ATS Server Cooling eBook

Select Technical Articles on the Application of Air and Liquid Cooling for High-Powered Computer Servers



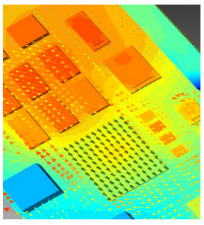
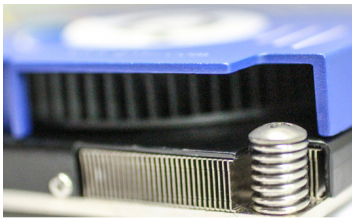

89-27 Access Road, Norwood, MA 02062 | T: 781.769. 2800 ats-hq@qats.com | www.qats.com

ATS ADVANCED THERMAL SOLUTIONS, INC.
Innovations in Thermal Management®

Table of Contents

Thermal Management of Servers (Parts I and II) 3

Effective Cooling of High-Powered CPUs on Dense Server Boards 9

Meeting the Thermal Management Requirements of High-Performance Servers 11

Application of Air Jet Impingement in a 1-U System 16

Thermal Performance of Heat Sinks with Heat Pipes or Vapor Chambers for Servers 20

Thermal Solutions for Servers

Part One – Aisles and Cabinets

A major issue with servers is the heat they generate and the heat that surrounds them. This heat has to be managed to ensure proper server performance. Even server installations with integral cooling systems may have to contend with higher power electronics or higher volumes of hot processors than they were originally designed for.

Server rooms, enclosures, racks and PCBs are all typically crowded, which poses challenges to cooling solutions at every level. And, in most cases, multiple solutions are needed. The good news is that for each of these applications, a range of cooling methods is usually available.

At the macro level these methods include cabinet positioning, ventilation fans, blanking panels to maximize cold air flow, portable air conditioning units, and skirts and barriers to direct cold air. At the micro level, liquid cooling integrated with heat sinks and spreaders provide effective chip level cooling.

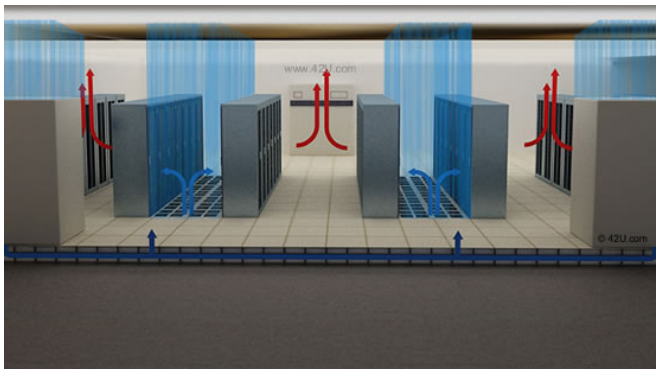

Figure 1. Hot aisle/cold aisle room layouts are widely regarded as the starting point for heat management and energy efficiency in server rooms. [42U Data Center Solutions]

For server rooms, laying out hot aisles in between cold aisles has helped solve many heat issues and remains in common use. Typically, air from the hot aisles is captured by CRAC (computer room air conditioning) units, cooled, and then distributed to the cold rows via perforated flooring systems.

Portable air conditioners, such as the ClimateCab AC units by Black Box, attach directly to racks or enclosures to directly cool components. Many sizes of combined cabinet/AC units are available, including Black Box's range of NEMA 12 server cabinets with M6 or tapped rails and various cooling capacities.

Figure 2. AC Units attached to server enclosures save money by cooling the cabinet and not an entire room. ([Black Box Network Services]

Server-filled data centers are often built as add-ons to existing facilities. The cooling system provided for these custom-added centers can be less reliable than the central HVAC installation, resulting in periodic down-times and longer-term outages. Portable air conditioners are available to provide extra cooling to heat sensitive areas during system outages or when

extreme conditions occur. These units can protect equipment – and inhabitants – from overheating when the temperature demands exceed the capacity of a building's climate system.

