The Design of Electrostatic Loudspeakers

Frank Verwaal
June 2011
Rev. 3

1 Preface

Electrostatic loudspeakers are interesting things – both from a technical and a musical point of view.

On the internet you can find testimonies of the achievable sound experience that make your mouth water. Indeed some ESL's have a naturalness that is hard to achieve with a conventional electro-dynamic loudspeaker.

I believe that this remarkable performance is due to three properties of the ESL.
First, the mass of the diaphragm can be made so small that the electrostatic force is almost directly driving the air.
Second, due to the directivity of an ESL a bigger part of the sound power reaching the listener comes through the direct path and a smaller part comes from reflections.
Third: because of the light diaphragm and the smooth phase response, a well designed ESL can reproduce a square wave input signal of, say 300Hz, as a square wave, as we can show with a good microphone and an oscilloscope.

Another attractive point of ESL's is that it is possible to home build them – and I do not mean to mount purchased drivers in a cabinet that you build yourself. No, you can literally assemble the drivers yourself, from relatively simple materials.

For several decades now, a deep understanding of ESL's has been available from Peter Walker [1, Walker] and Peter Baxandall [2, Baxandall].
This enables ESL builders to design, and therefore know, the sensitivity, frequency response and directivity of the loudspeaker before starting the actual construction.
Unfortunately this readily available theory is not widespread; neither among ESL manufacturers, nor among amateur ESL builders.
Because of this, ESL design is still largely a matter of trial and error.

With this book I have tried to describe all the effects that determine the performance of an ESL without using any mathematics. This should make the material accessible for all ESL builders.

If you want to design your own ESL you cannot avoid to use some formulas to predict the frequency response, sound pressure level and directivity.
These formulas are available in the appendices, together with some deeper background insight.
In the interest of brevity, I have in most cases not included the derivations of the formulas. This does not have to stop you from using them. Just plug the intended values of the design parameters of your ESL in the formulas and evaluate them to see what happens. If you don't like the result you can change the design.
Admittedly, this design approach is still iterative, but at least the iteration loop does not involve building an ESL and scrapping it if it is not what you wanted.

The theory presented here allows us to write computer programs for the simulation of the sound pressure and frequency response of an ESL.
I have made three such programs available through my website [3, Verwaal]: *ESL.pas*, *ESL.m* and *ESL_ladder.m*.

2 Contents

3 Revision History

Revision	Date	Changes
1	June – August 2009	Document created
2	September 2009	Corrected some errors Corrected some typo’s Fixed ESL_ladder.m
3	June 2011	Corrected some typo's

4 Introduction

This book introduces the reader to the electrostatic loudspeaker.
The structure of the book is as follows:

Chapter 5 describes the basic operating principle.
Chapter 6 defines some quantities associated with sound, such as sound pressure, particle velocity and sound pressure level.
Chapter 7 discusses the differences between flat sound waves, cylinder shaped waves and spherical waves.
Chapter 8 discusses the advantages of operating the loudspeaker with a constant polarizing charge.
Chapter 9 forms the core of this book.
It dives into a multitude of mechanisms that influence the frequency response.
This chapter provides the basic insight needed to design a loudspeaker with a flat frequency response in a desired frequency range.
Chapter 10 discusses the sensitivity of the loudspeaker or, to put it simply, how loud it will sound.
Chapter 11 discusses how the loudspeaker behaves in the listening room.
This material also provides clues on how to choose the frequency response and the directivity of the loudspeaker.
Chapters 12 and 13 are design oriented. Chapter 12 takes the approach to minimize path length differences to the point that their influence on the frequency response disappears. This leads to conflicting requirements on panel size, which we can meet by a three-way design.
Chapters 13 and 14 explore other solutions to prevent the unwanted effects from path length differences.
Chapter 15 describes the limitations and trade-offs of the step-up transformer.
Chapter 16 discusses construction details and the choice of materials.
Chapter 17 gives some considerations regarding safety.

5 Operating Principle

Most loudspeakers have some moving part that transfers its vibrations to the surrounding air where they become sound waves.
In a “normal” electrodynamic loudspeaker (Figure 1 and Figure 2) this moving part is the *cone*[1].
The cone moves because it is driven by the equivalent of an electric motor, consisting of a coil (the *speech coil* or *voice coil*), which is suspended in the magnetic field of a strong magnet. The force that such a coil feels when we send a current through it is known as the *Lorentz force*[2].
Because the cone is driven in one central point it must be rigid to ensure that it moves uniformly. As a result the mass of the combination of speech coil and cone cannot be made very small. Therefore, a large part of the force generated by the speech coil is never transferred to the air but is used to accelerate and decelerate this mass instead.

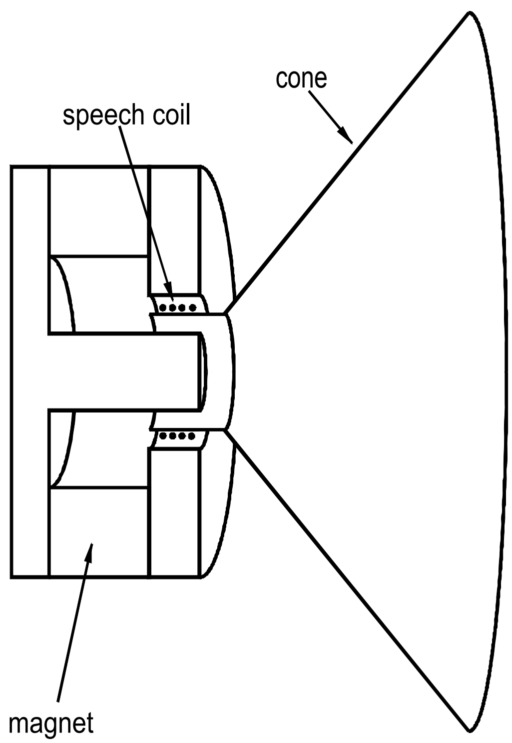

Figure 1 Electrodynamic loudspeaker (driver).

[1] In tweeters, this part is often not cone shaped but dome shaped.

[2] Although the French, of course, call it the *Laplace* force.

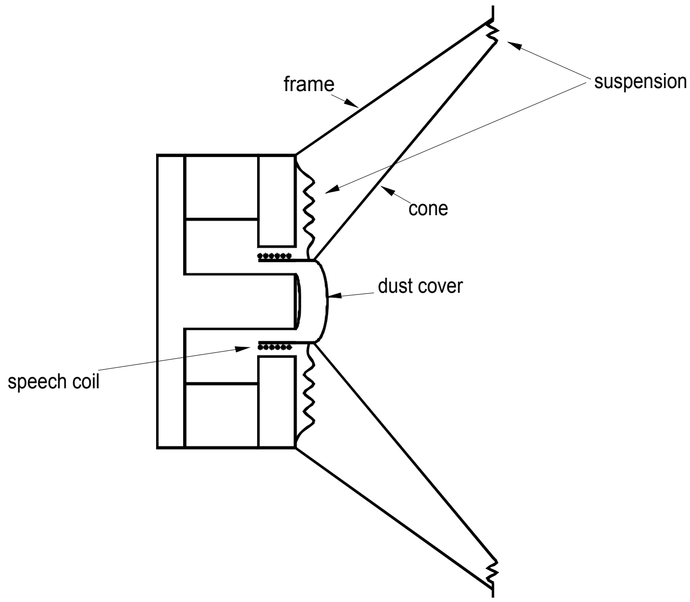

Figure 2 **Cross section of an electrodynamic driver.**

In an electrostatic loudspeaker (Figure 3), by contrast, the moving part is a very thin sheet, typically made of polyester film. We call it the *diaphragm*.
Because the force that drives it is distributed over its entire surface there is no need for it to be rigid. As a consequence it can be extremely light – so light that its mass is negligible. The electrical force driving the diaphragm therefore almost directly drives the air.

The driving force in an ESL is not the Lorentz force but the attractive or repulsive force between electric charges, known as the *Coulomb force*[3].

Everybody has experienced the attractive force between electric charges when observing how certain plastics, after they have been rubbed with a piece of cloth, attract dust particles and small snippets of paper.
The electric charge that the rubbing has placed on the surface of the plastic cannot move[4] because the plastic is an isolator. This is why we speak of *static electricity*.

[3] No objection from the French here.

[4] That is to say: not very fast. It will leak away eventually though.

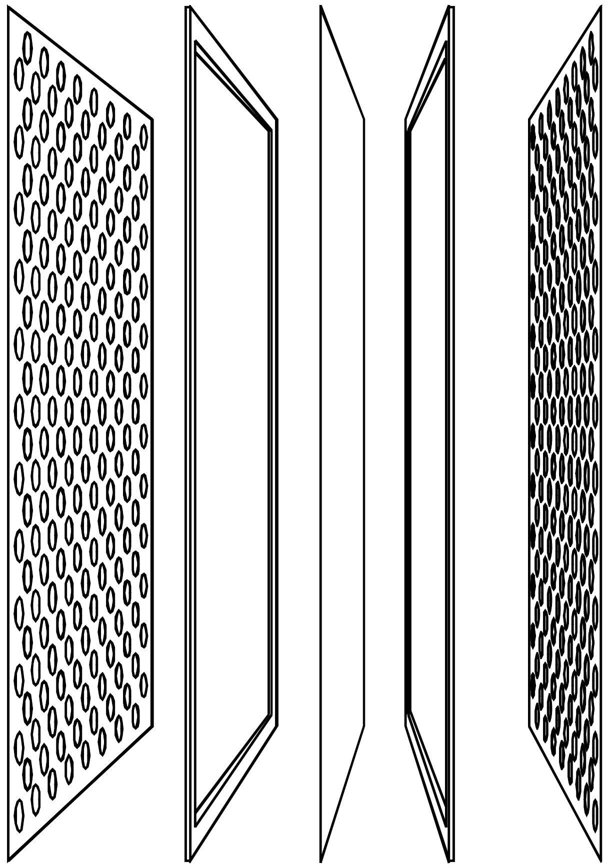

Figure 3 **Exploded view of an electrostatic loudspeaker.**

In an electrostatic loudspeaker, we use the Coulomb force to move the diaphragm.

For that purpose we place a charge on the diaphragm by connecting it to a voltage source that delivers a dc voltage of several thousand volts[5]. We call this the *polarizing voltage*.
We suspend the charged diaphragm between two perforated conductive stator plates, called the *stator electrodes*, as shown in Figure 3.
The perforation of the stator plates allows the sound generated by the moving diaphragm to escape.
Now, we connect the stator plates to a high voltage audio source.
We typically obtain the high audio voltage from a normal audio amplifier in combination with a so called *step-up transformer*, although we can also build a special purpose high voltage audio amplifier.
The high audio voltage that we apply to the stator plates generates an electric field between these plates, which pulls on the diaphragm charge.

[5] Of course, for this to work we must make the diaphragm electrically conductive.

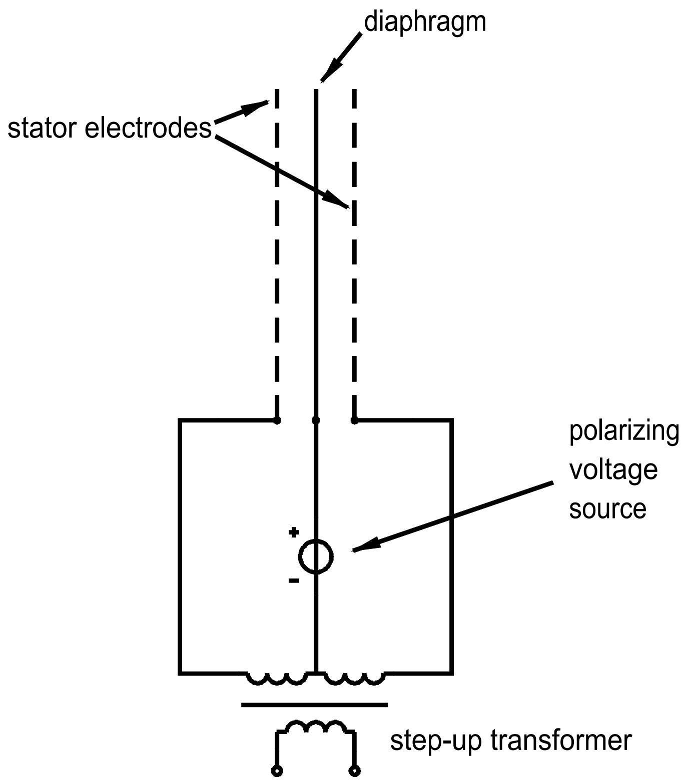

Figure 4 ESL with step-up transformer and polarizing voltage source.

Figure 4 shows a typical arrangement of stator plates, diaphragm, polarizing voltage source and step-up transformer.
The diaphragm is suspended between the two stator electrodes.
To keep it in its place it must be stretched.
The idea is that this stretching does not impede its movements.

Figure 5 shows a slightly modified arrangement designed to achieve two things:

1. a constant drive *current* on the stator plates instead of a constant drive *voltage* and
2. a constant polarizing charge on the diaphragm.

The reasons why this is desirable will be explained in the chapters 8 and 9.
We achieve a constant drive current by adding series resistors between the step-up transformer and the stator electrodes.
We achieve a constant polarizing charge by adding a high valued resistor between the polarizing voltage source and the diaphragm.

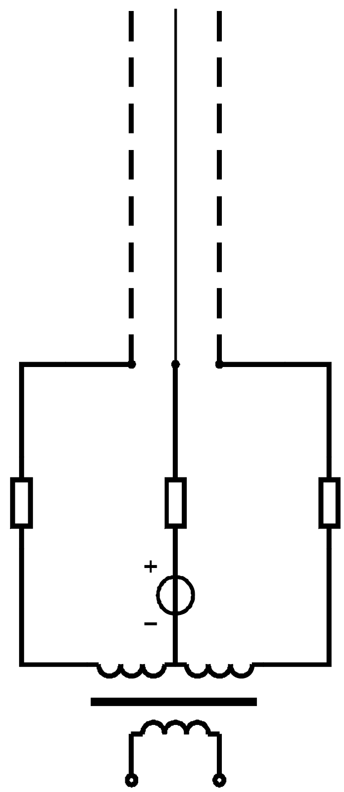

Figure 5 Current driven ESL in constant charge mode.

6 Acoustical Quantities

6.1 Frequency

Sound consists of vibrations in the air, which travel through that air as waves.
The propagation velocity of sound waves in air[6] is about 343 m/s.
The number of cycles per second of the vibration is called the *frequency*.
The unit of frequency is the hertz (Hz). One hertz equals one cycle per second.

In our youth we can hear sounds with frequencies between roughly
20Hz and 20kHz.
As we grow older we gradually loose the highest octave, and that reduces the frequency range to 20Hz – 10kHz or perhaps even somewhat less.

6.2 Sound Pressure and Particle Velocity

Sound waves consist of two components:

1. Fluctuations of the pressure about the equilibrium value that is formed by the atmospheric pressure. We call these fluctuations the *sound pressure*.

2. Air moving from the high pressure area's to the low pressure area's.
   We call this the *particle velocity*.

Each of these two sound components can and will generate the other.
This interaction between the two components of the sound allows it to propagate through the air as a wave.

The human ear is predominantly sensitive to the pressure fluctuations.

The unit of pressure is the Pascal (Pa), which is equal to one Newton (N) per square meter (m).
One Newton is about the force that gravity exerts on a large apple (the joke is from Peter Baxandall [2, Baxandall]).

The atmospheric pressure is approximately 100,000 Pa (10^5 Pa), which equals the weight of one hundred thousand apples on a square meter[7].

Appendix I discusses another important acoustical quantity: the *volume velocity*.

---

[6] At room temperature and atmospheric pressure.

[7] Or about 1kg per square centimeter.

6.3 RMS Sound Pressure

We need a way to express how loud a sound is.
For this we use the *rms value* of the sound pressure.
The letters rms are an abbreviation of *root-mean-square*; which is the familiar measure of the magnitude of an ac voltage. It is suited to represent the magnitude of other things than voltages too; in this case we use it to express the magnitude of the sound pressure.
The rms value of a pressure that varies sinusoidally with time is equal to approx. 0.707 times the peak value.

6.4 Acoustical Impedance

As it turns out, in a flat wave (see chapter 7 for a discussion of flat waves) the sound pressure is always proportional to the particle velocity.
The same is (approximately) true for waves that are approximately flat, for example sphere shaped or cylinder shaped waves at a large distance from the source.
The proportionality constant depends on the atmospheric pressure and on the mass of a cubic meter of air[8] and amounts to 406.4 Ns/m^3.
We could call this the *acoustical impedance* of the air.

6.5 Sound Intensity

Consider a sound wave that crosses some imaginary rectangle (window), which is perpendicular to the direction in which the wave propagates.
The *sound intensity* is the power that the sound carries through that window per unit of area of that window. It is therefore a *power density* and has the unit W/m^2.
It is equal to the product of the rms sound pressure and the rms particle velocity[9].
Because in a flat wave these two are proportional to each other, a two times larger sound pressure is accompanied by a two times larger velocity and hence carries four times more power per unit of area.

---

[8] Incidentally, the mass of a cubic meter of air at atmospheric pressure and room temperature is approximately 1.18kg. More than you expected, is n't it?

[9] Here we assume that the pressure and the velocity have the same direction and are in phase. This is generally not the case, but it is true in a flat wave.

6.6 The Decibel

The human ear can deal with a large range of sound pressures.
The softest sound that a healthy ear can hear is 20 μPa_{rms} at about 1000 – 3000 Hz.
The loudest sound that the human ear can handle is determined by the threshold of pain. It occurs at an rms pressure that is about five million times larger: 100 Pa_{rms}.
Note that this is still only one thousandth of the atmospheric pressure.

Such a large range presents practical difficulties.
For example, if we try to make a plot of sound pressures versus frequency on a linear vertical scale, we can see only the top decade, which is the range from 10 Pa to 100 Pa. The lower pressures fall on the horizontal axis.

We solve this by using the *decibel* (dB).
The decibel is a measure of the *ratio* of the signal powers of two signals.
The dB is a logarithmic measure, which means that a pair of two sound levels having a certain *magnitude ratio* always differ by the same number of dB's, irrespective of their absolute magnitude. If we plot sound pressures in dB, equal pressure ratios are always separated by the same distance in the plot.

As an example, when one sound pressure is 10x larger than another, they differ by 20 dB, no matter how loud each sound is individually.
When one sound pressure is twice as large as the other, they differ by 6 dB.
It is useful to memorize these relations.

6.7 Sound Pressure Level (SPL)

The *Sound Pressure Level* is the rms sound pressure expressed in dB with a reference level of 20 μPa_{rms}. The SPL therefore expresses how much louder a sound is than the softest sound that you can hear.
A sound pressure of 20 μPa_{rms} corresponds to an SPL of 0 dB_{20uPa}; the threshold of pain of 100 Pa_{rms} corresponds to 134 dB_{20uPa}.

To get an idea of what a certain SPL means subjectively, the following table (found on the internet) gives the SPL of some familiar sounds[10].

Typical sounds	**Typical music**	**SPL, dB_{20uPa}**
Threshold of pain		135
Artillery, 100 yards	Cannon (peaks)	130
Pneumatic chipper		125
Riveter, nearby		120
Loud car horn, nearby	Very loud rock (peaks)	115
	Very loud classical (peaks)	110
		105
Inside NY subway	Very loud classical (avg)	100
	Loud classical music	95
Heavy truck		90
Inside motor bus	Moderately loud classical	85
Noisy traffic, corner		80
Noisy office	Soft popular music	75
		70
Business office	Soft classical music	65
Conversational speech		60
		55
Private office	Very soft music	50
Background noise, city home		45
		40
Background noise, suburb		35
Library		30
Background noise , country night		25
Whisper, leaves rustling		20
Good recording studio		15

[10] This table leaves us with the question of how you determine the SPL of a non-stationary signal, such as a cannon shot. You cannot determine the rms value of such a signal and that is a necessary step towards determining the SPL.
I do not have an answer to that question.

7 Flat Waves, Cylindrical Waves and Spherical Waves

7.1 Introduction

It will be useful later on to know something of flat waves, cylindrical waves and spherical waves.
This chapter briefly discusses their properties.

7.2 Flat Waves

The Shape of Flat Waves
A flat wave is a wave in which the crests and troughs form parallel planes, perpendicular to the direction in which the wave propagates.
A flat wave is generated by a vibrating wall of infinite size[11].
Figure 6 gives an illustration of such a sound radiator and of the waves it generates.

SPL at a Distance
A flat wave has the peculiar property that its sound intensity and therefore also its sound pressure does not decrease with the distance that the wave travels.
This is not hard to understand: the acoustical power that is transmitted by one square meter of the wall must at some distance from the wall pass through a surface area of the same square meter[12].

Vibrating Wall of Finite Size
Of course a vibrating wall of infinite size is physically not possible[13].
A practical vibrating wall of finite size still generates waves which are approximately flat in the vicinity of the wall. At larger distance, however, the waves start to fan out and gradually bend into a spherical shape. The distance over which the waves can be considered approximately flat depends on the size of the sound radiator.

A flat vibrating wall of infinite size is a good description of a very large ESL.

---

[11] A wall of infinite size is of course not practically possible.
The problem with physics is that ordinary every day situations are often too complicated to analyze, while idealized situations involving simple shapes such as spheres and infinitely long and thin lines are mathematically much simpler to describe.
The joke[b] goes that the owner of a race horse asked a physicist what to feed his horse to make it go faster. The physicist replied: that is a complicated problem, but I do have an exact solution for spherical horses in vacuum.

[b] Joke by my colleague Marcel van de Gevel; I don’t know here he got it from.

[12] This assmumes that the air has no losses, for example due to friction. At audio frequencies and at distances of only a few meters, that is a reasonable assumption.

[13] One of its problems is that it generates an infinite amount of power.

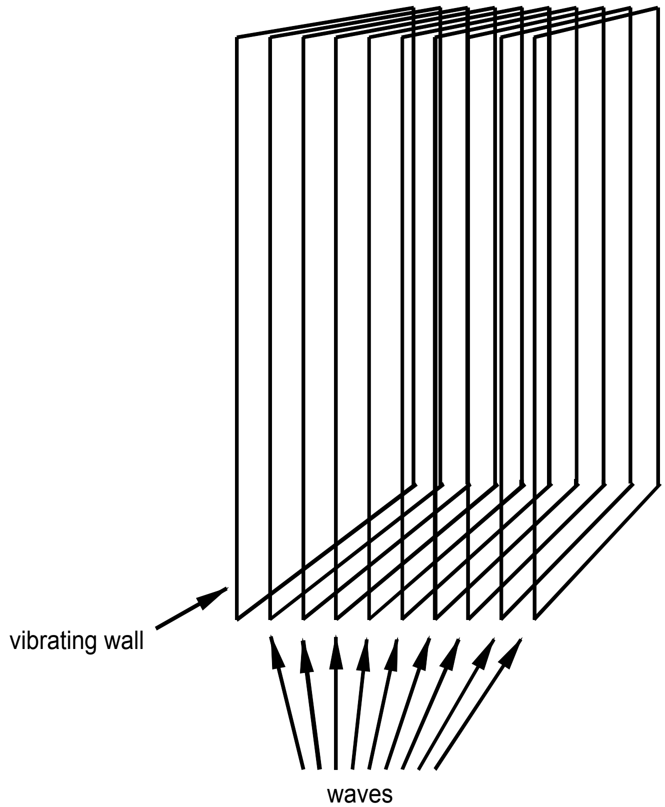

Figure 6 Flat waves emanating from a vibrating wall.

Relevance
Flat waves are generated by a flat ESL, at least in the direct vicinity of the diaphragm.
At larger distance the waves fan out and become spherical.
Flat waves are therefore what the listener to a set of electrostatic headphones hears.

A flat ESL of such large size that the listener finds himself in the region where the waves are still flat would be a solution to a certain problem we face in the design of an ESL. This is the problem of undulations of the frequency response due to destructive interference caused by path length differences.
We will discuss the effects of path length differences in detail in chapter 9, sections 9.7 and 9.8.

Unfortunately, for this solution to work the size of the loudspeaker must be impractically large. We will investigate this in some more detail in section 13.5.

7.3 *Cylinder Shaped Waves*

Line Source and Cylinder Waves

A cylinder shaped wave is a wave in which the crests and troughs form co-axial cylinders.
Cylindrical waves are generated by a pulsating cylinder of infinite height.
We call such a source a *line source.* Note that a line source does not have to have a zero radius.
Figure 7 shows a section of such a cylindrical radiator and the waves it generates.

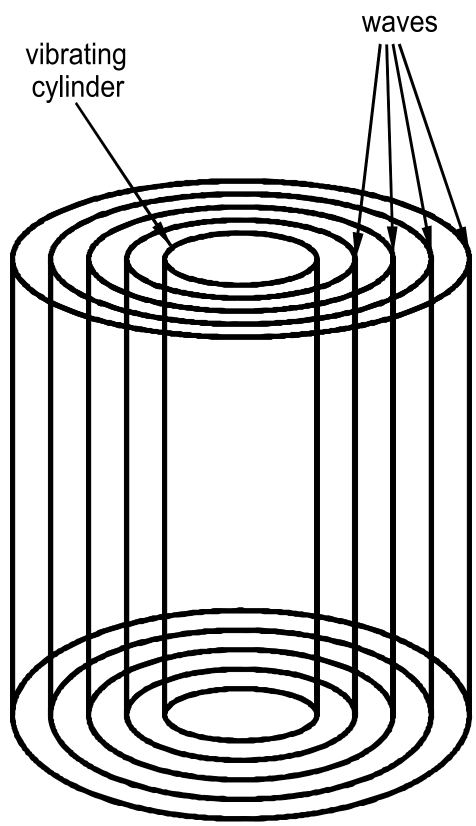

Figure 7 **Vibrating cylinder, creating cylindrical waves.**

Curvature of the Waves

Due to the curvature of the waves, the particle velocity and the sound pressure are not in phase, nor do they have the ratio of 406.4Ns/m^3 that we normally find in flat waves.
As the waves spread out, however, their curvature decreases and at a sufficiently large distance they are almost flat. From that point on the velocity and pressure are again approximately in phase and have again approximately a ratio of 406.4Ns/m^3.

Decrease of the SPL with Distance

As the cylinder shaped waves spread out, the area of their cylinders increases proportionally with the distance that the wave has traveled.
Therefore, the power that the source radiates per meter of its height spreads out over an increasing area. As a result the sound intensity is four times smaller when the distance is four times larger. It follows that the sound pressure is two times smaller at that distance[14].

Conclusion:
The SPL of a line source drops with 3 dB each time that the distance doubles[15].

Of course a cylinder shaped source of infinite height is not physically possible[16].
A practical cylinder source of finite height will still generate waves that are approximately cylindrical as long as the distance is small compared to the source's height.
At larger distances, the waves fan out and eventually become spherical.

We can realize a line source of infinite height by placing a line source of finite height between floor and ceiling. Reflections in the floor and the ceiling will extend the source to a virtual infinite height.

Relevance

A line source of infinite height is a solution to the problem of the undulating frequency response due to destructive interference caused by path length differences.
We will discuss the effect of path length differences in detail in chapter 9.
We will discuss the line source solution in more detail in section 13, section 13.6.

Many ESL designs have a small width but a large height. They are (approximately) line sources of finite height.

---

[14] This relation holds at such a large distance that the ratio between pressure and velocity has become approximately $406 Ns/m^3$.

[15] At least at sufficiently large distance.

[16] If only because it would radiate an infinite amount of power.

7.4 Sphere Shaped Waves

Point Source and Spherical Waves
A sphere shaped wave is a wave in which the crests and troughs form concentric spheres.
Spherical waves are generated by a pulsating sphere.
We call such a source is a *point source*. Note that a point source does not have a zero radius.

A point source has no preferred direction; it radiates sound equally in all directions.
We call such a source *omni-directional*.

A point source generating a sinusoidal waveform is characterized by its frequency and its strength. The strength is given by its rms *volume velocity*[17].

Figure 8 illustrates a vibrating sphere and the waves it generates.

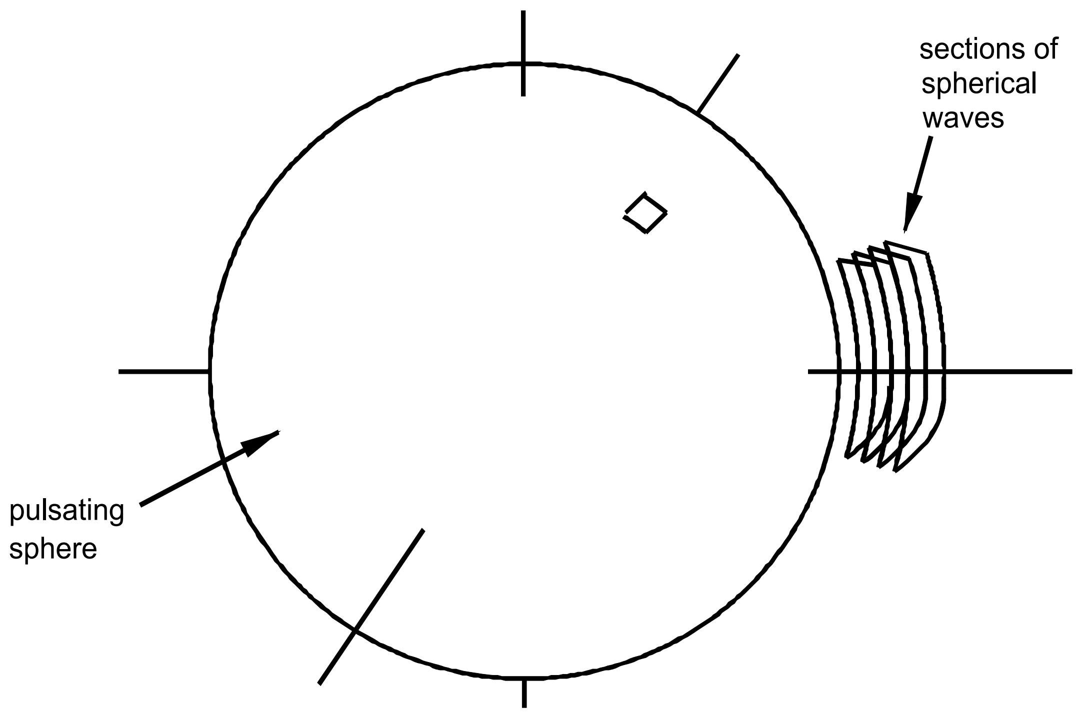

Figure 8 Pulsating sphere, creating spherical waves.

[17] The volume velocity is the product of the rms surface velocity of the surface and its surface area. It effectively tells you how much air the source displaces per second. See appendix I for more details.

Curvature of the Waves – Proximity Effect.
Due to the curvature of the waves, the particle velocity and the sound pressure of sphere shaped waves are not in phase, nor do they have the ratio of $406.4 Ns/m^3$.
As it turns out, this causes at close distance and low frequencies a particle velocity that is larger than you would expect based upon the sound pressure.
In other words: in close proximity the velocity has too much bass.
This is known as the *proximity effect.*
The proximity effect plays a small role in the frequency response of ESL's and a big role in the frequency response of electrostatic headphones.
It is also important when using a type of microphone[18].

As the waves spread out, their curvature decreases and at some distance they become approximately flat. Beyond that point the velocity and pressure are again approximately in phase and have again approximately a ratio of $406.4\ Ns/m^3$.

Decrease of the SPL with Distance
As the sphere shaped waves spread out, their area increases with the square of the distance that the wave has traveled.
Therefore, the power that the source radiates spreads out over an increasing area.
As a result the sound intensity is four times smaller when the distance is two times larger.
It follows that the sound pressure is two times smaller at that distance[19].

Conclusion:
The SPL of a point source reduces by 6 dB each time that the distance doubles[20].

---

[18] When we place a microphone that is sensistive to particle velocity, rather than sound pressure, too close to a point source, the proximity effect causes an emphasis of the bass tones. The mouth of a singer or a speaker behaves as a point source.

[19] This relation holds at such a large distance that the ratio between pressure and velocity has become approximately $406\ Ns/m^3$.

[20] This relation is valid at sufficiently large distance, where the ratio between pressure and velocity is approximately $406\ Ns/m^3$.

Relevance
The waves of any loudspeaker of finite size will at a sufficiently large distance become spherical.

Further, an ESL designed to behave as a point source would be a solution to the undulations of the frequency response due to path length differences.
These will be discussed in sections 9.7 and 9.8
Loudspeakers designed along these lines are discussed in chapter 13.

Point sources play a crucial role in the derivation of the frequency response of an ESL (Walker's equation), which we will discuss in section 9.2.

A point source (pulsating sphere) is a good model of a cone loudspeaker mounted in a closed or semi-closed box. This model is valid for frequencies that are low enough that the size of the box is small compared to the wave length.

Finally, the mouth of a person that speaks or sings behaves as a point source, at least at low frequencies.

8 Constant Charge mode of Operation

In the set-up of Figure 4 we bring the polarizing charge on the diaphragm of an ESL by connecting it to a high voltage source.
A disadvantage of this set-up is that the movements of the diaphragem change its capacitance, and thereby also the polarizing charge. This causes distortion.

As an example, suppose that a double bass and a flute sound together.
The bass tone causes relatively large diaphragm excursions, and the resulting modulation of the diaphragm charge modulates the sensitivity of the loudspeaker.
As a result, the bass tone modulates the loudness with which the loudspeaker reproduces the flute.
We call this kind of distortion *intermodulation distortion* and it sounds ugly.

A solution is to lock the charge onto the diaphragm.
We could do this by disconnecting the diaphragm once it is fully charged, but this has the problem that the polarizing charge will slowly disappear due to leakage currents.

A better solution is to connect the diaphragm to the polarizing voltage through a large valued resistor. Figure 5 shows this series resistance. Ignore for the time being the resistors that were added in series with the stator electrodes.
The resistor is so large that by the time a tiny percentage of the polarizing charge has managed to leak off the diaphragm, it wants to get back on it.
We call this the *constant charge* mode of operation.

This leaves one problem.
When the diaphragm does not move uniformly, as it never does, the polarizing charge, which is now nicely locked onto the diaphragm, redistributes with every diaphragm movement.
This happens because those parts of the diaphragm that move closer to one of the stator electrodes have a larger capacitance and therefore attract a larger portion of the polarizing charge.
As a result, the double bass modulates the radiation pattern of the flute.
This is intermodulation distortion in another guise.
It seems as if we made little progress by using the series resistance.

The solution is to use a diaphragm that has a poor conductivity.
The poor conductivity makes it impossible for the charge to move around fast enough to keep up with the diaphragm movements.

For this purpose we must increase the *sheet resistance* of the diaphragm.
The sheet resistance is the resistance of a square piece of sheet[21].
A sheet resistance between 10^9 and 10^{11} Ω will prevent the polarizing charge from redistributing all the time [2, Baxandall, page 114].
A larger value creates the risk that too much polarizing voltage is lost across the diaphragm resistance due to leakage current.
With such a low conductivity diaphragm, we no longer need the series resistance of Figure 5.

[21] It is not difficult to see that the resistance of a square piece of sheet is independent of the size of the square.

9 Frequency Response

9.1 Introduction

In this chapter we investigate what determines the frequency response of the loudspeaker. We will find a large number of effects that we must take into account.
Once we have indentified them all – and once we can quantify them – we can design a loudspeaker with any desired frequency range and a flat frequency response in that entire range.
We will also indentify the trade-offs involved in such a design.
We will along the way also find methods to predict the sensitivity and the directivity of the loudspeaker.

9.2 Frequency Response with Voltage Drive and Current Drive

Voltage Drive
A large but slowly diminishing majority of ESL builders believes that the stator electrodes of an ESL must be voltage driven, just like in Figure 4.
By *voltage drive* we mean that the drive *voltage* is independent of frequency.

We will see in this chapter that the SPL at some distance from the loudspeaker depends not on the drive voltage but on the drive current.

As the input impedance of an ESL is almost a pure capacitance[22], a constant (frequency independent) drive voltage results an increase of the drive *current* with frequency at a rate of 6 dB/oct. Therefore also the SPL increases at the same rate.

[22] The input impedance of the loudspeaker can of course not be purely capacitive or else it would be impossible to deliver any electrical power to its terminals.
As a result it could not radiate any acoustical power.
The deviations from a pure capacitance are small, wich is due to the small transduction coefficient of an ESL.
A more accurate description of the input impedance appears in [2, Baxandall] and in [3, Verwaal].

Walker’s Equation and Current Drive
In 1980, Peter Walker[23] published a paper [1, Walker] in which he presented a refreshingly simple way of calculating the sound pressure of an ESL and how it depends on frequency.

Walker arrives at the conclusion that the frequency response of an ESL *on the axis* of the loudspeaker at *a sufficiently large distance* is flat when it is *current driven*[24].
By *current driven* we mean that the drive *current* is independent of frequency.
As the input impedance of the ESL is still capacitive, the result of driving the loudspeaker with a constant *current* is that the drive *voltage* drops with frequency at a rate of 6 dB/oct.

The consequence of the above is that a voltage driven ESL will have a far field frequency response that rises at a rate of 6 dB/oct (Figure 9).

When we drive the loudspeaker with a voltage that drops with frequency at a rate of 6 dB/oct, the current will be frequency independent and the far field axial frequency response will be flat.

The equation that describes the relation between drive current and sound pressure is known as *Walker's equation*.
Those who are not scared off by an equation can find Walker’s equation in appendix II, along with a sketch of how he derived it.

---

[23] Peter Walker was the founder of Quad, famous British manufacturer of audio equipment.
He is also the designer of the famous Quad ESL57 and ESL63 electrostatic loudspeakers.

[24] Peter Baxandall [2, Baxandall] discusses this in some more depth than Peter Walker did.

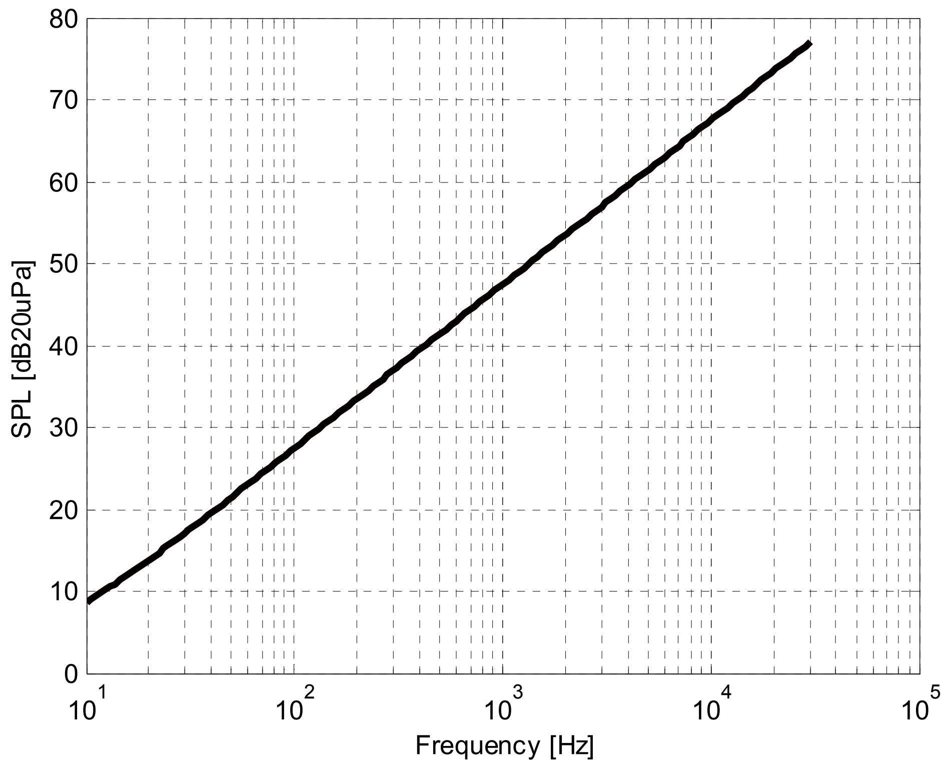

Figure 9 Axial Far field frequency response of a voltage driven ESL.

Far Field Conditions
Walker's equation is valid under a certain condition.
That condition is that the distances from the listening position to all points on the diaphragm are approximately equal[25]. See Figure 10.
Under this condition, the sound from all points of the diaphragm reaches the listener with the same path delay and therefore with the same phase.
When this condition is met, we are in the *far field*.
As you can see from Figure 10, the far field condition is fulfilled if the listening distance is sufficiently large and/or the diaphragm is sufficiently small.

