Reproduced from
Studio Sound, July 1990

Why do equalisers sound different?

Michael Gerzon claims to have no definitive answers to this question but does offer interesting conjectures and food for thought

Every equaliser has its own sound, and for some time it has been suspected that this is not simply a matter of their amplitude responses. It has been believed that the ears are insensitive to the phase response of equalisers but now some people are starting to claim that, really, amplitude response is quite unimportant and most of the subjective effect of equalisers is due to their phase response.

Certainly, phase response is relevant, as noted by Phil Newell in his series on monitor loudspeakers[1]. He notes, correctly, that adjusting the polarity of speaker units to maximise flatness of frequency response can often give a much more coloured result than the other polarity, which can give a sharp dip in the amplitude response but a much smoother phase response. However, this does not prove (and neither did Newell claim) that phase response is much more important than amplitude response, only that both are important and must be carefully related to one another for the best sound.

Let's look at what determines the subjective sound of an equaliser. Unlike others, I do not claim to give definite answers. My aim, rather, is to make some tentative conjectures, report some rules-of-thumb that have often been used with some success and to raise some questions so people can give some intelligent thought to the problem and maybe eventually find some answers.

We will rapidly enter the treacherous areas of hi-fl subjectivism, however, unlike the woolly-minded approach of many in the hi-fl press, I believe that ultimately one needs no magic pseudo-science to explain the mysteries.

The problem with most 'objectivists' who demand measurable reasons for subjective differences, is that they are very narrow-minded about what kind of measurements they will consider. They often demand that measurements can easily be done on conventional audio test setups. We shall see that it is highly unlikely that some of the most audibly important aspects of equaliser response can be measured either in the amplitude or the phase response but that we shall probably have to look elsewhere.

Now this is very near heresy. It is a standard mathematical result in the mathematical theory of linear filters that the behaviour of any such filter is completely determined if one knows its amplitude and phase response. This is no longer true if the filter has nonlinear distortion – and many subjective differences between equalisers are believed to be due to nonlinear distortion effects. However, I claim that even if one has a perfectly linear filter and measures both its amplitude and its phase response, one will still not, from these measurements alone, be able to predict its sound.

I am a mathematician, and I do believe the theoretical mathematical result that the filter response is completely specified by its amplitude and phase response. The key words in the above are 'if one… *measures* its amplitude and phase response'. The point is that real-world measurements are never exact, and what we shall see is that incredibly small changes in amplitude and phase response, supposedly quite 'negligible' according to objectivist ideas, can actually have large audible effects. This is not to say that these effects cannot be measured, only that measurements of amplitude and phase responses are not the way to do the measurement.

If I prove to be right in my claims, we shall have to stop thinking of filters purely in terms of their amplitude and phase responses but will have to find other new ways of looking at them.

The evidence

In the late 1950s, H D Harwood at the BBC made a discovery whose importance is still not fully recognised. In investigating the performance of loudspeakers, he discovered that low-level delayed resonances severely coloured the reproduced sound even if these resonances were 40dB below the main speaker response. At first sight there's nothing very world-shaking about that. But consider what the effect of such a delayed resonance is on the amplitude and phase response; 40dB down means a signal whose amplitude is only 1% of the main signal. This means that the amplitude response must vary only between 99% and 101% of flat, *ie* within ±0.1 dB. The effect on phase response must similarly be within 1/100 rad, *ie* within ±0.6°. In other words even in the late 1950s Harwood showed that variations in phase response of around only 1° and in amplitude response of ±0.1dB produced audible colouration.

I am not claiming that all possible variations of phase and amplitude responses of this magnitude will produce colourations, only that specific variations produced by delayed resonances.

With the improvements in audio technology since that date, it would probably now be safer to tighten up these figures by a factor 10 – in other words to suspect delayed resonances 60dB down of colouration, even though they would cause amplitude variations of only ±0.01dB and phase variations of ±0.06°.

