Running on Empty

by Phil Ward

We've all experienced it. The new component that sounds raw and ungainly at first but as the hours, days or maybe weeks of use pass by, transforms into the clean, clear, fast, open and musical hi-fi we dreamed of - or at least gets much closer than we ever thought possible. Running-in, burning-in, settling-down, call it what you like, has become a "given" and second nature in the hi-fi business. Without it, press reviews would perhaps be shorter and both manufactures and retailers might have a harder time selling product; "Yes sir, it does sound a little bright at first but that's to be expected. It'll settle down after a little while", is a pretty useful line. Tell me you've never heard something like it?

But what actually goes on when a product runs-in? Now I'm not remotely qualified to comment on or investigate running-in effects in cables or electronic components - I don't really understand the science. But in the case of loudspeakers, the transformation in sound we so often read about and the mechanisms that are suggested as responsible, should be apparent from the electro-acoustic measurement of a speaker before and after run-in. So in the second half of this piece I'm going to present a set of measurements of a pair of B&W CDM1 NTs - a speaker that B&W in the product Owner's Guide describe thus; "The performance of the speaker will change subtly during the initial listening period. If the speaker has been stored in a cold environment, the damping compounds and suspension materials of the drive units will take some time to recover their correct mechanical properties. The drive unit suspensions will also loosen up during the first hours of use. The time taken for the speaker to achieve its intended performance will vary depending on previous storage conditions and how it is used. As a guide, allow up to a week for the temperature effects to stabilise and 15 hours of average use for the mechanical parts to attain their intended design characteristics."

For the test, both speakers' impedance and frequency response characteristics were measured before run-in, and then one subjected to three weeks of use at high level for around 10 hours a day (the second speaker was stored in a similar environment but far enough away that it wouldn't pick up any acoustic energy). Both were then measured again with the non-run-in speaker providing a control reference. Non-linear characteristics, harmonic distortion for example, were not measured. This was partly due to the difficulty of doing so with the equipment available and partly due to the fact that, by their very nature, the change mechanisms under investigation, if significantly present, would very likely show themselves in the linear domain.

But now I've gone and done it. I've used the "measure" word in Hi-Fi Plus - the empyrean citadel of music-first subjectivism. I can hear already the muttering about "measuring the wrong thing" or "measurements don't tell the whole story". Well in another context that may be so, but remember these measurements are not aimed at trying to quantify the sound of a speaker. They are aimed simply at revealing change - the measurement system is "transparent". So don't fret. I'm not proposing we put the clock back 25 years, stop listening to any music and simply rely on a box of measuring tricks. I'm simply proposing that, if for example and as I've heard suggested, running-in alters the high frequency resonance of a metal dome tweeter due to its dome material work-hardening, the effect should be visible in a measurement comparison before and after.

"But why would we want to do that?" I hear you ask. "Surely we all know that speaker running-in effects occur so why waste time proving things we already know?" Well my short answer is that I've never been entirely convinced that running-in, at least in speakers, is as clear-cut as it seems. Undoubtedly there can be big changes in the subjective characteristics of a speaker as it warms-up with use. The causes of those changes are pretty easy to predict and pretty easy to measure. But they are reversible. As a speaker stops playing and cools, barring perhaps a little thermal hysteresis, its electro-acoustic characteristics will return to, as near as makes no difference, their room temperature state. However, the proposition that during a running-in period there are irreversible changes in the material structures and mechanisms of a speaker has always seemed to me a little far fetched, and a little too, well, convenient. Especially when there are plausible, although less convenient, psychoacoustic explanations for the run-in phenomenon that we've all experienced.

My doubts arise from three different directions. First, in very many years of working with speakers I've little direct measured evidence for the commonly proposed running-in mechanisms. The only one that I've definitely seen is the reduction in the suspension stiffness of a bass driver - which in turn reduces its fundamental resonance. But even this I doubt is typically of much, if any, significance. The very nature of drivers in boxes, where the enclosed air (in both reflex and closed box) is, 99 times out of a 100, by far the dominating stiffness parameter in defining the low frequency response shape (and transient performance), means that, all other things being equal, driver suspension stiffness is a relatively unimportant parameter. In other words, the low frequency alignment (by which I mean both steady state and transient behaviour) is fundamentally insensitive to driver suspension stiffness. If this wasn't the case, manufacturers would have a much harder time producing consistent product. Driver suspension stiffness is one of the most difficult parameters to control in manufacture (apart from the fact that it varies with temperature) and a typical driver will have ±15% tolerance on free-air fundamental resonance - mostly down to suspension stiffness variation. However ±15% is acceptable because once a driver is mounted in a box the stiffness of the air inside dominates. If you don't believe me, think how stiff a driver in a closed box feels compared to the same driver in free-air.

