



The Sound of Science

July 2017

Preface

We have always believed that the most sophisticated sound receiving device of all time is the truly incredible human ear, and we both acknowledge and are challenged by how this wonderful instrument can detect the most minute of sonic differences that we currently struggle to measure. Exhaustive listening tests therefore remain an essential element in turning a very good cable in to an exceptional one.

But to create a very good cable in the first place, requires an understanding of the science of cable design. It was for this reason that in 1995 we published our very first 'white paper' on speaker cable design titled 'The Genesis Report'. This established design principles that still hold true today, some 22 years after its original publication

Over the years, we published additional white papers covering HDMI & digital audio and this new report 'The Sound of Science' consolidates all of this research in to one publication.

I believe it makes fascinating reading.



Bob Abraham
Co-founder QED
July 2017

Introduction

Around the world, debate continues among both audiophiles and videophiles as to which cable designs offer the best performance.

QED, an established market leader in this field, has won many awards for its high-performance audio and A/V cable products. Key to this success has been the application of strong engineering principles, coupled with extensive test and measurement of both our own and competitors' products.

Listening tests are vital: QED engineers are all too aware that measurements alone don't tell the whole story. But if any cable introduces measurable errors and distortions it obviously cannot convey the signal accurately. QED believes that cables should be as accurate, transparent and neutral as possible, and it is with this credo that our cable development is undertaken, based on measurement and guided by exhaustive listening evaluations.

Here we re-visit some of the areas covered in previous QED cable reports and add three new sections, one about biwiring loudspeakers, one about HDMI cables, and one about digital audio optical cables.

The cable's role

Ideally every cable should transfer a signal between two items of equipment with zero loss and distortion. In the real world this is not possible because subtle changes occur in the signal and these may result in readily perceived changes to sound or video quality. The degree of signal degradation is determined directly by the design of the cable.

Maximising real world cable performance requires an understanding of the signal transmission process and the engineering tools available, to ensure that the signal arrives in the best possible condition.

Some theory and interesting facts

Cables appear at first sight to be very simple components in both their construction and operation. Once you look at them in detail, though, they become more complex and a lot more interesting.

Conductors

Metal conductors such as copper work very well. These provide free electrons which are not locked into the metal's atomic structure and can therefore move. There are an incredibly high number of free electrons available: in just one cubic centimetre of copper there are approximately 8.5×10^{22} – almost 100,000 million million million.

The movement of free electrons is random, normally with the net effect of zero current flow. But when a potential difference – a voltage – exists across the conductor, an electrical current flows in response.

Copper is an excellent conductor because its unbound electrons have a large mean free path of about 100 atomic spacings between collisions. The electrical resistivity of a conductor is inversely related to this mean free path, so copper's resistivity (its resistance to current flow) is low. Silver has lower resistivity still – the lowest of all the metals – but copper is far more abundant, and therefore cheaper, and almost as good, so it is used more extensively.

Unfortunately all metal conductors possess finite resistance which opposes the movement of electrons and causes some of the electrical energy to be dissipated as heat. Only superconductors have zero resistance, but a superconductor that works at room temperature has yet to be discovered.

How does the signal travel and how fast?

Scottish mathematician and theoretical physicist James Clerk Maxwell (1831-1879) provided the answer. Signals in fact travel down cables as an electromagnetic wave, and this wave travels at a very high velocity – close to the speed of light. If the metal conductor was perfect virtually all the energy would be transferred from one end of the cable to the other by this means.




Figure 1. Magnetic and electric field alignment in an electromagnetic wave

The electromagnetic wave consists of electric and magnetic fields that are spatially at right angles to each other and to the energy flux or power flow along the axis of the cable (Figure 1). Electromagnetic wave theory further shows that when conductor pairs are used, the majority of the energy propagates in the space between the conductors (Figure 2).




Figure 2. EH field pattern around a conductor pair

These fundamental principles apply to all types of signal transmission, whether low frequency audio or high frequency video, digital or analogue. Free electrons are required for the electrical conduction of signals but signal energy travels principally in the fields in and around the conductor.

The importance of dielectrics

Since the space around the conductor carries signal energy, it's obvious that the material in that space is an important component of cable design. The dielectric space around the conductor has a direct effect on the propagation of the electromagnetic wave along the cable, particularly in terms of its speed, energy storage and energy loss.

Dielectric loss increases linearly with frequency. In the non-conducting dielectric, the electrons are bound to the atoms but when an electric field is present, the orbit of the electrons becomes distorted as the negatively charged electrons become attracted to the positive conductor. Displacement of these electrons requires energy, which is drawn from the signal and ultimately dissipated as heat within the dielectric. The better the dielectric, the lower the loss.




Figure 3. Variation of magnetic field strength through and beyond a solid and hollow conductor (not to scale)

We know that electromagnetic fields around the conductor radiate for some distance, rising to a maximum at the conductor-dielectric boundary (Figure 3). This means that objects and materials (other than the dielectric) which are present within the electric and magnetic fields will also have a direct effect on the propagation and power loss of the signal.

The importance of characteristic impedance

This area is often misunderstood. For audio signals (20Hz to 20kHz) travelling over a short distance (some metres), the characteristic impedance of the cable has virtually no impact on the energy transfer between the transmitter (amplifier) and receiver (loudspeaker), and therefore is not a consideration in speaker cable design. This changes once the wavelength of the electromagnetic wave traveling down the cable becomes shorter than the cable length. Designing a high performance cable that carries a radio frequency (RF) video signal consequently requires considerable care to ensure that signal energy is not lost.

At low frequencies the characteristic impedance – which is determined by the physical dimensions of the cable’s conductors, their spacing and the relative permittivity of the dielectric – is a complex quantity but at high frequencies it becomes a resistance.

Ideally the cable should transfer all the electromagnetic wave energy into the receiver, without any energy being reflected back. If the wave is partly reflected, the reflection will interfere with the incoming signal. To prevent this and ensure reflection-free signal transmission it is necessary for the transmitter output impedance and receiver input impedance to have the same value as the cable’s characteristic impedance.

The effect of having mismatched impedances is readily observed using Time Domain Reflectometry (TDR) in which a very short pulse is sent down the cable. Any mismatch in terminating impedance causes some of the electromagnetic wave energy to be reflected back after a short delay. If the cable is unterminated then virtually all the energy is reflected back. Using TDR we can establish how well a cable is manufactured and how well-engineered (matched in impedance) its connectors are (see Figure 4).




Figure 4. TDR traces showing energy reflection from the end of a cable

When designing cables for carrying high frequency signals the connector impedance is critical. When we developed QED’s TTV range of aerial cables, TDR analysis showed that many so-called ‘high end’ connectors did not have the correct impedance, causing a significant fraction of the signal energy to be reflected back to the source. When signal levels are particularly low, as is the case with antenna feeds, the outcome is poor picture quality and fewer stations in the channel listing.

**Conductor
Chemistry**

A conductor's chemical composition is very important – the presence of trace elements such as silicon, magnesium and phosphorus decreases electrical conductivity. The most widely used copper for cables is electrolytic tough pitch (ETP) copper, which consists of extremely high purity metal plus oxygen in the range of 100-650ppm (parts per million) concentration. Oxygen is used as an alloying element and also as a scavenger as it reacts with most of the impurities in the copper. Adding around 0.02% oxygen to ETP copper increases its conductivity.

Oxygen-free copper (OFC) is produced primarily for its ability to be heat treated without embrittlement, for ease of use when welding and brazing. The acronym for oxygen-free high thermal conductivity copper, OFHC, is a registered trademark of Phelps Dodge Specialty Copper Products. OFHC is a highly refined grade of copper that contains almost no oxygen or other impurities and this is the grade employed exclusively in all QED cable products. Certified oxygen-free high thermal conductivity copper contains a minimum of 99.99% copper making it the purest metal in common use.

If the temperature and time constants are carefully controlled in an oxygen-free environment while the copper is manufactured the grain structures in the material can be reduced so that the purity of the copper increases to as much as 99.999% (termed 'five nines'). The vast majority of QED cables, both in our speaker cable and interconnect ranges, utilise five nines copper. The IACS (International Annealed Copper Standard) gives a percentage ranking for conductivity of OFHC at around 102.4% whereas ETP coppers fall typically within the range 100-101.5% IACS. Although linear crystal coppers approaching a purity of six nines have been made in small quantities by various continuous casting or Czochralski processes, it remains prohibitively expensive for use in commercial products and, in the absence of any official IACS designation, at ambient temperatures has no conductivity advantage over its five nines counterpart. OFHC is particularly well suited to cryogenic treatment, in which the copper's temperature is slowly reduced to below -190°C and then gradually returned to ambient, and QED takes advantage of this property by cryogenically pre-treating its high-end cable range. It is thought that the improvement in physical properties that occurs when copper is cryogenically treated results from the elimination of dislocations in the material's microstructure which might otherwise impede the free movement of electrons throughout the crystal lattice.

