

Enhancing Thermal Performance in Embedded Computing for Ruggedized Military and Avionics Applications

Darren Campo, Jens Weyant, Bryan Muzyka
Advanced Cooling Technologies, Inc.
1046 New Holland Avenue
Lancaster, Pennsylvania, USA, 17601
Phone: (717) 295-6821
Email: darren.campo@1-act.com
Web Page: www.1-act.com

ABSTRACT

Embedded computing systems used in many military and avionics applications are trending toward higher heat fluxes, and as a result performance is being hindered by thermal limitations. This is intensified by the high ambient conditions experienced by today's modern warfighter. In many applications liquid cooling is replacing air flow through chassis for both thermal and environmental benefits. Although liquid cooled solutions prevent contaminants from being introduced into the electronics cage, it does introduce leak risks, particularly with line replaceable units that are consistently swapped in and out of the card cage. For this reason it's preferred to attach the liquid cooling loop externally to the card cage and often at the base. As a result, there are several components in the thermal path that contribute to the overall thermal resistance including the thermal interface material, conduction card, wedge lock interface, and card cage wall. This paper outlines a series of passive thermal improvements which are easily integrated into legacy, or existing, systems and can provide a 3-4x increase in dissipated power.

The first area for improvement is the conduction card. Heat pipes may be incorporated into conventional conduction cards to significantly reduce thermal gradients as heat is transferred from the electronics to the wedge lock connection at the card cage wall. Likewise, heat pipes may also be used to enhance the thermal performance of the card cage wall by embedding them directly into the chassis itself, or by bolting on heat pipe embedded heat spreaders. Both solutions are particularly amenable to retrofitting and new designs. Embedding heat pipes within the conduction cards and cage provides significantly higher effective thermal conductivity with minimal size, weight, and cost consequences, especially when compared to copper and annealed pyrolytic graphite alternatives.

KEY WORDS: thermal, embedded computing, heat pipe, electronics cooling, high thermal conductivity

NOMENCLATURE

A	area, m ²
t	thickness, m
T	temperature (°C)
k	thermal conductivity (W/m-K)

Greek symbols

ΔT temperature delta (°C)

INTRODUCTION

Embedded computing systems continue to be extremely valued by the military for wide spread use in the field. As the demand for highly capable electronics increases, program requirements are dictating cutting edge computing architecture whose performance is bottlenecked by thermal management.

Power density is ever increasing on these boards with increased performance from shrinking chip die sizes. [1] The chips are very sensitive to heat, with max case constraints providing little margin over the 71°C ambient temperatures in which the systems are designed to operate per Mil-STD-810G, Method 501.5. [2] Military standards also necessitate a sealed ruggedized enclosure to protect against salt-fog and dust, putting a priority on effective conduction spreading to move the waste heat from the chips through several interfaces to reach an exterior air or liquid heat sink on the chassis wall.

Engineers must also minimize size, weight, and power (SWaP) when designing electronics cooling systems. [3] Again, efficient spreading is needed to optimize surface area within geometry and weight restraints. Natural convection heat sink solutions are oftentimes preferred to forced air convection or liquid cooling. The latter two options allow for smaller, lighter heat sinks with lower thermal resistance, but come with risk-adders such as dust-clogged fans, power loss, and liquid line leaks. The preference for natural convection puts even more strain on a limited temperature margin. A tight cost budget may also eliminate heavy copper or pyrolytic graphite in favor of a less conductive, dense, and expensive material such as aluminum.

Thermal management systems must juggle all of these requirements simultaneously, which proves difficult as many design criteria are often working against one another. Heat pipes have been identified as an effective tool for achieving many of these goals in diverse applications. A diagram of heat pipe operation is shown in Figure 1. Heat pipes are passive two-phase heat transfer devices that attain low thermal resistance by exploiting the thermophysical properties of a working fluid. The working fluid operates inside a sealed envelope, typically copper for electronics cooling applications. A wick structure lines the interior wall to pump the liquid along the length of the pipe via capillary force. The copper tube is drawn under a vacuum, thus the fluid inside the pipe exists at the boundary between its liquid and vapor state,




Figure 1: Heat Pipe operation. Left: axial cross-section view of fluid movement between the evaporator and condenser. Right: Radial cross-section view of the materials that comprise a heat pipe.

known as the saturation curve. Any amount of heat input through the copper wall vaporizes the liquid that is contained in the wick structure in the region known as the evaporator. The higher local pressure of this region drives the vapor and its sequestered heat to the colder, lower pressure region of the pipe known as the condenser. Here, the vapor condenses into liquid on the copper wall, passing off its latent heat of vaporization in the process [4]. The wick structure then pumps the liquid back to the evaporator to absorb more heat. This process loops continuously as long as a temperature gradient exists between two points along the length of the heat pipe. A well-constructed heat pipe will have just a 2-5°C temperature drop across its length.

