

Low Profile Heat Pipe Heat Sink

& Green Performance Characterization for Next Generation CPU Module Thermal Designs

Marlin Vogel, Guoping Xu*, David Copeland*, Sukhvinder Kang**, Brad Whitney**, George Meyer***, Kenya Kawabata****, Matt Connors*****

Electronic Cooling Solutions and *Sun Microsystems and **Aavid Thermalloy
and *** and Celsia Technologies**** Furukawa and *****Thermacore
email: mvogel@ecooling.com

Abstract

Increasing thermal demands of high-end server CPUs require increased performance of air-cooling systems to meet industry needs. Improving the air-cooled heat sink thermal performance is one of the critical areas for increasing the overall air-cooling limit. One of the challenging aspects for improving the heat sink performance is the effective utilization of relatively large air-cooled fin surface areas when heat is being transferred from a relatively small heat source (CPU) with high heat flux. Increased electrical performance for the computer industry has created thermal design challenges due to increased power dissipation from the CPU and due to spatial envelope limitations. Local hot spot heat fluxes within the CPU are exceeding 100 W/cm², while the maximum junction temperature requirement is 105 C, or less.

The CPU power dissipation continues to increase and the number of CPUs per server continues to increase for next generation servers. This has resulted in increased data room energy costs associated with supplying additional power to the server, and also cooling the server. Typically in the past, if two heat sink technologies met the thermal performance requirements along with meeting the reliability performance requirements, the least expensive technology would be utilized. In the future, heat sink thermal performance specifications will consider including the impact of energy cost savings attained through reduced server air flow rate requirements if utilizing a superior heat sink technology warrants a potential increase in heat sink cost.

Keywords

heat sink, heat pipe, vapor chamber, embedded heat pipe, data center, energy costs

1. Introduction/Background

In order to meet the next generation CPU thermal requirements with a phase change heat sink, two heat sink technologies and their associated prototypes will be described. Each of the heat sink technologies utilize internal liquid-to-vapor phase change [1] to efficiently spread the local CPU power to the air-cooled fin structure. The two passive phase change heat sink technologies are: multiple embedded heat pipes; and a vapor chamber / heat pipe design. CFD analysis results indicated that an optimized all-metal heat sink would not meet the sink-to-air thermal resistance while also meeting the pressure loss and maximum mass requirements.

Analysis was carried out with a commercially available computational fluid dynamics software [2] to optimize an all-metal heat sink design for System "A". The optimized heat sink had a copper pedestal, 4mm thick copper base, and 22mm tall aluminum fins with a fin spatial density of 23 fins per inch. The analysis results showed that the minimum CPU sink-to-air thermal resistance attained is 0.12 C/W while meeting the specification requirements shown in 1.1.

1.1 Heat Sink Thermal Design Requirements

3-sigma high sink-to-air thermal resistance: 0.09 C/W

CPU heat source size: 25mm x 25mm

heat sink air pressure loss:

200 Pa (System "A")

100 Pa (System "B")

heat sink air flow rate:

60 cfm (System "A")

60 cfm (System "B")

air flow direction:

bottom-to-top, opposite to gravity (System "A")

front-to-back, perpendicular to gravity (System "B")

heat sink spatial envelope:

35mm height x 300mm width x 220mm flow length
(System "A")

37mm height x 312mm width x 83mm flow length
(System "B")

mass:

1800 grams (System "A")

700 grams (system "B")

altitude: sea level

heat pipe-orientations: (see Figure 1.1)

non-gravity assisted fluid return (System "A")

opposite to gravity within lower region, gravity assisted

within upper region fluid returns (System "B")

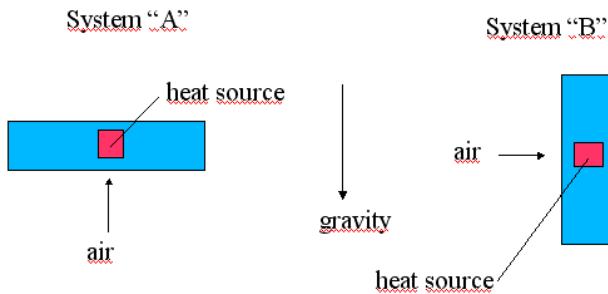


Figure 1: Heat pipe orientation and air flow direction for System "A" and System "B"

