

Introduction to Dynamic Loudspeaker Design

March 4, 2014

A loudspeaker represents a way of converting electrical signals to sound signals. All speaker do this by having the electrical signal exert some sort of force on a "diaphragm", a relatively large, more or less flat piece of material which is made to vibrate by the force applied to it. There are a number of ways of having the electrical signal exert forces on the diaphragm, but by far the most common is the "dynamic" loudspeaker, in which an electrical current flowing through a coil of wire which is immersed within a magnetic field is used.

It was already discovered by Ampere about 200 years ago, that if you had a magnetic field and a wire carried a current through that field, then that wire would feel a force on it proportional to the current in the wire, and proportional to the strength of the magnetic field. That force was in a direction perpendicular to the wire, and also perpendicular to the direction of the magnetic field.

Thus, in a dynamic loudspeaker, a wire is wrapped around a cylinder many times (to increase the force since each piece of wire will feel the same force since each carries the same current.) That cylinder is then immersed in a strong magnetic field (Good loudspeakers use some of the strongest permanent magnets available). The amplifier then, by raising and lowering the voltage, drives more or less current through the coils of the loudspeaker, producing more or less force on the diaphragm.

1 Efficiency

As we learned in the notes on efficiency of radiation, for any vibrating body, like a loudspeaker, the efficiency with which the motion of the air right next to the body is converted into sound depends strongly on frequency and on the size of the body. Below the so called knee frequency, (the frequency at which the wavelength is approximately twice the diameter), the efficiency of that conversion gets lower and lower, the further the frequency is from the knee frequency– this fall off away from the knee frequency is at 6dB per octave (a factor of 4 (6dB) for each halving of the frequency (octave)).

Above the knee frequency, the efficiency is unity, but the sound tends to get more and more beamed straight ahead from the speaker.

Except in cheap speakers, manufacturers tend never to use a speaker in the range of frequencies above the knee frequency to minimize this beaming effect. People usually want the sound to be the same whether sitting straight ahead of the speaker or more off to one side.

However, this leaves the real problem of the falling of the efficiency with frequency below the knee frequency. To compensate for this, the manufacturers use the natural resonance of the cone of the speaker. Above the resonant frequency, the response of the cone of the speaker to the force gets smaller and smaller the higher the frequency. In fact, the velocity of the cone falls at 6dB per octave as you get further away from the resonant frequency. We can see how we can put this resonant response together with the efficiency to get a flat response for the speaker. If you look at frequencies which lie between the resonant frequency of the speaker cone, and the knee frequency for that speaker cone, then as one goes up an octave, the response of the cone drops by 6dB due to the distance from the resonance frequency. However, the efficiency with which that motion of the cone goes up by 6dB because of the greater efficiency as you get nearer the knee frequency. Thus the drop of 6dB due to resonance is exactly compensated by the 6dB rise due to increased efficiency, making the actual sound produced by a given electrical energy fed in the same at those frequencies. This means that for signals between the resonance and the knee, the amount of sound output for a given energy input is constant. Above the knee frequency, the falling of the response due to the resonance is no longer compensated by the efficiency since the efficiency is one for all frequencies higher than the knee. Thus the sound output for a given energy input falls at frequencies higher than the knee. Similarly, for frequencies lower than the resonance, the motion of the speaker to a given input energy again falls at 6dB per octave away from the resonance. Similarly the efficiency also falls at 6dB per octave, implying that in this range the sound falls at 12 dB per octave. Ie, below the resonance frequency of the speaker, the amount of sound produced for a given energy gets smaller and smaller the lower the frequency. Ie, a dynamic loudspeaker has poor response both above the knee frequency and below the resonant frequency.

Note that the efficiency of the speaker in converting electrical energy tends to be very low. Throughout the frequency range of usefulness, the speaker is either very inefficient in converting the motion of the air at the diaphragm to sound, or is moving very little because one is so far from the resonant frequency of the diaphragm. Typical good speakers tend to have efficiencies of much less than 1%. (ie over 99% of the electrical energy fed into the speaker is wasted as heat). At about a meter from a good loudspeaker, the loudness of the sound tends to be around 90dB for each watt of energy fed into the speaker. 90dB at a meter is about 10^{-3} watts of sound produced for each watt of electrical energy fed in.

