

New Approaches To Audio Measurement... Or, some measurements that matter – and why some don't!

Introduction And Overview

In late September, at the UK Hi-Fi Show (Whittlebury Hall) and a week later at the Rocky Mountain Audio Fest, an unusual presentation was unveiled to the public and industry alike. Making use of both audio demonstration and measured results, it showed the results of a recent research project that sought a new and more relevant approach to audio measurement. But it was something else that made this event unique: the results being presented were the product of a combined effort, a shared, cooperative venture between two competing hi-fi companies and a third, independent research company with a background in advanced signal processing as applied to defense projects. It's an unusual story, but more than that, understanding the central players and their specific contributions is critical to appreciating the results, how they were arrived at and their wider significance. Like so many stories, it started as one thing but ended up as entirely another...

Early days and starting out...

This story starts at two hi-fi shows, one in Manchester, the other in Denver. As editor of Hi-Fi+ magazine, I was conducting a series of public demonstrations that showed the importance of cables and racks to the performance of hi-fi systems. Each demonstration consisted of three identical sets of electronics, arranged so that they could drive the same speakers: the only differences between the systems were the power and signal cables, equipment supports and the use of a Quantum Qx4 power purifier. Supplying the cabling for the third (and best sounding) system, as well as helping with the demonstration itself, was Steve Elford of Vertex AQ. But, as impressive as these demonstrations were (and you can read a detailed account in the associated download) the real Eureka moment came at the end of a day – we just didn't realize it at the time. Sitting around and musing on the day's events (and resting up a bit – doing these involved demos can be pretty exhausting) we mused on the fact that, as big and as obvious and as musically important as the differences we'd just been demonstrating were, no one had yet managed to measure them successfully – a stunning indictment of the current state of audio measurement, as well as its focus.

Such observations are normally confined to the realm of fantasy or "What if..." but this one lodged in fertile ground, because Steve Elford as well as running the audio arm of Vertex AQ, is also an ex-RAF technical officer with a military consulting role. Some months later, working on a defense project, he came across Dr Gareth Humphries-Jones of Acuity Products, a Doctor of Applied Mathematics and a specialist in signal processing algorithms as applied to high performance sonar and radar systems – and a man with a problem: the latest sonar system on which he was working was failing to meet spec. At which point Steve pointed

out the performance benefits in hi-fi terms of cables and equipment supports. Now, whilst it's fair to say that Gareth was skeptical to say the least, he agreed to attend an upcoming demonstration, similar to the ones we'd carried out in Manchester and Denver. Even as a non-audiophile the results were so obvious that they piqued his interest and suddenly our "What if..." had grown a set of academically and professionally impeccable researchers with considerable experience and a whole toolbox of advanced techniques developed in a parallel field. In the meantime, I had left the magazine and joined Nordost with responsibility for their marketing. Put all those things together and you get a unique convergence of opportunity: not only do you now have not one but two hi-fi companies with an interest in developing a measurement protocol capable of reflecting previously unmeasurable effects, but you have a separate and independent research body with their own interest in doing so, as well as the most unusual factor of all, the link between them.

First Steps...

From the start we wanted Acuity to undertake independent research, unfettered by the preconceptions that govern the hi-fi industry. With that in mind we simply supplied them with an initial test rig designed to allow them to examine the effects that Gareth had already experienced in the demonstration that he'd witnessed. This test-rig consisted of the following items:

- 1x CD Player (an older UK built model which cost about £4000 when new)
- 1x 3-shelf Quadraspire rack
- 1x Vertex AQ Roirama AC cord
- 1x Nordost Vishnu AC cord
- 1x Vertex AQ Kinibalu Platform
- 1x Quantum Qx4 Power Purifier


We also supplied a second Vertex power lead (without acoustic absorption) and another Nordost Vishnu to be used between the Qx4 and the CD player.

Initially, Acuity set out to discover whether there was any measurable difference between the performance of the player simply sat on the Quadraspire rack and once an audiophile mains lead, equipment support and Quantum unit were employed. With no predetermined path they approached the problem very much as they would a sonar system – but with one significant difference: in this instance the use of a CD player and the possibility of bit perfect copying of disc content meant that they had access to what they term "truth data", a known input signal, rather than a whole ocean of sound.


The initial approach was deceptive in its simplicity – essentially a basic sum and difference technique – but with no audio prejudice to lean on it also adopts a radically different perspective on the problem. Traditional audio measurement approaches are based in the frequency domain, using FFT technology and the steady state test tones it relies on. Acuity, with their complex signal processing background looked straight at the time domain, the area which is most revealing when it comes to the systems they normally work with. What’s more, given their status as truth data, far from being a daunting prospect, the use of actual musical signals as test material was seen as significantly simplifying the problem.

