Design Considerations for a Class A Amplifier Enclosure

Patrick K. (alias EUVL)

In summer 2010, Jan Didden invited me to write an article for Linear Audio. I was somewhat flattered, to find myself in the company of industry names like Pass and Self, but wasn't sure of the subject to write about. Then, folks at the DIY Audio forum (www.diyaudio.com) found out that I have some pretty Class A power amplifier enclosures, as well as access to good quality machining facilities. I agreed to design a Class A power amplifier enclosure and have it manufactured in a limited quantity. Of course, this was just one possible implementation and others may have different requirements. Therefore, I decided to write about the design considerations for such an enclosure; by going through the design process, using this enclosure as an example, the article hopefully provides guidance to those who wish to dimension and design their own.

Estimating total heat dissipation

The most important function of a Class A power amplifier enclosure is without doubt the provision of heat sinking to keep the temperature of the power transistors within their reliable safe operating region. The amount of heat to be dissipated determines the size of the heat sinks, which in turn determines the size and shape of the enclosure. Thus, the first thing one should work out is the total amount of dissipation.

The dissipation of a Class A amplifier itself is relatively easy to determine. It is simply rail to rail voltage times total bias. Take the example of the EUVL F5X, the rails are ±16V, and total bias per channel is 4A [1]. Total dissipation per channel is thus 128W; 256W for 2 channels.

That looks almost too easy, but we need to add a few more heat sources, such as rectifier diodes. To be on the safe side, one should allow at least 1V forward voltage per diode. So for one bridge per channel of F5X, the total diode dissipation is 8W.

With such high current draw, it is not practical to provide sufficient filtering capacitance to flatten the rectified waveform without excessively high ripple currents. So some form of power supply filtering is desirable. The simplest of those is the well-known CRC after the rectifiers. But a resistor produces voltage drop, and heat. For effective 1st order filtering, let's assume an RC corner frequency of say 5Hz. Assume also that each C is 44000uF, R needs to be 0R68, with a voltage drop of 2.7V and a dissipation of 11W per R, or 22W per channel.

But if you are prepared to drop 2.7V, you also have other options. One example is e.g. a MOSFET based follower ´regulator´, or more widely known as capacitor multiplier (**Figure 1**). Let's take a look at the datasheet of the 2SK3497 UHC-MOSFET. The transconductance at 4A is about 10.5S, i.e. the output impedance of the follower regulator is 0.1R. The Vgs at 4A is about 2.45V at room temperature, but dropping to 2.25V at a case temperature of 100°C. Such a high tempco is normally a sign of trouble (thermal run-away) if the device is used in amplifier circuits, as the bias would increase drastically as the device warms up. For the follower regulator application however, it merely means that the dropout voltage decreases by 0.2V with temperature rise. The regulator MOSFET itself gets less heat (dropout voltage x bias current), and the minute rail voltage increase to the amplifier itself is of no significance. The saturation voltage Vds is below 1V at 4A, so that the incoming ripple voltage can be up to 2.2V pk-pk, i.e. equal to the MOSFET Vgs at bias. For the negative rail, one can just use the complementary device, 2SJ618. They are both easily available for reasonable price. On paper at least, the follower regulator provides a better performance than a CRC filter. (But I am also aware that some people swear by passive unregulated power supplies for power amplifier applications).

So, for the example of a stereo enclosure for 2 channels of F5X, the total dissipation amounts to 294W. But we're still not done. Your friendly domestic power company is allowed to fluctuate the mains voltage within a certain margin for load regulation. There might be some tolerances on the transformer voltage under load, etcetera. It is perhaps wise is allow for at least 10% extra, and we end up with something around 330W for 2 channels, or 165W per heat sink.

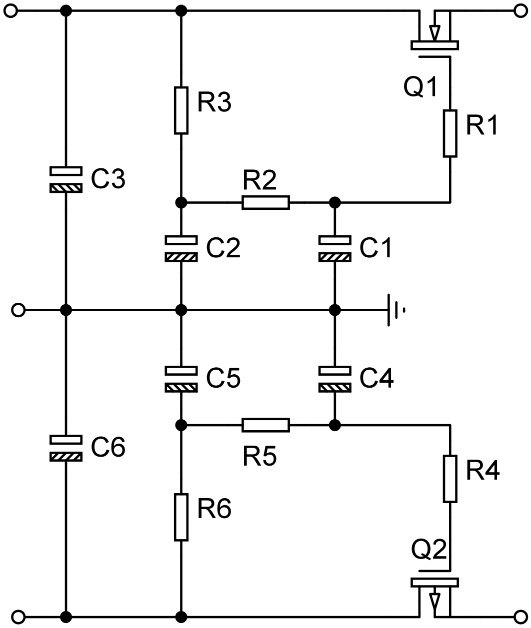

Fig 1 - typical UHC MOSFET follower regulator

Heat sink selection

You may wonder why there is so much discussion about heat sink selection. After all, most heat sink catalogues contain detailed technical sections, and the sinks are normally accompanied by an Rth vs. length curve. Just read the value off, multiplied by the dissipation, and you're done. Right ? Wrong !

