

## Meter Madness

Mike Rivers

In May 1999, the classic VU meter celebrated its 60th birthday. It served the industry well, and when properly interpreted, can still be a useful tool to indicate loudness. Today's digital recording processes have brought us around to taking a hard look at the inadequacies of both the traditional VU meter and the peak meters for level management.



The classic VU meter has long been the common audio level indicator. VU stands for Volume Unit, and the meter indicates how loud something sounds. The traditional VU meter is analog, with a logarithmic (decibel) scale that has zero marked about 2/3 of the way up scale (this is what we mean when we say the meter “reads zero”, not where the pointer rests when the equipment is turned off), with high resolution (generally 1 division per dB) between 0 and +3, and progressively lower resolution (more dB per scale division) going downscale from zero.

These days, the mechanical meter is often replaced by a column of LED's, with individual lights taking the place of the meter scale ticks. On a computer-based audio workstation, often a meter is simply drawn on the monitor screen. The Cool Look factor is lower, but it's a lot cheaper than using mechanical meters.

In this article we'll review the characteristics of the traditional analog VU meter and also focus audio level measurement and management in the digital domain, the concept of headroom, and how loudness and audio level are related (and how they're not).

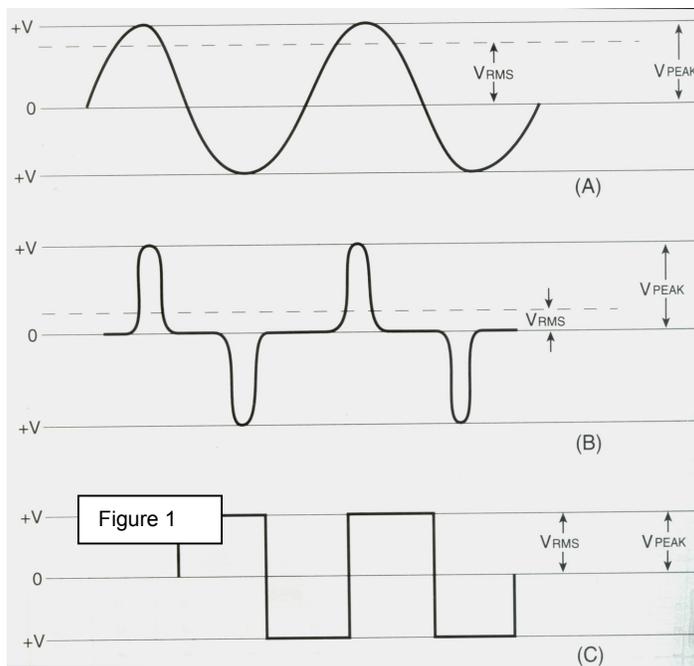
### What's a VU?

The Volume Unit meter was originally designed to help broadcast engineers keep the overall program level consistent between speech and music. There's a well defined standard for the mechanical rise-and-fall response characteristics of the pointer (to which few of today's VU-looking meters comply). A standard VU meter actually responds a little too fast to accurately represent musical loudness but fast enough to show movement between syllables. Unscientific as it sounds, the response of the VU meter was tailored so that the pointer movement looks good when it's indicating speech. It was easy to tell at a glance whether speech or music was going out over the air, whether it was about at the right level, and when something wasn't working. Engineers learned that brief excursions up to the +3 dB top end of the scale rarely caused distortion in the analog equipment of the day, nor did they sound too loud.

## How is Audio Level Measured?

Once sound is converted to electricity, we can represent the audio level by measuring the alternating electrical voltage. A symmetrical signal such as a sine wave - a single-pitched note with no overtones or distortion - spends equal time on the positive and negative sides of 0 volts. The numerical average of the positive and negative voltages over many cycles is zero - not a very useful measurement when we want to know how loud a sound that voltage represents.

A measurement that corresponds closely to how we perceive loudness is called the RMS (root mean square) average. The RMS value is calculated by squaring the voltage at each point on the waveform (remember, the square of a negative number is a positive number), summing the squares, dividing by the number of points measured (the more points, the more accurate the calculation), then taking the square root of that quotient. For a sine wave, the RMS value is 0.707 times the peak value. (Figure 1a). You've probably heard of this numeric relationship



(it's the square root of 2) between peak and RMS before. Understand, though, that this numerical relationship holds true only for a pure sine wave.

