An opportunity to examine the circuit diagram *and to listen* to Peter Walker's new loudspeaker was too good to miss . . .

QUAD ESL 63

It has never been our policy to publish review articles on audio products — and we are not going to start now. This does not mean that we, as critically-eared engineers, do not have our opinions on the standard of engineering demonstrated in any particular product. Therefore when the manufacturer provided us with a copy of the circuit diagram and an *opportunity to listen* to a Quad ESL 63 we did not put up any resistance . . .

Having said that, we believe that the following article, about the rather intriguing innards of this unusual loudspeaker, will interest many of our readers. After a more general introduction on electrostatic loudspeakers, to set the stage, we shall go into more detail on the ESL 63. Finally we shall discuss the way the ESL 63 behaves as a doublet and how it interfaces with listening rooms.

The electrostatic loudspeaker

An electrostatic loudspeaker is in several ways the counterpart of the ubiquitous moving-coil loudspeaker. Firstly in the strictly theoretical sense: the force exerted on an ESL diaphragm is proportional to the applied *voltage* rather than to the *current*, whereas it is the *current flowing in the ESL motional impedance* rather than the *voltage across it* that is proportional to the net diaphragm velocity. Another aspect is that an ESL diaphragm, carrying only a more or less static charge (where else could the name come from?), can be very large in size and yet light — much lighter in fact than the mass of air that it moves. Figure 1 shows a cross-section of a modern electrostatic drive-unit. The diaphragm is a very thin flexible thermoplastic film, stretched tightly to provide a restoring force directed towards the central equilibrium position. It must be possible to charge the diaphragm reliably without the individual charges being able to distribute themselves unevenly during excursions — otherwise the force, which of course acts on the charges would be unevenly distributed as well. This calls for an extremely high surface resistivity. The 3-micron polyester diaphragm in the ESL 63 uses electronic conduction, obtained by doping it with one donor atom to about ten million non-conductive atoms. (This is incidentally something of a technological breakthrough and it was in fact Quad's toughest development problem).

The fixed plates, perforated to allow the passage of air, actually form a parallel-plate capacitor. The AC voltage applied across this capacitor sets up the signal field and the driving force is simply the total charge times this field strength. When the diaphragm moves it will cause a charge-displacement in the plate-circuit and we see that the rate of diaphragm movement (velocity!) must be proportional to the rate of charge-displacement (current). In practice the perforated plates are also of a thermoplastic material, with a conductive layer printed on the outside; this greatly simplifies high-voltage insulation problems and in the ESL 63 also enables patterned electrodes to be used. (See figure 2.)

One more interesting aspect is that of the damping of the fundamental resonance. In an electrostatic drive unit this is the resonance between the diaphragm compliance (due to the earlier-mentioned restoring force) and the mass of air in the immediate vicinity (plus of course the very small mass of the diaphragm itself). Musically oriented readers may note that this *air-mass-load* ('inertance' in the acoustical scheme of things) is the same mechanism as that requiring end-correction in the tuning of an organ pipe. Electrical damping of this resonance can be achieved by regulating

1

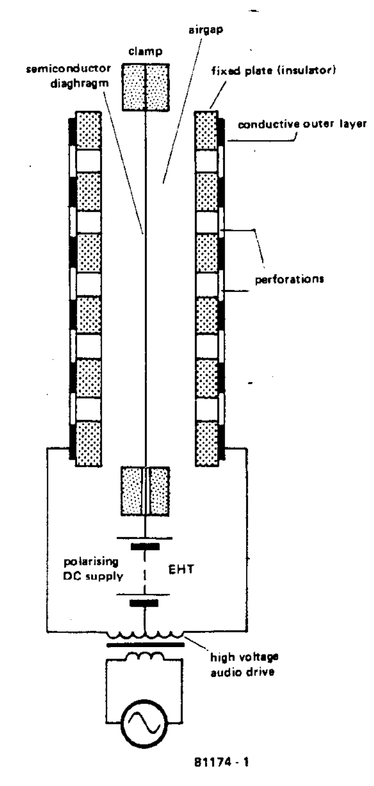

Figure 1. Cross section of a typical electrostatic drive unit with external circuit. The apparent simplicity is misleading.

the motional current (in the moving-coil case, the voltage). In the ESL 63 there is also internal acoustical damping by means of an airflow resistance.

