Ayre Acoustics QA-9 USB A/D converter

Since its founding in 1993, Colorado-based Ayre Acoustics has made its name with amplifiers and preamplifiers based on truly balanced, solid-state circuitry that didn't use the ubiquitous panacea of loop negative feedback to produce linear behavior. Their first digital product was the D-1x DVD player, reviewed for *Stereophile* by Paul Bolin in February 2003, which offered unusually good video performance. The D-1x was followed by the C-5xe and DX-5 universal players, respectively reviewed by Wes Phillips (July 2005) and Michael Fremer (December 2010). But the most intriguing digital product to come from Ayre was the QB-9 digital processor. Reviewed by WP in October 2009, the QB-9 has just one input, USB, and uses Gordon Rankin's proprietary Streamlength code to give asynchronous operation, which in theory offers the best jitter suppression. "The QB-9 *isn't* a computer peripheral," said Ayre's marketing manager at that time, Steve Silberman. "It makes computers real high-end music sources"—a statement with which WP agreed.

Then, at the 2011 Rocky Mountain Audio Fest, I was given a dem of what I at first assumed was a QB-9—until I saw the bar-graph level meters to either side of the sample-rate display. This was the QA-9, a high-performance analog/digital converter, housed in the same chassis as the QB-9 and intended to allow audiophiles to make rips of their LPs of the highest quality possible.

The QA-9 operates at sample rates up to 192kHz, outputting 24-bit data either by a USB 2.0 connection or by AES/EBU. (The two aren't operational at the same time.) The basic version costs $3950; a Pro version, which includes DSD and Word Clock outputs on transformer-coupled BNC jacks, costs $4750. Given my extensive experience of both domestic and professional A/D converters, which has convinced me that the most critical process in digital recording is the initial analog/digital conversion—nothing downstream can put right whatever was done wrong in that conversion (footnote 1)—I asked for a production sample of the QA-9.

Design Philosophy

Via e-mail, I asked Ayre's founder and designer, Charlie Hansen, why a manufacturer exclusively known for its domestic high-end audio products had ventured into a field dominated by pro-audio companies.

"If we build, say, a great preamplifier, then one audiophile will be able to bliss out on the music. On the other hand, if we build a great A/D converter, then literally millions of people could benefit from the improved sound. Several professionals own Ayre equipment, usually for their home systems. Somewhat typical but most notable is Rick Rubin [record producer and founder of American Recordings and Def Jam Records]. He says that Ayre makes the best-sounding digital playback gear, and he has been after us for years to build equally good-sounding recording gear.

"To my knowledge, nobody has ever built an ADC with fully discrete, fully balanced, zero-feedback analog circuitry. In my experience, the single most important factor in getting good sound in digital audio is with the *analog* circuitry, so that was reason enough to start the project. Of course, we at Ayre are never content to rest on our laurels, so we also incorporated a slew of other innovations."

The heart of the QA-9 is the AT1201, a two-channel A/D converter chip that is new to me, from a company also new to me: Arda Technologies. The chip's datasheet describes it as an "advanced multi-bit sigma-delta" converter that will operate up to a sample rate of 384kHz with an astounding dynamic range of 124dB.

"Arda does consulting work for big contractors. They are audiophiles, and in their spare time designed and built a new A/D chip that clearly outperforms anything from the competition. Specifically, the out-of-band noise is ridiculously low, and this gave us a lot of freedom to 'think different' (as Apple ungrammatically said a couple of decades ago)."

What did Hansen mean by "think different"?

"Ever since DSD was announced, it has received uniformly positive comments on the sound quality. . . . [T]o be honest, we were basically baffled by Sony's original marketing material. We couldn't figure out how the system worked, let alone why it sounded good. But over the years, not only did more and more information leak out, but our understanding of digital audio at Ayre grew rapidly.

"In short, the reason that DSD sounds as good as it does is because there is *no* filtering done on the record side. The playback side requires a filter (per the Scarlet Book specs), but compared to the brick-wall filters used in typical PCM products, this is a much, *much* gentler affair. When we developed the QB-9 we spent nearly four months performing listening tests on digital filters. It was clear that both the anti-aliasing (record) and reconstruction (playback) filters had a very strong influence on the sound of digital audio (remember how the 'non-oversampling' DACs were quite popular for a few years?), and so we wanted to really understand what was going on. . . .