The topic is complicated because there is no strict definition of how Rth is applied. Some manufacturers define their Rth at 75°C temperature rise (dT), some at 80°C, some even 90°C, and Rth is not constant with dT.

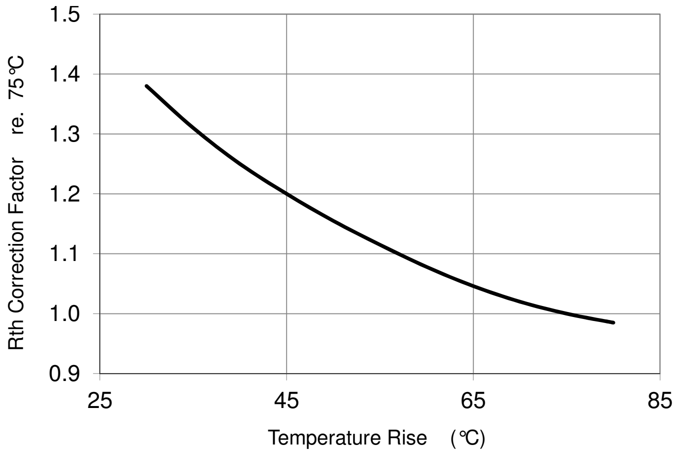

Figure 2a: Rth correction factor versus temperature rise

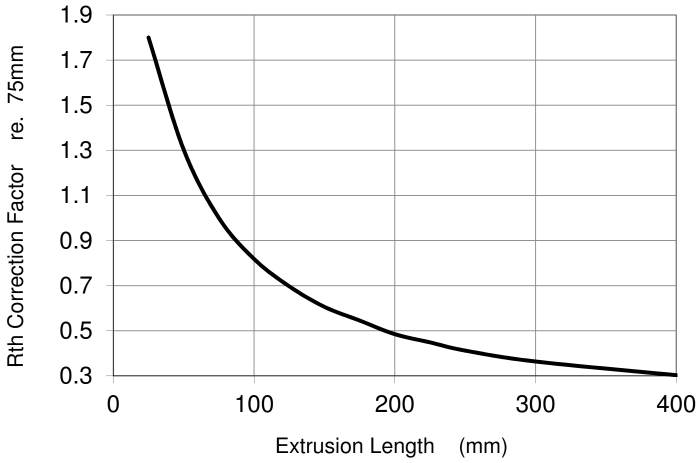

Figure 2b: Rth correction factor versus heatsink length.

Let's start with one particular brand to get a feel for the issues. One well known manufacturer for heat sinks extrusions is Aavid, and they have a quite extensive technical section. For example, they publish a set of correction tables for *their* heat sinks at their webpage [2]. We need to apply 3 correction factors; one for extrusion length, one for dT (starting with an estimate), and one for forced air cooling velocity (ignore if using natural convection, i.e. passive, no fan cooling).

The figures given by Aavid are illustrated in **Figure 2**. **Figure 2a** shows how the Rth of a heatsink referenced to 75°C temperature rise should be corrected for other temperature rise figures. **Figure 2b**, on the other hand, shows how the Rth figure changes with extrusion length. Note however, that these curves may differ significantly between manufacturers, although the general trend is similar. E.g. the correction factor of a Conrad heat sink is listed as 1.33 for dT = 30°C, whereas it is 1.26 for Aavid, a difference of 6%.

Looking at the length correction factor of Aavid, there is still quite some performance to be gained by increasing extrusion length beyond say 150mm. However, this is only a simplified story, for the length correction factor curve not only depends on the length, but also on fin geometry. The reason for that is that at increased length, another dominant factor comes into play namely the resistance to natural convection air flow between the fins. This flow resistance *increases* with increased length, but it also *decreases* rapidly (to the 3rd power as a first approximation) with increased fin pitch.

Figure 3 illustrate this well. While the Rth is still having a useful drop beyond 150mm length, the one with the 10mm pitch almost flattens after 150mm. Both of these sinks have 40mm net fin height, so that the only variable here is really fin pitch.