Look at the string of narrow pulses in Figure 1b. The voltage spends much more time at zero than at some positive or negative value. Even though the peak value is the same as for the sine wave above it, the RMS value of this waveform is lower. It will read lower on a VU meter, and would sound quieter than the sine wave.

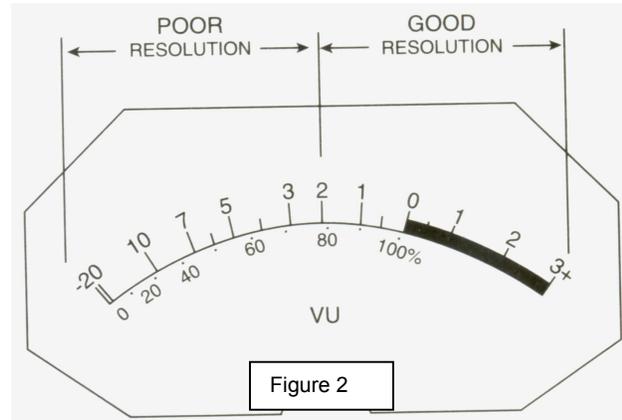
Since a square wave (Figure 1c) spends all its time either at the peak positive or peak negative voltage, its RMS and peak amplitudes are equal.

## The Classic VU meter

Today a VU (or pseudo-VU) meter is most often used simply to indicate an impending overload, but that was never part of its original function. An analog tape recorder is typically calibrated so that 0 VU on its meter corresponds to the signal level required to produce the chosen reference fluxivity level (185

nanowebers/meter is standard) on tape. At 10 dB above reference fluxivity, you'll have a measurable but tolerable amount of total harmonic distortion (THD), around 1% for a professional recorder (this is what's usually referred to as "analog tape compression" or "analog warmth"). That 10 dB or so between 0 VU and the point where THD reaches 1% is the headroom. We don't watch a recorder's VU meter to tell how loud our recording is, but rather to assure that we don't exceed the available headroom.

But is the VU meter a good headroom indicator? Only if you're working with program material that's fairly consistent in level and doesn't require a generous allowance for surprise peaks. Take a close look at the VU meter scale. (Figure 2) While the meter scale has a total range of 23 dB, fully half (the top half) of the scale represents only 5 dB. This is good resolution for measuring steady tones (like when calibrating the recorder or setting levels within a system), but pretty wasteful when working with recording media that's capable of handling a dynamic range of better than 65 (or for digital recorders, 90) dB. There's virtually no usable resolution below -10 VU. If we assume that we have at least 10 dB of headroom above 0 VU (that we can't see on the meter), the VU meter really only tells us where a signal is within about a 13 dB range.



Why such a compressed scale? Practicality. Perceived loudness is a logarithmic function. It takes more than twice the signal voltage for something to sound twice as loud. A linear scale that represents loudness just wouldn't look right. Remember that one of the design criteria for the VU meter was that the pointer response looked good for speech.

It's no surprise that modern, highly compressed music will shoot the meter pointer well up scale and it'll stay right there until the fadeout. On uncompressed material with a wide dynamic range, there's plenty of completely audible stuff down below the -20 mark on the VU meter, but an inexperienced engineer who trusts the meter rather than his ears can be misled into thinking that anything that barely moves the meter is too soft.

Today, there are a many meters, both mechanical and of the LED or LCD ladder style with a scale which looks like a VU meter but doesn't meet the VU standards. These are useful for establishing steady state calibration levels when setting up a system, but don't accurately represent either loudness or headroom. LED ladder meters are often found on digital equipment, but it's rarely specified whether they're measuring an analog voltage or the number representing the

amplitude of a digital sample. In either case, the garden variety meter doesn't provide the same dynamic response of a real VU meter. It can show you average, and sometimes peak level, but it won't tell you much about apparent loudness.

