

Chapter 7: *Power System Thermal Design*

Long life and reliable operation are key requirements of many electronic systems. Excessive heat dissipation in components can cause them to degrade or fail sooner than would normally be expected. Good thermal design results in power converters with lower component temperatures, higher reliability and longer life.

Power System Thermal Design

7.1 Introduction

Thermal design is one of the major keys to a successful power system. As reliability is directly linked to component operating temperature, knowledge and control of component temperatures, both inside and exterior to power converters, is absolutely required in order to ensure system reliability. The nature of the power cooling system is changing as more and more applications use some form of distributed architecture. In the days of centralized systems, all of the power converters were located in one place, and cooled with either self-contained or dedicated external fans, which were optimized for the particular application, and designed by specialists. Today, in DPA systems, on-board DC/DC converters are used and scattered throughout the system. They usually do not have any customized dedicated cooling mechanism, but rather utilize the standard cooling paths inherent in the system's packaging design. These paths, which can include conduction through the PCBs, and free or forced convection, are typically shared with the load circuits themselves. This new cooling environment, along with circuit board designers not always being familiar with DC/DC converters, has resulted in thermal problems being one of the most common trouble spots with first time DPA developers.

In this chapter we will cover the thermal design concepts that would be required to design a successful DPA power system. A primer on heat transfer will review the physical mechanisms used in electronics cooling and summarize the thermal properties of materials used in electronics packaging and cooling. Next, we will provide a description of the typical thermal environments encountered in electronics systems along with generally accepted limits associated with the thermal design of present-day electronic equipment. Power system thermal design techniques will be discussed, including the importance of efficiency, derating the output power of converters, airflow management, conduction and convection cooling

Power System Thermal Design

guidelines, usage of heatsinks, alternative cooling techniques, thermal measurements and simulation. The thermal aspects of power converter selection will be discussed, with emphasis on internal component temperatures and the trade-offs between open frame and encapsulated designs. Finally, we will provide some examples of system thermal design using actual Artesyn converter products.

7.2 Heat Transfer Primer

The most commonly used mechanisms for heat transfer in electronics design are conduction and convection. In this section, we will review the basic operation of each of these mechanisms and present the mathematical relationships needed for their analysis and modeling. We will also briefly describe the less commonly used transfer mechanisms of radiation and phase change. The basics of the terminology used in airflow analysis will also be described.

Conduction Cooling - Conduction is a heat transfer mechanism used in every electronics design. Even if a system is designed for convection cooling of the circuit boards, conduction is still the dominant heat transfer mechanism within the component devices and on the circuit board. This is especially true of power electronics, where concentrations of heat are developed in components such as power silicon and magnetics. This heat must be transferred via conduction to the component case, the circuit board or a heatsink before it can be handled by the system-level cooling mechanism(s). Consequently, all electronics designers must be conversant with the techniques of thermal conduction and its analysis.

Conduction is defined as the transfer of heat through a static (non-moving) material. Heat will flow through the material if there is a temperature differential across it. The heat flow will be in the direction from the higher temperature to the lower temperature. The rate of heat flow will depend upon the temperature differential and the thermal conductivity of the material. The overall relationships are analogous to current flow in an electric circuit, where temperature differential is equivalent to voltage differential, thermal conductivity is equivalent to electrical conductivity, and the rate of heat flow is equivalent to electrical current. These relationships are shown in Figure 7.1.




Figure 7.1 - Dynamic Response

The thermal conductivity, K , of the material will determine the amount of heat flow for a given temperature differential or the needed temperature differential for a given heat flow. There are many different materials used in electronics systems, and they have vastly different thermal conductivities, as shown in Figure 7.2. The conductivity units in the figure are $W/in\ ^\circ C$. They can be easily converted to the metric units of $W/m\text{-}^\circ K$ by multiplying by 39.37. The table is arranged in rough order of conductivity. Note the large variation between

air and copper - over 14,000 to 1. Both of these materials are used in most electronics packaging systems, along with several materials with intermediate thermal conductivities. For reference purposes, the Thermal Coefficient of Expansion (TCE) of some of these materials is also shown.

Materials	Thermal Conductivity - K W/IN °C	TCE ppm / °C
Still Air	0.007	-
Epoxy Glass	0.01	15-45
Thermal Grease/pads	0.01 - 0.04	-
FR4PCB - z axis	0.01	-
FR4PCB - 1oz x,y axis	0.59	-
FR4PCB - 2oz x,y axis	1.13	-
FR4PCB - 3oz x,y axis	1.55	-
Distilled Water	0.028	-
Glass	0.028	8.5
IMS Insulation	0.03	25
Thermal Epoxy	0.02 - 0.03	20
Titanium	0.02 - 0.03	-
Stainless Steel	0.66	16.2
Alumina Al ₂ O ₃	0.70	5 - 7
Eutectic Solder	1.27	21
Silicon-Electronics Temps	2.5 - 4.0	7.6
Aluminium Alloys	2.5 - 5.5	23.6
Beryllium Copper	3.0	-
Brass	3.5	19
Beryllia BeO	5.0	6.8
Gold	7.6	13.2
Copper	9.9	16
Silver	10.6	17
Diamond	17.5	1.2

Figure 7.2 - Thermal Properties of Materials used in Electronics

Thermal conductivity (and its reciprocal thermal resistivity) is a property of the bulk material. To determine the actual amount of heat flow in a specific situation, the physical dimensions of the conducting material must be known. The physical dimensions, along with the thermal conductivity, will allow calculation of the **thermal resistance** for the heat flow path in question. The thermal resistance, R_{th} , is given by the expression:

$$R_{th} = \frac{L}{K \cdot A} \text{ } ^\circ\text{C/W}$$

Equation 1

where: L = Length of thermal path - in

A = Cross-sectional area of thermal path - in²

K = Thermal Conductivity - W / in °C

The resulting thermal resistance is equivalent to the resistance value of a resistor in an electrical circuit. We can now calculate the temperature rise, ΔT , when a given amount of power (heat) flows through this specific thermal resistance:

$$\Delta T = Q \cdot R_{th} \text{ } ^\circ\text{C}$$

Equation 2

where: Q = Heat Flow in Watts

Figure 7.3 shows an example of how the above expressions can be used in a power electronics application. A DC/DC converter has a heatsink attached with a thermal pad sandwiched between them to eliminate thermal gaps. There are 12 Watts dissipated internal to the converter, and it is assumed that all of this power is dissipated through the heatsink and is distributed uniformly over the area of the thermal pad. From Figure 7.2, we find that the typical thermal conductivity of this material is 0.03 W/in °C. It is 2 inches on each side giving an area of 4 in². The pad thickness is 0.02 inch, which represents L, the length of the thermal path, in equation 7.1. Using these values, the thermal resistance is calculated to be 0.17 °C/W. Equation 7.2 is then used to obtain the temperature rise from the 12 watts dissipation, giving a ΔT of 2.0 °C.

Power System Thermal Design

Thus, if the heatsink temperature were 50°C, the DC/DC baseplate temperature would be 52 °C. As this example shows, calculation and analysis of conductive heat flow is computationally straightforward due to its static and linear nature.




Figure 7.3 - Example of Conductive Heat Flow

Convection Cooling - Convection is the most widely used system-level heat transfer mechanism in electronic systems. Convection is defined as the transfer of heat from the surface of an object to a moving fluid (either gas or liquid). If the fluid flow is created by gravitational forces on the fluid as its density varies, it is referred to as natural convection or free convection. If the fluid flow is created by external means such as fans or blowers, it is referred to as forced convection. In either case, it is not a static system, but is dynamic in nature due to the movement of the fluid. The most commonly used fluid is air.

Figure 7.4 provides conceptual representations of convection cooling systems. In the natural convection systems, the airflow mechanism is as follows. As heat is transferred to the air in contact with the surface of the converter, the air temperature rises and its density

decreases. The reduced force of gravity on the less dense air will cause it to rise in relation to the surrounding air. Its place will then be taken by cooler air from below. This process will then produce an overall flow of air from the bottom to the top of the surface being cooled. If the enclosure has openings at the top and bottom, the airflow will consist of cooler external air flowing through the enclosure. If the enclosure is sealed, a self-contained recirculation pattern will develop and the temperature of the internal air will stabilize at a value determined by the heat transfer through the walls of the enclosure to the external environment.




Figure 7.4 - Convective Heat Transfer

Equation 7.2 applies also to convection systems, except that the thermal resistance, R_{th} , is more difficult to determine. R_{th} for convective systems is defined as:

$$R_{th} = \frac{1}{h \cdot A} \text{ } ^\circ\text{C/W}$$

Equation 3

where: A = Surface area of object being cooled
 h = Convective Heat Transfer Coefficient

The heat transfer coefficient, h , will have the units of $W/in^2 \text{ } ^\circ C$. Combining equations 7.2 and 7.3, we can express the rate of heat transfer, Q , as:

$$Q = h.A.\Delta T$$

Equation 4

where:

ΔT = Temperature difference between the object and the cooling fluid.

