much more effective, particularly at the lowest standing wave.

Figure 34 shows some real measurements from a prototype of Reference 5 with no wadding and with optimised wadding in the enclosure. These measurements are taken with the microphone very close to one of the low frequency drivers to clearly show the standing wave effects, the dip at 35Hz is due to the port. Note that even before the wadding is added to the enclosure the behaviour is very good, this is due to the carefully designed enclosure and driver locations. Once the wadding has been added the overall response is smoother, the dip at 400Hz (corresponding to the first standing wave) is suppressed and there is little change to the low frequency response save for a slight shift in the port tuning frequency.

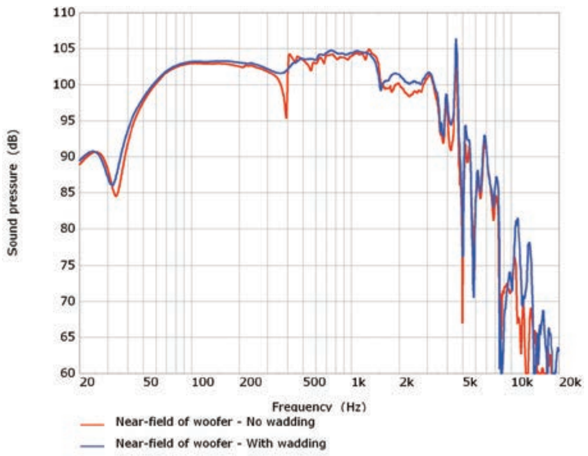

Figure 34. Nearfield measurements of an LF driver from a prototype of Reference 5 showing the response with and without wadding added to the enclosure.

4.3.2 Cabinet vibration control

To reproduce the most realistic listening experience, the sound needs to come purely from the drivers and not from the cabinet. Any resonant vibration from the cabinet panels will add unwanted distortion to the music. This means that the cabinet walls of a loudspeaker ought to be as inert as possible. The construction of The Reference cabinets has always been a strength of the range. The new models maintain this tradition and are constructed in thick high density wood the extensive internal bracing. Conventional internal bracing aims precisely to stiffen and support the inside structure as well as each panel. Figure 35 illustrates the effect of adding a pair of braces, crossing behind the driver, compared to a box without any bracing. The braced box has raised the resonance from 600Hz to 1.6kHz but the amplitude of the peak still remains the same. Adding stiffness to the cabinet will always have this effect, to reduce the severity of the resonance it is necessary to add damping.

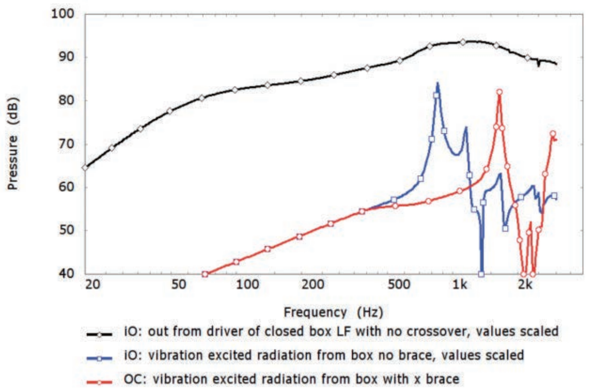

Figure 35. Closed box FEA predicted output from diaphragm and walls with and without x-brace.

During the development of the KEF LS50, it was found that by adding material with high mechanical resistance and low stiffness between the walls, baffle, driver and brace a dramatic amount of damping could be added to the cabinet. Figure 36 shows the same loudspeaker previously shown in Figure 35 using this damped bracing configuration. The peak at 1.6kHz is dramatically reduced in amplitude by about 30dB. The Reference uses this technology throughout the range in order to minimise the cabinet vibration.

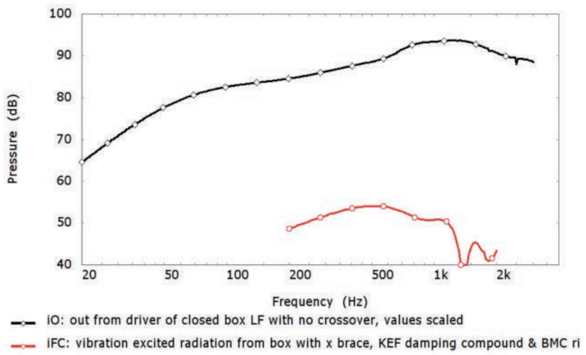

Figure 36. Closed box FEA predicted output from diaphragm and walls with constrained layer of damping material between brace and walls.

Figure 37 shows a cross section of Reference 5, the cabinet walls are supported by extensive internal bracing. These braces are connected to the cabinet walls using a layer of damping material. The drivers are also braced by the internal structure and connected using a layer of damping material to further control the cabinet vibration. The front baffle of the cabinet is constructed in an exceptionally strong laminated aluminium and resin composite, this adds a great deal of strength and mass to the cabinets. The baffle also forms part of the damping arrangement and is connected resiliently to the rest of the cabinet using high loss pads.

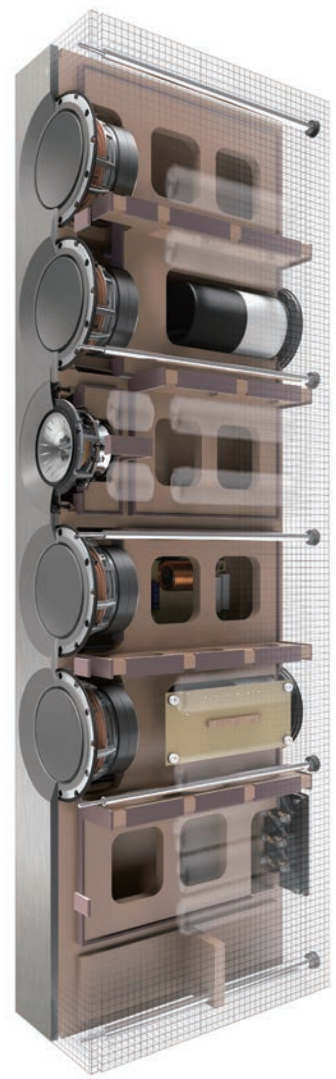

Figure 37. CAD section of Reference 5 showing extensive internal damped bracing.

4.4 Driver details

4.4.1 Uni-Q Driver

The Reference uses a specially developed 11th generation Uni-Q driver array. This particular Uni-Q is only used in The Reference and incorporates many features found in that of the flagship KEF Blade loudspeaker. Figure 38 shows a CAD rendering of the new driver, with the midrange motor shown in isolation on the right. In particular, note how open the rear of the driver is, this is done intentionally in order to maximise the rear venting of the midrange. On the right hand image the large aluminium ring on the top of the motor system is visible, another identical ring is buried in the magnet system to help to control the midrange distortion. The design of this driver is outlined in more detail in the next few paragraphs.

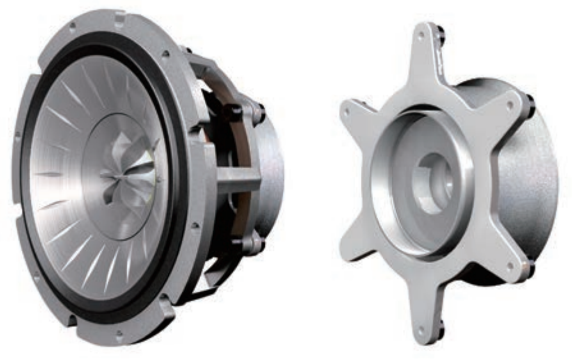

Figure 38. CAD rendering of The Reference Uni-Q, right hand image shows midrange motor in isolation.

The Tweeter

The tweeter in the new Uni-Q is closely based on that of the KEF Blade. It uses a powerful neodymium motor system with a ring magnet and a copper cap to reduce distortion. The rear acoustic design of the tweeter is critical for low distortion performance. The new Reference tweeter uses a large central vent which carries the rear acoustic radiation gently away from the back of the dome. More information can be found in Appendix V.

The dome itself is constructed from aluminium and uses KEF's patented stiffened dome design. This is a unique technology which enables the 25mm dome to extend into the ultrasonic bandwidth, more information can be found in Appendix IV.

The acoustic design around the tweeter dome is absolutely critical to the performance of the 11th generation Uni-Q design. It has taken most of KEF's 20 plus years of experience in designing the Uni-Q to fully understand and fully optimise this area.

The first stage in getting the optimal performance from the tweeter is to match the shape of the tweeter dome to the surrounding waveguide. If these two parts are not matched then large response irregularities can appear in the tweeter output. More information is given in Appendix III. Once the waveguide and dome geometries are matched, then the waveguide can be fully utilised to control the tweeter dispersion and to match this to the midrange driver. When combined with the optimal dome shape technology, the overall tweeter response is - unlike a conventional baffle-mounted tweeter - completely free from off-axis nulls. This difference can be clearly seen in Figure 39. This optimal geometry is patented, and hence only found in KEF Uni-Q loudspeakers.

