Level Practices (Part 1)

Part I: The 20th Century
Dealing With Peaks

Overs, levels, and headroom, how to get the most from your equipment

Digital recording is simple--all you do is peak to 0 dB and never go over! And things remain that simple until you discover one plugin or processor telling you a signal peaks to -1 dB while another meter (e.g. in your DAW) shows an OVER level, yet your digital outboard processor tells you it just reaches 0 dB! This article will explore concepts of the digital OVER, machine meters, loudness, and take a fresh look at the common practices of dubbing and level calibration.

Section I: Digital Meters and OVER Indicators
Manufacturers often have to pack a lot in their product, therefore compromising on meter design and accuracy to cut costs. A few outboard machines' meters 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 costly illuminated segments), using inaccurate calculations and/or time constants or by just not translating the values right to the visible meter. As a result, there may be a -3 point and a 0 dB point, with a big no man's land in between and the values not being represantative for the signals momentary peak-level. The manufacturer may feel he's doing you a favor by making the meter read 0 even if the actual level is between -1 and 0, or by setting the threshhold of the OVER indicator inaccurately or too conservatively (longbefore an OVER actually occurs). Even if the meter has a segment at every decibel, on playback, the plugin or DAW may not be able to tell the difference between a level of 0 dBFS (*FS = Full Scale)* and an OVER. 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's only one way around this problem. Get a calibrated digital meter. Every studio should have one or two. There are lots of choices, from Dorrough, DK, Mytek, NTT, Pinguin, 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. A good digital audio meter can also distinguish between 0 dBFS and an OVER.

The Paradox of the Digital OVER
If digital levels cannot exceed 0 dB (by definition, there's nothing higher), then how can a digital signal go OVER? One way a signal can go OVER is during recording from an analog source. Of course the digitally encoded level cannot exceed 0 dBFS, but a level sensor in an A/D converter causes the OVER indicator to illuminate if the analog level is greater than the voltage equivalent to 0 dBFS. If the recordist does not reduce the analog record level, then a maximum level of 0 dB will be recorded for the duration of the overload, producing a nicely distorted square wave. There is a simple (digital) way of detecting if an OVER had occurred, even on playback--by looking for consecutive samples at 0 dB, which is a square wave. A specialized digital meter determines an OVER by counting thenumber of samples in a row at 0 dB. The Sony 1630 OVER standard is three 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 very conservative standard--most authorities consider distortion lasting only 33 microseconds (three samples at 44.1 KHz) to be inaudible. Manufacturers of digital meters often provide a choice of setting the OVER threshold to 4, 5, or 6 contiguous samples, but in this case it's better to be conservative. Even 6 samples is hard to hear on many types of music, so if you stick with the 3-sample standard, you'll guarantee that virtually all audible OVERs will be nipped in the bud, or at least detected! Once you've used a good digital meter, you'll never

want to go back to the built-in kind.
In the diagram below, a positive-going analog signal goes OVER in the area above the dotted line.

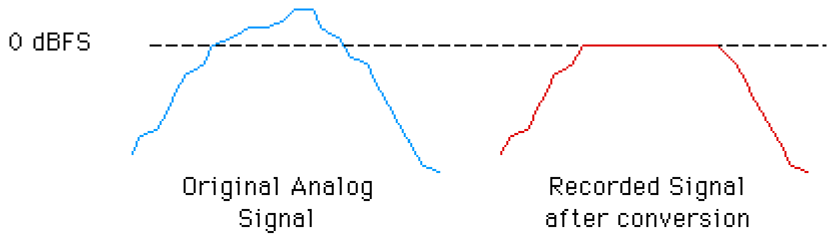

Using External A/D Converters or Processors
There is no standard for communicating OVERs on an AES/EBU or S/PDIF line. So if you're using an external A/D converter and feed the signal into any machine, the OVER indicator there will probably not function properly or at all. I advise ignoring the indicator if it does light up, unless the manufacturer confirms that it's a sample counting OVER indicator. They'll probably reveal that it's an analog-driven level detector. 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 on the feeded machine.

When making a digital dub through a digital processor you'll find most do not have accurate metering (be sure to read The Secrets of Dither before using any digital processor). Equalizer or processor sections can cause OVERs. Contrary to popular belief, an OVER can be generated even if a filter is set for attenuation instead of boost, because filters can ring. Digital processors can also overload internally in a fashion undetectable by a digital meter. Cascaded 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.