In the digital world – jitter and crosstalk

Jitter is a term used to describe time variation in the arrival of a digital signal at the receiver. All metal cable links have a finite bandwidth which attenuates the high frequency components of the digital signal. A transition delay (that is, delay in the receiver detecting a 1 or 0) occurs which varies depending on the pattern of the digital signal. Digital equipment normally has engineered-in tolerance to jitter but it is not always effective, depending on the frequency of the incoming jitter. In such cases cable-induced jitter is transferred to the recovered clock signal, as is the case with an S/PDIF digital audio signal. For long cable lengths the digital signal 'eye pattern' closes and reduces in size, ultimately resulting in a failure of the link (Figure 5).




Figure 5. Digital data signal eye-patterns

When more than one signal channel is transmitted down a single cable, as is the case with digital video, 'crosstalk' – the unwanted transfer of energy from one channel to another via inductive or capacitive coupling – must also be considered. Individual screening of signal pairs can improve crosstalk performance, but this can cause a reduction in the available bandwidth. The effect of this is a reduction in the pulse rate time or slew rate, which degrades signal transfer. Various design techniques can be used to improve performance, such as increasing the conductor pair twist rates, ensuring that conductor pairs have different twist rates to minimise coupling, and using better-performing dielectric materials.

Loudspeaker cables

Because cable is a link within a system its assessment must be considered within the context of its two partners – in the case of loudspeaker cable, the amplifier and loudspeaker. The loudspeaker cable is, in effect, an extension of the amplifier's circuit, equivalent to adding extra components to its output with the electrical properties of resistance (R), capacitance (C), inductance (L) and conductance (G).

Most power amplifiers ensure fidelity by comparing their output signal with their input using a technique called negative feedback (NFB). Any error appearing at the output of a NFB amplifier is corrected by the amplifier automatically applying the opposite error at its input. As can be seen from the diagram in Figure 6, an NFB amplifier can only correct errors that appear at the point of feedback take-off. Errors at the loudspeaker input due to the influence of the cable remain uncorrected. (A few NFB amplifiers have been designed to take their feedback from the loudspeaker terminals in an attempt to counteract cable effects, but this configuration is very rare.)




Figure 6. Block diagram of a negative feedback power amplifier connected to a loudspeaker

We can assess a cable's performance by comparing its input (at the amplifier output) to its output (at the loudspeaker input). Any difference amounts to degradation of the signal.

**Real effects
on system
performance**

Terms used to describe the subjective effects of cables range from positive ones such as 'transparent', 'coherent', 'tight', 'detailed' and 'rhythmical' to negative comments such as 'grainy', 'loud', 'forward', 'twangy' and 'smeared'. Our research has shown that some of these can be explained by analysing electrical measurements made with the cable connected to a real loudspeaker load.

The two graphs in Figures 7 and 8 are frequency responses, measured using two different cables. The upper trace in each was taken at the amplifier's output terminals and the lower curve at the loudspeaker's input terminals. There are clear differences in performance between the two cables.




Figure 7. Frequency responses at either end of a low-resistance loudspeaker cable




Figure 8. Frequency responses at either end of a high-resistance loudspeaker cable – compare with Figure 7

Figure 7 was obtained using a ribbon cable with very low loop resistance, while Figure 8 shows the effect of a solid-core cable. The ripples in the lower curves, measured at the loudspeaker input, are due to the frequency-dependent variations in the impedance of the loudspeaker system within the audio bandwidth affecting the voltage drop across the cable. The loss in the cable is the difference between the upper and lower curves. Clearly, there is a greater loss in the solid-core cable (Figure 8) due principally to its greater DC resistance. Here the frequency response of the loudspeaker is modified by as much as 0.8dB at 200Hz.

Given this clear evidence that low cable resistance is necessary to ensure the flattest possible frequency response into real-world loudspeakers, it's surprising that there was ever a trend away from low-resistance cables to higher-resistance solid-core types.

Why not just use a large cross-section cable to reduce resistance?

In traditional large cross-section cables, whether solid-core or multi-stranded, high frequency signals are forced to travel towards the outside of the conductor. Less and less of the available cross-sectional area is used as the signal frequency increases. This is called 'skin effect'. It means that at high frequencies the resistance of the cable appears to be much higher than it does for low frequencies, and this has a detrimental effect on the fidelity of the sound you hear. QED AirCore™ Technology solves this problem by creating a hollow tubular conductor geometry through which all frequencies can pass with equal ease.

AirCore™ Technology also features a special Litz-like construction designed to obviate an equally damaging sonic problem known as 'proximity effect'. Where two conductors are laid side by side and carry current in opposite directions, the alternating magnetic fields built up around each conductor tend to reinforce current flow in the parts of the conductors which are nearest to each other and to cancel current carried in the side of the conductors furthest away.

In a normal figure-of-eight speaker cable – even one below the size where the skin effect becomes a problem – proximity effect will cause the resistance to rise by up to 25 per cent within the audio band, which affects linearity and therefore fidelity. QED Supremus cable avoids this problem entirely. Each of the sixteen silver-plated, 99.999% oxygen-free solid copper strands which comprise the conductors are individually insulated by a nearly invisible layer of enamel. This material was chosen because it has extremely good insulating properties when deposited in a very thin layer and so can electrically separate each strand within the conductor bundles. Because each strand is twisted around a hollow central polyethylene core, no strand remains at the outside or inside of the overall conductor along its entire length, thus evening out the current density and keeping the resistance more uniform through the audio band.




Figure 9. Resistance versus frequency for QED Supremus speaker cable versus a conventional cable

Figure 9 shows how QED Supremus speaker cable (green trace) compares with a standard cable (red trace) of the same cross-sectional area. Above about 1kHz the standard cable exhibits rising resistance due to the proximity effect whereas the QED Supremus cable does not exhibit any resistance increase until beyond 10kHz.

Inductive Effects

The effects of inductive reactance on AC signal phase shift for a number of cables we've tested is shown in Figures 10 and 11. The higher the cable inductance, the greater the effect on phase shift.




Figure 10. Loop inductance (μH per 10 metres) for 13 different cables




Figure 11. Cable induced phase shift at 20kHz for the same cables as Figure 10

Examining the geometry of each cable sample revealed that the majority of the multi-stranded cables tested were inherently inductive. The inductance of a cable depends on the area of the conductors, their relative spacing and the permeability of the surrounding media. (High permeability materials, such as iron and ferrite, are used to increase inductance, in wound inductors for instance.) In cables, the wider the conductor spacing, the greater the inductance. Many of the multi-strand loudspeaker cables on sale feature conductors that are widely spaced, some by more than three times the conductor diameter, resulting in higher values of inductance.

Averaging the inductive effect across our sample cables gave an effective phase shift of 0.42 degrees per metre. So for a 10 metre cable length there would be 4.2 degrees of phase shift. At present, the audibility of nonlinear frequency-dependent phase shift is uncertain, although amplifiers that exhibit poor phase response are often criticised as being 'grainy'.

Response peaking due to inductance and capacitance

Another effect of inductance is high frequency amplitude loss due to increased cable impedance (inductive reactance increases with frequency). High cable inductance can also be responsible for a rise in voltage at the loudspeaker terminals caused by interaction between inductive and capacitive reactances resulting in a damped resonance. This can be a problem with electrostatic loudspeakers, which present a higher capacitance load than conventional moving coil loudspeakers.

An example of resonant peaking is shown in Figure 12, the two response traces having been measured at either end of the cable (amplifier output and loudspeaker input). Here the increased impedance of the cable at frequencies above the response peak results in considerable loss of signal level when combined with the amplifier's own roll-off.

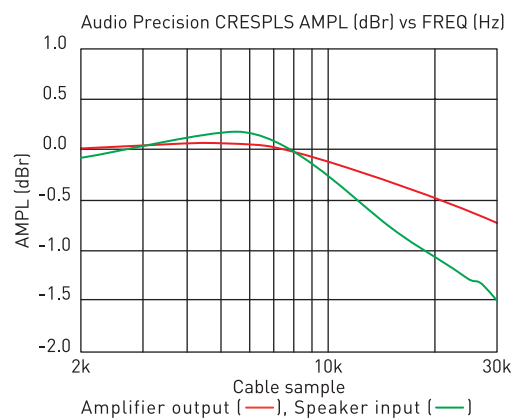


Figure 12. Frequency responses for a high-inductance cable feeding a capacitive load

Dielectrics

Loudspeaker cable conductors are sheathed with insulation, otherwise known as dielectric, to prevent shorting. This inevitably introduces additional losses because, as previously described, all dielectrics absorb some energy. The dielectric loss is sometimes referred to as the dissipation factor or $\tan \delta$ and this increases with frequency. As a general guide, the higher the dissipation factor at a given frequency, the greater the power loss within the dielectric.