The working fluid is selected based on the operating temperature range dictated by the application. For many electronic cooling applications, water is the primary choice. This is because operating temperature range of a copper water heat pipe matches nicely with the allowable operating temperatures of many electronics. Typically, heat pipe operation is driven by the capillary limit, defined as the maximum power for which the wick structure's capillary force can overcome the sum of the liquid, vapor, and gravitational pressure drop. [5,6]

Heat pipes can move significant power across long linear distances in embedded computing devices and are much more effective at spreading heat compared to bare metals due to the two phase heat transfer within the pipe. The effective thermal conductivity of a heat pipe can be several orders of magnitude higher than aluminum (167 W/m-K) and copper (380 W/m-K), in the range of 10,000-100,000 W/m-K depending on length. Large distances from chip cases to chassis walls are very challenging for solid metal conduction. However, the heat removal is much more manageable with heat pipes due to their two-phase heat transfer despite acting like any other solid-solid conduction interface. Savings in this area of the thermal path allows for weight and temperature margin to be distributed to other components of the cooling assembly.

Heat pipes are also very flexible from a design integration perspective. They can be flattened or bent along 3 axes to conform to most form factors. This makes them very adaptable for use in multiple stages of the thermal network. Most chassis configurations consist of electronics boards

mounted to conduction trays that are thermally and mechanically locked along the edges into in card guide arrays on the chassis wall. The dense packing of electronics boards does not allow for local, dedicated heat sinks in close proximity to the heat generating components. Heat pipes can be embedded in the conduction card without increasing assembly size to route heat from the chip locations to the card edge to promote heat transfer. Now the thermal sacrifices made in favor of using a lighter and cheaper base material like aluminum can be regained with heat pipe integration, all with negligible change in weight. They can also enhance spreading through the base of a fin stack on the chassis wall or conduction from the card guide rails to the wall that interfaces with a liquid cooling line. Separate heat pipe embedded plates can be bolted onto the wall separately, making them useful for integrating into legacy systems and future designs alike. A heat pipe embedded aluminum plate can have an effective thermal conductivity ranging from 550 to 1,200 W/m-K. This 3-6x improvement in thermal performance combined with minimal weight gains makes heat pipe solutions ideal when designing with SWaP in mind.

The following sections will illustrate the benefits of heat pipe integration with real-world case studies showing the design challenges that can be overcome with heat pipe embedded thermal solutions.

ENHANCING THE CONDUCTION CARD

Heat generated on electronics boards in ruggedized enclosures typically pass through multiple thermal components before reaching the ultimate heat sink. The initial conduction spreading off of the chip itself is a critical step in the process because high localized heat fluxes will result in the highest temperature gradients. Board components are thermally joined to a conduction card via a thermal interface material, such as a gap pad. The conduction card must efficiently spread the heat from the component location to the edge, where it interfaces with the chassis rails and held in place with a wedge-lock clamp to provide a thermal and mechanical joint fixture.

Thin conduction frames are preferred because they allow for more electronics boards in a smaller, lighter overall chassis package. A thin form factor will also have a higher spreading resistance, which compounds the issue of long linear distances

between the components and cooling rails. Computing is constantly progressing in terms of power and capabilities but size is standardized, forcing the need for creative solutions since adding more thermal material is not an option.

When engineering effective cooling solutions, design teams will evaluate baseline all-aluminum cards to determine where the hot spots are occurring and how much improvement will be needed. Board components like CPUs and FPGAs typically have max case temperatures anywhere between 85-100°C, occasionally other supporting chips may be capable of withstanding case temperatures up to 120°C. Considering they operate in ambient temperatures from 55°C or higher, the baseline solution is usually not robust enough.

