

KNOWLEDGE ALLIANCE BRIEFING PAPER

NEW AUDIO MEASUREMENT TECHNIQUES AND THEIR IMPLICATIONS

03 FEBRUARY 2011

Introduction

Background

Since mid 2008 a collaborative project has been underway to examine the possibilities of producing a new measurement strategy, processes and tools for the measurement of high-fidelity audio systems. This collaborative work is being carried out jointly by 3 companies; Vertex AQ Ltd and Acuity Products Ltd, both based in the UK, and Nordost Corporation based in the USA.

The basis of this work is the premise that high-quality audio accessories make a clearly demonstrable improvement to the sound quality of hi-fi, but conventional measurement techniques fail to identify the changes. The object of the exercise was to explore new approaches to the measurement of audio using defence systems processes - to discover whether these processes could produce measured evidence of the sonic impact of accessories.

Furthermore, if such defence systems techniques could be used to reliably measure changes brought about by accessories, the implications could be significant for our understanding of audio systems in much broader contexts, such as the often considerable differences between conventional measurement figures for electronics and the reality of the sonic qualities in listening tests.

The National Audio Show (UK) and Rocky Mountain Audio Fest (USA) Presentations 2010.

At NAS and RMAF 2010 we presented the latest position on the measurement project. But those presentations were by necessity simplified down to a 40 minute PowerPoint presentation and Q & A session; all in a show context.

This document is an expanded version of the presentation. It contains more written explanation of the work carried out within the project. It sums up both the two distinct stages in development to date, as well as our latest thinking regarding the real-time error mechanisms present in hi-fi systems. Finally, it also discusses the possibility of producing a widely available, software-based measurement tool using this approach.

Phase 1 - the early work - 2008/09

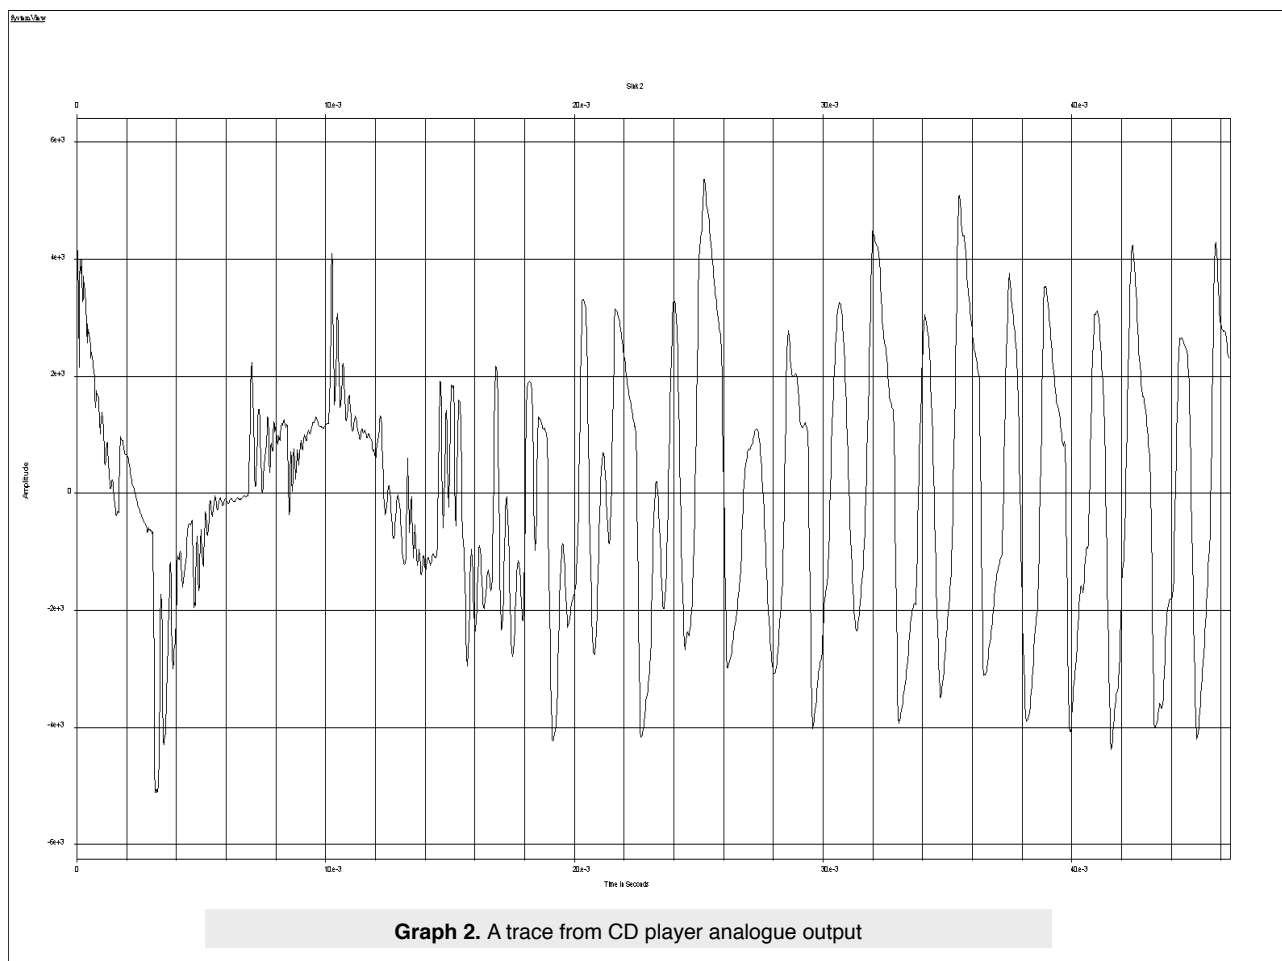
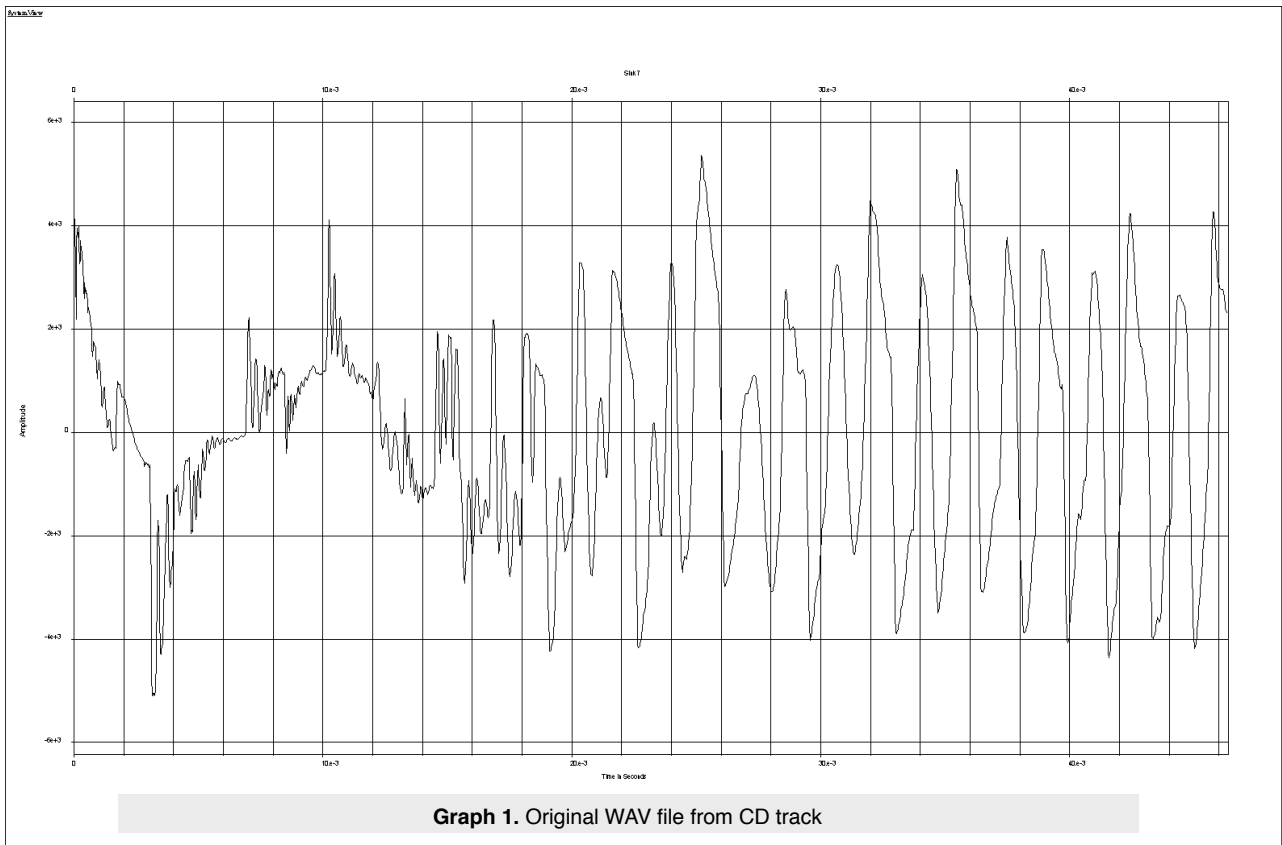
After initial agreement between Nordost, Vertex AQ and Acuity Products to commence this work, Acuity were provided with 3 integrated CD players and some Vertex and Nordost products, and asked to see if they could detect any difference with the CD player output using the specialist products over stock items (mains leads, supports and Quantum QRT were used at this early stage).

