


Understanding the *dCS* Ring DAC™

An introduction to D/A conversion and the unique technology
found inside all *dCS* DACs



Synopsis

Digital to Analogue Converters (DACs) are a fundamental part of all digital audio systems. Their design has a profound impact on how we experience music and how an audio system performs.

In this paper, we explore the basics of digital to analogue conversion and the different types of DACs available today, before going on to examine the *dCS* Ring DAC™ – the unique digital to analogue conversion technology found inside all *dCS* DACs.

We'll explore the Ring DAC's design, and demonstrate how it differs to conventional Ladder DACs, while also showing how the Ring DAC aims to resolve issues that can lead to distortion, resulting in a class-leading performance that reveals the finest of musical details and resolves all aspects of sound.

Digital audio: the basics

To understand how the Ring DAC works, it helps to understand the basics of digital audio and the methods used to capture and store sound.

Sound is analogue, and it's created when varying amounts of pressure cause air particles to vibrate and bump into one another – a process that produces longitudinal waves. This is much like what would happen if you asked two people to stretch out a slinky between them and had one person push the slinky forward. Their push would cause a 'ripple' to pass through the slinky, pushing each coil forward and compressing it into the next. Each time a new coil is pushed forward, the previous one would retract back, or 'rarefy'. This wave of compression would move through the slinky until it reached the other end.


This same process happens with sound. When a person speaks, their vocal cords excite and push the air surrounding them back and forth, and this creates longitudinal waves in the air. When these longitudinal waves reach an endpoint – the human ear – the changes in air pressure are translated into electrical signals that the brain perceives as sound.

The whole purpose of recording music is to take these variations in air pressure and store them in such a way that they can later be reproduced by a transducer such as loudspeakers or headphones, to enable a listener to hear the original audio event as it happened.

Today, this is usually done using one or more microphones and an analogue to digital converter (ADC). A musical performance is captured with microphones, which convert the kinetic energy in air particles into electrical energy (a voltage). An ADC is then used to convert this voltage into a format that can be stored by computers or streamed over the internet. The ADC looks at the voltage that is coming in from the studio equipment, such as a microphone or mixing desk, and determines how high that voltage is, storing it as a group of binary digits (1s and 0s) called a 'word'.

PCM, bit depth and sample rates

The most prevalent format for encoding digital audio is known as PCM, or Pulse Code Modulation. There are two key variables in PCM: the sample rate (how frequently samples are taken) and the bit depth (how many bits, or 1s and 0s, are in each audio sample word).



This diagram shows how an analogue sound wave can be represented with 16-bit 44,100 samples per second PCM encoding

Bit Depth

Bit depth refers to how many bits can be used to describe the absolute position of a sound wave in digital audio recording. A lot of digital audio formats, including CDs, use a bit depth of 16. This means that the ADC can have one of 65,535 possible values at any given point.

It is generally accepted that the human ear can perceive equivalent to 20 bits of dynamic range, which equates to around 140dB (the upper limit being the threshold of pain). CD audio, in its 16-bit format, will achieve around 96dB of dynamic range (the difference between the loudest and quietest volumes that can be sampled). Through the use of dither, the addition of low-level noise to the signal, this dynamic range can be increased beyond 120dB, a significant improvement. Moving to a hi-res format like 24-bit, this dynamic range increases to 144dB – assuming the equipment is actually capable of working in true 24-bit.

It is a common misconception that 24-bit audio simply records louder and quieter sounds than is possible with 16-bit audio, but this is not the case. Instead, the same range of loudest to quietest is measured, but with 24-bit sampling it is done with considerably more steps than with 16-bit. This means the absolute value of the waveform at any given point can be much better represented.

“It is a common misconception that 24-bit audio simply records louder and quieter sounds than is possible with 16-bit audio, but this is not the case.”

Imagine for a moment trying to measure the height of a particular window on a skyscraper. In one case, you can only measure in increments of 1 metre. If the window is 10.7m high, you could round down to 10m, or up to 11m, but in either case there would be a degree of error.

