Revision

1

BENCHMARK MEDIA SYSTEMS, INC.

AD2404-96 and AD2408-96 – MULTI-CHANNEL 96 kHz ANALOG TO DIGITAL CONVERTER

Operating Manual

AD2404-96 & AD2408-96 - ANALOG TO DIGITAL CONVERTERS

Operating Manual

5925 Court Street Road
Syracuse, NY 13206-1707
Phone (315) 437-6300• Fax (315) 437-8119
www.benchmarkmedia.com

Table of Contents

Quick Start Guide 6

- Quick System Overview 6
- Power Supply Connection 7
- Understanding the Status LEDs 7
- How to Connect the Audio Interfaces 8
- How to Set the Front Panel Switches 9

Design Philosophy 12

- Design Goals 13
- Design Methodology 14

System Overview 15

- AD2008 Function Blocks: 15
- Analog Input Stage 15
- CS5396 Converters 16
- AES Input / Clock Recovery 16
- AES Outputs 16
- Digital System Control 17
- Digital Level Meters 17

Jumper Settings 18

Using the Meters 19

- Meter Time Constants 19
- Fast Decay Time-Constant 19
- Slow Decay Time-Constant 20
- Peak Hold Function 20
- Selecting Meter Range 20
- What is “FSD”? 20
- How is “FSD” Different from “0 dBFS”? 21
- How is “FSD” Different from “Over”? 21
- The Fallacy of “Over” Indicators 21
- Do “Over” Indicators Have a Place? 23

Jitter 24

- What is Jitter? 24
- When is Jitter Low Enough? 25
- The Tape Machine - A Jitter Analogy: 26

Signal Interconnect 27

- Analog Audio Inputs 27

Jumpers 29

- Option Jumpers for Revision 1.XX 29
- Digital I/O Jumpers 30
- Miscellaneous Jumpers 30

Dither, Word Length Reduction, and Noise Shaping – a Tutorial 31

- Why is Word Length Reduction Necessary? 31
- Quantization Noise 31
- How does WLR create Quantization Noise? 32
- Can we Reduce Quantization Noise? 32
- How does Quantization Noise Sound? 32
- Rounding vs. Truncation 33
- Dither 33
- What is a “Noise Floor”? 35
- Word Length Reduction Techniques 35
- Truncation without Dither 35
- Flat Dithered WLR 36

Dithered Noise-Shaped WLR 37

“IR” Word Length

Reduction System 40

IndexError! Bookmark not defined.

Chapter

1

Quick Start Guide

For those of us who hate to read manuals first.

This guide will take you thorough a typical setup of the AD2408-96 and will direct you to other sections of this manual if your installation has special requirements.

Quick System Overview

The AD2408-96 is an eight-channel 24-bit analog-to-digital converter specifically designed for applications requiring eight or more perfectly phased, low jitter, low distortion, A to D conversion channels. The AD2408-96 has two **fully independent four-channel converter modules** that share a common digital-audio reference input. The AD2404-96 is a four-channel version of the AD2408-96, and uses only one of these converter modules. This manual applies to both products.

The converter operates over a **sampling frequency range of 24 to 100 kHz**, and provides **24, 20, and 16-bit word lengths**. The 20 and 16-bit word lengths are re-dithered from 24-bits using **selectable noise-shaping curves** or TPDF dither. All noise-shaping curves are optimized for the selected sample rate, and for the number of bits being removed.

A **digital-audio reference input** jack allows the AD2408 to be locked to an external digital-audio reference signal. This jack also serves as the input for Digital-to-Digital processing.

Digital-to-Digital processing functions include Word Length Reduction (re-dithering), Sample Rate Conversion, and, conversion between "Single Cable" and "Dual Cable" interface modes.

All digital inputs and outputs support sample rates from 28 kHz to 100 kHz using the **"Single Cable" interface mode** (also know as "AES3 Multichannel Mode"). In this mode, a channel pair is transmitted on a single cable, at a frame rate equal to the sampling rate.

A **"Dual Cable" interface mode** (also known as "AES3 Single Channel Double Sampling Frequency Mode") allows 88.2 and 96 kHz recording on 44.1 and 48 kHz equipment. In this mode, two successive samples of a single high sample-rate channel are transmitted in place of a pair of low sample-rate channels. Two cables are required for two channels.

Each conversion channel has a fully digital **multi-function 9-segment LED meter**.

Power Supply Connection

The AD2404 is available with either an internal international power supply, or a pair of 9-pin D-sub connectors for connection to an external supply.

Internal Supply: **Check voltage selector switch for proper AC input voltage.** If voltage setting needs to be changed, remove power cord, open the access door, and then remove voltage selector cam. **Do not rotate cam while it is inserted in the housing, this may cause damage to the voltage selector switch!** Reinsert the cam with the proper voltage selection facing outward. Close the access door, insert the appropriate IEC terminated power cord, and then apply power.

External Supply: **Connect the converter to the external DC supply using either nine-pin D-sub connector on the rear panel, then apply power.** A second D-sub connector is provided for daisy chain wiring. For power supply requirements, specifications, options, and pin assignments, see page ????

After applying power, the "POWER" LED will light, and the "LOCK" LED will flash up to eight times indicating that automatic calibration is in progress.

Understanding the Status LEDs

The "POWER" LED is located at the far left of the front panel. It is driven from the +5 V digital supply and indicates that the 4-channel converter module is receiving power. To avoid accidental shut down, there is no power switch on the AD2404. The converter is designed for continuous duty.

The "PHASE LOCK (Slave)" LED is located on the front panel above the "POWER" LED. **This light will turn on whenever the AD2404 is locked to a reliable digital audio reference of the correct frequency.** If the light is off, the AD2404 is operating as a master sync generator using an internal crystal reference. **A flashing light indicates an error condition.** For more information see page ???

How to Connect the Audio Interfaces

Analog Audio Inputs may be balanced or unbalanced. Input reference level is +4 dBu. Headroom is 20 dB. In other words, 0 dBFS (the full-scale digital clip point of the converters) will be reached when the input levels reach +24 dBu. **Make sure your signal source can achieve levels of at least 24 dBu without clipping.** An optional "Variable Gain" motherboard is available. Converters equipped with this option have jumpers and gain trim pots which can be changed to select 0 dBFS clip points ranging from +4 dBu to +24 dBu (see page ?????).

Digital Audio Interfaces are available as either an **XLR equipped AES3 ("R" option)** or as a **BNC equipped SMPTE 276M ("B" option)**. The "R" option provides one 110-ohm AES3 digital audio input, and four 110-ohm AES3 digital audio outputs. The "B" option provides one 75-ohm SMPTE 276M digital audio input, and eight 75-ohm SMPTE 276M digital audio outputs. All **outputs use professional status bit formats**.

> Note: **All four-channel converters and all "B option" converters have two sets of digital outputs**. These additional outputs provide some additional features. In most modes of operation, the **primary outputs** have adjustable word lengths, while the **auxiliary outputs** have fixed 24-bit word lengths. This allows simultaneous 24 and 16 bit recordings from the same converter. Also, if the "Dual Cable" modes are used, there are enough digital outputs to provide 8 channels of conversion. "Dual Cable" operation on eight-channel "R" option converters is supported but is limited to 4 channels. (see page ?????)

The 110-ohm AES3 interface is by far the most popular interconnect for professional digital audio equipment. However, the 75-ohm SMPTE 276M interface is rapidly gaining popularity in applications that require long transmission distances. The SMPTE 276M interface utilizes 75-ohm coax and can easily achieve transmission distances of 1000 feet without cable EQ, and 3000 feet with cable EQ. AES3 and SMPTE 276M use identical data formats. Note that SMPTE 276M is a formal standard that references AES3-id (an earlier information document issued by the AES). Consequently, SMPTE 276M interfaces are often called AES3-id. SMPTE 276M is slightly more specific than AES3-id, but the two interfaces are fully inter-operable. **Use 110-ohm digital audio cable for AES3 connections, or 75-ohm coax for SMPTE 276M connections.** Incorrect cable impedances may increase jitter, may cause data loss, and will reduce the maximum transmission distances. Standard analog audio cable should be avoided.

