www.brightnoise.nl

From microphone to loudspeaker

frans.sessink.nl

What is wrong with digital audio?

CD signal from scratch
Digital audio for CD is specified for two channels (left and right stereo), each having 16 bits word length and sample frequency 44.1 kHz. Perfect AD conversion (no errors on integral and differential linearity and brick-wall 20 kHz low-pass filter before the sampling) allows audio signal to noise ratio of 16 * 6.02.. dB + 1.76 dB + 10 * LOG10 (22.05 kHz / audio_BW) or 98.519 dB in 20 kHz audio bandwidth. (See http://en.wikipedia.org/wiki/Dynamic_range).

The input of the AD converter must have an additional dither signal (white noise or spectrally shaped noise) to be sure that low-level signals are digitized in a correct way. Alternatively the cumulative quantization error can be used in a feedback system to have at least a first-order noise shaping. The addition of dither is at the cost of the measured audio signal to noise ratio but gives an improvement of the subjective sound quality. The crackling sound with low-level input signal, a result from an error signal that is correlated to the momentary audio signal due to the amplitude quantization, has changed into smooth-sounding bright noise.

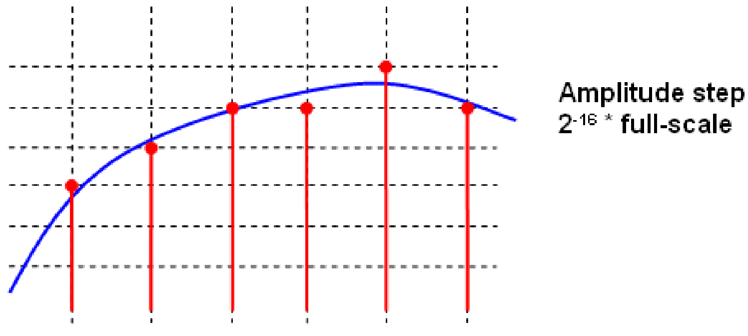

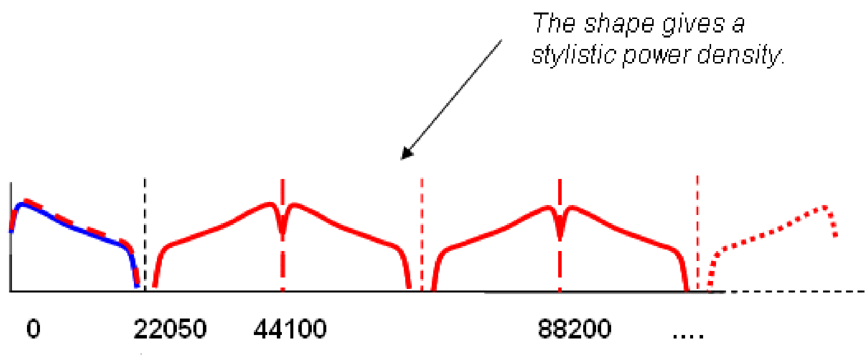

Reproduction in a flash DAC with this 44.1 kHz sample frequency results in the calculated audio signal to noise ratio. When processed without further errors, the distortion is equal to or hidden in the quantization noise; now at the output present as signal plus noise. Main problem in this virtual reproduction chain is the spectral repetition of the signal, causing high-energy above the audio bandwidth of 20 kHz. The audible signal is a perfect copy of the input signal at the coding side; as good as it can be with the 16 bits quality limitation. There is no need for more than 16 bits DAC resolution; a specification of 18 bits or whatever is although an indication that the linearity of the 16 MSB-side bits is sufficiently accurate.

Nothing wrong with this signal within the given boundaries. CD program material is still the best available quality there is and many audio freaks are happy with the 16 bits. See www.audio-life.nl (http://www.audio-life.nl/index.php?pid=28) or www.audionote.co.uk
Human ear response is a sufficient linear low-pass filter.