Thus, as you can see, if you have all the design freedom (e.g. machine your own heat sink or use custom ordered extrusions), you can tailor the heat sink profile to suit your required geometry. E.g., in a tower design using a heat sink of 420mm width x 50mm depth x 500mm height, the optimum fin pitch is about 17mm.
Most of us are not so fortunate to get customised profiles, and are stuck with standard profiles on offer. Many of these, especially those with width > 300mm, have a fin pitch of 10mm. In those cases, there is little point in choosing a heat sink much higher than 150mm. This in turn leads to the commonly seen "flat pack" enclosure designs.

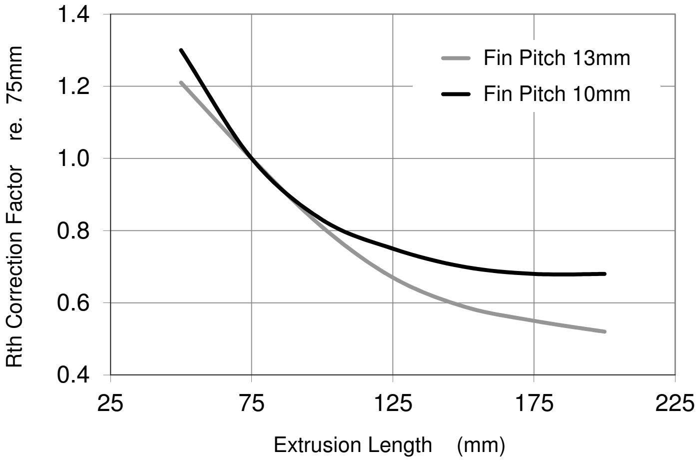

Figure 3: Rth correction factor dependency on fin pitch

As already explained, you cannot just use the Rth figures provided in the catalogue without applying all correction factors. And comparing Rth figures between manufacturers has to be done with great care. Once this is mastered, choosing the right heat sink becomes less of a daunting task.

So now let's look at the other side of the equation, the allowed temperature of the devices to be mounted on the heatsink. Almost all transistor manufacturers quote a maximum working junction temperature of 125 ~ 150 °C. Some manufacturers, while listing these numbers, put in a small note saying that for highest reliability, the design operating temperature should be substantially lower. A detailed description of the failure modes of semiconductors can be found in chapter 2 & 3 of the Toshiba Semiconductor Reliability Handbook [3,4], but even there no single magic number can be found which says make or break. The only figure which can be considered useful is Figure 2.2 in Chapter 3. In general it is wise to allow some 20% margin on the maximum operating power (or 10% voltage / current). This leads to a maximum junction temperature of 100°C as a design goal.

Let's consider again our F5X example. We have a total of 4 power MOSFETs, each dissipating 32W nominal. The maximum junction temperature allowed is 100°C. We also need to consider:

1. the temperature drop from junction to case;
2. the temperature drop across the insulator between the TO247 package and the heat sink.

For 1, most modern TO247 power transistor packages have an Rth of about 0.83 °C/W. At 32W, this equals 27°C. A good insulator such as Keratherm 86/82 has a specific thermal resistance of 35°C.mm^2/W, or 3.5°C for a TO247 with 32W; this is the value for 2. above. The maximum local heat sink temperature at the MOSFET case should thus be not higher than 70°C. As I will show later, the local temperature variation across a heat sink can be as high as 20°C. Thus a calculated *average* heat sink temperature of 60°C would make a good rule of thumb in heat sink dimensioning.

The need for actual thermal measurements

Even with all correction factors applied, the Rth figures still do not give you the entire picture. All manufacturer figures assume a uniformly distributed heat input to the entire back surface of the sink. This is hardly the case in reality. Thus, most heat sink suppliers would always recommend a 3-dimensional thermal simulation, to which most hobbyists do not have access, or thermal measurements using realistic heat sources and real heat sinks under actual working conditions.
Such a thermal measurement has been carried out for the F5X-design using a Conrad MF35-151 heat sink. To recap, the manufacturer's Rth with a 40°C dT correction factor applied is 0.26°C/W. The measurement was carried out by bolting 6 TO247 MOSFETs directly onto the heat sink, and applying a total dissipation of 135W, uniformly distributed between the MOSFETs. The actual steady state temperature on the heat sink is shown in **Figure 4**. The ambient temperature was 22.1°C.