It should not have been a surprise in the mid 1970s when it was discovered that turntables had a 'sound', since measurements had long revealed bumps in their amplitude responses associated with delayed resonances in their mechanical system of a magnitude already identified by Harwood as being subjectively important in loudspeakers. However, until that time, most audio engineers had ignored the tiny bumps and kinks in their measurements of record frequency response as being too small to be audible, despite the prior evidence of Harwood's work.

In the mid 1980s, a second piece of evidence – that the ears could easily pick out tiny deviations in amplitude and phase response – emerged, in connection with digital filters. In an early attempt at digital noise suppression, Roger Lagadec, then at Studer, investigated a multiband digital noise gate that split the audio signal into 512 bands, noise-gated the bands separately and then put them back together again. Although this was very effective in reducing noise, it was discovered that there was a disturbing audible colouration, even if the noise-gating action was switched off. It was discovered that this colouration was due to the amplitude response of the filtered bands added together again not being quite flat – there was a ±0.1dB ripple in the frequency response. It was found that to remove the upsetting audible colouration required this ripple to be reduced to around ±0.001dB. In this case, all the digital filters had linear phase responses – so only the amplitude response could be blamed for the colouration.

So, from a variety of directions, we have been finding that tiny ripples in the amplitude and phase responses can have important subjective audible effects, causing obvious or even gross colourations. And yet, one can also produce large deviations in amplitude or phase response that cause almost no colouration, as we shall see. It becomes evident that simply looking at the magnitude of deviations from flatness of either amplitude or phase responses tells us very little about the subjective result. The ears are responding to something else entirely. But what?

Rules of thumb

Frankly, we don't yet know for sure how the ears I determine the degree of colouration of an equaliser. Over the years, however, many indirect clues have been found as to what is happening. If we cannot yet say who committed the murder and how, at least we can list some of the suspects and their possible motives.

One of the oldest suspects, whose guilt is believed in by many audio designers, hi-fi buffs and even professionally respected recording engineers, is known, in the best tradition of spy novels, by a single code letter: Q. The rule of thumb used by many designers is simple: it will sound coloured if the filter has a Q much larger than 0.5. I am indebted to Tom Holmes, formerly of Philips Hi-Fi Labs, many years ago for pointing out that Q smaller than about 0.6 appears to sound uncoloured, and larger than 0.6 appears to sound coloured.

Q is shorthand for 'quality factor'. In recording work, Q is encountered mainly as a number on parametric equalisers – with high Q corresponding to sharp peaks or dips and low Q to broad peaks or dips. However, the concept of Q applies to all filters, even to all-pass filters, which have a flat amplitude response but a nonlinear phase response. It is not intended here to give the mathematical theory of Q – if you are familiar with the theory of filter design you'll already know – however, an understanding of the meaning of Q can be given. Although conventional theoretical accounts of Q look at the frequency behaviour of an equaliser (*ie* its amplitude and phase response as a function of frequency) it is actually easier to explain the idea in terms of the time response of a filter.

All filters smear out in time any sharp momentary impulse fed into them and it is a commonplace of filter theory that the behaviour of a linear filter is determined entirely by its impulse response, *ie* its output waveform when fed with a sharp impulse. All analogue filters (as well as those digital filters that use digital feedback in their realisation – known technically as 'recursive' filters) not only affect the initial shape of the impulse waveform but also affect the way the impulse response eventually dies away to nothing. It is the nature of this decaying part of the impulse that has to do with Q.

A low Q filter invariably has the final part of the decay of its impulse response die away smoothly without any oscillation or change of signal polarity. If, however, the Q exceeds 0.5, then the final part of the decay oscillates about zero (Fig 1). It appears that the ears are sensitive to such oscillations in the decay part of the impulse response, even at very low levels. Such oscillations are often termed resonances, the frequency of the resonance being the frequency of this final decaying oscillation.