Now imagine the same situation but this time, you are able to use increments of 0.2m. The window is again 10.7m high. You are still unable to measure the exact height of the window, but being able to round to 10.6m or 10.8m brings you much closer to the actual value.

This is, in essence, what happens when increasing the bit depth of digital audio. You are able to measure the absolute value of the waveform with much greater precision, which has the effect of reducing what is known as quantisation noise in the audio.

Quantisation noise is the audible noise which is generated by the error in the measurement. Measuring a 10.7m window as 11m, for example, would result in an error of 0.3m. In digital audio, this kind of error creates negative audible effects.

When working with hi-res audio, each additional bit which is added to the bit-depth of a signal halves the quantisation error, quarters the error power, and thus reduces quantisation noise by 6dB.


Sample Rates

If the human ear can only hear up to 20,000Hz, is there any reason to use sample rates higher than 20,000Hz? As it happens, yes. One of the most important aspects of digital audio is the Nyquist Theorem, which specifies that the digital audio samples need to be taken at a minimum of twice the highest frequency one is trying to record in the original analogue audio.

As the upper limit of human hearing is widely accepted as 20,000Hz, digital audio needs to be sampled at at least 40,000Hz to be able to reproduce the full range of human hearing. For reasons that will be discussed later (related to the digital filtering inside a digital to analogue converter), full range recordings are sampled slightly higher than this, with CD audio being sampled at 44,100Hz. The rate at which these samples is taken is referred to as the sample rate, which defines how many samples are used per second.

Running digital audio at higher rates also allows for gentler anti-aliasing filters to be used. Higher sample rates and gentler filtering mean that filters will have less impact on the audio, with fewer effects like pre- and post-ringing impacting sound quality.

These two numbers, the sample rate and the bit depth, are what define PCM audio. The display of a *dCS* product playing back PCM data will show 24/192 when playing back a PCM stream with 192kHz 24-bit data.



This diagram shows how an analogue sound wave can be represented with 24-bit 176,400 samples per second PCM encoding. The sample rate being higher than CD audio above allows for a greater representation across the X axis of this graph, whereas the higher bit-depth allows for the exact amplitude of the wave to be more accurately represented with each sample – the Y axis.

Pulse Density Modulation (DSD)

Unlike PCM audio, where the ADC sampling process takes the absolute value of the analogue voltage coming in to it at any given point, Pulse Density Modulation (PDM) works based on the time between two samples dictating whether the wave is increasing or decreasing in amplitude.

If the samples are closer together, the wave is increasing in amplitude. If they are further apart, the amplitude of the waveform is decreasing. The absolute value of the waveform is not known per se when looking at an individual sample, as it would be with PCM but put together, the samples produce a good representation of the original waveform.

Understanding the *dCS* Ring DAC™

The caveat with this method is that the ‘dynamic resolution’ (the amount of information about the amplitude which is stored in any one sample of the audio) is incredibly low, being 1 bit, so the samples need to be taken at a much higher rate than with PCM audio.

Where PCM typically samples at 44,100 samples a second, DSD works at a minimum of 64 times this rate, around 2,800,000 samples per second. This process of encoding digital audio creates a lot more noise. This is due to both the low bit depth (which at 1 bit creates more quantisation noise) and the higher sample rate (essentially, turning things on and off at a much higher rate creates noise). In order to make the format usable, data has to be noise-shaped to clear the quantisation noise out of the audio band into the ultrasonic region (above 20kHz), where it cannot be heard.

The result is near 24-bit performance in the audio band (0 – 20kHz) and a signal bandwidth that extends beyond 100kHz. The price for this 1-bit approach is a very large amount of noise in the ultrasonic region (20kHz – 1.4MHz), but this is not normally heard as a noticeable background noise. This method of digitally encoding music is what is used in the format Digital Stream Direct (DSD). This format of 1-bit conversion is the basis of Bitstream Sigma-Delta Digital to Analogue Converters.