Designers often place the high power components near the edge of the conduction card to reduce thermal gradients within aluminum frame. The high heat flux has a short path to the card edge, but is still relatively high which increases ΔT through the thermal mechanical interface. The clamps used to hold the cards to the chassis often result in a 70:30 split of heat rejection, which compounds the thermal bottleneck at the clamping interface [7]. The benefit of adding heat pipes into the conduction cards can be two fold. First, the conduction gradient from the heat source to the card edge is reduced. Secondly, the heat can be spread along the edge of the card, thereby reducing the heat flux through the joint and reducing the ΔT associated with the multiple metal to metal interfaces at the joint. A sample conduction card with integrated heat pipes is shown below in Figure 2.



Figure 2: Standard 3U-size conduction card with integrated heat pipes. The heat pipes are able to take waste heat from chips anywhere on a board and spread it as uniformly as possible along the card edge to reduce the heat flux into the cooling rails.

Shown in Figure 3 is thermal simulation of a 6U aluminum conduction card with liquid cooling at the edge rails of the chassis. Heat loads are applied on various pedestals that are located all over the board. Contact resistances are applied at the interfaces between components like the card, chassis rail, and wedge lock to accurately capture the thermal path from the card to the cooling loop. The cooling loop wall is assigned a fixed temperature based on the maximum ambient dictated from the design requirements.




Figure 3: Thermal simulation of an aluminum conduction card. Areas of the card with the most difficult conduction path to the cold edge rails could not keep critical components in safe operating temperature ranges. The max temperature recorded on the card was 137°C.

In this scenario, a majority of the components were failing thermally as temperatures were 110°C and higher in a large region of the board. As expected, the center of the card is heating up due to the long conduction distance from the edge cooling rails. The card has variable thickness, with the greatest thickness occurring where the pedestals are machined. However, most other regions of the board that were key conduction paths to the edge are as thin as 0.070" or even less.

With the known heat loads, max case temperatures and the card's thermal hot spots, heat pipes may be designed for integration into multiple areas of the board. The flexibility of heat pipe design is shown here as 3mm diameter heat pipes are flattened to 0.070" to conform to the thickness of the conduction frame and bent to avoid standoffs and through holes on their way to running along the edge of the card for isothermal heat dissipation into the cooling rail. A total of 6 heat pipes were implemented without increasing weight or thickness of the card. The pipes can be permanently fixed in the card with solder or epoxy to provide a reliable thermal interface that also is robust enough to handle other mechanical requirements such as shock and vibration.

A follow up conduction model is simulated with identical boundary conditions as the baseline model, except this time the heat pipes are an active contributor to the heat removal. The heat pipes are separated into 2 solid materials: the wall material which accounts for the thermal resistance of the heat pipe envelope, wick structure, solder interface with the frame, and evaporative/condensation heat transfer; and the vapor space, which is assigned a high thermal conductivity value to achieve equivalent solid conduction performance as the real-world thermodynamic cycle. Figure 4 shows the results of the heat pipe embedded model with the previous baseline model shown for comparison. The maximum temperature is reduced by 28°C, which is a 54% reduction in total ΔT , and the card itself is much more isothermal.

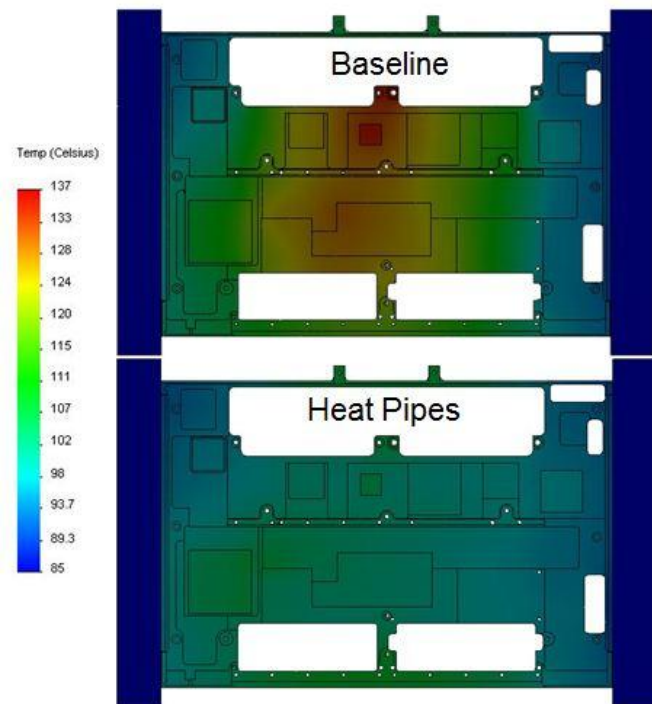


Figure 4: A comparison of thermal models between an aluminum conduction card (TOP) and the same card with heat pipes embedded in the aluminum (BOTTOM). Max temperature was reduced from 137°C to 109°C with embedded heat pipes.