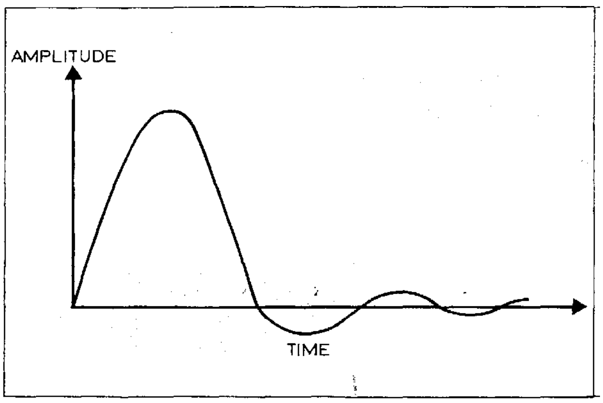

Fig 1: Impulse response of filter with Q greater than 0.5

Now Q can be thought of as a measurement of how rapidly the amplitude dies away per cycle of oscillation. The bigger the Q, the smaller the decay (in dB) per cycle, and the more cycles are gone through before a given degree of decay (in dB) occurs.

The effects of this oscillatory decay are audible even if the early part of the impulse response of a high-Q filter is such that it has an absolutely flat 1 amplitude response. It is perfectly possible to design such high-Q 'all-pass' filters (Fig 2 gives a typical circuit) and generally they sound much more coloured than low-Q all-pass filters, even if the latter are designed to cause many hundreds of degrees of phase shift.

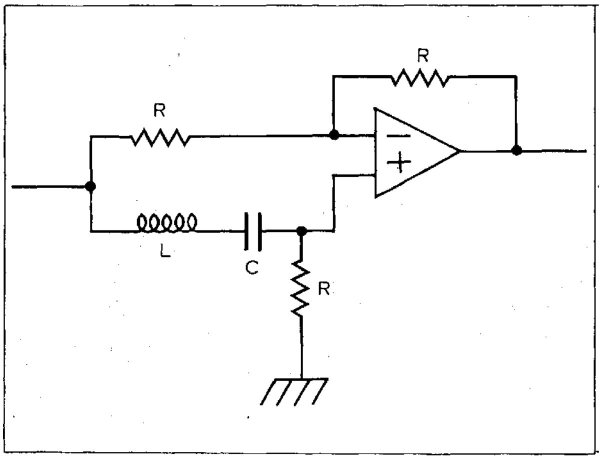

Fig 2: Circuit of high Q all-pass filter with flat amplitude response

It is significant that, in order to achieve a flat frequency response with a rapid crossover between speaker units, most multiway speakers are actually designed to have an all-pass response and if a high Q all-pass is chosen in order to make crossover rates more rapid, such speakers are likely to sound coloured[2].

The magic audibility threshold Q=0.6, rather than the strict no-resonance figure Q=0.5, suggests that the ears will actually tolerate some oscillation in the decay but the amount involved is surprisingly small. For a Q of 0.6, successive cycles of the oscillation are attenuated by about 80dB compared to the previous cycle. This means that, for a Q of 0.6, alternate polarity swings are about 40dB below the previous swing of opposite polarity.

The Q=0.6 threshold was theoretically suggested on quite different grounds derived from the behaviour of conventional 12dB/octave lowpass filters. The 'maximally flat' such filter, *ie* the one with the flattest amplitude response in the pass band, is known as a Butterworth filter and has a Q of 0.71. In much audio work requiring an uncoloured sound, preference has been for use of a lowpass filter with a maximally flat phase response, which is termed a Bessel filter, which has a Q of 0.58. Bessel filters have a slower high frequency roll-off but their 'smoothest possible' phase response has been found to give a subjectively superior sound. The empirical threshold Q of 0.6 for low colouration is very close to the Q of the Bessel filter having maximally flat phase response.

Anecdotal evidence of the importance of Q arises from experiments conducted by Philips in the early '70s with the then-new Dolby A noise reduction system. Dolby A is a multiband system using Butterworth filters to separate the frequency bands. It was noted by many engineers that Dolby A gave some subjective colouration, so Philips' engineers tried replacing the Butterworth filters with Bessel filters. They indeed found that such 'Bessel Dolby A' had a much lower audible colouration than standard 'Butterworth Dolby A'. The only problem was that it was incompatible with the already-standard Butterworth Dolby A system, so it proved to be impractical to introduce the Bessel version into studio use.