Cabinets

Loudspeaker cabinets as we know them exist mainly to achieve useful bass output from the moving-coil driver, with its

2

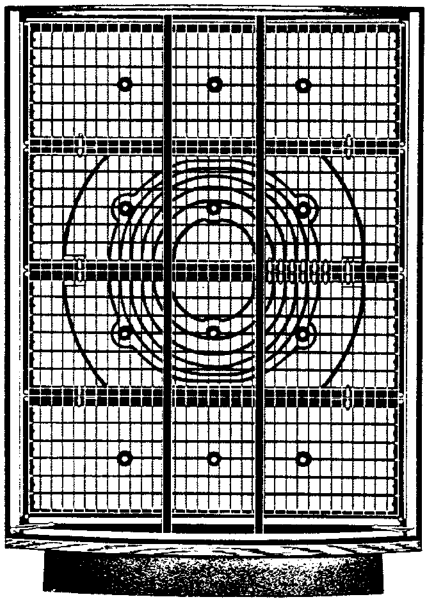

81174 - 2

Figure 2. An illustration of one of the ESL 63 fixed plates, showing the separations as dark rings between the various driving sections.

3

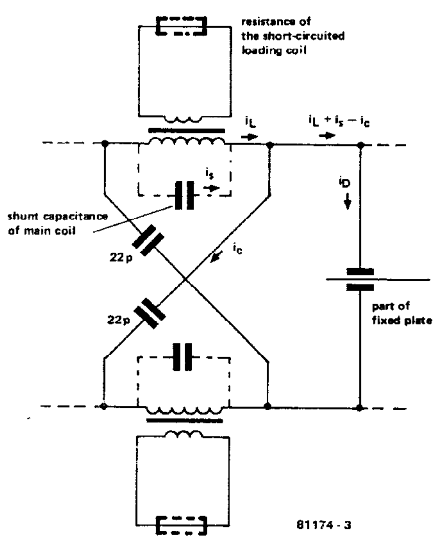

81174 - 3

Figure 3. One section of the delay line. If $i_c > i_s$ the circuit is an all-pass network; with $i_c = i_s$ it becomes an LC line.

relatively small diaphragm area. They evolved from typical open-backed radio-set cabinets, when it was realised that destructive interference by the phase-inverted backwave was reducing the net bass pressure in the room. Not counting bass-horns — which are acoustical matching-transformers and tend to be enormous — there are two basic ideas: either absorb the backwave in a closed box, or use it to drive a phase-shifting resonant system that will provide extra bass output (bass-reflex and some 'transmission-line' - labyrinth - systems). As we see it, two properties of electrostatic drivers make it unnecessary and undesirable to mount them in cabinets. *Unnecessary* because they can easily be made so large in radiating area that the destructive interference problem is only significant at the very lowest musical frequencies. There it can be overcome by bass-boosting — without distorsion — since the drive system itself is linear. (The system must of course be able to make large enough excursions: cubic-metres-per-second of volume-velocity remain cubic-metres-per-second.) *Undesirable* because boxes, by their nature, operate with considerable internal pressures. Such pressures would tend to come straight out through the electrostatic diaphragm.

Delay-line matching

One of the problems with electrostatic drive is the reactive current flowing in the plate-to-plate capacitance at high working frequencies. Peter Walker pointed out, as long ago as 1954, that this problem can be solved by using several loudspeaker sections as the shunt-elements in an LC delay-line. British patent number 1 228 775, published in 1971, explained how this matching arrangement could be deliberately used to control the radiation pattern of a large-surface electrostatic loudspeaker. In May 1979 an AES paper finally made it clear what was coming.

A section of the delay-line used in the ESL 63 is shown in figure 3. It is rather intriguing. With the diagonal capacitors it *looks* like a first-order all-pass network. It probably is, although at least part of the capacitor-current will null that through the winding stray-capacitance (shown dashed). The apparently-shorted secondaries in fact apply damping to the inductors, either to make the all-pass section behave itself on transients or to provide amplitude taper along the line (or both). The delay per section is 24 μsec, which corresponds to a path-length difference in air of just over 8 mm.

The essential point about this method of matching is that the unwanted *acoustical reflection* at the edge of the finite diaphragm will appear as an *electrical* reflection on the line. It can therefore be eliminated by a simple electrical modification to the line.