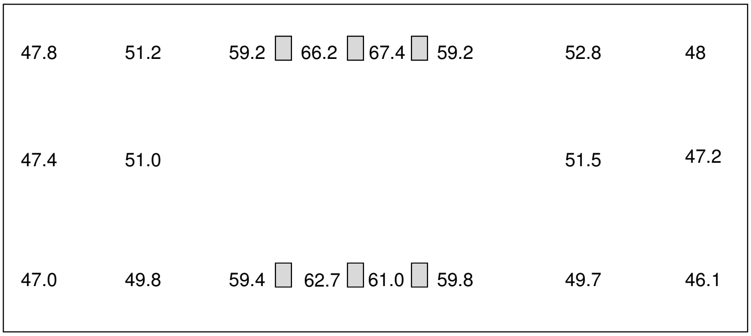

47.8	51.2	59.2	▯	66.2	▯	67.4	▯	59.2	52.8	48
47.4	51.0								51.5	47.2
47.0	49.8	59.4	▯	62.7	▯	61.0	▯	59.8	49.7	46.1

Figure 4: Steady-state temperatures across heatsink with six heat sources as shown (measured)

As could be expected, the top is hotter than the bottom, and temperature rise is highest in the middle. For the average sink temperature of 56.8 degrees, Rth = (56.8 – 22.1) / 135W = 0.26 °C/W. Worst case, highest temperature point 67.4 degrees, Rth = (67.4 – 22.1) / 135W = 0.34 °C/W . (The first number is identical to that given by the manufacturer !!). In this particular setup, if we include the additional dissipation due to rectification and power supply regulation, the MOSFET junction temperature can be calculated as:

22°C + (330W x 0.26°C/W / 2) +10°C + 3°C + 27°C = 105 °C

A bit on the high side, perhaps. The additional surfaces of the front and rear panel will help to alleviate the situation. For example, if we mount the rectifiers and the regulator MOSFETs directly on the front panel, then the amplifier MOSFET junction temperature becomes

22°C + (282W x 0.26°C/W / 2) +10°C + 3°C + 27°C = 99 °C, only a slight improvement.

The average heat sink temperature is then 59°C with 22°C ambient.

The heat source distribution used in the experiment is, from the thermal point of view, far from optimum. By rotating it 90° about the centre of the heat sink, for example, the heat sources are wider apart laterally to make better use of the width of the sink. Furthermore, we can place the heat sources further away from the top than from the bottom edge, thus giving more dissipation area around the sources in the vertical direction. Both of these factors will lead to a lower peak temperature, even though the average temperature will remain the same, as the following simulation result shows. The maximum temperature difference over the heat sink is reduced from 20°C to 7.3°C just by clever but realistic placement of the heat sources.

Figure 5: Reduced temperature differences with optimised heat source placement (simulated)

Just to confirm the accuracy of the simulation, the same configuration was measured experimentally after allowing the heat sink to reach steady state after 1 hour (**Figure 6**):

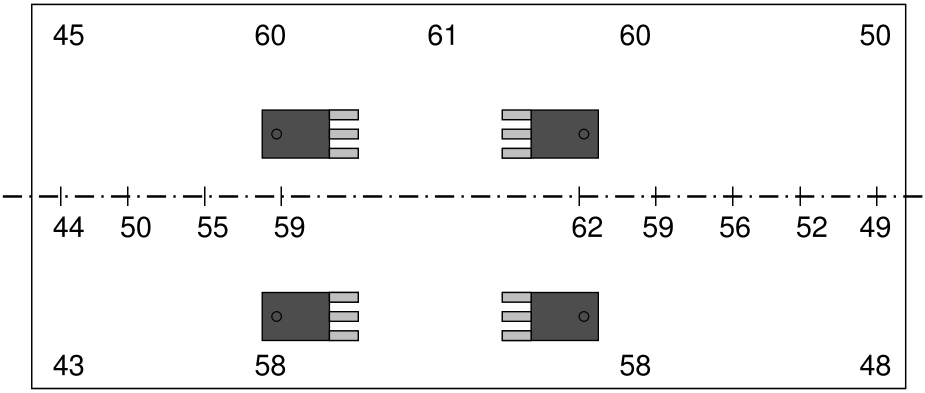

Figure 6: Reduced temperature differences with optimised heat source placement (measured)

The lower temperature measured at the free ends of the heat sink are contributed to extra heat sinking effect of the front and rear panels of the case.

