Tech Info

Dither in the ADI-1 / ADI-8 PRO

Remarks about the Need for Dither

»Technical Information Index

Preamble

Whenever 20-bit or 24-bit digital audio is copied to a 16-bit medium (for instance), the word length is reduced by discarding the lower bits. Truncation causes low-level components of the signal to distort. To combat this 'quantization distortion' noise at a level corresponding to the least-significant bit - or below - is added to the signal before truncation, randomly modulating the signal. This process is called 'dithering'.

One common misconception is to think that the added noise only modulates the LSB (least significant bit) of a signal, equating the amplitude of the noise itself with the amplitude of the modulation it causes. The truth is that ALL bits are modulated.

All A/D converters have built-in dither. It is a property of the converter chip, manifesting itself as an unavoidable noise floor. A/D converter chip manufacturer's use this noise - on purpose - to linearise the chip's behavior at lower signal levels.

To illustrate how dithering works fig. 1 shows a 1 kHz sine-wave at -80 dBFS. The signal has been generated digitally at 20-bit resolution, and is also analyzed at 20-bit. The FFT shows zero distortion, i.e. there are no peaks above the noise floor.

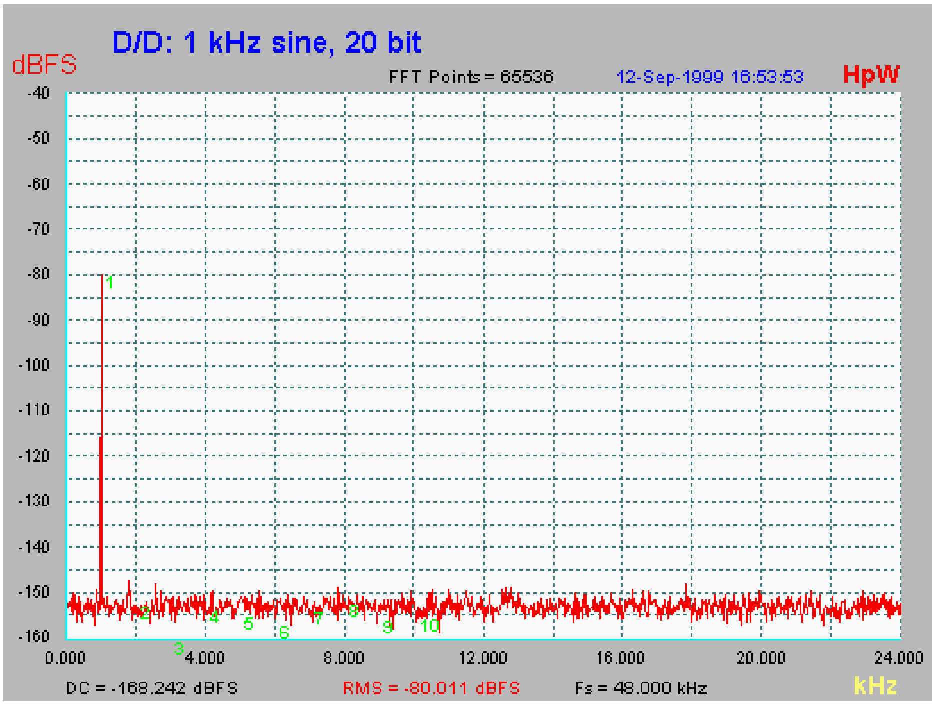

Fig. 1

The noise floor here is a result of quantization error in the 20-bit signal. Because the resolution of any digital signal (as opposed to analogue) is finite, even 24-bit digital signals show such linearity errors. Some dither was added to the test signal by the audio generator, to smooth the noise floor a little (for presentation purposes only.)

The analyzer was then set to 16-bit word length. The analyzer truncates the digital signal (by 4 bits), leading to the result in fig. 2.

Many harmonics have appeared as peaks across the entire frequency spectrum. These peaks are high enough above the noise floor to be heard as distortion.