High voltage audio

The signal voltage applied to the ESL 63 drive system can peak at over ten *kilovolts*. This is necessary to achieve field-strengths near the ionisation (flash-over) limit, across a total airgap wide enough to permit sufficient diaphragm movement at low working frequencies. Designing an audio transformer to actually do this, over the full frequency range and at low distortion, must have been an interesting exercise . . .

In fact the ESL 63 has two identical quite hefty transformers, wired with their secondaries in series. Apart from enabling the bulky things to be fitted neatly inside a shallow base, this offers an elegant way of reducing the leakage inductance and stray capacitances that

set the upper limit to a transformer's bandwidth.
A point worth noting here is that there is no reason why an iron-cored audio transformer should in any way degrade the performance of the circuit in which it is used. On the contrary, use of a transformer is often the best way – if not, as here, virtually the only way – of doing the job.
Figure 4 plots the modulus of the input impedance. Perhaps surprisingly, it does not look too different from that of a typical conventional loudspeaker. Maintaining the diaphragm charge in spite of leakage – and extra losses due to localised airgap ionisation – requires the application of EHT to the (semi) conductive surface. The voltage should be so high that it will produce a polarisation field strength in the two airgaps (i.e. between the diaphragm in rest and each fixed plate) equal to half of the ionisation threshold. In the ESL 63 this is about 5.25 kV, corresponding to about 2 kV/mm. The diaphragm charge is of course proportional to this polarisation field strength.
The EHT generator is shown in figure 5. It is a classic Cockcroft-Walton cascade rectifier with one little addition: the AC feed is roughly stabilised by VDRs to make the EHT more or less independent of mains voltage fluctuations. A neat further detail is that the charge is delivered via a capacitor-bridged neon lamp. This, together with the 1 OM-ohm resistor and the very much higher leakage resistances, forms another classic circuit: the flashing-neon relaxation oscillator. The number of flashes per second is proportional to the rate at which charge is delivered to the diaphragm. Presumably this originated with Peter Walker's need to keep an eye on what was happening. One can almost hear him say: '. . . extremely sensitive and much better than clumsy meters with all those kilovolts around.'

4

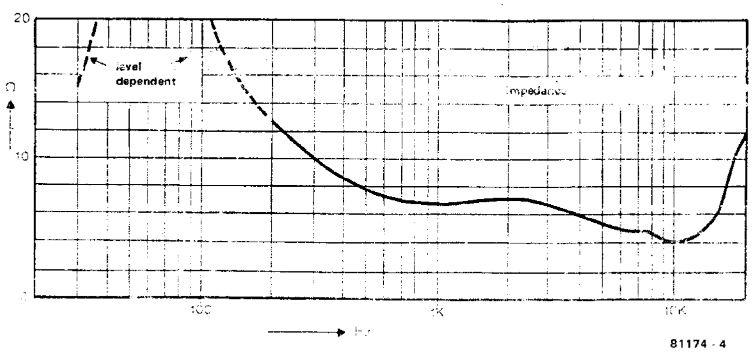

Figure 4. A plot of the input impedance-modulus.

Protection

An electrostatic loudspeaker is fundamentally linear up to the point at which ionisation occurs in one or other of the airgaps. As soon as that happens you have a few milliseconds left to get the drive voltage off, otherwise a flashover will permanently damage the system. The protection circuit must therefore operate almost instantaneously, then hold long enough for the ions to cool down. Figure 6 shows how this is achieved in the ESL 63. The high-frequency noise radiation that accompanies the onset of ionisation is picked up by the antenna – a length of wire running around the high voltage circuitry – and detected by T3. Noise above a certain level is a reliable indication that a potentially dangerous situation is building up. When this occurs the 555 timer will trigger, firing triac T1. Then . . . power amplifiers beware. This loudspeaker hits back . . .

The breakover diode T2 and triac T3 transfer the firing of T1 to the audio input in the event of the mains power being off. The circuit is therefore foolproof.
This arrangement will on its own protect the loudspeaker against accidental overloading. The power amplifier obviously must have a well-designed short-circuit protection system, even when it is incapable of delivering an excessive output voltage. (We noticed a shutdown apparently caused by the lady of the house plugging in an electric kettle!)
We have just seen that the protection system operates drastically and without warning. This could be very disconcerting to a listener who is concentrating on a loud musical passage: one instant the loudspeaker is performing superbly well, the next instant it is silent. The ESL 63 therefore also has a soft input-clipper, designed to start giving an audible distortion-warning

