RECORDS TO LIVE FOR!

DAVE ALVIN
BLASTERS GUITARIST SPEAKS

stereophile

ONLINE AUTHORITY: STEREOPHILE.COM

CH PRECISION C1.2
THE SOUND OF LIVE MUSIC

⬆ DYNAUDIO FOCUS 10 WI-FI SANS COMPROMISE

⬅ MOFI SOURCEPOINT 10 ANDREW JONES GOES CONCENTRIC

⬅ ROTEL AT SIXTY: DIAMONDS GLITTER

FEBRUARY 2023
$8.99 US
$10.99 Canada

JIM AUSTIN

CH Precision C1.2

D/A PROCESSOR

If you're reasonably handy, you can probably build your own digital-to-analog converter. It won't cost much, and if you're careful, and knowledgeable enough to understand and follow some rather technical instructions, or if you have patience enough to follow advice from a few different online discussion forums—and the judgment to distinguish the good advice from the bad—then the DAC you make may end up sounding very good.

So it's no surprise that you can buy very good Chinese-made DACs that measure very well, very cheaply. Those Chinese DACs are probably designed by first-rate engineers, and while extracting maximum technical performance from a good DAC chip requires care and attention, it isn't rocket science.[1]

What, then, is the point in paying tens of thousands of dollars for a D/A converter?

It's a reasonable question, one that every DAC shopper must answer for themselves. Is extremely low measured jitter, noise, and distortion all that matters in a DAC? Is it sufficient assurance that it will sound "perfect," as good as a DAC can sound? Or is it possible to take this basic technology further, despite what the measurements show? It's easy enough to find people who are quite happy with their $1k DAC and smugly confident that they're getting the best possible sound. But in perfectionist audio (and certainly in this magazine), it's axiomatic that progress is *always* possible, that you can *always* do better, and that measurements—at least the easy and obvious measurements, such as S/N ratio, distortion level and profile, and Miller-Dunn J-Test jitter—don't tell the whole story. And if you listen with trained ears through topnotch audio systems well set up, it's frankly hard to miss the improvement in sound achieved by expensive DACs produced by companies committed to achieving the best possible digital sound.

And if you disagree? Then you just saved yourself a ton of money.

The CH Precision C1.2 D/A Controller

I'm sitting back in my lightly chewed IKEA chair, listening to Benjamin Grosvenor's performance of the Liszt B-minor sonata, S.178, recorded in Queen Elizabeth Hall at London's South Bank Centre. It's from Grosvenor's album *Liszt*, and it's streaming from Tidal (24/96 MQA, Decca). I'm listening on a system most would consider very good; it certainly isn't cheap. It includes the Wilson Alexx V loudspeakers, two Burmester 218 amplifiers (each bridged for mono, in for review[2]), the Pass Labs XP-32 preamplifier, and not-quite top-level cabling by Nordost and AudioQuest.[3]

The source of this music is the new CH Precision C1.2 D/A Controller ($43,000 as equipped), aided at the moment by a complete CH Precision digital front-end: the X1 power supply ($20,500), the T1 clock ($24,500), and the D1.5 transport ($49,500 but not currently in use). I've set the volume to what I'd expect to hear if I were sitting in the first few rows of the concert hall—and indeed, the sounds I'm hearing *could* be emerging from a Steinway on the stage of a good concert hall.

Well, to be completely honest: not quite. This is a very good performance and well-recorded, but, while the highs I'm hearing have an appropriate, crystalline "ping," the lower-midrange keystrokes seem ever so slightly dulled; a touch of transient bite is missing. There's also some congestion on the loudest passages, a sense that the piano's case is filling up with sound and distorting a little, some-

1 Yet, a look at some of JA1's measurements reveals that commercial implementations of common DAC chips *often* fall short of a chip's potential.

2 According to *Stereophile* policy, reviews must be performed in a well-known room, mainly on well-known equipment, so I have already listened extensively—for several weeks—on my reference Pass Laboratories XA60.8 amplifiers. See my review at stereophile.com/content/pass-laboratories-xa608-monoblock-power-amplifier.

3 You'll find my reviews of the Wilsons and the Pass Labs preamps at stereophile.com/content/wilson-audio-specialties-alexx-v-loudspeaker and stereophile.com/content/pass-laboratories-xp-32-line-preamplifier, respectively.

SPECIFICATIONS

Description Upsampling D/A processor with volume control, remote control for basic functions, Android-only app for settings and operation. Conversion type: Linearized R-2R using four PCM1704 chips per channel, operating at 24 bits, 705.6kHz & 768kHz, DSD via DoP or direct conversion up to DSD128. Bypassable volume control operates in 0.5dB steps. Adjustable channel balance, switchable mono (summing) and absolute phase selection. Standard inputs: AES3, S/PDIF, TosLink, CH-Link HD (proprietary). Optional inputs: Asynchronous USB, Ethernet, analog (XLR and RCA). Class-A output stage with zero global negative feedback. Output levels (FS, RMS): 255V (single-ended/BNC or RCA) or 5.1V (balanced/XLR). S/N ratio: >120dB. THD + N: <0.001% FS below 22kHz, B-weighted. Optional clock-sync board. 800 × 480 pixel, 24-bit RGB AMOLED display.

Dimensions 17.3" (440mm) L × 6.3" (133mm) H × 17.3" (440mm) D. Weight: 44lb (20kg).

Finish Silver.

Serial number of unit reviewed 0Y9F1401. Manufactured in Switzerland.

