Arcam AV40 Audio-Video Preamplifier – Technical Report

Introduction

A sample (understood to be z-stock, serial number unknown) was recently loaned by its owner to an independent US audio reviewer, Amir Majidimehr, who owns both an audio dealership near Seattle and a website, Audio Science Review (ASR).

The ASR report was not favourable and highlighted a number of apparent deficiencies it its technical performance, especially in the digital domain.

We took a sample (serial number AN0327 HJ000015) with software v1.28, which is the same SW release as the unit reviewed by ASR. An Audio Precision AP2700 analyser, which can generate analogue audio and S/PDIF + optical digital audio streams, plus a Quantum Data 780D (QD) generator for HDMI audio were used.

The analysis was limited to 2-channel audio, as was the ASR review.

Our findings are arranged to be in the same order as those in the ASR report, and the two reports can be read side by side.

Findings and Commentary

HDMI Audio DAC measurements

First graphic – the ASR results are broadly correct in terms of distortion *at these settings.* ASR set the audio output to be 4V rms for a full scale digital input (0dBFS), so requiring a volume setting of 78, whereas the design output is intended to be less - around 2V rms – to match Arcam power amplifiers (volume 72 or 73). The AV40 uses a number of analogue multiplexers to route its signals internally and their distortion increases with signal level. There is a multiplexer on the L/R channels after the volume control and while this does not overload at 4V rms it does exhibit more distortion including a trace of high order harmonics. This is what he saw and in normal usage at around 2V maximum output this does not occur.

However, the noise floor observed by ASR should have been substantially lower and unchanging. We believe ASR fell into a rather subtle trap when performing their measurements. In the AV40 the audio inputs and outputs are not connected directly to chassis ground (unlike the digital sockets, which are, for EMC reasons). You should not connect these 2 grounds together or you get a ground loop in which some of the current driving noisy modules like WAN, Bluetooth, Avatar etc finds its way back to the power supply through the audio ground return path. This adds a low level “chattering” noise which appears on the audio outputs, irrespective of volume control setting.

Unfortunately, when you simultaneously connect an HDMI or coaxial digital signal to the AV40, from the same box that is used to analyse the audio, the connected screens in the audio cables can create an unwanted ground loop that adds noise. This does *not* occur in an audio system with class 2 (non grounded) chassis – whether downstream power amplifiers or upstream digital sources – so is not a problem in the field, only on the test bench.

Because the HDMI generator that we used (the Quantum Data) is a separate class 2 device and not part of an AP system, we did not have this issue except with co-axial S/PDIF (and took corrective steps in this case).

Our repeat of this test is shown below:

The red data is the exact test ASR did – a full scale (0dBFS) 1kHz digital signal via HDMI, using the AV40’s balanced outputs at volume 78. Output signal level was 3.83V rms. Only the right channel is shown as it has the worst distortion. This is the red data.

Firstly - note the *much* lower noise floor across the whole audio band as this test avoids the ground loop mentioned above which we believe occurred during the ASR measurements.

Otherwise the distortion spikes are comparable with ASR’s results. THD+N (total harmonic distortion plus noise) in my unit at 10Hz-22kHz analyser bandwidth was 0.00200% on this (R) channel, 0.00145% on the other (L) channel. These are materially lower than the ASR results and give an effective dynamic range of 94dB (R) and 97dB (L).

The green data drops the digital input level by just 1dB, leaving the volume at 78, so the output falls to 3.42V rms. This drops the THD+N to 0.00148% (R) and 0.00102% (L), increasing dynamic range to 97dB and 100dB respectively. Note also the disappearance of the high order harmonics on the right side of the graph.

The magenta data keeps the same -1dBFS digital input level and simply reduces the volume by 5dB from 78 to 73, so the output is 1.92V/channel. THD+N falls to 0.00099% (R) and 0.00091% (L) putting the dynamic range at over 100dB. This is inside ASR's green "good" range in their comparison chart.