It is clear that the **convection coefficient, h** , is an important part of determining the cooling effectiveness. The convection coefficient is dependent upon a number of factors related to the properties of the cooling fluid, including:

- Thermal Conductivity, K
- Viscosity
- Density (function of temperature and altitude)
- Specific Heat
- Velocity
- Turbulent or Laminar Flow

The heat transfer physics of this convective system is a much more complex computational environment than conductive cooling since many of the effects of the above variables are non-linear and interactive. Consequently, empirical methods are often used to measure and characterize the heat transfer effectiveness of the most common fluid, air, under conditions typically encountered in electronic systems. We will present here some general comments on the relative importance of the most important factors along with expected ranges of h for various situations. More specific values of converter thermal resistance as a function of airflow (and h , if required, from equation 7.3) can be found in the datasheets for power converters.

Many of the variables are essentially locked in when the cooling fluid and the operating environment (external temperature and altitude) are selected. These variables include thermal conductivity, viscosity, density and specific heat. The remaining variables (A , ΔT , velocity and turbulence) are more controllable by the system designer. Area and air temperature are especially powerful in determining the overall heat transfer effectiveness. The surface area and the temperature of the cooling fluid have a direct and linear influence on the heat transfer as shown in equation 7.4. As a consequence, they are both very important system design parameters when attempting to achieve superior thermal performance.

The velocity of the cooling fluid is also important, but does not have a linear influence on the convection coefficient, h . The relationship is:

$$h \propto \sqrt{\text{velocity}}$$

Equation 5

Consequently, the airflow velocity will eventually have diminishing returns in terms of cooling effectiveness as shown in Figure 7.5. The physical explanation of this is that, when the airflow velocity becomes high, all of the heated air is swept away from the surface and replaced with unheated air. Once this occurs, there will be a finite rate of heat transfer depending on the unheated air temperature. This effect is also shown in Figure 7.5.

Power System Thermal Design

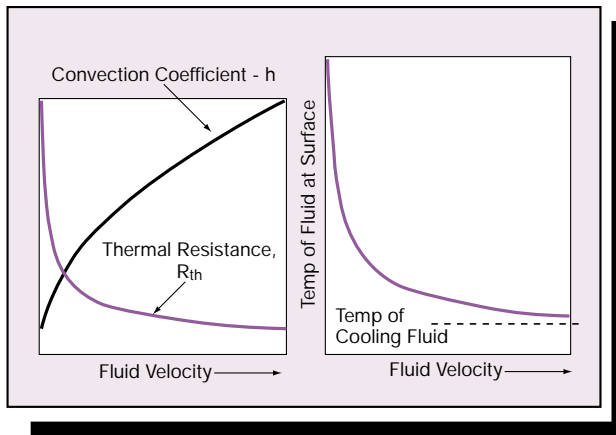


Figure 7.5 - Effect of Fluid Velocity

The amount of turbulence in the fluid flow is also important, with turbulent flow being significantly more effective in heat transfer than laminar flow. The transition between laminar flow at low flow rates to turbulent flow at higher flow rates will occur at a point determined by the properties of the cooling fluid. For air at temperatures and altitudes encountered in electronic systems, the transition will be in the range of 150 to 200 linear feet per minute (lfm). The units used to specify airflow will be discussed in a following section. Most designers of forced convection systems will strive to obtain turbulent flow due to its increased effectiveness.

Most convection cooling analysis assumes an **operating altitude** near sea level, but altitudes encountered in some applications will require adjustment due to the decreased air density. Avionics equipment is one obvious example. But even earth-borne applications can experience altitude effects. As a rough guideline, an altitude of 3,000 meters will decrease heat transfer by 20% and a 10,000 meter altitude will cut the cooling effectiveness by half.

As we have seen, natural or free convection is conceptually treated the same as forced convection, except with low laminar flow rates. The value of the convection coefficient for air will

depend on many of the above factors, but will typically be in the range of 0.003 to 0.016 for natural convection systems and in the range of 0.006 to 0.32 for forced convection systems for a flat plate. If the surface area of the object being cooled is known, the above estimates can be used to predict the range of thermal resistance from equation 7.3. For example, for a square plate two inches on a side, the area will be 4 in², and the thermal resistance will be in the range of 16 to 83 °C/W for natural convection and 0.8 to 42 °C/W for forced convection.

To summarize, the effectiveness of convective heat transfer can be enhanced by:

- Increasing the surface area
- Reducing temperature of cooling fluid
- Increasing velocity of cooling fluid
- Increasing turbulence of cooling fluid

Radiation Cooling - Heat can also be transferred by radiation from a hotter body to its cooler surroundings. In electronics systems, the "surroundings" is usually the ambient air in proximity to the component or circuit board. Heat transfer via radiation is unique in that there does not need to be any fluid in contact with the object being cooled, and therefore it is the only mechanism we have discussed that will work in a vacuum. The heat energy is transferred via electromagnetic radiation, mostly in the infrared portion of the spectrum. The amount of power transferred will be a function of the effective surface area, the relative surface emissivity of the object, and the difference between the object and its surrounding's temperatures to the 4th power. The equation for calculating power transfer via radiation is:

$$Q = A \cdot e \cdot \sigma \cdot (T_0^4 - T_s^4) \text{ W}$$

Equation 6

Where:

Q = Power Transferred in W

A = Effective Surface Area of hot object

e = Surface Emissivity

s = Stefan-Boltzmann constant = $3.66 \times 10^{-11} \text{ W / in}^2 \text{ } ^\circ\text{K}^4$

T_O = Temperature of hot object - $^\circ\text{K}$

T_S = Temperature of Surroundings - $^\circ\text{K}$

It is important to note that the effective surface area is usually less than the total surface area of the object. Surfaces that are shielded or obscured from the surroundings will not contribute - the effective surface area is essentially the "line-of-sight" area as seen by the surroundings. For a heatsink, for example, much of the area between the fins may not contribute to radiative cooling. The **surface emissivity**, e , is a factor between zero and one that defines how effective the surface is at radiating heat. A perfect black body has an emissivity of one. Some emissivity values for typical surfaces found in electronics systems are shown in Figure 7.6. Note that for paints and anodizing, the color used does not affect the emissivity, while the surface reflectivity does.

Materials	Surface Emissivity, e
Polished Silver	0.02
Polished Aluminum	0.04
Bright Tin	0.04
Rough Aluminum	0.06
Dull Nickel Plate	0.11
Stainless Steel	0.28
Steel	0.55
Oxidized Steel or Copper	0.78
Anodized Aluminum	0.80
Glossy Paint	0.89
Flat Paint	0.94

Figure 7.6 - Surface Emissivity of Electronics Materials

As an example, let's calculate the amount of power transferred by radiation from a large stud-mounted power diode in a free convection environment. The surface of the diode is at 100°C (373°K) and the heat is being radiated to the surrounding air

which is at a temperature of 50°C (323°K). The effective surface area is 0.5 in^2 . From Figure 7.6, the surface emissivity is 0.11 for nickel plating. From equation 7.6 we obtain:

$$Q = (0.5) (0.11) (3.66 \times 10^{-11})(373^4 - 323^4) = 17 \text{ mW}$$

We can see that for the typical temperatures encountered in electronics, the effect of radiative cooling is rather minimal. It is most effective when the hot object is significantly hotter than its surroundings - a condition opposite to what we usually try to achieve, since very hot components have reduced reliability. The effect will be even less when forced convection is used, since this will reduce the surface temperature of the object. For natural convection systems, radiative cooling can be a significant factor at very high altitudes where the air density is very low. It is used in the design of avionics and aerospace equipment and is the major external heat transfer mechanism in space applications. For most land-based forced convection systems, it contributes perhaps 5% at most to the total heat transfer. Due to this limited contribution, most power system designers choose to ignore the radiation heat transfer mechanism and use whatever benefit it provides as an extra design margin.

Phase Change Cooling - When a fluid changes from its liquid state to its vapor state, it absorbs a tremendous amount of energy, called its **latent heat of vaporization**. This absorption of energy occurs without any increase in the temperature of the fluid! If a hot object being cooled supplies this energy, a very efficient cooling system can result, with performance orders of magnitude superior to conduction and convection. The specific performance will depend upon the fluid used along with other factors. Some examples of electronics cooling technologies that utilize a phase change mechanism are heat pipes and evaporative spray cooling, which will be discussed later in this chapter.

Power System Thermal Design

Airflow Physics - Forced convection cooling depends upon the generation of a moving air stream. Here we will summarize the relationships that exist between airflow, air volume, pressure, and system airflow impedance along with some of the design trade-offs involved with providing the airflow.

In order to do any analysis of a forced air cooling system, the amount of air going through the system must be determined. This is referred to as the **volumetric flow rate**, and will typically be expressed in terms of **cubic feet per minute (CFM)**. Fans and blowers used in electronic systems are rated in terms of their CFM capability. Unfortunately, determining the system airflow is not as easy as looking at the rating on the fan(s). The fan rating is predicated on zero backpressure on the fan. In actual systems, there will be some restriction to free airflow that will create a pressure drop through the system and force the fan to overcome a positive pressure. The restriction of the airflow is determined by the system **airflow impedance**. The impedance will create a pressure drop, expressed in inches of Hg or mm of Hg that will increase as the airflow through the system increases. The system airflow impedance must be measured or estimated by analysis before the actual system airflow can be determined. The manufacturer of the fan or blower will publish a curve, such as the one shown in Figure 7.7, that defines the CFM rating as a function of pressure drop. If the system airflow impedance curve is plotted on the same axes, the intersection of the two curves will determine the actual airflow, in CFM, of the system. In the example shown in the figure, if the fan rating at no backpressure were 100 CFM, the actual system airflow would be perhaps 65 CFM.