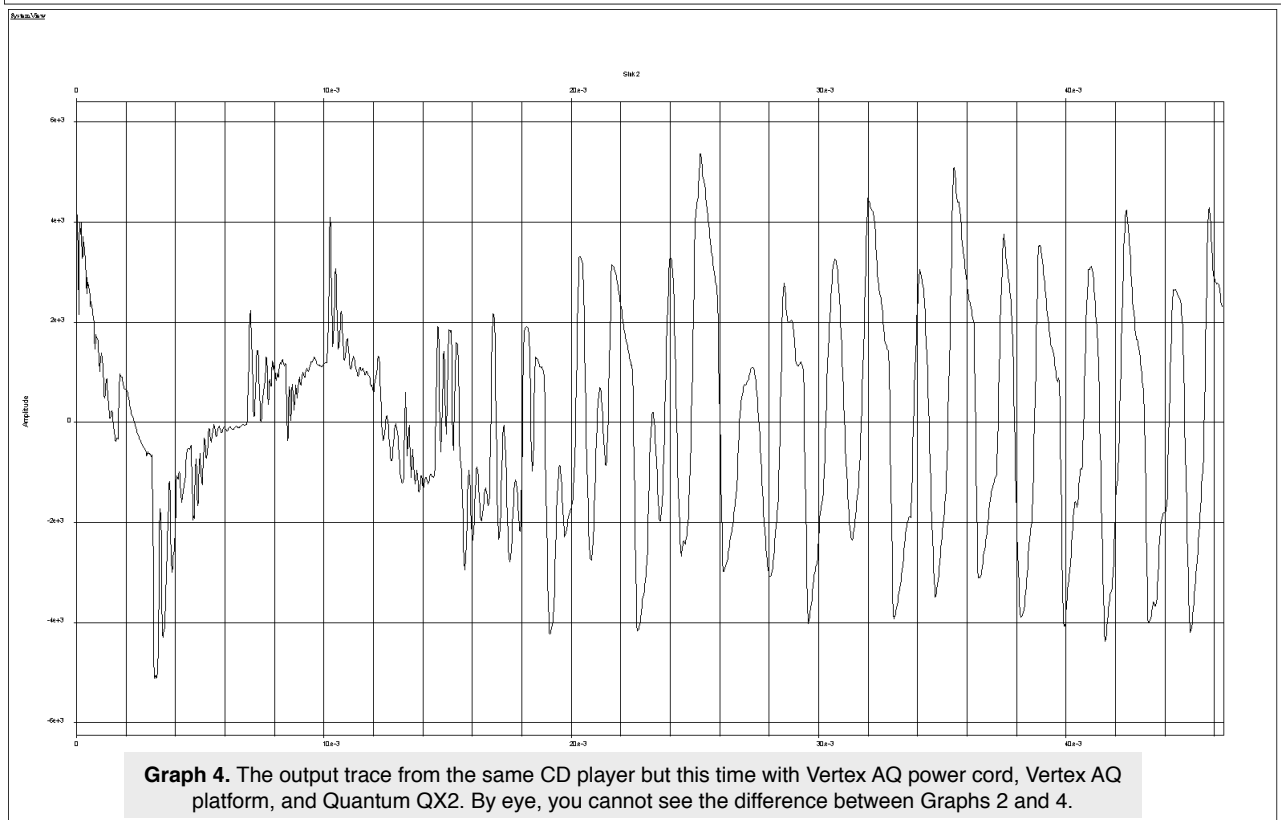
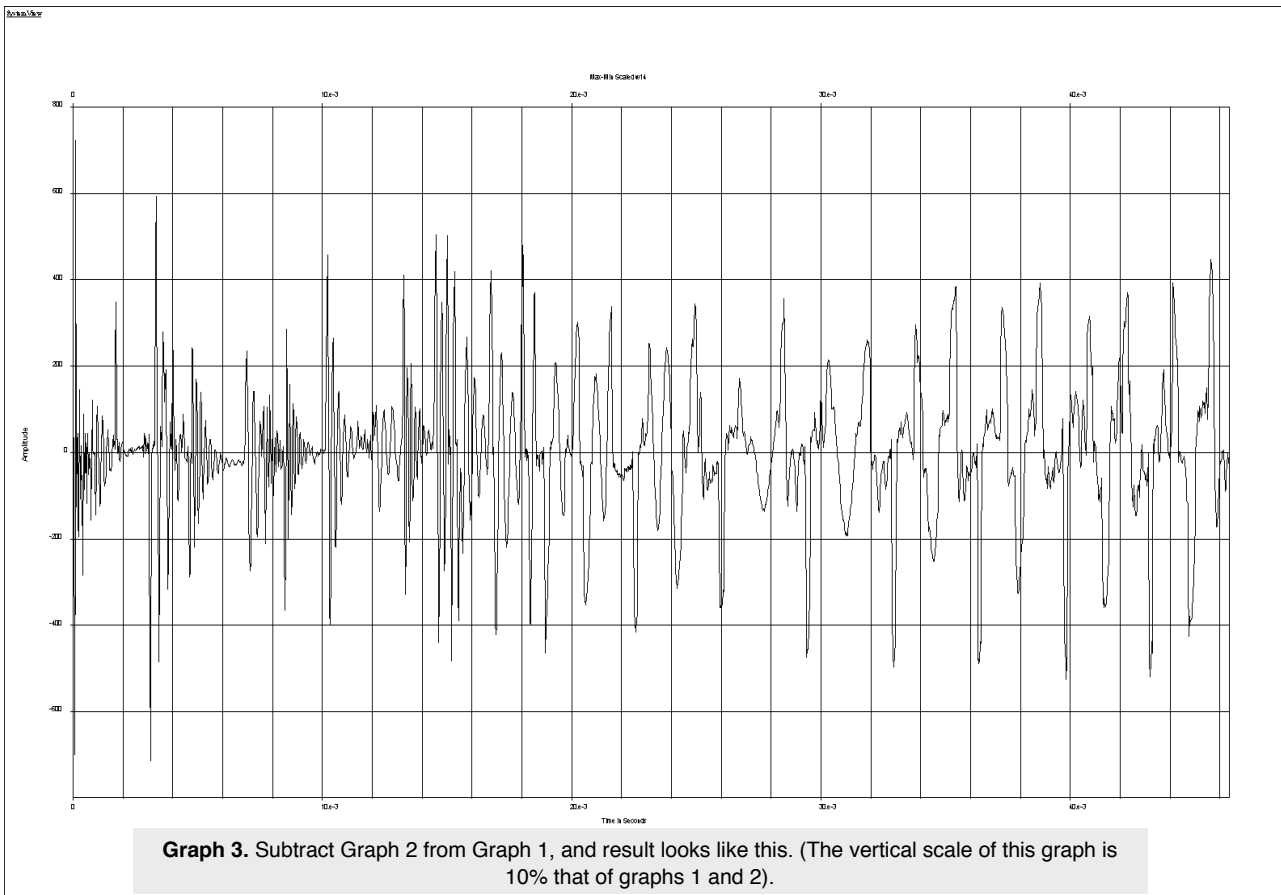
Readers should note that all testing and analysis is carried out entirely independently by Acuity. No Nordost or Vertex staff are present during testing or assessment of the results, or indeed the design of the testing protocols and algorithms. This “air gap” is critical to the credibility of the process and is maintained as a validation requirement by Acuity under the same procedures it employs for its defence customers.

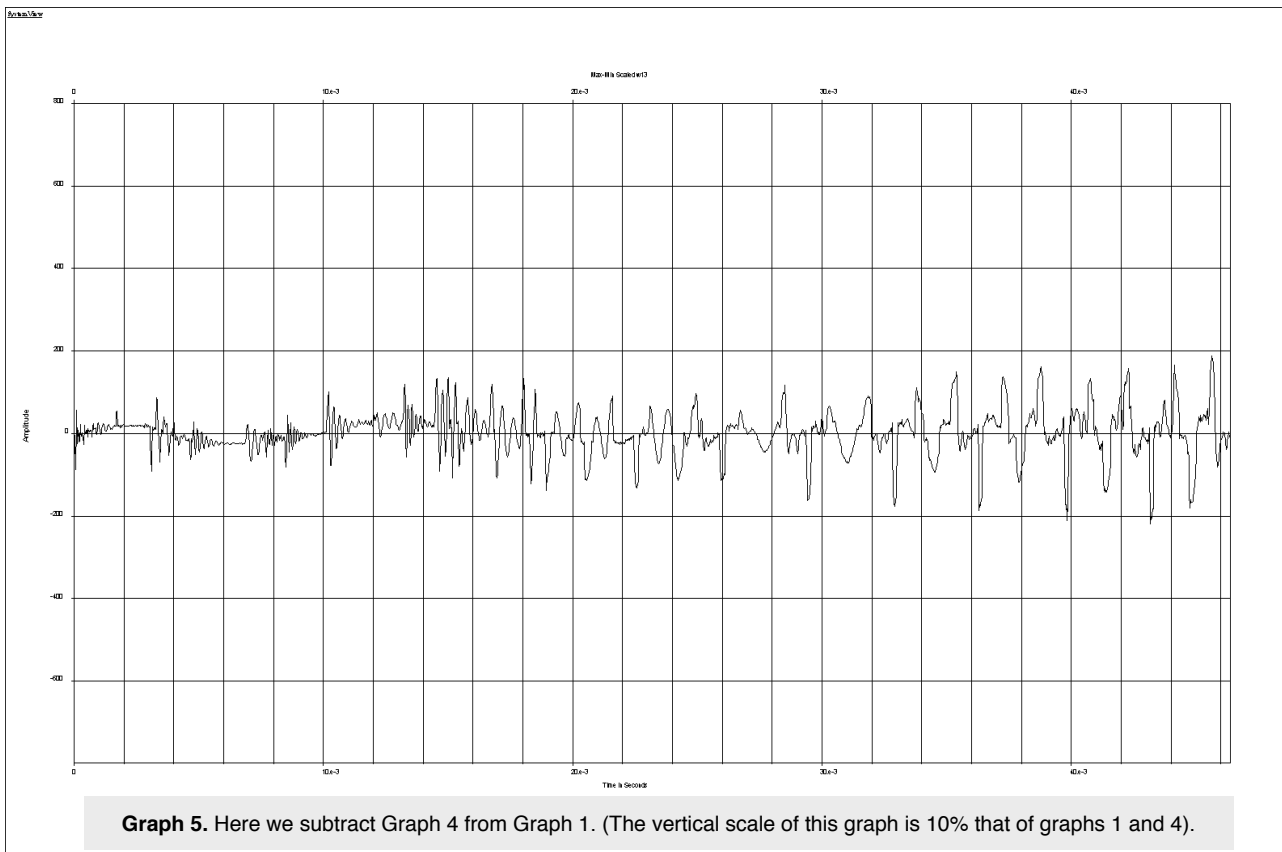
The first simple analysis was carried out by examining WAV files and then using a subtraction algorithm to get a difference output. The process was conceptually simple and comprised the following steps:

- Select a specific portion of a real music track.
- Copy this portion of the track into a PC. This becomes the reference data (WAV file).
- Play CD track in CD player, taking analogue outputs back into PC using a good quality sound card.
- Compare the CD output WAV file with the reference data using a simple alignment and subtraction method to produce a difference plot.
- Make a change to the accessory fit (ie mains lead), whilst leaving the rest of the system untouched, and repeat with the same track.
- Again compare the CD output WAV file with the reference data to produce a second difference plot.
- Compare the before and after difference plots.