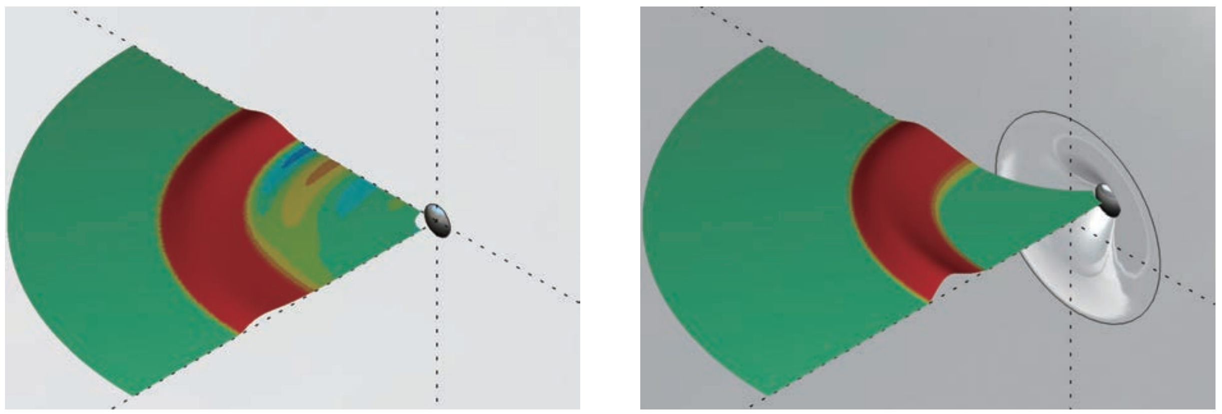

Figure 39. Comparison of a HF transient radiated by a conventional 25mm dome on baffle (left) and a optimally shaped 25mm dome in a waveguide (right).

The tangerine waveguide is another key piece of technology that is used by KEF to optimise the high frequency performance of the Uni-Q driver. A dome is close to the ideal shape for a loudspeaker diaphragm, particularly when placed in an optimal dome and waveguide configuration. However, a dome tweeter does not quite have the correct surface velocity to be the perfect radiator – when a signal is played through the tweeter the dome moves in one axis. Because the angle of the dome surface relative to this motion is greater towards the edge of the dome, this means that the surface-normal velocity is lower towards the dome perimeter. Ideally the surface-normal velocity would be constant over the entire dome surface. This difference is illustrated in figure 40.

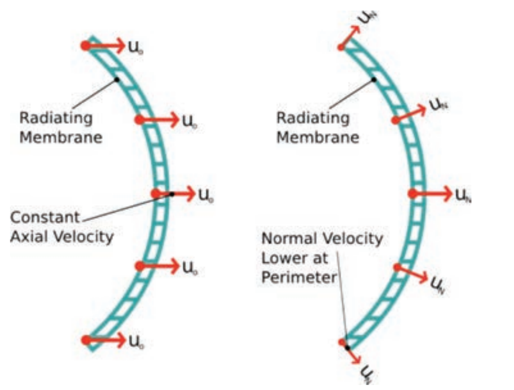

Figure 40. Comparison of tweeter dome axial and surface normal velocity.

The purpose of the tangerine waveguide is to correct for this difference between the real normal velocity and ideal normal velocity by directing the output from the dome into a small chamber under the tangerine waveguide and then to control the expansion of the sound into the waveguide through specially shaped channels. The design of the channels is a very complex process and required extensive use of FEA computer models in addition to involved mathematical analysis. A still from one of the computer models is shown in Figure 42 along with the final production part.

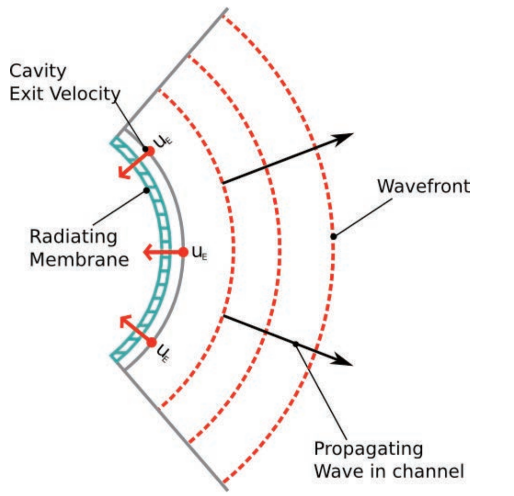

Figure 41. Illustration of the target behaviour of the tangerine waveguide.

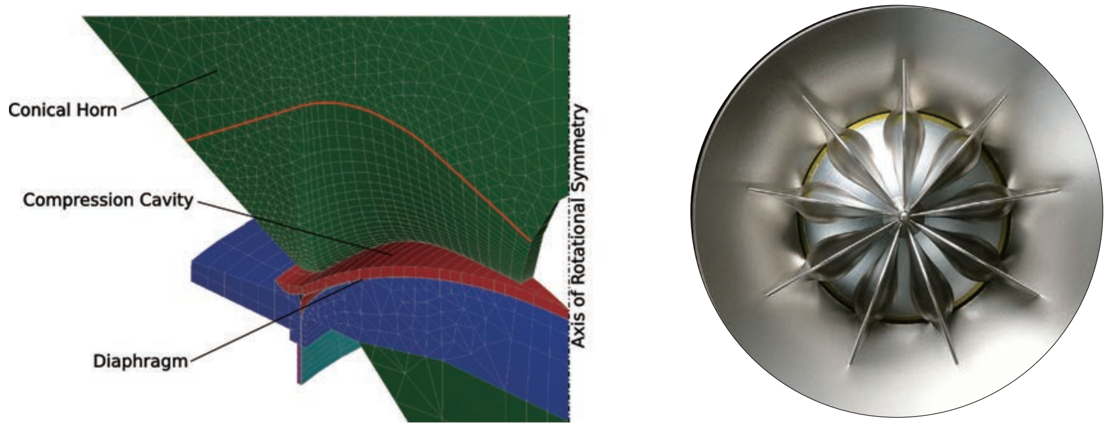

Figure 42. FEA computer model of the tangerine waveguide (left), final production waveguide (right).

Figure 43 shows the frequency response of a prototype of The Reference tweeter measured in a test waveguide. Also shown in this figure is a computer model of the same device. Note that the two responses are in very good agreement. The tweeter response is exceptionally smooth without any trace of resonance or interference. The sensitivity of the tweeter is greater at the lower end of the device due to the waveguide which actually assists in the acoustic coupling of the tweeter dome. In the 2kHz region, which is a critical area in most music material, the sensitivity is much higher than a conventional tweeter. In the crossover this downward tilt is corrected but the beneficial result is that the tweeter runs "cool" as less power is delivered to the tweeter at the lower end. For example at 2kHz almost one quarter of the power is fed to the tweeter, this is a huge advantage in terms of the maximum output and the linearity of the treble.

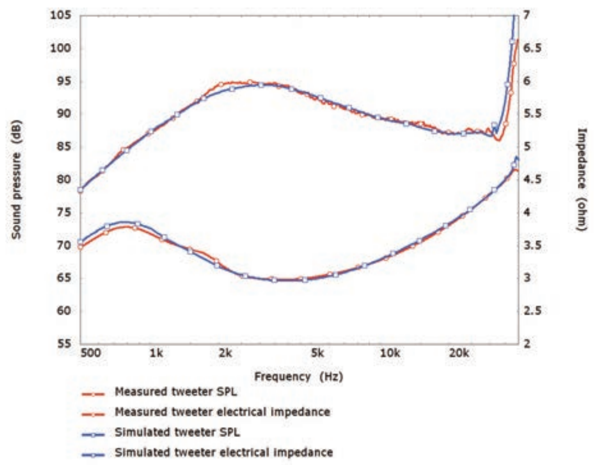

Figure 43. Frequency response of a prototype of The Reference tweeter measured in a test waveguide

In the full loudspeakers, the tweeter "waveguide" incorporates several physical parts and it is absolutely critical that these are designed to disturb the tweeter output as little as possible. The tangerine waveguide, the midrange cone, the midrange surround, the shadow flare and the baffle all form a single continuous waveguide surface. At each junction between two parts on this surface the gap is carefully controlled to minimise any effect and the parts meet exactly tangentially. The midrange surround uses a Z-Flex design so that it too forms this surface as smoothly as possible (see Appendix VI for more details).

The Midrange

The midrange driver has a difficult task: it must bridge the gap between the low and high frequencies, covering the range over which our ears are most sensitive. It is therefore exceptionally important that a midrange driver is well behaved above the working range, otherwise they might conflict with the tweeter output. The shape of the midrange driver cone, surround and surrounding objects are very important. In addition, the midrange driver cone experiences much higher vibrational forces than a bass driver cone because if the range over which it operates.

The Uni-Q uses a 5inch aluminium cone driver for the midrange output. The size of this driver is carefully optimised so that the dispersion of the tweeter and the midrange matches as well as possible at crossover. Aluminium is a very good choice of material for a loudspeaker diaphragm as it is both stiff and light and also easy to form into a complex shape. The stiffness and low mass mean that the cone operates as a rigidly body over the entire midrange region. This is unlike some competitor drivers where the cone operates in resonance over much of the working bandwidth. The benefit of a rigid body operation is that the sound radiated from the cone is free from irregularities due to differences in how the cone vibrates at different frequencies, providing the best possible coherence to the radiated sound. The rigid operation is also very important in controlling the dispersion of the driver. The target is for the dispersion of the midrange driver to slowly narrow monotonically as frequency increases. When a cone enters breakup the dispersion characteristic changes and typically becomes wider, this sudden change in dispersion is not desirable.

The issue with metal diaphragms is that when they do enter breakup, as they have little internal damping, very large irregularities in the response occur. These can be easily 15dB or more in magnitude. This is large enough to be a problem even if breakup occurs well above the crossover frequency. It is possible to add damping material directly to the cone to control these resonances, however, this is not a good solution as this direct damping application is very heavy and this results in a driver with low sensitivity.

KEF use a unique technology called cone neck control to avoid the traditional breakup problems of metal cone drivers. With cone neck control the cone is not rigidly connected to the voice coil of the driver. A resilient high damping link is used to connect the two parts together. This link is carefully designed and fine tuned with the help of computer modelling so that within the band of the driver the force from the voice coil is fully transferred to the cone. Above crossover, however, the resilient link begins to flex and to damp the cone motion. The effect on the driver response is quite dramatic. The breakup peak from the driver is reduced by around 15dB and the driver response is considerably smoother. The penalty is a small mass increase in the moving parts, albeit much less than using a direct damping approach. Figure 45 shows the modelled frequency response of a midrange driver with and without cone neck control technology.

