's gradual saturation curve produces 3rd and 2nd harmonics, much gentler on the ear than the higher order distortions of solid state amplifier clipping.

But it's a different story when the peak-to-average ratio of raw, unprocessed digital audio tracks can be 20 dB. Adding 20 dB to a reference of +4 dBu results in +24 dBu, which is beyond the clipping point of many so-called *professional* pieces of gear, and so doesn't leave any room at all for a *cushion*. If you adapt an active balanced output to an unbalanced input, the clipping point reduces by 6 dB, so the situation becomes proportionally worse.[12] Dual-output consoles that are designed to work at either professional or semi-pro levels can be particularly problematic. To meet price goals, manufacturers often compromise on headroom in professional mode, making the so-called semi-pro mode sound cleaner! It is an unpleasant surprise to discover that many consoles clip at +20 dBu, meaning they should not be using a professional reference level of +4 dBu (headroom of only 16 dB and no cushion). Even if the console clips at +30 dBu (the minimum clipping point I recommend), that only leaves a 6 dB cushion when reproducing

**MYTH:**
*+4 dBu is always the best level to use for 0 VU with balanced analog electronics.*

music with 20 dB peak-to-average ratio. That's why more and more high-end professional equipment have clipping points as high as +37 dBu (55 volts!). To obtain that specification, an amplifier must use very high output devices and high-voltage power supplies. Translation—better sound (all other things being equal), and also higher cost due to the need for more robust power supplies and devices.

These robust output drivers that have this kind of headroom sound better if they can deliver a clean high level into a 600 ohm load, which means they can probably handle long cable runs with their high capacitive loads. Long runs should probably be balanced, but since many mastering studios have small ground-loop areas, we often use custom-made unbalanced equipment, which often has simpler, quieter circuitry.

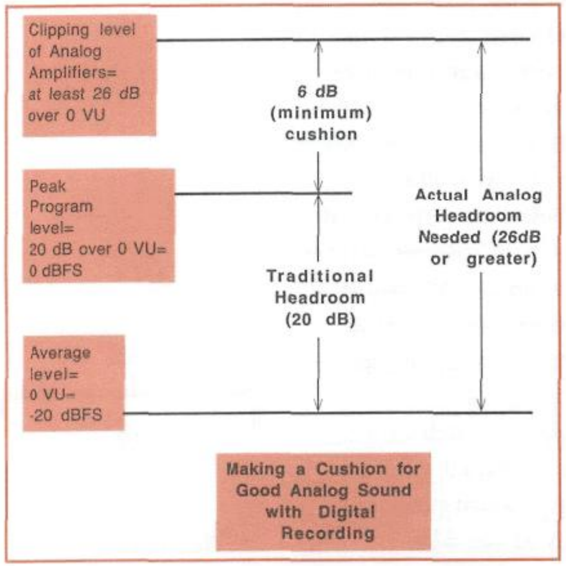


One of the most common mistakes made by digital equipment manufacturers is to assume that, if the digital signal *clips* at 0 dBFS, then it's OK to install a (cheap) analog output stage that would clip at a voltage equivalent to, say, 1 dB higher. This almost guarantees a nasty-sounding converter or recorder, because of the lack of cushion in its analog output section and the potential for 0 dBFS+ levels.

How can we increase the cushion in our system, short of replacing all our distribution amplifiers and consoles with new ones? One way to solve the problem is to recalibrate all the VU meters. SNR will not be significantly lost if we set 0 VU= 0 dBu or even -4 dBu (not an international standard, but a decent compromise if we don't want to throw out equipment), and things will sound cleaner in the studio. Once we've decided on a standard analog reference level, we calibrate all analog-driven VU meters to this level. At left is a diagram describing the concept of *cushion*.

## IV. Gain Staging—Analog and Digital

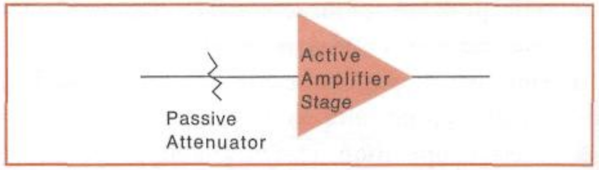


*In the top device, signal enters a passive attenuator and exits through an active amplifier stage. This circuit effectively has infinite input headroom. The bottom device's input headroom is determined by the headroom of the input amplifier.*

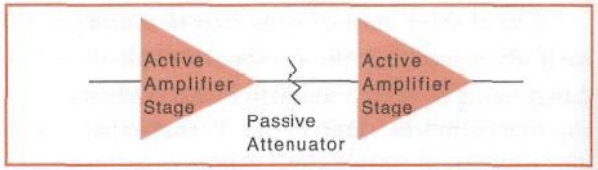


### Analog Signal Chains

Now that we know how to choose an analog level, it's time to chain our equipment together. To really get a handle on our equipment, we should determine its internal structure. The above figures represent two possible internal structures. All structures are variations on this theme.