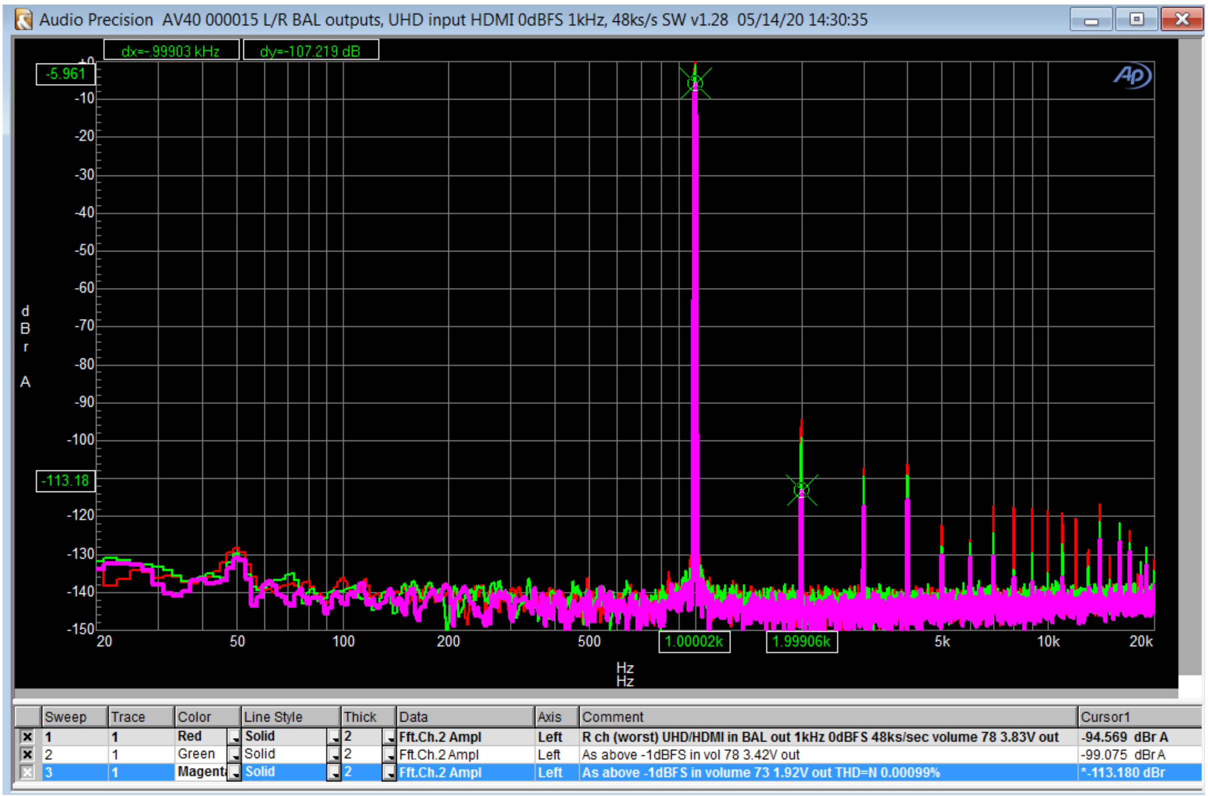

With reference to 0dBFS in and volume 78, ASR in the third chart reports the residual noise as -94.7dB or just under 16 bits. We have to assume this is an unweighted measurement which is severely distorted by an unwanted ground loop. This measurement is normally taken as A-weighted, which takes account of the ear's lack of sensitivity at low and high frequencies. It typically measures about 3dB better than white noise.

Using an A-weighted filter, with the digital signal generator muted, we measured 11.2uV rms at each balanced output, when full scale output was 3.83V rms. Unweighted in a 10Hz to 22kHz band the noise was 15.7uV rms

This equates to the noise floor being -110.6dB (A-weighted) or -107.7dB unweighted, referenced to full signal at 3.83V rms. This is 18.5 bits (A-weighted) or 18 bits unweighted, a decent result *and a full 13 to 16dB better than reported by ASR*.

Intermodulation distortion

The next ASR graph measures intermodulation distortion to the SMPTE spec, in which a strong low frequency tone is mixed with a weaker higher frequency one in a 4:1 ratio. In measurement these frequencies are removed and any others (the intermodulation products) remain as the residual, this is then graphed out. The ASR graph does not show which frequencies were chosen, so we used some default AP ones, 500Hz and 3kHz, and measuring the AV40 via Toslink (optical S/PDIF), we took the simple FFT shown below.

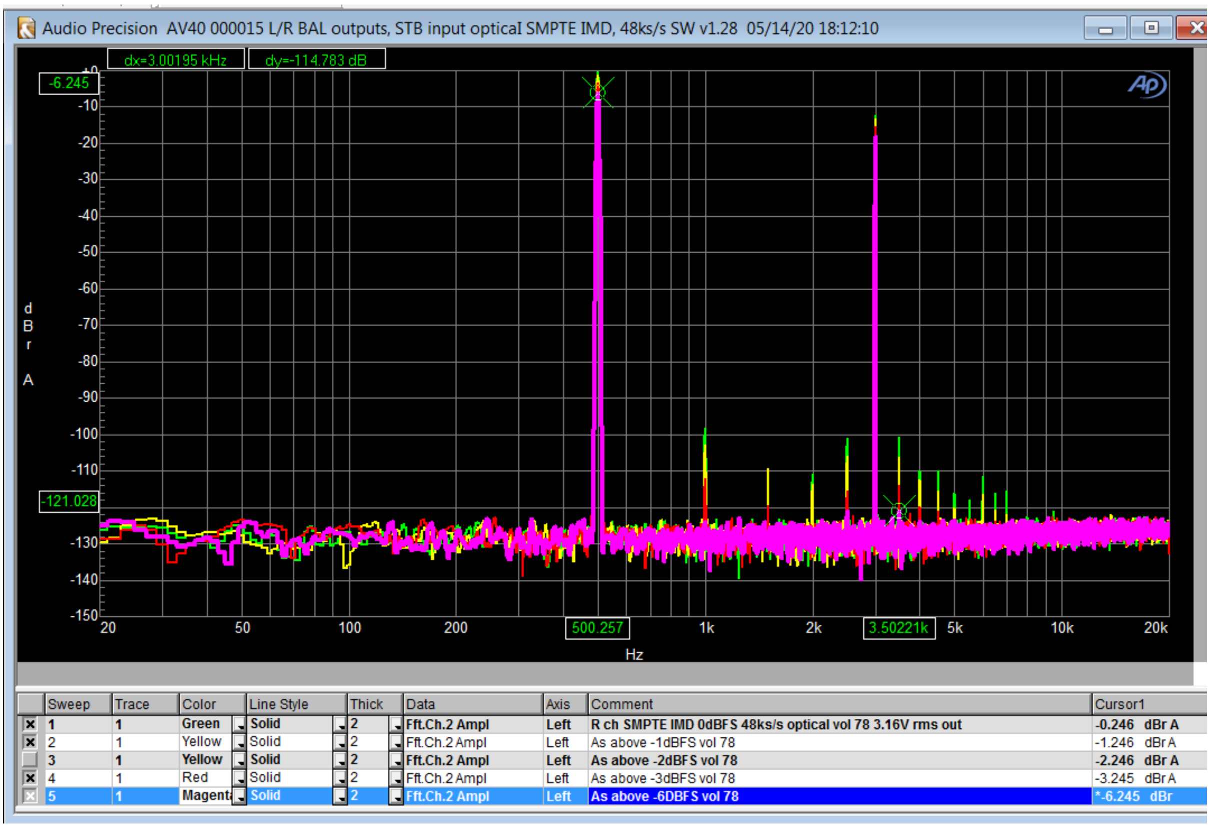

The input levels are 0dBFS (green), -1dBFS (yellow), -3dBFS (red) and -6dBFS (magenta). We calculated the distortion from the 3 largest residuals:1kHz, (second harmonic of 500Hz) and the 2 IM sidebands at 2.5kHz and 3.5kHz, using an rms sum. We ignored the noise floor at these high levels (*also, there is no ground loop!)*. It is clear all distortion products are in the noise at -6dBFS input.