The Digital Audio Reference Input accepts either **professional or consumer** status bit formats. The PLL automatically supports 1:1, 2:1 and 1:2 frequency ratios when a fixed sample rate is selected. The PLL supports 1:1 or 2:1 frequency ratios when variable sample rates are selected. On "R" option converters, the reference input is a female XLR connector terminated with 110 ohms. On "B" option converters, the reference input is a BNC connector without an internal termination. This high impedance BNC input allows looping of a single reference input to multiple converter frames. **The BNC reference input on "B" option converters must be externally terminated using the supplied BNC "T" connector and 75-ohm BNC terminator.** (see page ?????)

A reference input is required when:

- The converter will be operating in a variable speed mode ("VAR 1:1", or "VAR 2:1").
- More than four channels must be phase locked together.
- The converter is to be locked to an external studio reference.
- Any "Digital-to-Digital" mode is selected.

If a reference is required, but no reference input is detected, the "LOCK" LED will flash rapidly. If a reference is present, but lock has not been achieved, the "LOCK" LED will flash slowly.

> Note: Using an external reference will not degrade or alter the low-jitter performance of the AD2408 converter at any of the fixed sample rates. **The jitter performance of the fixed sample rate modes is maintained even when the reference has very high levels of jitter.** However, the "VAR 1:1" and "VAR 2:1" **variable speed modes have wide frequency ranges** that preclude the use of the final VCXO stage of the AD2408 PLL. **If a variable speed mode is used, it is important to provide a low-jitter reference signal.** Use the fixed sample rate settings whenever possible.

How to Set the Front Panel Switches

The "SAMPLE RATE" Switch is a twelve-position rotary switch, which is located on the left side of the front panel. This switch selects the sample clock frequency, the PLL (Phase Locked Loop) mode, dual or single cable interface modes, and digital-to-digital functions. **For typical A/D conversion applications, select one of the fixed "Single Cable" sample rates ("44.1", "48", "88.2", or "96").** This will set the free running frequency of the converter, and will allow automatic PLL operation if a reference of the proper frequency is applied to the digital audio input. When a fixed sample rate is selected, the PLL automatically supports lock ratios of 1:1, 2:1, and 1:2. For example, 96kHz conversion can be locked to either a 48 or 96 kHz AES digital audio reference. Use the "VAR 1:1" or "VAR 2:1" for non-standard sampling rates. "VAR 1:1" locks in a 1:1 ratio to the reference sample rate. "VAR 2:1" locks in a 2:1 ratio to the reference sample rate. For more information see page ????.

> Note: Converters will enter a **calibration mode** whenever the "SAMPLE RATE" switch is rotated. Calibration is completed within 5 seconds.

The "WORD LENGTH" Switch is a twelve-position rotary switch, which is located to the right of the "SAMPLE RATE" switch. It **selects 24, 20, or 16-bit word lengths at the main outputs**. All 16 and 20-bit Word Length Reduction (WLR) modes are TPDF dithered prior to noise shaping.

There are four word length reduction modes ("TPDF", "NS1", "NS2", and "NS3"). All forms of word length reduction raise the noise floor of a digital transmission system. TPDF dither is spectrally flat, it is not shaped, and it will sound noisier than the noise-shaped modes. "NS1", "NS2", and "NS3" are noise-shaping modes which are psycho-acoustically optimized to take advantage of the ear's low-level sensitivity curve. "NS1", "NS2", and "NS3" will sound 6, 12, and 18 dB quieter than "TPDF" respectively. "NS1", "NS2", and "NS3" can provide 17-bit, 18-bit, and 19-bit performance respectively at a 16-bit word length. At 20-bits, "NS1" and "NS2" provide 21-bit and 22-bit performance respectively. **Use the maximum word length that is compatible with your digital audio equipment. Maintain 24 or 20-bit word lengths as long as possible. If WLR to 16-bits is required, do so at the latest possible point and time. Avoid processing 16-bit signals. In general, NS2 will produce the quietest 20-bit signal, and NS3 will produce the quietest 16-bit signal.** The other WLR settings are optimized for special circumstances. See page ???? for details.

> Note: The **auxiliary outputs** (available only on 4-channel "R", and 4 or 8-channel "B" converters) are always **fixed at 24-bits**. These outputs allow simultaneous 16 and 24-bit outputs from the same converter.

The "METER" Switch is a three-position toggle switch located to the right of the sync switch. It selects one of **three digital meter functions** and may be switched at any time. The down position sets the meter scale to **6-dB steps**. The center position sets the meter scale to **1-dB steps**. The up position sets the meter scale to **1-dB steps with peak hold**. Moving the switch to the center position will clear the peak hold. Note that the bottom of the 1-dB scale is expanded and includes a -20 dBFS LED that can be used to set the input level relative to a 0-dB house reference. **Start with the meter**

switch in the down position, as this will make it easier to verify that signals are present at the analog inputs.

Chapter

2

Design Philosophy

Our Goals and Design Methodology

The AD2404R was designed by John Siau, Allen H. Burdick, and Ralph Henry at Benchmark Media Systems, Inc. It is carefully engineered to reliably provide the highest possible audio transparency. We have not tried to add "warmth" or "color" to the audio. Instead, we have attempted to produce a piece of equipment that sounds as close to a piece of wire as possible.

Converters are often viewed as digital products, and are often designed by digital hardware and software engineers who may lack experience in analog audio. The importance of the analog circuitry is often overlooked, and the resulting defects can easily go undetected on the bench. Unfortunately, our ears often detect these defects. Our ears have a dynamic range of about 130 dB, we are able to hear tones that are 20 or 30 dB below a 20 to 20 kHz white noise signal, and we are able to hear multiple tones of various amplitudes simultaneously. Many audio measurements are only capable of measuring the tone having the highest amplitude.

Because bench tests have often failed to detect audio defects, some have discounted their value and have attempted to rely primarily upon listening tests. We feel that listening test are necessary to agree that ultimately the way a pUnfortunately, listening tests inexact,

There are a number of factors that can contribute to poor converter performance. Noise, THD, IMD,

under both ideal and adverse operating conditions. For example, the converter is designed to tolerate RF interference, static discharge, common-mode interference, line noise, high jitter reference signals, and we have enclosed it in a heavy gauge chassis.

One of our goals was to create a reliable 24-bit converter with the highest performance available. In addition, we have endeavored to create one of the most complete and useful feature sets available in an A to D converter. We have leveraged our experience with single-chip FPGA based digital processing to add features without adding the cost, size and noise penalty of additional hardware. All processing within the FPGA is synchronous with the converter clocks, and digital to analog crosstalk is below measurement limits.

The AD2408 feature set includes one of the most advanced and transparent word length reduction systems available. We began our design process by carefully, measuring, and evaluating the currently available word length reduction systems. We identified several specific opportunities for improvement, and enlisted the design services of the Audio Research Group, at University of Waterloo. Special thanks are in order to Robert Wannamaker, Stanley Lipshitz, and John Vanderkooy for their research in the field of noise-shaping, and to Robert for his custom designed noise-shaping curves which are specifically tailored to maximize the performance of the AD2408 converter. All design curves are based upon the most recent psycho-acoustic data, are individually optimized for sample rate, and are matched to our converter's dynamic range.

Design Goals

- Highest level performance of any 24-bit design
- Low per-channel
- Consistent and repeatable performance
- Low jitter even when locked to a high jitter reference signal
- No performance degradation when phase locked
- High quality word length reduction
- High density package
- Useful level meters
- Sample, frame, and block accurate phase locking of any number of channels
- RF and static discharge immunity

At Benchmark Media Systems:

> Reliability is: The ability to achieve and consistently maintain specified performance under a wide range of adverse operating conditions.