The overall frequency response of the midrange driver is shown in Figure 46. The frequency response is very smooth and well controlled to well above the crossover frequency of approximately 2.5kHz.

Figure 44. Cone neck control resilient link between the cone and the voice coil.

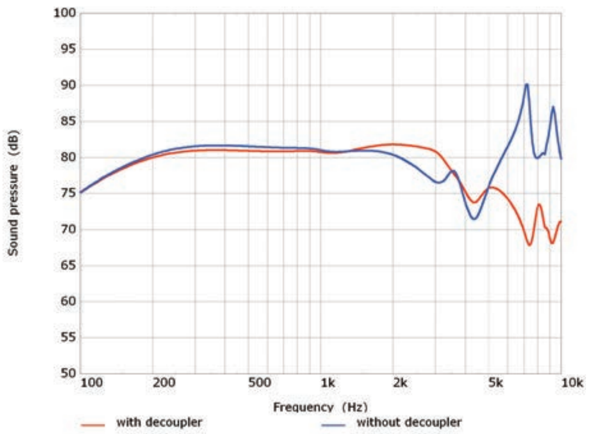

Figure 45. Midrange driver response with and without cone neck control technology.

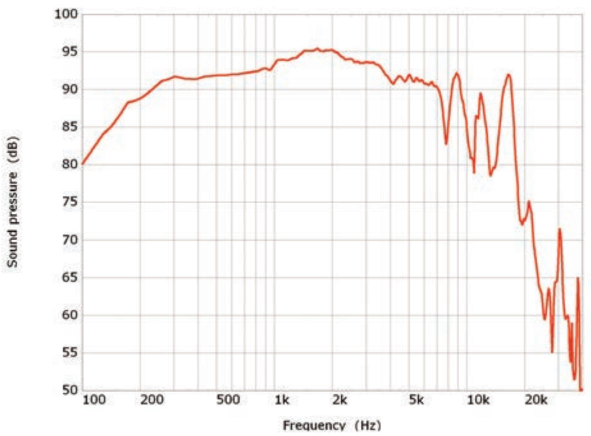

Figure 46. Axial frequency response of the midrange driver measured in 2pi at 1v/1m.

4.4.2 Low frequency drivers

The low frequency driver is shown in section in, right. Computer modelling was used extensively in the design of this driver in order to achieve the best possible performance. Rather than a conventional cone diaphragm, the driver uses a shallow aluminium disc. This gives the driver a much lower profile when placed in the loudspeaker baffle and ensures that it has the minimum effect on the response of the midrange and tweeter.

The rear of this disc is supported by a vented coupler. This is a unique technology to KEF that was first developed for the Blade loudspeaker. The vented coupler serves two purposes, firstly it allows very free movement of air away from the driver as the diaphragm moves. With a conventional driver, air can often be trapped inside the voice coil and this can lead to losses and distortion at high levels. At the centre of the motor system there is a large venting hole that further aids the movement of air at the rear of the driver. The geometry of the vented coupler is fine tuned using computer analysis so that it connects with the aluminium disc at a “nodal” position. A nodal position is a region on a structure where a particular resonance has no effect. By choosing a nodal position to drive the aluminium disc the first resonance of the disc is suppressed.

The voice coil on the bass drivers is particularly large compared to other drivers of a similar size. This is very important as at high levels a great deal of heat must be dissipated from the coil. A large voice coil has a fundamental advantage in this respect as for the same power input it will not become as hot, simply because it has a greater thermal mass. Secondly, the larger area aids the dissipation of this heat into the motor system and surrounding metal work. The overall result is much less power compression. The larger voice coil could present a problem in terms of mass, however to overcome this issue the coil uses aluminium wire.

The driver uses an undercut pole and an overhung voice coil design. Both of these features are to maximise the excursion capability of the driver for low distortion output even at high levels. The surround and suspension of the driver are also carefully designed with this in mind. In order to control the distortion in the upper bass and lower midrange region the motor system incorporates large aluminium Faraday rings above and below the magnetic gap.

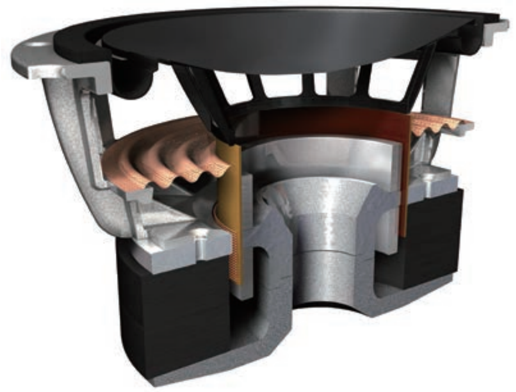

Bass driver cross section

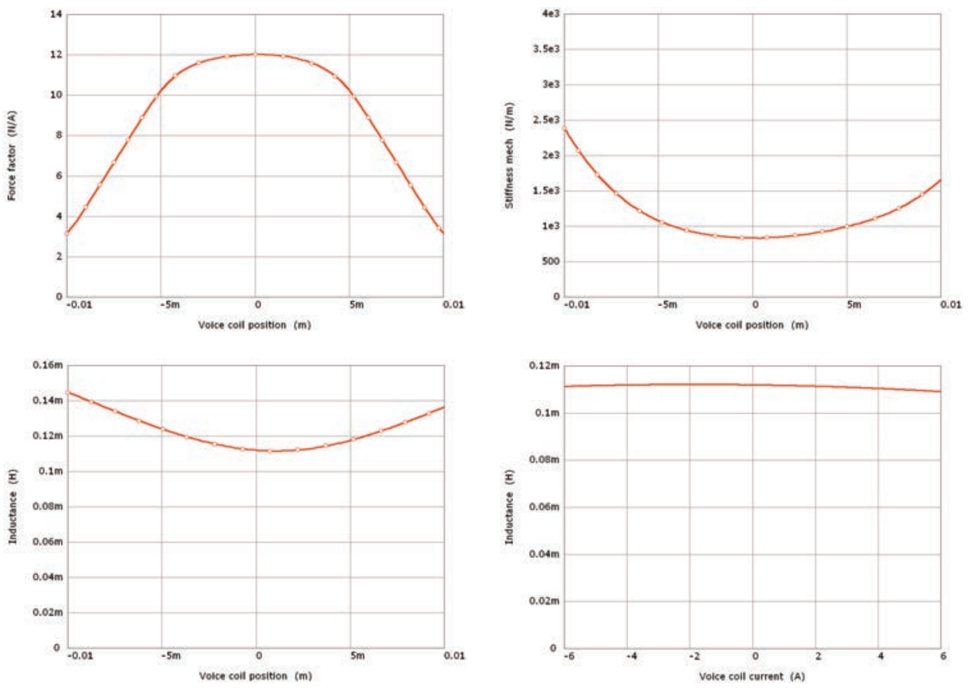

Figure 47. Non-linear LF driver parameters measured with Klippel analyser system, 16Ohm driver version.

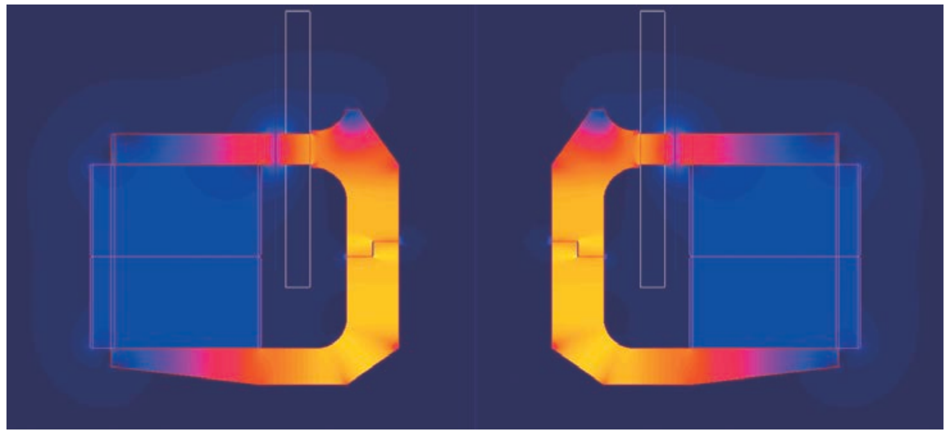

Figure 48. Static magnetic FEA simulation of LF driver motor with flux density shown as colour.

4.4.3 Motor system design

The Reference system and driver design has specifically focused on achieving the lowest possible distortion performance. Though there are several sources of non-linearity in loudspeakers, it is the driver motor system design that is most critical in achieving low distortion in the midband region where the diaphragm motion is small. At KEF, advanced FEA computer analysis methods are used to design the motor system and to analyse the distortion performance. For the Uni-Q driver, due to the close proximity of the mid and high frequency motors, the two magnetic systems must be designed to work together. Figure 49 shows a static magnetic FEA analysis of the MF and HF motor systems, the magnetic flux density due to the permanent neodymium magnets is shown. The field of each motor is concentrated in the magnetic gaps, there is little stray field. The undercut design on the midrange driver can be clearly seen, as can the large central venting hole used to accommodate the rear loading of the tweeter.

