

**Thermal Management of
High Power in Small Spaces: Myths
and Misconceptions Challenged**



Thermal Management of High Power in Small Spaces: Myths and Misconceptions Challenged

Introduction

As power dissipation of electronics increases in parallel to the miniaturization pressures on these same volumes, previous rules of thumb and myths involving system level thermal management and cooling need to be challenged to ensure a valid overall packaged system solution that meets the overall function and environmental requirements. These misconceptions and myths arose when today's decision makers were just entering the electronics packaging industry. The capability of the electronics, their packaging density and power dissipations are much greater today, so these previous rules of thumb no longer apply. These misconceptions are that: (1) Natural Convection is always available to be used; (2) Thermal performance is only an issue for the chassis design; (3) Power is the only board information that is required for chassis design; (4) Heatsink and Thermal Interface Materials (TIM's) can solve any thermal problem. (5) CFD is sometimes thought of only being colored pictures for directors but it can be a valuable tool for thermal management. The source of these myths and misconceptions is discussed and the challenging of these myths and misconceptions provides a path to understand today's thermal management issues and how to ensure that they are addressed up-front in product planning.

Myth #1: Natural Convection is always available

When approaching new systems thermal designs, the vast majority of the initial discussions revolve around the need to cool the electronics without the

use of an air moving device, a fan. This desire from system architects is derived from limited power budgets and Mean Time Between Failure (MTBF) initial allocations; fans are always a problem in these areas. The problem is that natural convection cooling has very low performance, and most of today's embedded electronics is just too hot. If natural convection cooling is a hard requirement, this needs to drive the up-front payload power dissipation budget and the envelope dimensions of the system.

Myth #2: Thermal performance is only an issue for the chassis design

System designers have a tendency to save the thermal issues to well after the payload selection process is completed. Valuable design information pertaining to device power dissipations and locations, air flow paths and overall thermal performance selection is sometimes missing, buried or distilled into temperature data only.





Usually this detailed information is available for the board products, but must be solicited from the engineering archives from most board vendors to make it available for chassis design. Data sheets for commercially available payload cards are very sketchy on their thermal performance characteristics. Table 1 shows some examples of typical board specifications that cover thermal performance of these cards.

to the design details. This practice results from the early days of standard based cards such as VME and cPCI having low power dissipation and higher temperature parts due to the processors available at that time. Figure 1 shows the exponential growth in processor power dissipation of processor generations as their capacity has also increased.

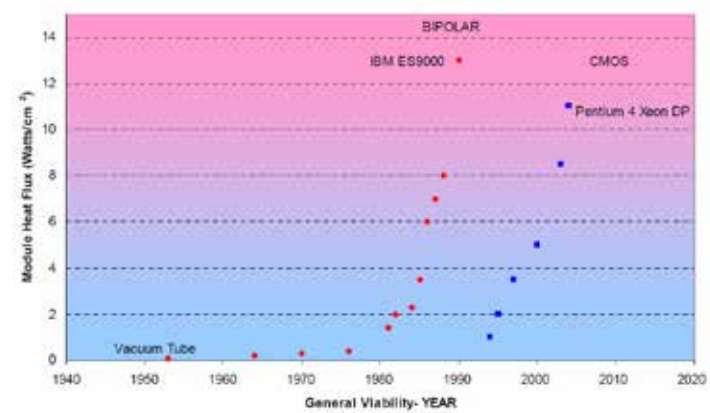
Table 1: Historical Board Thermal Data Summary

Typical Circuit Card Thermal Data
100 CFM or 400 LFM when operating at 71° C
-40° to 85° C
0° to +70° C, operating
Operating Temperature: 0° to +55° C
Operating: 0 to +60° C
Convection Forced Air
300 LFM min @ 560° C @ 500 MHz
0° to 55°C Forced Air Convection
No Specification Provided!

Myth # 3: Power is the only board information that is required for chassis design

This statement is an assumption that creeps into the design process for selection of payload card for population into a system. The electrical performance is paramount during system architecture for power supply sizing, and the packaging concerns of power and temperature effects on devices are left

Figure 1: Processor Power Over Time
Processor Power Growth Versus Time



This growth in power dissipation impacts the fluid flow properties of card layouts since the devices are getting larger and the heat sinks to cool these hotter devices are getting larger for forced convection cooled cards. This all increases the flow impedance, which lowers the total amount of air flow that is available to cool this specific card, and impacts the chassis selection of the air moving devices.

This also applies to conduction cooled assemblies designed in accordance with IEEE 1101.2, ANSI/VITA 30.1, ANSI/VITA 47 and VITA 48, where module temperature specifications end at the module edge and do not include wedge clamp interface losses.