Now Dolby A is a reciprocal system, *ie* one whose decoding nominally exactly undoes its encoding, so that any audible effect of the filters was evident only in the small residual decoding errors due to imperfections in the tape path. Yet, despite the small magnitude of these errors, the difference between the Butterworth (Q=0.71) and Bessel (Q=0.58) systems was still easily audible. It is notable that later noise reduction systems introduced by Dolby Labs tended to avoid high Q filters in the critical mid-frequency bands of the audio range.

Work in connection with surround-sound encoding and decoding systems at Philips in the mid 1970s confirming the finding that Q must not exceed 0.6 to avoid audible colouration could be extended from lowpass filters to allpass filters, *ie* the effect was not amplitude response.

Not that easy

OK, so Q is the reason for the 'sound' of equalisers? Low Q is uncoloured, high Q over 0.6 is coloured? If only life were that simple. This rule of thumb does seem to work reasonably well over quite a wide range of analogue filter designs but it is far from infallible.

The demonstration that Q is not the crucial factor behind audible colouration comes from the highly coloured digital filter discovered by Lagadec and others at Studer in the mid '80s. The coloured filter had a flat phase response and the amplitude response consisted of 512 ripples of ±0.1dB uniformly spaced across the whole audio band (*ie* at about every 50Hz). This gives an impulse response as shown in Fig 3, *ie* the main impulse is surrounded by just two smaller impulses each about 46dB down, one preceding the main impulse by 20ms and one following it by the same amount.

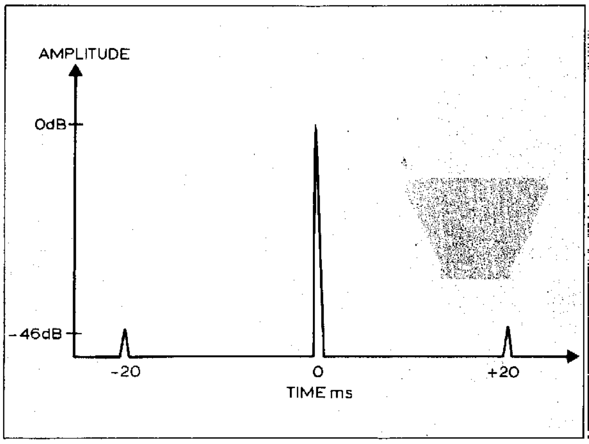

Fig 3: The impulse response of the coloured filter discovered by Roger Lagadec

Now this filter has no decay whatsoever, *ie* its Q equals 0. It only has a single discrete pre-response at low level and a single low-level post-response. Yet its colouration is highly audible. Detailed investigations showed that the main cause of the subjective colouration was the pre-response of the filter (*ie* the part before the main impulse) and that the pre-response had to be held below –80dB to avoid becoming obviously audible. Minus 80dB is merely one part in 100 million of the total signal energy and in the past, many 'objectivists' would have howled with derision at the thought that such tiny residues of error could possibly be of any audible importance. Much of what we now know about audible colouration by filters and equalisers is consistent with the conjecture (informed guess!) that what matters is not amplitude or phase response, but the *low-level* behaviour of the impulse response well away from the main transient. This is certainly not to say that amplitude or phase response, or the high-level behaviour of the impulse response, are unimportant or have no effect but that their main effect is often a relatively benign change of tonal quality, to which the ear can easily adapt, rather than obvious colouration that remains obvious even after time for adaptation is allowed.

The conjecture just made is probably not wholly true, *eg* a 12dB/octave treble boost will sound pretty ghastly despite having a good decay behaviour but for moderate and relatively smooth changes in amplitude and phase response, this hypothesis is at least a reasonable starting point for explaining why some filters sound more coloured than others.

Transient effects

Although we have conjectured that, on the basis of available evidence, low-level effects well away from the main impulse may be largely responsible for colouration, we have not yet specified precisely what kind of low-level effects are important. After all, we have already noted that a smooth non-oscillatory decay is generally relatively harmless.