5

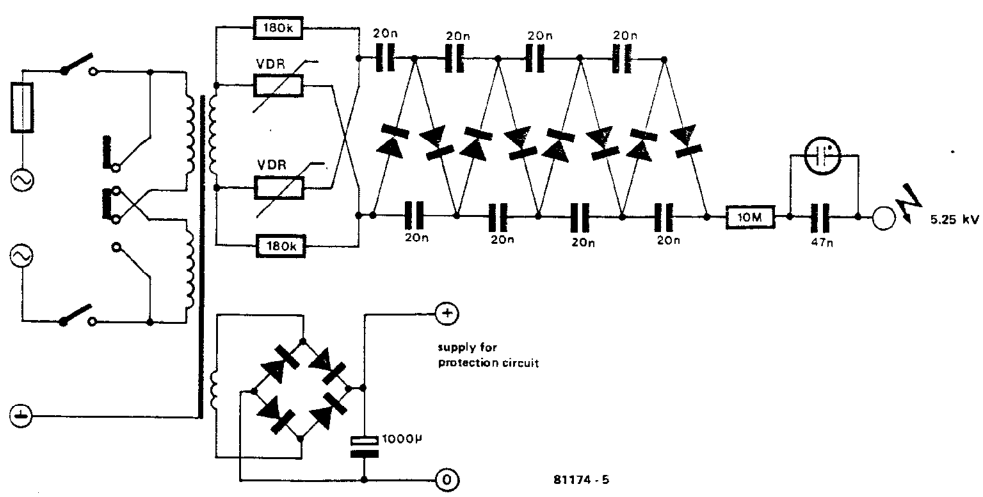

Figure 5. The power supply and EHT generator of the ESL 63. Note the EHT stabilisation by VDR's. The flash rate of the neon indicates the rate at which charge is delivered to the diaphragm.

6

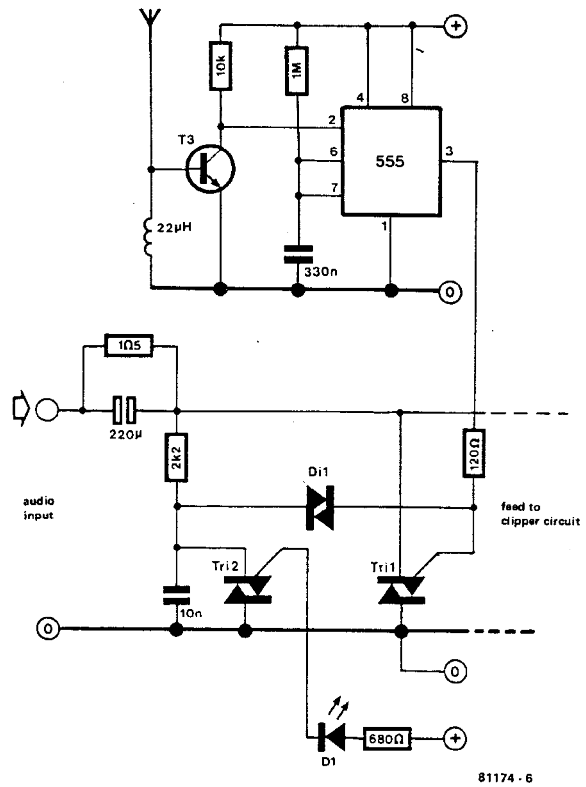

Figure 6. The emergency brake. Normal operation is with triac T3 held on by the power supply. The 555 timer is triggered by ionisation detector Tr3 and then fires T1 to short-circuit the audio input. With power off T1 can be fired via T2 by an excessive input signal.