1.2 Description of Prototypes

Embedded Tower Heat Pipe Heat Sink designs:

The Supplier A embedded heat pipe [1], [2] heat sink prototype is shown in Figures 2 and 3. The design was optimized by the supplier through the use of internally developed design tools as well as a commercial CFD software tool. The prototype supplier indicated that the performance advantage comes from design methods that balance the internal and external heat pipe geometry to minimize the intrinsic temperature drop in the heat pipes to distribute the heat over the base of the heat sink and the joining processes to minimize the interfacial temperature drops to get the heat into and out of the heat pipes. Other embedded designs have a lower performance because they

do not successfully achieve this balance. The Supplier B embedded heat pipe heat sink is shown in Figures 4 and 5.




Figure 2: Supplier A Embedded Heat Pipe Heat Sink (System "A") – top view.




Figure 3: Supplier A Embedded Heat Pipe Heat Sink (System "A") – bottom view.




Figure 4: Supplier B Embedded Heat Pipe Heat Sink (System "A") – top view.




Figure 5: Supplier B Embedded Heat Pipe Heat Sink (System "A") – bottom view.

Vapor Chamber Heat Pipe Heat Sink design:

The Supplier C vapor chamber [3], [4], and [5] heat pipe heat sink is shown in Figures 6 and 7. The vapor chamber is a 3-dimensional heat pipe located in the heat sink base and is a relatively new technology that became commercially available during the mid-1990s, as compared to traditional unidirectional heat pipe technology that has been available for over 25 years. An aggressive development effort was carried out by the prototype supplier which allowed the wick thermal resistance to decrease by 50%. This provided a competitive edge over other heat sinks that incorporated vapor chamber technology. Vapor chamber allows consistent extremely flat (no gaps) interface to heat sink. Pedestals designed to be included in the vapor space of the heat sink. This prevents having to conduct through large amounts of copper before being dissipated. Power is scalable to higher levels. With additional power input, a vapor will increase in resistance less than a conventional heat sink. Compared to embedded heat pipe heat sinks, vapor chamber heat pipe heat sinks are much less stiff and considerably weaker. In order to prevent excessive flexing and possible deformation, a stiffener was added to the bottom of the heat sinks. This featured slots in the air baffle which permitted a small amount of airflow to cool components on the board beneath the heat sink. The Supplier D vapor chamber heat pipe heat sink is shown in Figures 8 and 9.




Figure 6: Supplier C Vapor Chamber Heat Pipe Heat Sink (System "A") – top view.




Figure 7: Supplier C Vapor Chamber Heat Pipe Heat Sink (System "A") – bottom view.




Figure 8: Supplier D Vapor Chamber Heat Pipe Heat Sink (System "A") – top view.




Figure 9: Celsia Vapor Chamber Heat Pipe Heat Sink (System “A”) – bottom view.

1.3 Data Center Cost Characterization

Heat sink designs which exhibit sink-to-air thermal performance margin with respect to the specification requirement will potentially allow the data room energy costs to be reduced due to requiring a reduced air flow rate. The reduced air flow rate combined with the inherently reduced pressure loss allows required electrical input power for the systems fans to be decreased. Aero power conversion efficiency for present server air movers ranges from ~15% to 40%, yielding a heat dissipation load on the data room which is equal to 60% to 85% electrical input power required by the server’s air movers. Reducing fan aero performance requirements allows energy cost savings to be attained through reduction of the fan electrical input power and through to reduction of the heat dissipated by the fans resulting in a reduction of the heat load on the data room cooling system.