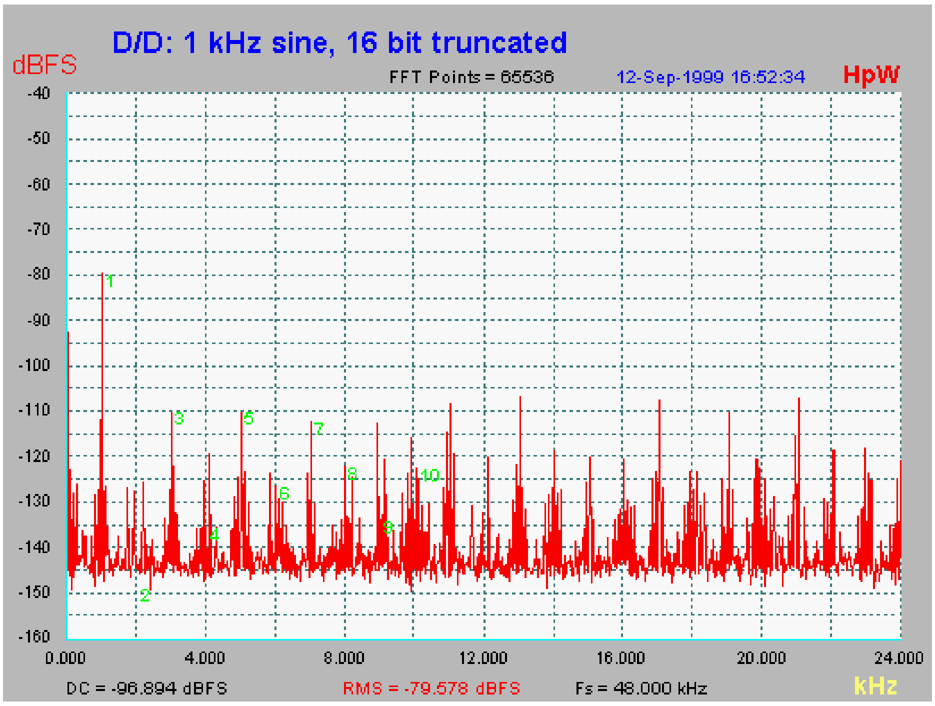

Fig. 2

The noise floor (which used to be -152 dBFS) is now about -143 dBFS, because 16-bit quantization error is higher than 20-bit quantization error. Also, there is a slight DC offset here. The level of the 1 kHz signal itself remains unchanged at -80 dBFS.

The level of the test signal (-80 dBFS) lies well within 16-bit range. We thought it might be interesting to find out what would happen if the test signal were below the 16-bit range.

We therefore repeated the test with a 1 kHz sine-wave at -110 dBFS. Fig. 3 shows the original signal...

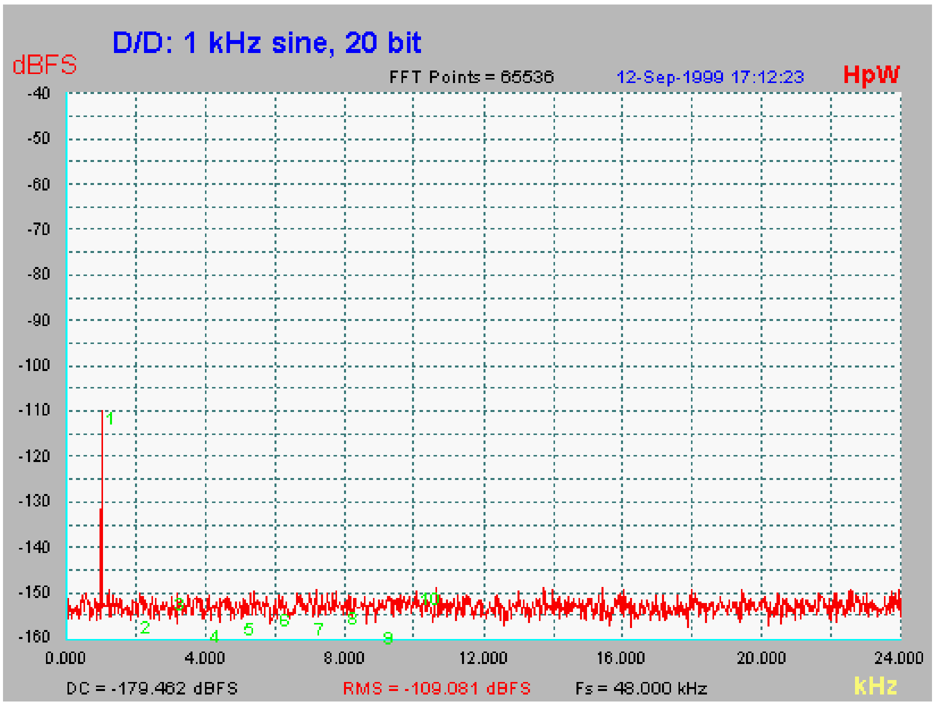

Fig. 3

...and fig. 4 shows the results of reducing word length to 16-bit. The sine is still present (!), but its level has been changed to about -95 dBFS.

There is a relatively smooth noise floor at around -135 dBFS (where it belongs - considering the current analyzer settings.) Again, there is a slight DC offset here, as well as several audible harmonics, though they are much less prominent than in the previous test.