You might be tempted to even out the temperature immediately around the heat source further with the aid of heat spreaders. However, bear in mind that such spreaders are only effective if they cover a large area, and have substantially less lateral thermal resistance than the heat sink base plate, which in this case is 8mm solid aluminium. Thus, an effective spreader has to be at least 12mm aluminium, or 6mm copper. And it is a misconception to think that bolting a heat spreader on is equal to adding thickness to the heat sink base plate monolithically. This is certainly not the case, due to the finite flatness of the interfacing surfaces, and hence the necessity to use a conformal interfacing medium, such as thermal wax. Such a medium has easily 10x higher thermal resistance than the parent metal itself. Only by soldering a heat spreader of the same material directly onto the heat sink you will realise the full benefit.

Fine points to note

So far so good. But you can easily lose more performance if you are not careful. For example, the calculated values here apply to heat sinks with vertical fins free from obstructions at the top and the bottom. Turning the heat sink 90° such that the fins now become horizontal will increase its Rth by 10~15%. Similar loss in performance is to be expected by placing the heatsink horizontal (exceptions are heat sinks that are designed for such orientations, such as finger coolers for CPUs), and by using clear anodising instead of black anodising. Another surface treatment sometimes found on heat sinks is powder coating. It is hard to believe that there is no performance penalty by using powder coating instead of anodising. A typical anodising layer is somewhat less than 10µm thick consisting of porous aluminium oxide with a thermal conductivity of about 16 W/m/K. A typical powder coating is somewhere between 50~125µm thick with a conductivity of 0.2 W/m/K. While its contribution to the resistance of the total thermal path might not be dominating, it is still a factor of 400 worse than anodising. One cannot help to think that the advocates of powder coating might have a hidden agenda as powder coating is much more effective to hide cosmetic defects in extruded aluminium [5].

Other thermal design issues

For optimal deployment of the heat sinks, it has already been mentioned that the fins should remain unobstructed. The heat sink back surface, as well as inner surfaces of the front and rear plates, also constitute additional heat radiating surfaces. These should therefore also receive minimum obstruction. The best way to do so is to leave the bottom plate clear from the heat sinks. There is no real reasons to fear large openings at the bottom unless you have small house pets (lizards, mice and cockroaches included), small kids who could craw under the 12mm gap between heat sink and floor, and the like. The top plate should optimally be just a wire mesh or perforated plate with >50% opening ratio.

But they would vibrate, you might say. They would, since they are thin and light. But because they are highly perforated, the vibration is likely to come from mechanical coupling through a heavy amplifier enclosure rather than acoustic pressure waves bouncing back and forth in your listening room. And because they are highly perforated, they are also not effective as vibrating membranes which will generate disturbances to your listening experience. In this respect, wire mesh is actually a better choice, as the internal friction between the large number of wire crossings is highly effective in damping any vibrations. In any case, magnetic materials should be avoided for the enclosure. They can easily be set into vibration when placed sufficiently close to the stray magnetic field of the power transformer(s).

There are many well known insulators one can use between the MOSFET and the sink – aluminium oxide, kapton, mica, even beryllium oxide which is now banned worldwide due to toxicity. The tricky bit is that they all come with different standard thicknesses, and many of them require transition layers on both sides, such as a thermal conductive wax. The proprietary Keratherm 86/82 [6] has also been mentioned above.

Table 1 below shows the temperature difference across different insulators, using the footprint of a 2SK1530 (21 x 26 mm), and 32W dissipation.

Material	Thickness (m)	Thermal Conductivity (W/m/K)	Transition layers (Aavid Ultrastick)	°C/W	dT @ 32W
Al_2O_3	0.001	25	2	0.13	4.26
BeO	0.001	330	2	0.07	2.10
Kapton	0.00005	0.24	2	0.44	14.13
Mica	0.00005	0.4	2	0.29	9.25
Keratherm 86/82	0.00025	6.5	0	0.07	2.25

Table 1: Thermal insulator temperature gradient for different materials

As one can see, the best performer is BeO, now unobtainium. This is closely followed by Keratherm, not because it has the best thermal conductivity, but because it is compliant and does not require transition layers (thermal wax such as Aavid Ultrastick [7]). Only then comes alumina, then mica, and Kapton is actually pretty useless despite its high price and glamour factor.

There are more and more such ceramic filled elastomer insulator sheets appearing on the market in the last few years. So Keratherm is by no means the only or the best. But the above illustrates clearly that the complete picture is more important than single material parameters.