2. Experimentation Test Setup for Heat Sink Performance Validation

Heat sink thermal performance and pressure drop validation conducted at Sun Microsystems were performed in a wind-tunnel system as shown in Figure 10. The test setup consists of a fully ducted air flow channel, a heater block under the heat sink, an airflow test chamber to measure air flow rate, and an air-mover. The duct cross sectional area is the same as that of the heat sink to create fully ducted air flow. Two pressure taps and two thermocouples located upstream and downstream of the heat sink were used to measure the static pressure and air flow inlet and outlet temperatures respectively. The airflow test chamber located downstream of the air flow channel used calibrated nozzles to accurately measure the air flow rate. The blower is operated at the suction mode to move air from the flow channel to the air flow test chamber. The volumetric air flow rate is calculated from measured pressure difference across the nozzle




Figure 10: Performance Validation Test Stand

3. Heat Sink Performance Test Results & Discussion

All of the shown prototypes met the thermal specification requirements. Each of the suppliers provided several samples that met the initial required sink-to-air thermal resistance (0.12 C/W, 3-sigma high). The required sink-to-air thermal resistance decreased from 0.12 C/W to 0.09 C/W in the later stages of the system development effort. This required designs that utilized embedded heat pipes to quickly carry out a significant development effort in order to meet the finalized sink-to-air thermal resistance. As a result, the number of samples submitted for validation was limited to a small sample size for the embedded heat pipe heat sinks. Six additional suppliers were provided the thermal specification and either indicated that they could not meet the spec requirements, or provided samples which did not meet the spec requirements.

The heat sink pedestal which interfaced with the CPU has a spatial envelope of 55mm x 55mm x 10mm for System A. The limited spatial envelope for the pedestal presented a significant design challenge for routing the heat pipes embedded in the base into the pedestal and minimizing the pedestal material thickness located between the CPU and the heat pipes. In addition to the CPU pedestal, System “A” contained two pedestals that interfaced with multiple components, that had a total power dissipation equal to 25% of the power dissipated by the CPU. It was found that when powering on the non-CPU components attached to the non-CPU pedestals yielded an effective increase of approximately 0.003 C/W for the vapor chamber design CPU sink-to-air thermal impedance, while yielding an approximate increase of 0.1 C/W for the embedded heat pipe design.

A significant development effort was carried out on the System “B” heat sinks by the suppliers that provided the above System “A” samples. However, the development effort for the System “B” heat sink design and the associated validation test results were not finalized prior to ending the development program, therefore the reported results are for System “A”, only.

4. Data Center Cost Trade-Off Results

The design complexity of embedding heat pipes into the System “A” pedestal, along with the adverse impact of the non-CPU component power dissipation on the effective CPU sink-to-air thermal impedance allowed the vapor chamber design to attain a significant thermal performance

margin as compared to the embedded heat pipe design. This potentially allows the required server air mover flow rate to be reduced for systems containing the vapor chamber heat sinks as compared to the embedded heat pipe heat sink design.

Today's server air mover power supply requirements are approximately 5% to 15% of the total power supplied to the server. The higher air mover power is required at the maximum ambient air temperature supplied to the server. The relationship between data room energy cost and the server air mover power requirements for continuous operation is shown in Eq. 1 and 2. Eq. 1 represents the energy cost associated with the electrical power supplied to the server air mover. Eq. 2 represents annual data room energy cost associated with cooling the power dissipated by the server air movers.

$$(Eq. 1) \quad ESPC = SP \times V \times \Delta P \times (1/\eta) \times C \times CR \times NH$$

where,

- ESPC = electrical power supply cost [=] \$
- SP = server supplied power [=] Watts
- V = ratio of required air volumetric flow rate to SP
[=] non-dimensional
- C = energy conversion constant = 0.117
- ΔP = required air mover pressure head [=] in.H2O
- η = ratio of air mover output aero power to supplied electrical power conversion [=] non-dimensional
- CR = cost rate of electricity [=] \$ per kiloWatt-hour
- NH = time of operation [=] hours

$$(Eq. 2) \quad CPC = ESPC \times (1 - \eta) \times (1/ COP)$$

where,

- CPC = energy cost of cooling the heat dissipated by the air movers [=] \$
- COP = ratio of data room heat removal rate to the electrical input power required by the data room cooling system [=] non-dimensional