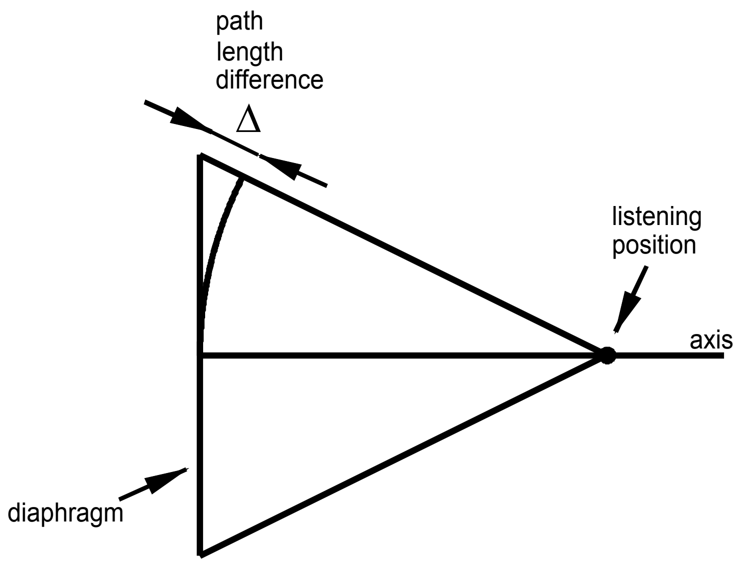

Figure 10 **Different distances to different parts of the diaphragm.**

In order to decide if the path length differences are sufficiently small, we must compare them to the wave length of the sound wave.
The differences are sufficiently small if they are a tiny fraction of a wave length[26].
The wavelength of a 20 Hz bass tone is about 17 m; the wavelength of a 20 kHz tone is about 17 mm.
We see, therefore, that for Walker's relation to be valid at high frequencies you need a much larger listening distance and/or a much smaller diaphragm than at low frequencies.
We will come back to this in sections 9.6 of the current chapter and in chapter 12.

[25] Appendix II explains why this condition must be met.
In short: if we replace the listener by a point source, the distance must be long enough that the sound waves from the source are approximately flat at the location of the loudspeaker.
Under that condition these waves make the diaphragm of the ESL move in an uniform way.

[26] A good estimate would be 1/4 to 1/8 of the wave length.

To get a rough estimate of whether or not the far field condition is fulfilled in a particular situation, we could take as a criterion that the path length differences must be smaller than a quarter of a wavelength. Two contributions from different locations on the diaphragm having such a path length difference will arrive at the listener with a phase difference of 90^o. At such a phase difference, their sum will be 3 dB smaller than it would have been if they had been in phase.
While it is not difficult to derive a formula for the path length difference given the size of the loudspeaker and the listening distance, you could also make a drawing to scale, similar to the one in Figure 10, and simply measure the path length difference Δ.

Acoustic Short Circuit
Unlike most electrodynamic loudspeakers, which usually have their drivers mounted in a closed or semi-closed box, the ESL of for example Figure 5 is an open design.
An ESL can, in principle, also be of a closed design.
For reasons that would carry us too far, however, a closed ESL design does not work out well.

The open design of an ESL brings the following problem:
The sound pressure emanating from the front and the back of the ESL have opposite phases[27].
This gives rise to the question of whether the sound coming from the back side can travel around the loudspeaker and cancel the sound radiated from the front.
This is sometimes called *acoustic short circuit*[28] or *phase cancellation*[29].

The answer to this question is that this acoustic short circuit indeed happens, and hat Walker's relation already takes its effect into account.
Appendix II explains why this is so.
In fact, if we would add a baffle to the loudspeaker to reduce the effects of the short circuit, Walker's equation would no longer be valid[30].

[27] We call this a *dipole* loudspeaker or a *doublet* loudspeaker.

[28] In fact, the path from back to front has a certain acoustic impedance because of the surprisingly large mass of the air that has to move from back to front.
Therefore, the short circuit is not really a short circuit.

[29] We do not like the term *phase cancellation*, because another effect, which we will discuss shortly, can also be called phase cancellation.
This second effect is the cancellation due to different path lengths.
A term that can mean two things is ambiguous and should be avoided.

[30] We would get a larger bass response than Peter Walker's relation predicts.

Example of how Walker's Relation can be used

Suppose that we have a loudspeaker with a distance between the stator plates of 3 mm.
This means that the distance from each stator to the diaphragm is 1.5 mm.
Suppose that the polarizing voltage is 5000 V.
Further suppose that the drive current is 1 mA_{rms}.
Then Walker's relation (appendix II) predicts that the sound pressure on the axis at a distance of 3 m is 0.515 Pa_{rms}.
This corresponds to an SPL of 88.2 dB_{20uPa}.
According to the table in chapter 6 this SPL is slightly more than that of moderately loud classical music and slightly less than that of a heavy truck.

Discussion

We see from the above example that Walker's relation allows us to predict not only the frequency response but also how loud the loudspeaker will sound.

The Paradox of the Small Loudspeaker

Surprisingly, we did in the above calculation we not use the size of the loudspeaker.
Indeed, as also Peter Walker points out, the size of the loudspeaker does not enter into the equation. Neither does the frequency.
The weird conclusion is that a loudspeaker of 10 cm × 10 cm produces *at the same drive current* the same sound pressure as one of 1 m × 1 m, even at bass frequencies.
This seems like a paradox.

The solution to that paradox is that the two loudspeakers indeed produce the same SPL *provided that they are driven by the same current*.

It is not possible to maintain at bass frequencies in the small loudspeaker the same drive current as in the large one. The reason is as follows.

The impedance of both loudspeakers is almost purely capacitive.
The smaller loudspeaker has a 100× smaller area than the big one, so its stator capacitance is also 100× smaller. Therefore its input impedance is 100× larger.
In order to get the same drive current to flow into it, we need a 100× larger drive voltage.

If, for example, the big loudspeaker needs 1000 V_{rms} to get the current to flow into it, the small loudspeaker needs 100 kV_{rms}.
Obviously the impedance, and therefore also the voltage needed to get the current to flow into the ESL, increase towards the bass frequencies.
This is why the sound pressure of the small loudspeaker is limited in the bass range.

In the end, therefore, the world is as we intuitively know how it is should be: we need a large diaphragm area to get a decent bass range at a decent SPL.

Reduction of SPL with Distance

Walker's equation predicts that the SPL drops by 6 dB each time the distance doubles.

If you think about it, this makes sense.
At a sufficiently large distance the waves of the loudspeaker are spherical with the loudspeaker in the center of the sphere. As the waves spread out, the radius of the sphere increases. The acoustic power that the loudspeaker radiates distributes over the area of the sphere. That area quadruples as the radius doubles.
Therefore the power density (in Watt per square meter) is divided by four as the distance doubles. That reduction of the power density corresponds to a reduction of the sound pressure by a factor two.
This is a reduction of SPL by 6 dB.
See also section 7.4.

9.3 Realization of Current Drive

To realize a constant (that is: frequency independent) drive current, we can consider the use of an amplifier having a current output.
The output of such an amplifier behaves as an ideal current source.
Connecting the stators of the loudspeaker to an ideal current source, however, creates the problem that the dc voltage across the ESL is not fixed.
This means that a pure current drive is not possible: for low frequencies the current source behavior must always change into a voltage source behavior in order to fix the dc voltage[31].

A practical way to realize a constant drive current at audio frequencies is to take a high voltage audio source and place a large resistor R in series.
The voltage source could be a normal audio amplifier with a voltage output, followed by a step-up transformer, or it could be a custom designed high voltage amplifier.
In practice we split the series resistor into two resistors, each with value $R/2$, to preserve the symmetry of the circuit.
Figure 11 shows the circuit arrangement.

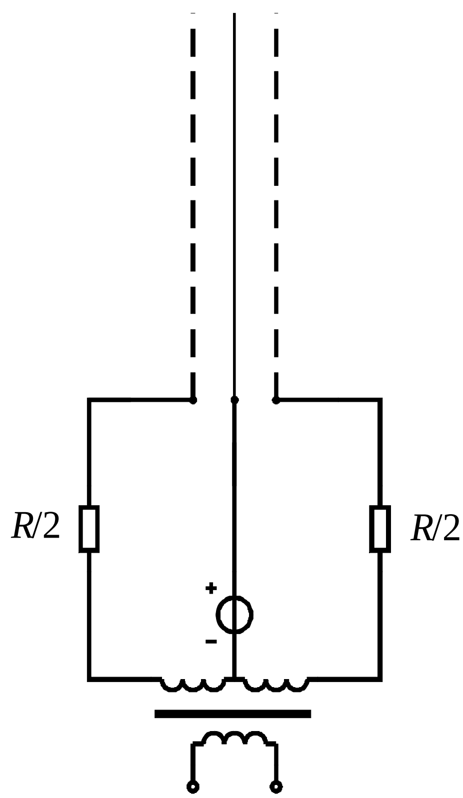

Figure 11 ESL with step-up transformer and series resistors R/2.

[31] But every disadvantage comes with an advantage: we do not have to design a special amplifier with a current output.

For audio frequencies, the impedance of the loudspeaker, which is almost purely capacitive, is low enough that the current is completely determined by the transformer output voltage and the total series resistance R.

For low frequencies, the impedance of the ESL increases to the point that it becomes larger than the resistance R. Below that frequency the total impedance doubles with each octave that the frequency drops. Therefore the drive current rolls off at a rate of 6 dB/oct and a constant drive current can at those low frequencies no longer be maintained. Due to the decreasing drive current, the frequency response of the loudspeaker drops as the frequency decreases.

Appendix III gives the formula that describes the -3 dB corner frequency of the roll off.

We try to design the loudspeaker such that the corner lies below the lowest frequency we want the ESL to reproduce.

At dc, the plates connect through a resistance to the step-up transformer and that nicely fixes the dc voltage across the plates to zero volt.

Example.
Suppose that the amplifier delivers a voltage of 10 V_{rms} and that the step-up ratio of the transformer is 1:100. The transformer output voltage is then 1000 V_{rms}.
If the total series resistance is for example 1 MΩ (2×500 kΩ), the drive current for a sufficiently high frequency is 1000 V/1 MΩ = 1 mA_{rms}.

Suppose further that the area of the loudspeaker is 1 m^2 and that the distance between the stator plates is 3 mm.
Disregarding the holes in the plates the stator capacitance is then 3 nF. You will find the formula that you need to calculate the stator capacitance in appendix IV.
A large portion of the stator area, say 50%, consists of holes.
The metal area lost to these holes leads to a reduction of the capacitance, but the reduction is less than proportional to the lost area.
The actual capacitance will therefore be somewhere between 1.5 nF and 3 nF.
Let’s assume that the stator capacitance ends up being 2 nF.
The frequency where the drive current starts to roll off is determined by the RC product and is 80 Hz. The formula is in appendix III.

2-way or 3-way design
We see here that a reasonable bass response is possible with a loudspeaker of 1 m^2
This size is too large, however, to satisfy the requirement of sufficiently small path length differences up to the highest audio frequencies.
For that we need a smaller loudspeaker size of perhaps 7 cm × 7 cm
[2,Baxandall, page 146].
These two conflicting requirements can be reconciled by using a two- or three-way design.
We will discuss this in more detail in section 12.2.
Other solutions will be discussed in chapter 13.

The Trade-off between size, bass response and SPL

A loudspeaker of 1m^2 is a large structure to have in your living room.

In order to improve the WAF[32] we could reduce the size of the loudspeaker but this will be at the expense of the bass frequency range.

If, for example, we half the area (so it will be approximately 70 cm $\times$ 70 cm), we also half the stator capacitance. Therefore the bass range, which was originally 80 Hz, will now be limited to 160 Hz.

The SPL above 160 Hz does not change.

We can get the original range of 80 Hz back by doubling the value of the series resistance so the RC product is again what it was before.

This, however, halves the drive current so we will loose 6 dB of SPL.

We can get the original SPL back by doubling the ratio of the step-up transformer to 1:200.

There are, however, practical limitations on the feasible step-up ratio.

This will be discussed in chapter 15.

In short: we see that there is a trade-off between bass frequency range, loudspeaker size, SPL and drive voltage.

[32] WAF = Wife Acceptance Factor.

9.4 Off-Axis Response – Directivity of a Dipole Loudspeaker

We have seen that Walker's equation allows us to predict the sound pressure and the frequency response of an ESL on the axis of the loudspeaker.
We now want to do the same for off axis listening positions.

We want to be able to predict the off axis response first of all because it determines the tone balance of the reverberation field.
But also, as we will see later, being able to predict the off axis response of a small loudspeaker is a necessary step towards calculating the response for loudspeakers that are too large for the path length differences to be negligible.
That is: it is a first step towards calculating the response at closer distances than the far field distance.

First of all, with off-axis listening positions, the maximum loudspeaker size for which the path length differences are negligible is now much reduced, as we can see from Figure 12
Further, the path length differences no longer vanish with increasing listening distance, as we can also see from the figure.

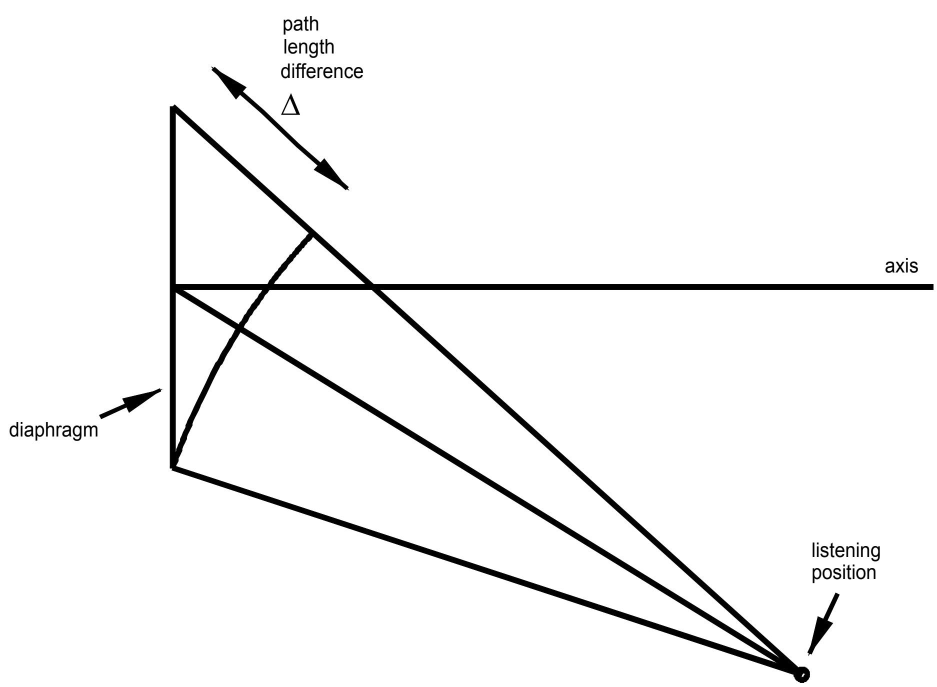

Figure 12 Path length differences with off-axis listening position.

We now assume that the loudspeaker is so small that path length differences are negligible.

As it turns out, we can easily adapt Walker's equation such that it is valid for off axis listening positions as well.
Detailed analysis shows that we just have to multiply the sound pressure that the equation predicts by a number that depends on the angle that the listening position makes with the loudspeaker axis[33]. This number is smaller than unity for all off-axis listening positions and drops with increasing angle.
Appendix V gives a sketch of the reasons why this is so.
Figure 13 shows a plot of the reduction of the sound pressure for off-axis listening positions.

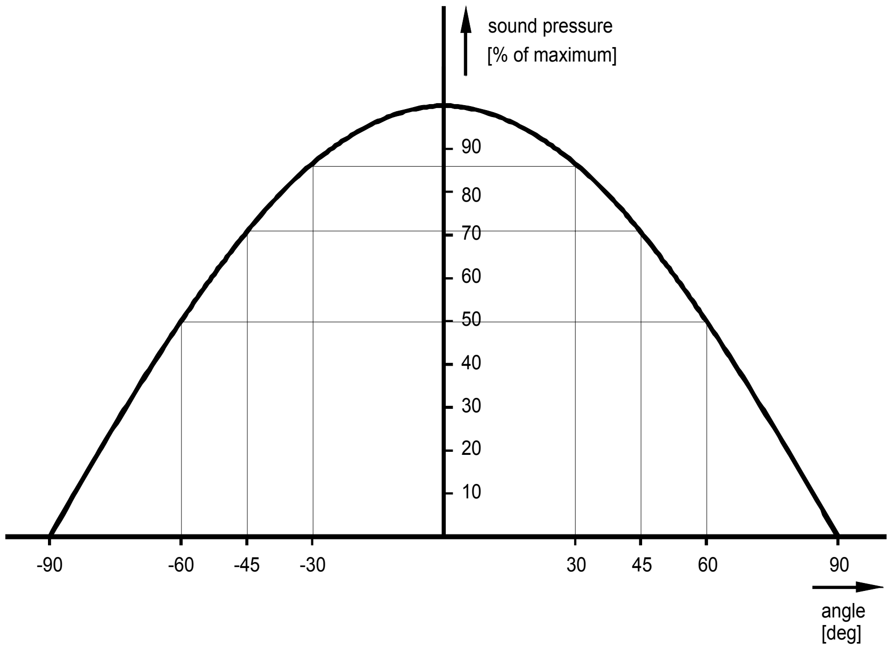

Figure 13 **Reduction of sound pressure for off-axis listening positions as a function of the angle that the listening position makes with the axis of the loudspeaker.**

As we can see from Figure 13, the sound pressure drops to zero at an angle of 90^0. This means that by the side of the loudspeaker, in the plane of the diaphragm, the sound pressure is zero. This makes sense if we consider the symmetry of the situation.
The particle velocity, however, is at those locations nonzero as we can see by looking again at the symmetry.

Because the human ear is mainly sensitive to pressure, we expect to hear almost nothing when we place ourselves to the side of the loudspeaker.
When you try this in a typical listening room you will find that you still hear the sound from the loudspeaker. What you hear is the sound bouncing off the walls. We call this the *reverberation field*. To really verify that the SPL of the direct sound at that location is zero you would need to do the experiment outside or in a room which is free of reflections (an *anechoic chamber* or *dead room*).

[33] In fact, it is the *cosine* of that angle.

If we convert the reduction factor shown in Figure 13 to dB, we get the picture of Figure 14.

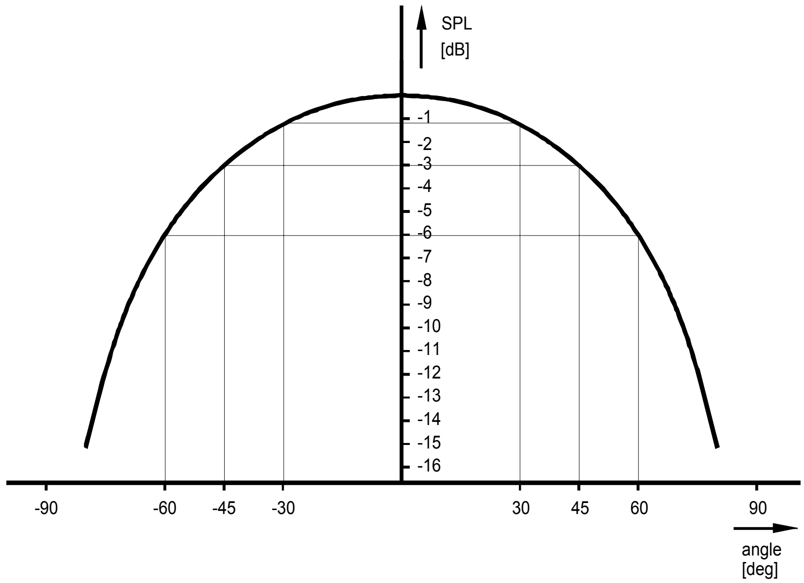

Figure 14 Reduction of SPL at off-axis listening positions in dB.

If we plot the SPL in a polar plot as a function of the angle, we get a diagram known as the *directivity plot*.
The directivity plot of the ESL has the shape of the figure of eight (see Figure 15).

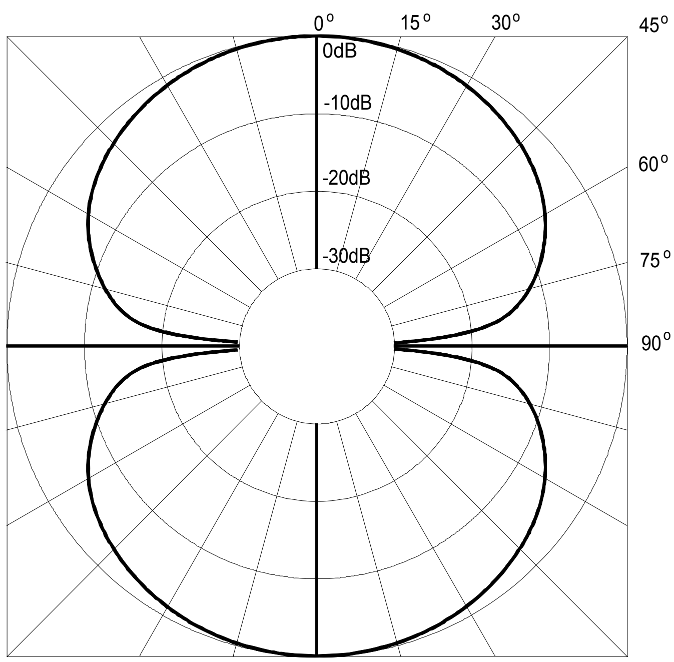

Figure 15 Directivity plot of a dipole loudspeaker (at low frequencies).

9.5 Proximity Effect

The proximity effect is a phenomenon that increases the bass response of an ESL at close listening distance.
At the typical distances we encounter in a practical listening situation (a few meters), it is a small effect.
We describe it here for completeness, and because it is of course not negligible when we are dealing with electrostatic headphones.

Appendix VII gives a sketch of the mechanisms that give rise to the proximity effect.
See also section 7.4.

Figure 16 shows the response of a small loudspeaker at a various distances.
At 3 m distance the frequency response starts to increase below 18 Hz, which is too low to be of any help for the bass response.

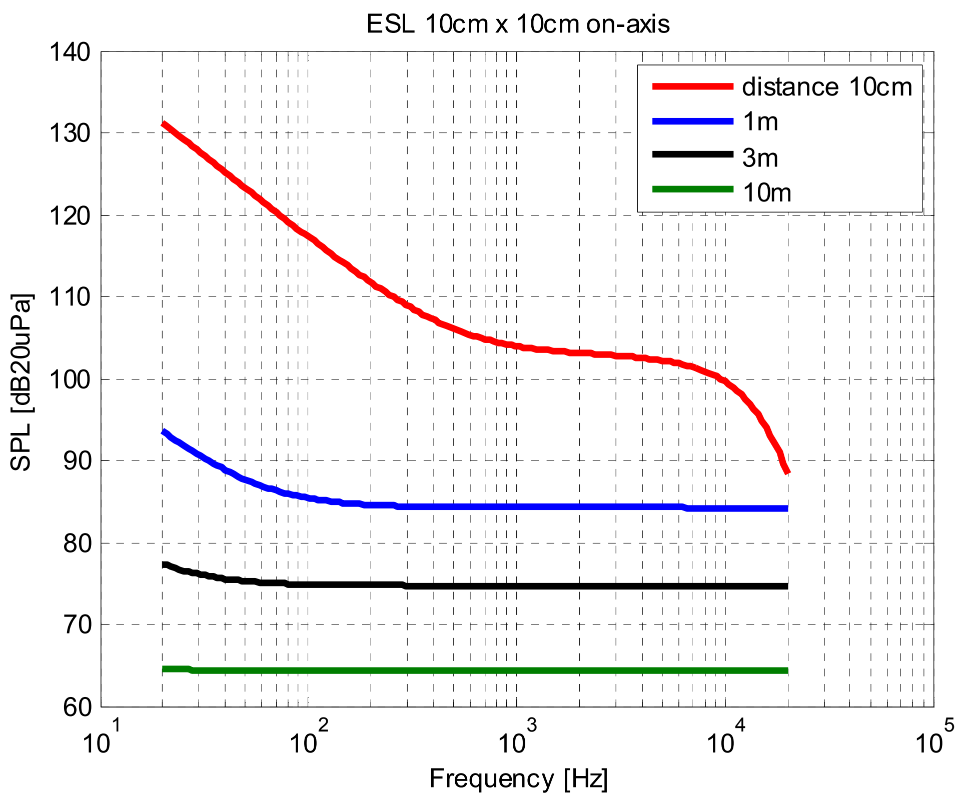

Figure 16 Proximity effect.

9.6 Generalizing Walker's Equation

The Limitation to Small Loudspeaker Size

Peter Walker's relation is valid for listening positions on the axis of the loudspeaker at a distance that is large enough to make the path length differences to all points of the diaphragm negligible.

To decide if the path length differences are negligible, we must compare them to the wave length. Therefore, if we want to use Walker's equation, we need for the high frequencies a much smaller diaphragm, or a much larger listening distance, than at bass frequencies.

In short: Walker's equation is limited to on-axis listening positions and to either large listening distances or (very) small loudspeakers.

In the previous section we already extended Walker's equation to off axis listening positions.

In this case, the path length differences do not vanish with increasing distance so a large distance does not help.

For off axis listening positions, therefore, Walker's relation is limited to (very) small loudspeakers, regardless of the distance.

The use of Superposition to Calculate the Response of a Large ESL

In this section we want to generalize Walker's equation further, so that it can predict the response in any point in space, both axial and off axis, and both far away and close by.

To achieve this goal, we must get rid of the condition that the path length differences are negligible. For this purpose we can use the *superposition theorem*.

Appendix VI discusses that theorem in some more depth.

To calculate the sound pressure of a large ESL in an arbitrary point in space we divide in our mind its surface area into small sections, each of them small enough that it satisfies the conditions under which Walker's equation is valid[34].

We will call these small pieces of area *surface elements*.

Superposition now says that we can find the total sound pressure at any desired listening position in space by adding[35] the individual contributions from all surface elements.

This route to a generalized version of Walker's equation was also suggested by Peter Baxandall [2, Baxandall, page 141, 142).

---

[34] The approximation becomes exact in the limit case where the size of the surface elements becomes vanishingly small.

[35] For those readers who are familiar with calculus: in the limit case where the size of the surface elements becomes vanishingly small, the summation changes into an integral.

Because the listening position is for most surface elements not on the axis of that element, we must for each surface element take the angle between the listening position and the axis of the surface element into account as we did in section 9.4.
Further, because the distances from the surface elements to the listening position are not equal, we must take the corresponding delay and the associated phase shift into account.
Finally, because we want to find a relation that is also valid at close range, we must take the proximity effect into account.
If you are interested: appendix VIII shows the details of the calculation.

Walker's original equation is simple enough that we can evaluate it with the use of an electronic calculator, by filling in the quantities such as drive current, distance, polarizing voltage and diaphragm-to-stator distance.
The addition of all contributions from all surface elements, however, is too much to be done by hand, all the more because of the need to keep track of the individual phase differences. It is, however, not difficult to write a computer program to do the calculation. Three such programs are available free of charge from my website [3, Verwaal][36].

The Approximation in the Superposition Approach.
Before we continue, it is worth while to point out an implicit approximation in the above line of reasoning.

Using the results from section 9.4 we can calculate the response of each current driven surface element.
When we re-assemble the whole loudspeaker from its surface elements, we connect their stators in parallel. When we drive the complete ESL by a current, the assumption is that all surface elements take the same portion of the total drive current.
At first sight this seems a reasonable assumption because all surface elements have the same area and therefore have the same capacitance.

We have overlooked, however, that the diaphragm of each surface element moves due to the sound generated by all others. In other words: each surface element acts as a microphone that listens to the sound from all other surface elements.
As there is no reason to suppose that all surface elements will experience the same particle velocity of the air, this will make the distribution of the total available current over all elements unequal.

[36] The Pascal program *ESL.pas* (compiled to *ESL*.exe) is limited to one rectangular panel and does not take the influence of the cross-over filter into account.
The Matlab script *ESL*.m can simulate a loudspeaker consisting of several panels, driven from a cross-over filter. It allows some freedom in the relative placement of the panels.
The Matlab script *ESL_ladder*.m does the same as *ESL*.m, but with a different type of cross-over filter.

Having said that, we must recognize that the redistribution of current does not occur in the absence of a polarizing voltage, as in that case the capacitor microphones do not work.
As we crank up the polarizing voltage, they begin to work and they work more efficiently as we increase the voltage. As it turns out, however, the effect is still very small by the time we reach the maximum polarizing voltage that the breakdown field strength of air will allow. Even at the maximum practical polarizing voltage, the current of a surface element is dominated by the stator capacitance, not by the movements of its diaphragm. The approximation we unknowingly made therefore turns out to be justified[37].

Sanity Check
The generalized Walker equation is valid in all points in space.
As a special case, it is therefore also valid in very close proximity, say one millimeter, from the diaphragm.
In chapter 10 we present a very direct way of calculating the sound pressure in those locations. The simplicity of that line of reasoning makes it almost immune for errors.
We can now verify that the modified Walker equation predicts with very good accuracy the same sound pressure. This is a good sanity check.
The modified Walker equation splendidly passes this test.

Many Spherical Waves can add up to a Flat Wave (or a Cylinder Shaped Wave)
It increases our insight to make the following observation.
As all the surface elements are infinitely small, they all produce spherical waves.
We know, however, that the waves of a flat ESL in the direct vicinity of the diaphragm are flat.
Apparently, the infinite number of spherical waves combine (super-impose) in such a way that a flat wave results.

At first sight this might seem like an unexpected coincidence.
On second thought, however, we realize that the surface elements are arranged symmetrically (i.e. as an infinite array in a flat plane) and therefore they can only add up to a wave that has the same symmetry.

In a similar way the spherical waves of all the surface elements of a line source combine to form a cylinder shaped wave.

Note, finally, that when the individual spherical waves combine, they do so taking into account their relative magnitudes and phases as determined by their path length differences. These magnitudes and phases are also influenced by the proximity effect, which we discussed in section 9.5.
All these effects work together in a subtle way to produce a flat wave or a spherical wave, depending on the type of source.

---

[37] What we have said here comes down to recognizing that an ESL has a low *transductioncoefficient*. Compare this to an electro-dynamic loudspeaker with such a weak magnet that its input impedance is completely dominated by the speech coil resistance and not by the velocity of the cone.
See [2, Baxandall] and [3, Verwaal] for a discussion of transduction factors.

9.7 *Phase Cancellation due to Path Length Differences*

We consider now an ESL with an on-axis listening position and we consider the way that the contributions from all surface elements add up.
Due to the different distances that the surface elements have to the listening position, their contributions have different phases.
When the phase differences are small (at a low frequency or at small path length differences) they add up constructively.
In this case, the modified Walker equation must of course give the same result as Walker's original equation[38].

When the phase differences become large (at closer distance and/or higher frequency) some contributions may have phases opposite to some others. They then combine destructively and they may partly cancel.
We could call this *phase cancellation*[39] or *destructive interference*.

The path length differences tend to produce cancellation at high frequencies and therefore we can expect in a large loudspeaker at a short listening distance a reduction of the high frequency response.

Comb Filter

We have now seen that the response of a large ESL is the sum of many contributions, each having a different delay. This is similar to a *comb filter*, which we will therefore now discuss briefly.

If you add a signal to a delayed copy of that signal, you get what is known as a *comb filter*. Figure 17 gives a block diagram of such a filter. At each frequency where the delay is equal to half the cycle time plus an integer multiple of it, the two contributions cancel. As a result the frequency response has notches at regular frequency intervals.
Hence the name "comb filter".
Figure 18 shows the frequency response of such a filter on a logarithmic frequency scale.

---

[38] This is another sanity check for the modified Walker's equation and again it passes the test.

[39] The term phase cancellation is also used to describe the acoustic short circuit, see section 9.2. A term that can mean two different things is confusing and therefore we will avoid using it.

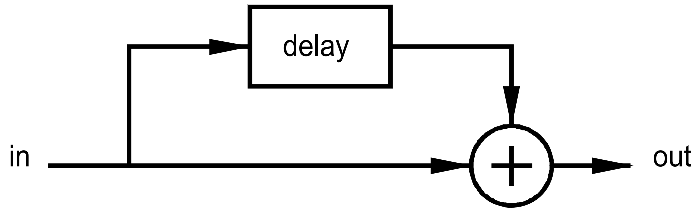

Figure 17 **Comb filter.**

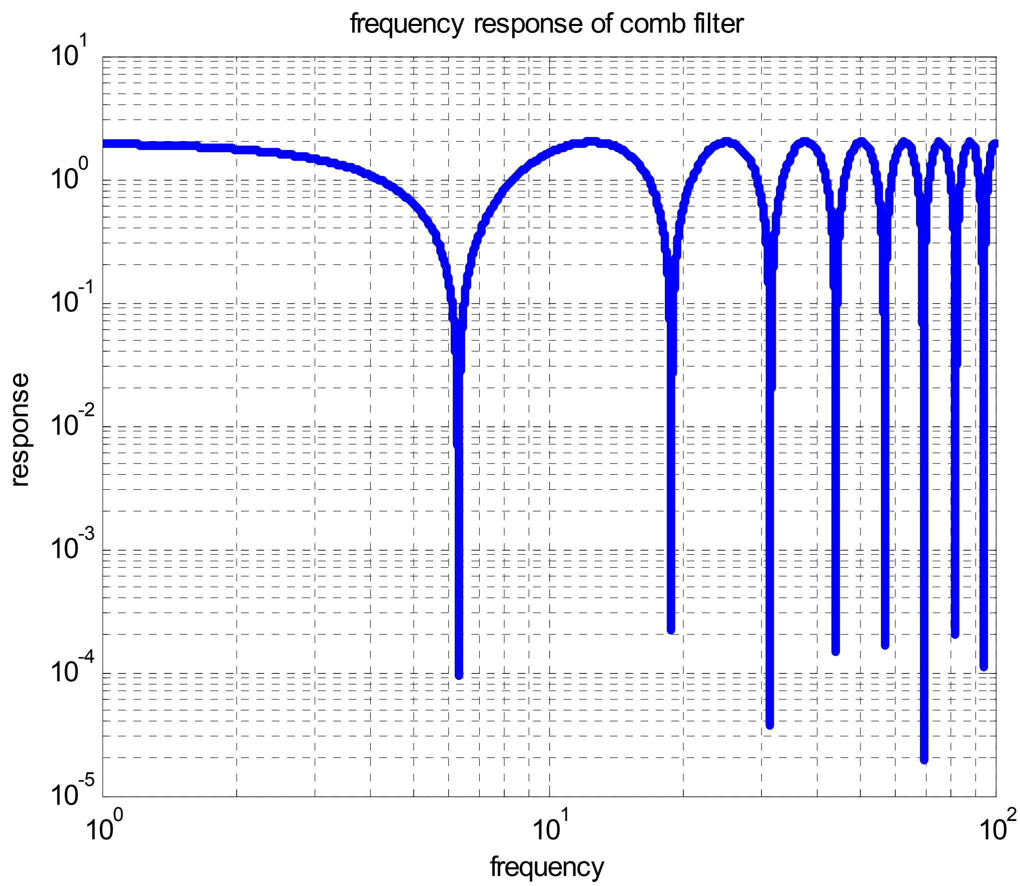

Figure 18 **Frequency response of a comb filter.**

Undulations of the High Frequency ESL Response due to Path Length Differences
Consider now the on-axis frequency response of an ESL.
An ESL that is large enough for path length differences to become important is similar to a comb filter, except that this time a multitude of delayed signals add up instead of just two.
We can therefore expect at high frequencies a sort of comb filter response with undulations between certain minima and maxima.
Further, because some contributions add destructively instead of constructively, we expect a reduction of the high frequency response.

As an illustration, Figure 19 shows the axial sound pressure versus frequency of an ESL of 1 m × 1 m at distances of 1 m, 3 m, 10 m and 30 m[40].
You clearly see the effects of the path length differences at high frequencies.
You also see that the effect of the path length differences becomes smaller as the listening distance increases.

For comparison, Figure 20 shows the same curves for a loudspeaker of 10 cm × 10 cm. You see that even at the shortest distance of 1 m and the highest frequency of 20 kHz the effect of cancellation due to path length differences is hardly noticeable.

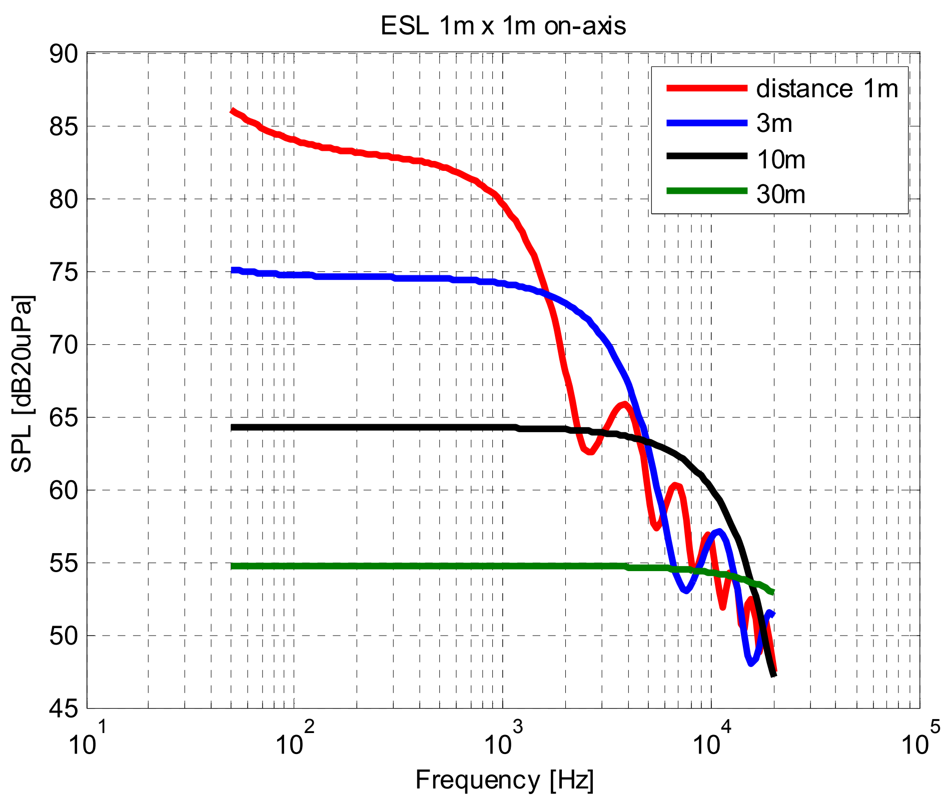

Figure 19 On-axis SPL of an ESL of 1 m x 1 m at distances of 1, 3, 10 and 30 m

[40] Like all frequency response plots in this book, this is simulated using the modified Walker's equation.

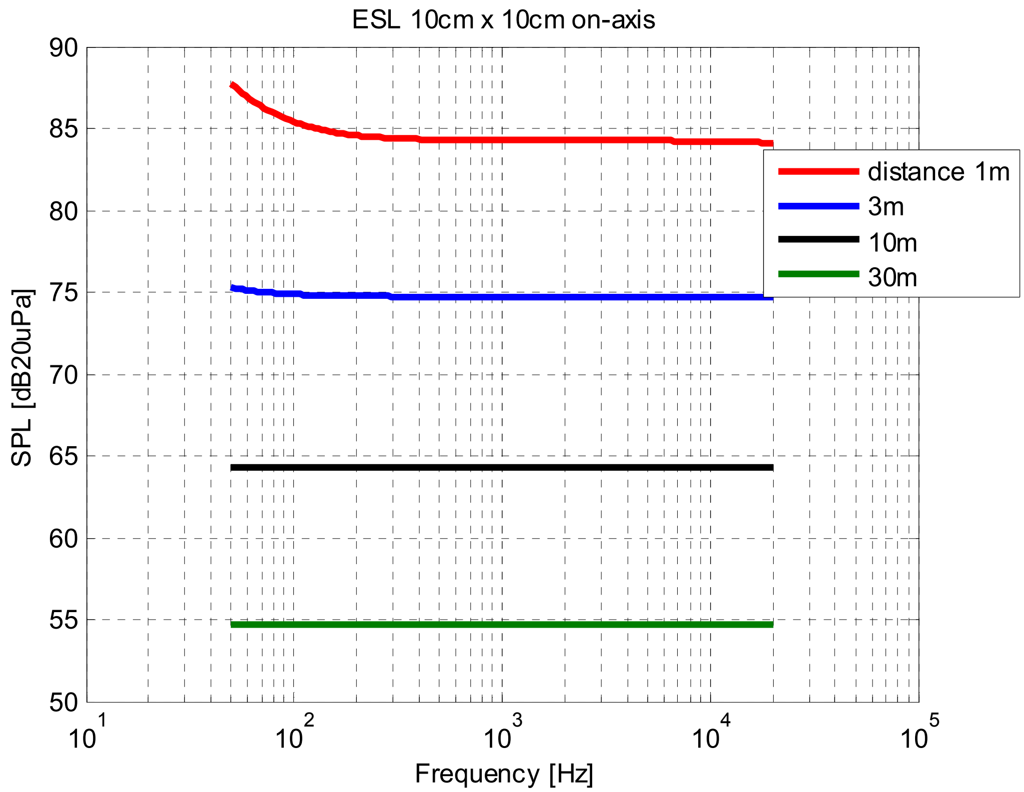

Figure 20 **SPL of an EXL of 10 cm x 10 cm at distances 1 m, 3 m, 10 m and 30 m.**

Discussion – Two- or Three-way Design

We have now seen that path length differences tend to reduce the frequency response at high frequencies and also cause at the high side of the audio spectrum undulations of the frequency response. As both effects are undesired, we must find design solutions that avoid these artifacts. One solution is to try and avoid the path length differences. This approach will be explored in chapter 12.
There are also other solutions; we will discuss these in chapter 13.

9.8 Off-axis Path Length Differences - Directivity

When we move the listening position away from the axis, we can of course still use the generalized Walker equation to predict the response of the loudspeaker.

