

LIVE SOUND

I N T E R N A T I O N A L



INSIDE

KNOW YOUR NOISES

Examining the overlooked factor behind unexpected meter readings.

VINTAGE SOUND QUALITY & REPLICATING IT IN MODERN MIXING

KEY INTERCOM SYSTEM COMPONENTS, FEATURES & FUNCTIONS

A REAL WORLD GEAR LOOK AT WIRELESS MICROPHONE SYSTEMS





KNOW YOUR NOISES

The overlooked factor behind unexpected meter readings.

by *Merlijn van Veen*

In a May 2025 video, Dave Rat (renowned FOH engineer and founder of Rat Sound Systems) highlighted a puzzling metering mystery. Inside his digital mixing console, he summed two internally-generated white noise signals – first identical, then uncorrelated.

According to theory, the peak level should rise by +6 dB for identical signals and only about +3 dB for uncorrelated signals. Yet in both cases, the console's peak meter showed the same +6 dB increase.

Meanwhile, his external analog meters behaved as expected, clearly reflecting the 3 dB difference, and when that analog sum was routed back into the digital console, its meter tracked correctly too. So why

did the initial test – conducted entirely within the digital domain – appear to fail?

TL;DR

The apparent discrepancy, and the resulting confusion, arises from engineering shortcuts used in signal generation and measurement. These shortcuts – including uniformly distributed noise generation, rectifying average-responding meters, and sine-wave-based calibration conventions – simplify implementation but alter how peak level, average level, and discrete digital samples relate to the continuous analog waveform.

As a result, measurements made under these conditions can appear inconsistent when compared across devices operat-

ing under different assumptions. While these shortcuts rely on approximations of signal behavior and measurement quantities, they introduce differences that can produce misleading results in precise measurement contexts. This article exists to clarify those discrepancies and their underlying causes.

The apparent disagreement arises because, in the digital domain, uniformly distributed white-noise samples are summed. After passing through the D/A converter, however, the resulting analog waveform exhibits different statistical behavior and produces different peak characteristics.

Moreover, the peak value of a digital signal is defined solely by its largest individual sample value, regardless of whether those maximum samples occur in isolation or in succession. In contrast, the analog signal at the output of the D/A converter may reach a higher instantaneous amplitude when multiple large samples occur consecutively, due to the interpolation inherent in the conversion process.

This means that while each piece of equipment was operating correctly within its own design framework, the combina-

tion of differing assumptions, implementations, and measurement conventions produced results that appear contradictory but are, in fact, entirely consistent when viewed in context.

This also illustrates a broader principle: peak meters are not reliable indicators of perceived loudness. Their proper role is to monitor proximity to clipping, not to assess subjective or statistical loudness characteristics.

Ultimately, this reflects a fundamental measurement principle, expressed succinctly by long-time audio professional Ivan Beaver: “If you don’t know the answer before you start to measure, how do you know you’re getting a good measurement?”

In this instance, without a full understanding of both the signal characteristics and the measurement methods employed, it becomes difficult to properly interpret – let alone reconcile – the observed results.

CREATED EQUAL?

White noise can be generated in different ways – and those methods matter. For example, you might create it from a sequence of coin flips (a Bernoulli distribution [1]), by rolling dice (a Uniform distribution [2]), or by sampling from a Gaussian (a Normal distribution [3]).

Each approach produces a signal with different statistical characteristics – specifically, their crest factor [4] – even though they all sound equally “white” to both you and me, as you’ll hear in this video.

Bernoulli white noise has a crest factor of 1 (or 0 dB), uniform white noise has a crest factor of $\sqrt{3}$ (approximately 4.8 dB), and Gaussian white noise typically exhibits a crest factor around 4 (about 12 dB). The latter is explained shortly. Click here to download the audio files.

GAUSSIAN NOISE

The crest factor of Gaussian random noise is theoretically infinite, as the distribution is unbounded. In practice, however, extreme values are rare. The frequency with which a given amplitude threshold is exceeded can be estimated using the survival function [5] of the underlying distribution, with contributions from both tails accounted for explicitly.