The dielectric constant (ϵ_r) and loss ($\tan \delta$) for several popular speaker cable insulation materials are shown in the table below:

Insulator	ϵ_r	Approx. $\tan \delta$ at 10kHz
polyvinylchloride (PVC)	4.0 - 8.0	0.01 - 0.05
polypropylene (PP)	2.25	0.0004
polytetrafluoroethylene (PTFE)	2.10	0.0002
low density polyethylene (PE)	2.20	0.0002
high density polyethylene (PE)	2.30	0.0003
foamed polyethylene (PE)	1.69	0.0001
air	1.0006	virtually 0
vacuum (for reference)	1.0000	0

The results of dissipation measurements on a range of loudspeaker cables are shown in Figure 13. It reveals a surprisingly wide spread of results.




Figure 13. Cable dissipation factor, measured at 1kHz, for the same 13 cables as Figures 10 and 11

All dielectrics also possess a property known as permittivity. The lowest permittivity dielectric (apart from a vacuum) is air and this introduces the lowest loss of any practical material. The greater the permittivity, the greater the power loss and also the higher the capacitance. This is because permittivity is a measure of how easily the dielectric 'permits' the establishment of the electric field, which is the very cause of the cable's capacitive effect. Conversely, the lower the permittivity of a dielectric (the closer it is to a vacuum) the lower will be both the loss and capacitance.

Capacitance is also governed by the spacing and diameter of the conductors. The greater the gap between any two conductors in a given dielectric, the lower the capacitance (the reverse being true of inductance). Designing a cable with low inductance and low capacitance is particularly difficult when poor quality

dielectrics are used. The majority of lower-priced cables and many we tested, use PVC dielectrics, which cause inherently greater capacitance and dielectric losses. Whatever is done with the conductor spacing and diameter, such cables are at a distinct disadvantage, having greater capacitance or inductance or both.

Effects of Capacitance

In theory, the capacitance of loudspeaker cable should have little effect on system performance because the cable is driven by a very low source impedance, typically fractions of an ohm for most power amplifiers. Although the capacitance forms a low-pass filter when connected to this impedance, its effect on frequency response within the audible range is typically minuscule. More insidiously, unduly high cable capacitance in a speaker cable may indicate poor dielectric quality and high dielectric losses.

Some esoteric cables employ a number of separately insulated paralleled wires to form the two conductors. With certain geometries and lower-grade materials, this can cause capacitance to rise to a high level. One cable we tested had a capacitance of 1375 pF compared to the average sample capacitance of 500 pF for a 10 metre length.

Another factor to be considered with speaker cables is amplifier stability. In some cases a little extra capacitance on the output of an amplifier can make it oscillate, overheat and even self-destruct. More usually the amplifier may oscillate momentarily at a frequency above the audible range during operation and show no obvious symptoms. Well-designed feedback amplifiers normally have a good gain/phase margin, which ensures that small extra phase shifts due to increased load capacitance do not cause such problems. But some commercial designs do not have sufficient gain/phase margin for unconditional stability and it is these which can cause problems with long lengths of higher-capacitance cables. Additionally, inductance is likely to be low for high-capacitance cables, leading to a further reduction in the stability margin. Even if the amplifier is not outwardly unstable, sound quality can still suffer, with harsh and forward-sounding results as the amplifier operates on the edge of instability. Figure 14 shows an example of instability due to high-capacitance cable, visible as 'ringing' on a square wave signal.




Figure 14. Square wave 'ringing' with a high-capacitance cable

**Capacitance
versus inductance**

With a single pair of conductors in a given dielectric, reducing their spacing reduces inductance and raises capacitance, while increasing their spacing has the opposite effect. Consequently it is often assumed that it is not possible to reduce inductance without increasing capacitance. Indeed, this has almost become embedded in audio folklore. But comparisons we have carried out between different conductor layouts, even with similar total conductor cross-sectional area (and hence similar DC resistance), have shown that it is possible to achieve low inductance and low capacitance, even with the same dielectric material, simply by rearranging the conductors (see the table below). These results illustrate the profound effects of cable geometry.

parameter	Qudos	Profile 8 (inner/outer)	Profile 8 (left/right)
capacitance (pF/m)	35.6	37.2	21.3
loop inductance ($\mu\text{H}/\text{m}$)	0.55	0.39	0.54
loop resistance ($\text{m}\Omega/\text{m}$)	14.0	15.0	15.0
characteristic impedance (Ω)	118.5	104.7	159.6

Qudos and Profile 8 are two discontinued QED speaker cables. Standard Qudos comprised two bunches of 79/0.2 in a figure-of-eight layout. Profile 8 had eight bunches of 19/0.2 in a flat layout. The cross-sections and therefore DC resistances of the two were about the same and LDPE insulation was used in both cases. Therefore any differences in inductance and capacitance are due to geometry. Profile 8 can be configured in a number of ways. The table shows results with the inner four and outer four conductors used as a pair and also with the left four and right four used as a pair. Compared to standard Qudos, Profile 8 connected using inners and outers shows a significant reduction in inductance and actually slightly less capacitance, which is counter to the rule of thumb often quoted. Conversely, the left/right configuration gives similar inductance to Qudos but with capacitance almost halved. The geometry also affects the characteristic impedance as shown, although this is only of academic interest.

(Note: One of the major benefits of AirCore™ and X-Tube™ technology is that it pretty much halves the inductance of a cable given the same figure 8 geometry with and without the AirCore™. Because the electrical field acts towards the centre of a conductor, if this is replaced by air or a non-conducting material, then the self inductance is automatically reduced by up to half.)

Acoustic crosstalk

One subjective effect often noticed by listeners when certain speaker cables are used is an increase in soundstage width. At first glance it is difficult to see how this could happen considering the high electrical isolation between stereo channels.

We thought one explanation might be that acoustical coupling occurs between left and right channels via the loudspeakers themselves. Ideally, the left speaker should convey only left-channel signals and vice versa. Each channel of the amplifier ideally applies electromagnetic braking to its loudspeaker, preventing it from being moved by acoustic waves from the other speaker. It achieves this braking or damping due to its low source resistance – but in practice the cable resistance intervenes, increasing the source resistance ‘seen’ by the speaker and reducing damping. So, each loudspeaker can vibrate in sympathy with the (delayed) output from the other, with subsequent narrowing of the soundstage. If this were true, low-resistance cables would give a wider stereo image.

Though this might sound far-fetched, the spectral analyses of loudspeaker terminal voltage shown in Figures 15 and 16 show just this effect. The peaks marked with a cross are due to movement of a non-driven loudspeaker’s diaphragm in response to test tones reproduced by another loudspeaker nearby. The 500Hz peak is reduced by about 10dB when the non-driven speaker is connected using a lower-resistance cable (Figure 16).




Figure 15. Spectral analysis of the voltage across the input terminals of a connected but non-driven loudspeaker when another loudspeaker nearby is driven with test tones at 500Hz and 3kHz




Figure 16. Repeat of Figure 15 but with a lower-resistance cable connecting the non-driven loudspeaker to the amplifier

Transient performance

When subjected to a sudden change in signal amplitude, resonances can be excited in a loudspeaker, particularly the fundamental resonance associated with the speaker's bass roll-off. Control of this resonance, and others, is exerted by electrical damping. Extraneous movement of the drive unit voice coil causes it to generate an electrical current which flows through the loudspeaker cable and amplifier's output impedance. As this current opposes the unwanted motion, it should be as large as possible, which in turn means that the cable resistance needs to be as low as practicable.

The two waveforms in Figure 17 show, respectively, the amplifier's output voltage and the voltage at the speaker terminals on just such a transient signal, in this case using a high-resistance cable. At around 2.4ms, when the amplifier output drops to zero, the speaker voltage continues negatively, then rises and 'overshoots' positively before settling. This represents unwanted diaphragm movement. Figure 18 shows the same loudspeaker's behaviour with a lower-resistance cable: the improvement is clear.