Let's look at device clamping next. For Keratherm-like insulators a tensile strength of 10 N/mm^2 is specified. Even though compressive strength is normally higher, it is a good indicator for the maximum amount of pressure one can apply. For a 2SK1530, this converts to 5000N !! A single M3 bolt in high tensile steel can exert 3800N. But if you want non-magnetic stainless steel, then it reduces to 1750N. But does such a high pressure help? Bergquist shows [8] that depending on insulator types, there is performance to be gained by applying pressure as high as 1.3 N/mm^2 when using elastomeric insulators. It is less so if a thermal wax is used as a transition layer. E.g. only 0.6 N/mm^2 is recommended for Aavid Ultrastick. This is probably due to the fact that the thermal wax will change to liquid phase at some 70°C, and thus will then adapt itself to the exact shape of any gap between device and sink.
The problem with single-screw mounting is that the pressure is applied at the wrong place, i.e. not directly under the silicon substrate. It is normally (and understandably) at the top of the package above the chip itself. Thus, possibly less than half of the applied pressure is seen by the chip, and the pressure decreases rapidly with distance from the mounting hole. Therefore, even if you don't intend to apply the full 1750N, it is still beneficial to deploy an additional clamp directly over the silicon substrate, as shown in **photo 1**. The bottom surface of the clamp should have a slight curvature, or a thin elastomeric layer in between, to avoid excessive stress concentration which might in turn cause premature failure of the transistor encapsulation over lifetime.

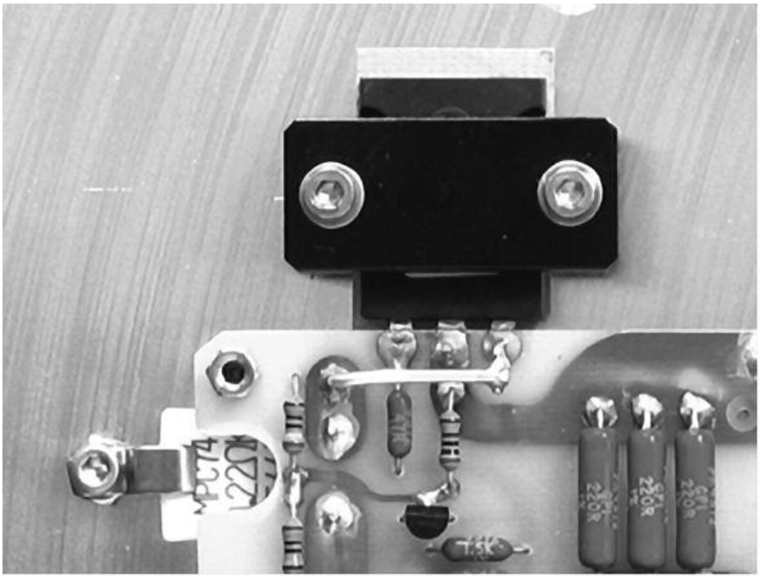

Photo 1: More uniform pressure distribution by use of a clamp.

Overall Dimensions and Enclosure Proportions

Beauty is in the eye of the beholder. And this applies equally well to audio gear. But for audio gear, the issue is complicated by the WAF (Wife Acceptance Factor).

Having said that, there is no escaping the fact that high power Class A amplifiers need large heat sinks, and people expect ´fat´ enclosures psychologically. Also high end gear always has thick front panels; thin front panels always look "cheap", irrespective of overall size.

So it is then a matter of proportions. An informal survey of commercial products suggests that the ratio of width to height should be between 2.2 to 2.7, and the front plate thickness should not be less than 4% of the width. The depth should be around 1x to 1.2x of the width. Too much depth means long heat sinks and it is difficult to spread the heat adequately.

Since the depth and the height are already given by the selected heat sink, the overall dimensions can be arrived at quickly. One only needs to make sure there is sufficient space for other components, notably the power transformer(s). And these should not come too close to the amplifier PCBs because of their stray magnetic fields. Transformers with a built-in magnetic shield help; Faraday cages in Mu-Metal are even better.

In the case of the F5X diyaudio enclosure design (**photos 2, 3, 4**), the inside dimensions are 250 W x 344 D x 140 H. This will take 2 F5X PCBs, 2x 500VA toroidal transformers (for dual mono), 16 power supply capacitors of 30mm diameter x 50mm height, plus enough space for wiring, slow start, protection circuits, and even voltage regulators with associated heat sinks.