True performance is only achieved if "real world" results consistently match tests conducted in a controlled laboratory environment.

Design Methodology

- Computer aided circuit analysis
- RF design techniques
- Proper grounding techniques
- Extensive shielding and isolation of digital signals
- Extensive power supply isolation
- Extensive testing on AP System Two
- Simulation of adverse operating conditions
- Low level, and high level listening tests
- ESD and RF immunity testing
- FCC emissions testing

Chapter 3

System Overview

An Inside View of the AD2408

The AD2408 has five major function blocks. In combination, these blocks allow the AD2008 to set new performance benchmarks. These outstanding benchmarks include; low-noise, low-distortion, low-jitter, and high-jitter immunity.

AD2008 Function Blocks:

- Analog Input Stage
- CS5396 Converters
- AES Input / Clock Recovery
- AES Outputs
- Digital System Control
- Digital Level Meters

The first four blocks are in the audio "critical path". As such, these blocks are designed to achieve maximum audio transparency. The fifth block controls the meters, the PLL, the digital I/O, and the user interface. While this fifth block is not in the audio "critical path" is certainly could cause interference with the "critical path" if it were inadequately isolated from the audio. Finally, the digital level meters are designed to efficiently convey useful signal level information.

Analog Input Stage

The analog input stage is designed to provide the ultimate in transparent and uncolored audio. To this end, the input stage is designed to provide frequency response which is flat to better than +/- 0.01 dB from 10 to 20 kHz. Frequency response extends well beyond 200 kHz. Phase non-linearity is less than ???? degrees from 20 Hz to 20 kHz.

A proprietary analog preprocessing circuit reduces the odd harmonic distortion that is typical of most analog-to-digital converters. The result is that the THD+N of the AD2408 is limited by white noise rather than by distortion products. A second and perhaps more important achievement is that IMD is reduced. This is important since IMD products are not harmonically related to the fundamental, and are therefor much more objectionable than harmonic distortion.

With the proliferation of digital audio equipment, computers, and wireless technology, our studios and recording venues are rapidly becoming high RF environments. For a number of years, Benchmark has been building Microphone Preamplifiers and Audio Distribution Amplifiers for use in broadcast facilities. Many of these facilities are located at or near transmission sites and must perform flawlessly in high RF fields. The AD2408 brings this technology to digital audio equipment. To prevent RF interference, the AD2408 has two stages of passive RF filtering prior to the active section of the input amplifier. The active section itself is wide-band and RF-stable. The result is freedom from the "unexplainable" grunge and distortion that can result when RF interference signals cause an input stage to oscillate or clip.

CS5396 Converters

The CS5396 analog to digital converters

AES Input / Clock Recovery

AES Outputs

The standard cable sets supplied with the AD2004R provide access to the 110 ohm balanced AES/EBU outputs. The 110 ohm interfaces are isolated with high quality shielded transformers. Transient voltage protection is provided by Schottkey diodes. The outputs will withstand direct hits from 8000 volt static discharge, accidental connection to phantom power, and short circuits. The AES/EBU outputs are best suited for transmission distances of less than 1000 feet. Specially designed 110 ohm digital audio cable should be used. It is absolutely essential to use properly terminated digital audio cable when transmission distances approach 1000 feet.

Two additional 1 Vpp, 75 ohm unbalanced, coaxial, SMPTE 276M digital outputs are provided at the 26 pin D-sub connector. These outputs support long transmission distances; 1000 feet without cable EQ, and over 3000 feet with cable EQ. These 75 ohm digital outputs are rise time limited (per SMPTE 276M-1995) to allow the distribution of digital audio using non-clamping NTSC or CCIR video distribution amplifiers. In addition, the outputs are designed with accurate 75 ohm source impedance which extends above the bandwidth of the output. This unique Benchmark design eliminates the possibility of high frequency standing waves in the coax cable. The are protected from transients using Schottkey diodes and the output filter itself.

The 110 ohm and 75 ohm outputs are available simultaneously and are isolated from each other. Thus 2 outputs are available for each channel pair. The 75 ohm output could be used to drive a long cable while the 100 ohm output could be used for local monitoring. Both are driven from the same active electronics, and therefore, the local monitor provides full confidence that the remote feed is active.

Digital System Control

Digital Level Meters

Chapter

4

Jumper Settings

Setting Analog Input Level Preset Jumpers.

The AD2004R has jumpers for selecting preset input levels, and for enabling the front panel level controls. Presets allow setting the 0 dBFS clip point at +28 dBu, +24 dBu, and +22 dBu. Use the preset levels whenever possible. The preset levels are matched to an accuracy of better than +/- 0.3% and are not subject to the possibility of potentiometer noise.

Table 1

0 dBFS	REFERENCE LEVEL (20 dB headroom)	A	B	C	Cable from Front Panel Pots
+16.5 dBu	-3.5 dBu	OFF	OFF	OFF	P#0 **
+22 dBu	+2 dBu	ON	OFF	OFF	P#0 **
+24 dBu	+4 dBu	OFF	ON	OFF	P#0 **
+28 dBu	+8 dBu	OFF	OFF	ON	P#0 **
Variable (+18 dBu to +28 dBu)	Variable (-2 dBu to +8 dBu)	OFF	OFF	ON	P#1 **

** # = channel number

Chapter

5

Using the Meters

9 Segment LED Meters.

The AD2004R has a nine segment LED meter for each of the four audio channels. The meters are fully digital and respond to both positive and negative going peaks. Thresholds are determined by digital comparators, and therefore are exactly matched between channels. Dual time constants extend the on-time of each LED so that peaks having a duration of only one sample can be displayed and measured accurately. The "FSD" clip indicator is accurate to one quantization level.

Meter Time Constants

The meters on the AD2004R have instantaneous peak response. In other words, the amplitude of a single sample will read accurately on the meter. However, the response time of the human eye is much too slow to allow us to see an LED light up for one sample of a 44.1 kHz or 48 kHz clock. Therefore, the meters on the AD2004R incorporate a decay time constant which extends the on-time of all LEDs, and a second slower time constant which extends the on time of the highest LED triggered by a peak. The result is that any event having a duration of only one audio sample is easily observed.

Fast Decay Time-Constant

The fast time-constant is active in all meter modes. Its purpose is to compensate for the relatively slow response of the human eye. Here is how it works: If any segment of the LED meter turns on, it and all of the segments below it, are held on for 375 samples (or 7.8 mSec). This 7.8 mSec on-time is just long enough to make the LED clearly visible to the human eye. At the end of the 7.8 msec delay, the first time constant releases its control of the meter segments. If a higher peak should occur during the 7.8 msec interval, this new peak will be displayed, and the 7.8 msec timer will restart. Thus no peaks are ever missed, and all are visible.

Slow Decay Time-Constant

The slow time-constant is active whenever the peak hold function is off. After the first time constant releases control of the meter segments, a second 0.5 sec time constant will continue to keep the highest illuminated LED lit. All meter segments below this LED will continue to display peaks using the 7.8 msec time constant. At the end of 0.5 sec, the highest illuminated LED will shut off, and the LED below it will turn on. This will restart the 0.5 sec timer. Thus, peaks will decay at a rate of one meter segment every 0.5 seconds. However, if at any time, an audio peak occurs which is higher than the one being held, the new peak will be held and the timer will restart. Again, no peaks are lost, and all are visible.

Peak Hold Function

Moving the meter control switch all the way up enables the peak hold function. This sets the slow time constant to infinity. The highest peak will be held, all segments below the peak level will be controlled by the fast (7.8 msec) decay time constant.