These graphs show the error of a typical CD player on the lab bench playing real music.







You will see, looking at the 2 subtraction plots, that the error difference has reduced by at least 50% (the graph scales are identical). Interestingly, further analysis showed that the error products were not particularly associated with the amplitude of the signal, but seemed to be more closely linked to the gradient and complexity of the music.

Important points at the end of this work

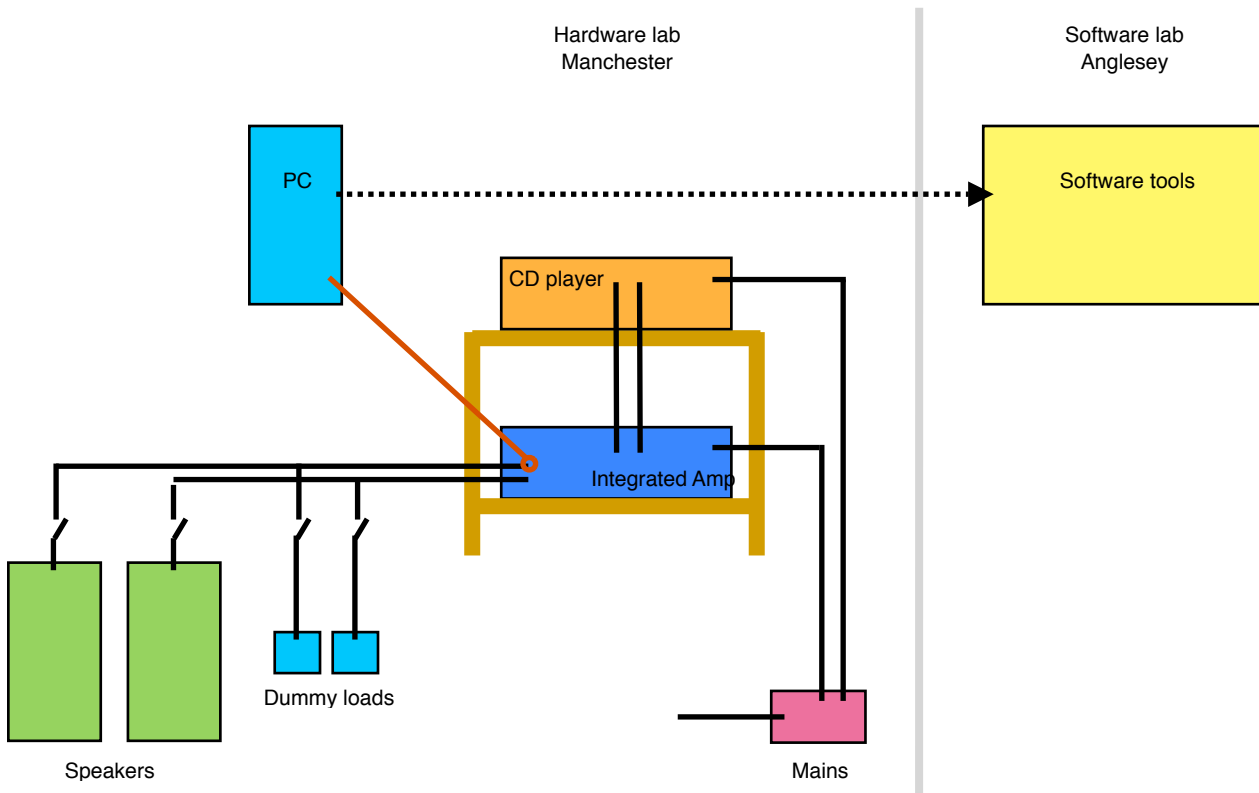
Right from the start, the measurement approach quickly established 2 significant elements to the thinking; the use of real music as opposed to test tones, and the measurement of the output in the time domain as opposed to the frequency domain.

It had always been Vertex AQ's and Nordost's contention that simple measurement techniques were somehow missing a significant element of hi-fi performance. And this possibility was quite quickly supported by the scientists at Acuity Products. In summary, conventional measurement techniques do not really stress the system under test (as is the case when playing real music), nor do they tell you how the output might change in the time domain. As we progress with more detail about the measurement process, the shortcomings of conventional measurement techniques should become readily apparent.

Phase 2 - 2009/10

If the work carried out in Phase One of the research project might be termed “Proof Of Concept”, the next stage was all about refining the process and improving both the accuracy and our understanding of the results.

Hardware test configuration



Above is a schematic of the basic layout of the hi-fi and test system at the Acuity lab in Manchester, England. The data analysis and software development is carried out in Anglesey, North Wales. Note that the test signals are picked up from the analogue outputs from either the CD player or amplifier analogue outputs and digitized by a high-quality sound card in the PC.

The system as a whole is not currently calibrated in any way. The results at this time are always comparative. The PC and sound card are a constant and changes are only made to the hi-fi elements of the system.

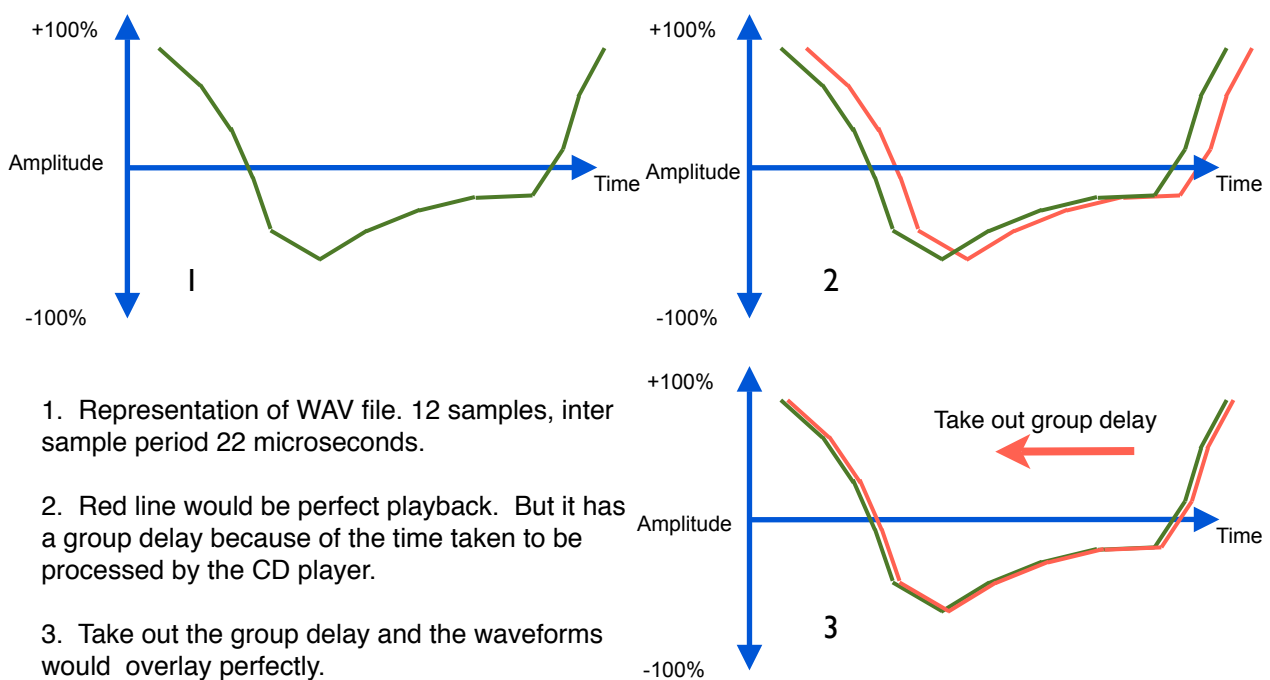
The principles of comparing 2 complex analogue waveforms

In the first year's work we were comparing 2 waveforms. These waveforms are analogue in their origin - they are the music encoded, via the recording process, on a CD. By its nature the left and right channel information is each a complex, but single composite waveform made up of all the musical information of the recording. In its raw state you cannot determine which parts of that waveform might be vocals or a violin - indeed when we listen to anything, either recorded or live, it is only a composite waveform that reaches the ear, and it's the brain's sophisticated decoding processes that turn it into anything recognizable or intelligible. It is this process of the brain decoding and placing in order

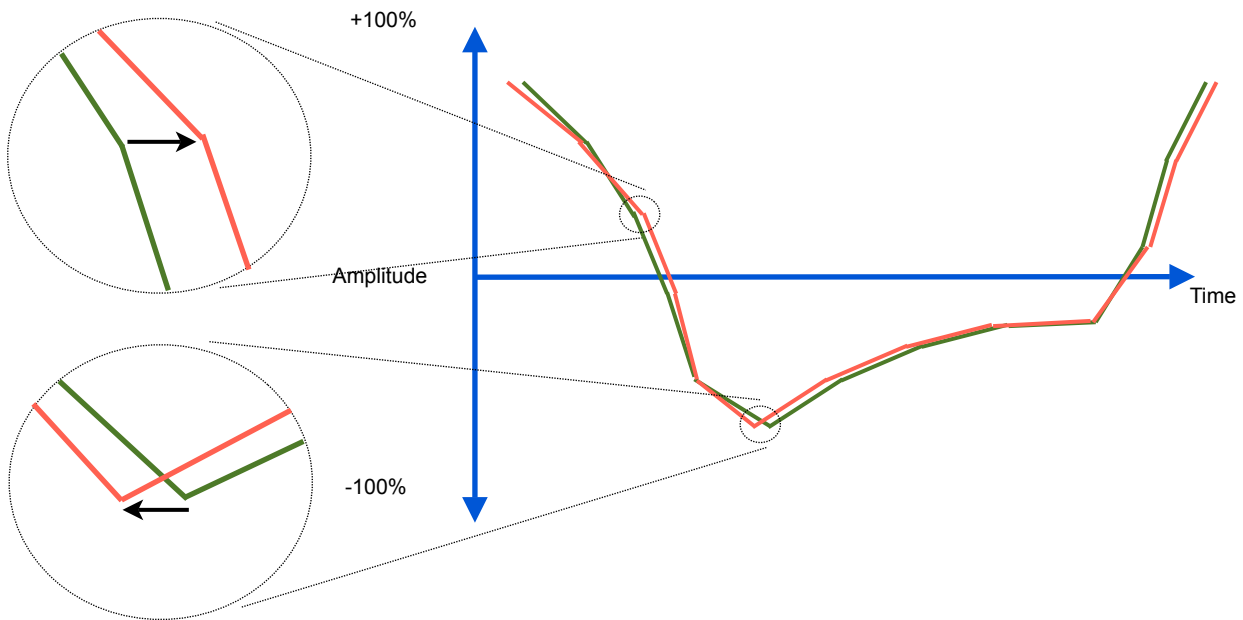
what it receives that is key to our listening pleasure - and the ability of the brain to do that easily, or not, can make all the difference between a good hi-fi and a bad one.