Price $36,000 for stereo version with a single HD input board. Dual-mono version is $77,000. Options: Digital input board, $2500; asynchronous USB input board, $3000; Ethernet streaming board, $6000; analog input board (with one balanced and one unbalanced input), $2500; clock synchronization board, $1500. As equipped, $46,500. Upgrade from C1 status is $4000.

Approximate number of US dealers 7. Warranty: Three years, parts and labor.

Manufacturer CH Precision Sàrl, ZI Le Trési 6D, 1028 Préverenges, Switzerland. Tel: (41) (0)21-701-9040. Web: ch-precision.com.


DSD 64
DAC
CH LINK HD
10.0000 MHz
C1.2
SACD
7 ▶ 0:01:07
D1.5
10.0000 MHz
BNC 1
INTERNAL
LOW
T1
OUTPUTS
ON
SLOT LEFT
ON
SLOT RIGHT
X1
GIULINI
IN CONCERTO

thing I've noticed in live performances but not this much.

Despite these minor flaws, this system is delivering a spectacular experience. The piano has real grunt—more than makes it to my listening seat at most of the piano performances I attend[4]—and lots of high-end sparkle. Decay, of notes high and low, is natural and even.

But what of those flaws I heard? Should we blame them on the CH Precision digital front-end? No. It's clear that the fault lies in the way the piano is miked, which trades transient clarity for low-end impact.

How to build a CH Precision DAC

If your goal is to make a DAC that's better than one you can make with a very good DAC chip, the way to do it is to start with a concept. You need a theory for how to proceed, or, as baseball commentators like to say about hitting, you need a thoughtful, fundamentally sound *approach*. It helps, of course, if your theory is correct, and if it's just plain wrong you're in trouble. But for reasons I think will soon become apparent, your theory need not be *precisely* on the money. CH Precision's approach—shared generously with me by Florian Cossy and Thierry Heeb—is based on the notion that timing is everything. Getting the frequency part right is easy enough. What's hard is getting things right in the time domain.

Both Cossy and Heeb are engineers. Heeb is the digital guy. In addition to being the "H" in "CH," he's a senior researcher at the University of Applied Sciences and Arts of Southern Switzerland, specializing in DSP for audio. Cossy—the "C" in "CH"—is the company's CEO; his engineering expertise is on the analog side.[5]

The first step toward understanding why timing matters in a D/A converter—or why it makes sense to *assume* it matters beyond mere 1s and 0s—is to recognize, as Heeb told me months ago in a Zoom conversation, that in audio, a digital signal is best thought of as analog. "Even if the signals or the electrical signals are supposed to be digital, basically just two levels, a zero and a one, as soon as you get into an electronic board, they are actually analog signals, current or voltage flowing through components. That is especially true, for instance, for clock signals. If you just consider clock signals as being a shift between two values between zero and one, you don't really get what clock is. The most important point in clocking is in the time domain"—well, duh—"with finite resolution. Basically, it boils down to an analog signal again."

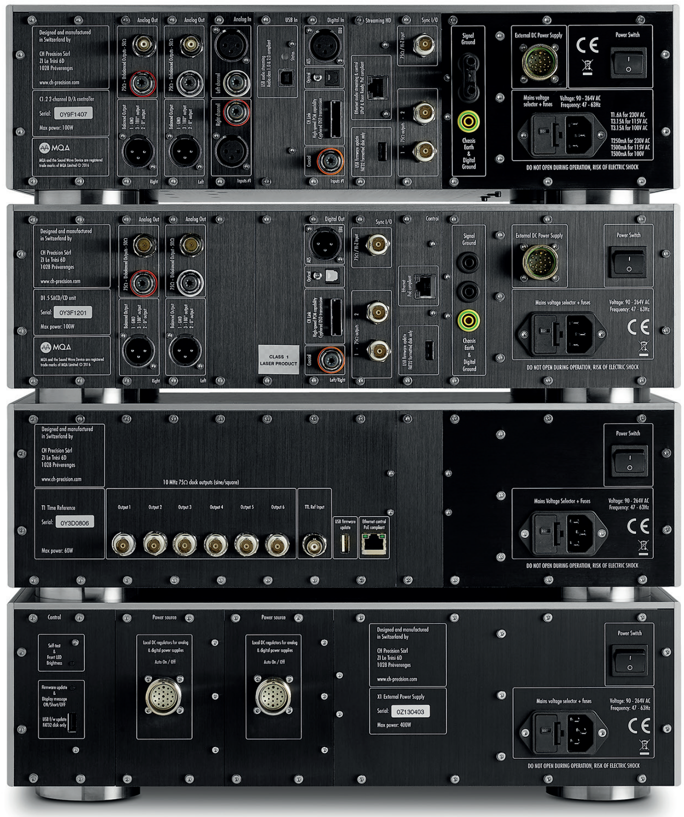

4 Although not, I'm thinking, at Manhattan's newly rebuilt Geffen Hall. I've attended two shows there now. Though a little bit dry, that hall has serious grunt.

5 Also, of course, "CH" stands, in Latin, for "Confoederatio Helvetica," or Swiss Confederation—for Switzerland.

MEASUREMENTS

I measured the CH Precision C1.2 with my Audio Precision SYS2722 system,[1] repeating some measurements with the higher-performance APx500. The external power supply and clock weren't available for the testing. The C1.2's coaxial and optical S/PDIF inputs and AES3 input accepted data sampled at all rates up to 192kHz. Apple's AudioMIDI utility revealed that the C1.2 accepted 16- and 24-bit integer data via USB sampled at all rates from 44.1kHz to 384kHz. Apple's USB Prober app identified the C1.2 as "CH Precision USB Audio 2.0" from "CH Precision" and indicated that the USB port operated in the optimal isochronous asynchronous mode.

