TECHNICAL WHITE PAPER
Prepared by Gannon Kashiwa

The Pro Tools **48-bit Mixer**

This technical article provides detailed information about how the Pro Tools mixer operates. In so doing, we will demonstrate its summing characteristics and explain how a 48-bit "clean" mixer functions within the 24-bit TDM (Time Division Multiplexing) environment. By providing some 'behind-the-scenes' information about mixing and summing in Pro Tools, we hope to shed light on a few myths about mixing 'in the box' with Pro Tools, as well as provide you with a better understanding of the mechanics of summing signals.

Though the tasks of any digital mixing system are the same—combining audio streams without clipping while also retaining their low level detail—the approaches used can be very different. This article provides a description of the approach taken in Pro Tools and the benefits of using this system.

NUMBER OF BITS	RESOLUTION	DYNAMIC RANGE
2	4	12 dB
3	8	18 dB
4	16	24 dB
8	256	48 dB
12	4,096	72 dB
16	65,536	96 dB
24	16,777,216	144 dB
32	4,294,967,296	192 dB
48	281,474,976,711,000	288 dB
56	7.20575940379 E16 (add 16 0s)	336 dB

Figure 1: Bits, resolution, and amount of dynamic range

WHERE ARE THE BITS AND HOW ARE THEY USED?

Let's first define some terms that we'll be using in this discussion.

- **Bit:** A binary digit. In digital audio, the word length of the system indicates how many bits are applied to recording the sound and used for calculating changes such as level, EQ, dynamics, etc. A single bit represents the smallest change in the signal. Large digital words provide more discrete values so the changes represented in the smallest or least significant bit (LSB) can be very slight. Bits are also used to describe the dynamic range of the system. A single bit represents about 6.02 dB of dynamic range.
- **Resolution:** The number of discrete values available in a digital system. In figure 1, you can see the number of discrete steps each number of bits yields. More discrete steps allow for very fine changes in the signal to be faithfully reproduced.
- **Decibel:** A logarithmic scale for representing the ratio of two amounts of power. Here, we're using the decibel as the unit of measurement for dynamic range—the difference between the highest power signal and the smallest power signal of a recording.

Now, let's take a look at how Pro Tools takes signals from the analog world, converts them to digital signals, and follow them through the system to where they emerge again as reconstructed analog audio.

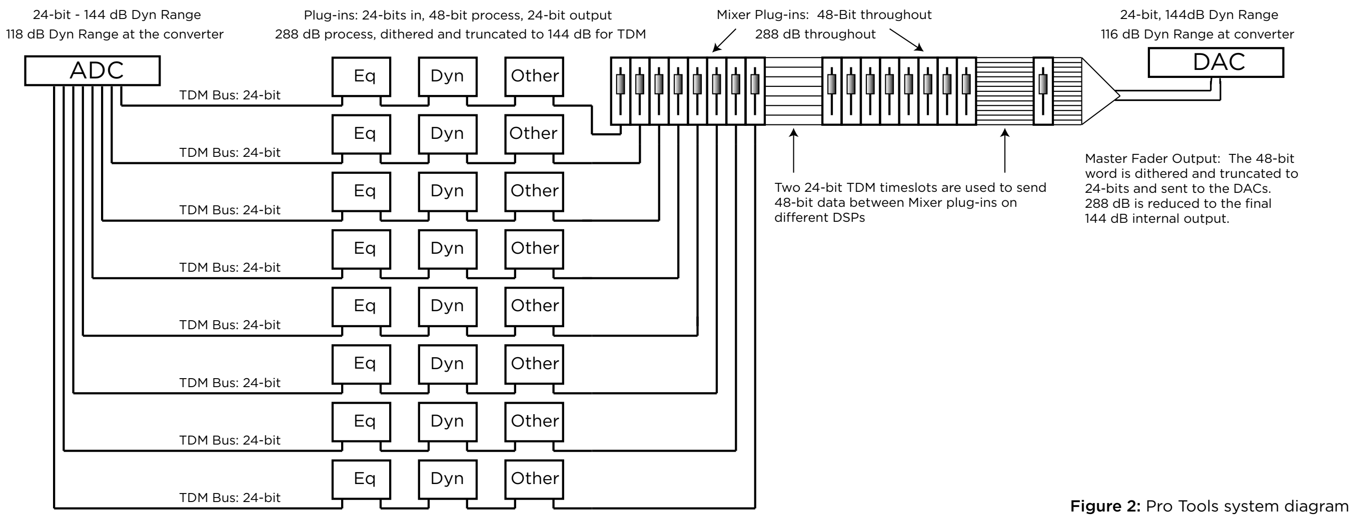

Figure 2: Pro Tools system diagram

We'll start with the analog signal arriving at the A/D converter; for the purpose of this discussion, we'll use the Digidesign 192 I/O as the converter and assume we're working with 24-bit audio sampled at 48 kHz.

The analog audio enters the system as 24-bit audio samples which have a theoretical dynamic range of 144 dB (remember 1 bit = 6.02 dB dynamic range). This, of course, is limited to converter performance. The Digidesign 192 I/O's analog to digital converters deliver a respectable 118 dB of dynamic range to work with. It's worth noting that when comparing specifications of audio interfaces, the dynamic range is defined as the difference between the peak power levels of the loudest audio signal minus the power level of the noise floor or the equipments' self-noise. The very best performing converters available today deliver about 120 dB of dynamic range (and are considerably more expensive than a 192 I/O), so 118 dB is a very good starting point for audio entering the system. So, now we have a tidy stream of 24-bit audio samples to work with.