We saw earlier the principle of capturing and comparing 2 waveforms. And remember that the approach here is to have a reference waveform in a PC, then play the same track in a CD player (and through an amp too perhaps), take the analogue output of the player and compare it to the original, back in the PC.

Lets look at the principles of this process in more detail. The illustrations below represent a a very short part of a musical passage (one channel), expanded to the point where you can discern the individual sample points. The green trace is the reference trace, the red one the output from the CD player. The first step is to remove the fixed time error (or group delay) that is introduced by passing the signal through the player itself, thus time aligning the two traces.



But in reality, even when you take out the group delay, the trace from the playback does not exactly overlay the reference curve. No matter what you do to try to overlay it exactly, it will not fit all the way along the length of the trace. The playback curve has been distorted of course, it is no longer the same shape as the reference curve. What you actually get looks more like this:



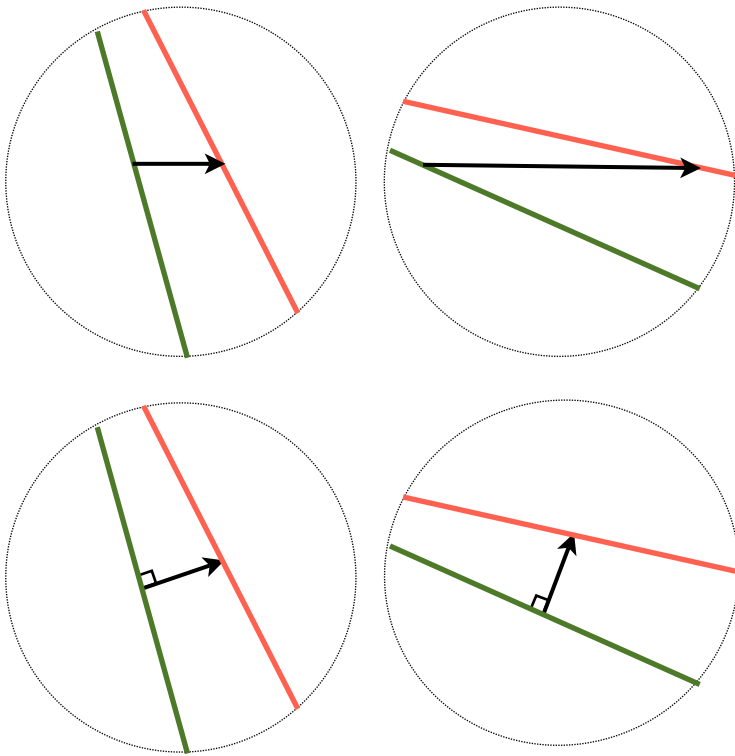
Follow the traces and you will see that sometimes the output signal leads the reference, and sometimes it lags. This is typical of the output trace derived not just from a CD player, but actually, from any electronic component. The degree and precise nature of the error may vary, but each block of circuitry (whether it's a hi-fi component, a sonar system or a radar) will inevitably introduce its own associated errors.

Improvements from last year - 1 - change of measurement vector technique

Last year we were measuring horizontal displacement only, reflecting errors in the time domain. In reality, errors also occur in the vertical axis, or amplitude domain. And the relative impacts of these errors in either time or amplitude are in turn affected by the gradient of the line.

So this year, the measurement algorithm was changed considerably. Now, the software measures perpendicular displacement (closest approach), reflecting the cumulative error in both horizontal and perpendicular axes. If we are to achieve our goal of producing a measurement system that avoids the pitfalls of the conventional techniques we discussed earlier, then we must have information that represents both the amplitude and time domain.

In the diagram below, the top 2 circles show how the measurements were taken in the first year's testing, and the lower 2 circles shows how the algorithm takes the measurements now.



Last year measurement was in the horizontal axis only. Now the software algorithm measures Perpendicular Displacement (PD).

The units of measurement with this method are more complex in their derivation but they are mathematically robust.

Improvements from last year - 2 - alignment

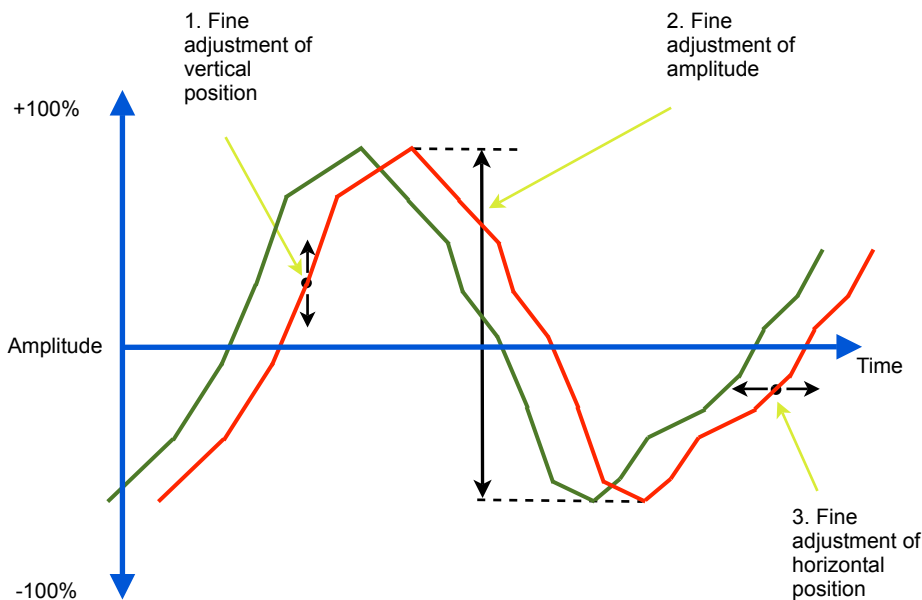
The accurate alignment of the reference and replay traces is a fundamental requirement of this whole process. After the first years work, which was in essence proof of concept, the alignment process was refined significantly. The process below describes how we ensure the results are a true representation of the differences in the 2 curves, and not just the measurement of an alignment inaccuracy.

The alignment process is carried out in 3 stages:

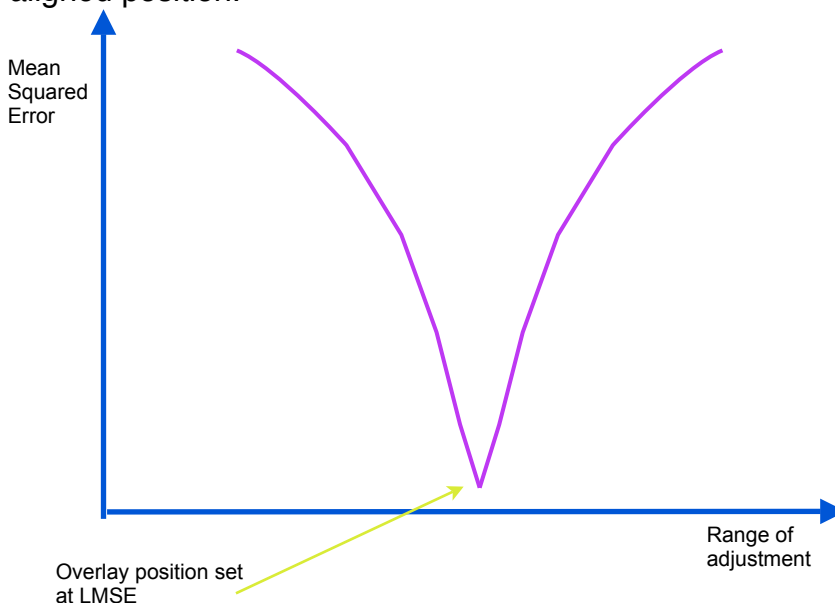
1. An appropriate piece of music is chosen from a CD and copied into the PC (a reasonably dynamic, wide bandwidth piece of music). At the start of this piece a timing mark is added to the WAV file in the form of a single large value amplitude spike (one sample). This marked track can then be burnt onto a CD-R for playback in the CD player. Remember that there is no consequence if the burning process adds some error to the recording, because always in the end we are only making comparative judgements of the final results when we make accessory equipment changes to the system later. The playback is captured as a WAV file from the PC sound card and using the mark on the reference file and the mark on the playback file, the 2 are aligned to within approximately 1 sample width (+/- 11 microseconds for red book CD)
2. Next, the power in the replay trace has to be exactly matched. And in this year's work we were also measuring the output of a power amp, so considerable attenuation and matching is required in this case. This process of power matching is done with an amplitude squared technique. By squaring the amplitude it all becomes positive integers, then the level of the playback trace is adjusted until the power (area under the graph between 2 selected sample points) is identical.

3. Then finally, the vertical and horizontal position of the replay trace is finely adjusted using a Least Mean Squared Error (LMSE) technique. This now uses the PD algorithm to run along the 2 traces, and measure the perpendicular error. Some of these will be positive, and some negative, so again they are squared. Then the replay trace is displaced a little, and the PD and error squaring run is done again. Each time this is done there will be an overall summed error value, and the trace is moved around until the summed LMSE is at the minima. This is then deemed to be where the traces are correctly aligned.

Conceptually, the fine alignment process does this:



Each time the relative position of the 2 traces is moved, the LMSE run gives an overall value, and this can be plotted on a results graph as shown below. The results of all the runs produce a very clear 'V', or notch result, and so the point of the notch is the correctly aligned position.