Practice Safe Levels
When recording to digital from an analog source, if you have an external digital meter set to 3 samples, then trust its OVER indicator and reduce gain slightly if it illuminates during recording. If you've been watching your levels prior to generating the OVER, chances are it will be an inaudible 3 sample OVER. You won't lose any meaningful signal-to-noise ratio, and you'll end up with a cleaner recording, especially when sending it for mastering. At the mastering studio, a file which is too hot can cause a digital EQ or sample rate converter to overload. There are ways around that, but not without complicating the mastering engineer's life.

Section II: 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, watch the levels to make sure they don't overload, 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, 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 PART II, we will discuss loudness in more detail, but let's summarize 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 the meter can be fooled (see PART II).

Did you know that an analog and digital recording of the same source sound very different in terms of loudness? Make an analog recording and a digital recording of the same music. Dub the analog recording to the digital domain, peaking at 0 dB. The analog dub will sound about 6 dB louder than the all-digital recording! That's a lot. This is because the typical peak-to-average ratio of an analog recording is about 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?). That's why pop producers who record digitally may have to compress or limit to compete with the loudness of their analog counterparts.

The Myth of "Normalization"

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. This is not a serious problem esthetically, 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. 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.

Judging Loudness the Right Way

Since the ear is the only judge of loudness, is there any objective way to get a handle on how loud your CD will sound? The first key is to use a single D/A converter to reproduce all your digital sources. That way you can compare your *CD in the making* against other CDs, in the digital domain. Judge plugins, CDs, workstations, and digital processors through this single converter. Another important tool is a calibrated monitor level control with 1 dB per step settings. In a consistent monitoring environment, you can become familiar with the level settings of the monitor control for many genres of music, and immediately know how far you are (in dB) from your nearest competitor, just by looking at the setting of the monitor knob. At *Digital Domain,* we log all monitor settings used on a given project, so we can return to the same setting for revisions. In PART II, we will discuss how to use our knowledge to make a better system in the 21st Century.

The Moving Average Goes Up and Up...

Some of the latest digital processors permit making louder-sounding recordings than ever before. Today's mastering tools could make a nuclear bomb out of yesterday's firecrackers. But the sound becomes squashed, distorted and usually uninteresting. Visit my article on Compression for a more detailed description of the loudness race. While it seems the macho thing to do, you don't have to make your CD louder than the loudest current CD; try to make it sound *better*, which is much harder to do.

Section III: Calibrating Studio Levels

That concludes our production discussion. This next section is intended primarily for the maintenance engineer. Let's talk about alignment of studio audio levels. Stick around for a fresh perspective on level setting in the hybrid analog-digital studio.

Marking Tapes

dBm and dBv do not travel from house to house. These are measurements of voltages expressed in decibels. I once received a 1/4" tape in the mail marked "the level is +4 dBm." +4 dBm is a voltage (it's 1.23 volts, although the "m" stands for milliwatts). The 1/4" tape has no voltage on it, it doesn't have any idea whether it was made with a semi-pro level of 0 VU = -10 dBv or a professional level of +4. Voltages don't travel from house to house, only *nanowebers per meter* on analog tapes, and *dBFS* on digital tapes. That doesn't diminish the importance of the analog reference level you use in-house. It's just irrelevant to the recipient of the tape. Just indicate the magnetic flux level which was used to coordinate with 0 VU. For example, *0 VU=400 nW/m at 1 KHz*. Most alignment tapes have tables of common flux levels, where you'll find that 400 nW/M is 6 dB over 200 nW/m. Engineers often abbreviate this on the tape box as *+6dB/200*.

Deciding On an In-House Analog (voltage) Level

Just use the level provided by your console manufacturer, right? Well, maybe not. +4 dBv (reference .775 volts) may be a bad choice of reference level. Let's examine some factors you may not have considered when deciding on an in-house standard analog (voltage) level. When was the last time you checked the clipping point of your console and outboard gear? Before the advent of inexpensive 8-buss consoles, most professional consoles' clipping points were +24 dBv or higher. A frequent compromise in low-priced console design is to use internal circuits that clip around +20 dBv (7.75 volts). This can be a big impediment to clean audio, especially when cascading stages (how many of those amplifiers are between your source and your multitrack?). In my opinion, to avoid the "solid state edginess" that plagues a lot of modern equipment, the *minimum* clip level of every amplifier in your 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. This means at least +30 dBv (24.5 volts RMS) if 0 VU is+4 dBv.