And the other mechanisms one hears proposed? Well, irreversible changes in the material characteristics of a cone, suspension or surround material, if they are responsible for the significant swings in performance one sees described, would surely produce obvious measurable effects in frequency response or spectral decay? If so, I've never seen them. Similarly, the previously mentioned "work-hardening" of a metal tweeter dome ought to change its high frequency response shape, especially its fundamental break-up resonance, pretty significantly. Again, show me the resonant frequency shift and/or level change that should occur if the material qualities change? Now it may be that these effects were there all the time, but as I wasn't specifically looking for them, they passed me by. And that, perhaps, is what this exercise is all about.

A second reason for doubt is a suspicious, smoking-gun of inequality in the perception of running-in that I've never been comfortable with. It's this. Why do run-in products always sound better? What possible reason could there be for running-in effects always improving the sound? How many times have you read, or heard; "Well, it was fabulous as soon as I turned it on but coloured, bright and unlistenable after a couple of weeks"? I recall reading something along those lines just once, in this very publication, a few issues ago, but that's once for running in resulting in a worse sound, against countless numbers for the opposite. If the mechanisms behind running-in are not understood and fully taken into account in the design stage, and with respect to the loudspeaker industry, I don't believe they really are, then one would expect the subjective effect of the mechanisms to be sometimes positive and sometimes negative (and sometimes to make no difference at all).

The only plausible explanation I've heard for one sided result of running-in is that designers tend to work on one development "mule" which, through the development process, becomes run-in. All the development decisions are then made on a run-in product so one would expect subsequent production units to sound best after a significant period of use. There's perhaps some currency in this, but in my experience it also reveals a decidedly romantic, rose-tinted view of modern product development processes and I don't find it convincing. Pressures of time to market now mean that development processes have to be run in parallel, not in series. Gone, or at least rapidly going, are the days of one engineer endlessly tinkering with a precious prototype until its ready to be presented to an expectant world. The larger scale loudspeaker companies in particular, will most likely be working with a rolling batch of development prototypes - older ones being left behind and new ones being introduced as, for instance, off-tool samples of new parts become available. The "designer works with run-in product" theory also raises the question of why, when an engineer starts development with virgin product, he or she is not sent down the wrong development route in misguided attempts to improve the performance when all that was needed was to leave it playing for a while?

And the third reason that I doubt. Well apart from the multitude of external influences on the performance of a hi-fi system (mains quality, ambient temperature and humidity, background noise level, mood, reversible thermal changes) that make irreversible running-in effects so hard to quantify, there's also the complexity and adaptive nature of our own hearing to stir into the mix. Open any psychoacoustics text book and you'll get some idea of the complexity inherent in our perception of, and response to, sound. There are well documented adaptive phenomena, both long and short term, that could very plausibly be responsible for the perception of running-in. And while I have little evidence for or experience of genuine mechanical running-in effects, I do have experience of listeners adapting their response to a system as they become used to it. A tonal balance that is unacceptable at first can, over time, become acceptable as the ear and brain adapts. If that's not "running-in" I don't know what is. I'm not saying that running-in doesn't exist and isn't a completely valid response to a new product, just that it may well not reside in the equipment but in the listener.

As any prudent scientist will argue: if you have an unexplained phenomenon, the source of which might reside within either a complex or a simple mechanism, look for it in the complex one before you add unwarranted complexity to the simple one. A speaker is typically a roughly assembled mechanism of wood, metal, plastic and glue. The human brain is the most complex mechanism our planet has to offer (yes, yes, mine naturally excluded). Go figure.