Reference motor systems use a mix of underhung (tweeter) and overhung (low frequency/midrange) voice coil designs to ensure that the the force from the motor system is not modulated with the motion of the driver diaphragm. In addition both motor systems incorporate carefully designed conductive regions. The tweeter has a copper sleeve on the inside of the magnetic gap, the low frequency and midrange use large aluminium rings above and below the gap. These conductive regions are inductively coupled to the voice coil and they are used to cancel the self-inductive behaviour of the voice coil. The placement and size of the rings are fine tuned using computer analysis to ensure that the voice coil inductance is not modulated with the voice coil position. In addition, the cancellation of the voice coil inductance itself is very effective in reducing distortion. This is because the magnetic moment generated by the voice coil will, if not controlled, modulate the permanent magnetic field. For example Figure 50 shows the magnetic flux density generated in the steel parts near the midrange motor gap due to the voice coil current, with the conductive rings in place the magnitude of this flux is significantly lower. Figure 51 shows the corresponding inductance for two voice coil positions.

The result of this attention to detail in the design of the drivers and motor systems in general is that the distortion performance of the loudspeakers is exceptional. Figure 52 shows the percentage harmonic distortion of a prototype of Reference 5 at a 90dB/1m output level. The mid-band distortion is less than 0.07%.

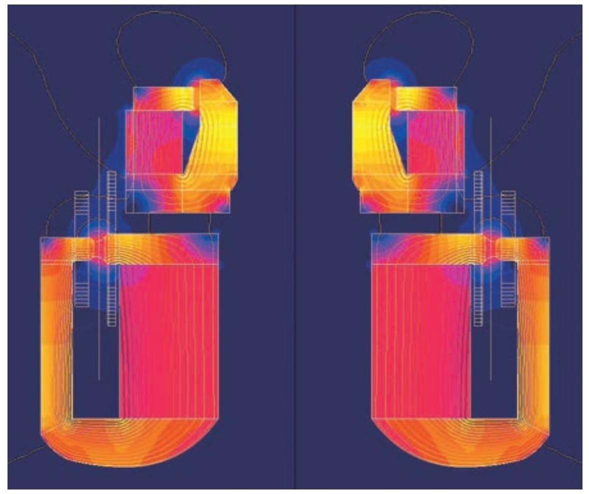

Figure 49. FEA magnetic analysis of Uni-Q motor showing flux density due to permanent magnets.

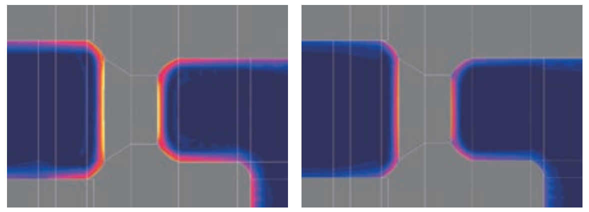

Figure 50. Flux density modulation of midrange motor system gap area at 2000Hz without conductive regions (left) and with conductive regions (right).

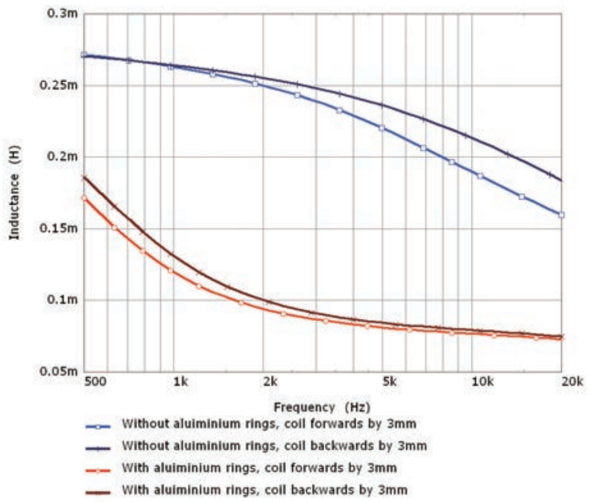

Figure 51. Midrange voice coil inductance versus frequency for motor system with and without aluminium rings.

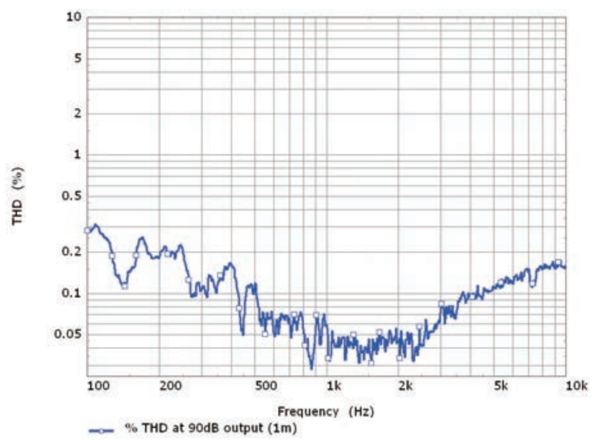

Figure 52. Measured THD of Reference 5 prototype as a percentage of the output SPL at a level of 90dB/1m.

4.5 Crossover design

The crossover is the electrical circuit that divides the input signal and sends appropriate parts to the relevant drivers in the loudspeaker system. The Reference loudspeakers are passive designs, this means that the electrical crossover circuit is not powered and the components must perform at the high power levels that are sent to the loudspeaker by the driving amplifier. Crossover design can be a complex subject and there are plenty of strong opinions of how it should be approached. The approach taken to the design of The Reference crossovers are outlined in the following sections.

4.5.1 Crossover component distortions

Passive crossovers are constructed from simple electronic components such as resistors, capacitors and inductors. Real resistors, capacitors and inductors are not ideal, and all introduce some level of distortion to the system. Selecting the correct types of each component can greatly affect the performance of the crossover, and choosing the wrong components can result in a crossover that introduces a significant amount of distortion into the loudspeaker system.

For The Reference, extensive objective testing of individual components were carried out to identify those with the lowest distortion. This study turned out some extremely surprising results. In many cases there was little correlation between the size, cost or published specification and the measured component distortion. For example, Figure 53 shows the distortion levels of three identical value inductors of a similar size, cost and DC resistance, but of different constructions. As can be seen, the worst inductor had 10 times more distortion than the best.

Figure 54 shows a comparison of the measured distortion for some different capacitors. Again, there were significant differences between different types. Using this testing allowed very quick identification of the best components for further investigation in listening tests.

This testing approach was also applied to fully constructed crossovers, for example Figure 55 shows the measured distortion of a prototype Reference 5 high frequency crossover circuit. Both curves are for the same circuit, but one has been built with a single poor component choice, whilst the other has been made using only low distortion components. This shows that a single bad component can introduce a large level of distortion, whilst correctly selecting the components can reduce the distortion generated by the crossover to a negligible level. By using this methodology through the whole crossover network, the result is a lower distortion and a better sounding loudspeaker system.

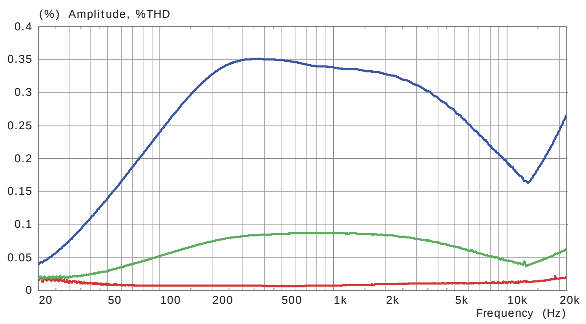

Figure 53. Measured component percentage THD at 20V input level for three inductors of the same inductance, similar cost and similar resistance but different constructions and manufacturers.

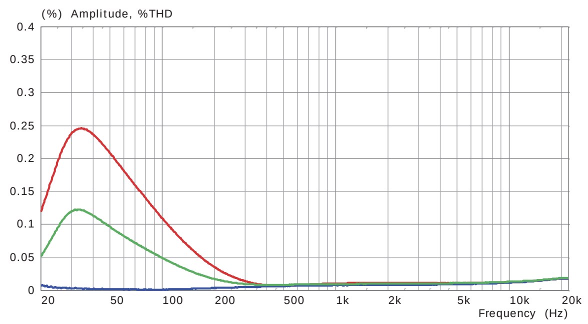

Figure 54. Measured component percentage THD at 20V input level for three capacitors of the same capacitance, but different constructions and manufacturers.

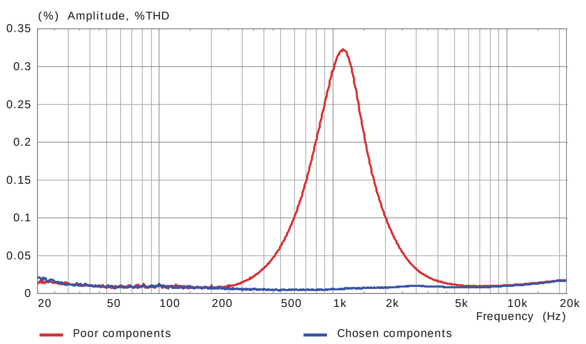

Figure 55. Distortion measurement of full HF crossover at 20V input level with two different component selections, note the component values and filter responses are identical.

4.5.2 Crossover filter order

Crossover filters have been widely studied and there are numerous publications on different approaches to their design. The most widely known are the classical Butterworth and Linkwitz-Riley types. Both of these families of filters can be implemented in either active electronics or with passive electronics[8]. These filters have the characteristic that when the output of a complimentary high-pass filter and low-pass filter are summed the resulting signal has a perfectly flat frequency response. In a loudspeaker the high-pass filter would feed the signal to the tweeter and the low pass filter would feed the signal to the woofer. The order of a filter determines how much attenuation is present in the filter stop band. The first four of these classical crossover designs can be seen in Figure 56, all filters result in a crossover frequency of exactly 1000Hz and result in a flat summed response. The effect of increasing filter order can be clearly seen - the 1st order high-pass filter provides 20dB of attenuation to the tweeter at 100Hz, whereas the 2nd order high-pass filter provides 40dB of attenuation to the tweeter at 100Hz.