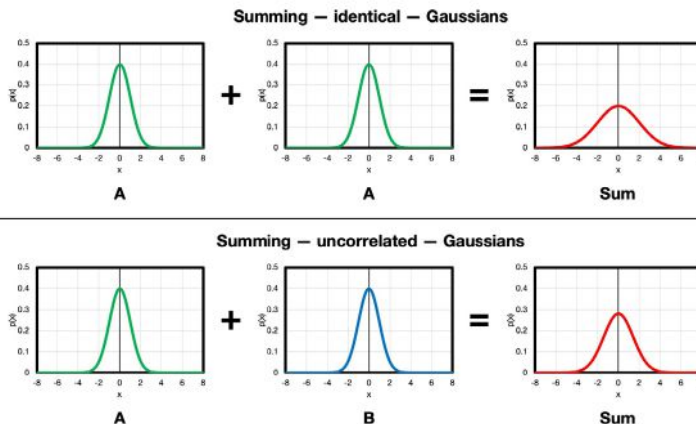


Figure 1

For example, in white Gaussian noise generated at 48 kHz:

- a sample exceeding +12 dB above the RMS level is expected about once every 300 milliseconds;
- a +15 dB event may occur once every 18 minutes;
- and a +18 dB peak might appear only once every 335 years.

Because it’s impractical to design around these extremely rare events, the audio industry has reached a general consensus: treating Gaussian noise as having a 12 dB crest factor. While there are some exceptions and differing practices, most follow this guideline, a convention I’ll adopt for the remainder of this article.

SUMMING GAUSSIAN DISTRIBUTIONS

Figure 1 shows the probability distribution function (PDF) [6] for Gaussian noise. Note the characteristic bell-shaped curve. When summing identical Gaussian noise signals, both RMS and peak levels rise by +6 dB; for uncorrelated Gaussian noise signals [7], the increase in both RMS and peak is only +3 dB. In either case, the 12 dB crest factor remains functionally the same.

What about the white noise generators in consoles and digital audio workstations. Which distribution do they use?

UNIFORM NOISE

The easiest way to generate white noise is to use a random number generator that produces values uniformly distributed between two equal but opposite extremes, where each value is equally likely. It’s reminiscent of rolling a die, where each number from 1 to 6 has an equal chance of appearing – except that our “die” has significantly more faces.

Unlike Gaussian noise, which has a high crest factor due to its rare outliers, this signal maintains a much lower crest factor – just under 5 dB ($\sqrt{3}$ on a linear scale). So what happens if we sum these together?

When summing identical uniform noise signals, both RMS and peak levels once again rise by +6 dB, and once again, no crest factors were harmed during the making of this summation.

Figure 2 shows the probability mass function (PMF) [8] for two dice with identical outcomes and a pair of uncorrelated dice. For discrete random variables, such as dice, probabilities are described by a PMF, which plays a role analogous to the probability density function (PDF) for continuous random variables.

Note in Figure 2a that when one die is forced to take on the value of the other, the face values are simply doubled. By contrast, when summing uncorrelated uniform distributions, the resulting distribution changes from uniform – char-

acterized by its rectangular shape – to triangular (Figure 2b).

Reminiscent of the board game Catan (Settlers of Catan), where you roll two dice and take the sum. While each die on its own follows a uniform distribution, their combined total does not.

It's much rarer to roll a 2 or a 12 than it is to roll a 7, simply because there are fewer possible combinations that produce the extremes and many more that add up to values near the center.

TRIANGULAR NOISE

The crest factor of a triangular distribution [9] (Figure 2b) is 3 dB higher than that of a uniform distribution, increasing from just under 5 dB to just under 8 dB ($\sqrt{6}$ on a linear scale).

Note in Figure 2 that sums for dice with identical outcomes and sums for uncorrelated dice share the same extreme values, namely 2 and 12. Thus, the peak limits of the uncorrelated sum are identical to those of the sum of identical signals. However, the large number of combinations that produce intermediate amplitudes causes the RMS value to increase less relative to the peak (Figure 2b). Consequently, the crest factor increases by 3 dB.

While rolling a pair of ones or a pair of sixes is six times less likely with independent, uncorrelated dice than with dice having identical outcomes, the probability remains non-zero.

Because extreme values remain possible even though they occur less frequently, triangular noise still produces near-maximum peaks with regularity. At a 48 kHz sampling rate:

- a sample typically falls within 0.5 dB of the distribution's extremes once every 7 milliseconds;
- and within 0.1 dB, approximately every 160 milliseconds.