Of course I wrote the first part of this article knowing the results of the experiment. If the set of measurements following revealed any of the evidence we might expect from irreversible changes in the speaker's materials and mechanisms then I would have been writing a very different, and I admit, somewhat surprised piece. But the fact is, there were no changes between before and after run-in that were greater than either the expected experimental error, or perhaps more significantly, greater than the typical pair match or production consistency. So we can't even argue that the fine detail or minor changes are significant. If they were, it would be all but impossible to manufacture consistent product. Each speaker produced would sound as different as is typically reported is the case between virgin and run-in.

Of course I may have got unlucky and chosen, in the CDM1 NT, a speaker that needs no running in (in which case B&W wasted some words in the Owner's Guide). But in the cold light of day, even a relatively advanced product like the CDM1 NT is really not so different in terms of mechanisms and materials from countless other speakers. Or maybe a suggestion I heard while researching this piece; that running-in is only significant for drivers with less than perfectly aligned components (imagine a cone set at slightly the wrong height. Perhaps this would introduce stresses in the suspension that are "massaged" out with use?) is the reason that a blank was drawn? Either, or both, of these reasons might be significant. But even while I appreciate that it's impossible to prove a negative, the absence of evidence for any "mechanical" running-in in this case is pretty strong.

The measurements were carried out over two sessions by loudspeaker measurement consultant Phil Knight and myself. Measuring equipment was a Toshiba PC running MLSSA Version 9.0, with custom made pre and de-emphasis filters and a custom made driving amplifier. For the acoustic measurements a calibrated B&K measuring microphone and custom made power supply were used with the mic at 1m from the speaker aligned with a point between the bass/mid driver and tweeter. The measuring environment was a large, quiet, empty room (a middle of nowhere village hall if you must know) that enables measurement accuracy and confidence down to 200Hz. Curves showing low frequency data were produced by generating a mathematical model of the speaker's low frequency parameters from measurement of its impedance

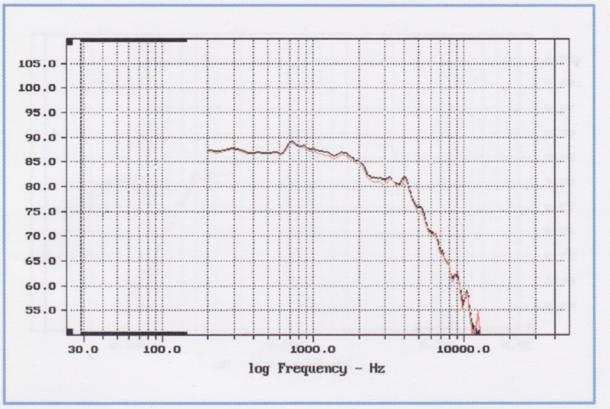

Figure 1.

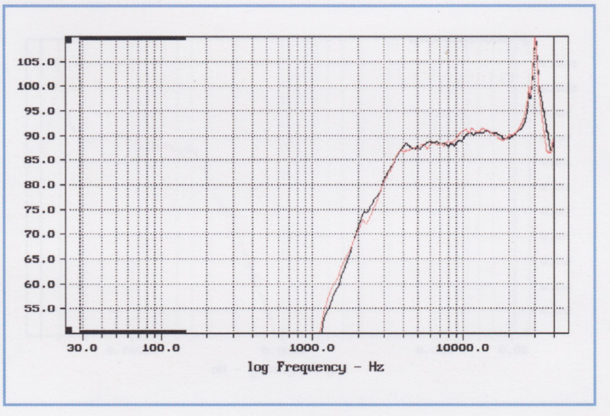

Figure 2.

(this low frequency response modelling technique is explained in the side-bar).

Figures 1 and 2 respectively show the bass/mid driver (200Hz to 12kHz) and tweeter (1kHz to 40kHz) frequency responses through the crossover (the bi-wire terminal panel allowed tweeter and bass units to be measured separately) for each speaker before running in. These curves illustrate the baseline pair match of the speakers. Both tweeter and bass/mid drivers are generally within ±0.5dB through the passband although the tweeter match goes a little haywire in the high-pass roll-on region (this is partly because tweeter suspension stiffness, like bass unit suspension stiffness, is hard to control in manufacture and partly because the steep slope of the curve exaggerates Y-axis difference). ±0.5dB is however a pretty good pair match. You'll see precious few speakers achieve better.

