MAAT

LINpro & LINsurround

User Manual

MAAT Inc.

MAAT Incorporated

101 Cooper St

Santa Cruz CA 95060 USA

More unique and essential tools and tips at:

www.maat.digital

Table of Contents

Installation & Setup

MAAT LINpro and LINsurround are based on the same plug–in. A mono/stereo license enables LINpro, while a multichannel license enables LINsurround. When preparing for installation of your MAAT LINpro and LINsurround, we recommend that you quit all applications prior to proceeding. Also, after downloading the installer and prior to installation, let any anti–malware measures you may have running scan the installer. Then, temporarily disable all anti–malware measures once scanning is complete.

To install, simply double click on the downloaded Installer for your particular operating system. The installation process will guide you through the install procedure. If you have an older OS version, restarting your host may be required. Don't forget to reënable anti–malware measures or simply restart your host.

Licensing

MAAT software is licensed through our dongle–free, cloud–based license control framework. If your hosts are net–connected, our licenses "float," making it easy to move from one machine to another. All you have to do is close all MAAT stuff on one machine before you open anything on another. If you're going to use the laptop in an environment where Internet is spotty or non–existent, you can also go offline temporarily for up to 30 days.

To license your MAAT purchase, open yours or instantiate an example in your preferred plug–in host. When licensing, we recommend you open a new session in your DAW just for this purpose.

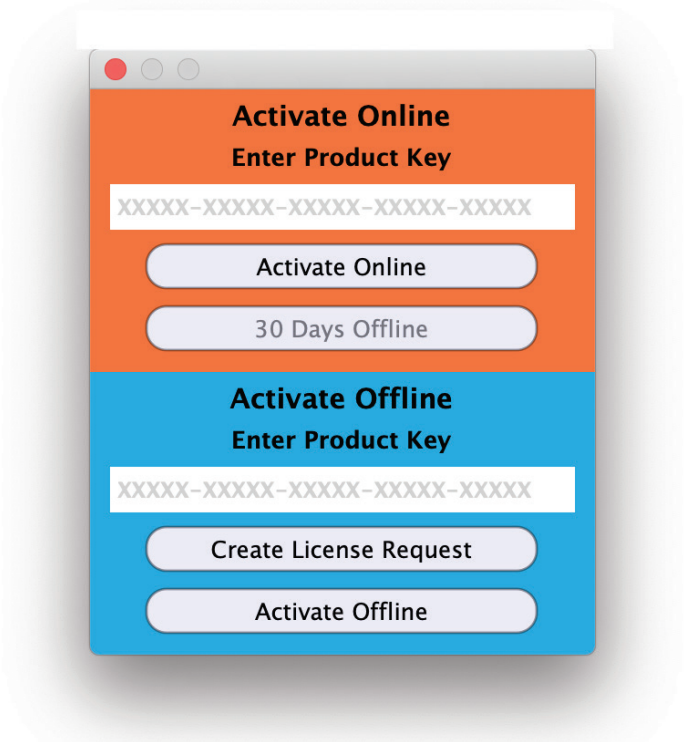

Figure 1: The initial Activation dialog

Online Activation

Once the Activation dialog appears, paste or type in your Product Key supplied at the time of purchase. Then click on Activate/Deactivate Online and follow the prompts.

> NOTE: The Online process requires a live internet connection. You must have an active internet connection for activation.

A feature of our cloud–connected system is that, if you close a DAW session that contains MAAT plug instances, then open that session on another host, the license will "follow" you onto the new host as long as you have live Internet.

You can also return or park your license on our licensing server, and pick it up later on another host. Once your product is licensed, click the Activate/Deactivate Online button to deactivate your license, then repeat to reactivate.

Offline Activation

If, for security reasons you have an offline host computer or you plan on working without a reliable Internet connection, it is best to go with an offline license. There are two offline options, temporary offline and full offline.

Temporary Offline

If you plan on working without a reliable Internet connection, we designed the temporary offline option just for you. The temporary offline process also requires a live Internet connection, but only during activation.

A temp offline license has two unique features:

- It can remain offline, without an Internet connection, for up to 30 days
- It auto–renews its lease whenever it does make a server connection

So, if your Internet is spotty or unreliable, or you're heading temporarily to a location where an Internet connection isn't available, the temporary offline option let's you work untethered to the 'net for up to a month.

By clicking the "30 Days Offline" button, your online license will become temporarily offline for a maximum of 30 days. If, however, you use the product while connected to the Internet at all during that time, the "lease" duration is automatically reset, extending another 30 days.

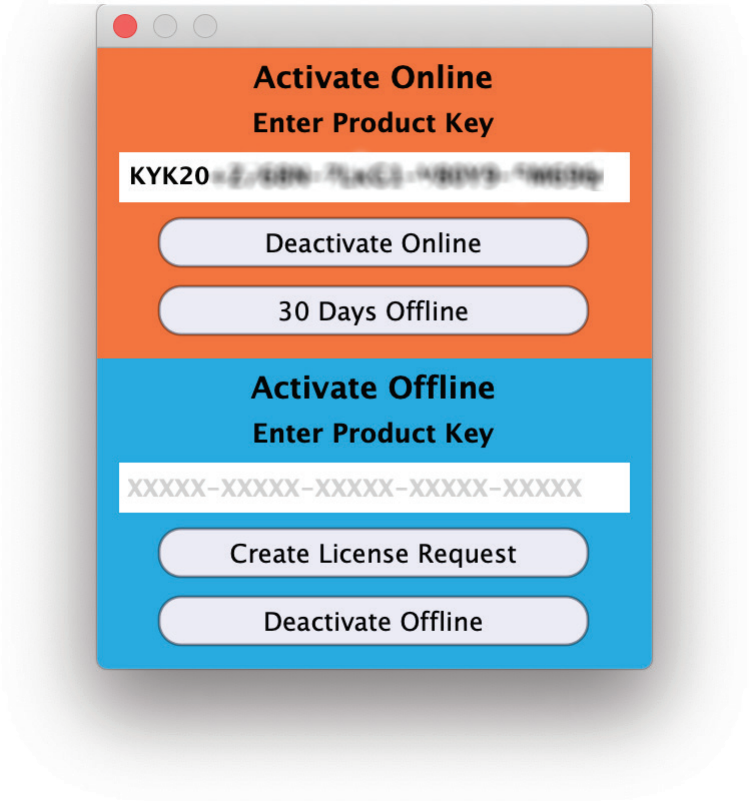

Figure 2: The 30 Days Offline activation option

Full Offline

This will license your MAAT product without the need to periodically contact our licensing server. If you anticipate working without the Internet for more than a month at a time or, for security purposes, your studio machine is "air gapped" and has no Internet, take a moment to surrender your online license, and replace it with an offline version.

NOTE: A FULL OFFLINE LICENSE DOES NEED TO BE EXPLICITLY RENEWED ONCE A YEAR.

To prepare for offline activation, you will need:

1. Your Product Key!

2. Some method to move your License Confirmation file to the offline host machine.

That latter requirement can be accomplished in any number of ways, but a USB flash drive is often readily available and most convenient. You will also need your Product Key that was supplied to you via e–mail when you purchased your MAAT product.

The offline activation process starts with generating a License Request, with a "maatr" file extension, which is submitted on–line

to our server or via e–mail to MAAT Support <support@maatinc.zohodesk.com>. A License Confirmation file, with a "maatc" extension, will be returned to you, which will activate your product.

To proceed, click on the Activate/Deactivate Offline button, and follow the prompts.

In a year's time, 365 days from the date the maatr file is created, a full offline license must be renewed using the same maatr/maatc method.

Introduction

Thanks for purchasing MAAT's LINpro re-dithering plug–in, and congratulations on committing yourself to maintaining and preserving linearity in music!

We built LINpro and LINsurround to be used during post–production, when redithering is most needed. This is especially true during the mastering stage, where it is quite common to play back a digital stream and send that to one or more analog devices. Once the signal passes through the analog equipment, it is re–digitized and further processed in the digital domain. These steps require re-dithering and we provide quite a bit of control over that process so the engineer can either use the optimal dither for your release format, or a custom setting for a subjective "sound."

Today, all audio workstations allow high quality production processes: recording, mixing and mastering, normally with at least 24–bit final output quality. However, with the exception of iTunes and a few other services, lossless distribution and storage format is often the 16–bit linear PCM format widely used for CD production and lower quality lossless streaming. If optimally achieved, 16 bits are sufficient to deliver an excellent quality end product, far better than any lossy coding format.

Quick Start

We know, reading is old school and such an easy tool to use as LINpro would seem to be self evident, wouldn't you think? However, before you begin your explorations with this tool, we want to make a few things clear so you will get the best use out of your "LIN" purchase. Thanks for struggling through this part!

Dithering Basics

Before we dive into details, here are a few very important concepts to remember:

- Dither prepares the finished audio for down–stream processing by the consumer
- When in doubt, TPDF is the universal default choice
- DACs usually do not dither

We will elaborate more shortly but, when in doubt, use High Passed Triangular Probability Density Function or HP–TPDF dither. It will func-

tion properly in all delivery and playback scenarios. Note that DACs or Digital–to–Analog Converters usually do not redither. So, for DACs with digital volume controls, the audio quality will not be optimal when the volume control is used.

To produce 16–bit formatted files, high resolution source audio files, minimally 32–bit floating point or 24–bit fixed, have to be reduced to a 16–bit word length. After this operation however, referred to as "re–quantization," a whole range of problems arise:

- the information beyond the 16th bit is lost — fine, low-amplitude audio detail beyond a 16–bit quantizing level disappears
- the noise floor increases to a 16–bit quantizing level
- for large amplitude signals, harmonic and intermodulation distortions appears
- for small signals, nonlinear distortions rise dramatically — small signals are precisely where modern analog devices are generally linear
- for small signals, stereo image instability appears
- at normal levels, unnatural falling off of reverberation occurs — smoothly fading into the noise floor becomes impossible

Especially for low level signals, the quantization errors correlated with and modulated by the original signal can be quite noticeable. Re–dithering technology helps to decorrelate the quantization error and consequently convert these discrete, unpleasant distortions into white spectrum noise. Proper re–dithering reduces or eliminates many of the problems listed above.

The decorrelation process is performed by adding a random signal to the high resolution audio signal prior to the re–quantization process. This random signal is called dither. Dither performs linearization of the nonlinear characteristic of the quantizer.

Re-quantized and dithered content sounds more pleasant, one can say "more analog." The price of this improvement is a noticeably higher background noise level. However, it is white in character and constant in level. Especially at low signal amplitudes, where the error is correlated to input signal, quantization noise sounds subjectively much more unpleasant than the even higher amplitude uncorrelated white noise produced by proper dithering.

The remarkable effect of a properly implemented dither algorithm is that a 16–bit signal can clearly reproduce tones at -110dB. After re–quantization without dither, there is simply no output for input signals below -96dBFS. All low amplitude detail is lost.

Why Shape?

However, with only dithering, we still have no influence on the level and audibility of the quantization noise. Therefore, in addition to dithering,

we use noise shaping algorithms. Noise shaping modifies the whole distortion spectrum: re–quantization plus dither noise. The noise power in the system cannot be removed but it can be distributed so the noise after shaping is significantly less audible. It is even possible to partially achieve an SNR or signal–to–noise ration below the original quantization noise floor. It is important to mention that the input signal is transmitted transparently through properly implemented dithering and noise shaping processes up to the re–quantization point.

In general, it is possible to use noise shaping without dither but a significant correlation between successive rounding error samples occurs; distortion and signal–dependent noise modulation. Therefore, we recommend always using noise shapers with proper triangular dither.

Shaping The Noise

Because the ear's sensitivity to low amplitude broadband noise is not uniform with frequency, we use noise "weighting" curves. Weighting, as mentioned above, redistributes the power spectrum of the noise. Dithered re-quantizing noise shapers approximate the inverse low amplitude noise audibility curve (in the old days, audibility curves were referred to as "Fletcher–Munson curves") allowing one to achieve the least audible noise penalty. Again, this leads to the counterintuitive conclusion that psychoacoustically optimized noise shapers allow for 19–bit audio performance from a 16–bit standard media delivery format. Signals with amplitudes as low as -120dB can be clearly detected.

The Interface

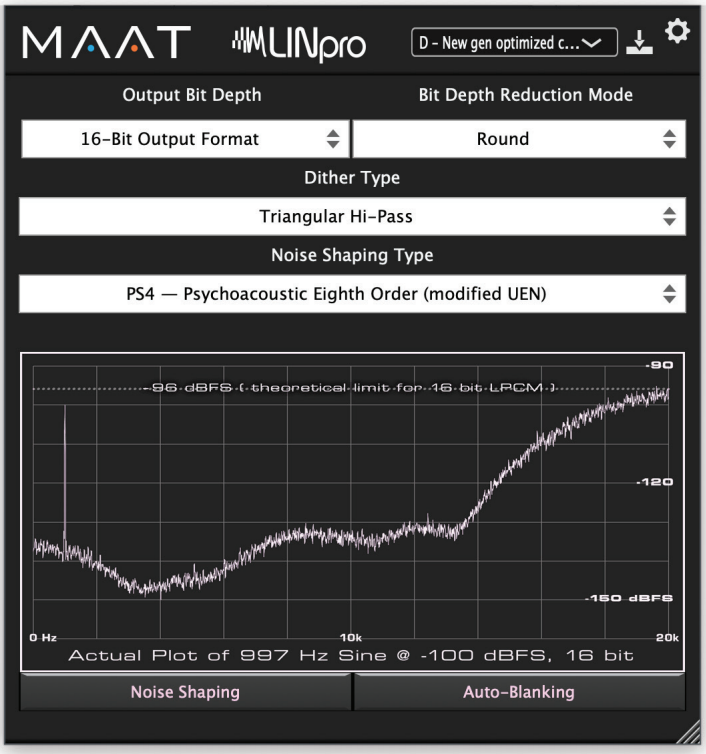

Figure 3: LINpro's user interface

The straightforward user interface contains a series of menus and parameter controls. All of LINpro's functionality is wrapped in a visually unobtrusive user interface. The plug–in is also very "light weight," demanding a minimum of CPU resources so it won't slow down your host. We've worked hard to make this an exceptionally useful and usable tool and we hope you'll dig it.

Dither Details

In the digital audio world, dither is noise added to reduce distortion. What? Adding noise to decrease distortion?! Yup, let's dig in some more…

Dither is an interesting and fundamental aspect of digital audio. Although not usually thought of as "DSP," re-dithering is definitely signal processing and so should be part of the overarching DSP umbrella. Back in the Distant Past, re-dithering was an integral part of our original digital telephone system, and much of the pioneering research on the subject was performed at Bell Labs.

According to information theory, you should add a pinch of noise to your digital audio every time you change it. Why is that? Well, take a look again at figure 3 below; that blocky looking mess that used to be a smooth and continuous low amplitude sine wave.