One example of these AC units is the Office Pro 63 portable cooler from MovinCool. This is a high end, high performance cooler. It provides 60,000 Btu/h cooling performance which can readily cool large spaces such as server/telecom equipment rooms. The Office Pro 63 has a programmable controller for continuous operation, and an automatic condensate pump.

MovinCool also produces the CM12, a ceiling-mounted packaged air conditioner. It fits into an existing drop-ceiling to save space in server rooms with limited floor space.

Figure 3. The Office Pro 63 Portable Cooler Provides 60,000 Btu/h Cooling to Lower Temperatures in Electronic Equipment Rooms. [MovinCool]

Figure 4. The 5KK30 Air-Cooled Portable Air Conditioner with 29,000 Btu/h Cooling Capacity. [Koldwave]

A range of specialty portable coolers can provide spot cooling in troublesome server hot spots. This capability can remove the need and expense of installing a permanent air conditioning system. These devices are economical and efficient, cooling only the area or object which must be cooled. Portable coolers require minimal installation and are easily moved from one room to another. Cooling air is delivered instantly.

For multiple hot spots, the air-cooled portable 5KK30 system from Koldwave delivers cooling air to different locations via three flexible air outlets. The quiet-running unit is designed for ambient temperature and is simple to install and operate. It has a cooling capacity of 29,000 Btu/h.

The engineers at Tripp Lite recommend the use of blanking panels to fill unused rack spaces. This forces cold air through equipment and prevents hot air from recirculating through open spaces. Snap-in 1U blanking panels install easily and the 1U size fills empty rack spaces evenly. Tripp Lite suggests installing brush strips, gaskets and grommets to block air leaks around cable channels and other gaps.

Another technological approach is the use of vortex enclosure cooling systems. These work by maintaining a slight pressurization inside a cabinet to keep electrical and electronic components clean and dry. Most vortex systems are thermostatically controlled to maintain enclosure temperatures within a specified temperature range.

The core of this technology is vortex tubes, mechanical devices that separate compressed gas into hot and cold streams. Air emerging from the cold end can reach -50°C, while air emerging from the hot end can reach 200°C. The tubes have no moving parts.

EXAIR Cabinet Cooler systems use vortex tube technology to create a cold air outlet flow which is pumped into an electronic cabinet. As air is pushed into the cabinet the Cabinet Cooler system also provides its own built-in exhaust so there is no need to vent the cabinet.

This creates a positive purge on the cabinet which will keep out dirt, dust and debris.

Figure 5. The ETC Dual Cabinet Cooler System Minimizes Compressed Air Use and Produces 20F Air for Cabinet Cooling. [EXAIR]

Every server installation is unique, and from its components to its cabinets to its room environment there are likely to be heat issues. The above discusses some of the more common thermal solutions in use today at the aisle and cabinet level.

References

1. http://www.42u.com/
2. https://www.blackbox.com/
3. http://movincool.com/
4. http://koldwave.com/
5. http://exair.com
6. http://www.cableorganizer.com/computer-cabinets/rack-fans.htm
7. http://www.cableorganizer.com/tripp-lite/smartrack-air-conditioning-unit/
8. http://www.airpacinc.com/blog/bid/82753/5-Simple-Server-Room-Cooling-Solutions
9. http://www.chilleddoor.com/discover-chilled-door
10. http://www.electronicsprotectionmagazine.com/main/articles/increase-rack-cooling-efficiency-and-solve-heat-related-problems/

Part Two – Racks, Cards, and Cabinets

When the components inside a server get too hot, onboard logic may turn them off to avoid damage to the server. But not all components are protected like this. Along with the effects of an unscheduled server shutdown, if and when a hot server does not power down, its output could be compromised, and its lifespan could be shorter than expected.

While much infrastructure and wide area cooling is employed in server installations, there is often a need for additional cooling of individual cards or select components on some cards. There are different ways to provide this localized cooling. Some are proven and traditional, and other ways are less conventional but have shown to be effective.

The following are some current examples of cooling solutions at the rack, card and individual component level.