Assuming the industry range for the parameters shown in both Eq. 1 and 2, the combined total annual energy cost associated with the server air movers is \$17,000 to \$300,000 for a data center that is supplying a MW of power to the servers. The server air mover annual energy costs suggest that there may be an opportunity to utilize advanced heat sink technologies with an associated increased cost, and have the increased cost compensated by the reduction in the data room energy costs.

Table 1 shows the estimated time required to recover a \$100 cost difference between heat sink technologies for 4 case studies. The results are based on assuming that thermal performance characteristics of the lower cost heat sink technology is similar to the embedded heat pipe design

shown above, and also assuming that the higher cost technology has thermal performance characteristics similar to the vapor chamber design shown above. Case A represents combining the industry's highest data room cooling performance efficiency associated with traditional computer room conditioners (CRAC) with the highest server cooling performance efficiency associated with air-cooled servers, and associated lowest cost electricity rate for to the data center.

	Case A	Case B	Case C	Case D
Electricity Cost per kWh	\$0.05	\$0.075	\$0.10	\$0.15
PUE (COP)	1.3 (5.0)	1.5 (2.5)	2.0 (1.6)	3.0 (0.45)
Fan Efficiency	50%	35%	25%	15%
Module Air Flow Rate per CPU (cfm)	31 to 53	42 to 71	62 to 106	124 to 213
Air Flow By-Pass	0	25%	50%	75%
Module/Server Pressure Loss per CPU and In-Line Components	0.30" to 0.73" H2O	0.55" to 1.5" H2O	0.82" to 2.25" H2O	0.82" to 2.25" H2O
Energy Cost for Fan Electrical Input Power & Cooling of the Fan Power Dissipation per Heat Sink & In-Line Components	\$0.44 to \$4.0	\$6.44 to \$29.5	\$31 to \$557	\$299 to \$1419
Period of Time Required to Recover \$100 Cost Difference per Heat Sink	28.1 years	4.3 years	0.2 years	0.1 years

Table 1: Energy cost characterization and results

Case D represents combining the industry's lowest data room cooling performance efficiency associated with traditional computer room conditioners (CRAC), with the lowest server cooling performance efficiency associated with air-cooled servers and utilizing the highest cost electricity rate for the data center. Cases B and C are considered to represent the typical performance for the

data room cooling systems combined with the typical performance for the server cooling system.

The highest energy efficient server cooling design, Case A, assumes that the air movers exhibit the highest aero power conversion efficiency in the industry; assumes that the air is perfectly ducted through the heat sink and that the air flow rate requirement for the components located upstream and downstream the CPU is equal to the air flow rate through the CPU heat sink. Case A also assumes that the total system pressure loss is equal to the CPU heat sink pressure loss.

The lowest energy efficient server cooling design, Case D, assumes that the air movers exhibit the highest aero power conversion efficiency in the industry; assumes that 75% of the air flow rate bypasses the CPU heat sink due to a components in parallel flow, or in series air flow requiring air flow in addition to the CPU heat sink. Case D also assumes the pressure loss of the components in series air flow maximum total system pressure loss of 2.25" of H₂O, and assumes that non-CPU components pressure loss is in series air flow exhibits turbulent loss characteristics.

The results in Table 4.1 shows that for Case A, over 28 years would be required to recover a \$100 cost differential between heat sink technologies. The results also indicate that only 0.1 years would be required to recover a \$100 cost differential for Case D. Cases B and C show that the recovery time is 0.2 to 4.3 years.

4. Conclusions

All of the heat sink suppliers met the thermal specification requirements. The sink-to-air thermal requirements and development schedule were very aggressive. Several suppliers were not able to meet the specification requirements. The CPU pedestal design combined with the impact of the non-CPU pedestals in System "A" design allowed the vapor chamber designs to attain increased performance margin as compared to the embedded heat pipe design.