Figures 3 and 4 again show a bass/mid unit and tweeter match but this time it's the same speaker measured twice a month or so apart. These control curves demonstrate the basic confidence inherent in the measuring process. Generally the accuracy is again within ±0.5dB although, again, there's some inconsistency in the tweeter roll-on region - temperature differences between the two measuring sessions and, again, the "slope effect" are most likely to blame. The value of

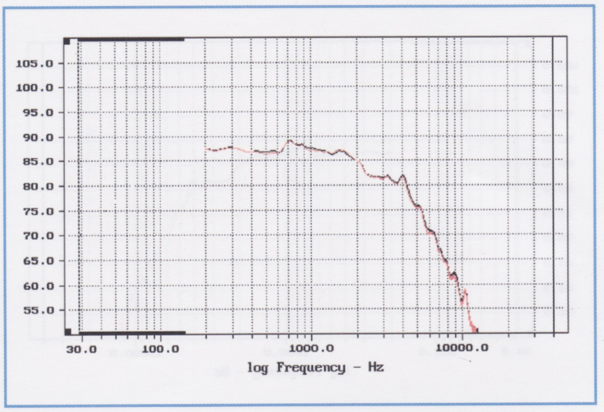

Figure 3.

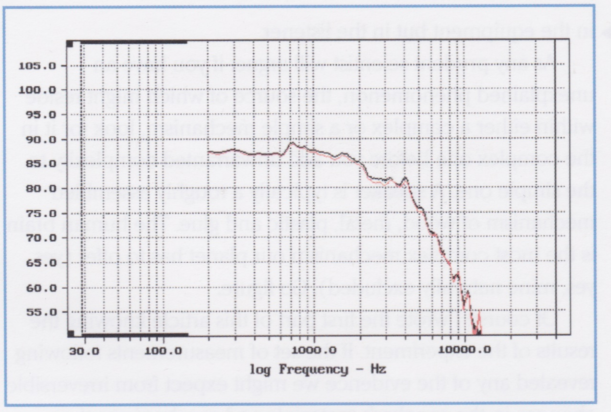

Figure 3b.

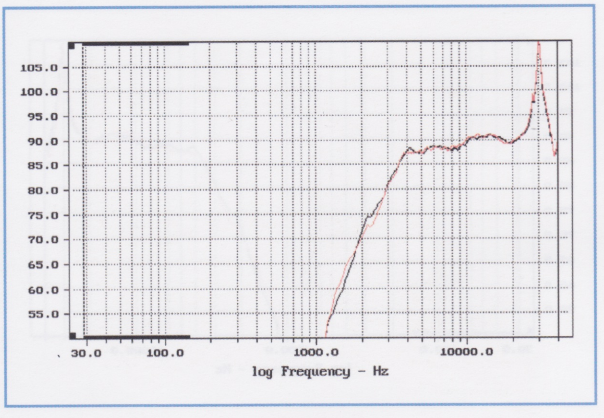

Figure 4.

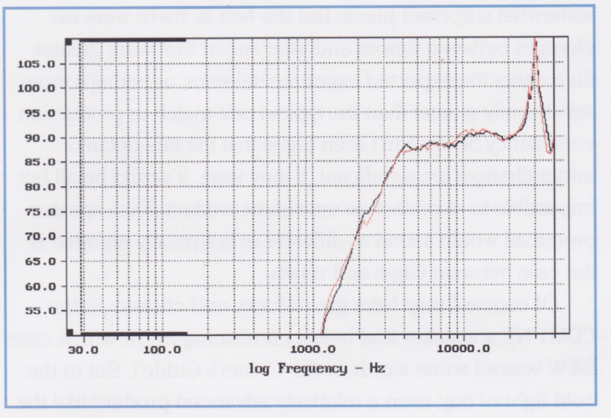

Figure 4b.

these curves is that they provide confidence that when we come to compare the before and after measurement for the run-in speaker, the changes revealed will be real and not measurement artefacts.

Figures 3b and 4b show the same information as Figures 3 and 4 but here the data is presented with one curve referenced to the other. Presenting the data in this form makes it far easier to appreciate changes (the vertical scale is now 1dB rather than 5dB) and is the method I've used to present the data for before and after run-in - Figures 5 and 6.

