The Gentle Art of Room Correction.

BOOTHROYD STUART MERIDIAN

The advent of high-powered DSP (digital signal processing) systems has led manufacturers to attempt to equalise their systems to compensate for problems with the listening room acoustics.

However, simply throwing DSP power at the problem is not necessarily the most effective approach. *Richard Elen* looks at how Meridian Audio has addressed the problem with its Room Correction system, included in the latest version of its high-end digital home theatre processor.

Meridian's 800 optical disc player and 861 processor represent the company's flagship products at the high end of the consumer audio and home theatre industry. The 861 now ships with Meridian's unique Room Correction system, including a new DSP card and system software.

No listening room is perfect, and even in the professional environment, where a room can be purpose built for recording, mixing or mastering, with extensive acoustic treatment, there have to be adjustments to the signal being replayed to compensate for deficiencies in the room.

When it comes to the home listening environment, these problems are significantly magnified. Few but the most expensive listening rooms are purpose-built, and in most cases the room is used for other purposes along with the entertainment system.

The path to creating a good listening room requires careful design – from choice of room dimensions, to making sure that walls don't face each other, to the best selection of acoustic treatment. However, in most cases there are serious restraints placed on the design of a listening room. You will likely be obliged to use an existing room, for example, and there may be architectural or other features that you have to work around.

This article looks at the acoustic problems you will most likely encounter in a listening room, and how Meridian has chosen to help ameliorate them, by means of Meridian Room Correction.

You'll also find out how *not* to solve problems, and why other current approaches are not the most effective methods for dealing with common acoustic problems in the listening room.

Why 'room EQ' doesn't work

The traditional approach to cleaning up room acoustics in the professional field – both in the studio and for live sound reinforcement – and more recently in the home environment, has been 'room equalisation'. Here, a pink noise source is used to map the response of the room. 'Pink' noise has a constant power per octave. To our ears, this sounds like a 'natural', even noise.
The noise is played back through the

speaker system and captured by a microphone, whereupon it is analysed. Often, analysers used for this kind of work will generate a plot of level against frequency that indicates where peaks and troughs in the response lie. This is known as the 'magnitude response' of the loudspeaker-room interface.

A typical method of approaching these problems is to use an equaliser to simply invert the curve – where there is a peak of, say, 3dB at a certain frequency, the equaliser provides a trough to cancel it out; where there is a dip in the response, an equivalent compensatory boost is given to the signal. Technically, this is known as 'equalising the magnitude response'.

Originally, and particularly in live sound applications, a graphic equaliser would be used to apply EQ at around the frequencies indicated by the analyser.

There are several problems with this approach, not least the fact that it doesn't actually work very well.

To begin with, it aims to arrive at a 'flat' response, and if this were actually achieved, the sound would be extremely boring, sucking all the life out of the room.

In addition, some equalisers often used for this purpose (such as the graphic EQ mentioned earlier) produce all kinds of undesirable effects when they are added to the signal path, such as ringing and phase errors which can be worse than the problem they are trying to fix – although modern DSP-modelled equalisers can be designed to avoid some of these challenges. For similar reasons, it is also problematical to try and perfect the room impulse response (the way it responds to brief sounds or 'impulses').

Why you can't invert the response curve

More important, however, is the fact that simply inverting the full-bandwidth loudspeaker-room response (amplitude and phase) has a number of other problems, some of which are insurmountable.

Listening rooms come in all shapes and sizes, and many have distinct acoustic problems.

BOOTHROYD STUART MERIDIAN

Visualising the size of sound

We can hear frequencies from a few Hertz (cycles per second) to about 18kHz, although as we get older we tend to lose the ability to hear higher frequencies, and if we are aging rock musicians we may be missing some higher-mid frequencies too.

Sound travels at about 345 metres per second in air at 25° Celsius. This can be altered by both temperature (lower temperature reduces the speed somewhat) and humidity. The frequency of a sound and its wavelength are related according to the formula:

$$\lambda = v / f$$

(where f is the frequency in Hz, λ is the wavelength in metres, and v is the velocity of sound in metres per second).

This means that the *wavelength* of a 200Hz wave is about five feet (or 1.725m).

However, once you get up to a frequency of 1kHz, the wavelength is only about a foot (0.345m to be precise).

As a result, you can easily see that trying to control the response of a room at higher frequencies will affect only a tiny area. Move your head and you'll miss it.

That's one reason that Meridian Room Correction only operates below 250Hz.

You can also see how the effect of objects, and the effectiveness of acoustic treatment, in a room decrease as the frequency gets lower. An absorber will be only half as effective at 30Hz as it is at 60Hz.