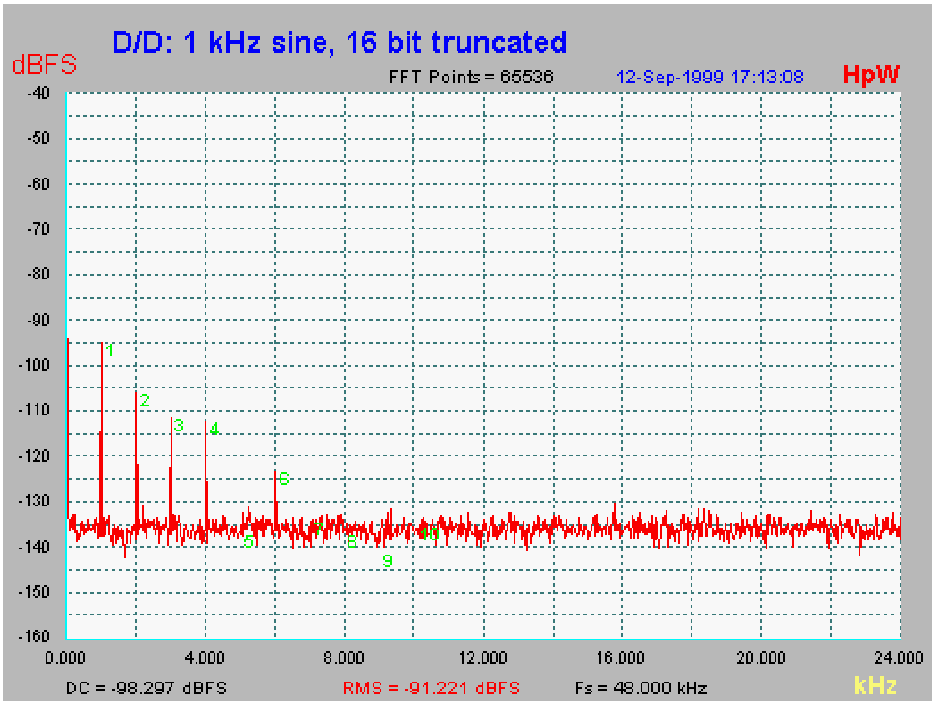

Fig. 4

You might have come across similar information about digital theory before - there are several sites scattered around the Web, many of which also include practical test results. However, their purpose is often to prove that A/D converters without dither would produce horrendous distortion, and therefore sound terrible. This is actually (willful?) deceit, as truncation of a digital test signal has little or nothing in common with truncation of an A/D converter's signal, and certainly doesn't lead to such 'catastrophic' results as seen in fig. 2. This is mainly thanks to inherent noise in the converter, as can be seen in the following measurements using an ADI-1.

ADI-1: Dither on chip

We often get inquiries from anxious customers, asking how using a 20-bit ADI-1 as a high-quality A/D converter for a 16-bit DAT recorder could possibly be an improvement over a 16-bit one. Or - because the ADI-1 lacks dither - whether it could actually degrade the result. In a test about the ADI-1 by Sound On Sound, the author Hugh Robjohns even claimed that the ADI-1 'cannot be used as an input device for CD-R, Minidisk or DAT machines because there is no option for correct dithering down to 16 bit resolution'. Maybe he should have tried it first...

All A/D converters generate noise to some degree or other. The level of inherent noise in our ADI-1's 20-bit ADC is about -100 dBFS (RMS unweighted), and is therefore slightly below the theoretical resolution of a 16-bit system (rule of thumb: 16 x 6 dB = 96 dB.) On the other hand, spikes don't count much in RMS evaluated results. Peak levels in the noise signal are about -90 dBFS (DIGI96 owners can check both these figures using the DIGICheck test program - see the screenshot to the right.)

You may already understand what is happening here: On the one hand, inherent noise generated by the ADI-1 is considerably lower than in similar 16-bit converters. On the other hand, it is still enough to cause a dither effect when the signal is truncated to 16-bit (just like any 'intentional' dither would.)

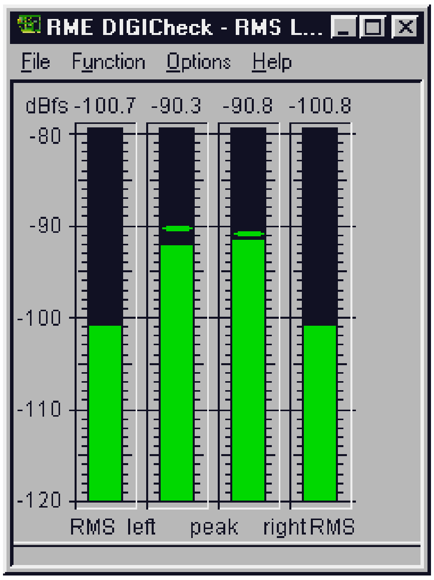

Fig. 5

The big question now: Is this 'natural dither' effect (caused by inherent noise in an ADC) anywhere near as good as an intentionally applied, well-defined amount of dither from an external source? The answer: It is!

After listening sessions suggested lower noise (and no distortion) from the ADI-1 system than from systems using conventional dither, we wanted to test this phenomenon objectively. We found the reason immediately: Inherent noise in the ADI-1 is equivalent to a dither level of 0.7 LSB. Standard levels here are normally 1 LSB (which guarantees zero distortion but has a signal-to-noise ratio of about 92.5 dB RMS unweighted due to the added noise), or 0.5 LSB (where the S/N is over 95 dB RMS unweighted using a high-quality DAC, but distortion is not totally wiped out in low signal levels.)

With inherent noise of about 0.7 LSB, the ADI-1 has found the golden mean quite by chance: maximal Signal to Noise ratio at lowest distortion, as a by-product and 'on chip'. To complete the story, here are the results of A/D conversion in the ADI-1 with and without truncation.

Fig. 6 shows a 1 kHz analogue sine converted to 20-bit digital at -80 dBFS. The noise floor level is about -134 dBFS.

