DAT330 – Principles of Digital Audio
Cogswell Polytechnical College
Spring 2009

Week 3 – Class Notes

Digital Audio Reproduction

4. Digital Audio Reproduction

Reproduction Processing

The reproduction signal chain accepts a binary coded signal and ultimately reconstructs an analog waveform. Reproduction circuits must minimize any detrimental effects of data storage or transmission. However, in order to maximize the data density, the physical fidelity of the stored or transmitted channel code waveforms is allowed to deteriorate. In this state, the signal is not as “clean” as the original data. Nonetheless, data can be recovered without any deterioration by using a waveform shaper circuit (identifies transitions and reconstructs a valid data signal).

Synchronization pulses and other clocking information in the bitstream are identified, by timebase correction circuits, to synchronize the playback signal, delineate individual frames, and determine the binary content of each pulse. Timing errors, such as jitter, contained in the playback signal are corrected by phase-locked loops (PLL) and data buffers. PLL circuits are used to reclock the channel code (compares input and self-reference to correct phase differences). Data buffers are memory into which data is fed irregularly as it is received. However, data is clocked from the buffer at an accurate controlled rate; ensuring precise data timing.

Modulated audio data (MFM, EFM, etc) is demodulated to NRZ code. As such, the audio data becomes once again a readable binary form that is ready for further processing. Demultiplexing restores the parallel structure of the data from the serial bit input.

Audio data and error-correction data are identified and separated from other peripheral data. With the error correction data, the signal is checked for any errors that have occurred during the encoding process. Afterwards, the data has to be de-interleaved in order to reassemble the signal data. Then, the data is checked for errors by using redundancy techniques, or parity. If for example, the calculated parity does not match the parity read from the bitstream, a error has occurred. Error correction algorithms are used to calculate correct values or, if the error is too large, error concealment or interpolation methods are used to reconstruct the lost data.

After this entire process has been done, the data is ready for D/A conversion.

Digital-to-Analog Converter

The D/A converted determines how accurately will the digitized signal be restored to the analog domain. The DAC requires great precision; for example, with a ±10 V scale, a 16-bit converter will output voltage steps of 0.000305 V, while a 24-bit converter output steps measure 0.00000119 V.

Traditional DAC show many of the same errors found in ADC. The resolution of the converter is determined by the absolute linearity error and the differential linearity error. Errors for high signal levels tend to be uncorrelated, but for lower levels the error is correlated.. As a result, low signal levels will be distorted.

Consider, for example, a signal at -80dBFS will pass through only six or seven codes in a 16-bit quantizer. If half of those codes are missing the overall performance will be that of a 14-bit quantizer. The bits and their associated errors will be signal dependent. Thus, the error is correlated with the audio signal. Resulting in harmonic/intermodulation distortion and noise.

In reality, 16-bit conversion is insufficient for 16-bit data. Meaning that they must have a greater dynamic range than the audio signal itself. Also, the accuracy of the DAC is measured by its linearity and not by its bit number.

Weighted-Resistor DAC

A simple ladder, or weighted-resistor, DAC accepts an input digital word and converts it to an output analog voltage or current. This type of DAC contains a series of resistors and switches. A switch is used for every input bit, while the corresponding resistor represents the binary value associated with that bit. A reference voltage generates current through the resistors. A digital 1 closes a switch and contributes a current, and a digital 0 opens a switch and prevents current flow. An operational amplifier sums the currents and converts them to an output voltage. Thus, low- value binary words (with many 0's close few switches) will result in a low voltage, while high-value words (with many 1's close many switches) and will produce a high voltage.

We can calculate the output voltage depending on the input digital word with the following equation:

$$Vout = -Vref \sum_{1}^{n} \left(\frac{bn}{2^n}\right)$$

where *n = 1,2,3,...*

For an 8-bit word, our equation will be:

$$Vout = -Vref \left(\frac{b1}{2} + \frac{b2}{4} + \frac{b3}{8} + \frac{b4}{16} + \frac{b5}{32} + \frac{b6}{64} + \frac{b7}{128} + \frac{b8}{256}\right)$$

This DAC design is rarely used in practical applications because of the complexity in manufacturing the resistors with sufficient accuracy.