There are further developments into DSD audio, whereby higher and higher rates are used. The original rate, referred to as DSD/64 or Single Speed DSD, runs at 64x the rate of CD audio. DSD/128 or Double Speed DSD runs at 128x CD audio rates, and so on for DSD/256 and DSD/512.

DSD files, even at the standard DSD/64 rate are large. The data rate is 5644.8 kbps for 2-channel stereo.

“There are two factors to consider: can the [DAC] perfectly reproduce the original amplitude of the wave when it was recorded ... and can it do it at exactly the right time?”

D/A Conversion

DACs – Digital to Analogue Converters – are a crucial part of almost all modern music playback setups, in one form or another. They play a vital role in helping to translate the original musical performance of an artist to a listening experience for the end user. The fundamental concept of a DAC is to translate digital audio – whether it is streamed from Spotify or Tidal, stored on a DAP or played from a NAS – into an analogue voltage which is used to drive a transducer like loudspeakers or headphones.

When making this digital to analogue conversion, there are two factors to consider: can the converter perfectly reproduce the original amplitude of the wave when it was recorded (in other words, can it output the right voltage), and can it do it at exactly the right time? Whether the converter can reproduce the correct voltage comes down to the DAC circuitry itself, and whether it converts the sample at the right time comes down to the system’s clocking.

Digital audio is stored in binary format (1s and 0s) as a series of ‘samples’. As discussed earlier, the number of consecutive binary digits that are used to represent the original sound wave is called the bit depth. 16-bit audio, for example, has 16 consecutive binary digits, all either 1 or 0. A DAC needs to translate this binary number to an analogue voltage, and that voltage is what drives transducers to produce sound. A DAC achieves this using a series of current sources – electronic components that each generate an amount of analogue voltage.

Ladder DACs


One of the most common approaches to D/A conversion is to have one current source always working for one of the digital audio bits exclusively. For example, one current source will always be following what the first bit in the digital audio signal is doing. Another current source will always be following what the second bit in the digital audio signal is doing, and so on for as many current sources as are needed. As the current sources go on, the amount of energy they must generate gets smaller and smaller (it halves for each consecutive current source).

When looking at a diagram of how these components would be laid out, it looks a lot like a ladder, hence the informal name these types of DACs have been given – Ladder DACs. To ensure that the voltage generated by each current source is incrementally smaller the further down the chain they are, resistors need to be used between current sources. The values and layout of these resistors gives name to the two prominent types of Ladder DACs – R-2R DACs and Binary Weighted DACs.

One very important distinction to make early on – the *dCS* DAC (the Ring DAC) is not a Ladder DAC. We'll explain the core differences between the Ring DAC and Ladder DACs in a later section of this paper.


R₂R DACs

R-2R DACs (a subset of Ladder DAC) use one of two resistor values to control the amount of voltage generated by each current source. Resistors of value R are used between each current source section, and resistors of value 2R are used on each current source. If a particular bit in the audio signal goes high (a 1 instead of a 0), the corresponding switch is enabled, and that current source output goes high. The outputs of all current sources are then fed to a summing bus, which provides the overall output of the DAC.



Binary Weighted DACs

In Binary Weighted Ladder DACs, resistors of decreasing values are used to create increasingly small steps in power generated by current sources. If the first resistor has a value of R , the next would be $2R$, then $4R$, then $8R$, $16R$, and so on for as many steps as required. This hierarchy of resistor values is what gives this approach the Binary Weighted name.



Margins of error

The main drawback with both the R-2R and Binary Weighted DAC approaches comes from the fact that resistors, like all electronic components, have an element of error in their values. For example, a gold tolerance resistor guarantees the resistance of the component will be within 5% of its stated value. This means that for the resistors used in a Ladder DAC, the current generated by that section of the DAC could be either lower or higher than needed. The key point here is that a Ladder DAC uses the same current source for a given bit in the audio signal every time, meaning the error is exactly the same every time the bit goes high. Here, the errors in the component values are correlated to the audio signal. This results in an audible linear distortion of the signal, adding in unwanted harmonic components.