At this point we enter the realms of conjecture and hypothesis in a big way. The following ideas are suggested as useful to equipment designers and others in getting good results or avoiding bad ones. These ideas are not pure guesswork – they are constrained by a lot of existing psychoacoustic knowledge and know-how – but neither are they gospel truth. No doubt, with time and experience, these ideas will be refined and exceptions identified.

Everything that is known about the way the ears perceive transients suggests that, all other things being equal, a pre-response (*ie* before the main impulse) in a filter will have more audible effect than a similar mirror-image post-response after the main impulse. This is not just consistent with Lagadec's findings on his digital filter, but is also consistent with the Haas Effect, whereby transient sounds tend to be preferentially localised by the transient arriving at the ear first, with later transients (up to about 40ms later, when separate echoes are heard) playing a reduced role. This is also consistent with the physiological effect of forward inhibition or temporal masking, whereby the perception of stimuli tends to suppress or reduce the sensitivity to the perception of stimuli following immediately afterwards.

This is not to say that, in some circumstances, later stimuli cannot also alter the perception of those immediately preceding them. Such backwards inhibition effects are well documented in the experimental psychology literature but generally, conventional forward inhibition is a stronger effect.

From another point of view, it is not implausible that the ears notice pre-responses much more strongly than post-responses, since pre-responses (*ie* effects before the cause) are rare in nature. This is not to say that they can't happen. The classic example is sound being picked up from a distant performer by a microphone on a stand on a non-rigid floor. Sound travels much faster through solids than through air, so sounds travelling through the floor and up the microphone stand to the microphone arrive before the main sound arriving through the air.

Generally, only bass frequencies arrive via the floor transmission routes but the characteristic bass pre-response is audible and, once recognised, can be heard as a characteristic colouration.

So one perhaps unexpected moral of our discussion is the need to take precautions to minimise the transmission of sounds to the microphone via the floor (or ceiling or walls, etc). This can be done by suspending microphones in a shock mount or via appropriate cables, by decoupling microphone stands from the floor by suitable compliant damped floor coverings and by using microphone types (notably some omnis) that are relatively insensitive to vibrations transmitted to their bodies.

The particular undesirability of pre-responses is especially relevant to digital filters and equalisers. Although it is not absolutely impossible to design analogue filters that have pre-responses, it is jolly hard. The filters have to be non-minimum-phase, and to have substantial pre-responses must be very complicated. Such complication is much easier to achieve in the digital domain, where memory (and hence pre delay) is cheap and plentiful.

The classic example of pre-responses in digital filters is something many people (including a previous editor of this magazine) had claimed is always desirable – namely filters having a linear phase response. A filter with a linear phase response suffers from no phase shift at all (other than an overall constant time delay, which we can ignore). This seemed like Nirvana compared with the awful phase shifts suffered by the analogue minimum-phase brickwall filters widely used with non-oversampling A/D and D/A converters.

A little thought, however, shows that linear-phase filters might not always be as desirable as they might seem. Linear-phase filters have an impulse response that is time-symmetric: their pre-response is the mirror-image of their post-response (Fig 4). The reason for this is that they have, by definition, no phase shifts and so behave in exactly the same way whether one looks into the future time direction or the past time direction.

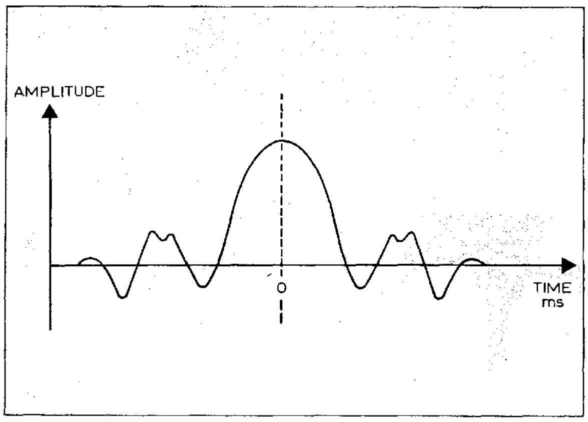

Fig 4: Typical impulse response of a phase-linear filter. Note the mirror-symmetry around time zero

But, being unnatural and of greater audible effect than post-responses, such pre-responses could well have substantial audible side-effects that would be heard as audible colouration – as Lagadec found with his filter.