The calculated distortion results are: 0dBFS -95.1dB (ASR -80dB)

-1dBFS -98.7dB (ASR -81dB)

-3dBFS -101.8dB (ASR -85dB)

-6dBFS -110dB (ASR -86dB)

In practice the -6dBFS input result was noise limited and was probably around -101dB, or -107dB ref full scale and the level of in-band unweighted noise floor. The results are comparable with the other product shown on the graph and some 15 – 17dB better than measured by ASR, *because there is no ground loop*.

Jitter measurements

Jitter is the result of minute timing errors in clocking a digital signal in a D/A converter. On a high resolution FFT it shows up as skirts, a broadening near the bottom of individual signal spikes. At high levels it can result in sonic degradation and a smearing of musical instruments in stereo space. The amount of jitter is directly proportional to an audio signal's frequency, so tests are often carried out at ¼ FS (12kHz in a 48ks/sec audio stream).

The jitter test involves superposing a 1-bit 250Hz low frequency signal at exactly digital zero, so that the data periodically transitions from 0111111..11 (1 zero and 23 1s in a 24-bit signal) to 10000000..00. If there are problems then this jump tends to show them up. In a properly designed system with 24-bit resolution these spectral spikes, at 250Hz and higher odd order harmonics, will be buried in the noise. In real world systems a plain 12kHz high level tone will give the same result.

At 16-bit resolution the spikes will however show up exactly as in the ASR graph. This implies that the AV40 is somehow limited to 16-bit resolution, as was clearly stated in a very critical manner in the ASR review.

Whilst it is possible that there is a bug in the DSP code that causes this, we were not able to reproduce this graph using the jitter test signal from the AP2700 when set to output 24-bit audio. Nor did we see any evidence in any other tests of such a serious bug.

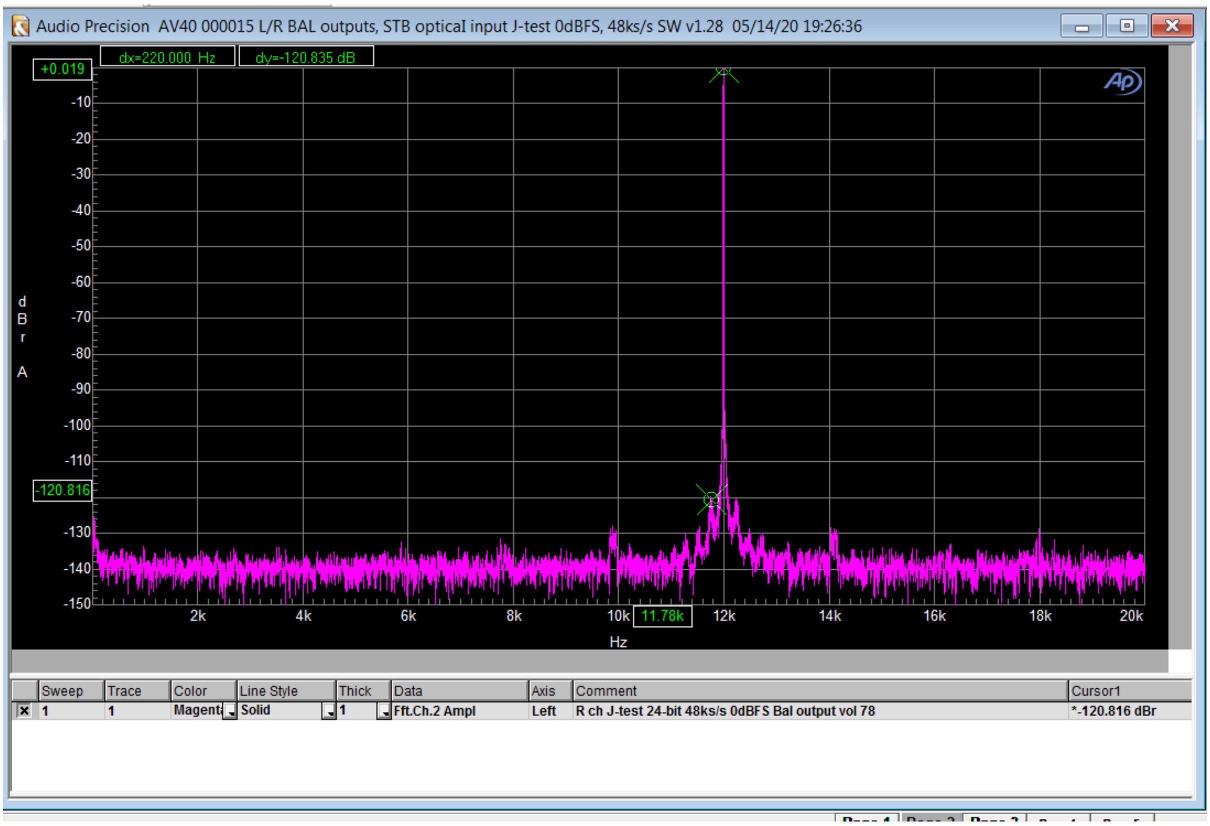

As an experiment we then changed the digital output of the AP2700 from 24-bit resolution to 16-bit and re-ran the jitter test as shown below; the 16-bit signal is shown in red and the proper 24-bit signal above is still occasionally visible as the magenta trace in the background.