Figure 17. Voltage waveforms at either end of a high-resistance cable driving a closed-box speaker with a 400Hz toneburst

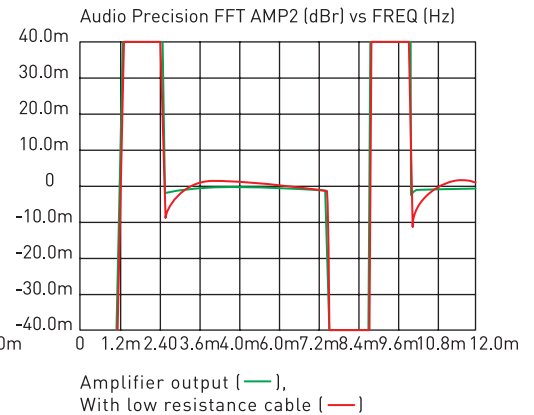


Figure 18. Repeat of Figure 17 but using a low-resistance cable

Cable inductance also increases the impedance between amplifier and loudspeaker and our measurements have shown this to have further harmful effects on transient performance. The loudspeaker's complex mechanical/electrical system works best when the cable has as low an impedance as possible over the whole frequency range – not just at low frequencies where spurious cone movement is controlled by DC resistance.

Cable-Induced Distortion

Loudspeaker cable electrically ‘distances’ the speaker from the amplifier in several ways, affecting frequency response, damping and stereo separation as already described. In addition, distortion at the loudspeaker input terminals is significantly higher (particularly second harmonic) than at the amplifier output.

This effect depends largely on the cable’s DC resistance. Shown in Figures 19 and 20 are plots of second harmonic versus frequency. The top trace in each case shows the distortion at the loudspeaker, while the bottom traces show distortion at the amplifier output. In Figure 19, obtained using high-resistance cable (0.065 ohms per metre), distortion is roughly three times that in Figure 20, obtained with low-resistance cable (0.004 ohms per metre).




Figure 19. Second harmonic distortion versus frequency at either end of a high-resistance speaker cable

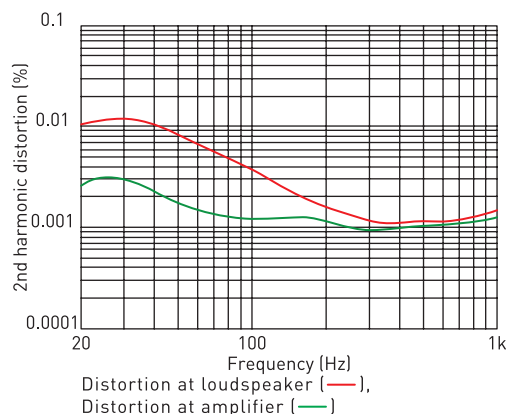


Figure 20. Repeat of Figure 19 but using a low-resistance cable

Figure 21 shows the effect on second harmonic distortion of using different loudspeakers with the same cable. Note that the cable itself does not cause distortion (its DC resistance is essentially linear), rather its presence prevents the amplifier’s negative feedback from accurately correcting distortion generated by nonlinearities within the speaker system. Connecting the amplifier directly to the speaker showed distortion accurately corrected to within a small percentage of the amplifier’s non-loaded condition.

A loudspeaker cable’s influence on distortion at the speaker terminals is not limited to low frequencies. Our tests have shown that mid- and high-frequency distortion figures are significantly worsened by increased cable inductance, which causes cable impedance to increase with rising frequency.




Figure 21. Repeat of Figure 20 but using a different loudspeaker

Multi-strand vs single-core distortion

It has been claimed that multi-strand cables introduce 'diode' effects due to current jumping between strands within the cable and thereby crossing very many metal/oxide/metal junctions. This is sometimes said to be caused or worsened by skin effect as a result of it pushing the current towards the conductor surface at high frequencies.

Making the assumption that the current does 'jump' as suggested, we put signal in via one strand of a multi-strand cable and measured the signal from a different strand. Even using the Audio Precision AP1 test set to its maximum capability, no increase in distortion compared to using all the strands could be found, as shown in Figure 22. This graph overlays curves for the two cases – they are so close that they could as easily be repeat tests of the same measurement. It seems that inter-strand diodes do not exist, or if they do they are shorted out by the many good conductors pressed together over the cable's length.

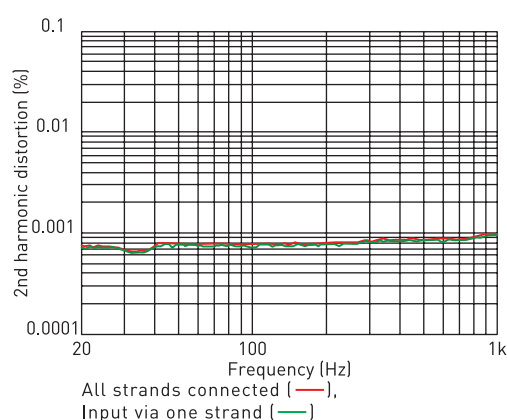


Figure 22. Total harmonic distortion versus frequency for a multi-strand speaker cable measured one strand to another and using all strands (traces overlaid)

Characteristic impedance

This parameter is crucial in high-frequency transmission lines as it determines the source and load impedance that must be used to prevent unwanted electrical reflections and standing waves occurring within the line. To work correctly, a transmission line must be terminated at both ends by a resistive load equal to the characteristic impedance.

Loudspeaker cables are not transmission lines because even the longest are only a fraction of the electrical wavelength of the highest audio frequency. In any case, they cannot be terminated at both ends with the correct impedance. An 8 ohm amplifier output impedance, for instance, would ruin loudspeaker electrical damping while seriously compromising frequency response.

Directionality

Measurements to test for loudspeaker cable asymmetry in the samples all of which were unscreened and some of which were directionally marked by their manufacturer, revealed little to suggest the existence of directionality. Blind listening tests also revealed that listeners were unable to discriminate a cable's direction.

The lay of the cable, on the other hand, was found to have a measurable influence on performance, so to be reliable, any listening or measurement tests would require identical cable positioning for each direction. We were careful not to choose cables that were engineered with screening in this test as some of these are wired asymmetrically

Conclusions

Although there will always be those who remain sceptical about the importance of loudspeaker cables, our research clearly indicates that system performance can be improved or degraded by the loudspeaker cable. Our testing has also established a fair degree of correlation between measured performance and sound quality.

Our findings can be summarised as follows:

1. DC resistance

Low cable resistance is of paramount importance if high sonic performance is to be attained, but it should not be achieved at the expense of other crucial parameters. High cable resistance results in several undesirable consequences: frequency response aberrations, impaired transient response, increased distortion at the loudspeaker terminals and reduced inter-channel separation.

All cables exhibiting high resistance measure badly in these areas and subjectively their performance is highly dependent on the partnering loudspeakers. The forward midrange presented by these cables correlates closely with their effects on frequency response. High cable resistance also reduces dynamic impact with heavily-scored music.

2. Inductance

Cable inductance is a prime cause of high-frequency attenuation and phase shift in loudspeaker cables. Inductance causes cable impedance to rise with frequency, reducing output in the very upper frequency range, sometimes preceded by response peaking. In addition, inductance increases distortion at the loudspeaker terminals and degrades the loudspeaker's overall transient behaviour. Low inductance is required to achieve a flat frequency and phase characteristic, low distortion and good transient response.

3. Skin effect

These are of minor significance in loudspeaker cables of moderate cross-sectional area but become important in cables with larger conductors where, together with greater inductance, they result in greater high-frequency signal loss.

4. Insulation quality

Dielectric dissipation factor has proved to be a very strong indicator of sound quality in our listening tests. Most of the better-sounding cables we have tried use superior dielectric materials whereas PVC-insulated cables give the worst sound quality. Cables which measure badly for dielectric loss appear less able to reveal subtle detail, losing some of the atmosphere revealed by cables with superior dielectrics.

5. Consistency of performance

Speaker cables interact both with the amplifier and the loudspeakers. Consequently, some cables give varied results in different systems. Those which have proved to perform most consistently are those with minimal inductance, capacitance and resistance. Unless an amplifier relies on inductance to maintain stability, keeping the speaker cables as short as practically possible optimises performance. High cable capacitance is best avoided because it can result in amplifier instability, which can degrade sound quality and negatively impact amplifier reliability.

6. Directionality

Despite many manufacturers marking cables directionally, we have found no evidence under controlled conditions that speaker cables are directional. But we have found that merely constructing a cable differently can affect its inductance and capacitance, which may have an impact on sound quality.

7. Solid-core vs stranded cables

Solid-core conductors have been introduced on the basis that, if made thin enough, a solid conductor will show less variation in loss at high and low frequencies than a thicker, stranded conductor. Our research suggests that it is more likely to be the insulation and geometry of many solid-core cables which are responsible for their generally higher performance than stranded conductors. In any case, paralleling

up conductors, whether solid or stranded, reduces inductance, which has a far greater influence than skin effect.

The stranded cables we've tested tend to have higher inductance and dielectric loss than many solid-core counterparts which generally use separately-insulated wires (resulting in lower inductance) and higher-quality dielectrics (resulting in lower leakage losses). We have found no evidence to support the contention that stranded cables suffer from distortion due to diode effects between strands.