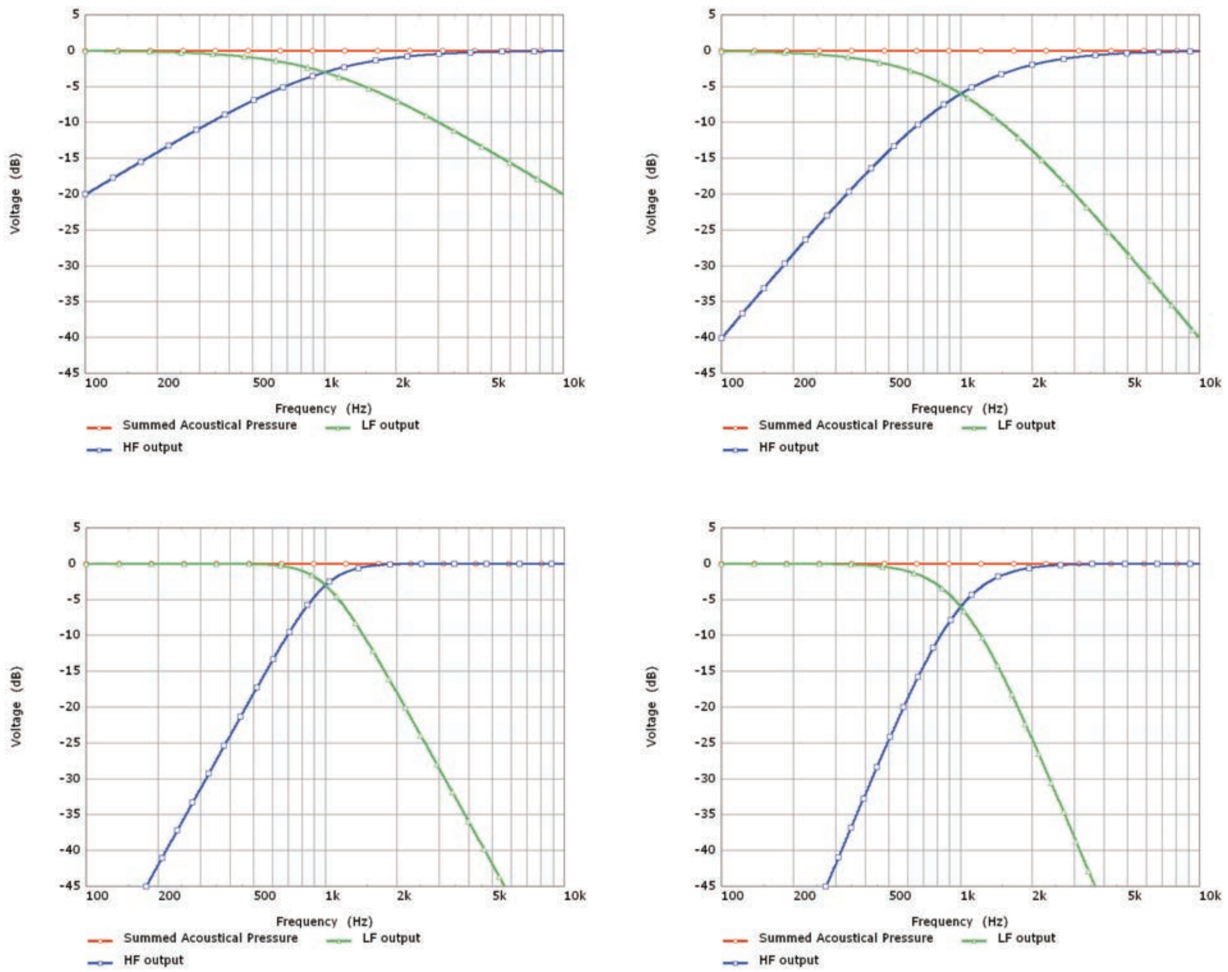

Figure 56. Classical crossover filter arrangements, from top left to bottom right, 1st order Butterworth, 2nd order Linkwitz-Riley, 3rd order Butterworth, 4th order Linkwitz-Riley.

8 This is only true given certain constraints on the load, typically in text books the component values required to achieve the classic filter types are given for a purely resistive load.

In terms of a loudspeaker one might think that a higher order will always be preferable, because the tweeter can be protected better from low frequency signals and the woofer signal can be curtailed before any high frequency cone breakup. However, there are several considerations that are not shown by the frequency level response. Firstly, a greater number of components are required for a higher order crossover and this makes it much more difficult to design a passive circuit that is completely transparent and distortion free. Secondly, with increasing filter order the summed signal becomes increasingly smeared in time. One way to look at this time smearing is in terms of the relative transmission delay of the crossover to different frequencies, this information is shown in Figure 57. It can be clearly seen that the time smearing is worse for the higher order crossovers. Note that this time smearing is inherent to the type of filter rather than the realisation of the filter. This means that irrespective of whether the filter is passive active or digital, a Butterworth 3rd order crossover filter will result in the same group delay characteristic.

Polarity inversion of the some of the drivers in a multi-way loudspeaker in order to avoid interference dips at the crossover frequency is commonplace. This is necessary because of the time smearing of the crossover filters and the drivers themselves. This can be a source of confusion. For example the classical 2nd order Linkwitz-Riley crossover requires one of the driver polarities to be reversed, whereas the 4th order Linkwitz-Riley does not. The time smearing of a second order crossover means that at the crossover frequency the low frequency signal is delayed by half a wave period and the polarity reversal is necessary to achieve proper summation between the two drivers. The 4th order crossover has more time smearing and the low frequency signal is delayed by a full wave period and hence no polarity reversal is required. There is no advantage to the fact that tweeter polarity is not inverted in the 4th order system, this is simply a consequence of the greater time smearing.

The 1st order Butterworth crossover filter has the interesting property that there is no time smearing. This fact, combined with it's simplicity, has made it very popular among audiophile loudspeakers. However, the lack of time smearing is something of a fallacy. While it is absolutely true that complimentary 1st order Butterworth filters will sum to a result without time smearing, in these classical crossover types the

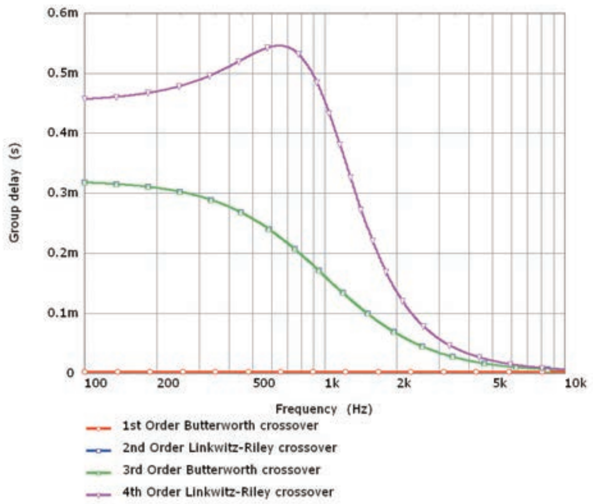

Figure 57. Group delay of classical crossover filter arrangements. Note that the group delay of the 2nd order Linkwitz-Riley and the 3rd order Butterworth crossover are identical

raw response of the loudspeaker drivers is completely neglected. In reality the raw driver responses must be taken into account. Figure 58 demonstrates this. The left figure shows the raw response of a theoretical tweeter with an exceptional response of -3dB at 300Hz. The centre and right hand figure show the overall response and group delay when a 1st order Butterworth crossover is used with this tweeter and a perfect woofer. Even with this exceptional tweeter response the overall frequency response is no longer flat and there is significant time smearing evident in the group delay plot. The reason that this approach did not work is that the response of the drivers must be taken into account when designing the crossover, it is the overall transfer function of crossover filter plus driver that counts. Unless the drivers themselves have no roll-off then it will not be possible to achieve the theoretical zero time smear of the 1st order Butterworth crossovers. In practice this is simply not possible.

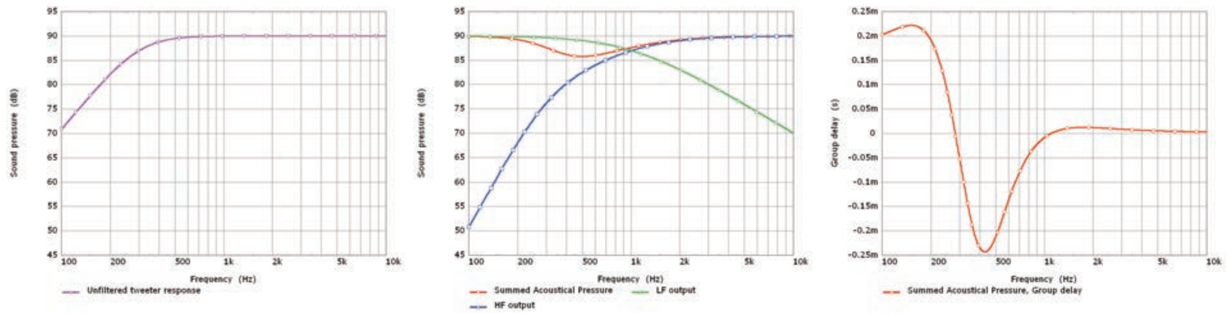

Figure 58. Raw tweeter response (left), 1st order Butterworth crossover using this tweeter (centre), overall group delay (right).

It is possible to achieve a reasonable response using this tweeter and a first order crossover, by inverting the tweeter polarity and manually optimising the filters, as shown in Figure 59, however note that the group delay is now comparable to a second order crossover.