Off-axis listening positions have larger path length differences than on-axis positions. Not only are the differences larger; they no longer vanish with increasing distance either. As a result, the high frequency response at an angle from the axis is reduced.
As an illustration, Figure 21 shows the frequency response of a loudspeaker of 30 cm $\times$ 30 cm at 3 m distance on the axis and at 3 m distance under a 45^o angle.

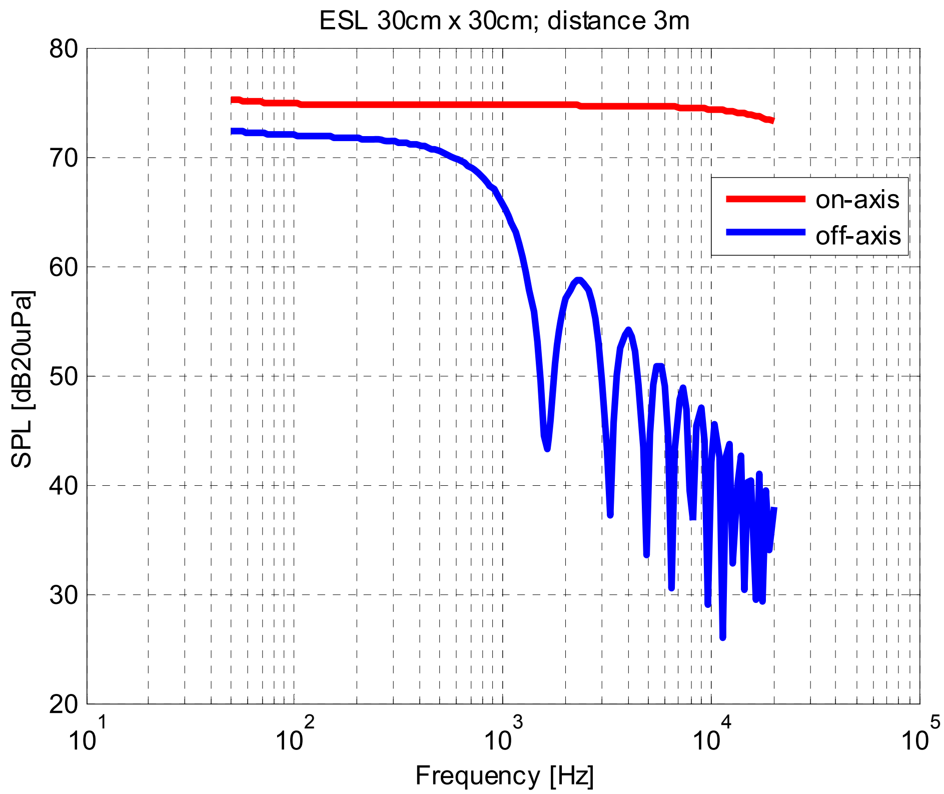

Figure 21 **Frequency response of an ESL of 30 cm x 30 cm at a distance of 3 m - on-axis and off-axis under 45^o.**

This shows that when the frequency increases, path length differences increase the *directivity* of the loudspeaker.
Note that the reduction of the off axis SPL at low frequencies, as shown by the plot, is due to the (much smaller) low frequency directivity of a dipole source, as already discussed in section 9.4.

Figure 15 shows the directivity pattern of an ESL at low frequency, which has the familiar figure of eight shape.
By contrast, Figure 22 shows the pattern of the same loudspeaker at a higher frequency.

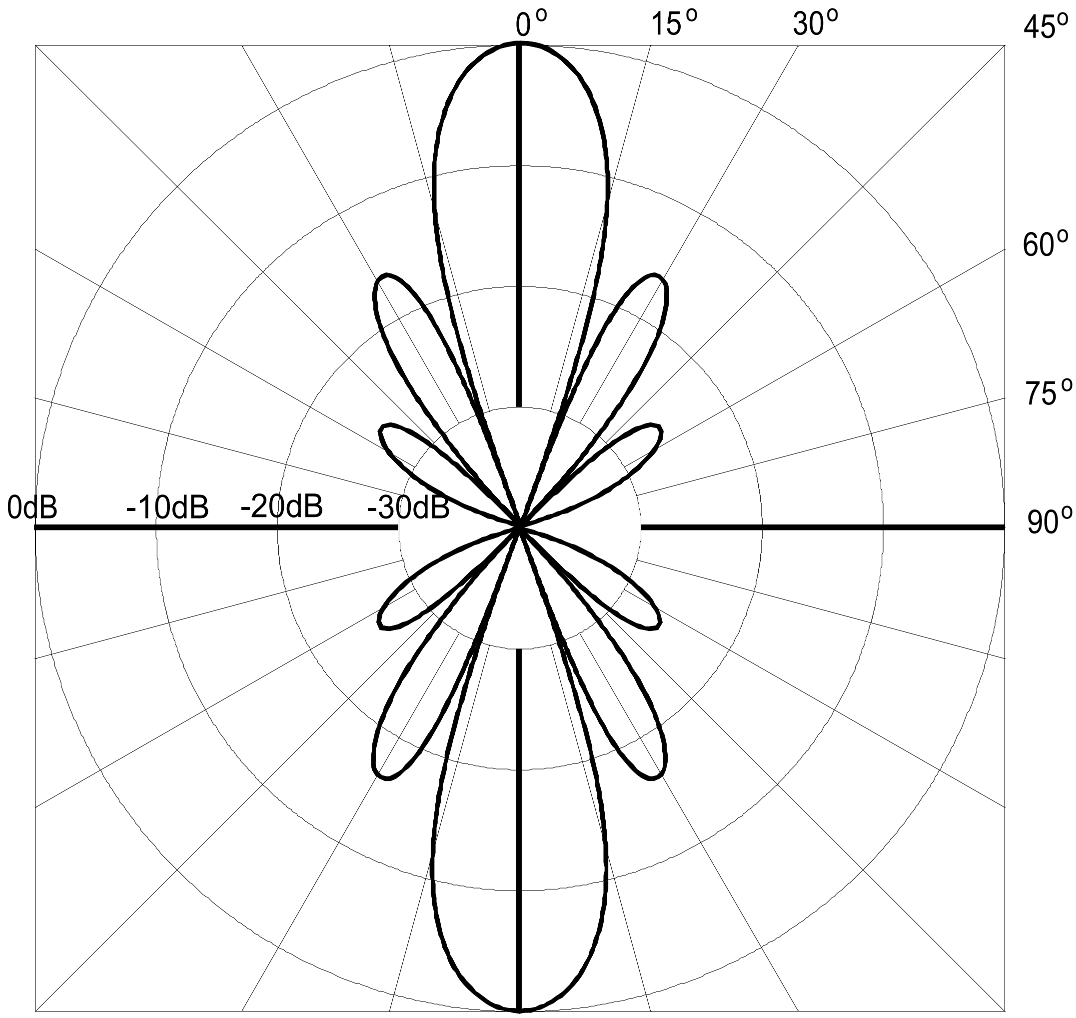

Figure 22 **Directivity pattern of an ESL of 1 m x 1 m at 1 kHz.**

If this already happens at a frequency of 1 kHz, we can imagine that at 5, 10 or 20 kHz the effect will be even more dramatic.
If the diaphragm size is too large, an ESL can be almost as beamy as a flash light.
While a certain amount of directivity can be beneficial (see section 11.3), too much of it is undesirable.
For this reason, the tweeter part of an ESL must not be too large.
Baxandall says that the optimum size is about 7 cm × 7 cm [2, Baxandall, page 146][41].

[41] In fact, Baxandall speaks of a round tweeter panel with a 7cm diameter.

9.9 Diaphragm Mass

We have assumed all along that the mass of the diaphragm is negligible, and that is also a necessary condition for Walker's equation to be valid.

That assumption is only justified if the diaphragm is sufficiently thin.
The diaphragm is often made of polyester (brand names: for example *Mylar, Melinex* or *Hostaphan*), having a mass density of about 1390 kg/m^3.
In that case the maximum thickness that allows us to neglect the mass is 4 μm.

A thickness of 6 μm can be used as well, but it will lead to some loss of response at high frequencies. That must then be corrected by an equalizing filter between the pre-amplifier and the main amplifier.

Figure 23 shows the axial far field response of a current driven loudspeaker with a too heavy diaphragm.

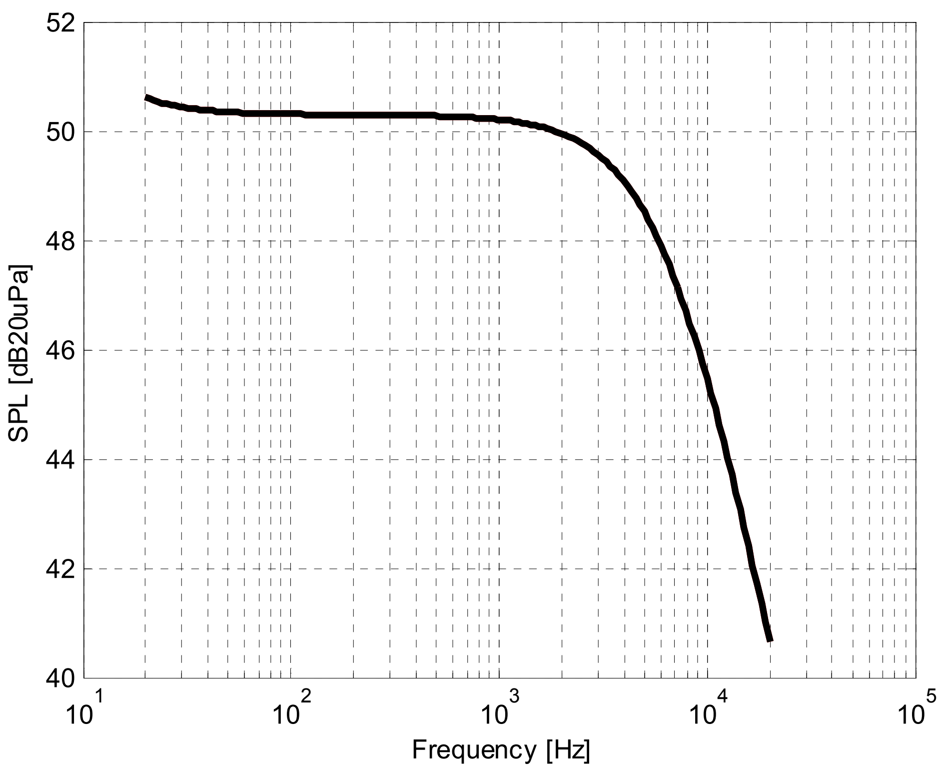

Figure 23 Response of a current driven loudspeaker with a heavy diaphragm.

9.10 Air Mass in Stator Holes

We said before that the stator plates must be perforated to allow the sound generated by the diaphragm to escape.
Walker's equation is valid under the assumption that the stator plates are completely acoustically transparent.

While it is not possible to exactly satisfy that requirement, we can certainly try to approximate it.
For that purpose we must keep the stator plates as thin as possible while having the largest possible fraction of the area covered with holes.
Let's call that fraction the *openness*.

If the stator plate is too thick or the openness is insufficient, the mass of the air in the holes will reduce the high frequency response in a way similar to Figure 23.

When the stator electrodes are made of perforated sheet metal, it is best to choose a thickness of no more than 1 mm.
In the Quad ESL63 the stator electrodes are made of printed circuit board (PCB) material. While standard PCB's have a thickness of 1.6 mm, the ESL63 uses material with a thickness of 1.3 mm.

The holes can be round or square. Round holes can have a diameter of 3 to 5 mm and square holes can have a side of 3 to 5 mm[42].
The openness should be at least 40%.

Some ESL builders construct the stator electrodes as a grid of isolated electricity wires. This has the advantage that the problem of stator isolation is automatically solved.
The diameter of the wires can be 1 mm to 3 mm, corresponding to a conductive cross sectional area of 0.15 mm^2 to 1.5 mm^2.
The distance between the wires can be such that the openness is 50 - 60%.

Some believe that the sound has less difficulty travelling around the round wires than it does travelling through the sharp edged holes in stators made from perforated plates

---

[42] Smaller size holes should be avoided when the stator plates will be painted for isolation, as the paint tends to fill the holes and close them.

9.11 Diaphragm Stability and Diaphragm Bass Resonance

Stability of the Diaphragm Suspension
The charged diaphragm, suspended between the stator plates, has the tendency to move away from the desired quiescent position in the middle of the gap and stick to one of the plates.
You can compare this to an iron nail suspended by a string between two magnets[43].

To keep the diaphragm in its place in the middle of the gap, it must be stretched.
That way the electrical force trying to pull it out of its place is counteracted by a mechanical force that tries to restore it to its intended position.
The restoring force is an elastical one: it is proportional to the excursion.
We can describe the strength of this force by a spring constant (stiffness) or by a compliance constant.

The amount of stretching needed to get a stable diaphragm depends on the polarizing voltage, the distance between the diaphragm and the stator and the unsupported width of the diaphragm.

As a guidance, you can assume that the following combination of parameters will work:

Distance between diaphragm and stator:	1.5 - 2 mm
Polarizing voltage:	max 5000 V
Unsupported diaphragm width:	20 cm
Diaphragm thickness:	4 - 6 μm
Amount of stretching:	1 - 2 %

[43] Metaphor by Roger Sanders [12, Sanders].

Bass Resonance
The stretched diaphragm will show a resonance in the bass region (Figure 24).
The frequency of that resonance is determined by the stiffness of the suspension and the resonating mass. Contrary to popular belief, that mass is not the mass of the diaphragm but the mass of the air that moves with it.
The mass of 1 m^3 of air at room temperature and atmospheric pressure is surprisingly large: about 1.18 kg. The diaphragm mass is very small compared to the air mass.
In vacuum the resonance frequency would be determined by the diaphragm mass and it would be somewhere in the kHz range.
In air the resonance is in the bass range and the frequency depends on the amount of stretching, the unsupported diaphragm width and the size of the loudspeaker.

You can demonstrate this resonance by mounting a stretched diaphragm on a small wooden frame of approx. 20 × 20 cm. You can use heat shrinking to stretch the diaphragm.
If you hold the stretched diaphragm close to your ear and you gently hit the frame with the soft side of your fist, you will hear the bass resonance.

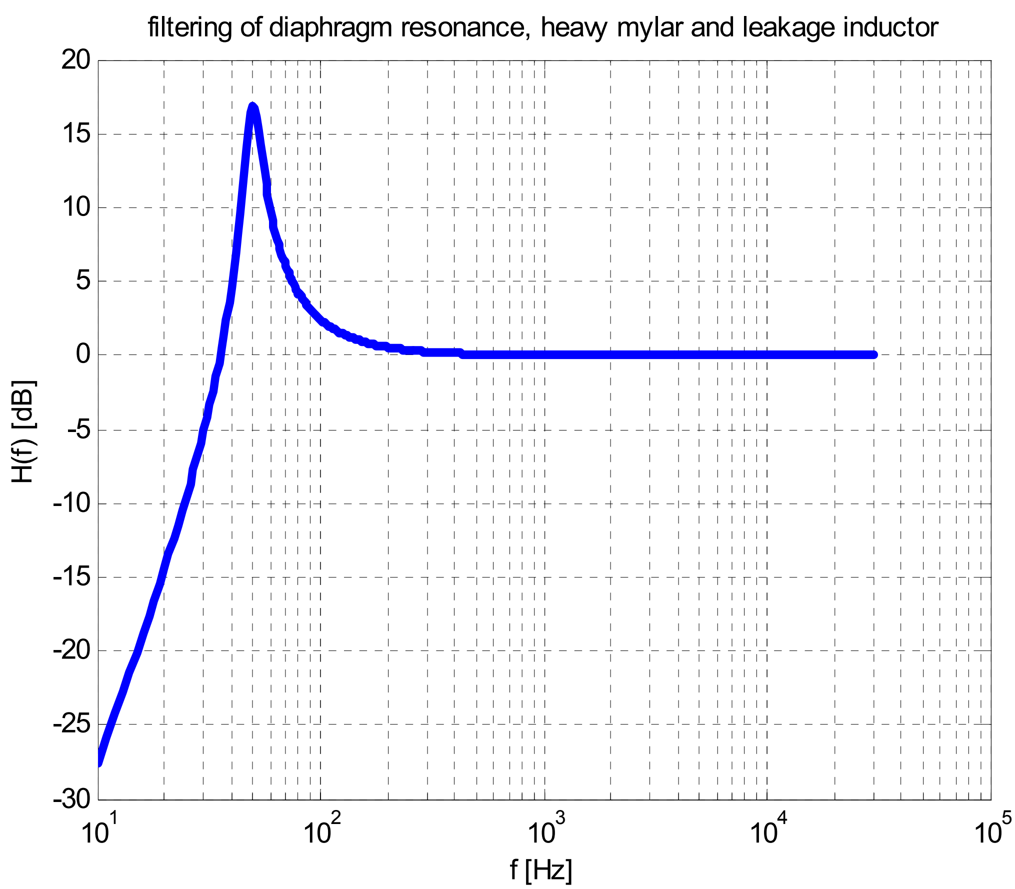

Figure 24 **Peak in the bass response due to poorly damped diaphragm resonance (Q = 7).**

Multiple Resonances

A mass-spring system having only one spring and one concentrated mass has just one resonance frequency.
The diaphragm of the ESL, by contrast, will be supported in many places and therefore behaves as several springs. Further, the resonating air mass is not concentrated in one point but is instead distributed over a certain region of space.
As a result, we will not get one resonance peak in the bass frequency range but several of them.

Damping of the Bass Diaphragm Resonance

Without further measures, the diaphragm resonance can be very pronounced.
It will result in a strong but ugly bass reproduction.
To improve this, the resonance must be damped.
The way to do this is to add a damping cloth (fabric) to the back stator.
The cloth acts as a resistance to the air flow. That resistance converts the resonant energy stored in the mass-spring system into heat and thereby removes it from the system.

The damping cloth does not have to be thick, it just has to be tightly woven to present the right amount of resistance.
It can be glued to the back stator plate, but we must take care that the glue does not cover the part of the cloth over the holes as that would make the cloth air tight.
Applying such a damping cloth is much easier in loudspeakers using perforated plate stators than in loudspeakers using stators made out of stretched wire.

While we can calculate (at least approximately) the required value of the damping resistance, this does us little good because cloth with a specified resistance is not for sale.
We therefore need to experiment to determine the right amount of damping.

The Use of the Diaphragm Resonance to Extend the Bass Frequency Range
It is of course possible to damp the resonance to a point that it is completely gone. We can, however, make good use of the resonance to extend the bass range by approximately one octave. For that purpose, the resonance frequency must be well controlled and the damping must be such that a quality factor (Q-factor) of about two remains. The right amount of diaphragm stretching and the right amount of damping must again be determined by experimentation.

As an example, Figure 25 shows the frequency response of a loudspeaker having a bass roll-off at 100 Hz. The blue trace shows the same loudspeaker with a diaphragm resonance at 50 Hz with a quality factor Q = 2.

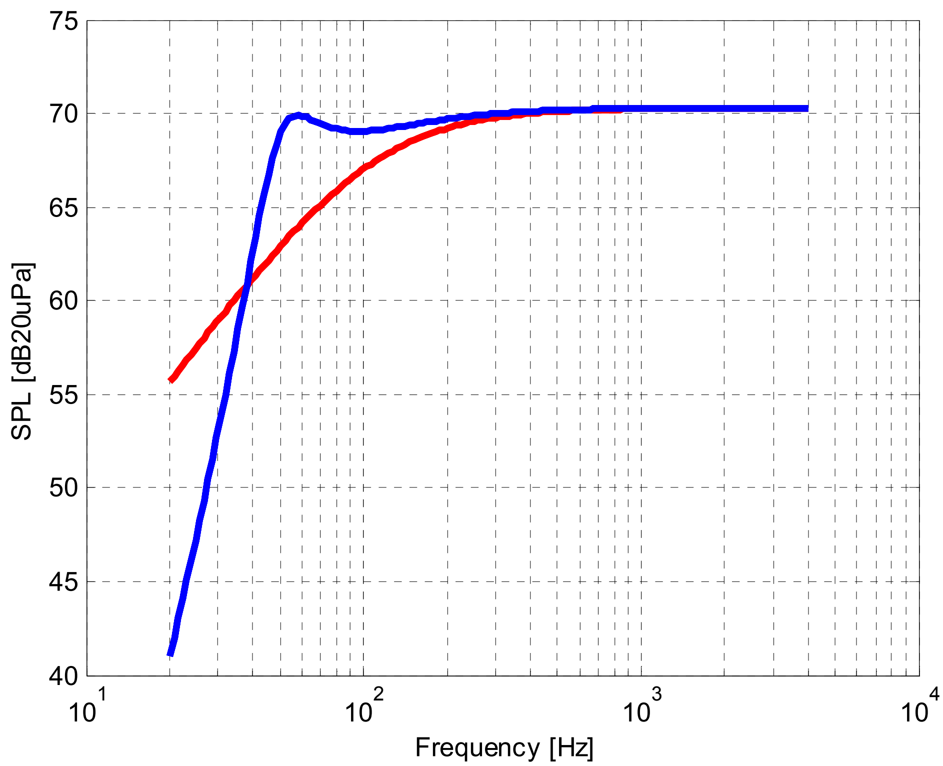

Figure 25 **Extension of the bass frequency range by using the diaphragm resonance.**

Damping of the Resonance by Voltage Drive

In a cone loudspeaker a similar resonance occurs, which is due to the combined mass of the cone and the speech coil and the stiffness of its suspension.
This resonance is usually damped by driving the loudspeaker from a voltage source.
This short circuits its terminals so that any resonant movement of the cone results in a current in the speech coil. As this current must overcome the resistance of the speech coil, it generates heat in that resistance. That heat corresponds to resonant energy that is removed from the system, with damping as the result.

It is a common misunderstanding that the same happens in an ESL when the terminals are short circuited by voltage drive conditions.
It is true that with short circuited terminals a resonant movement of the diaphragm results in a current flowing between the short circuited stator electrodes.
There is, however, no resistance in the circuit to convert the stored resonant energy into heat. Therefore all the resonant energy remains in the system and no damping occurs.

When driving the loudspeaker from a voltage source with series resistance there is indeed a very small amount of damping because the movement of the diaphragm causes a small current in the resistance which thereby converts some of the resonant energy to heat. This produces very little damping, however, because the loudspeaker, used as a microphone, does not produces enough electrical output[44].

[44] That is: at realistic values of the polarizing voltage, the *transduction coefficient* of the loudspeaker is too small to provide sufficient damping.
This is similar to an electrodynamic loudspeaker having a very weak magnet.
In that case too, short circuiting the terminals of the loudspeaker will result in only a small amount of damping.

9.12 Diaphragm Resonance in the Mid Region and Damping

Mid Range Diaphragm Resonance

The stretched diaphragm has also multiple resonances in the mid frequency range.
These occur because waves can travel across the surface of the diaphragm and reflect against the diaphragm support (the spacers).
Due to these reflections standing waves occur and due to the distributed nature of the system the resulting resonance frequencies are manifold.
This will result in an undesired coloration of the sound reproduction.

Damping of the Mid Range Diaphragm Resonances

We can damp these mid range resonances by the same damping cloth we already need to control the bass resonance.
There is, however, a difference.
While damping of the bass resonance allows a distance between the damping cloth and the diaphragm of up to a decimeter, the mid range damping is only effective if the cloth is no more than a few millimeters away from the diaphragm.
This is in part due to the much smaller wave length of the mid range frequencies.
We can easily rialize such a short distance in an ESL using perforated plates as its stators, but it is much harder to achieve in an ESL that uses stretched wires as its stator electrodes.

Experiment

You can demonstrate the mid range resonance frequencies using again a small wooden frame with a stretched diaphragm glued to it.
To excite the mid range resonances, you gently tap the diaphragm with a rod of uncooked spaghetti. You can clearly hear a colored response, even though you will not hear one distinct tone.
When you place the frame on the seat of an office chair (I assume the seat is made of cloth), and repeat the experiment, you will find that the coloration is gone, provided that the distance between the cloth and the diaphragm is sufficiently small (< 5mm).

9.13 Spacers and Parasitic Baffling

Walker's equation assumes that there is no baffling.
We cannot avoid all baffling, however, due to the necessary spacers.
For the bass frequencies we can expect that parasitic baffling results in a slight increase of the response.
For the mid range and high frequencies, it could happen that the spherical wave, when it has expanded to the point that it reaches the edge of the panel, reflects against the spacers. This could again result in some high frequency undulations.
For this reason it is wise to keep the spacers, as well as supporting structures, as narrow as possible.
It would help to cover them with damping material.

9.14 Response of a Poorly designed ESL

Taking all of the effects described in this chapter into account, we can design an ESL with a flat frequency response. An example of such a design appears in chapter 12.

The current section describes the response that we can expect when we do not take the proper measures.

1. Using voltage drive will produce a response that increases with 6 dB/oct. This will result in a lack of bass and a thin sound[45].

2. A strong diaphragm resonance will partly compensate for the lack of bass. Because of the high *Q* of the resonsance, the bass response will be ugly.

3. If you use a diaphragm of 12 μm or even 25 μm polyester, the 6 dB increasing slope of the response will come to an end somewhere in the mid frequency range due to the diaphragm mass. Disregarding any other effects (such as path length differences) this will make the frequency response flat above the corner frequency.

4. If we have a one-way design, path length difference and phase cancellation will further reduce the high frequency response. the loudspeaker will also be very beamy.

5. If we connect such a large loudspeaker directly to the step-up transformer, a large resonance peak may occur in the high tones region due to resonance of the transformer's leakage inductance with the capacitive load of the ESL.

Figure 26 shows an example of the axial frequency characteristic that might result.
We see that there is a dramatic lack of bass due to the voltage drive.
The diaphragm resonance somewhat compensates the loss of bass.
The heavy diaphragm and the path length differences tame the surplus of high frequencies.
Around 10 kHz, the transformaer resonance compensates the loss of high frequencies.
Above that resonance the loudspeaker gives up.

Due to the strong response in the mid frequency range, such a loudspeaker can make an impression with copper instruments and female singing voices, but it is not of any practical use.

[45] Unless, of course, the loudspeaker is so large that we find ourselves in the region where the sound waves are still flat.
At a realistic loudspeaker size, this will not be the case.

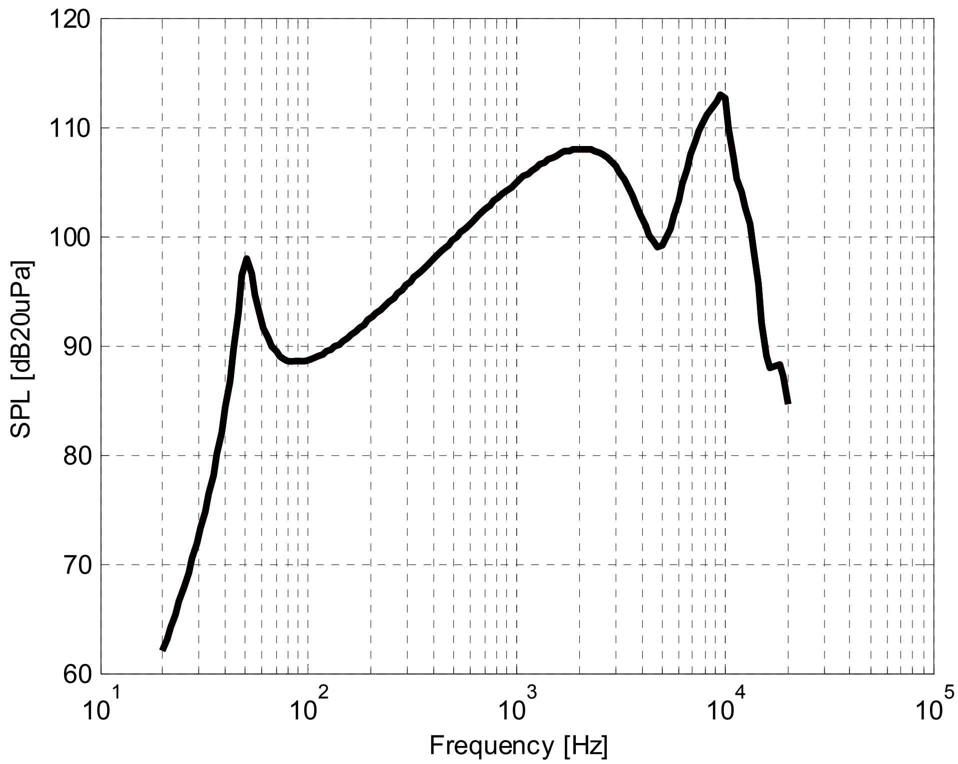

Figure 26 **Voltage driven ESL of 1 m x 1 m; distance 2 m; diaphragm thickness 12 um; diaphragm resonance at 50 Hz with Q=7; transformer resonance at 10 kHz with Q=7.**

10 Sensitivity

The previous chapter focused mainly on the frequency response of the electrostatic loudspeaker.
In the current chapter we want to take a look at how loud the loudspeaker will sound.

Sensitivity
It is customary to express the *sensitivity* of a loudspeaker as the SPL it produces at 1m distance[46] when driven by a voltage of 2.828 V_{rms}.
In case you wonder why people chose such a strange value for the voltage:
it corresponds to a power of 1 W in 8 Ω.

A typical sensitivity for an electrodynamic loudspeaker is 88 - 90 dB_{20uPa} at 2.828 V_{rms} and 1m distance. The Quad ESL63 has a sensitivity of 86 dB_{20uPa}.
It must be our goal to design an ESL with a similar sensitivity.
This will turn out to be quite a challenge.

The maximum Electrostatic Force on the Diaphragm
First of all, the electrostatic force on the diaphragm per unit of area is limited by the breakdown field-strength of air.
When we exceed that field-strength we get an electric discharge that we call an *arc*.

The breakdown field-strength depends a bit on the distance between diaphragm and stator, see Figure 27. With a typical distance of about 2 mm it is about 4 kV/mm.
See [2, Baxandall page 121].

The voltage budget we have before we get arcing must be distributed between the polarizing voltage and the audio voltage.
Baxandall [2, Baxandall, page 123] shows that we get the largest possible electrostatic force if we choose the polarizing voltage equal to half the breakdown voltage.
The peak signal voltage must then be twice as large, because it works on a twice as large air gap.
For a 2 mm distance between diaphragm and stator, we must therefore choose a polarizing voltage of 4000 V and a peak audio voltage of 8000 V_{peak}.

[46] This is customary for conventional cone loudspeakers, where the SPL drops by 6 dB each time the distance doubles.
If we want to express the sensitivity of a loudspeaker that does not have such a behavior (for example a line source), it gives a more fair comparison if we calculate or measure the SPL at 3 m distance and extrapolate back to 1 m assuming the 6 dB per factor of two in distance.

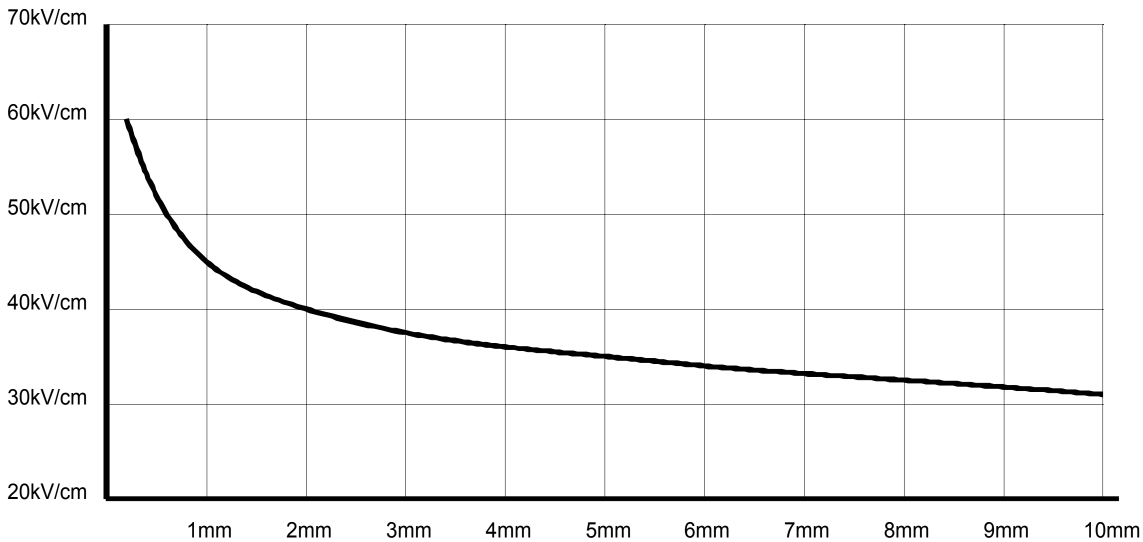

Figure 27 Breakdown voltage of air as a function of the gap size.

Baxandall [2, Baxandall, pages 121 - 123] also shows that the peak force on the diaphragm per unit of area (1 m^2) is in that case 70 N, independent of the width of the air gap[47].
He further explains that one newton is the force by which the earth gravitational field pulls on a large apple.

Baxandall thinks that due to practical limitations it will not be possible to attain this theoretical limit and that the best we can do is about 50 N/m^2.

Appendix IX gives the formula for the electrostatic force.

---

[47] That is: if you choose a larger air gap, you can increase the polarizing voltage and the signal voltage such that the field strength is again close to the breakdown voltage. You will achieve in that case the same electrostatic force on the diaphragm.
Here the slight reduction of the breakdown voltage with increasing gap size, as presented by Figure 27, has been neglected.

Maximum SPL for Electrostatic Headphones

Due to Newton's third law of motion (action equals minus reaction), the electrostatic force must be balanced by the force exerted on the diaphragm by the sound pressure. The force needed to accelerate the diaphragm mass has here been neglected.

Due to the symmetry of the situation the sound pressures at the back and the front of the loudspeaker are equal in magnitude, though of opposite sign.
Therefore, the maximum sound pressure we can get in the direct vicinity of the diaphragm is 25 N/m^2 or 25 Pa.
Assuming a sine shaped waveform, the rms value is 17.7 Pa, which corresponds to an SPL of 119 $dB_{20\mu Pa}$. This is the sound pressure of a riveter at close distance.

As the direct vicinity of the diaphragm in practice means a distance of a few centimeters at most, this result tells us what maximum sound pressure we can get from a pair of electrostatic headphones.

We see from the above that electrostatic headphones will produce a flat frequency response when driven by a constant (frequency-independent) voltage.

Appendix IX gives the expression for the sound pressure in the direct vicinity of the diaphragm.

Sensitivity of an Electrostatic Loudspeaker

To calculate the sensitivity of an electrostatic loudspeaker, using Walker's equation, we must first choose the size of the diaphragm. For a bass panel, the area is usually taken between 0.3 m^2 and 1 m^2.

Let's take as an example a bass panel area of 0.5 m^2.
Let's further assume a diaphragm-to-stator distance of 2 mm.
Not counting the area lost to the holes in the stator plates, the panel then has a capacitance of 1.1 nF (see appendix IV).
We assume that the holes reduce the capacitance to 800 pF.
In accordance with Baxandall's calculations we take the polarizing voltage 4000 V and the peak audio voltage 8000 V.
We must now choose the series resistance, which is a trade off between bass frequency range and SPL.
Let's aim for a bass frequency range of 50 Hz and let's assume that we can get one octave for free by using the diaphragm resonance, as discussed in section 9.11.
Then the cut-off frequency of the RC time must be 100 Hz. With a capacitance of 800 pF we need a resistance of 2 MΩ, where we have used appendix III to perform the calculation.
The peak drive current is then 8000 V/2 MΩ = 4 mA.

Using Walker's equation (appendix II) we find a peak sound pressure at 1m distance of 3.712 Pa, which corresponds to an rms value of 2.62 Pa_{rms}. This corresponds to an SPL of 102.4 $dB_{20\mu Pa}$. This is the sound pressure inside the NY subway.
At 2 m distance the SPL is 6 dB less.

Now to find the sensitivity for a drive voltage of 2.828 V_{rms} we must choose the step-up ratio of the transformer. The maximum that is practically feasible without compromising the frequency range of the transformer is about 150× (see chapter 15).
This transforms the 2.828 V_{rms} audio voltage to a secondary voltage of 424 V_{rms}.
With the series resistance of 2 MΩ this results in a drive current of 212 μA_{rms}.
This produces an SPL at 1 m distance of 200 mPa_{rms}, corresponding to 80 $dB_{20\mu Pa}$.
We see that the loudspeaker is 10 dB less sensitive than a typical cone loudspeaker and 6 dB less than the ESL63.

This numerical example shows that it is a challenge to get sufficient sensitivity out of an ESL.

11 Interaction with the Room

11.1 Introduction

This is a good time to discuss the interaction between the loudspeaker and the room. From the study of that interaction we can draw some conclusions about how to design the loudspeaker.
We will split this topic in two parts: bass frequencies and above bass frequencies. The discussion of bass frequencies will produce clues about the best loudspeaker placement.

11.2 Bass Frequencies

Omni Directional Cone Loudspeaker
We consider first the cone loudspeaker.
At frequencies that are low enough to make the loudspeaker box small compared to the wavelength the loudspeaker behaves as a point source.
A point source is omni-directional; it radiates equally in all directions.
As we will see shortly, this has consequences for the way it behaves in the room and for the best placement of the loudspeaker.

Directional Dipole Loudspeaker
As we have seen, a dipole loudspeaker has at bass frequencies a mild directivity, as expressed by the "figure of eight" shaped directivity diagram. See section 9.4.
Due to this directivity, a dipole radiates for the same axial SPL 3× less power than an omnidirectional loudspeaker. See for example [3, Verwaal] for the details of the calculation. That factor of 3 corresponds to 4.77 dB.

The directivity has the benefit that for the same subjective loudness of reproduction the neighbors hear 4.77 dB less noise.

Another advantage of the dipole loudspeaker is that the power ratio of the direct sound and the reflected sound changes in favor of the direct sound.
Again, the improvement is 4.77 dB.
A larger relative contribution from the direct sound contributes to a subjective sense of "directness" of the sound. This advantage gets stronger at higher frequencies, where the directivity of an ESL increases.

The directivity of an ESL also has an effect on the way it sounds in a room having resonances in the bass range, as we will discuss shortly.

Room Resonances
If we picture a typical listening room as a large shoe box, standing waves can exist between floor and ceiling and between each pair of opposite walls[48].
The corresponding resonance frequencies are determined by the distance between the parallel surfaces: if an integer number of half wavelengths fits between them we have a resonance. We also call these resonances *room modes*.

Room modes in the bass range result in boomy and inaccurate bass reproduction.
If you play back a dry beat on the bass drum you want it to sound, well, dry.
In a room with strong resonances, however, these will store part of the energy from the drum and they will only slowly give off that stored energy.
The result is that you still hear the drum hundreds of milliseconds after it was hit.

You can compare this to the rumbling echo's of the thunder in a thunderstorm.
If you are really close to the place where the lightning bolt strikes, you notice that the sound is sharp and dry and is very limited in time. It does not rumble.
At larger distance, the energy of the thunder gets smeared out in time, causing the rumble.

Bass Traps
If you really want to, there is a lot you can do about the resonances in your room.
What you have to do is to mount *bass traps* against all walls.
These typically consist of plywood panels of a carefully chosen thickness, which are mounted at a carefully chosen distance (a few centimeters) from the wall. The mass of the panel forms together with the air in the gap a resonator that is tuned to the frequency of the room mode[49].
Some damping material must be inserted in the gap.
The resonator "sucks" the energy out of the standing wave and converts it to heat in the damping material.
If you search the internet you will find more details.

[48] Probably standing wave patterns can also exist in other directions than the three perpendicular to the walls and the floor, but that does not really change the picture.

[49] Because the Q-factor of the resonator is low, it does not have to be accurately tuned to the room modes. The trap works in a frequency band of perhaps one octave. Usually you get the best results by placing traps that are tuned to frequencies one octave apart.

Spatial Distribution of Pressure and Velocity in the Standing Waves

Like all sound waves, the standing waves of the room modes consist of both pressure and particle velocity.
In the standing waves of the room modes these are spatially distributed in a characteristic way.

Close to the walls the particle velocity must of course be zero because the air has nowhere to go.
At the lowest possible resonance frequency, where one half wavelength fits between the walls, the velocity maximum occurs in the middle, right between the walls.
See Figure 28.

Based on symmetry considerations it is not difficult to see that the sound pressure will be zero in the middle and have maxima close to the walls. See Figure 28.

Figure 28 shows the distribution of particle velocity and pressure.

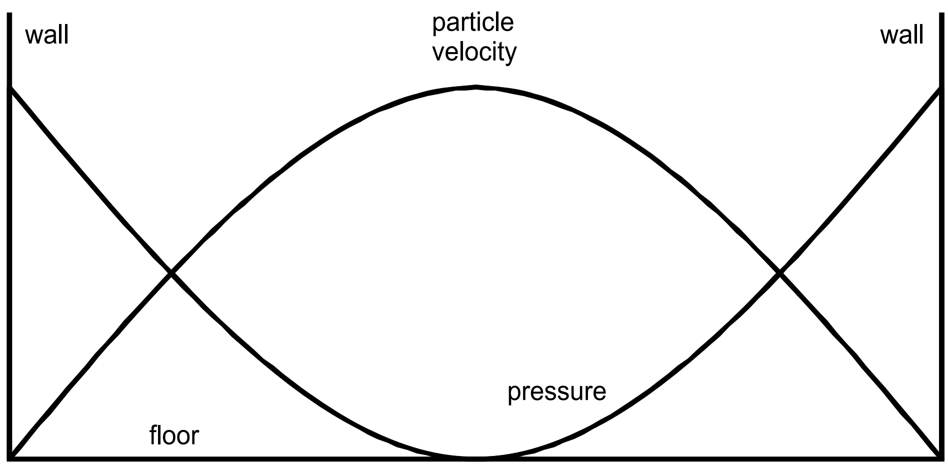

Figure 28 Spatial distribution of particle velocity and pressure in a standing wave.