R-2R Ladder DAC

This DAC circuit is similar to the weighted-resistor design, but instead of using resistors and switches, it uses two resistors per bit. At each node of the ladder, the current splits resulting in current flows through the switch resistors that are weighted by binary powers of two. If a current *I* flows from the reference voltage, *I/2* flows through the first resistor switch, *I/4* flows through the second resistor, *I/8* flows through the third resistor, and so on. The R-2R ladder equation is in fact very the same as that of the weighted-resistor. R-2R networks are preferred because they are easily manufactured.

Zero-cross Distortion

This type of distortion occurs at the zero-crossing point between the positive and negative polarity portions of an analog waveform. When a resistor ladder DAC is switched around its MSB to reflect a polarity change, say from 1000 0000 0000 0000 to 0111 1111 1111 1111, an internal network of resistors must be switched. Current fluctuations and variations in bit switching speeds create nonlinearity and glitches. Waveforms, such as sinusoids, continually oscillate between positive and negative polarity. The zero axis is crossed repeatedly as the MSB is turned on and off. This error is noticeable for low-level signals because the fixed-amplitude glitch is proportional to the signal level. Zero-cross distortion can be reduced by calibrating of the MSB. Sign-magnitude configuration switches between positive and negative outputs can minimize this distortion type. Complementary DAC for each waveform polarity can also alleviate this problem.

High-bit DAC

Higher bit resolution DAC have been introduced in PCM reproduction systems in order to provide a greater playback fidelity for 16-bit recording. Because 16-bit converters cannot fully decode a 16-bit signal without any error, a higher bit converter (18- or 20-bit) will reduce any error and will improve reproduction. Thus, in order to realize the full potential of audio fidelity, the digitalization and processing steps must have a higher dynamic range than the final recording. For example, 16-bit words, CD audio quality, are better converted with 18-bit converters since they provide better amplitude resolution by ensuring a fully linear conversion of the 16-bit signal. Because an 18-bit DAC has four times as many output levels as a 16-bit converter, any nonlinearities will be correspondingly smaller, and the increased quantization word length at this stage provides a larger SNR. Quantization artifacts are also reduced. It could be said that higher bit conversion is analogous to oversampling. No new information is created, rather, the existing information is used more efficiently.

Output Sample-and-Hold Circuit

The S/H circuit on the output samples and holds the signal output from the DAC, and is used to remove irregular signals called switching glitches. Output S/H also compensates for a frequency response anomaly called aperture error.

DAC can generate erroneous signals, or glitches, which are superimposed on the analog output voltage. Because data input to the DAC needs time to stabilize to the correct binary levels, some input bits do not switch states simultaneously, a momentary value larger than desired will create a voltage spike. If these glitches pass through the output, they will be perceived as distortion.

Thus, an output S/H circuit can deglitch the DAC output signal by acquiring the output voltage only when the circuit has a stable output condition. That correct voltage is held by the S/H circuit during the interval when the DAC switches between samples. The output signal is now a glitch-free PAM signal.

When correcting the aperture error (high frequency attenuation) the S/H circuit can use different approaches. Aperture error is created by the width of the output samples. The narrower the pulse width, the less the aperture error. Ideally, the DAC output should be an impulse train that correspond to the original sample points. However, in practical applications the PAM waveform is made up of pulses with a width of one sample period. The spectrum of theses finite-width pulses attenuates high frequency content. This output spectrum differs from the original input spectrum.

The frequency response follows an ideal lowpass function (sinc function). If the output pulse width is equal to the sample period, the frequency response is equal to zero at multiples of the sampling frequency.

If the output pulse width (controlled by the S/H circuit) is narrower with a shorter hold time, the attenuation at the half-sampling frequency can be decreased. A pulse width of one-quarter the sampling period is considered to be optimal because a shorter hold time degrades the SNR of the system.

Another solution to the aperture error, is to provide a high frequency boost just before the signal reaches the S/H circuit. This high frequency boost is offset by the aperture error and thus, produces a flat frequency response. Alternatively, a pre-emphasis high boost can be applied at the systems input and afterwards a de-emphasis at the S/H output will produce a flat response.

