



EXOGAL Computational DAC and PowerDAC Technology: The Resounding Difference

Jeff Haagenstad
4/24/2019

STATE OF THE ART, PRE-2013


Most people with an interest in audio have heard of Digital to Analog Converters (DACs) and Amplifiers. They've heard of Ring DACs, Ladder DACs, R2R DACs, and Delta/Sigma DACs. They've heard of Class A, Class AB and Class D amplifiers. (Spoiler Alert: We don't use any of those!) They may not know the exact technical differences between them, the relative advantages and disadvantages of each of them, or the reason that manufacturers make the design choices that they make. But there is a degree of common knowledge that creates a fundamental basis for everyone who participates in audio as a business or a hobby.

This paper is purely concerned with digital technologies, and thus analog technology will only be brought into the discussion where appropriate.


This paper also attempts to make the concepts simple enough for someone who is not an Engineer, a Mathematician or a Physicist to understand. Please forgive us if we take a few liberties as we attempt to simplify the discussion!

A DAC's essential purpose is to take music that has been previously digitized (quantized) from the original analog music and return it to the analog state to allow your ears to correctly perceive the music as it was originally recorded.

Let's say you start the Quantization process with this original analog waveform:




and digitize it into this digital representation:



Under accepted sampling theory, this digital representation contains all the necessary information to *recreate* the original audio waveform.

Now obviously, this is not a great representation of the original waveform and if you simply played it back, as is, it would sound terrible for two reasons:

1. All of those sharp edges. The human ear hates sharp transitions. The ear would hear those sudden jumps between levels as a harshness. A person hearing it might also experience it as physical pain.
2. Amplifiers, both analog and digital, are not good at rendering sudden “vertical” transitions from one voltage to another. The electronic components are subject to inertial forces that cause the signal to go above and below the intended level. This is called overshoot and undershoot. The combination creates a phenomenon known as ringing. You can see it in this drawing:



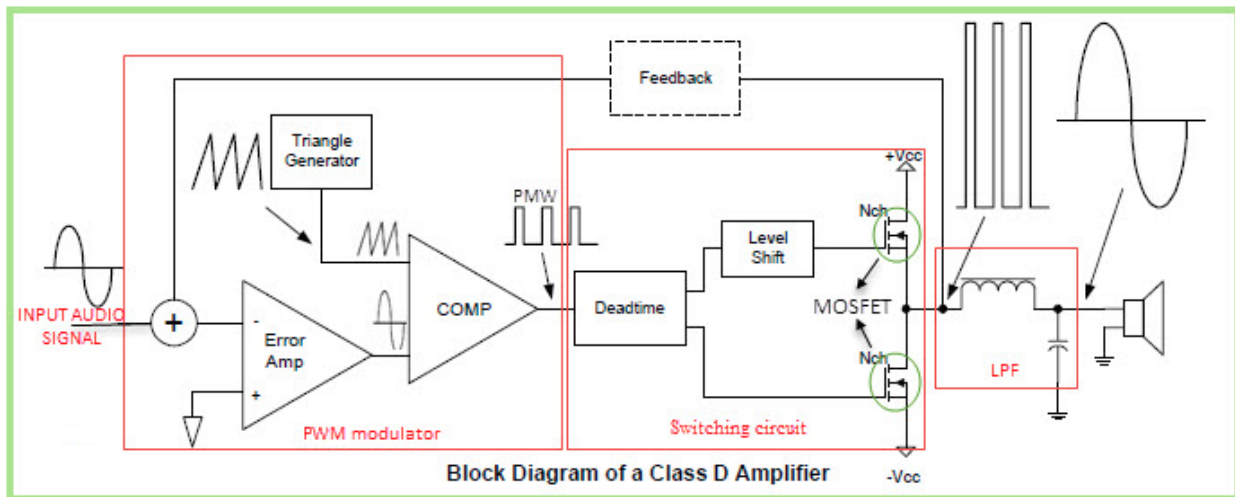
In the real world, electronic components simply cannot create instantaneous transitions, and so it's virtually impossible to perfectly render the square edges in these signals.