The C1.2's output impedance was a usefully low 64 ohms, balanced, 73 ohms, RCA unbalanced, and 49 ohms, unbalanced BNC, all values consistent from 20Hz to 20kHz. With the C1.2's gain set to its maximum, the output level with a full-scale 1kHz tone was the specified 5.1V for the balanced output—0.5dB lower than that of the C1.2's predecessor, the C1—and 2.54V for both types of unbalanced output. Reducing the maximum gain by an indicated 12dB reduced the output level by exactly 12dB.

All the outputs preserved absolute polarity, which can be seen in fig.1. This graph indicates that the C1.2's reconstruction filter is a very short, linear-phase type, with just one cycle of ringing on either side of the single high sample. This type of time domain–optimized filter is associated with a very slow low-pass function, which can be seen with the magenta and red traces in fig.2, taken with 16-bit white noise at –4dBFS. The output doesn't reach full stop-band attenuation until an octave above the audioband! With a full-scale tone at 19.1kHz (blue and cyan traces), an aliased image at 25kHz lies at –12dB and other aliased images can be seen between 60kHz and 70kHz. Distortion harmonics of

1 See stereophile.com/content/measurements-maps-precision.

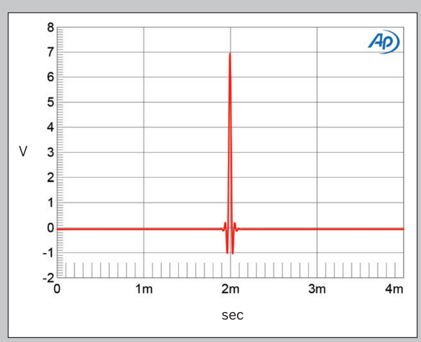

Fig.1 CH Precision C1.2, impulse response (one sample at 0dBFS, 44.1kHz sampling, 4ms time window).

I'll just throw this in: In the physical world, *music* happens in the time domain. True, we *do hear* frequency—as pitch, and combinations of frequencies at chords, or as vocal or instrumental timbre—but, strictly speaking, those musical signals exist only as a function of time: In your ear canals, there is only one level of pressure at an instant of time, and it changes.

The frequency domain is, strictly speaking, a mathematical abstraction.

There are two things (at least) that lead to time-domain errors: timing randomness—also known as jitter[6]—and an intrinsic lack of precision in D/A conversion, which Heeb (and others) call time smearing. Time smearing is the same concept that MQA is intended to address—they too call it time smearing—and, indeed, CH Precision's approach to dealing with that phenomenon seems quite similar to MQA's approach. In a comment published in my review of the CH Precision D1.5 transport/player[7], Heeb said, "Time smearing is basically if you put a single pulse through the system, if you have a filter with a very long impulse response, that single sample will extend over a large number of samples." The goal, then, is to shorten the impulse response so that the musical content of an input sample extends over as little time—over as few samples—as possible. How is that achieved? With an approach to conversion that's quite different from the approach outlined by the foundational document of digital audio, Shannon's theorem.

Shannon's theorem says that if certain conditions are met, the output of an A/D–D/A sequence can exactly match the input, mathematically. But that's not true in the real world, under any real-world circumstances, because the conditions are unphysical. They do not exist. For example, the basic mathematical function Shannon employed for sampling and reconstruction—the sinc(x) function—goes from minus infinity to plus infinity in time, which in the real world never happens. ("There is no energy in the signal before the instant where the musician starts playing," Heeb wrote

6 I've been hearing for years, from digital designers, that jitter can affect sound at far lower levels than previously thought—and that the effects of jitter are manifold: It's not just the edginess heard, for example, on the jitter tracks on *Stereophile Test CD 2* that affect imaging precision, subjective tonal balance, and other aspects of musical presentation.

7 See stereophile.com/content/ch-precision-d15-sacdcd-playertransport.

measurements, continued

the 19.1kHz tone are extremely low in level, however, with the third lying at just –97dB (0.0014%).

With 44.1kHz data, the C1.2's output was down by 3dB at 20kHz (fig.3, green and gray traces), which is typical of a B-spline–based reconstruction filter. The responses with data sampled at 96kHz and 192kHz followed the same basic shape, but with the –3dB frequency proportionally higher. Neither the frequency responses nor the superb channel matching changed at lower volume-control settings. Channel separation (not shown) was also superb, at >120dB in both directions below 3kHz and still 113dB at the top of the audioband. The low-frequency noisefloor (fig.4) was very clean, with no power supply–related spuriae present.

Fig.5 shows the C1.2's balanced output spectrum with a dithered 1kHz tone at –90dBFS with 16-bit data (green and gray traces) and with 24-bit data (blue and red traces). With the 16-bit data the noisefloor is that of the dither; with 24-bit data the noisefloor drops by around 20dB, which suggests a high resolution between 19 and 20 bits. However, a regular series of harmonics is present with the 24-bit data, with the odd-order harmonics the highest in level. (The unbalanced outputs behaved identically, other than with lower levels of the even-order harmonics.) This behavior implies that the lowest significant bit is being truncated. With undithered 16-bit data at exactly –90.31dBFS, the three DC volt-

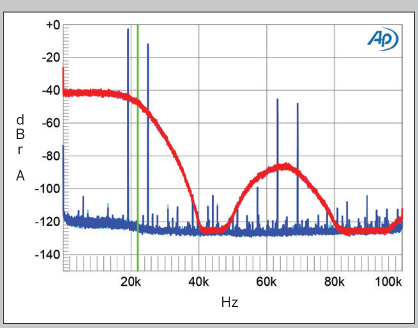

Fig.2 CH Precision C1.2, wideband spectrum of white noise at –4dBFS (left channel red, right magenta) and 19.1kHz tone at 0dBFS (left blue, right cyan), with data sampled at 44.1kHz (20dB/vertical div.).