As your ears are at bass frequencies only sensitive to pressure, you can check the validity of the above by playing some music that you know will wake up the room resonances. You then move your head through the room to find the locations where the "boom" is most prevalent. These are the pressure maxima.
You are likely to find these near the walls and even more near the corners of the room (the picture of Figure 28 is of course a one-dimensional simplification).
The corners where the floor and two walls meet are the worst locations.

Coupling of an Omnidirectional Loudspeaker to the Resonance Modes

Based on symmetry considerations, we expect that an omni-directional loudspeaker will not couple to the room mode if we place it in the pressure minimum in the middle of the room.

An omni-directional loudspeaker couples most strongly with the room modes when we place it in a pressure maximum.
The best recipe for inaccurate boomy bass is therefore to place a cone loudspeaker close to the wall, close to the floor and particularly in a corner.

Because the omni directional loudspeaker has no preferred direction, its orientation with respect to the room does not affect how much it excites the room mode. If we turn it by 90^o we will still get the same amount of boom.

For the reasons explained above, a good audio equipment shop will always advice you to place the loudspeakers a generous distance away from the wall and from the corners.

Coupling of a Dipole Loudspeaker to the Resonance Modes

Based on similar symmetry considerations, a dipole loudspeaker will not couple to the room modes if we place it at the location of a pressure maximum, which is also a velocity minimum.
Therefore, in order to minimize the boom bass, we would like to place an ESL close to the wall and perhaps even in a corner.
Placing it too close to a wall will, however, not work well because the ESL radiates sound from the back as much as from the front.

When we place and ESL close to a velocity maximum of a room mode it will couple with that mode, but only if the particle velocity of that room mode has a component along the axis of the loudspeaker.
We see therefore that contrary to the omni-directional cone loudspeaker the direction in which the ESL faces matters.

Based upon the above, a dipole loudspeaker is three times less likely to couple to a room mode as an omni-directional loudspeaker.

Bass Emphasis by Reflection in the Floor

If we place an ESL some distance above the floor, the sound bouncing off the floor will create a virtual image of the loudspeaker below the floor, see Figure 29.

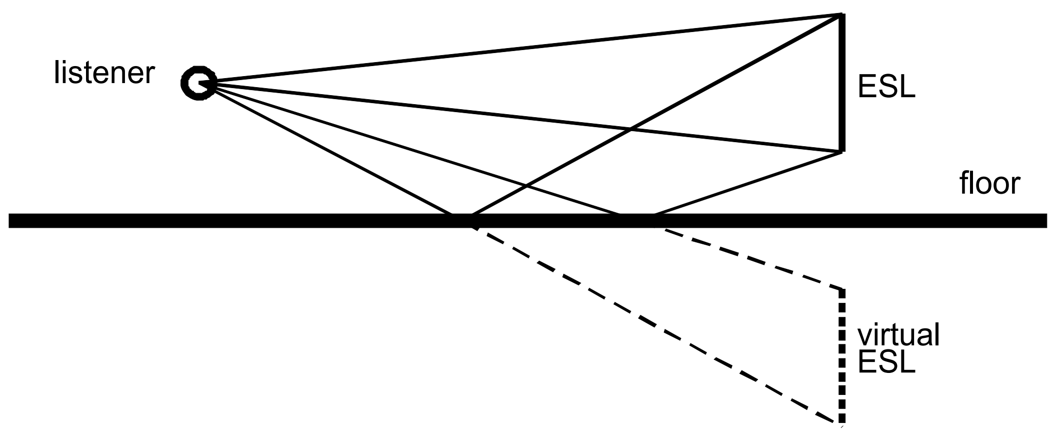

Figure 29 ESL with its reflection in the floor.

For high and midrange frequencies, the sound of the actual loudspeaker and its image will combine with all possible phase differences due to the different path lengths.
As a result, their sound pressures add up by power. This doubles the power of the sound pressure, which is equivalent to a 3 dB increase of the SPL[50].

For low frequencies, however, the difference of the path lengths is negligible compared to the wave length and therefore the two contributions have the same phase.
As a result, the sound pressure doubles, which is equivalent to a 6 dB increase of the SPL.

The net result is that the reflection in the ground emphasizes the bass response by 3 dB compared to the higher frequencies.

If we decide to put a high pole carpet on the floor, this will prevent the reflection of high- and midrange tones, but not of the bass tones. In that case, the bass emphasis is 6 dB compared to the high and midrange tones.

If we design an ESL we must correct for these effects by reducing the free space on-axis bass response by 3 - 6 dB.

[50] As these so called early reflections disturb the reproduction of mid-range and high tones, it would be better to damp the reflection by a high pole carpet.

Emphasis of Bass by the Left and the Right Stereo Channels
A similar effect takes place when we have two stereo channels.
In most cases the mid- and high frequency content of the two channels will be more or less different, but the bass is in most cases the same.
The mid- and high frequency content again adds by power, both due to the difference in content and due to the path length differences.
Because at bass frequencies the path length differences become unimportant, the bass content adds by pressure.
This again leads to a 3 dB benefit of the bass compared to the mid- and high frequencies.

If we design an ESL we can therefore reduce the free space on-axis bass response by another 3 dB.

Correction of the Bass Response
We must correct the free space axial bass response of the loudspeaker to compensate for the bass emphasis due to ground reflection and the combination of the two stereo channels. If we fail to do so, the loudspeaker will sound too bassy.

11.3 Above Bass Frequencies

Increased Directivity at Higher Frequencies
At higher frequencies the directivity of an ESL inevitably increases due to the size of the panel and the resulting path length differences.
These path length differences are of course the same as for the bass frequencies, but as they must be compared to the wave length they only become important at higher frequencies.

Directivity up to 1 kHz
You might think that the off axis response is unimportant, just as long as the axial response is flat. Of course you would want to keep the directivity low enough that you could move your head a bit without ending up outside the main beam.

This, however, is not the whole story.
The off axis radiated sound still reaches the listener, but indirectly through reflections and reverberation.
If the off axis sound lacks the high tones due to excessive directivity, the indirect sound will be too bassy and the entire sound image will be wrong.
According to Baxandall [2, Baxandall, pages 143,144] it is undesirable that the directivity of the loudspeaker increases by any significant amount up to 1 kHz.

According to Baxandall, a loss of high tones above 1 kHz in the reverberation field does not do so much damage because there is some attenuation of those frequencies in the reflection path as well[51].

Directivity above 1 kHz
We cannot avoid that the directivity above 1 kHz increases.
This is not necessarily a bad thing as it emphasizes the contribution of the direct sound.
That lends an appealing directness to the sound reproduction.
There is, however, an optimum amount of directivity.
According to Peter Baxandall, the tweeter part of a flat loudspeaker should therefore not be larger than 7 cm × 7 cm [2, Baxandall, page 146].

[51] Other sources claim that it is desirable to keep the off axis frequency characteristic more or less flat up to perhaps 5 - 10 kHz.

12 Three-way Design Approach

12.1 Avoiding the Adverse Effects of Path Length Differences

In chapter 9 we have found that path length differences lead for high frequencies to destructive interference, causing a loss of the high tones and undulations of the frequency response at high frequencies.

To avoid these undesired effects, the most obvious design strategy is to avoid the path length differences. There are, however, also other possibilities, which we will explore in chapter 13.

In the current chapter, we take the design approach to avoid the path length differences that would result from the use of large panels.
Or, to put it more precisely, we try to make the differences so small that we can disregard them.

With this approach we cannot avoid a multi-way design, because for a decent bass response we need a loudspeaker that is too large to keep the path length small enough for the highest frequencies.

With a two- or three-way design, the loudspeaker will consist of different panels with sizes that are optimized for the frequency range that they must handle.
Equivalently, some designs use one big panel with stators that are divided into different segments for the different frequency ranges. Splitting up the area of an ESL in this way is often called *segmentation*.

The question we must now first answer is whether we can limit ourselves to a two-way design or if we need a three-way design after all.

12.2 Three-way Design

We cannot avoid a 3-Way Solution
We will show that we cannot escape the use of a three-way design.

We have seen in section 11.3 that the directivity of and ESL is optimal with a tweeter size of about 7 cm × 7 cm.
In chapter 9 we saw that a bass panel needs to have an area of 0.3 to 1 m^2.
Let's assume for the moment a bass panel of 70 cm × 70 cm.

The area of the tweeter is then 100× smaller than the area of the woofer.
Assuming that the woofer has a bass frequency range of 100 Hz (not counting the extra octave we can get by using the diaphragm resonance), we must expect[52] that the lowest frequency that the tweeter can reproduce is 10 kHz.
Due to path length differences, the large bass panel cannot work satisfactorily up to 10 kHz.
It appears therefore that we cannot escape the use of a mid range panel.
We sometimes refer to the bass panel, mid-range panel and high frequency panel by the names *woofer*, *squacker* and *tweeter*.

12.3 Relative Placement of the Panels

Based on the above we will have a 3-way design, with a woofer panel of about 70 cm × 70 cm, a tweeter panel of about 7 cm × 7 cm and a squacker panel with a size somewhere in between.

The question arises how to place these panels relative to each other.
The most obvious choice might be to place the three panels side by side as in Figure 30.
Experience and (computer) calculations have shown that this does not work well.
At the cross over frequencies the path length differences between the two panels that must both equally contribute is too large. The effect that this has on the frequency response can be simulated with the Matlab simulation program *ESL.m*.
An example of the effect of this kind of panel placement appears in section 12.6.

[52] This assumes that all panels use the same diaphragm-to-stator distance and the same polarizing voltage.

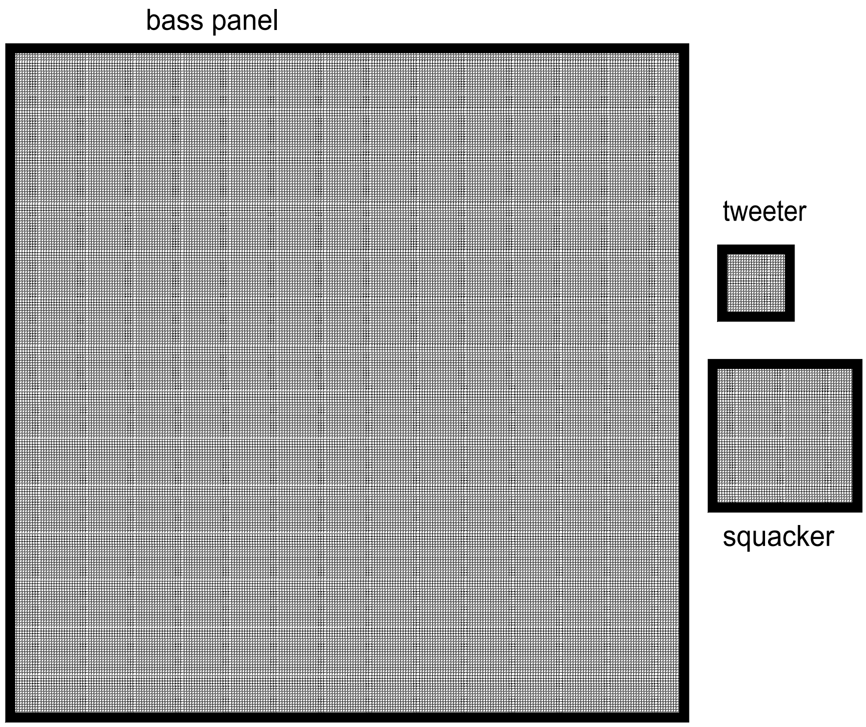

Figure 30 **Bass panel, squacker and tweeter placed side by side.**

As a different solution, we could consider placing the squacker in front of the bass panel and the tweeter in front of both (Figure 31). That way all panels have the same axis. This does not work either, because even a path length difference of 5 cm at a cross over frequency of 300 - 400 Hz between the bass panel and the squacker is already noticable.
At the cross-over frequency between the squacker and the tweeter things are of course even more dramatic.
The Matlab simulation script *ESL.m* allows you to simulate the effect of this panel placement.
An example of the effect it has on the frequency response appears in section 12.6.

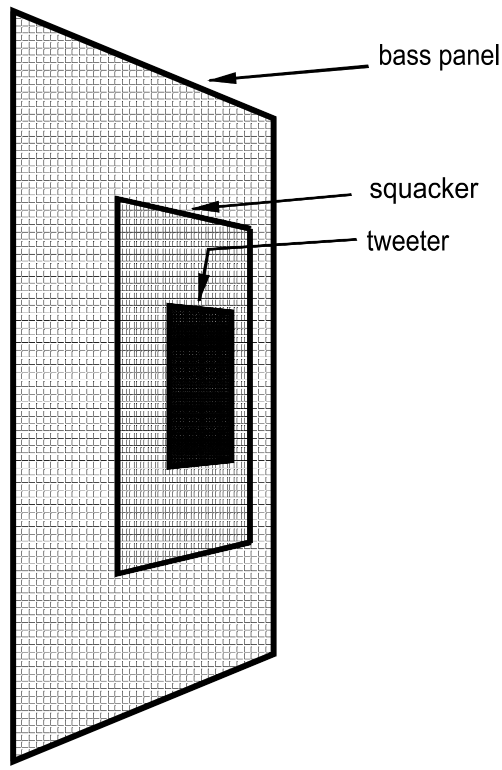

Figure 31   Coaxial placement of the panes in front of each other.

The best solution is to leave a hole in the bass panel in which we can mount the squacker.
Similarly we must leave a hole in the squacker to make room for the tweeter.
See Figure 32. That way, all panels can be at the same time co-planar and co-axial.
Such a segmentation lends itself particularly well to realization of the stators using perforated PCB material.

For the purpose of simulating these kind of loudspeakers, the Matlab simulation program *ESL.m* allows us to define rectangular panels with rectangular holes in them.

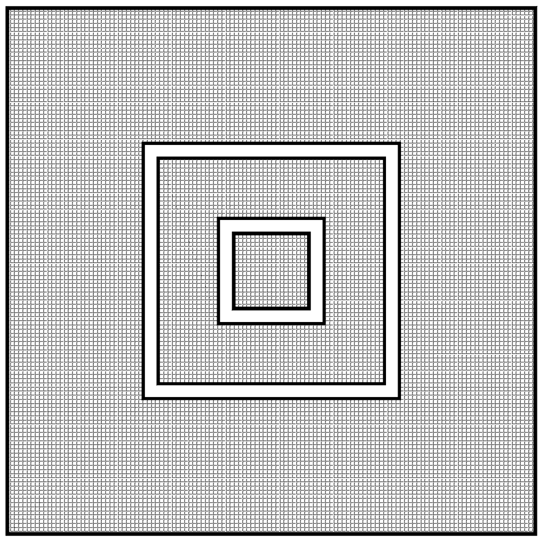

Figure 32 Panels co-planar and coaxial.

If we want to use perforated sheet metal for the stators, we may decide for ease of construction to split the bass panel in two equal parts and mount the squacker between these. See Figure 33.
This arrangement is somewhat similar to the configuration proposed by d'Appolito for cone loudspeakers [10, d'Appolito].

The Matlab simulation program *ESL.m* allows you to define a bass panel with such a large rectangular hole in it that it is effectively cut in two.

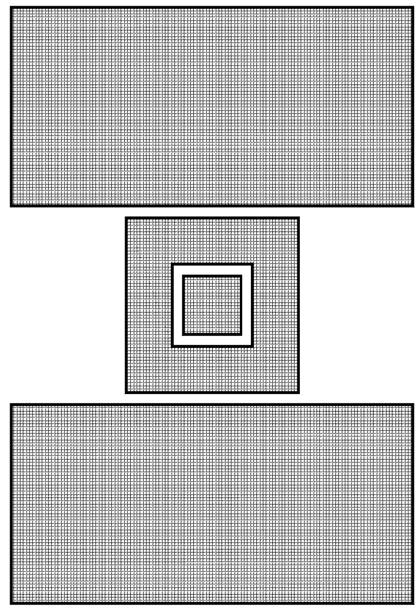

Figure 33 **Co-planar panel placement with two bass panels in a symmetrical arrangement.**

12.4 Choice of the Cross-over Frequencies

Natural sounds, including human voices, have most of their power in the frequency range 300 - 4000 Hz.
It should not come as a surprise then, that the human ear is most sensitive to errors in the sound reproduction in that frequency band.
Our sound perception is much more forgiving for errors that occur outside that frequency range.

Indeed, it has been long known that the frequency band between 300 Hz and 4000 Hz is the most important one for conveying intelligible speech. This is why the telephone bandwidth is limited to a similar frequency range.

In order to get a loudspeaker that sounds as natural as possible, we should therefore strive to avoid cross-over frequencies in the 300 - 4000 Hz range.
This means that this entire band must be handled by the squacker.
This is yet another reason why a three-way design is to be preferred over a two-way design.

As we have already seen, a 7 cm × 7 cm tweeter has difficulty handling frequencies below 10 kHz. A squacker capable of handling the range from 300 Hz to 10 kHz, however, is not a realistic thing to ask either.
We will therefore have to compromise a bit on the 300 Hz cross over frequency between woofer and squacker.

12.5 Cross-over Filter

Current Drive

To achieve the desired current drive conditions for the panels, we can drive them from a voltage source (e.g. the secondary voltage of a step-up transformer) with resistors R in series.

Neglecting for the moment any parasitic capacitance working in parallel with the panels, all resistors must be of the same value so that all panels receive the same drive current[53].

To preserve the symmetry of the circuit, each resistor must be split in two equal parts of value $R/2$. Figure 34 shows the circuit diagram.

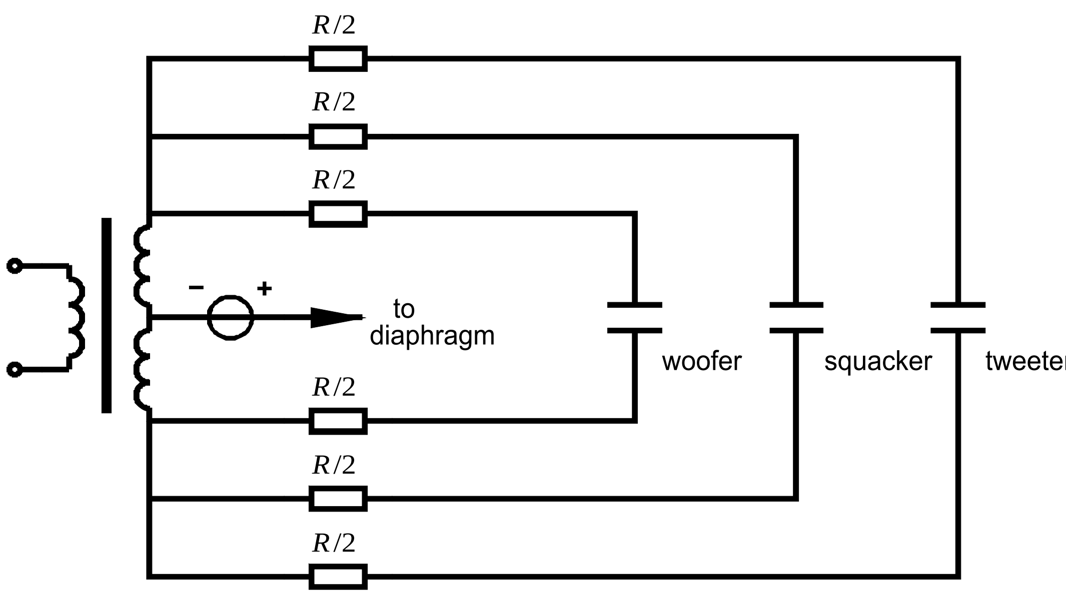

Figure 34 Three way loudspeaker with series resistors.

[53] This assumes that all panels have the same distance from diaphragm to stator and the same polarizing voltage.

Filter Action

The question now arises how to achieve the necessary cross-over filter action.

Fortunately, each panel has already a natural frequency limitation at the low frequency side. This low side cut-off frequency is determined by the RC product of the series resistor and the panel capacitance, see appendix III.

What remains is to limit the frequency range of the bass panel and the squacker at the high frequency side.

We can achieve this by again splitting the resistors of value $R/2$ into two resistors of value $R/4$ and connecting a shunt capacitor across between the tap points

Figure 35 shows the resulting circuit, where the two filter capacitors are $C_{filt,W}$ and $C_{filt,Sq}$.

The values of the capacitors must be chosen such that the bass panel and the squacker panel are limited to the desired cross over frequencies.

A small inconvenience of this solution is that high voltage capacitors of good linearity are needed.

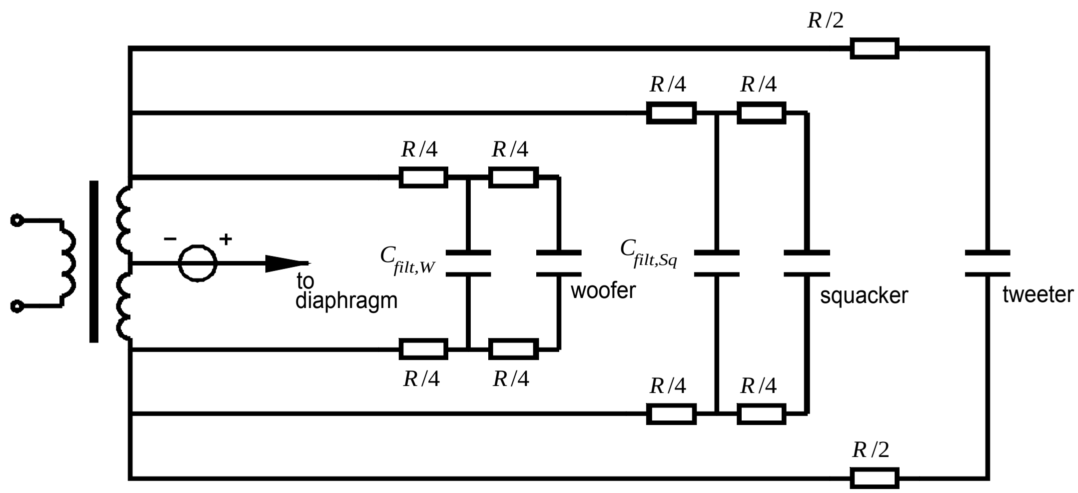

Figure 35 Cross-over filter.

The simulation program *ESL.m* allows you to simulate the frequency response of the loudspeaker including this cross-over filter network.

Parasitic Capacitance
In the above presented analysis, we disregarded all parasitic capacitances.

Consider for the moment a loudspeaker using stators made of perforated sheet metal. The front and the back stator are glued together with the diaphragm in between using isolating spacers. See chapter 16 for construction details.
The stator area where the spacers are located does of course not contribute to the active panel area. It does, however, have a capacitance, which is even increased by the relative dielectric constant of the spacer material.
This parasitic capacitance works in parallel with the active panel[54].
Figure 36 shows the circuit diagram including the parasitic capacitances.

The spacer area is the largest relative to the active area in the smallest panels. Therefore, the tweeter suffers most from parasitic capacitance, followed by the squacker.

The parasitic capacitance eats away a part of the drive current of the panel.
To compensate for this, we must reduce the series resistance such that the RC product is again the same as it was intended without the parasitic capacitance[55].
We must increase the shunt filter capacitor of the cross-over filter by the same factor, such that again the RC products are as they were before.

The Matlab simulation program *ESL.m* allows you to specify for each panel the parasitic capacitance so that you can verify that the corrective measures in the dimensioning of the cross over filter work out as intended.

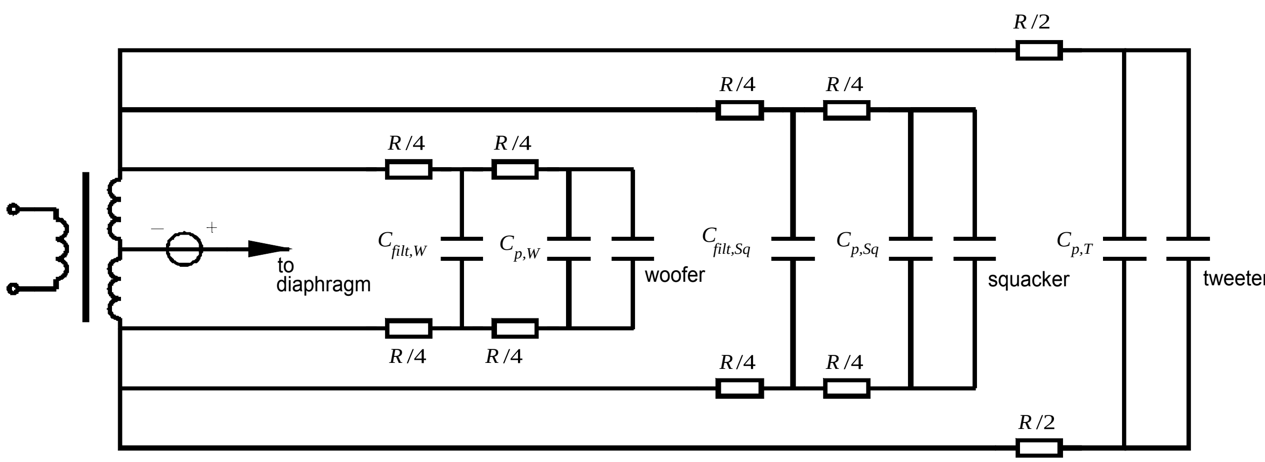

Figure 36 Circuit diagram including parasitic capacitances.

---

[54] The parasitic capacitance can be much smaller if the stators are made from perforated PCB material. We must then remove the copper at the location of the spacers.

[55] We can easily see that this restores the drive current to the panel to the originally intended value, because it restores the drive voltage to its original value.

Ladder type of Cross-over Filter

A popular type of cross-over filter is the ladder circuit shown in Figure 37.
Figure 37 shows the circuit for a two-way system, but it is typically used for loudspeakers with much more segmentation – up to perhaps 6 segments.
This filter has the advantage that the capacitance of the tweeter forms the filter capacitor for the woofer. This avoids the use of special high linearity high voltage capacitors.

A disadvantage is that it is not possible to obtain equal drive currents for the panels: it can be easily seen that the bass panel gets less current than the tweeter[56].
This property makes this type of cross-over filter more suited for line source loudspeakers, which turn out to need just such a distribution of drive currents, see section 13.6.

Only if the tweeter has a large parallel capacitance (either parasitic or intended) can we achieve equal drive currents in the tweeter and the woofer.

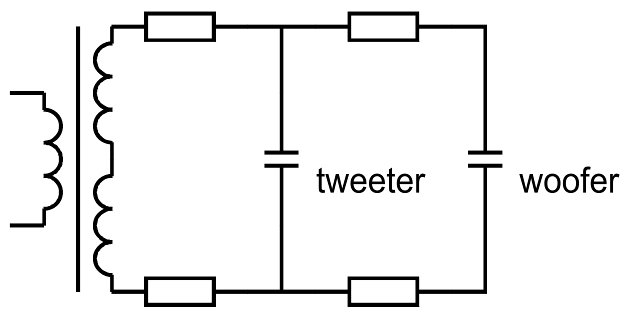

Figure 37 Ladder type of cross-over filter.

[56] To see this, assume fist that in the bass frequency region the woofer becomes very low impedant, while the tweeter is high impedant. Next, assume that in the high tone frequency region the tweeter becomes very low impedant as well.

12.6 Example of a Three-way Design

Introduction
By way of example we will now dimension a three way design as described in the previous section.

Distance between Diaphragm and Stator
In order to get as much sensitivity as possible we choose the distance between diaphragm and stator rather small: 1.5 mm.

Polarizing Voltage and Step-up Ratio
With a diaphragm-to-stator distance of 1.5 mm and an estimated breakdown field strength of 4 kV/mm, the estimated breakdown voltage is 6000 V.
In accordance with the considerations given in chapter 10, getting the maximum possible sound pressure from the loudspeaker then requires a polarizing voltage of 3000 V and a peak audio voltage of 6000 V.

To get the best possible sensitivity, we choose a step-up ratio of the audio transformer of 150×. This is about the highest step-up ratio we can realize by using two transformers of 1:75. See section 15.12.
We expect that a high quality audio amplifier can produce a peak output voltage of 40 V. With these parameters we can obtain a maximum secondary voltage of 6000 V_{peak}.

Size of the Panels
We choose the following panel sizes:

- Tweeter: 7 cm × 7 cm
- Squacker: 18 cm × 18 cm
  with a square hole of 8 cm × 8 cm in the middle for the tweeter to fit in.
- Woofer: two panels of 20 cm × 40 cm each, above and below the squacker.

This brings the total woofer panel area to 0.16 m^2.
This is a bit small, but for this example I wanted to design a compact loudspeaker.
The price to be paid is of course a rather limited bass range.

The panels will be placed as in Figure 33.

The hole of 8 cm × 8 cm in the squacker panel is rather small for the tweeter measuring 7 cm × 7 cm. This is probably only possible when using stators of perforated PCB material, where the segmentation is realized by etching a suitable pattern in the copper.

Estimated capacitance of the panels

We can easily calculate the capacitance of the panels from their dimensions.
We estimate a capacitance reduction due to the holes in the plates by a factor 0.8×.
Assuming that the stators are made of PCB material, the construction can be such that there is almost no parasitic capacitance. We will therefore at first neglect the parasitic capacitance.

Using the formula given by appendix IV, the estimated capacitance of the panels, including the loss due to the holes is

Panel Capacitance		
Panel	**Size**	**Capacitance**
Bass	two panels of 20 cm × 40 cm	378 pF
Squacker	18 cm × 18 cm with a hole of 8 cm × 8 cm	61.4 pF
Tweeter	7 cm × 7 cm	11.6 pF

Cross over filter

We use a cross over filter like the one shown in Figure 35.
It takes some tweaking using the Matlab simulation program *ESL.m* to get the overall response flat. It is true that when using equally valued series resistors the response of each of the three panels is the same in the asymptote.
The corner frequencies, however, are too close by to be neglected and therefore we arrive at slightly different values for the resistors.
The resistor- and capacitor values found in this way are:

Panel	Panel Capacitance	Series Resistance R1	Series Resistance R2	Filter Capacitor
Woofer	378 pF	1.35 MΩ	1.35 MΩ	700 pF
Squacker	61.4 pF	1.6 MΩ	1.6 MΩ	77 pF
Tweeter	11.6 pF	1.9 MΩ	1.9 MΩ	-

Figure 38 shows the axial frequency response of the loudspeaker at a distance of 10 m.
The simulation was done at a primary input voltage of 2.828 V_{rms}.
The SPL at 10m distance is 54 dB_{20u}, which extrapolates to 74 dB_{20u} at 1m distance.
Compared to the 86 dB sensitivity of the Quad ESL989, we are 12 dB short.

We see that the cross-over frequencies are 560 Hz and 3 kHz.
The lowest cross-over frequency is a bit higher than desired. We can get it lower by increasing the series resistance or increasing the size of the squacker panel.
The first option reduces the sensitivity, which is already on the low side.
The second opton will make the loudspeaker more beamy in the mid frequency range, which is not desirable either.

The bass frequency range is 85 Hz, which is not too bad.
We can improve it by making use of the diaphragm resonance.
A resonance at 50 Hz with $Q = 1.8$ brings the -3 dB frequency down to 47 Hz (Figure 39).

We can get a higher sensitivity by reducing the resistor values, for example by a factor 2, but the cross-over frequencies will increase accordingly and so will the bass cut-off frequency.
To get the original cut-off frequencies back, we can increase the size of all panels by a factor 2 (tweeter 10 cm × 10 cm;squacker 25 cm × 25 cm and bass panels 4 × 20 cm × 40 cm).
We must then also double the value of the filter capacitors.
The price we must pay is a dramatic increase of the directivity.

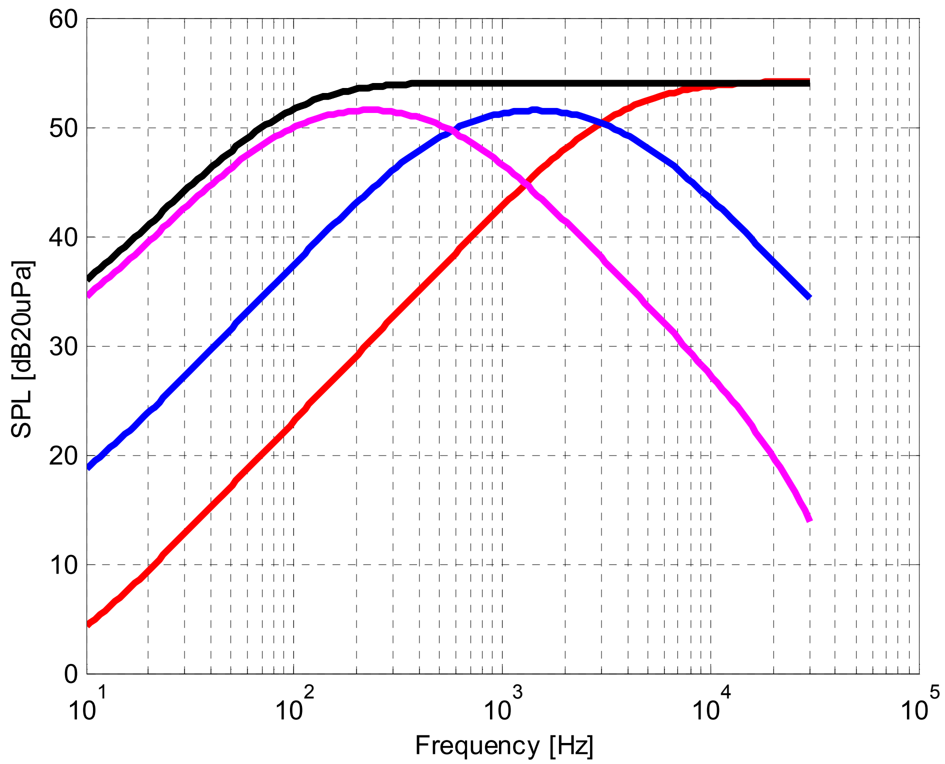

Figure 38 **Frequency response of three-way loudspeaker.**

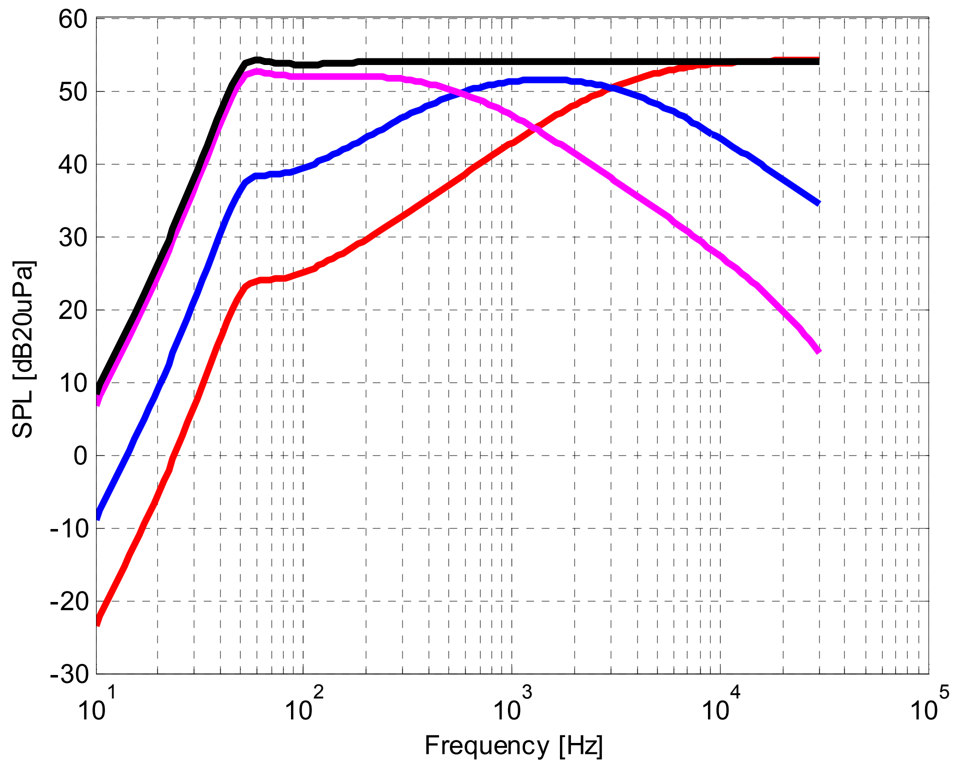

Figure 39 **Frequency response when making use of the diaphragm resonance to increase the bass range.**

As we said before in chapter 11, it is important not to let the directivity increase significantly with frequency up to 1 kHz, because this would disturb the tone balance of the reverberation field.
Simulation of the off-axis response at an angle of 45^o shows that at 1 kHz the response is only 2 dB down. See Figure 40.

The parasitic capacitance that works in parallel with the panels has here not been taken into account. While building the stators from perforated PCB material allows us to avoid most of the parasitic capacitance that would otherwise exist at the spacers, some extra capacitance due to wiring cannot be avoided.
The simulation program *ESL.m* allows us to specify the (estimated or measured) capacitance so that we can compensate by a different dimensioning of the cross-over filter.

Finally, we must be aware that we have neglected the influence of listening room, and especially that of the floor. Because of this we can expect that the loudspeaker produces too much bass. Tuning of the component values of the cross-over filter, based on listening in the actual room, will be needed to correct this.

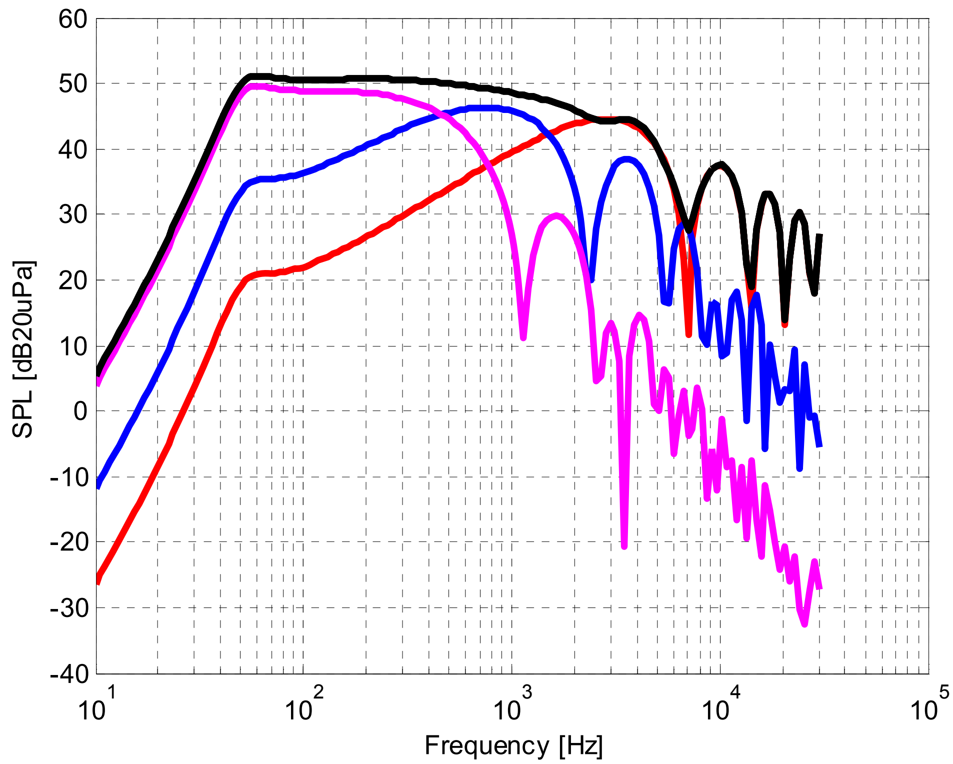

Figure 40 Off-axis frequency response at an angle of 45°.

Placement of the Panels relative to each other

This is a good time to demonstrate the influence of the relative placement of the panels. The simulations we did so far were based on a placement similar to Figure 33.

For a realistic comparison we will first simulate the response at a more realistic listening distance of 3 m. Figure 41 shows the simulation result, where we zoomed in a bit to show the 1 dB ripple in the relevant frequency band.