"The relatively recent advent of using the personal computer as a way to store and play back music files has changed the game entirely. Now we can buy, store, and play back music at quad-sample rates [176.4 and 192kHz] (or even higher), if so desired. But the problem is that digital audio equipment is still designed by digital audio engineers, traditional engineers, who, if they are using a 192kHz sampling rate, are going to make their equipment with flat frequency response to 96kHz and then just put a brick-wall filter on it. Fortunately, I am not a traditional digital audio engineer! I therefore have the freedom to ask 'Why?,' and I do that a lot!

"So for the QA-9 in Listen mode, we decided that the goal was to make the converter operating at the quad-sample rate to perform more or less like a perfect 30ips analog tape machine. The frequency response is down about –3dB at 50kHz, but it goes down to around 1Hz with no 'head bumps' to worry about. There is *zero* wow and flutter, and the distortion and noise are about an order of magnitude better (*ie*, 20dB) than the best analog tape machines.

"Yes, the anti-aliasing filter (in Listen mode) is down only about 20dB at the aliasing point of 96kHz, but is that *really* a problem? Do any recordings actually have any meaningful amount of musical energy up at 172kHz, where there would be aliasing that would fold down into the audioband?

"Of course not! There aren't many instruments with significant amounts of energy above 90–100kHz. Nor are there many microphones with any significant response this high. Or mike preamps. Or mixing boards. Or whatever.

"So the first thing that we did was to use a completely different type of digital filter at the output of the delta-sigma converter. (Every audio ADC chip made for the last 20 years has been a delta-sigma type, as the competing successive-approximation devices died off long ago.) Instead of using the normal low-pass Finite Impulse Response (FIR) filter to turn the output of the delta-sigma DAC into PCM, we use a moving-average filter. This doesn't just allow for *improved* transient response, but actually *perfect* transient response. And since we are starting with a 256Fs [11.2896MHz] sample rate, we don't have the problems exhibited by 'non-oversampling' DACs that also have 'perfect' transient response. There is no pre-ringing, no post-ringing—no ringing whatsoever. (This is in the Listen mode at both the quad and double sampling rates. It is not possible to use this trick at the single sample rate, where instead we use a more conventional FIR low-pass filter, but of course this is a slow-rolloff design to minimize ringing, and also is a minimum-phase design, so that all of the ringing occurs *after* the transient, with no unnatural 'pre-echo' before the transient occurs.)

Footnote 1: I haven't forgotten Meridian's implementation of Peter Craven's "apodizing filter," which replaces the original A/D converter's acausal ringing at the Nyquist Frequency with post-impulse ringing at a very slightly lower frequency. But there are many other ways for A/D converters to misbehave.

Ayre Acoustics QA-9 USB A/D converter Page 2

"So now we are starting to get somewhere—perfect transient response, combined with zero-feedback, fully discrete, fully analog circuitry."

Hansen had mentioned that the QA-9's response was flat to around 1Hz.

"Another thing that we wanted to avoid that is ubiquitous with ADC chips is the brick-wall *high-pass* filter that is intended to keep DC out of your system. Now, it is good to eliminate DC for a variety of valid reasons, but adding a brick-wall filter will create ringing no matter whether it is a low-pass or a high-pass filter. The Arda ADC chip carries its own similar DC filter, but we chose instead to implement ours from 'scratch' in the Field Programmable Gate Array (FPGA) instead.

"If you tried the conventional approach of using an FIR, you would need several bazillion taps to make a filter that operated that low in frequency. So instead, what people normally do is use an Infinite Impulse Response (IIR) filter, where the output signal is recirculated through the filter to get the desired response. This has the side effect of turning the filter into a minimum-phase type, which would normally be good, but in this case one does not want to introduce phase shift in the bass range. Also, I do not like to use IIRs, as the inevitable errors accumulate as the signal recirculates through the filter. The goal is that the signal is smaller than one LSB by the time the error is equal to one LSB, but it doesn't always work as planned.

"At any rate, we had to come up with a different technique to avoid these problems, so we just did the same thing at the low end that we did at the high end. There is a moving average that samples a few thousand data points and averages them together. Ideally they will all average to zero, but if they do not (*ie*, there is a DC offset), then the module slowly starts adding or subtracting numbers so that the average comes back to zero. By doing this slowly, it avoids any clicks or pops. By using a moving average, it is (once again) transient perfect, with no ringing, pre-ringing, overshoot, or pre-echo. Essentially what we have done is exactly emulate the way an analog servo works and implement it in the digital domain."