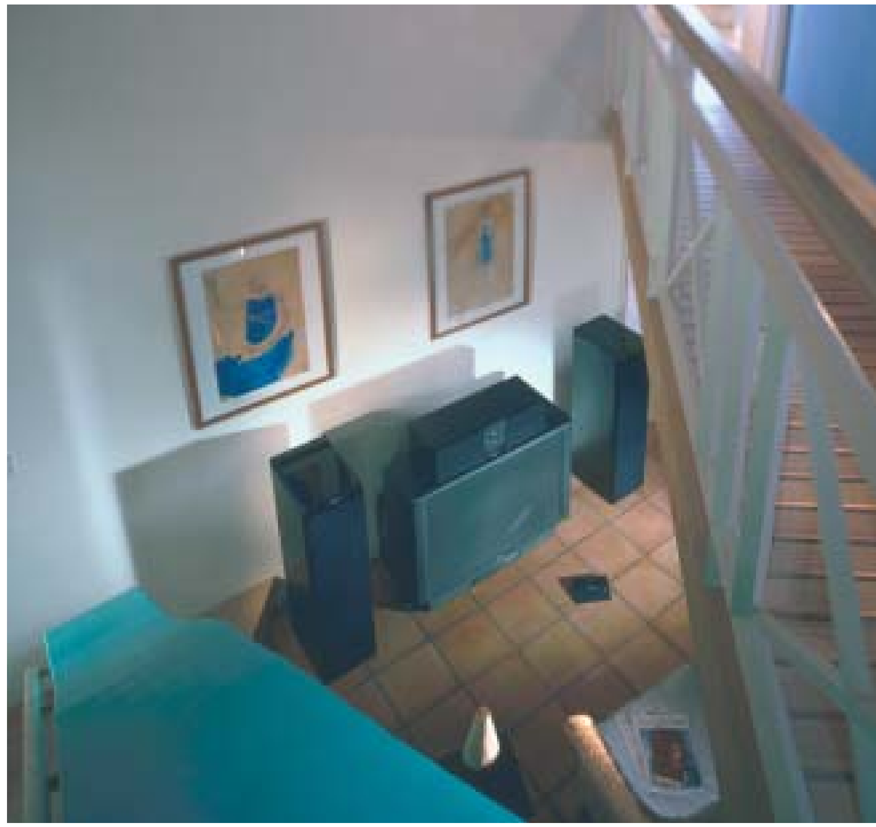

Fundamentally, it is actually impossible to accurately invert the loudspeaker–room response. There are some technical reasons for this – room responses are generally what are known as 'non-minimum phase'.

To change a non-minimum phase response into a linear phase response requires that signal starts to emerge from the filter before the sound has arrived at its input. This requires the filter to know what is going to happen in the future. Dr Who could probably design filters like this, but no Earthly engineer can.

Not only that: if you were to successfully create an equaliser that handled non-minimum phase responses, it would introduce a number of undesirable effects such as audible pre-echoes and latency (where it takes some time for the signal to pass through the system), leading, for example, to lip-synching problems with video.

And the problems only start there. The magnitude response of a speaker in a room will have significant dips as well as peaks. If you try to compensate for the dips by boosting at those frequencies, you can easily run out of headroom in the system. Imagine you have a 100W amplifier. To cancel out a 10dB notch, you would need to trade that in for a 1000W amp.

In addition, we are very sensitive to the effects of room resonances and remember, a dip in one place in the room may be a peak in another.

Indeed, equalisation like this only really works for one listening position. It is quite possible that in making the room sound better to one listener, we might make it dreadful for the person sitting next to them. In fact, attempting to improve the listening in one spot in this way can make the sound significantly worse for everyone else in the room *(see sidebar, left)*.

Evidently, this approach is not a satisfactory one.

Considering room modes

A better technique is to look at 'room modes'. These are commonly known as 'resonances', and occur as a result of the setting up of 'standing waves' in an acoustic environment.

Standing waves are created by sound waves travelling through space, encountering a boundary such as a wall, and being reflected back on themselves. The original and the reflected waves interfere with each other and, where they reinforce, produce 'antinodes', areas of louder sound (and corresponding 'nodes', areas of quieter sound, where they interfere destructively). These areas stay in the same place in the room for a given frequency and do not move about – hence the name 'standing waves'.

