White Paper | September 2006

radisys®

Successful Carrier Card Design for COM Express: Introduction to Thermal Design Considerations

Israel Dubin, Manager, Electro-Mechanical Power Platforms Group,
Radisys Corporation

Overview

Gordon Moore, one of Intel's founders, predicted in 1965 that the number of transistors placed on a computer chip would double each year. Over time, the pace has slowed a little and the prediction has been updated accordingly, but it still generally holds true, now doubling approximately every two years. Intended as a rule of thumb, this prediction became a guiding principle for the industry to create faster and denser powerful semiconductors at proportionately less cost.

Moore's Law deals with the growth in chip circuit density. It leaves the practical problems of dealing with the effects of this observation, including thermal management, for design engineers to solve. Because the COM Express™ specification combines significant numbers of high density integrated circuits into a tightly packed module that's typically embedded into compact products, engineers need to carefully consider thermal management as they develop their designs.

Through understanding of a few basic principles, engineers can improve the reliability of their designs and prevent premature failures when their products go into the field. This white paper provides an overview to help systems designers successfully manage the thermal aspects of their COM Express applications. By paying attention early in the design to potential "hot spots" and considering airflow during their design and prototype phases, designers can limit thermal management problems and assure reliable application operation in the field.

Because every thermal design has a unique set of interdependencies, examples in this white paper are only for the purpose of illustrating general points and not as recommended solutions for any specific design.

CONTENTS

Executive Summary *pg. 2*

Tools of Thermal Design *pg. 2*

Thermal Management for COM Express Design *pg. 4*

Defining Thermal Solution Parameters *pg. 4*

Thermal Interface Material *pg. 5*

Conclusion *pg. 10*

Executive Summary

Following Moore's Law, advances in microprocessor technology have enabled system developers to provide increased processing power in ever smaller, more highly integrated devices. Because of this, we now have microprocessor technology embedded or able to be embedded in nearly everything from automobiles to telecommunications equipment. Today we have cell phones that send and receive music, pictures and video and even television programs. Lightweight laptops house the power of yesterday's supercomputers. Handheld PDAs and Blackberries keep us on schedule and up on email. There's the "does everything" Treo Smartphone with laptop-like capabilities—email, television, personal videos, music, audio player, Web surfing and even GPS mapping. Things Moore probably never dreamed of holding in the palm of his hand when he helped found Intel are now commonplace.

A frequently cited benefit of reduced processor size is reduced power consumption. But this is really only true at the level of the individual devices within the chip. The overall power reduction has not progressed at the rate of miniaturization. As processors got smaller, we've packed more heat generation into that smaller space and the devices have gotten hotter. With this higher performance, thermal management needs to be addressed, especially in embedded computers that operate in warm environments inside a closed chassis or other densely packed enclosures.

Tools of Thermal Design

Thermal design is often complex. Successful thermal designs can require understanding of thermodynamics, fluid mechanics and heat transfer through different mediums. Thermal behavior can be subtle, even counter-intuitive, making it difficult and time-consuming to comprehend. Ignoring the different components of thermal solutions can adversely and dramatically affect device reliability and system performance.

When energy is transferred as a result of a temperature difference, it's called heat. In electronic designs, heat is a direct result of power consumption. It is transferred in three ways (or in a combination of these): conduction, convection, and radiation.

- Conduction transfers heat when the surface of a solid medium contacts another and both are at different temperatures.
- Convection is the exchange between a moving fluid (air propelled by a fan for example) contacts a solid surface with a different temperature causing a thermal-energy exchange.
- Radiation transports heat as electromagnetic waves that travel at the speed of light.

Managing System Airflow

Thermal management in most of our environments usually depends on two things: heat transfer from the parts, often augmented with heat sinks and the airflow necessary to carry the heat off and then away from the system.

System designers must consider the requirements of all of the components on a COM Express module and its carrier card when determining the airflow required for their application. The amount of airflow depends on the ambient temperature (TA) surrounding the product, power dissipation of the devices, and the amount of temperature increase or rise the components or environment can tolerate. The placement of a module in a system chassis, the ability to both deliver and remove the required airflow and limitations of the application environment can all have dramatic impact on the design and selection of the components making up the solutions.

