Metrum Acoustics’ Digital/Analogue-Converters (DACs) are praised in the international hifi-community and press for their musical presentation. Especially the purity of the sound, due to the absence of digital artefacts (“rubbish”), the natural sounding voices and instruments, the beautiful stereo image and the noticeably good separation of the instruments are consistently praised. Because of the unique qualities that define the DACs of Metrum Acoustics the listener is intimately included with the music. Therefore the music transforms from “something which is happening in the background” to a compelling and realistic experience.

Another aspect that should not remain unmentioned is the highly competitive quality-price-ratio. The DACs from Metrum Acoustics do not cower from the competition. In truth, they can easily compete with significantly more expensive DACs!

The question remains: what makes the DACs from Metrum Acoustics so unique? To answer this question, the designer Cees Ruijtenberg has written the following text, where he explains his design philosophy and why his DACs are different from “normal” DACs.

Design Philosophy Metrum Acoustics NOS DACs

Introduction

In 2008 I decided to start an investigation in order to map the sound characteristics of different DACs. The performance of these DACs was largely unconvincing, though I could not yet tell what the underlying causes were. My dissatisfaction was fed by my frequent visits to concerts and the seeming impossibility to bridge the gap between the *live* performances and digital recordings. For many years already I had been interested in the technical side of recording music and I had from time to time made some. Due to that I was in the possession of a decent set of microphones, a mixer and a studio-tape recorder. Despite the barely acceptable signal-to-noise ratio and the limited dynamics (≤ 60dB) that the recording equipment could handle back then, the experience felt “real” and involved you in the music. That is why I understand vinyl-lovers, despite the pops, noise and scratches that are inherent to this medium. I wanted to recreate this sense of realism and feeling with digital equipment en I started my investigation in the hope of recreating it. Because of my work and experience in the field of electrostatic loudspeakers I have many contacts there and they gave me the opportunity to listen to a so-called NOS (*NonOverSampling*) homemade DAC. Despite its easily apparent shortcomings I heard something there which had been missing all along in other DACs: the emotion, the involvement with music and the experience was suddenly back. I wondered was this the road I should take? From that moment on I shifted my investigation to studying the defects of contemporary digital equipment, which is generally based on the widely used method of “oversampling”. As a designer of electronics I possess advanced measuring equipment and with my roughly 35 years of experience I should be able to pursue this. Despite this, it kept surprising me that in a system, as a direct consequence of oversampling, many artefacts are shown that really should not be in the audible range.

Ryohei Kusunoki

Because I was cautiously optimistic about the sound characteristics of NOS DACs, I actively started searching for any knowledge which was available on this subject. At various fora, it soon became apparent that I was not alone in searching for the “real” music experience. I also came across an article from Ryohei Kusonoki who explains the often used oversampling methods in a comprehensive manner. The article shows that, when ones vision is based on experiencing live music, it seems odd to rely on oversampling. This is why I set aside all of my knowledge concerning digital registration and imaging and decided to follow my heart.

Modern Techniques

Both hobbyists and professionals who are convinced of the NOS-principle have to make do with old DAC-chips, that were once developed by Philips and which hit the market in the 80's. For the time, they were truly amazing DAC-chips. Companies such as 47 Labs, Zanden, Audio Note and Abbington Music Research, who are all convinced of the validity of the NOS-principle, are, due to the lack of more modern components, forced to use old chips such as the TDA1543 or the TDA1541. They are right in doing so, for the more modern chips on the market have grown increasingly complicated and are burdened in most cases with FIR-filters that, though they make oversampling techniques possible, make it impossible not to use them.

Because I had, as an electronics-designer, often developed products for industrial purposes I had gained much knowledge of industrial components. Especially in the fields of process and medical engineering, DAC-chips are used without the earlier mentioned FIR-filters. Would it be possible to utilize these chips for audio-products? Which characteristics should I focus on in order to improve on the "old" TDA1543 or TDA1541?