What about the Measure filter setting?

"If you are in a situation where you need totally flat frequency response to the highest possible frequencies (*eg*, when recording bats or dolphins), the Measure position allows for that. For those filters, we use the moving average to bring the sample rate down from 256Fs to either 8Fs or 4Fs, and then use an FIR low-pass filter to do the final division by two. While these are minimum-phase, slow-rolloff types, they still will introduce a cycle or two of ringing, while the Listen filters are absolutely perfect in regard to transient response."

Setup

The QA-9 has a single pair of XLR input jacks. I used RCA-to-XLR adapter plugs to connect the single-ended output of the Liberty phono preamp that Michael Fremer reviewed last May. Reducing the input level by seven clicks gave a maximum level of –3dBFS with all the LPs I played using this preamp and my Linn Arkiv B cartridge. Otherwise, setup was as simple as plugging a USB cable into my Mac mini or MacBook Pro and selecting the QA-9 as the default input device. (A couple of times, I did need to restart the computer for it to recognize the QA-9.) A rear-panel DIP switch selects the Measure or Listen filters, turns the Pro version's Word Clock and DSD outputs on or off, and sets the AES/EBU sample rate. When the QA-9 is connected to a host computer via USB, its sample rate can be set with the chosen recording software or with the system control panel.

My review sample was of the Pro version of the QA-9, which includes both Word Clock and DSD outputs and can transmit DSD data to the host computer using the DSD-over-USB 2.0 protocol developed by a team of independent engineers—see "Method for transferring DSD Audio over PCM Frames Version 1.1," by Andreas Koch (Playback Designs), Andy McHarg (dCS), Gordon Rankin (Wavelength), and Michal Jurewicz (Mytech). Charlie Hansen tells me that v.1.2 of the DSD-over-PCM standard for USB is in development, and that Ayre will offer, he hopes, a free update when that standard is finalized.

However, as I do not currently have a DSD-capable DAC—our review sample of the dCS Debussy is awaiting an upgrade—I didn't experiment with the QA-9's DSD capabilities. That will have to wait for a Follow-Up. However, my dCS 904 and Metric Halo A/D converters all successfully locked to the QA-9's Word Clock output, as did the Logitech Transporter DAC. (A word-clock output is essential for pro-audio use, as it enables other converters to be synchronized to the QA-9, to allow multitrack recording. However, as the QA-9 lacks a word-clock input, those other converters can't be QA-9s.)

A Hiccup

I had made only a few recordings with the first sample of the QA-9 when I ran some preliminary measurements. To my dismay, even though all 24 bits in its data output were active and the sound quality of the files it made was far from disappointing, the noise floor was higher than I'd expected (though still below the LPs' groove noise). At 48 and 96kHz, the converter appeared to be operating with just under 18 bits of resolution; at 192kHz, it was actually equivalent to only 16 bits. Concerned that I had a faulty sample, I contacted Ayre about this problem, who asked me to send it back to them to be checked out. (I wasn't giving Ayre special treatment; my policy is to let manufacturers know if we have a review sample that is broken or doesn't appear to be working correctly.)

It turned out that the first 20 QA-9s had suffered from both a firmware problem and what Ayre referred to as "subtle assembly errors." To their credit, Ayre recalled all 20 QA-9s from customers and updated them free of charge. At the end of June, they sent me a sample that was representative of current, corrected production. All of my auditioning comments and the published measurements refer to this revised sample.

Sound Quality

An odd subtitle, given that the QA-9's output is a stream of digital data. The quality of those data can be assessed only by using a D/A converter, which will introduce its own sonic signature. I examined the QA-9's sonic performance in two ways: 1) by ripping LPs to 24-bit/192kHz-sampled AIFF files via USB, using Ayre's recommended Vinyl Studio app ($29.95 for Mac OSX and Windows XP, Vista, 7); and 2) feeding the QA-9's AES/EBU output to the D/A processors I had on hand, using a 192kHz sample rate when possible.