It must be noted that the output S/H circuit can also introduce transition error on its own. Slow transitions from sample to sample can introduce incorrect intermediate values into the staircase voltage. Transition errors can be perceived as distortion in the signal. Although these errors can be removed by the output filter, these products can internally beat with the sampling frequency to produce in-band distortion. Ideally, to avoid distortion products this circuit must change as quickly as possible from hold to sample mode. In practice, is impossible to achieve the necessary high slew rate, or rate of voltage change per unit of time. A solution to this problem is to modify the S/H circuit.

Output Lowpass Filter

The last filter in an audio digitation system is a lowpass filter known as the anti-imaging filter. This filters has two function: remove all frequencies above the Nyquist frequency, and reconstruct the DAC output PAM staircase into a smoothly continuous waveform. The PAM staircase is an analog waveform, but contains modulation artifacts (high frequency components) created by the sampling process not present in the original signal. The output LPF removes the high frequency content in the staircase and smooths it out into a continuous waveform. This filter is also known as a smoothing filter.

Impulse Response

If the input of a system is an input, the output of the system is called the impulse response. A system can be completely characterized by its impulse response. For example, a filter can be described by its impulse response in the time-domain or by its frequency response. The Fourier transform relates time and frequency domains. Note that multiplying in the frequency domain is equivalent to convolving in the time domain.

An ideal LPF has a brick-wall response in the frequency domain, while its time domain representation is given by *sin(x)/x.* This function is called the sinc or cardinal sine function.

This ideal LPF output filter removes the high frequency content in order to smooth out the output samples, however, detailed analysis shoes that the waveform is reconstructed by this filter. The sampling theorem dictates that an ideal brick-wall filter is needed to exactly reconstruct the output waveform from its samples. When the samples from a bandlimited input signal are processed (convolved) with a sinc function, the bandlimited input signal represented by the samples will be exactly reproduced.

When a sample goes through this ideal LPF, the resulting response is a sinc response. If the filter's cutoff frequency is limited to half the sampling rate, then the sinc curves pass through zeros at multiples of *1/fs.* So, if multiple audio samples pass through this filter, the resulting waveform is the delayed summation of all of the individual sinc responses from each sample.

Thus, the summed response of the impulse response of all samples has a value at each sample while the rest of the response, given at *1/fs,* will be zero. Resulting in a continuous reconstructed waveform. This output waveform is theoretically identical to the bandlimited input waveform.
The reconstruction is effective regardless of where the sample was taken with respect of the waveform's timing.

Digital Filters
Digital filters have replaced analog brick-wall filters because they do not introduce phase shifting and distortion. A digital filter can be a digital circuit or algorithm that accepts audio samples and outputs audio samples. Filtering in this case is done by altering the samples (simulating an ideal LPF and reconstruction). Digital filters remove the high frequency images prior to DAC. Afterwards, an analog filter removes the remaining high frequency contents of the signal. In most cases, a finite impulse response (FIR) oversampling filter is used because of its less complex design. Filter architectures typically employ a series of delays, multipliers, and adders.
Oversampling filters have two tasks: to re-sample and interpolate. Let's observe this process:

1. The input signal is sampled at f_s and has images centered around the sampling frequency.
2. Resampling begins by increasing the sampling frequency. This is done by inserting zero-valued samples between the original samples at some interpolation ratio.
3. The oversampling sampling frequency now equals the interpolation ratio times the input sampling frequency. The resulting spectrum is the same as the original input spectrum.
4. This interpolated data is filtered by an ideal LPF with a cutoff frequency of $f_s/2$. So far, the original data sampled at f_s is indistinguishable for the oversampled data.
5. The fixed sample values occurring at the oversampling frequency comprise the LPF coefficients.
6. The oversampled input values are convolved with the impulse response coefficients. The resulting filter's output is an interpolated digital signal with images centered around multiples of the oversampling frequency f_a.