8. Metallurgy

Electrical conductivity is slightly superior for cables utilising high-purity copper. Greater improvements to conductivity can be achieved with silver-plated copper or pure silver conductors. Generally we have found that a cable's geometry and dielectric material are more significant than conductor metallurgy in determining its sonic performance.

Conclusion

All our research confirms that the most accurate and consistent-sounding loudspeaker cable will have minimal DC resistance, inductance and capacitance combined with low loss dielectrics. Cables having a small cross-sectional area in an attempt to avoid the skin effect have higher DC resistance, with sonically obvious harmful consequences.

QED's engineers have shown that the 'rule' relating high inductance to low capacitance, and vice-versa, is oversimplified. Capacitance and dielectric losses can be reduced by using high-quality insulation material (low-density polyethylene) and by minimising insulation wall thickness and designing narrow webs (consistent with mechanical integrity), thereby increasing the ratio of air to solid dielectric. By optimally orientating multiple parallel stranded conductors around a central non-conductive core, e.g. Aircore™ and X-Tube™ Technology, QED has been able to reduce both inductance and capacitance below that of a single conductor pair of the same DC resistance. Use of stranded conductors of sufficient cross-section keeps DC resistance low.

The result is a range of low-loss, transparent-sounding loudspeaker cables of superior performance. The correlation between insulation and sound quality has also influenced the design of QED's interconnect cables, where the use of foamed LDPE insulation to increase the air/solid ratio in the dielectric helps maximise sound quality.

Biwiring: an exploration of the benefits

Biwireable loudspeakers have been available since at least the late 1980s. (You can easily tell if a speaker is biwireable because it will have four connection terminals on the back rather than two.) The idea is that instead of using just one run of cable to each speaker, a separate cable is used for each of the two pairs of terminals, once you have removed the shorting wires or plates that normally connect them. This means that you will use twice as much speaker cable than if you used the normal single-wire method, which has prompted cynics to use the phrase 'buy wire' to mock the idea. If you have not biwired speakers before, you might be tempted to try because of what you have read in the hi-fi press or online – but you may also wonder if the extra outlay is well spent.

Proponents of biwiring point to the obvious sonic benefits they hear and cite the fact that speaker manufacturers fit the extra terminals as proof that there must be something in it. Detractors argue that manufacturers are merely maximising the marketability of their products and point out that there is little published evidence to prove that biwiring makes any audible difference. Biwire enthusiasts, meanwhile, theorise that separating the low- and high-frequency signal currents can eliminate distortions caused by interactions between them.

Biwiring, then, is a controversial subject. Is it a worthwhile sonic upgrade or just a clever marketing ploy by the cable companies to encourage you to buy twice as much cable? We decided to do some research to see if an authoritative answer can be given to this question.

How are normal loudspeakers connected to an amplifier?

In Figure 23 a single speaker cable is shown connecting the output terminals of one channel of an amplifier to a loudspeaker. The speaker is a non-biwireable two-way type, having a woofer to reproduce low frequencies (LF) and a tweeter to reproduce high frequencies (HF). A passive electrical circuit (represented by the two boxes labelled 'HF network' and 'LF network') is used to filter the signal from the amplifier so that only the low frequencies are passed to the woofer and only the high frequencies to the tweeter.




Figure 23. Traditional single-wire method of connecting an amplifier to a loudspeaker

The filter networks cannot create sharp cut-off points at the chosen 'crossover' frequency but instead apply more gradual attenuation with considerable overlap. In Figure 24 you can see the point where the two response curves cross (green for the woofer, red for the tweeter) is at about 2kHz, at which frequency the two drivers contribute equally to the acoustic output. If the LF and HF networks are correctly designed then the combined output of the two drivers (black curve) remains essentially flat through the crossover region.




Figure 24. Frequency response curves for a two-way loudspeaker

How are biwired speakers connected to an amplifier?

Figure 25 shows how a typical biwired speaker is connected to an amplifier. It has four terminals instead of two, one pair for the HF network and one pair for the LF network. Two speaker cables are used to connect the HF terminals and the LF terminals separately to the amplifier output terminals.

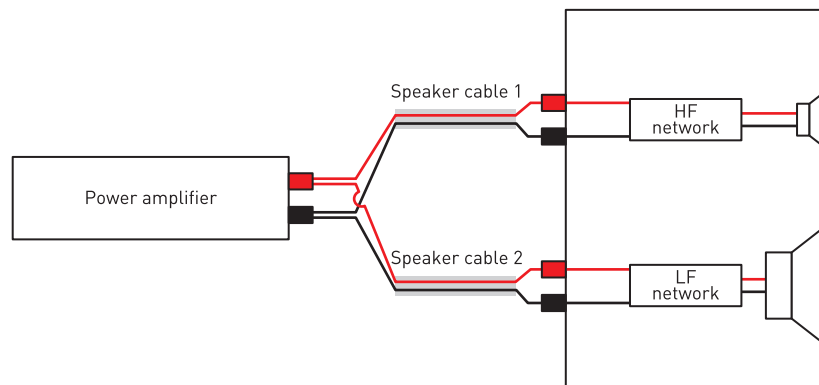


Figure 25. Biwired connection of an amplifier to a loudspeaker

How does signal current flow in single-wired and biwired speakers?

With a single-wire connection, signal current for both the drive units flows along the same cable. With a biwire connection, the low frequency and high frequency signal currents are separated. At high frequencies the impedance of the LF network is high and so the HF current is low; at low frequencies the impedance of the HF network is high and so the LF current is low. It's not a matter of the low and high frequencies 'knowing which way to go' – it is just a natural consequence of the biwire connection.

What are the benefits of separating low- and high-frequency signal currents?

Wherever there is nonlinearity in a system – and loudspeaker drive units are nonlinear – there will be intermodulation distortion introduced in the form of sum and difference frequencies. Unlike some types of harmonic distortion, intermodulation distortion is by nature audibly objectionable. If separating low and high frequency signal current by biwiring reduces the loudspeaker's intermodulation distortion then we can expect it to have an audibly beneficial effect.

If intermodulation distortion is reduced by biwiring, how can we measure it?

Now we are close to proving that biwiring has some genuine sonic benefit. We have accepted that the two cables carry different signal currents and that by keeping LF and HF currents separate we may have prevented the generation of intermodulation distortion. If it was possible to measure intermodulation distortion in a single-wired speaker connection and to demonstrate a reduction of that distortion in the same speaker when biwired, we could make a really good case for the audible benefit of biwiring being real.

Some work on this subject has been published by Jon Risch¹. He used a split-band test signal comprising 10 tones divided into two frequency bands, one low and one high, starting at 100Hz and 5kHz respectively. The frequencies of the tones were carefully chosen so that any intermodulation components which are added are minimally masked by harmonic distortion. Using the same methodology, we generated a similar test signal and conducted our own investigation to see if we could measure any differences between single-wired and biwired speaker connections.

First we created the multitone test signal and burned it to a CD. Figure 26 shows a spectral analysis of the CD player's output when playing the test CD. The high frequency tones can be seen as five distinct peaks in the graph from 5kHz to about 9kHz at around -10dB; the low frequency tones appear as five clear peaks from 100Hz to a little less than 200Hz, also at around -10dB. Note the low level of noise and distortion between the two bands.




Figure 26. Spectral analysis of the split-band multitone test signal, measured at the CD player output

Then we used a current probe to measure the spectra of the signal currents in a single-wired connection and in each cable of a biwired connection (woofer cable and tweeter cable). Results for the single-wire connection (blue trace) are shown in Figure 27, overlaid on the black trace of the test signal from Figure 26. Intermodulation products have significantly raised the distortion level, at some frequencies to only 40B below the level of the test tones.




Figure 27. Spectral analysis of the output of a current probe measuring cable current in the single-wire connection (blue trace), overlaid on the trace from Figure 26

Compare this with the measurement taken from the tweeter cable in the biwired connection, tested under the same conditions, shown as the red trace in Figure 28. First, note that the LF tones have been attenuated by over 30dB which proves that low- and high-frequency currents have indeed been separated in the biwired connection. Second, see how intermodulation distortion has been significantly reduced below the upper set of tones but still within the passband of the tweeter.




Figure 28. Spectral analysis of the output of a current probe measuring cable current in the tweeter cable of a biwire connection (red trace), overlaid on the single-wire trace from Figure 27

Now let's look at the spectrum for the woofer cable (green trace in Figure 29). The high frequency components have been attenuated this time, again proving that the LF and HF signal currents are separated in a biwire connection. There has been a less dramatic reduction in the level of the intermodulation products but there is still some improvement.