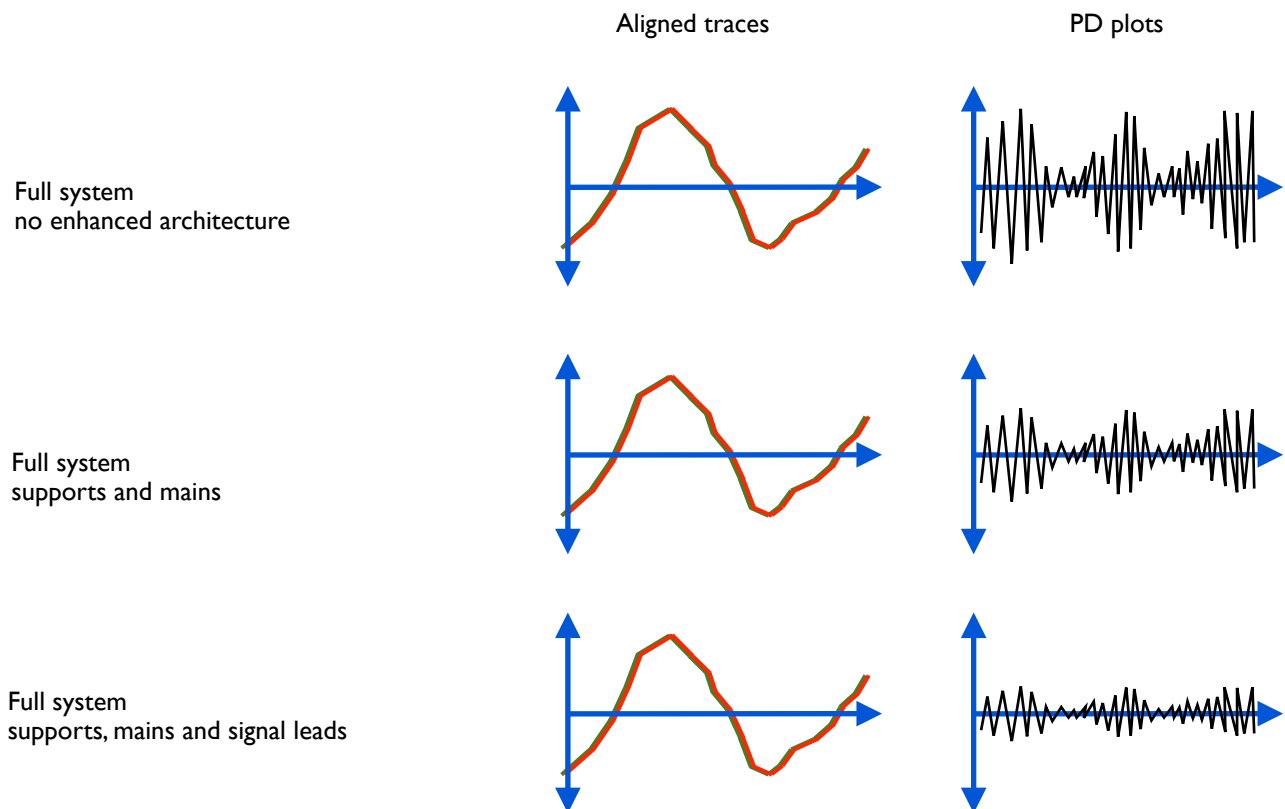
The height of the lowest point above zero on the LMSE plot represents the real error that is within the system. The average value at this point is remarkably consistent at each of the different system configurations.

Validation

Once the LMSE point is determined, the processing algorithms then run through the track sample (approx 2000 samples at Red Book CD 44.1 KHz), and measure the PD at every sample point, and plots it on a graph. We can produce a great many of these graphs with this PD process, in varying states of 'good' or 'bad' accessories. But clearly, now we were confident that the measurement process was robust mathematically, we had to ask ourselves how we validate these results against our real world listening experience with such equipment.

But the validation of the results came about just that way - through doing many runs in different configurations. Because as Acuity carried out these runs, and grouped them according to the condition of the set-up, they were remarkably consistent in their overall error values. With no system enhancements the PD plots were always large, and at various states of enhancement, the PD plots were reduced consistently by an easily determined amount. This exactly matches our expectations of the sort of improvements such products would bring in a real listening situation

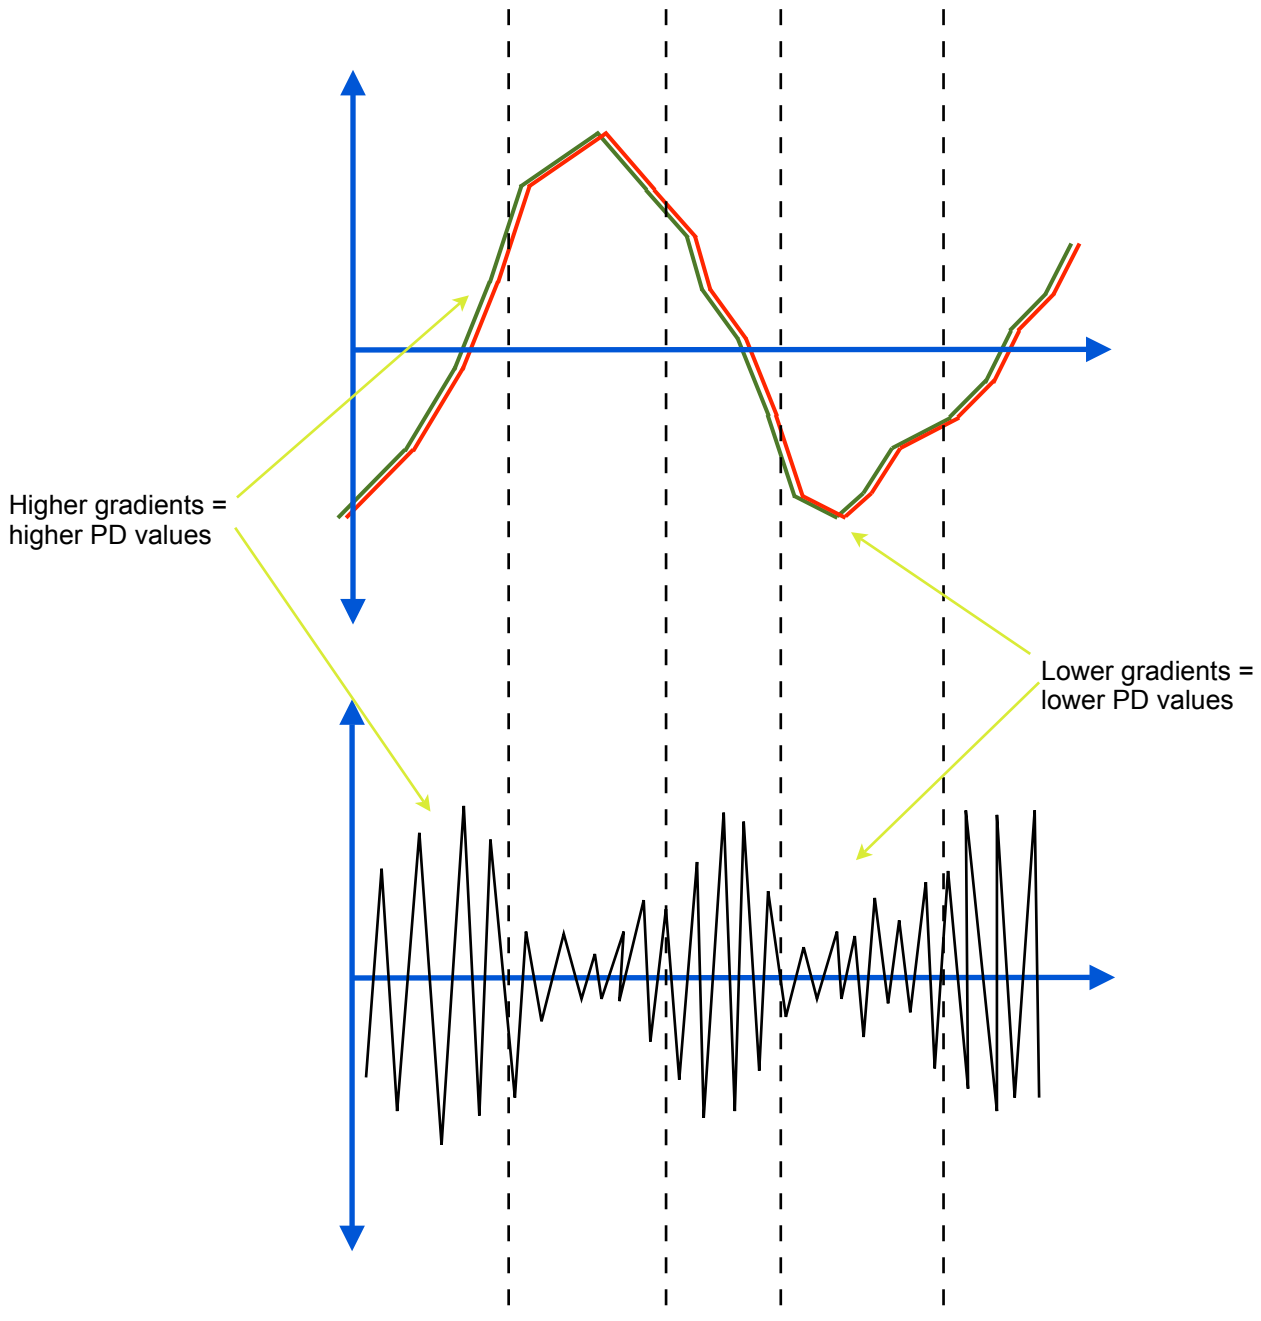
The principle of this is shown below. We are using illustrations at this point because they are easier to read than the complex plots produced by the measurement process, but we have many, many examples to support these conclusions.



The next part of the validation was even clearer evidence of the correlation between the PD plots and the complexity of the musical signal (see below). Again this directly reflects

our hi-fi experience. How many times have we clearly heard a system get progressively worse as the music gets more challenging?

In the diagram below we show how the value of the PD plots clearly reflects the complexity of the music. And in particular, it is the steepness of the curve's gradient which has the clearest link to the amount of error created. In other words, the more rapidly the input signal changes, the greater the PD error.



What we are seeing here is two things: the way in which the distortion mechanisms within a hi-fi system relate directly to the demands of the signal being played, and the way in which specific approaches can impact on those mechanisms. Because we are dealing with external accessories, we can add or remove them at will, but the affects could be achieved internally, by adapting the construction or design of an amp or CD player itself say. The

accessories employed in this testing program directly attack microphony, RFI artifacts and AC power quality. The results (measured and heard) are unequivocal, but also show how the results of the measurement program could extend beyond the realm of system set up and tuning, into the area of electronic product design.