7

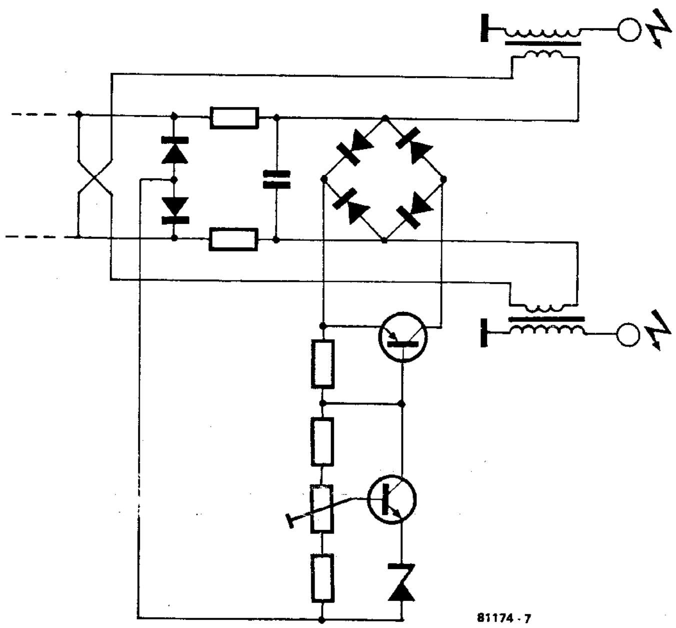

Figure 7. Part of the input circuit showing the soft clipper that reacts to peaks over 40 volts.

about 3 dB below the shutdown level. How this works can be seen in figure 7. The signal is fed to the audio step-up transformers via small series resistors. Tr2 acts in a preset-adjustable voltage-threshold circuit. When the input voltage peaks exceed 40 volts, in either direction of swing, Tr1 will be turned on, drawing extra current through the series resistors. This will cause an audibly non-linear reduction in primary drive voltage.

The clipper circuit could presumably be disabled without affecting the safety-margins in any way. We would in fact have preferred an optical warning device — also driven from a monoflop — instead of a headroom reducing clipper. How about a switchable option?

FRED

The acronym FRED — full range electrostatic doublet — is another Quad original. Let us now try to understand how it works. Peter Walker illustrates his principle of radiating an expanding wavefront from a flat diaphragm, by means of truncated delay-line matching, with figure 8. Others have interpreted this as producing a *virtual* point source about 30 cm behind the loudspeaker as viewed by the listener. The real ESL 63 situation seems to be more complex.

An *acoustical doublet* consists of two equal sources of opposite sign, each small compared to the wavelength and spaced a similarly small distance apart. The net pressure at a distance from the doublet that is large compared to the *spacing* is accurately described by a cosine function (the cosine of the angle between frontaxis and observer direction). This is shown in figure 9. If we've got it right, what Peter Walker has done is to arrange the delay and amplitude shading of the diaphragm drive (with the entire diaphragm behaving as a phased array of *real* sources) to maintain the doublet-like axial lobes throughout essentially the entire working range, even though the array quickly becomes large compared to the wavelength. The radiation patterns of figure 10, showing only a slight sharpening at 8 kHz, illustrate just how well he has succeeded . . .

This is not only a fine loudspeaker: it is a fascinating piece of applied physics.

Listening . . .

Two aspects of the loudspeaker-room interface that can cause trouble in a domestic listening-environment are *early reflections* and *standing waves.*

A simple case of an early reflection problem is shown in figure 11. The loudspeaker is standing on a highly reflective floor, so that the listener receives sound along the direct path *and* indirectly off the floor. The reflection will arrive perhaps one or two milliseconds after the direct wave and with rather less intensity. The most effective approach is to imagine that the floor is