Next, the 24-bit signal is placed on the 24-bit TDM bus and is transported to DSPs for plug-in processing. TDM enables each sample period to be "sliced" into 512 "timeslots" so many signals can be moved at the same time—it's a very fast and wide bus and that's how many signals can be handled with extremely low latency in Pro Tools|HD. At the DSP level, many of today's plug-ins process audio using "double precision" math. Double precision means just that; instead of 24 bits, the results of calculations are carried out and stored as 48 bits. This means that more of the "remainder after the decimal point" that represents the sound is maintained during digital signal processing, which, over multiple processing steps, can translate into higher sound quality. A double precision plug-in works on the signal with 48-bit processing and then applies dither before reducing the result to 24 bits to be placed back on the TDM bus.

To illustrate why more precision is beneficial, here's a simple example in normal decimal math. Take two numbers representing the amplitude of a signal and some amount of gain change being applied to that signal:

0.96 (original signal amplitude) x 0.612 (negative gain) = 0.58752 (resulting signal amplitude after gain change)

You can see that multiplying the numbers produces a new number with more digits than the original number. If you were limited to our initial resolution, the result would be rounded to 0.59. Now, imagine how many similar processes occur in even the simplest mix. Gain changes, EQ, panning, and mixing tracks—all of these are equations that produce longer results. It's easy to see that the numbers (in terms of precision) grow very quickly. In double precision plug-ins, these numbers are calculated out to 48 binary places. When binary numbers are used, 48 bits provide for 281,474,976,711,000 discrete numbers. You might think about resolution like this: Say the range of a theoretical system is 0 to 10 and you have 8 bits to work with, you have 256 possible values between 0 and 10. Apply 48 bits to that same range and you have 281,474,976,711,000 possible values; much finer detail can be represented.

A BIT OF DITHER

There is a dithering stage in most double precision plug-ins and one final dithering stage at the post master output of the summing mixer. Dither is noise with very specific properties[i] added to the signal in order to de-correlate the noise floor from the original signal so that when length reduction occurs, the resulting waveform does not contain any harmonic distortion or noise floor modulation artifacts. This allows the plug-in and the mixer to perform very precise calculations while maintaining the low level detail when handing the signal off to the next process. Within the Pro Tools TDM mixing environment, when the dithered mixer is used, dither is added to the 48-bit signal before it is reduced to 24 bits to be placed back on the TDM bus and sent to the DACs. Dither is not added on a track by track basis, as this would produce unwanted accumulated noise. Dither is only added once at the Master Fader output of the mix bus so the total system dither noise is 3 dB at -144 dB.

"Not so fast!" you say. "You just added noise to my signal and truncated my word from 48 bits to 24 bits—I've been robbed!" Well, it's true that within this system topology, the signal must be 24-bit in order to pass from DSP to DSP via the TDM bus, but it's important to understand that the level of the dither added to the signal is around -144 dB, which is below the noise floor of the converters. Consider that our ears can deal with a dynamic range of around 120 dB (from the threshold of audibility to the threshold of pain). This gives you an idea of how low "-144 dB" is in terms of audibility—you'd have to have your monitor system cranked up to extremely high levels to even hear any sound that might exist in a noise floor this low. The same could be said about the loss of the lower 24 bits. The sample values represented by the lower 24 bits in a 48-bit word are between -144 dB and -288 dB, so the dither only affects the signal at -144 dB or lower.

THE PRO TOOLS 48-BIT MIXER PLUG-IN

Just as EQs, dynamics, and reverbs are plug-ins, the Pro Tools Dithered Mixer is a plug-in as well. It differs from most other plug-ins in that it can grow to span multiple DSPs and as it grows, it passes signal from DSP chip to DSP chip *at a full 48 bits* instead of dithering and truncating back to 24 bits. It does so by using *two 24-bit TDM timeslots per connection.* This enables the mixer to maintain an internal dynamic range of at least 288 dB from beginning to end.

The extra 24 bits in the system are used to provide channel faders with additional dynamic range *above and below* the original 24-bit word, and it guarantees that the same fidelity is maintained when adding more inputs to the mix bus. In the Pro Tools +12 mixer, 9 bits are reserved for levels above 0 dBFS, providing 54 dB of headroom. This is enough headroom to allow 128 tracks of full code, correlated audio (imagine sample-aligned sine wave source files) to be summed with all faders at +12 without clipping the "input side" of the bus. It also provides enough bits below the 24-bit word to allow channel faders to be placed at nearly -90 dB before they stop contributing a full 24 bits to the mix.

Since in the real world, audio signals are almost never exactly correlated in this way, you can mix a far larger number of inputs within the Pro Tools mixer without clipping the "input" side of the mix bus at all. You can clip the "output" side of the mixer, but that's what Master Faders are for—they allow you to trim your final output level to avoid clipping the DAC or 24-bit digital output when your mixed signals leave the Pro Tools mix environment. This is analogous to an analog console mix bus, where you trim the master bus with a master fader to avoid clipping the output circuitry in the console.

Similar to analog mixers, the Pro Tools mixer is comprised of individual input channels and a summing stage. At the input stage, each channel's 24-bit word is multiplied by 24-bit gain and pan coefficients to create a 48-bit result. The new 48-bit word contains the original 24 bits "shifted" lower in the 56-bit register to allow for headroom and "footroom" below unity gain, enabling channels to be turned down without losing precision. Specifically, it's possible to pull any channel fader down to -90 dB and its signal still retains 24 bits of precision. As channel faders are pulled down, there is a loss to the lower bits of the newly extended 48-bit word which represent signals down to about -240 dB—but a full 24 bits of precision is maintained down to -90 dB.