Selecting Meter Range

The AD2004 has two meter ranges; one with 6 dB steps, and one with 1 dB steps. The 6 dB scale allows monitoring for signal presence, as well as coarse adjustment of levels. The 1 dB scale allows highly accurate adjustment of digital levels. In addition, the bottom two steps of the 1 dB scale are expanded 4 dB and 10 dB steps. ranges have expanded step sizes at the bottom end of the scale. More specifically, the 6 dB scale is expanded to the scale is expanded at the low end. In either range, a light will not light until the appropriate threshold is reached. In other words, the -1 dBFS light will remain off until a digital code equal to or exceeding - 1 dBFS is encountered. Therefore, a signal at -1.01 dBFS will read -2 dBFS on the meter. This guarantees that the step between -2 and -1 is the same size as the step between -1 and FSD.

What is “FSD”?

“FSD” stands for “Full Scale Digital” and has a very specific and slightly different meaning than “0 dBFS”, or “Over”. It is important to understand the difference.

The “FSD” indicator will light whenever the minimum or maximum digital code of the converter is reached for a duration of one sample or more. No other digital codes will ever cause the “FSD” indicator to light.

More specifically, AES3 and SPDIF digital audio transmission systems use twos-complement notation so that both positive and negative voltages can be represented.

In 20-bit twos complement hexadecimal MSB first notation, codes 7FFFF and 80000 will cause the FSD indicator to light.

How is "FSD" Different from "0 dBFS"?

Ideally there is no difference, but in practice, 0 dBFS has often been used to describe any digital code which is very close to full scale. "FSD" is used to describe only the minimum and maximum digital code.

An "FSD" meter must have knowledge of the digital word length in order to work properly. A 16-bit "FSD" code will not register "FSD" when feeding a "FSD" meter that is expecting a 20-bit word length. The reason for this is that the 16-bit word will have 4 trailing zeros appended to it (to make it a 20-bits word), and the 20-bit "FSD" meter will interpret this as a level which is 16 codes below full scale. Fortunately testing for non-changing trailing bits easily solves this problem.

How is "FSD" Different from "Over"?

An "Over" indicator (as specified in the Sony 1630 OVER standard) will only light if a minimum or maximum digital code is reached for three or more consecutive samples. One or two consecutive full-scale digital codes are not considered an "Over". Some "Over" meters deviate from the Sony standard and allow the selection 4, 5, or 6 contiguous full scale codes before indicating an over.

The Fallacy of "Over" Indicators

"Over" indicators were developed based upon two assumptions:

1) A digital clip cannot be heard if it has a duration of three or less consecutive samples.

2) Because a signal may reach the maximum or minimum digital code for one or two samples without exceeding it, and because it is important to preserve all available quantization levels (digital codes), it is necessary to assume that such an event does not necessarily indicate a clip.

While assumption 1 may be true for an isolated non-repetitive event, it does not hold true in the real world where we are recording music. We are not are not usually recording random noise but instead are recording complex combinations of musical tones. Individually, each musical tone is a highly repetitive waveform. When combined, these musical tones form a waveform that is at least somewhat repetitive and often very repetitive. In such a case, it is possible, and quite probable, that a clip which has a duration of one or two samples will be repeated a number of times. Such a clip may never produce three consecutive full-scale codes, but is often very audible.

The fallacy of assumption 2 is that there is virtually nothing to be gained by preserving 2 quantization levels, but much to be lost by reducing our ability to detect clipping. The "advantage" of "Over" detection is that the minimum and maximum digital codes can be used to carry the audio signal without generating a clip indication. Without "Over" detection, these two digital codes essentially become illegal because they will generate a clip indication whenever they are used. But how much of an advantage is there in preserving these two codes? In a 20-bit system, we have 1,048,576 unique digital codes (or quantization levels) to work with. If we reduce this by two, we still have 1,048,574 codes to work with! In other words, we reduce our headroom by only 0.000016 dB! Even in a 16-bit system, the headroom reduction is still only 0.00027 dB! Furthermore, if we consider the proability of reaching but not exceeding the minimum or maximum code, we discover that it is very unlikely that a clip has not occurred. Remember that the maximum code represents an infinite number of quantization levels above clip, and only one level below clip. Therefor, it makes far more sense to assume a clip has occurred whenever a minimum or maximum code occurs.

In summary, on "Over" meter can allow an additional 0.000016 dB of headroom before clip indication, but we lose our ability to detect every audible clip. In fact, severe and highly audible clipping can occur without ever generating three consecutive full-scale codes.

Still not a believer? Try this simple test: Feed a 10 kHz tone into any A to D converter at a level equivalent to +3 dBFS. As shown in the graph below, peak voltages will reach almost 1.5 times clip level, clipping will cause severe harmonics at 20 kHz, distortion will exceed 14%, it will sound horrendous, but an "Over" (3 consecutive full scale codes) never occurs. **This test illustrates that "Over" indicators may ignore audible clipping!**

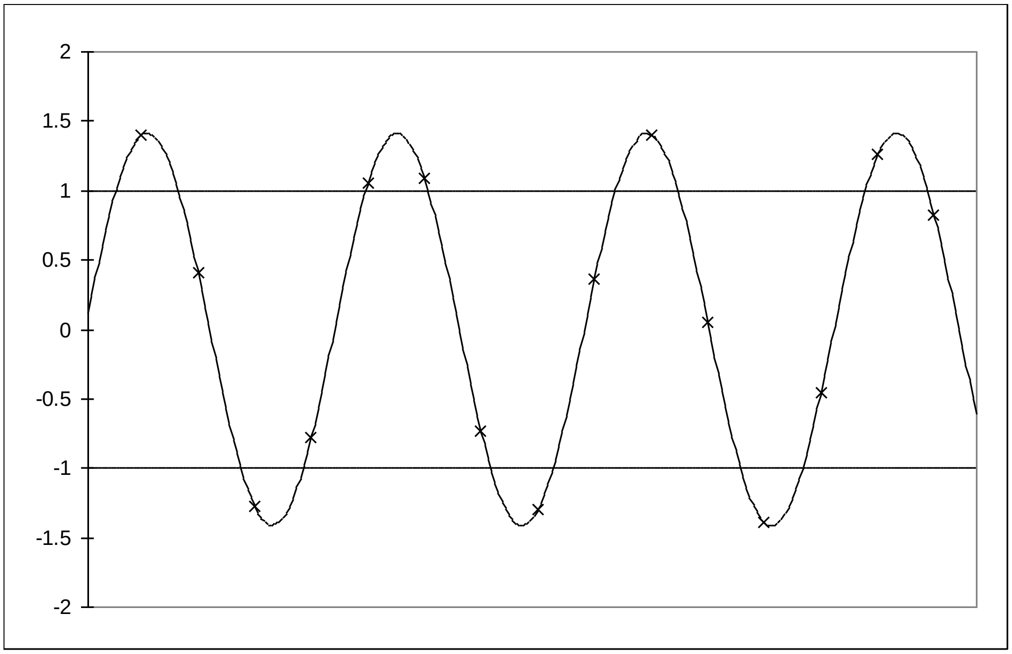

Figure 1 10 kHz Tone at +3 dBFS (1 and -1 represent clip levels)

Do "Over" Indicators Have a Place?

Perhaps

In spite of our best efforts to produce clean audio, there is always a demand for a CD that is louder. The truth is, 16-bit masters that are created using "Over" meters may end up 3 dB hotter than masters that are created using FSD metering. But didn't we just say that using an "Over" meter increases the headroom by only 0.00027 dB? Yes but, what really happens, is that the new recording often has 3 dB clipped off of the highest peaks. And yes this clipping can probably be heard. And yes the CD will sound louder. But sometimes that is all that seems to matter.

A suggestion: Keep the 20-bit originals clean by using a FSD meter. Create a final 20 or 24-bit mix entirely with FSD meters. Then and only then, transfer to 16-bits (using an appropriate dither process), and adjust the levels using an "Over" meter. This way, if the clipping proves objectionable, you can still go back to the clean mix and repeat the transfer. Better yet, avoid the use of an "Over" meter entirely whenever you are not being pressed to achieve maximum loudness.