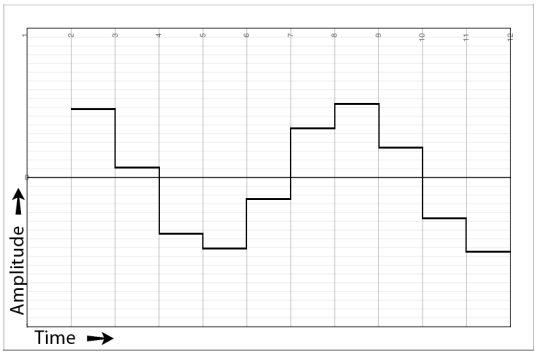

Figure 4: A digitized sine wave

Because a digitizer can only choose one discreet sample value during each sample (time) period, not half of one and half of another, adding a touch of noise helps to nudge the system to land on the statistically "right" choice. Injecting a hint of noise "linearizes" a digital conversion, reducing distortion. A properly designed digital system needs to add noise, called re-dithering, every time you change the gain (multiplication), or mix signals together (addition). That includes EQ, effects and pretty much anything else we typically do to our client's music!

The ideal amount of dither noise is about half of the step size of the digitizer, though it differs depending on the type of noise being used. In the case of good ol' linear PCM on which most all of our music is based, is 0.5 times 6dB or 3dB is correct for rectangular dither (RPDF). For triangular PDF, the ideal is about 4.7dB. Either way, half a bit is a very small amount of additional noise.

Without proper re-dithering, the distortion that results from any add or multiply is "correlated." This means it bears an audible relationship or subjective connection to the signal, the music. Trouble is, it's distortion, so it's a nasty sounding parody of the music and is usually considered to be quite unpleasant to listen to. When properly re-dithered, the signal is linearized, resulting in the ability to hear clean music within the noise floor at less than one bit of amplitude. In other words, that LSB, or Least Significant Bit, now can contain meaningful music "under" the noise.

There are several common types of dither, with the simplest being broadband noise, known as rectangular probability density function (RPDF) dither. Mathematically, this RPDF dither has uniform or equal probability of including any frequency within the audio band from 20 to 20 kHz. So, it has equal power at all frequencies and sounds "white." Because of its power distribution, it is quite audible, even at low amplitude and so should be used with caution. More common is triangular PDF or TPDF, which has less low frequencies and more high frequency energy. Because our hearing is less acute at higher frequencies, TPDF dither is generally less audible than RPDF dither.

As mentioned earlier, dither can linearize a digital audio system, significantly reducing distortion for low amplitude signals, and increasing the signal resolution well below the quantizer step size. Dither does add some additional noise. However, even if the wideband noise has about the same total power as the signal, the high Q filtering action of the basilar membrane allows our ear to clearly resolve sine waves well below, "buried" in the noise.

Shaped Dither

Back to shaped dither…A fairly modern invention is "shaped" dither. Shaping refers to the power spectrum of the dither signal, which is tailored or shaped for even less audibility. Noise shaping is a technology which reduces the audibility of noise resulting from dithering and re-quantizing. Noise shaping is accomplished not with equalization, but by processing the signal created by comparing the high resolution input and re–quantized output.

Our hearing mechanism is optimized for the frequencies around which the human voice is centered so, higher frequency sounds are less audible than lower frequencies. The more you shape the dither noise, shoving the dither energy toward higher and higher frequencies, the less audible it becomes…sort of. When used in larger than normal quantities, as with several generations of processing, or when set to extreme shaping, the dither signal itself becomes obnoxious, and is a bit like the sound of nails on a blackboard.

Noise shapers can modify the distortion spectrum after re–quantization in a way such that the signal–to–noise ratio can be improved over the theoretical quantization noise floor. The noise floor can be optimized to allow for the least possible audible noise penalty in a system. To quote a 1992 AES paper by Robert A. Wannamaker, the "…object (of shaped

dither) is to minimize the total perceived output noise power. For this purpose, a new perceptual frequency-weighting function is introduced which provides closer matching to the ear's measured sensitivity at high frequencies…" Noise shaping shifts the noise energy from high audibility regions where the ear sensitivity is highest to areas with lower sensitivity.

There are a variety of different shapes, from simple to very sophisticated, all based on psychoacoustics. An advanced, ninth order noise shaper can reduce the perceived output noise power by about 17 decibels. With some noise shaping schemes, performance equivalent to 19 bits of precision are possible within a 16–bit delivery format.

NOTE: Even if you decide to use un–shaped dither for your final deliverable, shaped dither is useful for monitoring prior to delivery.

Since shaped dither moves the subjective dither noise away from our most sensitive hearing region, you can used shaped dither during post–production to best perceive the subtleties of the music. Once all post and mastering is finalized and the content is ready for delivery, then you can choose the best dither method for the distribution formats and scenarios.

Some examples of shaped dither brands are Apogee's UV22HR, iZotope's MBIT+, the Power Consortium's POW-r, and Sony's SBM II, (Super Bit Mapping Two). MBIT+ and POW-r are adjustable, which is useful only when applying dither as an effect. The above mentioned branded redithering schemes also require a license, increasing cost. We have evaluated and improved on those legacy methods for LINpro.

Limitations

When processing is performed in a DAW, dithering should be applied at the very end of the deliver chain. In actuality, this is re–dithering since the ADC or analog–to–digital converter self–dithers the signal during initial digitization.

Notice I said "at the end…" This is important in that dither should only be re–applied when the content is moved out of your DAW, either by an AES/EBU or other streaming spigot, or as a file or set of stems.

In the 1990s, manufacturers began developing re-dithering schemes that provided the linearization without the audible addition of noise. The first of the resulting products was Apogee's original UV-22, mentioned above, an effective if brute force approach. A handful of better performing and better sounding spectrally shaped re-dithering products followed from several manufacturers.

Unlike the early days of digital audio, many, if not most consumers now listen to music through a largely digital signal chain. The DAC may be in their phone, embedded in their headphone cable or inside the "smart speaker" on the kitchen table. As legendary mastering engineer Bob

Olhsson mentioned when asked about his re-dithering choices, “In most cases, people are never going to hear (my product) without some kind of downstream digital processing.” Theses days, only audiophiles have integrated amps or stand–alone preamps with analog gain staging. Re-dithering with anything other than TPDF does not “translate” well in our modern, all–digital world, as shaped dither increases distortion rather than reducing it.

With that fact in mind, one can also think of dither as just another effect, a “color” paintbrush in the DSP toolbox. No re-dithering at all produces a crunchy aspect that one well known mastering engineer employs as part of his signature sound, while more that a half step adds a veil of extra noise that may be perfect for your particular client.

Our point is: rely more on your ears than your eyes. Take some time to listen to the different dither choices available in LIN, and learn what they do to a piece of reference music played through reference consumer systems. It’s easy to do…grab some audio with which you’re familiar, and gain adjust it down by 8 bits or more, that’s at least 48 decibels. Then, instantiate LINpro and turn your monitor gain up enough to hear what the different settings are doing.

To make your dither choices very audible, increase the amplitude. This is easily done by reducing the Output Bit Depth setting. Since 8–bit digital audio holds only 48dB of dynamic range, LINpro’s output at that setting is around 48dB below full scale or 0dBFS. We also give you an 8–bit output word length setting to help you differentiate between the effect each setting has on your material.

Once you’re familiar with the setup and the character of the dither generator, try a much larger gain reduction so you are “down in the mud,” where dither really counts. Be careful as this requires a huge amount of monitor gain or, use LINpro’s 8–bit mode. Then, compare TPDF to rectangular and noise–shaped in an all-digital playback scenario compared to that much rarer all–analog playback situation. Remember to gain-reduce the digital playback chain so the playback engine processes the signal one more time. Then, level match the analog chain and start comparing.

In Use

Optimized Choices

The collection of controls include output bit depth, bit depth reduction method, dither type and noise shaping type. Noise shaping bypass and auto–blanking round out the settings choices.

In general, dither and noise shaping types can be chosen in any combination, but not all combinations are recommended. To avoid sub–optimal results, we only allow triangular and triangular high-pass dither in combination with any noise shaping type. For comparison purposes, a no–dither choice is available. In this case, rounding only at the request-

ed bit width is applied. Truncation, which causes significant re–quantization distortion and DC offset, is included for educational purposes.

New To Dither?

8-bit mode is also implemented for training purposes, to provide easily perceived experimentation with dither types and noise shaping choices. Dithering to 24 bits is useful when moving from 64–bit floating point to fixed point. It also helps with 32–bit floats, but only just. The 32–bit floating point format has only a 24–bit mantissa, and therefore there's not that much information to preserve.

> NOTE: When not using noise shaping, the use of triangular high-pass dither is a recommended practice.

If all this dither madness is new to you, we suggest you start out in 8 bit mode, with dither and noise shaping disabled. Set up a project just for training your ears, and add several tracks of different genres. In your DAW's Master section, reduce the gain of your clips by 48dB. Then, insert LINpro and, downstream from that, insert/dial up gain makeup of +48dB. All that will allow you to better hear the result of your LINpro manipulations…

Again, start out in 8 bit mode, with dither and noise shaping disabled. Listen to how the program material "sounds." Get used to that, then enable truncation and again listen to how *that* sounds. Now, switch from truncation to rounding. Note that dither is still disabled…Once you have a solid sense of how those different settings affect sound quality, enable the simplest possible dither setting; TPDF with no shaping. Compare that effect to high passed TPDF. Next, compare the various shaped noise options…With all that aural experience, you are now ready to "dive into the deep end"!

Choice Is Good

All of LINpro's psychoacoustic noise shapers have been precisely optimized for the 16 bit/44.1 kHz PCM format, also known as "Red Book". They profit from inherently lower quantization noise in the audible frequency range, and are also applicable for 48 kHz deliverables. Of course, dithering and 1st order noise shaping work perfectly with any sample rate and any playback chain, analog or digital. For deliverables at 88.2 kHz and higher sample rates, use the default TPDF dither with 1st order noise shaping.

Output Format	Bit Depth Reduction	Dither Type	Shaping	Comments
8 bit	No Change	None	None	
16 bit	Round	TPDF	NS_1	1st Order Noise Shaping
24 bit	Truncate	HP–TPDF	NS_2	2nd Order Noise Shaping
		RPDF ‡	NS_3	3rd Order Noise Shaping

	PS_1	Psychoacoustic 1
	PS_2	Psychoacoustic 2
	PS_3	Psychoacoustic 3
‡ – Psycho-acoustic noise shaping is not available	PS_4	Psychoacoustic 4

Table 1 : Setting choices

Output Bit Depth

For your deliverable word length, 8, 16 or 24 bits are available. An 8 bit setting provides a very obvious audible result, making it easy to arrive at an optimal setting for your content.

Bit Depth Reduction Mode

In the words of our Head of Development, this setting, "reaches into the bits of every float and deciding what to do with them." Effective precision determines what the *really* least significant bits contain. If the Output Bit Depth is set to any value other than 24, then this menu decides the contents of the defined LSB and all bit cells past that.

The choices are Round or Truncate. Of those settings, only Round(ing) will produce optimal results. Truncate will result in loss of all information in the lower order bits. Rounding without dither is not recommended because the lowest amplitude information is lost.

The Round setting rounds to the desired word length. At this setting, the specified LSB represents a single bit rounding of all the lesser bits past it, those carrying information about the signal that are lower in amplitude. If Output Bit Depth is set to 16, then the 16th bit will contain a value that results from rounding the 17th through 24th bit. The 17th through 24th bit will then be "padded" with zeros or "null data." Round is the default as rounding the excess bits produces the most linear result.

The Truncate setting simply "chops off" or truncates all bits past the LSB for the specified word length. If Output Bit Depth is set to 16, then the data contained in the 17th through 24th bit is discarded, and those bits are then filled with null data (zeros).

For a slightly "crunchy" sound, use the Truncate setting in conjunction with Output Bit Depth set to 16, plus TPDF type dither and zero dither gain. The zero gain setting means no dither noise will be applied and, because of truncation, LINpro will simply toss out the 17th through 24th bit. This means the 16th bit will not contain any low amplitude information, as it would if dither had been injected. With dither, information in the 17th through 24th bit would have preserved. The crunchiness is caused by the low amplitude loss of detail generated when truncation is employed.

Dither Type

Four choices: that include Rectangular or Triangular Probability Density (RPDF/TPDF), plus none or High Passed TPDF.

NONE	simple truncation, no dither; not recommended because of asymmetrical quantization error
TPDF	dither with triangular probability; spectrally flat, additional 1.8dB noise floor, but no noise modulation, highly recommended especially with noise shaping
HP–TPDF	dither with triangular probability and noise shifted towards high frequencies; results in reduction of audibility, universally applicable, highly recommended especially when noise shaping is not used
RPDF	dither with rectangular probability; historical, not recommended because the 3dB higher noise floor is modulated by the input signal

NOTE: When dither is disabled or set to NONE, then noise shaping (see below) is also disabled.

Noise Shaping Type

Shaped noise is not broadband noise that has been EQ'd. Shaped noise is derived from the method used to generate the noise, not how the noise is post–processed. Unlike plain TPD noise, which affords you an extra bit's worth of signal underneath the noise floor, shaped noise can add two bits or more of perceptually clean signal embedded in the noise floor.

NONE	No noise shaping; resulting noise spectrum is more audible than a shaped dither.
NS1	First order, non–modified error feedback; results in a continuous shift of noise to higher frequencies, allows perceptual preservation below the theoretical quantization noise floor for low frequency range (up to ≈3 kHz) with a slight increase in noise level at higher frequencies where audibility is much lower

Starting with 2nd order error filtering, we enter the range of weighted noise shapers. They allow adjusting noise according to modern psychoacoustic research. Audible noise can be reduced extensively by about 17dB, almost 3 more bits. After re–quantization of a high resolution signal, perceptual audio quality equivalent to 19 bits can be retained within only 16 bits of word length.

The most common curve in psychoacoustics is A–weighting. However, this curve is not satisfactory to approximate the low amplitude sensitivity of human hearing. Since dither is applied to optimize the LSB or Least Significant Bit, a better choice is to employ E–weighting,

which is based on the inverse of the ISO 226 15 phon equal loudness curve modified for diffuse field.

NS2 — 2nd order error noise shaping; approximates the modified E–weighting curve; offers almost two bits of improvement in apparent SNR, in the range of approximately 3-4 kHz, the highest loudness sensitivity

NS3 — 3rd order error noise shaping; a closer approximation to the E–weighting curve then 2nd order

Psychoacoustically optimized noise shapers have been designed to modify the noise spectrum to be least audible across the entire audio spectrum.