Photo 2: FX5 enclosure design, top-front view

Photo 3: FX5 enclosure design, inside top view

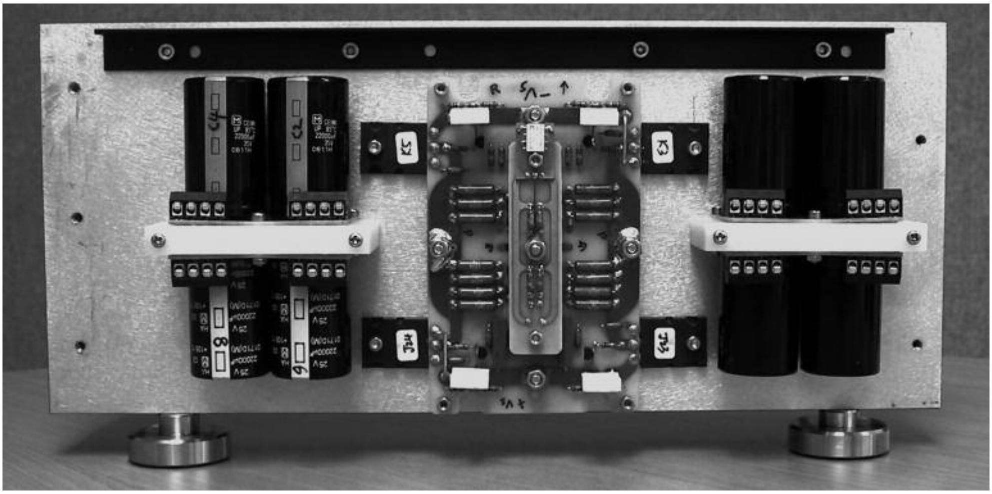

Photo 4: FX5 enclosure design, one heatsink with mounted amplifier channel

However, one of the most impressive front panels the author has come across belongs to our Editor –it can only be described as a piece of art *(it* is *a piece of art! See [9] - Ed.).*

Minimising vibrations

We already briefly described the use of perforated plates for top (and bottom) to minimize acoustically induced vibrations due to sound waves in the listening room. The other major source of vibrations is the power transformer. While over-dimensioning can reduce humming, the high pulse current

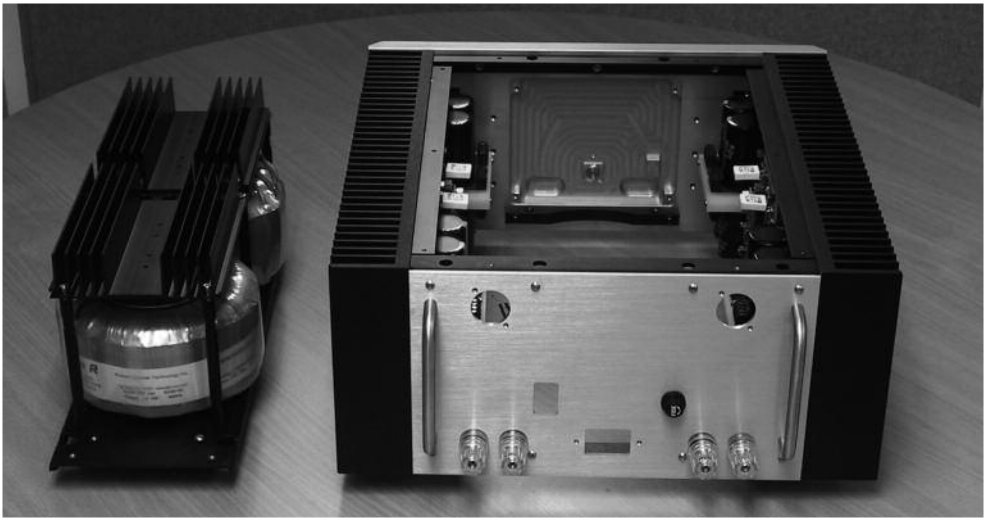

Photo 5: Vibrationally insulated power transformer sub chassis

experienced by all Class A power supplies still causes mechanical vibration at mains frequency. This vibration should also be decoupled from the enclosure, especially those made from thin iron plates. The best way to do so is to place the transformer(s) on a separate cradle with its own feet (**photo 5, 6**). The cradle is normally bolted to the amplifier enclosure during transportation. Once it is put in place, the bolts can be gradually loosened such that the cradle is lowered onto its own feet, leaving a physical gap between the cradle and the enclosure. *(This is similar to how turntable chassis are isolated from the environment in operation but bolted to the enclosure for transport – Ed.)*. Both the cradle and the enclosure feet are lined with elastomeric material of some sort which provides vibration isolation and damping at 2x mains frequency. Vibration from the transformer now has to go through 2x such elastomers before it can reach the amplifier enclosure, the table or rack being an additional sink for such vibration energy.