It was clear to me, that modern DAC-chips offered several advantages compared with the "old" DAC-chips. Since the 80's great headway has been made regarding switching noise, the speed and conduction of MOSFET-switches, which are used in so called R2R-ladder circuits and finally the linearity of earlier mentioned circuits was improved. Outside of these improvements little changed. However, it seems that these characteristics are important for, among others, the ease and naturalness with which the music is portrayed, improving the desired "real" feel of the music. Speed and bandwidth are features which are also of importance in amplifiers. This is also why op-amps seem to cause many problems in audio circuits. Terms such as "slew rate" and "open loop gain" apparently weigh far heavier here than Harmonic distortion does. It makes sense that the electronics that can portray music well are generally derived from radio frequency engineering, where the frequencies can be as high as 500 megahertz. When using components that allow such a bandwidth, completely different design techniques are required in order to build a product that is both good and stable. Examples of these are the designs made by Nelson Pass.

Despite the general availability of many measuring instruments in the past few years, it has only become apparent that measurements alone will rarely show if a product will be pleasant to listen to. Are the methods used then insufficient? No, absolutely not. In the design phase many problems can be solved by using relevant tests. But, for example, when you are testing for the spaciousness of the sound stage, current tests fall short. Then our own ears can be used as an additional instrument.

Bandwidth and speed are important factors in electronic circuits in the field of audio. This stems from the fact that signals which run through "old-fashioned" semi-conductor circuits can become distorted when tainted with high-frequent components. The technical term for this phenomenon is a "slewing-induced distortion". The circuit is in such a case not able to adequately follow the input signal. Our own hearing appears to be very sensitive to this: it often leads to a sound stage which is experienced as "flat".

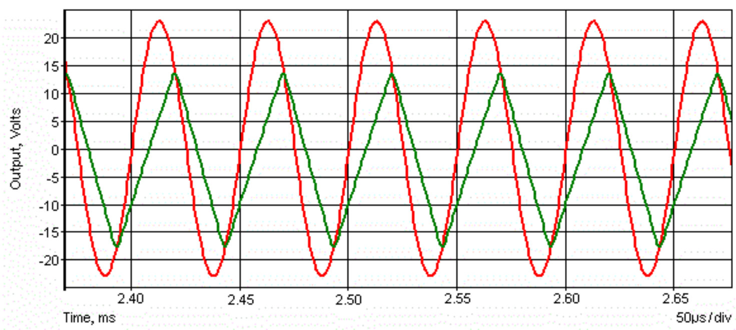

Limited slewing rate. The green line (output signal) shows the inability of the amplifier to follow the input signal (red)

The current output of a DAC-chip is one where a variety of audio- and high frequency signals are prevalent. The switches that control the R2R-laddercircuit in the DAC-chip, create high frequency switching noise. This noise, no matter how low, is present in the audio signal as an extra component and makes that standard components, such as op-amps, no longer function correctly. Despite that, many manufacturers use op-amps as a current/voltage converter, even though they are not suited for the job. It is important that the electronics that are used to convert the current (from the DAC-chip) to voltage, are both swift and have sufficient bandwidth.

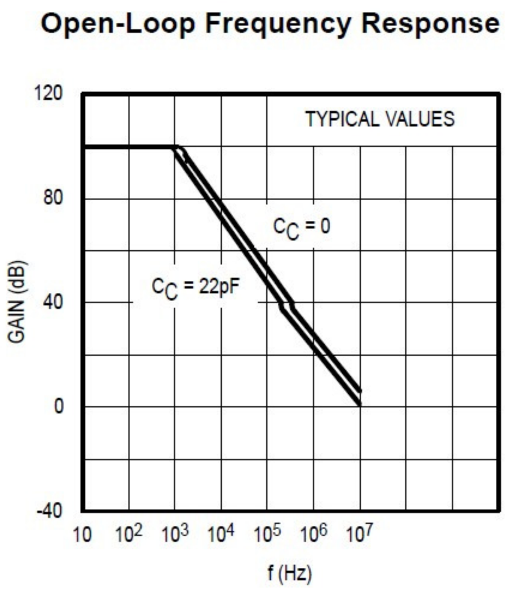

Open-loop-bandwidth of the NE5534 audio op-amp which stops, without feedback, at 1kHz already.

From 2008 to mid-2010 I sought a DAC-chip which had the characteristics which I described earlier. Eventually I found the correct DAC-chips and used them in the Quad, Octave, and the Hex DACs. That this choice was the correct one was made clear from the number of positive reviews we have received worldwide.

Perception

It is still remarkable, that though a sinus-shaped signal with a sampling rate of 44.1 kHz looks relatively choppy, our hearing does not experience it as such.