Figure 7.7 - Determination of System Volumetric Flow Rate

The general expression for the maximum amount of heat that can be removed from a system by forced air convection is:

$$Q = sh \cdot w \cdot \text{density} \cdot \text{CFM} \cdot \Delta T \quad W$$

Equation 7

where:

Q = Heat removed in Watts

sh = Specific heat of air

w = Energy capacity of water = 1 calorie/gram °C

density = Density of air

CFM = Airflow rate

ΔT = Temperature rise from input to output

Note from the above equation that, for equivalent heat removal, the CFM must increase if the air density decreases - i.e. at high altitude.

If we assume that the system is at sea level, equation 7.7 can be simplified by inserting the specific heat and density of air at one atmosphere of pressure, and converting to the most commonly used units, yielding:

$$Q = 0.615 \cdot \text{CFM} \cdot \Delta T \quad W$$

Equation 8

where:

Q = Heat removed in Watts

CFM = Airflow rate in CFM

ΔT = Temperature rise from input to output in °C

For example, if the system airflow is 100 CFM and the allowable temperature rise in the cooling air is 20 °C from inlet to exhaust, the maximum cooling capability of the system will be 1.23 kW.

While the CFM will determine the maximum possible heat that can be removed from the system, the heat must be transferred from the components or sub-assemblies to the moving air stream before it can be removed from the system. This will be determined by the rate of the airflow right next to the component being cooled, a variable that is referred to as the **linear flow rate**. It is essentially the speed of the air as seen by the component, which will be influenced by the amount of cross-sectional area that the air is flowing through. The linear flow rate, also known as **air velocity**, will be expressed in units of feet per minute in English units or m/s in metric units. The ft/min unit is often referred to as **linear feet per minute**, or **lfm**. A linear flow rate in lfm can be converted to m/s by multiplying by 0.005. In other words, 1 m/s is about 200 lfm. Linear flow rate is an important system parameter because it will be used to calculate the heat removal capabilities of components and heatsinks as described in following sections of this chapter.

The linear flow rate for a given system is calculated by dividing the system volumetric flow rate by the effective cross-sectional area of the airflow. Using English units, this results in:

$$\text{lfm} = \frac{\text{CFM}}{A} \frac{\text{ft}}{\text{sec}}$$

Equation 9

where:

lfm = average air velocity

CFM = system volumetric airflow

A = effective cross-sectional area of system

See Figure 7.8 for an example of how the air velocity is calculated. Here the system is a rectangular card cage containing 12 identical cards with the air flowing perpendicular to the plane of the drawing. The effective airflow area is the difference between the enclosure area and the area taken by the edges of the cards. The enclosure is 6" by 12" giving an enclosure cross-sectional area of 0.5 ft². The end of each card is 0.5" by 5" for a cross-sectional area of 0.017 ft² for each card and total of 0.21 ft² for the 12 cards. The effective system cross-sectional area is the difference between the enclosure area and the card area, or 0.29 ft².




Figure 7.8 - Determination of System Linear Flow Rate

Assuming that the system volumetric flow rate is 100 CFM, we can now calculate the air velocity, in lfm, from equation 7.9:

$$\text{lfm} = 100 \text{ ft}^3/\text{min} / 0.29 \text{ ft}^2 = 345 \text{ ft}/\text{min}.$$

This is an average air velocity, since the absolute velocity at each location in the system will vary somewhat as a function of the distribution of the airflow. The average velocity assumes

Power System Thermal Design

that the airflow is equally distributed through all of the open area of the system, which would only be possible with very detailed attention to air plenum design and baffle placement.

Finally, it is useful to understand the ramifications of the system CFM on the fan design and other system variables. Figure 7.9 summarizes the effects of changing the volumetric airflow upon other variables. The data in the figure is normalized to the ratio of CFM change. For example, if the CFM is doubled, the system pressure drop will increase by a factor of four, the fan input power by a factor of eight, and the acoustic noise by approximately 15 dB.

Technique	Effect of CFM Change $\Delta\text{CFM} = \text{CFM2}/\text{CFM1}$
Fan Speed	Speed $\propto \Delta\text{CFM}$
Fan Input Power	Power $\propto \Delta\text{CFM}^3$
System Pressure Drop	Power $\propto \Delta\text{CFM}^2$
Acoustic Noise Increase	Noise $\propto 50 \text{ Log}_{10}\Delta\text{CFM}$ dBA

Figure 7.9 - Effects of CFM

7.3 The System Thermal Environment

Most electronics systems, other than military and automotive applications, operate in fairly similar environments - the same types of environments that are inhabitable by humans. There are, consequently, some generally accepted ranges or limits for variables such as external air temperatures and allowable temperature rises inside the equipment. It is also useful to define the various locations in the system where temperature measurements are made.

Figure 7.10 shows how we will define the most commonly used temperature measurement locations in the system. The **external temperature** is the temperature of the room or area in

which the equipment is located. This will be the starting point for many of the thermal calculations needed for the power system design. **Ambient temperature** is always measured in the air rather than on a solid surface. It is the local air temperature next to the component or sub-assembly being evaluated. Note that there can be several different ambient temperatures of interest in one system. In Figure 7.10, for example, we have shown ambient temperatures for the input air cooling the power supply heatsink, the circuit board ambient behind the power supply and the circuit board ambient in an area free of any heating effects of the power supply. **Surface temperature** is a very useful measurement, since it forms the basis of conductive heat transfer calculations. More and more power converters utilize a surface temperature measurement as the means of verifying thermal performance. It too can be relevant at more than location in the system, the power supply baseplate temperature and the surface temperature of PCBs, for example. The final major classification of temperatures is the **junction temperature**. This is the actual operating temperature of the interior of a component - the junction in the case of a semiconductor. It is also a key parameter for components such as capacitors and magnetics. It is important since it will determine the failure rate and ultimate reliability of the component. The junction temperature will be influenced by the surface temperature and the ambient temperature.

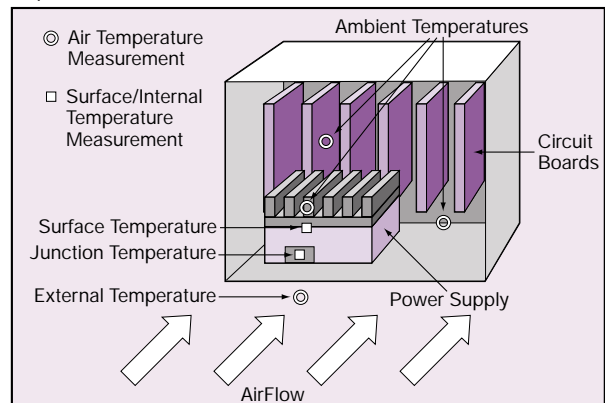


Figure 7.10 - Definition of Temperatures

Perhaps the most severe environment encountered in telecom and datacom equipment is that experienced by remote telecom equipment that must operate in buildings without air conditioning. The external temperature can be up to 60 °C in such situations, and the equipment is designed for free convection to eliminate the reliability and service issues associated with fans. Allowing for an internal temperature rise of up to 20 to 30 °C inside the equipment results in local ambients above 70 - 80 °C. As a consequence, these systems use large enclosures with plentiful space surrounding the circuit boards, resulting in low overall power density.

If higher density is desired, forced-air cooling is almost always used. With forced air, significantly more heat can be removed from the enclosure with the same external temperature and less temperature rise, as determined by equation 7.8. It is important to remember that free or natural convection does not imply the total lack of airflow. As was pointed out earlier, natural convective air currents will be set up. The airflow generated by this mechanism can range from 0.1 to 0.3 m/s, or 20 to 60 lfm. Most forced-air systems are operated between 100 and 1000 lfm. As was seen in Figure 7.5, increased air velocity begins to have diminishing returns after a certain point, helping to define the upper air velocity limit along with required fan power and acoustic noise. The resulting useful range of air velocity for typical systems is shown in Figure 7.11.

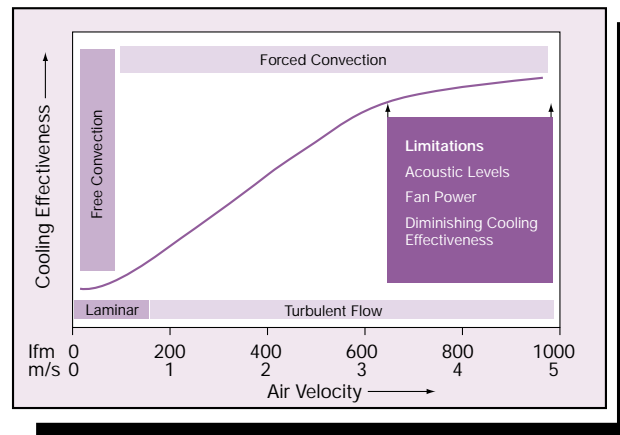


Figure 7.11 - Airflow Rates in Typical Electronics Systems

7.4 The System Thermal Design Techniques

The power system designer can significantly enhance the reliability of the end product by paying attention to the thermal design. In this section, we will provide some general guidelines for successful thermal design as well as more specific suggestions for enhancing both conduction and convection heat transfer mechanisms. We will also briefly summarize the status of less conventional cooling techniques such as heat pipes and thermoelectric cooling.