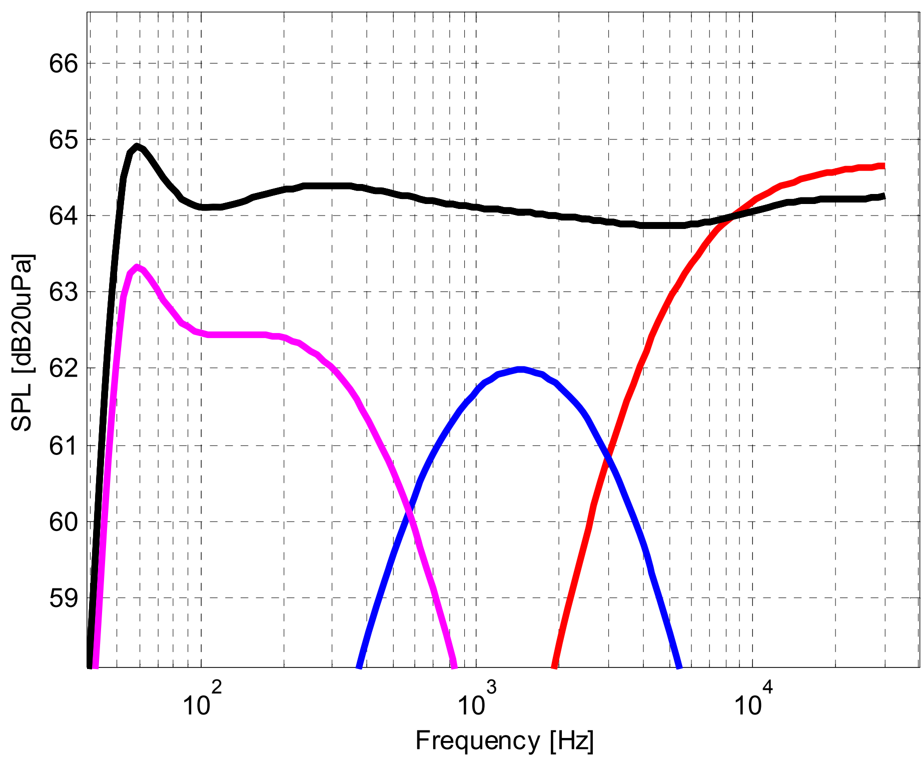

Figure 41 Frequency response at a distance of 3 m.

If we place the woofer, squacker and tweeter co-planar side by side horizontally, the center of the tweeter is 49.5 cm offset compared to the center of the woofer.
Figure 42 shows the frequency response for this situation, where the three panels have now been taken square with sizes 40 cm × 40 cm, 16.2 cm × 16.12 cm and 7cm × 7 cm. We see that above 5 kHz the response is severely degraded.

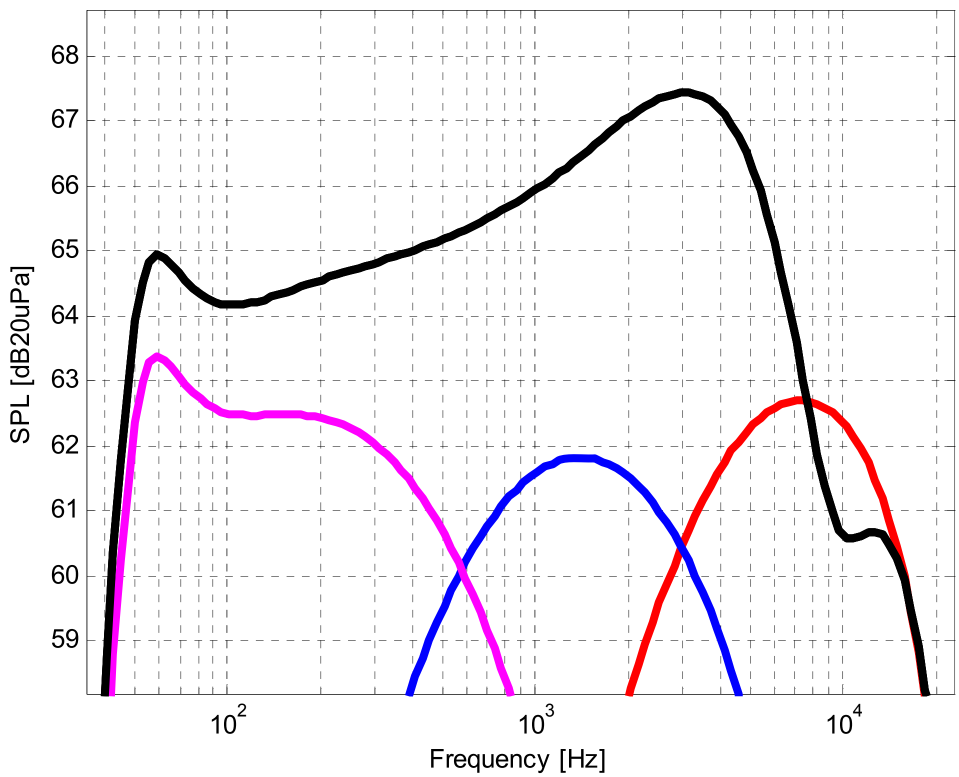

Figure 42 Frequency response with the panels placed side by side horizontally.

Now let us try to place the three panels co-axial in front of each other with the woofer in the back and tweeter in the front.
The distance between the panels will be 5 cm.
Figure 43 shows the rather terrible result.

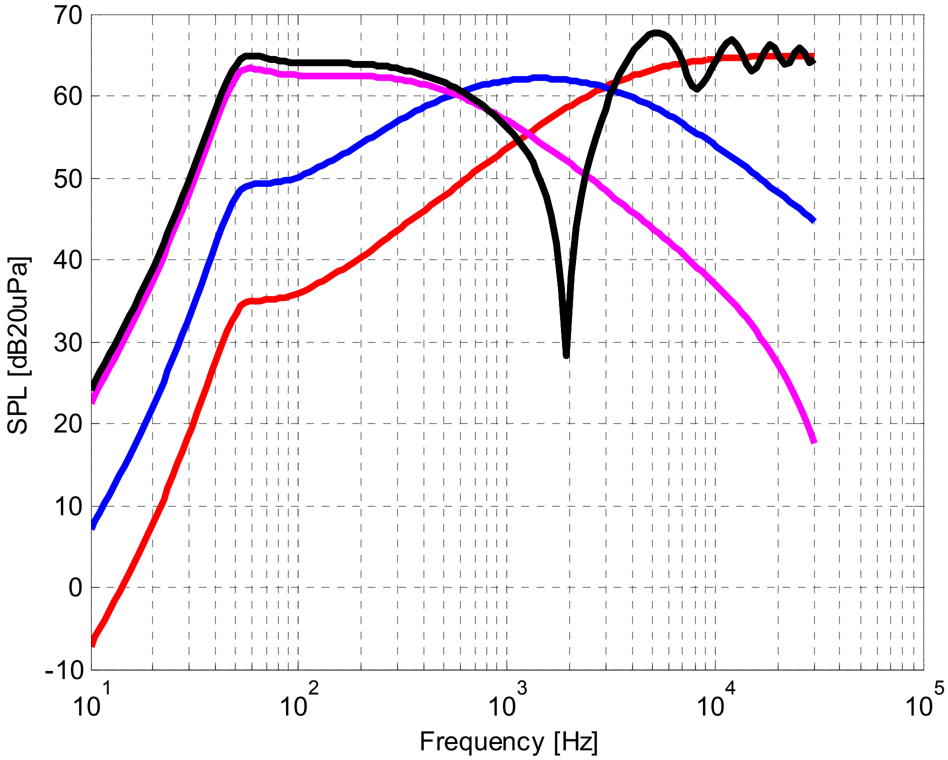

Figure 43 **Coaxial placement with the tweeter 5cm in front of the squacker and the woofer 5cm behind the squacker.**

Some improvement is possible by placing the tweeter and the squacker coaxial in the same plane with the woofer 5cm behind them.
The tweeter and the squacker can form one panel with segmentation using stators from PCB material.
Figure 44 shows the result. It looks a lot more acceptable, although we still see a ripple of 5 $dB_{p\text{-}p}$. This is for a loudspeaker a small error.

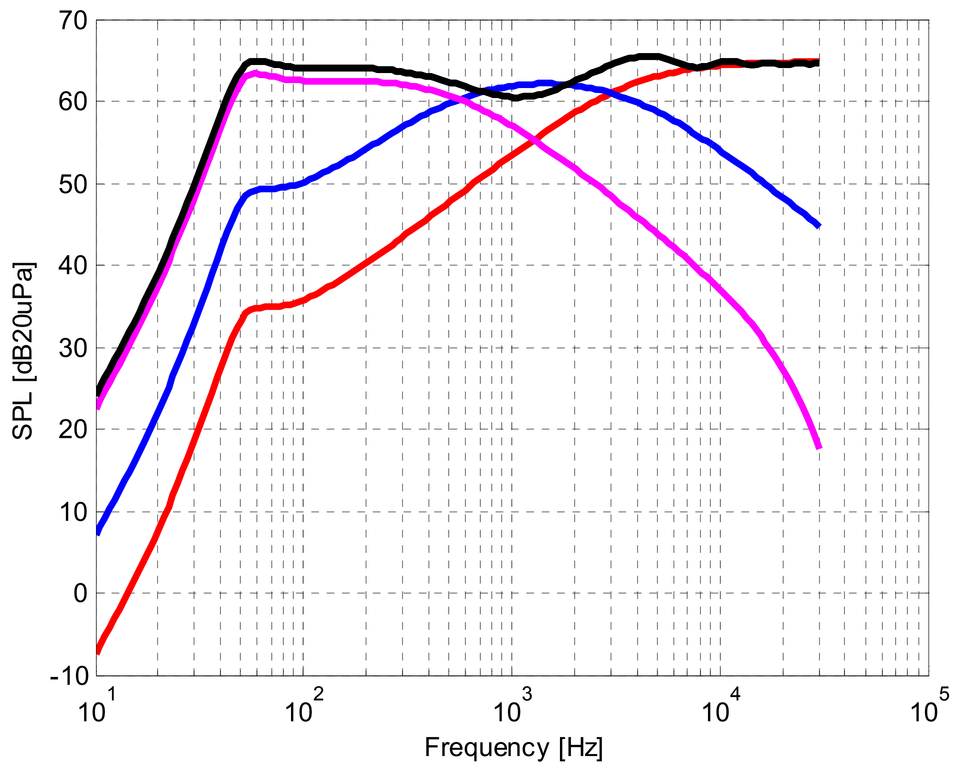

Figure 44 **Placement of squacker and tweeter in the same plane, 5cm in front of the woofer.**

13 One-way Design Approach

13.1 Introduction

In the previous chapter we sought to combat the negative effects of the path length differences by making them small without abandoning the idea that the loudspeaker should be flat.
Following that path we found that the panel must be so small that it can only serve as a tweeter. A separate panel will be needed to reproduce the bass frequencies.
As the bass panel is too large to work up to the tweeter cut-off frequency, we need a mid range panel (squacker).
This forced us to adopt a three-way design.

In the current chapter we explore alternative solutions.
Some of these aim to avoid the path length differences by using a curved (concave) design.

Others deliberately increase the path length differences by using a design that is curved the other way (convex) such as to approximate a pulsating sphere (point source). This increases the path length differences but they end up generating sphere shaped waves.

Yet other solutions aim to generate flat waves by using a very large flat loudspeaker or cylinder shaped waves by using a line source. Here, again, the path length differences are large but the contributions from all surface elements produces the desired flat or cylinder shaped waves.

Most of these designs circumvent the need for a cross over filter and are therefore one-way designs.
In some cases, however, segmentation and cross-over filtering can still be helpful.

The following sections discuss briefly some of these possibilities.

13.2 Concave Loudspeaker

One obvious solution to avoid the path delay differences is to use a concave loudspeaker (forming a cut-out section of a sphere) with the listener in the center of the sphere (Figure 45).
This solution ensures that the distance from the listener to all points of the diaphragm is the same. Therefore no path length differences limit the size of the loudspeaker and so we can stick to a one-way design.
As a bonus we can expect at the listening position a larger SPL due to the focusing effect.

A disadvantage of this approach is that we must expect a non-flat off-axis response. Therefore the reverberation field will contain not enough of the high tones, causing the loudspeaker to sound too bassy. This is probably not too serious, however, because of the large contribution of the direct path.

A big disadvantage is that the directivity will be so high that the loudspeaker will sound well only in one spot of the room.

A practical difficulty is of course that we cannot keep a concave diaphragm in place by stretching it. We can possibly solve this by approximating the curved shape by a large number of small square flat panels. It is not difficult to work out how small these panels must be in order to keep the remaining path length differences sufficiently small[57].

An intermediate solution is to make the loudspeaker large and concave in the vertical direction while keeping its width down to perhaps 7 – 10 cm.

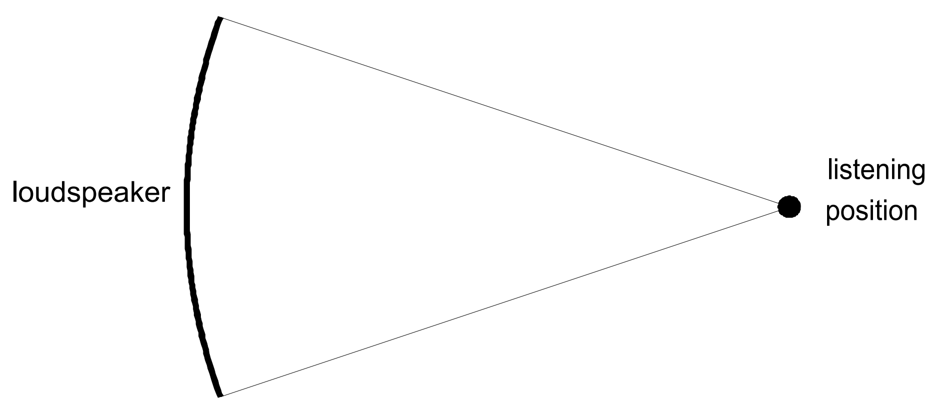

Figure 45 Concave loudspeaker.

[57] It would be best if the remaining path length differences were smaller than 4 mm (which is approximately one quarter of the wavelength at 20 kHz). This would probably require a too large number of panels, so that we must accept a larger residual path length error.

13.3 Convex Loudspeaker

A quite opposite approach to avoid the adverse effects of path delay differences is to try and approximate a point source (pulsating sphere)[58].
Even though there are now large path length differences, the contributions from all parts of the diaphragm combine in such a way that sphere shaped waves result.
The loudspeaker is omni-directional[59] and the frequency response is flat.
We can again use a one-way design.

Because it is not practical to realize a complete sphere, we must content ourselves with a small part of it.
This way we end up with a convex loudspeaker (Figure 46).
The sound seems to be coming from a virtual point behind the loudspeaker at the center of the partial sphere. The loudspeaker will no longer be omni directional.

Realizing only a part of the sphere causes truncation errors, which will re-introduce ripple of the high frequency response.
This can be avoided by gradually tapering off the drive voltage (or the polarizing voltage) towards the edges.

The construction challenges of a convex loudspeaker are similar to those of a concave loudspeaker.

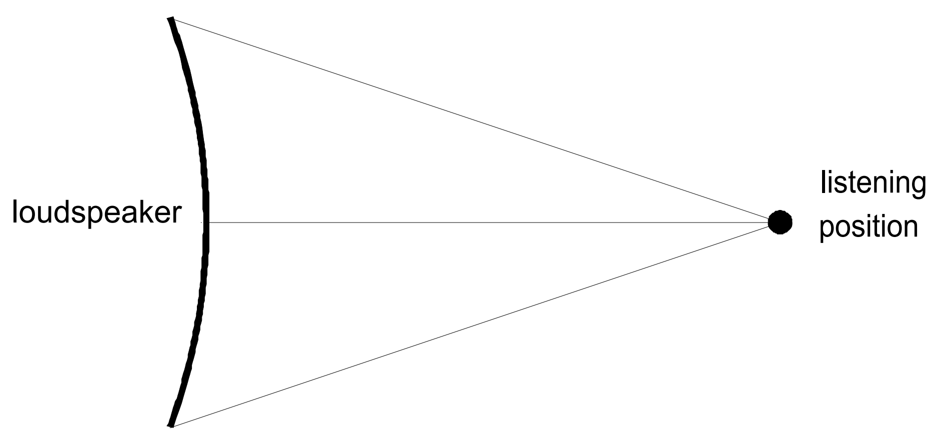

Figure 46 Convex loudspeaker.

[58] A problem that must be solved is what to do with the sound radiated at the back side towards the center of the sphere. This is solved when we build a loudspeaker that only forms a small part of the sphere.

[59] This means that we loose one of the advantages of the ESL: its directivity.

13.4 Simulated Point Source using a Delay Line

A solution proposed in 1980 by Peter Walker [1, Walker], and used ever since by Quad[60], is the simulated point source.
It uses a flat loudspeaker of roughly square or rectangular dimensions, which is divided in concentric circular segments (Figure 47).
The segments are driven by a delay line, with the outer segments delayed compared to the inner segments.
This simulates a point source having its center some distance behind the loudspeaker.
Again, the response of the outer segments must be tapered off, at least at the high frequencies, to avoid truncation errors.

Quad conveniently realizes the segmentation with concentric circular segments by using stators made from perforated PCB material.
It turns out to be a good choice to reduce the width of each next ring, in such a way that the areas of the segments are equal.

This design allows some control over the directivity of the loudspeaker, similar to the way that we can control the directivity of a conventional three-way ESL design by choosing the sizes of the tweeter and the squacker panels.

Of course, this design does not avoid the segmentation and the delay line has taken the place of the cross-over filter.
Building and correctly dimensioning the delay line is difficult and this makes this design approach less attractive for home-built loudspeakers.

Quad ESL63

Figure 48 shows the segmentation of the ESL63.
The loudspeaker is built from four panels of about 40 cm × 20 cm.
The mid two panels are segmented in the above described way.
The segments have equal area's and therefore equal capacitances.
The delay per segment is 24 μs, corresponding to a path length difference of 7 mm in air.
The simulated point source has its center approx. 30 cm behind the loudspeaker.
The polarizing voltage is 5.25 kV

[60] Starting with the ESL63 and subsequently the ESL989 and its successors.

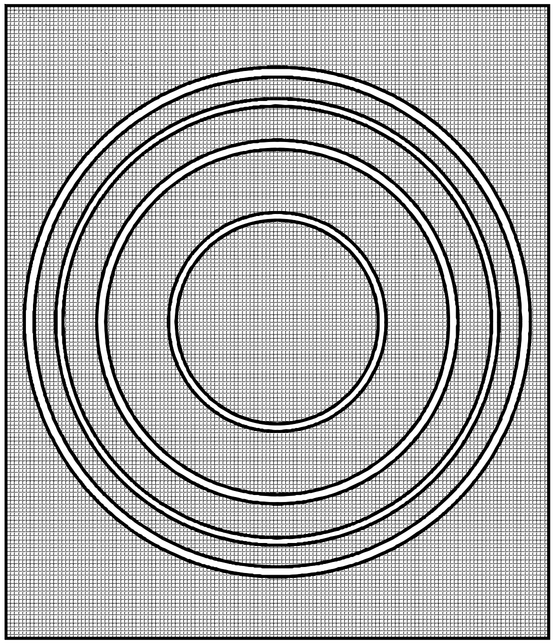

Figure 47 Segmentation in concentric circles to simulate a virtual point source with its center behind the loudspeaker (drawing not to scale).

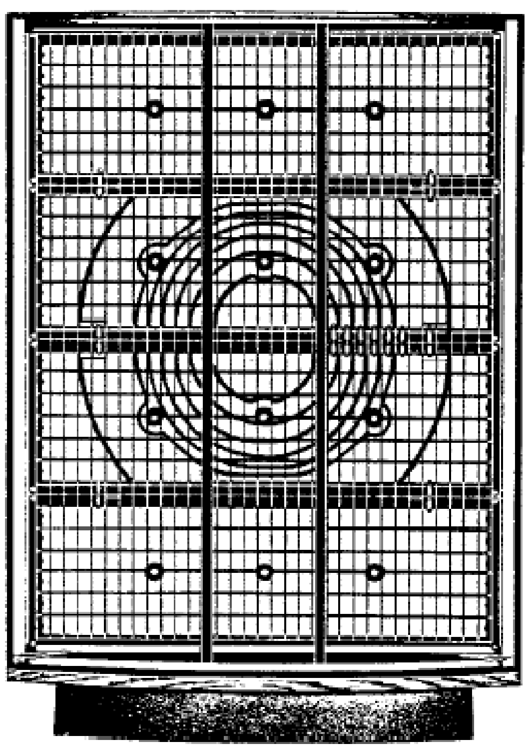

Figure 48 Segmentation of the ESL63.

13.5 Very Large Loudspeaker - Flat Waves

Yet another approach to the challenge of realizing a flat frequency response in the presence of different path delays is the use of a very large flat loudspeaker.

As we have seen in chapter 10, the frequency response of a voltage driven flat loudspeaker in the immediate vicinity of the diaphragm is flat without the need for segmentation or a cross-over filter.
At close proximity, the sound waves are still flat[61].
This approach works well for electrostatic headphones.

We can try to make the loudspeaker so large that the region of flat waves extends to the listening position, which will typically be at a distance of a few meters.

Such a loudspeaker has the property that, in the vicinity region where the waves are flat, the SPL does not decrease with the distance[62].

We can expect that for this to work, the size of the loudspeaker must be considerably larger than the listening distance.

Simulations with the Matlab simulation program *ESL.m* confirm this; a size of 8 m $\times$ 8 m is not even large enough[63] (see Figure 49).
An added complication is that voltage drive of such a large panel is extremely difficult to realize. This is therefore not a very practical solution.

We see from the simulation that the achievable SPL is very high.

[61] Even though the path length differences are very large, the contributions from all surface elements combine in such a way that flat waves result. The frequency response is flat when the loudspeaker is voltage driven.

[62] That is: up to the distance where the waves are no longer flat and start to fan out.

[63] We could use a loudspeaker that covers an entire wall of the room and rely on reflections in the floor, the ceiling and the two adjacent walls to virtually extend the loudspeaker to infinite size.
This is still not a practical solution.

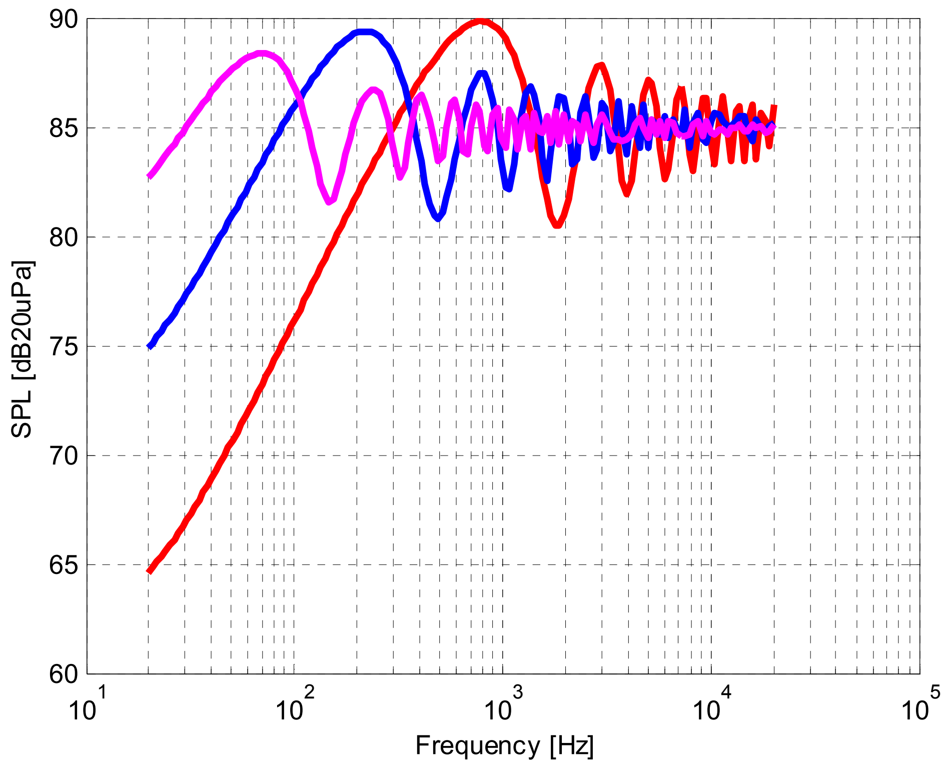

Figure 49 SPL of large flat loudspeaker at 3m distance.
Sizes are
red: 2 m × 2 m,
blue: 4 m × 4 m,
magenta: 8 m × 8 m.
Input voltage 2.828 V_{rms}; step-up ratio 150×.

13.6 Line Source - Cylindrical Waves

Introduction
Another approach towards coping with path length differences is the line source ESL. This solution holds real promise.
Because it lends itself well to realization using stators made out of isolated electrical wires, it is popular among amateur ESL builders[64]. For these reasons we will give a little more attention to this design approach.

Variation of SPL with Distance
We have seen in chapter 7 that the SPL of a line source decreases by 3 dB each time that the distance doubles. This compares to 6 dB per doubling of the distance for a point source design, such as the three way design of chapter 12.

One-way Line Source
In principle, we can realize a line source by building a high and narrow ESL.
The height must ideally be infinite, but a loudspeaker that extends from floor to ceiling works equally well, because reflections in both extend the virtual height to infinity.
The width should be no more than 7 cm to avoid excessive horizontal directivity.
With a typical ceiling height of 2.60 m, this results in a diaphragm area of 0.18 m^2.

Voltage- or Current Drive; Equalizing Filter
A voltage driven line source has a frequency response that increases with frequency at a rate of 3 dB/oct. Figure 50 shows a simulation result that demonstrates this property.
See also Baxandall [2, Baxandall, page 166].

Because the input impedance of the loudspeaker is almost purely capacitive, the current increases with frequency at a rate of 6 dB/oct when the loudspeaker is voltage driven.
It follows that driving the loudspeaker with a constant current causes a frequency response that drops with frequency at a rate of 3 dB/oct. See Figure 51.

For a flat frequency response, we must apparently drive the loudspeaker by a voltage that drops with frequency at a rate of 3 dB per octave or a current source that increases with frequency at a rate of 3 dB per octave.
See Baxandall [2, Baxandall, page 166].

[64] Fikier [4, Fikier] describes a basic construction.

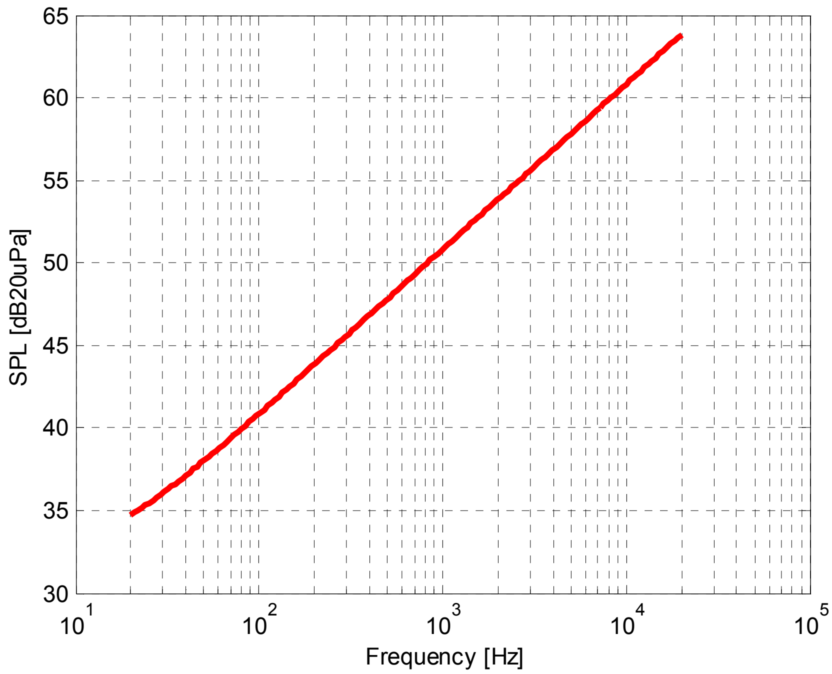

Figure 50 **Frequency response of a voltage driven line source.**
For this simulation, a width of 2 cm and a height of 200 m was used.

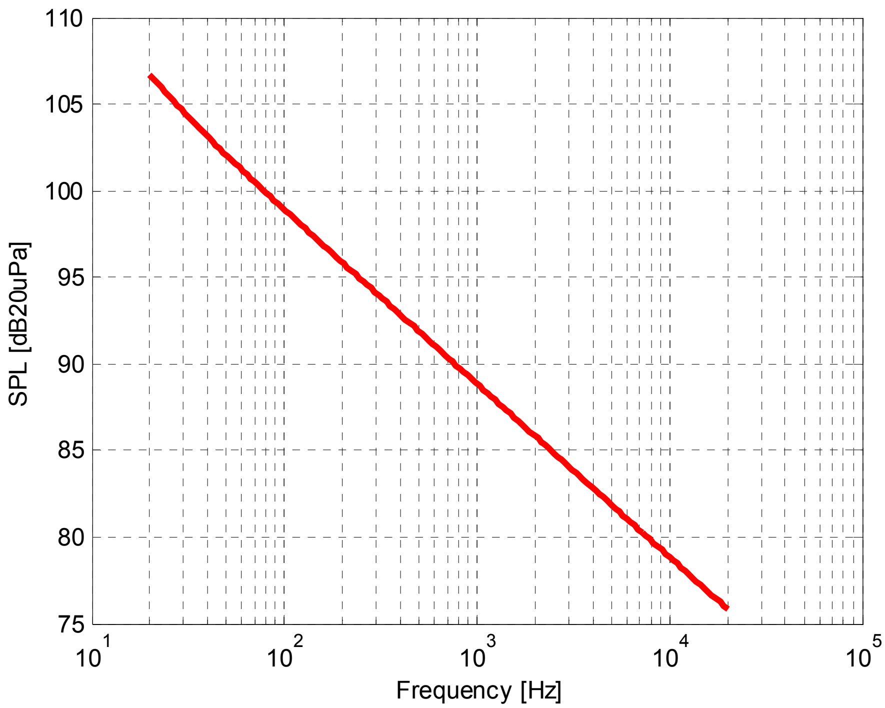

Figure 51 **Frequency response of a current driven line source.**
For this simulation, a width of 2 cm and a height of 200 m was used.

We can realize current drive conditions using a step-up transformer with a series resistance; voltage drive conditions exist when we omit the series resistance.
In the latter case, the loudspeaker presents a difficult load impedance to the amplifier. Further, the resonance caused by the panel capacitance and the leakage inductance of the step-up transformer must be damped, see section 15.5.

In case of current drive, we can achieve the required +3 dB/oct slope by inserting a suitable filter between the pre-amplifier and the power amplifier. Figure 52 gives an example of such a filter.
In the case of voltage drive, we can achieve the required -3 dB/oct slope by inserting a different filter[65] (Figure 53).

The following table gives the component values of the filters.
It is wise to take for the resistors types with 1% tolerance and for the capacitors types with 5% tolerance.

Resistor	Value [Ω]	Capacitor	Value [F]
R_A	10k		
R_B	10k		
R_1	1k		
R_2	1k	C_2	4n7
R_3	1k8	C_3	10n
R_4	3k9	C_4	22n
R_5	8k2	C_5	47n
R_6	15k	C_6	100n
R_7	33k	C_7	150n
R_8	68k	C_8	330n
R_9	120k	C_9	390n

Figure 54 shows the frequency response of the equalizer filter for the voltage driven line source[66].
Figure 55 shows the same response where the slope of -3 dB/oct has been removed such that the remaining error is more clearly visible.

Note that the increased bass response that we expect due to reflections in floor and ceiling in the case of a point source ESL do not occur in this case.
Or rather, they have already been taken into account.

[65] Note that both filters are inverting.

[66] In fact, it shows the impedance of the RC feedback network.

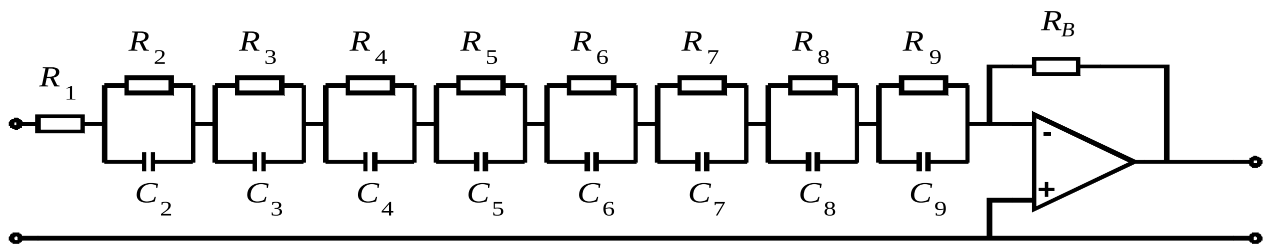

Figure 52 **Equalizing filter between pre-amplifier and power amplifier to realize a +3 dB/oct slope of the drive current.**

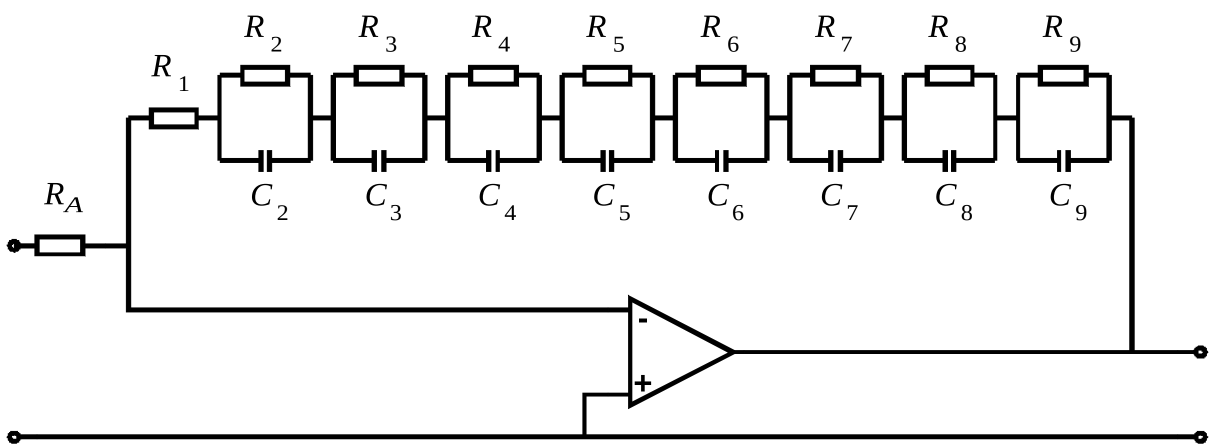

Figure 53 **Equalizing filter between pre-amplifier and power amplifier to realize a -3 dB slope of the drive voltage.**

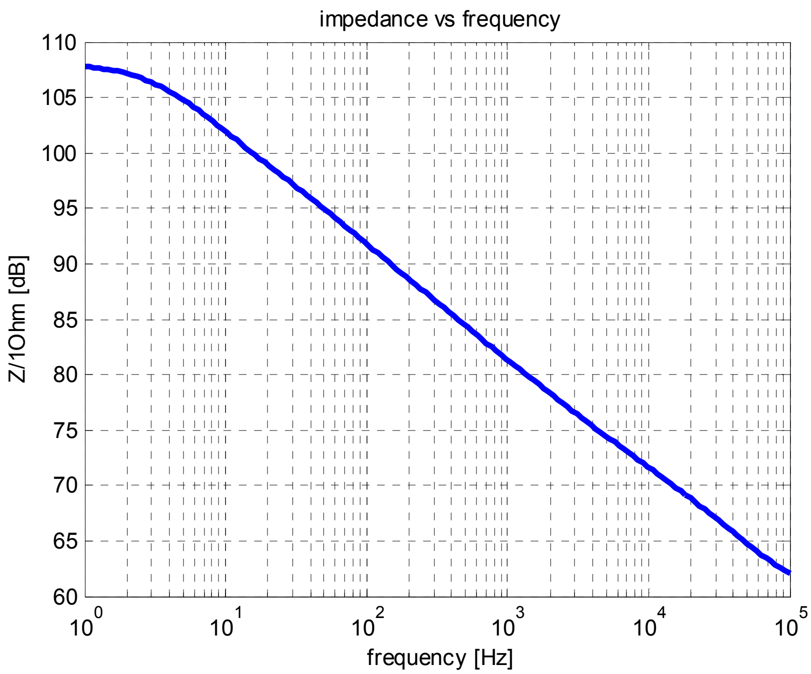

Figure 54 **Frequency response of the equalizing filter for a voltage driven line source. (in fact, the impedance of the feedback circuit).**

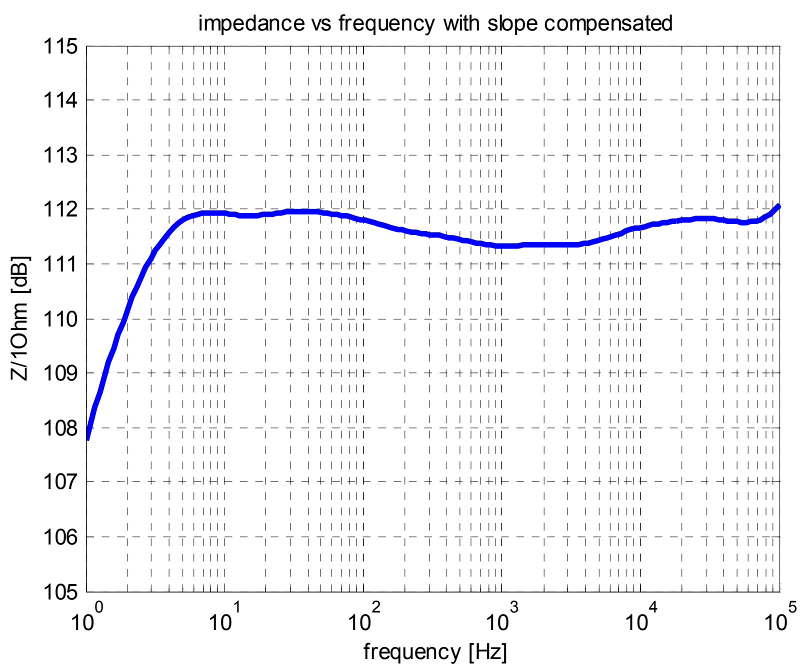

Figure 55 **Frequency response of the equalizer for the voltage driven line source, with the nominal slope removed.**
The filter for the current driven ESL has exactly the opposite ripple.

Extended Bass Frequency Range using Voltage Drive

A big advantage of the line source with voltage drive is that due to the -3 dB/oct slope of the equalizer, the large drive voltage needed to reproduce bass tones is only needed when these bass tones actually occur.

Compare this to the three-way design of chapter 12, where the large drive voltage needed for bass reproduction is also needed for mid-range and high frequencies.

A line source with voltage drive can therefore play quite a bit louder when there is not a strong bass, for example when you play singing accompanied by acoustical guitars.
For playing pipe organ music, you must then decrease the SPL to the level that is the normal limit for the three way design.

This advantage of the voltage driven line source turns into a similar disadvantage if we opt for the current driven version.

As an example, I have simulated the SPL that you can achieve at 3 m distance as a function of frequency.
I used a width of 7 cm; polarizing voltage of 5000 V and peak drive voltage of 7500 V.
We can obtain such a drive voltage for example from an audio amplifier capable of delivering a peak voltage of 50 V, combined with a step-up transformer with 1:150 step-up ratio[67].
The diaphragm-to-stator distance has been taken 2 mm; this is a bit more than usual because we expect that the narrow panel width will result in larger diaphragm excursions.
I have assumed a reduction of the panel capacitance due to the holes of a factor 0.8, which gives a panel capacitance of 124 pF per meter of panel height.

Figure 56 and the following table gives the results:

Frequency	**SPL at 3m distance [dB20uPa]**
20	82
50	85
100	88
200	91
500	95
1000	98

At 50 V_{peak} the ESL63, having a sensitivity of 86 dB at 1 m distance will produce at 3 m distance an SPL of 20*log10(50/(sqrt(2)*sqrt(8))) + 86 – 20*log10(3) = 98.4 dB[68].
This line source can do the same for all frequencies above 1000Hz.

[67] Of course the amplifier will in this case be loaded with a transformed load capacitance of 7 μF. This corresponds to 1.1 Ω at 20 kHz.

[68] That is: it would if it could handle 50 V_{peak}.

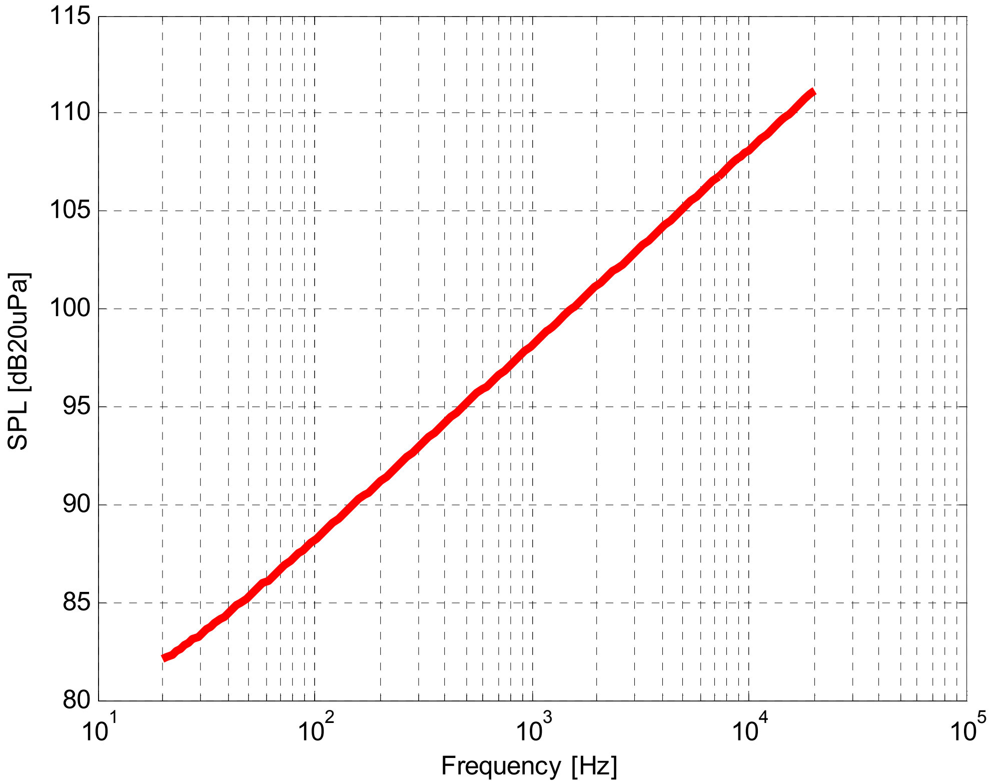

Figure 56 **SPL of a floor to ceiling ESL, width 7 cm, d = 2 mm, Vp=5000 V, drive voltage 7500 Vpeak.**

13.7 Segmented Line Source

We may want to build a loudspeaker of smaller height than floor-to-ceiling, in which case we must accept that the design deviates from an ideal line source of infinite height.
As we will see, this will lead to a ripple in the frequency response.