Traditional DAC and amplifier technologies attempt to smooth these harsh transitions using time delays and filtering techniques which are intended to put some "slope" into those vertical transitions which eliminates the harsh transitions and reduces the overshoot and undershoot and thus the ringing effects.

On the plus side, these techniques do go a long way towards cleaning up the output signal and making them more pleasing to your ears. On the minus side, these techniques can also distort, muffle, smear or eliminate high-frequency transient sounds. (Think of an instrument like the Triangle, where virtually all of the tones in the notes of a triangle are high-frequency transients.) Many of these high-frequency components are lost or altered in unnatural ways, and the recreated sound doesn't sound like a perfect recreation of the original.

No matter the DAC technology used by most music DACs on the market, they emphasize *smoothing* and *filtering*, rather than *actually* attempting to recreate that original waveform.

So those are the DAC effects. As if that isn't bad enough, digital amplifiers cause their own set of problems. Here is a block diagram of a typical Class D amplifier:



The typical Class D amplifier does what any amplifier does: It takes an input signal and increases the amplitude of that signal so that it is at a useful level. In the case of an audio system, that tends to mean driving a pair of speakers or a pair of headphones.

Unfortunately, there are two main problems with Class D amplifiers (lots of people will say there are *many* problems, but we're only concerned with the 2 largest!) These are:

1. Gain Stage Noise
2. Feedback Noise



Gain Stage Noise is created in several ways. The modulator may not properly synchronize with the incoming signal, creating synchronization errors. The Comparator may have an inaccurate reference or may be inefficient, leading to amplitude errors. The Switching Circuit may be too slow, which creates delay and smear and can wipe out transient signals.

Feedback Noise can be generated when the feedback circuit is inefficient or improperly designed, or the summing circuit may not properly mix the signals which allows the comparator to stay balanced and function properly.

Many types of noise can be created in these two electronics sections which can lead to sounds that are perceived as a hiss, a hum, “muddiness” (where the sound is muffled) or other unpleasant colorations. A common complaint of Class D amplifiers that are very well-designed to eliminate all of these noise sources is that they sound “clinical”, which means they sound too clean, unlike how the sound would naturally occur. Sometimes, you just can’t win!




THE EXOGAL COMPUTATIONAL APPROACH

EXOGAL was founded by people who had been involved with digital audio, almost since its inception. Their experience at Wadia Digital and other companies had shown them that while that 2013 “State of the Art” was fairly good, there was a lot of room for improvement.

Primarily, quantized digital data was intended to allow the *re*-creation of the original analog waveform, but every digital audio company simply took that digital data, created a rough waveform, smoothed it, and sent it off to an amplifier. While, ironically this works amazingly well, there are fundamental limits to the audio quality which just cannot be overcome in this approach. Basically, every existing DAC architecture had the same fundamental flaws. We realized that it was time to take a different approach, which we would eventually begin calling “Computational”. But we will get to that.

What EXOGAL realized in 2013 was that analog sound was digitized using a process called “successive approximation” where it is sampled fast enough to catch the tiniest detail the human ear could hear. (This is normally where the author of a paper would go on about the Nyquist Theorem and other sampling theorems. And lots of readers would get stuck here and argue endlessly about the effectiveness of each sampling theorem, etc. and then totally miss the point of this paper. God Bless the Internet!) We’re not going to get into any of that. The recording industry has established these standards, and most of them are a reality we just have to live with. Moving on...

Once EXOGAL realized that we and everyone else had previously missed the point entirely, we took a look at what we could do about it. This led us to the realization that the only *perfect* approach was to reverse the successive approximation process. However, regardless of what anyone tries to tell you: Once information is lost, you cannot recover it. (Remember this point! We will come back to it many times!) Many DACs attempt to interpolate the lost data by looking at the amplitude difference between one sample and another and breaking that gap into a series of much smaller steps:




In this example, the gap between the three blue levels is filled in by dividing it into eight equal steps, increasing by $1/8^{\text{th}}$ the amplitude with each step.