In figure 1 I have presented the analysis of one speaker. We note that another key feature is that the Q of the speaker cone itself must be carefully designed. If the Q is too large, then the speaker will have a heightened response at just

that resonant frequency. If sounds near that frequency are produced, the sound will be boomy. The optimal Q is about .5, ie the speaker cone must be very well damped. Fortunately, because of the high motion of the speaker at this frequency, the speaker itself tends to produce a voltage in the coils which opposes that of the amplifier itself, lessening the current that the amplifier can deliver at that frequency. Ie, the amplifier itself acts as form of damping for the cone, meaning that the mechanical damping does not have to be quite as accurately designed as might otherwise be the case. This amplifier damping works only if the magnet in the speaker is very powerful however- ie is of help only for expensive loudspeakers.

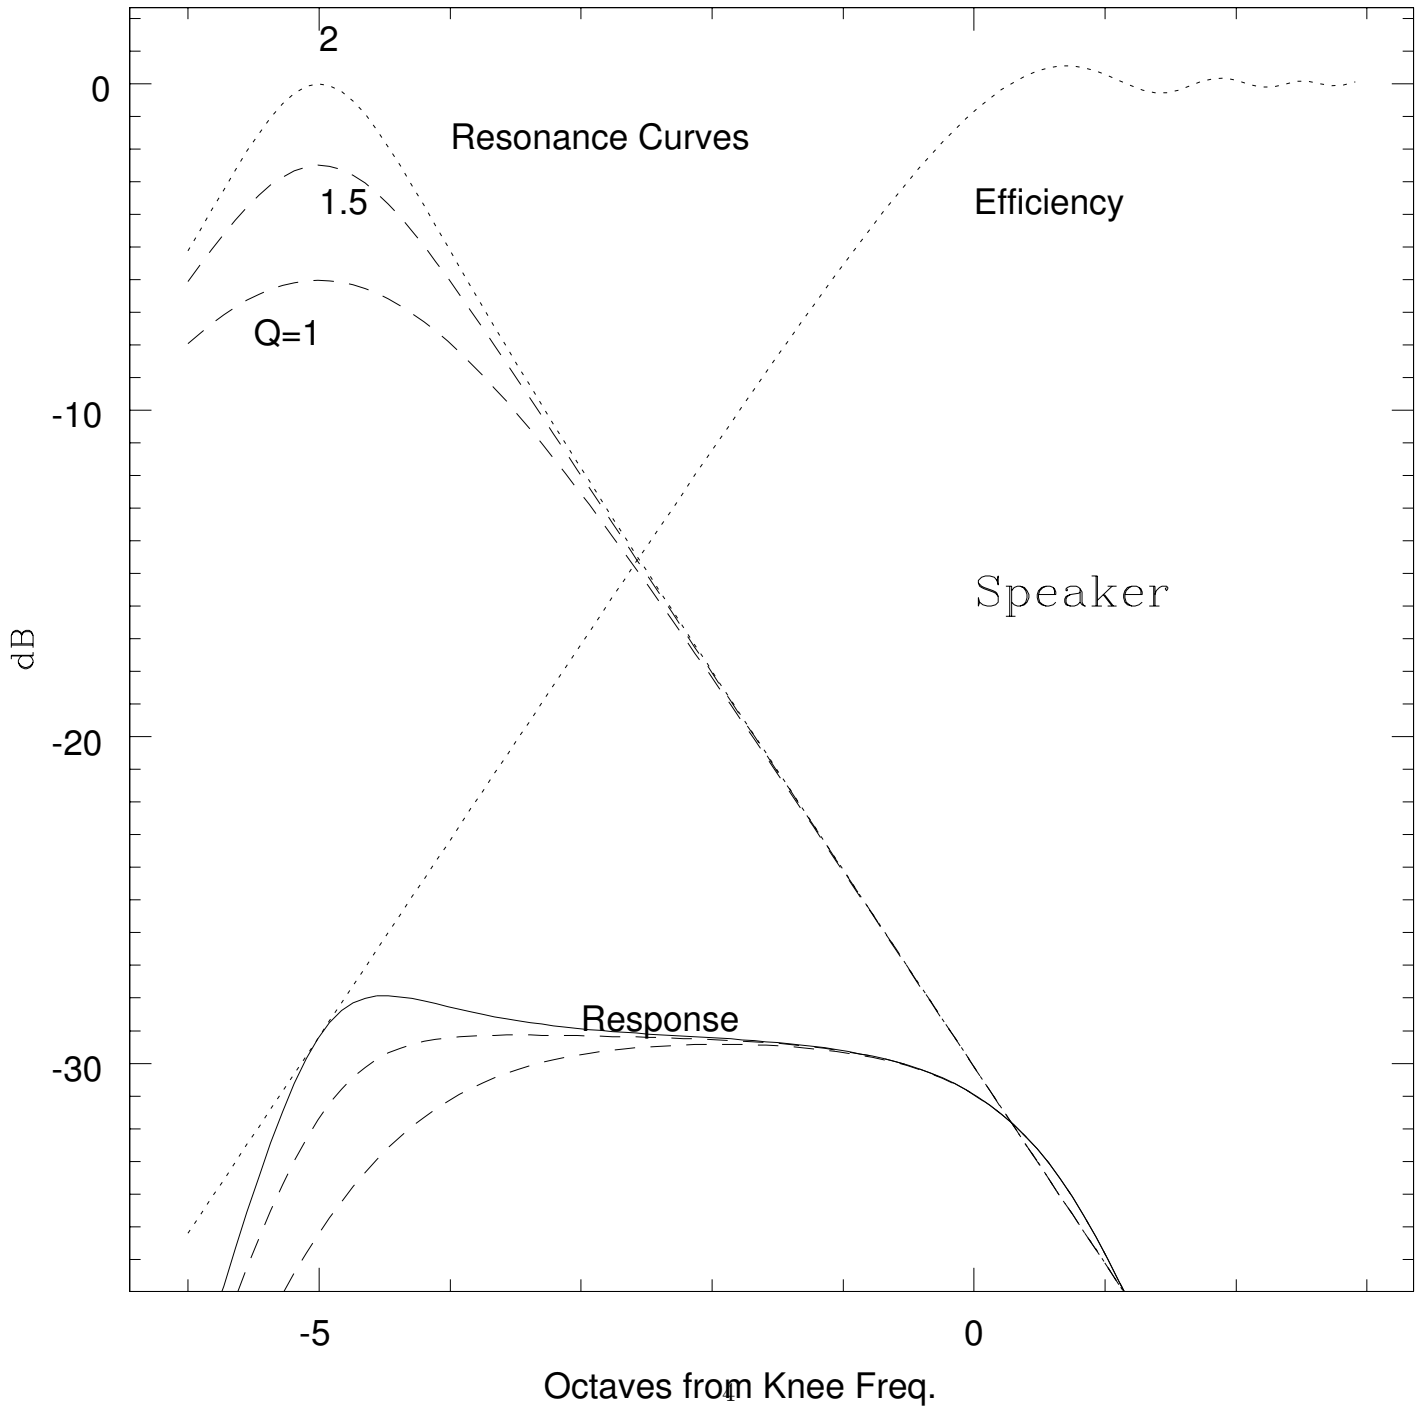


Figure1– Theoretical low frequency Speaker Response curves

Small speakers tend to high high resonant frequencies (small mass) and thus are poor at reproducing low frequencies. Good speakers for bass notes tend to be very large, to have a low resonant frequency. They also have a low knee frequency.

One way of trying to raise the knee frequency in a larger speaker is to try to design the cone so that only the inner part of the cone vibrates at higher frequencies. since the vibrating part is smaller, the knee frequency goes up. It is however extremely difficult to do this without introducing extra resonances into the cone – resonances like those of a drum head. These extra resonances will suddenly cause parts of the cone to move a lot for a given electrical input– ie will produce peaks in the sound output at those specific frequencies. This tactic thus tends to be used only on very cheap speakers. (You can often tell by seeing an extra little cone attached to the inside of the bigger cone near its center.)