The frequencies at which resonances occur can be calculated relatively easily for simple rooms. The formula for the fundamental (lowest) resonant frequency is given by:

$$f = v/2 \times d$$

(where v is the speed of sound in air, 345 metres per second at 25° C, rising a

BOOTHROYD STUART MERIDIAN

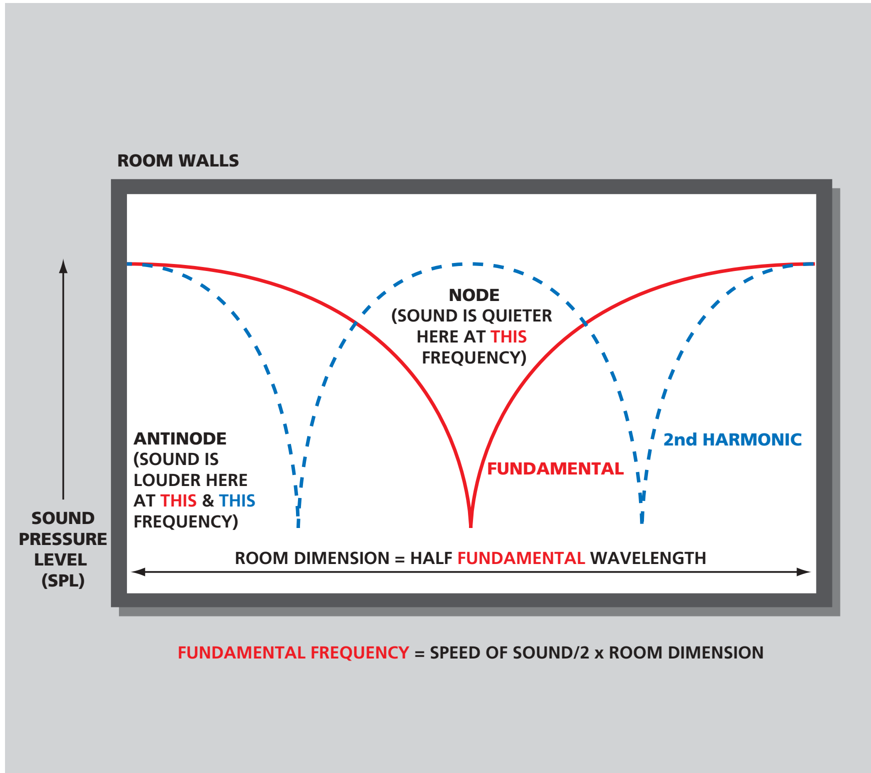

Fig. 1. Room modes, or resonances, are caused by standing waves. These are created by sound waves travelling across a room, and being reflected back from the opposite wall, to and fro. The frequencies at which the room resonates are determined by the distance between the walls. This diagram shows the fundamental (red), the lowest (and most significant) frequency at which the room will resonate, and the second harmonic, at twice the frequency (blue). The sound appears louder at any resonant frequency where the air vibrates the most – these points are called antinodes. There is an antinode at each wall, at all frequencies, where the wave is reflected back and interferes constructively. The sound is quieter where the vibration is least – at nodes. The depth of the nodes is dependent on the acoustic treatment of the room – a better room has nodes that are not as deep. This is because acoustic treatment reduces reflections, and therefore there is less cancellation between the reflected waves.

little with temperature, and d is the room dimension in question).

This sounds fairly straightforward, until you remember that the pattern of standing waves in a room involves three dimensions – height, width and length – each with their own standing wave patterns. The situation is further complicated by the fact that standing waves are not only created at the fundamental frequency, but at various multiples of it (harmonics). Furnishings further complicate matters: depending on their size and construction they can make things better, or worse, or have no effect whatsoever. It's rather difficult to tell just by looking at them.

You can probably already see, however, that there are some fairly basic things you might like to think about when setting up a listening room. For example, you might imagine (correctly) that a room in which two or three of the dimensions have a common factor, such as the width being the same as the height and half the length, might produce difficulties, because there would be common resonances between different pairs of walls.

Luckily, standing waves only really cause a problem at low frequencies. At higher frequencies, the harmonics are often close together and overlap, and do not affect the sound as obviously, as the intensities of the harmonics are far less than that of the fundamental, and are even more dependent on listener position. This is one reason why trying to correct room problems at high frequencies can produce more problems than it solves.

Another problem at higher frequencies is that while humans can tell the difference between direct and reflected sound, a measuring microphone cannot. We are trying to correct the *room*, not the loudspeakers, which we assume are already fairly well optimised. At high frequencies, though, it's hard for the system to tell.

The need to avoid messing with higher

BOOTHROYD STUART MERIDIAN

frequencies comes into play in another respect too. Stereo and surround localisation is an important part of the listening experience: it is important that we do not affect the phase or level of sounds above 200–250Hz as these are vital in our perception of the stereo or surround image. We also want to avoid placing the speakers as if in a different room to the listener.