There are however some problems with the repeated spectral energy. One is that the loudspeaker must be able to handle the high-frequency signal amplitude in the range above 22.05 kHz. Imagine a smoking tweeter when the audio amplifier output power is sufficiently high. The tweeter power handling is not rated for the full energy popping up around 44.1 kHz. The second problem is that the audio amplifier must be able to handle the input signal without adding distortion. When there is a non-linearity in the chain then this will cause an unwanted detection or, in other words, an unwanted conversion from inaudible high-frequency signals into audible output signal. A simple explanation for the detection phenomenon is the use of an amplifier with high feedback and insufficient slew-rate and power bandwidth. *The higher the feedback, the more tragic the result.*
Asymmetric slew-rate limitation is the worst there is.

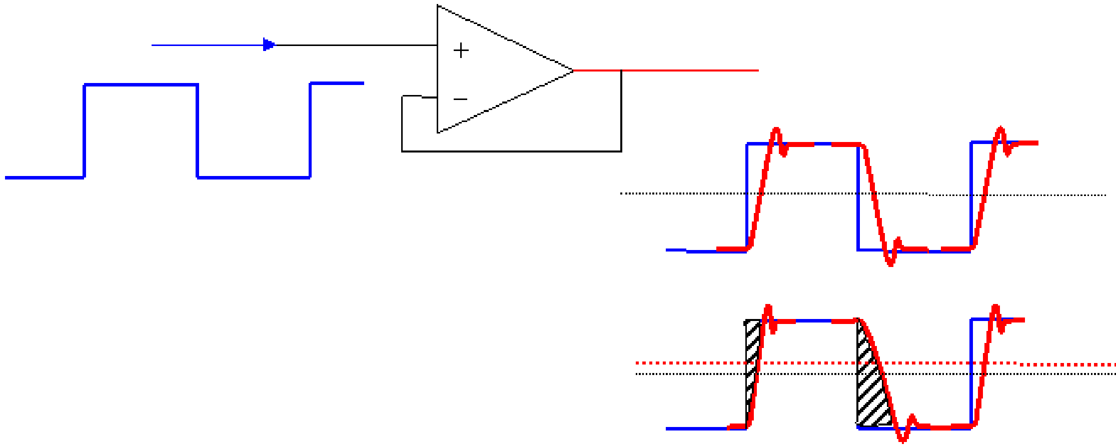

The output signal of a DAC contains a high-frequency spectrum and also a very fast settling quantized level signal. Most high-performance DA converters are settling within a microsecond. The attempt to average the DAC output signal with an operational amplifier is a challenge. Some designs use a transformer in advance of further processing to have a linear averaging and at least some low-pass filter function without the above mentioned detection phenomenon.

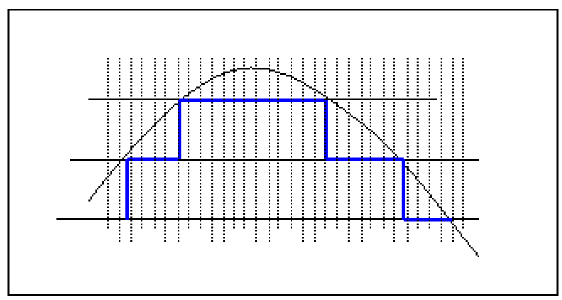

To force the linear averaging of the DAC output signal some designs use vacuum tubes or special op-amps with high power bandwidth (high gain-bandwidth product AND high slew rate!).
Example with a valve can be found at http://www.audionotekits.com/dac2_1.html
Of course: a transistor in common-base or FET in common gate is also a good and linear current-to-voltage converter. When you know what you are doing this can be a perfect choice!

Analog active or passive filters with pass band up to 20 kHz and sufficient stop band attenuation starting at 22.05 kHz are problematic for reason of complexity and left/right channel differences. In most systems available on the market there is digital filtering for the attenuation of the unwanted frequencies above the Nyquist frequency 22.05 kHz. Such filters make only sense in over-sample systems.