General Design Guidelines - The first and foremost recommendation for successful thermal design is to minimize the amount of heat that needs to be removed! This is best accomplished by using efficient power converters. The latest generation DC/DC converters utilizing synchronous rectification will drastically reduce the amount of power dissipated in the converter, resulting in reduced demands on the system's cooling mechanisms. An increase in converter efficiency from 85% to 90%, for example, will reduce the dissipated power by 37%! The system designer should take advantage of any efficiency

Power System Thermal Design

increases that are available. Even if there is a slight price premium for such converters, it is more than compensated by reduced need for system cooling hardware and enhanced reliability of not only the converter, but of the entire system due to lowered operating temperature.

The next most effective overall strategy is to minimize the ambient temperatures internal to the equipment. This goal is achieved by paying attention to several design criteria such as a conservative packaging design that does not push the envelope on overall system power density, maintaining clear and clean airflow paths for either free or forced convection cooling, eliminating dead-air pockets, and providing adequate air velocity. Using multiple board-mounted converters as opposed to a larger centralized converter can also help achieve a more balanced thermal environment in the system and eliminate areas of heat concentration and resulting hot spots. The recent trend towards DPA should be seen as a positive step towards more manageable and predictable system thermal designs.

Converter Thermal Derating - The power derating curve published in datasheets by converter manufacturers is a useful tool for determining if the converter is suitable for usage in a given thermal environment. These curves plot the permissible power output vs. the converter's ambient operating temperature in the system. The curve is predicated upon not exceeding an allowable maximum temperature internal to the converter. Figure 7.12 shows an example of a thermal derating curve. In this example, the converter is rated to operate up to a maximum ambient temperature of 100 °C. In a natural convection environment, with an ambient temperature of 70 °C, and maximum output power the internal converter temperature will be at its allowable maximum. At operating ambients in excess of 70 °C, the converter's output power will need to be reduced to maintain the same maximum internal temperature. The derating curve defines the required extent of this reduction. For example, if the operating ambient is 80 °C, the curve shows that

the converter output power must be reduced to 75% of maximum. If the converter full power rating was 50 W, the derated value would then be 37.5 W at 80°C ambient.

Maintaining full power output up to 70 °C is actually quite good performance for a converter. The operating temperature range can be extended even further by using forced-air cooling. Derating curves for forced air cooling at various air velocities are often plotted in the same graph. In the example, the curve for 1.5 m/s airflow is given. With 1.5 m/s airflow, this converter will operate at full power output up to an ambient temperature of 90 °C. Thus, it would not need to be derated for this system's 80 °C operating ambient.

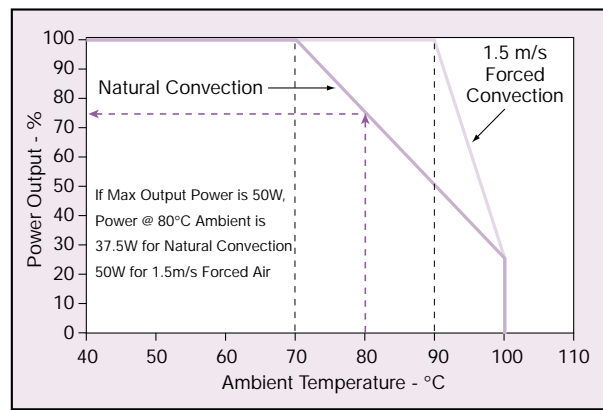


Figure 7.12 - Using Converter Thermal Derating Curve

Selection of Cooling Method - There are no hard and fast rules for selecting the cooling mechanism(s) for the system. In many cases, they will be pre-determined by the prior selection of load circuit packages, the system high level packaging design (card and board sizes, etc.) and the power architecture and converter choices. Figure 7.13 presents a general depiction of the most commonly used heat transfer technologies as a function of the average areal power density in the system. Free convection is useful up to 1 W/in² and forced convection up to 10 W/in². At power densities above 10 W/in², the need for some form of higher performance cooling technology such as phase change or cold plates is indicated. There are some caveats to keep in

mind when using these guidelines. First, the upper limits for each technology are just that, and most conservatively designed systems will be operating well below the maximum. Also, most systems will have widely varying power densities in different parts of the system. Low power PCBs, for example, could be well under 0.1 W/in^2 , while the heatsink on a high power converter is at a much higher areal power density. Also note that the areal power density is based on effective surface area, so that heatsinks will have a larger area than their PCB board footprint. Also keep in mind that almost all systems actually use a combination of heat transfer mechanisms. Conduction is very important at the component and circuit board level even in systems that rely on natural or forced air convection at the higher assembly level.




Figure 7.13 - Practical Range of Heat Transfer Methods

Conduction Suggestions - All systems, even if convectively cooled, will utilize conduction for heat transfer at the component and PCB level. There are design practices that can optimize this important heat transfer path. First of all, understand the converter thermal design. Refer to the datasheets and application notes provided by the converter supplier to determine the heat removal path(s) designed into the converter. With the trend toward non-encapsulated power converters, more and more designs are depending upon some heat conduction through the pins to the PCB in addition to the normal convective

thermal interfaces. Sometimes certain converter pins will be identified as significant heat transfer paths. If so, make sure that they are connected to heavy PCB traces. In fact, using maximal PCB traces under a power converter is always a good design practice. Figure 7.14 shows some general layout suggestions that will accomplish these objectives. Note that it is almost identical to Figure 5.15, which suggested layouts that would minimize EMI problems. This is indeed good news, as the same layout will benefit the system in regard to both EMI and thermal design.

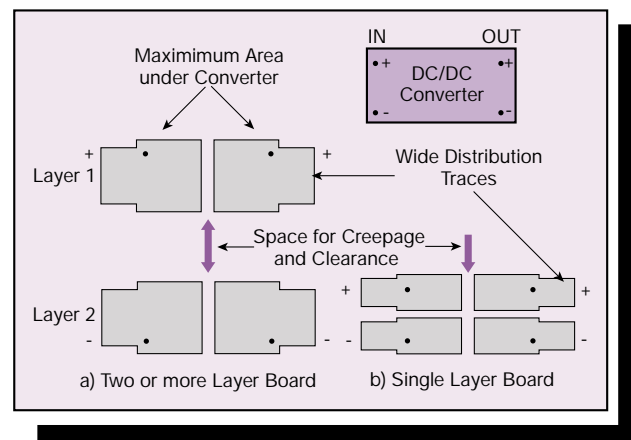


Figure 7.14 - PCB Layout for Increased Conductive Heat Transfer

The use of maximal PCB trace area under the converter and for the input and output power distribution is important because the PCB contains copper, which has excellent thermal conductivity. As shown in Figure 7.2, the thermal conductivity of copper is roughly 1000 times greater than that of glass epoxy. Consequently, the PCB will have excellent thermal conductivity in the x and y (horizontal) planes if it contains sufficient copper. This conductivity can be used to advantage to remove heat from the power converter and spread it out into the PCB where it can be effectively removed by convection mechanisms. The usage of 2 oz or 3 oz copper for the power planes will enhance this heat transfer path even further.

The next most common opportunity for enhancing thermal

Power System Thermal Design

conductivity is with the assembly of heatsinks to power converters. Even though the heat removal from the heatsink is via convection, the path from the converter to the heatsink depends upon conduction. No surface is perfectly flat or smooth, and there will be some air gaps between the converter and the heatsink if the heatsink is directly attached. Air has a thermal conductivity roughly 4 orders of magnitude lower than metals, so any air gaps will significantly reduce the heat transfer performance of the interface. This problem is attacked by placing a heat transfer pad between the converter and the heatsink to fill the air gaps and replace them with a material with a higher conductivity than air. Thermal pads, which are widely available commercially, will offer at least a ten times improvement in conductivity relative to air. In addition to better thermal performance, their usage will drastically reduce the production variation from unit to unit that would result from dependence on the flatness characteristics of the converter and heatsink.

Free Convection Suggestions - The design of free convection systems avoids the complications of fans and their control systems, but still demands attention to airflow control. The general mindset when doing a free convection design must be to be conservative and keep the overall system power density low. Free convection designs run into trouble when either too much dissipation is attempted in a given volume or when the lack of clean airflow paths result in localized hot spots. The best free convection designs feature open layouts with clean paths for airflow around components with any significant power dissipation. The cabinets should be well vented at both the top and bottom so that external air can flow freely through the enclosure.

Components or sub-assemblies with high power dissipation should be located near the top of the enclosure. This will prevent the heated air from these components from rising the ambient temperature of all the circuitry above it. This hot air

would result in lower average component temperatures and higher reliability. Another variable that the system designer can use to influence the performance of a natural convection system is the orientation of the circuit boards. It is very important to place the circuit boards so that the plane of the boards is vertical. If the board plane were horizontal, it would block the formation of air convection currents and substantially reduce the cooling capability of the system. For a given circuit board area, the thermal performance will be better if the longer dimension of the board is oriented horizontally rather than vertically, as seen in Figure 7.15. The reason is that there is less "pre heating" of the air in the horizontal configuration. This can be a useful technique to keep in mind, if your system has flexibility with regard to board packaging. Figure 7.16 summarizes some of the "dos and don'ts" of natural convection system design.