Liquid Cooling for Racks

Internal

Direct Contact Liquid Cooling (DCLC) is a cooling solution from CoolIT Systems that uses the thermal conductivity of liquid to provide concentrated cooling to select surface areas, such as server racks and their components. DCLC systems can replace or minimize the need for system air conditioners and fans, and instead allow higher rack densities. Any server in any rack can be liquid cooled by DCLC.

An example of the DCLC technology is CoolIT's CHx40 Liquid-to-Liquid heat exchanger for single cabinets. The 2U CHx40 module distributes clean, treated coolant, and can manage 40 kW of processor load per rack, or 50 servers per rack.

The benefits from this system are increased density availability for CPUs and GPUs. CoolIT Systems server modules can cool any combination of CPU, GPU and memory components, with customization available for VR, ASIC, FPGA and other devices. Servers with these modules remain hot-swappable and simple to service.

Figure 1. A Liquid-to-Liquid heat exchanger module for cooling board components in cabinets. [CoolIT Systems]

The Fujitsu Liquid Loop Cooling technology uses a hybrid cooling model. To remove heat, there are six pumps for each processor, which circulate coolant over hot spots and back to an air-cooled radiator. The cooling system is sealed and maintenance-free.

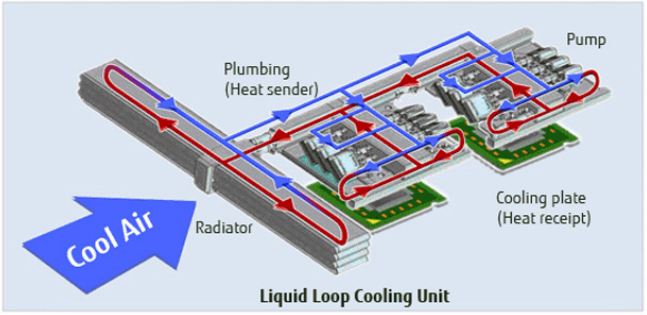

Figure 2. Liquid Loop Cooling Allows Heat Sinks to be Smaller and Located Away from Hot Spots. [Fujitsu]

In order to cool hot spots with air, heat sinks and fans must be located at the hot spots. However, this can result in components like memory or IO being further away from the processors. The space between a CPU and memory chips largely affects memory access time, degrading overall performance. In liquid cooling, heat sinks can be located away from hot spots because heat is moved via cool liquid tubes to heat sinks. This allows processors and other components like memory and IO to be close enough to reduce access time.

By not only eliminating server design restrictions, but also integrated memory controllers to the processor, memory access time is reduced to one-fifth of the previous SPARC server. The Liquid Loop Cooling system has reduced heat sink sizes. Further, this efficient cooling helps reduce the noise which may be created by lower fan rotation.

External

A not as new, but quickly developing method for thermally managing racks of server cards is by submersion cooling. Typically, a series of boards and all their components are immersed in a non-conductive liquid. The liquid absorbs the component heat and transfers it away from the card.

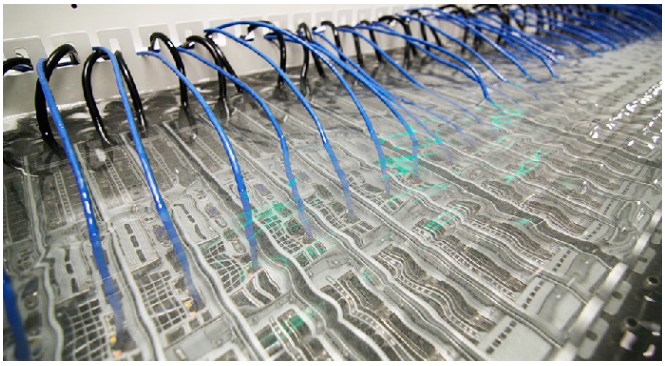

Figure 3. ElectroSafe coolant flowing over servers in a CarnoJet System also protects components from environmental contaminants. [Green Revolution Cooling]

An example is the CarnoJet System from Green Revolution Cooling which has been proven to provide an overall increase in a server's performance. The system features inert ElectroSafe coolant, which keeps server components an average of 20°C lower than obtainable by an air-cooled environment. This improved cooling allows for faster processor

clock speeds and higher density racks. Servers immersed in the coolant are protected from dust, moisture, and oxidation, which are major sources of equipment failure. The coolant is constantly filtered and in motion. This provides a rinsing effect that prevents any kind of particulate accumulation.