Accordingly, the peak-hold feature of a peak meter refreshing at, for example, 50 Hz, is very likely to capture a near +6 dB peak during each 20-millisecond interval.

Therefore, when summing uniform noise signals, peak measurements tend toward +6 dB, regardless of whether the

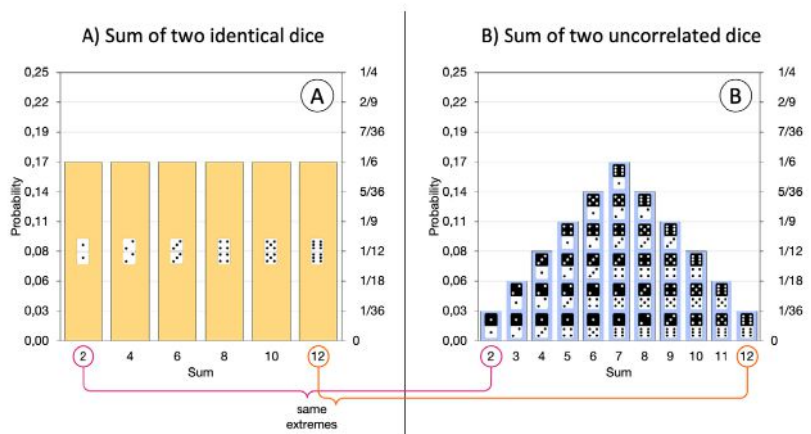


Figure 2

PDFs for *n* "dice" centered at zero




Figure 3

signals are identical or uncorrelated. This is evident from the dice sums in Figure 2, whose extremes coincide, and explains why the digital mixing console's peak meter shows the same level in both cases.

MORE DICE PLEASE

Once we consider more than two dice, illustrating every possible discrete PMF with dice pictograms, as in Figure 2, quickly becomes impractical. Instead, we can gain a qualitative understanding by considering the corresponding continuous PDFs (Figure 3).

The Central Limit Theorem (CLT) [10] states that the sum (or average) of a large number of independent, identically distributed random variables – such as rolling “n” dice – tends toward a normal (bell-shaped) distribution, regardless of the original uniform distribution of each die.

As more dice are added, the distribution of the sum becomes smoother and more symmetric about the mean, with middle sums occurring more frequently. Note the progression from uniform (rectangular) to triangular and then to increasingly bellshaped distributions, as shown in Figure 1, as the number of dice is increased.

More generally, whenever two uncorrelated signals whose crest factor is less than 12 dB – such as those with a uniform amplitude distribution – are summed, the crest factor of the resulting signal tends to increase as the amplitude distribution becomes more Gaussian.

The CLT will play a central role later in this article in explaining what happens to our noise signals once they are passed through a digital-to-analog converter (DAC). So why can't the digital mixing console's persistent peak meter reading

be reconciled with analog meters as well as our hearing? In other words, “A man with one meter knows what he’s measuring. A man with two meters is never sure.”

PEAK LEVEL IS NOT LOUDNESS

Of all the meters encountered in the digital domain, the peak meter is ubiquitous for a reason: hard clipping is unforgiving.

However, peak levels – unlike RMS levels – do not convey loudness, whereas metrics that do, such as loudness units relative to full scale (LUFS) [11] or equivalent continuous sound level (Leq) [12], are all proportional to RMS levels. Human hearing integrates energy over time and is therefore more closely aligned with RMS-like metrics than with instantaneous peaks.

Further, to reconcile what flavor of noise we are dealing with in the analog domain, we must examine one of its defining signatures – besides noise color – namely its crest factor, derived from both the signal’s peak level and its RMS level.

Both the analog mixing console and the analog Dorrough 1200B stereo signal test set employ average-responding meters, which does not preclude them from measuring signal peaks, as we will discover shortly. However, only the Dorrough is fast enough to also monitor peak levels. Furthermore, careful attention must be paid to the scale of each meter.

METER SCALES

The Dorrough 1200B is unique within this setup in that it provides one segment per decibel, whereas the other meters in this setup display anywhere from one dB per segment to as much as 30 dB per segment, making them ill-suited for accurate level readings (**Figure 4**).

Finally, not all averages are created equal, and they can bias our conclusions. While the analog meters in this setup are average-responding, they are not so-called *True RMS* meters.