For the same card in a conduction cooled module configuration large temperature rises (ΔT 's) are associated with this dry metal to metal contact area between the card back edge and where the module interfaces with the chassis wall. This dry contact area has the following typical temperature impedance characteristics vs. the contact pressure generated by the module wedge clamp; these curves show the limitations of increasing pressure beyond a certain

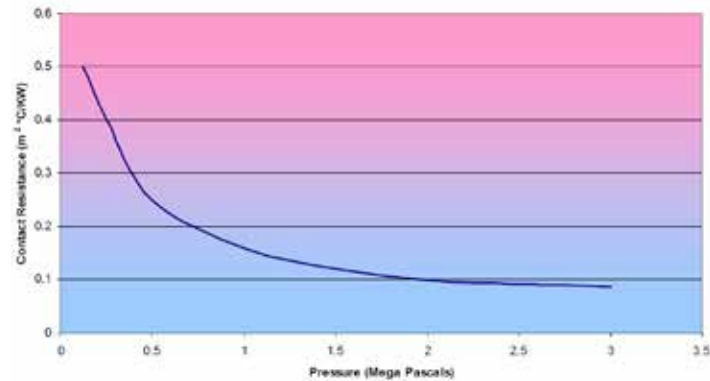




point. Different types of wedge clamp have vastly different performance. Not knowing the style and thermal performance of the wedge clamp (mounted to the module) makes assessment of the maximum chassis wall temperature problematic for today's hotter modules.

performance charts that are readily outputted from CFD programs. The colored temperature plots tell a story of the temperature profiles, but are often perceived as being more important than the reporting of the estimated junction temperatures for the critical devices.

Figure 2: Dry Contact Thermal Impedance
Contact Resistance 7075 T6 Ground Plate Surface Finish 4




Myth #4: Heatsink and Thermal Interface Materials (TIM's) can solve any thermal problem

Yes, hot components can be cooled by aggressive application of high efficiency heatsinks and highly conductive Thermal Interface Materials (TIMs), but their application in solving a device thermal problem late in the design cycle can cause other system-level problems upstream or downstream from this solution. More aggressive heatsinks have increased fin count, increased fin thickness, or increased turbulence - all of which increase the heatsink's flow impedance (pressure drop). This can cause problems at the system level, e.g. if air-moving devices cannot handle the higher pressure required, or adjacent modules are much lower pressure.

Myth #5: CFD is sometimes thought of only being Colored Pictures For Directors

Computational Fluid Dynamics (CFD) software is used to solve three-dimensional heat transfer problems, since textbook heat transfer correlations are for simpler geometries and test cases than the complex topology of modern circuit cards. Results are often presented using colored system

Figure 3: CFD Velocity Profile



Challenging the Myths

In order to ensure good thermal management, we need to challenge each of these myths to provide value for the system design. This is accomplished by thoroughly understanding the thermal and power dissipation issues involved with cooling embedded electronics and applying thermal design techniques across the range of packaging levels from die to platform level.

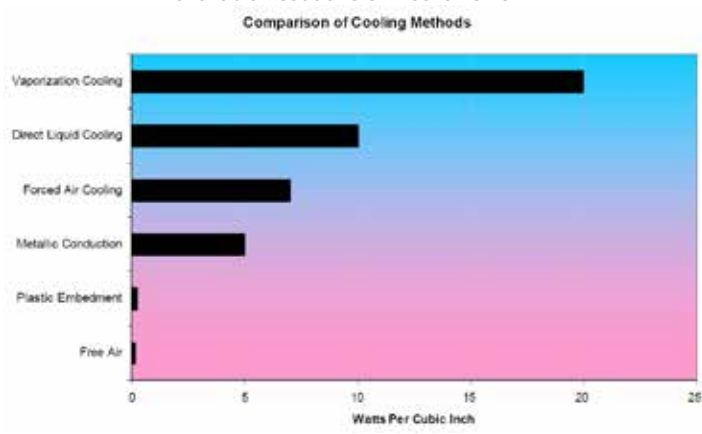
Challenge to Myth #1: Natural Convection is always available!

Natural convection cooling is very rarely utilized in high power systems. When we get down to the basics of power dissipation and physical size to determine the power density (Watts/cubic inch)



and compare this value against known capabilities for the various heat transfer methods, the limitations of natural convection (free air) cooling become obvious. For example, guidance from MIL-HDBK-251, Figure 4 shows a comparison of cooling methods for relative power densities of the system.

Figure 4: Typical power densities compared to available heat transfer mechanisms



Let's look at why natural convection cooling has these limitations.