0 VU is an arbitrary voltage level. It's whatever is "normal" for the device on which the meter is installed, at the point in the circuit that it's measuring voltage. A meter that indicates input or output level is generally calibrated according to one of a several industry sort-of-standards. The most common one is that 0 VU represents a voltage level of +4 dBu, about 1.23 volts RMS. There's a good story here. See the sidebar at the end of this article for more details for where this odd number came from.

Most modern mixers are calibrated to this convention - when the output meter reads 0 VU, it's putting out +4 dBu. But this isn't always the case. For many years, Mackie believed this was too confusing so they calibrated their meters so that 0 VU = 0 dBu. The "semi-pro" mixers popularized from the 1980s were generally calibrated for 0 VU = -10 dBV (10 dB below 1 volt, about 0.32 volts). Professional recorders are usually calibrated so that their meters read 0 for an input level of +4 dBu. Are you beginning to see what "madness" is in the title of this article? Wait! There's more!

## Digital Metering

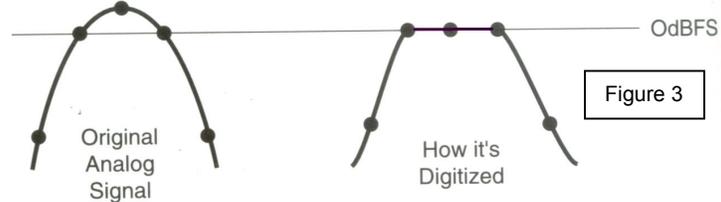
A digital meter (which is *not* a VU meter at all) has a scale with 0 dB all the way at the top. This represents full digital level - all the bits on, regardless of how many bits the system uses. We refer to this maximum level as 0 dBFS which is dB relative to the full scale digital value (all bits on). Unlike with a VU meter, there is no, zero, zilch headroom above 0. You can't turn on more bits than the system's word length. It's up to you, the engineer, to decide where to set the analog-to-digital converter's input gain to allow as much headroom as you'd like as long as the signal level to the converter never exceeds what would be digitized to 0 dBFS. While today's digital systems typically allow internal processing word length greater than what goes in or comes out, a 24-bit analog-to-digital converter is flat out when that 24<sup>th</sup> bit is turned on.

Perceived loudness is the same whether the source is analog or digital, but we view headroom and operating range differently in the two worlds. In the analog world, there's a range over which distortion increases gradually, and we include at least a portion of that range in our useful working range. In the digital world, distortion is nearly independent of level until all the available bits are used up, and at that point, it's (literally) all over.

The meters on early DAT recorders measured the signal level on the analog side of the converter. The "Overload" indicator turning on when the input voltage

reached that produced full scale output from the digital-to-analog (D/A) converter. This tells you that there was likely an overload when recording, but since the recording can never exceed full scale, “overs” metered on the analog side will never show up on playback.

Modern digital recorders and all software DAW meters monitor the actual digitized value. Any input voltage which exceeds the 0 dBFS level will be digitized as 0 dBFS, leaving a nice flat topped waveform looking quite unlike the original. (Figure 3) While the digital level never exceeds 0 dBFS, clearly there’s something wrong, about which we would like to be warned.



The conventional way of determining a potential digital overload condition is to count the number of consecutive 0 dBFS samples, turning on the “Over” indicator when that number exceeds a preset threshold. The Sony PCM-1630, the original “official” digital standard recorder from which all CD glass masters used to be cut, indicates an Over any time it sees three consecutive full scale samples. This is reasonable since we can assume that the first of a consecutive string of 0 dBFS samples occurred while the waveform was on the way up, and the last occurred on the way down, therefore it must have tried to go over the 0 dBFS level some time in between.

Three consecutive 0 dBFS samples for an Over might be considered to be pretty fussy. It could mean that only one sample tried to go over the limit (the ones on either side being exactly full scale). At 44.1 kHz, one sample is too brief to be detected audibly by most listeners, with most forms of program material. Some digital level meters offer a choice of the number of contiguous full scale samples before turning on the Over indicator, so you can be more or less conservative, depending on the music you’re recording.