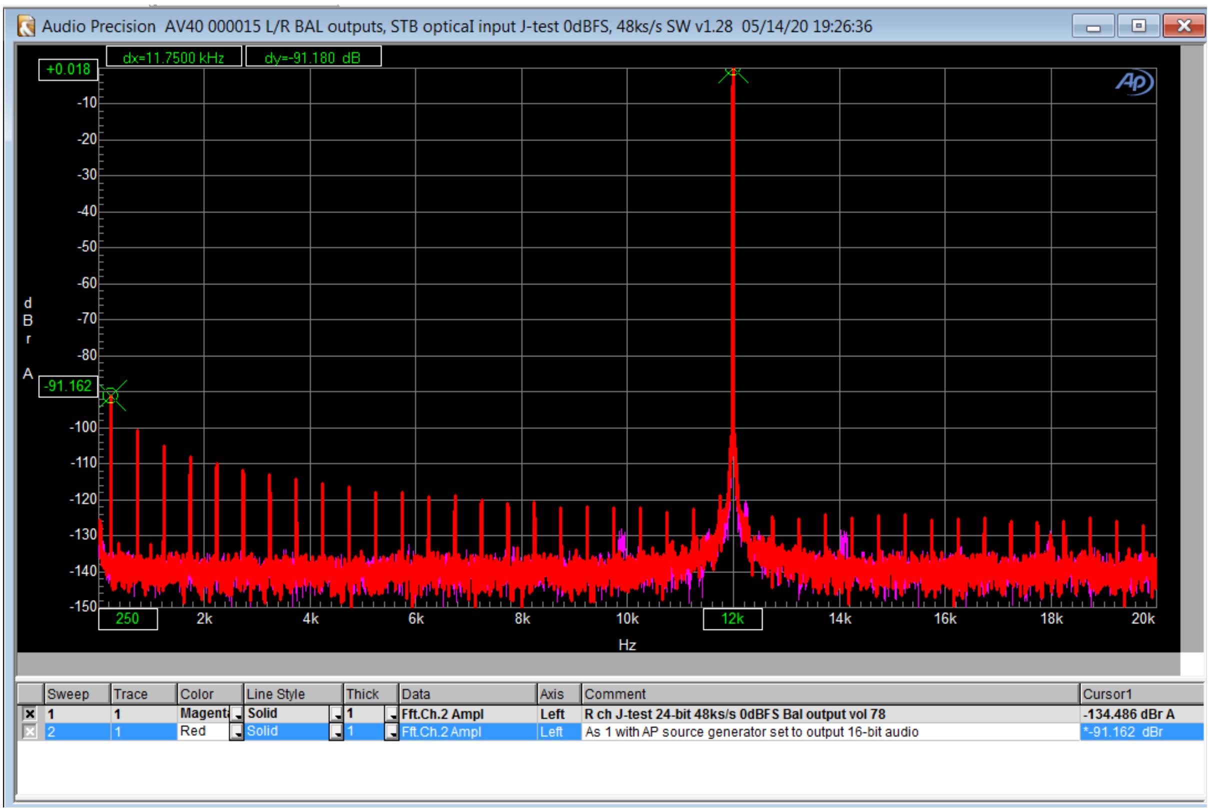

The 24-bit result is pretty good with low level close-in sidebands at about +/-220Hz at -120dBr – these will *not* be audible. The copious amounts of low frequency jitter using Toslink (optical S/PDIF) shown in the next ASR graph are completely absent here.

The 16-bit test is uncannily close to the ASR graph (which also shows increasing noise below about 5kHz – *the previously mentioned ground loop*). The left hand marker is at 250Hz – the toggle frequency of the jitter test – and it is the same amplitude as seen in the ASR graph.

The simplest explanation would be that the Audio Precision generator used by ASR was inadvertently set to 16-bits for this test. This would also explain the following test for DAC linearity, which implied the AV40 had less than 16 bits of resolution.

We ran the same test for DAC linearity but with the 1kHz source set to 24-bit audio. This version did not subtract the original signal, so the perfect result would be a straight-line graph going from (0dB, 0dB) in the top right-hand corner to (-130dB, -130dB) in the bottom left-hand corner.

The red curve was with a 22Hz-22kHz filter in the analyser; below about -105dBFS input internally generated white noise in the AV40 started to intrude on the output.

The green graph inserted a 1kHz bandpass filter to remove most of this noise – good linearity was now visible to about -120dB (20 bits) before noise took over.

DAC reconstruction filter properties

ASR shows a graph of the noise output from the digital path between 18 and 30kHz.

We were able to reproduce this very closely.

The point is that the AV40 uses a minimum phase FIR (finite impulse response) digital reconstruction filter for reasons of sound quality. Unlike normal IIR (infinite impulse response) linear phase filters, this one has no pre-ringing in its impulse response. However its out-of-band attenuation is not as good as a typical IIR filter, meaning its image rejection is also not so strong. (The digital image reflection of say 10kHz at a 48ks/sec sample rate is 24+(24-10) kHz = 38kHz, plus another at 24+(24+10) kHz = 58kHz). These can be observed in later graphs.

With the AV40's filter these out of band signals are not audible and are at a sufficiently low level to cause zero problems downstream. However if you then put high level high frequency signals through the processor and measure the "distortion" with a wide band (80-90kHz) analyser pass band, you will of course record these ultrasonic image signals as distortion. At higher sample rates (88.2ks/s and up) the effect largely disappears.

This is what the next ASR graph shows, using a sample rate of 48ks/s. It is correct but, in our opinion, meaningless in the real world. We ran the same 0dBFS frequency sweep with 80kHz, 30kHz and 22kHz low pass filters – you can see the ultrasonic “distortion” associated with high frequency input signals largely go away.