This approach does simplify the filtering somewhat and makes the steps smaller which are less uncomfortable to your ears, but it doesn't actually help you re-create that original signal. And interpolating between two points is not *recovering* previously-lost information!

No, to try to reverse the successive approximation quantization process, it was going to require something fairly complicated: Calculus. And to further complicate things, a dedicated processor on which to run the algorithms.

As we said, EXOGAL realized that the data didn't just represent the original music as-is. It provided the information to re-create it. Or rather, in our case, a very close facsimile of the original data. If you simply play "Connect the Dots" with the samples, then you may leave out a lot of information. For example, in 2013, the state-of-the-art created output waveforms that looked like this:




When you remove the green "data" and just connect the dots that those data points represent, you can see the gaps between the resultant waveform and the original black analog waveform. And these gaps exist even if you interpolate between the samples because the yellow line is, in essence, an *infinite* interpolation between samples. Those gaps are errors that can be heard by the human ear. All of that filtering and interpolation of the old State of the Art is really just a "sanding-off" of the sharp edges created when you play Connect the Dots!

Our algorithms start with the exact same digitized data, but we essentially treat it as a series of points on a mathematical graph. The real difference, and it's significant, is that we apply Calculus-based, Digital Signal Processing (DSP) algorithms to the data and to dramatically oversimplify it,

we apply Calculus functions on the data to reconstruct the waveform. In Calculus, it's all about interpreting slopes and rates of change. Remember the old problem: "If you throw a ball that weighs X , at an angle that measures Y , with force Z , how high will the ball go, and where will it land?" That math problem takes straight-line forces and calculates curves to obtain a result.


So, where an older technology DAC reconstructs the waveform by connecting the dots with *straight* lines, we connect them with *curves*. Each segment from sample to sample has a curve of its own, and the steepness, depth and sharpness of the curve depends not only on what's happening between two adjacent samples, it also depends on what's happening in the adjacent samples. When the curve segments are connected, they re-form a waveform that is a closer *representation* of the original signal.



Again, the black line is the original analog waveform, the blue line is our newly-created waveform. As you can see, they are not a perfect match, but they are much closer to the original. The most important factor is that now the resulting waveform is a continuous series of curves, with no sharp edges, no smearing, no noise, no ringing, etc. We can then apply much less intrusive, less invasive and less destructive filtering to this waveform which keeps it largely untouched and unaltered.



If you compare the unfiltered waveforms reconstructed by older DAC technologies (yellow) and our computational DAC (blue), you can see how different they are, and it really highlights the losses that occur during the peaks and valleys, especially if there are sharp transients in the original sound:



We want to draw your attention to the curve at the top of the right-most peak in the example. We just told you that data, once lost, *cannot* be recreated. How do we put that curve back in there? This is where the magic of the math comes in. Remember, we know the slope and rate of change between *each* sample. When we know this for *every* gap between samples (we call this the interstitial space) we can look at how the curves that fill the gaps that come before and after the gap we're currently processing will need to bend. And we can bend that interstitial curve appropriately, so the segments fit together seamlessly. When you do this over the entire signal, the radius of that peak becomes apparent in the math.

Since our reconstructed waveform is dramatically closer to the original, we can apply filtering with a much lighter touch which has less impact on the overall signal. Remember that a filter affects *all* the data, even though its *intent* is to only affect the bad parts of the signal – the jagged or ragged bits.



THE COMET DAC

The Comet, as a product, is a DAC / Preamp much like any other DAC / Preamp you can buy. It has most of the same inputs, outputs and features as any other DAC / Preamp on the market. Except that instead of relying on any of the pre-made chip solutions, or FPGA implementations of cookbook DAC algorithms mentioned in the first paragraph of this paper, we use Computational technology based on our custom-designed processor and our calculus-derived algorithms. Thus, the Comet's DAC section is unlike any other manufacturer's DAC.