If we still want to achieve sufficient diaphragm area for bass reproduction, we must increase the width.

To avoid excessive directivity, we must then apply segmentation in the horizontal direction, such that all segments have the same height, equal to the height of the loudspeaker. The narrowest segments are used for the highest frequencies.
This kind of segmentation is very well compatible with stators built from stretched wires.

The cross-over filter can in this case be the ladder filter of Figure 57, where the capacitances of the segments themselves are used to achieve the filter action.
As already mentioned in section 12.5, such a ladder filter has the property that the lower frequency segments get a smaller drive current than the higher frequency segments, but this fits well to the requirements of the line source, where the high frequency drive current must receive a pre-emphasis of 3 dB/oct.

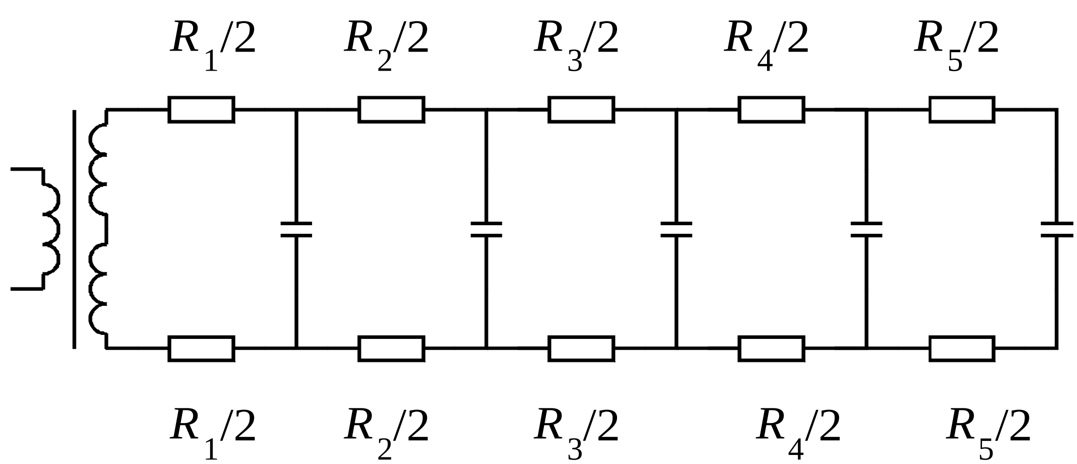

Figure 57 **Ladder type of cross-over filter.**
The capacitors are the segments of the ESL, high to low frequency from left to right.

We will take as an example the segmentation as used by Dutch ESL builder Edo Hulsebos, with the following parameters:

The segmentation:	18-12-6-2-2-2-6-12-18 wires = 78 wires
This way we arrive at 5-way design.	
Wire pitch:	4 mm
The total width will then be 78 × 4 = 312 mm.	
Loudspeaker height:	1.50 m
step-up ratio:	150×
Input voltage:	2.828 V_{rms}
Polarizing voltage:	5400 V
Distance between diaphragm and stator:	2.7 mm
Listening position:	axial, 3 m distance

The dimensioning of the cross-over filter is shown in Figure 58.

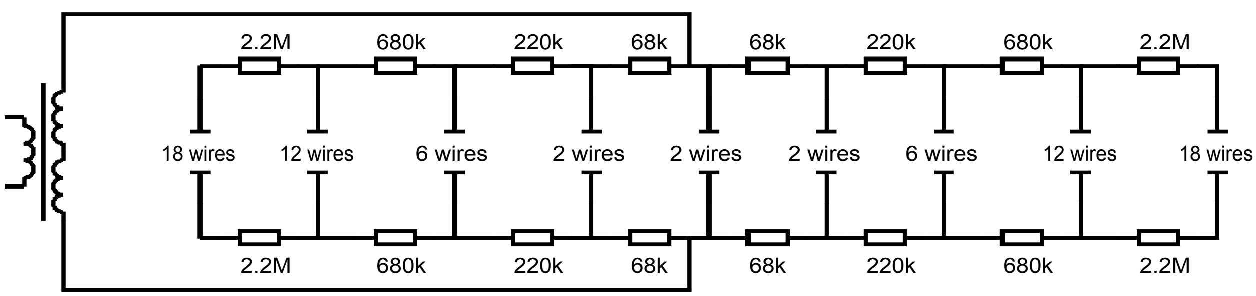

Figure 58 Segmentation of Edo's loudspeaker.

We simulate this loudspeaker using the modified Matlab script *ESL_ladder.m*.
The simulation does not take into account the reflections in floor and ceiling.
It also neglects any baffling.
For the simulation, we take the equal segments left and right of the center segments together.

To allow comparison to Edo’s simulation results, we take the segment capacitance as predicted by the formula for the flat plate capacitor (appendix IV) without any reduction to account for the openness of the stators.
The segment capacitances are then

2 wires	19.7 pF
4 wires	39.3 pF
12 wires	118.0 pF
24 wires	236.0 pF
36 wires	354.0 pF

Figure 59 shows the simulated segment voltages with the l.f. value set to 0 dB. The figure shows the filter action that the capacitances of smaller segments have on voltages of the larger segments.

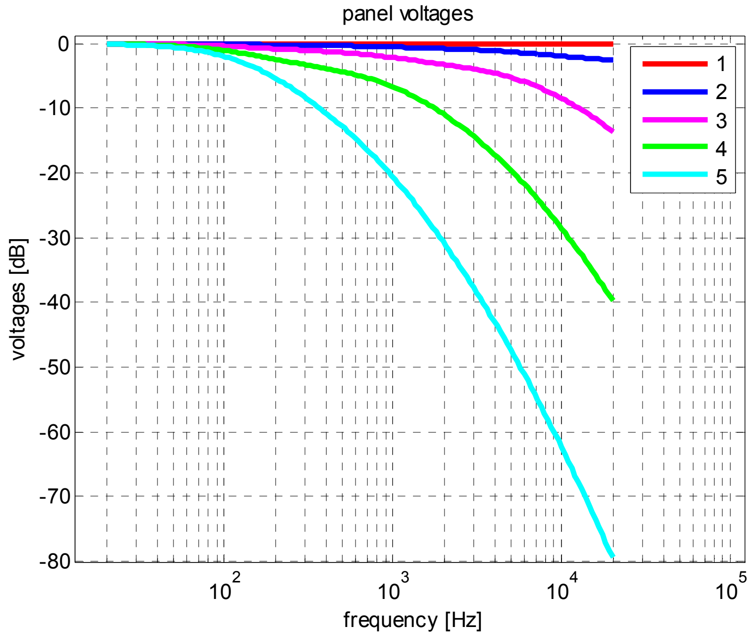

Figure 59 **Segment voltages, relative to l.f. value.**

Figure 60 shows the simulated segment currents, where we see that indeed the maximum currents for the bass segments are smaller than the maximum currents for the high frequency segments.

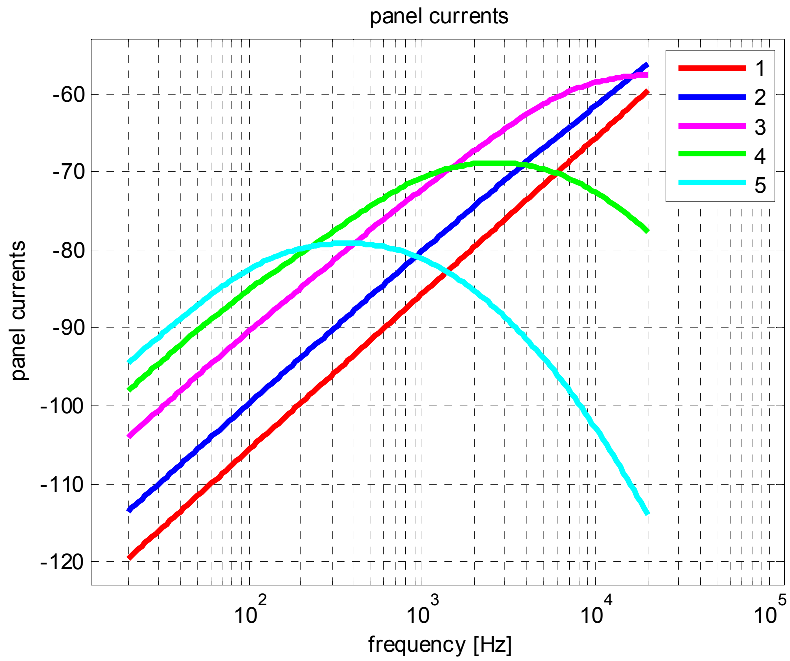

Figure 60 **Segment currents.**

Figure 61 shows the simulated response.
The order of the traces from l.f. to h.f. is: cyan, green, magenta, blue and red.
The black trace shows the sum of these five responses.

We see that the response as a lack of bass frequencies, which is due to the finite height of the loudspeaker.
We also see some some ripple, which is the result of the truncation error due to the finite height.

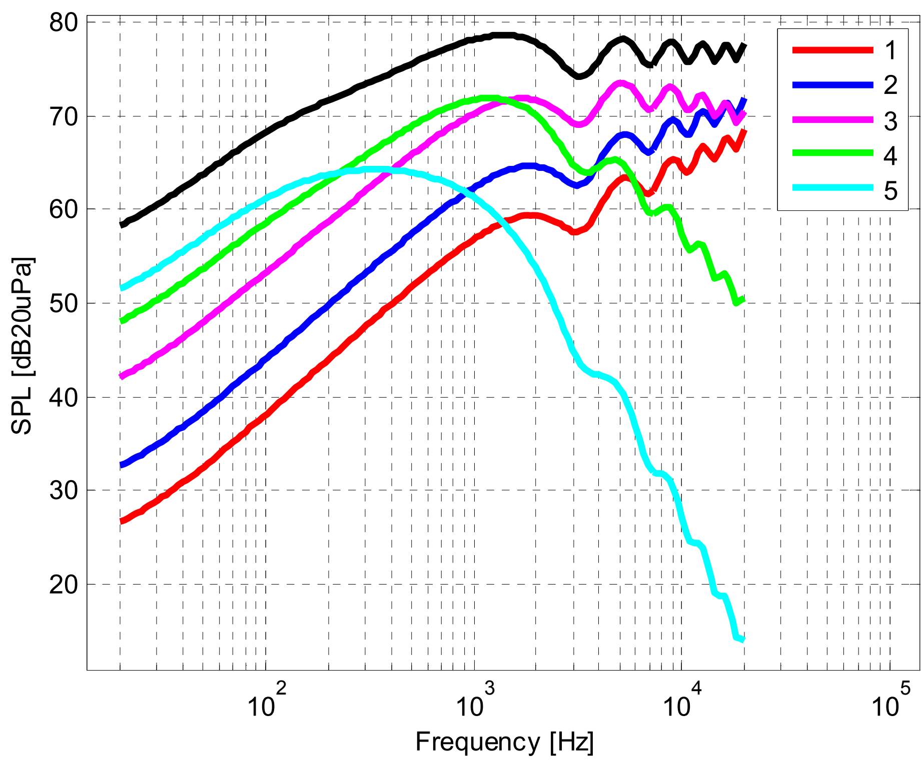

Figure 61 Frequency response of truncated line source with segmentation.

If the loudspeaker height would be from floor to ceiling, virtual images in floor and ceiling would effectively extend it to infinite height.
To investigate this situation, we will simulate the response we get when the length of the loudspeaker is extended to such a large value that it approximates infinity.
For that purpose we use the parameter *heightmultiplier* and set it to a value of 200.
This increases the loudspeaker height to 300 m while keeping the drive current of the panels per unit of height the same.
Figure 62 shows the simulated response.
Figure 63 shows the same plot zoomed in.
This curve fits to within a few tenths of a dB with Edo's simulation.

Note that the l.f. frequency range, which is about 100 Hz, can be extended down to 50 Hz by using the diaphragm resonance, as explained in section 9.11.

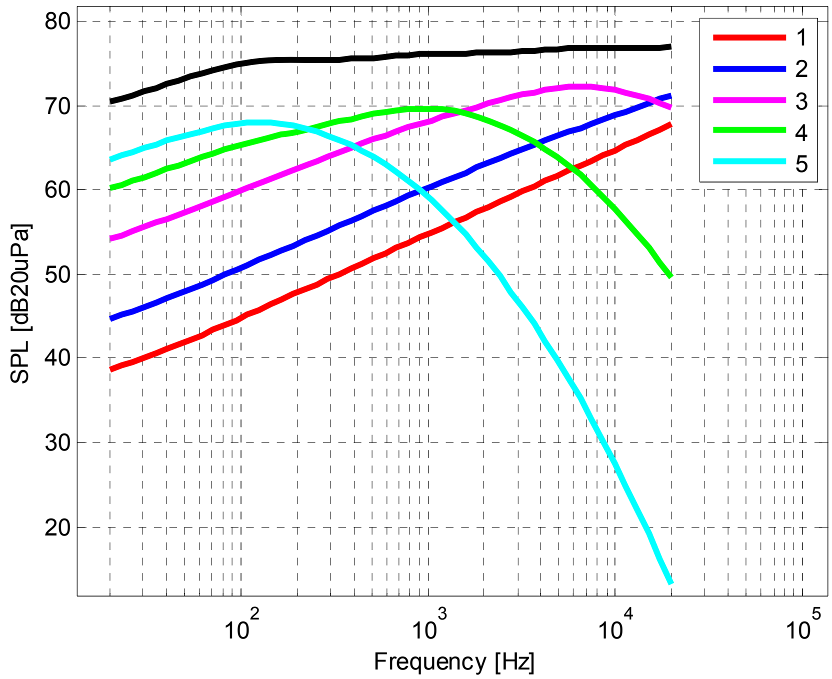

Figure 62 **SPL of segmented line source when the height is extended to 300 m.**

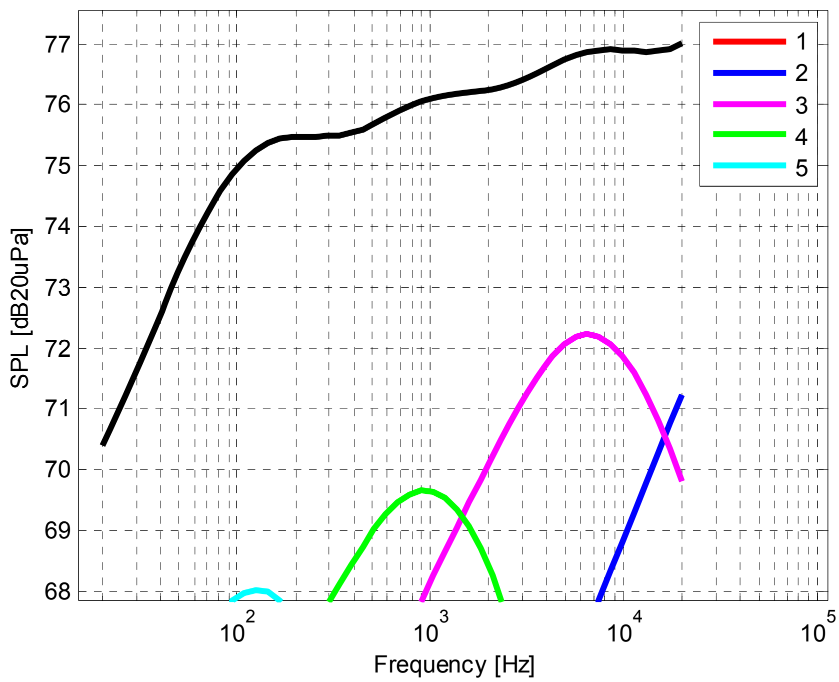

Figure 63 **SPL of segmented line source when the height is extended to 300 m zoomed in.**

Finally we will do a third simulation with this loudspeaker.
Because the loudspeaker height of 1.50 m is less than the ceiling height, the reflections in floor and ceiling will in reality form an infinite array of virtual loudspeakers with gaps in between.
We will simulate this situation, where we restrict ourselves to two images in the floor and two in the ceiling.
We assume a ceiling height of 2.50 m and we place the loudspeaker 30 cm above the floor.
The listening position will be 3m away from the loudspeaker, 1m above the floor.
Figure 64 shows the loudspeaker between floor and ceiling, together with four of its images.

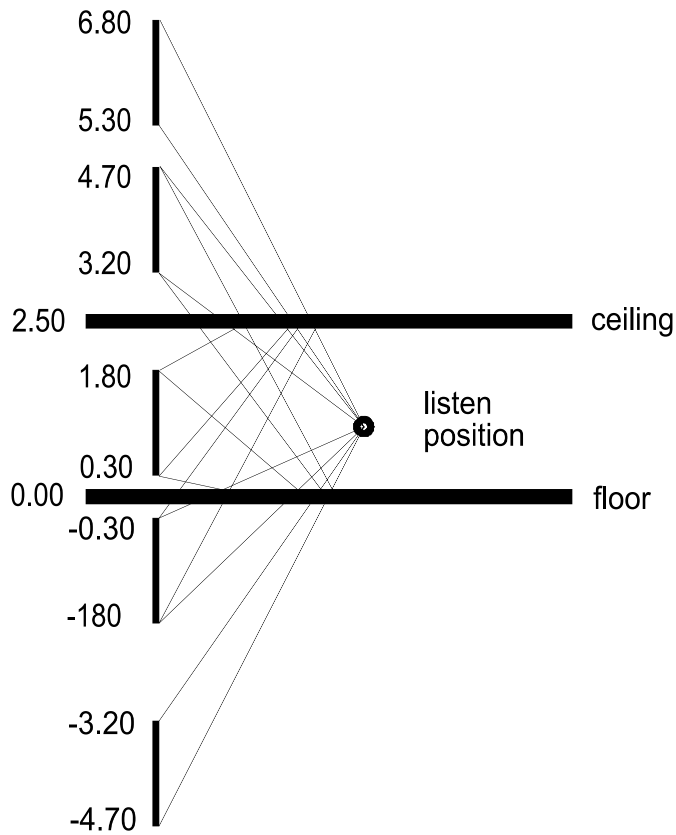

Figure 64 **Loudspeaker with 1.50m height between floor and ceiling. Four reflections included.**

Figure 65 shows the simulated SPL.
We see that this result is somewhere between the free space response of Figure 61 and the response of the infinite height line source of Figure 62.
The rapid ripples are due to the longer distances between the listening position and the edges of the images in floor and ceiling. Due to these long path lengths it takes only a small frequency change to fit an extra wave length on this path length.
The notch at 190 Hz can be explained by a path length difference of
$0.5 \times (343 \text{ m/s} / 190 \text{ Hz}) = 0.90 \text{ m}$ between the real loudspeaker and the first image in the floor. Based on Figure 64 this seems about right (the distance between the center of the loudspeaker and the center of its image is about 60 cm).

Note that again any baffling used in the loudspeaker has been disregarded.

While this simulation cannot pretend to give accurate results (because of the limited number of image loudspeakers and because of the omission of the baffling), it does give an idea of the trend.

Finally, note that the smooth response of the 3-way design (Figure 38) does not take any reflections in floor and ceiling into account. When these are added, we can expect ripples similar to the ones in Figure 65.

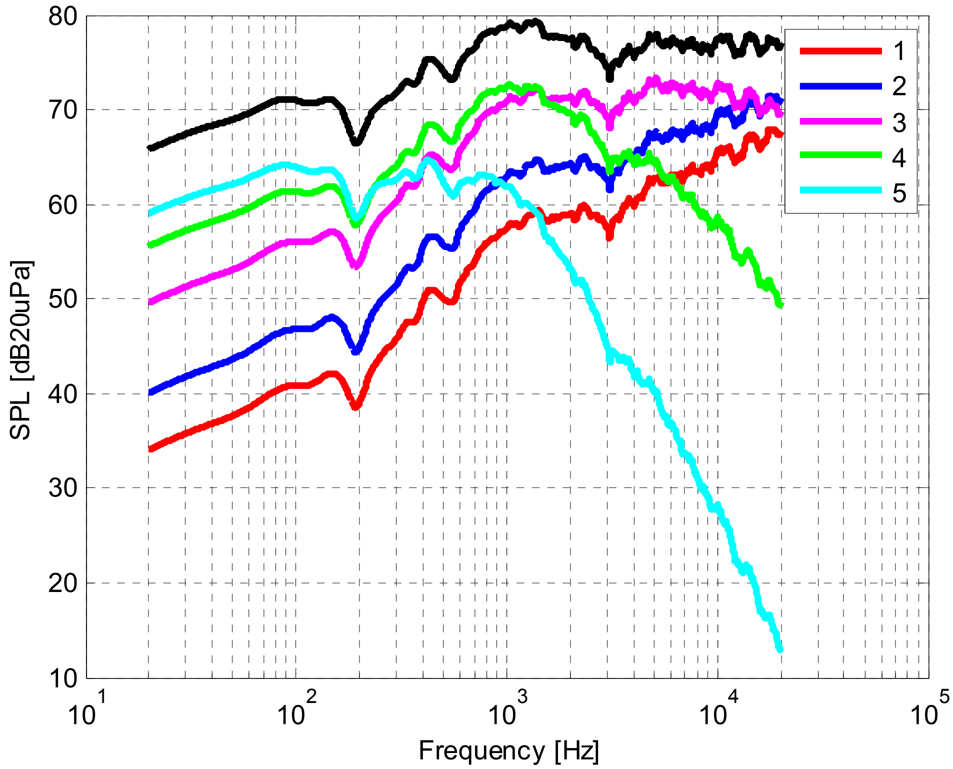

Figure 65 **ESL of 1.50 m height between floor and ceiling.**
Four reflections taken into account.

13.8 Curved Line Source - Cylindrical Waves

A third approach to the approximate realization of a line source is to build a tall loudspeaker that is curved horizontally (Figure 66).
This is the approach taken by the commercial ESL manufacturer Martin Logan.
Ideally, the loudspeaker should have the shape of a half cylinder.
In practice, less curvature is used and to avoid truncation errors the response should ideally be tapered off towards the edges.

Again, building a curved loudspeaker presents additional construction challenges.
One solution is to stretch the diaphragm only vertically.

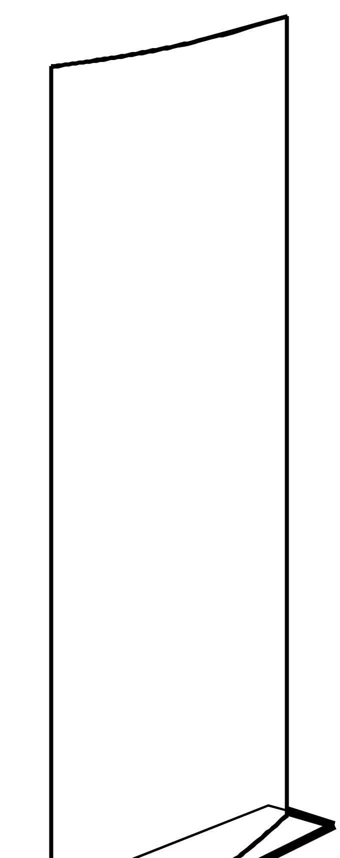

Figure 66 Curved loudspeaker.

14 Hybrid ESL

Up to now we have discussed only full range ESL designs.
A full range ESL is an ESL that handles the entire audio frequency range.
As we have seen in chapter 12 and chapter 13, it is difficult to reconcile the different requirements on a full range ESL.
In particular, it is difficult to reconcile the large area needed for bass reproduction with a small enough area needed for acceptable directivity at the higher frequencies.
We found that in the end we have a bulky loudspeaker and we still must compromise on the sensitivity.

On top of all this, the design challenges of a full range step-up transformer are considerable. See chapter 15.

Some designs solve these problems by adopting a hybrid solution, where the mid range and high frequencies are handled by a small ESL and the bass reproduction is done by an electro-dynamic woofer.

Roger Sanders [12, Sanders] is one of the advocates of the hybrid solution.

15 Step-up Transformer

15.1 Introduction

The step-up transformer is one of the most difficult parts of the loudspeaker to properly design. There are a large number of parasitic effects, which we will discuss in this chapter; and therefore there are a number of painful trade-offs to be made.
Many of these trade-offs become easier when we design a step-up transformer only for the mid-range and high frequencies; that is: when we build a hybrid loudspeaker.

No matter whether you buy transformers from a commercial party or from a fellow amateur or whether you decide to build transformers of your own: the transformer is likely to be the most expensive part of your loudspeaker.
Expect to spend a few hunderd Euro’s per piece.

15.2 Step-up Ratio

The step-up transformer consists of a primary winding and a secondary winding (the latter one with a center tap) on an iron core.
The *turns ratio* or *step-up ratio* is the number of turns of the secondary winding divided by the number of turns of the primary winding.
If for example the primary winding has 50 turns and the secondary winding has 5000 turns we have a step-up ratio of 100×.

The step-up transformer has a number of limitations that depend on its design parameters. We will discuss them in the following sections.

15.3 Ideal Transformer

Consider first an ideal transformer, which is driven at the primary side and loaded at the secondary side.
The ideal transformer has two simple properties:

1. The secondary voltage is the primary input voltage times the turns ratio.
2. The primary current is the secondary load current times the turns ratio.

These two properties are valid at all frequencies, and also at dc.

We will represent the ideal transformer by the schematic symbol of Figure 67.
The dots in the schematic symbol denote the polarity of the windings.

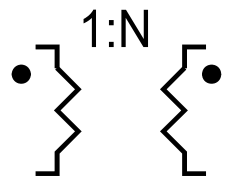

Figure 67 Ideal transformer.

The ideal transformer does not exist in real life. It is only an image of the mind. Nonetheless, it is useful because it can help model the behavior of the practical transformer.

Practical transformers deviate in their behavior from the ideal transformer. Figure 68 shows the schematic symbol for the practical, non-ideal, transformer.

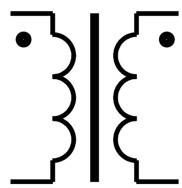

Figure 68 Non-ideal transformer.

In the following sections, we will investigate how a practical transformer differs from an ideal one.

15.4 Magnetizing Inductance

Magnetizing Inductance
An ideal step-up transformer would present no load (i.e. infinite impedance) to the amplifier output when the secondary winding is open circuited.

In reality, the primary winding has under these circumstances an inductive input impedance, which loads the amplifier.
We call this the *magnetizing inductance*.
The impedance of this magnetizing inductance drops as the frequency decreases.

As a result of the inductive input impedance, a drive voltage on the primary winding causes a current in that winding. We call this the *magnetizing current*.

If we want to build a model (an equivalent circuit) of the non-ideal transformer having this magnetizing inductance behavior, we can add an inductor at the primary side of an ideal transformer.

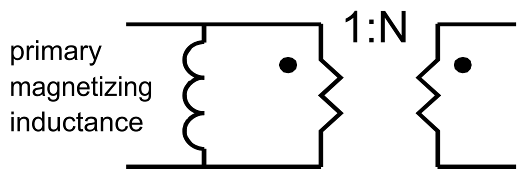

Figure 69 Magnetizing inductance added to ideal transformer.

As it turns out, we have the choice of placing the magnetizing inductance at the primary side or at the secondary side of the transformer.
If we place it on the secondary side, we must multiply the value by the square of the turns ratio.
So if the primary magnetizing inductance is 1 mH and the turns ratio is 1:100, the secondary magnetizing inductance will be 10 H.

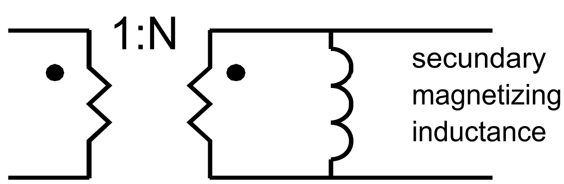

Figure 70 Magnetizing inductance on the secondary side.

The finite magnetizing inductance is one of the reasons that we cannot use a practical transformer at dc.

Magnetizing Inductance in the Step-up Transformer
In our ESL, the primary inductance of the step-up transformer must be high enough that the magnetizing current does load the amplifier too much.
If we want the impedance of the primary inductance to be > 8Ω (the typical load impedance for which many amplifiers are designed) at 20 Hz, the inductance must be > 0.064 H.

Another complication is that the primary inductance becomes a short circuit for dc.
If the amplifier has even the smallest dc offset at its output this will lead to a very large short circuit current, which could damage the amplifier or interfere with its operation.
It could also saturate the transformer core. See section 15.9 for a more detailed discussion of saturation due to amplifier offset..

Fortunately there is always some resistance in the circuit, composed of the copper resistance of the primary winding and the output resistance of the amplifier.
If the output offset voltage of the amplifier is very small, as it will be in a high quality amplifier, the resulting dc current will be within reasonable limits.
This is illustrated by Figure 79 in section 15.9.

Suppose for example that the offset voltage is 10 mV.
Suppose further that the total resistance is 100 mΩ.
Under these conditions the resulting dc current is 0.1 A.
As the amplifier is designed to deliver several ampères to an 8 Ω or 4 Ω load, that is entirely acceptable[69].

The primary inductance depends on the properties of the iron core, its dimensions and on the number of turns of the primary winding.
In fact, if we double the number of turns, the inductance quadruples.
Appendix X gives the formula for the inductance.

It seems, therefore, advantageous to take a large number of turns for the primary winding. There are, however, reasons for not wanting a large number of turns, as we will see shortly.

[69] That is: from the viewpoint of the amplifier.
Whether the transformer is also happy under these operating conditions remains to be seen (sections 15.9 and 15.13).

15.5 Leakage Inductance

Leakage Inductance
The amplifier output behaves, at least approximately, as a voltage source.
An ideal transformer, when driven by a voltage source, will behave at its secondary output also as an ideal voltage source.

In reality, the secondary output behaves as a voltage source in series with a small inductance, called the *leakage inductance*.
The leakage inductance exists because not all of the magnetic flux generated by the primary winding goes through the secondary winding and vice versa. Part of the flux leaks away, hence the name of the leakage inductance.
Toroid transformers[70] generally have less leakage flux and therefore they have smaller leakage inductance.

We can model the leakage inductance by adding an inductance in series with the secondary winding of an ideal transformer (Figure 71).

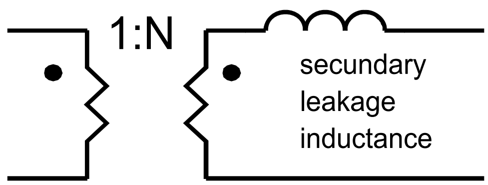

Figure 71 Leakage inductance at the secondary side.

As it turns out, we can also model the leakage of flux by placing a leakage inductance at the primary side. We must then reduce its inductance by a dividing the value by the turns ratio squared.

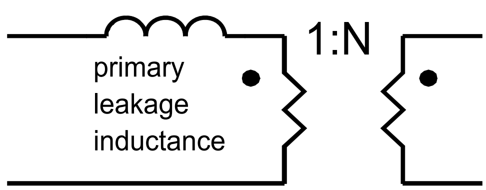

Figure 72 Leakage inductance at the primary side.

[70] Toroid transformers are the ones having a doughnut shaped core.

If we want to model both the magnetizing inductance and the leakage inductance, we can place each on the primary or the secondary side, so we have four possibilities. Figure 73 shows one of these.

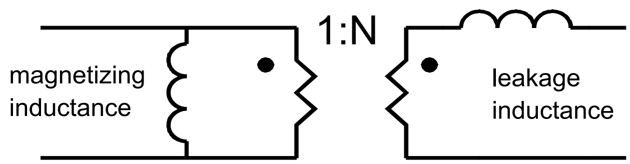

Figure 73 **Transformer with magnetizing current and leakage inductance.**

Leakage Inductance in the Step-up Transformer

The leakage inductance limits the capability of the transformer to drive a capacitive load. Such a load may consist of an ESL without series resistance, or it may be the capacitance of the wiring from the transformer to the loudspeaker.
Also the construction of the transformer itself will result in a parasitic load capacitance. For these reasons we want the leakage inductance to be as small as possible.

When the transformer is voltage driven (or short circuited) at the primary side, the load capacitance forms a resonant LC circuit with the leakage inductance.
It is important to have the resonance frequency of this parasitic resonator well above the audio frequency range.

If the load capacitance is external to the transformer, the resonance can be damped by using a series resistance of a value equal to the impedance of the capacitive load at the resonance frequency (Figure 74).

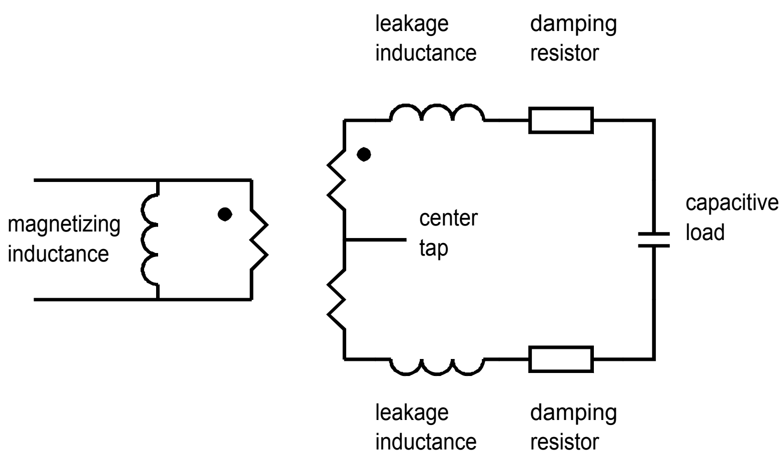

Figure 74 **Step-up transformer with leakage inductance, capacitive load and damping resistors.**

The leakage inductance depends on the number of turns.
If we double the number of turns, the leakage inductance quadruples.
Here we see a reason why we do not want to increase the number of turns without limit.

15.6 Parasitic Capacitance

A practical transformer has parasitic capacitance between the turns of the secondary winding.
As the secondary winding is often placed over the primary winding, one side of the secondary will also have a parasitic capacitance to the primary, see Figure 75.
If the construction of the transformer is symmetrical, the capacitances from both sides of the secondary winding to the primary are equal.
The construction of the transformer must be such that the parasitic capacitance is minimal. Even so, it remains large enough to be significant.

When the transformer is voltage driven (or short circuited) at the primary side, the parasitic capacitance forms a resonant LC circuit with the leakage inductance.
It is important to have the resonance frequency well above the audio frequency range.

It is unfortunately not possible to add damping to this resonance.

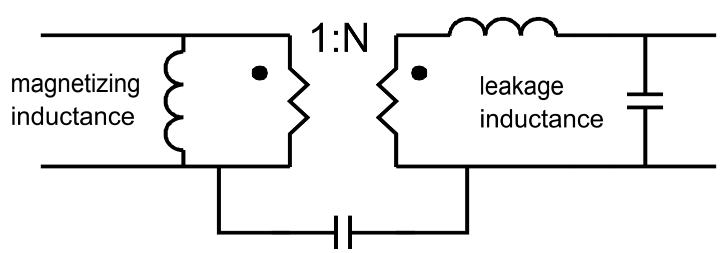

Figure 75 Transformer with parasitic capacitances.

In a step-up transformer for an ESL, it is important that the construction is optimized such that the parasitic capacitance does the least amount of damage.
For that purpose, the center tapped secondary winding must be made of two separate windings, identical but counter wound, such that the capacitance to the primary winding is connected to the center tap (see Figure 76).

It is easy to see that reversing the polarity of both windings would have the effect that the parasitic capacitance ends up on the outer sides of the secondary, where it will do maximum harm (Figure 77).

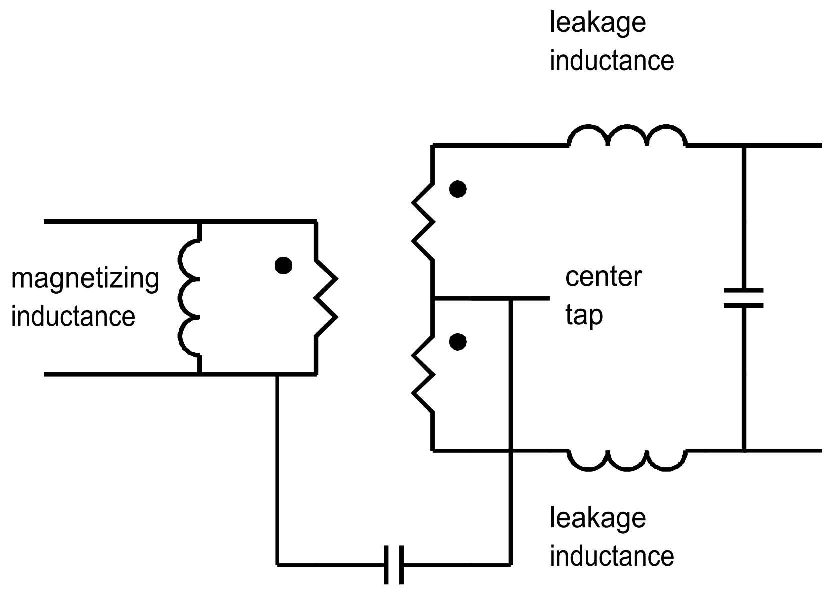

Figure 76 **Step-up transformer with optimal construction.**

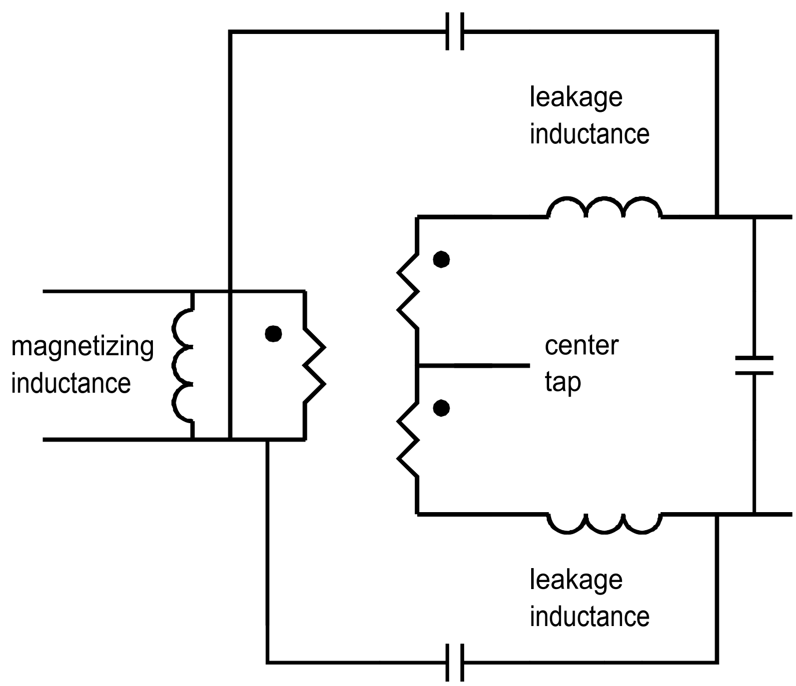

Figure 77 **Step-up transformer with wrong construction.**

15.7 Input Impedance for High Frequencies

In practice the input impedance of the primary winding, even with an open circuited secondary, tends to drop towards the highest audio frequencies.
The value can become (much) lower than the 4 - 8 Ω that the average audio amplifier is designed to drive.

An important cause of the drop of the high frequency input impedance is the parasitic capacitance between the turns of the secondary winding. This capacitance transforms back to the primary side and in the process it gets multiplied by the turns ratio squared.

Eddy currents in the core can also be responsible for a high frequency drop of the input impedance. These effectively work as a parasitic secondary winding with high resistivity, which is short circuited.

A third possible contribution is that the magnetic properties of the iron could change at higher frequencies.

We would like to use a transformer that maintains an unloaded input impedance of 4 Ω or higher over the entire audio frequency range[71].
A lower impedance might damage the amplifier by drawing larger currents than what it was designed for. Even if the amplifier survives, it may produce distortion if it has to drive a too low load impedance.
As most step-up transformers have a much lower unloaded input impedance at 20 kHz, it is important to verify that the amplifier can handle this load[72].

---

[71] Even if the magnitude of the impedance is sufficiently high (> 4 - 8 Ω), its phase angle may cause a load current that is phase shifted compared to the voltage. Some amplifiers may not be able to deliver such a phase shifted current.