PS1 — Psychoacoustically optimized 5th order error noise shaping; follows an improved E–weighting curve based on a 15 phon equal loudness curve but with even better approximation of the low and high frequencies then is possible with the modified E–weighting curve

PS2 — Psychoacoustically optimized 9th order error noise shaping; based on the F-weighting curve, PS2 is a further refinement of the modified E-weighting curve, perfectly fitting splined ISO data in the range of 40 Hz to 16 kHz. This curve can provide up to 17dB of perceived improvement relative to an unshaped spectrum, almost 3 more bits in the SNR

PS3 — Psychoacoustically optimized 8th order error shaping; newer generation based on the ITU 486 weighting curve with modified response above 12 kHz according to recent observation on the ear's sensitivity to high frequencies

PS4 — Psychoacoustically optimized 8th order error shaping; newest generation based on sensitivity to noise than to pure continuous tones, this curve is referred to as uniformly-exciting noise (UEN) at threshold which describes the minimum audible SPL of critical–band–wide noise on frequency axis. The curve was further modified to represent the power spectrum of UEN at threshold in a free field

Noise Shaping Enable/Disable

This button disables noise shaping for comparison purposes. Instead of using the Noise Shaping Type menu, you can turn noise shaping on or off to instantly hear the effect.

Auto–Blanking

The blanking option, enabled by default, automatically mutes the dither generator after a preset 11 milliseconds of silence. This prevents you from injecting noise into what would otherwise be "digital black" or null data.

Delivering From Floats To 24 Bit Fixed

What settings are best when delivering from single or double precision floats to 24 bit fixed? Noise shaping typically is used to perceptually "push" the noise floor below what/where it would be with dither alone. As mentioned previously, noise shaping with dither at the 16th bit can move the effective floor lower, over the entire audio passband, by up to ~3 bits. So what does that do for you if you're dithering at the 24th bit? 24 bits encodes 144 dB of dynamic range, which is well beyond either audibility or reproducibility on normal playback gear at normal playback volumes.

With HRA releases, if the data is just truncated at 24 bits rather than properly dithered to 24 bits, correlated noise will result. Vicki R. Melchior, current chair of the AES High Resolution Audio Technical Committee, states that there is "…studio evidence that the truncation noise can become audible because it will be correlated to the signal. It's very hard to show with listening tests, but multiple mastering engineers have heard it." Melchior recommends rounding to 24 bits, adding HP-TPD dither at the 24 bit but not bothering with noise shaping.

Preferences

The gear icon in the upper right corner opens the General pane and Info pane. The General pane contains one button that, when clicked, saves the current plug–in state as the default state going forward. The Info pane contains credits, an Activate/Deactivate License button, and a button to open this user manual.

Presets

To the left of the Preferences gear icon is a User Preset menu and a right flanking Save button. The Save button allows you to store a state, while the Preset menu allows you to recall any factory or user–saved settings. If an existing preset is modified, an asterisk appears after its name.

Holding down the option or alt key switches the Save button to a Delete or remove (preset) button. Simply select a preset from the menu. Then, holding down the modifier key, click the button.

Preset A: Recommended for audio material of average dynamic range; converts re-quantization noise to subjectively more pleasant "analog" noise

Preset B: Enhanced by smooth wideband noise shaping with greatest effect at lower frequencies; works properly for any audio formats

Preset C: Enhanced by a complex psychoacoustically optimized curve (tones hearing threshold); recommended for re-quantization of complex program to the 16-bit Red Book Compact Disc format; can subjectively retain up to three additional bits of original within a 16-bit deliverable

Preset D: Enhanced by a new generation, complex psychoacous-

tically optimized curve (noise hearing threshold); recommended for re-quantization of high resolution 24-bit material to the 16-bit Red Book format; can subjectively retain up to three additional bits of original within a 16-bit deliverable, greatly increasing perception of very low amplitude detail, reverberation and spaciousness

LINpro presets are stored as cross–platform XML files. For migrating presets from one host to another, you'll find them at…

macOS	/Applications/MAAT/LINpro/Presets
Windows	C:\Program Files\MAAT\LINpro\Presets

Specifications

System Requirements

- macOS: 10.9 and newer, 64 bit only
- Windows: 7 and newer, 32 & 64 bit
- 4GB RAM minimum

Supported formats

- 44.1 to 384 kHz sample rate
- AAX, AU, VST 2, VST 3

Updates

Please always use the latest version of the software! You can find your current version on the Info Tab of the back panel. You can download the latest version simply by visiting:

http://maat.digital/support/#installers

For optimal security and stability, you should always stay up to date with Operating System revisions, and we keep up with compatibility changes to our products. We also continue to optimize for reduced CPU load, and this very user manual gets its own improvements.

To stay up to date with the latest version and product releases, please subscribe to our occasional newsletter. You'll find a opt–in form on our Contact page, or sign up for a copy of our free and very handy 2BusControl plug–in which will also subscribe you. Don't worry, we know you are busy so we only send out an average of 10 or 12 mailings a year.

Support

For product support, please visit:

https://www.maat.digital/support

License Central

License Central, located in MAAT's Shared directory, is a free utility that validates, repairs and logs your MAAT licenses. It also displays what MAAT products you have installed and assists in downloading both updates and demos.

License Central lists all of our products, and shows you:

- If a product is installed and what licenses you have [CLOUD, SUBSCRIPTION, TRIAL, OFFLINE and TEMP OFFLINE]
- What versions are installed, and version installers are available for download

It also provides:

- A one–click download of an update or a demo
- A button to activate or deactivate any of your licenses
- A copyable list of all your Product Keys
- One–click generation of the MAAT diagnostic report

It also fixes license issues automagically.

In Use

When you launch License Central, it scans your host computer and then attempts to match up each product found with a Product Key on our license server. That scanning happens in real time, and the process is shown as a progress bar. When quitting, this process happens in reverse.

Once all the licenses are validated, a list on the left displays all MAAT products, including License Central. Products that are not installed are grayed out, and installed products are displayed in high contrast. To the right of each product entry is a check mark, which indicated that the product is licensed.

Selecting a product entry displays information about the installed version and if an update is available for download.

Controls

There are two persistent buttons along the top. At upper left is Check for Updates, which refreshes the list of installed products, and pings our server again for currently available versions to download. It also recreates your Product Key list.

At upper right, the Activate/Deactivate button brings up the familiar blue and orange MAAT license window, allowing you to deactivate or return your license to the Cloud, and to switch to a 30 day temporary offline license. If you have requested a 365 day full offline license, use the Create License Request function in the MAAT license window, and have received your ".maatc" confirmation file, you can also take your

license offline. Finally, you can also paste in a Product Key to activate a new license.

> NOTE: Switching to 30 day temp offline requires that your license be already activated. 365 day full offline licenses require a license confirmation file as supplied by MAAT's support department.

If any product is out of date, a third Download Update button will appear when that product is selected. This button is an express method for updating you MAAT products.

At bottom left of the License Central window is an Auto–Refresh check box, which is selected by default. Deselecting that will prevent License Central from frequently scanning your host for changes including new installs. It starts a scan every 3 seconds, waiting for completion of that scan, then begins a new 3 second countdown before starting the next scan. If you find that License Central is interrupting you workflow while scanning, then disable Auto–Refresh.

Preferences

In License Central, the Preferences' gear icon brings up the Preferences window with four tabs.

Paths

For those who set up alternate directories for the management of their plug–ins in a DAW, two alternate plug–in paths can be defined. These directories will be also included when License Central searches.

Keys

This tab lists all of your Product Keys that are known to our license server. Please take a moment to copy all of your Product Keys, paste them into a text or word processor file, and print out a hard copy. Your Product Keys are your proof of purchase and, without them, your purchase will not work and we cannot provide product support.

Troubleshooting

The Troubleshooting tab has one button, which generates a Diagnostics Report. This report, automatically written to the Desktop, gathers useful, non–personal information about your machine's state that our support team can use to diagnose any problems.

Info

As with all our products, the Info tab lists credits, copyright information, and the version number you are running.

Share The Love

Would you like to help us in our quest for better sound quality? Help support MAAT by sharing the love…Like us on Facebook!

https://www.facebook.com/maatdigital/

and Twitter too:

https://twitter.com/maat_digital

Please tell your friends and colleagues about us. We really appreciate it, and thank you for supporting better quality audio.

Support

For product support, please visit:

https://www.maat.digital/support

Would you like to help MAAT? Help support MAAT by sharing the love…Like us on Facebook:

https://www.facebook.com/maatdigital/

and Twitter:

https://twitter.com/maat_digital

and tell your friends about us!

Credits

Product Concept, Design & Management:

OMas

Programming:

Adam, Agent of Shield

UI:

Friedemann “I love it!” Tischmeyer

Thanks Dr. Christoph Musialik for his contributions in shaping the underlying engine on which LINpro is based. Thanks also to all the scientists who, 30 years ago, performed the fundamental research work that led to the development of all the sophisticated dither and noise shaping theories on which later investigators built, including ours. In turn, these pioneers paved a path to much better audio quality at shorter word lengths: Gerzon, Craven, Lipshitz, Vanderkooy, Wannamaker, and many others who have continued their investigations into this topic.

Appendix One — The Math

For a deep dive into the basic theory of dithering, look no further than Robert A. Wannamaker's 2003 Doctoral thesis entitled, "The Theory of Dithered Quantization". A PDF version is available for your profound elucidation here:

http://www.robertwannamaker.com/writings/rw_phd.pdf

Appendix Two — The Curves

Though our understanding of how dither is perceived by the human auditory system has evolved over time, the fundamentals have not and were first described in the modern sense in 1961. I was first exposed to dither as a high schooler, where the gas gauge in my car required mechanical dither, aggressive tapping on the gauge face, to yield an accurate reading. Later, I dove in deep while working for Sonic Solutions, manufacturer of the first professional DAW. Thank you, James A. Moorer for that!

For your deep dive, I suggest yuou start with, of all things, the Wikipedia page:

https://en.wikipedia.org/wiki/Dither

As a follow up, here are some loosely related additional writings from prominent audio practitioners you might find of interest…

Eelco Grimm once wrote:

"What's the difference between genuine dither and noisy lower bits?

The amplitude probability density [or power distribution versus frequency — Omas]. 'Genuine' dither has a

probability density function from a particular, rather restricted family. Sripad and Snyder proved that the Fourier transform of the PDF (Probability Density Function) of the dither must go to zero at periodic distances from the origin. Practically, this means that proper dither is the sum of one or more random signals having rectangular PDFs of width Q/2 (where) Q is the LSB size. The number of that are summed (which) determines how many moments of the truncation error are de–correlated from the input signal. We use two in audio because that de–correlates the power (as well as the mean) of the truncation error. Adding two independent rectangular PDF dither components yields a triangular PDF (TPDF) dither signal.

Arbitrary noise doesn't meet these requirements. It still helps, but if you listen to only the truncation error signal, you'll hear traces of the original program."

Bob Katz once wrote:

"My usual solution to 'cover up the distortion' problem is to use one of the many different flavors of dither I have available to me. One of them usually does better than others. One just has to be aware of the tradeoff in apparent low level ambience and resolution. You never get something for nothing…the hotter the dither, sure, you might hear less distortion and get a 'smoother' sound, but the low level ambience slowly leaches away."

In response to Bob, Daniel Weiss wrote:

"There is only one right dither amplitude for any given dither flavour. Dither is not 'covering' distortion, it simply makes the quantization error to be uncorrelated from the audio signal and thus there is no quantization distortion anymore - provided the dither level is high enough. Dither level may be higher than necessary, but that does not 'cover' more distortion, it simply adds more noise.

I don't see why manufacturers make the dither level variable by the user. [We at MAAT removed that option in LINpro and LINone] Makes no sense - except maybe when the signal already has some noise floor which is

suitable for dithering, i.e. a relatively high noise floor." [A noisy ADC would be a good example of that; confirmed by David Josephson & Jim Johnston. Johnson adds that, "Yes, delta-sigma convertors have relevant dither issues." — OMas]

Marc Lindahl once wrote:

"If you have 'shaped' dither, then you can run into problems with compromised headroom…Also, 'perceptual' shapes can go from imperceptible to perceptible if 'stacked'…other than that there's probably more danger from stacking quantization noise, since it's kinda like harmonic distortion, and can quickly emerge from the 'mud'…"

Appendix Three — The Curves

These are the same illustrations that appear within the plug–in. They are not MATLAB simulations, they are actual plots generated by LINpro's optimized dither engine. Notice that, in all but the first graph, the -100dBFS input signal or excitation (the vertical spike at 997is preserved even though the word length has been reduced to 16 bits.

A 16–bit word can only encode or store 96dB of dynamic range. As mentioned earlier, any information quieter than -96dBFS *should be lost* when reducing the word length to 16. The first plot, an empty one, shows that all information residing in the 17th through 24th bit in a 24–bit word is gone *without proper dithering*. For all other illustrations in this appendix, the input signal is preserved.