The issue with this is the fact that the larger current sources (correlating to the more significant bits in the audio signal) have the same margin of error as the smaller ones. In the case of a 24-bit ladder DAC, a 1% error in the most significant bit (MSB, or the largest current source) would be larger than the entire seventh bit, and 104dB louder than the 24th bit. The MSB needs to be accurate to 0.000006% to allow for 24-bit resolution.

Understanding the *dCS* Ring DAC™

One further issue Ladder DACs suffer from is Zero Crossing Point Distortion. Given that each current source has a potential correlated error associated with it, what happens when, for example, in a 16-bit DAC, we go from reproducing an amplitude of 32,767 to 32,768? The DAC changes from having the first (most significant) bit low and the following 15 bits high, to having the first bit high and the following 15 bits low. This is called the Zero Crossing Point. The size of the errors associated with each current source / bit here – specifically the fact that the sum of the 15 errors with 32,767 and the one error with 32,768 – are both very large compared to the least significant bit (LSB). This means that the change from 32,767 to 32,768 in the DAC can be much bigger than one LSB. The result of this is linear distortion, which is extremely undesirable.

The solution to the issues posed by the linear distortion of a Ladder DAC is to remove the link between the original signal and the physical resistor value errors associated with specific sample values.

The Ring DAC

How can the issues described previously with Ladder DACs be resolved? What would a DAC designed from the ground up to effectively de-correlate errors in the DAC itself and remove the resulting distortion look like? That is where the *dCS* Ring DAC comes into play.

The Ring DAC is the proprietary DAC technology found inside all *dCS* DACs. On the surface, the Ring DAC may look like a Ladder DAC. There is a latch and a resistor for each current source, and these current sources are fed to a summing bus. The key difference between the Ring DAC and Ladder DACs, however, is that the Ring DAC uses current sources of equal value. This is what is known as a ‘unitary-weighted’ or ‘thermometer coded’ DAC architecture.

Another crucial difference is that the Ring DAC, unlike Ladder DACs, does not use the same current source(s) for the same bit every time. There are 48 current sources within the Ring DAC, all of which produce an equal amount of current. The Field Programmable Gate Array (FPGA) controlled nature of the Ring DAC allows the sources to be turned on and off in such a way that any component value errors are averaged out over time. Firing the same bit three times on the Ring DAC might give one output slightly high, the next slightly low, the next somewhere in the middle, as opposed to outputting the sample slightly high every time, or slightly low every time (as seen in a Ladder DAC).

It takes a considerable amount of signal processing power and know how to optimally operate a thermometer coded DAC, but the benefit with this approach is that it almost entirely removes the linear distortion from the signal (keeping in mind that the highly artificial distortion many DACs produce is very noticeable to humans and has a negative impact on perceived sound quality).