The Comet is designed to render digital audio specifically using the approach we just described in the previous section (which we won't belabor again!). The Comet's audio processor is based on a multi-core, custom Digital Signal Processor (DSP) built on an Altera Field Programmable Gate Array (FPGA). The DSP takes the digital data stream from the selected input, combines other factors like digital filter parameters, desired volume settings, etc. and the DSP processes the data stream through our algorithms while incorporating those parameters directly in the computation. Notice the specific use of the words "data stream"! The Comet treats the *digital data* as a *series of data points*, rather than as a simple audio stream where the stream itself is a rough approximation of the desired audio output signal.

This difference is critically important. For the audio stream approach to work properly, timing and clocking must be absolutely precise, and jitter (or error in the sampling of the audio stream) must be carefully eliminated or avoided. The DAC must ensure that *every* sample is gathered at precisely the same point bit to bit throughout the audio track. If not, sampling errors will creep into the stream, and they will appear as audible artifacts in the output stream. However, since the Comet treats the digital data as data, timing impacts are minimized. As you will see in some of the Comet's reviews where measurements were performed, jitter is very, very low. This completely negates the "need" for Galvanic Isolation, external re-clockers, clock regenerators, etc.

Also, by using "Volume" as a parameter in the calculation, there are no losses of dynamic range caused by the Comet's volume control. Many digital volume controls essentially work by "throwing away" the least significant bits of the audio data and shifting the higher-order bits down to fill the gap. This has the effect of reducing amplitude, but it also decreases dynamic range and causes a loss of data which ultimately lowers the audio quality. The Comet's approach dynamically rescales the data points in the stream, minimizing the bits that are lost on the low end, while avoiding clipping on the high end. This creates a much cleaner, more realistic sound throughout the volume range.

The output of the Comet's audio processor is the recalculated analog wave that was described in the previous section. This analog wave fully incorporates the listener's chosen volume settings, and all that remains is to send this wave to the Comet's analog section where the output level is matched to a level and impedance expected by an audio power amplifier (analog or digital.)



The Comet performs this by using the last stage of the Texas Instruments DAC chips as a driver for the Op Amps that ultimately send the signals to the analog output connectors. Note that we do not use the DAC portion of the DAC chips – only the output stage!



THE ION PowerDAC

First of all, what is a “PowerDAC”?

This term has been floating around the audio world for decades and used by many companies. In general, over the last twenty years or so, the term “PowerDAC” has been inaccurately used to describe some digital amplifiers that use Pulse Width Modulation (PWM) as their basis for amplification.

However, when the term was originally coined (and we believe it was coined by Wadia who tried to deliver the first, true PowerDAC, failed at it, and then pivoted to PWM technology) it was intended to describe *the control of analog power in the digital domain*. Unfortunately, digital electronic components that were robust enough and fast enough to switch large amounts of power very quickly didn't really exist at the time. Eventually, engineers gave it up as impossible and switched over to PWM which is now the basis of many Class D and other digital amplifier technologies. Fundamentally, PWM works pretty well, however PWM also gives a digital amplifier the acoustic characteristic that many people, especially lovers of vinyl records and tube amplifiers, will describe as “clinical”. In this case, clinical means: too clear, no warmth, no subtlety, “brute force”, “in your face”, etc. Fans of digital amplifiers will counter this complaint by saying that the sound is just simply clean and crisp (...duh - like it's supposed to be!) In truth, both arguments have some merit. But in fact, digital amplifiers, while generally delivering a very clean sound, also deliver artifacts caused by digital noise sources and in an analog world with analog ears, this digital noise sounds unnatural.

To most people, however, these digital noise sources aren't readily apparent because their audio systems are most likely built from lower quality components with a lot of good design corners having been cut. These tradeoffs are masked by the lower source quality of music sources like MP3 recordings. It's hard to hear the deficiencies if A) all you've heard, and B) all you're likely to hear is a low-fidelity recording, C) played through a low-quality amplifier and cheap speakers and D) it's a higher priority for you to have a wide variety of music as opposed to a smaller library with higher fidelity. It's a tradeoff of quality for convenience. Well-informed audio consumers know that this compromise need not be made!

