Loudspeakers and Not-so-loudspeakers.

J. Newell 2017.

Most people, when choosing a loudspeaker, have great interest in knowing if it will be loud enough for the intended purpose. When a project is budget conscious and must meet mandatory specifications, it is essential that the specifier of the equipment chooses the most appropriate unit for the purpose. Manufacturers often provide this important information on product specification sheets and many products are chosen based on this data. This seems a simple process, yet many systems-engineers have fallen at this hurdle at some point in their career. So why is this simple process so fraught with danger? Well it all comes back to that same old evil: marketing. If we are to be kind we could call them untruths, but it is almost universal that the specifications of how loud a loudspeaker will go are simply not true. In fact, it is often far from true; and occasionally written to deceive.

But surely, it is simply a question of a manufacturer putting a loudspeaker in a measurement-rig and turning it up until it goes no louder, or fails, then recording that result. Well, yes; that would be a relatively certain way of knowing *some* of how it would behave, but it is also not very useful as a measure for what the loudspeaker would do. The loudspeaker may well be sounding abominably horrible at that level, and the output may well not be what we would consider useful. Manufacturers have generally settled on a “cheat” figure because it gives the best-looking number, and as we all know, in selling speakers, “louder is better”.

Unfortunately, manufacturers do know that the specifications they publish are not truthful, or at least in the spirit of truthfulness. Why is this? Well, what manufacturers do is usually take a half-space response and a measured-on-axis sensitivity figure, recorded over some often-arbitrary bandwidth, and apply a simple calculation intentionally ‘mis-assuming’ that this all rises linearly with power-input, all the way to the AES program maximum power rating, before catastrophic damage. And that is it! That is all they usually do. However, as manufacturers know, there are significant factors that make this calculation incorrect, and not just by a few decibels here and there - it can often be as much as 10dB, or more, away from the specified number. This is huge, and certainly enough to cause a system to fail a compliance-test to a point where tweaking will not bring it into compliance.

So where is it going wrong?

Where should I start? Let us start with our old enemy, power-compression. Power compression is a nasty horrible thing, it steals our valuable amplifier power away from us as we go louder. Power compression is a simple function of voice coil temperature increasing with input power, and voice coil resistance increasing with temperature. As the voice coil resistance increases, the coil presents less load to the amplifier (draws less current) and therefore consumes less power. The amplifier then needs to present a proportionally higher voltage to deliver any increase in power, and therefore, output, but this further increases the temperature and resistance and requires even more input to proportionally achieve any more output. The more you push into the driver, the more it will proportionally suffer from power compression. As a direct-radiating loudspeaker often has an efficiency below 10%, what is not converted to acoustic output mostly gets converted to heat in the motor mechanism, or lost in unwanted mechanical output. Power compression is not simply related to power input, different loudspeakers behave very differently depending on their electrical, thermal, and mechanical properties. It is not a constant figure that can be simply applied to all measurements.

If we look back 20 or 30 years we will see that most high-performance loudspeakers are largely the same as modern high-performance loudspeakers in most mechanical aspects, with the exception of one major property, power handling. The biggest advances in recent years has not been in complex, more-linear motor-technology, but in adhesives and materials. Modern loudspeakers have adhesives that can tolerate far higher operating temperatures, and are made from materials that do not burn or warp until they are substantially hotter. Suspension adhesives have advanced greatly, as have the materials that the suspensions are made from. Other proprietary technologies, like the dual rear-suspensions were used by Gauss Loudspeakers in the 1970s, have become open technologies benefiting many products, and to some minor degree, cooling technologies have advanced.

In most aspects however, the big advance that allows us to drive 2 kW into a couple of grams of copper coil is simply the thermal tolerance of the motor materials; not cooler running coils. We have not really built a better transducer, rather we actually have built a better heater. As the overall moving-mass of many of the high-power loudspeakers has increased due to the more robust materials used, often the sensitivity has fallen relative to their older counterparts. We often find ourselves confronted with a 2 kW loudspeaker that is no more capable of any more output than a 200 W loudspeaker of 1987. (The latter being a lighter, more efficient loudspeaker with a nominal sensitivity in the low 100 dBs, and the former being a heavier, more robust loudspeaker with a nominal sensitivity in the mid 90 dBs, with more power compression.