Chapter 6

Jitter

What is it, and what does it do to my audio?

Jitter is often misunderstood, and can be difficult to measure. In certain circumstances, jitter can be benign, but in others, jitter may cause sonic artifacts. These artifacts may be far more serious than most people realize. It is extremely important to understand where and when jitter is a problem, how to measure it, and what can be done to prevent it.

What is Jitter?

Jitter is time-base error. More specifically it is a measure of how early or late a digital transition occurs. These digital transitions may be the rising and falling edges of system clocks or of digital audio signals. Jitter in conversion devices such as; Analog to Digital Converters (ADCs), Digital to Analog Converters (DACs), or Asynchronous Sample Rate Converters (ASRCs), will cause phase modulation of the audio signal. Jitter between two fully digital devices will not usually cause phase modulation but may cause bit errors if the jitter levels are unusually large. However, jitter at these digital-to-digital interfaces can often pass through a system to a conversion device were jitter can cause phase modulation.

What Causes Jitter?

Jitter can be caused by poor circuit design, bandwidth and noise limitations of digital transmission systems, electromechanical variations in record and playback devices, and even the physical spacing of the optical pits on a CD. The bandwidth limitations of the AES/EBU transmission system guarantee that jitter on the interface will exceed levels that are acceptable for 20-bit data conversion. This does not mean that the AES/EBU interface should be abandoned, it simply means that a jitter free clock must be recovered from the interface before the clock is sent a conversion device. Low-jitter clock recovery can be achieved using proper Phase Locked Loop (PLL) design techniques.

What is a PLL?

A PLL or Phase-Locked-Loop is the electrical equivalent of a mechanical flywheel.

Even if we were to use high bandwidth distribution of a 256X (11.2896 MHz) digital clock, it would be difficult to achieve jitter levels low enough for 20-bit data conversion without using a well designed PLL. However, very few PLL designs have achieved RMS jitter levels in the sub 100 psec range. To make matters worse, many ADC and DAC devices cannot achieve the necessary performance when operating in a master clock mode.

Is There a Cure?

Yes. The AD2004 has a very unique multi-stage PLL which sets new benchmarks for low jitter clock recovery. When locking to an AES/EBU reference having jitter as high as 5 nsec, the internal PLL will produce a phase locked clock having jitter of 12 to 16 psec RMS. With a 10 kHz , -1 dBFS test tone, the sum total of all jitter induced sideband energy in the digital audio output of the AD2004 is -123 dBFS or lower when locked to a moderately jittery AES/EBU reference. This is near the theoretical limit of non-dithered 20 bit audio. The AD2004 is a rare example of such a device. It has the ability to reduce jitter by 50 dB (a ratio of 316 : 1). It can lock to any signal having jitter as high as 5 nsec RMS (5 billionths of a second) and still produce a clock which has jitter below 16 psec (16 trillionths of a second) at the converter. However, as we shall see, it is only necessary to

Are there Different Types of Jitter?

Do "Jitter Killers" Really Work?

Can Compact Disks Contain Jitter?

When is Jitter Low Enough?

At first glance, it seems obvious that low jitter amplitude is important. However, as we look closer, we discover that the frequency (or spectral distribution) of the jitter is at least as important as amplitude. Next, as we look throughout the entire digital recording and playback system, we discover that certain parts of the digital chain are extremely sensitive to jitter while others appear to be nearly immune to jitter. But, upon further investigation, we start to discover that devices can interact with each other and cause "strange phenomena". These "phenomena" are really nothing more than design defects which cause sonic artifacts in the presence of jitter. More on this latter. First lets step out of the world of ones and zeros consider something which we can put our hands on:

The Tape Machine - A Jitter Analogy:

Jitter is the digital equivalent of wow and flutter. An analog to digital converter (ADC) is the digital equivalent of an analog tape machine in record mode. A digital to analog converter (DAC) is the digital equivalent of an analog tape machine in playback mode.

Analog and digital systems are both subject to time base errors: In analog recording, temporary tape speed variations will cause sections of a tape to pass a head a little too early or a little to late. In a digital ADC or DAC, jitter can cause a sample clock pulse to occur a little too early or a little too late. In either case, the results are the same: Audio signals will be temporarily shifted to slightly higher or lower frequencies until servos can make corrections. In analog recording, a large flywheel is attached to the capstan to reduce rapid fluctuations (flutter), and the capstan motor speed is controlled by a carefully designed servo which reduces slower variations (wow). Typically, the servo will compare the speed of the capstan to a crystal reference in order to accurately control the average speed of the capstan. If a second tape machine is slaved to the first machine, the servo in the second machine may use a time code pulse as a speed reference. Similarly, in high quality digital systems, a clock signal is either supplied directly from a crystal oscillator (master mode), or by phase locking a voltage controlled crystal oscillator (VCXO PLL) to an external clock reference (slave mode). In both analog and digital systems, time-base errors which occur during recording will create permanent frequency fluctuations in the audio signal. Time-base errors which occur during playback will have the same sonic effect as time-base errors which occur during record. The sonic artifacts caused by time-base errors in record and playback modes are additive. However, there is one very important difference. Playback time-base errors can be eliminated by fixing the playback device. Record time-base errors can only be fixed by repeating the recording session! This difference underscores the necessity of having low jitter in the ADC.

Now, lets extend the tape machine analogy to address interface jitter. he AD2004R has jumpers for selecting preset input levels, and for enabling the front panel level controls. Use the preset levels whenever possible. The preset levels are matched to an accuracy better than +/-

Chapter

7

Signal Interconnect

Detailed interface information..

The AD2004R has a 10-pin pluggable barrier strip for the analog audio inputs, a 6-pin Molex connector for the power supply input, and four BNC connectors for the digital audio interfaces.

Analog Audio Inputs

Analog Audio Inputs must be balanced. Factory preset input reference level is +4 dBu. Headroom is 20 dB. In other words, 0 dBFS (the full scale digital clip point of the converters) will be reached when the input levels reach +24 dBu. Internal jumpers can be changed to select other 0 dBFS clip points ranging from +16.5 dBu to +28 dBu (see Table 1 on page 18). Barrier strip connections are marked on the rear panel. Pre-wired XLR cable sets are available for the AD2004R.

Connector pin assignments:

1) **Channel 1 Negative**

2) **Channel 1 Positive**

3) **Chassis Ground**

4) **Channel 2 Negative**

5) **Channel 2 Positive**

6) **Channel 3 Negative**

7) **Channel 3 Positive**

8) **Chassis Ground**

9) **Channel 4 Negative**

10) Channel 4 Positive

Digital Audio Interfaces comply with SMPTE 276M. They are 75 ohm, 1 Vpp, professional format coaxial interfaces on BNC connectors which provide outstanding performance when long transmission distances are required. If your other equipment uses a different interface, see page ????? for special wiring instructions.

The Digital Audio Reference Input accepts either SMPTE 276M or consumer SPDIF formats. Two BNC connectors are provided ("Input" and "Loop"). The two BNC connectors are wired in parallel and are not terminated internally. **If you are using a digital audio reference to phase lock a single AD2004 in a slave mode, connect this reference to the "Input" connector, and connect a 75 ohm BNC terminator (supplied with every AD2004R) to the "Loop" connector. If you are not using a digital audio reference, store the BNC terminator on the "Loop" connector.** If you are installing multiple AD2004R converters, you can daisy chain a single digital audio reference through multiple AD2004R devices. Be sure to connect a 75 ohm terminator to the "Loop" connector of the last AD2004R on the daisy chain. For more information see page ????.

Chapter

8

Jumpers

Setting Jumpers on the AD2404-96 Converter Board.