Several factors generally determine system airflow or the ability to use it:

- Fan capacities and flow directions
- Chassis/case size & design

- Locations and shapes of chassis/case air intake and exhaust vents
- Board component placement
- Placement of other cards, add-ins and cables within the flow paths

A system may have a strong airflow, but still contain “hot spots” within the chassis. Poorly positioning fans, components, chassis structure, cables, etc., can restrict, block and redirect the airflow within the system, often creating unexpected hot regions. To avoid some of these, ensure space is provided to locate items requiring minimal cooling, well away from those that do. (For detailed examples of this, see the “Real World Cooling Applications” section later in this document.)

Deficiencies of Existing SIP Control Protocols

The other key driver for MSML is the weakness in the existing SIP-based protocols available for control of media servers. There are two such protocols available, Netann (now RFC 4240) [NETANN] and MSCML {MSCML].

Heat Sinks

A heat sink cools a device by increasing the surface area of the part it’s connected to. It effectively reduces the thermal density of an assembly by allowing energy to dissipate over a larger surrounding area. That is, it allows thermal energy to flow from the hotter device to the cooler heat sink or radiator and escape into the yet cooler ambient air of the environment. Heat sink selection often requires careful consideration of material selection, fabrication methods, and mechanical limits in size and mass as well their impact on the air stream moving through and around them.

Heat Pipes

Heat pipes move heat. They acquire heat from a source and move it efficiently to a location where it can be more readily disposed of. For example, a heat pipe can work well where space is insufficient for a heat sink at the heat source or where no airflow is available for the device mounted heat sink. In these cases, a heat pipe can be used to move the heat to a location where there is room for a radiator and where the heat can be carried away from it by the air stream. Successful heat pipe applications require consideration for how the heat can efficiently be transferred to and from the heat pipe itself.

Thermal Interface Material

Thermal Interface Material (TIM) is used to fill the surface imperfections between heat sources and heat sinks or heat pipes. The material works by displacing air which would otherwise act as an insulator between the parts and replacing it with a more thermally conductive material. There are many different types of these materials and proper selection is often a matter of trade-offs between thermal performance and other characteristics of the material, both electrical and mechanical, often over a product’s operating lifetime.

Fans

Fans are often needed to move the air through a heat sink or system chassis in sufficient volume to carry away the energy. However, faster fans aren’t always an option and don’t always mean better cooling. Higher capacity fans can increase system noise levels beyond acceptable limits for the application and can also introduce unexpected flow management problems within the system.

Adaptive Cooling

Adaptive cooling gives designers the flexibility of real-time temperature control. It uses hardware and often software to monitor specific devices or environmental conditions and then modify the operation of system components. Responding to dynamic changes in component operation or ambient conditions allows fan speeds to increase as necessary while keeping noise to a minimum when not. Under appropriate control, processor speeds can be adjusted to reduce power and entire sections of an electronic assembly can be turned off or disabled to reduce both power consumption and system temperature rise. These tools have their own power consumption and dissipation needs too and of course add another level of complexity and cost to system design.

Thermal Management for COM Express Design

Although designers can benefit from the highly integrated, densely-packed and high performance COM Express modules to get to market quicker, they must still do some work to provide ways to remove the heat dissipated by the module and its surrounding assembly so it can provide its best performance.

Proper thermal management ensures that the temperatures of all components for a COM Express-based system are within specified functional limits. As the module and its support assemblies run hotter, their effective life is reduced. Operating a module (or any device) for extended periods outside of its specified range will eventually degrade its performance and create reliability issues for the end user.

If device temperatures go too far past the specified ranges, operating characteristics may change. For example, excessively high (or low) operating temperatures can affect a device's internal timing. Once that device's timing changes it may disrupt the operation of associated circuits, halt the system's operation entirely or damage a module permanently.