Up-sample or over-sampling filters
Over-sampling (repeat values or insert zeroes) of the 44.1 ksps (kilo-samples per second) 16 bit signal with a higher (integer ratio, synchronous!) frequency gives the same signal spectrum at that higher sample frequency. Example: 8 times over-sampling (352.8 kHz). A digital low-pass filter, FIR or IIR, can be inserted. The output spectrum for the example value of 8 times over-sample plus a filter system with perfect attenuation is zero to 20 kHz wanted spectrum plus the aliased spectrum of this signal, now starting at 352.8 kHz – 20 kHz = 332.8 kHz. In-between

there is no spectral energy. After a DAC (that can handle the 352.8 kHz sample frequency without settling problems and without frequency-dependent inaccuracy!) there is need for a perfect-averaging current-to-voltage converter and a simple low-pass without distortion. May be there are active filters with op-amps that can be used, but just to be sure ***use passive components in the first filter sections***.
Most op-amps are not even aware at the output that the input is highly overloaded with transients!

So what is wrong with up-sample filters? The input word length is 16 bits and the coefficients are ***for example*** also 16 bit. The output of each multiplication has word length 32 bit in this example.

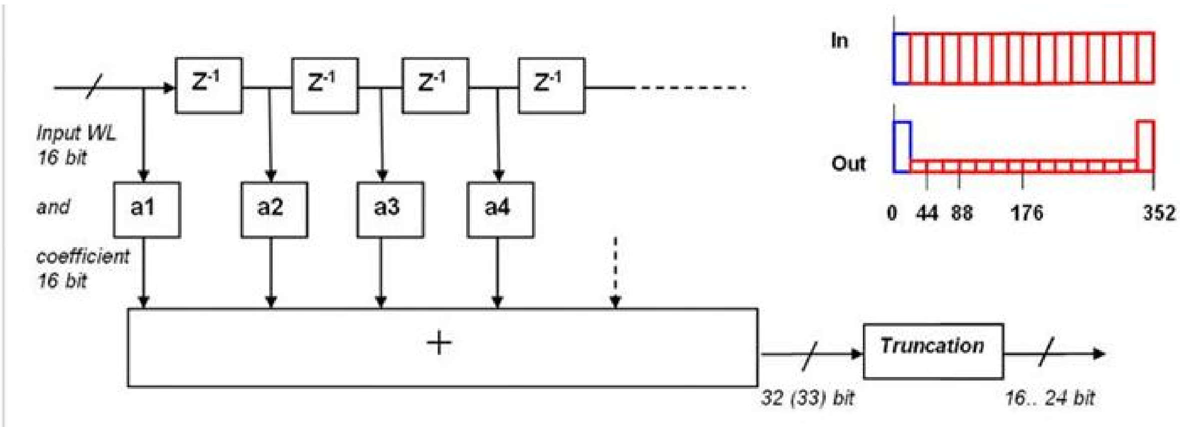

Input 16 bit word length, coefficient 16 bit word length, output signal 32 bit word length.

Due to group delay properties of the filter even 33 bits can be required for linear signal handling!

Now the problem: common use is truncation of the output word length to a value slightly more than 16 bits, say 18 or 20 bits or even 24 bits. Where in the chain is this truncation executed? In some chip sets the outputs of each multiply action is already truncated at the input of the filter accumulator. Such a truncation makes the filter transfer described by frequency-dependent amplitude and phase or delay also and unwanted dependent from the signal amplitude. Example: coefficient value 2^{-10} (that is 1 / 1024) and output word length 20 bits results in MPY (coefficient a1, a2 etc.) output amplitude zero when the input signal is below a certain level. In this example the FIR tap signal just disappears for filter input level signal value 2^{-10} (related to full-scale) or -60 dBFS. ***So the filter transfer is not constant at all when playing music!*** The signals on left and right are not equal in amplitude and phase (or delay) and the filtering is with a dynamic music or speech signal slightly different for the two stereo channels! Measurement with a stationary sine wave signal will not reveal this unwanted phenomenon. Listening to a stereo low-level recorded violin solo however will expose the phenomenon.