Ayre Acoustics QA-9 USB A/D converter Page 3

My first rips were of these LPs: *Astaire* (English Mercury 9109 702), a 1979 album from the English singer and pianist Peter Skellern in which he imaginatively arranges songs made famous by the elegant dancer for his multitracked voice, piano, rhythm section, and English brass band; "Die Tänzerin," from German singer Ulla Meinecke's *Wenn Schon Nicht für Immer, dann Wenigstens für Ewig* (German RCA 426124); Sly and the Family Stone's *Fresh*, from 1973, which features Doris Day reprising "Que Sera, Sera" (Epic ECD 69039); the first album from British blues band Steamhammer (1969, English CBS 63611), which features on second guitar Martin Quittenton, later to achieve fame as the co-writer of Rod Stewart's hit "Maggie May"; that audiophile classic, the late Radka Toneef's reading of Jimmy Webb's "The Moon Is a Harsh Mistress," from *Fairytales* (Odin LP03); *The King's Singers Sing Flanders & Swann and Noël Coward* (1977, EMI EMC 3196); and a favorite of John Marks, the 1970 recording by the King's College Choir and the Jacques Orchestra, conducted by Sir David Willcocks with narrator John Westbrook, of Ralph Vaughan Williams's *An Oxford Elegy*: a setting of two poems by Matthew Arnold, in purported homage to fellow-composer Gustav Holst (HMV ASD 2487). I list these as examples of recordings I used to love but haven't listened to in decades, and which the QA-9 pushed me into digging from the darker recesses of my LP shelves.

I tried a variety of sample rates with these LP rips: 44.1kHz was very good, but didn't capture the essence of the original LPs' sounds; 96kHz was better; but there was no doubt that with a 192kHz sample rate I could not distinguish between the LP and the digital rip. And believe me, I tried. I A/B'd the two versions until blood came out of my ears and I was heartily sick of this music I hadn't heard for, in some cases, decades. When, in *An Oxford Elegy*, John Westbrook declaimed "Come, let me read the oft-read tale again . . ." for what must have been the tenth time, I felt like screaming "No! *Don't* read it again!"

With the 192kHz rips, the LP's surface noise floated free of the music in a manner similar to how it does with analog playback, making it easier to ignore it. At 44.1kHz, the surface noise was integrated into the music, increasing its annoyance. While 24/192 rips are profligate with hard-disk space—the 22-minute Vaughan Williams is a 1.5GB AIFF file—this was the only way to go with the QA-9.

Comparisons
I have used three two-channel A/D converters to make the recordings that have been released on the Stereophile and other labels over the past 12 years: two dCS 904s, joined in 2003 by a Metric Halo ULN-2. The dCS is the final version of one of the first 24-bit, high-sample-rate converters, the dCS 900, which was introduced in 1996. The 904 takes balanced line-level signals and uses a 5-bit flash converter embedded in a sigma-delta loop; it can output DSD data via an SDIF-2 connection, as well as 24-bit LPCM up to 192kHz over single and dual-AES links. The Metric Halo is a well-regarded two-channel A/D and D/A converter with low-noise microphone preamps and a headphone amplifier. It uses a 24-bit sigma-delta converter from AKM, their 4393 chip. Although it has AES/EBU and S/PDIF inputs and outputs, the ULN-2 is primarily used to feed 24-bit data at sample rates up to 96kHz to a host Macintosh computer via FireWire.

I'm pleased with the sound of the recordings I've made with these converters; it was to see how a new A/D design from a leading high-end manufacturer compared with them that triggered my doing this review.

Because I have no easy way of storing the dCS 904's dual-AES/EBU 192kHz output as a computer file, and as the Metric Halo is limited to 96kHz, I digitized LPs with these converters at 96kHz. I then redigitized the LPs I'd sampled at 192kHz with the QA-9 at 96kHz. Although the converters were limited to a peak level of between –3 and –6dBFS when the LPs were being ripped, I normalized the peak level of all the files to –0.1dBFS with BIAS Peak, so that there were no loudness differences. Although the dCS offers a choice between brick-wall and slower-rolloff anti-aliasing filters, I used the brick-wall.

First up was the Metric Halo ULN-2 at 96kHz. In "Die Tänzerin," Ulla Meinecke's voice sounded slightly harder at high levels than the QA-9, with a little less palpability to the image. The repeat echoes for which, in the 1980s, this track became renowned as an audio-show demo track were as audible through the ULN-2 as they'd been through the QA-9, but paradoxically, the soundstage overall was a little flatter. Low frequencies were very similar with the two converters, the punchy bass interjections of the electronic keyboard in the song's coda sounding equally weighty—though if I had to swear to it, the Ayre's lows were marginally tighter overall.