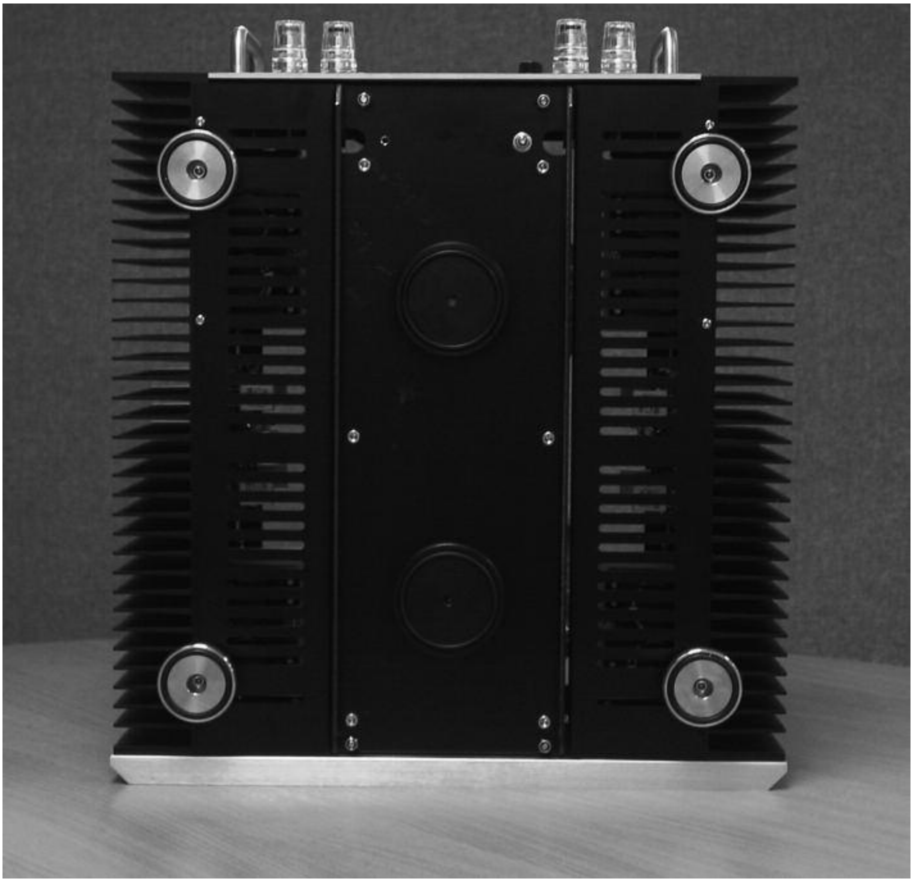

Photo 6: Power transformer sub chassis in transport position

Editors' note:

The thermal simulations have been performed on-line at www.r-tools.com. *Sample thermal simulations are available at no charge.*
I am indebted to Zaher Aboumourad *of* Mersen Canada Toronto Inc *(creaters of R-tools) for kindly providing the thermal maps in this article and on the cover.*

References

1. http://www.diyaudio.com/forums/pass-labs/183362-f5x-euvl-approach.html
2. http://www.aavidthermalloy.com/technical/correct.shtml
3. http://www.semicon.toshiba.co.jp/eng/shared/reliability_pdf/bde0128g_chap02.pdf
4. http://www.semicon.toshiba.co.jp/eng/shared/reliability_pdf/bde0128g_chap03.pdf
5. http://www.alphaanodising.co.uk/technical%20info.html
6. http://www.kerafol.com/en/innovation-in-thermal-management/products-solutions/standard-films/keratherm-red.html
7. http://www.aavidthermalloy.com/products/options/ultrastick.pdf
8. http://www.bergquistcompany.com/pdfs/techLibrary/SilPad_Overview.pdf
9. http://www.linearaudio.nl/beenthere-1.htm ; also see http://www.helgijoensen.no/velcome.html

10 http://www.aavidthermalloy.com/technical/index.shtml

11. http://www.aavidthermalloy.com/technical/papers/pdfs/select.pdf
12. http://www.fischerelektronik.de/fischer/uploadfischerfcool/Fischer/A01_E.pdf
13. http://www.fischerelektronik.de/fileadmin/fischertemplates/download/Katalog/f.cool.e.pdf
14. http://www.r-theta.com/products_extrusion.asp
15. http://www.r-theta.com/rtools_front.html
16. http://www.r-tools.com/
17. http://www.federalpacific.com/literature/drytrans/10transformernoise.pdf
18. http://www.pmeweb.co.uk/audio/psu.pdf