Utilizing the above heat sink performance characteristics to carry out data room energy cost characterization showed that for today's typical data rooms and server cooling designs the time period required to recover an assumed \$100 cost difference between heat sink technologies could be less than 1 year, potentially justifying using a higher cost heat sink technologies in future server designs. The energy cost results also indicate that there is potential to incorporate energy cost considerations into future server CPU heat sink thermal specifications. Although the this study focused on air-cooled heat sink technologies, similar cost studies may be applicable for alternative server thermal technologies that result in a reduction of required server air flow rates. Alternative server thermal technologies potentially include CPU package design, thermal interfaces, self-contained liquid-cooled CPU modules [7], [8], [9], [10], [11], [12], [13], liquid-cooled servers [14] and liquid-cooled power electronic systems [15].

References

1. M. Vogel, D. Copeland, A. Mastow, S. Kang, B. Whitney, M. Connors and Joe Marsala, "Passive Phase Change Tower Heat Sink & Pumped Coolant Technologies for Next Generation CPU Module Thermal Designs", 2008, IMAPS Advanced Technology Workshop on Thermal Management.
2. Flomerics – Flotherm Version 8.2, 2009
3. Wu, R., " Heat-pipe type radiator and method for producing the same", United States Patent 6435266, 20, 2002.
4. Cheung, C., Moore, M., Prosperi, R., "Channel Connection for Pipe to Block Joint", United States Patent Application No. 20010050165, December 13, 2001.
5. Grubb, K., "CFD Modeling of a Thermo-Base Heat Sink", <http://www.Thermacore.com/papers.htm>.
6. North, M., "Advances in Heat Sinks, Cold Plates, and Heat Spreaders", ASME Interpack Panel Session, 2005.
7. Kang, S., Miller, D. and Cennamo, J., "Closed Liquid Cooling for High Performance Computer Systems", Proc. ASME Interpack, 2007
8. Upadhyaya, G., "Active Micro-channel Cooling System for High Heat Flux Processor Cooling Applications", Cooligy White paper, 2006.
9. Upadhyaya, G., Zhou, P., Horn, J., Goodson, K., Munch, M., "Electro-Kinetic Micro Channel Cooling System for Servers", Proceedings of IITHERM 2004, Las Vegas, pp. 367-371, June, 2004.
10. Upadhyaya, G., Zhou, P., Horn, J., Goodson, K., Munch, M., "Electro-Kinetic Micro Channel Cooling System for Desktop Computers", Proceedings of SEMI-IITHERM, San Jose, pp. 26-29, March, 2004.
11. Upadhyaya, G.; Munch, M.; Peng Zhou; Horn, J.; Werner, D.; McMaster, M., "Micro-scale liquid cooling system for high heat flux processor cooling applications", Proc. Semiconductor Thermal Measurement and Management Symposium, IEEE Twenty-Second Annual IEEE, pp. 116 – 119, March, 2006.
12. Upadhyaya, G.; Zhou, P.; Goodson, K.; Munch, M.; Kenny, "Closed-loop cooling technologies for microprocessors", Electronic Devices Meeting, Proc. IEDM '03 Technical Digest, IEEE International, pp. 32.4.1 – 32.4.4, Dec., 2003.
13. Marsala, J., "Pumped liquid cooling system using a phase change refrigerant", US Patent 6,679,081, January 20, 2004.

14. M.J. Ellsworth, Jr., L.A. Campbell, R.E. Simons, M.K. Iyengar, R.R. Schmidt and R.C. Chu, "The Evolution of Water Cooling for IBM Large Server Systems: Back to the Future", 2008, Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, pp. 266 – 274.
15. D. Saums, D. Levett, J. Howes and J. Marsala, "Vaporizable Dielectric Fluid Cooling for IGBT and Power Semiconductor Applications", 2008, IMAPS Advanced Technology Workshop on Thermal Management.