Designing optimal solutions for COM Express-based products is a complex problem. A designer must consider many parameters for a high-performance, low-cost thermal solution including:

- Surrounding ICs
- Maximum power of all components
- Ambient and case/housing temperatures
- Module geometry and orientation
- Airflow directions
- Heat source dissipation surface
- Heat sink profile, fin geometry and orientation and mass
- Material (aluminum, copper or alloy)
- Manufacturability
- Heat sink attachment mechanism and therma interface material (grease, dry-film or thermal tape)

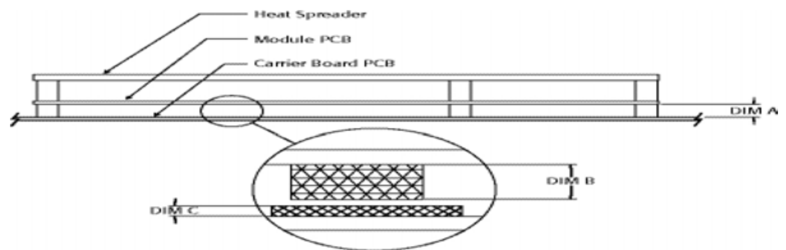

Figure 1. *The tight tolerances of the COM Express specification mean that engineers much pay close attention to removing heat from their designs.*

Defining Thermal Solution Parameters

To size a heat sink for a given device requires solving a number of equations describing the thermal characteristics of the targeted device to be cooled. For illustration, we'll focus on a processor that might be used on COM-E modules.

First, you must know how much power you will be dealing with. That is how many Watts the device will dissipate. For current Intel processors, this is usually expressed as thermal design power (TDP). For example, a specific 2.0GHz Pentium-M family may have a TDP of 27W. Designing with that TDP will permit a given solution for all members of that particular processor group or family.

Second, you need to know the maximum junction temperature (Tjmax) of the device. If the product will operate in cold environments, the minimum junction temperature (Tjmin) will need to be considered too to determine if supplementary heating is required. Using the same 2.0GHz Pentium-M, we know that is a minimum of 0°C and maximum of 100°C.

Third, you need to know the ambient temperature of the air that will surround the device.

As figure 2 illustrates, the total thermal resistance is the resistance from the device junction to the ambient air. It's expressed as R ja (or Rja). TA is the ambient temperature (60° C in this example).

R ja = Tjmax - TA / TDP
or (100°C-60°C) / 27.0W = 1. 48°C/W

Starting with the Rja of 1.48°C/W, we need to subtract the thermal resistance component (or components) not within the board designer's control. In these examples, that fixed component is the resistance between the device junction and the exterior of the device case where our primary thermal solution will connect, or Rjc.

Once you've removed the Rjc component of the thermal resistance, what's left is the total resistance that your heat sink and its thermal interface to the processor case (or die) can have for any given airflow. Using the example of a 2.0GHz Pentium-M, Rjc = 0.94°C/W, leaving 0.54°C/W for the thermal interface and heat sink.

If, for example, the thermal interface has a resistance of 0.10°C/W, then we need a heat sink that has a resistance of no more than 0.44° C/W.

Thermal Interface Material

The thermal interface material (TIM) selected, whether adhesive tape, silicon pads, gap pads, or thermal grease, must have a thermal resistance consistent with the requirements for the other thermal resistance components. However, the proper selection is more complex than just picking the material with the best heat transfer characteristics: Each material has its strengths and weaknesses and so you must select the right material for the application.

You must control the thickness to maintain proper clamping pressure between your thermal solution and the device you're trying to cool. Otherwise you can warp an assembly creating other problems or even crack and destroy a device.

Consider the degradation of the TIM material in the application's environment to ensure the TIM will perform adequately during the life of the product. Will it dry out and shrink over time? Will it "pump" out due to thermal cycling over time?