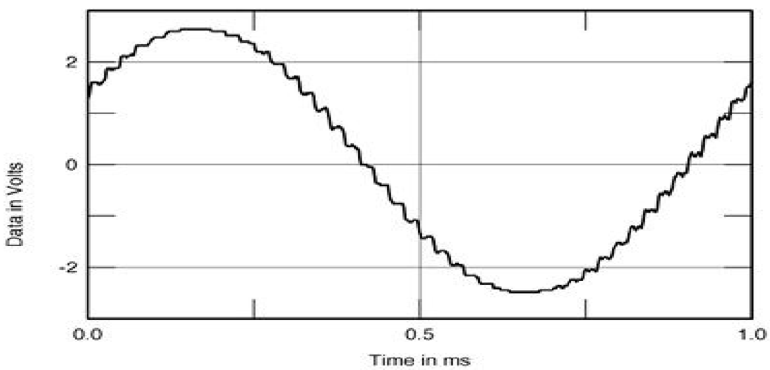

Sinus of 1 kHZ with a sampling rate of 44.1 kHz.

Some things will be enhanced because (assuming CD-quality) in one second merely 44.100 analogue values can be presented; so this does not add to the accuracy.

When looking at an oversampling DAC the values in between will, due to the FIR-filter, be calculated and filled in. A FIR-filter should (theoretically) give a better result, in practice however, it turns out it that there are no audible differences.

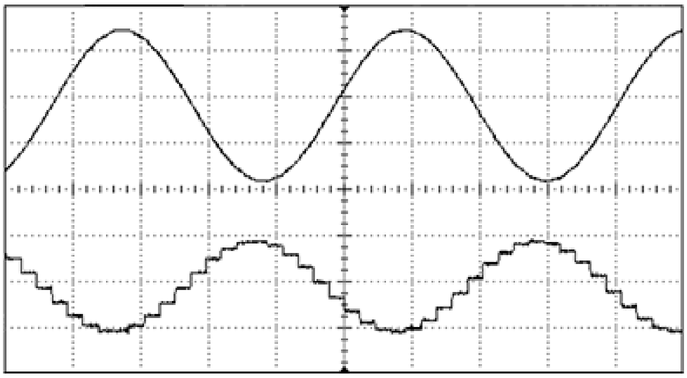

Below: the NOS-signal, Above: the signal after passing a FIR-filter

A fun experiment: let a NOS DAC (through an amplifier and loudspeakers) generate a sinus. Then record it with a microphone en show the result with the help of an oscilloscope. The picture below shows the result.

