

Heatsinking Requirements of Class D Amplifiers

Class D amplifiers' known high efficiency often leads to a gross underestimation of the cooling required. This document outlines an approach to estimating the cooling necessary for a given choice of amplifier module and power rating.

1.1 Design for a cool amplifier

As mentioned, it is important to apply proper cooling to our high power SMPS and amplifier modules. A proper thermal design depends on multiple factors. This document aims to give you more insight into this subject.

Let's start:

How well a cooling arrangement works is expressed as "thermal resistance Θ ". This number indicates how much warmer (in Kelvin or degrees Celsius) the heat sink becomes compared to ambient temperature when asked to dissipate a certain amount (in Watts) of thermal power. Thermal resistance is expressed in Kelvin per Watt. The lower thermal resistance becomes, the more power you can get rid of for a given allowable temperature rise.


This gives us the following formula: $\Delta t = \Theta * P$.

Where Δt is the temperature rise of the heatsink, Θ is the thermal resistance and P is the heat generated by the module (also P_{loss}). A good rule of thumb: $P_{loss} = P_{idle} + (0.06 * P_{out})$. This will be explained more thoroughly later in this document.

A sensible maximum heat sink temperature would be 80°C and a conservative ambient temperature would be a toasty 35°C. This gives us a thermal design guideline of $\Delta t = 45K$. It is important to remember that ambient temperature refers to the temperature 'around' the module, meaning the air temperature inside your casing.

1.1.1 Interpreting datasheet graphs

Each new datasheet now contains a chapter dedicated to heat-loss. In this chapter one will find two graphs; one illustrating the heat-loss vs output and one illustrating efficiency vs output. An example of these graphs is given below. These graphs are taken from the NC400 datasheet.



1.1.2 Evaluation using the excel sheet

If the graphs are not present, or your setup is a little different, you can use the spreadsheet called "losscalc.xls" from the Hypex website. The underlying formulae are given in the annex to this document. In this paragraph we explain how this spreadsheet can be used to evaluate your setup and gain insight about this subject.

The user is required to plug a few data sheet values into the spreadsheet:

- Rated amplifier power (e.g. 250W)
- Rated load (e.g. 4 ohms)
- Idle loss (e.g. 3.8W)
- Efficiency at rated power (e.g. 92%)
- Rated SMPS power (e.g. 600W)
- SMPS idle loss (e.g. 7.5W)

In addition the user should specify how many amplifier channels are used (e.g. 2) and the actual load used for thermal evaluation (typically 8 ohms for agency testing).

Amplifier loss calculation	SMPS loss calculation	Combined result
<small>Bruno Putzeys, Hypex.</small>		
Fill out the items marked in yellow		
Rated Power 250 W	Rated Power 600 W	
Rated Load 4 Ohms	Number of simultaneously driven amps 2	
Idle loss 3.8 W	Idle loss 7.5 W	Combined idle loss 15.1 W
Efficiency at rated power 92%	Efficiency at full power 90%	
Actual Load 8 Ohms		
Power loss at 1/8th Pr 4.36 W	Power loss at 1/8th Pr 11.48 W	Power loss at 1/8th Pr 20.20 W
Power loss at 1/3rd Pr 5.29 W	Power loss at 1/3rd Pr 16.86 W	Power loss at 1/3rd Pr 27.45 W

Figure 1: Typical worksheet screen shot

The spreadsheet returns graphs for amplifier, SMPS and combined power loss and efficiency. Also, the precise figures are given for 1/3rd and 1/8th of rated power (into the test load).

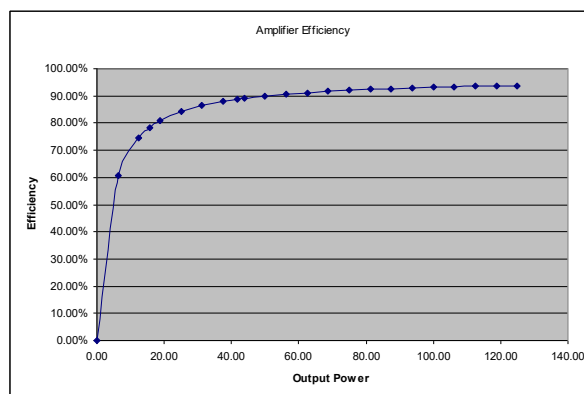


Figure 2: Amplifier efficiency extrapolated from data sheet values.