At the bass end, the resonant frequencies are fairly easy to identify and tend to be quite separate.

What are the effects of low frequency resonances on the sound? Here are some of them:

- Some frequencies are emphasized or de-emphasized – if you listen to a sweep tone, or notes played by a bass guitar, for example, you might notice that some frequencies or notes are much louder than the others, and some are much softer.
- Some notes will hang on or ring much longer than others, making them stand out.
- Pitch changes can occur during the decay of a note – this will be particularly noticeable when the note is close to, but not at, a resonant frequency.
- Short notes appear to change pitch.
- Echoes occur, where a single note is turned into two or more shorter notes. This may be particularly noticeable when the note is between two closely spaced resonances.

Unlike the entire loudspeaker–room response curve, which cannot be inverted accurately to provide a room equalisation curve because it is non-minimum phase, controlling room modes is more practical because they *do* exhibit minimum phase behaviour. In addition, while the magnitude response of real rooms – ie, rooms which do not have perfectly reflecting walls – may vary a great deal in different locations in the room, the decay time of a resonance can be measured successfully almost anywhere in the room, and does not vary much with listener location.

In addition, the filters you build as a result, which will be designed to control the reverberation time at specific frequencies, will work anywhere in the room too.

Meanwhile, in the measurement phase of the operation, detecting room modes in this way is less affected by variations in the sensitivity of the microphone at different frequencies – even a cheap mic will work fine.

Filtering and Resonance

Successfully handling room resonances requires two things: an ability to detect the room modes and the ability to build filters to control them. By 'controlling' a resonance, we mean flattening it out, so that it is less 'peaky' and less pronounced. This is known as reducing the 'Q' ('quality') of the resonance, and it can be shown that reducing the Q is equivalent to reducing the decay time of the room at that frequency – or, more accurately, producing a filter that, taken with the room resonance, produces a shorter resulting decay time at that frequency than the resonance on its own. The type of filter required to do this job is a notch filter – this approach to room correction does not employ peaking filters (for example to deal with dips in the room response by increasing the reverb time), as they can introduce loss of headroom and other problems.

The reverberation time of a room is usually quoted in terms of its 'RT_{60}'. This is the number of seconds it takes for a sound (typically a short-duration 'impulse') to decay by 60dB.

So the aim of our room correction system is to identify important low-frequency room modes – and their Q – and construct notch filters to reduce the

Meridian Room Correction was first introduced for the 861 Digital Processor, whose card-based structure makes it easy to add new features and hardware. The 861 is shown here with the covers removed, on top of an 800 Optical Disc Player.

excessive decay time due to resonance at those frequencies to something more like the overall decay time of the room. This is known as the 'Target RT_{60}'.

Of course, introducing a filter to control room resonance means that the direct signal path is filtered as well as the excitation of the room, so it's important to investigate the overall impact of the filter on what a listener will hear in the room. In practice, the effect is very positive: the effects of the resonance are significantly reduced and minimal other effects are introduced in their place. It can also be demonstrated that cascading a number of filters (to handle more than one room mode) does not have any untoward effect on the sound as a whole, and a number of filters can be used, each controlling one room mode.

However, it is a fact that what we are doing is to modify the signal fed to the speakers so as to smooth out the excitation of the room at the bass end. There will inevitably be some change in the shape or envelope of initial transients at the filter frequencies. What impact will this have?

The fact is, not very much. Only very specific frequencies are involved: these are below about 200Hz, where there is little in the way of stereo or surround localization. In practice, any low-frequency transient modifications are masked by what is happening at higher frequencies.

Meanwhile, it's a fact that the undue emphasis given to certain notes by uncorrected room resonances has a significant *negative* impact on the listening experience. Applying appropriate room correction gives the listener a much improved impression of a recording: it is especially easier to 'listen into' the acoustic of the original performance.

Real rooms, real modes

When applying correction to real listening rooms, what kind of rooms do we need to consider? Well, the average home theatre or music listening room has a floor area in the range 10–50 square metres and a room volume of 25–250 cubic metres. The average distance from the listener to the front loudspeakers is in the region of 2–5m.

BOOTHROYD STUART MERIDIAN

The power of processing – multiplied

Meridian's Room Correction system brings additional digital signal processing power to bear in a system that already possessed an astonishing level of DSP capability.

The previous version of Meridian's flagship 861 processor, for example, used an AB00 DSP card, with a pair of Motorola 56002 chips delivering the ability to process 60MIPS (million instructions per second), in addition to another 300MIPS on the main CD20 DSP card.