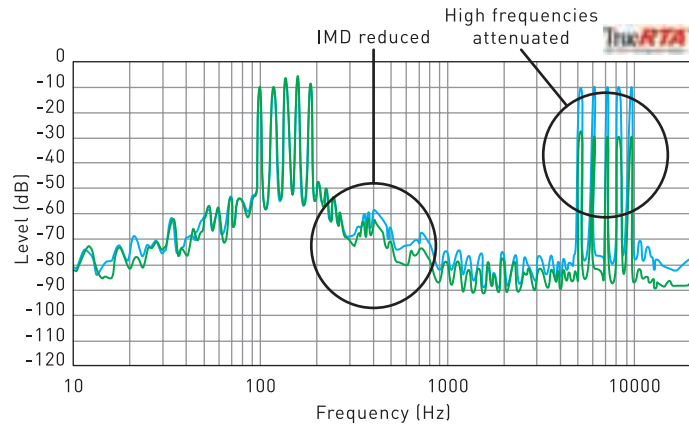


Figure 29. Spectral analysis of the output of a current probe measuring cable current in the woofer cable of a biwire connection (green trace), overlaid on the single-wire trace from Figure 28

Conclusions

If your speakers have four binding posts then you can biwire them but you will need twice as much cable. Low-frequency and high-frequency signal currents will then travel in separate cables. Our measurements show that this reduces intermodulation distortion caused by nonlinearity in the speaker system.

In light of this evidence it's sensible to conclude that where the opportunity exists and funds allow, biwiring should be explored as a means of improving the performance of any suitable high fidelity loudspeaker system.

Reference

¹Risch, Jon M, "A New Class of In-band Multitone Test Signals", Paper 4803, Audio Engineering Society 105th Convention, September 1998

Notes

The test signal CD was played using an all-in-one CD player/amplifier set at half volume, into a pair of floorstanding speakers. Current probe measurements were recorded using a Tascam US144 USB audio interface and analysed using TrueRTA. The single-wire measurements were taken with two runs of speaker cable connected in parallel (speaker terminal shorting connectors in place); the biwired measurements used the same arrangement but with the shorting connectors removed.

HDMI cables

Some writers in the general press and on the internet suggest that it's not worth upgrading to high-quality digital audio and HDMI cables. As a leading manufacturer of cables and accessories for more than 40 years, we have a wealth of experience in this area and, at our cable assessment and design facility, we own some of the world's most advanced HDMI digital measurement equipment. As well as using this to develop QED cables, we also use it to evaluate other HDMI cables on the market. So here's our (well-informed!) contribution to this debate.

We agree: a lot of HDMI cables are the same!

We are shocked by some of the cables we measure, particularly when premium cables prove to offer little or no measurable improvement over lesser models from the same manufacturer. On these occasions we find ourselves in agreement with our more sceptical journalist friends. Just because a cable looks flashier or is priced higher than another cable, it is not necessarily better. It may just be standard cordage disguised with some 'bling'.

Misconceptions

There is a common misconception that digital data transfer in HDMI is perfect, that noise immunity and error correction work together to ensure that data received by the sink equipment (the receiver) and subsequently shown as a picture on the TV screen is always the best it can be.

In fact there is no such error correction in an HDMI link. Rather, errors that creep in due to logic level transitions in the physical layer (the cable) are minimised by encoding the data stream so as to minimise those transitions. At the sink (TV) a simple test is used to verify the stream and it is then converted into a viewable picture, complete with any cable-induced timing errors.

These timing errors, known as jitter, can affect the perceived quality of the picture so it is important to keep them to a minimum. This requires that the cable introduces as little 'skew and slew' (clock skew and slew rate limitation) into the signal as possible, which can only be achieved by the use of carefully controlled cable geometry and top-quality materials.

Just ones and zeroes? Yes, but there is more to it than that

"Digital signals are either on or off and immune to noise – so as long as there is a signal the other end of an HDMI cable that can still be read by the TV then the received signal is perfect, isn't it? How can the type of cable make any difference to the received signal – expensive HDMI cables are a waste of money, right?"

You'll often see this argument deployed but it relies on the erroneous assumption that any errors or noise introduced by the transmission of digital signals along a cable are masked by the natural noise immunity of the signal's discrete logic levels or repaired by error correction algorithms in the receiving equipment. The true situation turns out to be much more of a grey area.

What is inside an HDMI cable and what signals does it have to carry?

HDMI stands for High Definition Multimedia Interface. It is a way of conveying encrypted digital video and audio signals from a source (eg a DVD or BD player) to a sink (a TV or AV amplifier) and is capable of bandwidths that far exceed that of standard definition TV signals. The High Speed standard, which is the maximum bandwidth described in the HDMI 1.4 specification, uses 10.2 Gigabits per second transmission for a deep colour picture resolution of 48-bit 1080p60 (1080-line progressive scan at 60 frames per second). This requires the HDMI cable to convey

10.2 billion discrete logic levels every second in order to meet the High Speed HDMI standard.




Figure 30. Internal construction of an HDMI cable in cross-section

An HDMI cable incorporates four twisted pair cables (two insulated wires twisted into a helix) to carry four data channels used for Transmission Minimised Differential Signalling. Figure 30 shows the internal construction of a typical HDMI cable with the TMDS lines (label 4 and its equivalents) and several other conductors that are used for communication between the source and sink, transmission of remote control signals and video encryption in the form of High Definition Content Protection, all of which work at much lower data rates than the TMDS lines. There are also conductors for power and ground.

TMDS

There is a limit to how quickly the signal can be switched from logic level 1 at 0.6V to logic level 0 at -0.6V due to cable capacitance. The twisted pair conductors act like the two plates of a capacitor and when the signal changes from a low to a high voltage or vice-versa this capacitance must be charged or discharged before the required voltage level can be attained. This takes a finite time and has the effect of rounding off the square edges of a perfect digital signal. If the data rate is too fast, the cable can still be charging up from a zero to a one when it's time to discharge from a one to a zero again, and the receiving equipment may not see any logic transition at all.

In order to prevent this, the serial data stream that makes up the HDMI signal is split into three parallel signals that are sent down three separate twisted pairs within the HDMI cable. This reduces the maximum data rate each conductor must be able to carry to a third, so the bandwidth requirement is reduced to 3.4 Gb/s for a High Speed cable. These twisted pairs are known as TMDS lines. TMDS stands for Transition Minimised Differential Signalling, the mechanism of which is explained in more detail in the next section. Because it is the logic level transitions that cause errors in digital signalling, these are minimised using a simple algorithm. The 'differential' part of the interface description comes from the fact that the transmission line is made up of a twisted pair and separate ground wire. This hot, cold and screen geometry is used in analogue signalling too, where it is known as a balanced line and used to cancel the effect of external electrical interference.

Figure 31 shows the signal that appears at the output of the HDMI source equipment on each of the three TMDS lines. This signal is assumed to be perfect and is, as yet, unaffected by the electrical characteristics of the HDMI cable.

Known as an eye diagram, this measurement is used to test the HDMI cable for compatibility with the HDMI standard. It is a composite density map of a million samples taken over a few seconds of a random bit pattern running at the 3.4 Gb/s data rate defined by HDMI.org. There are three such data lines which add up to 10.2 Gb/s altogether – the maximum data rate for HDMI and that defined as “High Speed” by the HDMI standard. The width of the graticule (or window) represents a time period of just 571ps (571 millionths of a millionth of a second), enough to encompass two logic level transitions or three bits of data. All the possible transition combinations are present: signal sections which remain at logic one across the whole graticule; signal sections that stay at logic zero across the whole graticule; and signal sections where one or more high/low or low/high transitions occur.




Figure 31. Eye pattern of a High Speed HDMI output signal

The colours in the diagram represent the density of the sample points, with blue representing only a few samples and red the most. It can be seen that the signal is very rarely caught in the transition between zero and one, with most of the samples being found in the discrete levels of either logic one or logic zero. The purple area top and bottom and in the middle are the “don’t know” areas where the receiving equipment would not be able to interpret the incoming signal as a one or a zero. If any of the samples fall into this area then there is an error in the data.

In addition to the three data lines within an HDMI cable there is a fourth similar shielded twisted pair which carries the master clock signal. In HDMI the data is split up into chunks of 10 bits and because of this the master clock only has to run at a tenth of the data rate, i.e. a maximum of 340MHz in a High Speed HDMI cable. This helps to keep the clock accurate because the lower its frequency, the less its waveform is affected by cable parameters.

In order for the sink device to read the digital data from all three of the data channels, it must be told when to look at the incoming signal. This is the purpose of the master clock but because it runs at one-tenth of the data rate, the TV must multiply the clock frequency by 10 and then realign it with the data window in each of the three TMDS data lines. This reproduced bit rate clock does not have a guaranteed phase relationship with any of the three data lines. In addition the

HDMI specification allows a certain amount of skew between any two data lines. Because of this the clock phase must be adjusted individually for each data line to sample the incoming serial bits correctly. This involves aligning the clock rising edge to the middle sampling window of each data line. This is done by delaying the data by varying amounts based on the incoming data stream and how much it is skewed or jittered from the ideal $\times 10$ data clock that has been recovered. If the jitter is too large then errors can creep through.