The AD2404R converter board has jumpers for user options, for test functions, and for adaptation to various mother boards. The options list is subject to change and expansion as new software versions are released. Please check to see which software revision you are using before changing jumpers.

Option Jumpers for Revision 1.XX

DC Filter: A digital high-pass filter is available to remove DC offsets from the digital audio signals. The filter is a first-order high-pass filter and has the following characteristics at a 48 kHz sample rate:

At Fs = 48 kHz:	
Frequency Response:	-3 dB at 1.8 Hz
	-0.036 dB at 20 Hz
Phase deviation:	5.3 degrees at 20 Hz
Passband Ripple:	None

The filter response is a function of sample frequency, and scales linearly with sample rate. **The DC filter is enabled when a jumper is installed between pins 13 and 14 of header "P5".** The state of this jumper is read at boot-up and whenever the sample rate selection switch is rotated. The filter should not be required for most applications, and is not installed at the factory.

Digital I/O Jumpers

These jumpers adapt the four-channel converter boards to a variety of different motherboards and system configurations. Removal or incorrect placement of these jumpers will disable one or more of the digital outputs. It should not be necessary to change these jumpers. The state of these jumpers has an immediate effect on the digital outputs.

	JP1	**JP2**	**JP3**	**JP4**	**JP5**	**JP6**
AD2404-96 X	1 TO 2	1 TO 2	1 TO 2	1 TO 2	1 TO 2	1 TO 2
AD2404-96 B	2 TO 3	1 TO 2	2 TO 3	1 TO 2	2 TO 3	2 TO 3
AD2408-96 X	1 TO 2	1 TO 2	1 TO 2	1 TO 2	1 TO 2	1 TO 2
AD2408-96 B	2 TO 3	1 TO 2	2 TO 3	1 TO 2	2 TO 3	2 TO 3

Miscellaneous Jumpers

Enable/Test Jumper: This jumper must always be installed between pins 1 and 2 of P4. Removal of this jumper will have no immediate effect while the system is operating. However, the system will not reboot if this jumper is missing.

Chapter 9

Dither, Word Length Reduction, and Noise Shaping – a Tutorial

What is noise shaping, … when, why, and how much should I use?

Dither, and Word Length Reduction (WLR) are necessary processes in most digital audio systems. Unfortunately, these processes add noise to our digital audio signals. Noise Shaping is one method of reducing the audibility of this added noise. However, noise shaping can add additional noise power without benefit if it is used improperly. The decision to use noise shaping, and the selection of a particular noise shaping curve do not have to be accomplished via trial and error. This tutorial should provide the recording professional with a basic understanding of these processes, and provide some practical techniques for successful application of WLR using dither and noise shaping.

Why is Word Length Reduction Necessary?

The number of bits in a digital word (sample) increases whenever we process a digital signal. If we add two 16-bit digital numbers, the result requires 17 bits. If we multiply two 16-bit numbers, the result can require up to 32 bits. Very simple audio processing functions can create very long digital word lengths. Consequently, we have to decide what to do with all of these extra bits.

Quantization Noise

If we truncate or round off the least significant bits of any digital signal, we will add quantization error (or noise) to that signal. Quantization noise is an unavoidable fact of life for all digital systems. In a linear encoding system, it is impossible to reduce digital word lengths without creating additional quantization noise. Quantization noise is not something that we intentionally add to a signal; quantization noise is caused by round-off or truncation errors.

How does WLR create Quantization Noise?

Quantization errors are added by a truncation or rounding process because of the difference between the numeric value of the input and output samples. If the input signal is rounded off, then each output sample will have a quantization error equal to +/- ½ LSB (1 LSB peak to peak). The significance of this 1 LSB error signal is a function of word length. At 24-bits, one LSB represents a very small portion (1/16,777,216 or –144.5 dBFS) of the full-scale range of our digital system. However, at 16-bits one LSB represents a much larger portion (1/65,536 or –96.3 dBFS) of our full-scale range. **The output word length determines how much quantization noise any single WLR process will generate.**

Can we Reduce Quantization Noise?

The only way to reduce quantization noise in a linear encoding system is to use longer word lengths. At 24-bits, the quantization noise is so low (-144 dBFS) that it doesn't even come close to limiting the performance of the finest digital audio equipment available. However, at 16 bits, the quantization noise is –96 dBFS, and it can very easily limit the system performance. **In most cases, quantization noise becomes a permanent part of a digital signal.** Converting a short 16-bit word length to 24-bits will not reduce the quantization noise already encoded into the signal.

Quantization noise increases every time we add an additional WLR step. At 16-bits, two cascaded WLR steps will elevate the quantization noise from –96 dBFS to –93 dBFS. Every time we double the number of WLR processes (at a given word length), we increase the quantization noise by 3 dB. Unfortunately, the audio path through a digital mixer usually requires many cascaded WLR processes. Every gain change, E.Q., or mix process will require at least one WLR process. If our digital mixer uses 16-bit processing, we can expect substantial levels of quantization noise. The quantization noise produced by one 16-bit WLR process is equal to that produced by 256 cascaded 20-bit WLR processes, or 65,536 cascaded 24-bit WLR processes. **Avoid equipment that processes audio using 16-bit internal word lengths, the results are likely to be 14-bit quality or worse.**

How does Quantization Noise Sound?

Quantization noise can take on various forms depending upon the WLR technique. Simple truncation or rounding will create quantization noise that is almost entirely comprised of signal related distortion products. This type of distortion is non-musical and is easy to hear because it is not harmonically related to the input signal. More sophisticated WLR techniques can create white or shaped quantization noise that is almost totally unrelated to the audio signal. In addition, noise-shaper can move quantization noise into frequency bands that are harder to hear. **When properly used, the best noise shapers can achieve near 20-bit psycho-acoustic performance at 16-bits.**

Rounding vs. Truncation

At first it may appear that rounding would produce smaller quantization errors than truncation. If we round, the errors in the output words will be uniformly distributed between +½ and -½ LSB (of the shortened word length). If we truncate (ignore the extra bits in the input words) then the errors in the output words will be uniformly distributed between +1 and -0 LSB. Note that in both cases, the error signal has a uniform distribution and peak to peak amplitude of 1 LSB. The only difference is a ½ LSB DC offset. This very small DC offset has no significance in most audio systems, but could be removed later if desired. **It is important to understand that in an audio WLR system, truncation and rounding produce identical results. Rounding does not reduce quantization noise, and it does not alter the audibility of that noise. Rounding has a digital processing cost of one addition operation per sample while truncation is free.** All of the WLR systems described in this tutorial (including "truncation") may or may not include a rounding operation.

What is Dither?

Dither is a noise signal that is typically added to an audio signal prior to quantization. In a WLR process, dither is added to prevent undesirable forms of quantization noise. The dither signal is usually a white noise signal, and is usually generated digitally.

The quantization noise in a WLR process is the direct and unavoidable result of shortening the digital word length. The output word length of the WLR process determines the average power of the quantization noise. **Dither does not alter the quantization noise power, but it does alter the character of the quantization noise.**

Undithered WLR processes can subject low-level signals to 100% distortion. Dither randomizes quantization errors and breaks the correlation between the audio and the errors. **Dither eliminates distortion caused by the WLR process.**

Undithered WLR processes can cause noise modulation. Without dither, the audio signal can modulate the amplitude of the quantization noise. While the average noise power remains unchanged, the instantaneous noise amplitude is a function of the input audio signal. Dither insures that the quantization error at any sample is always determined at random. **Dither prevents noise modulation.**

Dither will preserve low-level signals that would be destroyed in an undithered WLR process. If dither is not applied before truncation, low-level signals will be distorted, and may disappear and reappear in response to other audio content. However, a properly dithered 16-bit WLR process can preserve very small signals that were captured at high resolution (24-bits). Incredibly, a Flat Dithered 16-bit TPDF WLR system can preserve tones as low as –120 dBFS. These –120 dB tones can be heard, distortion free, below the –91.6 dBFS white noise floor of the 16-bit WLR system. Signals are often not masked by the system noise until they are more than 30 dB below the noise floor! Noise shaping can further extend this small signal resolution by reducing the masking effects of the dither and quantization noise.