Other benefits of submersion cooling include cost savings on large-scale air-conditioning systems and their maintenance, and on power requirements. But despite many advantages, liquid cooling still faces obstacles to becoming more widely used. The concept itself can cause apprehension with some engineers. But at high densities of server hardware, liquid cooling may be a smart solution.

Fan Cooling

There are many active server rack accessories to help improve cooling. Among them are enclosure blowers, rack air conditioners, and cooling fans. Here we'll look at some fans.

Fans can be used to cool rack set ups where larger AC units may not be practical. Internal fans typically take up as little as 1U of rack space and can install directly into a rack or enclosure.

Figure 4. A rack mounted fan unit that provides 90 cubic feet of air per minute to cool hotspots in a server cabinet. [Data Comms Direct]

For lower density racks, fans can be a viable solution when increased cooling is needed. Fans facilitate the movement of ambient air through the rack. They can be mounted on a cabinet's ceiling, sides, and doors, or incorporated directly in the rack as slide out fan trays. To improve the effectiveness of any fans, proper airflow management is needed. This may need the use of blanking panels, gaskets or sealing tape.

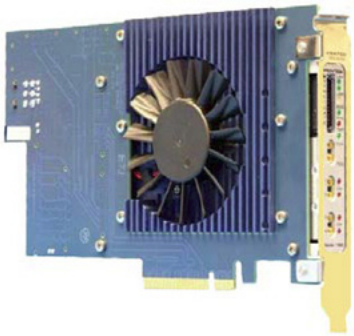

Figure 3. A fan-equipped heat sink attached through a cutout in the PCIe carrier card, removes heat effectively from the hot components on the inside surface of the attached XMC module.

Figure 5. An integral fan helps cool hot components on a PCIe card. [Pentek]

Many server processors require active heat sinks – where flowing air or liquid enhances the sink's cooling performance. A common example is incorporating a small fan within a sink's cooling fins. As an example, a PCIe carrier card, found in some server schemes, has a cutout opening directly under the warmest components.

This allows the use of a fan-equipped heat sink that removes heat effectively from those components on the inside surface of an attached XMC module. The embedded fan forces air across the fins for active cooling.

Sensors can be helpful for finding dead spots in cooling airflow. They can help determine if moving a fan or using a different fan with a higher CFM rating improves cooling results.

Since air flow is hard to visualize, the best solution is often found by trial and error.

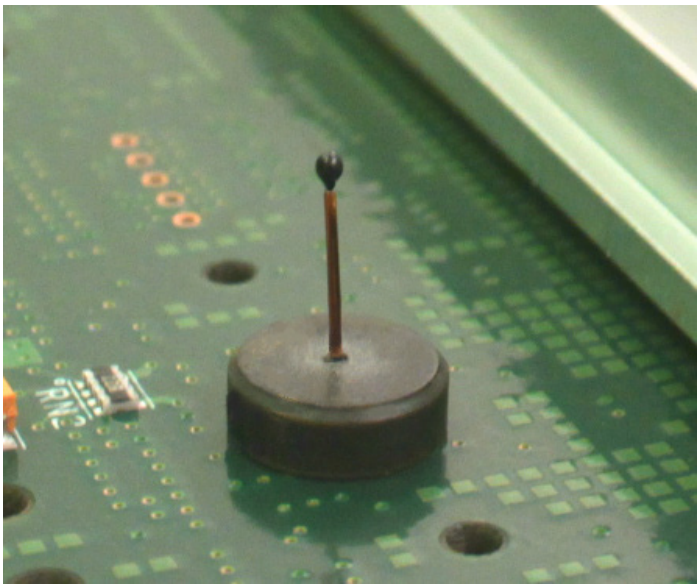

Figure 6. A Candlestick sensor measures air temperature and velocity at different locations on server cards. [Advanced Thermal Solutions, Inc.]