Instead what is done is to use different sized speakers for the different frequency ranges. One may for example use a very large speaker, with a resonant frequency of say 50Hz, and a knee frequency of say 300Hz, to reproduce only the vary lowest bass notes. One will then use a much smaller speaker – say 10 cm across, to produce the frequencies between the 300 Hz, to say 2KHz.(This is called the midrange). One then adds a tiny (one or at most two centimeters across) to handle the highest frequencies from say that 2kHz to 20kHz (this is called the tweeter).

In other cases, one may have one speaker from say 100Hz to 1kHz, and the second from 1kHz to 10kHz. This makes the manufacture of these speakers more difficult if they are to handle the broader range of frequencies.

If one is to do this, one must make sure that the two speakers are well matched, so that their overall efficiency in converting electrical power to sound is the same. (One does not want speakers in this the high frequencies are produced much more loudly than the low, or vice versa).

In some of the better loudspeaker combinations, one will try to make sure that the speakers are used only over frequency ranges well within the range between the resonance and the knee, where the response is most uniform. Ie, one will direct the electrical energy at the speaker only for those frequencies where one knows that the speaker can handle it well.

What can destroy the best laid designs of a speaker manufacturer?– the brief answer is other extraneous resonances. While the main resonance of the cone is useful, as described above, other resonances are not. For example it is very easy for the walls of the cabinet in which the speaker is enclosed to have resonant frequencies at which they vibrate a lot when driven by the sound pressures coming off the back of the speaker. Those vibrations of the cabinet will either add or subtract from the sound produced by the main speaker cone, depending on the relative phase of sound produced by that resonance vibration of the wall as compared with that of the main speaker. Ie, the walls of the

speaker should be designed to vibrate as little as possible (ie made of a stiff and heavy material), and any resonances should be as strongly damped as possible. For example in a set of speakers I made, the inside walls are covered in about 1 inch of roofing asphalt shingles embedded in tar and braced with steel bracing, both to stiffen the walls, add mass to the walls (both of which make it hard for any pressures inside the speaker to actually move the walls) and to damp them (because of the viscosity of the tar.) Furthermore, the inside is stuffed full of fibreglass (actually wool is probably better, so stuffing sheepskin rugs in would probably be better except for the moth factory this might make) to absorb the higher frequency sounds emitted by the backs of the speakers.

There is one speaker design in which resonances can be useful. The lowest frequency that the speaker itself will produce efficiently is at the resonance of the speaker cone. By introducing something which has a resonant frequency somewhat below this one can increase the bass response below the that lowest resonance of the speaker itself. Alternatively, one can use a main speaker with a higher cone resonance, and still get the same bass response. The main techniques used to get this extra resonance are either to add a second non-driven speaker with a lower cone resonance (eg by making it more massive), or more popularly, by designing the speaker to have a "coke bottle" resonance at that lower frequency. By putting a hole in the speaker cabinet, and by designing the size and length of the hole to make sure that that resonance has a frequency which is just slightly (less than an octave) below the cone resonance of the speaker, one can slightly (less than an octave) extend the low frequency response of the system. The cost is that for even lower frequencies (well below the resonance of the cone), one now has the sound from the back of the speaker through the hole cancelling that from the front, meaning that the speaker response at really low frequencies falls off even faster.

Also, if that extra resonance is not extremely well designed (ie unless it has just the right Q), it can have a response which is too high, and gives the bass a very tuned and boomy sound (ie certain notes are produced very efficiently but other notes only a few tones away are produced much more poorly). Ie, it is more difficult to design a good bass reflex (as such speakers which use that "coke bottle" resonance to increase the bass) than an "acoustic suspension" (plain speaker in a box) speaker. However, since the temptation to use the bass reflex design to allow a more extended bass while using a cheaper, or smaller speaker is great- ie bass reflex tends to get used as a cost saving measure- such speakers tend to be worse in their bass response, as they tend not to be carefully designed and thus tend to produce boomy basses. (Some customers aid in this by hearing that boomy sound and thinking it means that the speaker has a good bass response. It is only with repeated and careful listening that they realise that the speaker indeed does produce an exaggerated response, but only over a very narrow range of frequencies.)