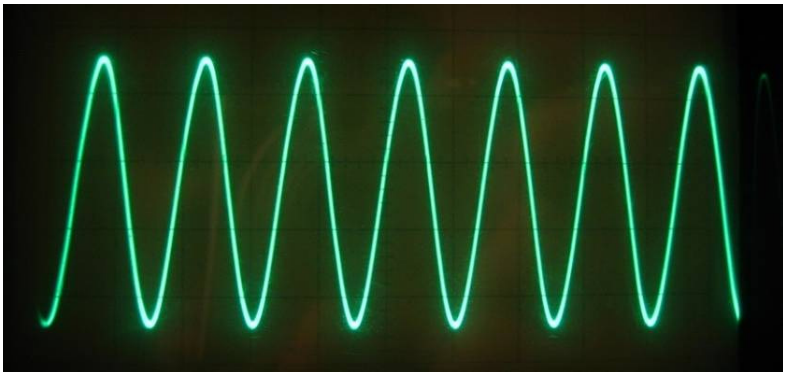

Sinus from NOS DAC to a loudspeaker, recorded with a measurement microphone

The choppiness that was originally present in the signal has completely disappeared. The reason for this, is that the signal has already passed several filters before it reaches the microphone. In the case of the Metrum Acoustics DAC there is a mild 70kHz (first-order) filter just before the signal is sent to the output-terminal. The pre-amplifier has, in order to prevent slewing-issues, also got a mild filter. This is generally a 100kHz with 6dB per Octave filter. Then the signal goes to the power amplifier where this is repeated once more. At this point the signal has already passed 3 filters. Even if the filtering is taking place at a relatively high frequency, the sharp edges of the sinus shaped wave will have vanished by now. That is not the end of it though! The signal is now offered to the loudspeakers. A loudspeaker is, in and of itself, a filter and will, at it's best, pass through the signals from the amplifier in the frequency area between 20 Hz and 35 kHz after which the curve drops dramatically. Depending on the characteristics of the tweeter the signal will eventually deteriorate with about 12 to 24 dB per octave. Where the first filters were mild and combined filtered away about 18 dB per octave, the tweeter alone nearly doubles the amount of filtering. Even then we have not finished. The measurement microphone that we used to record the signal kan register, at most, 35 kHz. The question now is, of course, whether this experiment was fair. If we are only testing a DAC, then that is justified. We are, however, testing for listening experience. Truthfully, the filtering used by the microphone is not quite sharp enough to simulate our ears. This deserves some explanation.

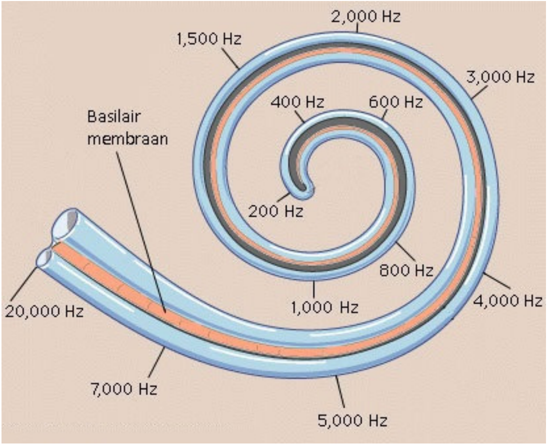

Cochlea with the basilar membrane

Pictured above is a part of our ears. It concerns the cochlea with is located in the inner ear. Every spot in the cochlea is sensitive to a certain specific frequency. The highest frequencies are registered at the front while the lowest are registered at the end of the cochlea. The basilar membrane, which is part of the cochlea, functions as a base for 15 to 20 thousand hair cells. Each of these hair cells are connected to a nerve which is connected to the brain. This is only a (very) short summary of the working of the ear because in truth, the ear is infinitely more complex. We however will limit ourselves to the sensory cells (hair cells) that each have describe a limited range of frequencies in such a way that they all overlap. The sum of all of these determines the range of our hearing, which on average is between 20 Hz and 20.000 Hz. With 20.000 hair cells their bandwidth is very small and thus very selective. This means that for one specific hair cell, except the specific frequency that it is tuned to, a signal gets filtered out at 40 dB per octave. This means that our hearing on this level behaves much like a band-pass filter, comparable to the filters found in CD-players. Implementing filters of this level in a CD-player can thus be described as overkill.

Applied testing methods

NOS DACs seem to score badly on tests, when you look purely at test results. This is partially correct, because the signal passes fewer filters before it is present at the output-terminal (see the above mentioned). Introducing a sharp filter at the output-terminal of a NOS DAC should improve the test results when it comes to harmonic distortion and noise characteristics. As a footnote, these are tests that are prescribed and generally agreed on by the industry. The problem with these distortion tests though, is that when measuring the amplitude-domain only static signals are used, such as sinus-signals. As soon as an impulse is used for testing this, the results are vastly different. At the bottom of this page is the transient response of a NOS DAC. The faint rounding at the top of the signal is due to the earlier mentioned mild filter at 70 kHz.

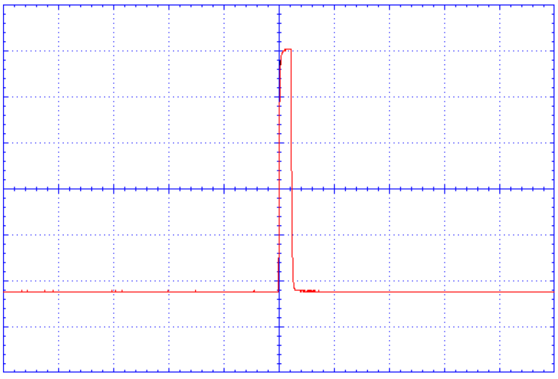

On the next page is the exact same impulse, but now passed through a DAC that utilizes oversampling.