The test rig consisted of the CD player, sat on the top shelf of the Quadraspire rack and power using a standard IEC AC cord. A bit perfect copy of a particularly energetic musical passage (on the basis that it’s exactly this sort of signal that causes hi-fi systems most problems) was made on a PC and then the disc was replayed on the CD player, the same passage fed from the analogue outputs of the player, back into the PC, via a high quality sound card, where it could be compared to the original data. With both sets of data in the PC, they could now be compared, allowing Acuity a window onto the performance of the CD player. By overlaying one signal on the other and subtracting it, any deviation from the original caused by the replay process will be revealed in the residual error that remains. It sounds simple, but in practice the accurate alignment and subtraction of those signals requires considerable expertise – which is exactly where Acuity’s extensive experience and signal processing “toolbox” comes in.


Reproduced below, you’ll see two traces, one generated by the original data and the other by the same track replayed through the CD player. Initially they look very similar – exactly as they should. But once overlaid, the differences are actually significant, as represented by the third trace. This is reproduced to a larger scale in order to give you a clearer view of what is happening, but in numerical terms it represents peak error levels of around 10%. No wonder our systems don’t sound much like live music!



Original wav file track

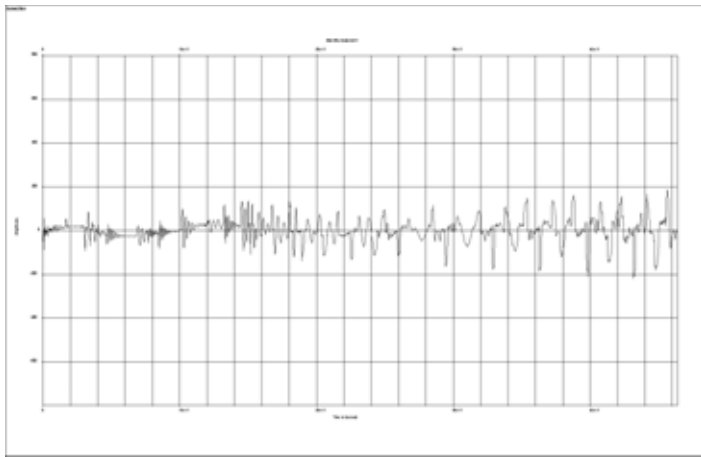


Signal from analog output of test player



An actual 'difference' trace - wav minus untreated player

What Acuity did next was to place the player on the Vertex platform and hook it up using the audiophile power cords and the Quantum Qx4, just to see if these items affected the output signal. Repeating the process outlined above produced the new error plot seen below. Compare that to the original output and you can see something approaching a 50% reduction in gross error! Clearly, these things are materially impacting on the performance of the CD player. In fact, so much so that Acuity initially questioned the validity of the results, but repeated testing showed them to be completely consistent.




An actual 'difference' trace - wav minus player with platform, AC cord and Qx4

The significance of this result should be neither over nor underestimated: on the one hand, break down the test into individual steps and you get the first unequivocal data that demonstrates the effect of AC power cords, equipment supports and Quantum QRT technology; on the other, the fact that there is a difference doesn't mean that that difference is significant. Further investigation was clearly required, even if just showing a measurable effect from cable and support accessories that are actually outside the signal path is itself a significant step forward.


Re-examining The Frequency Domain

Whilst Acuity were confident that it was the Time Domain that would prove the most fruitful source for investigation, they didn't want to neglect the Frequency Domain. With specific, audible effects to examine, they were confident that meaningful and helpful results could be obtained. One of the things that we'd discussed with them in terms of system performance indicators, was the ability of better systems to hold low-level detail such as small percussion motifs, apart in an otherwise thunderous orchestral crescendo.


With this in mind, Gareth created a small, synthetic signal consisting of a 12kHz peak and a series of "harmonics" at 1K intervals, each descending in level. Repeating the test protocol once again, but looking at the results from the synthetic "ting" in the frequency domain, you'll see clear differences between the performance of the player with and without the accessories.




Test wav file




Untreated player output



Output with Platform, AC cord and Qx4




Untreated player output



Output with Platform, AC cord and Qx4

In particular, the noise floor below and especially above the test tone has dropped significantly, while the “skirts” between the peaks have also dropped. Less obvious is a small but significant increase in the peak values themselves. In fact, if we represent these changes graphically in a simplified form they’d look rather like this.