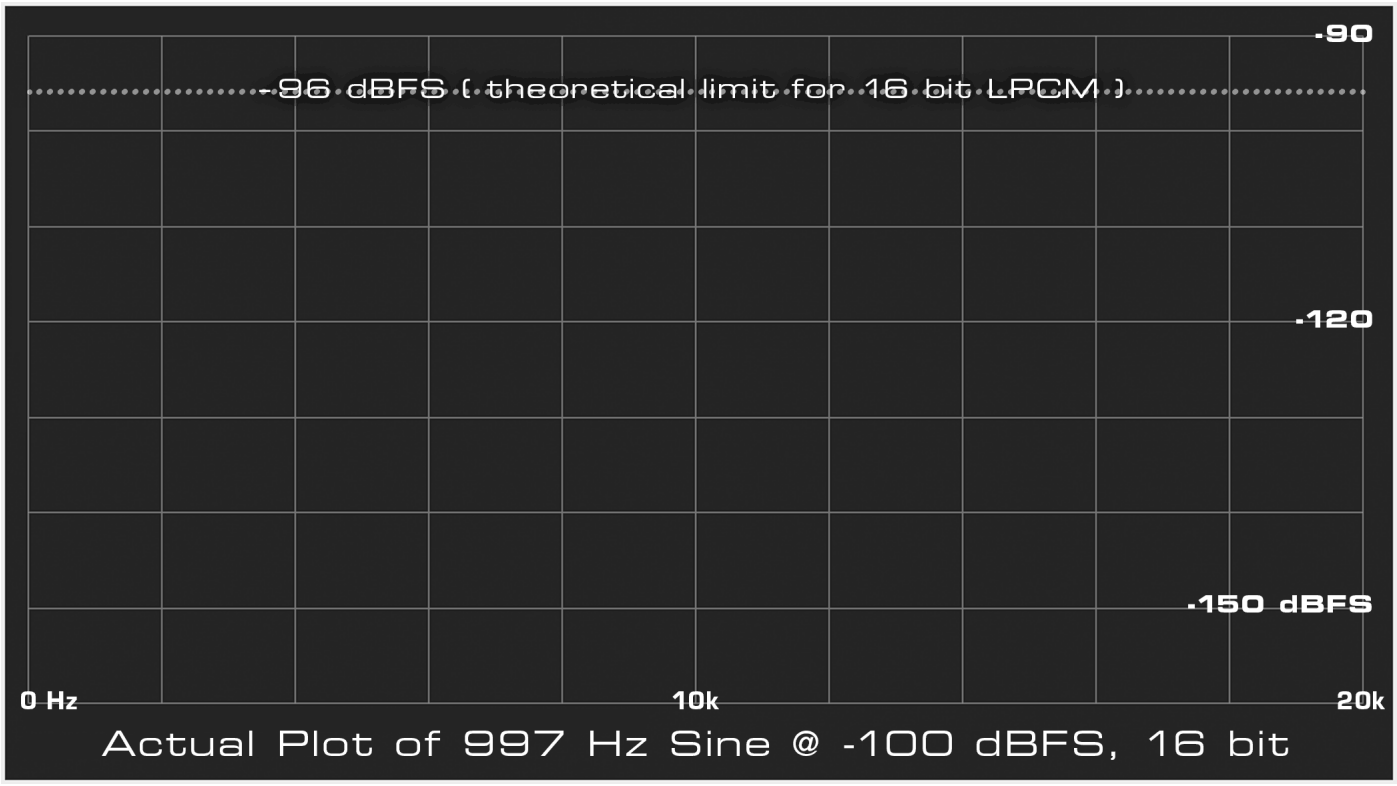

Figure A1 — No dither, no noise shaping = no low level information preserved

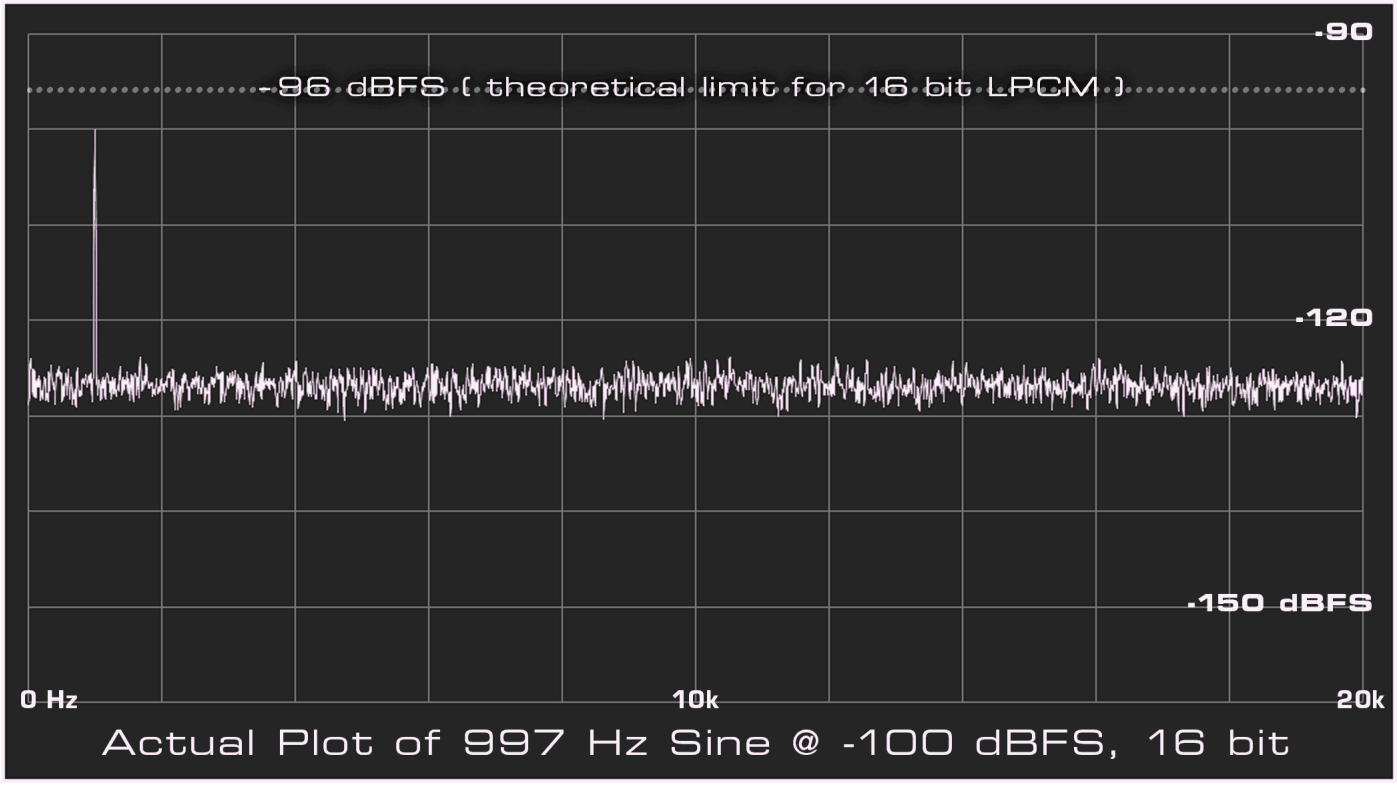

Figure A2 — Rectangular dither, no noise shaping

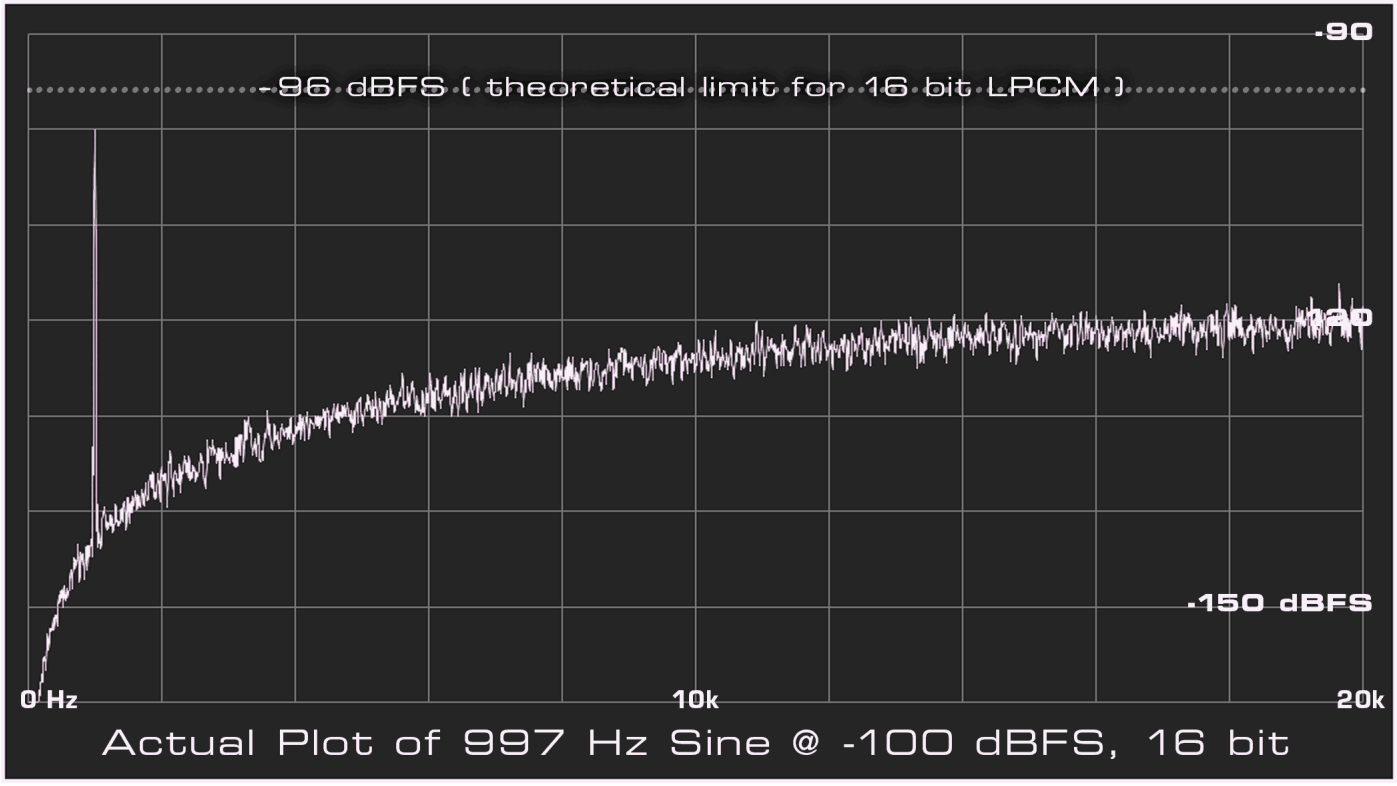

Figure A3 — Rectangular dither, NS1

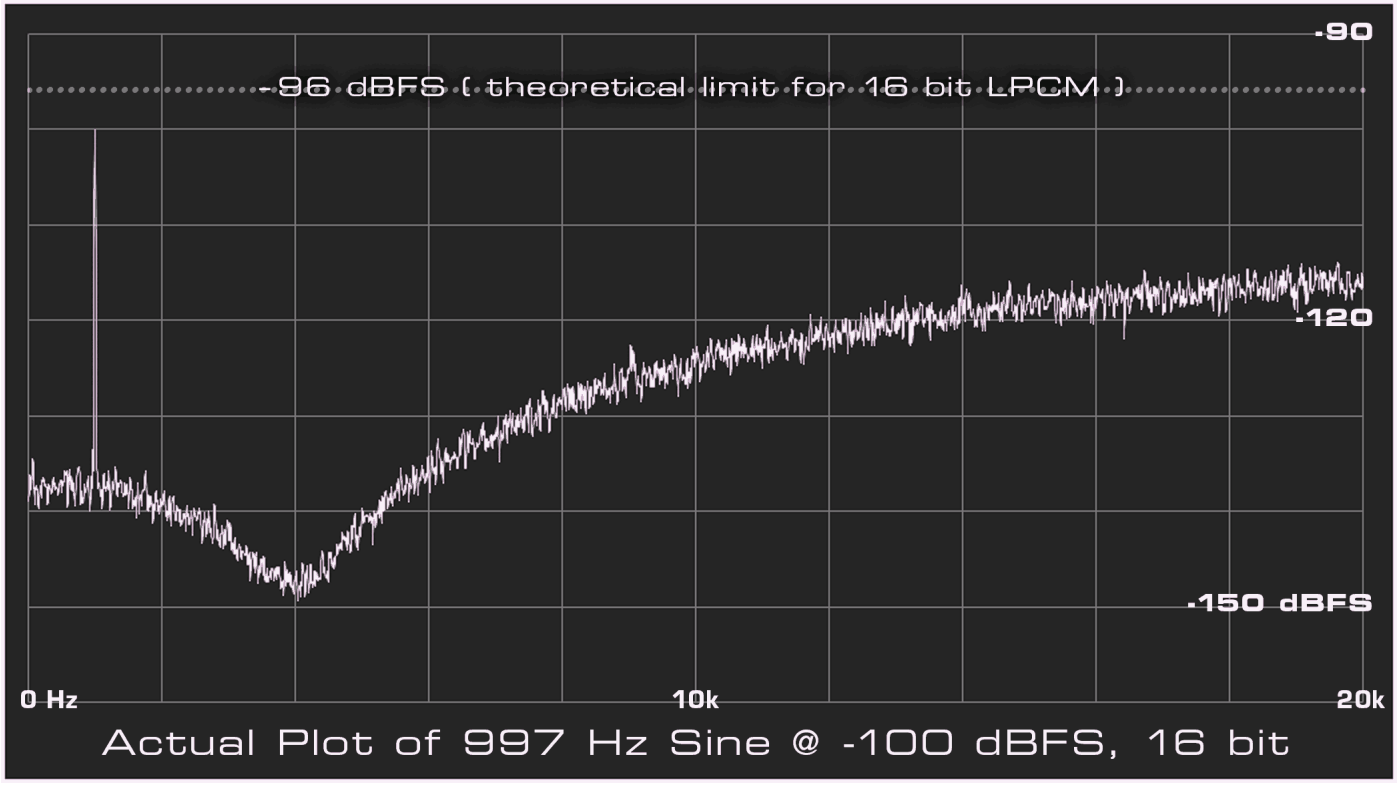

Figure A4 — Rectangular dither, NS2

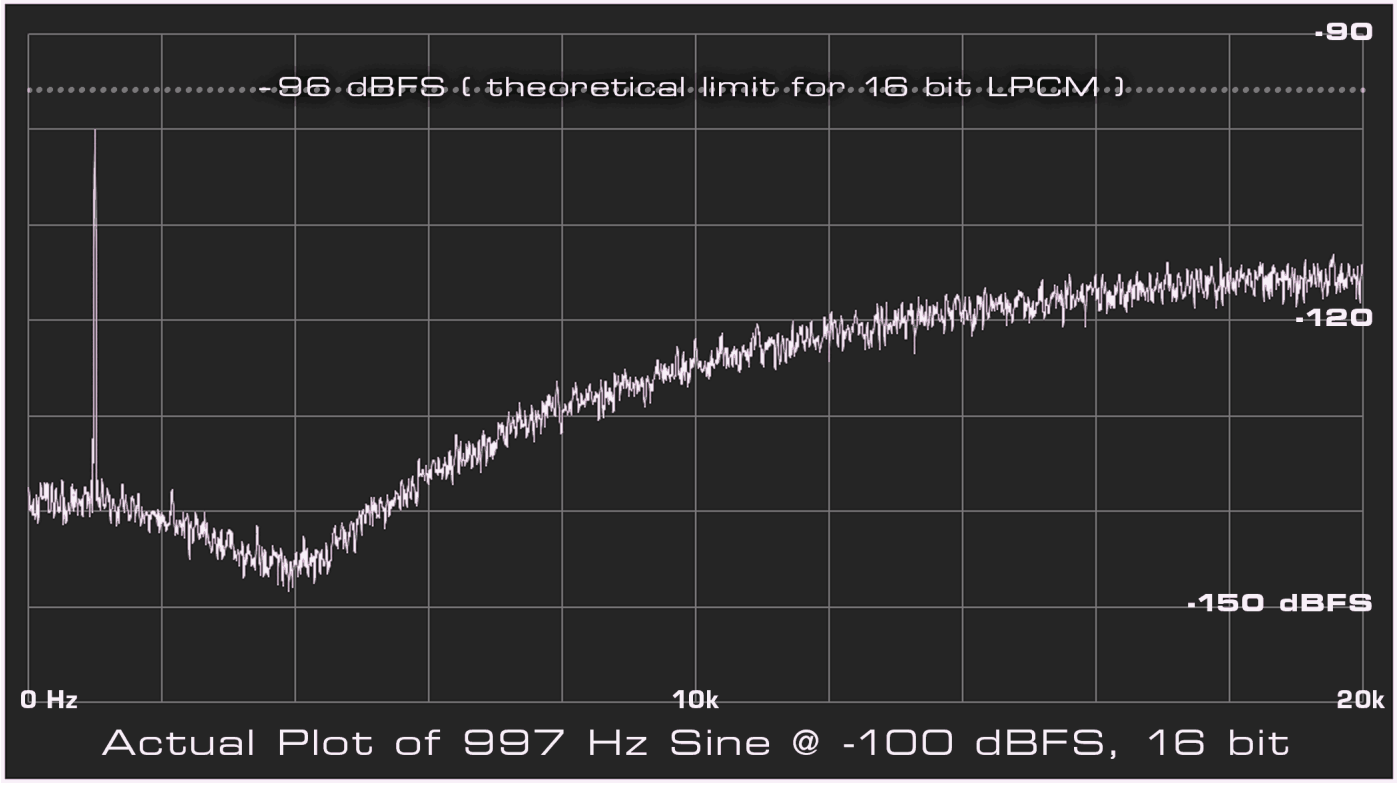

Figure A5 — Rectangular dither, NS3

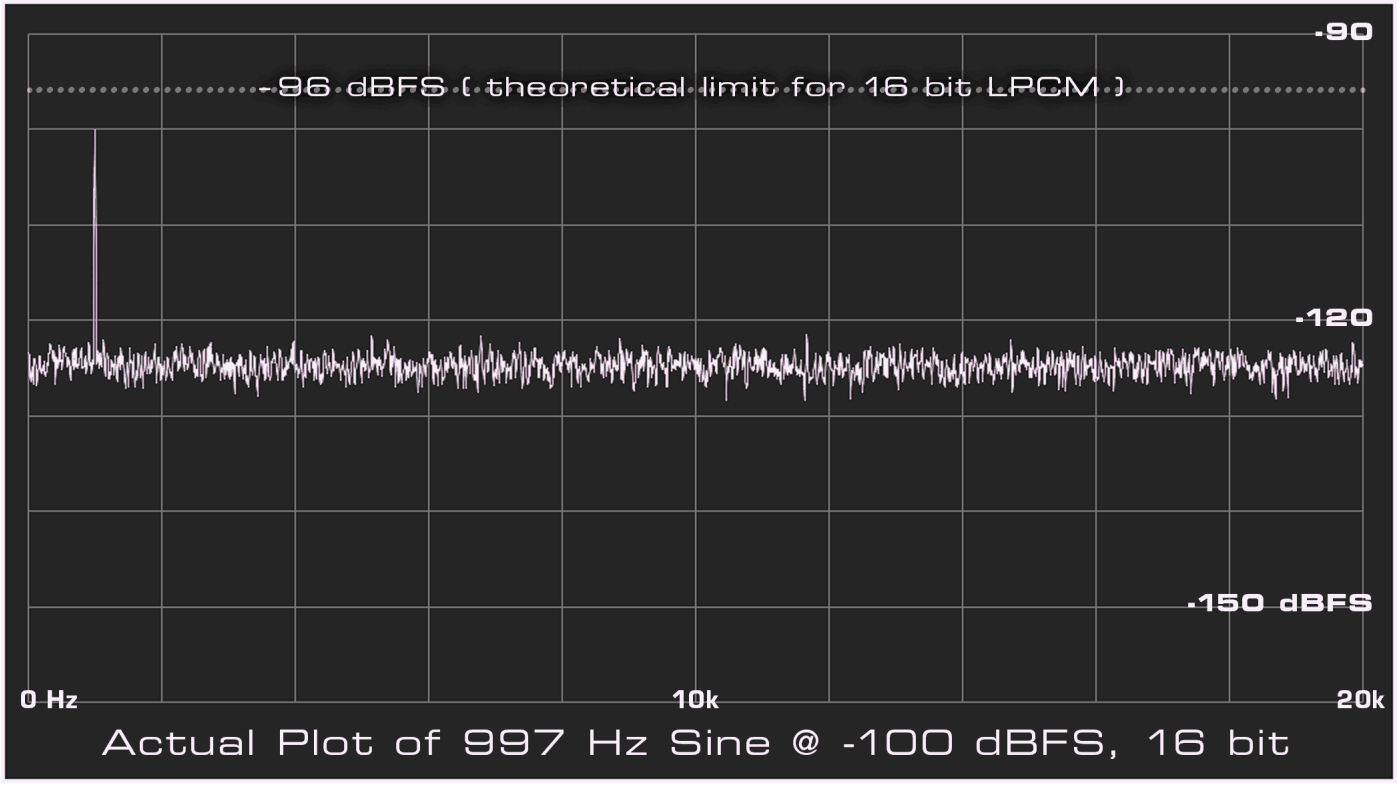

Figure A6 — Triangular dither, no noise shaping

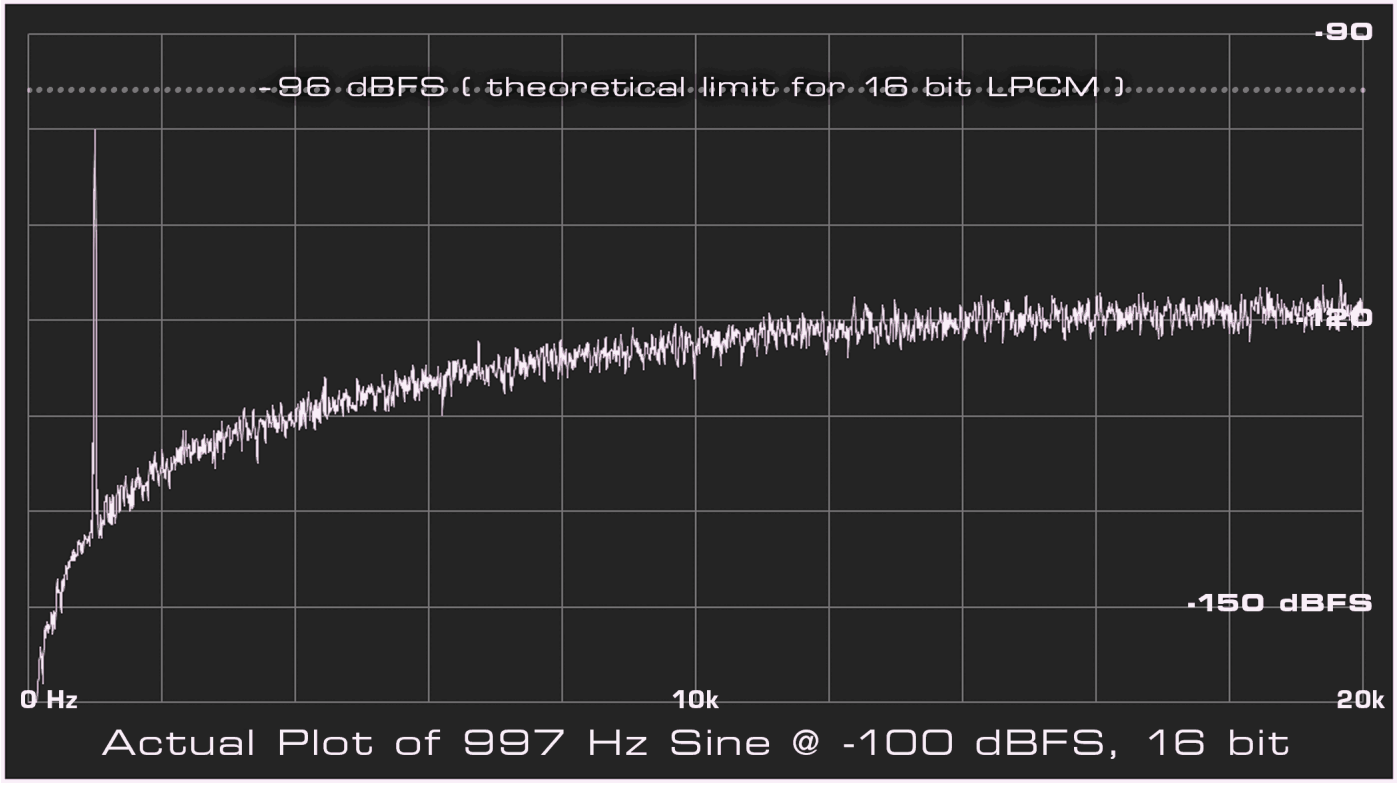

Figure A7 — Triangular dither, NS1

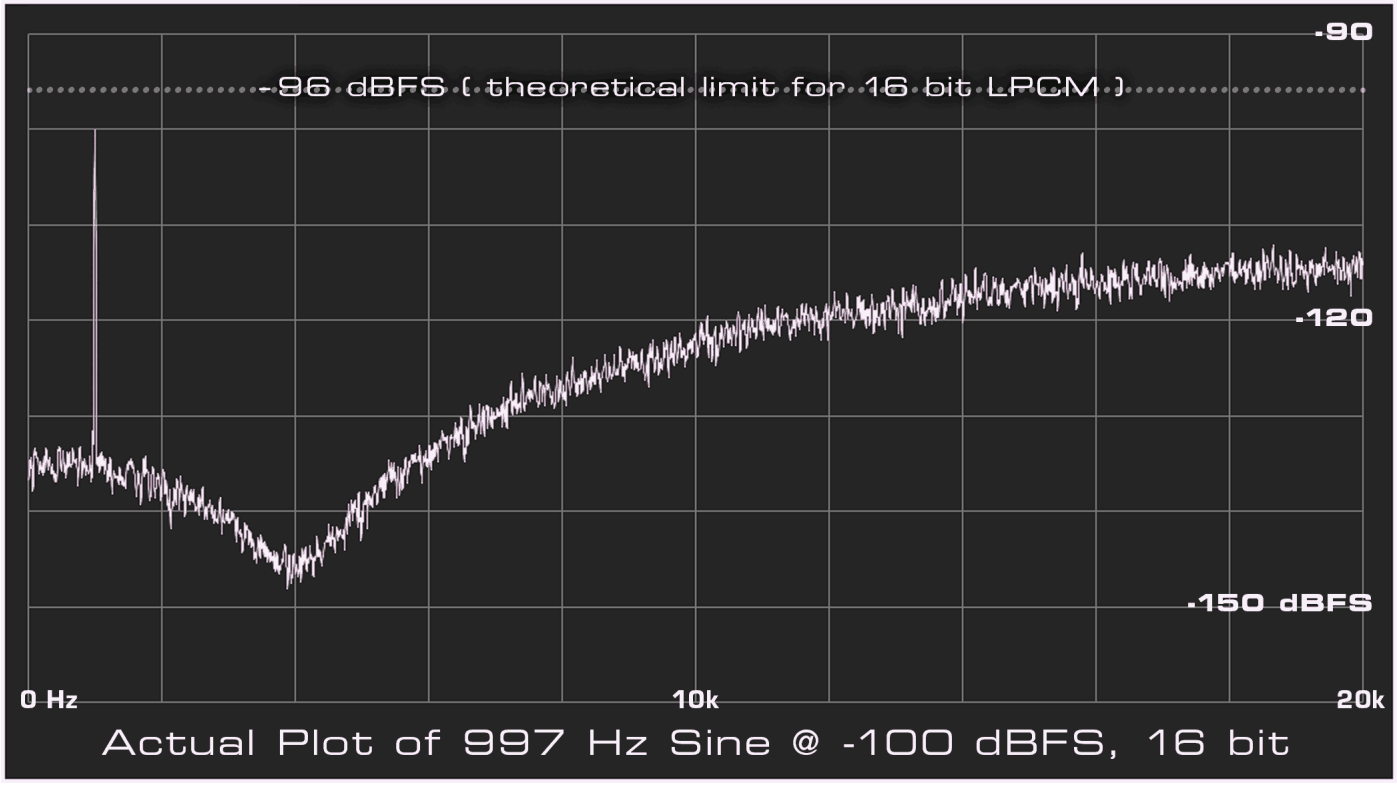

Figure A8 — Triangular dither, NS2

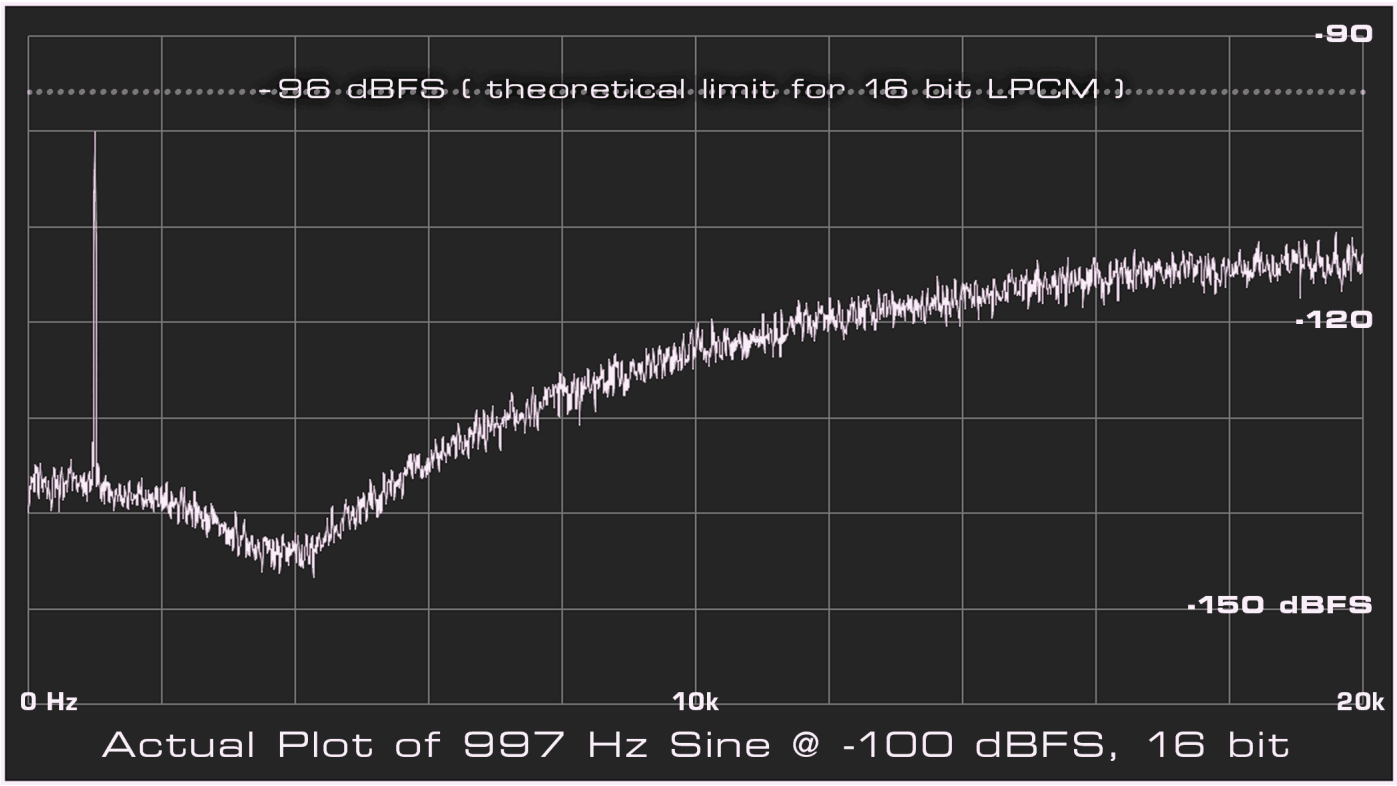

Figure A9 — Triangular dither, NS3

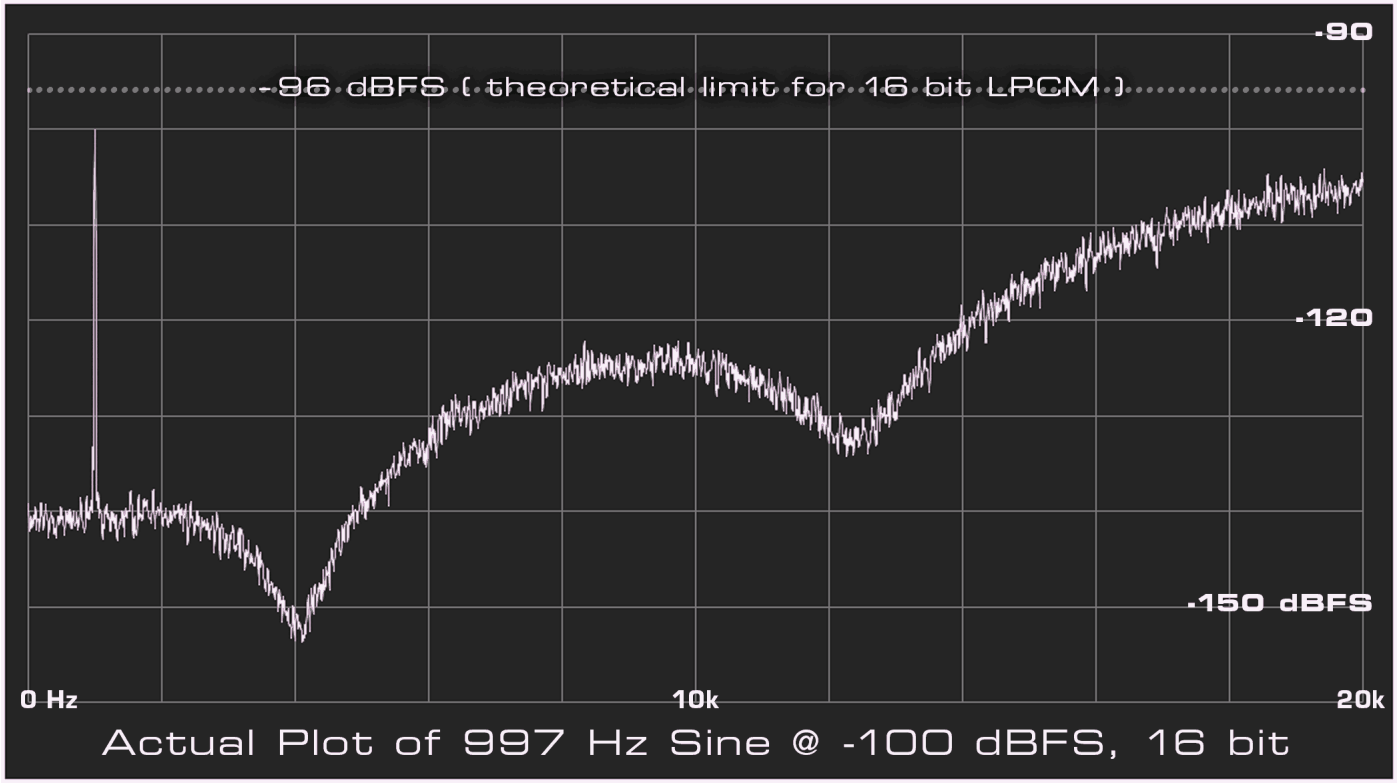

Figure A10 — Triangular dither, PS1