It's important to understand that while some lower bits of the 48-bit word are truncated when reducing individual channel levels, the quantization (or rounding error) that occurs when they are lost adds about *one millionth of a dB* to the total quantization error. This is determined by adding the quantization noise from the single precision quantization to the quantization noise of the double precision quantization. For example, when the mixer represents -144 dB, the actual value is more like -143.99999 dB and it is quantized to -144 dB. The difference is about one millionth of a dB and is astronomically close to the "ideal" mixer—one that doesn't ever lose bits internally.

At the summing stage, Master Faders perform in a similar way to channel faders in terms of resolution. The final gain stage is in effect a master fader, which is always present; and, when the Pro Tools UI displays a Master Fader, the user is given a "handle" to adjust the range of the final gain stage output. As a built-in part of the mixer, Master Faders don't consume any additional DSP resources, so they are the best way to adjust the final output of an internal bus or external pair without losing precision. See figure 3 for a detailed illustration.

SHIFTING EXPLAINED

Let's take a closer look at how the Pro Tools mixer uses bit shifting to optimize headroom and "footroom" while maintaining maximum transparency. Bit shifting describes the movement of a binary word within a register, and is used to increase or decrease gain. It can be an explicit instruction (arithmetic shift left or right) or it can be accomplished by multiplying by a coefficient with power of 2. A negative power of 2 causes a "right shift" (i.e., gain reduction) and a positive power causes a "left shift" (i.e., gain increase). Shifting by multiplication moves the number by the magnitude of the power to which it is raised (e.g., 2 to the 4th power produces a left shift of 4 bits).

In the Pro Tools mixer, input signals are multiplied by coefficients representing pan and gain as they enter the mixer, which produce 48-bit words from the original 24-bit words (a 24-bit number multiplied by a 24-bit number produces a 48-bit number). This initial multiplication operation places the original 24-bit word roughly in the middle of the 56-bit register with leading and trailing zeros making up the remainder of the 48 bits. (This example is simplified for this discussion; other gain and pan coefficients will yield more complex results.)

Following the initial multiplication operation, inputs are summed and right shifted again. Then, the final result of the mixer—the output sent to the DAC or plug-ins on the Master Fader—is the dithered 24-bit result shifted left by the inverse of same amount as the accumulated right shift, thereby returning the original signal level to the DACs.

The process looks something like this:

- ***Figure 3, row A*** illustrates a single, 24-bit, full-code input before it enters the mixer plug-in.
- ***Figure 3, row B*** shows the same input as it is placed in the mixer. Unity gain and pan left has been applied. It is the same signal within a 48-bit word "shifted" 4 bits to the right.
- ***Figures 3, rows C and D*** show 128 full-code tracks summed at unity and +12 respectively, consuming some of the headroom bits.
- ***Figure 3, row E*** shows the output of the mixer as all channels are summed to a pair of outputs. The 48-bit word is explicitly right shifted by 7 bits to pull it out of the overflow area and make up the +12 gain stage.
- ***Figure 3, row F*** illustrates the position of the 48-bit word at the output of the mixer after a Master Fader has been set to -42.1 dB. The effect is to pull the output of the mixer out of clipping and restore it to 0.0 dB.
- ***Figure 3, row G*** shows the 24-bit mixer output from row F as it appears at the DAC. The 48-bit word is dithered before reducing it to 24 bits, thereby preserving low level detail. The result has been shifted to the left by 11 bits to reproduce the original input signal with its 24 bits intact.
- ***Figure 3, row H*** illustrates the use of "footroom." This is a 48-bit word representing a hard panned channel fader pulled down to -42.1 dB. Again, the original 24-bit word remains intact and is right shifted. Notice that there are still 8 bits of footroom left. The 15 bits below unity gain are used to preserve low level detail down to -90 dB. The original 24-bit word does not suffer any loss until a channel fader is moved below -90 dB.

Pro Tools +12 dB Stereo Dithered Mixer

	8 Overflow Bits								Middle 24-Bits																								Lower 24-Bits																							
	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
A									0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1																								
B									0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
C						0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0			
D				0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0					
E									0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
F									0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0
G									0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1																								
H									0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0
I									0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0
J									0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1																								

Figure 3: Pro Tools +12 Stereo mixer from the 56-bit register perspective

- ***Figure 3, row I*** illustrates the mixer output with 128 full-code inputs summed at -42.1 dB and a Master Fader set at 0. The mixer output result is nearly identical to the initial single channel at unity gain as well as 128 channels summed at unity with a Master Fader set at -42.1 dB—all three operations result in almost exactly the same output.
- Finally, ***figure 3, row J*** shows the 24-bit output as it is presented to the DAC. As in row F, the result has been shifted to the left by 11 bits, which restores the original signal.

You can see that the initial multiplication operation on the original 24-bit word produces a new 48-bit result that contains the original word plus some headroom and "footroom." Inputs are then summed with other inputs and another shift occurs, placing the original 24-bit signals back in the middle of the 56-bit register. Except at the final summing stage, the only bits lost in the shifting process are those representing audio at about -240 dB down, and the result is a quantization error of about one millionth of a dB. A final left shift returns the dithered 24-bit signal representing the mixer output to the DACs.

The benefit is that not only is a tremendous amount of headroom for summing signals provided, but even more accommodation is made for preserving the low level details in the audio for when channel faders are set far below unity.