Figures 5 and 6 illustrate pretty clearly that there were no changes in frequency response from 200Hz to 40kHz between virgin and run-in that were greater either than the accuracy across the two measurement sessions or greater than the basic pair match of the speakers. Ironically, the changes in the control speaker (Figures 3b and 4b) are

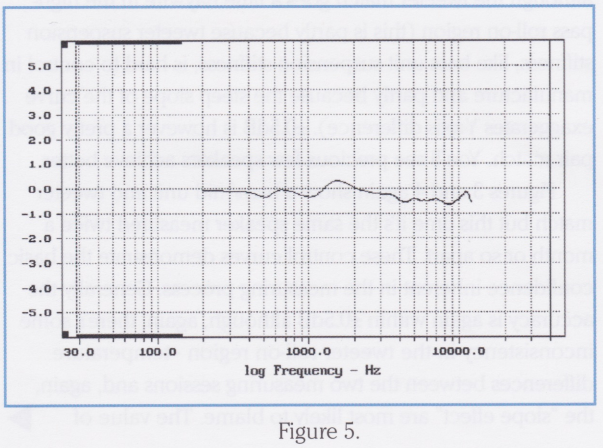

Figure 5.

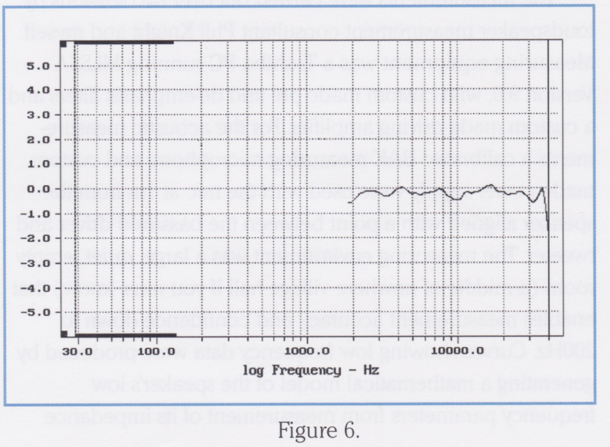

Figure 6.

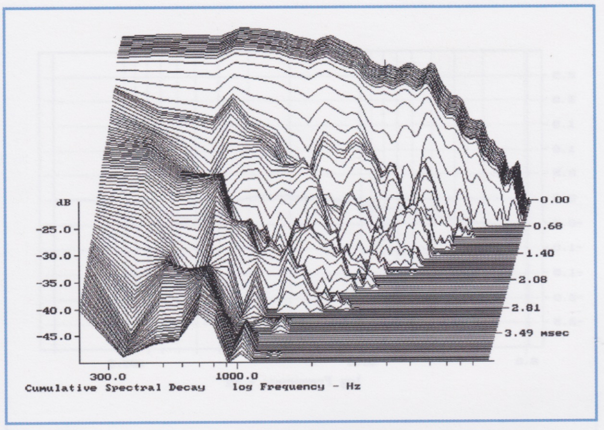

Figure 7.

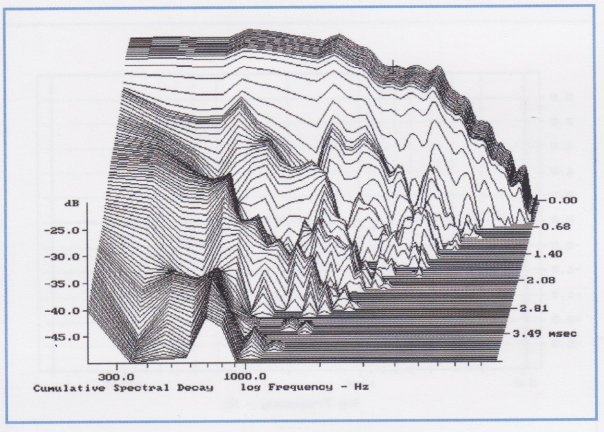

Figure 8.