Note that I am not claiming that extended pre-responses automatically give a coloured sound, only that the risk of such colouration is higher than for similar post-responses. It is ultimately a question of trying out a given filter response empirically and listening for colouration.

To take an extreme example, I would expect that a time-reversed high.Q filter (one whose impulse response is the time-reverse of that of a high-Q filter) to be highly objectionable and a phase-linearised version of a minimum-phase high-Q filter to be only a little less objectionable, due to such filters having extended 'pre-ringing' in their impulse response. In particular, one would expect digital phase-linear graphic and parametric equalisers to have worse audible colouration than a well-designed analogue or minimum-phase digital equaliser designed more conventionally. Experience of phase-linear digital equaliser products tends to bear out this increased audibility of colouration.

On the other hand, a well-designed pre-response with no sudden sharp changes in level (even at very low levels) or oscillations, and with smooth gradual increase, is expected to be much less objectionable subjectively, as is a pre-response whose ripples and oscillations lie well outside the audible frequency range. I would expect, for example, carefully-designed phase-linear highpass filters for cutting out low frequency rumble noises or for bass speaker equalisation to sound considerably better than current minimum phase highpass filters. The latter's severe phase distortion produces very audible tonal and dynamic colouration even when present in small amounts and of low Q. The key words in the above are 'carefully designed', avoiding a badly shaped pre-response.

In general, it would be expected that digital filters somewhere between the minimum-phase behaviour of analogue filter designs (*ie* with the minimum phase shift consistent with the actual amplitude response and with no pre-response) and linear phase might sound better than either. Designing a carefully-tailored pre-response to minimise audible colouration is, as yet, uncharted territory but once they get their teeth into it, I would expect designers of digital equalisers to start coming up with some subjectively interesting products.

I suspect that, as in other areas of audio, there will be no unique 'best' phase response for a given frequency response but, rather, different choices of trade-offs among different subjective virtues and defects. We can look forward to countless future arguments about which of many competing approaches is really best.

Post-responses

So much for the *terra incognita* of digital filters with pre-responses. Conventional analogue filters with only a post-response can still have substantial audible colouration effects, as was recognised by Harwood in the 1950s. How can we find out what kind of post-responses sound 'nice' and what do not? Although Q can be a useful guide, it does not tell us everything we need to know. It is possible to design filters with a high Q that don't sound too bad, and equally possible to design low-Q filters that are pretty awful or even downright unlistenable.

As everywhere else in audio, there is probably no single magic number that guarantees the goodness or badness of a particular filter response. The key to all this probably lies in learning to understand and analyse filters in a large number of ways, *ie* not placing all one's eggs in one basket.

Traditional frequency responses do tell quite a lot, although they certainly don't tell whether or not an equaliser is guaranteed to sound good. The presence of a broad band of emphasis over a range of frequencies will often give a general indication of the balance of the sound, although not its subjective quality or frequencies of audible colouration. However, areas of raised response, *eg* peaks, tend to 'stick out' as audible colouration, although narrowband dips can often be virtually inaudible.

In general, frequency responses give one little information about the audible 'smearing' of transients and phase responses (or the closely related measurement of group delay) may not be that much more helpful. Sudden peaks or dips in the phase response or group delay can be a symptom of audible transient smearing or colouration but on its own, phase or group delay response contains too little information to allow the resulting sound to be reliably predicted. We have seen that it is also necessary to look at the impulse or time response of filters and that very tiny discontinuities or oscillations in the time response can be audible – especially if they are spaced from the main impulse response by a few milliseconds (before or after). Thus it seems advisable to examine the tails of the impulse response 'under a magnifying glass' if looking for possible symptoms of colouration. It might even be useful to look at the impulse response processed by being fed through a compressor with a very fast time constant, in order to bring up low level artefacts to make them visible.

Neither the response in the frequency nor the time domains alone are adequate. The eye (used to assess measured data) is no good at picking out the frequencies at which trouble is occurring from examination of the impulse response.