See the FFTs at the end of the article for a clear picture of how the mixer accomplishes the examples above with absolutely no distortion and what dither actually looks like when added.[ii]

FACT CHECK

Why is this so important? Because there is a lot of debate about the summing accuracy of the Pro Tools mixer. This topic is discussed at great length in several audio forums and there are almost as many opinions as there are posters. Unfortunately, there are also many misconceptions surrounding the issue, so we're showing exactly how the Pro Tools mixer works to help people better understand it. This is not an effort to persuade people to mix "in the box," but to clear up the notion that there is something inherently deficient or sonically incorrect by doing so. We want our users to use Pro Tools to make the best sounding mixes they can *in whatever way they choose!*

DOES ADDING INPUT CHANNELS REDUCE DYNAMIC RANGE?

One theory put forward is that each time you double the number of mixer channels, you lose 1 bit of dynamic range (i.e., you must reduce the channel volumes by 6 dB, and therefore lose a bit of precision off the bottom). This would be completely true if the mixer's bit depth was the same as the word length of the system because there is no additional space for the numbers to grow.

A 24-bit mixer summing 24-bit words would drop bits in the following way:

- ***Figure 4, row A*** shows the original, full-code input channel as it appears at the mixer. There is no initial multiplication performed, so it is at full level at both the DAC and the mixer.
- ***Figure 4, row B*** shows 128 summed channels of full code 24-bit audio. A Master Fader would have to be used to pull the output level down by -42.1 dB to pull it out of clipping, but the penalty in this system is a loss of 8 bits.
- ***Figure 4, row C*** shows an individual channel reduced by 42 dB. Since there are no "footroom" bits available, the lowest bits are simply cut off and the result is a 16-bit word.

In the Pro Tools 48-bit mixer however, this is not the case. By shifting the original 24-bit word to roughly the middle of a 56-bit register, there is enough headroom to sum 128 full-code tracks at +12 without internally clipping, as well as enough low level resolution to pull channel faders down to -90 dB without losing any of their original bits.

24-Bit Mixer

24-Bits | *Truncated Bits*

	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
A	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
B	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	
C	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	*0 0 0 0 0 0 0 0*
D	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	*0*	*0*	*0*	*0*	*0*	*0*	*0*	*0*	

A A single 24-bit full code input in a 24-Bit mixer.

B 128 full code inputs summed at unity gain with a master fader used to pull the result out of clipping. 8-bits are lost and result in 16-bits of data.

C 1 full code channel with its channel fader set at -42 dB. Again, 8 bits are lost to truncation and result in 16-bits of data.

D The truncated word with its fader pushed back up to 0 dB. The truncated bits are lost and the result is much different from example "A"

Figure 4: 24-bit mixer

Pro Tools +12 dB Stereo Dithered Mixer

	8 Overflow Bits								Middle 24-Bits																								Lower 24-Bits																							
	55	54	53	52	51	50	49	48	47	46	45	44	43	42	41	40	39	38	37	36	35	34	33	32	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
A									0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1																								
B									0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Figure 5: 24-bit full code signal as placed in the 56-bit register in the Pro Tools Mixer

CAN CLARITY AND DETAIL BE IMPROVED BY SUMMING EXTERNALLY?

Another misconception is that there is clarity and preservation of detail provided by summing externally—mixing to stems and using an analog mixer to recombine them. It's very common in this workflow to apply other sonic treatments from hardware effects like analog compression and EQ—which can be subjectively pleasing. But these should be recognized as elements that are consciously added to the mix, not correcting for a deficient digital mixer. Consider that *any signal that is recorded, panned, summed, or processed (i.e., all signals) is handled by the Pro Tools mixer.* All sub-mixes as well as individual track outputs are handled by the mixer, whether summed internally or sent straight to external outputs.

As with any high-quality digital mixing system, the Pro Tools mixer is designed to provide a result that is sonically transparent—that is the reason this level of precision is provided.

ARE CERTAIN FADER VALUES MORE TRANSPARENT THAN OTHERS?

A third misconception is that any fader values other than unity and even multiples of 6 dB degrade the signal by introducing rounding errors to achieve the desired gain. This belief is probably based on the assumption that 6 dB movements on a fader invoke a bit shift, and is not true because a pure shift can only be done with an *arithmetic instruction* or a coefficient with a perfect power of 2. You'd need a value of 6.0205999132796239042747778944899 (even this is an approximation) to produce a pure shift through multiplication; and, as we described earlier, the Pro Tools UI provides for one decimal place of precision. All other gain values entered from the user interface except unity are scaled and result in 48-bit results, which are handled by the mixer in the same way.

We should take a moment to distinguish between the result of lowering a fader within the context of the internal mixer (where 48-bit precision is maintained) and *lowering the level of a signal that is being sent to a DAC output.*

As with any digital audio system (not just Pro Tools), if you lower a fader assigned to an individual DAC output you can lose resolution and dynamic range, as you are simply reducing the number of bits representing the output signal in comparison to the DAC's theoretical 24-bit resolution. You would be throwing away dynamic range that the DAC can represent if you, say, pulled your fader down to -14 dB and left it there so that peaks hitting the converter were far below full scale. That would decrease your signal-to-noise ratio significantly, so optimizing signal levels going to individual DAC outputs is important to preserving signal quality—you want to get as much resolution out of the DAC that it can offer before clipping to take maximum advantage of its dynamic range performance.

Of course, the inverse is also true: you must lower the output level of the internal mixer to avoid clipping the DAC—that wouldn't be good either! Additionally, if your downstream devices cannot accept the signal level from the DAC when it is fed a healthy signal, you should trim down the analog output level from the DAC, or lower the input level on the receiving side, for best results.