Component Level Cooling

A new active heat sink solution for server CPUs, including the Intel Xeon (Haswell and Skylake), AMD Opteron series and Cavium ThunderX ARM CPUs. The thermal solution integrates a heat sink, heat pipes or vapor chamber, and blower. The dualFLOW™ heat sink is designed in low profile for 1U and 2U rack servers.

It has rows of embedded heat pipes in its base to help transfer large volumes of component heat into a dense fin field.

Figure 7. A DualFLOW heat sink with heat pipes embedded in its base and a PWM-enabled blower on top. [Advanced Thermal Solutions, Inc.]

The blower on top draws air inward, across the many fin surfaces, then pushes this warmed air out and away.

The dualFLOW™ heat sink can be mounted onto PCBs using PEMs or screws inside springs. Both kinds of hardware allow adjusting so the heat sink is mounted for optimum performance. Its initial mounting force is 9.2 PSI.

A thermal grease interface assures excellent thermal transfer from the component to the heat sink's base. For added cooling performance, a PWM controllable fan can be mounted above the heat sink fin field.

References

1. https://journal.uptimeinstitute.com/a-look-at-data-center-cooling-technologies/
2. http://www.coolitsystems.com/index.php
3. http://www.fujitsu.com/global/products/computing/servers/unix/sparc/technology/ecology/llc.html
4. https://www.blackbox.com/
5. http://www.datacenterjournal.com/whats-stopping-liquid-cooling/
6. http://www.pentek.com/pipeline/22_2/cooling.cfm
7. http://www.dcdi.co.uk/product/25/19-racks/136/fans-rack-cooling/617/3u-19-rack-mounted-3-fan-unit
8. http://www.qats.com/Products/Instruments/Temperature-and-Velocity-Measurement/Sensors/Candlestick-Sensor
9. http://www.cableorganizer.com/computer-cabinets/rack-fans.htm
10. http://www.electronicsprotectionmagazine.com/main/articles/increase-rack-cooling-efficiency-and-solve-heat-related-problems/

Effective Cooling of High-Powered CPUs on Dense Server Boards

The main goal of electronics thermal management is to efficiently remove enough heat from a device's active region so that it stays within its rated temperature. Providing effective cooling presents different design challenges, not all of which involve the chip itself. Some thermal challenges are related to the system in which the chip resides. A common example is cooling a device positioned on a crowded printed circuit board (PCB). The congestion of components restricts airflow and space, which makes the use of many conventional cooling devices difficult.

Optimizing PCB for thermal management has been shown to ensure reliability, speed time to

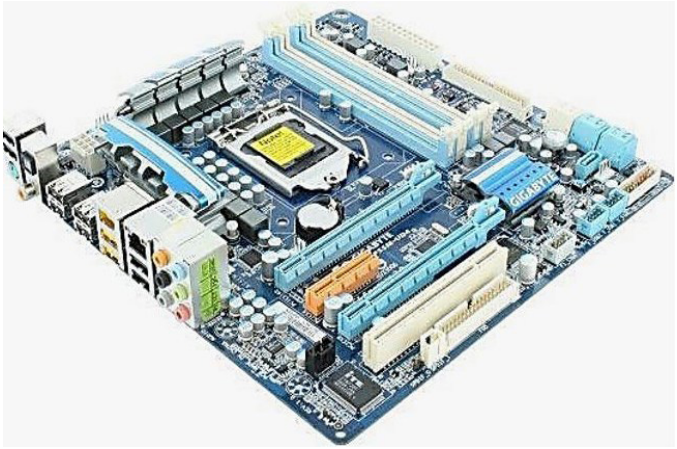

Figure 1. A dense motherboard from Gigabyte Technologies featuring an Intel P55 chipset. [1]

market and reduce overall costs. With proper design, all semiconductor devices on a PCB will be maintained at or below their maximum rated temperature. Applying thermal management can sometimes be problematic for dense boards employing fine pitch devices. (Pitch is the space between the center of one BGA ball to the center of the next one.)