As amplifier power has become cheaper (at least in terms of the cost of the hardware) we have become more cavalier with its use. More worryingly, many modern, high-tech, lightweight amplifiers are often not even capable of maintaining sustained output to the power that they are rated. It is common to see footnotes to specifications, quietly warning that maximum output is time limited to a very short period indeed.

There have been limited ventures into resolving the power-compression issue, and the Community Professional Loudspeakers, AirForce system, was one notable example. It was a true monster of a system, with drive units that were both fluid and force air cooled. It contained drive units with high-power blower units mounted in the rear of the cabinets, run from multi-stage program-adaptive power supplies. I had the luck to work with this system, as well as on its continued system-development for a period of time. It was incredible to hear what a system virtually free of power compression is capable of.

Is this really such a big problem for a short period of operation?

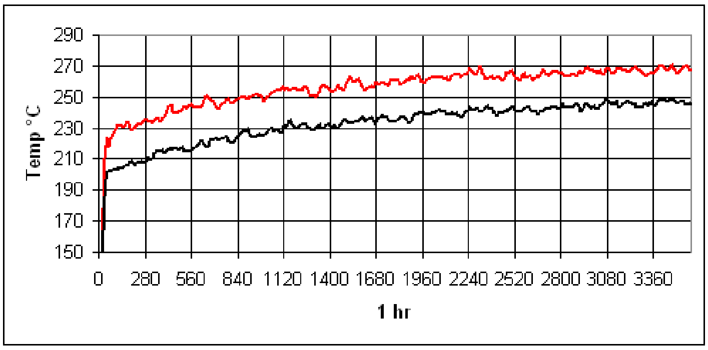

Figure 1. Heating effects with time at an unspecified power (Image 18 sound). The initial rise being voice coil heating, the slow further rise being motor structure heating resulting in less cooling of the coil.

Some question the dynamics of power compression and some claim it takes a while for this to happen, but this is simply not the case. It is exceptionally quick, and a function of some simple laws of physics. It is a matter of a few seconds, certainly not minutes, before any loudspeaker receiving high power levels is well into power compression. How quickly something heats up is a simple function of the power applied, its thermal conductivity, its mass, and its ability to dissipate that heat.

A loudspeaker voice-coil is, by design, a highly conductive material, and being copper or aluminium it will transfer heat very rapidly indeed. Most loudspeaker voice-coils are of only a few grams of mass, so there is little thermal transfer latency in the materials themselves, especially when they are, in their entirety, the source of the heat. We will not see a voice coil take much at all time to saturate with heat. In a high-power loudspeaker of, say, 1 kW, the amount of energy applied is very high indeed, and, considering that the motor efficiency is very low, we see an overwhelming amount of energy driven into a small conductive mass, with only a small surface area to dissipate this heat to somewhere else.

Most loudspeaker voice-coil formers are made from thermally insulating materials, such as fibreglass, Nomex, Kapton, polyimide, or paper, with only aluminium being a thermally conductive material (aluminium, however, being one of the least popular materials due to many of its negative properties). Most coils are wound on one side of these formers, and thus only one side of the coil is able to efficiently dissipate heat. Inside-outside wound coils have twice the dissipation surface.

We also see significant effects within the motor structure itself. Our only hope is that the voice-coil heat is mainly transferred away by the air, but sadly, in a loudspeaker driver, this is not a great help. Most loudspeaker motors do not have a good *flow* of cold air to the outside, most simply pump the same old air back and forth, close to the motor inside a cabinet. The actual air *flow* is therefore commonly very low, and worse still, air itself is not a good conductor of heat, and we need to get a lot of it away from something very small very quickly.

Some loudspeaker manufacturers have turned to ferrofluid for improved thermal conductivity, and in some cases this can improve matters, mainly in lower-powered high-frequency drivers. Ferrofluid is, however, not without its own problems. In high powered loudspeakers, the voice coil temperatures can thermally overload the ferrofluid. In this case the coil exceeds the boiling-point of the ferrofluid carrier liquid, and the liquid can boil off leaving a thick, sticky, crystallised deposit in the coil gap. In tests on an 800 W low-frequency driver, I managed to vaporise the ferrofluid carrier in just over 5 seconds using music-program at high level. On lower powered midrange units, I observed over time that the high-mid response would deteriorate, even though the temperatures did not greatly exceed the thermal limits of the ferrofluid in such a dramatic way that the woofer did. There, we took to replacing the ferrofluid periodically. Ferrofluid is only useable at temperatures of up to 150° C for long periods, yet many loudspeaker voice coils can reach as much as 300° C before failure.