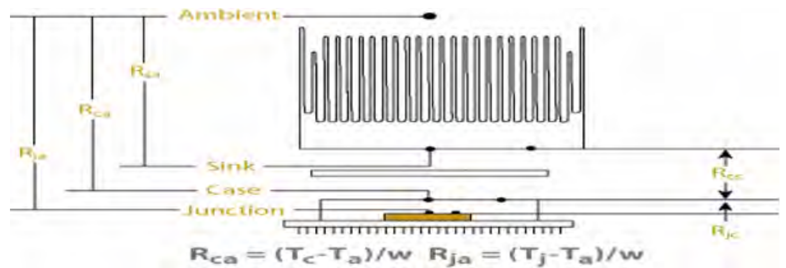

Figure 2. *Calculation using the maximum junction temperature of the device with the maximum ambient temperature computes the thermal resistance, determining how much thermal resistance is left for the thermal solution.*

The application method, operating characteristics and if necessary, transportation environmental conditions must all be considered to ensure that the material is not damaged or degraded either during delivery to the initial assembly point or as part of a finished assembly. How will the material be protected during handling? If the temperature in transit goes to high, will the material melt and flow off the carriers before final assembly?

Failure to properly observe these parameters and select the proper material can adversely affect the reliability of a product in the field and increase its manufacturing and warranty cost.

Assembly Power

A designer must also examine overall assembly power. Usually the processor TDP represents only 35-40 percent of the heat needing removal from a single board processor assembly. All of the heat needs to be dealt with, including any excess heat, to maintain an appropriate operating temperature.

Returning to the 27W processor example, we can estimate that there is actually closer to 72W to remove from this assembly—assuming there are no other sources of heat within the enclosure such as a power supply or peripheral devices. This energy must be managed to maintain the desired temperature rise. For our example, let's allow a 10°C air temperature rise from the inlet to the outlet of the package.

Knowing the total power and the desired temperature rise, we can solve for total flow volume (in cubic feet per minute or liters per second) necessary to maintain the air temperature at the desired level.

CFM=(3160* KW)/TF	LPS=(830* KW)/TC
CFM=(3160*.072)/18=12.64	LPS=(830*.072)/10=5.97

Real-World Cooling Applications

If finding the right thermal solutions were as simple as just solving a few simple equations, we'd just rush out and buy off-the-shelf heat sinks with thermal interfaces that meet the requirements described for our devices, including fans with finger guard and air filter to cool the system.

While such parts are available, their ability to provide the needed performance depends upon a number of other variables beyond the control of OEM heat sink and fan vendors.

It's up to the system engineers or integrators to evaluate the conditions of their specific environments to determine the actual performance and suitability of specific parts to the intended tasks.

Forced Air Convection Cooling Example

Forced convection cooling with an air-cooled heat sink depends upon the velocity of the air passing through or across the heat sink surface.

In simple terms, forced-convection cooling "scrubs" thermal energy off of a hot part and transfers it to a cooler medium where it can be carried away by air molecules.

When the airflow is slow, molecules in the airflow come in contact with a processor's hot surface longer allowing it more time to transfer energy. However, if the air is carried away too slowly, the processor can overheat.

Moving the air faster puts more molecules in contact with the surface to absorb more heat. However, moving the air over too fast reduces the time in contact and therefore the amount of energy transferred at a given volume while possibly increasing system noise.

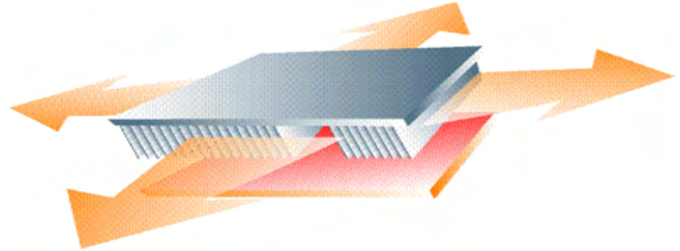

Figure 3. *In forced-convection cooling, the moving air molecules "scrub" heat away from a hot processor.*

The result is a typical heat sink performance plot showing a still increasing but decelerating rate of heat dissipation with a heat sink as air velocity increases. In many PC applications, the air velocity "sweet spot" is 200-800 feet per minute (fpm) or 1-4 meters per second. Above 800 fpm, the slope of the performance curve usually flattens and in many cases only 2-3 degrees of improvement is to be gained by going faster.

Fan Types and Cooling Issues

Most fan cooling systems use axial-flow fans, so named because the inlet and exhaust flow direction are both in-line with the axis of the impeller. Axial fans are very low cost and readily available. They're great for moving a large volume of air at relatively low noise levels—provided the physical resistance to the flow is relatively low.