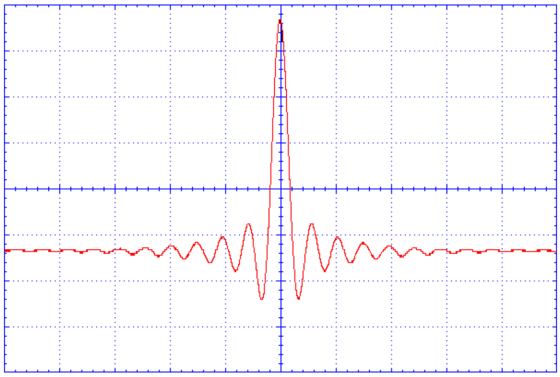

Because our hearing naturally functions as a strong filter, our brains tend to interpret the signal from the NOS DAC as if it has passed through a FIR-filter. This is due to the limited bandwidth of our hearing. Looking at the picture on the top of this page, we can wonder how the eventual picture will look if another equivalent filter is added by our hearing. It is well-documented by both musicians and authorities in the field of audio, that especially percussion instruments suffer from this effect. It is therefore not unfounded when NOS DACs are claimed to sound the most natural of all the alternatives. Because at the same time the test results for all NOS DACs fall short, the question can be raised whether the correct tests are being done to accurately gauge their quality. All measurements are, after all, performed without the benefit of any filter.

Typical oversampling dac (above) versus Metrum Acoustics dac (below)

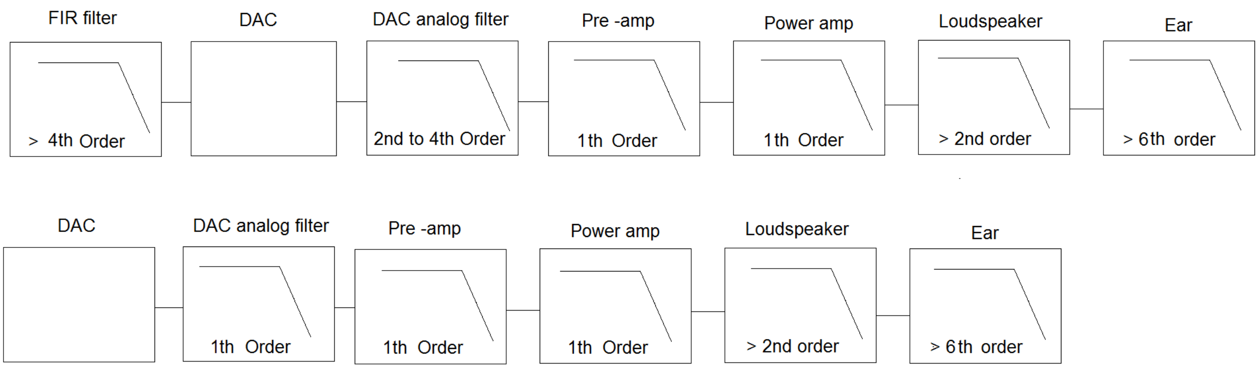

The FIR filter gives an extra link in the chain and contributes, aside from a strong filter, extra oscillation with impulses

Jitter

Jitter has been a point of interest for quite a while. Reducing jitter to extremely low levels is seen as essential for an improved sound stage. Without starting the discussion about how much jitter can still be perceived, it should be mentioned that there are different types of jitter. Because it remains large unclear in what way jitter is measured a short explanation is useful. The jitter in question is nothing more than a deviation in the time domain. You can compare this with the length of seconds for a clockwork: sometimes these are 0.9 and sometimes they are 1,1 seconds long. On average they are about one second long but the individual seconds differ in length. This is comparable to a DAC, where music consists of a series of (digital) samples. The samples enter the DAC-chip and are synchronized by the clock. The accuracy of the clock is therefore essential. Not only over longer periods of time, but each pulse should be exact in length.

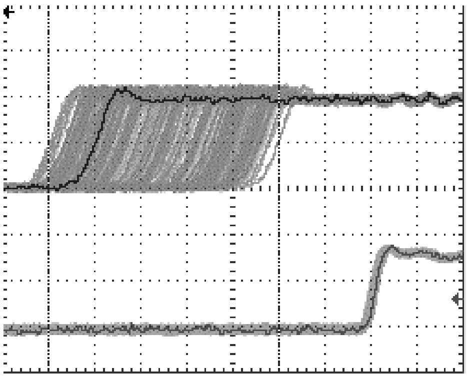

At the top of this page we have magnified two clock pulses, that jump from a low to a high level. The pulses have been observed over a longer period of time, and were stacked on top of one another. The bottom clock jumps at almost the exact same moment, while the top clock shows more random behavior. Sometimes the jump is early, sometimes late. This behavior is known as jitter. Aside from the fact that jitter can turn up at many points in the audio chain, the detectibility of the jitter is key to the performance of the DAC.