Version 4 replaced the AB00 card with an EF20, featuring a pair of Motorola 56367s running in double precision (48-bit) mode, providing an additional 300MIPS and allowing all processing to be carried out at high resolution. This all adds up to a total of 750MIPS (including 150MIPS of DSP on the IE42 digital input card for de-jittering and upsampling) – no less than four times as much DSP power as any other comparable consumer audio product.

The new DSP card also provides additional RAM and 8 megabits of Flash memory for ultra-high-speed loading of DSP instructions.

All this additional processing power makes it possible to develop room correction filters with unmatched precision and fidelity. The filters are ultra-linear, and ultra-low-noise thanks to 48-bit double precision processing.

Even though most digital systems only offer 24-bit inputs and outputs, internal processing must have significantly longer word-lengths than the I/O in order to handle the extra bits produced by DSP operations. The simple addition of two binary numbers can add an extra bit, and there are many millions of these operations being carried out every second. In addition, operating in double precision dramatically improves the sound of a DSP system, making it noticeably more transparent.

An interesting discovery is that, unlike a recording studio control room or a movie theatre, this average speaker-to-listener distance is generally larger than the 'critical distance' for the room. The critical distance is the point at which the levels of direct and reverberant signals reaching the listener are equal. Even so, there will be a minimum of 2ms between the arrival at the listening position of a direct sound and its reverberation in the room reaching the listener.

It should also be noted that because of the comparatively small size of any home listening environment, it is not necessary to control *really* low frequency modes (below about 15Hz) because most rooms are too small to have any – and most speakers don't go down that far either!

Listening rooms generally have longer reverb times at lower frequencies than they do at higher ones. This is because the acoustic effect of carpets, furnishings, curtains and so on decreases with frequency – another reason why we do not need to control room resonances above about 250Hz.

In general, listening to a recording is more enjoyable if the room in which it is played back has a shorter reverberation time than the place in which the recording was made (assuming a fairly natural approach to recording). This enables the recorded acoustic, such as that of a cathedral or concert hall, to 'overlay' the listening room acoustic and essentially mask it. This becomes particularly apparent when replaying surround-sound recordings, in which the 'envelopment' you experience in a natural performance environment such as a concert hall is recreated in your listening room. However, you do not want the reverberation time to be too short: if it drops below about a third of a second, we tend to find the room sounds unnaturally 'dead': in addition, hi-fi loudspeakers are not designed to give their optimum performance in dead rooms.

The comments above concerned the overall reverberation character of a listening room. In the case of room modes, however, the resonances found in a room will have reverberation times that are significantly longer than those experienced at other frequencies. It is not unusual for the reverb time of a room at a resonant frequency to be two or three times as long as the room as a whole.

How you derive information on the room modes requires a lot of intense calculation on the basis of a fairly simple measurement technique, as does the creation of filters to ameliorate them. In the case of Meridian Room Correction as implemented in Meridian products, these calculations are hidden from the listener and we will not discuss them here.

If you are interested in reading a more technical description of the theory behind Meridian's Room Correction system, we invite you to read the paper, *The Loudspeaker–Room Interface – Controlling Excitation of Room Modes*, by Rhonda Wilson, Michael Capp and Robert Stuart, presented at the Audio Engineering Society 23rd International Conference, Copenhagen, Denmark, in May 2003.

Using Meridian Room Correction

There have been some attempts at room EQ in the consumer, as opposed to the professional, audio industry, but by and large they have attempted to invert the loudspeaker–room response, which as we have seen is not a satisfactory approach. They have also often attempted to control the entire audible frequency range, which is, as we have seen, unnecessary at very low frequencies (below 15Hz) because there are no resonances there; while above 250Hz or so, attempts to modify the response can cause more problems than they solve.

BOOTHROYD STUART MERIDIAN

Such systems, in addition, have generally provided only the filtering capability and it has been left to the user or installer to measure the room response (a process that often requires expensive equipment and some expertise), calculate the equalisation required, and input a set of complex parameters into the equaliser.

By contrast Meridian Room Correction (a combination of an additional DSP card and new software – Version 4 or later – in the case of the 861 digital processor) implements the room-mode-based correction theories outlined earlier in a manner that is easy for the owner or installer to perform. The measurement of the room and the building of the appropriate filters can be performed largely automatically, with only minimal work required on the part of the user. However, advanced users and installers can fine-tune the system manually for the best results in difficult environments.

Meridian Room Correction is implemented in conjunction with the equipment setup procedure, which is carried out with the aid of a personal computer. An extension to the Meridian Configuration program is used to lead the user through an initial microphone setup procedure, and then proceeds to perform a series of measuring tests. From the results of these tests, a filter profile is constructed, which contains a specific series of filters for each channel – up to a total of 60 filters for an entire multichannel profile.