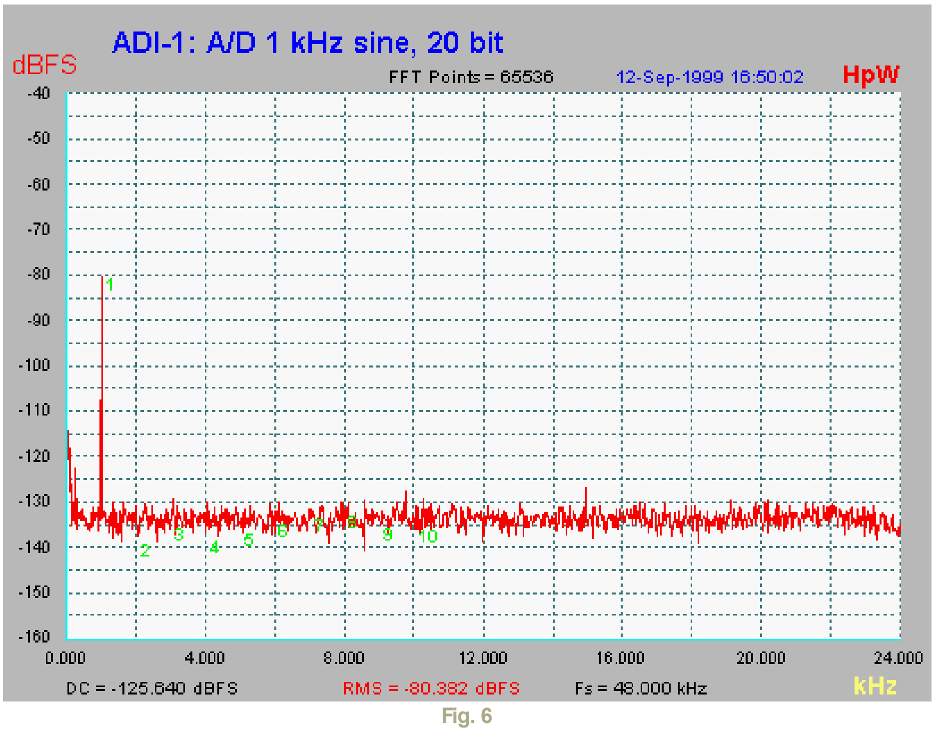

Fig. 6

Fig. 7 shows the same signal converted to 16-bit. Discarding the bottom 4 bits results in a (4 dB) higher noise floor and the same slight DC offset we saw before. The distortion level here is very low (-120 dBFS) and it only rises above the noise floor at frequencies over 12 kHz, so it is practically inaudible.

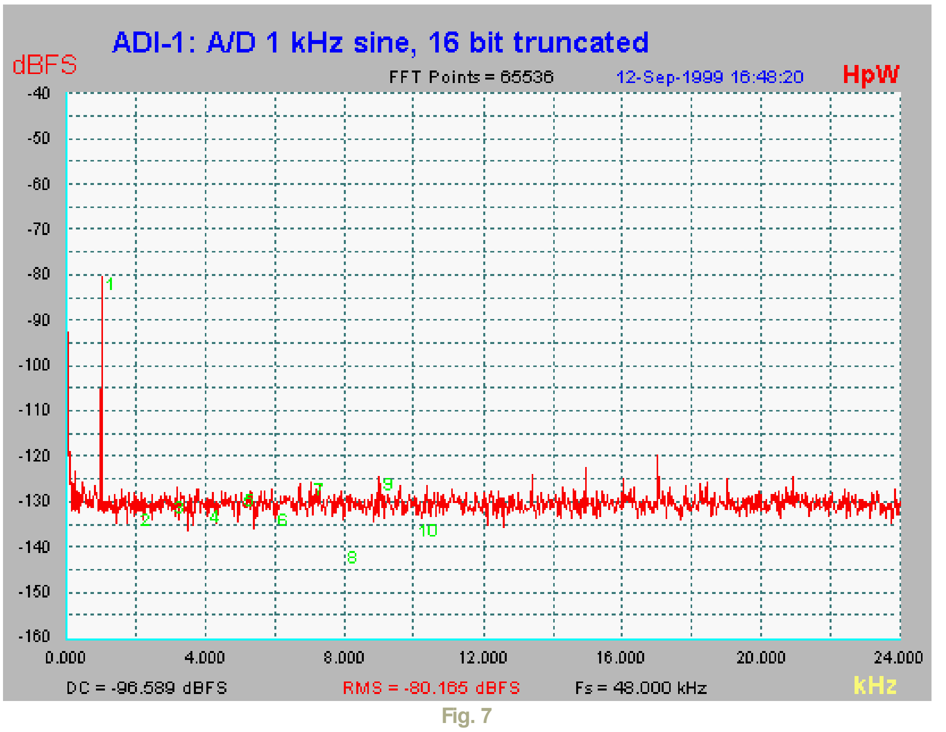

Fig. 7

Fig. 8 shows a 1 kHz analog sine converted to 20-bit digital at -110 dBFS. Again, the noise floor level is about -134 dBFS.