The distance of the baseband and sidebands is lager with oversampling. A gentle low-order LPF can be used to remove these images without producing phase shifting or other artifacts.
The oversampling ratio is given by:

$$R=\frac{fa}{fs}$$

f_a is the oversampling frequency and f_s is the input sampling frequency. Oversampling requires an insertion of *(R-1)* zeros per input sample. The new spectral images are found at multiples of $(R \times f_s)$. Lowpass filtering creates intermediate sample values through interpolation.

FIR Oversampling Filter

Finite impulse response filters use a finite number of points in their impulse response. For example, a four-times oversampling FIR digital filter will generate three intermediate steps between each input sample. The filter consists of a shift register of 24 delay elements, each delaying a 16-bit sample for one input sampling period. Thus, each sample remains in each delay element for a sample period before it shifts to the next delay element. During this time each 16-bit sample is tapped off the delay line and multiplied four times by a 12-bit coefficient stored in ROM, a different coefficient for each multiplication. Each tap has a different set of coefficients. The four sets of coefficients are applied to the samples in turn, thus producing four output values. The 24 multiplication products are then summed four times during each period, and are output from the filter. The *sinc(x)* filter characteristic determines the value of the interpolated samples. Each 16-bit word is passed to the next delay, where the process is repeated. The computed impulse response is a 28-bit word (16+12 = 28).

The sample products are summed and then a weighted average is obtained. Four times as many samples are present after oversampling, and the interpolation values are calculated by the filter. This type of filter design is called a transversal filter.

The interpolated waveform reconstruction by the four times oversampling filter is as follows: the input samples are treated as sinc(x) impulses placed at the center of the filter. Their maximum filter response amplitudes are equal to the original sample amplitudes, and the width of each impulse response is determined by the response of the filter (usually the cutoff at the Nyquist frequency of the input samples). The summation of each response contribution gives the interpolated samples. Each of the three interpolated samples is formed by adding together the four products.

A brick-wall filter must bandlimit the output spectra. Oversampling filtering, the images between 20 kHz and 156.4 kHz are suppressed. Although in this case a four-times oversampling filter was employed, two-times, or eight-times oversampling digital filters are commonly used.

DAC can convert an oversampled waveform far more easier, because the successive changes in amplitude are smaller in an oversampled signal. The slew rate, rate in variation of the output waveform, is lower. Also, less ringing and overshoot, reduces intermodulation distortion. Digital filters tend to be more stable, have linear phase, and are less susceptible to temperature, age, etc.

Noise Shaping

Oversampling also provides a decrease in audio band quantization noise, because the total noise power is spread over a larger frequency range. The quantization noise floor can be calculated by *10log(R)*.

Noise, or spectral, shaping is accomplished by sigma-delta modulation. Noise shaping can significantly reduce the in-band noise floor and can be effectively applied to both ADC and DAC. If the quantization errors are independent, or uncorrelated, the noise spectrum is white; if the nature of the error dependency is selected, then the noise spectrum can be shaped to any desired frequency response,while keeping the audio signal unchanged. Noise shaped is produced by algorithms that select the desired statistically dependent errors.

A simple noise shaper works by taking in the data words (28-bit) from the filter and rounds and dithers them to crease the MSB of 16-bit words. The remaining 12 LSB are delayed by a sampling period and subtracted by the next data word. The result is a shaped noise floor in decreased in the audio band. However, the out-of-band noise is high in frequency and is less audible, but can be attenuated by the output filter. Noise shaping must be used in low-bit converters so that the noise floor is minimized to acceptable levels. Usually, noise shaping and oversampling are used in conjunction.

Output Processing

After digital filtering is applied, the data is converted to analog form with a DAC. The aperture effect of an output S/H circuit is manifested by a null at $R \cdot f_s$, thus suppressing the oversampled image. By slightly boosting the high frequencies in the digital filter, the attenuation can be compensated. The remaining band around this image can be removed with an analog filter.

Alternate Coding Architectures

While linear PCM is considered to be the classic audio digitization architecture, other digitization methods can offer other advantages, as well as disadvantages. Linear PCM systems represent a fixed scale of equal quantization intervals that map the analog waveform. The quantizer's word length determines the number of quantization intervals available for encoding.