8

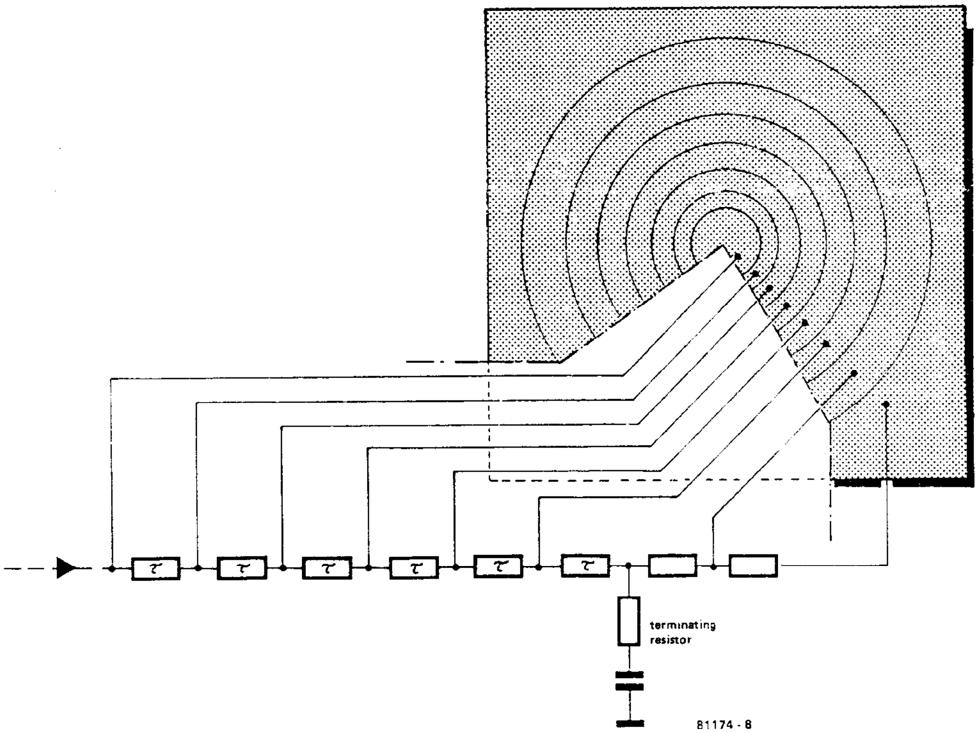

Figure 8. This illustration shows how the ESL 63 radiates a curved wavefront from a flat diaphragm using delay line matching. The delay-line taps are connected to the concentric areas of the drive unit fixed plates.

transparent with the reflection being produced by an *image* of the loudspeaker underneath it. The extra distance travelled by the delayed wave will correspond to one or more half-wavelengths of some musical frequency (usually in the mid-range). The problem is that destructive interference will cause a partial cancellation of the pressure due to the direct wave, resulting in fairly broad response-dips centred on frequencies corresponding to odd numbers of half-wavelengths. The whole-wavelength situation, with the two waves more or less in-phase, will result in broad peaks that are usually even more unpleasant than the dips.

The image-method can be applied to more complex situations provided one uses enough images to account for *all* the troublesome reflections. Note, by way of an example, that a double-bounce reflection would be 'radiated' by an *image of an image.* We believe it was this effect that we ran into on our first hearing of an ESL 63, in a room with a tiled floor. The loudspeaker was well clear of walls — but it 'honked' slightly. The effect disappeared when we either moved the loudspeaker into a carpeted room or else raised it off the floor on a makeshift standard. (An empty milk-bottle crate, actually. Optional extra?)

The carpet presumably attenuated the reflection sufficiently; the effect of the milk-crate is more instructive: refer again to figure 11. We started by noting that the ESL 63 radiates as a doublet,

9

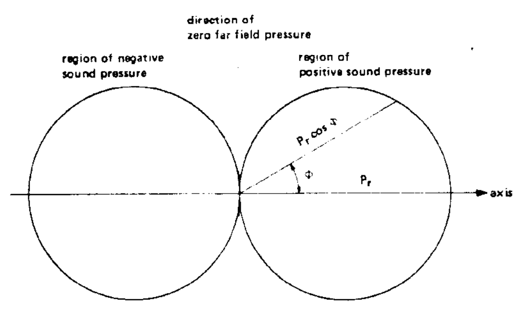

Figure 9. The so-called radiation pattern of a doublet.

with axis horizontal and about 50 cm off the floor. Now fit in the listener, with his ears about 1 metre off the floor and about 3 metres away from the loudspeaker. The vertical angle subtended by the direct listener-path at the loudspeaker is . . . etc. etc . . . This produced the puzzle that the milk-crate operation is . . . etc. etc . . . This produced the puzzle that the milk-crate operation could only have reduced the reflected wave pressure by a couple of decibels. Then it dawned. Moving the image source about 35 cm further below the floor had caused its listenerpath to be intercepted by a coffee-table!

A standing wave as a room-interface *problem* should be distinguished from a standing wave behaving itself as part of the reverberation process. Any standing wave mode is characterised by its *natural frequency* and its degree of *damping.* The problem only appears when a single lightly-damped mode, or a closely-spaced group of such modes, occurs in isolation. Musical tones from an instrument or a loudspeaker, particularly sustained tones, that happen to be near the natural frequency can build up a forced vibration of distressingly high amplitude. After the tone stops the mode will decay more or