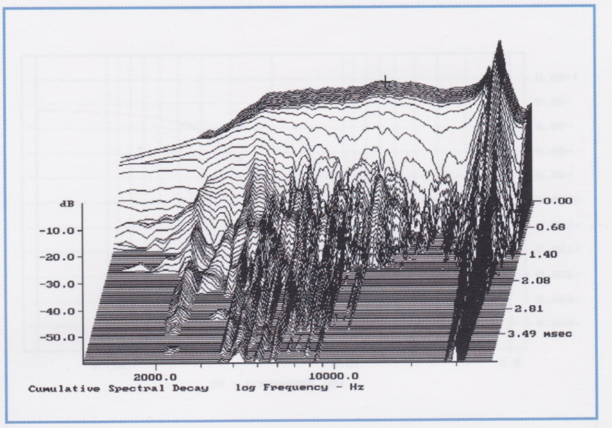

Figure 9.

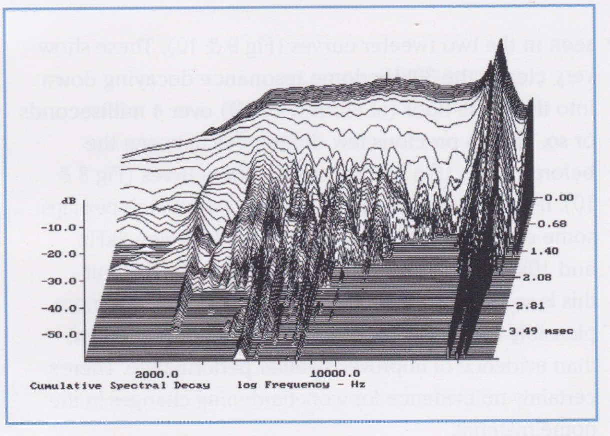

Figure 10.

slightly greater than the changes in the run-in speaker. No, we didn't get the speakers confused. Most likely we were just a little more accurate with speaker and microphone positioning on one speaker than the other.

Of course simple frequency response measurement tells us little about the time domain behaviour of a speaker. Perhaps running-in changes only reveal themselves in the energy decay behaviour (although I'm at a loss to explain why that might be). Figures 7, 8, 9 and 10 show the Cumulative Spectral Decay curves before and after running-in of the bass/mid units and tweeters (Note: The vertical scale on the tweeters is compressed compared to the bass/mid units simply so that the tweeters +20dB resonance peak could be included).

A CSD curve effectively illustrates how well a speaker stops. Imagine that the speaker is reproducing a wide band noise signal that is instantaneously turned off. The rearmost curve is the steady state frequency response, while time is displayed running from back to front. So the curves towards the front illustrate the speaker's energy decay characteristics with frequency. A good illustration of how a CSD curve illustrates energy decay can be

Low Frequency Response Modelling

Rather than make any attempt to measure directly the low frequency performance of the CDM 1NTs, which for accuracy down to the low Hertz would demand a really vast empty space, the low frequency curves were generated from a mathematical model of the speaker's low frequency parameters. Neville Thiele and Richard Small described in the 1970s how a speaker's low frequency characteristics could be described using analogue electrical filter theory. Their work showed that the mass of the various moving parts, stiffness of the suspension, mechanical and electrical damping, and the acoustic load presented by the enclosure could be seen as analogous to electrical inductance, capacitance and resistance components in an electrical circuit - an electrical circuit with an input and an output. If the analogous circuit is arranged appropriately, its output is the acoustic response of the speaker and enclosure, while its input will have exactly the same impedance as the real speaker. So if you measure the low frequency impedance of a speaker, set up a computer model of its analogous circuit (there are standard circuits for closed box, reflex, coupled-cavity, etc.) and then run a computer optimisation routine to find the unique set of component values that "fit" the measured impedance, you end up with an accurate (to within typically ±0.25dB), mathematically generated low frequency response curve.

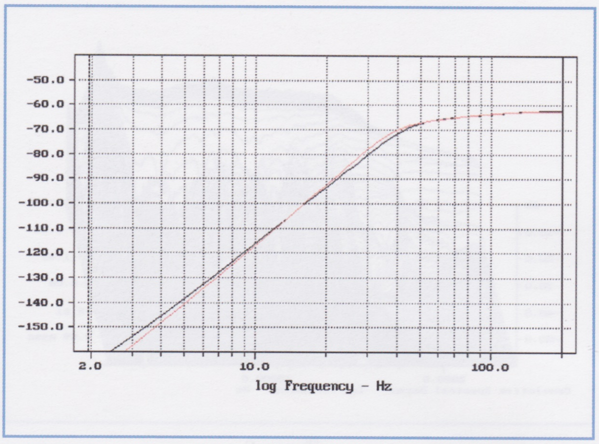

Figure 11.