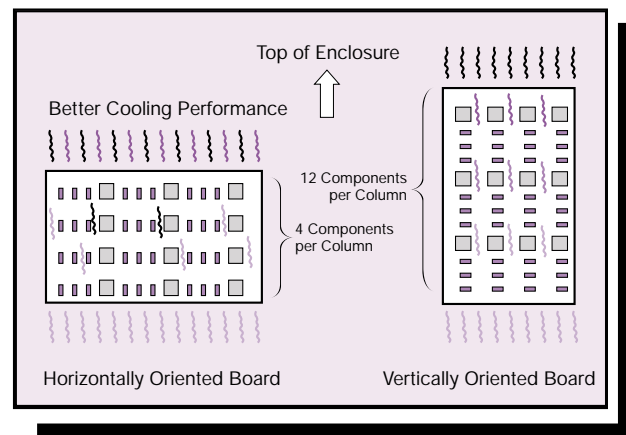


Figure 7.15 - Effect of PCB Orientation




Figure 7.16 - Some Natural Convection System Design Guidelines

Forced Convection Suggestions - Many elements of forced convection design are similar to those of free convection. The system layout must be configured so that there are no "dead air" spaces or areas of restricted airflow. This is accomplished by using air plenums and ducts to route the cooling air over all of the system components with significant power dissipation. Airflow can be provided by either placing the fan at the air input to the system and pushing air through the system, or by placing the fan by the exhaust and using it to pull air through the system. Either of these approaches will work. We discussed earlier how turbulence in the air stream helps to increase the effectiveness of convective cooling. In general, the "pushing" approach will provide more turbulence than the "pulling" alternative. In some systems, this could be a deciding factor in choosing the location of the fans. Decisions must also be made about the degree of centralization for the system's fans. Using one or two larger centralized fans will require additional mechanical design for the ducting and airflow routing mechanisms. The alternative is to use a more distributed approach, with dedicated small fans that are used for "spot cooling" of individual components with the highest power dissipation. Many systems will use a combination of these approaches to obtain the correct balance between cooling system cost and design complexity.

Minimizing the system's airflow impedance will have big benefits in terms of enhancing its overall thermal performance. As we saw in Figure 7.7, reducing the airflow impedance will increase the CFM through the system for any given fan. This increase in CFM will then increase the linear flow rate across the components being cooled as calculated in equation 7.9. Using physically large inlet and outlet air vents and efficient air filter designs will reduce the system impedance. Of course the internal routing of the airflow is also important, with attention to eliminating any significant restrictions. In addition to increasing airflow, these design details will also reduce the acoustic noise level of the cooling system. The fan selection itself can also help to reduce the acoustic noise. For a given CFM rating, fan noise will be reduced by either:

- Using a fan with a smaller diameter blade (lower blade tip speed)
- Using a lower speed fan

In addition to improved acoustic performance, lower fan speed also reduces the input power to the fan. This, along with the present availability of low cost **brushless DC fan** motors, has increased the popularity of using **variable speed fans**. The speed of a brushless DC motor can be varied by changing the input DC voltage to the motor. A 12V fan, for example, will run a maximum speed with an input voltage of 12V, but can be easily adjusted to lower speeds by reducing the input voltage. Running the fan at a lower speed (assuming that it meets the cooling demand) has several advantages:

- Lower acoustic noise
- Reduced fan input power
- Less dirt accumulation in the air filters
- Increased fan reliability due to less bearing wear

The brushless design of the fan also provides enhanced reliability due to the elimination of the brushes, which are

Power System Thermal Design

subject to wearout mechanisms. Brushes can also be a source of EMI, so the brushless design is superior in that regard also. Brushless DC fans are widely available with nominal input voltage ratings of 12V, 24V and 48V.

There are two basic control approaches for regulating the speed of a variable speed fan, as shown in Figure 7.17. With **open loop control**, the temperature of the inlet air to the equipment is sensed and used as an indicator of how fast to run the fan. At high external temperatures, the fan would be run at full speed and at reduced speed at lower values of external temperature. The control scheme is simple in that there is no feedback from any thermal variables inside the equipment. The disadvantage of this simplicity is that the control system cannot compensate for events that occur within the system such as a clogged air filter or changes in the power dissipation with load changes.

Closed loop control is much more versatile. In a closed loop variable speed fan control system, the temperature sensing is done either at the equipment air exhaust or at some critical component internal to the system. This sensed temperature is then compared to a reference and the fan speed varied to force the sensed temperature to the desired value. This approach has several advantages over the open loop design:

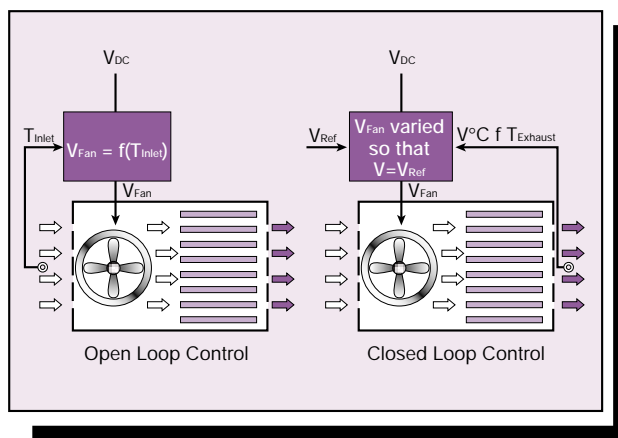


Figure 7.17 - Variable Speed Fan Control Techniques

- Compensates if thermal load is added to or removed from the system
- Compensates for changes in altitude
- Compensates as air filters become dirty
- Compensates for partial fan failures (multi-fan systems)

The only disadvantage of closed loop control is that, like all feedback systems, it will require some attention to stability considerations to insure that the thermal response time constants are such that there will be no undesired oscillations in the system temperatures. This is not normally a problem because the control loop will have a much faster response time than the heat buildup in the system.

The fan control circuitry itself is fairly straightforward, and can be packaged in a number of ways. There are off-the-shelf fan control boards available that can be wired to a brushless DC fan and to a remotely located temperature sensor. If you prefer a discrete approach, several of the analog IC manufacturers now make fan control ICs that can be added to your system control boards. This will require more design effort, but can result in a very low cost solution. There are also fans available with built-in control circuitry. This is a very easy solution, but may provide less flexibility and adjustability than the other alternatives.

Many brushless DC fans provide a "tach" output that can be used by the system to verify that the fan is running. The tach output will be a series of pulses as the fan rotates. It is sometimes tempting to use this signal as a system "temperature fail" indicator. If the functionality of your system is especially critical, it is recommended that you also include an actual temperature sensor for this purpose. Using the tach signal only verifies that the fan is running (and its speed). If there is a serious fault such as a completely clogged air filter or an object placed over the inlet or exhaust venting, the fan will run at full speed but the temperature inside the equipment can reach unacceptable

levels. The tach signals from the fan are, however, helpful in diagnosing fan failures and in generating alarms before the system temperatures actually reach dangerous levels.

Heatsinks - Heatsinks have long been associated with power converters, and not always in entirely favorable contexts. They tended to be large, heavy and required manual assembly processes that added cost during manufacturing. Heatsinks do serve a valuable function, however. They allow for the spreading of heat from components with high power dissipation so that it can be assimilated by the system's convective cooling mechanisms. The good news is that heatsinks are needed less often now than in the past. The recent trends toward "open-frame" DC/DC converters, higher conversion efficiencies and usage of lower power board-mounted converters have all contributed to eliminating the need for a heatsink in all but high power converters. In fact, there are now products available from Artesyn (EXB50 series, for example) that will provide a board-mounted 50W output converter that will operate at full load at 70°C ambient in a natural convection environment without any heatsink. In spite of these advances, there are still needs for heatsinks in high power applications, so we will review here their basic properties and provide some tips for their successful application. An example of heatsink selection for usage with an Artesyn DC/DC converter is given in the Thermal Design Examples section later in this chapter.

Aluminum is the dominant material used in heatsinks. It has achieved its popularity because it is lightweight, inexpensive, machinable with standard tooling and is easily extruded due to its ductility. Standard heatsinks are extruded in long sections and then cut to different lengths so that different part numbers are available for different cooling needs. The extrusion and cutting processes result in fins that run in either the short dimension or the long dimension of the heatsink. The system airflow has to be in the same direction as the heatsink fins if the heatsink is expected to perform its function. Consequently,

some converter manufacturers will offer or recommend two different heatsinks for the same converter, with fins running in different directions to accommodate the system airflow direction. This option is referred to as "**heatsink orientation**" or "**fin orientation**". For natural convection systems, the heatsink fins should be oriented in a vertical direction so that the convective air currents from the bottom to the top of the system will result in an air velocity in the same direction as the heatsink fins. There are some heatsinks designed to eliminate the need for two different fin directions. These designs use either protruding "pins" instead of fins or mechanically machined slots into the extruded fins to provide for airflow paths in the orthogonal direction to the fins. These heatsinks are referred to as "**pin type**" or "**slotted**" heatsinks. They are more expensive than plain extrusions and also offer reduced performance compared with a finned heatsink in its preferred airflow orientation. There are also "**serrated fin**" heatsinks available. These will have small undulations on the surface of the fins that marginally increase their surface area. The increased surface area will not have a corresponding reduction in their thermal resistance in most applications, however. This is because of the air boundary layer that forms along the fin results in a "dead air" space that is deeper than the serration height for the airflow velocities encountered in most electronic systems.