I think that this is part of the explanation for the poor stereo image reproduction from many CD-players.

Apart from this signal-level dependent filter characteristic there is something in the truncation of the output signal of the full-adder that combines the tap signals. Even when the multiplier has full 33 bit (ringing, overshoot) word length there can be a problem. For the above given values of 16 bit signal and 16 bit coefficient the word length should be more than 32 bit since the filter shows some ringing or overshoot in the impulse response. The truncation to 16 or 18 or 20 or 24 bit gives a new quantization error, correlated to the audio signal. Without dither addition at this point in the chain the audible effect is a crackling sound with low-frequency components, correlated to the audio signal. The measured amplitude of this additional crackling noise can be 30 dB below the noise floor from the recording and ***can still be audible as unwanted signal.*** Bright white noise accompanied by crackling noise.

This word length truncation is a new quantization in the system and must also have a dither addition before that quantization process; the truncation The already present input noise is not covering the new quantization noise at all!
Alternative solution for dither is an error signal accumulator system that assures an average-zero error over longer time. This is a first-order noise shaping and is probably a better solution.

Flash DAC or sigma-delta DAC
A flash DAC converts the audio to a value at the moment of a clock edge. Clock timing error gives the correct sample value at the wrong moment. For low-frequency output this is not a problem, since the previous and next samples have almost the same value. For higher audio frequencies the error can be calculated via the slew rate of the audio signal and the rms-jitter of the clock edge. The assumption used here is that "correct value at wrong moment" can be expressed in "wrong value at correct moment" via the slew-rate of the signal. Maximum audio frequency of 20 kHz (period time 50 µs) and full-scale amplitude of 16 bit converted in 2 volt peak-to-peak or 700 mV rms sine wave output signal gives slew rate of 2 * π volt / 50 µs = 126 kV/sec or 126 mV/µs. The error signal is a sampled sequence, 44.1 KHz, with momentary error values equal to "slew rate * jitter_rms". The outcome is a Nyquist band from zero to 22050 Hz that contains the jitter-caused extra unwanted energy. The assumption is only valid for jitter according Gaussian white noise distributed situations. Spurious FM-tones on the clock result in a different effect, not at all comparable with noise.
To guarantee an audio noise floor of 110 dB below full scale, that is 12 dB headroom for safety, the noise floor is allowed to be (0.7 V rms - 110 dB =) 7 µV rms. Jitter in the order of (7µV/ 126 kV/sec =) 55 picoseconds rms causes this worst-case situation. That can be done but S-PDIF clock regeneration with this specification requires some extra effort in the design. Attention: this is for white noise. Some people claim to hear effects from discrete-frequency jitter modulation with lower level. The human ear is a narrow-bandwidth FFT analyzer with unsolved mysteries.

In the near future you can find here on this internet site an explanation of DAC settling time effects (exponential settling versus slew-rate limited settling) on experienced audio quality. Settling behavior is in my opinion more important than jitter or DAC linearity.

For a sigma-delta DAC the situation is different. The one-bit data stream is calculated in a numerical machine and the feedback path for the DAC assumes correct and constant values for "high" and "low" or +1 and -1 "unit", the outputs of the quantizer. The average of the bit stream values is only correct when the energy for a +1 or a -1 value is converted into equal amounts.