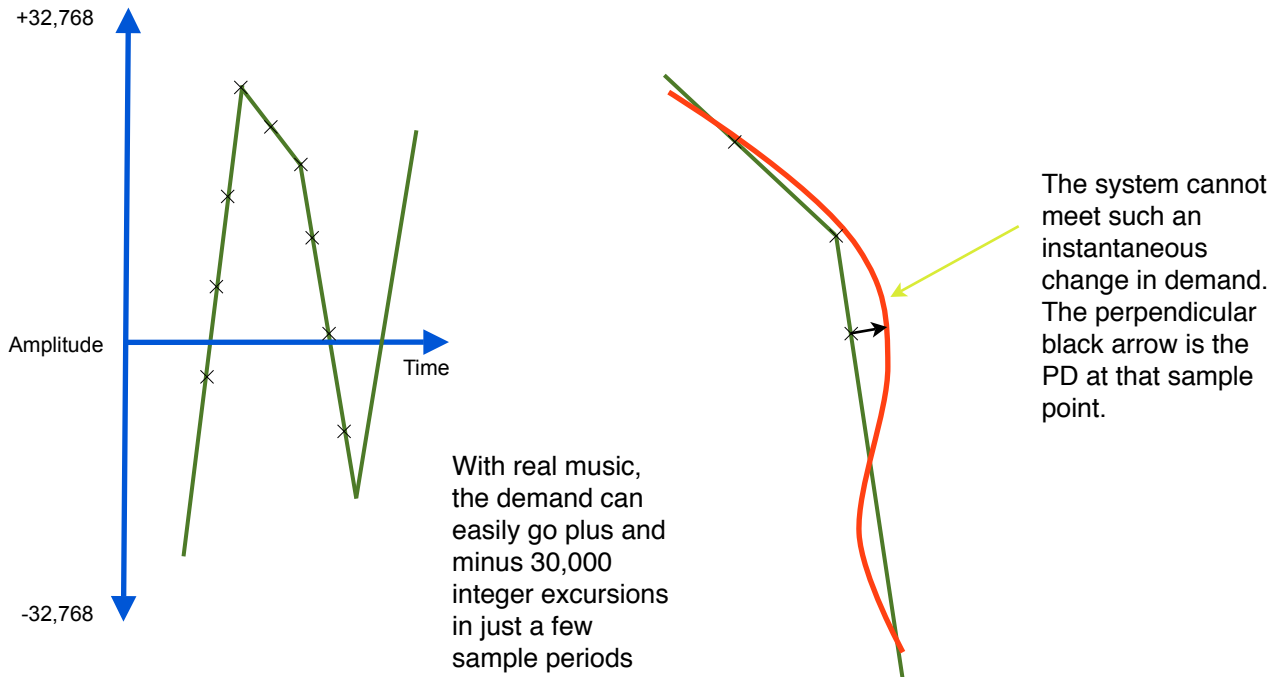
Tracking Error?

But its still a little difficult to grasp in our mind what is really going on. We have become used to visualizing things in the frequency domain (even though they may be of limited value), but it is instructive to look at attitudes to systems engineering in other sectors (such as defence) where we now know that they have a more time domain orientated viewpoint.

They start from the assumption that no system can be perfect - perfection is an unattainable dream. And so every system will fail in some way to output the signal that has been fed into its input. The system will never respond quickly enough, it will be under damped or over damped, it will act like an oscillator to a certain extent and ring - and so on. That is the nature of all systems. And all of these effects can be considered in a systems design, but at the end of the day, they always have to be traded-off. But in hi-fi it seems, some of the factors which have an effect (such as microphony) are in general, not being traded-off at all. They are being ignored (and not seen in the measurements either of course!).

But how does a defence engineer visualize all this in the time domain. Well its surprisingly simple actually. As the input signal rapidly changes, the system fails to keep up with that change to a certain extent. A good analogy is a dog chasing a rabbit. The rabbit is light and can make sharp changes in its course. The dog is much heavier and does not have those powerful back legs of the rabbit, so as it chases, it undershoots and overshoots the rabbit's track.

And thats it! This is what our hi-fi systems are really doing - they are failing to accurately track the input signal. What comes out is undershooting and overshooting the input signal. And with our new techniques we are measuring the difference between where the track should be, and where it actually is, at each sample point. This is what the PD measurement really represents.



With the development of our measurement techniques we always instinctively used real music because it was dynamic - and we know through listening that complex music will more readily trip up a system. But now what we do when we look at the results in the time domain is revert back to comparing them as 16 bit WAV files. This actually makes our job at this stage relatively easy now, because the Acuity scientists can take PD measurements at each sample point on the trace, and quantify them as actual integer values.

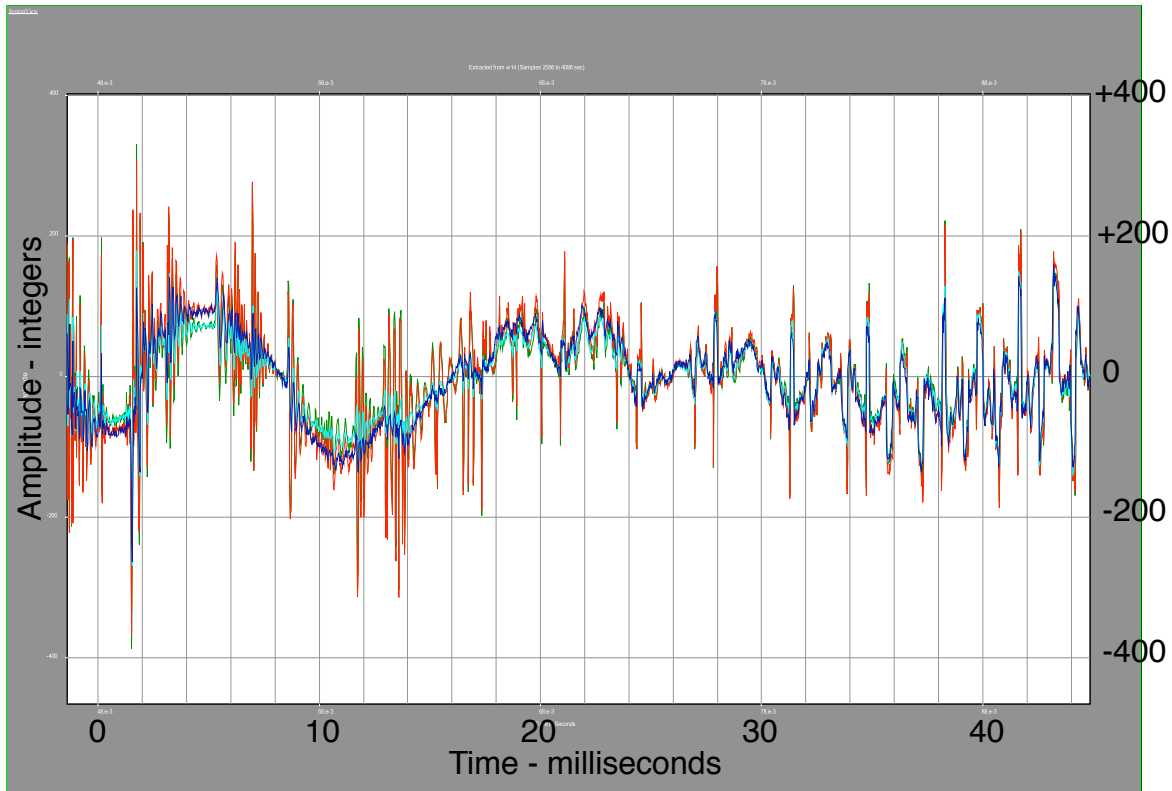
In the diagram above we show how you might visualize tracking error. But first it is important to remember that with real music, the WAV file can easily go plus and minus 30,000 integers in just a few sample periods (maximum with 16 bits (actually 15 bits and a sign bit) is $\pm 32,768$). And when we now quantify the PD errors on a dynamic track we easily see PD errors in the order of 100 to 150 integers. That's an error value equivalent to 6 or 7 of the least significant bits of our WAV file word!

Now some people get confused at this point - they imagine a big vertical spike popping up out of a smooth sine wave, and say things like "I can't imagine at all that a system can be doing this". So you must remember the context of the dynamic signal, and that these values are an undershoot or overshoot in real time. And then just remember your listening experiences - a badly set up system will be mistracking by these amounts - that's why it sounds so rough and disjointed. Then go through that system with a good infrastructure methodology and it will be transformed - the tracking error will be reduced dramatically. And that is what we have done in this measurement work of course - and we clearly see these massive reductions.

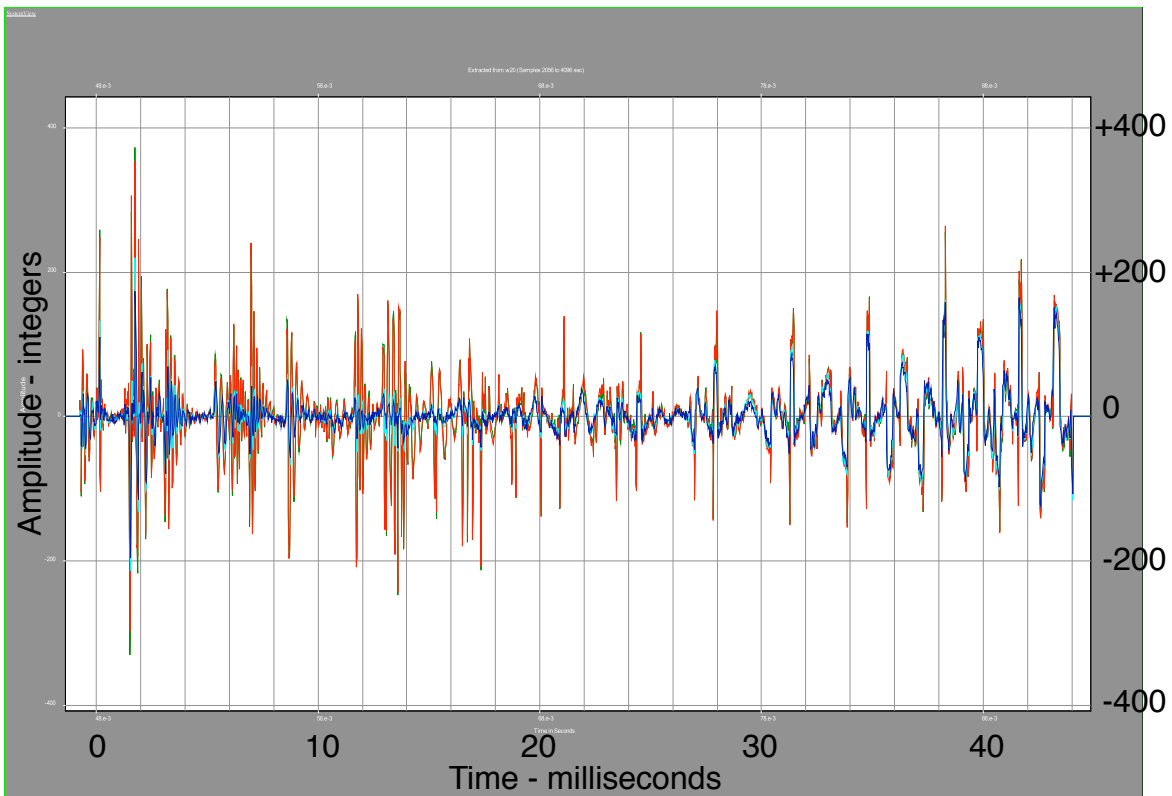
Now for some real data

Below is an example of the graphical output that the algorithms now produce. This shows a 4-color plot. You will see that there are 2 distinct characteristics to this graph. There is a large 'S' shaped excursion that is happening in a slower time period, and there is a fine and rapid positive and negative excursion superimposed on this.