DOES LOWERING THE MASTER FADER COMPROMISE RESOLUTION?

This brings up another important point: while you can sum hundreds of "real world" tracks in the Pro Tools mixer without clipping the "input side" of the mix bus, you can easily clip the "output" side of the bus (when you go to a 24-bit DAC or 24-bit digital output, or submix signals to an Aux Input). That's why it's so important and useful to utilize Master Faders—they utilize no DSP, provide metering that lets you know if you're clipping on the way to your destination, and allow you to trim the output level of your sub- or final-mix *without losing any resolution* (because resolution is preserved by the 48-bit mixer). Because Master Faders preserve resolution when trimming, you don't have to trim down all of the contributing faders to the mix to obtain the optimal output level. This is just like a well-designed analog console—you trim down the master fader to avoid clipping the output circuitry in the console while the mix bus headroom allows the input faders to stay in their sweet spot.

CAN DYNAMIC RANGE BE ENHANCED BY MIXING TO STEMS?

There's also the notion that you can increase dynamic range by mixing to stems and using an external summing mixer. Consider the analog audio is leaving the Pro Tools mixing environment via a 192 I/O's digital to analog converters, which deliver about 116 dB of dynamic range. From there, you are subject to the performance of the downstream analog devices, which may affect the sound in a subjectively pleasing way by imparting their own non-linear characteristics, but there isn't a way to increase dynamic range—it just isn't physically possible.

DELAY COMPENSATION AND OUTPUT MONITORING

Phase problems have also been an issue in the past, so it should also be mentioned that since Pro Tools version 6.4, the entire mixer uses Automatic Delay Compensation which ensures that signals remain phase coherent, no matter what routing techniques and plug-ins are used. In versions of Pro Tools before 6.4, double bussing and using different plug-ins and analog inserts could create phasing issues that, if not manually compensated for, would wreak havoc in the mix. By enabling Automatic Delay Compensation, the entire mixer dynamically adjusts itself to accommodate channel to channel differences caused by processing and inserts.

Even though lowering faders to very low levels is possible without losing bits, the system works best with analog volume control before connecting to your monitoring system. Avoid connecting the outputs of the 192 I/O interfaces directly to the monitors and pulling down channel faders or using a Master Fader like a volume control to achieve normal listening levels, as this causes the converters to operate at very low levels where their non-linearity is most pronounced, producing mixes with extremely low levels (see the info about optimizing DAC levels above).

The main point to remember is that not only can you mix with confidence completely within the Pro Tools environment, or 'in the box,' but you can produce absolutely stunning mixes that are not possible in any other way. We're not trying to persuade you to work one way or the other—your workflow may include a 'hybrid' mixing environment, and that's great. Just be sure you are comparing apples to apples, attribute what you hear to the correct components, and know that the Pro Tools mixer is in no way degrading the sound you work so hard to get.

This is probably a good place to wrap up. Thanks for sticking with it! Of, course there are many factors that affect the quality of digital recording and we'll talk more about them in upcoming articles.

One forum you should check out is www.3daudioinc.com. Lynn Fuston created a set of CDs called the "Awesome DawSum". Even though it was done in 2003, it's still very appropriate and ear opening. Lynn and a team of volunteers carefully made comparable mixes of a tune on 29 different platforms, including analog mixers, digital mixers, and DAWs. Listen for yourself and compare the tracks, then look at the key to reveal what was what; the results are fascinating and not *at all* what you'd expect!

Please send your feedback on this article and requests for other discussion topics to techpapers@digidesign.com

---

[i] Lipshitz, S.P., R.A. Wannamaker and J. Vanderkooy, "A Theoretical Survey of Quantization and Dither," *J. Audio Eng. Soc.,* vol. 40, 1992 May, pp. 355-375.

[ii] **SOME ARTWORK**

The following is a set of FFT plots capturing the results of these tests. The source audio was a pure full-scale 24-bit sine wave at 997 Hz sampled at 44.1 kHz. This test tone is special in that it is exceedingly pure. It just contains a pure sine wave with a white noise floor. (For those who want to know, the plots are 32k point FFT with a Blackman window and 16 averages.) The setup bussed the test audio to 128 Aux Inputs, to another bus with a Master Fader assigned it, then to another audio track and recorded. The resulting files were then analyzed in MatLab and FFTs created from the resulting files.

First, here is an FFT of the original audio file. You'll notice the level of the noise is around -180 dB. That's because every time you double the number of sample points, you reduce the noise measurement by 3 dB. In this case we have 32,768 sample points or 2 to the 15th power and 16 averages or 2 to the 4th. You subtract the averages from the doublings which gives you 11—11 x 3 = 33 dB difference between the plot and actual noise floor. The average noise floor of the test signal is around -147 dB.