NOT ALL AVERAGES ARE EQUAL

The simplest and most cost-effective way to create an average-responding meter is to apply full-wave rectification to the signal and then measure its mean value.

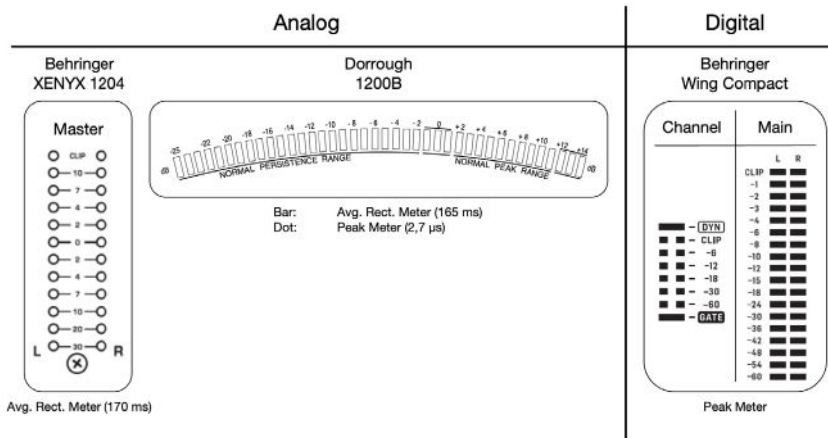


Figure 4

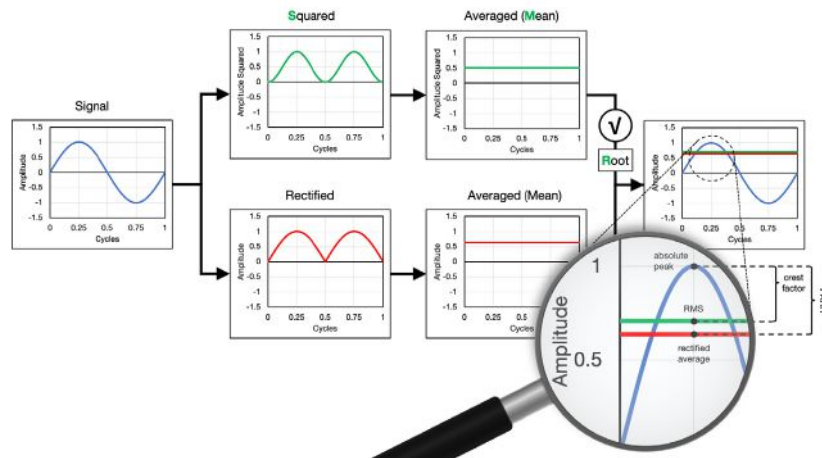


Figure 5

This tells us – on average – how far the waveform’s amplitude strays from zero (**Figure 5**).

However, unlike RMS, which requires squaring and is more difficult to implement in analog circuitry, this metric is not proportional to signal power and therefore does not directly convey loudness. Except for a few edge cases, a rectified average is biased to show a smaller value than the true RMS value because rectification does not preserve the relationship between amplitude and signal power. The magnitude of this bias depends on the signal.

The peak-to-average-rectified ratio (PARR) shown in Figure 5 is analogous to the crest factor (the peak-to-RMS ratio) and is always higher, except for signals

such as square waves. Please note that I use the abbreviation *PARR* (with two Rs) to avoid confusion with *PAR* (peak-to-average ratio, with one R), which is commonly used in the audio industry to mean the peak-to-RMS ratio, also known as the crest factor.

For clarity, I will refer to the crest factor explicitly as the peak-to-RMS ratio throughout this article and will avoid using the term peak-to-average ratio altogether.

PROPERTIES OF VARIOUS SIGNALS

The table in **Figure 6** summarizes RMS-normalized metrics for the noise signals discussed in this article, as well as for a sine wave. As stated earlier, simply because a meter is average-responding,

COPYRIGHT (C) MERLIJN VAN VEEN. ALL RIGHTS RESERVED

COPYRIGHT (C) MERLIJN VAN VEEN. ALL RIGHTS RESERVED

do not assume that it cannot measure signal peaks; this depends on the meter's responsiveness.