A bare aluminum conduction has a thermal conductivity of 167 W/m-K, but the same aluminum card with embedded heat pipes may have a bulk conductivity of 500-1,200 W/m-K, with higher local conductivities in the area of heat generation where the heat pipes are installed. This is also higher than copper cards, which have thermal conductivity of 380 W/m-K and are much heavier and more costly to machine.

For this case, heat pipes were able to surpass thermal goals without requiring big changes to the mechanical or electrical design that was already in place. No active power is required for operation and each heat pipe is reliable enough to last up to 20 years in the field, likely outliving most other assembly components. This however is just one, although very crucial, step in the thermal resistance network. There are still other interfaces between here and the ultimate heat sink. It is also not the last in which heat pipes can aid thermal performance.

ENHANCING THE CHASSIS

The electronics chassis that houses an array of VME or VPX cards provides a barrier between the extreme outside elements and the sensitive components. An example is illustrated below in Figure 5. Reliability and risk management usually dictate that any heat sink, air or liquid cooled, must act on the exterior wall of the housing. The previous section discussed the thermal path from the board to the chassis rails. This section will focus on how heat pipes can be implemented to spread heat evenly through the chassis walls to interface with the next higher level cooling assembly.



Figure 5: Boards mounted to conduction cards are slid into chassis and locked in place with wedge-locks. The heat must be spread from the rails along the walls to effectively use the natural convection fins that are machined or cast on the external faces.

Heat pipe implementation into the chassis can follow a similar approach to the conduction cards by enhancing conduction spreading based on where hot spots are being observed and where the heat sink is located. Using the example from Figure 5, it is seen that the top face of the enclosure accounts for roughly half of the available heat sink volume. However, the heat at the rails is in closest contact with the other half of the volume, split between the two vertical walls. Normal solid metal conduction would incur huge temperature gradients before the fins on the top of the chassis could dissipate the heat to the air.

In similar cases, liquid cooling may only be accessible on the bottom base plate. In both of these scenarios, heat pipes can be embedded in the chassis to bypass the thermal penalties one would get with typical solid conduction. Figure 6 below shows an example in which straight heat pipes are vertically integrated into the chassis wall to move the heat to the proper enclosure face for heat removal.




Figure 6: Heat pipes can be fully embedded in chassis walls to move heat from the horizontal rails down to a liquid cooled base plate.

By allowing the heat pipes to do the bulk of the heat transfer, the metal chassis can be directed for ruggedization and does not need to be tasked with thermal management by conduction, allowing for a lighter and thinner chassis.

While fully embedding heat pipes in the walls is the preferred option for reducing size and weight from the start, this option would result in costly redesign iterations for legacy systems. Unlike the conduction cards inside the chassis, the external chassis walls usually have more available volume. Thus, separate bolt-on heat pipe embedded plates become attractive low-cost options for achieving thermal goals in existing assemblies.

Below in Figure 7 is a cutaway view of an aluminum card cage assembly for housing electronics cards of varying power mounted on the internal rails. The medium for cooling is a liquid cooled base. A bolt-on heat pipe embedded plate was considered to provide a robust thermal path between the rails and the base.




Figure 7: Cutaway view of a ruggedized electronics box that uses a heat pipe spreader plate to enhance conduction to a liquid cooled base plate.

For this simulation, heat loads were applied directly to the card rails and a fixed temperature was assigned to the bottom surface of the base to represent the coolant loop. The maximum card rail temperature was known from the design requirements. Initial thermal modeling without the spreader plate shows temperatures exceeding the maximum allowed on the rails. The simulation provided two options for moving forward to a successful design: increase the wall thickness of the aluminum to lower spreading resistance and incur significant weight penalties in the process, or use a separate bolt-on heat pipe plate without changing the chassis geometry. Advanced Cooling Technologies designed a bolt-on spreader plate to meet thermal design goals. The comparative thermal models are shown in Figure 8.