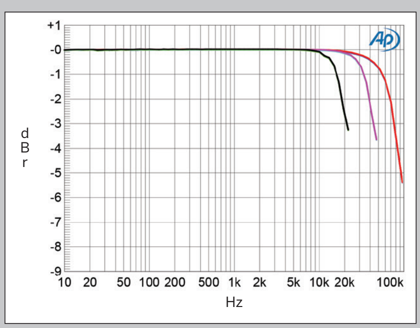

Fig.3 CH Precision C1.2, frequency response at –12dBFS into 100k ohms with data sampled at: 44.1kHz (left channel green, right gray), 96kHz (left cyan, right magenta), and 192kHz (left blue, right red) (1dB/vertical div.).

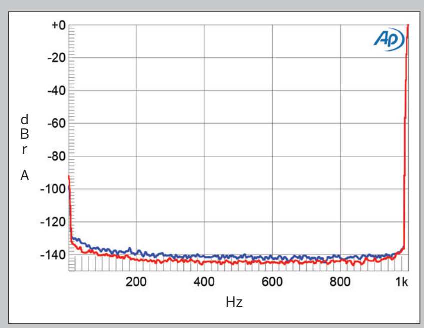

Fig.4 CH Precision C1.2, spectrum of 1kHz sinewave, DC–1kHz, at 0dBFS (left channel blue, right red; linear frequency scale).

to me by email.) Anyway, CH Precision would not *want* to use a sampling/reconstruction "kernel" that's infinite in duration, because, well, that's *a lot* of time smear.[8] "We prefer to use splines, which have a much more compact support,[9] which makes it so that when the sample goes in, what comes out has, in our case, [no more than] 100µs of pre-ringing and post-ringing," Heeb said. A particular spline can be used to represent music locally; a long series of overlapping splines can represent a whole song or symphony.

In my review of the D1.5, I found it to be a transport of obvious quality. I also found it to be, with its two monophonic D/A converter cards, an excellent player of CDs, SACDs, and MQA CDs.[10] Good as it was, though, those DAC cards are limited implementations of the CH Precision conversion approach. The C1.2 is an outright assault.

The C1.2 upsamples everything (except, according to the Roon Signal Path display, MQA data, which makes sense) to 16 times the base rate: 44.1kHz data and its multiples are upsampled to 705.6kHz; 48kHz data and its multiples are upsampled to 768kHz. This, though, is not your mother's upsampling. In performing this upsampling, the C1.2 does something that was common in the early digital era but that's surprisingly rare these days (so maybe it *is* your mother's upsampling): It keeps all the original data points, interpolating new samples between them. Other approaches, most notably asynchronous sample-rate conversion, obliterate the original stream completely (except the very first sample) and replace it with a completely new datastream. The time series described by the new stream may be very close to the old stream; nevertheless, this strikes me as an interesting point, philosophically and perhaps sonically: How can you claim the original spectrum is perfectly recreated (it's not) when all the data have different values?

The "base"-model C1.2 doesn't include a USB input, but you can get one ($3000). CH Precision's USB input card is a bit different from others. While it does reclock incoming data—that's the advantage of an asynchronous, isochronous USB interface, in principle—it does not resample. Even via the USB input, the original samples are preserved.

At the end of this chain of conversion technologies is something surprising: a DAC chip. Not just any DAC chip, but one that was an important step forward for digital audio when introduced—in 1998. It is Burr-Brown's PCM1704 R-2R ladder DAC chip, four per channel. Why do it this way instead of laying out an actual R-2R ladder with resistors, as several much cheaper, excellent-sounding Chinese imports do? I asked that. "The fact that it is a monolithic chip makes it both consistent and wonderfully accurate to work with, something that a discrete ladder cannot achieve even with the highest precision resistors," Cossy answered. He also wrote, "Even though it is an 'old' chip, it more than meets current requirements."

8 Modern sampling theory long ago abandoned the idealized notion of perfect reconstruction. An example of this is the use of a reconstruction kernel (a spline function, say) that differs from the one used for sampling (perhaps a sinc(x) function). "The key question is, how do the sampling and reconstruction kernels combine?" Heeb wrote in answer to another question. "In other words, what is the result of a reconstruction kernel applied to a sampling kernel on a unit pulse? If the result is close enough to identity (in a given frequency band and a given time space), then different kernels can be used with no apparent drawback." So, wise designers long ago stopped being slaves to Shannon's theorem, favoring instead an approach that attempts to minimize error and to *shift* error to where it does the least harm. This, I believe, is why there's more than one legitimate approach to D/A conversion—and why it remains an unsolved problem. There are various legitimate approaches—valid assumptions as to *where* the inevitable error does the least sonic harm.

9 "Compact support" is a mathematical term that means that, outside a certain finite range, the value of the function is zero.

10 Although as an MQA-CD player, I had nothing to compare it to. It was the only MQA-CD player I've ever auditioned.

measurements, continued

age levels described by the data are well resolved (fig.6), and high-frequency noise is extremely low in level. With undithered 24-bit data at the same level, the result was a well-formed sinewave (not shown).