Natural convection heat transfer, sometimes mistakenly referred to as "Free Convection" is described as a change in the temperature of the air particles adjacent to the warm surface of the system dissipating heat. The temperature of these air particles is increased, thus changing their local density and causing these higher temperature fluid particles to become more buoyant in comparison to the fluid particles away from the warm surface. This is the only mechanism that drives convection current that can be used to cool electronic equipment without the use of fans. This is a good application case, but it comes at a cost of having lower power and requiring a large surface area to conduct the heat energy from.

Table 2 shows some examples to underscore what you can and cannot do with natural convection cooling. We are all familiar with the simple light bulb, 100 watts in this case. In all normal light fixtures, a fan is not required to keep this item cool (although it does get very hot), so let's look at the details.

Table 2: Comparison of Heat Dissipating Devices

	Light Bulb	BGA Package
Power dissipation	100 W	25 W
Surface area	106cm ² (bulb surface area)	1.96 cm ² die area
Heat flux	0.9 W/cm ²	12.75 W/cm ²

From review of the details shown in Table 2, the power dissipation per unit area, heat flux for the BGA package is almost 15 times higher than that of a 100 Watt light bulb. The conclusion here is that it is important to review heat flux before selection of the cooling technique, and then assess the application of extended surface heat sinks and heat spreader technologies, including phase change micro-channel heat spreaders or heat pipes to get the heat energy transferred into the ambient environment.

Also, keep in mind that the above illustration is based on room temperature and Mean Sea Level conditions. It is important to understand that the viability of Natural Convection further degrades as the air density decreases due to temperature or altitude effects. Lower power systems require close evaluation in order to be successful using natural convection, but system architects of high power systems should stay away from this myth for the reason stated above.

Challenge to Myth #2 and 3:

Thermal performance is only an issue for the chassis design!

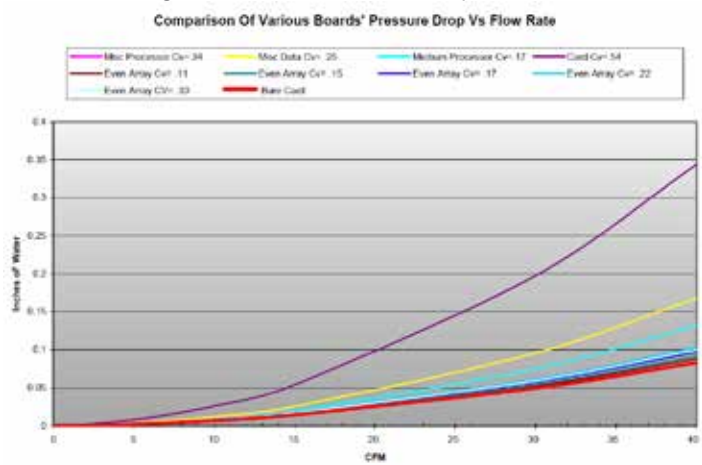
Power is the only board information that is required for chassis design!

To adequately thermally manage circuit cards, the mass flow heat transfer laws require knowing the mass flow rate through the chassis and inlet-to-outlet fluid temperature rise across the chassis. It is the flowing of the fluid that transfers the heat energy from the power-dissipating devices to the ambient environmental conditions. Power dissipation is the



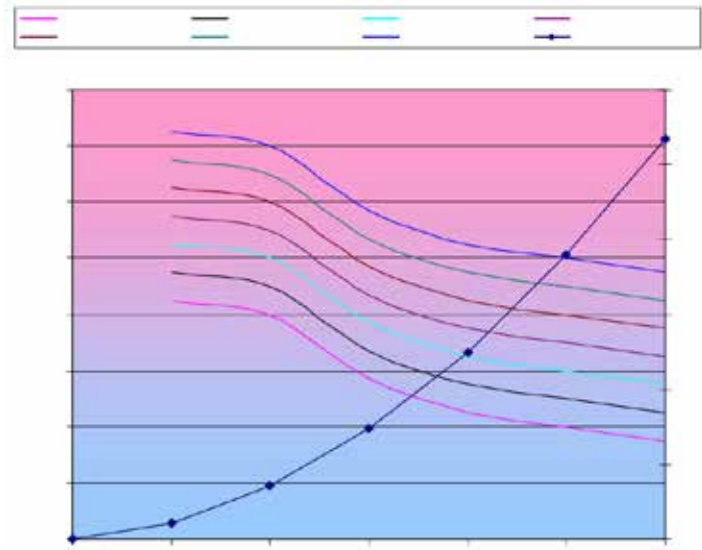
information that is known for most circuit cards, but other design information is necessary in order to select air movers and set the final system level power consumption of the system. From a chassis design standpoint the characteristic flow impedance and device layout information are design characteristics that are needed for all air-cooled systems, but especially for small form factor systems. Figure 5 shows that from a fluid flow standpoint, as you place more components on a circuit card and stack multiples of these cards in a flow stream, the amount of pressure to develop the flow rate through the card increases quickly. This is a normal characteristic of all fluid carrying elements, from the simple water piping in everyone's home to high performance jet engines.