But that's not why we're here. So, to repeat the statement in the prior paragraph, a true Power DAC is intended to *control analog power in the digital domain*. While other manufacturers, including Wadia at this point, continued to pursue other digital technologies, EXOGAL's founders determined that electronic component technology had finally advanced to the point where analog power at high levels now *could* be quickly and reliably controlled by digital circuits and that a true PowerDAC was now fully possible.




So, while the Comet already re-calculates an analog wave that closely approximates¹ the original source's sound, that wave must be turned into power to drive a pair of speakers.

Already knowing the limitations of PWM, we still wanted to quickly and precisely control the analog output while avoiding the pitfalls of PWM and also, as we soon learned, the forces of Physics.

We were correct that the electronics technology had advanced sufficiently. We could indeed drive a high-current power supply from ground (zero volts) to maximum voltage in 10 microseconds (μs). That's the equivalent of 100 kilohertz (100 kHz) or almost *five times* the accepted limit to human hearing, which is around 20 kHz. Why that's fantastic! All we have to do is set the output device to the proper voltage and Whoosh! Up it goes! Well, as we soon discovered, "up it goes" becomes a literal term when you are driving a high voltage at high currents in a nearly-instantaneous manner. There's this little thing called "Electromotive Force" or EMF, which is essentially inertial momentum in an electrical signal.

If you recall a couple of sections ago, we were describing the conditions of overshoot and undershoot. We used this drawing:



That ringing you see at the leading and trailing edge of the square wave is essentially caused by overshoot and undershoot. In Physics terms, the EMF of the signal wants to keep going in the same direction it's already heading. When this is done at high power, the ringing will exceed the ratings of the output device and the results can be spectacular. In this case, a one-inch long plasma jet, out of both the top and the bottom of the device.

While very exciting and festive, it wasn't terribly useful. It quickly became clear that raw Physics wasn't our friend, and everyone knows that you can try to fight Physics, but the results are usually not good. So, it became clear that we had to get friendly with Physics and make it help us instead of trying to hurt us.


¹ Important Legal Disclaimer: EXOGAL makes no claim to *perfectly re-create* the original analog signal! We are very specific that we are *re-calculating* an *approximation* of the original wave, and that the *differences* between the two are so tiny that they are, for all practical purposes, inaudible!

The great thing about the Comet's algorithm design was that we were already calculating all the twists, turns, peaks and valleys of the audio signal, so we had *a priori* (advance) knowledge of where the signal needed to be, relative to where it was. Using this information, we created tracking algorithms that we could use to create counter forces that we could exert on the rapidly rising and falling signals and negate the EMF that was creating these large spikes.


We gave these counter forces very catchy and clever names: the Upper Counter Force (UCF), and the Lower Counter Force (LCF). While the Comet is busy calculating the basic analog waveform, The Ion is also taking that data and calculating 1) how much power will need to be driven to the speaker terminals, 2) how far and fast the power change is from moment to moment, and 3) calculating precisely how much UCF and LCF is required to damp and constrain the power signal's EMF and keep it from going too high or too low. Essentially, the UCF and LCF are the guardrails that keep the output signal where it belongs: safely in the lane.



So, to revisit our square wave with ringing, the UCF and LCF create damping forces that act on the main output signal and create the guardrails to keep it where it is supposed to be:



In real time, the UCF and the LCF track the output voltage and keep the output signal where it belongs:



(UCF in Red, LCF in Blue. Gap exaggerated!)

Now that we have the ability to switch power at very high speeds, and also constrain the output to nearly exactly where it belongs, it becomes possible to take our highly-precise analog wave and drive power to the speakers with incredibly high transient response. This fast-transient response allows the music's detail to fully project out of the speaker, almost completely without the muting and smearing effects of filters, and without the digital noise sources that make digital music sound unnatural.