Many of the chassis card rails were exceeding maximum allowable temperatures due to the long path for the waste heat to take to the heat sink. The bolt-on spreader solution was able to drastically reduce temperatures by allowing the heat to bypass conduction through the chassis wall. It was determined that the wall thickness required to maintain safe operating temperatures would have added 10lbs of additional weight to the chassis. The heat pipe spreader is able to achieve the same

result with a solution that is under 5lbs total per chassis by bolting 2 plates onto the external card rail walls.

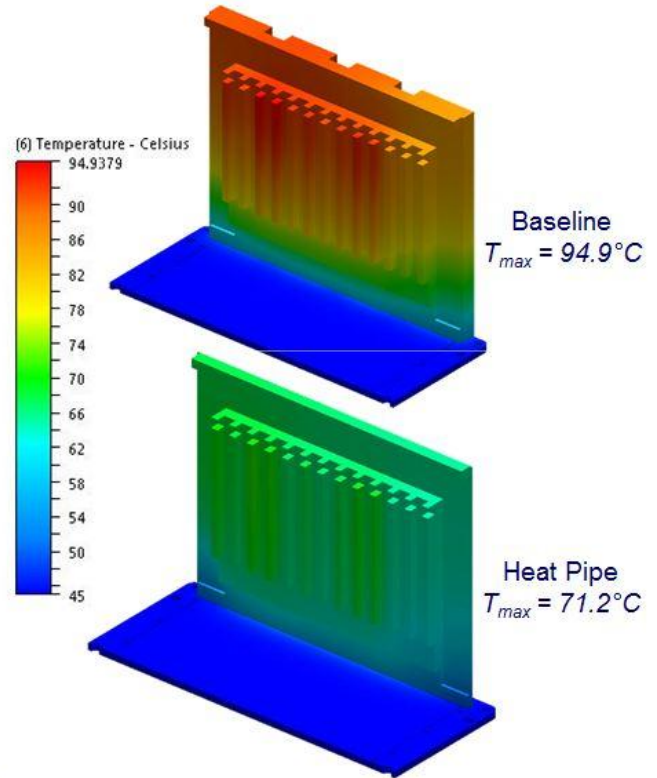


Figure 8: Comparison between an aluminum chassis (TOP) and the same chassis with a heat pipe spreader plate attached to the external wall (BOTTOM). Conduction gradients through the wall to the liquid cooled base are greatly reduced.

CONCLUSIONS

Heat pipes provide multiple approaches for enhancing the thermal performance of ruggedized electronics chassis. Standardized configurations of electronics cards in chassis are not optimally set up for uniform spreading to occur throughout the heat sink area. Direct heat pipe integration or heat pipe bolt-on solutions give suppliers great flexibility to meet thermal requirements without compromising other design goals like size, weight, or device capabilities.

Effective cooling solutions are heavily relied upon for keeping today's computing equipment operating in the toughest environments. The performance of these computing systems is largely thermally limited, making any temperature margin extremely valuable. Ruggedized systems in particular pose unique difficulties for constraining SWaP in military electronics chassis. Integrating heat pipes is a proven method for enhancing heat removal in virtually all stages of the thermal path to stay on top of these challenges.

References

- [1] D. Turner, "Case Study: Designing for Rugged Thermal Management to Meet Military Needs," Military Embedded Systems, 9 December 2010. [Online]. Available: <http://mil-embedded.com/articles/case-management-meet->

- military-needs/. [Accessed 11 January 2013].
- [2] United States Department of Defense, *Mil-STD-810G, Department of Defense Test Method Standard Environmental Engineering Considerations and Laboratory Tests*, October 31 2008.
- [3] M. Smith, "Case Study: Rethinking Rugged Subsystem Computing Design to Accomodate Military Budget Cuts," *Military Embedded Systems*, 4 September 2012. [Online]. Available: <http://mil-embedded.com/articles/case-military-budget-cuts/>. [Accessed 11 January 2014].
- [4] Faghri A., *Heat Pipe Science and Technology*, Taylor & Francis, Washington, DC. 1994
- [5] S. Garner, "Heat pipes for Electronics Cooling Applications," *Electronics Cooling*, September 1996. [Online]. Available: <http://www.electronics-cooling.com/1996/09/heat-pipes-for-electronics-cooling-applications/>. [Accessed 11 January 2014].
- [6] D. Reay & P. Kew, *Heat Pipes: Theory, Design, and Applications*, 5th Edition, Burlington, MA, 2006.
- [7] "Rugged Printed Circuit Board Hardware" Birtcher Products, 2010, Vol. 2, p.81