Digital Over indicators are all about statistics and perception, though. While it might not be musically interesting, you can record a full scale level square wave accurately and the Over indicator will never go off.

Time marches on. We’ve learned through experience and better listening that A/D converters sound better below their maximum level. With modern 24-bit systems, a single full scale sample is usually show as an overload.

The influx of 96 kHz A/D converters raises an interesting question – should we double the number of FS samples required to indicate an “over”, since that would represent a flat-topped waveform of the same duration as three samples at 48 kHz. The jury’s still out on this one (it’s more like nobody ever asked), but modern convention is to indicate an “over” when one sample reaches 0 dBFS.

Something to be aware of with software metering is that refreshing the meter on the screen is usually a low priority task for the program. Their meters are useful for setting initial levels, but they may run a bit behind once the “tape” starts rolling. DAW meters often lag far enough behind the A/D conversion to let you get into trouble if you’re working close to the edge. When depending on an on-screen meter for level monitoring while recording, consider keeping your peaks several dB below full scale. Particularly with 24-bit converters, the lower signal level won’t hurt the signal-to-noise ratio enough to worry about, and if your recording will be tweaked by a mastering engineer, it’ll leave a little breathing room for digital level adjustment or equalization.

Meter lag due to the computer being busy with higher priority tasks isn’t really a problem today since modern computers have quite ample processing power.

Many video production houses have adopted a more stringent convention, primarily to assure that they don’t run out working space when mixing various elements of a project. The SMPTE standard for nominal digital level is -20 dBFS, but rather than allowing peaks up to 0 dBFS, they don’t want to see anything higher than -10 dBFS. This tosses away about a bit and a half of resolution and raises the noise floor slightly, but this isn’t terribly important considering that most playback is from a standard television set speaker, and it absolutely assures no “overs”.

In the early days of CD manufacturing, if your digital master tape contained any “overs,” the mastering house simply rejected it, telling you to fix your problems before cutting the glass master. They didn’t want to be blamed for your distortion. Now that *controlled* distortion is an accepted production tool in popular music, judicious application of digital limiting and compression can keep the Over light from coming on while effectively holding the level within a fraction of a dB below full scale throughout the entire song.

What’s the difference between that and digital clipping? Control. If you don’t like how the limiter sounds, you can change it, but while someone might intentionally (and hopefully, sparingly) uses digital clipping as an effect, it’s totally inappropriate as a leveler.

### **But How LOUD is it?**

Loudness is a function of both the recorded level and how far the listener has the volume turned up. In the past few years, commercial (and following in their footsteps, independently produced) recordings have been in a race to make each one sound louder, when played at the same volume setting, than the last. A whole segment of our industry has sprung up as a result of our inability to train the user to get up and adjust the volume, but this is just a fact of life today.

But what's the nominal operating level of a digital device? The simple answer is that it's whatever you say it is. Many digital recorders and A/D converters with good metering resolution (a lot of lights) have a mark on the scale somewhere between -20 and -12 dBFS. This is the *recommended* operating level, Everything from that point on up to 0 is headroom, but you can leave as much or as little as you want (or dare). For theater sound, Dolby Labs has established the standard of 85 dB SPL (acoustic sound pressure level, C weighted, slow response) for a digital level of -20 dBFS. In a theater with a properly calibrated playback system, one which actually plays back -20 dBFS at 85 dB, most of the audience can hear quiet dialog and they don't get blown out of their seats when the car chase ends in an explosion.

With the coming to being of home theater systems, Dolby recognized that sound mixed for 85 dB plus headroom in a loud theater doesn't translate well to small speakers and small rooms. 20 dB is a lot of dynamic range, and 105 dB SPL is more than most listeners can tolerate, so they recommend lowering sound pressure level to 79 dB for -20 dBFS.