And we were feeling pretty proud of ourselves that we'd made this work. However, shortly thereafter, we made another discovery: not all speaker drivers are capable of, or willing to be driven at really high frequencies or at really high volume. And we discovered this by listening to some music at a fairly high volume, when suddenly two puffs of smoke emanated from the ribbon tweeters on our speakers. And literally before the smoke cleared, there were no high frequencies to be heard at all. We consulted with Sandy Gross of GoldenEar Technology who helped us to understand that some ribbon tweeters don't like frequencies above 30 kHz, especially at high power. Turns out many regular dome tweeters don't like it much, either! OK, so we can simply roll off frequencies above that and we're good. Except that solution doesn't get at the real issue: It's not just frequency that can fry a driver. Power can do it as well.

We were already calculating the frequencies and output power of our output signal and creating counterforces to keep everything where it belonged. And we realized that it was a relatively simple matter to monitor that output signal and keep track of how the speaker was likely reacting to the signal. And by doing this measurement, we realized that we could sometimes identify situations



where the speaker may be drawing too much power because of a non-linear impedance curve or some other factor. We then developed more algorithms that let us predict some conditions that may be damaging to the speaker so that we could shut down the PowerDAC, in many cases before the speaker or the speaker's crossover was damaged. The protection circuit of most amplifiers are designed to just protect the amplifier. In the Ion, we can now prevent some, *but not all*² situations where the listener would damage their expensive speakers.

But that's just a pretty good start. In 2019 we added a circuit called HyperDrive to the Ion which increases the regulation granularity of our internally-derived high voltage power supply. HyperDrive allows us to essentially drive "closer to the guardrails" which gives us a performance boost in terms of output power, improved frequency response across the dynamic range of the PowerDAC, and better protection from anomalous conditions such as bad Mains power, etc.

We continue to research and experiment with PowerDAC technology, and the future will see higher output power and even greater acoustic performance.

Remember: we're out here, blowing things up, so you don't have to!

² It is not scientifically possible to detect *all* of these circumstances. While we know that we can detect *some* of them, we have no idea if this comprises 1% or 99% of the possible scenarios. Bottom line: You still need to be careful to use properly rated speakers with our PowerDAC!



CONCLUSION

Hopefully now you have a better understanding of EXOGAL's core technologies, why they are different than anything else on the market or in the industry, and why we believe they are a better solution. Most importantly, why they give a listener a much truer listening experience.

We believe this holistic, bottom-up design approach achieves a much better result than simply tweaking an off-the-shelf chip reference design. This approach was fine five years ago, but the needle has moved – and we moved it! We design our products to fully *solve the problem*.

We also hope that you'll understand why adding corrective measures like jitter reduction devices, or external re-clockers, filters and the like aren't necessary and will likely cause more degradation than improvement. Similarly, you should now understand why we don't support third-party streaming technologies claiming to restore a lower-resolution track to a much higher resolution. As we've repeatedly said: *information, once lost, cannot be re-created*. These "psycho-acoustic" tricks may sound different, but that does not imply they are better, or that everyone perceives them in the same way. We believe that even on a lower-resolution track, we can re-create a better analog waveform than any of these "post-processed" techniques can hope to create. We suggest you try it for yourself!

Most importantly, our products aren't just engineered for innovation and performance, we also make them compact, attractive and affordable (and hopefully you think they're attractive and affordable, too! Of course, we're engineers and scientists. What do we know about style and affordability?)

They are simple to connect, simple to use and they fit into a modern lifestyle. They are designed and implemented in a way that allows us to continue to innovate and to offer affordable upgrades to our customers who took a chance and made an early commitment to our products and our technology.

We believe that our products based on these core technologies truly represent the Future of Audio in a way that no other manufacturer can match.

FOR MORE INFORMATION

EXOGAL, LLC
4657 Aspen Ridge Circle
Eagan, MN 55122 USA

exogal.com
sales@exogal.com
651-964-0698

Copyright Information © 2019, EXOGAL, LLC, All Rights Reserved