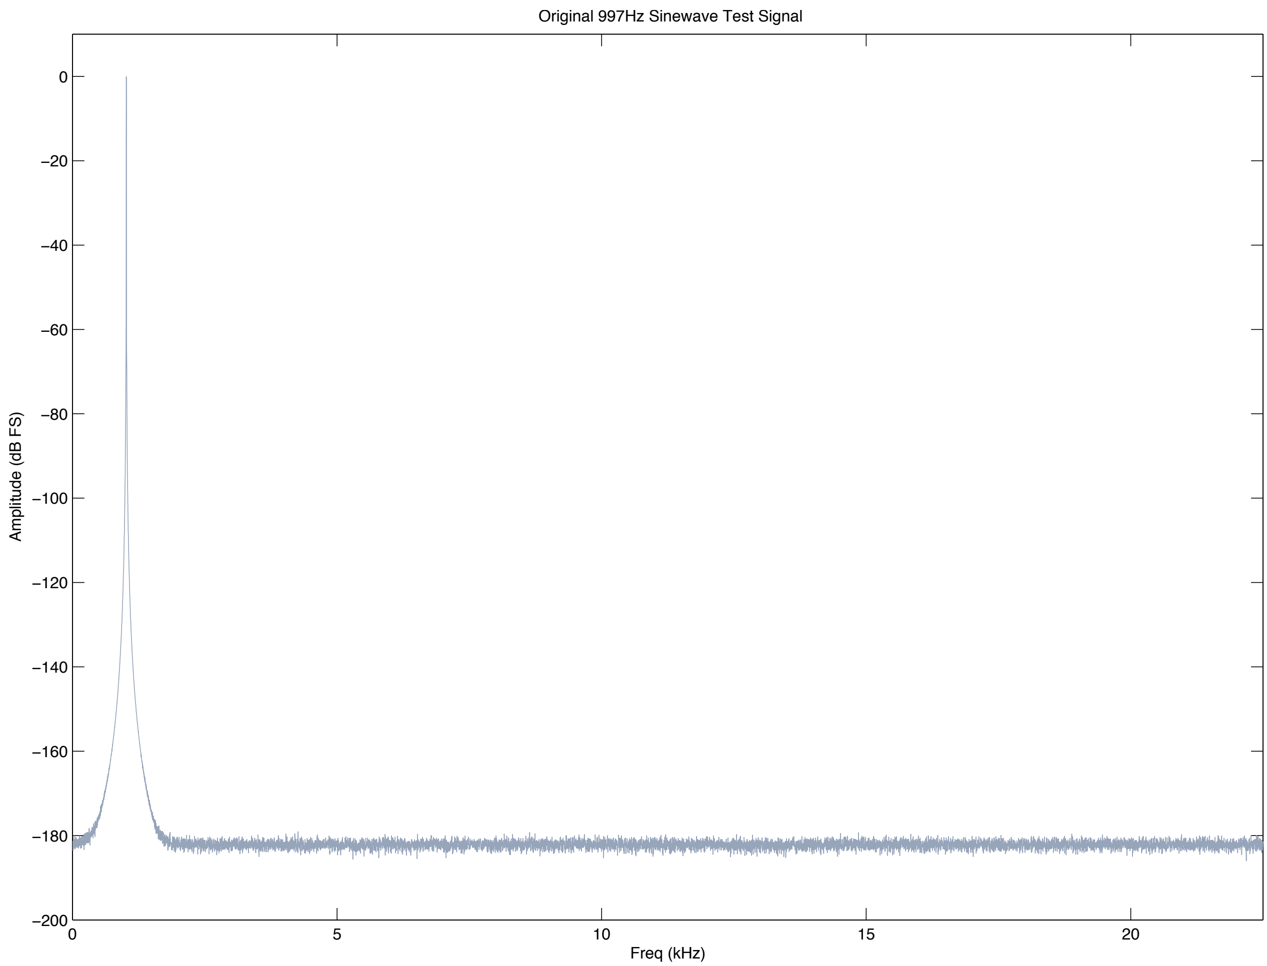

Next is the result of busing the test signal to 128 Aux Inputs set at unity gain, through a Master Fader set at -42.1 dB and re-recorded on an another audio track. You can clearly see that no distortion was added, although the noise floor has increased slightly by about 3 dB which corresponds to the 1-bit dither added at the final summing stage.