Of course nobody wants to admit to turning anything down, so let's play with the numbers a little and see if we can make up that level. If we start out with the standard Dolby calibration of -20 dBFS = 85 dB SPL, we'll reach 79 dB SPL at a digital level of -6 dBFS. Since we've compressed the bejeezus out of all the tracks, we don't have a lot of dynamic range so we can run a little hotter. Raising the nominal operating level from -20 to -14 dBFS will jack everything up by 6 dB, getting our peaks up to 0 dBFS. Now, 0 dBFS still corresponds to 85 dB SPL, but instead of being able to go 20 dB over the nominal level before running out of bits (the headroom that we've allowed), we now have only 14 dB of headroom. We keep our fingers crossed that it's enough.

How do we deal with the reduced headroom? By reducing the ratio of peak to average level (this is called the crest factor) so that the average can be higher without the peaks trying to exceed the maximum allowable level. 20 dB peak to average ratio is typical for uncompressed music of almost any genre, which is why we start out with a nominal operating level of - 20 dBFS.

When judging apparent loudness, the ear responds to the average rather than peak level. In a typical pop song, the snare drum will be the loudest individual sound, and its peaks tend to ride a few dB above the rest of the mix. This is what sounds right to our ears, and it's those peaks that determine the maximum recording level. Let's say that at several points in the song when the drummer is feeling particularly enthusiastic, the peak level hits 0 dBFS.

Now let's mix our hypothetical song again, but this time we'll sit on that snare with a limiter or compressor so that despite the drummer's enthusiasm, his peaks still remain pretty level. You'll find that the meter no longer hits 0 dBFS, but the mix doesn't sound any quieter either. Now that we've taken care of those

“outliers”, we can boost the master level by a few dB so that get back those 0 dBFS peaks. Since we’ve just raised the overall level (including the average level) the mix sounds louder and we’re still within the digital limit. Voila! Instant mastering!

## **A New Meter and An Experiment**

In a presentation at the October 1999 Audio Engineering Society convention, mastering engineer Bob Katz of Digital Domain described the concept of monitoring at a calibrated level. He’s proposed a new type of meter scale that essentially displays headroom relative to a calibrated monitoring sound pressure level rather than the digital level. His meter scale looks like a VU meter in that it’s calibrated both above and below 0 VU, but unlike a VU meter, it has a linear scale. With Bob’s K-System, when mixing for maximum dynamic range, the scale goes up to +20 VU with 0 VU representing a monitor level of 83 dB SPL. For typical pop music, full scale is +14 VU, while for heavily compressed “radio ready” music, full scale is +12 VU. The K-system recognizes that different genres and different listening environments require different approaches to level management without destroying dynamic range.

While it may seem that the the broadcast mix, the one that we want to appear the loudest, has the lowest full scale level, in reality, it’s really the loudest. With 0 VU corresponding to 85 dB SPL, with the smallest amount of range between 85 dB and maximum, a mix running around 0 VU will keep the average volume higher than that of the widest dynamic range mix which allows 20 dB headroom above 0 VU. If this sounds a bit too much like smoke and mirrors, you can read Bob’s full proposal (he’s continually updating it) on his web page, <http://www.digido.com> and fill in some of the mental blanks.

I thought I’d give this a try at home to see how this works in practice. Feeding my DAT pink noise from my trusty Minirator, I set the record level so that DAT’s meters read -20 dBFS. Using my Radio Shack sound level meter set for C weighting, I adjusted the control room monitor level for a reading of 85 dB at my normal listening position. I put on a tape with a pop music mix and it was so loud it nearly blew me out of my chair! Fortunately I didn’t blow anything and the system still had headroom - that’s good). A tape of an orchestral recording was much more civilized, with crescendos being a bit loud, but in general a healthy level.

OK, this was back around 2001. If you want to try this today, you can find a pink noise WAV file on line, load it into your DAW, and check the SPL with a free iPhone app.

I then re-calibrated the monitor level pot so that -20 dBFS pink noise gave me Bob’s recommended “compressed pop” 77 dB SPL, played the pop music tape again, and darn if it didn’t sound just about right, good and loud, but not window-rattling - about as loud as the crescendos of the classical music, but in your face

all the time, not just on loud sections. That's what seems to be the goal of modern popular music mastering.