This spreadsheet can be used for a variety of "what if" scenarios. For instance it can be shown that using 180W rated modules to build a 50W rated amplifier will result in remarkably disappointing efficiency. Or, say that a design already runs warm at idle, this calculation shows that the thermal design is clearly deficient. The above example of 2x UcD250 + 1x SMPS600 has an idle dissipation of 15W. If the thermal design is such that it already runs uncomfortably warm, be reminded that the dissipation at 1/3rd of rated power is going to be about 27W. Somehow halving the idle losses may reduce the perceived problem but the dissipation at 1/3rd rated power is still 20W which the thermal design clearly can't handle. In other words, powering the box up and evaluating the soundness of the thermal housekeeping by holding one's hand on the lid is quite dangerous.

2 Selecting a heatsink

2.1 Required heatsinking.

What becomes rapidly clear is that the SMPS and amplifier modules need some way of getting rid of heat. The myth that “it’s class D so it doesn’t need cooling” is nothing more than that – a myth which is sometimes kept alive by people simply turning palpating amps running idle. Some class D modules, particularly those optimised for idle current rather than sound quality will indeed perform well in this test but unless their efficiency at higher power is phenomenal the thermal performance at elevated power will be disastrous.

The required thermal resistance works out as:

$$\theta_{ha} = \frac{T_{h,max} - T_{amb}}{P_d} = \frac{\Delta T}{P_d}$$

$T_{h,max}$ is the maximum temperature you want to allow. There is a maximum $T_{h,max}$ given in the datasheet. T_{amb} is the ambient temperature (also given in the datasheet) and P_d is the power dissipation calculated earlier. For heatsinks that are not externally accessible, it is acceptable to take $T_{h,max}$ from the data sheet. Quite often though heatsinks are mounted externally to optimize airflow. In that case keep in mind that under normal operating conditions (1/8Pr) IEC60065 restricts permissible ΔT of externally accessible parts to 40°C.

Given the large surface area of the heat spreader (the metal plate on the amplifier module), thermal resistance between it and the heatsink can be ignored. This does not count for the UCD700 and UCD2K heatsinks..

2.2 Using a COTS heatsink

In the simplest case a commercial off-the-shelf (COTS) heatsink will be bolted on the module, in which case the thermal resistance can be looked up in the data sheet. Two caveats apply. One, this value can only be relied upon when air flow is unrestricted. Secondly, it’s precisely the heatsink temperature rise which drives the convection that cools the heatsink so the relation is not at all as linear as the concept of thermal resistance would have you believe. Heatsinks are specified with a temperature rise of 75K, so one should apply a correction based on the desired temperature rise:

$$\theta_{ha,75K} = \theta_{ha} \cdot K_{corr}(\Delta T)$$

ΔT	K
30°C	0.80
40°C	0.85
50°C	0.90
60°C	0.95
70°C	0.98
75°C	1

The corrected value is the one you’ll be looking for on a data sheet for a heatsink.

2.3 Using the cabinet as heatsink

In active speakers it is common practice to bolt the amplifier modules on a plate that closes the back of the enclosure. In that case the 40 degree temperature rise limit applies and a good rule of thumb is 3.5 to 4W per square decimetre (0.22 to 0.25W/sq.in). The lower end of this range corresponds to a large plate with a compact heat source (e.g. a single amp module) in the middle. Usually though several amplifier modules are spread around the plate, making it more efficient. If the amplifier is mounted in a separate subenclosure, this figure can be improved further using vents that admit airflow.

2.4 A word on heatsink colour

Some products may intentionally or unintentionally get used in direct sunlight. In that case, black is not a good colour.

A myth that's repeated distressingly often even by manufacturers of heatsinks is that it helps if the heatsink is black because that would improve heat radiation. This is a confusion of terminology: black at visible wavelengths is not the same as black in the far infrared.

Heat radiation and absorption of radiation are converse processes and indeed a body will be better at radiating heat if it is "black". But black here means: absorbing all radiation *at that particular wavelength*. The peak wavelength of blackbody emission at 50°C is 9µm. Visible wavelengths are in the range 0.4µm...0.7µm. For a heatsink at 50°C to be an efficient radiator it doesn't matter whether it *looks* black but whether it's black in the far infrared. The oxide layer on anodized aluminium is strongly absorptive and hence emissive at 9µm, irrespective of whether any pigments were added. Those only affect optical wavelengths. Anodizing a heatsink improves its efficacy but colour is irrelevant there.