Axial fans have a shortcoming requiring designer consideration. They are sensitive to turbulence at their inlets. Designers must allow a smooth entry for straight flow into the fan for it to operate within the curves presented in the data sheets. Disrupting the flow by forcing the fan to draw air in from other than a straight coaxial path or by allowing obstructions that interfere with the entry path dramatically drops the fan's airflow and can unevenly distribute air at the fan outlet.

As a general rule, the smaller an axial fan, the thinner its venturi (its body). A thinner venturi typically means the fan will be more sensitive to inlet turbulence and has less ability to handle flow resistance than its thicker brother.

Radial fans and other fan types trade volume for pressure in varying degrees. Regardless of the type of fan used, all are affected to some degree by resistance to airflow and require suitable clearances to operate.

Like any rapidly moving fluid, air wants to move in straight lines and not bend or go around corners. Anything in its path it must flow around. Each part, each cable, each bend, each curve represents a resistance to airflow. Unless you consider the full airflow path of a design, you can create unanticipated turbulence and potentially hot spots.

Every surface adds drag on the air stream passing over it or is a "brick wall" the air stream has to flow around changing the pressure or resistance to flow within the system. Besides redirecting flow, if this cumulative resistance gets high enough, it can effectively stall a fan. If the pressure equals the capacity of the fan, then the blades spin, but move no air through those parts of the system.

In smaller enclosures where usually smaller, less powerful fans must be used, this problem worsens. Smaller fans usually move less air to start with and are more sensitive to resistance. A less powerful fan produces less air volume, delivers that volume more slowly and is more sensitive to resistance in its flow paths.

Even a simple round wire finger guard at the inlet of an axial fan can reduce the performance of a small fan by 20 percent or more. Imagine what could happen with a solid enclosure wall in close to a fan inlet. This is a frequent problem for motherboards with fans mounted on top of a processor heat sink placed parallel and close to the enclosure top cover as sometimes found in low-profile servers.

To add to these worries, introducing more turbulence also contributes to fan and system noise. It can even put additional strain on the motor bearings so that they run hotter and wear out quicker.

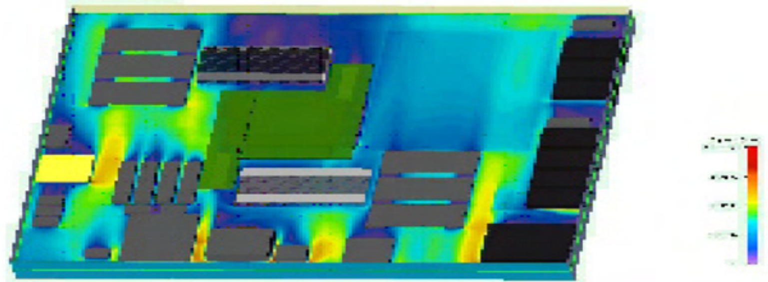

Figure 4. *Example of a design showing airflow blocked by upstream parts, then bypassing downstream parts, lowering air velocity through the board center and allowing the downstream parts to run hot.*

Figure 4 shows PC board parts placement initially driven only by the electrical design rules. These resulted in an obstructed air-flow. The high velocity airflow (bottom) starts out evenly distributed across board entry; but it is redirected to both sides lessening the airflow (blue/dark blue) in center areas. Most of the airflow bypasses critical parts allowing some of the parts in center to overheat.

Axial Fan General Cooling Guidelines

Without putting some engineering into a design, it's hard to make air go where you want. Here are a few guidelines that can help:

- Keep entry paths to fan inlets straight and clear.
- Allow at least 1.5 times the fan-body thickness for clearance between the fan inlet(s) and any upstream obstructions.
- Keep air filters, guards and other obstructions as far away as practical from the fan inlet(s).
- Be aware that partial obstructions over part of an inlet may reduce flow or create a "dead spot" downstream of the fan outlet inline with the obstruction.
- Axial fans tend to create "swirl" patterns when air enters and exits them. These swirls interact with enclosure walls and can create weak flow regions next to the walls.