The red trace in fig.7 plots the error in the balanced output level as a 24-bit, 1kHz digital tone steps down from 0dBFS to –140dBFS. (This graph was taken with the left channel's output; the right channel behaved identically.) The amplitude error starts to increase below –80dBFS, which is associated with the harmonic distortion seen in fig.5. I understand that the C1.2 uses parallel pairs of PCM1704 DAC chips; the behavior in figs.5 and 7 might be due to the DAC pairs not being perfectly matched in low-level linearity. (Achieving good low-level linearity with R-2R ladder DACs is always difficult,[2] which is why many designs use sigma-delta chips where this is not an issue.[3])

The C1.2 offered very low levels of harmonic distortion, with the third harmonic the highest in level at –100dB (0.001%, fig.8). Though other harmonics are present, these all lie at lower levels. These harmonics didn't increase in level when I reduced the load to 600 ohms, but the low-order harmonics decreased slightly in level when I reduced the signal level by 3dB. Fig.9 plots

2 For an example of excellent low-level ladder-DAC linearity, see figs.11 and 12 at stereophile.com/content/holoaudio-may-level-3-da-processor-measurements.

3 See stereophile.com/content/pdm-pwm-delta-sigma-1-bit-dacs-john-atkinson.

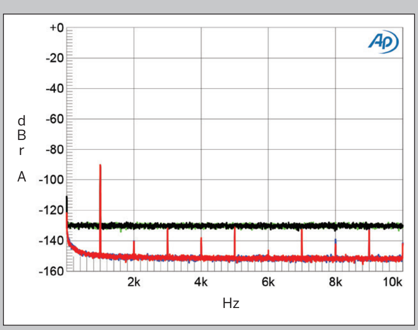

Fig.5 CH Precision C1.2, spectrum with noise and spuriae of dithered 1kHz tone at –90dBFS with: 16-bit data (left channel green, right gray), 24-bit data (left blue, right red) (20dB/vertical div.).

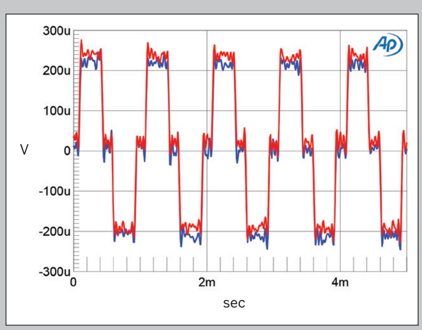

Fig.6 CH Precision C1.2, waveform of undithered 16-bit data, 1kHz sinewave at –90.31dBFS (left channel blue, right red).

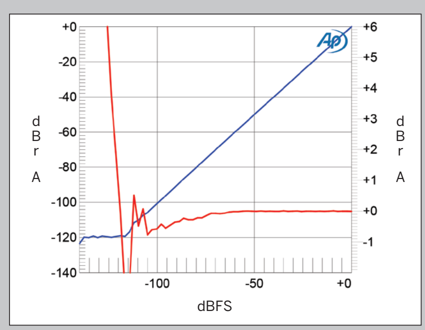

Fig.7 CH Precision C1.2, left channel, 1kHz output level vs 24-bit data level in dBFS (blue, 10dB/vertical div.); linearity error (red, 0.5dB/small vertical div.).

You wouldn't expect CH Precision to use a boring old volume control,[11] would you? Well, they don't. Instead, the C1.2 utilizes a hybrid analog/digital control, which combines three large analog steps (via relays) with smaller digital domain steps.

The C1.2 from the outside in

Everything I've written up to now was true of the earlier C1 DAC (except maybe the part about the volume control; I'm not sure about that). So, what's new in the C1.2? What has changed?

First, though, an aside on naming. Why name two products released so close together so differently? The D1.5 came out months before the C1.2. Why not call them both "1.2" or "1.5"?

At CH Precision, the model-number increment indicates upgradeability. Physically, the C1.2 is very similar to the C1. The ".2" designation indicates that if you own a C1, you can upgrade it to .2 status—for $4000. The D1.5, though, is so different from its predecessor, the D1, that it's not physically possible to upgrade the older to the newer: The D1.5 uses a different transport, with a different door height. Not to worry though: CH Precision offers a guaranteed buy-back in cases like that.

One very visible change in the C1.2 is the introduction of MQA support, which was not present in the C1. The C1.2 supports full decoding—unfolding and rendering—to frequencies up to 768kHz, including data from MQA-CDs played back on the D1.5 transport over the proprietary CH interface as well as MQA data streamed from Tidal.

I asked the two engineers how MQA is handled in the C1.2—with an off-the-shelf chip, perhaps? Not hardly. The C1.2 detects whether a datastream is MQA or not then sends it in one of two directions, toward the MQA algorithm (MQA data) or toward the PEtER upsampling algorithm (everything else). MQA data is interpreted in silico using a software library provided by MQA.

Maybe the biggest news with the C1.2 is a new MEMS (microelectromechanical systems)–based clock, which is shunt-regulated (roughly, designed so that noise is shunted to ground) and temperature-compensated for improved accuracy.

Processing power has increased by a factor of four. The most obvious impact of this change is on the expanded range of input formats supported; the C1.2 now supports all of them, from a file or silver disc. The more significant impact of this increased computational power is more precise upsampling calculations. That's possible in part because the computational space has been expanded to 32 bits fixed-point (not floating).

It has become clear over time that one key to achieving the best possible digital sound is to keep the signal path free of noise. So—this is new also—the C1.2 turns off all processing channels that aren't currently in use, in order to lower system noise.