## Those Pesky Reference Levels

The real meter madness associated with digital level measurement is when you have devices with different reference (volts for 0 dBFS) levels in the same system. To make it more befuddling, you don't always know a device's reference level unless you measure it. A lot of gear is specified only with a nominal analog reference level, but without telling you its corresponding digital level. Further, many digital I/O devices have no input or output level controls, so you can't calibrate the reference level to match other gear in your system – you have to take whatever calibration the manufacturer gives you. It's bad enough when there's a standard and not everyone follows it, but in this case there's no standard for the analog level equivalent to 0 dBFS.

This leads to the common complaint of “my mixes aren't hot enough” or conversely, “it plays much too loud.” A dirty little secret is that budget priced A/D converters (whether it's a stand-alone converter or integrated into another device) tend to be a bit on the less sensitive side. It's more likely that +4 dBu going in will give a digital level in the –20 dBFS ballpark than –12 dBFS. The reason is that with less gain on the front end, they're digitizing a lower quiescent noise level, so they can advertise a lower noise floor on the digital side. You need to hit the input pretty hard in order to get to 0 dBFS.

Sometimes you can crank up the output level of the source, for example a mixer, and sometimes you can't, like with a microphone. If you have some adjustment room, you need to be careful that you don't push the mixer feeding the A/D converter into clipping before the converter reaches maximum digital level. This can occur if you have a mixer with little headroom (a greater risk with older semi-pro mixers than modern ones), or if you have a mismatch of nominal analog operating levels between the mixer and recorder. If you're using a recorder with +4 dBu input sensitivity together with a mixer with a nominal operating level of -10 dBV, the only way you'll be able to get the recorder's meters to approach zero on peaks is to run the mixer's output level meter well above zero, which will seriously compromise, and may exceed, the headroom in your mixer.

### 21<sup>st</sup> Century Update

I've started packing a handheld digital recorder to make documentation recordings when I'm running the PA system, and I've run into a devilish problem. While these recorders have a record level control which will make the meters read on-scale, it's on the digital side of the A/D converter. A pro PA mixer's output can easily overload the input stage of the recorder (which has no level control) and what *feeds* the A/D converter will be clipped. The solution here is to put an attenuator in line with the signal feeding the recorder, or use a separate output on the mixer to feed the recorder, and bring its level down to the point where the input stage is safely below clipping.

## **Keep Your Eye On The Ball (and not on the meter)**

Once you've properly calibrated the gear in your studio, meters can tell you a lot about what's happening, but don't become a slave to them. Use your meters as tools and you'll spend more time recording and less time worrying about whether you have your levels set properly.

## **The VU Meter – A big sidebar**

There are standard specifications for how a VU meter should behave, and all the good ones (like on the front of your grandfather's Ampex) conform pretty closely to those specifications. Problem is that you can't always tell by looking at a meter how closely it complies to the standard. The specification defines several important characteristics:

- Sensitivity - How many volts it requires to deflect it to the zero mark
- Impedance - How it loads the circuit to which it's connected
- Rectifier characteristics - How it responds to various waveforms
- Ballistics - the rise and fall time of the meter pointer
- Scale - It has sufficient resolution to be useful

All of these things go toward indicating how loud we perceive a sound. Unfortunately, none of those characteristics help to determine peak levels, and that makes it tricky to use a VU meter to indicate level going to a digital device. Let's look at those characteristics and see how meters differ.

## **Sensitivity and Impedance**

When Western Electric designed the VU meter, they were concerned with driving miles and miles of telephone lines, which, like a loudspeaker, required a certain amount of power (voltage x current). They used 0 dBm (0 dB referenced to 1 milliwatt) as their standard unit of power. But for interfacing modern audio devices which draw negligible current (and hence require negligible power) we're interested in voltage.

Since the standard load impedance in a telephone system is 600 ohms, we defined the audio reference voltage as that which produced 1 milliwatt of power into a 600 ohm load and called it 0 dBu. This is 0.775 volts into any load impedance. It would have been convenient to build our VU meter so that it would indicate zero with an input of .775 volts, but that presented too many design conflicts. In order to achieve that sensitivity with purely a mechanical meter (no

electronics), it would have to have an impedance of around 400 ohms, which was low enough to drag down the output of the circuit it was measuring.