What has actually happened here is that the area under the graph has remained constant while the energy it represents has become more focused into the peaks themselves, tracing the original signal much more accurately. Look at the associated table (which breaks down the effects by item) and you can see that accurate measurements of the actual graphs display a 2.5dB drop in noise-floor below the test-tone, a 5dB drop above it and around a 0.5 dB increase in peak values (not insignificant given that this is measured on a logarithmic scale). But what is even more interesting is just how closely these results actually mirror the audible effects of the accessories in use. Compare the CD player’s performance, with and without those accessories and what you hear when they’re in use is an increase in focus and presence to vocals, greater dynamic range, a lower noise-floor, more space around and between instruments and greater low-level separation and detail – all of which makes perfect sense of the measured results. This close correlation is extremely unusual in the field of audio measurement and also easily demonstrable, as our recent seminars at the UK and RMAF shows revealed. And of course, once again, we’re showing that AC cords, equipment supports and AC quality have both an audible and measurable effect on audio performance.

Configuration	Typical noise reduction below group (dB)	Typical noise reduction above group (dB)	Typical increase in peak values (dB)
Averaged Player with platform	-1.1	-2.2	+0.3
Averaged Player with AC power cord	-0.9	-1.8	+0.2
Averaged Player with Qx2	-0.5	-1.5	+0.2
Averaged Player with platform, AC power cord and Qx2	-2.5	-5.0	+0.4




Once More Into The Breach... Or, back to the Time Domain

The results from the Frequency Domain experiments had proved interesting in terms of what was happening to the signal, but cast little light on why. Faced with this conundrum, Acuity tried various different approaches to unravel the mechanism at work, but it was not until they returned to the time domain, but this time on a sample-by-sample basis, that they started to make real progress. If you compare the test signal with the error signal, you might well expect the peak errors to coincide with the peak levels in that signal – but they don't. In fact, there is a correlation, but it's to the rate of change in the signal.

If you take a single peak or trough from the time domain printout of the original signal in Figure 1, and zoom right in on it so that you can see each sample individually (each "corner" in the plot) it would look something like this:




This example shows 12 samples with the standard 22 microsecond interval of Red Book CD.

Take the signal from the player's analog output and overlay it and the result should be a perfect facsimile, but displaced slightly to the right by the time taken for signal itself to pass through the player. If you removed this group delay (effectively dragging the plot back to the left) the original signal and the player's output should overlay almost perfectly.



MAKING THE CONNECTION


Take an actual example and at first glance the two curves do look pretty similar. But Acuity wrote a complex algorithm that allows them to precisely time-align individual samples in a data stream (one of their areas of particular expertise, as already mentioned). Once the two plots are overlaid it becomes clear that the player's output differs quite significantly from the original signal, leading it at some points, lagging at others. It is also obvious that, where the rate of change in the signal is low (the plot is "flat") the two curves line up pretty well. But as soon as the gradients increase, so does the error.



MAKING THE CONNECTION

Running a second algorithm allows Acuity to measure those errors (plus and minus) against time, producing a readout that looks like this:


Here, each red line represents an error, the longer it is, the greater that error. If the player's output matched the input signal, the result would be a single horizontal line but as you can see, in the case of our test player, the deviation is significant. How significant? Well, the vertical graduations on the plot are 10 microseconds each, and as you can see, there are plenty of examples in which peak-to-peak errors of over 40 microseconds (or two complete sample periods) occur!



Results In The Real World... Measurements you can hear!



With this new set of digital tools, we could at last look much more closely at what was happening to the performance of the player as Acuity changed the components in the test rig. By now, the results were sufficiently clear to actually demonstrate them in public, which was first done at the Whittlebury Hall hi-fi show in the UK (at the end of September 2009, and again at RMAF in Denver, a week later – although in the latter case UPS failed to deliver the actual test player to the show, meaning that we had to use a stand-in). The time taken to run the analytical algorithms makes real-time analysis impossible, but in order to show visitors the potential in this new approach to audio measurement, we devised the following demonstration.

Firstly, we borrowed a current machine of good quality and tested it in exactly the way described above. Thanks are due to DCS for the extended loan of a Puccini player, a machine that sells in the UK for a little under £10,000. The plot on the right, shows the measured results for that player – and you'll note that even with the stock power cord and support, the results are already significantly better than for the original test machine, peak-to-peak errors generally being well below the 20 microsecond mark.