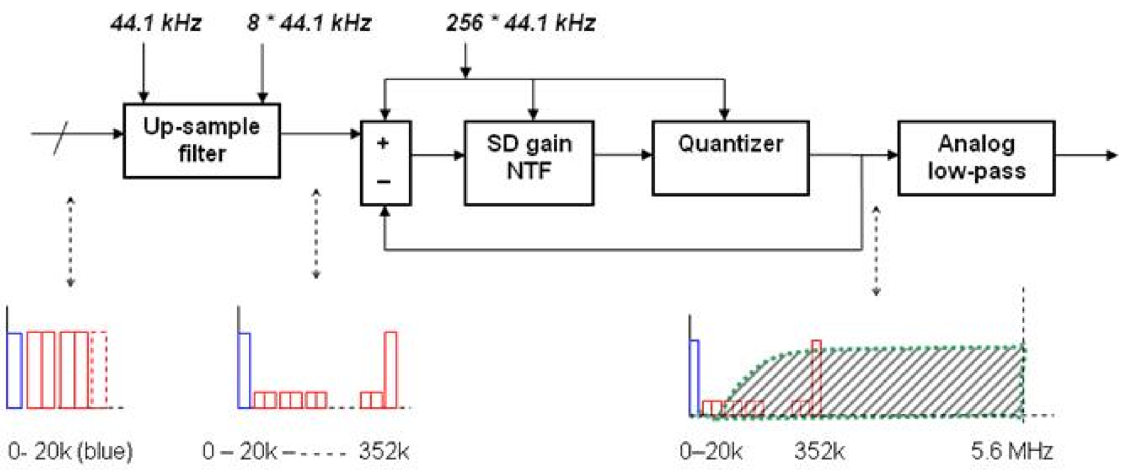

The bit stream sequence is for example 11.289 MHz (256 * 44.1 kHz). The spectral energy is always (signal plus noise) 1 unit rms in 5.6 MHz Nyquist bandwidth. The wanted-audio output (full-scale) is about 0.4 unit rms, since a sigma-delta can not be used up to its limits. Signal to noise of 110 dB is reached when the jitter causes a signal of 0.4 unit rms – 110 dB = 1.2 µ-unit in 20 kHz, corresponding to 19 µ-unit in the full Nyquist bandwidth. The conversion of the high-frequency noise into an increased low-frequency audio noise floor is a simple multiplication, or for radio engineers identical to frequency conversion. The result from 19 µ * sample clock period time or 19E-6 * 88.6 nanoseconds = 1.68 ps. This is the maximum allowed rms-jitter of the clock signal. Expressed in phase noise this corresponds to -153 dBc/rtHz for the wideband noise floor of the 11.289 MHz oscillator. This is 10 to 20 dB better than standard crystal oscillator performance. Hope your best is good enough!
The calculation for a multibit sigma-delta DAC is different; a multibit sigma-delta DAC is less sensitive for jitter.

The challenge in design is to regenerate a clock signal from S-PDIF to crystal-reference-stability with the mentioned jitter.

Keep also in mind that the bit stream is multiplied with the supply voltage in the DAC. Where common balanced circuitry cancels the supply voltage, here the supply rejection ratio is nil! And even worse: high-frequency noise on the supply line is converted into audio frequency output signal in this radio-like frequency converter function.

Be aware of the fact that the output spectrum of a high-performance DAC like the TI/Burr-Brown http://pdf1.alldatasheet.com/datasheet-pdf/view/84737/TI/PCM1794.html is not given by the transfer of the up-sample filter (figs 1 and 3 in the specification). On top of that filtered audio signal there is the sigma-delta shaped noise: high energy above 100 kHz! The low-pass filter to remove the sigma-delta noise is the extra 3 op-amps per channel output I-to-V converter/filter. Do you trust the speed of the op-amps in the application diagram from the PCM1794?

CD player read errors, error correction and concealment
The compact disk player reads all information once. The error recovery can repair few read errors. When the error recovery has reached its limits then there is a graceful degradation that masks the missing information. This is called concealment. See http://www.digital-recordings.com/cdcheck/cdmore.html for a basic explanation. The better way for a reproduction chain is an exact copy of the data and storage on a hard drive or an SD-card. The BER of such a device is much lower, when the original CD recording is copied with Exact Audio Copy. Store the uncompressed data (*.wav) and use a media player. The exact copy of the CD data can be