The aluminum extrusion process has limitations with regard to the range of acceptable fin geometries. The fins need to have a minimum thickness and minimum separation between fins as well as a maximum height to be extrudable by a cost-effective manufacturing process. Because of this, a physically small heatsink can only contain a fixed number of fins with a limited height, which will limit its effective surface area and cooling capacity to a maximum value. Heatsink manufacturers have circumvented this limitation by using the aluminum extrusion for only the bottom portion of the heatsink and using thinner sheet materials such as steel and brass for the upper finned portion. The fins can be either individual pieces of these materials or one

Power System Thermal Design

piece that is folded accordion fashion into a series of pleated fins. These fins are then attached to the aluminum bottom section either mechanically or by means of a brazing process. These types of heatsinks are referred to as "**bonded fin**" or "**folded fin**". They offer increased surface area for a given footprint. Because of their reduced fin spacing, they need well-focused airflows of reasonable velocities to force air through them, and are intended for high-performance applications. They will also be more expensive than conventional extrusions, and are not frequently used for power converter applications.

The key to successful design with heatsinks is to know the expected airflow in the system. You cannot rely on an average airflow value derived from the fan CFM and the total system airflow area. It will be much more useful to measure the actual linear airflow in lfm or m/s in the area where the heatsink will be mounted. Then you can use the heatsink manufacturer's graphs of thermal resistance vs. airflow to make an intelligent selection of the particular model you need. Don't overlook the possibility of localized sources of airflow. Mounting a small fan so that it blows directly on the heatsink can often be much easier than depending upon ducting the system's air supply and will also result in higher airflow velocities. In fact, there are now heatsink/fan combinations available from heatsink suppliers that integrate both together into one assembly. These are sometimes called "**active heatsinks**".

If you are using a heatsink, pay attention to the interface between it and the power converter. This interface represents a series thermal impedance, and is composed of a material with much worse conductivity than either the heatsink or the power converter. Without any thermal interface material, the gaps in the interface will be filled with air, and you will be dependent upon the flatness and smoothness of the heatsink and converter surfaces to maintain thermal contact. The usage of a thermal interface material will improve the situation, but flatness and smoothness of the surfaces are still very important. As we saw

in Figure 7.2, even the best thermal interface materials, while better conductors than air, are orders of magnitude worse than metals, so the intent should be to minimize the required thickness of interface material. The larger the surface irregularities in the mated devices, the thicker the interface material must be, and the higher the overall thermal resistance from the converter to the outside world.

You have a wide choice of interface materials. Thermal grease, which was one of the original solutions, can still be a viable choice. It actually has better thermal performance than some of the "higher technology" more expensive solutions. It is growing out of favor, however, due to manufacturing concerns. It must be applied in carefully measured quantities, usually manually, and is messy when it is forced out of the interface when the heatsink is torqued down. It also has no adhesive or retention properties. Available alternatives to thermal grease include:

- Thermal Adhesives
- Pre-applied thermal materials (by heatsink manufacturer)
- Double-sided thermal tape
- Thermal Pads
- Coated Aluminum Foils
- Paraffin-based materials that melt when the interface heats up

The thermal conductivity of all the above solutions will be in the same ballpark. Therefore, the selection of the best one for a given application should be made based upon other considerations such as the manufacturing process used or the need for electrical isolation. Better thermal performance can be achieved if there is no need for electrical isolation, due to the generally smaller thickness. Whatever material or technique you use, make sure to use the recommended torque and assembly instructions supplied by the manufacturer. The suppliers of thermal interface materials also offer very useful technical

literature that will help you make an appropriate selection.

Alternative Cooling Technologies - While some form of conduction and convection are used in almost all systems, newer technologies are sometimes used in selected applications. Although they will not have widespread application in most systems, it is good to have an awareness of their general capabilities and properties. As was highlighted in the primer section, heat transfer by means of phase change of the cooling fluid can be especially effective. The two most widespread examples of phase change cooling for electronic equipment are heat pipes and evaporative spray cooling.

Heat pipes are low pressure sealed systems containing a fluid and linking a cooled object to a thermal sink. The fluid is boiled (evaporated) at the object end and the resulting vapor travels through the system to the thermal sink at the other end, where the vapor is condensed. This can transport a significant amount of heat from the object to the heatsink without any temperature rise of the fluid. The pressure inside the system is lower than atmospheric and adjusted so that the fluid's boiling point is in the vicinity of the operating temperature of the object being cooled. Thus even a fluid such as water is capable of being used at temperatures well below its boiling point at sea level. Other fluids in common usage include ammonia, methanol and acetone. As the name implies, heat pipes are usually configured in the form of a pipe or tube, with diameters between 3 and 15 mm. So that the heat pipe will work in any physical orientation, the interior is filled with a porous wicking material that will return condensed fluid back to the source end by means of capillary action. See Figure 7.18 for an illustration of the construction.




Figure 7.18 - Heat Pipe Technology

The heat pipe tubing can be bent into different shapes which gives the advantage of being able to snake it through the system to a remotely located heatsink. As such, it is a very effective means of cooling hot spots such as high power processor chips. Heat pipes have been used to cool chips in fault-tolerant super-mini computers and are now finding frequent usage in laptop computers by moving the heat from the processor chip out to the external enclosure where it can be dissipated to the air. They are very light in weight, low in volume and highly reliable due to the lack of moving parts. They also have the advantage of generating no audible noise. They are presently more expensive than conventional heat transfer methods, but the price is on a downward trend. The other disadvantage is that once they are removing their maximum amount of heat, they become "saturated" so that additional heat input at the source end will result in a rapid temperature rise of the object being cooled. In other words, they become non-linear, unlike conventional conduction and convection systems, which will show a linear increase in temperature as the cooling demand increases. This characteristic requires that the system designer have a very high level of knowledge about the maximum power to be dissipated and the saturation characteristics of the heat pipe.

Power System Thermal Design

Heat pipes have a unique characteristic in terms of their apparent thermal conductivity. To a first approximation, the performance of a heat pipe will be independent of its length. That is, if a pipe of length L can transport Q watts from an object at temperature T_1 to a sink at temperature T_2 , the same type of pipe with length $2L$ will have the same heat carrying capacity given that the object and sink temperatures are the same. For a solid cylindrical metal conductor of the same dimensions, the longer conductor would have a larger temperature rise for the same amount of power, Q , since the thermal conductivity of the metal (per unit length) is a constant. The net effect is that, for the heat pipe, longer devices will look as if they have higher thermal conductivity. Heat pipes will show apparent thermal conductivities between one (short pipes) and four (longest pipes) orders of magnitudes higher than copper!

The same phase-change principle used in the heat pipe can be used more directly with evaporative spray cooling. Here, a high dielectric strength cooling liquid is sprayed directly onto the chip or device being cooled. The liquid is then vaporized and condensed as part of a closed-loop system. This approach offers extremely efficient heat removal, but is more expensive and complex than a heat pipe due to the need for specialized system packaging, atomizers, and fluid containment. Evaporative spray cooling will be used mostly for very high power processor systems or for high power equipment in electrical generation facilities and sub-stations.

Another cooling method that utilizes refrigeration is **liquid cooling**, or **cold plate cooling**. With this approach, the devices being cooled are attached to metal plates. Internal to these plates are cooling coils containing a circulating cooled liquid. The liquid flow is generated by a closed-loop refrigeration system and pumps, and the heat removed is dissipated remotely near the refrigeration system. This results in a very high performance cooling system. The interface at the source end is a high conductivity metal plate. This makes it easy to design

predictable cooling paths for even high power loads. Liquid cooling is used often in high-end mainframe processors and in military electronics. The disadvantages include expense, the need for a large refrigeration and pumping system, and reduced overall reliability due to the cooling system complexity.

Systems designed to operate with very low junction temperatures (below $-40\text{ }^{\circ}\text{C}$) use **cryogenic cooling** for heat removal. This is done to achieve ultimate semiconductor performance and minimal switching speeds. Refrigeration systems capable of cooling any reasonable heat load to this level are extremely expensive and also very inefficient. It can require over 50 W of input energy to the cooling system to remove each watt of heat from the device being cooled.