But if certain layout guidelines are not followed and considerations are not given to a PCB's thermal performance, the device and the overall system can suffer from sub-par performance and reliability in the field. [2]

Today's circuit boards are often assembled with increasing density with the goal of making smaller, lighter systems, or to provide more processing power in demanding applications

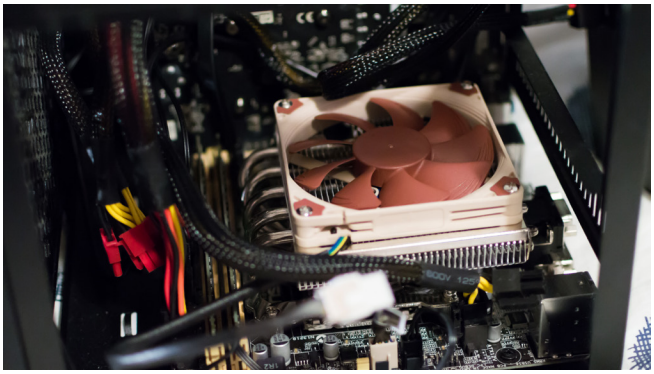

Figure 2. Crowded boards have limited space from where chip cooling air can be drawn. (Wikimedia Commons)

such as data centers and IoT (Internet of Things) applications. PCB designers must use proven layout techniques to ensure effective thermal performance for the board and its components.

Part of the trend toward higher density boards is related to the industry's adoption of increased server density. This means increasing the power of the chips, putting more chips per rack unit, and filling up the racks as much as possible. Rack power has transitioned from a normal of about 4 kilowatts to 70 kilowatts per rack.

High current electronic components like microcontrollers can generate a significant amount of heat. To keep the board temperature lower, it is usually best to mount these components near the center of the board. Heat can diffuse throughout the board and the temperature of the board will be lower.

Many components in this situation, such as GPU, will require a dedicated cooling system, such as a fan sink. But simply installing a fan sink on top may not provide the needed level of cooling. It is good practice to quantify

system flow bypass on the fan sink, and to also consider the proximity of components neighboring the fan sink. The mass airflow rate is the true measure of available coolant, along with the air velocity.

Obstructions in the intake or exhaust of the fan (e.g. neighboring components) must be carefully considered as their presence will impact the performance of the fan sink. The size and position of adjacent components can impact the fan's performance. [3]

One new and effective solution for cooling hot components on congested PCBs is the quadFLOW™ CPU cooler from Advanced Thermal Solutions, Inc. (ATS). The liquid-free cooler features a high-power blower that draws in air from four different directions. So, while proximate components may block local air in a couple of directions, the quadFLOW™ fin fields will pull in air from the other directions to make sure that the component is being cooled.

Figure 3. The QuadFlow CPU cooler draws air from four sides, passing it through cooling fins and expelling warm air. (Advanced Thermal Solutions, Inc.)

quadFLOW™ coolers are just 29 mm tall, so they will fit into standard 1-U racks and there are several options for base material (aluminum, copper, or vapor chamber) depending on performance, weight, or cost requirements. [4]

Before applying any thermal management hardware, the smartest engineering activity may be investing in various PCB design services. These include CFD studies on boards at the CAD stage to wind-tunnel testing of actual or dummy boards in conditions that simulate air distribution in real-world applications. Services are available for characterizing boards using research-quality instruments, temperature and air velocity sensors, and PCs.

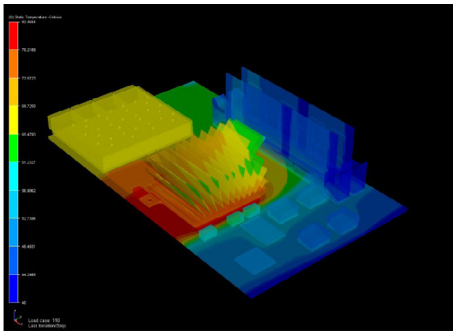

Figure 4. FloTHERM image reveals hotter and cooler regions on a PCB. (Advanced Thermal Solutions, Inc.)