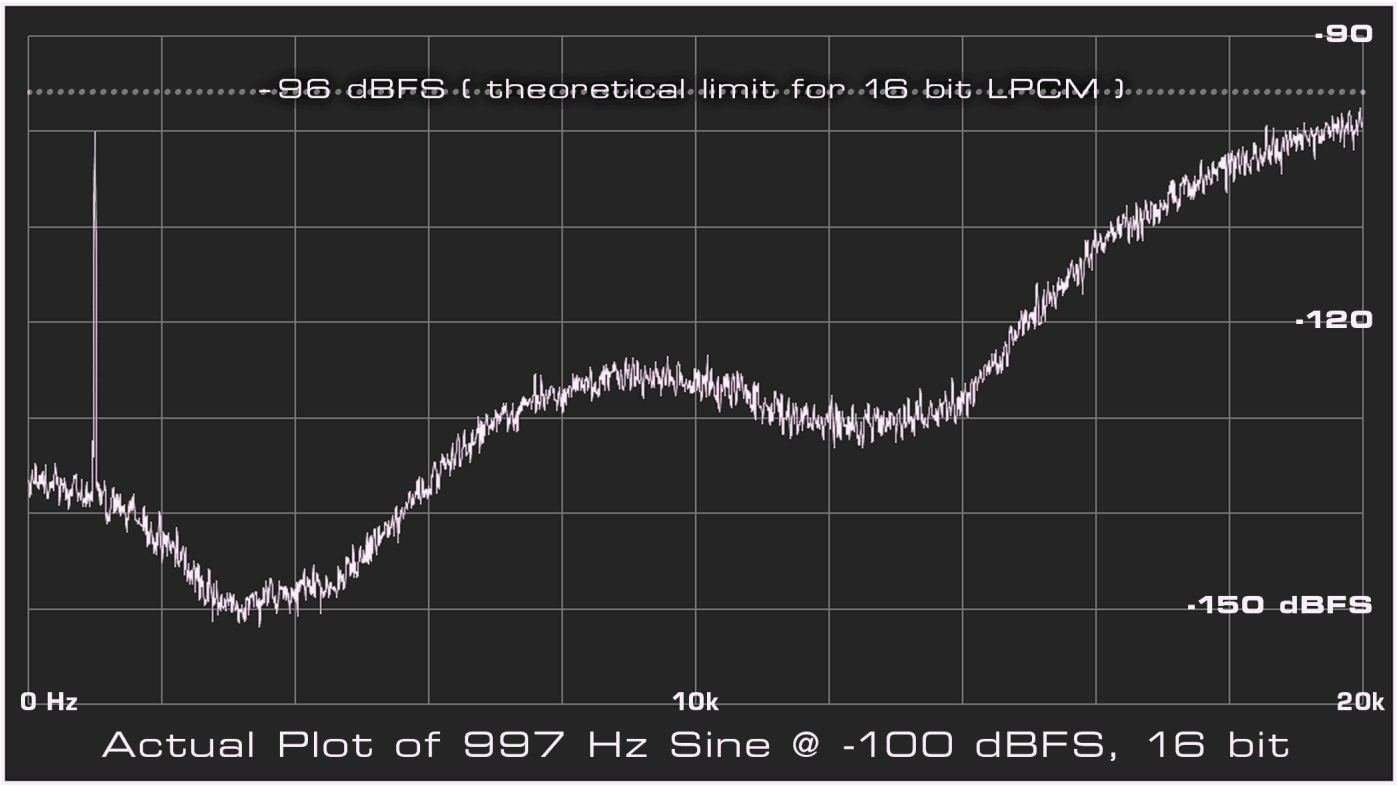

Figure A11 — Triangular dither, PS2

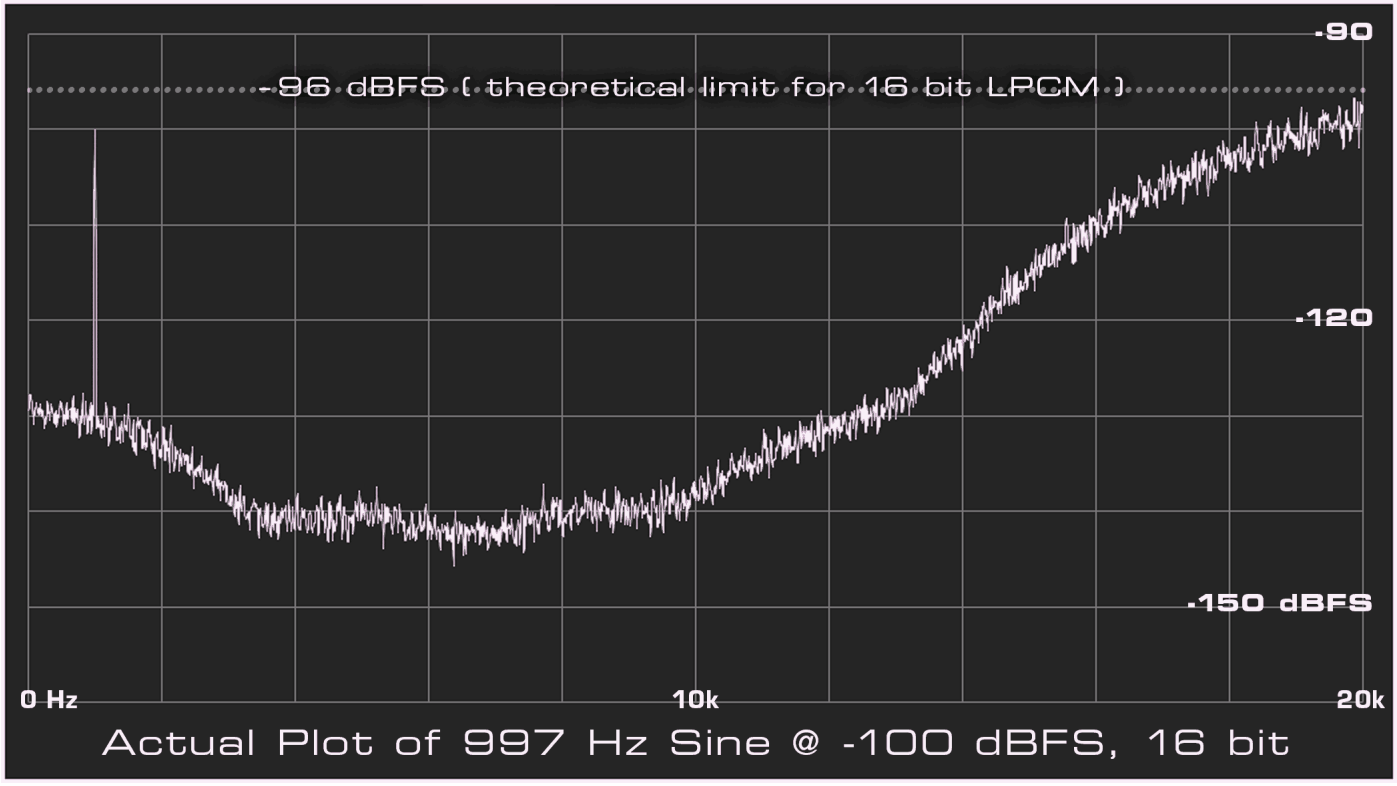

Figure A12 — Triangular dither, PS3

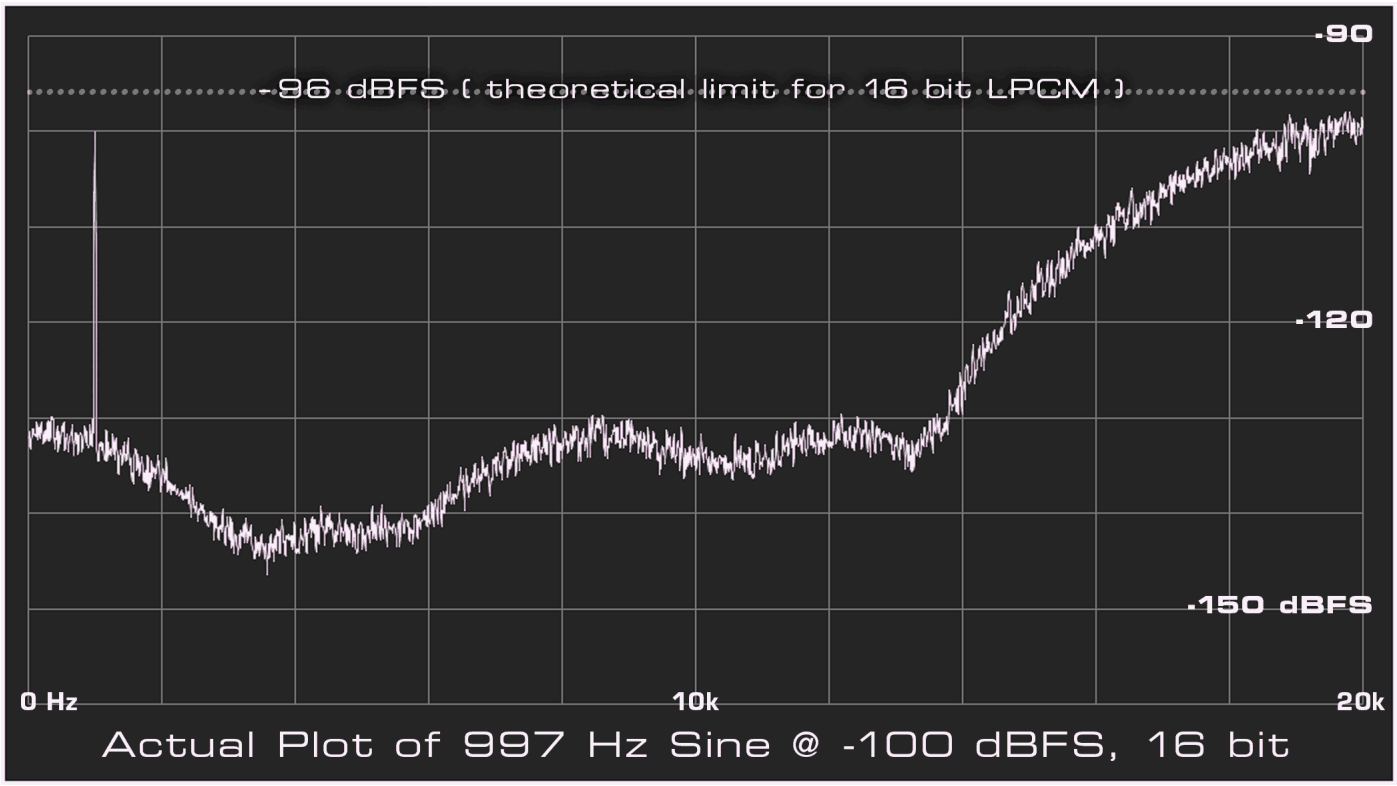

Figure A13 — Triangular dither, PS4

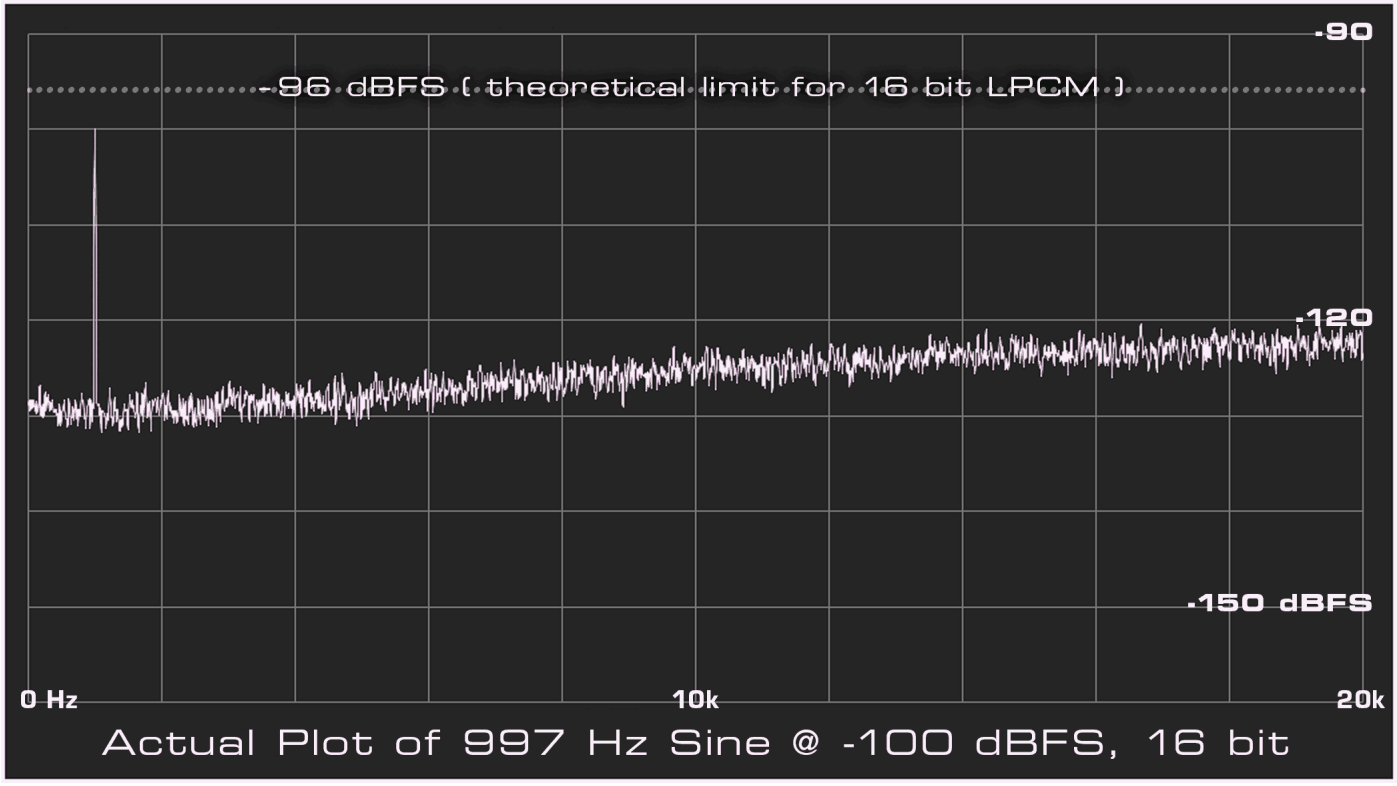

Figure A14 — High passed triangular dither, no noise shaping

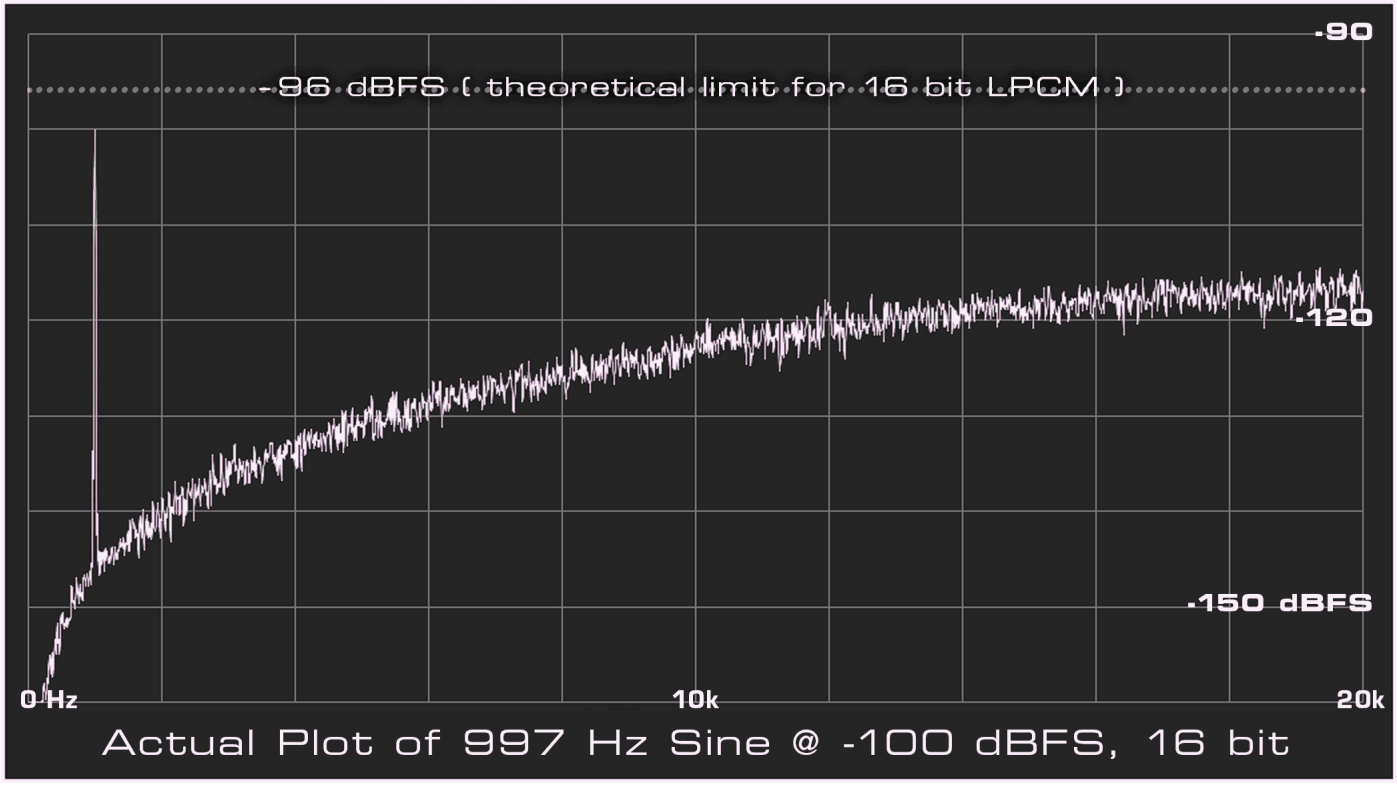

Figure A15 — High passed triangular dither, NS1

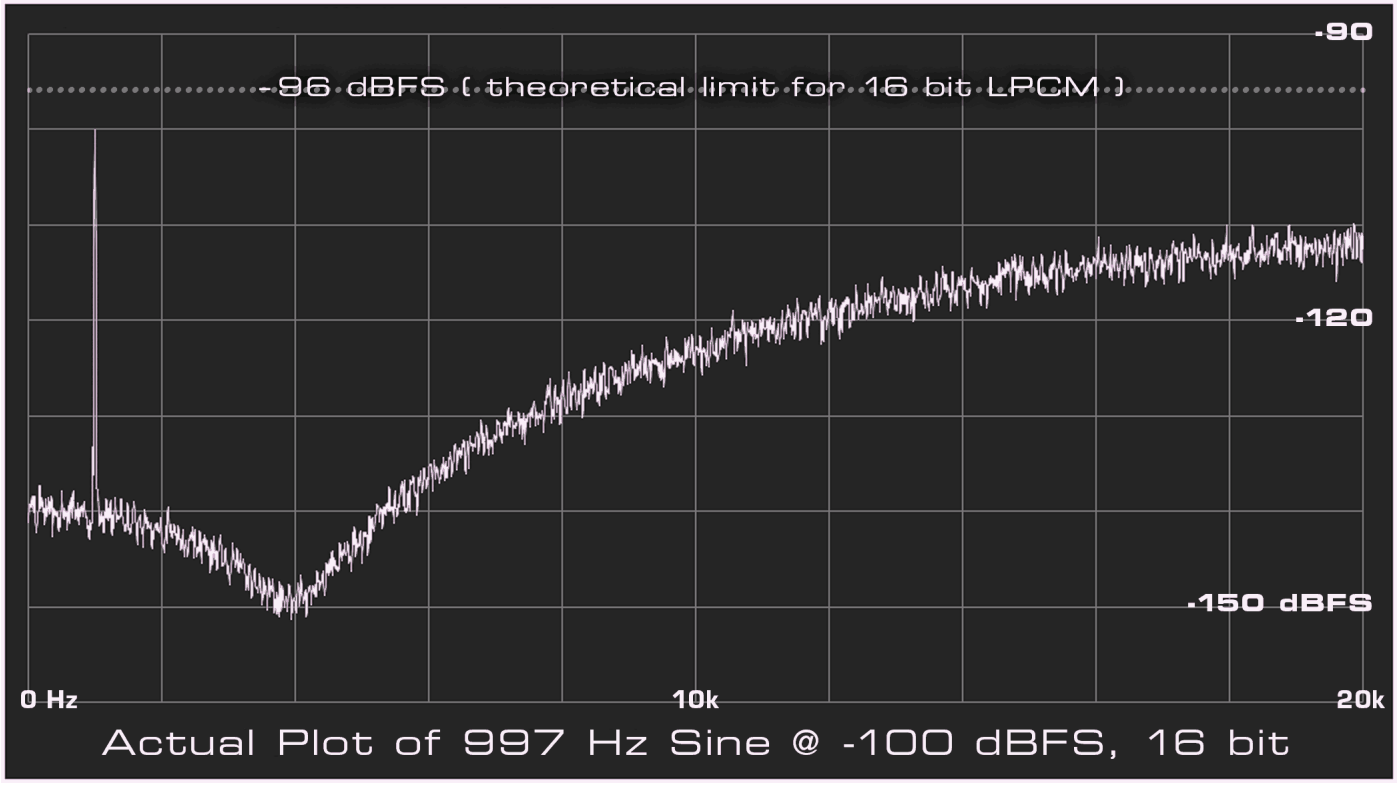

Figure A16 — High passed triangular dither, NS2

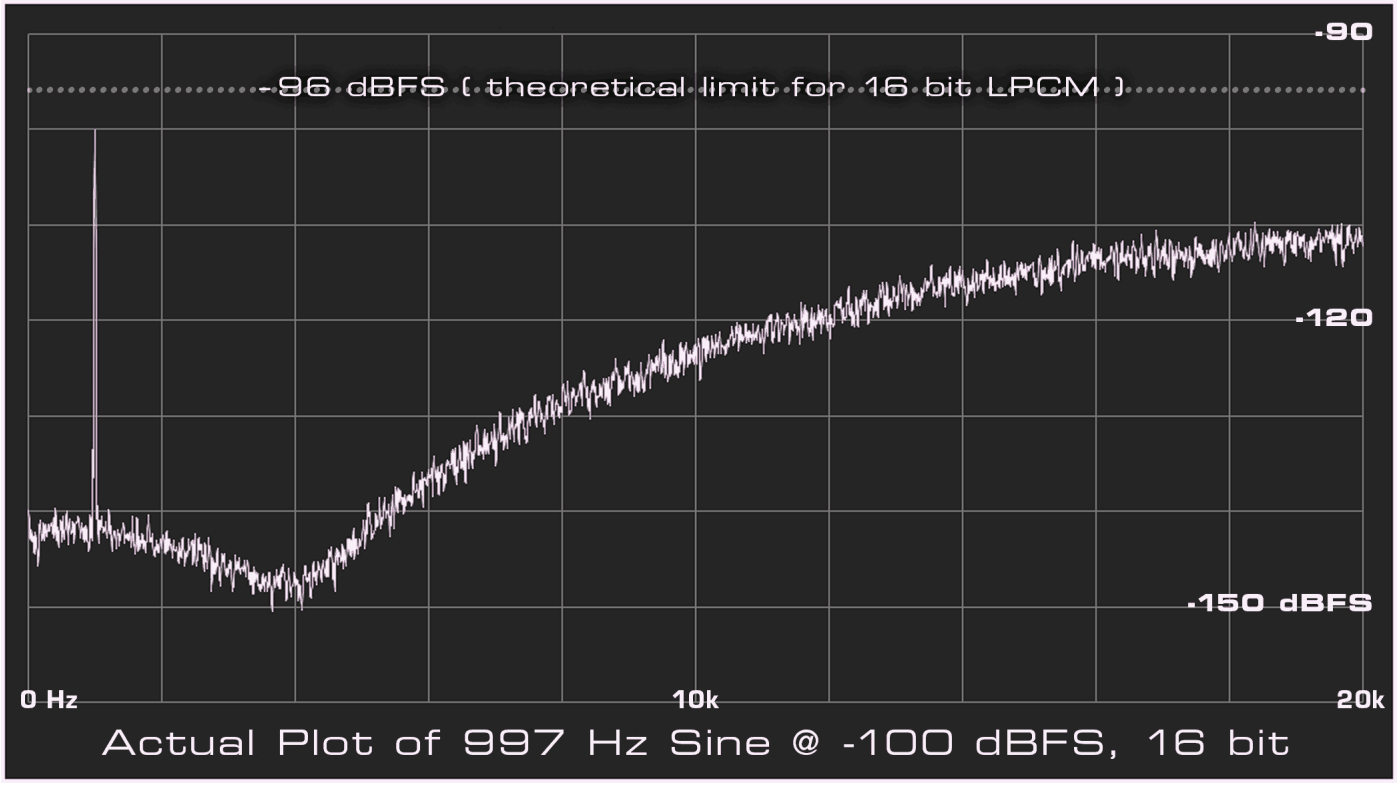

Figure A17 — High passed triangular dither, NS3

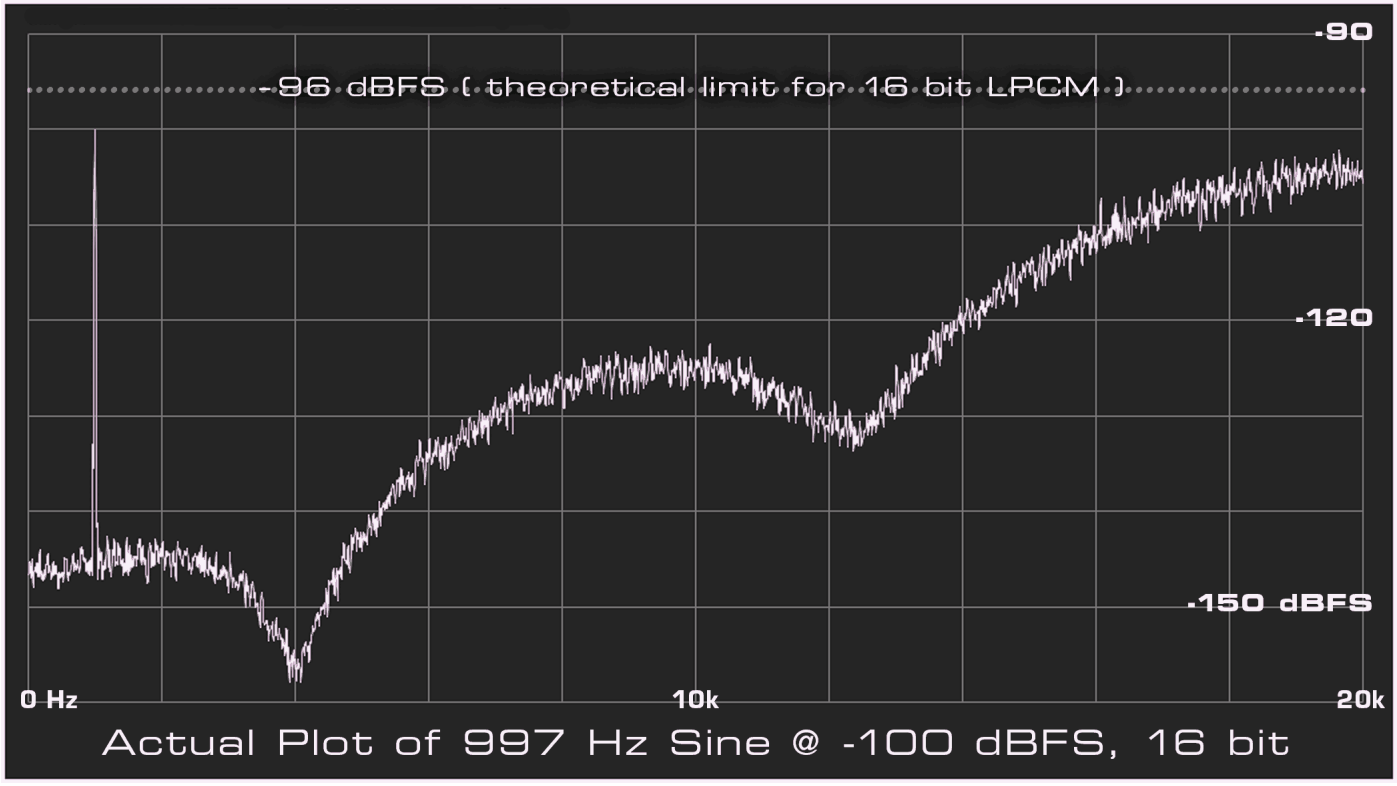

Figure A18 — High passed triangular dither, PS1

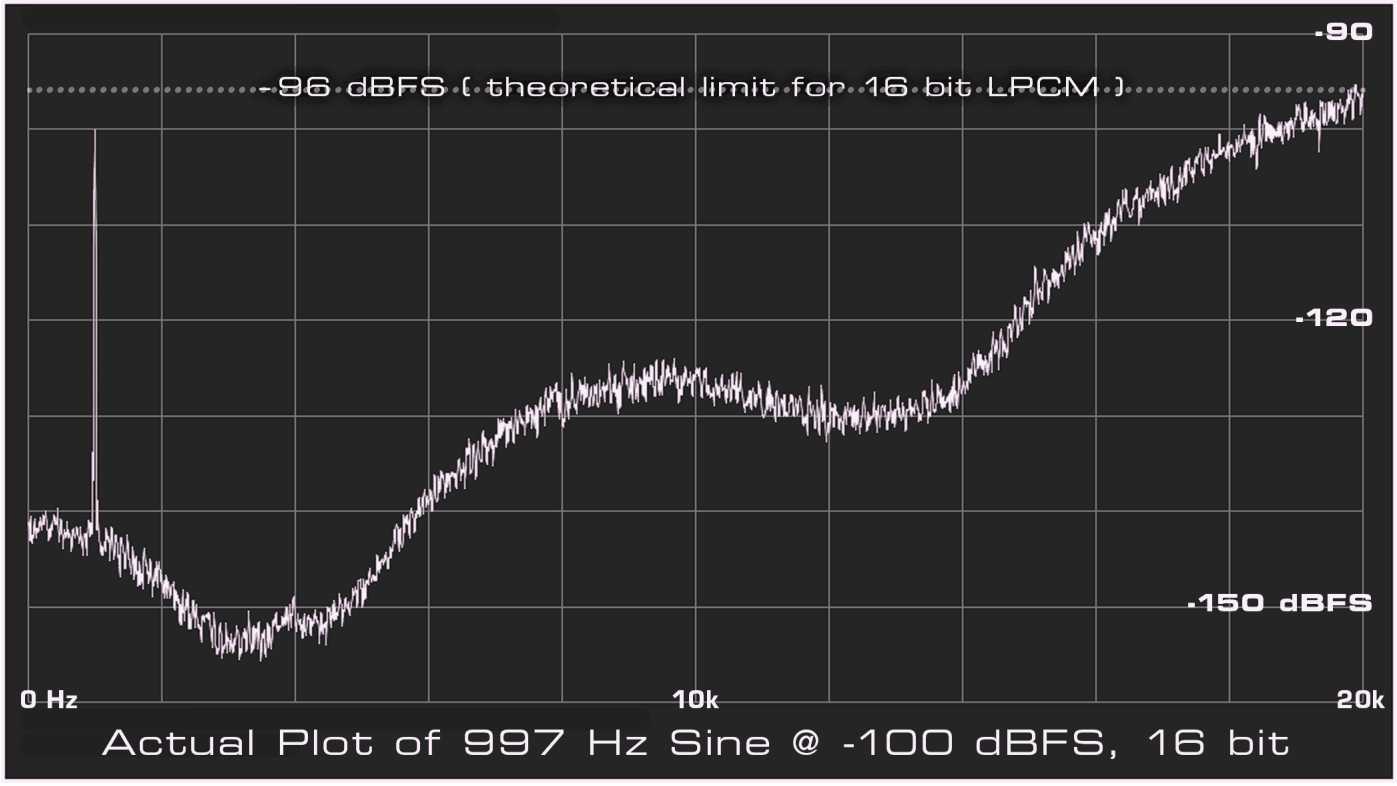

Figure A19 — High passed triangular dither, PS2

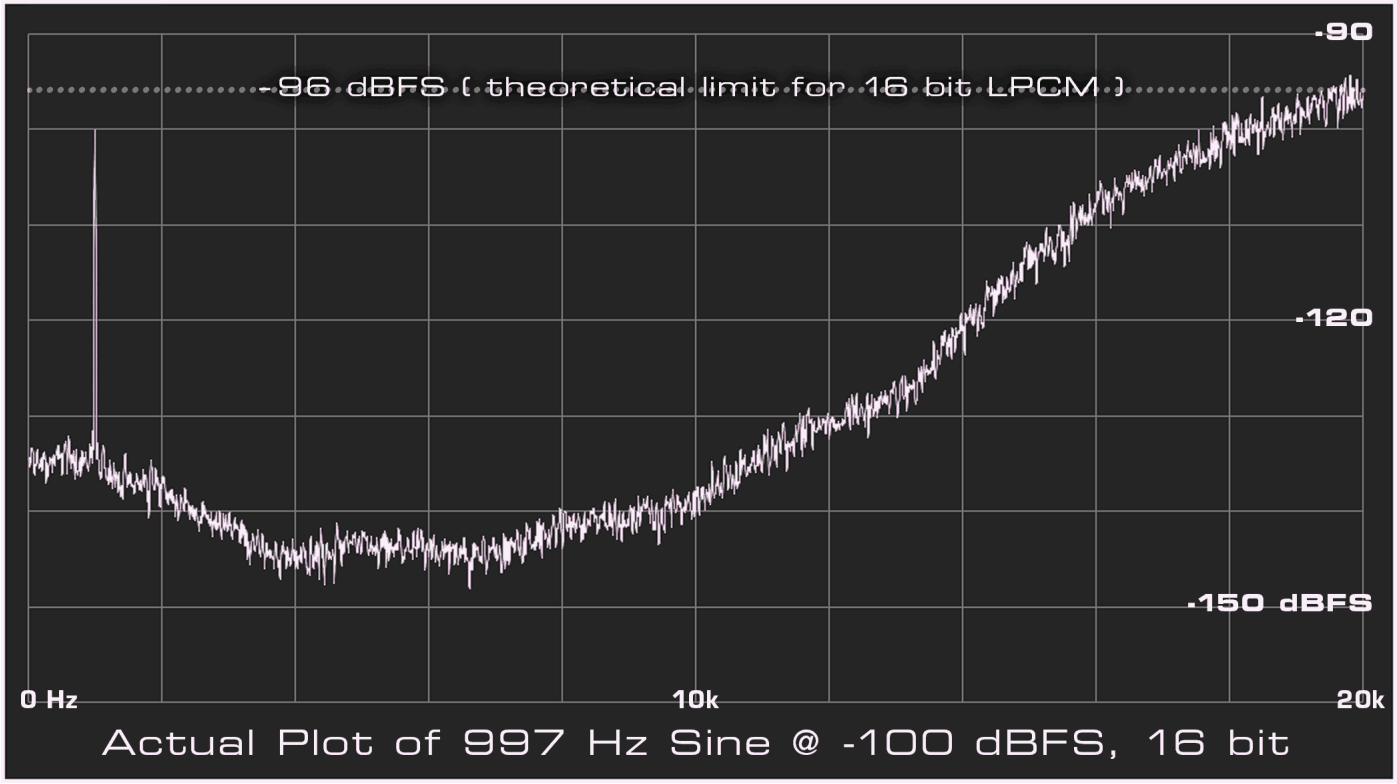