10

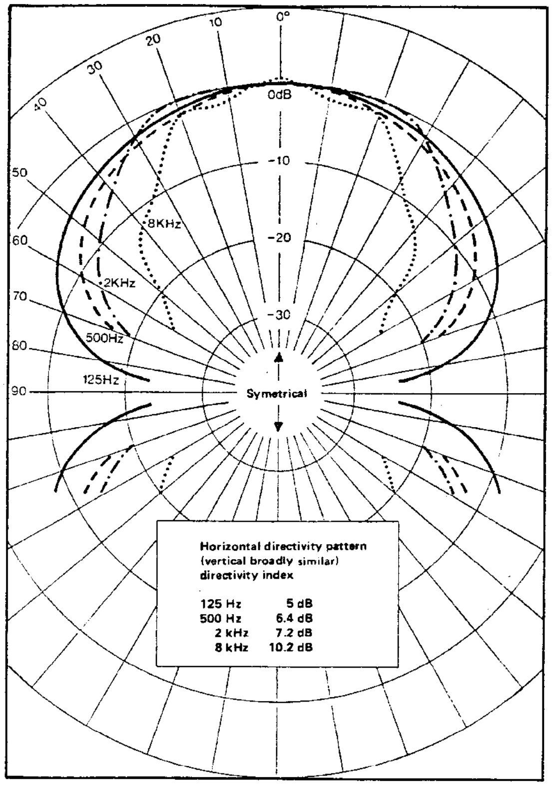

Figure 10. The measured radiation patterns of the ESL 63. Note that the patterns appear flattened in comparison with figure 9 due to the use of a logarithmic amplitude scale.

11

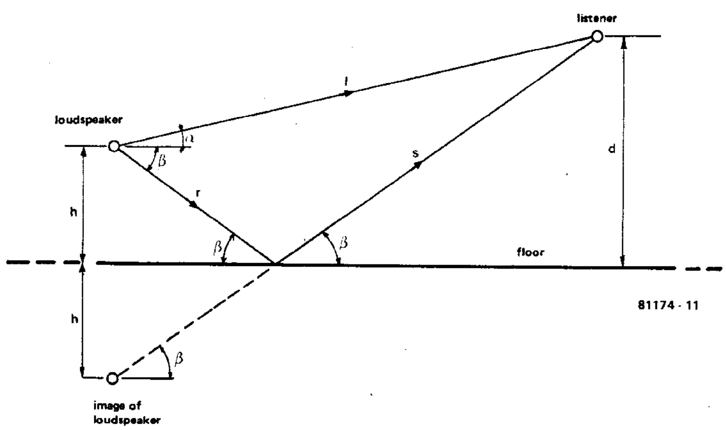

Figure 11. An illustration of the image source method applied to the simple case of one reflection off the floor.

less slowly at its natural frequency. Beats can occur if two or more modes are excited together (by the same tone) and then decay independently.

The advantage of a doublet in this room interface problem is that its output is in the form of a particle-velocity along the axis. It will therefore only couple to modes that have a significant component of particle-velocity along the axis at the doublet position. This is probably one of the reasons that electrostatic doublets are considered bass-weak, even when like the ESL 63 they are not. The room simply doesn't give the expected (liked?) booming response.

Combating the room-boom problem with a doublet starts by feeding it with a low-level sinewave (at the 'right' frequency, obviously!) and then moving it — and yourself — around, until the offending mode is identified. Then try to find a position or orientation for the loudspeaker that will sufficiently weaken the unwanted coupling.

To woof or not to woof . . .

Many years ago, when all available loudspeakers had a single nominally full-range cone driver, some innovative soul introduced the concept of the *tweeter.* The next step was to shift the working range of the main driver downwards in frequency, then call it a *woofer.* Later again came the *squawker* — an epithet that never really stuck — and now we have ultra-widerange systems that add a *supertweeter* and, sometimes, a *sub-woofer.* The *separate* subwoofer systems that are becoming available should more logically be called *rumblers* or *thumpers* . . .

The question now is: granted that the ESL 63 delivers its bass output the hard way — into the relatively unfavourable doublet airload — does it need help?

Peter Walker says quite emphatically 'no'. For one thing the ESL 63 bass response is more extended and better-controlled than that of its predecessor. For another, its doublet radiation pattern gives a more aperiodic interface to a typical listening room than available auxiliary bass systems do — because *they* are omnidirectional radiators.

On the other hand, people who insist on realistic (?) reproduction of organ pedal — or for that matter heavy trucks and underground trains — may disagree. That of course is up to them . . . (They might of course try listening *right in close* to a doublet. Proximity effect will provide quite an amount of bass boost!)