The final technique we will mention is **thermoelectric module** or **TEM** cooling. This is a solid state cooling technology that utilizes the Peltier effect. The semiconductor junction formed will generate a voltage (similar to a thermocouple) when heated. Conversely, if a DC voltage is applied to the junction, one end of the semiconductors will be cooled and the other heated. Reversing the direction of the applied voltage will reverse the direction of the temperature change. If several of these junctions are connected in series electrically, while maintaining a parallel connection of the hot and cold sides, a device can be fabricated capable of removing significant heat (tens of watts is very practical). Figure 7.19 shows the typical construction of a TEM. The TEM has several advantages. It is very small, light and reliable as well as inexpensive. It has no moving parts and generates no audible or electrical noise. Controlling the input voltage and current can easily regulate the degree of cooling. It can even be easily changed to a heater by just reversing the polarity of the input voltage. It also can cool a circuit to below the local ambient temperature, something that cannot be done by conduction or convection. It does have some disadvantages. Due to its construction, the sink for the removed heat is physically very close to the object being cooled. This will

require some other form of cooling on the sink in most applications. It is also inefficient - typically 30% - so the removal of one watt of power will require the expenditure of three watts or so of electrical energy. These characteristics make it ideal for localized spot cooling, but not for the main cooling technology of an entire system. The characteristics of the alternative cooling techniques we have discussed are summarized in Figure 7.20.




Figure 7.19 - Thermoelectric Module

7.4 Thermal Measurement and Analysis Tool

Following the design guidelines presented above will require some degree of knowledge about your specific system application, such as ambient temperatures and linear air velocities internal to the equipment. This information can be acquired by either measurement of a prototype system or by computer simulation and analysis or a combination of the two techniques. There are tools available for both measurement and simulation that will make this task much easier and result in reliable data.

Thermocouples are still a good way to capture temperature data within the system. They can be attached to the "hot spots" identified by the manufacturer of open frame power converters in order to determine the actual converter operating temperature in the system application. This information is vital for insuring that the converter is operating within its ratings and in predicting the converter reliability in the specific system. It is also

Technology	Heat Transfer Mechanism	Advantages	Disadvantages
Heat	Phase Change	Very Efficient Sink Located Remotely Small, Light Weight High Reliability No Elec/Acoustic Noise	\$ Non-Linear at Capacity
Evaporative	Phase Change	Effective High Power Cooling	\$\$
TEM Easy Feedback Control Small, Lightweight No Elec/Acoustic Noise	Conduction	Bi-Directional High Reliability Can Cool to below Ambient	Very Low Efficiency Sink Close to Source
Cryogenic	Refrigeration	Low Device Temperature	\$\$\$ Extremely Inefficient High Complexity/Reliability
Cold Plate	Conduction	Effective High Power Cooling Predictable Conductive Heat Path Remote Heat Sink	\$\$ Fluids, Plumbing, Pumps Size/Weight of Cooling Equipment

Figure 7.20 - Characteristics of alternative Cooling Technologies

Power System Thermal Design

recommended that thermocouples be used to measure the temperatures of heatsinks used with power converters and the ambient air temperatures at the inlet to power converters or their heatsinks. The easiest way to measure airflow is with either a hand-held **moving-vane anemometer** or a **hot-wire anemometer**. Either of these methods will allow you to "probe" for the absolute airflow value at different points in the system such as at the inlet to heatsinks and across board-mounted power converters. With the data collected on air velocity, ambient air temperatures and converter/heatsink surface temperatures, you will be able to make confident predictions about the thermal performance and reliability of the complete power system.

Thermal simulation and analysis should also be given consideration as a design tool. With PCs becoming more powerful, there is now very effective software available that can do three-dimensional thermal analysis in a cost-effective fashion. High-end workstations and expensive software are no longer required. The latest PC-based Computational Fluid Dynamics (CFD) software will allow you to model your system before the prototyping stage and get a head start on predicting the expected internal temperatures and airflows. The colored graphical outputs of these programs are a powerful tool to help with visualizing the airflow patterns and temperature gradients within your system.

7.5 Thermal Aspects of Power Converter Selection

Now that you have given some thought to the overall thermal design of the power system, you will need to select the actual power converters. This selection can make or break the overall thermal design. You will want to select efficient converters with a conservative internal design philosophy in order to arrive at a

robust and reliable system. In this section we will discuss some of the converter thermal design concepts that you should be aware of when making your decision on a converter supplier.

First of all, refer to the converter manufacturer's literature to determine the assumed cooling environment for the device. Some converters are designed for primarily conduction cooling, others for free convection and still others for only forced air convection. Many of the latest converters will utilize a combination of conduction and convection, spreading some heat out into the PCB for removal and some heat directly to the ambient air. You should, of course, select a converter whose cooling approach matches that of your system. All of Artesyn's board-mounted DC/DC converters rated at 50W or less are very suitable for either free or forced convection environments with little or no derating. Higher power versions of the BXB series will generally need forced convection cooling if you do not want to significantly reduce their power output capability.

The most important converter specifications from a thermal design standpoint are the converter's **efficiency**, the **maximum converter operating temperature** - usually a component surface or converter baseplate temperature - and the **thermal resistance** from the converter to ambient air. These three numbers will determine the amount of output power you can extract from the converter in a given system thermal environment or the thermal environment (ambient temperature and airflow) that will be required to obtain the desired amount of power output. The efficiency rating is used to calculate the actual amount of power dissipated by the converter - power that will need to be removed by the cooling system. Higher values of efficiency and maximum converter temperature and lower values of thermal resistance will result in better thermal performance. Examples of making the appropriate calculations are given in the last section of this chapter.

It is also important to select converters that were designed with

conservative **thermal derating** guidelines, with maximum component temperatures at full converter load held well below the maximum rating for the component for high long-term reliability. The types and **quality levels** of the internal components will also influence their ability to operate at elevated temperatures. Unfortunately, this type of information does not usually appear in datasheets or even in most application notes. You will be depending, somewhat, on the reputation and credibility of the supplier and on their willingness to share this type of thermal design information. Artesyn will share this type of thermal derating information, upon request, so that you can better understand the soundness of the design and develop confidence in the expected reliability predictions.

One major recent change in the newer versions of DC/DC converters has been to use an "**open-frame**" type of design rather than a totally encapsulated package. There are many benefits to this trend, including cost, packaging density, manufacturing process compatibility, and removal of failure modes. But one of the major advantages is the resulting thermal environment. With an encapsulated or potted converter, heat from the components with high power dissipations spread throughout the entire converter so that, to a first approximation, all of the components are operating at about the same temperature - a so-called **isothermal design**. While adding some thermal mass and inertia to the converter, this approach has a major disadvantage. All components are not capable of operating at the same maximum temperatures. Modern power semiconductors, for example, are rated for operation with junction temperatures up to 150 °C, while other components such as optocouplers and some types of capacitors may only be rated for 100 - 105 °C. In an isothermal design, there are two choices:

- Operate some components above their intended temperature rating
- Operate some components well below their normal operating range

The first choice is obviously not a good idea because of reliability concerns. The second choice, while reliable, results in a very inefficient design in terms of power density, cost, footprint and useful power output. The open-frame approach, if carefully designed, results in a much better overall design trade-off. Each type of component can operate within its intended temperature range rather than at one extreme or the other. For example, the data shown in Figure 7.21 are from an Artesyn 10W open-frame DC/DC converter operating at an ambient of 35 °C. Note how the open-frame design allows each component to operate at an appropriate temperature while still allowing for conservative derating for good reliability. This ability to "fine tune" the thermal design is only one advantage of the open-frame approach. Others include:

Component	Temperature Rating - °C	Temperature Operating - °C	Derating %
MOSFET	150	74.5	50.3
Diode	150	66.7	55.0
Planar PCB	130	44.0	66.0
Inductor	110	57.5	47.7
Film Capacitor	110	51.8	52.9
Opto Coupler	105	52.5	50.0

Ambient Temperature = +35°C

Figure 7.21- Example of Component Derating in Artesyn DC/DC Converter

- Direct component access to cooling air
- Both sides of converter can dissipate heat - up to twice the effective surface area for enhanced cooling ability
- Smaller "thermal shadow" behind the module. Reductions in ambient of components behind the converter of over 5 °C are typical.
- Rough surface increases turbulence and enhances heat removal
- Elimination of encapsulant reduces TCE stresses during

Power System Thermal Design

thermal cycling and thermal shock. Results in increased operating temperature range and higher reliability.

- Higher power output capability for given footprint (50% - 80%)
- Elimination of heatsink in many applications
- Lower profile and reduced card-to-card pitch
- Easy to measure converter component temperatures in system environment.

We have only presented a very brief overview of the thermal design issues and challenges associated with high density DC/DC converter design. To do a proper thermal design of such a product requires both highly specialized and highly experienced design personnel. That is one reason why there is no widely established market for completely custom high-density converters - the required design cost would be too high. This is one of the major advantages of using standardized converters. You can purchase an Artesyn standard DC/DC converter and be assured that all of the thermal design issues have been carefully addressed, and that you have one less item to worry about.

7.5 Thermal Design Examples

In this section we will present some examples of how to accomplish the thermal design for some typical systems incorporating Artesyn power conversion products. While the assumptions used here will be different than the specific situation in your system, the intent is to show meaningful examples that represent commonly experienced conditions. With the examples here and the detailed information published in Artesyn's datasheets and application notes, similar analyses for your system should be possible.

SXA10 in Natural Convection - The SXA10 is a 10W open-

frame DC/DC converter that is designed to operate in natural convection (as well as forced convection) telecom and datacom applications. In this example we will use the SXA10-48S3V3, which provides up to 2.75 A of 3.3 V with a nominal 48 V telecom input bus. Its minimum efficiency is 79%. This converter conducts some of its dissipated heat through its pins to the PCB, and the applications literature recommends a PCB layout to accommodate this. It is important to follow these layout guidelines for maximum thermal performance. Most of the heat removal, however, will be via convection from the substrate of the converter. As stated in the datasheet, the best natural convection thermal performance results when the converter is mounted in a vertical orientation. This will occur automatically if the system's PCB boards are mounted vertically - this is assumed for this example.