An important point to note is that the time smearing from a crossover is dependent upon the crossover frequency. This is because the group delay introduced by a crossover is proportional to the wave period at the crossover frequency. This means that time smearing is much more problematic with low frequency crossovers. This is a significant reason why a 3-way design approach has been taken with The Reference, in order to avoid a very low crossover frequency necessary on a 4-way system. It is also one of the reasons why it is extremely difficult to crossover convincingly between main loudspeaker and a subwoofer.

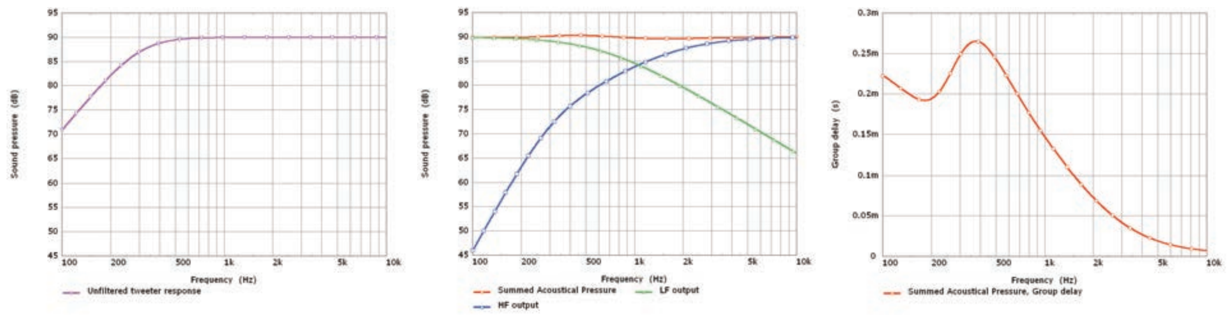

Figure 59. Raw tweeter response (left), manually optimised 1st order crossover using this tweeter (centre), overall group delay (right).

Using digital filters it is possible to achieve crossover filtering without any time smearing and with a high order roll off. The filters required to achieve this exhibit an unusual time domain impulse responses, an example is shown in Figure 60, with a pre-ringing before the main body of the impulse. There is some evidence to suggest that this pre-ringing is responsible for audible artefacts [13] [14]. Provided that the output from the drivers arrives at the listener in sync, the pre-ringing of the filters will cancel in the summation between the high and low frequency sections of the crossover. However, with most loudspeakers having separate treble and midrange drivers the relative arrival time of the tweeter and midrange signal is not constant and varies with listener position. Even with the listener located in a position where the sound arrives in sync from the two drivers the reflections from the room boundaries are likely to contain traces of the pre-ringing artefact.

With conventional loudspeaker systems the choice of crossover filter has a large effect on the dispersion of the loudspeaker. The dispersion of The Reference models, due to the Uni-Q driver and careful design of the low frequency directivity, is largely unaffected by the filter choice. This allows much more freedom to design a crossover with as little compromise as possible.

The Reference crossovers use a combination of 1st and 2nd order electrical filters. These filters do not follow any classical type but rather are carefully designed, using a combination of computer modelling and listening tests, to ensure good overall summation and equalisation of the natural driver responses. In some places a Thiele style notch is incorporated into some of the filter sections [15]. The emphasis has been to try and use the simplest possible filters that are able to fully control the drivers. Filters with an order greater than two have been avoided because of the additional components that they add to the signal path (see Figure 61)

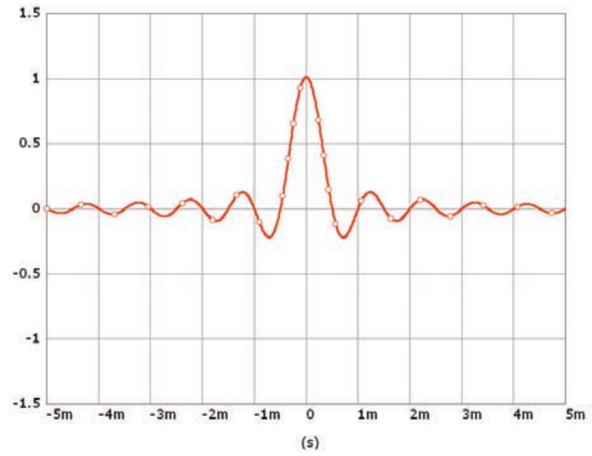

Figure 60. Example of a linear phase filter impulse response.

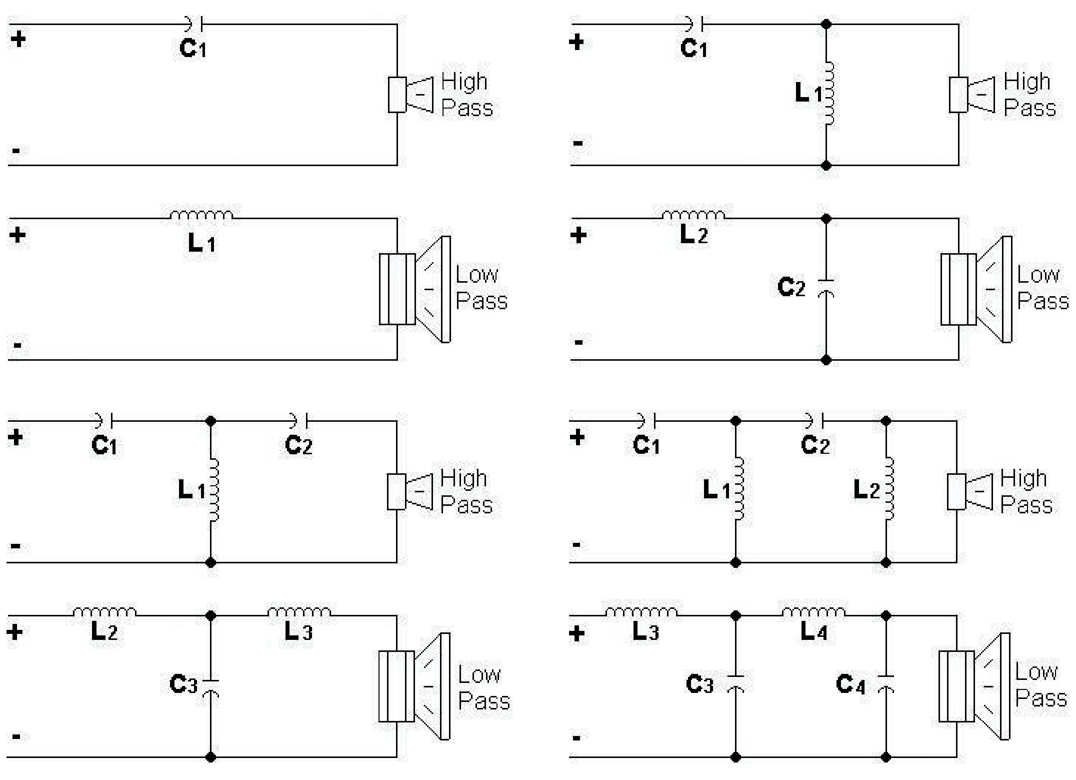

Figure 61. Passive crossover component layouts for 1st (top left), 2nd (top right), 3rd (bottom left) and 4th (bottom right) orders

4.5.3 Impedance conjugation and amplifier load

The Reference loudspeakers make use of impedance conjugation networks in the crossovers. Conjugate networks are required because real drivers have an electrical impedance that is not constant with frequency. Figure 62 shows a typical driver's electrical impedance magnitude compared to a resistor. With a resistive load, a passive filter circuit can achieve any of the classical filter responses exactly. Figure 62 demonstrates this for a 350Hz 2nd order Linkwitz-Riley crossover, with a resistive load the correct filter response is achieved. Also shown is the response achieved with the same passive filter loaded with the driver shown in Figure 62, due to the electrical interaction between the passive filter and the driver, the filter response target is not achieved. It is possible to get a slightly better match by manually optimising the crossover components, the result is shown in Figure 62. However, the version with the driver load still shows some response irregularity as well as attenuation of the very low frequency response.

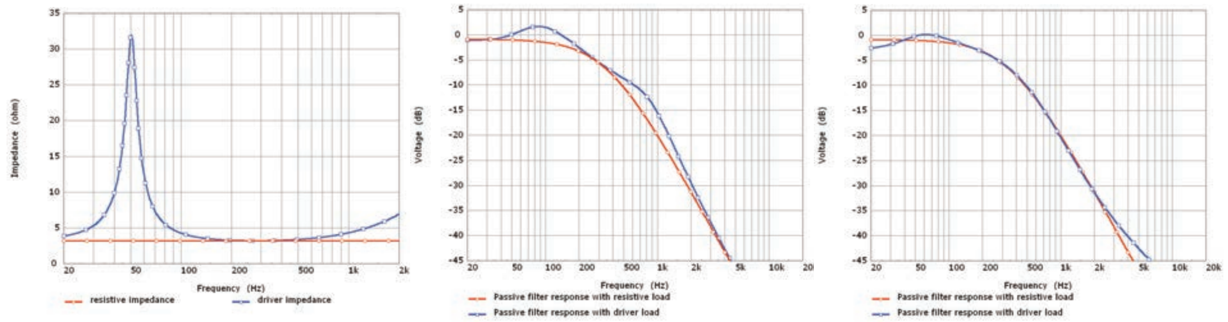

Figure 62. a. Comparison of electrical the impedance of a real loudspeaker driver and a purely resistive impedance (left).
b. Passive low-pass filter response with each load, filter designed to give 350Hz Linkwitz-Riley response with a 3.2 Ohm load (middle).
c. Passive low-pass filter response with resistive and real loudspeaker driver load, filter designed to give best match to 350Hz Linkwitz-Riley response in each case (right).