What else has changed? That hybrid analog/digital volume control has changed—but I'm not clear on whether it's completely new or just tweaked. Also, the display screen is better, although you won't notice it at first. It's capable of higher resolution than what you see most of the time. It's still not a thing you'd watch a movie on, but it looks pretty spiffy when you put the C1.2 in preamp mode and change the volume.

The C1.2 is modular, and when you consider all the options, remarkably flexible. It can utilize any common digital input, including Ethernet, plus CH Precision's proprietary data link, which resembles I^{2}S and supports data-transmission rates up to the highest rates you'll commonly see. With the analog input card, you get two sets of analog inputs. Since it has a very good volume control, you could make it the central component of your audio system.

11 The volume control can be difficult to locate among the C1.2's many menu options. It's hidden in the "Factory" menu, presumably because it's such a fundamental choice: whether to use the C1.2 just as a DAC or also as a preamplifier.

measurements, continued

the spectrum of the C1.2's balanced output with an equal mix of 19kHz and 20kHz tones, the 24-bit signal peaking at 0dBFS. The use of a slow-rolloff reconstruction filter results in high-level aliased images of the tones at 24.1kHz and 25.1kHz, but actual intermodulation products are very low in level.

Finally, I tested the CH Precision's rejection of word-clock jitter with 16-bit, undithered J-Test AES3 and TosLink data. Other than those closest to the Fs/4 spectral spike, the odd-order harmonics of the LSB-level, low-frequency squarewave are very close to the correct levels (fig.10, sloping green line), and no other sidebands are present. With 24-bit J-Test data (not shown), a single pair of sidebands was still present at ±229.6875Hz, but the random noisefloor lay at a very low –147dB.

Other than the possible truncation of the 24th LSB and the mismatch of the DAC chips' low-level linearity, both of which were present with the earlier C1, the CH Precision C1.2 offers generally excellent measured performance. The C1.2's behavior is dominated by its use of a reconstruction filter optimized for time-domain performance, with its very slow ultrasonic rolloff.—**John Atkinson**

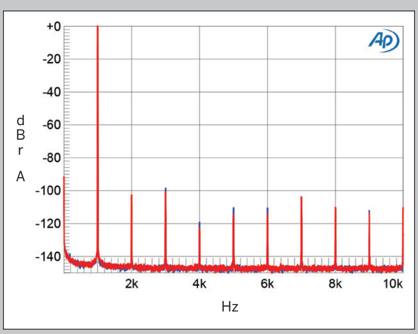

Fig.8 CH Precision C1.2, spectrum of 1kHz sinewave, 24-bit data, at 0dBFS, DC–10kHz, into 100k ohms (left channel blue, right red; linear frequency scale).

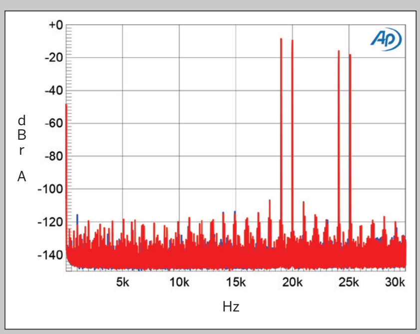

Fig.9 CH Precision C1.2, HF intermodulation spectrum (DC–30kHz), 19+20kHz, 24-bit data, at 0dBFS into 100k ohms (left channel blue, right red; linear frequency scale).

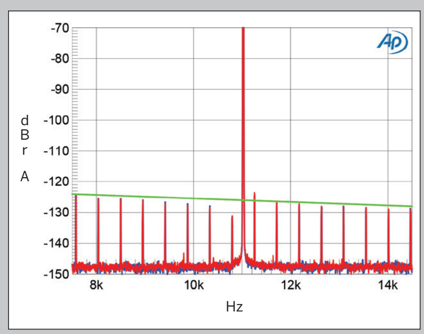

Fig.10 CH Precision C1.2, high-resolution jitter spectrum of analog output signal, 11.025kHz at –6dBFS, sampled at 44.1kHz with LSB toggled at 229.6875Hz: 16-bit undithered AES3 data (left channel blue, right red). Center frequency of trace, 11.025kHz; frequency range, ±3.5kHz.

Listening

What do we look for—rather, listen for—in a digital music source? Or, for that matter, any audio source, or any audio system? "Tastes vary" may be the most important answer to that question, but I hope there are values we can all agree on: rich timbres and textures, vivid colors, images that seem solid and real, commanding bass, airy highs. Some will insist more than others that the sounds our systems produce be true to the source, although that can be hard to determine. (Loudspeaker designers, a suggestion: Don't release a high-end speaker that can't accurately reproduce common piano, like the sound of a Steinway Model D in a good hall. I hear a few loudspeakers that can't do that at every audio show I attend.)

Recently, over lunch with a small group of *Stereophile* writers, I shared my belief that one crucial thing in experiencing reproduced audio is a constant sense of *surprise*. Heads nodded. When a hi-fi system does harm to the music, it often takes the form of *homogenization*, making sounds seem more similar to each other, hence more ignorable and less surprising. Dynamic compression, for example, reduces contrast between loud and soft sounds, which tends to make music less surprising (especially with so-called microdynamics) and so, less real.[12] Homogenization of every kind makes music less interesting and puts us to sleep. In contrast, a constant stream of pleasant surprises, which come through when the uniqueness of every sound is preserved, makes us look up and smile with delight at the music even when we're not paying close attention. That's a big part of what keeps me coming back.