Dummy or working PCBs can be tested in isolation or installed in their own packaging domain. Computational simulations can be made of engineered designs using computational fluid design packages such as 6SigmaET, FloTHERM and CFDesign.

These services are available from ATS, whose engineers can design board layouts to improve cooling airflow in dense systems. Natural airflow can be enhanced to individual hot components and to active cooling systems that rely on airflow for effective performance.

Often, these studies head off more expensive cooling solutions by showing that minor changes to component layouts or to the volume of airflow will resolve thermal problems. [5]

References

1. https://www.techpowerup.com/103375/gigabyte-unwraps-latest-p55-series-motherboards

2. https://www.embedded.com/design/configurable-systems/4395845/Ultra-fine-pitch-devices-pose-new-PCB-design-issues

3. https://www.hpe.com/us/en/insights/articles/why-youll-be-using-liquid-cooling-in-five-years-1710.html

4. https://www.qats.com/cms/2013/06/21/how-system-flow-affects-fan-sink-performance/

5. https://www.qats.com/Consulting/PCB-Board-Layout

Meeting the Thermal Management Requirements of High-Performance Servers

High-performance servers are devices specially designed to handle large computational loads, a huge amount of communication signals, fast data processing, etc. Due to their task-oriented nature, high-performance servers must have high reliability, interchangeability, compact size and good serviceability.

To achieve high computational speed, high-performance servers generally have dozens of CPUs and memory models. They also have dedicated date process modules and control units to ensure seamless communication between CPUs and parallel data processing ability. To reach higher speeds, the power dissipation of high–performance CPUs has been increasing continuously in the past decade for its use in high-performance servers.

Cooling dozens of kW servers brings a unique challenge for thermal engineers. To deal with the ever-growing high heat flux issue in high-performance servers, it will need the cooperation of electrical, mechanical and system engineers to solve the problem. The job to remove the high heat flux from CPUs to ambient requires chip level, board level and cabinet level solutions.

Wei [1] described Fujitsu's thermal management advancements in their high-end UNIX server PRIMEPOWER 2500. The server cabinet is shown in Figure 1. Its dimensions are 180cm × 107cm × 179cm (H×W×D) and has a maximum power dissipation of 40 kW. The system configuration of PRIMEPOWER 2500 is shown in Figures 2 and 3. It has 16 system boards and 2 input/output (I/O) boards installed vertically on two back-panel boards. The two back-panel boards are interconnected by six (6) crossbars installed horizontally.

To cool the electrical components inside PRIMEPOWER 2500, 48 200-mm-diameter fans are installed between the system board unit and the power supply unit. They provide forced air cooling for system boards and power supplies. In addition, six 140-mm-diameter fans are installed on one side of crossbar to cool the crossbar boards with a horizontal flow. The flow direction is shown in Figure 3. Each system board is 58 cm wide and 47 cm long.

There are eight CPU processors, 32 Dual In-Line Memory Modules, 15 system controller processors, and associated DC-DC converters on each system board. The combined power dissipation per system board is 1.6 kW at most.

Forced air-cooling technology is commonly used in computers, communication cabinets, and embedded systems, due to its simplicity, low cost and easy implementation. For high-performance servers, the increasing power density and constraints of air-cooling capability and air delivery capacity have pushed forced air cooling to its performance limit.