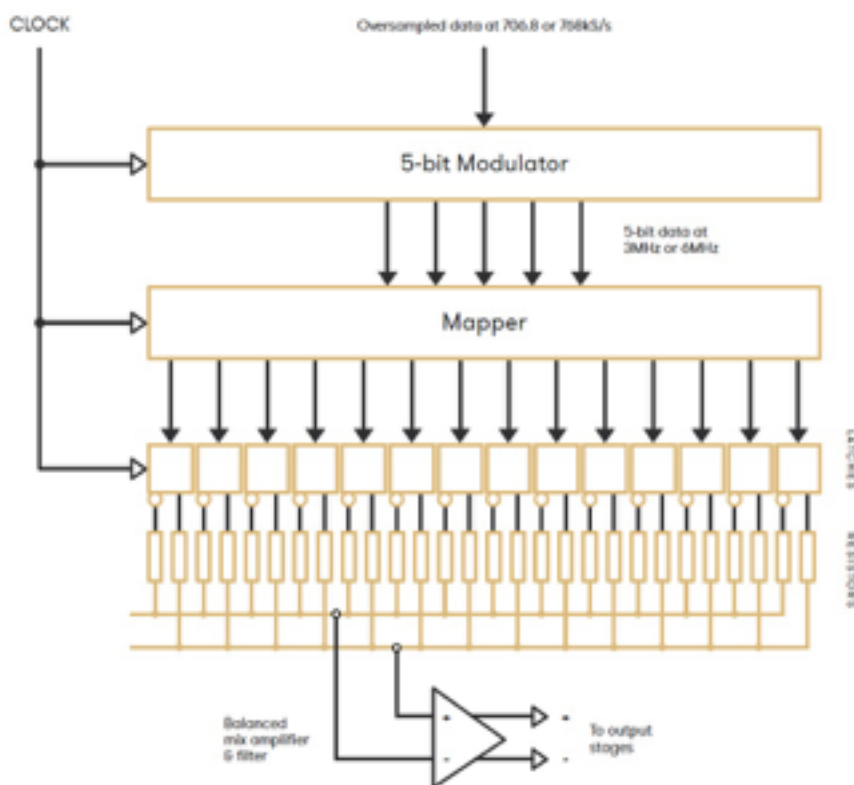
The Ring DAC process may be thought of as decorrelating errors. Background noise (an uncorrelated error – one which is not linked to the audio signal itself) is very prevalent in nature, whereas artificial distortion (a correlated error) is not. This results in the Ring DAC having class-leading distortion performance, particularly at lower signal levels. What this means, in listening terms, is that more fine detail can be resolved and heard in the audio.

The nature of exactly how the Ring DAC decides which current sources need to be turned on or off at any given point to generate the correct signal is dictated by a highly sophisticated set of rules defined in the *dCS* Mapper.

“The Ring DAC process may be thought of as decorrelating errors “


Understanding the *dCS* Ring DAC™

While it might appear to be random, it is the culmination of three decades of continuous work, resulting in a carefully calculated set of patterns used to minimise noise, distortion and crosstalk while primarily keeping the highest degree of linearity by averaging out the contribution of components that fall out of specification over time. Improvements to the Mapper over time have allowed for a lower noise floor to be achieved, while maintaining the signature linear sound associated with the Ring DAC. The Mapper is what allows for the noise created by the Ring DAC to be pushed outside of the audible band of frequencies and then filtered out.



This diagram illustrates the basic layout of the Ring DAC

The Mapper works at 5 bits, so PCM data which arrives at the Ring DAC is first oversampled to 706.8kHz or 768kHz. This is then modulated to 5 bits at a rate between 2.822MHz and 6.144MHz (depending on the unit, settings and content sample rate) and fed into the Mapper, which distributes this signal to the current sources in the DAC.



This diagram illustrates the output of the Modulator within the Ring DAC, modulating incoming digital audio signals to a 5-bit high rate format ready for conversion to analogue

Bitstream Delta-Sigma DACs

One key common factor between the previously discussed DAC architectures is that they all feed the DAC with PCM digital audio. This could be one of several bit depths, but will usually be at least 16-bit. Bitstream Delta-Sigma DACs, however, use only 1 bit.

Trying to reproduce a complex waveform with a single on-or-off signal might seem an odd approach to take. However, the key difference here is that the 1-bit signal does not try to determine the exact amplitude of the wave at any given point (this is how Ladder DACs work, and is the basis of Pulse Code Modulation audio). Instead, the timing from one bit to the next indicates whether the waveform is increasing or decreasing in amplitude, and by how much.

The useful aspect of this approach to D/A conversion is that it removes all of the errors caused by component value errors (resistor values) in the current sources as they are all self-referential: they are either on or off. The problem, however, is that a large amount of quantisation noise is generated that must be noise-shaped to fall outside of the audible frequency range. This requires extremely high sample rates, and the faster the source is turned on and off the more noise is generated. In addition, because the only known quantity about the bitstream is how long it is on or off for, jitter becomes an issue. There is a trade-off between how fast a system can be run and how much jitter and noise is introduced as a result.

Given these factors, one-bit systems are fairly uncommon today. Even DSD bitstream systems decimate the signal into a multi-bit format for DSP processing such as mixing and EQ.