With "Puttin' On the Ritz," from the Peter Skellern LP, the Ayre's high frequencies sounded smoother than the Metric Halo's, but without being in any way mellow or rolled off. The small inflections in Skellern's phrasing were slightly better developed than with the Metric Halo; once I'd heard that, going back to the Ulla Meinecke rip revealed the same difference between the QA-9 and the ULN-2.

Turning to the dCS 904, again with 96kHz files, the converters sounded much more similar than had the QA-9 and ULN-2. The low-frequency keyboard jabs in "Die Tänzerin" sounded identical through the dCS and Ayre, as did Meinecke's voice, whereas the ULN-2 had emphasized fricatives a little. After repeated comparisons, especially using headphones, I felt that the reverberation accompanying the finger snaps and occasional handclaps that so effectively punctuate this song had a slightly more integrated relationship with the direct sounds through the QA-9, leading to a slightly more developed sense of image depth than with the dCS 904. With the Skellern, again, the converters were very close, but with a slightly better developed sense of the recorded space through the Ayre. I also felt that the dCS 904 was not so kind when it came to dealing with the inevitable end-of-side distortion.

One peculiarity: In Sly and the Family Stone's "Que Sera, Sera," when Doris Day reaches the end of the first line —"When I was just a little girl, / I asked my mother, what will I be?"—she flattens the final word. With the 96kHz Ayre transfer, it was just a little more easy to distinguish the pitch of that blue note against the accompanying chord on the Rhodes piano and organ than with the dCS. I have no idea what that means, other than to note that listening to the original LP, there was never any ambiguity about the note's pitch.

To put these comparisons into context, the Metric Halo costs $1695 and includes a wealth of DSP functionality; the dCS 904 cost around $7000 when last available. For Ayre's QA-9 to be not only competitive but to excel the expensive dCS converter's performance, even by a small degree, while costing $3950 or $4750, makes it a great value.

Summing Up

Why rip an LP when a CD is available? I'm not going to get into the whole "LPs sound better than digital" matter —*pace* Mikey Fremer. But CDs from before the mid-1990s were generally recorded and mastered using A/D converters that were only optimistically described as having 16 bits' resolution. It was only at the end of the 1980s, with the advent of the first true high-resolution A/D converter, designed by Robert Adams of dbx and marketed by UltraAnalog, that CDs had a true 16 bits of resolution. (Yes, the word length on a CD has been 16 bits since the crystallization of the format in the famous "Red Book" in 1981, but the least-significant bits of the 16 back then were often quantization artifacts or random noise.) As a well-recorded LP can have resolution greater than CD's 16 bits in the region where the ear is most sensitive, and can have a recorded bandwidth greater than the CD's 22.05kHz, transferring an LP to digital with something like Ayre's QA-9 makes a lot of sense.

On the other hand, and particularly with classical LPs, getting rid of the inevitable ticks is a royal pain. I have always kept my LPs in good condition, and I both clean them with a carbon-fiber brush and get rid of static with a Zerostat before I rip an LP. But there are always some ticks, and classical music, with its wider dynamic range, is less good at masking them. I resolve to eliminate only the worst problems (using the waveform pencil tool in BIAS Peak Pro 7) when I clean up the file (footnote 2), but after I've done so, lower-level ticks that hadn't bothered me before now sound louder, and my obsessive-compulsive tendency urges me to make another pass. And then another.

Given that I believe you should rip LPs with your loudspeakers muted (almost all turntables and tonearms are microphonic, to a greater or lesser degree), and adding the cleaning-up required, ripping an LP becomes a lot like work—even without taking into consideration that I also normalize the 24/192 file after the click removal and prepare a downsampled 24/48 version of the tracks to play on my iPod Classic. It has become something I don't relish having to do.

But when recordings you love have never been issued on a good-sounding CD, it makes sense to rip them with Ayre's QA-9—it's the closest thing to a truly transparent audio component I have encountered. Even though I've purchased many A/D converters over the years—a Manley ADC, two dCS 904s, a Metric Halo ULN-2 and MIO2882, an E-MU 404, an M-Audio Transit—I am buying the review sample. There are *many* LPs waiting to be ripped.

Footnote 2: Pure Vinyl's Rob Robinson points out that if you rip the LP "flat," with no RIAA equalization—which his program can apply when the file is played back—the tick will last only a couple of samples and is much easier to delete. Ripping with RIAA, as I do, results in what was originally a straight up/down pulse being stretched into something that both lasts longer and both under- and overshoots the waveform's timeline.