Display Data Channel

The Display Data Channel or DDC (Figure 30) uses two single conductors within the cable. It forms an I²C communications bus for the Enhanced Extended Display Identification Data (EDID) protocol which is a way that the source and the sink can communicate with each other and negotiate the data rate for the video signal. For instance, a DVD player set to output 1080p will announce this to the TV which will tell the DVD player to go ahead and set itself to receive 1080p.

Consumer Electronics Control

The Consumer Electronics Control or CEC (Figure 30) uses a single conductor and common ground within the cable. CEC allows the user to control up to 10 different devices over the HDMI link using only one remote control. For instance, it can be used to control the volume of an AV amplifier using the TV remote handset codes.

How is the video signal transmitted by HDMI and how are errors dealt with?

The TMDS signal transmission method used by HDMI is the same as that found in the DVI standard, with which it is backwards compatible. The idea is that errors are minimised, rather than corrected, by encoding the signal in such a way that transitions between logic high and low are kept to a minimum. This is achieved by performing a simple XNOR or XOR operation on each successive bit of the pixel data and its predecessor. Two control bits are then added to each pixel data byte to form a 10-bit video character. One of these extra bits is added to indicate which of the two logic functions was used. Additionally, if the previous character contained a lot of ones and the next character will also, the system can invert that character to maintain a mean zero DC offset in the signal, so a second control bit is added to the signal to indicate whether the character has been inverted or not. In this way the 10-bit pixel data character will contain a maximum of five transitions. Each pixel data character is bookended by a control data or blanking character of 10 bits which is used to indicate the boundaries of the pixel data character. There are three such channels in the HDMI signal. All this is clocked out on a fourth differential transmission line at between 74.25 – 340MHz depending on the video standard, with one clock pulse per data character so that the clock rate is one tenth of the bit rate.

When the signal is decoded the other end there will be no picture until the receiver has synchronised with all three video data streams and this is achieved by detecting the blanking periods and synchronising them with the one-tenth-bit-rate clock signal. The blanking signals are only distinguishable from the video data because the number of transitions in the 10 bit character is higher (therefore once the video signal has been corrupted enough that the pixel data is indistinguishable from the blanking data synchronisation is lost and the video will cease or not even start). If the incoming signals can be synchronised, the data is then decoded according to the settings of the two control bits to restore the correct number of data transitions.

Audio transmission

There is a third type of character in the HDMI firmament and this is known as the Data Island Period and it is here that the audio signal resides. The audio data is encoded using a variation of TMDS known as TERC4 or TMDS Error Reduction Coding and this data is also encoded using the BCH error coding method described in the next section.

Error detection

There is no description in the HDMI 1.4 specification of any form of error detection or correction process for the video part of the signal. Both the clock and data are recovered from the incoming signal based on the assumption that the bit error rate is less than 10^{-9} , i.e. one error in a billion bits (or four errors per second in a 1080p60 signal). This assumption is made because the source electronics, cable and sink (receiver) must all conform to the eye mask test. The audio data, by contrast, is subject to error detection and correction, in a similar process to CD.

There is time to implement error correction for the audio because the audio data rate is very much slower than that of the video data rate. Audio information is not transmitted in a linear way but using a process known as the Cross-Interleave Reed Solomon Code or CIRC. The data is sent in a different order to that in which it is played, and is re-sorted after being retrieved. If a poor cable link causes data to be corrupted then it can cause a long period of data corruption called a burst error. As the audio data is not sent linearly, when it is reassembled the burst error is spread among many frames and can thus be more easily dealt with by the parity checking error correction system. In parity checking an extra bit is added to each data word to indicate whether all the bits added together give an odd or even result. When these data words are assembled together in a grid-like fashion the parity bits on horizontal and vertical axes indicate the precise position of any bit errors, which can then be toggled to correct them. In this way CIRC can correct errors of up to 3500 bits in length and, if correction isn't possible, can conceal by interpolation transmission errors of up to 12,000 bits in length.

The upshot of all this is that a cable can affect a video signal by introducing errors, and affect both audio and video by introducing timing errors. Not only can the cable add jitter to affect the timing of the three TMDS data streams, but at or near the eye pattern mask bit errors could increase to visible levels before the synchronisation is lost, and these will be displayed uncorrected as 'spangly' pixels. In practice most properly designed and certified cables are working at speeds far below those that would take them near bumping into the mask but a 7m cheap cable transmitting 1080p60 might struggle to convey a distinguishable blanking period and bit error rate could climb well before this happened. Even with faster cables the subtle differences in timing and skew between the three TMDS lines and the clock will all be displayed uncorrected. It is even possible that this may be perceived by some viewers as a drop in picture quality.

The expectation that there will be digital error correction on the HDMI transmission line is flawed because by the time it is ready for transmission to the sink, the signal is assumed to be perfect. Rather, error correction takes place when digital data is read from the DVD or BD. This is just as it is with CD: once decoded, the signal may be transmitted via S/PDIF interface to an external DAC. There is no error correction in this transmission system, and neither is there one in the HDMI system. The only exception to this is the error correction applied to the audio part of the HDMI signal only. This will correct bit errors caused in transmission but cannot deal with timing errors caused by jitter.

Many TVs have post-transmission error detection of some description which can detect a spangly pixel and 'correct' it. This is done by interpolation, the incorrectly coloured pixel being detected and changed to match its near neighbours. Of course, this type of error correction can be wrongly as well as correctly applied, adding noise to the picture. The algorithms work far better on a still picture than a fast-moving one, so the correction can be seen operating when a still frame immediately follows a fast-moving sequence. For this reason most video fanatics leave post-transmission error correction switched off.

What is jitter and why does it affect picture quality?

Figure 31 was taken from a signal generator running at the High Speed data rate for HDMI of 10.2Gb/s. The eye pattern demonstrates negligible jitter as the rise and fall times of the signal are fast and consistent. Compare this with Figure 32, taken at the far end of a cheap 1m cable conveying the same signal.




Figure 32. Eye pattern of a marginal pass HDMI cable

You can see a significant increase in the spread of the data points measured by the eye pattern analyser in both level and time. The absolute voltage level of the signal varies by up to 350mV but more significantly the zero crossing point of the signal varies in time by up to 125ps.

These variations in the level and in the timing are known as jitter, and there are various causes of it within the cable. The first is the impedance of the cable and its complex interaction with the signal. The cable has capacitance and for the voltage level to swing from logic low to logic high, or vice-versa, this capacitance must first be charged up or discharged and this takes finite time. If the time allowed for the logic level transition is short compared to the time the cable takes to charge up then rounding of the square wave takes place. You can clearly see this in Figure 32. If the capacitance is sufficiently high the logic level may not reach its intended state before the next transition and so a "don't know" logic level is transmitted.

The purple areas of the eye pattern denote voltage levels which are perceived by the receiving equipment (usually a TV) as a "don't know" logic state. As long as none of the sample points falls inside these purple areas the cable is a pass for HDMI and can bear the High Speed logo. As long as the TV receives its clock signal exactly in the middle of the timing window the logic level it reads will be firmly a 1 or a 0, even though there is significant spread in the actual analogue level it sees. This is the beauty of digital encoding – the receiving equipment is able to distinguish the data even in the presence of noise as there are only two discrete logic levels that it needs to detect. In HDMI, a voltage is decoded as binary 1 as long as it lies between 70mV and 780mV. Unfortunately, as noted earlier, the

master clock must be recovered from the data and is only corrected every tenth bit. A data signal displaying 120ps of jitter will cause a corresponding variance in the recovered clock. This means that it is possible that the TV will get its clock signal up to 60ps early or late which is enough to make it see data that is at a “don’t know” level and thus cause a bit error. The HDMI spec says this is OK as long as the bit error rate is below 10^{-9} , or up to four bit errors per second in a Full HD (1080p60) picture. Noticeable pixel errors will start to creep in before sync and eventually the entire picture, is lost.

Timing errors are caused by this mechanism too and the signal can also be affected by crosstalk from adjacent signal carriers. Within the HDMI cable there are three differential signalling lines plus a master clock signal, so there is plenty of scope for this kind of crosstalk to take place. Finally, because the transmission lines utilise differential signalling there can be slight differences between the two conductors in each twisted pair that introduce timing errors known as intra-pair skew. This can be caused because the two conductors are slightly different lengths or because the dielectric (insulation) is inconsistent or the twist rate varies. All these geometrical and electrical parameters must be carefully controlled within the cable if jitter is to be minimised. Some manufacturers have recognized this problem and incorporate upsampling and local master clocks in their equipment to try to reduce or eliminate jitter, but there is inconsistency between manufacturers so it is important to try to preserve the signal integrity along the cable as a first line of defence, before any post-transmission recovery is added or required.