MAKING THE CONNECTION

Then we started adding the power cord, support and power purifier, measuring the effect in each case. But rather than simply showing people the printouts, we took the exact same player, along with the test rig of cables, Kinibalu platform and Quantum unit to the show – and hooked them up to a system. That way, when we added the power cord, visitors could hear for themselves exactly what happened to the sound – and we could show them the results, graphically, at the same time, projected onto a screen. So, adding the power cord produced a significant reduction in overall error:




In fact, that reduction in misplaced samples amounts to 36%. Repeating the process with the Kinibalu support reduces the error by a further 15%, although bear in mind that this result is cumulative – the Kinibalu had less to work on because the power cord had already improved things.



MAKING THE CONNECTION

Finally, we switched on the Quantum unit and once again, we were able to measure a further decrease in error – this time, another 11%, resulting in an overall reduction in timing error with all three upgrades in place, of 52%.




Of course, the beauty of this demonstration is that rather than simply telling people what was happening, we were showing them – in musical terms – just how important these changes are. And naturally enough, having gone all the way up the ladder of improvements, one step at a time, we then went right back to the starting point – all in one go, which was pretty sobering. The sonic changes made at each point in this process were both musically significant and perfectly apparent.


Wider Implications...

So far, this testing and research effort has been limited solely to the realm of CD replay. We have not looked at high-definition media, analog playback or systems as a whole, embracing amplifiers and speakers. The next stage will be to incorporate amplifiers and the signal cabling (interconnects and speaker cables) into the test rig, although each variable will make the process far more complex and difficult. However, we can already report certain significant implications as a result of this work.


Of course, any such test regime can also form a basis for comparison and with that in mind we tested a third machine on the basic rig with a standard power cord.



Older UK-built player, ~£4,000.00



Current Japanese player, £250.00



dCS Puccini, £9,500.00

MAKING THE CONNECTION


The error plots for these three machines make interesting reading. On the left is our aging audiophile player, a machine that ticks all the audiophile boxes when it comes to solid casework and substantial power supplies. On the right is the dCS Puccini. But in between is a current mainstream player, built into a flimsy, bent metal case and with a retail price of £250. And yes, it doesn't just measure better than the player on the left, it sounds better too. Which starts to suggest that the emphasis we place on certain aspects of physical and electrical design might need to be reexamined. It's not that casework or power supplies are unimportant – but the way they are executed clearly is. What is potentially even more interesting is what happens to the error plot of that £250 player if you put a decent power cord, support and Quantum unit on it... Yes, Acuity are working on that and we'll have results soon.

But even more significant is a second relationship that Acuity discovered within the time variance errors. The plot below shows a close up section of the error readings for the test player. This is plotted in Blue.

The Red plot is the same section of signal, once the power cord, platform and Quantum have been added. What is remarkable about this result is the extent to which the two traces correlate. Whilst they do not overlay or trace each other precisely, the peaks correspond to a remarkable extent. What is more, this result is repeatable, the same musical extract played in the same machine and under the same operating parameters, gives exactly the same results every time you run the test, irrespective of geographical location, even when the tests are conducted months apart.

That means that the mechanism responsible for the errors cannot possibly be random, removing jitter in its various forms from the equation. Let's just repeat that, because as soon as you talk about timing errors, the audio community immediately thinks "Jitter": The errors we are measuring cannot be attributed to jitter, because they are not random and jitter is! Instead, they are related to the load placed on the system, the dynamic demands imposed by the musical signal. And that means that any test regime that relies on test tones as opposed to musical signals, won't see these errors. Yes, a sine wave varies, but it has nothing like the dynamic or frequency complexity that real music imposes – and that means that this is a major distortion mechanism that has remained all but invisible to existing measurement techniques.

As we suspected, we are far from the only people working in this realm, and a number of companies came forward after the demonstrations, informing us that our results echo theirs, achieved through different approaches. Whilst many of us have been aware that there are structural issues with the way that systems reproduce music, our listening experiences telling us what makes a difference if not why, it now seems that much greater emphasis on the time domain will be critical in advancing our understanding of the way in which electronics and mechanics effect the musical capabilities of our systems – and the way in which inadequacies in existing measurement techniques have held us back. For years we've been told that if we can't measure it we must be imagining it. The time has come to rewrite that dictum; if we can hear it but we can't measure it, then we should be thinking about getting some better measurements.



Nordost Corporation 200 Homer Avenue Ashland MA 01721 USA

Tel: +1 508 881 1116 Fax: +1 508 881 6444 Email: info@nordost.com Website: www.nordost.com