The second plot shows a de-trended plot, removing the long-term errors, so that an average measurement can be taken of the short-term results.



Basic and full Vertex configurations (dummy loads)
 Green - CD Output, Red - Amp Output,
 Lt Blue - CD Output + Vertex, Dk Blue - Amp Output + Vertex



Basic and full Vertex configurations (dummy loads). Removal of long-term trends.

Green - CD Output, Red - Amp Output,
Lt Blue - CD Output + Vertex, Dk Blue - Amp Output + Vertex

But why do we go to the bother of de-trending the plots? Well for a very important reason as it turns out. The changes to the accessories only effect the short-term errors! When we change the accessories, do a run, de-trend the results and measure them, we clearly see significant reductions in the errors.

We'll come back to the long-term trend later.

Data from specific test runs

Vertex AQ tests

A full setup test was carried out to compare stock power chords, signal leads and with no dedicated supports, verses a full Vertex AQ setup (using the lowest price 'standard' items from the Vertex range) . Exactly the same CD player, amp and, in this case, dummy loads were used.

But we will also use this example to show more detail of the latest way we are showing the measurements. And as we said earlier, we now describe the absolute error values in integers, which are adjusted to match the Red Book CD format. In this example the musical excursion was $\pm 11,000$ integers, and the averaged errors are as high as ± 150 integers. Note that the amplifier output is attenuated back down so that its integer range is again identical to the scale of the WAV file.

Configuration	Integer value error - CD output	Integer value error - amp output	Approximate musical excursion value of test track	percentage error - CD output	percentage error - amp output
Basic system	+/- 140	+/- 150	+/- 11,000	1.27%	1.36%
Full Vertex	+/- 65	+/- 85	+/- 11,000	0.59%	0.77%
Reduction	75 (53%)	65 (43%)	-		

Configuration	Integer value error - amp output	Error expressed as bit depth	Integer range for this bit depth
Basic system	+/- 150	approx 7 LSB	+/- 128
Full Vertex	+/- 85	approx 6 LSB	+/- 64

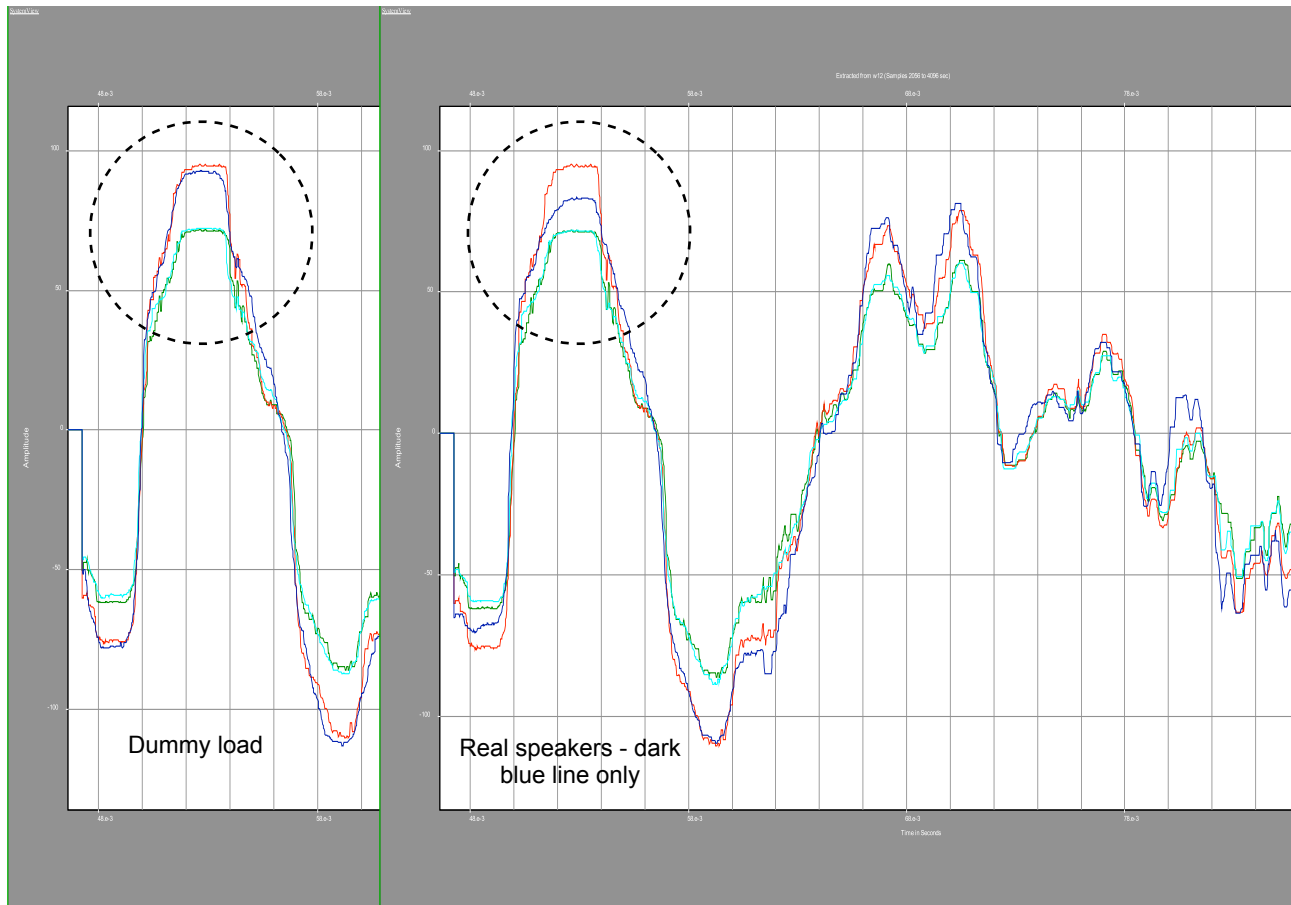
What is interesting here from a context point of view is how the untreated basic system has an overall averaged PD error at the amplifier output of 1.36% when it is playing that track (which has an excursion value of +/- 11,000). But when the Vertex components are put in place, playing exactly the same piece of music, the averaged PD error at the amplifier comes down to 0.77%. Thats almost a 50% reduction.

Another way of looking at it is shown in the second table. The untreated PD error is in the order of the last 7 bits of the WAV file information. With the treated system that comes down to the last 6 bits of the WAV file. Also, put this in context of some of the debate about digital errors from a theoretical point of view. Take quantization error for example, which can only be a maximum of +/- half the least significant bit, thats 0.5 of an integer. Our overall tracking error of the untreated system is 300 times that amount!

Now, whilst all of the examples shown reflect the impact of external accessories on a system - exactly the kind of products that both Vertex AQ and Nordost produce - that merely reflects our collective interest in this specific aspect of system performance. At the end of the day, he who pays the piper (and this is a costly and time consuming program) calls the tune. But it would be just as easy to focus the testing on other aspects of system performance - maybe the internal construction of a CD player, for example, or the precise placement and fixing of a mains transformer within a chassis - and the results would be just as telling.

Long-term Errors

Earlier in this document we mentioned the presence of a longer-term error structure which was clearly visible in the results. During the testing this year we began to discover more interesting things about these effects. Firstly, they are certainly not influenced by the infrastructure changes in the test system. But we did see some very clear changes when we changed the system between driving real loudspeakers, or the dummy loads.



In this diagram we show a change to the long-term errors in the system when the system either drives the loudspeakers or the dummy loads. Firstly we need to explain the diagram carefully.

- On both graphs, the green and light blue lines are the CD player output (with and without Vertex). And the red line is the output of the amplifier without vertex, and the dark blue line is the amp output with Vertex.
- On the left hand graph the red and dark blue lines (with and without Vertex) are both the results driving the amp into the dummy load.
- On the right hand graph the red line (without vertex) is still the amp driving the dummy load. Its the dark blue line that shows the Vertex equipped amp now driving the real speakers.
- So really we can disregard all but the dark blue lines here - and we are comparing the reduction in the peak of the dark blue line when driving loudspeakers, compared to the height of the peak when driving the dummy load.

Now Acuity had not really deliberately carried out this test like this. This was just a series of runs, and switching from loudspeakers to the dummy loads was primarily done to reduce noise disturbance to other laboratory workers! It was only when comparing results that this interesting effect was discovered.

But when more research was carried out about the behavior of systems, comparing these results with the performance of defence systems, a pattern started to emerge. The amp PD error output changing with its long-term trend from speakers to dummy load is because of the relationship between the amps fundamental characteristics such as bandwidth, damping factor, feedback topology and so on, and how these things react differently to a purely resistive dummy load, or a reactive speaker with its crossovers and driver coils.

And when the Acuity engineers considered all the other system's long-term trends, their view is that these are all set by the main engineering characteristics of the whole system. With the CD player for instance, this could be set by the type of processor used and the noise shaping algorithms employed. It could partly be the type of analogue output stage used in the CD player. But interestingly, these long-term trends still have a relationship to the music being played.