Jitter can also affect picture quality perception way before any bad pixels start to appear, a phenomenon investigated in ‘The Effects of Jitter on the Perceptual Quality of Video’, Claypool & Tanner, Worcester Polytechnic Institute USA 1999. In this paper the authors report “we found that jitter degrades perceptual (video) quality... and that perceptual quality degrades sharply even with low levels of jitter as compared to perceptual quality for perfect video.”

What does this degraded perceptual quality look like? Here is an extract from some research carried out at QED HDMI Labs in 2012: “In an effort to try to characterise this problem we artificially jittered a standard colour bar video signal until we could see a change in the video output from a TV. It was significant that errors in the picture were visible even though the HDMI link was still up and the cable was able to maintain a steady picture.”

The two eye diagrams in Figure 32 show the amount of timing error we were able to add to the cable before problems became apparent. The trace on the left




Figure 33. Eye patterns for an HDMI cable before (left) and after the addition of artificial jitter

shows that the cable itself added 81ps of jitter to the signal; the trace on the right shows that only 44ps of jitter needed to be added (125ps overall) to cause visible degradation in the video signal due to bit errors. A standard colour bar signal was sent to the TV and the transition between yellow and blue was seen to suffer

from errors, with pixels turning black. This can be seen in Figure 34, in the area encircled in red.

The evidence shows that timing errors introduced by a cable do find their way into the video signal and are display by a TV before it stops working altogether,




Figure 34. Picture errors introduced by jitter (black pixels in the encircled area should be yellow)

and that the jitter need only exceed about 125ps pk-pk for errors to be visible. This level of jitter has been measured by QED in several 1m cables claiming to be High Speed. Given this information, it is sensible to use a cable that introduces as little timing error as possible, to ensure that you are getting the best possible performance from your equipment. In practice this means choosing the shortest and fastest cable possible, even if you are only going to use it at half its theoretical maximum data speed.

Audio

Many reviewers cite an improvement in sound quality as being a differentiator between HDMI cables and our research does strongly bear that out. The audio information in the HDMI link is carried on the Data Island character and is subject to the same jitter problems introduced by the cable as the video data. We know that error correction eliminates bit errors as an explanation of improved sound quality; this just leaves jitter, which as we know can be effectively controlled by careful choice of cable. It does seem that many reviewers find it easier to hear a difference in the audio than see it in the video.

In order for a cable to achieve low jitter and high data transmission speeds it must exhibit low overall impedance and keep intra- and inter-pair skew to an absolute minimum. Cheap cables which are constructed with inferior materials might exhibit a pass at HDMI High or Standard Speed but as we have seen they will introduce enough jitter and bit errors to spoil perceived video quality. Cables that are manufactured using high-quality dielectrics and carefully controlled twist rates and geometries, and which are tested during development and manufacture to ensure that these parameters are consistently adhered to, will improve the jitter performance of the HDMI link and thereby improve perceived video quality.

Conclusions

Despite the inherent noise immunity of digital signals, an HDMI cable does not guarantee perfect transfer of the signal, a fact accepted by the HDMI organisation which specifies the allowed error rate.

In addition to reducing errors, a good HDMI cable introduces less jitter into the signal, and independent research concludes that even low levels of jitter can lead to a sharp decrease in perceived picture quality. To remove any doubt about this, we have shown that adding jitter to a cable induces visible picture errors.




Figure 35. Eye pattern of QED Reference HDMI cable which typically displays jitter of less than 25ps pk-pk

At QED we carefully control the materials and geometry used in the manufacture of our HDMI cables, ensuring consistent performance. And we can measure the difference in how our cables perform, as illustrated by comparing Figure 35 – the eye pattern for the QED Reference HDMI cable – to Figure 32, which is typical of a marginal pass cable. Our HDMI cable hierarchy is differentiated by a scientifically measurable improvement in jitter performance as you go up the range, from Profile HDMI with low jitter to Reference HDMI with ultra-low jitter. By investing in one of our premium cables you are getting an objectively measurable improvement over our 'entry level' cables, and that's a FACT. Proprietary technologies that enable the transmission of High Speed HDMI over longer distances – up to and including 15m – mean that, unlike those from many other manufacturers, long-length QED cables work with 3D and 4K signals.

Even if you are not convinced by the weight of picture-related evidence in support of good HDMI cable, don't forget that many users can more readily hear the improvement to audio quality than see the improvement to video quality.

Future-proofing is also important when discussing HDMI. If newer technologies come along which require higher transmission speeds than the current High Speed standard, many HDMI cables will fail to convey that signal properly. The QED Reference 1m HDMI cable uses unique technologies that allow it to convey data at speeds up to double the High Speed standard. So buyers can rest assured that if they purchase a QED Reference HDMI, with our Lifetime Guarantee, that it will, well ... last a lifetime.

Digital audio optical cables

Audio source equipment such as TVs and Blu-ray players are often connected to amplifiers or receivers using optical cables that carry digital audio. An optical link has the advantage over an electrical connection that it provides complete electrical isolation between equipment, eliminating ground loops and noise induced by large external voltages or currents. A high-end audio system is very easily degraded by connection to a source with a noisy ground.

For digital audio applications the interface adopted by the industry is that developed in 1983 by Toshiba and commonly called TOSLINK™ (short for TOSHIBA LINK). This uses a proprietary snap fit connection system and low power LED signalling using a red light source in the region of 650nm wavelength. Binary on/off signals can be output by the LED flashing at the appropriate intervals and the signal transmitted in this way usually conforms to the S/PDIF (Sony/Philips Digital Interface Format) standard. The system is only suitable for relatively low speed signalling over short distances of up to 10m or so.

Optical connection exploits the light guide principle. A thin transparent fibre is given a cladding material of lower refractive index so that light inside the fibre, incident at the fibre boundary below a certain critical angle, is reflected back into the fibre and thereby guided along it – even if the fibre has bends in it.




Figure 36. Total internal reflection within an optical cable (red lines). If a ray meets the fibre wall at too large an angle, it escapes the fibre (blue line)

Figure 36 illustrates this. Total internal reflection confines light to within the optical fibre, in a manner similar to looking down a mirror made in the shape of a long tube. Because the cladding has a lower refractive index, light rays reflect back into the core if they encounter it at a shallow enough angle (red lines). Any ray that arrives at the boundary at too high an angle escapes from the fibre (blue line) and this is responsible for losses along the cable.

Figure 36 also illustrates some of the major problems with conventional TOSLINK™ cables. The low power red light and large optical aperture used in the TOSLINK™ system means that the only type of cable that can be used is the multi-mode cable type. The reflected rays (red lines) can take paths of different lengths along the fibre depending on the original angle of incidence. As the numerical aperture of the connector allows for a large range of incident angles, light rays from the same pulse can exit the far end of the cable at different times. These different groups of light rays are known as “modes”. The output signal pulse, being an aggregate of different modes, begins to spread out, losing its well-defined shape. The need to leave spacing between pulses to prevent overlapping limits the bandwidth of the connection and therefore the rate at which information can be sent.

Because of the short connection distances and relatively slow data signalling speeds of TOSLINK™ it is possible to use very inexpensive plastic optical fibres (POF) made of an acrylic glass material known as PMMA (polymethyl methacrylate). These are usually 1mm in diameter to fit the TOSLINK™ connector and this large width only exacerbates the time smearing or jitter caused by the multiple modes within the fibre. Calculations show that a typical 1m optical cable with a PMMA fibre of refractive index 1.5 and a critical angle of 76 degrees has the potential to introduce jitter of up to 145ps into the S/PDIF signal. These plastic optical fibres (POF) account for virtually all commercially available digital audio optical cables available today.

The alternative adopted by QED in our Glasscore™ technology cable (Figure 37) is to employ multiple fibre bundles of ultra-fine borosilicate glass optical fibres (GOF) of no more than 50µm diameter each – thinner than a human hair – to make up the 1mm diameter necessary to fit the TOSLINK™ standard connectors. In the QED Reference Optical Quartz cable there are 210 such fibres, each of which guides the light along a much tighter path, so reducing the time disparity between different modes compared to a larger diameter single fibre. In this way data pulses are no longer elongated or smeared by the physical process of travelling along the cable.




Figure 37. QED Glasscore™ employs multiple 50µm borosilicate fibres

Reference Optical Quartz has a bandwidth of over 150MHz, greater than 10 times that of single acrylic glass fibres, and vastly exceeds the requirements for high resolution multichannel audio. In contrast to conventional optical cables, the bandwidth is completely unaffected by bending the cable.



WWW.QED.CO.UK