What is TPDF Dither?

Dither signals can be easily and accurately generated with digital hardware. For example, if 4 bits are to be removed from a signal, a 4 bit random number generator can be used to generate a uniform distribution of all of the numbers that can be represented by 4 bits (i.e. 0 through 15). This type of dither is known as **"Uniform Probability Distribution Function" (UPDF)** dither because all of the possible numbers (0 through 15 in this example) are generated with equal likelihood.

A pair of dice can be used to illustrate dither generation. A single die will return numbers between 1 and 6 with equal likelihood, and therefore it is an example of a UPDF number generator. If we sum two UPDF dither generators we produce dither at twice the amplitude, but with a distribution of numbers that is no longer uniform. Two dice thrown together will return numbers between 2 and 12, but the chance of rolling a 7 is much higher than rolling a 2 or a 12. The outcomes are no longer uniformly distributed. If we plot the probability of all outcomes from 2 to 12, we will see that they have a **"Triangular Probability Distribution Function" (TPDF)**.

When it comes to decorrelating quantization errors from the audio, it can be shown that TPDF dither is far more effective than UPDF dither. Both forms of dither add noise to the audio signal (above and beyond the added quantization noise). UPDF is added at a level equal to the noise of quantization, and therefor it increases the noise by 3 dB. TPDF is added at a level that is 6 dB higher than the quantization noise, and therefor it increases the noise by 4.7 dB. TPDF is usually chosen because the improved decorrelation justifies the additional 1.7 dB noise penalty. **Most Dithered WLR systems combine two separate and independent 1 LSB UPDF dither signals to create a 2 LSB TPDF dither signal.**

What is a "Noise Floor"?

The noise floor of a system is a measurement of the average output noise when no signal is present at the input. **The noise floor is not a threshold below which all sounds disappear.** In an analog system, or a properly dithered digital system, **tones may still be audible when they are 30 dB below a white 20 Hz to 20 kHz noise floor**. However, many of the most common audio measurements have no ability to measure what is happening below the noise floor. Fortunately, bandpass filters and FFT techniques can allow measurements that exceed the capability of our auditory system. These test techniques can identify and demonstrate the extraordinary demands that our auditory system places upon audio equipment. For example, it can be shown that the low-level performance of A/D or D/A converters is quickly destroyed by the presence of spurious tones and low-level non-linearities. These defects will go undetected in THD+N and Dynamic Range tests.

WLR will raise the noise floor of an audio signal. The level and audibility of this increased noise varies greatly depending upon the WLR technique. Also, the quality of signals that lie below the noise floor can be destroyed by improper WLR. Again, these low-level defects will go undetected in THD+N or Dynamic Range tests, with an FFT analysis. By the way, the human auditory system has a wonderful FFT analyzer built in, so beware!

Word Length Reduction Techniques

Truncation without Dither

Truncation WLR (without dither) is simplest WLR method. To truncate we simply ignore the least significant bits. AES/EBU and SPDIF interfaces are designed to support truncation without additional hardware. For example, a 24-bit AES/EBU signal connected to a 16-bit device will result in truncation of the 8 least significant bits. **Truncation comes for free, and consequently it often happens by accident! It is essential to know what word lengths are supported by each piece of equipment in you digital audio chain.**

Truncation is the least acceptable WLR system because all of the quantization errors take on the form of signal related distortion products. In fact, truncation is a wonderful way to create some really nasty and non-musical distortion.

If we truncate to 16 bits, we will create distortion at -96 dBFS (0.0016% for a full-scale signal). This may seem acceptably low, and consequently many have failed to recognize the problems caused by 16-bit truncation. To understand these problems, we must look at the numbers in greater detail: The –96 dBFS distortion signal has a fixed -96 dB relationship to full scale. This amplitude <u>does not</u> decrease when the input signal amplitude decreases. Consequently, a -60 dBFS signal will be subjected to distortion at -96 dBFS, a difference of only 36 dB (15.8 % distortion)! Remember that this is nasty sounding quantization distortion, it is not random noise. **This distortion will reach 100% when the input level is -96 dBFS.** Furthermore, signals below -96 dBFS may disappear and reappear due to severe intermodulation caused by quantization.

Curiously, the distortion caused by truncation will decrease if we increase the noise levels on the input signal. In fact, most of the distortion effects will disappear altogether if the input audio signal has a noise floor exceeding 1 LSB of the output word length (-96 dBFS for a 16-bits). The distortion is not masked by the added noise, instead, the added noise randomizes the quantization errors so that they are transformed from distortion to additional random noise.

In the recent past, 16-bit truncation often produced acceptable results. Early 20-bit converters had noise floors that were not much below –96 dBFS. When the converter noise was summed with the contributions of room noise, microphone pre-amp noise, it is unlikely that many recordings achieved noise floors below –96 dBFS. Consequently these signals were "self dithered" and truncation created little or no distortion. This has led some to discount the value of dither. However, a high quality 24-bit converter, a quiet room, high sound pressure levels, and properly set low-noise mic-preamps can produce a signal with a noise floor as low as –117 dBFS. **Truncate a low-noise signal to 16-bits, and you may not want to listen to it!** However, a 20-bit truncation of this signal would produce excellent results. The -117 dBFS noise floor provides a more than adequate dither signal for 20-bit truncation. This 20-bit truncation can only be improved upon by using noise-shaping techniques.

Given current state of the art recording equipment, truncation to 16-bits usually will add distortion, while truncation to 20-bits usually will not. As equipment improves, and noise floors drop below –120 dBFS, it will become necessary to add dither prior to 20-bit truncation.

Truncation will add distortion to signals that are inadequately self-dithered. Truncation will add white noise (but not distortion) to signals that are adequately self-dithered. Successful truncation can only be guaranteed if the input noise levels are well know. A dithered WLR process should be used when the input noise floor is not accurately known. The low noise floors of high-quality 24-bit converters can only be psycho-acoustically preserved if the dithered process includes noise shaping.

Flat Dithered WLR

Flat Dithered WLR systems require a random noise (or dither) signal of at least one least significant bit of the shortened word length (i.e. -96 dBFS when reducing to 16-bits). This dither is normally added to the audio prior to truncation. The added dither randomizes and decorrelates the quantization noise so that it no longer has any relationship to the audio signal. A dither signal of adequate amplitude will change the quantization errors from distortion to noise. More specifically, the dither randomizes the quantization error at each sample so that the amplitude of the error is no longer a function of the audio signal.

Most Flat Dithered WLR systems use digitally generated TPDF dither. **TPDF dither insures that all of the added noise is random, white, and free from both distortion and audio modulation.** The total noise contribution from a Flat Dithered TPDF WLR is 4.7 dB higher than the quantization noise alone. Therefor, the noise contribution of 16 and 20-bit TPDF dithered WLR are –91.6 dBFS and –115.7 dBFS respectively. Flat Dithered WLR systems using TPDF dither work well, but the noise added by these systems will usually sound louder than that added by Dithered Noise Shaped WLR systems. **Use a TPDF dithered WLR function if: the audio signal already has a high noise floor, the noise floor is well masked by continuously high program levels, the signal is to be compressed, or the signal is to be subjected to special processing such as frequency shifting.** Do not use TPDF dither for applications that demand the lowest possible perceived noise levels. Very low noise floors can only be preserved using a Dithered Noise Shaped WLR system.

Dithered Noise-Shaped WLR

Dithered Noise-Shaped WLR systems use a digital filter to shape the dither and quantization noise into frequencies that are harder to hear (above 20 kHz), and away from frequencies that are easier to hear (1 to 5 kHz). This shaping makes the noise harder to hear even though it actually increases the unweighted noise power.