In many cases, the degree of power compression is related to several factors, such as the power applied, the voice-coil size, the voice-coil surface area, the air-flow and the conductivity of the motor components. Depending upon the power of the loudspeaker and the thermal qualities of the motor, we can often see between 3 dB and, in poorly engineered systems, above 6 dB of loss of output related directly to power compression.

If a manufacturer does not factor-in measured power compression losses to their calculated maximum output figures, then they are hiding some truths. All manufacturers know their loudspeakers suffer from power compression. Some manufacturers even put these figures on their data sheets, but others ignore them completely.

Without knowing exactly what a loudspeaker's power compression figures are, we can usually play safe and deduct 6dB from their calculated maximum SPL figures to get closer to the actual maximum output.

This is not the whole problem: there is more at the other end of the scale.

What is the sensitivity of a loudspeaker, how is it measured, and how should it be measured? Again, we are in a fine mess. A loudspeaker simply does not have a sensitivity figure; but there is more to it than this. Loudspeakers have frequency responses and directivity responses. Both affect the output. First, we must know what bandwidth we need the loudspeaker to work in; and then over what area we need it to cover.

Let's look at the first part here. Bandwidth.

If we know what bandwidth we need the loudspeaker to work in, and what type of program we will be reproducing, we can look at a more accurate sensitivity figure. Will we need to equalise the loudspeaker flat? Will we have important program in a part of the bandwidth that is lower output than the headline sensitivity figure? If the answer to those questions is yes, we then cannot believe the spec' sheet sensitivity figures.

The answer to our sensitivity question is usually the lowest wide-Q sensitivity point on the frequency response plot within the range we want to use the speaker. If we are thinking of maximum output capacity and we add equalisation or raise program level over significant bandwidth of the unit, we still only have the same maximum power handling capacity, and we have used some of that up when flattening the response. We will therefore hit maximum power, sooner. Thus, if we add 10 dB at 100 Hz, we reduce the wide-band power-handling capacity by a factor of 10, *and our maximum wide-band output by 10 dB* if our program has required content around 100 Hz.

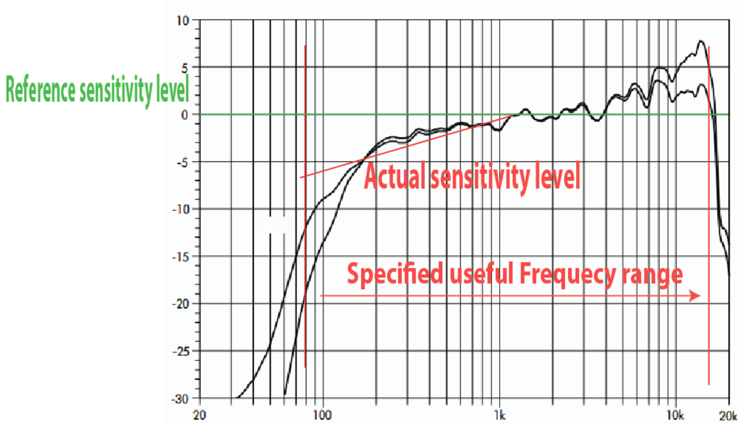

2 Actual sensitivity in the useable frequency band.

If we take the loudspeaker in the plot in Figure 2 and we want a relatively flat response over the 80 Hz to 1 6 kHz nominal bandwidth, we will have to equalise it to bring the red line up to the green line *(and we are being generous here in allowing a little LF roll-off)*. In effect, we will be adding up to 6dB of boost in the region where the input-signal demands most power from the system. At a low operating level, we have increased our system bass output, but as we have not increased our power handling capacity, by the time we get close to maximum output we hit system limits 6 dB sooner because of the extra 6 dB we added to the signal to flatten the bass. This is a situation where a high-sensitivity loudspeaker that needs equalising to become flat may well provide less maximum output than a less-sensitive loudspeaker with a flatter useable frequency response. In this case, our maximum output SPL figures actually mislead us. The quieter loudspeaker, on paper, may well be the louder loudspeaker in use.