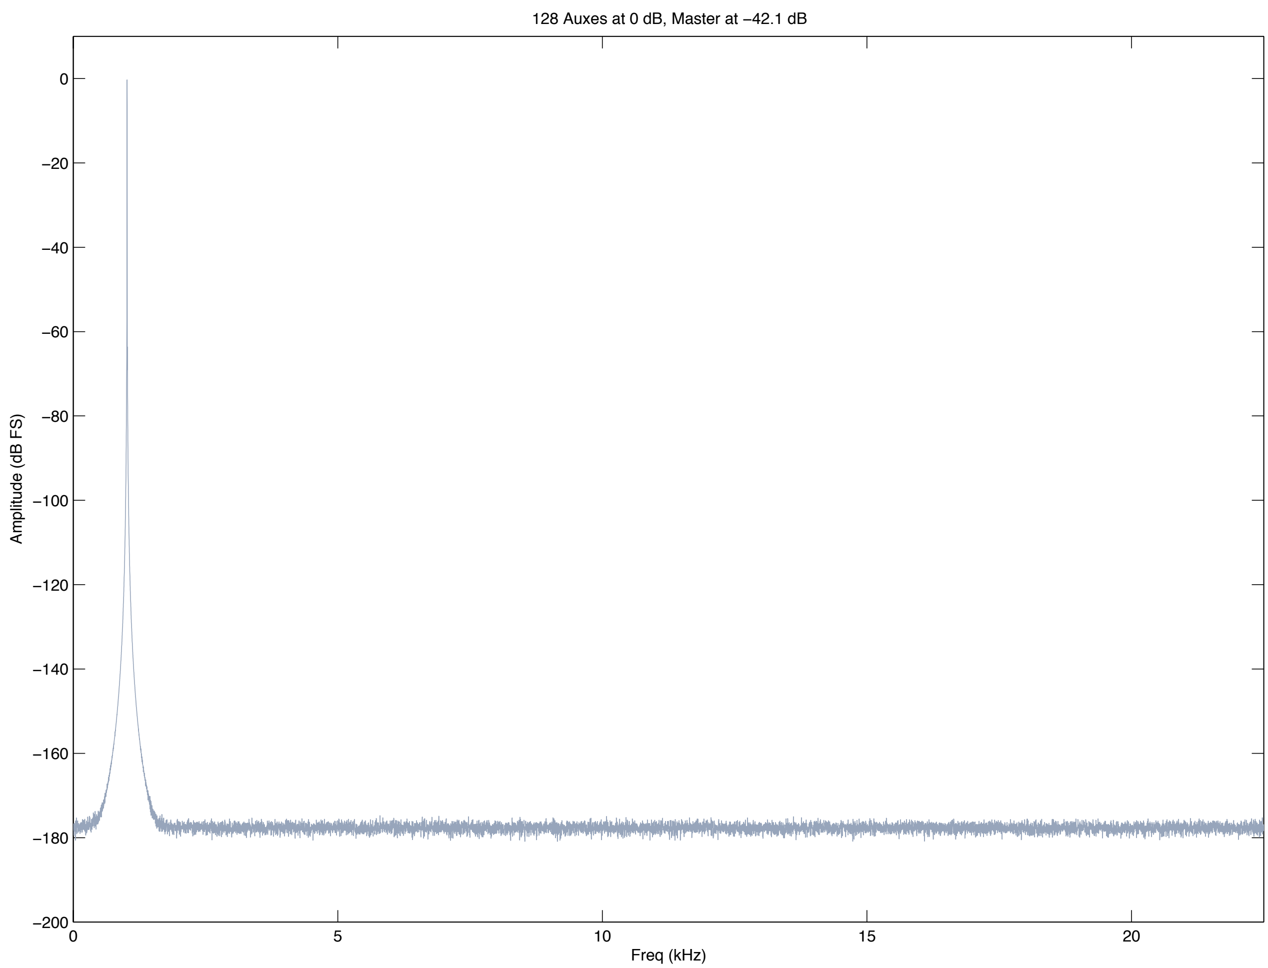

Next is the result of setting all 128 Aux Inputs at -42.1 dB and the Master Fader at 0.0 dB. The same result as above.

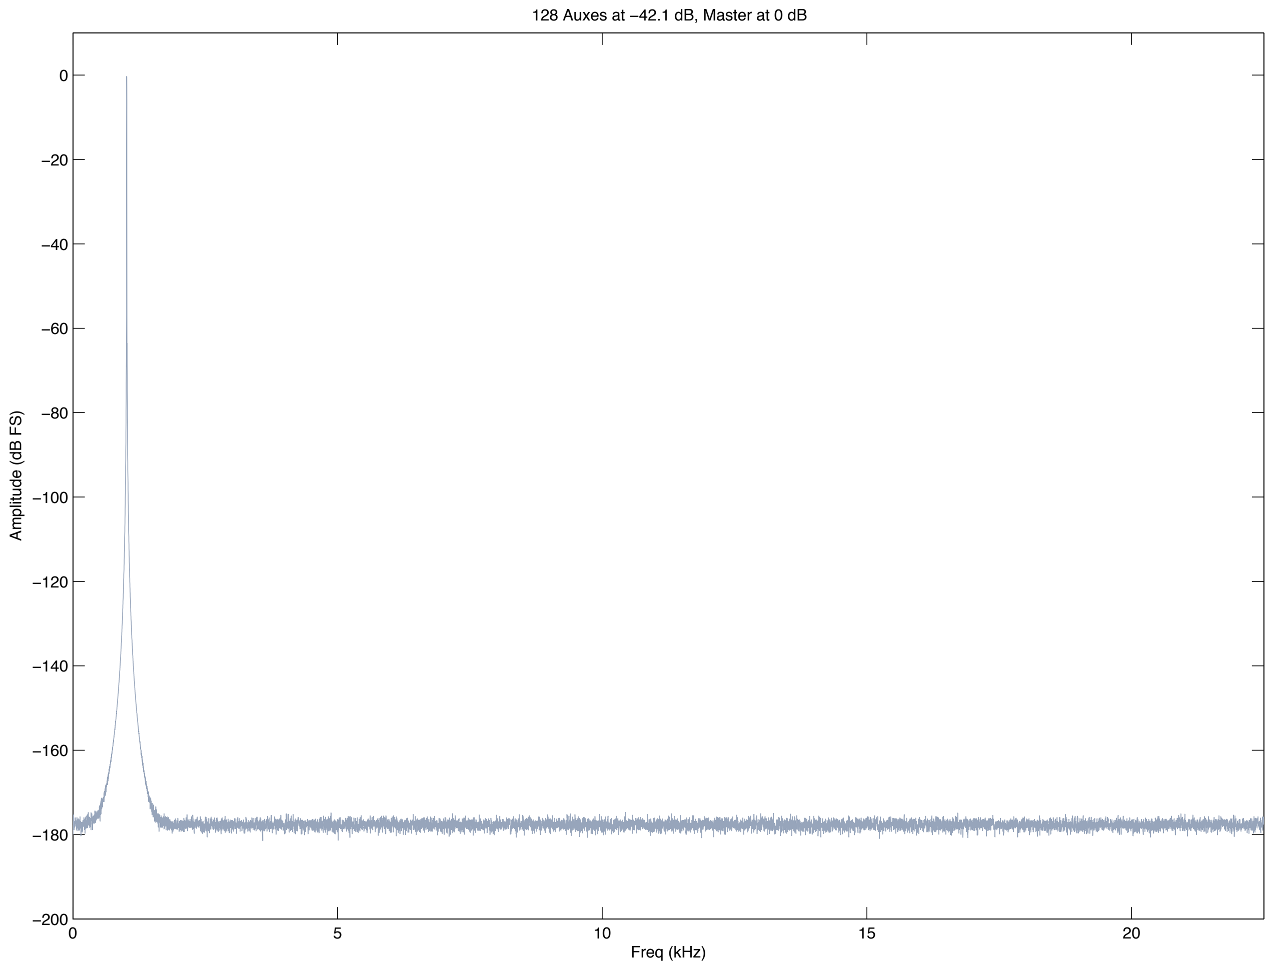

Finally, here is the output of the mixer with all 128 Aux Inputs set at +12 dB and the Master Fader set at -54.1 dB. The same result as above.