And of course these things cannot be changed at all by changing the system ancillaries. And that is what was found. The long-term trends are set by the electronics, and are not affected by the system set-up. But the short-term trends ARE very clearly a significant error mechanism, and they are very drastically reduced by the application of better infrastructure.

Conclusion thus far

We know that in defence engineering, every system is a compromise, we have to balance tradeoffs with bandwidth, impulse response and so on, and we tend to think of noise effects as random. However, we have never before seen hard proof that a large proportion of noise is not random, but instead is signal related. Nor that signal transfer is clearly being impaired by systematic microphony, RFI and other effects - a fact demonstrated by the application of our external "tools" that specifically attack those effects, and the creation of a measurement protocol that reveals their impact.

So to summarize, we now see two distinct error mechanisms

1. A short-term error structure that is significantly affected by changes to system ancillaries and set up. Introducing better ancillary equipment shows a substantial reduction in this short-term error – in line with our listening experiences and musically demonstrable results.
2. A relatively long-term error structure that seems almost impervious to ancillary variations – although clearly reflects the impact of the speaker load on amplifier performance. It seems likely that this error mechanism reflects major design criteria embodied in the equipment, eg, the noise shaping selected in CD players, amplifier topology or output devices etc.

But, in all cases:

- Both error structures relate directly to the programme material being used.

- The error structures clearly reflect changes in musical energy.
- The short-term error structure is reduced with the effective application of specific techniques.
- The results are completely reliable and repeatable.

What we have now is a measurement process that is reliable, repeatable and meaningful – these measurements correlate to an uncanny degree with our collective listening experiences.

We also have a process that has been entirely developed by an independent defence contractor (Acuity) with impeccable academic and professional credentials. This is a rigorous and scientifically valid approach to assessing audio performance in general – not just the fanciful creation of a couple of cable manufacturers!

Next steps

When we started this work it was a journey of exploration. We knew our products made a massive difference to the performance of hi-fi, but we knew we could not measure it effectively. Then by good fortune (the “other” side of Vertex AQ and its involvement with the UK defence industry) the link was made to Acuity. The like-minded interest of Nordost, together with the possible application of the research to the further understanding and development of sonar systems, allowed us to embark on a three-way, cooperative research project, taking advantage of Acuity’s independence and vast experience, knowledge and skill. Furthermore, Acuity were going to tell it like it is. They have their credibility with the UK MoD to uphold, and they apply very strict results validation processes before ANY information is released from their scientists. We could be sure that their work and results would be academically and scientifically rigorous; anything else would cost them too much.

The creation of this new measurement methodology based around real music and output in the time domain is a great step forward. The most significant thing of all is the ability to properly identify the deviations in output, in the time domain, by using the Perpendicular Displacement method. And because this is a very laborious process, Acuity's skill in creating automated algorithms to do the work has proven to be an essential cornerstone to the whole approach.

But now we have these algorithms working, the impact of their implications is extremely exciting. First and foremost, we now believe we can turn these complex algorithms into a software suite of measurement and analysis tools, available to all. To create the additional software to give a good user interface, and to output the results in an easily understandable manner is now, we believe, wholly achievable. It's just a job of work and could result in an affordable end-user tool for the audio enthusiast, as well as a more comprehensive set of tools for audio designers.

Just as important is the new perspective we are gaining on the way audio systems work with and react to real music, their infrastructure, set up and environment. We can now see just how large the impact of these factors is on performance; we have a measured expression of the huge gulf in performance between say, a well-tuned system using good electronics and one with bad electronics, poorly set up. And those numbers are derived directly from the music being played. Suddenly, a whole range of approaches to system set up and design, previously the subject of heated and unsubstantiated debate, can be

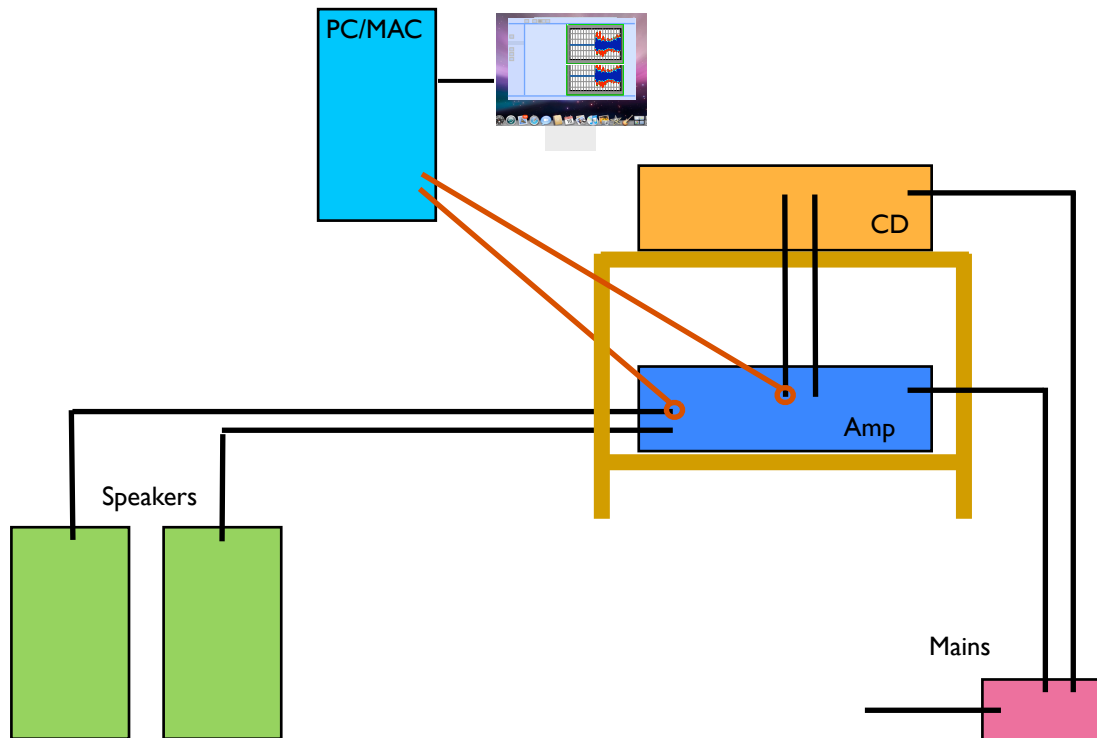
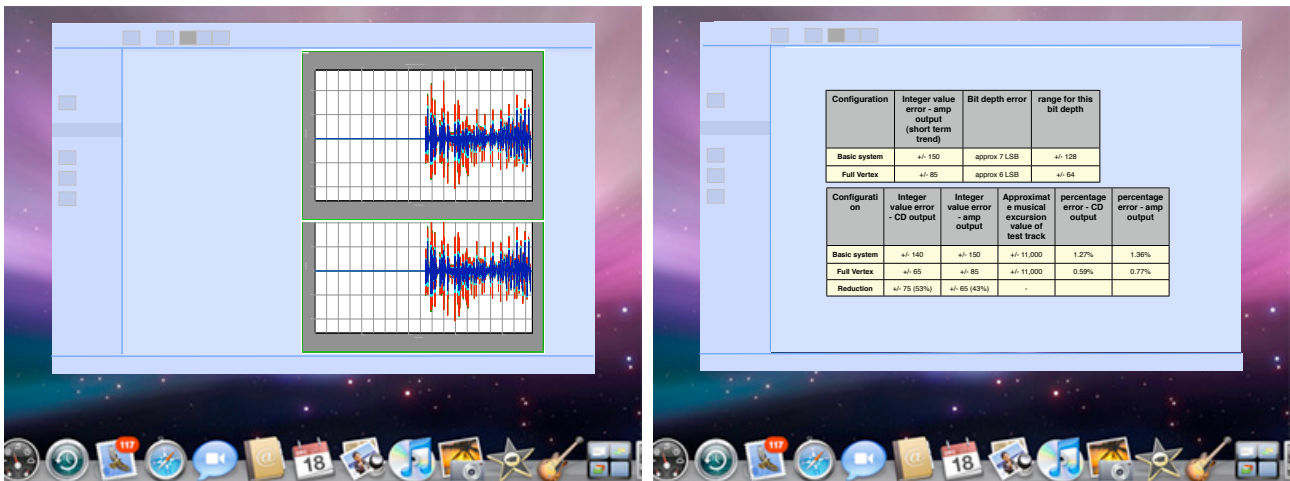
examined on a scientific basis. We need no longer reject audibly effective actions, simply because our previous measurement techniques were inadequate to the task of explaining them. And before there's a chorus of, "Well, you lot manufacture cables, so you would say that!", consider the following:

- Acuity don't manufacture anything for the audio industry. Any field in which they apply this methodology will be subjected to the closest possible scrutiny to defence standards.
- What is being proposed here is just a tool. Properly applied, it can be used in a number of different ways. And properly applied it could eliminate or drastically reduce the reliance of electronics on external measures such as support cones or platforms - products that we currently manufacture.

As audiophiles we all intuitively know one fundamental thing - basic published measurements (whether they come from manufacturers or magazines) whilst they might, for instance, guide us when it comes to matching the electrical characteristics of an amp and speakers, bear little or no relation to the musical quality of the resulting listening experience. We've all heard systems where the bits of kit measure with 0.01% distortion yet the system sounds bad; and we've all heard other systems, with higher levels of measured distortion, say 0.1%, which deliver extremely enjoyable musical results. How can this be unless our existing measurement techniques are missing something absolutely fundamental?

The work done by Acuity does NOT provide all of the answers. It does NOT provide any one critical answer. What it does provide is a wholly new perspective that significantly extends the usefulness of what we already know. This is an addition to existing measurement techniques, which gives us a new viewpoint both when it comes to system performance and how we can interpret and understand the results of what we already know.

Over the coming year, our aim is to design, develop, test and release a downloadable, end-user measurement tool based on the core work we have done thus far. The result (in functional terms at least) could look and work something like the diagram below.



In this example we are showing the input and output (one channel) connected up and the signal fed to a stereo sound card. The user makes changes such as supports or mains leads and sees the changes between the input and output signal with real music. So changes can be seen and heard - and the user can make more informed decisions about the equipment choices for the system.

A small hardware kit would be made available with test leads, attenuators and double connectors. The software would be web downloadable and supported by comprehensive guides and regular maintenance and enhancement updates.

END