How Much Headroom is Enough?

Have you noticed that solid-state equipment starts to sound pretty nasty when used near its clip point? All other things being equal, the amplifier with the higher clipping point sounds better, in my opinion. Perhaps that's why tube equipment (with their 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.

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*. In the days of analog tape, a 0 VU reference of +4 dBv with a clipping point of +20 dBv 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 today, where the peak-to-average ratio of raw, unprocessed digital audio tracks can be 20 dB. Adding 20 dB to a reference of +4 dBv results in +24 dBv, which is beyond the clipping point of many so-called *professional* pieces of gear, and doesn't leave any room 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 (all those headroom specs have to be reduced by 6 dB if you unbalance an amplifier's output). Be particularly suspicious of consoles that are designed to work at either professional or semi-pro levels. To meet price goals, manufacturers often compromise on headroom in professional mode, making the so-called semi-pro mode sound cleaner! You'll be unpleasantly surprised to discover that many consoles clip at +20 dBv, meaning they should never be using a professional reference level of +4 dBv (headroom of only 16 dB and no cushion). Even if the console clips at +30 dBv (the minimum clipping point I recommend), that only leaves a 6 dB cushion when reproducing 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 dBv (55 volts!). To obtain that specification, an amplifier must use very high output devices and high-voltage power supplies. Translation--better sound.

To summarize, make sure the clip point of all your analog amplifiers is at least 6 dB (preferably 12 or more dB) above the peak level of analog material that will run in the system. I call this additional headroom *the cushion*.

How can you increase the cushion in your system, short of junking all your distribution amplifiers

and consoles for new ones? One way to solve the problem is to recalibrate all your VU meters. You will not lose significant signal-to-noise ratio if you set 0 VU= 0 dBv or even -4 dBv (not an international standard, but a decent compromise if you don't want to throw out your equipment, and you have the expertise to make this standard stick throughout your studio). Try it and let me know if things sound cleaner in your studio.

Once you've decided on a standard analog reference level, calibrate all your analog-driven VU meters to this level. Here's a diagram describing the concept of *cushion.*

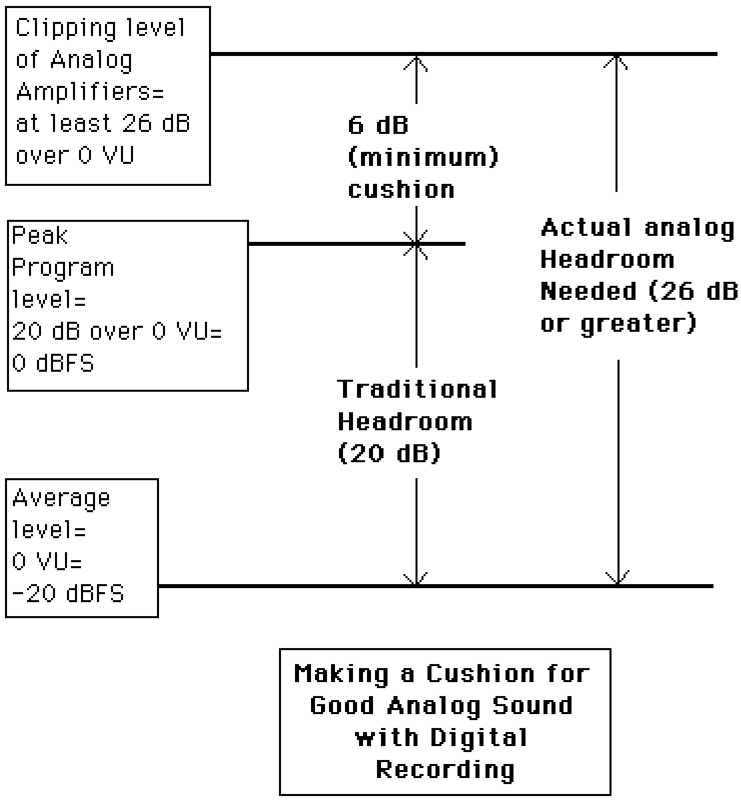

Dubbing and Copying - Translating between analog and digital points in the system

Let's discuss the interfacing of analog devices equipped with VU meters and digital devices equipped with digital (peak) meters. When you calibrate a system with sine wave tone, what translation level should you use? There are several de facto standards. Common choices have been -20 dBFS, -18 dBFS, and -14 dBFS translating to 0 VU. I'd like to see accurate calibration marks in digital recorders and DAWs at -12, -14, -18, and -20 dB, which covers most bases. Most of the external digital meters provide means to accurately calibrate at any of these levels.