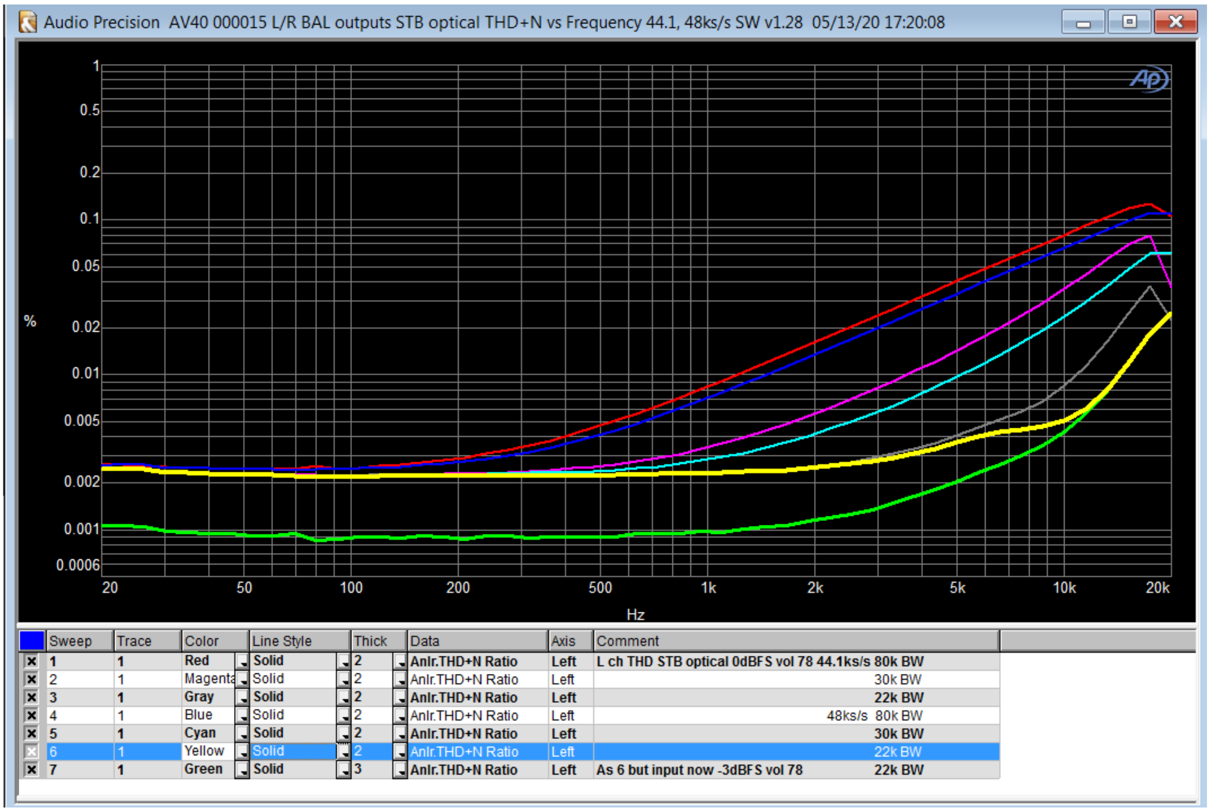

Sweep	Trace	Color	Line Style	Thick	Data	Axis	Comment
1	1	Red	Solid	2	Anlr.THD+N Ratio	Left	L ch THD STB optical 0dBFS vol 78 44.1ks/s 80k BW
2	1	Magenta	Solid	2	Anlr.THD+N Ratio	Left	30k BW
3	1	Gray	Solid	2	Anlr.THD+N Ratio	Left	22k BW
4	1	Blue	Solid	2	Anlr.THD+N Ratio	Left	48ks/s 80k BW
5	1	Cyan	Solid	2	Anlr.THD+N Ratio	Left	30k BW
6	1	Yellow	Solid	2	Anlr.THD+N Ratio	Left	22k BW
7	1	Green	Solid	3	Anlr.THD+N Ratio	Left	As 6 but input now -3dBFS vol 78 22k BW

We first ran the test at full level (0dBFS and volume 78) at both 44.1ks/sec and 48ks/sec with an 80kHz low pass filter – the red and blue lines respectively. The blue line is comparable with the ASR graph. The magenta and cyan lines are with a 30kHz filter, which will cut off harmonic distortion above 10-15kHz fundamentals but these are not the issue here – at high frequencies the digital images are the dominant factor.

The grey and yellow lines are with the 22kHz filter and show more attenuation of the digital images – at up to 6kHz residual harmonic distortion is still captured and is around 0.004% - an adequate result.

The green final curve is at 48ks/s 22kHz with the input signal reduced to -3dBFS. Mid-band distortion plus noise (THD+N) drops to 0.001% which is a very good result.

HDMI 10kHz 0DBFS wideband spectrum analysis to 96kHz

This is shown in the next ASR graph.

The fundamental is 10kHz at 0dBr (reference level); its harmonics at 20kHz and 30kHz are visible at -88dBr, with more going out to 90kHz at about -110dBr. The 38kHz and 58kHz signals are digital images described above, at -67dBr. Given the 0dBFS input and the high volume setting of 78 (4V rms output), all this is expected.

The other signals in the audio passband should not be there, nor the uplift in noise below 5kHz (most likely the previously mentioned ground loop). We were not able to reproduce these when I duplicated the test using the Quantum Data generator to drive the AV40's UHD HDMI input.