Other systems offer new mapping techniques that offer data reduction and more coding efficiency. However, in spite of these advantages, audio fidelity can be sacrificed.

Longer word lengths reduce the quantization error, but also increase the data bandwidth. Uniform PCM quantization is optimal for uniformly distributed signals, but most audio signals are not uniform. Alternative PCM systems can reduce the quantization error by using nonuniform quantization levels.

Floating-Point Systems

This type of systems use a modified PCM architecture design to accept scaling values. It is an adaptive approach that uses nonuniform quantization. The scaling factor is instantaneously applied from sample to sample. Floating-point quantizers creates its data words in two parts: the data value (waveform value and scaled amplitude) and the scale factor (quantization step size). The scale factor varies with the gain of the signal in the PCM ADC. Thus, the converter is used more effectively (increased accuracy) by boosting low-level signals and attenuating high-leveled ones.

The SNR in floating point systems is determined by the ADC's resolution. Changes in the signal will determine the changes in the gain structure, which affect the relative amplitude of the quantization error. For longer-word converters, a complex signal can mask the quantization error; however, for simple tones the error can be audible.

Block Floating-Point Systems

This architecture is derived from floating-point systems and offers the advantage of data reduction and its usefulness in bandlimited transmission and storage. Block floating-point systems have a linear fixed PCM ADC preceding the scaling factor. Only a short duration of the analog waveform is digitized and a scale factor is calculated depending on the largest value represented in the block. The data is then scaled upward so that the largest value is just bellow the full scale. This reduces the total number of bits needed to represent the signal. During decoding the data is properly rescaled.

Block floating-point systems avoid many of the audible artifacts introduced by instantaneous scaling and the gain error is reduced. This system performs well with audio signals.

Nonuniform Companding Systems

In this system, the quantization levels are spaced far apart for high-amplitudes and closer together for low-amplitude signals. This is accomplished by compressing and expanding the signal. If the signal is compressed prior to quantization, the smaller values are enhanced, and large values are diminished. Companding is done with a logarithmic function and a linear PCM quantizer is used. The compressed signals “see” the quantization intervals as uniform, the conversion is equivalent to one nonuniform step. At the output, an expander is used to inversely compensate for the nonlinearity of the reconstructed signal. In this manner, the quantization levels are more effectively distributed over the audio dynamic range.

The SNR is increased for smaller signals and can increase the overall dynamic range. However, noise is increased and is correlated; but, the signal amplitude tends to mask this noise.

µ-Law and A-Law Companding

Both systems are examples of logarithmic companding. µ-law encoding has nonuniform quantization step sizes increasing logarithmically with signal level. The signal is compressed prior to quantization and the inverse function is used for expansion. The value of 0 corresponds to linear amplification (no compression), or uniform amplification. Larger vales result in greater companding. A-law's quantization characteristic varies logarithmically.

Differential PCM Systems

This system relies on the fact that the measure of a waveform changes from sample to sample. This difference between samples, then requires fewer encoding bits. Differential systems exploit the correlation between sampled signals, and thus uses this redundancy to encode fewer bits in amplitude instead of encoding the entire waveform.

Differential PCM (DPCM) coding uses a type of predictive coding in which a predicted signal is subtracted from the input and the difference signal (error) is quantized. The coded value is the difference between the prediction and the actual signal. The decoder produces a prediction from the previous data. In this manner, the waveform is reconstructed sample by sample.

If the prediction is accurate and the error is small, predictive coding requires fewer bits to quantize an audio signal. The frequency response of the coded signal can be filtered to improve the SNR and the inherent noise is masked by the low frequency content.

Delta Modulation

Delta modulation is a form of DPCM that uses a high sampling rate so that only a 1-bit quantization of the difference signal is needed to encode the audio waveform. Positive and negative transitions in the quantized waveform are used to encode the audio signal. The coded signal amplitude decreases with frequency, so the SNR decreases by 6 dB/octave. Delta modulation offers excellent error performance because there is no MSB. Each bit tracks the difference between samples, this limiting the amount of error to that difference. If the encoding cannot track a complex audio waveform, transient distortion will occur. Degradation is possible, so error correction techniques are necessary.