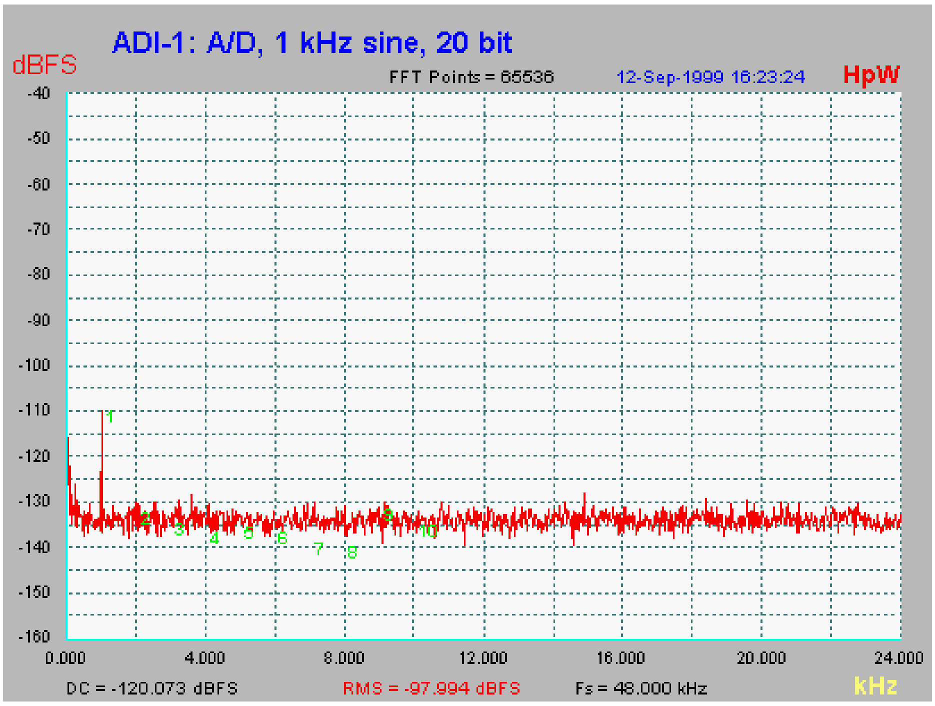

Fig. 8

Fig. 9 shows the same signal converted to 16-bit. Again, truncating the 4 least significant bits results in a (4 dB) higher noise floor and a small amount of DC offset. The signal level has shifted slightly due to the inevitable system linearity errors, but its level of -107 dB is very close to the original -110 dBFS. No distortion could be detected at all.

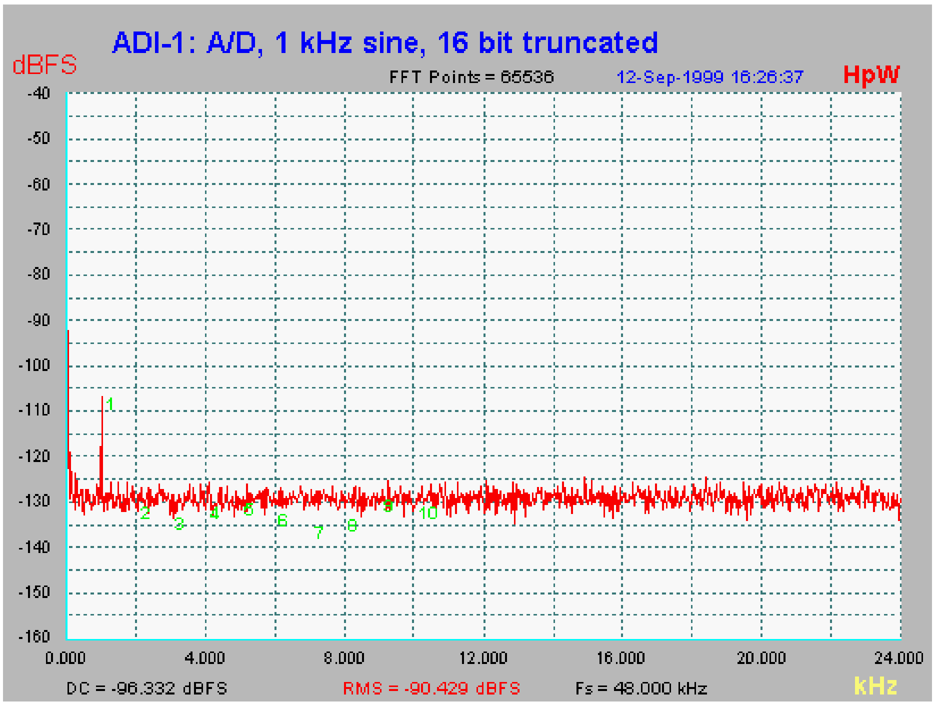

Fig. 9

Conclusion

Any additional dither in the ADI-1 would only push up the noise floor, and less distortion than zero is impossible. The ADI-1 is therefore an excellent substitute or companion to any devices with built-in A/D converters, such as samplers, soundcards, CD-R drives, Minidisk or DAT machines. In direct comparison to conventional 16-bit converters, the 20-bit ADI-1 wins because it can produce higher quality results while remaining completely compatible with the 16-bit format. Dither? Don't worry!

ADI-8 PRO

Maybe you are an ADI-8 PRO owner, and you have skipped the ADI-1 section above? Stop!! Please go back and read the entire Tech Info, because a lot of the above information also applies to the ADI-8.

First of all, a short summary:
Whenever 20-bit or 24-bit digital audio is copied to a 16-bit medium, the word length is reduced by discarding some of the least significant bits. This truncation causes quantization distortion in the lower-level components of the signal. To combat this distortion, noise at a level of the least significant bit (or below) is added to the signal - this is called 'dither'. And especially: External dithering is unnecessary if the sum of noise from the source as well as from the A/D converter is above a certain threshold.

The A/D converters in the ADI-8 PRO have a SNR of 112 dB (RMS unweighted.) The inherent noise is so low that it cannot dither the 16-bit truncated audio reliably. In other words: This noise is equivalent to dither below 0.5 LSB, basically not enough for dither.

Accomplishing dither in a device such as the ADI-8 PRO, which is completely FPGA-based (field programmable gate array), requires a great deal of time and effort. This is why we started out by checking just how 'bad' the effects of simple truncation in the ADI-8 really are, and discovered several interesting facts, which seem to have been kept secret, presumably to keep up appearances (and prices) for those devices which boast sophisticated dithering methods.

Remember (see the ADI-1 tests) that truncating signals outside the 16-bit range does not cause them to disappear altogether, but only changes their levels slightly. This effect - retention of information beyond what you would expect, looking at the system resolution (16-bit / 96 dB here) - is one of advertising's favorite arguments for dither, and is easy to prove. In our case however, the following applies: External dithering is unnecessary, the effect is an inherent property of all (!) modern A/D converters.