Our results are shown below:

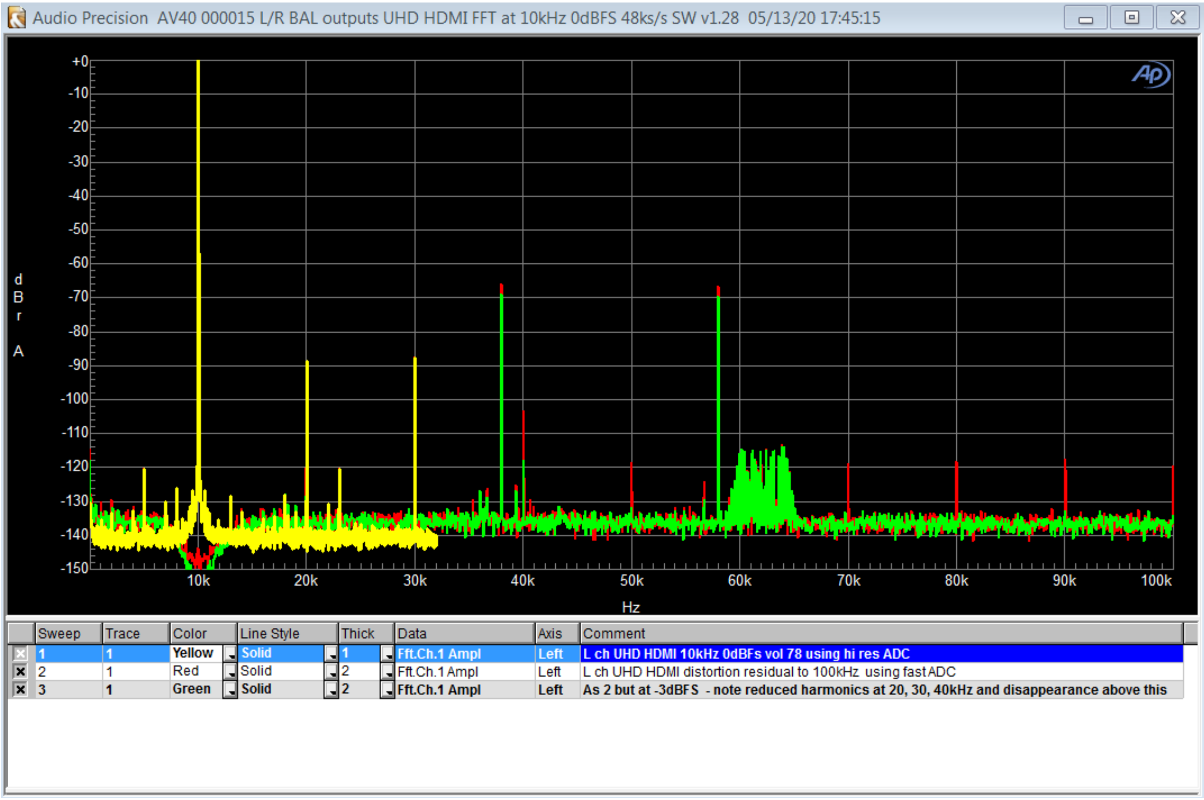

Some explanation is necessary. The AP2700 can only perform a high resolution FFT up to 32kHz; above that a medium resolution ADC is available to over 100kHz. By using the 2700's distortion analyser to remove the 10kHz fundamental it is then possible to "see" out to 100kHz with only a slightly higher noise floor.

The yellow line is the high resolution one – none of the spurious spikes in the ASR graph are present. There are 2 fixed tones at -120dBr, which are probably related to the digital filter but will be inaudible. Everything else in the <20kHz audio band is below -120dBr, including jitter sidebands.

The red line extends out to 100kHz. The splash at 60-70kHz is a switched mode power supply in the AV40 – at the level shown it is utterly innocuous. The higher order harmonics are visible at -120dBr out to 100kHz – again irrelevant.

The green line is where the input signal has been reduced from 0dBFS to -3dBFS. The high order distortion harmonics vanish into the noise floor and the second and third harmonics fall by about 8dB in absolute terms or 5dB in relative terms, so the THD almost halves.

Conclusion – Good performance as expected.

32 tone test signal

This is a complex custom signal generated inside the AP's DSP and designed to enable quick frequency response testing in a production environment, rather than for high resolution audio analysis. The AP allows you to program in any set of tones you like, in any quantity.

ASR claim its resembles "music" but it doesn't – there is far too much high frequency energy and the signal is simply not found in real life. Even a square wave (100% clipping of a sine wave) has proportionately less HF energy. We believe ASR's result - the digital filter in a FIR minimum phase fast roll off filter will cause a square wave to ring strongly and the AV40's DAC is set up to just avoid clipping this. It is perfectly possible this artificial signal will clip the DAC. It is marginal because backing off the signal 1dB is sufficient to restore normality.

Highlighting the clipped result as a problem seems to be a disproportionate response.

Dual tone of the same test

This is known as the CCIF intermodulation test and is very powerful. Two closely spaced high frequency tones, each at -6dBFS (i.e. reaching up to 0dBFS when the peaks add up in phase), are played through the system and the result displayed as an FFT. Second order non-linearity in the amplifier (which generates even harmonic distortion) will cause the HF tones to "beat" or intermodulate and output the frequency difference between them. Therefore 19 and 20kHz will produce a 1kHz beat tone. 4^{th} order non-linearity will generate 2kHz and so on.

Third order non-linearity will generate sidebands adjacent to the main tones, i.e. 18kHz and 21kHz in this case. Similarly a 5^{th} order non-linearity will cause tones at 17kHz and 22kHz, and so on.

The ASR curve is not very good. We ran this test and the result was pretty clean.