There's another thing, though, that tends to come up in conversations about digital sources: a sense of relaxation in the music, whether the music *encourages* or at least *allows* relaxation in the listener—or whether, conversely, it is itself a source of stress. Digitally reproduced music *can* be stress-inducing. (So, in a different way, can scratchy old LPs.)

It's a peculiar idea: that something so important in hi-fi, some of the most important stuff, is something we experience in some unknown way but don't directly, or consciously, hear. How else to explain bass that (as I wrote in my review of the CH Precision D1.5 transport/player) sounds "fundamental" (in the sense of the root word "fundament") and "seismic," when we all know the LF frequency response will measure the same? So it's not the intensity or depth of the bass response I'm hearing per se; rather, it's how I experience it—and something in the music causes that.

When Heeb, Cossy, and the C1.2 documentation talk about the importance of precision in upsampling calculations or of reducing timing errors, they're not saying that if we don't do those things we'll end up with awful jitter, gross errors that affect measured frequency response, or that transients will be dulled or artificially sharpened (although our perceptions of all those things may be altered). In executing their design brief, they are indeed producing a more accurate signal, but the most important subjective consequences are—let's say, indirect. The specific mechanism is unknown, at least to me, but when you get it right, you hear it. They're gaining something, but they're also getting rid of something that, when it's present, stands in the way of our ability to perceive music simply and directly, with low stress and nothing interfering. (Then again, I think much that we call distortion, measurable or not, is like that.)

Remember Grosvenor's B-minor sonata, which I started listening to earlier? After it ended, Roon Radio started up Khatia Buniatishvili's 2011 recording of the *Mephisto Waltz No.1*, also by Liszt. The sound got better—it's a much livelier recording. I'm listening to just two speakers, sitting 11' in front of me, but the sound is enveloping me in a way that the Grosvenor recording didn't, especially in the louder passages. That sense of piano-case overload I mentioned is absent from this recording. The perspective here is more distant than on the Grosvenor recording, yet I can clearly hear a difference in soundstage depth between the piano's high notes and low notes, sounds emerging where the hammer hits the strings—rather, where those sounds reflect off the piano's open lid. I'd say I'm sitting in row 20 or so—that's the aural perspective—so the piano is pretty far away, but the effect is very clear. And even the loudest sounds seem relaxed, stress-free.

Speaking of Buniatishvili: Not only does she make wonderful recordings of great music beautifully played; she also chooses superbly interesting repertoire. The next-to-last track on her album *Labyrinth*, from 2020 (24/96 FLAC, Sony Classical/Qobuz), is John Cage's 4'33". I won't say it's her best performance, but it's certainly her most *perfect*, the one with the fewest mistakes.

Over the last few months, I've listened to a lot of classical music, naturalistically recorded in a real space. (Is that choice of music affected by my current DAC? I wonder.) With such recordings, what I hear with the C1.2 is what acoustical instruments sound like, precisely rendered in space. The sense of that space, and of the sounds flowing through it, is expansive and relaxed; that expansiveness and sense of relaxation are somehow connected. Except when the pressures of getting the magazine out the door interfere with my state of mind, I am relaxed while listening.

In this issue, Jason Victor Serinus reviews Caroline Shaw's recording *The Wheel* (24/192 wave download, Alpha), with the French collective I Giardini; it's *Stereophile*'s Recording of the Month. It's *my* Recording of the Month, too.

One track Jason didn't mention in his review is the second, "Gustav Le Gray," which, for its first half or so, is identical to—indeed, *is*—Chopin's Mazurka Op.17 No.4. After the halfway point, the mazurka comes unglued. "Gustave Le Gray," Shaw writes in the liner notes, "is a multi-layered portrait of Op. 17 #4 using some of Chopin's ingredients overlaid and hinged together with my own." Fascinating stuff. Through the C1.2 DAC, it—especially the piano, which is what I focused on the most—simply sounded *right*.

Just now, I needed a break from writing, and my six-month-old puppy Ella (who is responsible for this lightly chewed listening chair I'm sitting in) needed a break from not peeing, so we headed outside then south on Riverside Drive. At last night's dinner, a guest had mentioned Duke Ellington Blvd., also known as W. 106th St.

12 Although even a hall acoustic—I'm tempted to say especially a hall acoustic—can homogenize sound. Also: Used tastefully, dynamic compression is an essential tool for audio engineers.

My wife's grandparents lived there for a long time, on the northwest corner with Riverside Drive. Their apartment building was across from a beaux-arts mansion, which some—including one of my dinner guests—have said Duke Ellington lived in for a while. He didn't;[13] he lived around the corner at a more modest address (331 Riverside, it is said). But when we both needed a break, I put Ella, the new puppy, on a leash, and we headed toward 106th Street.

All this put me in mind of the Duke, so when we got back, I put on one of my favorite albums—an unusual one for Duke—*Jazz Party in Stereo*.[14] I usually listen to this record—this album—on vinyl. How would *Jazz Party in Stereo*, which is such a natural on vinyl, with its spacious soundstage, full of all sorts of percussive sounds, from timpani (aka kettledrum) to triangle, sound through fancy digital gear?

This is a ping-pong-y album. All those percussive sounds distributed across the soundstage, left to right and front to back, make a spectacular impression. Immediately, though, I noticed a lighter, smoother character to this highly percussive album, not in a good way. Is digital really this much worse than analog, even through a $43,000 DAC? And then I realized I was listening to a DSD file I bought some years ago (DSD64, Columbia). I know some people love it, but I have often found DSD to sound unnaturally smooth—it's one of those homogenizing influences I mentioned earlier.