Figure 5: Air-cooled Card Flow Impedances



More recent standard developments and major high capacity processor cards have acknowledged that more robust thermal performance data is required for design of chassis for aggressive application environments. Design layout data is making its way onto CFM analysis for both the prediction of device temperatures and also for the determination of the flow characteristics that assist chassis developers in selection of the correct fan selection for a coordinated system solution. This type of data can be in the form of pressure drop versus flow rate and also critical temperature devices versus flow rate, and are available from knowledgeable board vendors providing high-capacity circuit cards. This type of information has now become invaluable for the chassis designer and forms a basis for performing reliability assessments of the system while still in development.

Figure 6: Pressure/Temperature versus Flow Rate Example



For conduction cooled modules, module vendors are validating their designs to the requirements of ANSI/VITA 47 for conduction-cooled cards. This fixes the module edge temperatures as follows.

Table 3: ANSI/VITA 47 Conduction Module Edge Temperatures

Class	Minimum	Maximum
CC1	0° C	55° C
CC2	-40° C	55° C
CC3	-40° C	70° C
CC4	-40° C	85° C

In addition to the card edge temperature identified in Table 3, the wedge clamp vendor and model number and wedge clamp length need to be disclosed for completeness. This listing of the data of the maximum module edge temperature, wedge clamp vendor and model number, and length make up the minimum design boundary conditions sufficient for design of enclosures for cooling conduction cooled modules.



Challenge to Myth #4:

Heatsinks and Thermal Interface Materials (TIM's) can solve any thermal problem!

Heatsink and TIM manufacturers have risen to the challenge of providing highly capable designs and materials to meet the increased power dissipation in modern electronics devices. We now see in the industry heat sink forms that are no longer just simple extrusions. Folded fin, low flow bypass and active heatsinks are now available for use, but need to be integrated into the design as a planned event rather than a fix for an over-heated device.

Thermal management of modern electronics needs to have critical components identified and characterized for their detailed thermal performance. At preliminary board layout, identification of thermally critical devices and determination of their required thermal resistance will quickly identify when and where a heatsink and TIM will be required to be used for adequate thermal management. Device layout and air path planning on the card and through the chassis needs to be taken into consideration to ensure that the cooling air is being guided or directed to the devices needing the maximum amount of cooling air. Dense aggressive heat sinks with high surface areas have their value, but actually may increase the temperature of devices if the cooling air is bypassing the heatsink, since the fin spacing may be too dense for the fluid to travel through the fins. Careful air path planning and device placement will lead to selection of an efficient extended surface device (heatsink) to adequately cool critical devices. Layout and placement of devices with heatsinks becomes an important and iterative activity and closely ties to the air path developed with the chassis or enclosure that is being designed to carry these payload cards.

Challenge to Myth #5:

CFD is sometimes thought of only being Colored Pictures For Directors!

Computational Fluid Dynamics (CFD) analysis has risen to become an important tool in evaluation of the viability of electronic systems. The use of CFD is a benefit to all in determination of the

junction temperatures of critical devices, with both simple and complex device models. Determination of the device junction temperatures is where reliability assessments start, and the earlier these assessments are made available to the development team, the better organizations can improve system performance. CFD analysis, Handbooks (MIL-HDBK-251), Textbooks, Custom Spreadsheets and Flow Network Modeling are all valid tools and techniques to start and maintain a Thermal Management Architecture. Correct use of the tools can investigate and optimize available cooling options, ensure feedback to support system trade-off decisions and ensure that air movers are operating at acceptable operating points. Yes, the summary of these trade-off analyses may be in the form of a colored picture for a report or presentation, but these tools provide valuable insight into design tradeoffs. Detailed temperature predictions allow for interaction with the design team including EE, ME & Regulatory disciplines, to provide design feedback and to refine the thermal performance of the product under development.

Conclusion

Today's small form factor electronic systems, spanning the range from devices through circuit cards and chassis, need to be carefully evaluated for thermal performance. This is best accomplished by adhering to the basic premise of quickly breaking down the thermal problem into its basics. Once this is done, viable solutions can be identified and quick packaging tradeoff evaluations can be made. Investigating cooling options earlier in the design process can ensure a top down system approach for the thermal design, thus ensuring the development of a successful system using small form factor electronics in aggressive applications.