These interaction issues can be overcome by using impedance conjugation networks. This entails adding extra crossover components in parallel with the drivers to compensate for the natural driver electrical impedance. For example, figure 63 shows the input impedance for the same driver discussed above with the addition of a conjugation network to compensate for the impedance peak at driver resonance and for the rising impedance at high frequencies from the coil inductance. With both conjugation networks in place the electrical impedance of the speaker is almost perfectly resistive and consequently the response irregularity issues above completely disappear.

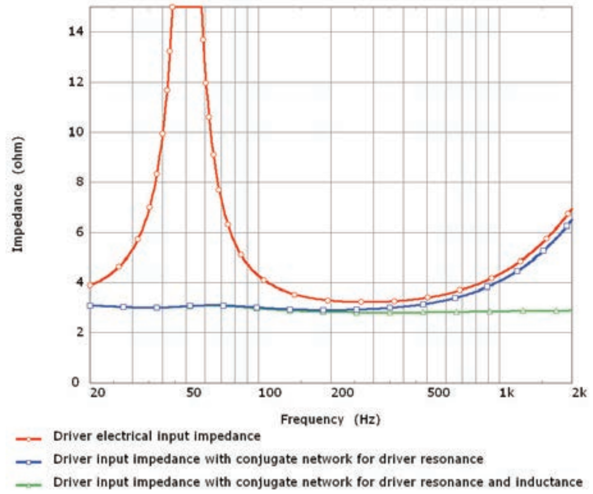

Figure 63. Typical loudspeaker driver electrical input impedance with impedance conjugation components

There is also a big change in the amplifier loading when an impedance conjugation network is added to a crossover. Figure 64 shows this amplifier load difference for a typical two-way closed-box loudspeaker system as an impedance magnitude and phase chart. This type of figure is widely used in various publications to aid consumers to asses the relative driving difficulty when comparing loudspeakers. Unfortunately this information does not lend itself to easy comparison as both low impedance magnitude and high impedance phase lead to difficult loudspeaker loads. For example, considering Figure 64, the impedance magnitude of the version with the conjugate network is lower at 60Hz whereas the version without a conjugate network has a significantly higher impedance phase angle. Another approach to looking at the impedance data has been suggested by Keith Howard [16], based on the work of Benjamin [17]. This approach presents the equivalent purely resistive load that would result in the same peak power dissipation in the output stage of a class B amplifier, thus combining impedance magnitude and phase into one figure of merit. This measure is called the "equivalent peak dissipation resistance" (EPDR), plotted in Figure 64 for the same two loudspeakers. From this plot it is clear that the loudspeaker without the conjugation network is a significantly more difficult load to drive. Not only does this mean that a loudspeaker with impedance conjugation may be used with a wider range of amplifiers, but it also means that the loudspeaker will perform better with any given amplifier because less power is dissipated in the amplifier output stage in trying to drive the awkward load.

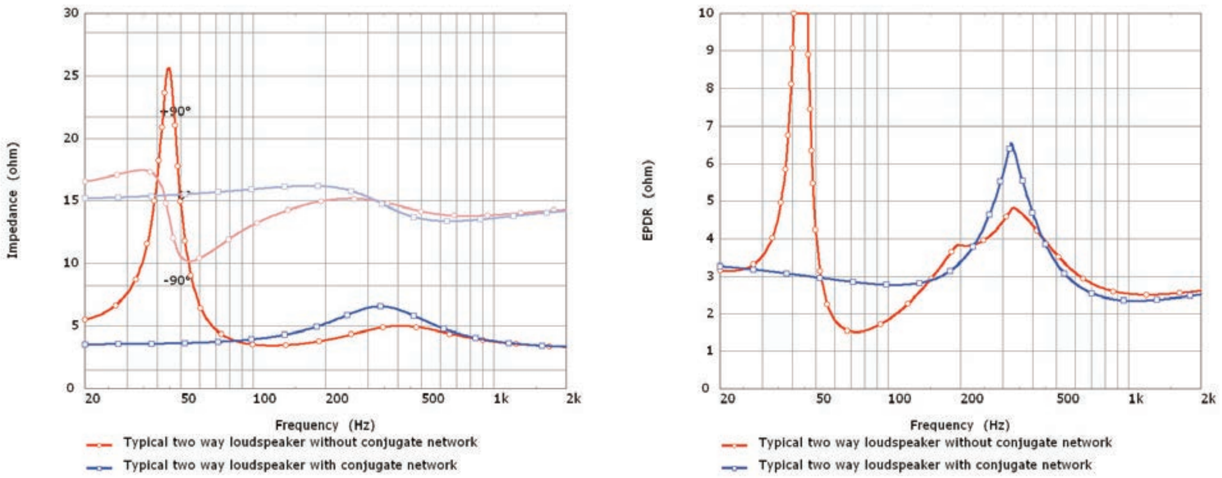

Figure 64. a. Typical closed-box two-way loudspeaker system electrical input impedance with and without a conjugate network in the crossover (left). b. Equivalent Peak Dissipation Resistance for the same two systems (right).

There is also a clear disadvantage to using impedance conjugation networks and that is that the number of components in the crossover is increased significantly. The impact of this is that there is more opportunity for the crossover to generate artefacts and distortion into the signal. For The Reference impedance conjugation has been used selectively where necessary to control the crossover to loudspeaker interface and to ensure that the loudspeaker load is reasonable for a good quality amplifier. By selective application of this approach it is possible to benefit from the advantages described above while not adding too many additional components to the crossover.

The result is a good balance between the different compromises whilst still retaining a relatively easy load for the amplifier. The EPDR of an early prototype of Reference 5 is shown in Figure 65 below in comparison to a competitor loudspeaker.

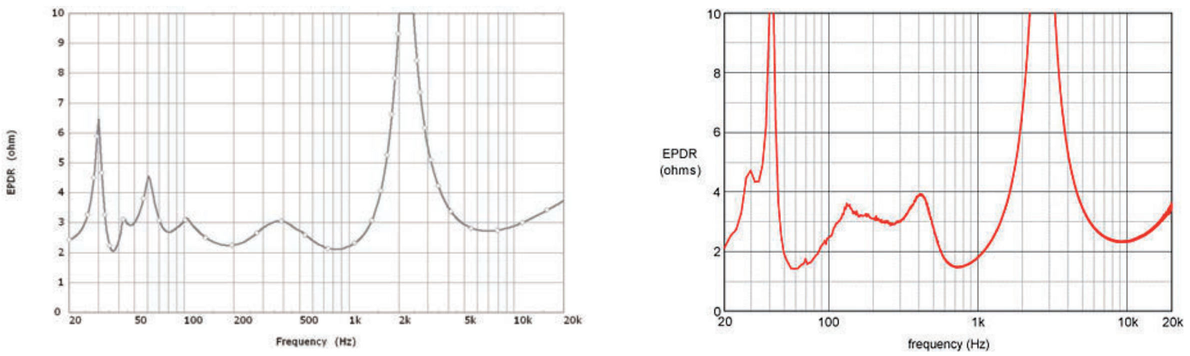

Figure 65. EPDR of Reference 5 prototype system (left) compared to competitor loudspeaker.

5 Voicing the loudspeakers

The crossovers were initially designed from measured responses of the individual drivers mounted in the final enclosure and based on the underlying approach outlined in section 4.5. Listening tests were used extensively to arrive at the final design for each model.

The listening sessions for the early prototype loudspeakers spanned several months and involved a large number of experienced KEF personnel, each with their own personal music tastes and preferences. The listening process is very much integrated with the technical engineering. Concerns raised during listening were rigorously investigated using various measurement techniques to try and identify any underlying physical cause. In numerous cases a clear correlation was found between objective data and subjective conjecture. This approach, though time consuming, is very effective. When issues are raised during listening it is often tempting to immediately change the loudspeaker voicing in an attempt to improve the loudspeaker's character. However, if an underlying problem exists then it will remain unresolved and the loudspeaker's balance will be compromised in order to disguise it.

Particular attention was paid to the component selection and the crossover layout. The short-listed components following the component distortion measurements described in section 4.5.1 were individually auditioned to ensure they did not limit the perceived sound quality. The capacitors for the higher-frequency section are vibration damped with mastic, to prevent sonic deterioration due to vibration. The crossovers are split into two different circuit boards and mounted apart from one another inside the loudspeaker. The crossover circuit boards are decoupled from the cabinet walls to minimise vibration transfer to the components. Each board only contains three inductors in order that they can be oriented in perpendicular planes to minimise cross talk.

Gratifyingly, as this systematic procedure progressed the preferred subjective performance merged with the balance that gave the best objective measurements. In particular the lowest distortion and the smoothest and flattest on and off-axis frequency response.

Appendix VII. Room modes and loudspeaker positioning

After the loudspeaker, the audio signal goes through one more process before reaching the ears: propagating across the room. This can have a significant effect on the overall performance of the loudspeaker, especially at low frequencies where room modes can play a prominent role in the clarity and evenness of the sound. No two rooms are identical, and the effects are difficult to predict. However, these low frequency problems can usually be significantly reduced through careful positioning of the loudspeaker and listener.

VII.I. What is a room mode?

Room modes occur at frequencies where the wavelength of the sound coincides with a dimension of a room, resulting in a standing wave resonance. Longer dimensions have lower frequency room modes than shorter ones. They are typically problematic below 200Hz, and occur in all 3 dimensions of the room. The standing waves formed in rooms are acoustically similar to those formed in the loudspeaker cabinet, as explained in Section 4.3.1.

Figure 76 shows the pressure of a room mode along a length of a room. The strength of the mode varies in a sinusoidal shape, being stronger at the walls due to the build-up of sound pressure. There is a node in the centre of the room where the pressure of the mode is zero. The pressures either side of the node are opposite phases, such that the total pressure in the mode remains zero.