DM fails to perform in high fidelity applications because of the trade-off between sampling frequency and quantization step size. High sampling frequencies increase the coded signal's bandwidth and quantization error is ever present. However, the high sampling frequency used in DM benefits with noise shaping.

Adaptive Delta Modulation

The quantization increment in ADM varies in size to overcome the transient response limitations of DM. The adaptive algorithm can achieve good tracking for rapid changes in the waveform. As a result, the transient response is improved and the SNR is increased without changing the sampling frequency or bit-rate.

Implementation of this system is difficult. Quantization noise increases with large signal increments caused by high-frequency and amplitude signals. Dither is difficult to inject because of the step size.

Companded Predictive Delta Modulation

CPDM varies the signal's amplitude prior to the constant step size delta modulator to protect against modulator overload (the sampling frequency is insufficient to track the signal's rise time, thus the differential is encoded instead of the original signal). Linear predictive filtering (coding) is used to reduce the quantization error. Digitally controlled amplifiers control the broadband gain. The bitstream controls the amplifiers in order to minimize the tracking error. These amplifiers continually adjust the signal over a large range to best fit the fixed step size of the DM. Overloads triggered by 1s and 0s ensure that the compression of the broadband gain adjusts the transients so that they do not clip at the modulator. Spectral compression can be used to reduce variations in the spectral content. Thus, the spectrum at the ADC is nearly constant.

Adaptive Differential Pulse-Code Modulation

This system uses predictive coding to achieve data reduction, as well as, combining the adaptive difference signal of ADM with the binary code of PCM. In general, the difference signal to be coded is scaled by an adaptive scale factor, and then quantized according to a fixed quantization curve. The scale factor is selected according to the signal's properties. Step sizes can be effectively varied directly with a varying step size, or by scaling the signal with a gain factor.

A linear predictor is used to output a signal estimate of each sample. This signal is subtracted from the input signal to yield a difference signal. This difference is the coded with a short PCM word (4-8 bits) and output from the encoder. In this manner, the signal is adaptively equalized, and the quantization noise adaptively shaped.

Time Based Correction

Modulation coding is used in storage and transmission channels to improve coding efficiency and make the data self-clocking. Successful recovery is limited by the timebase accuracy of the received clock. Timing errors caused by speed variations in the transport of an optical disc player, embedded clock instabilities in the data stream, and timing inaccuracies in the ADC or DAC oscillator clock can lead to data errors, noise or modulation artifacts in the converted waveform. Receiver circuits (for example, PLL) must minimize timing errors so that the data can be fully recovered.

Timing accuracy can be difficult in a digital environment, because of the present noise and interference. Also, timebase tolerances increase with word length.

Jitter

Jitter is defined as any deviation in the zero-crossing times of a data waveform from the zero-crossing times of a perfectly stable waveform. In an analog signal, timing variations can be perceived as instabilities in pitch.

However, a digital signal may cause bit errors in the bitstream or be indirectly audible as increased noise and distortion in the output analog waveform. Jitter is always present; its effect and tolerance depend on where in the signal processing chain the jitter occurs. For example, in the AD or DA conversion, low levels o jitter can induce artifacts in the analog output waveform.

Peak-to-peak jitter is the range in which each ideal transition is a period of variation or uncertainty in arrival time. Jitter in data can be generated in the storage medium, transmission channel, processing/regeneration circuits (ADC, DAC). Or it can occur as random variations in clock edges (white phase), variations in clock pulse (white FM), or to other events.

Spectrally, jitter can be described as the following: random jitter has a broadband spectrum (produces an increase in noise floor), periodic jitter appears as a single spectral line, FM jitter has sidebands. Jitter at frequencies less than the sampling frequency cause an accumulation of timing errors.

Eye Pattern

An eye pattern is shown as a display of superimposed successive transitions triggered from a stable reference clock, and timebase set to one unit interval. This pattern is used to interpret the quality of the received signal. It will reveal noise as the waveform's amplitude variations become indistinct, and jitter as the transitions shift about the time intervals of the code period. Other errors, such as peak shift, dc offset can be observed as well. Noise in the channel will tend to close the pattern vertically, while jitter closes it horizontally. The eye's width gives the percentage of the data period available to ascertain its logical value. The height, specified in volts, shows the maximum difference between these levels during the available time.