To properly test analog devices and determine their internal makeup, use a good clean monitor system, an oscilloscope, a digital voltmeter and a sine wave generator that can deliver a clean +24 dBu or higher (a tough requirement in itself). The first type of device has a passive attenuator on its input, which means that we can feed it any reasonable source signal without fear of overload. We can prove this by turning the generator up and attenuator down; if the output never clips within a reasonable range of the generator, then the device must have a passive attenuator on its input. Then, we disconnect the generator and listen to the output of the device as we raise and lower the attenuator. There should be no change in noise or hiss, and the output noise should be well below –70 dBu unweighted, preferably below –90 dBu A-Weighted. This also is an indication that the device has a passive attenuator on its input. If the output noise changes significantly at intermediate positions of the attenuator, then the internal impedances of the circuit are in question, or there may be some DC offset. The output noise of this device will be limited by the noise floor of its output amplifier. We determine the best **nominal operating level** of this device by taking the output clip point and subtract at least 26 dB for headroom and cushion.

The second type of device's input is an active amplifier stage, whose design is much more critical. It is very rare to find a solid state device built this way which that won't clip with >+24 dBu input. While raising the signal generator, turn down the attenuator to keep the output from overloading. If we hear clipping prior to the generator reaching +24 dBu, then the device has a weak internal signal path. The clip point determines the nominal analog input level, which should be at least 26 dB below this clip point. Then, to check if the device's internal gain structure is well balanced, we see if the output stage clips at the same point as the input stage or at a higher level.

When cascading analog gear, the signal-to-noise ratio and headroom of the cascade is determined by the weakest link, but by studying the internal structure of each piece, it may be possible to increase SNR of the chain by running higher levels at points in the chain that have higher clipping levels. With test tone and then music, listen closely to the noise floor and high level sound quality at the last device in the chain; if the output of the chain sounds good and reasonably quiet, then I don't worry about tweaking the chain. I was able to improve the signal to noise ratio of a tube-based tape recorder whose gain structure resembles the second device. The original manufacturer's conservative schematic specified nominal internal levels of –10 dBu at the output of the second active stage. But since the tubes distort at well above +30 dBu (headroom of 40 dB), I decided to run the attenuator higher and run levels of 0 dBu in the second stage. This improved amplifier signal to noise ratio from the second stage on, by 10 dB, without endangering distortion. The tube tape recorder still has 30 dB of internal headroom.

In an analog signal chain, raising the music signal level as high as practical as early as possible (within the limits imposed by headroom and clipping point of A/D converters) will improve the signal to noise ratio of the entire chain. Then, later in the mastering, we will reduce the signal level digitally in the digital chain that follows.

### Digital Signal Chains

**Headroom of the Chain:** It's a lot harder to grasp what's going on inside a digital signal chain, but we can test digital performance for headroom, clipping, and noise. Suppose we have a digital equalizer with several gain controls and equalization; we feed it a 1 kHz sine wave test tone at about −6 dBFS and turn up the 1 kHz equalization by 10 dB, observing that the output clips. Then we turn down the output gain control until the output is below 0 dBFS and verify by listening or FFT measurements that the internal clipping goes away. If not, then the internal gain structure of the equalizer does not have enough headroom to handle wide range inputs. We may be able to get away with turning down an input attenuator, but the early clipping indicates that this equalizer is not state-of-the-art. It is probably a first-generation fixed point unit and should be replaced. Modern-day digital processors have enough internal headroom to sustain considerable boost in early stages without needing an input attenuator, and clipping can be removed solely by turning down the output attenuator. The internal structure could be double-precision fixed point or floating point (see Glossary, Appendix 13); it's not easy to tell without asking the manufacturer. It is easy to be impressed by floating-point manufacturers' claims of hundreds of dB of headroom above 0 dBFS, but 24 dB or so internal headroom above 0 dBFS is probably enough; most well-designed fixed-point products have 24 or more dB internal headroom.