A common misconception is an 'analog' way of thinking projected into the digital domain: Discarding the least significant bits, any low-level components of a signal would be lost. The noise floor of a 24-bit signal at -112 dBFS would disappear completely when converted to 16-bit, digital zero would be the result. This is absolutely wrong. All parts of the signal which were originally below -96 dB will still be present in the 16-bit version (FFT proves this), but not at the original levels. In this context you could say it is a compression-type effect - levels which were further apart in the original have moved closer together.

In the early days of the digital era, A/D converters were DC-ridden, and parts of the signal really could be lost. However, the converter chips used in the ADI-1 and ADI-8 PRO include DC filters and automatic calibration, eliminating any DC offset.

This state of affairs is best understood by looking at DIGICheck's Bit Statistics page. It displays digital data in a binary form (as 'twos complement'.)

Fig. 10 shows that all 24 bits in the ADI-8 PRO 'wobble' i.e. they are active. After truncation to 16 bit, the lower 8 bits are dead of course, but the remaining 16 bits are still active. This means that there must be a signal present. Lo and behold: The base noise, which used to be at -112 dB, is now at -93 dB.

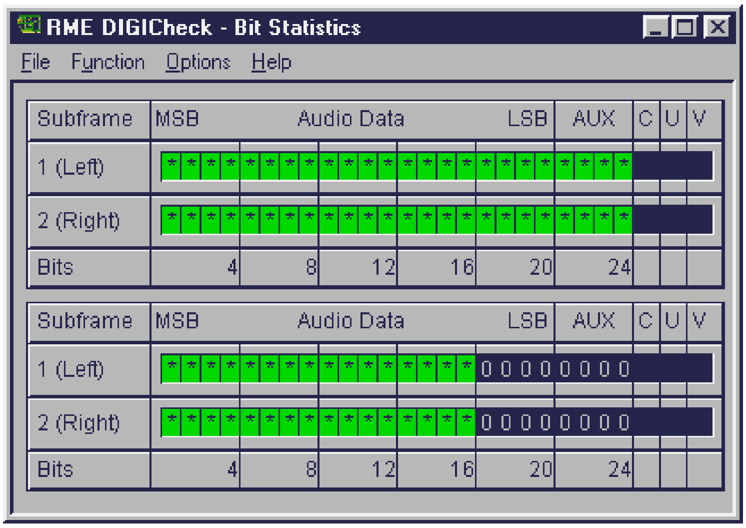

Fig. 10

In other words: Even with a minimum of dither (caused by 'basic' noise), the 'increased resolution' effect is retained. Of course there are some differences: The lower the dither level, the less accurate the transmission of the original level into the 16 bit signal will be. Linearity decreases, and formerly significant level changes in the region below -96 dBFS are compressed into only a few dB.

Fig. 11 shows a 1 kHz analogue sine digitally recorded as 24-bit, at a level of -80 dBFS. The noise floor is about -148 dBFS.

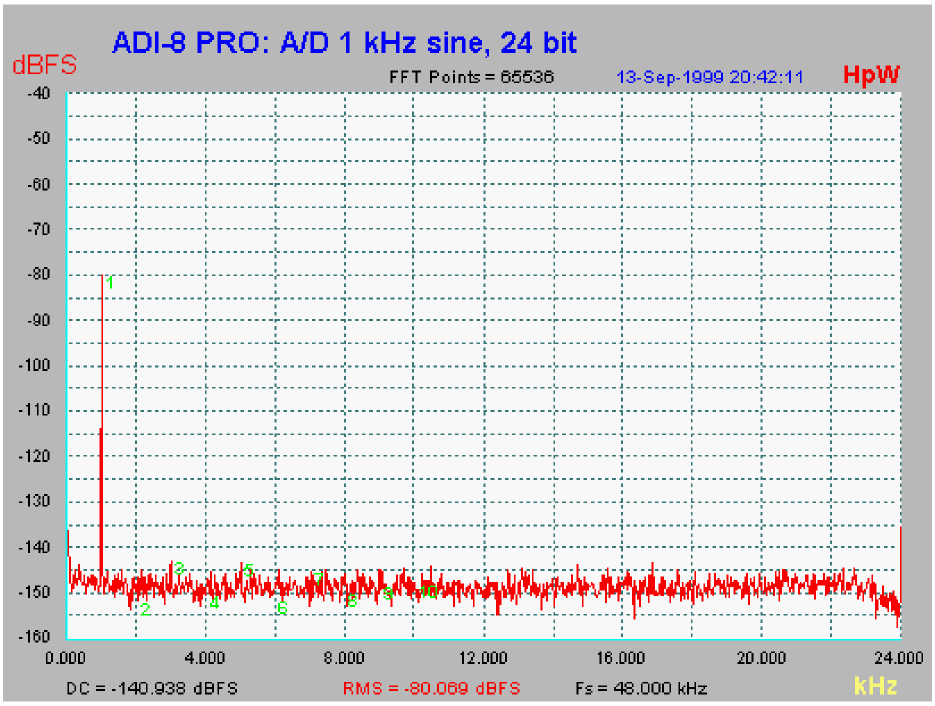

Fig. 11

Fig. 12 shows the result of truncating to 16-bit. As expected, the noise floor is pushed up to about -138 dBFS, and significant distortion appears across the whole frequency range. Note the distinct differences between this result and fig. 2.