Simultaneous time/frequency analysis

For this reason, over the years attempts have been made to analyse sounds and linear system characteristics simultaneously in both frequency and time, plotting the result as a graph over the two variables time and frequency. The most familiar example of this is the speech spectrogram, which plots the level of sounds at each frequency as a function of time.

All such attempts are compromises. After all, a frequency response analyses the response to sinewaves, and a sinewave (by definition) lasts forever. Any attempt to resolve behaviour in time reduces the ability to resolve frequency and *vice-versa.* If one blurs one's resolution in time to a minimum time interval Δt, and one's resolution in frequency to a minimum frequency range Δf, then a famous mathematical result (first used in Quantum Theory) known as the uncertainty principle asserts that one has to have

$$\Delta t.\Delta f \geq \tfrac{1}{2}$$

(This is the first and last mathematical formula in this article!), where Δt is the minimum time resolution in ms and Δf is the minimum frequency resolution in kHz. So if one has a time resolution of 5ms, then the frequency resolution can be no better than 100Hz – no good for examining details of the bass response.

Actually, as noted by Dennis Gabor (best known for his invention of holography, but who also worked in audio) back in 1946, the ears actually analyse the frequency content of sounds in time faster than suggested by the uncertainty principle by a factor of about 7. The seeming logical contradiction with the fundamental theoretical limit of time/frequency resolution is avoided by the ear's use of *a priori* or previously assumed knowledge of the nature of typical sounds but at the expense of getting the analysis 'wrong' when sounds not of the assumed form occur.

No one has yet succeeded in devising a method of simultaneous time/frequency analysis that beats the uncertainty principle limits on resolution using *a priori* information similar to the ear. Existing methods of analysis do not resolve enough detail in the two domains simultaneously to predict reliably how a filter will sound.

Nevertheless, several of the existing methods of time/frequency analysis do reveal some of the things that cause colouration: for example, both the techniques of *Time Delay Spectrometry* (TDS) invented by Richard Heyser, and earlier techniques of measuring frequency response after cutting off a first part of the impulse response reveal low-level delayed resonances. With these techniques, the initial frequency response may measure flat but the frequency responses associated with later times display a visible decaying resonant peak. However, while these methods have enough resolution to measure the grosser faults of loudspeakers, they still tend to mask the more subtle faults associated with many equalisers and also systems such as turntables.

There is an urgent need to refine existing methods of simultaneous time/frequency analysis to maximise the amount of fine low-level detail that can be seen. In computer-based analysis packages, this means carefully devising the filtering and 'windowing' used to minimise all discontinuity, resonance and aliasing artefacts, and using very high quality graphics to display the results on a very fine time/frequency grid. Otherwise the eye will not be able to resolve the required detail.

Even when this is done, analysis using several different trade-offs of time and frequency resolution will probably be needed, so details that occur predominantly in the time domain and in the frequency domain can both be examined.

Wigner distribution

A mathematically beautiful and elegant method of simultaneous time/frequency analysis was published by Eugene Wigner in 1932 (his application was to Quantum Statistical Mechanics, although its original application was apparently in another unspecified field). Despite its mathematical elegance, this Wigner Distribution has a lot of unwanted 'high frequency clutter' obscuring the wanted detail from the eye – for example, if two frequency components are present, the Wigner distribution also displays a spurious beat-frequency component at the average of the two frequencies. Despite its use in recent audio literature, where the Wigner distribution response of a number of filters and loudspeakers has been published, and despite the fact that in principle it contains all the information needed to understand a filter response, the large amount of clutter present makes it impossible for the eye to make out relevant details.

Nevertheless, the Wigner distribution may well form the basis of future improved methods of time/frequency analysis beating the uncertainty principle limit. This is because it can be shown mathematically (the methods of proof are buried deep in the Quantum Theory literature) that the normal methods of time/frequency analysis can be obtained from the Wigner distribution simply by blurring it with a suitable smoothing filter. (Technical note: this 2-dimensional smoothing filter has a response that is also a Wigner distribution.) Such blurring removes the unwanted clutter, at the expense of also blurring the wanted information. However, by using less drastic blurring than used to obtain conventional time/frequency analysis, much of the clutter can be removed without losing so much detail.