Whether an averaging meter can measure signal peaks depends on its integration time constant. The analog mixing console, judging from its electrical schematic, has a time constant of approximately 170 ms, which is too slow to track instantaneous signal peaks.

For reference, the Fast integration time used in sound level meters is 125 ms [12], whereas the now-deprecated Impulse time constant was 35 ms.

The Dorrrough 1200B, on the other hand, judging from its electrical schematic, processes the rectified signal using two separate meters. One employs a 165 ms time constant to indicate the average level, which Dorrrough refers to as the bar. The other uses a 2.7-microsecond time constant (2,7 μ s) for peak measurement, which Dorrrough refers to as the dot – in theory, fast enough to capture even instantaneous signal peaks, as modeled in **Figure 7**.

Note that for both uniform and triangular noise, the time-weighted rectified signal repeatedly reaches the expected peak level within a one-second interval. Likewise, the rectified average – although it underestimates the RMS level – should converge to its expected value after several time constants.

To test this theory, I purchased the Waves plugin version of the Dorrrough meter.

SOFTWARE VS HARDWARE

When I feed the Dorrrough plugin with uniform and triangular noise, I expect to observe PARR values of 6 dB and 9,5 dB, respectively, as indicated in the table in Figure 6. However, the plugin instead reports apparent values of 2 dB and 6 dB, which are roughly 4 dB lower than expected.

It turns out that Dorrrough meters are calibrated, according to the 1200B manual, using a sine wave such that the *dot* and *bar* illuminate the same segment, making that segment brighter. In other words, the Dorrrough meter under-reports instantaneous signal peaks by roughly 4 dB due to its sine-wave-based calibration.

	Sine wave		Uniform noise		Triangular noise		Gaussian noise	
	Lin	Log	Lin	Log	Lin	Log	Lin	Log
Peak	$\sqrt{2} \approx 1,41$	+3 dB	$\sqrt{3} \approx 1,73$	+4,8 dB	$\sqrt{6} \approx 2,45$	+7,8 dB	~ 4	$\sim +12$ dB
RMS	1	0 dB	1	0 dB	1	0 dB	1	0 dB
Average Rectified	$\frac{2\sqrt{2}}{\pi} \approx 0,90$	-0,9 dB	$\frac{\sqrt{3}}{2} \approx 0,87$	-1,2 dB	$\frac{\sqrt{6}}{3} \approx 0,82$	-1,8 dB	$\sqrt{\frac{2}{\pi}} \approx 0,64$	-2 dB
Crest Factor	$\sqrt{2} \approx 1,41$	+3 dB	$\sqrt{3} \approx 1,73$	+4,8 dB	$\sqrt{6} \approx 2,45$	+7,8 dB	~ 4	$\sim +12$ dB
PARR	$\frac{\pi}{2} \approx 1,57$	+3,9 dB	2	+6 dB	3	+9,5 dB	$4\sqrt{\frac{\pi}{2}} \approx 5,01$	$\sim +14$ dB
Dorrrough PARR _{noise} *	1	0 dB	$\frac{4}{\pi} \approx 1,27$	+2,1 dB	$\frac{6}{\pi} \approx 1,91$	+5,6 dB	$4\sqrt{\frac{2}{\pi}} \approx 3,19$	$\sim +10,1$ dB

* Calibrated w/ sine wave such that Peak = Average Rectified where peaks are attenuated by -3,9 dB

Figure 6

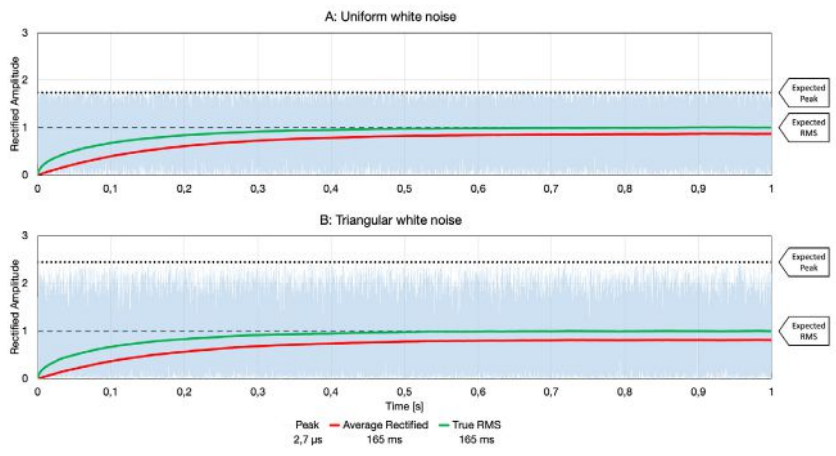


Figure 7

This is confirmed by the sine-wave measurement shown in the first row of **Figure 8** and prompted the inclusion of the final row in the table in Figure 6, labeled PARRsine, to denote the derated PARR values the Dorrrough meter is expected to report as a result of its sine-wave calibration.