The Meridian Room Correction effectively uses the Meridian processor to provide the 'front-end' of a powerful and highly accurate digital audio measurement system, which is capable of displaying the room response in a number of graphical ways. The computer then builds the filter parameters which, at the end of the procedure, are uploaded into the processor for implementation using the processor's prodigious Digital Signal Processing (DSP) power.

You begin by plugging a microphone – typically a low-cost basic sound pressure

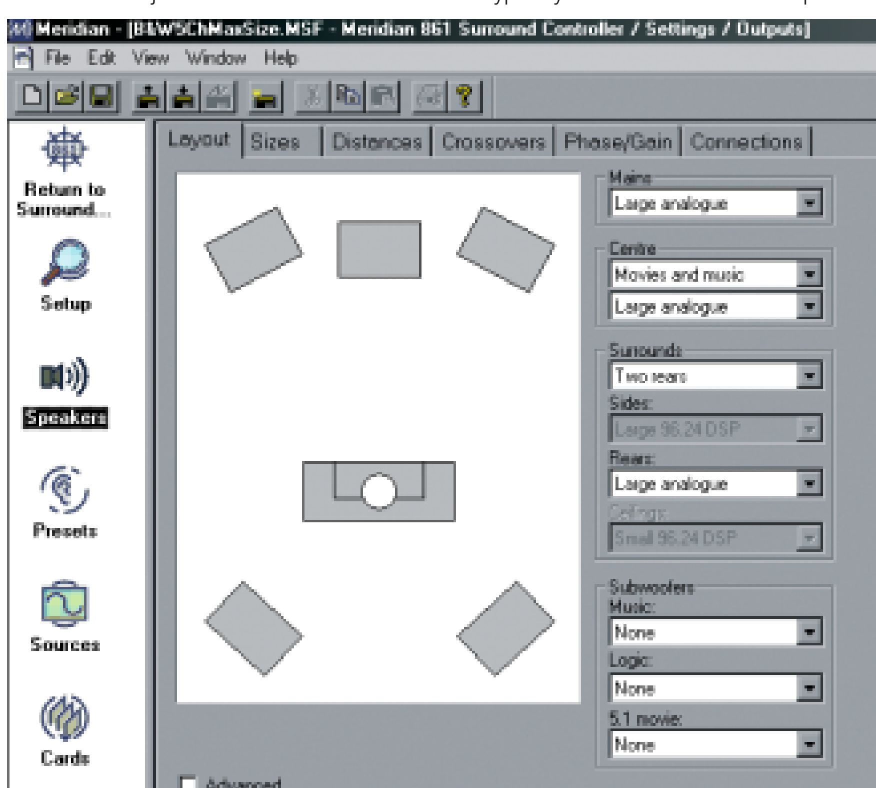

Meridian Room Correction will calculate filters for any stored loudspeaker configuration.

BOOTHROYD STUART MERIDIAN

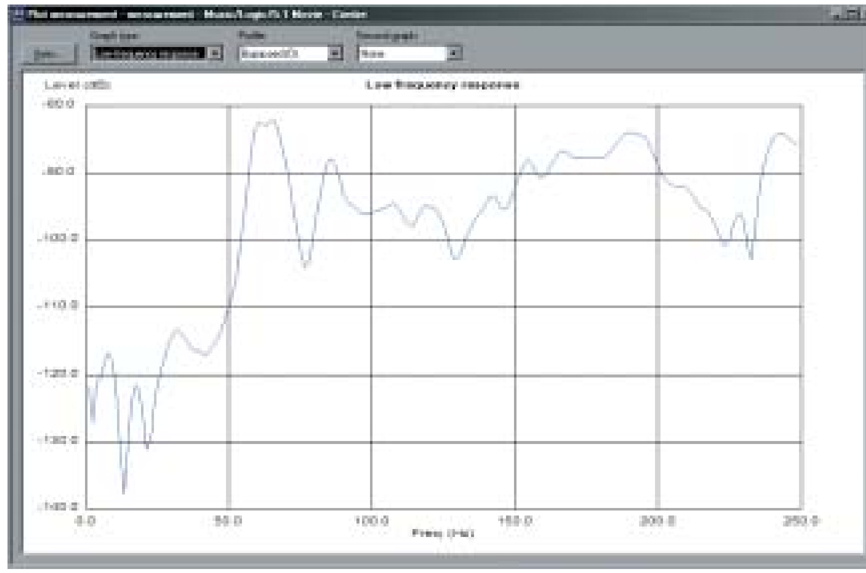

This plot shows the original low-frequency response of the room in graphical form, as measured by the Room Correction system.

level meter with a line-level output – into an analogue input on the Meridian processor. New Meridian 861 units include the special EF20 DSP card required for room correction; an IA00 analogue input card; and the new software, so everything that's needed is provided except the sound meter! The microphone and its connection is tested automatically before the program goes on to begin the room optimisation procedure, which can be performed automatically or, for the experienced user, manually. A special test signal is played through each speaker in turn and each channel in the speaker layout is measured. Depending on the number of channels and speaker layouts in your system, this may involve up to 24 tests and take up to 20 minutes to complete – but it's virtually all automatic: there is even an automatic countdown so that you can leave the room before the tests.