We will assume that the system application requires a maximum of 2.5 A of output current, or an output power of 8.25 W. With the 79% efficiency assumption, the power dissipated in the converter will be 2.19 W. This will be a useful number for entry into the system's thermal heat load budget and as an input to thermal simulation and analysis software. This number will be slightly conservative and "worst case" since it assumes the minimum efficiency rather than typical and also the maximum system current demand rather than typical. This additional design margin will result in improved reliability. The SXA10's thermal specification is based upon the ambient air temperature, so an estimate of the ambient temperature must be obtained, either from a thermal analysis software package, by measuring the ambient in a system prototype, or from experience with other similar products. Ideally, the ambient temperature estimate or measurement should be for the area directly below the converter as shown in Figure 7.22




Figure 7.22- SXA10 Thermal Design Example

We will assume that the maximum external temperature in this example is 50 °C and that there turns out to be a 20 °C temperature rise between the external temperature and the ambient below the converter, giving a value of 70 °C for the SXA10's ambient. The SXA10 is rated for full power output in natural convection up to an ambient of 85 °C, so it should operate very reliably in this application without the need for any derating. Verifying that the ambient is below 85 °C is actually the only thing needed to insure that this thermal design is valid. There are no complex calculations needed. If the environment were more demanding, with ambients approaching the 85 °C limit, there is another useful feature of the SXA10. Artesyn specifies a "thermal hot spot" on the main output diode tab of 110 °C. This point is easily accessible for measurement, thanks to the open-frame design. Attaching a thermocouple to this point is a very convenient way to verify that the converter is operating within its limits.

EXB50 in Forced Convection - The EXB50 is a family of open-frame high efficiency converters, with output powers up to 50 W. Due to the use of synchronous rectification, they are capable of full load operation up to 70 °C ambient in natural convection even at these high power levels. In this example, we will assume that the system has a very high ambient temperature in

the vicinity of the converter of 80 °C, requiring the use of forced convection cooling. We will use a EXB50-48S3V3 converter which is capable of 10 amps output at 3.3 V. The maximum system load is 8 A. The converter typical efficiency is 90%. The output power will be 26.4 W and the power dissipated in the converter will be 2.9 W. The EXB50 application note recommends the usage of 2 oz copper for the PCB traces and gives a suggested layout for enhanced thermal performance. It is assumed here that this type of layout is used.

First of all, let's determine the maximum output current that could be delivered with only natural convection. To make this determination, we can use the thermal derating curve from the EXB50 datasheet, which is shown in Figure 7.23. The curve shows that the converter must be derated above 70 °C ambient in natural convection. Using the system maximum ambient with this curve shows that the maximum output current in natural convection would be 6.6 A, less than the 8 A that is needed. So we next need to determine the air velocity in a forced convection system that is required to allow at least 8 A from the converter. We can use the curve of ambient temperature increase vs. airflow from the EXB50 application note for this purpose. This curve is also shown in Figure 7.23. To provide additional design margin, we will determine how much air velocity would be required for the converter to supply its full 10 A of output current. Then, operating at 8 A maximum, it will run at even cooler temperatures for higher reliability. Using the second curve in Figure 7.23, with an ambient temperature increase of 10 °C (full load to 80 °C ambient rather than 70 °C) we see that an airflow of 0.45 m/s (90 lfm) is required. This is a very modest airflow, and yet allows full load operation at 80 °C ambient. This outstanding thermal performance is made possible by the high efficiency and open-frame layout of the converter.

Power System Thermal Design

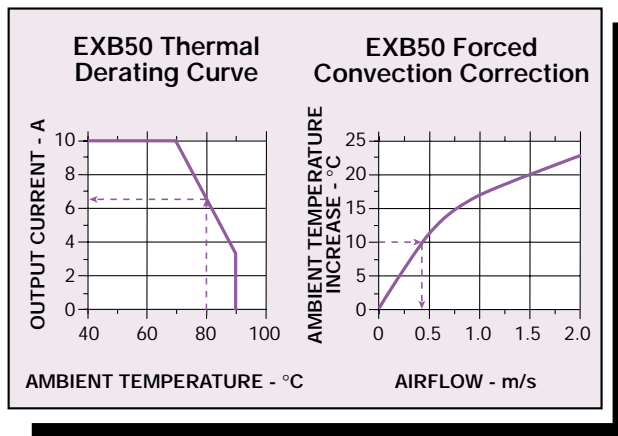


Figure 7.23 - EXB50 Thermal Design Example

The EXB50 applications literature also describes 4 "thermal hot spots" that can be used as monitoring points to verify the thermal performance of the converter under severe environmental conditions. Attaching thermocouples to these points can also provide valuable data for determining the long-term reliability of the converter in your particular system environment.

BXB100 with heatsink in Forced Convection - The BXB series is a family of high power converters designed for a forced convection environment. It is an enclosed converter module with an integrated aluminum baseplate. The converter is designed so that the dissipated power will be removed through the baseplate and optional heatsinks attached to the baseplate by means of screws into M3 tapped holes. There are three heatsink heights available so that the thermal performance can be tailored to the thermal constraints of the individual system application. The heatsink heights are 0.24", 0.45" and 0.95". We will use a BXB100-48S05 converter in this example. It will provide 5 V output at 100 W (20 A) at an efficiency of 83%. It is rated for a maximum baseplate temperature of 100°C.

This converter is installed into a forced convection system with an air velocity of 400 lfm (2 m/s). The temperature of the cooling

air (ambient) is 50°C. The converter supplies a system load of 75W. For enhanced reliability, the power system designer decides to limit the converter baseplate temperature to 90°C rather than the maximum 100°C rating. What heatsink should be selected for this application?

First, we need to calculate the power dissipated by the converter, which will be the amount of power that needs to be removed via the heatsink. Using the 75 W output power and the 83% efficiency, we find that:

$$\text{Input Power} = 75\text{W}/0.83 = 90.4\text{W} \text{ and}$$

$$\text{Dissipated Power} = \text{Input Power} - \text{Output Power} = 90.4\text{W} - 75\text{W} = 15.4\text{W}$$

Now that we know how much power we need to handle, and the baseplate and ambient temperatures, we can calculate the required thermal resistance from the converter baseplate to the ambient air:

$$\text{Thermal Resistance} = (T_{\text{Baseplate}} - T_{\text{Ambient}})/\text{Power} = (90 - 50)/15.4 = 2.6 \text{ }^{\circ}\text{C}/\text{W}$$

The Artesyn BXB applications literature provides data for the thermal performance of the three heatsinks as well as the bare module without a heatsink. These data are displayed in Figure 7.24. These curves show the thermal resistance of each heatsink (from baseplate to ambient and including the converter-heatsink interface) as a function of airflow velocity. Our system has a velocity of 400 lfm, so drawing a vertical line at 400 lfm will define the thermal resistance of each heatsink as it intersects the performance curve of that heatsink. We need a thermal resistance of 2.6 °C/W or less, so we find from the curves that the 0.45" high heatsink, with a thermal resistance of 2.4°C/W should satisfy our requirement.

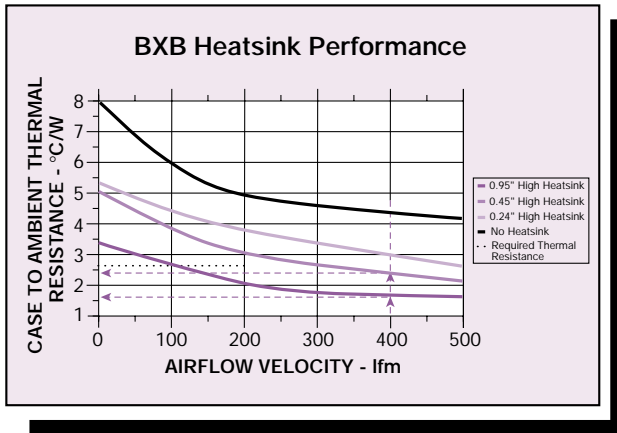


Figure 7.24 - BxB100 Thermal Design Example

Also note that the 0.95" high heatsink would provide significantly less thermal resistance - 1.65 °C/W. If the height of the heatsink were not a critical parameter in this application, the power system designer may want to consider using the larger heatsink due to its superior performance with very little additional cost. We can easily calculate the improvement in converter temperature with the 0.95" heatsink as follows:

$$\text{Temperature Rise} = (\text{Thermal Resistance})(\text{Power Dissipated})$$

$$\text{Temperature Rise} = (1.65 \text{ °C/W})(15.4\text{W}) = 25.4 \text{ °C}$$

$$\text{The baseplate temperature will then be } 50 + 25.4 = 75.4 \text{ °C}$$

As we will quantify in the chapter on Reliability, this decrease in converter baseplate temperature from 90 °C to 75 °C will provide a very significant improvement in power system reliability, and the use of the larger heatsink would be an excellent investment if the system packaging allowed for the increased height.