Most noise shaping systems use digitally generated TPDF dither. The dither is added to the input signal, and then the result is either rounded or truncated to a shorter word length. The discarded bits represent the error signal. This error signal is fed back to the input of the system through a digital filter. The result is that errors can be reduced at frequencies that are easy to hear, and shifted into frequencies that are hard to hear.

Many filter designs are possible, but most are based upon one or more low level hearing threshold study. These studies have been used to create weighting functions. These are often these are referred to as “F-weighting” functions, although there is currently no standard “F-weighting” function.

Some opponents of noise shaping have argued that there are a number of factors in the listening environment that can alter our low-level hearing threshold. It is important to keep in mind that these variations are relatively small, and that a single noise shaper design can provide substantial benefits over a wide variety of listening conditions.

A dithered and noise shaped 16-bit signal can psycho-acoustically approach the noise performance of a flat dithered 20-bit system. Noise shapers will always boost noise in high frequencies where it is more difficult to hear. If the noise shaping is too aggressive, high-frequency noise may cause problems in the rest of the audio system. Keep in mind that all of the high frequency energy created by the noise shaper must pass through the audio reproduction system and out the speakers. For this reason, most noise shaping systems limit high frequency noise boost to reasonable levels.

Multiple passes through an aggressive 16-bit noise shaper may produce a signal with excessive high-frequency noise. Also, the audible noise will increase by 3 dB every time the number of passes through a WRL device doubles. For this reason, it is important to limit the number of passes through a 16-bit Dithered Noise-Shaped WLR system. On the other hand, multiple passes though a properly designed 20-bit noise shaper, are not generally be a problem, because a 20-bit WLR process is 24 dB quieter than an identical 16-bit WLR process.

Example 1 – 24-bit ADC, 20-bit recorder, 16-bit distribution:

It is perfectly acceptable to use noise shaping to WLR a 24-bit converter for recording at 20-bits. This 20-bit recording could then be mixed at 24-bits and noise shaped back to a 20-bit master. The 20-bit master could then be noise shaped to 16-bits. The first two 20-bit noise shaping processes will have little or no effect on the final 16-bit product.

First, it is important to understand some basic properties of digital systems. A digital stopwatch can be used to illustrate these properties. First, that a single sample from a digital system has finite resolution. A digital clock can be used to demonstrate this concept. If my clock only has hour digits, it has very poor resolution, and I may only be able to tell you that it is after one o'clock and before two o'clock.

, suppose that I am waiting for an elevator, and there are lighted numerals telling me what floor the elevator is on. If the numeral 2 is lit, I know that the elevator is above floor 1 and below floor 3. This resolution is determined by the number of bits acquired at each sample. Imagine a postal scale tWLR can be achieved with noise shaping can AD2004R has jumpers for selecting preset input levels, and for enabling the front panel level controls. Use the preset levels whenever possible. The preset levels are matched to an accuracy better than +/- +/-

Chapter

10

"IR" Word Length Reduction System

A Detailed description of the AD2408 Word Length Reduction Features.

The AD2408 is equipped with a versatile Word Length Reduction (WLR) system. It has been carefully designed to preserve as much performance as possible when it becomes necessary to reduce the converter's word length from 24-bits to 20 or even 16-bits. Any word length reduction will add noise to the audio signal. However, we can choose how to add this noise.

At 44.1 kHz:

Output Word Length	Word Length Reduction Mode	Psycho-Acoustic Performance Relative to a TPDF Dithered Transmission System	F-weighted Dynamic Range	Unweighted Dynamic Range
16 bits	TPDF	16 bits	91.6 dB	91.6 dB
16 bits	IR1	17.4 bits	99.8 dB	87.4 dB
16 bits	IR3	18.1 bits	104 dB	79.4 dB
16 bits	IR5	18.4 bits	105.8 dB	72.9 dB
20 bits	TPDF	20 * , 19.8 ** bits	115.7 dB	115.7 *, 113 ** dB
20 bits	IR1	21.4 *, 20.3 ** bits	124 *, 116 ** dB	112.5 *, 112 ** dB
20 bits	IR3	21.7 *, 20.6 ** bits	128 *, 116.4 ** dB	103.5 dB

24 bits	None	24.8 *, 21 ** bits	144.5 *, 117 ** dB	139.8 *, 117 ** dB

* Digital to digital mode ** A/D mode (limited by converter noise floor).

Chapter 11

Modifications

Instructions for modifying the converter boards for special applications.

The AD2404 converter boards are designed to reach a full scale code at +24 dBu. A variable gain motherboard is available for AD2404 and AD2408 converter systems, and should be used when other full-scale clip points are required. However, it is possible to modify the gain structure of the converter boards. This modification must be done by a qualified electronic service technician. Improper procedures may result in damage to the converter boards and other equipment.

Hazardous Voltage Warning

Most converter chassis configurations have hazardous AC line voltages inside. The AC power cord must be removed before opening the chassis. After completing the modifications, a CMRR alignment procedure is necessary. This procedure will require application of AC power while the cover is removed. Exposed voltages exist at the rear of the AC power entry module. Use caution whenever AC is applied. The technician must replace the chassis cover and all fasteners after servicing.

ESD Warning

The converter boards contain sensitive electronic devices that may be damaged by electrostatic discharge. Use a wrist strap and an ESD safe workstation.

Selection of Components:

The modification will require the replacement of two resistors per channel (four per channel on rev B boards). The components required are 0.1% metal film resistors in 1206 surface mount packages and may be difficult to acquire. However, the components are available direct from Benchmark. The 0.1% tolerance is required to guarantee that the CMRR can be trimmed. It is also important to use metal film resistors, as these will produce lower distortion than thick film resistors.

Input Level at 0dBFS	Resistor “A” (Ohms) +/- 0.1%	Resistor “B” (Ohms) +/- 0.1%

24 dBu	1000	1020
22 dBu	1300	1330
20 dBu	1690	1740
18 dBu	2260	2320
16 dBu	3090	3160

Revision “B” boards will also require two 7.68 K 0.1 % resistors per channel to replace the 7.87 K resistors that were originally installed at the factory. Note: On Revision “B” boards, 1 M ohm resistors were often installed on top of the 7.87 K 1% resistors to facilitate common mode trim. These 1 M ohm resistors will not be required after installing the 7.68 K resistors because of the improved (0.1%) tolerance of the 7.68 K parts.

Special Instructions for Revision B:

1. Remove R101, R102, R201, R202, R301, R302, R401, and R402.
2. Replace the above with 7.68 K 0.1% metal film resistors.
3. Proceed to “General Instructions”

General Instructions for Revisions B and C:

1. Remove R104, R204, R304, and R404.
2. Replace with Resistor “A” (see Component Selection Chart).
3. Remove R103, R203, R303, and R403.
4. Replace with Resistor “B” (see Component Selection Chart).
5. Proceed to “Common Mode Trim Adjustment”

Common Mode Trim Adjustment:

The converter board has a Low Frequency Trim Resistor and a High Frequency Trimmer Cap for each channel. The resistor will need adjustment after gain modification. It is possible that the trimmer cap may need minor adjustment. The resistor should be adjusted first.

1. Apply a 200 Hz –12 dBFS Common Mode test tone to the channel 1 analog input.

2. Apply power.

3. Adjust R407 (Low Frequency Common Mode Trim) for minimum output level at the channel 1 digital output.

4. Apply a 10 kHz –12 dBFS Common Mode test tone.

5. Adjust C410 (High Frequency Common Mode Trim) for minimum output level at the channel 1 digital output.

6. Repeat for each of the remaining channels. Adjust R307 and C310 for channel 2. Adjust R207 and C210 for channel 3. Adjust R107 and C110 for channel 4.

7. Replace Chassis Cover and Screws.