Great – so the plugin behaves exactly as expected; we simply need to account for the 4 dB offset. Let's now see whether the hardware reports the same results.

Surprisingly, while the bar readings in Figure 8 are identical, all of the peak readings are higher. Why are the peaks inflated in the analog domain but not in the digital domain?

GET READY TO BE CONVERTED

To create an analog continuous waveform, the gaps between digitally generated noise samples must be filled in. This is where the digital-to-analog converter's (DAC) anti-imaging filter [13] comes into play: intermediate values are formed as weighted sums of many consecutive samples through convolution [14] with the filter. For noise, these samples are initially random and independent, with no dominant terms.

When many such random variables are summed, as in the dice example in Figure 3, the Central Limit Theorem (CLT) applies, and what began as a uniform or triangular distribution becomes increasingly Gaussian-like.

Figure 9 shows one possible realization of the waveform. Although infinitely many waveforms can pass through the sample values, the specific realization is uniquely determined by the anti-imaging filter. The resulting analog distributions are undeniably more bell-shaped, especially for the triangular distribution, which was already closer to a bell shape to begin with.

Note the newly introduced inter-sample peaks in the waveform approximations that exceed the noises' original bounds. These additional peaks inflate the crest factor and, subsequently, the PARR of the analog signals' versions.

Also note that I deliberately avoid the term "reconstruction," as the signals discussed here are highly synthetic and originate in the digital domain. If, instead, an analog signal were confined to below the Nyquist frequency (half the sample rate) and subjected to AD/DA conversion, then – in principle – no change to the signal would be expected, and the term reconstruction would be appropriate.

While uniform and triangular noise can be generated digitally, they cannot exist exactly in the analog domain. Once band-limited and filtered, their amplitude distributions inevitably become more Gaussian due to the summation of independent sample contributions.

As a result, the hardware Dorrough meter observes the same noise signals with higher PARR_{sine} values than the plug-in operating in the digital domain. If we compensate for the Dorrough's sine-wave calibration by adding 4 dB, the incoming noise signals – after passing through the DAC – must have had PARR values on the order of 9 dB and 11 dB for the uniform and triangular distributions, respectively.

This places them very close to the forecast values shown in Figure 9, bearing in mind that the DAC in Dave's digital mixing console may yield slightly different results.

	Waves Plugin	Hardware
Sine wave*		N/A
1x Uniform		
2x Uniform sum identical		
2x Uniform sum uncorrelated (triangular noise)		

* Calibrated w/ sine wave such that Peak = Average Rectified where peaks are attenuated by -3.9 dB