A single SPL figure is largely useless to tell how loud a loudspeaker will be. We need to look at the frequency response, find the lower points, decide if we need to correct that dip with EQ, and modify our figures by subtracting the EQ. If we do not EQ we must decide if the dip is in an important part of our program, and if this is so, we must take that dip as the sensitivity figure, not the overall sensitivity. [*It is worth noting that in some circumstances we can get away with boosting higher frequencies without impacting power handling as the* *program material is often lower in the high frequency region. One-way, single cone loudspeakers can often take large amounts of HF boost without adverse effect.*]

Who needs to hear the system?

So, on to part two: directivity. We have to be careful of where the sensitivity figure was measured, but it is usually perpendicular, on axis, at the point where the output is often highest. In order to see what we really get, we need to see a measured polar plot of the loudspeaker and look at how the on-axis response relates to the off-axis response. The majority of listeners will often be off-axis, so if we have to achieve a percentile coverage figure from a loudspeaker, we will need to know what the sensitivity is over a certain percentage of the directivity plot. If a loudspeaker strongly beams on axis, we could see as much as 6 dB discrepancy between specified sensitivity and, for example, 60% of the specified coverage area. Many loudspeakers have nominal coverage within a -6dB beam width, although a better loudspeaker may only have a 3 dB reduction, but very few indeed will have the full specified sensitivity over a large, useful percentage of their specified coverage.

We should beware, as beaming (narrowing of the directivity) often happens at higher frequencies, and many people assume this to be the case, but there are also other factors. If we take a simple two-way loudspeaker with a 15” woofer and a one-inch-throat horn, we have two points where it will beam. The most obvious being up in the very high frequency range. However, to get a 15” woofer to meet a 1” compression driver is no easy feat. There are few 1” compression drivers and horns that will go low enough to meet a 15” loudspeaker before it is well into beaming. This means that often we see a narrowing of directivity right in the middle of the critical vocal region, an if we are not careful we can find ourselves with low coverage over large areas of our intended audience: it may be covering great at 3,000 Hz, but with little coverage at 1,000 Hz, right where you need it.

So where are we at now, and what does all this mean?

Well, we were told that loudspeaker “X” had a maximum calculated output of **125 dB** at full rated power. Loudspeaker “X” is a value-engineered product for mass market install, so of reasonably inexpensive construction, but impressive power handling capacity. This is quite a common situation.

If we guess it has 5 to 6 dB of power compression when given the 1.5 kW it claims to handle, we can now take the calculated maximum output down to **119 dB**. Now, we look at the frequency response and find a typical average requirement that it needs 5 dB of responsible tailoring, here and there in the useful bandwidth, to get it to sound as we want. This EQ (either boosting dips or cutting peaks) will bring us down another few dB depending on the nature of the EQ. Our 119 dB is now looking more like **116 dB** - if we are kind.

If we now look at the directivity over the required area and see that our required coverage area is within the -5 dB coverage angles, we can now re-calculate our maximum output again, and our 116 dB has now become **112 dB**.

But beware, even at this lower figure there is no guarantee that the loudspeaker is playing clearly, or is free from distortion or other mechanical stress. Power handling figures are often simply endurance tests, and do not require the loudspeaker to remain linear in any useful way.

This is catastrophic!

This is not a minor miscalculation! Many systems specified based on data-sheet figures have failed to achieve anything close to the expected results in the field. For our purposes, over our required coverage area, equalised to sound correct and at full power, what we were told was a 125 dB-output loudspeaker is actually a 112 dB-output loudspeaker, at best. If we believed the published specification we were hoodwinked.

From a systems specifier's point of view, it is a minefield, and one that need not exist if only we had a truthful measurement done to an international standard, such as we have for power-handling.

All system specifiers and installers should really be well aware of this issue, yet many are not. In some cases, it can mean a life-safety system falling way short of what is required to function safely. However, in others cases it is 'just' a disappointed customer. Many times, I have seen very-skilled systems suppliers lose tenders which took into account the perceived excessive product-specifications, compared to other tenders which only met the requirements on paper. Those people managing the tenders were clearly blissfully unaware of the pack of untruths within the product specifications. Time and time again, the cheaper tender wins, but the installation then falls short at the commissioning stage, as everyone hides their heads in the sand and re-measures and re-re-measures the system until they can find some trick to make it "*pass*". It can be devilishly difficult to persuade purchasers that the paper specifications are simply not in the real world. After the system is in the air it is always too late to realise.