The picture clearly shows the easily measurable faults in the upper signal. Suppose each square stands for 100 picoseconds (ps), than the upper signal shows a deviation of at max 5 squares x 100 picoseconds = 500 picoseconds. The more stable the measuring equipment, the more reliable the results from the test will be. However, even measuring equipment will have inherent flaws which might negatively influence the results. When the results are very important the measuring equipment becomes as expensive as a midsize car…

There are also other techniques for measuring jitter, like spectral measurement. A stipulation for a good spectral measurement is, that a clearly defined signal must be used to measure by. Because the theoretical contents of a signal generate fairly low jitter rates, the jitter, that is directly caused by the hardware will show deviations in the spectrum.

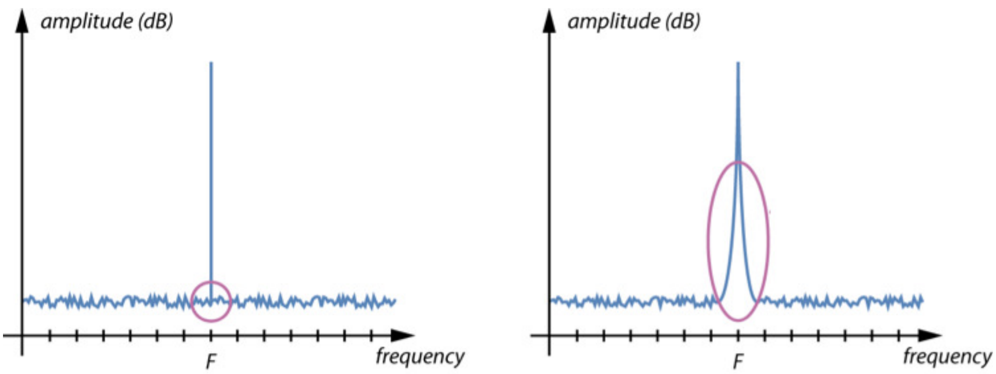

Results of a jitter test from a spectrum analyzer

The image above show two spectra. The left image is a signal without jitter: the signal is visible as a vertical needle. The image to the right shows a signal with a broad base, which is directly caused by jitter. Because jitter is a time-based deviation the spectrum analyzer will show jitter as higher and lower frequencies. This “swinging” signal causes the broad base and shows that a certain amount of jitter is present in the signal.

It is difficult to determine how much jitter is truly present. Julian Dunn († 2003), who was an important researcher in this field created a method in order to better map jitter. He designed a stimulus where the contests are chosen in such a way that in and of itself it should not give any jitter. Because the DAC has to reproduce the signal the effects of jitter that are caused by the hardware will be immediately visible in the spectral measurement.The method that Julian Dunn developed is called the “Jtest” The image below shows the Jtest applied testing a DAC.

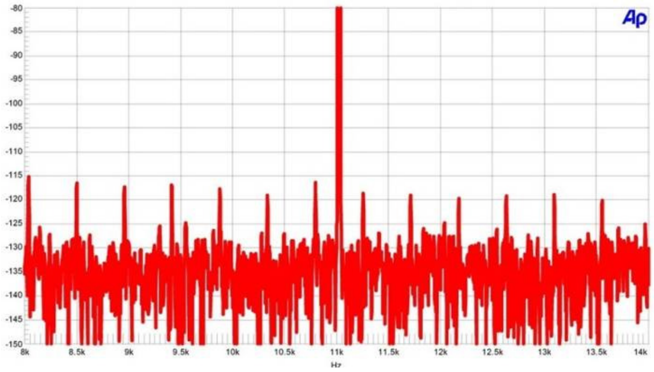

Jtest performed on a DAC

In an ideal situation the results would exclusively show the vertical needle. Because of jitter though on both sides of the needle many smaller needles are visible. By measuring the height (amplitude) of these little needles the amount of jitter can be calculated.