Interface Jitter versus Sampling Jitter

Interface jitter occurs in digital-to-digital transfers. Sampling jitter occurs when converting data into and out of the digital domain.

Interface jitter is of concern when it causes uncorrected errors in the recovered signal; the quality of the transmitted data can be monitored by error detection circuits at the receiver, such as a PLL. If the recovered data is error free, then interface jitter has not affected it. However, if any subsequent reclocking circuits do not remove the jitter, audible artifacts can result from sampling jitter at the DAC.

Sampling jitter affects the quality of an audio signal as it is sampled or resampled with this timing error. This type of jitter can be inducted by sampling clocks. The artifacts produced by sampling jitter will be heard as noise and distortion.

Jitter can occur at any stage of the signal chain and if not properly dealt with, jitter will accumulate in the throughput signal. Also, each device will contribute a small amount of jitter, as will the interface connections resulting in data errors or conversion artifacts. Badly jittered signals can be “cleaned” if retimed correctly with jitter attenuation.

Jitter in Storage Media

Storage media, such as magnetic tape and optical disc, can impose timebase errors in the output data signal because of speed variations in the mechanical drives they use. Speed variations in the transport by the eccentricities in the rotation of the capstan and spindle motors will cause the data rate to vary; the transport's speed can slowly drift or fluctuate rapidly. If this variation is within tolerance, no error will occur in the signal. Accurate clocks and servo systems must be designed to limit the mechanical speed variations. Input and output data must be buffered so the effects of data irregularities are absorbed.

Servo control circuits are used to read the timing information from the output data and generate a transport correction signal. Most of these circuits are PLL's. Jitter generated during transmission to an interconnection (electrical or optical) can be minimized if the data is reclocked though a buffer.

Jitter in Data Transmission

The receiver in a transmission chain must accomplish two tasks: data recovery and clock recovery. If the data is transferred but will not be generated (converted to analog) at the receiver, only data recovery is necessary. In this case, interface jitter is the cause of timing errors at the receiver. However, if the data is to be generated or requantized, data recovery and clock recovery are needed. High jitter levels may compromise the receivers ability to derive a stable clock reference needed for conversion. Low jitter tolerance depends on the particular converter design. PLL circuits, memory buffers, and sample rate converters are used to correct any timebase error present in the data stream.

Jitter transmission levels depend on cable characteristics. Cable induced jitter, or pattern-dependent jitter, is a modulation that depends on the data values themselves. For example, a pattern or 0s might produce more delay in transitions than patterns of 1s. The amount of modulation is a function of the cables bandwidth (high frequency attenuation). Because of this reason, some transmission protocols use a fixed synchronization pattern. In this way, the pattern is static and the pattern dependent jitter is removed.

For more complex installations, jitter protection is done with a master clock signal. In this configuration, jitter does not accumulate with each device, instead each device ignores the jitter in its input data and only accepts the clock information from the master clock reference.

Jitter in Converters

If not properly taken care off, jitter at the conversion points will produce degradation of the output waveform. Clock jitter at the ADC results in wrong sample values at the wrong time. Even if these samples are presented to the DAC with a jitter-free clock, the result will be the wrong samples at the right time. The amplitude of the error increases with frequency. Therefore, ADC designs must limit clock jitter by using a clock oscillator with very small jitter tolerances (usually less than 10 ps). At the DAC end, jitter will manifest itself as an increased noise floor and distortion at the output analog signal.

Jitter error can be observed by the jitter spectrum. Jitter bandwidth may extend beyond the Nyquist frequency. If not filtered, these components will fold back into the audio band and create noise and harmonic distortion.

High-quality audio equipment does not present serious jitter issues. Jitter based distortion is quite low, while the threshold of audibility is high. Meaning that at normal listening levels the effect of jitter is masked, even in the absence of a signal. Masking effects allow for high jitter and distortion levels to be acceptable.