Figure A20 — High passed triangular dither, PS3

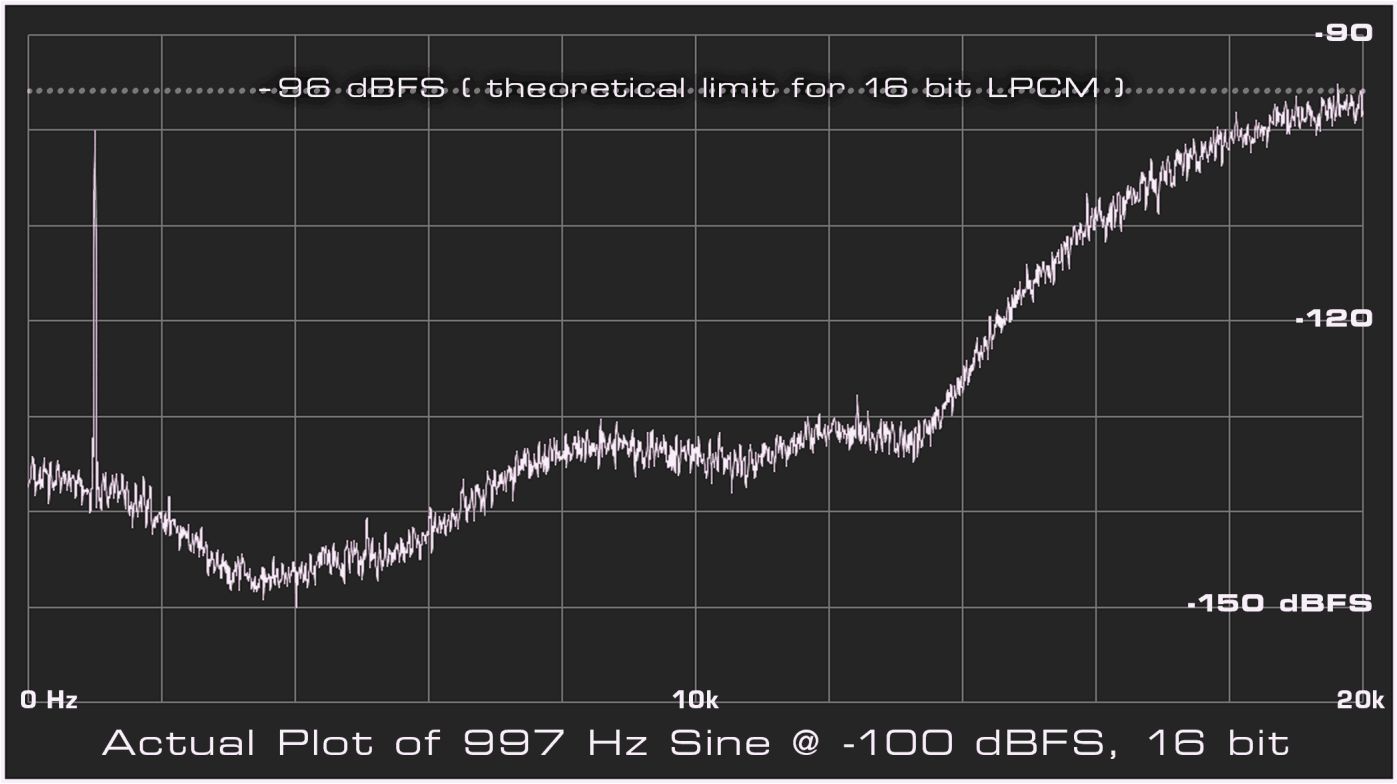

Figure A21 — High passed triangular dither, PS4

About This Manual

This manual was written in Adobe InDesign 17.3, and is set in Robert Slimbach's Minion Pro and Myriad Pro. The cover page is set in Aldo Novarese's modernist geometric Eurostyle. O.A. Masciarotte supplied the book *To Serve & Groove*, from which some material was adapted for this LINpro user manual.

This documentation, as well as the software described in it, is furnished under license and may only be used or copied in accordance with the terms of such license. The information in this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by MAAT Incorporated. MAAT Incorporated assumes no responsibility or liability for any errors or inaccuracies that may appear in this document.