- Obstructions in the air outlet path can create recirculation paths causing the airflow spin in a cyclic path around a point or series of points. A hot part in this recirculation pattern can overheat.
- Regions within an assembly with seemingly large amounts of airflow passing through can become hot because local obstructions have redirected the air around critical parts within the region.

Thermal Evaluation

Differences in motherboards, power supplies, add-in peripherals and chassis all affect the operating temperature of systems and the processors that run them. Consider thermal testing for a whole product to determine whether a specific COM Express board and carrier card combination has adequate airflow for the processors. (To begin determining the best thermal solution for your processor-based systems, contact your COM Express board or processor vendors.)

COM Express Board Stack

Figure 1 earlier in this document shows a portion of the mechanical specifications of the standard COM Express module. DIM A can be varied using different connectors. DIM B is defined by the specification and needs to be allowed for in any design. DIM C is under your design control since it is a function of the connector height (DIM A - DIM B.

Putting some numbers on the COM Express drawing shown in figure 1.

Typical board stack height is made up of:

- DIM A = 5mm = minimum carrier/module stacking height
- 2.0mm = module PCB thickness
- 5.2mm = tallest part on module side 1/top (SODIMM)
- Note: Pentium®-M Processor on top side = 2.85mm
- DIM B = 3.8mm = maximum part height on module Side 2/bottom
- DIM C = 1mm maximum part height on carrier side 1/top

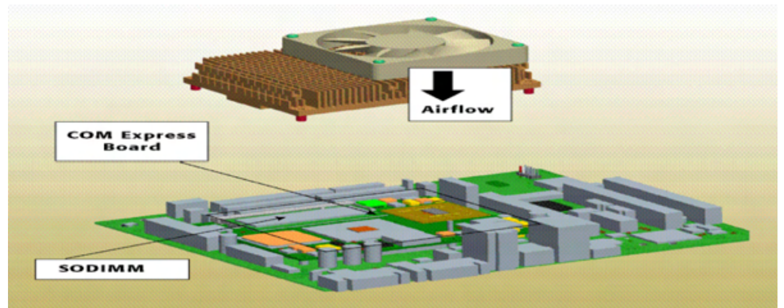

Figure 5. *COM Express module on a mother board showing a fan mounted directly to a heat sink and blowing directly down on to it.*

This gives a total stack height of 12.2 mm, (plus the thickness of the carrier board and its bottom side clearance requirements). Starting with the available space within the target enclosure and allowing for manufacturing tolerance, subtracting the 12.2mm leaves the height available for a thermal solution.

Using figures 1 and 5 to illustrate the installation of this board into a typical 1U tall enclosure, gives the following result:

- 44.45mm (1U nominal) = 43.7mm (1U actual available space)
- Less 2.4mm enclosure metal thickness and cumulative tolerance stack = 39mm available space
- 12.2mm for the COM Express module
- 1.6mm for the carrier thickness
- 4mm clearance beneath the carrier
- = 21.2mm available for the total thermal solution at minimum stack height

Placing a COM Express module on a single slot ATCA (advanced telecom computing architecture) carrier blade leaves the following space available:

- 30.48mm blade pitch
- = 21.33mm available side 1 height without a cover
- = 23.87mm with an insulating top side cover
- -12.2mm for the COM Express module
- = 9.13mm to 11.67mm available for the thermal solution

Forced-Air Cooling in the Real-World with COM Express

On the following pages, figures 6 and 7 show simulation examples of two different designs using the COM Express module and help illustrate the importance of understanding the subtleties of thermal analysis. Figure 6 shows a low density enclosure application with a thermal problem (figure 6a) and its solution (figure 6b). Figure 7 shows low profile, higher density design with its thermal problem (figure 7a) and solution (figure 7b).