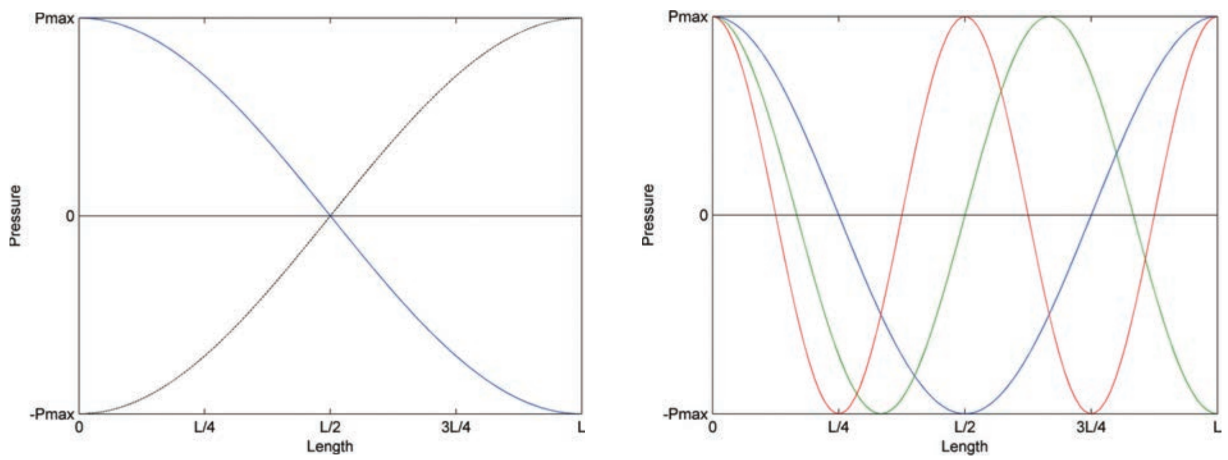

Figure 76. Modal pressure distributions between parallel walls for first (left) and second, third and forth modes (right).

Modes will also occur at integer multiples of the original resonant frequency (e.g. 20Hz, 40Hz, 60Hz, 80Hz…). At these frequencies an additional half sine wave can be fitted into the length of the room. The first mode will be half a sine wave (180 degrees), the second a whole sine wave (360 degrees), the third one and a half (540 degrees) and so on. Figure 76 shows the next 3 room modes, fitting between one and two wavelengths along the length. As can be seen the positions of the nodes and anti-nodes varies with the different resonances. When the loudspeaker and listener positions are moved around the room, the effect of the room modes will change depending on their proximity to a pressure maxima or minima.

The modes described so far are axial room modes, involving standing waves formed between a pair of parallel opposite walls. These are the strongest modal resonances in a room. If the room dimensions are integer multiples of each other, the modal resonances of each dimension will occur at the same frequency. This gives a low modal density with fewer but stronger modal resonances, which is more audible than a greater number of weaker resonances. The stronger resonances are harder to control, which is why shapes such as cubes are avoided when designing listening rooms.

Modes will also form between two pairs of walls such as the length and width, called a tangential mode. Similarly, modes can form across all three dimensions, called an oblique mode. Tangential and oblique modes are weaker than axial modes, but can reinforce resonances if they coincide in frequency. If room modes are strong in a room, then they can be damped using acoustic treatment, although large quantities are usually needed due to the long wavelengths at low frequencies.

VII.II. The effect of modes on loudspeaker responses

When a loudspeaker is placed in a room, if it is near to a pressure maxima of a mode (the point at which the magnitude of the pressure is at its largest value) it will excite it strongly. If the loudspeaker is placed near to a modal node (the point where the pressure is zero), then the mode will be excited much less. Excited modes result in resonant peaks in the response coinciding with the modal frequencies, and ringing in the time domain. Between the peaks nulls are often seen in the frequency response and occur due to destructive interference between adjacent room resonances. The depth of these nulls is greatly dependent upon the speaker position. In many ways these nulls can be more audible than the peaks as they tend to result in a perceived lack of bass.

Figure 77 shows a real in-room measurement. On the left of the figure a standard frequency response plot is shown, on the right a time frequency representation of the same data. Comparing the two, it can be seen that the peaks in the frequency response coincide with time domain ringing. Similarly, the dips in the frequency response result in dips in the time response. Note that the decay time associated with these irregularities is far greater than any of the other time smearing discussed in the body of this paper. These measurements are of a worst case scenario with the loudspeaker in the corner, where it excites all the room modes

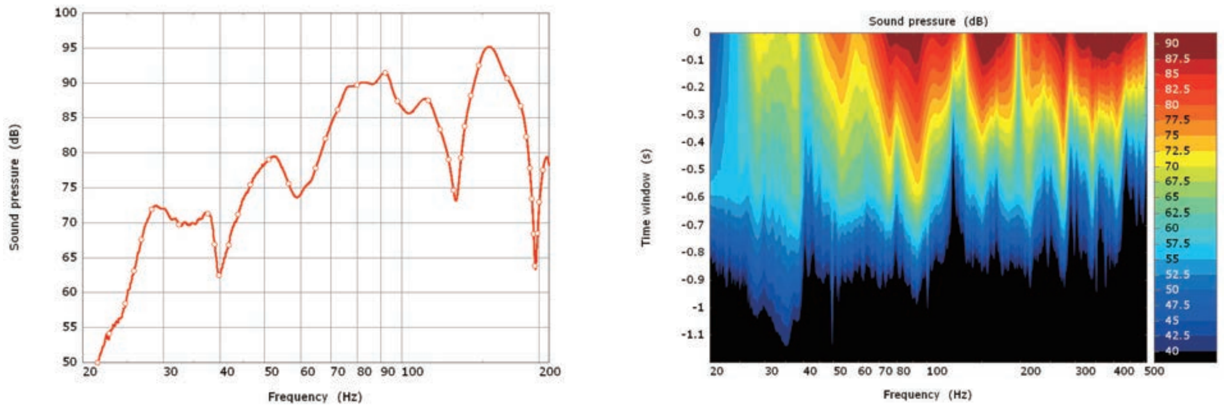

Figure 77. In-room loudspeaker response shown as frequency response chart (left) and short time Fourier response (right).

VII.III. Positioning

The effects of room modes can never truly be avoided. They can however be significantly reduced through careful positioning of the loudspeakers in the room. Small changes to the loudspeaker position can have a very large effect. For example figure 78 shows the response of a floor-standing loudspeaker measured from the listening position for two different positions in the room. The two positions are only 0.5m apart from one another yet the change in the response is as much as 7dB. The sound of the loudspeaker would be very different in these two locations. It is tricky to provide general advice that will work in all rooms, though there are some tips on speaker placement in appendix VIII. The best approach is to experiment with different loudspeaker and listener locations to find a layout with a balanced overall sound. In some rooms, changing the speaker location by as little as 10cm can make a big difference.

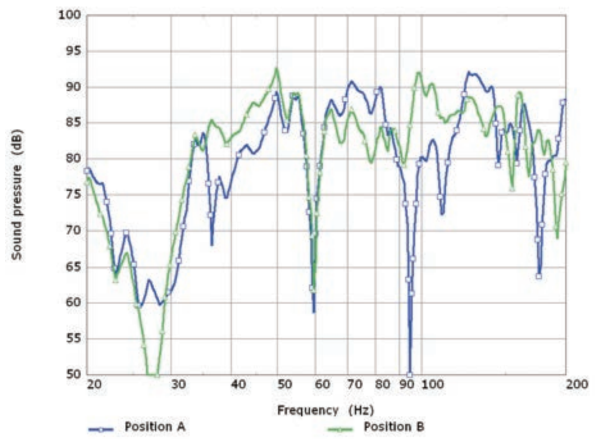

Figure 78. In-room loudspeaker measurements for two different loudspeaker positions 0.5m apart.

Appendix VIII. General loudspeaker positioning tips

- The speakers should be placed between 2 and 4m apart and the listener should be positioned exactly equidistant from the loudspeakers, and at approximately the same distance the speakers are apart.

- Ideally a wall should be directly behind the loudspeakers (the “front” wall). Both speakers should be the same distance from this wall.

- The distance between the speakers and the front wall is normally the most sensitive parameter for low frequency tuning. It is recommended that the listener fine-tune the loudspeaker to rear wall spacing to find the optimal distance. Normally this distance would be between 0.3m and 1.5m to the rear of the loudspeakers.

- Avoid placing the listening position in close proximity to any of the walls.

- Symmetry is very important for optimal stereo. Ideally the side walls, to the left and the right of the loudspeakers, should be at the same distance and of the same construction.

- The side wall has only a small effect on the timbre of a well designed loudspeaker provided it is not closer than around 1m. Ideally the side wall should be between 1 and 3m from the loudspeaker.

- If acoustically treating the room, it can be very helpful to add diffusion (uneven reflective surfaces) to the side walls. This helps to reduce timbral imbalance due to the side walls while maintaining a good stereo image.

- Toe-in can help to fine tune the stereo image and the perceived high frequency energy. Maximum high frequency energy will reach the listener when the tweeters are pointing directly at the listening position. With Uni-Q the balance does not suffer if you listen off axis, so using a flatter setup with less toe can be a very good option for lively rooms.

- Use a tape measure to ensure your spacings are exactly the same, the ear is very sensitive to arrival time.

- If it is not working, consider a complete change of approach. By using a different wall behind the loudspeakers you will most likely end up with a different listening position too. The listening position is just as important as the speaker positions.

- Bass traps can help a great deal to tame room mode resonances which will tend to improve bass overhang. However, they need to be physically large to have any effect on deep bass. Beware of poorly designed traps. A typical bass trap to effect frequencies below 80Hz needs to have a volume of approximately 0.5m3.