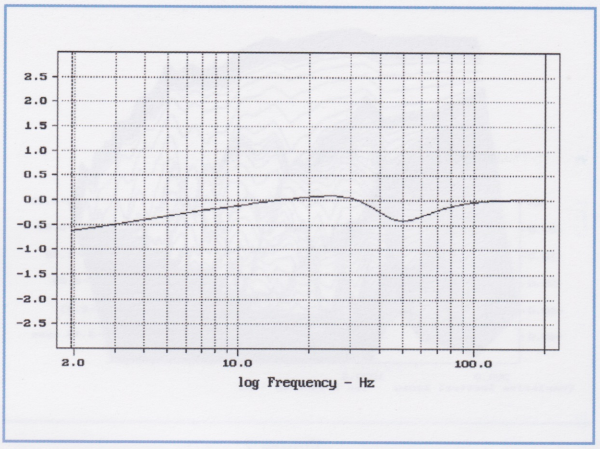

Figure 12.

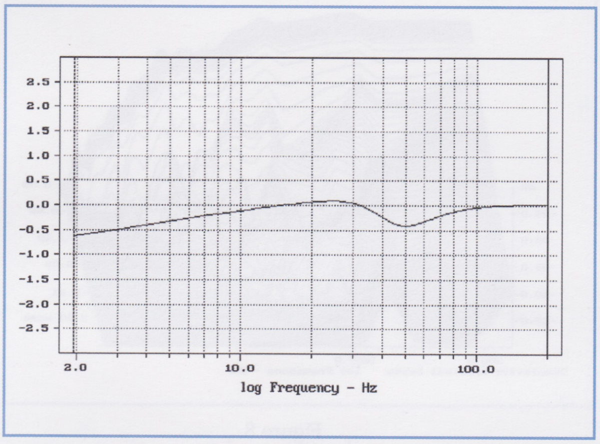

Figure 13.

seen in the two tweeter curves (Fig 9 & 10). These show very clearly the 30kHz dome resonance decaying down into the noise floor (at around -55dB) over 4 milliseconds or so. There's precious few differences between the before curves (Fig 7 & 9) and the after curves (Fig 8 & 10). If you squint at them for long enough there's perhaps some evidence of less low level hash between 2kHz and 10kHz in the "after" tweeter curve (Fig 10), but this is so far down towards the noise floor that it is more plausibly the artefact of a slightly noisier measurement than evidence of improved tweeter performance. There's certainly no evidence for work-hardening changes in the dome material.

The last three curves, Figures 11, 12 and 13, show the low frequency alignment and pair match of the CDM1 NT, the consistency between measurement sessions, and the before and after run-in change. Again, the difference curves (12 and 13) are presented with one measurement referred to the other.

Figure 11 shows the low frequency response shape and pair match of the speakers. There's a mild 1.5dB discrepancy between 20Hz to 40Hz. This is the port tuning region which suggests that a difference in the speakers port damping is to blame. The most likely cause is differently, arranged stuffing in the two enclosures. Although 1.5dB is a less good result than I'd have expected, it probably isn't particularly significant in practice. The real-world compromises of positioning in a domestic listening environment will probably swamp any mild bass pair match errors.

Figures 12 and 13 show the measurement consistency achieved on the control speaker and the change on the run-in speaker respectively. Again, and a little frustratingly, the control speaker shows slightly more change than the run-in speaker. Neither however exceeds 0.5dB till we get down to 2Hz. Almost certainly we're in the realm of random experimental error here, a conclusion re-enforced by the fact that while both curves show negative change around the port tuning frequency (ambient temperature change between measurement perhaps?), below 20Hz (where the output would be at -50dB) one shows positive and one negative change.

Finally, I wrote above that there's one mechanical running-in change of which I've had definite experience - a bass/mid driver's free-air fundamental resonance frequency falling. Well you can guess what happened on the run-in speaker can't you? It fell, yes, but by a somewhat paltry 1.98%. At least the control speaker only changed by 0.5%.

Thanks are due to B&W, Chris Binns, Phil Knight and Dr. Francis Rumsey for help and advice in writing this piece.