Except as permitted by such license, no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, recording, or otherwise, without the prior written permission of MAAT Incorporated.

Notice of Liability

The author and publisher have made every effort to ensure the accuracy of the information herein. However, the information contained in this document is provided without warranty, either express or implied. Neither the authors, MAAT Incorporated, nor its dealers or distributors will be held liable for any damages to be caused either directly or indirectly by the instructions contained in this book, or by the software or hardware products described herein.

MAAT INCORPORATED MAKES NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, REGARDING THE APPLE SOFTWARE. MAAT INCORPORATED DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE OF THE MAAT INCORPORATED SOFTWARE IN TERMS OF ITS CORRECTNESS, ACCURACY, RELIABILITY, CURRENTNESS, OR OTHERWISE. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE MAAT INCORPORATED SOFTWARE IS ASSUMED BY YOU. THE EXCLUSION OF IMPLIED WARRANTIES IS NOT PERMITTED BY SOME STATES. THE ABOVE EXCLUSION MAY NOT APPLY TO YOU.

IN NO EVENT WILL MAAT INCORPORATED, ITS DIRECTORS, OFFICERS, EMPLOYEES, OR AGENTS BE LIABLE TO YOU FOR ANY CONSEQUENTIAL, INCIDENTAL, OR INDIRECT DAMAGES (INCLUDING DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, AND THE LIKE) ARISING OUT OF THE USE OR INABILITY TO USE THE MAAT INCORPORATED SOFTWARE EVEN IF MAAT HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES DO NOT ALLOW THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATIONS MAY NOT APPLY TO YOU.

220622-2.11