When the measurements have been completed, the system begins to build the filters required for each channel in each loudspeaker layout. The filters are displayed graphically as they are constructed, and finally the parameters can be saved to the processor.

This is the way the system will generally be used. But you can go more deeply into the room correction process.

It might be helpful to take additional measurements, for example with the sofa in a different place in the room – perhaps you listen to music with a different furniture configuration to the one you use to watch a movie.

The Measurements panel gives you the opportunity to review and compare measurements as well as taking new ones. Measurements can be viewed graphically in three different ways, selected from a Graph Type menu.

The **Impulse Response** mode shows the response over time of the loudspeaker to a click. The plot shows a maximum just after time zero, with a rapid decay indicating the amount of reverberation in the room. Any additional peaks indicate resonances set up by the original impulse.

Another display shows the **Low Frequency Response** *(see screen shot above).* This is the loudspeaker–room response to a low-frequency sine wave swept between 0 and 250Hz. Ideally, this would be a smooth curve.

Peaks in the curve indicate resonances at the corresponding frequency, while dips indicate cancellations or absorption at the frequencies at which they occur.

BOOTHROYD STUART MERIDIAN

Most interesting of all are the **Waterfall Plots.** These three-dimensional images show the response of the room to a series of short sine wave pulses between 0 and 250Hz, and shows how the response decays with time at each of those frequencies. Ideally, a waterfall plot would have a smooth profile, and decay evenly to a low level within 0.5 to 0.8 seconds. Peaks indicate resonances, while dips indicate absorption at the corresponding frequency. Ridges extending towards you indicate reverberation at the indicated frequency. As you can see, these plots are particularly useful for indicating the presence of room modes (resonances).

It's also possible to bring up two waterfall plots on the screen at once, so you can make a comparison, for example, between the plot before and after applying room correction. The room correction aims to identify the strongest resonances and to reduce the decay time of each strong resonance to the average decay time for the room.

Advanced operation

In addition to the automatic procedure that builds filters for a profile – and there can be several profiles, corresponding to different speaker layouts or room configurations, for example – you can also manually build filters for a profile. Here, the target decay time can be set automatically by the program, or manually. Typically a decay time of 350 milliseconds is appropriate for listening at home, while 200ms gives a drier acoustic more suited to a recording studio control room.

Once filters have been built, they can be displayed either numerically or graphically, and edited – ideally after having saved a backup. Having selected a profile from the 'Active Profiles' list, clicking on 'Edit Filters' brings up a box with numerical information on the filters, including the filter number, centre frequency in Hz, gain in dB, and the filter bandwidth in Hz. Each parameter can be adjusted as desired. In addition, the