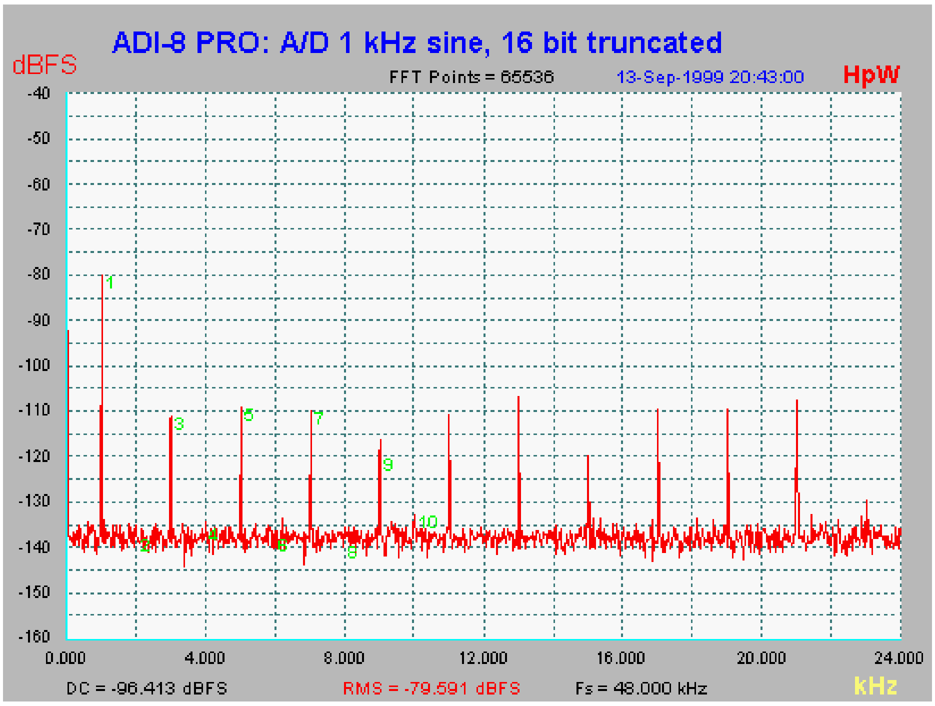

Fig. 12

The distortion you see here looks worse than it sounds. After all, the listener's ear has to battle with a constant and relatively high noise floor, out of which it's not only supposed to recognize the original signal, but also the newly generated harmonics. Although the distortion in this example is clearly audible (in case volume is pushed up enough to also hear the noise...), the whole test as such should be put into perspective: The A/D converter has been fed a (nearly) noise-free signal consisting of a pure sine-wave at 1 kHz. Such signals are normally only used to test equipment - music, speech etc. are much more complex signals which tend to mask out most low-level distortion.

However, it seems to us that the dynamic range (i.e. the SNR at the ADI-8 PRO's input) attainable in real conditions would be much more relevant:

- Example 1:
  The ADI-8 is being used as a multichannel converter, the source signals are coming from either the subgroups of a mixing desk, or directly from a high-quality microphone preamp. In reality, neither of these sources can deliver a signal as clean (of distortion and noise) as the one we used in the tests. Some figures: Typical noise level in a high-quality preamp is -127 dB, required amplification for the microphone is 30 dB, so the effective SNR is 97 dB. No distortion can be detected because the input noise already corresponds to an additional dither of 1 LSB. (DIGI96 owners can use DIGICheck to measure the true dynamic range at any time - many were in for a big surprise....)

- Example 2:
  The ADI-8 input is set to -10 dBV. As SNR drops (to 107 dB RMS unweighted), the 'on chip dither' effect cuts in, causing a marked reduction in distortion in the ADI-8. If you know just how low the SNR is in standard -10 dBV based equipment, you will not be worrying about the missing dither.
- Example 3:
  The ADI-8 is basically an exceptionally good A/D converter, and would be a worthy candidate to be used in dual channel mode as a line-level ADC. In this case however, the ADAT optical output would not be used recording to DAT but to a PC (and then to DAT.) The signal can be dithered and converted to 16-bit within the PC.

The following tests show that distortion is dynamic (and not constant as one might think) - distortion changes in level and distribution, depending on the strength of the signal.

Fig. 13 shows a 1 kHz analogue sine converted to 20-bit digital at -110 dBFS. The noise floor is around -148 dBFS.