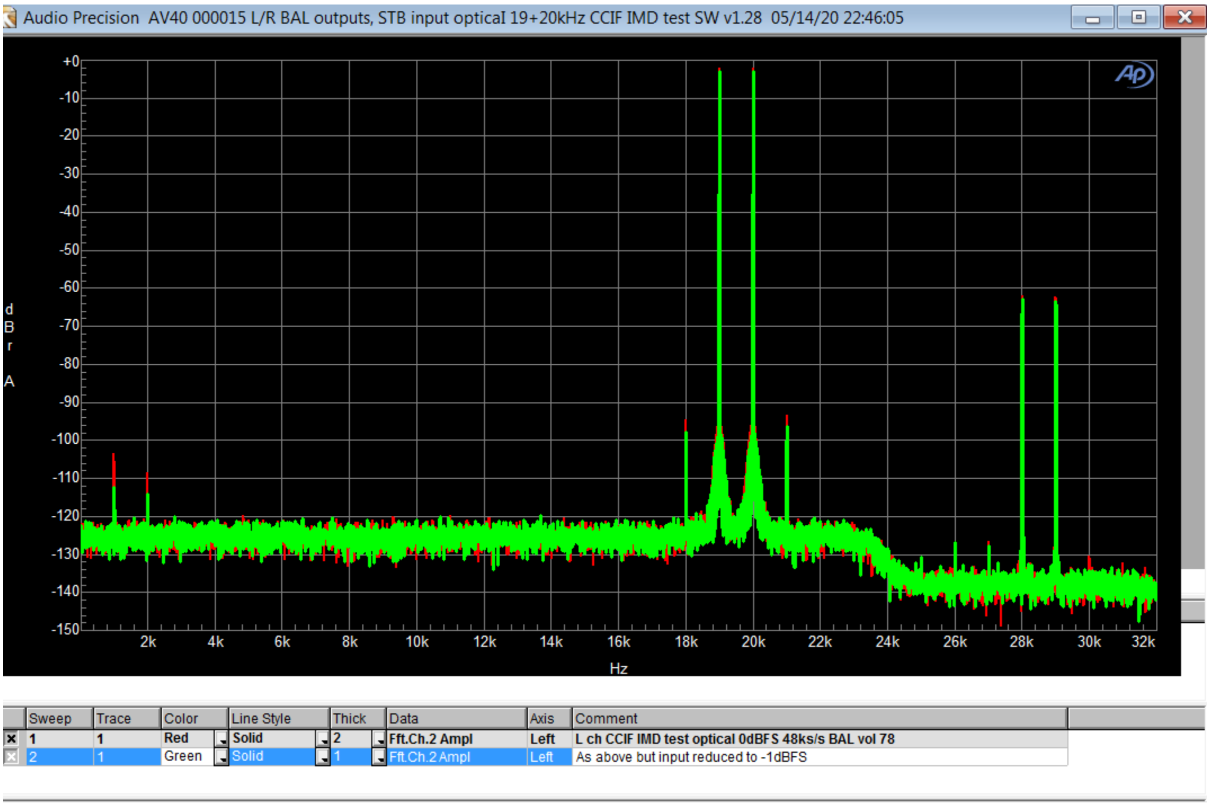

The red curve (mostly hidden) is the output at 0dBFS digital input; the green curve shows the effect of lowering the digital input signal by 1dB to -1dBFS. Note the considerable reduction in the 1kHz and 2kHz even order intermodulation products.

Analogue Input Audio Measurements

Whilst the AV40 will accept inputs of up to 4.5V rms on its analogue inputs without clipping, it is optimised for inputs of around 2V rms. This is the standard maximum output level for CD players for example.

ASR's graph/dashboard reflects his choice to feed 4V rms into the CD input of the AV40, which starts to stretch the analogue multiplexers. Such a high signal is seldom if ever encountered in normal audio systems.

We ran a measurement at 2V rms in and 1.91V rms (volume 72) out using the balanced outputs. The cursors show the second harmonic is -105dBr on the R channel (green) and -111dBr on the L channel (red). There are no high order harmonics.

THD+N is 0.00081% on the L channel and 0.00091% on the R channel with 22kHz analyser bandwidth, about 1/5th of the ASR figures. The noise floor is substantially less than in the ASR graph as there is no ground loop.