So, by time/frequency analysis using a carefully-smoothed version of the Wigner distribution, in future we may have the tools to see what filters and equalisers are doing in the time and frequency domains with more detail than was previously possible. Designers of the software packages for audio analysers need both to master the relevant mathematical tools and to design the required smoothing filters in the software to avoid the kind of discontinuity or oscillation behaviour we are looking for in the hardware audio filters and equalisers we want to analyse. In other words, the design of analyser software requires the same kind of skills required to design good audio equalisers.

The future

What an optimistic subheading! Actually, the future understanding of equalisers is still uncertain. What we do now know is that many low-level effects often ignored in the past are very important subjectively and that traditional methods of measurement and analysis are not yet refined enough to reveal their effects. One priority is to refine our methods of measurement and analysis to maximise the visibility of low level effects.

This means as much skill is required in the development of test equipment as has traditionally been applied to audio equipment. Nothing can be taken for granted. In particular, the filtering and 'windowing' on spectral analysis equipment needs to be much better behaved than has been the case until now. Much more attention is also required to the quality of the display of visual information, which should avoid the kind of steps, kinks and coarse grids of current displays, since one is actually looking for such discontinuities in an equaliser response as symptoms of its audible quality.

Meanwhile, the design of equalisers will remain an art, although I hope the questions raised here will help to inform the art and concentrate attention on potentially important factors – particularly in the design of digital equalisers.

One topic not covered is the role of circuit nonlinearities in the sound of analogue equalisers, or of 'rounding error' and requantisation effects in the design of digital equalisers. These are also important but would require several articles to themselves. In the above, I have assumed that the equalisers have been designed carefully enough to minimise such nonlinear effects but, sadly, this is often not the case in commercial products – particularly for digital equalisers.

A gloomy ending

One area of pessimism concerns the viability of using equalisers to compensate for defects in other equipment (microphones, loudspeakers and even multiple stages of bass roll-off in audio electronics). The problem here is that even very tiny residual errors in the frequency and phase responses may turn out to be almost as audible (or in some cases even more so) than the original unequalised errors. Equalisation may improve the tonal accuracy in such cases but it can (and often does) increase the audible colouration.

If this is right, we may be unable, ever, to 'fix it in the mix' properly, and this re-emphasises the importance of using the best and least-coloured sounding audio equipment at every stage of the audio recording chain. The best equaliser is no equaliser! Anything else may add useful creative pizzazz – and it is worth understanding what such creative equalisers are doing – but there are limitations to how far an equaliser can actually 'equalise' an already-coloured signal.

References

1) Philip Newell, 'Monitor Systems Part One: A Look at the Overlooked', *Studio Sound,* pages 46-53, August 1989

2) Philip Newell, 'Monitor Systems Part Five: Crossovers', *Studio Sound,* pages 60-67 (December 1989) and pages 48-53 (January 1990)

[KH note:

The Roger Lagadec and Dennis Gabor articles mentioned in the text but not referenced are

R Lagadec and T G Stockham, 'Dispersive Models for A-to-D and D-to-A Conversion Systems', Preprint 2097, 75th Audio Engineering Society Convention, 1984

D Gabor, 'Theory of Communication', Journal of the Institution of Electrical Engineers, 93, III, p429, November 1946

For reasons I can't explain, Gerzon's description of the impulse response of the Lagadec coloured digital filter doesn't correspond exactly with that described in the Lagadec and Stockham paper, which says, "In the case of the model above, which related well to a realistic (but, in retrospect, ill-designed) filter bank, the amplitude of the pre- and post-echoes is approximately 0.025, or, in other words, 32dB below the main pulse. The distance between echoes and main pulse is 2048 sampling intervals, or 40 milliseconds (0.04 seconds), and a pre-echo at –32dB preceding the 'main' signal by 40 milliseconds is of course quite perceptible, even with untrained listeners."]