[72] You can verify that the amplifier can drive the step-up transformer, or the entire ESL, by doing the following experiment.
Listen to the amplifier by a pair of high quality headphones.
Next, connect the transformer, or the entire ESL, to the amplifier output.
The ESL can be made silent by not supplying the polarizing voltage.
If the sound from the headphones does not change, the amplifier can handle the load.

15.8 Saturation by Bass Tones

When we drive the transformer by a primary voltage, the magnetic flux in the core swings up and down at the frequency of the driving voltage.
The flux swings are larger when the frequency is low, because the flux has more time to build up[73].
The flux swings are smaller if the number of turns of the primary winding is larger.
Appendix X gives the expression for the flux swing.

The magnetic flux that the core can handle is limited by the properties of the core, namely its physical dimensions and the material from which it is built.
When the flux gets larger than the maximum that the core can handle, we say that the iron *saturates*.
When the iron saturates, the core suddenly behaves as if it is not there and as a result the magnetizing inductance drops to almost zero.
This causes the magnetizing current to shoot up to large peak values. See Figure 78.
A saturating transformer can therefore damage the amplifier.
Even if the amplifier survives, the saturation current peaks will cause distortion.

Saturation does not set in very suddenly.
When the primary voltage approaches the saturation limit, saturation sets in softly and the magnetizing current becomes a bit distorted before the very large current peaks occur.

Saturation of the core limits the signal voltage that the transformer can handle at low frequencies.
A larger number of primary turns increases the voltage handling capability, but this goes at the expense of the high frequency properties of the transformer (leakage inductance and parasitic capacitance).
A better way to increase the signal handling is to use a core with a bigger cross sectional area.

Appendix X gives expressions for the maximum flux that the core can handle, the maximum number of Volt-seconds that we can impose on the primary winding and the maximum sinusoidal voltage we can apply at bass frequencies.

[73] More precisely: assuming a sinusoidal drive voltage, the flux swing is proportional to the area under a half sine wave. This area is voltage times time and therefore has the unit of Volt-seconds.
The saturation limit is set by the amount of volt-seconds that the transformer can handle.

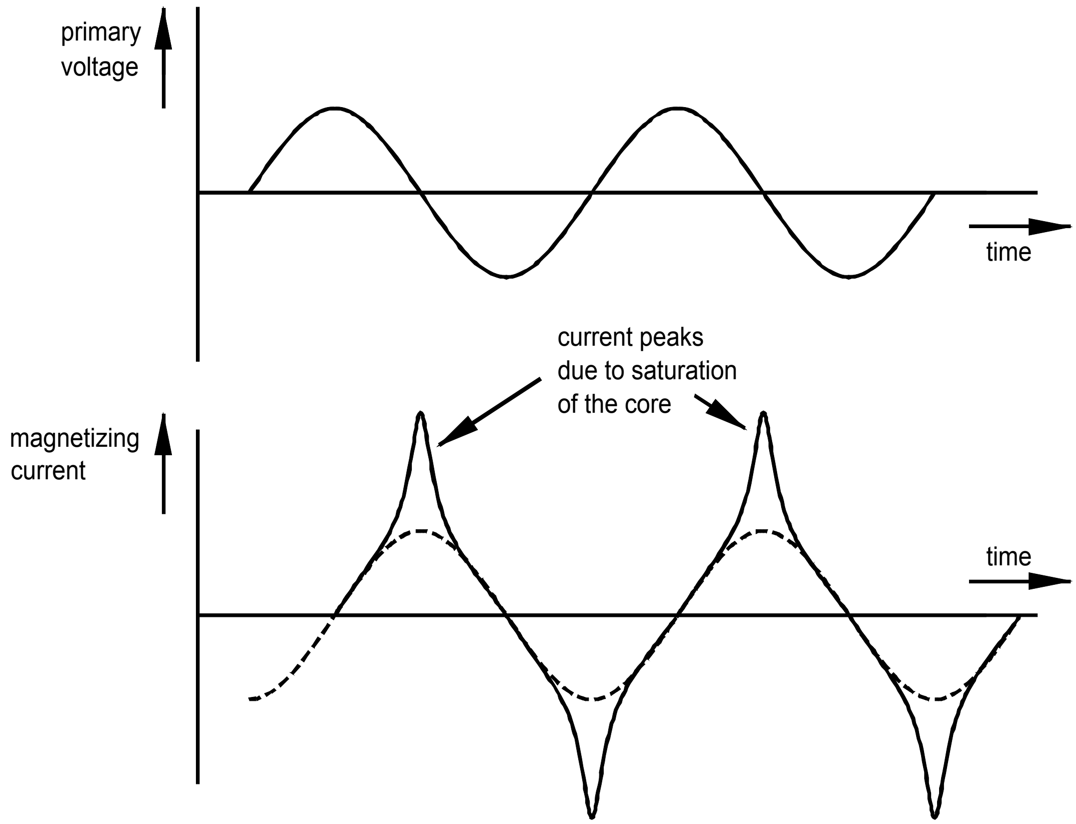

Figure 78 **Current peaks due to saturation of the core.**

15.9 Saturation due to Amplifier Offset

We have already briefly discussed the consequences of amplifier offset from the point of view of the amplifier (see section 15.4).
The amplifier has to supply a possibly large short circuit current.
Another consequence is that the number of volt-seconds in the dc offset grows to infinity as time goes by. This will cause any transformer to saturate within seconds.

Again, the series resistance (copper resistance and amplifier output resistance) in the circuit comes to the rescue.
As the magnetizing current grows, an increasing voltage drops over this resistance.
That voltage drop subtracts from the dc offset and soon reaches an equilibrium where the voltage drop is as large as the offset voltage.
At that point, no offset voltage is left to drive the core and the number of volt-seconds does not increase any further.

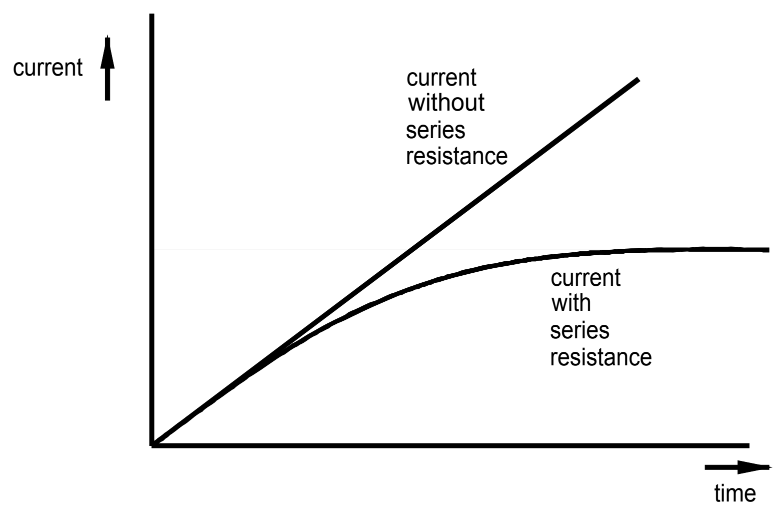

Figure 79 Current due to amplifier offset, with and without resistance in the circuit.

An air gap in the magnetic circuit dramatically reduces the magnetizing inductance, see appendix X for details.
While this sounds as a bad thing, it works to prevent saturation, because the same amount of volt-seconds leads to a much larger primary current.
In these circumstances the point where the entire offset voltage drops over the resistance is reached after a shorter time and the number of volt-seconds that occurs in this shorter time is smaller.

We see that an air gap helps to reduce the danger of saturation by a dc offset voltage.
As toroid transformers do not have an air gap, they are more vulnerable for dc offset than transformers made with an E-I core or a C-core.

Sometimes a series capacitor is used to block the dc voltage.
Due to the inevitable impedance that such a capacitor represents, the distorted magnetizing current causes a distorted primary voltage on the transformer.
Such a capacitor should therefore be avoided.

15.10 Properties of the Iron Core

The properties of the core depend on several parameters:

1. the properties of the iron
2. the cross sectional area of the core
3. the path length of the magnetic circuit
4. the thickness of the stacked iron sheets of which the core is built

Primary Inductance

The relative magnetic permeability of the iron is one of the factors that determine the magnetizing inductance.
That inductance also depends on the cross sectional area of the core, the length of the magnetic circuit and the number of turns of the primary winding.
The primary inductance also depends strongly on whether the magnetic circuit has an air gap and, if so, how big it is. See appendix X for more details.

Saturation

The saturation flux density of the iron together with the cross sectional area of the core determine the saturation flux.
The saturation flux determines, together with the number of turns of the primary winding, how much primary signal voltage the transformer can handle before the core saturates. See appendix X.

Distortion

A magnetic core consists of many microscopic volumes, called *Weiss domains*, which are fully magnetized. In a macroscopically unmagnetized core, the orientation and the size of the Weiss domains is random and such that their magnetization cancels on a macroscopic scale.

When the core gets magnetized by a voltage and a current on its primary winding, the boundaries between these domains shift such that macroscopically a net resultant remains.
In some iron mixtures that process happens more smoothly than in others.
In the latter case, distortion can occur, or the jerky movements of the boundaries becomes audible as a signal dependent noise.
Both the distortion and the noise are present only in the magnetizing current.

15.11 Maximum Feasible Turns Ratio of a Single Transformer

We have seen in this chapter that the choices to be made in transformer design are conflicting. A large turns ratio requires a large number of secondary turns and a small number of primary turns. A large number of secondary turns, however, limits the high frequency range of the transformer due to leakage induction and parasitic capacitance. At the same time, a small number of primary turns leads to a large magnetizing current and to saturation at strong bass tones.

The trade-off's that have to be made limit the feasible step-up ratio to about 100× for a transformer covering the entire audio frequency range (full range transformer).

For a hybrid loudspeaker, where the electrostatic part does not have to handle the bass, the saturation requirements are much relaxed and as a result a larger step-up ratio is possible without comprimising the high frequency properties.

15.12 A Higher Step-Up Ratio using Two Transformers

A higher step-up ratio can be obtained by using two transformers interconnected as in Figure 80.
The primary windings are connected in parallel and the secondary windings are connected in series.
The node where the two secondary windings are joined now serves as the center tap.

Note that the transformers have been combined in such a way that the parasitic capacitance between the secondary winding and the primary winding is connected to the center tap. If you connect the secondary windings incorrectly, one or both interwinding capacitances will be connected to the end of the secondary (Figure 81).

By using two transformers with a turns ratio of 1:75 we can in this way achieve a total step-up ratio of 150×.
Of course the primary input impedance of the two transformers work in parallel, and that halves the load impedance to the amplifier. We must make sure that the amplifier can drive the combined load.

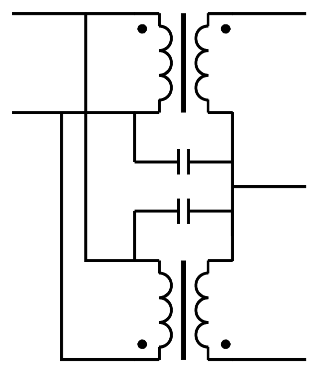

Figure 80 Two transformers correctly combined to get a higher step-up ratio.

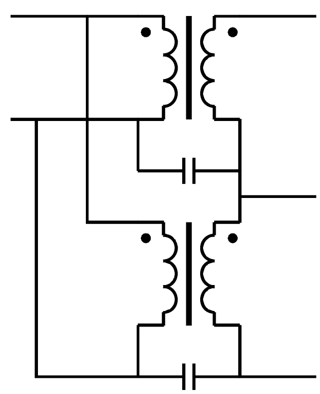

Figure 81 Two transformers incorrectly combined – the bottom transformer has its capacitance between primary and secondary windings in a harmful place.

15.13 Amplimo and Plitron

Menno van der Veen [6, van der Veen] has designed in 1998 two toroid step-up transformers for electrostatic loudspeakers, having step-up ratio's of 50× and 75×.
The Dutch company Amplimo [7, Amplimo] used to have these transformers in their program. They do not appear on their website anymore, but Amplimo might still be able to supply them.
The same transformers are still available from a Canadian company called Plitron [8, Plitron].

Plitron supplies the following technical data:

Type Number	**PAT4133-ES**	**PAT4134-ES**
Step-up ratio	1:50	1:75
Primary dc resistance	0.1 Ω	0.1 Ω
Secondary dc resistance	190 Ω	273 Ω
Saturation limit	17.9 V_{rms} @ 50 Hz	17.9 V_{rms} @ 50 Hz
Secondary inductance	719H	1600 H
Primary inductance, calculated from step-up ratio	288 mH	284 mH
Secondary Leakage Inductance	15 mH	22 mH
Secondary internal capacitance	700 pF	800 pF
Resonance frequency	32 kHz	25 kHz
Primary impedance at 20kHz	2.272 Ω	1.0 Ω

Discussion

Identical primary windings and identical cores
The specified secondary magnetizing inductance, when transformed back to the primary side using the turns ratio, results for both transformers in almost the same primary magnetizing inductance.
The primary dc resistance and the saturation voltage at 50 Hz are exactly the same.
We therefore suspect that both transformers have identical cores and identical primary windings.

Primary input impedance at 20 kHz
If we calculate the impedance of the secondary internal capacitance at 20 kHz and transform that back to the primary side using the turns ratio, we arrive at 4.5 Ω for the first transformer and 1.8 Ω for the second one.
As the specified input impedances are significantly lower, other effects must play a role too, or the secondary capacitance is larger than specified.

Resonance Frequency
The resonance frequency of the specified leakage inductance and the specified internal capacitance comes to 49 kHz for the first transformer and 38 kHz for the second one. As this is much higher than the specified resonance frequencies we might suspect that the actual internal capacitance is larger than specified.
Even so, the specified resonance frequency is above the audio frequency range.

Combination of two transformers to get a 1:150 step-up ratio
If two of the 1:75 transformers are combined in the way described in section 15.12 we can achieve a total step-up ratio of 1:150. The center taps of both transformers will in that case not be used.
The amplifier must be capable of driving a load that drops to 0.5 Ω at 20 kHz!

We do not know if the construction of the transformer is symmetrical, i.e. if the parasitic capacitance from each end of the secondary winding to the primary is the same.
If this is not the case, it will help to find out which secondary node has the highest capacitance and connect the pair such that these nodes are joined to form the center tap, as discussed in section 15.12.

Saturation voltage
The transformers saturate at a primary voltage of 17.9 V_{rms} at 50 Hz.
This corresponds to a peak voltage of 25.3 V_{peak}.
With a total step-up ratio of 1:150 the peak secondary voltage is 3795 V_{peak}.
The voltage handling capability is proportional to the frequency, so at 100 Hz the transformer can handle a primary voltage of 50.6 V and produces a secondary output voltage of 7590 V_{peak}.
The data sheet does not specify at what increase of the magnetizing current they consider the core to be saturated.

Saturation by dc offset voltage
We can expect that the transformers, being of the toroid type without an air gap, are susceptible to saturation by amplifier output offset.

Let's try to estimate how much offset they can handle.
We know that the transformers saturate at a primary voltage of 17.9 V_{rms} at 50 Hz.
This rms voltage corresponds to a peak voltage of 25.3 V.
With a magnetizing inductance of 286 mH we find that the peak magnetizing current is 281 mA_{peak}.
Therefore a primary dc current of 281 mA will saturate the core.
If we want to loose no more than 10% of voltage handling capability (corresponding to approximately 1 dB) to dc offset, the dc current must be no more then 28.1 mA.
At a dc resistance of 0.1 Ω the dc offset voltage of the amplifier must be no more than 2.81 mV. This is quite a restrictive number to ask of the amplifier.
This is therefore a real problem with this transformer.

The offset problem is often solved by placing a 1 Ω resistor in series with the primary winding. That way, we can allow the amplifier to produce 28.1 mV offset.
As the combined input impedance of two transformers at 20 kHz drops to 0.5 Ω we must bypass the resistor for high frequencies by a capacitor. See Figure 82.
If we want the capacitor to have an impedance of 0.5 Ω at 2000 Hz, it must have a value of 160 μF. Because of the large value, this will be a bipolar electrolytic capacitor.
We must make sure that the ESR (equivalent series resistance) of the capacitor is much smaller than 0.5 Ω.

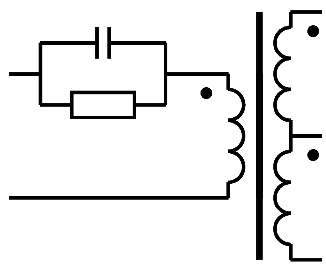

Figure 82 **Series resistor, bypassed by a capacitor can help to combat saturation due to amplifier offset.**

Even this solution for the offset problem is not without disadvantages.
At bass frequencies the shunt capacitor has negligible effect and the magnetizing current causes a voltage drop over the resistor.
Because the magnetizing current can be rather distorted due to saturation of the core, so is the voltage drop.
The remaining voltage driving the primary of the transformer is then also distorted.

The best solution, therefore, is to somehow control the amplifier offset to less than 2.8 mV.

15.14 DIY Step-up Transformer

The Dutch ESL forum [14, Dutch ESL forum] used to have a participant called Marc Schroeyers, who contributed very valuable information on how to construct your own step-up transformer. Unfortunately, Marc has withdrawn from the forum, but chances are that he is still listening in.

16 Construction Details

16.1 Introduction

In this chapter we discuss some practical construction details.
More construction details can be found in [4, Fikier], [5, Rehorst], [12, Sanders] and [13, ESL club].

16.2 Stator

Stator Size

The largest practical stator width without intermediate support for the diaphragm depends on the polarizing voltage and the diaphragm-to-stator distance.
When using a diaphragm-to-stator distance of 1.5 - 2.5 mm and a polarizing voltage of 4000 – 5500 V is about 20 cm.
With stators made from perforated plates it is therefore convenient to build the loudspeaker from elements of perhaps 40 cm × 20 cm. This is the choice that was made for the Quad ESL63.
A larger width, which is probably more convenient when using stators made from stretched electricity wires, requires that additional support for the diaphragm is provided in one or more locations along the width of the panel.

Stator Isolation

It is wise to isolate the stator electrodes.
First of all, this increases safety.
Secondly, when arcing occurs (as it surely will) the isolation limits the discharge current. Without such a current limitation a loud spark will result, which could damage the diaphragm.

The isolation of the stator surface should not be perfect, but that is not a problem because perfect isolation is not even practically possible.
With perfect isolation, the surface of the stator might charge up until it reaches the same electrical potential as the diaphragm.
This would reduce the effective polarizing voltage to zero.
Using a “leaky” isolation ensures that any charge that might build up on the surface will be drained to the metal underneath.

A too thick isolation layer increases the distance between stator and diaphragm, at the expense of the sensitivity of the loudspeaker.
The loss of sensitivity is not as large as we might think, because the electrically effective thickness equals the actual thickness divided by the relative dielectric constant (epsilon) of the material. The epsilon of PVC[74], for example, is 3.4. Therefore, an isolation thickness of 1 mm in PVC is equivalent to only 0.3 mm in air.

[74] The value depends on the type of PVC.

Stator of Perforated Sheet Metal

One way to realize the stators is to use perforated sheet metal.

The front and rear stators can be glued together with isolating spacers around the edge. The diaphragm will be glued between the spacers. See Figure 83.
You can let the spacers stick out to facilitate mounting of the panels in some kind of frame.

A disadvantage of stators from perforated sheet metal is that the area of the spacers does not contribute to the active loudspeaker area but it does have a parasitic capacitance. The parasitic capacitance works in parallel with the stator capacitance of the active part. If you calculate the area of the spacer, you will be surprised how much it is compared to the active area of the panel. On top of that the epsilon of the spacer material considerably increases the parasitic capacitance. We must therefore change the dimensioning of the drive resistors and the filter capacitors in the cross over filter to compensate for the parasitic capacitance, as discussed in section 12.5.

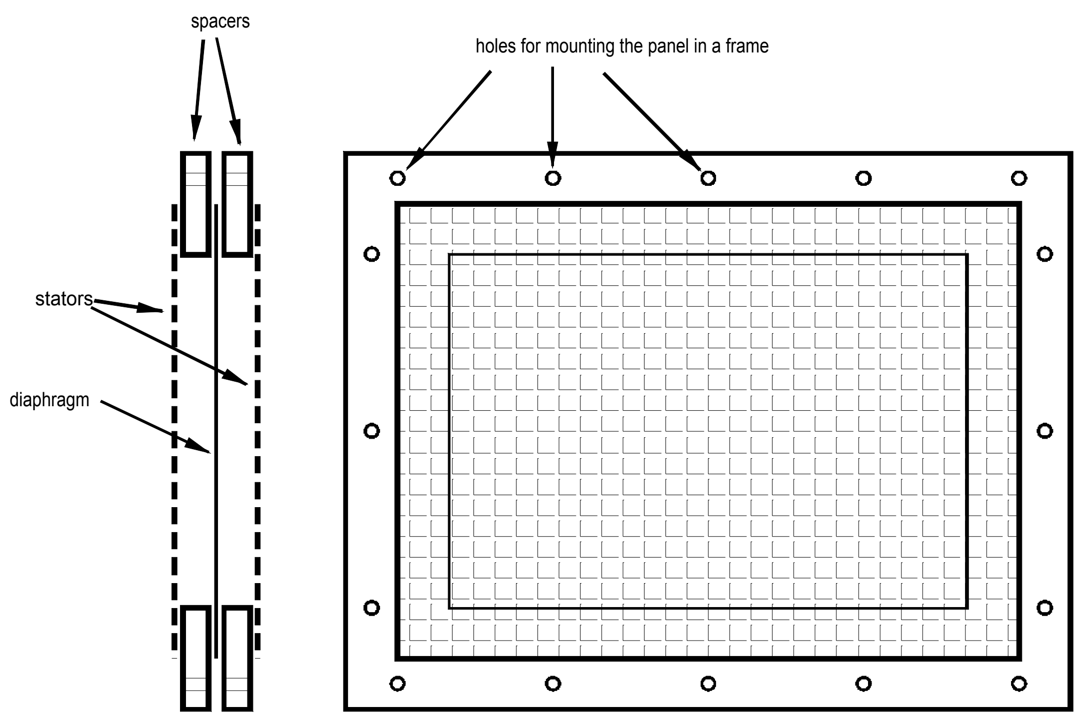

Figure 83 **Panel built from perforated metal plates. The spacers stick out to allow mounting in a frame.**

The most practical choices for the stator material are steel and aluminum.
Steel has the disadvantage that it must be protected against rust.
Aluminum has the disadvantage that we cannot solder a copper wire to it.
Therefore some other way has to be found to make reliable contacts to the stators.
Because a thin layer of oxide forms on the surface of the aluminum, that is more difficult than it seems. A poor contact can lead to distortion.

In order to maintain sufficient acoustical transparency, at least 40% of the area must consist of holes. A higher percentage (50 - 60%) would be preferred.
For the same reason, the plates must be as thin as possible; 1mm is a good choice.
The holes must be 3 – 5mm in size[75] and may be round or square.
The holes of the front and rear stators must be aligned.

The holes in the metal are made by punching.
This leaves the edges of the holes sharp on one side of the plate and rounded on the other side.
Because sharp edges promote arcing, the rounded sides of the stator plates must be on the inside, facing the diaphragm.

Perforated sheet aluminum of 1mm thickness is readily available from hardware stores in a size of 50cm × 100cm. The holes are square, 5mm × 5mm and have a 7.5mm pitch.
The holes cover 44% of the area.
This metal sheet is extremely floppy.
Fortunately the stiffness improves a great deal when panels with a width of about 20cm are built by gluing together a sandwich of a front stator and a back stator with spacers in between.

The metal sheet must be isolated by painting or powder coating.
This increases a bit the thickness.
When you buy the metal sheets, they are covered in a layer of grease, which must be removed before painting.
Take care that the diameter or size of the holes is not reduced too much by the paint.

A disadvantage of perforated metal stators is that it is difficult to build a squacker panel with a tweeter in the middle of it as in Figure 32 or Figure 33.
Any solution will end up having much area lost to the spacers.

[75] In fact, smaller holes can be used as well, but the risk increases that they will be closed by the isolating paint or coating.

Stator of Fiberglass PCB Material
Quad uses perforated Printed Circuit Board to construct the stator plates.
Instead of the standard thickness of 1.6mm Quad uses PCB plates of 1.3mm.
Quad has chosen to let the copper side of the PCB face outwards.
This way the epoxy of the PCB material acts as the isolator towards the diaphragm and the other stator.
Because the copper on the other side is not isolated, arcing can still occur and there is nothing to limit the current.

You could consider having the copper sides face the diaphragm and use the solder mask as isolation.
It would be possible to make the holes in the copper slightly larger than the holes in the epoxy. This allows the copper to be entirely covered by the soldering mask, including the edges.

Drilling the holes by hand is an impossible job.
The only feasible way is to have them drilled automatically by the PCB manufacturer.

A big advantage of PCB material is that it is easy to split area of a panel into different segments of any desired shape, and to use some of these for the tweeter and the others for the squacker.
Another advantage is that the parasitic capacitance can be practically zero by removing the copper at the edges where the spacers are.

Stator of Stretched Wires
It is possible to build stators as a grid of stretched electrical wires.
Stranded wire is easier to work with than solid wire.
Wire diameters from 1mm to 3mm are suitable.
By using isolated wire, the isolation of the stator is easily realized.

The openness of a wire stator can be high, up to 60% if desired.
Further, many ESL builders believe that the air flows more easily around the round wires than it does around the sharp edges of perforated sheet material.

A disadvantage of wire stators is that it is difficult to realize segmentation, other than tall vertical segments along the entire height of the loudspeaker.
ESL builders using wire stators therefore often take the line source approach.

Another disadvantage of wire stators is that it is difficult to attach a damping cloth to the rear stator.

Eddy Fikier [4, Fikier] describes in detail how to construct wire stators.

Stator Mass
Newton's third law (action equal minus reaction, also known as the impulse conservation law) says that the Coulomb force working on the diaphragm also works on the stators – but in opposite direction.
To keep the stators from moving, we must rely on their mass.

There are, however, two reasons why this is not as critical as you might think.

In the first place, the force on the diaphragm of an ESL is smaller than the force on the cone of an electrodynamic loudspeaker. That is because in the latter a big part of the force is needed to accelerate the heavy cone and speech coil.
For this reason, while the cabinet a cone loudspeaker working at large volume vibrates noticeable, the frame of an ESL hardly vibrates.

In the second place, if the stator vibrates a little, this has no consequences, at least if the stators are completely transparent. Because the transparency is not perfect, there will be an effect after all, but it will be small.

16.3 Spacers and Other Isolation Materials

For loudspeakers using perforated metal sheets as stators you need spacers made from some good isolator. For loudspeakers using stretched wire stators, you also need isolating spacers, unless the entire loudspeaker frame is made of some isolating material.

First of all: materials that are considered reasonable isolators at low voltages are not good enough at the high voltages we have in an ESL.
The polarizing voltage, for example, could easily be 5000V and we find a leakage current loading the high voltage generator of 0.1µA not negligible.
We must therefore have an isolation resistance in excess of 50GΩ.

This rules out the use of wood, plywood, hardboard, carton, paper etc.
We must use good isolators such as most plastics and pertinax.
Suitable plastics are polycarbonate (brand name plexiglas) or poly-acrylate (brand name lexan). PVC works well too.

Pertinax is a good isolator that was used extensively in the manufacturing of radio and television sets before plastics became available. Pertinax is made from layers of paper, pressed together with a phenol-formaldehyde resin to form sheets or plates.

Fiberglass printed circuit board works well too, but it is difficult to cut and the dust is a health hazard.

In the Netherland, sheets and plates of plastic material or pertinax are available from RU58 in the Hague.
Poly-acrylate plates can be obtained at a shop for do-it-yourself home improvement materials.

If the surface of the isolators is greasy, current will leak over the surface and this will not only load the source of polarizing voltage but it will also make a sizzling noise.
To avoid these unwanted phenomena we must clean the surface of the spacers with alcohol using a clean cloth. Acetone might also work, but some plastics will dissolve in it. In either case: be aware of the fire hazard.

A good way of avoiding leakage currents is the following.
When you apply the conductive coating to the diaphragm, keep 1cm away from the edge of the panel. That way the polarizing voltage does not reach the spacers.

If you cut stripes of isolation material to be used as spacers, be sure to remove the sharp edges with some abrasive paper to prevent them from cutting the diaphragm.

16.4 Diaphragm Material

The generally accepted diaphragm material is polyester.
The chemical name is polyethylene terephtalate (PET).
It is the same material that is used for the fabrication of soda bottles.

Polyester film is fabricated by a number of manufacturers:

Polyester Film	
Brand Name	**Manufacturer**
Mylar	Dupont Tejjin
Melinex	Dupont Tejjin
Hostaphan	Mitsubishi Polyester Film

The film is made in thicknesses of 4 μm, 6 μm, 12 μm and 25 μm.

In order to make the diaphragm mass negligible, the thickness should ideally be 3 - 4 μm.
A thickness of 6 μm is also suitable but will result in a 6 dB/oct roll-off above approximately 15 kHz. We can compensate this roll off by an equalizing filter between the pre-amplifier and the main amplifier. Because this makes the main amplifier work harder it is not a preferred solution.
Diaphragm thicknesses of 12 μm or 25 μm are only suitable for bass panels or for loudspeakers that are voltage driven.

Polyester sheet comes in a great variety, each having different properties.
Some are designed to heat shrink, which allows us to use a heat gun to stretch the diaphragm
Some have an etched surface to make them printable.
Some are tensilized, meaning that during manufacturing they are stretched to increase their strength or their Young modulus.
Stretching is often done only in the machine direction, this is the direction in which the production machine produces the film.

It is not easy to obtain small amounts of polyester film in the right thickness.
The 4 μm and 6 μm varieties are rarely used and are therefore not much produced.
If you can find it you are usually obliged to buy a roll of 2000 m.

Sometimes amateur ESL builders have managed to acquire a stock of the stuff and sell it in small quantities for a reasonable price.
In the Netherlands, Martin-Jan Dijkstra has for a long time supplied 4 μm and 6 μm Mylar to his fellow ESL builders, but I am unsure if he continues to do that till this day.
You will be able to find Martin-Jan through [13,ESL club].
Another possible source is [15, Twinstatic].

If you have obtained a film but you are unsure of the thickness, you can cut 100 pieces of, say, 3 cm × 3 cm and stack these to get a total thickness that you can measure with a sliding ruler.

16.5 Diaphragm Coating

Let's turn to the issue of giving the diaphragm the desired conductivity.
According to Peter Baxandall [2, Baxandall, page 114] the diaphragm should have a sheet resistance between 10^9 and 10^{11} Ω.
There are several ways to achieve this.

Nylon Paint
The Quad electrostatic loudspeakers are said to have a thin paint of nylon on their diaphragm. Nylon itself is an isolator but it attracts enough moisture from the atmosphere to get just the right sheet resistance.
There are anecdotes that these loudspeakers stop working under extremely dry conditions.
The use of nylon paint is not an option for amateur ESL builders because the solvents needed to create nylon paint are highly toxic.

Graphite
Another method to make the diaphragm conductive is the use of graphite powder.
This method is used commercially by Roger Sanders in the USA [12, Sanders].

Here is a sketch of how it works:
You tape the polyester film to a flat, smooth and clean surface (glass would be ideal) and make sure it is stretched a bit.
Next you clean it with alcohol or acetone.
Then you sprinkle graphite powder over it and rub that in hard with a cotton cloth.
You can obtain graphite powder from the hardware store where it is sold for the lubrication of locks. You need a fine grain quality to avoid that the large and sharp grains cut the film.
After rubbing, you vacuum clean away most of the excess graphite and remove the rest with alcohol.

You can test the conductivity by placing two coins on the surface and measure the resistance between them. I do not know what distance to use between the coins or what resistance you should obtain, but the internet will provide that information if needed.
If the resistance is still too high, you can apply another treatment with the graphite powder.

It is highly undesired to have loose graphite particles roaming around in the loudspeaker. It is therefore wise to carry out the above described procedure in another room than the one in which you assemble the loudspeaker.

More details on how to make the diaphragm conductive with graphite appear in
[4, Fikier], [12, Sanders] and [5, Rehorst].

Conductive Paint
Several people have invented some kind of paint that has the appearance of a water-thin liquid, which can be sprayed or applied with a brush.
The paint dries up to a conductive layer with the right sheet resistance.
The polyester film must be de-greased using alcohol or acetone before the paint is applied.
Again, Martin-Jan Dijkstra from the Netherlands has supplied fellow ESL builders with a paint like this of his own secret recipe. I am unsure if he continues to do that to this day.
Other sources of conductive paint can be found on the internet.

Soap
Some people have reported good results from the use of solutions based on soap.
A disadvantage is that these coatings tend to be sticky and therefore attract dust.

Anti-static Spray
You can experiment with anti-static spray that is sold in the stores in the household detergent department for the anti-static treatment of furniture.
In all likelihood the coating will not last for years, but it might be good enough for experimentation. You probably have to think of a way to prevent the spray from ending up on the surface of your spacers and other isolating parts.
Mark Rehorst [5, Rehorst] claims that he has found an anti-static spray that really works.

Wall Paper Glue
An unlikely solution that nevertheless works well – at least for experimentation – is the use of wall paper glue.
I mean the kind of glue that you prepare from a dry powder, which you sprinkle in water.
You prepare a solution that is a bit more diluted than it needs to be for wall papering and apply two or three layers to the stretched diaphragm.
Of course you clean the diaphragm with alcohol or acetone before you start.
The coating you get by doing this gets its conductivity by attracting moisture from the atmosphere.
A more detailed description appears in [4, Fikier].

There are two disadvantages to the wall paper glue coating.

First, I do not know how much mass it adds to the diaphragm.
You can find this out by taking 1m^2 of polyester film and applying the coating to it.
If you weigh it before and after you know the weight of the coating.

Secondly, the coating made of wall paper glue has a poor adhesion to the polyester film.
Fortunately it is not required to withstand large tearing forces but nonetheless I doubt that it will remain intact over a period years.
For your first experiments, however, it works surprisingly well.

Metalized Film
You may be able to obtain metalized polyester film.
This is used in window dressing and to treat window glass so that it keeps the heat out.

There are two reasons why this film is not preferred.
First, the conductivity is much too high.
Second, this kind of film has usually a thickness of 25 μm or more.

In spite of this, it may be good enough for your first experiments.

16.6 Connection to the Diaphragm

To make a connection to the coated diaphragm, many people use adhesive metal tape.
Both aluminum tape and copper tape are used.
Aluminum tape is available from the do-it-yourself home improvement shop, in the plumbing department.
It is wise to make several connections around the perimeter of the loudspeaker.

Many builders have reported that the connection to the diaphragm does not last long because the high voltage eats away either the coating or the metal of the adhesive tape.
Some have reported an improvement when using a negative polarizing voltage.

Personally I have always used paper tape – the masking tape used for painting.
At the voltages we are dealing with it is conductive, just like wall paper glue.
I lay a very thin copper wire on the side of the diaphragm and stick the tape over it.
Then, when I apply the coating (the paint from Martin-Jan Dijkstra) I paint several layers over the paper tape. I have never had any problems of loosing connectivity to the diaphragm.

16.7 Diaphragm Stretching

How much stretching is needed
In order to keep the diaphragm in its place, in the middle of the air gap between the stators, it must be stretched.
The amount of stretching influences the diaphragm stability and the resonance frequency. How much stretching you need depends on the size of the air gap, the width of the panel and the polarizing voltage. It also depends on the thickness of the polyester film.
Typically the film must be stretched by 1% to 2%. If you stretch it further, you no longer stretch it elastically but plastically. That is: if you release the tension it does not return to its original size.

You can monitor the amount of stretching by drawing a rectangle on the film, using a permanent marker, before you start stretching it.
When the rectangle has grown 1 – 2% in size, you have the right amount of stretching.

You will be surprised how much force you need to properly stretch the diaphragm.

Adhesive Tape
A method of stretching the diaphragm is by using a flat and smooth surface (glass would be ideal) and adhesive tape.
The idea is to lay the film on the surface and to tape it to that surface along the edges. Next, you place new pieces of tape at opposite sides and you apply them under a bit of tension. You work your way around the perimeter until you have reached the desired amount of stretching. Fikier [4, Fikier] describes this method in more detail.

I must say that I have never been able to get this method to work because the adhesive force between the tape and the flat surface was not sufficient.

Stretching Table

Many people build a stretching table, consisting of a flat and smooth surface with on each side a lath which can move. The two laths can be moved outward by turning some screws or bolts. The polyester film is attached to the laths by adhesive tape and is then stretched by turning the screws.
A stretching table like this can only stretch the mylar in one direction.

Figure 84 shows a sketch of such a table.
The laths on the left and right sides can be pulled apart by means of bolts (not shown).
I built this table from 18mm plastic coated fiber board.

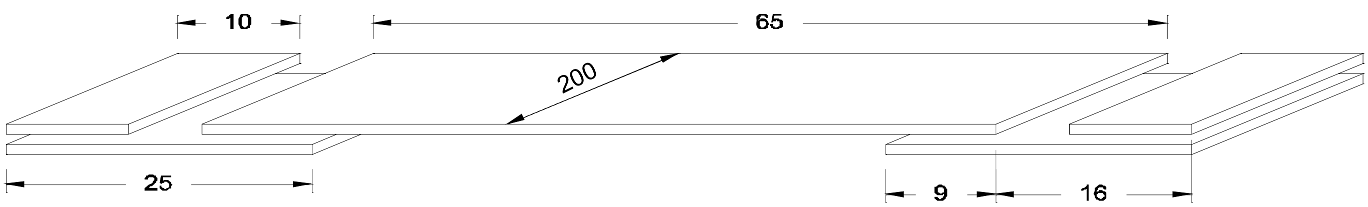

Figure 84 Stretching table.

Pneumatic Stretching Table

Mark Rehorst [5, Rehorst] has described a stretching table that uses a pneumatic tube (a bicycle tire) to do the stretching.
Figure 85 show the basic idea.
You place the diaphragm on top of the table and wrap it around to the bottom where you fix it with adhesive tape. Next you inflate the tube.
You can make the table round or rectangular.
Several people have reported using it with good results.

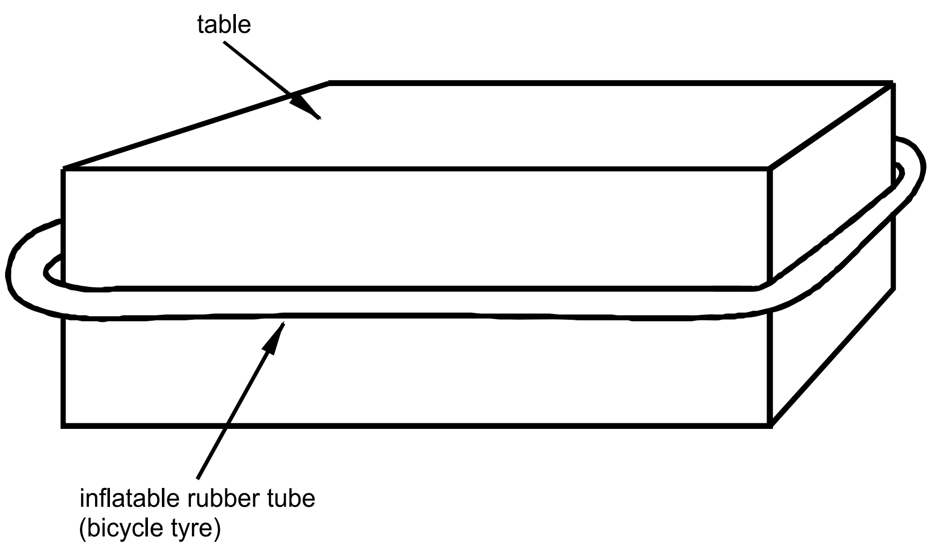

Figure 85 Pneumatic stretching table.

Heat Gun

Many people use heat shrinkable polyester film and stretch it using a heat gun.
You must be careful not to burn a hole in your diaphragm, as this does not only mean the loss of a stretch of film but also of the entire driver.
I have never used heat shrinking for my actual drivers, just for a small experimental set-up to investigate the diaphragm resonance.
I do not know how long the heat shrunken film will keep its tension.

16.8 Adhesives

You need glue to bond the diaphragm to the spacers and to bond the spacers to the stators.

Glue to bond to spacers made from pertinax, PVC, polycarbonate, lexan and PCB material is easy to find. In most cases 2-component epoxy glue will do the job.
This is available in the quantities you need from RU58 in the Hague or from Polyservice in Nieuwerkerk aan den IJssel (both in the Netherlands).
2-component poly-urethane glue is also suitable and is available from the shop for home improvement materials.

A bigger problem is to find an adhesive that will work on the polyester film.
I have seen many types of glue recommended, among which solvent based glue like bison kit and bison tix (probably only known in the Netherlands), 2-component epoxy resin, 2-component poly-urethane glue and even double sided adhesive tape.
I have myself tried all of these and some more.
In my experience non of these have a good adhesion to the polyester.

The solvent based glue simply does not stick. Besides, it does not cure unless at least one of the materials you try to bond is porous.

The 2-component epoxy hardly sticks. It does not require to be exposed to air to cure.

The 2-component poly-urethane is the best one that I have found.
The bond is not really strong, but at least there is a bond.
It does not require exposure to air to cure.
I used the type that is sold by Bison, in a blue box. It is available in do-it-yourself home improvement stores.
It is a rather thick paste and that makes it difficult to apply a thin and even layer.
The use of a glue-comb makes this task a little easier.