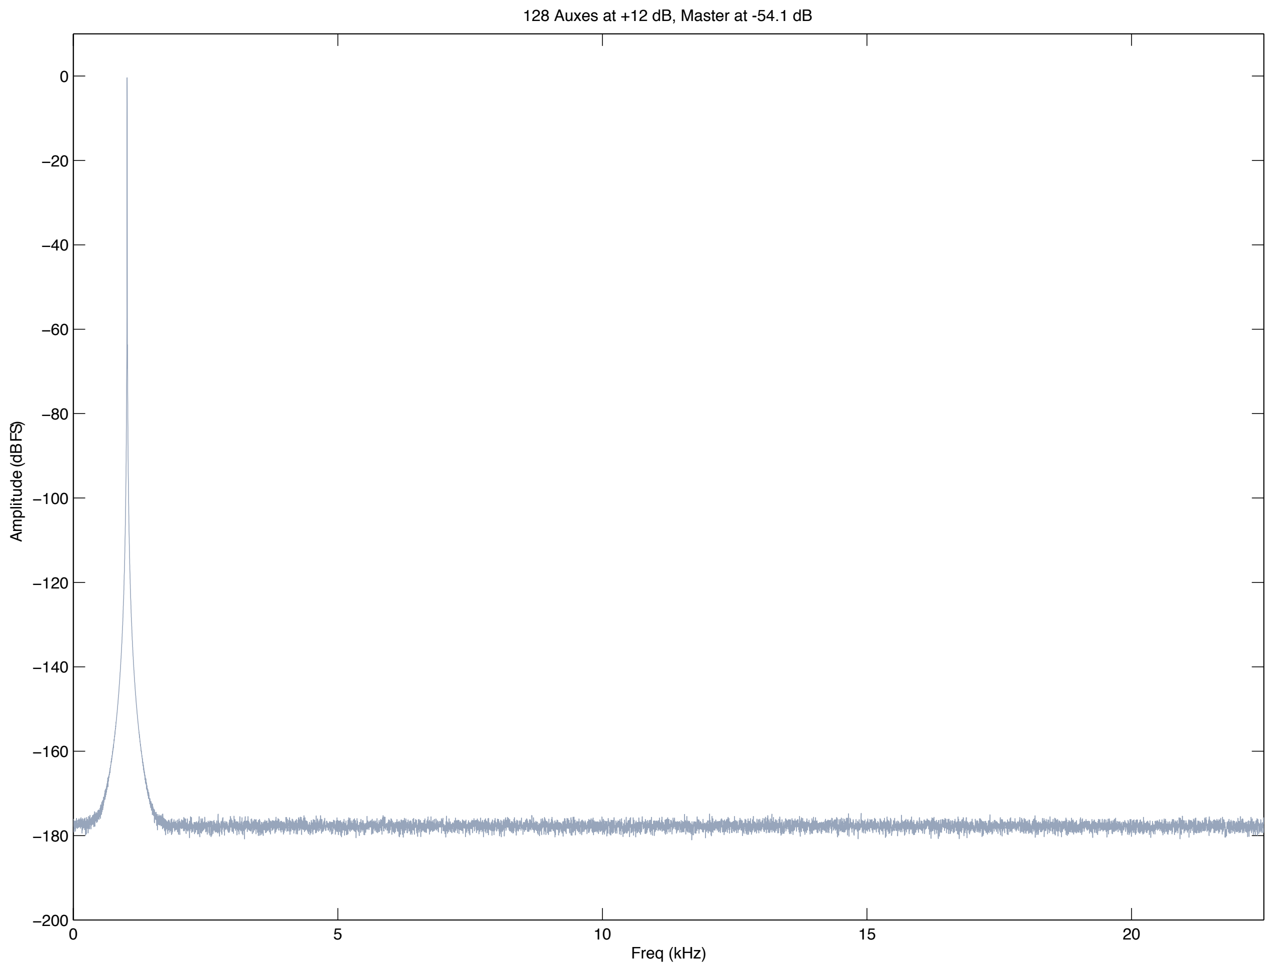

As you can see, this result—and those in the other examples—is a perfect sine wave with no clipping or distortion artifacts. There is a slight level difference in the examples due to less precision than is required in the user interface faders to represent the gain value required to obtain "bit-matched" outputs, but the output is completely free of distortion. In all cases, the actual measured dynamic range of the Pro Tools mixer output is 144 dB.

www.digidesign.com

DIGIDESIGN
2001 Junipero Serra Boulevard
Daly City, CA 94014-3886 USA
Tel: 650.731.6300
Fax: 650.731.6399

TECHNICAL SUPPORT *(USA)*
Tel: 650.731.6100
Fax: 650.731.6384

PRODUCT INFORMATION *(USA)*
Tel: 650.731.6102
Fax: 800.333.2137

INTERNATIONAL OFFICES
Visit our Digidesign Web site
for contact information