Figure 8

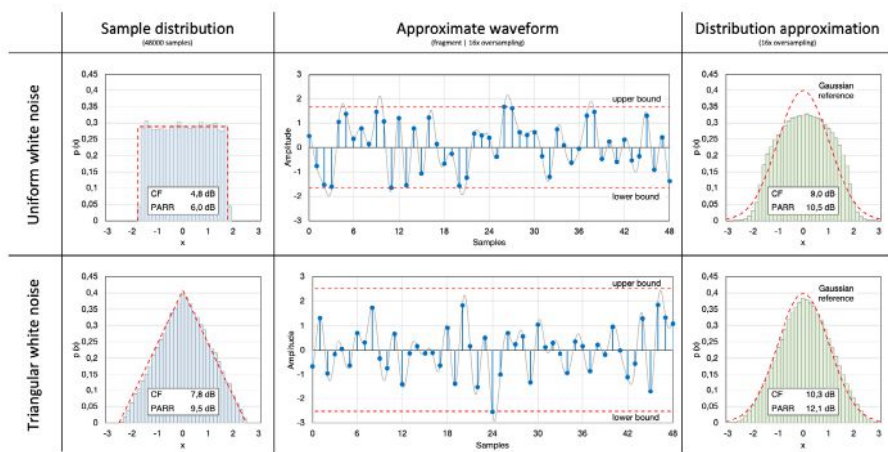


Figure 9

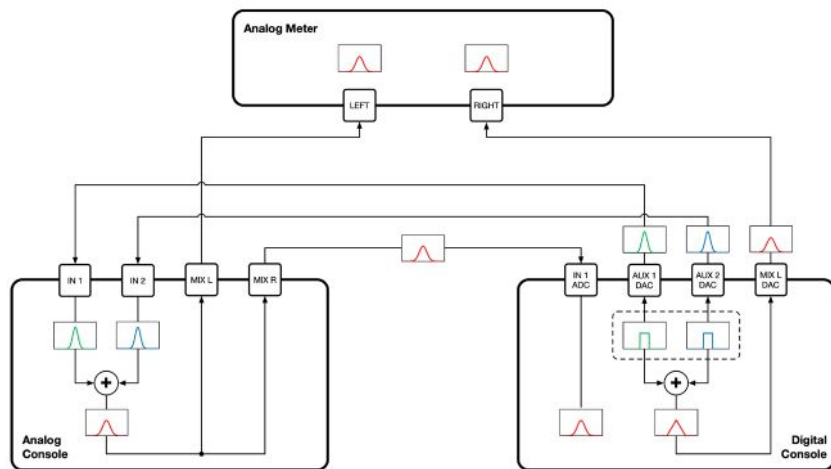


Figure 10



Watch On Youtube: <https://www.youtube.com/watch?v=xNSntbMb6EM>

Furthermore, additional analog-domain filtering may introduce further weighted integration and subsequent Gaussianization.

CONCLUSION

Figure 10 shows the signal flow in Dave's setup, with pictograms representing the various distributions to illustrate how the signals propagate through the system. As you can see, while uniform distributions were being summed in the digital domain, Gaussian-like distributions were – perhaps unknowingly – being summed in the analog domain due to DAC filtering. And for Gaussians, as noted earlier, both RMS and peak levels rise in lockstep.

In contrast, in the digital domain with unaltered uniform distributions, only the RMS level increases when switching from uncorrelated to identical signals — a change that the peak meter won't reflect, since the peak levels remain functionally constant (Figure 2).

The meter on Channel 1 of the digital console, connected to the analog console's mix bus, will track the Dorrough's peak meter because the signal has

already undergone the transformation that affects the noise's distribution and subsequent peak behavior.

For a counterexample, repeat the entire experiment, but this time, instead of using the digital console's internal (uniform) white noise generator, load two uncorrelated Gaussian white noise signals. In that case, even the console's peak meter – as well as the external analog meters – will reflect whether identical or uncorrelated signals are being summed.

Acknowledgements: I thank Dave Rat for indulging me when I shared this observation with him. True to form, Dave raised plenty of follow-up questions and concerns – given the prevalence of peak meters on digital consoles – about the practical implications of using them to monitor signal levels in a way that aligns with how humans perceive loudness. My answer was simple: don't rely on peak meters to judge loudness – use them to monitor proximity to clipping.

I also give special recognition to Dr. Roger Schwenke, honorary Mythbuster, for peer review, proofreading, and nudging me to make the article better. LSI

References

1. Wikipedia: Bernoulli distribution
2. Wikipedia: Uniform distribution
3. Wikipedia: Normal distribution
4. Crest Factor Part 1: Peak-to-Average Ratio
5. Wikipedia: Survival function
6. Wikipedia: Probability density function
7. Wikipedia: Sum of normally distributed random variables
8. Wikipedia: Probability mass function
9. Wikipedia: Triangular distribution
10. "But what is the Central Limit Theorem?" by Grant Sanderson of 3Blue1Brown
11. Wikipedia: LUFS
12. IEC 61672-1 "Electroacoustics - Sound level meters - Part 1: Specifications"
13. Wikipedia: Reconstruction filter
14. "But what is a convolution?" by Grant Sanderson of 3Blue1Brown

Based in The Netherlands, **Merlijn van Veen** (www.merlijnvanveen.nl) is a noted audio educator, and he also serves as senior technical support and education specialist for Meyer Sound.