As I mentioned earlier, NOS DACs don’t test well. The jtest is not different in this regard. On the internet the results of many jtests for NOS DACs which have been made under similar conditions can be found. Similar conditions are key if any comparisons between products need to be made. Miller Labs are often asked to perform the tests for reviews and such. These tests are always performed in the same way under similar circumstances. It is remarkable that every single NOS DAC performs poorly under the Jtest. Is every single NOS DAC designer using inferior techniques, or are measurement errors the problem here?

The next test has been performed in our own lab. We regularly investigate the characteristics of DAC-chips that could potentially be used in hifi-equipment. The research design tends to be highly experimental and multiple components are connected and replaced by others. This type of research is often aimed at linearity, distortion, switching noise and sound quality. We mentioned earlier that relatively swift electronics tend to sound much better than their more sluggish counterparts. This effect could be observed with current/voltage converters and other components. Surprisingly, Jtest results turn out far worse when the components become swifter (with deviations as large as 2 nanoseconds). When we measure the Jitter in the time domain however, we can see that in the earlier mentioned test the results stay under 30 picoseconds. The hypothesis that the result might be distorted by components in the signal that come from outside the audio range was eventually confirmed. When utilizing a current/voltage converter having a very limited open loop bandwidth and ditto slew rate the results drastically changed. The Jtest now gave almost the same result as the time domain test did. When listening to the music though, the sound stage had turned flat and lacklustre.

When using the spectrum analyzes for jitter tests it is striking that not just time deviations, but also amplitude deviations lead to a broadened base. Also signals from outside the audio range, or a 50 Hz hum caused by an inferior power supply, can give a distorted view. So can an (inaudible) signal of 88 kHz and a (equally inaudible) signal of 91 kHz give a difference frequency of 3 kHz, which in its turn modulates the measurement signal. That is why the Jtest can only be used on a DAC which utilizes sharp filters.

On the next page is a Jtest-result which shows a NOS DAC. The needles on both sides of the main signal show the amount of jitter. These are, by current standards, far to high but because they are caused by the lack of sharp filtering the result is negatively influenced. The same DAC shows an error in the time domain of only 20 picoseconds, which truly does not show in the result.

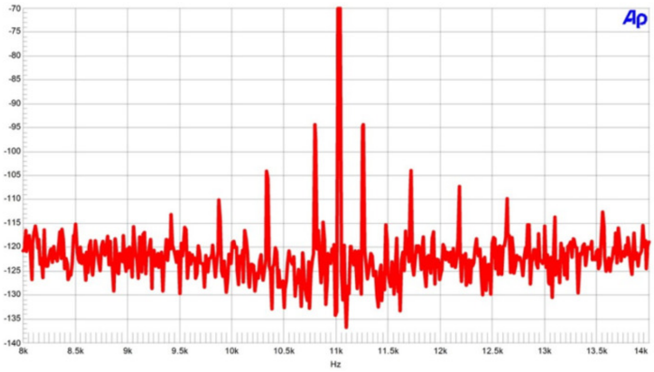

Because the basilar membrane behaves as a sharp filter for our hearing, it has become, in a certain way, part of the DAC. The measurements on a NOS DAC are therefore not measured at its logical end-point, but before the filter (our hearing). Herein lies the problem, which exists when comparing the actual time domain results and the way in which the Jtest interprets and shows results.

Conclusion

NOS DACS have been gaining in popularity for the past few years, mostly based on listening reviews. Especially people who regularly experience live music, appear to have a strong preference for this type of DAC. As Kusunoki had mentioned in his article, it is primarily the behavior in the time domain which gives oversampling DACs their "unnatural" quality. This shows in the way that percussion instruments sound too lacklustre and a sort of "excessive detailing", which causes certain instruments to lose their timbre and "warmth". The question whether we should follow our ears or the results of tests remains on the table.

The development of digital audio systems has not reached its zenith yet and we will certainly be confronted with new developments in the future. Certain is, that due to High Definition recordings the need for oversampling and sharp filters has lessened. How to approach the massive variety of CD's, with their low sampling rate of 44.1 kHz remains the question. To oversample or not to Oversample? Not oversampling seems to be the preference of musicians and audio-professionals, despite their "limitations".
Let your ears decide!

Cees Ruijtenberg
Metrum Acoustics
www.metrum-acoustics.nl