Silicone glue or silicone sealant is the only stuff that I could find that will form a really strong bond with the polyester. It is, however, almost impossible to apply a thin and even layer and it will not cure unless one of the materials is porous.

Mark Rehorst [5, Rehorst] recommends Scotchgrip 4693 or 4693H from 3M.
According to him it is readily available in the USA, it is cheap and it works like a charm.
Unfortunately I have not been able to find a supplier in the Netherlands.
Something that might be similar and might be available in the Netherlands is Scotchgrip 4475.

Some ESL builders use adhesive tape.
Some even use double sided adhesive tape to kill two birds with one stone: the tape serves at the same time as adhesive and as spacer.
My own experience is that adhesive tape will not hold the diaphragm tension very long.

16.9 Damping

For the damping of the diaphragm resonance, you need to attach a thin but tightly woven cloth to the rear stator.
You will have to experiment with different types of cloth to get the right amount of damping.
The best way to attach it is to spray the back of the rear stator with fabric glue and then apply the cloth.
Do not spray the cloth instead as this will render it impermeable.

Before you permanently attach the cloth, you may want to use some fixture that allows you to experiment with different types of cloth in such a way that you can easily remove it.

16.10 Cross Over Filter

The cross over filter described in chapter 12 consists of resistors and capacitors.
The resistors must be able to withstand the high audio voltages.
If they don't they may go up in fire or they may cause severe distortion due to dielectric breakdown or corona.
If you cannot find resistors with the right voltage rating, you can always connect several of them in series.

The capacitors must be able to withstand the voltage too.
If you use ceramic types, be aware that some types of ceramic dielectrics (especially the high epsilon types) are very non-linear and will therefore cause distortion.
You have to make sure that the capacitors are made from the linear type of ceramic dielectric.
Polyester film capacitors are less likely to produce non-linear distortion.
Preferably use capacitors with 5% tolerance.

Most ESL builders do not have capacitors in their cross-over filter.
Instead they connect their ESL panels as a ladder network using only resistors, as discussed in section 13.7, Figure 57.
Here the capacitance of the panels does the filtering.

For the construction of the cross-over filter you could design a dedicated printed circuit board. It works equally well, however, if you glue the resistors and the capacitors on a board of pertinax or PCB material using 2-component epoxy glue.
You can interconnect the components using rigid copper wire.

If you build the cross-over filter on a board in this way, you will have long wires to the panels. These add some parasitic capacitance, and this is especially undesired for the tweeter panel, which has a small capacitance of its own.
You can avoid this by placing the last resistors in the ladder network close to the panel, perhaps even solder them directly to the stator.
If you use a capacitor-less cross over filter of the ladder type you can mount the resistors directly between the panels to minimize the wiring capacitance.

16.11 Wiring

Some attention must be paid to the wiring of your loudspeaker.

Wire Capacitance

First of all, you must avoid having long wires from the step-up transformer to the cross-over filter as the step-up transformer will not like the extra capacitive load that they represent.

The capacitance of the wires leading to the panels must also be minimized as discussed in section 16.10.

Corona

Most PVC isolated wire is not rated for the high audio voltages from the step-up transformer. Even though catastrophic breakdown is not likely to occur, you can easily get leakage currents through the isolation, especially where two wires touch.
This leakage current is accompanied by an ionization of the air, which is called *corona* or *glow discharge*.
The discharge gives off a very faint blue light in the dark. This light is so weak that you may not be able to see it even in a totally darkened room.

The corona discharge is highly non-linear and this causes a nasty distortion.

You can avoid corona by letting all wires run in a plastic tube having a diameter of a few mm.
Such tube is available from stores for garden supplies (brand name: Gardena), where they are sold for the irrigation of the flower beds.

It is even better to have no air between the wire and the isolation.
You can achieve this by using a high quality coaxial RF cable of the type used to feed the antenna of high power RF transmitters.

16.12 Polarizing Voltage

For the polarizing voltage we need a voltage source that produces 2000 – 6000 V.

Output Current Capability

The voltage source does not have to be able to produce much current and indeed it would only be dangerous if it could. The load current is no more than the leakage current of the ESL and that should be well below 1 μA.

Mosquito Racket

A fast and cheap way to get a high voltage source is to buy one of those mosquito killers in the shape of a tennis racket.
You will need some equipment to measure the output voltage as this is not specified.

Cockcroft-Walton Generator

A suitable high voltage source is the Cockcroft-Walton generator.
It was developed by John Cockcroft and Ernest Walton to power one of the first particle accelerators back in the 20th century.
Figure 86 shows the circuit diagram. The diagram shows a version with only 7 sections, but for the polarizing voltage of an ESL you normally need more sections.

With diodes 1n4006 and capacitors with 630V rating, the circuit can handle an input voltage of $230V_{rms}$, which corresponds to a peak value of 325V.
The dc output voltage is equal to this peak input voltage multiplied by the number of sections.
Thus with 16 sections you will get an unloaded output voltage of 5200V.

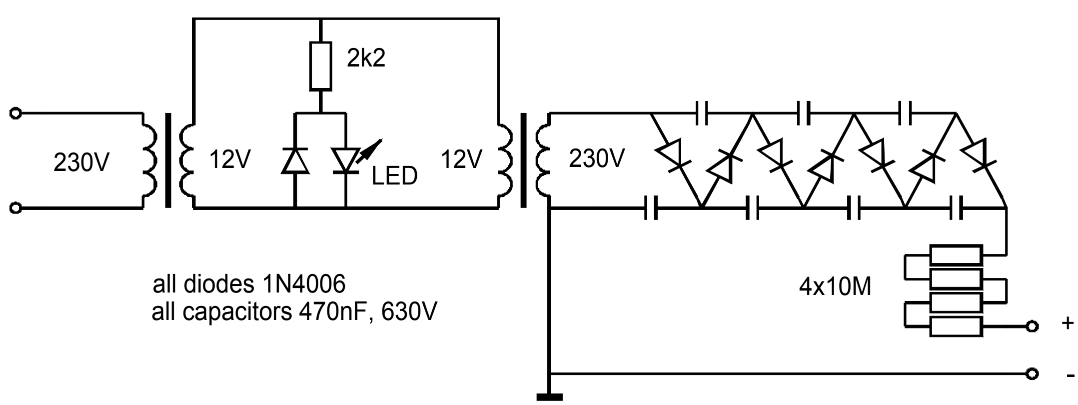

Figure 86 **Cockcroft-Walton high voltage generator.**

The output resistance of the generator is high and therefore it cannot supply large currents. That is a good thing because otherwise such a high voltage would be extremely dangerous.

The output impedance increases rapidly with the number of sections and if you go higher than 16 sections you may well find that the voltage does not increase as much as you had expected because the leakage currents load the generator too much.

A resistor of 40 MΩ in series with the output of the generator makes it less painful if you accidentally touch the output.
At the same time it prevents the destruction of all diodes in the generator in case you accidentally short circuit the output. That would be painful in another way.

For maximum safety, the generator should be isolated from the mains supply by a transformer that delivers the desired 230 V_{rms}. For this purpose you can also use a back-to-back connection of two small print transformers designed to produce for example 12 V from the mains supply voltage, see Figure 86.
You will find that the losses in these small transformers are so large that the output voltage is considerable less than 230 V.

Some ESL builders do not bother to use a transformer and drive the generator directly from the 230 V mains voltage. That has of course the disadvantage that the entire stereo set including the amplifier and the CD player will be live.
The danger can be mitigated by connecting a resistor of 1 MΩ in series with the “cold” output terminal.

If you decide to leave out the galvanic isolation, mount the entire circuit, including the two output resistors, in a plastic box.
If you live in a country having 110 or 127V mains voltage, then driving the generator directly from the AC outlet is of course not an option.

The voltage on the capacitors will be 2.828 times the rms input voltage, or 622 V.
The peak reverse voltage on the diodes will have the same value.
For the capacitors you can use 630 V types, although this leaves little headroom to deal with the case where the mains voltage is at its positive tolerance.
The 1N4006 diodes can handle a reverse voltage of 1000 V, which is more than adequate.

Fikier [4, Fikier] advices a value for the capacitors of 33 – 47 nF.
Such a low value, combined with the 50 Hz or 60 Hz frequency of the mains ac voltage leads to a very large output impedance and poor output current capability.
It also produces a large output ripple voltage due to the leakage current.
I always use 470 nF.

For the construction of the generator you can of course design a dedicated PCB.
It works equally well, however, to glue the capacitors upside down on a board of pertinax or PCB material and mount the diodes free hanging between them.

Output Ripple
The effect of output ripple is seriously under-exposed. That is: nobody talks about it.
It is important to have as little 50 or 60 Hz ripple on the high voltage as possible, as this will modulate the sensitivity of the loudspeaker and thereby blur the sound.
It is hard to give a specific number for the acceptable ripple but I would aim for an rms voltage of no more than one thousandth of the dc voltage. That means 2 – 5 V_{rms}.
That way the unwanted frequency components caused by the parasitic modulation are 66 dB or lower than the desired frequency components.
If you use a Cockcroft-Walton generator the ripple is smaller if you use larger capacitors.
Keeping the leakage current of the ESL down helps too.

If you use a battery powered mosquito racket, there will not be any 50 or 60 Hz ripple.
The high frequency ripple generated by the electronic generator is easily filtered by the diaphragm resistance and the stator capacitance.

16.13 Step-up Transformer

The important aspects of the step-up transformer have been discussed in chapter 15.
Step-up transformers are commercially available from Amplimo [7, Amplimo], Plitron [8, Plitron] and Sowter [9, Sowter].

You can also use output transformers for balanced tube amplifiers, connected in the reverse direction. You will, however, not be able to get more voltage out in the bass range than twice the supply voltage of the tube amplifier for which the transformer was designed.
Above that voltage the core will saturate.
Higher voltages are possible at higher frequencies.

You can also look for mains transformers, transforming 230 V to for example 12 V.
Also here you use the transformer in the reverse direction.
Note that at low frequencies the output voltage is limited by the saturation of the core to 230 V_{rms}.
Higher voltages are possible at higher frequencies.
You can double the output voltage by using two transformers together as discussed in section 15.12.
These transformers are of course not optimized for high frequency operation.
Because of these limitations, these transformers are only suited for your first experiments.

17 Safety

Because of the high voltages involved, safety is a point of concern when working with electrostatic loudspeakers.
Let's take a look at the high voltage sources.

Polarizing Voltage

The polarizing voltage is not as dangerous as you might think.
This is due to the limited current that the voltage generator can deliver.
You can, however, get a painful electric shock from it because the total capacitance in the circuit can discharge at once and provide a current spike which is much larger than the steady state current.
If you use a high resistivity diaphragm and place a 40 MΩ resistor in series with the output of the high voltage generator, the only remaining capacitance is that of the interconnecting wire. That is too small to generate a serious electric shock.

A high voltage generator lacking proper galvanic isolation from the mains power supply is rather dangerous as the entire audio chain can carry live voltages.
The danger of this set-up can be reduced by introducing two large valued resistors in both output wires of the generator as discussed in section 16.12.

Step-up Transformer

The output voltage of the step-up transformer is rather dangerous as the voltage is high enough to overcome the resistance of the human skin and it can deliver quite some current.

In most cases, current drive conditions for the loudspeaker will be created by using large series resistances.
These resistances limit the current to at most a few mA.
You can, therefore, get a painful electric shock by touching the front and back stators at the same time, but it will not be really dangerous.
Besides, you have to be pretty determined to touch both stators at the same time.
Finally, the stators will be isolated.
For safety, it would be best to mount the first resistors as close to the transformer as possible and place the transformer and the resistors in a plastic box.

If you decide to build a transformerless high voltage amplifier the situation becomes really dangerous. The power supply of such an amplifier will be several thousand volts and it can easily deliver enough current to kill you.

18 Literature

[1] New Developments in Electrostatic Loudspeakers.
P.J. Walker.
Journal of the Audio Engineering Society, 1980 November, Vol. 28, Number 11, pages 795 – 799.
Can be downloaded from [3].
Discusses among other things Walker's equation and give a sketch of its proof.

[2] Electrostatic Loudspeaker by Peter J. Baxandall.
One long chapter, devoted to ESL's, in Loudspeaker and Headphone Handbook, edited by John Borwick, ISBN no. 0-240-51371-1 or 0-240-51578-1.
The chapter can be downloaded from [3].
Extensive treatment of all theoretical aspects including Walker's equation.
Baxandall does not avoid the math.

[3] Website Frank Verwaal.
http://home.kpn.nl/verwa255 (mostly Dutch language).
Dutch text covering ESL theory and construction in 22 separate chapters.
Does not avoid the math.
Most of this was written without knowing of [2, Baxandall]
but fortunately everything agreed with [2].
Also the book you are now reading can be downloaded here.
Further [1], [2] and [4] can be downloaded.
Finally, the simulation programs *ESL.pas*, *ESL.m* and *ESL_ladder.m* can be downloaded here.

[4] Elektrostatische Luidsprekers, Zelf Bouwen of Zelf Kopen.
E. Fikier.
Uitgeversmaatschappij Elektuur BV, 1993.
Can be downloaded from [3].
The book in the Dutch language.
Lots of construction details.

[5] Website Mark Rehorst.
http://esl.rehorst.com; http://mark.rehorst.com/ESLs/index.html.
Description of construction and which materials to use. No theory. Easy to read.

[6] Website Menno van der Veen.
http://www.mennovanderveen.nl/nl/index.html (Dutch).
You can find here a discussion of the Amplimo/Plitron transformers of his design.

[7] Website Amplimo.
http://www.amplimo.nl/.
Unfortunately the information about the ESL step-up transformers is no longer there.

[8] Website Plitron.
http://www.plitron.com/.
Step-up transformers designed by Menno van der Veen [6].

[9] Website Sowter.
http://www.sowter.co.uk/.
British manufacturer of ESL step-up transformers.

[10] A Geometric Approach to Eliminating Lobing Error in Multiway Loudspeakers.
Joseph A. d' Appolito.
74th convention of Audio Engineering Soc., 1983 October 8-12. New York.

[11] Acoustics.
Leo L. Beranek.
McGraw-Hill Book Company Inc, 1954.
Standard work on acoustics. Leans quite a bit on the math.
Look here (or in [3, Verwaal]) to understand the behavior of point sources.

[12] The Electrostatic Loudspeaker Design Cookbook
Roger R. Sanders and Barry McClune.
ISBN 1-882580-00-1.
Easy to read. Lots of construction details.

[13] Website Dutch ESL club .
http://www.esl.hifi.nl/index.htm.

[14] Forum of Dutch ESL club
http://groups.yahoo.com/group/ESL-club.

[15] Twinstatic Audio
http://www.twinstaticaudio.nl/

19 Appendix I. Point Sources and Volume Velocity

The reason you might at some point become interested in what is meant by *volume velocity* is that it is the quantity that defines the strength of a point source.

Volume Velocity
Consider as an example a cone loudspeaker mounted in a box.
The cone vibrates and thereby has a time dependent velocity.
If we multiply that velocity by the area of the cone, we get the (time dependent) *volume velocity*.
As we can see, the volume velocity defines how much air the cone moves per unit of time.
We could also call this the air flow [m^3/s].
The volume velocity, as defined here, is time dependent.
In most cases, when people talk of volume velocity, they mean the rms value of this time dependent quantity.

Point Source
A point source is, contrary to what the name says, a source of a finite size.
It is a sphere with a small but finite radius, and that radius oscillates.
The rms volume velocity is the area of the sphere, multiplied by the rms surface velocity.
It defines the *strength* of the source.

Due to the symmetry, a point source is an omni-directional radiator. It radiates sounds equally in all directions.

A cone loudspeaker mounted in a closed or semi-closed box behaves to a good approximation as a point source for all frequencies where the size of the box is mall compared to the wave length.
Also the mouth of a person speaking or singing can be approximated as a point source.

See further [11, Beranek].

20 Appendix II. Walker's Equation and its Derivation

Walker's Equation
Walker's equation is

$$p_{rms} = \frac{V_p}{d} \frac{i_{rms}}{r} \frac{1}{2\pi c} \text{ [Pa = N/m}^2\text{]}$$

with

p_{rms}	the rms sound pressure [Pa]
V_p	the polarizing voltage at rest [V]
d	the distance between the diaphragm and a stator [m]
i_{rms}	the rms drive current [A]
r	the distance to the loudspeaker [m]
c	the speed of sound [343 m/s]
π	≈ 3.1416

The Derivation of Walker's Equation
We will not derive Walker's equation, but present a sketch of the derivation.

The derivation rests upon *reciprocity*.
Roughly speaking, a system is reciprocal if it has the same transfer (or gain, if you like) if we swap the input and the output.

> Sound propagation is reciprocal.
> Suppose that John is locked inside his car and due to some malfunction he cannot open the doors.
> The windows are closed too.
> George is standing outside the car and is trying to help John.
> They talk to each other by shouting through the closed window.
>
> Reciprocity now says that they hear each other equally loud.

To apply this to an ESL, we replace the listener (which is supposed to be on the axis of the loudspeaker) by a point source and we use the ESL as a capacitor microphone that listens to the sound of the point source.
We can calculate the output voltage of that microphone as it appears at its open terminals.
Reciprocity then says that the relation between the strength of the point source and the microphone output voltage is equal to the relation between a drive current of the ESL and the resulting sound pressure at the listening position.

It is not difficult to calculate the output voltage of the ESL when it listens to the point source.
The waves from the point source create at the location of the ESL a particle velocity.
If we know the strength of the point source, the magnitude of that particle velocity is well known [11, Beranek], [3, Verwaal].
Let’s assume that the diaphragm of the ESL is massless and that its movements are not restrained by its suspension. The diaphragm then just follows the velocity of the air.

For this to work, we must require that the size of the ESL and the distance to the point source are such that we can approximate the sphere shaped sound waves from the point source as being flat at the location of the loudspeaker. See Figure 87.
This requires that the loudspeaker is sufficiently small and/or the distance to the point source is sufficiently large (far field condition).

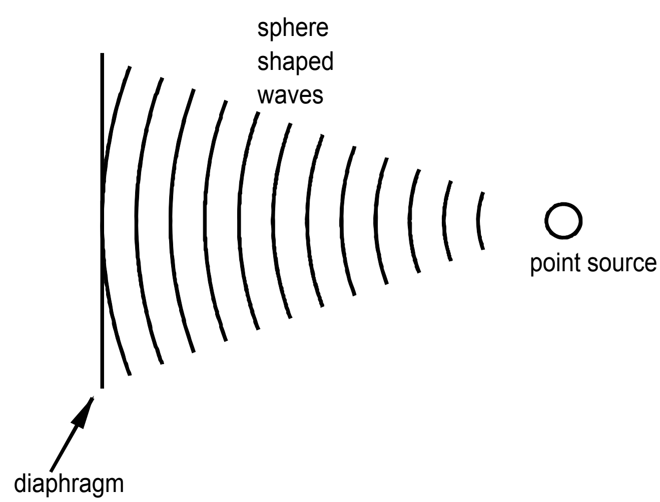

Figure 87 **ESL used as a microphone, listening to the sphere shaped wave from a point source.**

With that approximation the diaphragm moves in a uniform manner.
Under these conditions it is easy to calculate from the diaphragm excursions the output voltage that appears on the stator plates.

Note that apart from the restriction of the far field condition the size and the shape of the diaphragm do not matter.

Finally note that any influence of the acoustical short circuit has been implicitly taken into account.

The actual derivation of Walker’s equation appears in [1, Walker], [2, Baxandall] and [3, Verwaal].

21 Appendix III. The Cut-off Frequency of the Current Drive.

Consider an ESL panel with capacitance C_o, driven by a voltage source (for example a step-up transformer) through a total series resistance R_s (which will for reasons of symmetry be split in two resistors of $R_s/2$ each). See Figure 88.

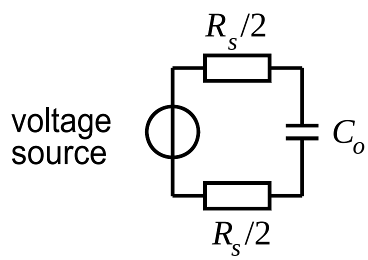

Figure 88 ESL panel, driven from a voltage source with series resistance.

The frequency below which the drive current cannot be maintained and starts to drop at a rate of 6 dB/oct is

$$f_c = \frac{1}{2\pi R_s C_o}$$

with $\pi \approx$ 3.1416.

22 Appendix IV. The Capacitance of a Flat Plate Capacitor.

Consider an ESL with area A [m^2].
Let the distance between the plates be $D = 2d$ with d [m] the distance between the diaphragm and a stator.
If we neglect the holes in the stator, the capacitance of the ESL is given by

$$\boxed{C_o = \varepsilon_o \frac{A}{D} \text{ [F]}}$$

with the constant ε_o given by

$\varepsilon_o = 8.85 \times 10^{-12}$ F/m.

23 Appendix V. Generalization of Walker's Equation to Off-axis Listening Positions.

We want to apply Walker's relation to a listening position that is off the axis.
Let the angle that the listening position makes with the axis be φ.

Following again the procedure we used in appendix II, we replace the listener by a point source and determine the voltage at the loudspeaker terminals as it listens to the sound of the source.

We still assume that the loudspeaker is small enough that the path length differences are negligible. This is equivalent to the requirement that the sound waves of the source make the diaphragm move in a uniform way.
Note that with the listening position off the axis this requires a much smaller diaphragm. We can see this from Figure 89.
Also note that increasing the distance now does not make the path length differences go away.

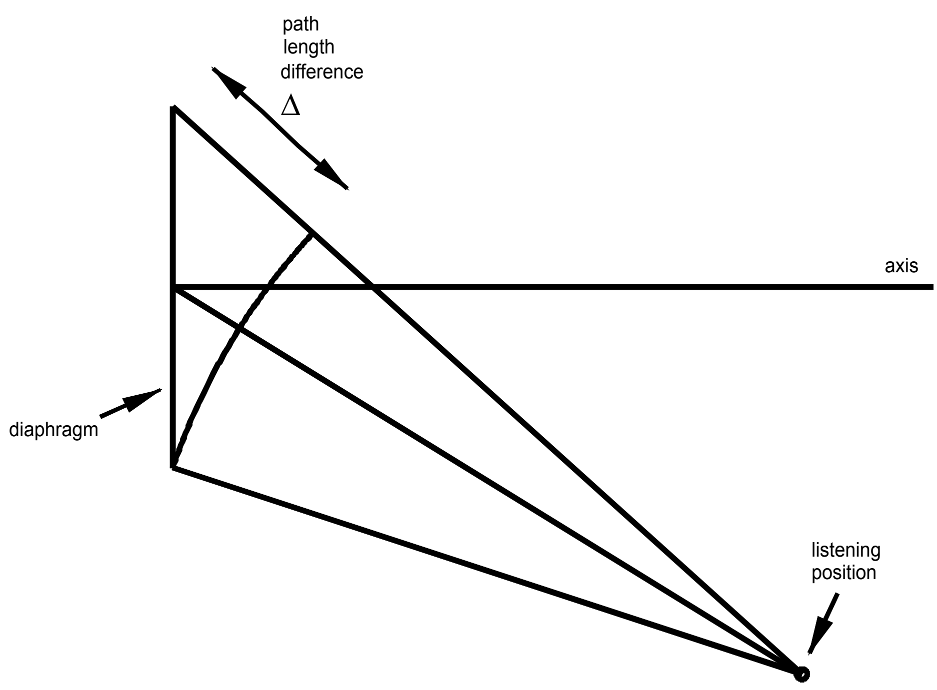

Figure 89 **Path length differences at an off-axis listening position.**

Walker's relation can now be found in a similar way as before.

There is one difference: the particle velocity of the air is now no longer perpendicular to the diaphragm. The diaphragm, however, can only move in the perpendicular direction.

We must therefore decompose the particle velocity in a perpendicular component and a component that is parallel to the plane of the loudspeaker.

The perpendicular component of the particle velocity is smaller than the particle velocity itself by a factor $\cos(\varphi)$

By reciprocity, also the sound pressure in the off axis listening position is reduced by the same factor.

When the angle of the listening position becomes 90° (so the listener is located by the side of the loudspeaker, in the plane of the diaphragm) the cosine becomes zero and therefore also the sound pressure is zero.
This result makes sense also because of the symmetry of the situation.
The particle velocity at that location, however, is not zero.

If you try this with an ESL in a typical listening room, you will still hear a sound.
That sound is indirect sound, reflected off the walls.

Conclusion:
The generalized Walker equation for small loudspeakers and off-axis listening positions is

$$\boxed{p_{rms} = \frac{V_p}{d}\frac{i_{rms}}{r}\frac{1}{2\pi c}\cos(\varphi)}$$

with

p_{rms}	the rms sound pressure [Pa]
V_p	the polarizing voltage at rest [V]
d	the distance between the diaphragm and a stator [m]
i_{rms}	the rms drive current [A]
r	the distance to the loudspeaker [m]
c	the speed of sound [343 m/s]
π	≈ 3.1416

24 Appendix VI Superposition and Linearity

Linearity
There is much confusion about what defines a linear system.
Some people believe that it involves an input/output relation that, when plotted on a piece of paper, produces a straight line through the origin.
Others believe that the straight line does not even have to pass through the origin.
Yet others say that a system is linear if it produces at its output no other frequencies than those that were offered at its input.
Neither of these definitions is the correct one, although some linear systems indeed do have these properties.

The correct definition is:

A system is linear if (and only if) it obeys the superposition principle.

Superposition
The superposition theorem applies to all linear systems.
The superposition principle can be split into two properties, called *homogeneity* and *additivity*.

Homogeneity
Consider a system with an input and an output.
The system is called *homogeneous* if multiplying its input signal by a certain number has the result that also the output signal gets multiplied by that same number[76].
Simply put: a two times larger drive causes a two times larger response.

Additivity
Consider a system with an input and an output.
Consider further two different drive signals that can be applied to the input.
Each generates at the output its own response.
The system is called additive if driving it with the sum of the two drive signals produces at the output the sum of the two individual responses.

If a system is *both* homogeneous and additive, it is said to be linear.

An ESL and the air through which the sound propagates together form a linear system because all the equations describing their behavior are (to a very good approximation) linear.

---

[76] As a special case: multiplying the input signal by zero also makes the output signal zero.

25 Appendix VII The Proximity Effect

The proximity effect is a result of the properties of a point source.
Peter Walker found the response of an ESL by using it as a microphone that listens to the sound (more specifically: the particle velocity) of a point source. See appendix II.

As it turns out, a point source has at close distance a velocity that has too much of the low tones in it. [11, Beranek], [3, Verwaal].
Then, by reciprocity, the ESL must have the same frequency response.

Why does a point source behave in such a way?
In a flat sound wave there is a fixed ratio between the sound pressure and the particle velocity. The sound waves around a point source are not flat but sphere shaped and the curvature of the spherical waves increases towards the source.
This results in a deviation of this ratio, and the deviation is stronger for lower frequencies and at shorter range.
As a result, the ratio between pressure and velocity changes for low frequencies, and we get more velocity in the bass range than we expected.

The reason that things are different when the waves are curved lies in the interaction between pressure and velocity. We have seen before that velocity leads to pressure and that pressure leads to velocity. This interaction creates a wave that propagates through the air. As it turns out, the curvature changes the way that velocity generates pressure. The precise mechanism that causes it to change is outside the scope of this book.

26 Appendix VIII Generalization of Walker's Equation

26.1 Introduction

In this appendix we want to show how we can generalize Walker's equation such that it can be used to predict the sound pressure at any point in space.
That includes off-axis positions and positions close to the diaphragm (down to one millimeter or less, if you desire).
This also removes the limitation on the size of the loudspeaker.

This appendix involves some mathematical calculations.
If you don't like these, you can go directly to the result in section 26.8.

26.2 Outline of the Calculation

We start with the original equation

$$p_{rms} = \frac{V_p}{d} \frac{i_{rms}}{r} \frac{1}{2\pi c} \tag{1}$$

This equation is only valid on the axis and at a large distance and/or a small loudspeaker size.
We circumvent these limitations by splitting up a large loudspeaker in many small loudspeakers, if needed the size of a postal stamp or smaller.
In the limit case the size of these small sub-loudspeakers goes to zero.
We call these small loudspeakers *surface elements*.

Because the surface elements are so small, we can remove the requirement that the listening distance must be large. It must be large compared to the size of the surface elements, which will be vanishingly small.

Using superposition[77] (see appendix VI), the total sound pressure then follows as the sum of the individual sound pressures generated by all surface elements.

In the limit case where the size of the surface elements becomes vanishingly small, the summation changes into an integration.

The following sections will follow the path outlined here to arrive a the general Walker equation.

---

[77] The superposition principle can be applied hiere because all equations describing the ESL and the sound propagation in air are linear[a].

[a] At least at sound pressures that are small compared to the atmospheric pressure. Even an SPL of 135dB satisfies that criterion, as it is 1000× smaller.

26.3 Complex Transfer Function

Complex Voltages, Currents and Sound Pressures: Steady State Analysis
First, we must recognize that the contributions from all surface elements have different phases, because they all have a different distance to the listening position.
As a result, when calculating the individual contributions, it is not enough to calculate their rms values or their peak values: we must also keep track of their phases.

A mathematically convenient way to do this is to switch to complex signal quantities. That is: the voltage and the current become complex, and so does the sound pressure.

Consider as an example a real valued voltage $V(t) = v_a \cos(\omega t + \varphi)$ with amplitude v_a and phase φ.
We can express this voltage as the real part of a complex signal

$$V(t) = v_a \cos(\omega t + \varphi) = \mathrm{Re}\left(v_a e^{j(\omega t + \varphi)}\right) \tag{2}$$

which we can reformulate as

$$V(t) = \mathrm{Re}\left(v_a e^{j\varphi} e^{j\omega t}\right) \tag{3}$$

Here j is the imaginary unit, with the property $j^2 = -1$.

We see that all the information about amplitude and phase is contained in the complex amplitude $v_a e^{j\varphi} = v_a \angle \varphi$.
Assuming that we know the frequency that we are talking about, this complex amplitude tells us all we want to know about $V(t)$.
For this reason we represent $V(t)$ by its complex valued amplitude $\boldsymbol{V} = v_a e^{j\varphi}$.
Notation: we use bold face to denote complex quantities.

If we want to find $V(t)$ back from its complex amplitude, acording to (2) we only need to add the time dependent factor $e^{j\omega t}$ and take the real part.

For reasons outside the scope of this book, this method of calculating is known under the name *steady state analysis.*

We have used bold face to distinguish the complex amplitude $\boldsymbol{V}$ from the original time signal $V(t)$ that it describes.
In general, when we do steady state analysis, we can omit the bold face without causing any confusion. This is because it is clear from the context that we are dealing with complex valued constants, not real valued time signals.

When performing steady-state analysis, we therefore use the symbols V, I and p to denote the complex voltage, current and sound pressure.

Complex Impedance

When we divide the dc voltage across a resistor by the current through that resistor, we get the resistance value:

$$\frac{V_{dc}}{I_{dc}} = R \tag{4}$$

The same is true for non-dc voltages and currents.
If, for example, we divide the rms voltage by the rms current, we again get the resistance:

$$\frac{V_{rms}}{I_{rms}} = R \tag{5}$$

If we have a two-terminal circuit containing resistors, capacitors and inductors, we can do the same with the complex valued voltage and current. The ratio is in that case a complex number, which we will call the *impedance*:

$$\frac{\boldsymbol{V}}{\boldsymbol{I}} = \boldsymbol{Z} \tag{6}$$

The impedance is therefore a kind of generalized complex resistance.
We have here maintained the notation in bold face to distinguish between the complex quantities and the real ones.

Complex Gain or Transfer Function

If we have an amplifier with a gain G, we can find that gain by dividing the rms output voltage by the rms input voltage:

$$\frac{V_{o,rms}}{V_{i,rms}} = G \tag{7}$$

This gain is a real number.
If the output voltage does not have the same phase as the input voltage, that information is lost in the gain G.
In order to fix this, we can define a *complex gain* as the ratio of the complex output voltage and the complex input voltage:

$$\frac{\boldsymbol{V}_{\boldsymbol{o},}}{\boldsymbol{V}_{\boldsymbol{i}}} = \boldsymbol{G} \tag{8}$$

were the quantities in (8) have now become complex valued.
The phase angle of that complex ratio represents the phase shift of the amplifier.

This complex gain is also known by the name *transfer function.*

Note that the transfer function is a complex constant as long as we confine ourselves to just one frequency.
When we start varying the frequency, we notice that the transfer function depends on frequency:

$$\frac{\boldsymbol{V_o}}{\boldsymbol{V_i}} = \boldsymbol{G}(f) \text{ or } \frac{\boldsymbol{V_o}}{\boldsymbol{V_i}} = \boldsymbol{G}(\omega) \tag{9}$$

where f [Hz] is the frequency and $\omega = 2\pi f$ [rad/s] is the angular frequency.

With the tools we developed in this section, we can now rewrite Walker's equation in terms of complex drive current $\boldsymbol{i}$ and complex sound pressure $\boldsymbol{p}$:

$$\boldsymbol{p} = \frac{V_p}{d} \frac{\boldsymbol{i}}{r} \frac{1}{2\pi c} \tag{10}$$

The factor

$$\frac{\boldsymbol{p}}{\boldsymbol{i}} = \frac{V_p}{d} \frac{1}{r} \frac{1}{2\pi c} \tag{11}$$

is now a *transfer function* or complex gain from drive current to sound pressure.

We see that the right hand side of (11), even though it is now *allowed* to become complex, has not yet done so. All the quantities in the numerator and denominator are real voltages, distances and speeds.
This is about to change in the following sections.

26.4 Wave Number

For convenieence, we will from now on drop the notation with bold face to denote complex quantities.

Expressions (10) and (11) are not entirely correct because the do not take into account that at a distance r from the loudspeaker, the sound will arrive at the listening position with a time delay r/c with c = 343 m/s the speed of sound.

This results in a phase shift $\varphi = -\omega r / c$.

We must therefore add to (10) a phase factor $e^{-j\frac{\omega}{c}r}$:

$$p = \frac{V_p}{d}\frac{i}{r}\frac{1}{2\pi c}e^{-j\frac{\omega}{c}r} \tag{12}$$

The ratio ω / c is known as the *wave number* and we denote it by the symbol k :

$$k = \frac{\omega}{c} \tag{13}$$

The wave number has the dimension of phase per distance and the unit rad/m.
Its physical meaning is as follows: it describes how fast the phase of a wave changes with distance.

We now arrive at

$$p = \frac{V_p}{d}\frac{i}{r}\frac{1}{2\pi c}e^{-jkr} \tag{14}$$

This is an intermediate version of Walker’s equation, which is still only valid on the axis and in the far field, but it now shows how the phase changes with distance.

26.5 Proximity Effect

If the size of the loudspeaker is small enough, (14) would remain valid at very small distances. When the distance becomes so small that the proximity effect kicks in (see section 9.5) we must modify (14) to take that effect into account.
For that purpose we must add a factor

$$\left(1+\frac{1}{jkr}\right) \tag{15}$$

which increases when the product kr becomes small.
From (13) we see that a small k occurs when the frequency is low.
Therefore, (15) starts to differ from unity at low frequencies and short distances.

With this modification, Walker's equation becomes

$$p=\frac{V_p}{d}\frac{i}{r}\frac{1}{2\pi c}e^{-jkr}\left(1+\frac{1}{jkr}\right) \tag{16}$$

This is an intermediate version of Walker's equation that shows us the phase delay due to the distance and takes the proximity effect into account.
It is still only valid on the axis and in the far field.

26.6 Off Axis Response

As we have seen in section 9.4 and appendix V, Walker's equation is still valid for off-axis listening positions if we add a factor $\cos(\varphi)$, where φ is the angle that the listening position makes with the loudspeaker axis.
We then get

$$p=\frac{V_p}{d}\frac{i}{r}\frac{1}{2\pi c}e^{-jkr}\left(1+\frac{1}{jkr}\right)\cos(\varphi) \tag{17}$$

This intermediate version of Walker's equation is valid off axis, provided that the loudspeaker is so small that path length differences are negligible.

26.7 The Sound Pressure of One Surface Element

We will now divide our large loudspeaker into many small surface elements, each of them small enough to let (17) be valid.
We assume that one surface element takes a portion of the total drive current

$$di = i\frac{dA}{A} \tag{18}$$

where A is the area of the loudspeaker and dA is the area of the surface element.[78]

Walker's equation for this single surface element then becomes

$$dp = \frac{V_p}{d}\frac{di}{r}\frac{1}{2\pi c}e^{-jkr}\left(1+\frac{1}{jkr}\right)\cos(\varphi) \tag{19}$$

which we can with (18) rewrite as

$$dp = \frac{i}{A}\frac{V_p}{d}\frac{1}{r}\frac{1}{2\pi c}e^{-jkr}\left(1+\frac{1}{jkr}\right)\cos(\varphi)\,dA \tag{20}$$

The small sound pressure dp is the contribution of surface element dA.

26.8 Summation and Integration

Now it only remains to sum all contributions (20) from all surface elements.
That way we find the total sound pressure

$$p = \sum_{all\ dA}\frac{i}{A}\frac{V_p}{d}\frac{1}{r}\frac{1}{2\pi c}e^{-jkr}\left(1+\frac{1}{jkr}\right)\cos(\varphi)\,dA \tag{21}$$

When we let $dA \to 0$ the summation changes into an integral

$$\boxed{p = \int_A \frac{i}{A}\frac{V_p}{d}\frac{1}{r}\frac{1}{2\pi c}e^{-jkr}\left(1+\frac{1}{jkr}\right)\cos(\varphi)\,dA} \tag{21}$$

This is then finally the general Walker equation.

Unfortunately the integral is not easy to solve analytically.
The Matlab script *ESL.m* solves it numerically by evaluating (21).

[78] Note that this assumes that the total drive current i gets evenly distributed over all surface elements. This is an approximation, but a very accurate one. See section 9.6.

27 Appendix IX The Electrostatic Force on the Diaphragm and SPL at Close Range

Consider an ESL with the following parameters

A	[m^2]	diaphragm area
d	[m]	distance between diaphragm and stator
V_p	[V]	polarizing voltage

The capacitance between the diaphragm and the two stator plates is

$$C_{diaph} = 2\varepsilon_o \frac{A}{d} \tag{1}$$

with $\varepsilon_o \approx 8.85 \times 10^{-12}$ F/m, the electrical permittivity of vacuum.
We have here neglected the reduction of the capacitance due to the holes in the stator.

The polarizing charge is then

$$Q_p = C_{diaph} V_p = 2\varepsilon_o \frac{A}{d} V_p \tag{2}$$

Now let's apply a signal voltage v_s between the stator plates.
This generates an electric field between the plates

$$E_s = \frac{v_s}{2d} \tag{3}$$

This electric field pulls on the polarizing charge with a force

$$F = E_s Q_p \tag{4}$$

which, by using (2) and (3) changes into

$$F = 2\varepsilon_o \frac{A}{d} V_p \frac{v_s}{2d} \tag{5}$$

This simplifies to

$$\boxed{F = \varepsilon_o \frac{A}{d^2} V_p v_s} \tag{6}$$

According to Newton's third law this force must be in equilibrium with the air pressure (sound pressure) on both sides of the diaphragm.
Due to the symmetry, the air pressures on both sides are equal in magnitude but opposite in sign.
Therefore, we find the sound pressure (which is force per unit of area) by dividing (6) by the area A and by 2.
We then find

$$\boxed{p = \varepsilon_o \frac{V_p}{2d^2} v_s \qquad (7)}$$

This is the sound pressure in the direct vicinity of the diaphragm.

This expression has direct relevance for the design of electrostatic headphones.

Another important use of this equation is that the immediate vicinity of the diaphragm is a special case of the generalized Walker equation, derived in appendix VIII.
We can therefore verify if the generalized Walker equation agrees with (7) for this special case (and it does!).
This is a good sanity check.

28 Appendix X Transformer formulas

The primary magnetizing inductance of a transformer is

$$L_p = N_p^2 \mu_o \mu_r \frac{A}{l + \mu_r d} \qquad (1)$$

where

N_p		is the number of primary turns
$\mu_o = 4\pi \times 10^{-7}$	[H/m]	is the magnetic permeability of vacuum
μ_r		is the relative magnetic permeability of the core material
A	[m^2]	is the crossectional area of the magnetic circuit
l	[m]	is the length of the magnetic circuit
d	[m]	is the length of the air gap.

The time integral of the voltage that the transformer can handle before saturation sets in is given by

$$\left(\int V\, dt\right)_{\max} = N_p A B_{sat} \qquad (2)$$

where
B_{sat} [Vs/m^2 = T] is the saturation flux density.

The number of volt-seconds in a half sine wave of amplitude V and frequency f is

$$\int_0^{\pi/\omega} V_{peak} \sin(\omega t)\, dt = \frac{2}{\omega} V_{peak} = \frac{1}{\pi f} V_{peak} \qquad (3)$$

As the flux density in the core is allowed to swing from $-B_{sat}$ to B_{sat}, the saturation limit will be reached when (3) is equal to two times (2).
It follows that the maximum voltage before saturation sets in is equal to

$$V_{sat,peak} = 2\pi f N_p A B_{sat} \qquad (4)$$