I switched to the MQA version, streaming (16/44.1 MQA/Tidal). The C1.2's front panel display turned green, indicating MQA Studio. This version was louder than the DSD version, so I turned it down a bit, matching levels by ear but only roughly. Restarting the track, I immediately noticed more grunt and heft in the drums, more sharpness—even harshness—in the high percussion (xylophone, vibraphones, glockenspiel, tambourine, triangle). At first, Jimmy Woode's bass sounded like it could be a kettledrum or some other percussion instrument, but over time its "pluck" emerged. Britt Woodman's trombone had real, blatty flesh. Duke's piano sound was very natural—one of the better-recorded jazz pianos I can remember from this era.

This is what this album sounds like. It's what the *record*—the LP—sounds like. I'd probably still put the vinyl on on a celebratory Friday night, but this sounds just as good, or—it pains me to say it—perhaps better. I'm listening at 10am on a Sunday morning, feeling exhilarated, nothing between me and the music.[15] Time to whip up a cocktail? It's 5 o'clock somewhere.

The CH Precision digital stack

This is a review of the C1.2 DAC, but I was privileged to hear that instrument in the context of the full CH Precision digital front-end, with the D1.5 transport, X1 power supply, and T1 clock. How much difference did all the fixins make?

Some difference, for sure, but I didn't find them necessary. As editor of *Stereophile*, I suppose I should be an absolute perfectionist, but the fact is, I have limits. Not infrequently, I hear sound that's totally satisfying, that I could happily, joyously, live with forever. I'm getting that with just the one box, the C1.2.

Sure, if money (and, importantly, space) were no object, I'd buy them all. I say "pretty sure" because money and space are indeed objects, so I can't *really* put myself in that position; I can only pretend.

In my review of the D1.5, using it as a player, digital conversions carried out by its dual-mono DAC boards, I found—this surprised me—the external clock made a big, meaningful difference. I did not find that to be the case this time, with the C1.2. I heard differences, subtle and difficult to describe, but none that substantially increased or decreased my pleasure in listening. The X1 power supply made a bit more difference, adding, I thought, a touch more flesh, more tangibility, to acoustic objects, but I could live without that, too.

Call me easy to please, but I'm willing to settle for just the $38,500 version (with the options I'd need installed)—although I'd also be tempted to include the analog input board for another $2500. I guess a $40,000 DAC—*this* $40,000 DAC—is good enough for me. So sue me.

If I were to buy both the D1.5 *and* the C1.2—and if money were no object, I would buy both, because it's nice to have the ability to play discs—I *would* add the T1 clock. And the power supply? Compared to the other components, it's pretty affordable. Might as well throw that in, too. If money were no object.

Summing up

I have little to add. There are other digital sources in this price class, from the three-letter companies, dCS and MSB. There undoubtedly are others in a similar price class—and one manufacturer at least is charging more. But I haven't heard any of those other sources in my system, so there's no way I can compare.

If there's a downside, it's the price. It would be cool if it cost a tenth as much, but then it would also be cool if I could fly. As I've often said and occasionally written, value is a question of values—and also of wealth. If you're richer than me, I'm okay with that. It's a decision each of us must make on our own.

The C1.2, both with and without its external clock and power supply, produced the best sound I've heard from a digital source—far better than far cheaper chip DACs that we've put in Class A+ on our list of Recommended Components. Which is a problem for *Stereophile*'s editor: Do we need to create a class A++? The CH Precision C1.2 gives new meaning to "turn it up to 11!" ■

ASSOCIATED EQUIPMENT

Digital sources CH Precision D1.5 transport/player (used as transport), X1 power supply, and T1 clock. Roon Nucleus+; Synology DS918+ 4-bay Network Attached Storage device with 16TB; Melco S100 Ethernet Dataswitch.
Preamplification Pass Labs XP-25 phono preamplifier, Pass Labs XP-32 line preamplifier.
Power amplifiers Pass Labs XA60.8 monoblocks, Burmester 218 stereo amplifiers bridged for mono.
Loudspeakers Wilson Audio Specialties Alexx V, Estelon XB Diamond Mk.2.
Cables Digital: AudioQuest Carbon & Cinnamon & Coffee (all USB); Nordost Valhalla 2 (Ethernet). Interconnect: Burmester (XLR), Nordost Valhalla 2, AudioQuest. Speaker: AudioQuest Thunderbird ZERO. Power: Burmester, Nordost Valhalla 2, AudioQuest Tornado High-Current C13, NRG-X3, and Monsoon C13.
Accessories PS Audio Power Plant P10 power conditioner, Butcher Block Acoustics RigidRack, IsoAcoustics, and Magico footers.—**Jim Austin**

13 The mansion does have a musical history though, sort of. Back in the 1980s and '90s, when my wife and I were visiting her grandparents in the apartment across the street, it was known as the Seagram building because Seagram heir Edgar Bronfman Jr. lived there. Bronfman would soon become Seagram's CEO and sell off valuable assets to make big bets in the music industry just as illegal file sharing was starting to shred it. He would go on to be CEO and chairman of the Warner Music Group and do a skillful job keeping the company afloat during some of its worst years.

14 There's a mono version, which is called just *Jazz Party*. Straight from the liner notes: "As the crowd gathered, Duke was on the phone calling his group of nine percussionists, and the studio lobby was filling up with kettle drums and xylophones. Chairs were set up for our unexpected audience, and Duke, with the innocent expression of a small boy who has just dropped a match into a gas tank, said, "Let's see what happens."

15 No, I'm not listening naked. That's not what I meant.