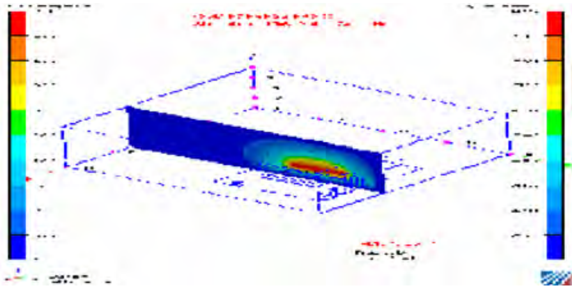

Figure 6a. *This is a natural convection example of COM Express using a 10W CPU limit. At 10W the CPU acan be held at a bit above 93° C in a normal 22° ambient environment. This is only about 7° below its limit. When the ambient temperature is raised, the processor exceeds its maxiumum ratings.*

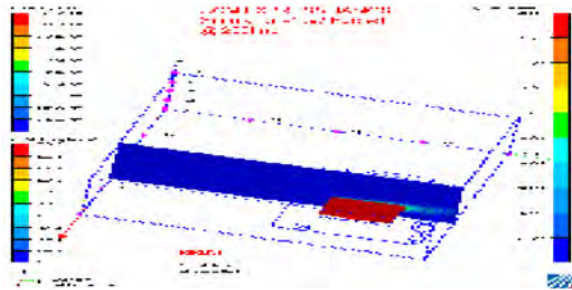

Figure 6b. *The same example with a relatively low 200 FPM airflow decreases the CPU temperature to 38° C at 22° C ambient. This allows plenty of margin even with the desired 60° C operating environment.*

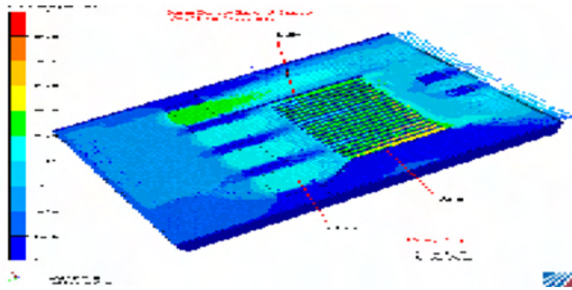

Figure 7a. *This application shows an existing 1U chassis with too much fan. There's too much airflow and most of the air was initially bypassing the processor. Adding baffles as shown in the figure made the system noisier as it forced too much air through the heat sink too fast and still left the processor hot.*

The solution to this problem was to use slower fans.

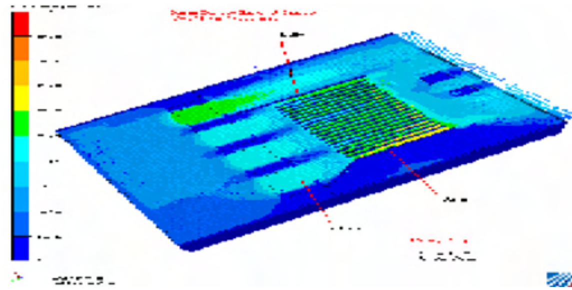

Figure 7b. *This is the same 1U chassis showing adequate cooling using 4 lower cost and quieter 10.7 CFM fans versus the 4 original and far noisier 25 CFM fans. This slower fan solution costs less, decreases noise and removes the processor heat. For later runs of the board, the fans were tweaked down o 6.8 CFM to still achieve >400 LFM at the CPU and a minimal temperature rise through the system.*

Conclusion

Today's smaller processors with higher power & increased thermal density in embedded systems require more heat dissipation and a more careful approach to thermal management. Systems designers need to consider the thermal aspects of their embedded designs, particularly COM Express designs because of their tight physical specifications. But even within such tight specifications, designers can achieve reliable thermal management for their applications by keeping power, airflow and temperature control requirements in mind. Designers have many options involving both hardware and software techniques to increase system reliability, create a more stable system for thermal management and increase the reliability of their product.

Corporate Headquarters

5445 NE Dawson Creek Drive
Hillsboro, OR 97124 USA
503-615-1100 | Fax 503-615-1121
Toll-Free: 800-950-0044
www.radisys.com | info@radisys.com

©2011 Radisys Corporation.
Radisys is a registered trademark of Radisys Corporation. Convedia, Microware and OS-9 are registered trademarks of Radisys Corporation. Promentum, and Procelerant are trademarks of Radisys Corporation.
*All other trademarks are the properties of their respective owners.
September 2006