How do you decide which standard to use? Is it possible to have only one standard? What are the compromises of each?

To make an educated decision, ask yourself: What is my system philosophy?

- Am I interested in maintaining headroom and avoiding peak clipping or do I want the highest possible signal-to-noise ratio at all times?
- Do I need to simplify dubbing practices or am I willing to require constant supervision during dubbing (operator checks levels before each dub, finds the peaks, and so on)?
- Am I adjusting levels or processing dynamics--mastering for loudness and consistency with only

secondary regard for the peak level?

Consider your typical musical sources. Are your sources totally digital (DDD)? Did they pass through extreme processing (compression) or through analog tape stages? Pure, unprocessed digital sources, particularly individual tracks on a multitrack, will have peak levels 18 to 20 dB above 0 VU. Whereas processed mixdowns will have peak-to-average ratios of up to 18 dB (rarely up to 20). Analog tapes will have peak levels up to 14 dB, almost never greater. And that's how the three most common choices of translation numbers (-18,-20, and -14) were derived.

Broadcast Studios

In Broadcast, *Practicality* is our object, simplifying day-to-day operation, especially if your consoles are equipped with VU meters and your recorders are digital. In broadcast studios, it is desirable to use fixed, calibrated input and output gains on all equipment. My personal recommendation for the vast majority of studios is to standardize on reference levels of -20 dBFS ~0 VU, particularly when mixing to 2-track digital from live sources or tracking live to multitrack digital. If you're watching the console's VU meters, you will probably never clip a digital tape if you use -20 dBFS as a reference.
For a busy recording studio that does most of its mixing, recording and dubbing to harddisc, standardizing on -20 dBFS will simplify the process. Recording studios who decide on -18 dBFS ~0 VU will run into occasional digital clipping. That's why I'm against -18 dBFS as a standard for recording studios using VU meters for recording.

If you standardize on a -20 dBFS reference, the more compressed your musical material, the more signal-to-noise ratio you seem to be throwing away, but this is not true. If your source is analog tape, you might throw away 6 or more dB of signal, but this is less important than maintaining the convenience of never having to adjust dubbing levels on equipment. Furthermore, the ear judges noise level by average levels, and if the crest factor of your material is 6 dB less, it will seem just as loud as the uncompressed material peaking to 0 dBFS, you will not have to turn up your monitor, and you will not hear additional noise. Remember: analog tapes typically sound 6 dB louder than digital tapes, if peaked to the same peak level.

A -20 reference is only a potential problem when dubbing from digital source to analog tape. In many cases, you can accept the innocuous 6 dB compression. We've been enjoying that for years when we mixed from live material on VU-equipped console direct to analog tape. When making dubs to analog for archival purposes, choose a tape with more headroom, or use a custom reference point (-14 to -18 dBFS), as the goal is to preserve transients for the enjoyment of future listeners. A calibrated peak level meter on the analog machine will tell you what it's doing more than a VU meter. For archival purposes, I prefer to use the headroom of the new high-output tapes for transient clarity, rather than to jack up the flux level for a better signal-to-hiss ratio.

If working in a broadcast facility which seems no live (uncompressed) material, then for the broadcast dubbing room, -14 is a good number (dubbing between analog and digital tapes). -18 is a safe all-around reference for all the other A/D/A converters in the broadcast complex, since most of the material will have 18 dB or lower peak-to average ratio, and occasional clipping maybe tolerated.

Mastering Studios

Mastering studios are working more frequently in 20-bit or 24-bit. In Part II, I suggest the 21st Century approach to mastering.

Analog PPMs

Analog PPMs have a slower attack time than digital PPMs. When working with a digital recorder, a live source, and desk equipped with analog PPM, I suggest a 5 dB "lead." In other words, align the highest peak level on the analog PPM to -5 dBFS with sine wave tone.