To solve that problem, the designers put a 3600 ohm resistor in series with the meter movement so it wouldn't significantly load a 600 ohm output, but the voltage drop across the resistor caused the meter to read about 4 dB too low. The solution (remember, this was in the days of gobs of headroom) was simply to crank up the voltage required to get the pointer up to 0. And so, the professional standard for audio equipment today is that a standard VU meter will read zero for a level 4 dB above the 0.775 volt (0 dBu) reference point, or about 1.23 volts.

## **AC/DC**

The type of movement used in the classic VU meter responds to direct current. Put AC into it and the pointer will just vibrate around the bottom of the scale. Since audio is by nature alternating current, the VU meter employs a rectifier to change the AC into DC in order to make the meter read an audio voltage. In the heyday of VU meters, there were only two types of solid state rectifiers, and the type that was the most linear, copper oxide, was chosen for use. It didn't give a true RMS output for any waveform, but at least all the meters were the same. Nowadays, rectification is usually done with an operational amplifier which gives more linear rectification than a copper oxide rectifier. Consequently a modern VU meter that's on the front panel of a preamp or mixer is likely to give a different reading than a traditional one when the waveform diverges from a simple sine wave.

## **Going Ballistic**

An ordinary voltmeter (if you can still find one with a pointer), though accurate for continuous simple waveform, doesn't do a very good job of responding to the rapid level changes and non-symmetrical waveforms present in audio. The VU meter's movement is designed so that its reading will correspond to a bench meter on steady state tones (so it can be calibrated), but on complex audio waveforms, will read somewhere between the average and peak value, making it work about like the human ear. If peaks are frequent and closely spaced, (for instance on a drum solo) the VU meter will tend to read a little higher than average level, but won't make it all the way up to the peak level. A proper VU meter takes about 1/3 of a second after a 0 VU level signal is applied for the pointer to reach the zero mark (by which time the signal level may have dropped a bit below its peak), while many electronic meters have near instant response. There are designs for electronic meters which approximate the mechanical response of a real VU meter, but it's not generally an advertised feature.

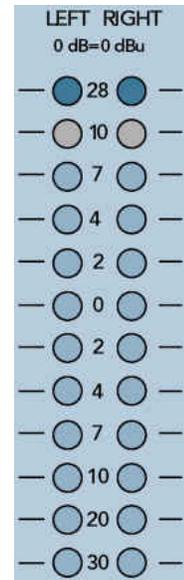
The VU meter ballistic makes it a useful tool because it approximates how we actually hear, but the engineer must understand that when it's necessary to know the actual peak level, a VU meter doesn't tell the whole story. The VU meter may read -6 during a drum solo, but the peaks when the stick hits the head may be 10 dB higher than that.

## Scale

The traditional VU meter has a standardized scale (even the color scheme is standard), and there's little difference between one and another. The scale runs from -20 to +3, with usable resolution over about a 13 dB range. But an LED meter has a light for each calibrated level, and they tend to cover a much wider range than a mechanical meter. The lowest indication is usually -30 or -40,



while the top of the scale goes to +20 or even as high as +28. The steps between the lower LEDs are generally fairly large, resembling the low resolution at the bottom end of the scale of a



mechanical meter. However, because of the extended scale above 0 VU, the LEDs up at the top tend to have fairly large steps between them as well. On the Mackie 1604 VLZPro meter shown here, the LEDs between -4 and +10 have roughly the same resolution as the working range of a mechanical meter, but there's only one step between +10 and +28. You can't tell, from this meter, where the level is when it's above +10, which, as the meter says, corresponds to +10 dBu.

You can probably live with this information gap when the mixer is feeding an analog device like a tape recorder or PA amplifier, but you're in trouble if you try to use this meter as a guide to setting levels with a digital system. On a typical A/D converter with +4 dBu corresponding to -20 dBFS, the second LED from the top (+10) corresponds to -14 dBFS. The top LED (+28) doesn't flash until you're 4 dB above full scale level. So in that critical range between -20 and -6dBFS, you really have only one useful LED indicator. One of these days they'll catch on.