Non-Oversampling / Non-Filtering DACs

There are some audio manufacturers who believe that the use of oversampling and the subsequent filtering has a net negative impact on the audio quality produced by a DAC. Their argument is that using a DAC which does not carry out the oversampling, or in some cases even the filtering process, produces a higher quality. This is supposedly since filters can produce some undesirable effects, such as pre- or post-ringing depending on the filter type, or phase issues introduced by the filter.

As a result, some DAC designs eschew filters altogether. The challenge this approach poses, however, is that the artificial Nyquist images created during the D/A conversion are still present in the audio spectrum. This becomes a real issue at lower sample rates (for example, with CD audio). In the below example, a 20kHz signal is reconstructed with its Nyquist image at 24.1kHz. The two frequencies 'beat' together, creating a waveform that looks similar to the image below.

As shown in the next graph, while the 24.1kHz tone is not audible in itself, a tone of 4.1kHz will be present if the amplifier and loudspeakers do not have good intermodulation performance. If the intermodulation was not present here, the 20kHz tone would be at the maximum amplitude at all times. As can be seen though, its amplitude is being modulated by the 24.1kHz tone. This is audible, and when compared to the performance that can be achieved when using correct filtering, is a questionable trade-off.

Understanding the *dCS* Ring DAC™



In the case of a Non-Oversampling DAC which does use a filter, this will likely have to take the shape of a 22.05kHz analogue filter. The lack of oversampling means that it isn't possible to employ digital filtering effectively (the digital filter would be working in the same space as the analogue filter), so the same sharp analogue filter must be used for all sample rates. As this filter is fixed, it's pass band cannot be raised when playing, for example, 192kHz content (as would be possible with an Oversampling DAC).

Most Non-Oversampling DACs or Non-Filtering DACs will almost invariably produce very high levels of distortion to the audio signal, degrading the original signal and irreversibly removing elements of the audio (effectively throwing away musical details). The distortion is often simply voiced to make it sound pleasing to the human ear (the ear has a fondness for certain types of distortion, hence distortion being a feature on guitar amplifiers), but the addition of this distortion in the D/A conversion stage irreversibly alters the musical performance. Most manufacturers would agree that a DAC should act as a transparent link between digital and analogue, and this simply is not the case with a Non-Oversampling or Non-Filtering DAC.

DAC Chips

Many high-end audio systems do not use a bespoke DAC design. Instead, they use a third-party chip to carry out D/A conversion. Partner this chip with a power supply and an output stage, and you have a usable product. However, whilst there are advantages to this approach (it's much quicker and more convenient to create a system using pre-assembled parts, rather than building a bespoke D/A conversion system from the ground up), there are some drawbacks, which it helps to be aware of when comparing DACs and chipsets versus proprietary designs.

In many cases, third-party chips come with a fixed set of anti-imaging filters that cannot be disabled. For audio manufacturers, this removes the ability to add bespoke filters to a system or select the filter which is most appropriate for a particular type of audio. This lack of flexibility can lead to undesirable performance in several areas. For this reason, chips with fixed anti-imaging filters are best avoided.

Using an off-the-shelf chip also presents a risk of obsolescence. If a chip's design and capabilities are fixed, then the performance of a system cannot be updated as new platforms and formats emerge. The solution to this problem is to use, as *dCS* do, an FPGA-based platform. As FPGAs can be reprogrammed and updated remotely, new features, functions and enhancements can be added over time via software updates, increasing the lifespan of a product and ensuring it remains at the forefront in terms of both features and performance.

dCS Insights & Innovation

Understanding the *dCS* Ring DAC™

dcsaudio.com | [@dCSonlythemusic](https://twitter.com/dCSonlythemusic) | info@dcsaudio.com

Data Conversion Systems Ltd.
Unit 1, Buckingway Business Park, Anderson Road,
Swavesey, Cambridge, CB24 4AE, UK



dCS
ONLY THE MUSIC