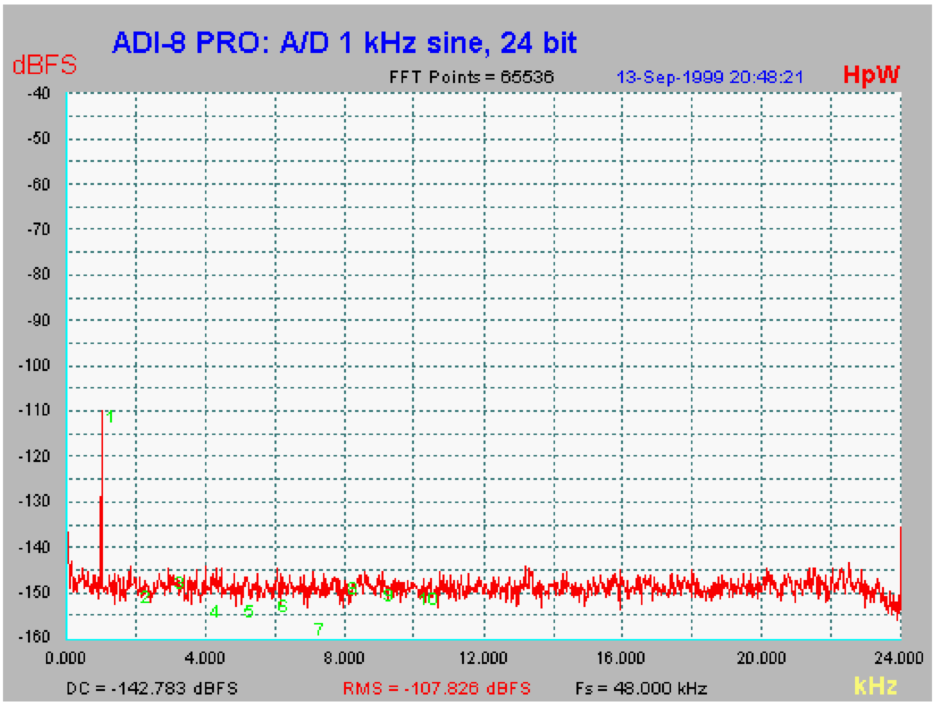

Fig. 13

Fig. 14 shows the results of truncation to 16-bit. The noise floor is pushed up to about -138 dBFS. The distortion consists of only one inaudible harmonic. It is masked by the noise, because of its low level and proximity to the fundamental. Compare this result to the one in fig. 4.

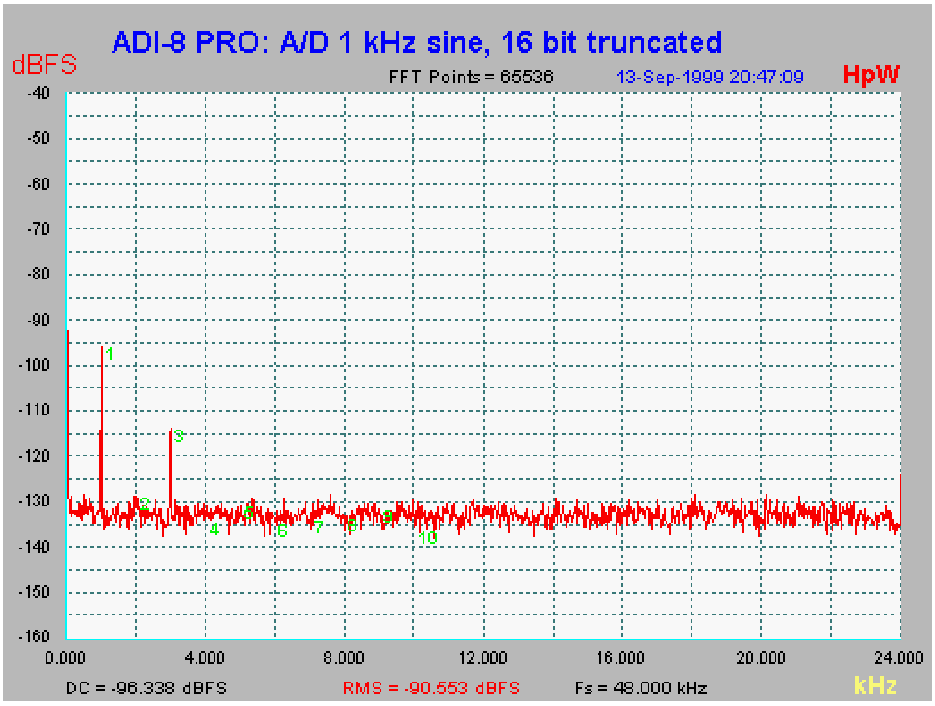

Fig. 14

To summarize: Audible distortion does not occur in the region below 96 dB, but only when the original signal is within the 16-bit range. The distortion depends on signal level, and is generally eliminated by noise from connected equipment and/or from the recorded signal.

Apart from results from the above tests, there are other good reasons why we (and you) can safely do without dither in the ADI-8 PRO:

- To retain the full dynamic range, ADI-8 PRO's Bit Split method is normally used when copying data to 16-bit media.
- Transferring to 20-bit (such as ADAT XT or Yamaha O2R) does not require dither, as the maximum dynamic range of the ADI-8 PRO is 'only' 18.6 bit (or 112 dB), so there is no loss in a 20-bit (120 dB) system. (In accordance with the ADI-1 findings above, 20 bits would still be enough for a dynamic range of up to 124 dB.)
- Tascam DA-38 or 98 owners can use the (often overlooked) built-in dither functions (please refer to the respective manuals.)
- To transfer data to a computer, 20 or 24-bit resolution can be used. Dither can be added at the very end of the chain here, i.e. after all editing and mixing has been done.
- The dynamic range of a signal source should actually be far above 100 dB. However, this is seldom the case in real life situations - most audio sources produce a considerable amount of inherent noise and/or hum.

Summary

To avoid being misunderstood here: We are not trying to persuade you that external dither is altogether pointless. Theoretically, the ADI-8 PRO could benefit from sophisticated dither or noise-shaping when transferring data to 16-bit media. In reality however, DC-free converters and limitations posed by real recording environments negate any advantages that dither might bring.

The purpose of this Tech Info is to dispel some common myths about dither. It is meant to supplement books, magazines and Internet documents. It is also an attempt to put some manufacturer's' statements into perspective - those whose main aim in life seems to be preaching the indispensability of dither, although this is not always true. We've already seen an example of such 'authoritative' statements in the section about the ADI-1 at the beginning of this text, and similar unscholarly opinions are not hard to find. When even reputable journals start publishing unfounded generalizations, then something obviously has to be done about it.