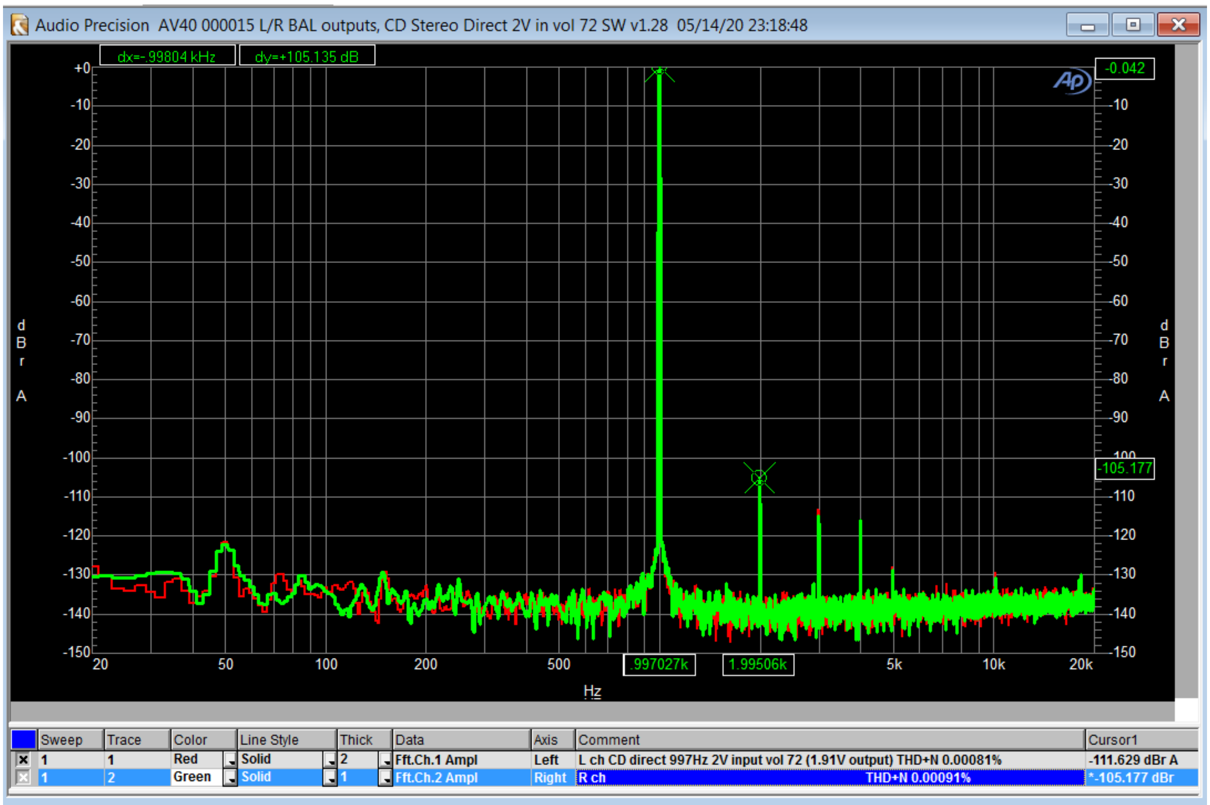

Running a THD+N vs frequency curve at these settings (2V in 1.9V out) produces the next graph, shown on the following page.

0.002% THD+N equates to -94dBr, about 7dB better than the ASR graph, even though the signal is 6dB lower. The dominant factor here is the noise floor of the balanced output amplifiers as mentioned earlier.

We also then limited the analyser bandwidth to 22kHz, which drops the midband THD+N to 0.0009% or -101dBr on the worst (R) channel. This would be about -106dBr (L) and -104dBr (R) at the unbalanced outputs because of the lower noise floor. Despite the complexity of the switching matrix in the AV40 preamp, these figures are approaching those found in many examples of high end stereo gear.

The residual noise floors of the outputs gave an A-weighted dynamic range of about 106dB for the balanced outputs (ref 2V out) and 112dB for the unbalanced outputs (ref 1V). This can be further improved by raising the volume control by up to 6dB or so, at the expense of a little more peak level distortion.

Taken overall, these results mean the AV40 meets its published specifications.

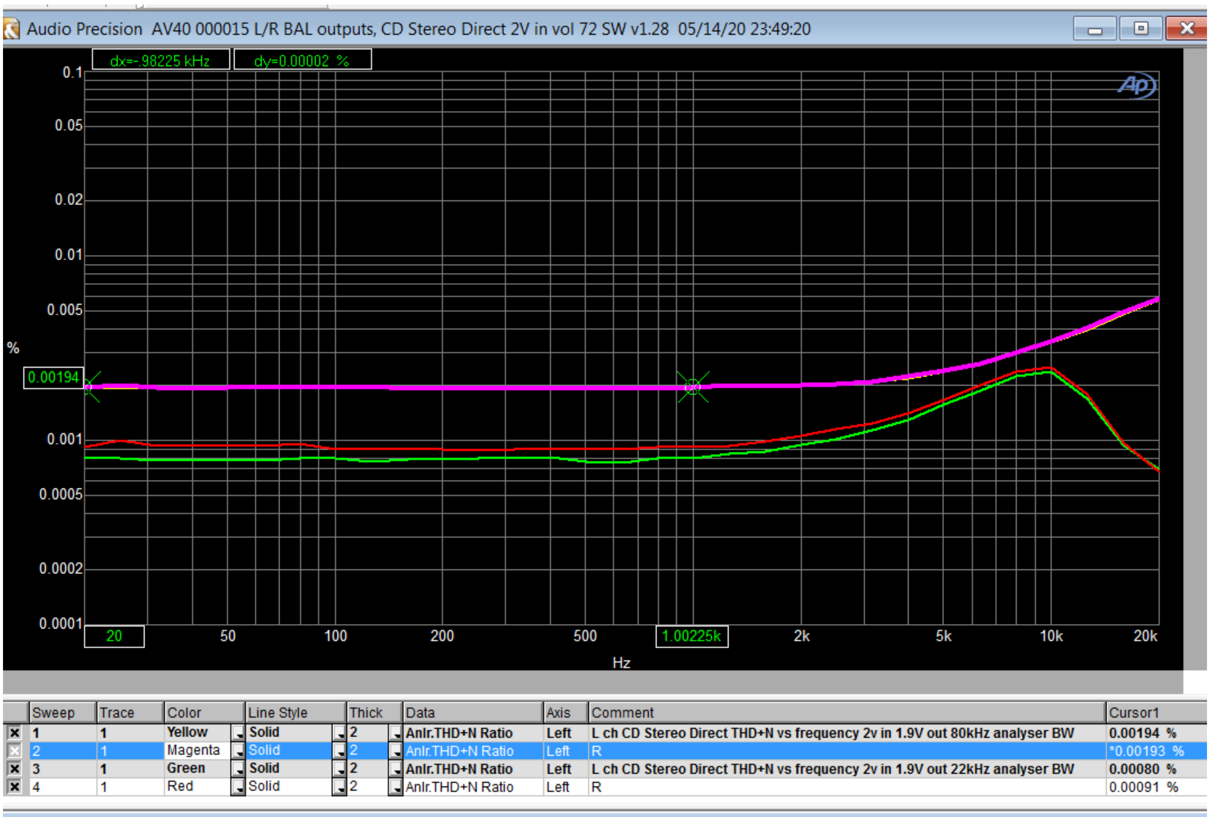

Conclusion

It was disappointing to see the manner in which Audio Science Review published data on the AV40 it reviewed.

Unfortunately, it is not easy to carry out reliable measurements on complex AV type products without having a thorough understanding of the product's architecture, design goals and limitations. We would have welcomed an approach by ASR prior to publication for the purposes of ensuring technical accuracy, as is the case with most reviewers. It is not in any way "underhand" for a reviewer to check their findings ahead of publication with a manufacturer and for the manufacturer to correct or show a valid explanation (as we believe we have done here).