Where colour becomes relevant is when you put the heatsink in direct sunlight. When you step into the sun and feel its warmth in your face, that's mainly visible light you're absorbing into your skin. So when you add black pigments in the anodising stage, you aren't improving its ability to radiate heat at 50°C *but you are significantly improving its ability to absorb sunlight*. Clear anodisation is by far better for sun-proofing a heatsink and it has no drawbacks at all. The same thing goes for paint. Virtually all paints are close to black in the far infrared. This absolutely includes *white* paint.

Clear anodization or white paint is without qualification the smart choice for speaker amp plates if there is even the slightest chance of the product ever getting used in direct sunlight.

2.5 Using thermal grease (or not)

Although the surface of a heatsink may look entirely flat, in many cases it is not. Thermal grease may be used to enhance the thermal conductivity between two parts, however thermal grease often has a higher thermal resistance than the material of the heatsink itself. A thick layer will therefore act as a thermal insulator instead of enhancing the thermal conductivity.

Depending on the application and the specific heatsink, one must consider if thermal grease is necessary or is even beneficial. A larger heatsink contact surface (like the NC400) may not benefit from thermal grease as much as a very small contact surface (like a TO-220 package).

We do not specifically advice whether or not to use thermal grease, but if you want to use thermal grease, we have some tips for you:

Generally the layer of thermal grease must be as thin as possible and evenly distributed over the contact area. Use your finger (with latex gloves) to spread the grease thin and evenly over the contact area. When a module is bolted tight and thermal grease is squeeze out from between the both surfaces, you often used to much!

If, for any reason, you need to return your module to us, please do our RMA division a favour and clean off the thermal grease you have applied! This can generally be done best using some alcohol wipes.

3 Annex 1:

3.1 Formulae for estimating amp losses from data sheet values

Spec sheets of class D amplifiers tend to give only two data points: idle loss and efficiency at rated output which is defined as:

$$\eta = \frac{P_{out}}{P_{in}} = \frac{P_{out}}{P_{out} + P_{loss}}$$

This is done based on the working assumption that we can linearly interpolate power dissipation between the idle and full power regimes:

$$P_{loss} = P_{idle} + \frac{P_{out}}{P_{rated}} \cdot P_{excess,rated}$$

Where $P_{excess,rated}$ is the difference between total losses at maximum power and the idle loss. In actual fact though, the direct cause of those losses is output current, not output power.

$$P_{loss} = P_{idle} + \left(\frac{I_{out,rms}}{I_{rated,rms}} \right)^2 \cdot P_{ex}$$

Where $I_{rated,rms}$ is the rms output current at rated power into rated load.

Rewriting this as:

$$P_{loss} = P_{idle} + I_{out,rms}^2 \cdot \frac{P_{excess,rated}}{I_{rated,rms}^2}$$

We see that the working assumption is equivalent to saying that the excess losses can be interpreted as a resistor through which the output current flows. Physical terms in this resistance are the power FETs' R_{dson} , trace resistance and coil resistance. The remainder is an abstraction of switching losses which isn't frightfully accurate but more than accurate enough for basing a thermal design on.

Thus:

$$P_{loss} = P_{idle} + I_{out,rms}^2 \cdot R_{loss}$$

Where

$$R_{loss} = \frac{P_{excess,rated}}{I_{rated,rms}^2} = \left(\frac{1-\eta}{\eta} - \frac{P_{idle}}{P_{rated}} \right) \cdot R_{load,rated}$$

3.2 Estimating SMPS losses from data sheet values

A similar reasoning goes for the SMPS. Power loss is linearly interpolated between idle and full power regimes. Note that the rated output power of the SMPS is actually "rated amplifier watts" i.e. taking into account the assumed efficiency of the attached amplifier.

The output voltage is more or less constant so we do not need to make a detour via the load current.

$$P_{loss} = P_{idle} \cdot \left(1 - \frac{P_{out}}{P_{rated}} \right) + \left(\frac{1-\eta}{\eta} \cdot P_{out} \right) \cdot \frac{P_{out}}{P_{rated}}$$