There is a way of increasing the efficiency of a speaker, and that is by not allowin the air just near the speaker to slosh around- ie directing the motion

straight ahead. If however one simply put a straight pipe onto the speaker, one would get exactly the same efficiency problems at the opening of the pipe to the air, and would have the additional problems caused by the resonances in the pipe itself. What one needs to do is to gradually flare out the pipe so that the exit to the room is larger (and thus has a lower knee frequency, and thus a higher efficiency) than the speaker. This is the philosophy behind horn loudspeakers, or "transmission tower" speakers. The problem is that at places where the horn flares out, the signal tends to get reflected. Thus if the rate of flare (ie the distance over which the expands in size by a factor of 2.7 in area size) is shorter than 1/6 of the wavelength of the sound, the horn just tends to reflect back the sound, and not allow it to escape. Since at 50Hz, the wavelength is 6m, few people could tolerate a horn many meters long in their living rooms. (The folded corner Klipsch Horns are an example where the corner of the room itself is used as part of the horn, and the horn itself is folded up to fit into a box of about a bit under a cubic meter in size.) Thus this option is not very popular except in very large halls. However, at higher frequencies, (above 200 or 300 Hz) this can be a popular way of making quite high efficiency speakers. (horn tweeter I recently bought has an efficiency in converting electrical energy to sound of about 25%, rather than the 0.1% of the typical closed box speaker).

Figures 2 and 3 of this document are some typical curves of the response of a particular loudspeaker (in this case an electrostatic 2 meter x 1meter x 10cm midrange/tweeter speaker as measured in an anechoic chamber (a very large room- 10's of meters in size in all directions in which the walls have been covered with over a meter deep sound absorber to ensure that no reflections from the walls can disturb the measurements). This speaker instead of using the force exerted on a current carrying wire in a magnetic field to drive the speaker, instead uses a mylar sheet covered in electrical charge between two very high voltage screens (the volatage on the screens is driven by the amplifier and is in the kilovolt range.) These stretched sheets of mylar have all of the modes one would expect of a drum. and in the speaker response curves one can see evidence of these resonances in the large array of dips and wiggles in the response curve. Because of the very light diaphragm these are very well damped by radiation into the air of the enegy, but they still leave a remnant in the small- 2-3dB - bumps and dips. The features at 10-20kHz are probably due to resonances in the air caught between the screens on which the voltages are applied.

Note that these are actually quite a good set of speakers. The variations in the responses (figure 2) are less than about 3 dB over the whole range of frequencies (except at about 10KHz, which is getting up to the limit of hearing of many people, and is an area where there is not much power in most music or speach anyway).

From : The Audio Amateur
 Issue # 2 1977
 David P Hermeyer. page 4-9

FIG. 12




Fig. 12. Twenty panels measured at 10 feet.

Figure 2- Frequency response curve of home made Electrostatic Loudspeaker

The same article (The Audio Amateur, Issue 2, 1977 David P Hermeyer, pp 4-7) also illustrates the use of the polar diagram to indicate the response of the speaker in various directions at various frequencies. The most important thing to note is the scale indicator- two radial divisions are 10dB. Ie, if the sound is less in a sideways direction by two radial divisions, that sound is less by 10dB (or a factor of 10). One division would be 5 dB, etc. A graph like this can indicate large differences in intensity in the various directions. Thus from the center to the outside edge here corresponds to 10 radial divisions or 50dB. Ie this graph could represent changes in intensity of a factor of 100,000. (Note that the center does NOT here mean zero intensity. The dB scale never gets to zero intensity which would correspond to minus infinity dB).




Figure 3- Polar response of Electrostatic Loudspeaker
Copyright W G Unruh

Fig. 10. Polar response pattern.