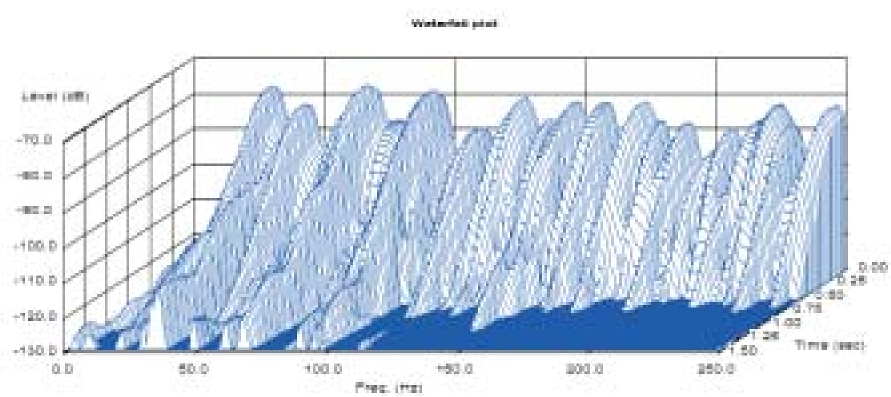

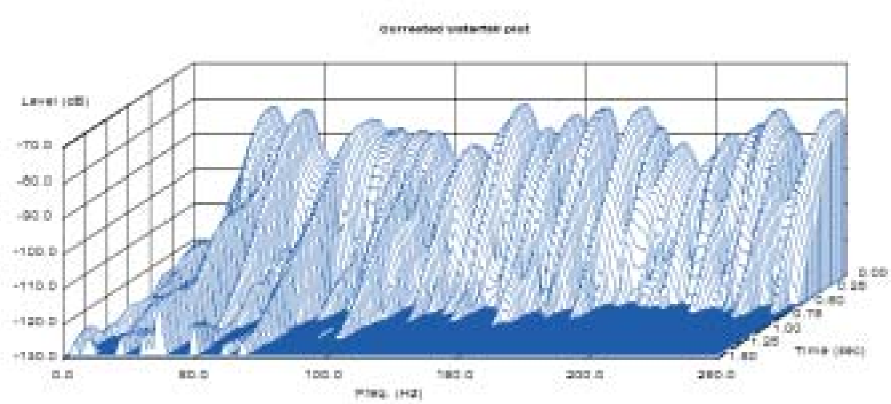

This 'before and after' 'Waterfall' plot of a fairly difficult room shows the frequency response of the room against time, and makes it easy to see the decay at different frequencies. The upper plot shows the original response; the lower plot after Room Correction has been applied. In particular, note how the peaks at several frequencies have been controlled, and the reduction in reverberation times at those frequencies, representing a significant improvement.

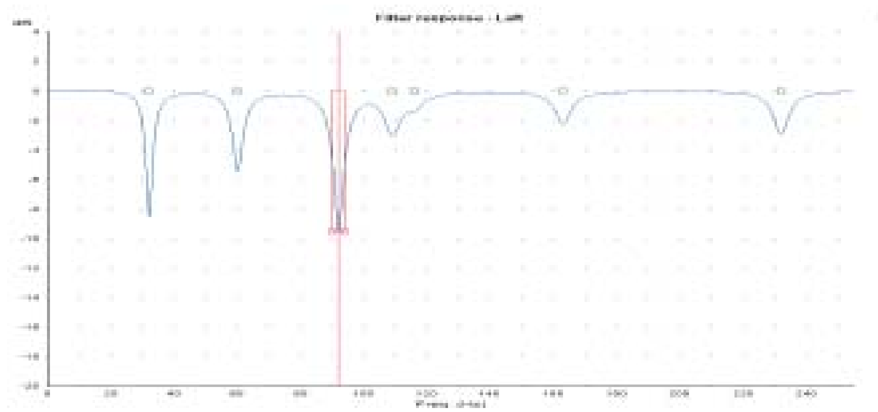

Graphical editing of filter characteristics within the Room Correction application.

list shows the Room Mode bandwidth and equalised bandwidth in Hz; the RT_{60} value in seconds, and the Target RT_{60} – in other words, the decay time that the filter is intended to produce at that frequency.

Editing, as well as display, can be carried out numerically or graphically. Clicking on the ‘Graphical’ tab brings up a display of the overall filter response. Each notch filter has a small square above it, and clicking on the square brings up an edit mode for the filter with which its centre frequency, width and gain can be adjusted by dragging. Filters can also be deleted if required.

The difference that Meridian Room Correction brings to a system is effective and immediately noticeable. But this does not mean that suddenly all your room response plots will be flat – as we have seen, that is not the idea. What Meridian Room Correction *does* do, however, is to bring the reverberation times of your room into line at important resonant frequencies.

While it can’t make a dreadful room into a good room, it *can* make a good room sound noticeably better. And in fact, the better the room acoustics, the less Meridian Room Correction will do, or need to do. But if you are a serious music listener, you will no doubt agree with our beta testers, who uniformly approved of the effect of Room Correction on their listening environments, with comments like, “Exciting”, “Much better sound”, “Improved bass”, and “Subtle but extremely effective”.

Meridian Room Correction can make a significant and worthwhile contribution to your listening pleasure, and help to make the performance of real rooms more amenable to the enjoyment of high quality music and audio.

Meridian Audio Limited
Stonehill, Stukeley Meadows,
Huntingdon, Cambridge PE29 6EX
United Kingdom
T +44 (0)1480 445678
F +44 (0)1480 445686

Meridian America Inc.
3800 Camp Creek Parkway
Building 2400, Suite 122
Atlanta GA 30331 USA
T +1 (404) 344-7111
F +1 (404) 346-7111

Further technical details, images, product reviews and company history are available from Meridian Audio or from our website,

www.meridian-audio.com

Version RC1b 160503
Written by Richard Elen
Designed by Allen Boothroyd & Richard Elen