**Distortion of the Chain and Individual Processor Levels: With a digital chain, we no longer have to consider the audio signal level between the various items of equipment**; raising the source signal in a 24-bit digital signal chain does not make a meaningful SNR difference, considering the inaudible (approximately -139 dBFS) noise of the chain.* No longer should we get hung up on having a low signal level; instead, consider every calculation as a source of quantization distortion. Instead of *optimizing levels*, what matters most in a 24-bit digital chain is to reduce the number of total calculations; give the job of gain changes and other calculations to the components with the highest internal resolution (e.g., those which would introduce the least quantization distortion or *grunge*). In fact, we should avoid raising the signal until it reaches a device with the cleanest-sounding gain control, even if the source audio level is very low. For example, if the workstation has lower resolution, we try to hold everything at unity gain in the DAW and reserve the gain changes or EQ for higher-precision devices later in the signal chain. In other words, pass a perfect clone (bit-transparent copy) of the source from the DAW onto the next device in line to do processing.

**Noise of the Chain:** The only significant noise floors in a 24-bit chain are not from the chain itself but from the original sources, including mike preamp noise. We are primarily concerned with the

* Each processor does add its own quiescent or idle noise, which is cumulative, but in a good chain rarely adds more than 3 to 6 dB to the −139 dBFS RMS noise floor.

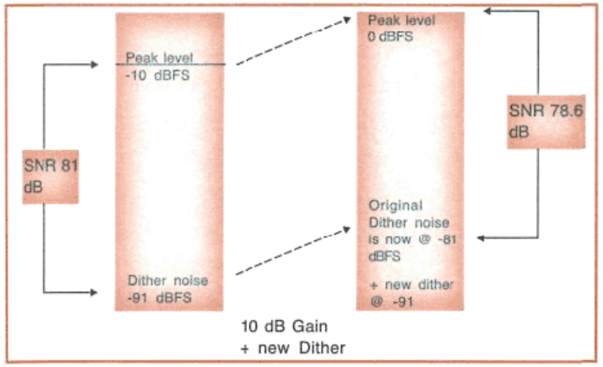


*A 16-bit recording with peak level low at –10 dBFS. When gain is raised 10 dB and redither is added, the original 81 dB signal to noise ratio is reduced by about 2.4 dB.*

impact of the summing of the higher level noises, and summing a new 16-bit dither with the source's dither noise can add a veil if the original was 16-bit.

Let's take an example of a 16-bit recording whose peak level is 10 dB low, as in the above figure. In mastering we may choose to raise its level by 10 dB and add 16-bit dither before turning it into a 16-bit CDR. This 16-bit recording's original 81 dB SNR is the difference between signal at –10 dBFS and dither noise at –91 dBFS.[13] When we raise the signal by 10 dB, both the original signal and the noise are raised equally, so the original signal to noise ratio is almost unchanged. However, the total SNR is the sum of the original dither which is now at –81 dBFS and the new dither which is at –91 dBFS. We ignore the insignificant noise of the gain processing, well below –130 dBFS, so the total is –78.6 dBFS, and the SNR of the source has been deteriorated by (81-78.6) or 2.4 dB. The more gain we apply to the source, the more distant the old noise will be above the added dither noise, and the smaller the new dither will seem when the two noises are summed. So, reconsider doing anything if you have to raise a signal by only a few dB, because the new dither will be very close to the old; if we perform no gain change and just add dither, the noise floor is raised by 3 dB. If we lower the gain, the new dither predominates over the old. Despite this degradation, many times we have to live with compromises in mastering, since we still receive 16-bit sources; and we are forced to adjust the level according to the esthetics of the album. I've had considerable luck reducing cumulative sonic veiling by using noise-shaped dither.[14]

The manufacturers of the Waves L2 claim that peak limiting allows raising level enough to be significantly above the dither noise, and thus increases the signal-to-dither ratio and resolution. But exercise caution, because to my ears the apparent noise improvement is more than offset by the degradation of sound quality (the limiter reduces transient clarity).

If we could avoid 16-bit dither, by producing an output at 24-bit that the consumer could use, then mastering processing and gain-changing can be performed with no significant penalty, with noise floor 48 dB below the noise of 16-bit. This is the promise of delivering higher wordlengths to the consumer and another reason to record in 24-bit in the first place.