Figure 1. PRIMEPOWER 2500 Cabinet. [1]

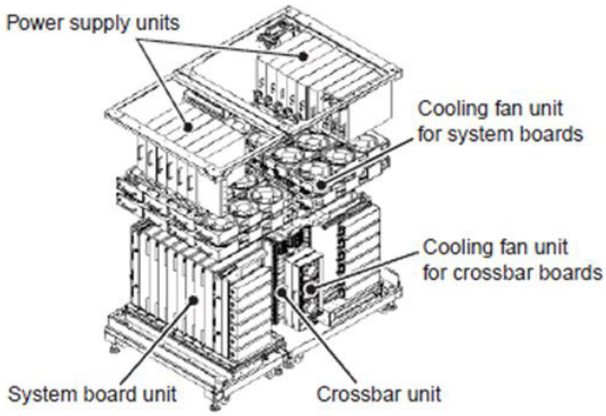

Figure 2. PRIMEPOWER 2500 System Configuration. [1]

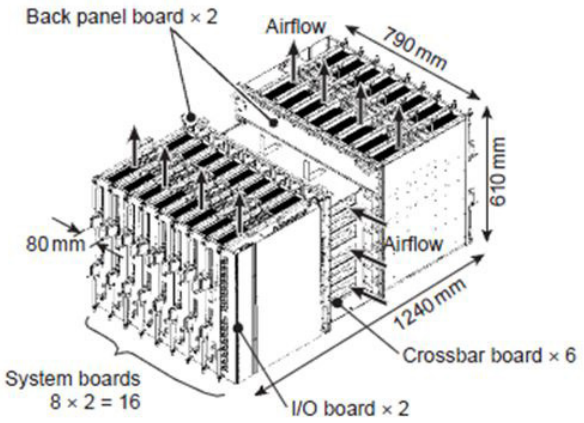

Figure 3. PRIMEPOWER 2500 System Board Unit. [1]

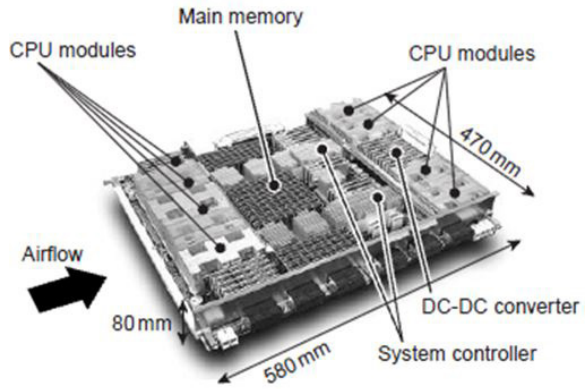

Figure 4. PRIMEPOWER 2500 System Board. [1]

For high power systems like PRIMEPOWER 2500, it needs a combination of good CPU design, optimized board layout, advanced thermal interface material (TIM), high-performance heat sinks, and strong fans to achieve desired cooling.

The general approach to cool the multi-board system is first to identify the hottest power component with the lowest temperature margin. For the high-performance server, it is the CPUs. For multiple CPUs on a system board, generally, the CPU located on downstream of a board or other CPUs has the highest temperature.

So, the thermal resistance requirement for this CPU is:

$$\mathbf{R}_{ja} = \frac{T_{\mathbf{j},\max} - T_a - \Delta T_a}{q_{\max}} \quad (1)$$

Where Tj,max is the allowed maximum junction temperature, Ta is the ambient temperature, ΔTa is the air temperature rise due to preheating before the CPU, and qmax is the maximum CPU power.

The junction-to-air thermal resistance of the CPU is:

$$\mathbf{R}_{ja} = \mathbf{R}_{jc} + \mathbf{R}_{TIM} + \mathbf{R}_{hs} \quad (2)$$

Where Rjc is the CPU junction-to-case thermal resistance, RTIM is the thermal resistance of thermal interface materials, and Rhs is the heat sink thermal resistance. To reduce the CPU junction temperature, it is critical to find intuitive ways to minimize Rjc, RTIM, and Rhs, because any reduction in thermal resistance is important in junction temperature reduction.

The CPU package and heat sink module of PRIMEPOWER 2500 are shown in Figure 5. The CPU package has an integrated heat spreader (IHS) attached to the CPU chip. A high-performance TIM is used to bond the CPU chip and IHS together, see Figure 6. The heat sink module is mounted on the IHS with another TIM in between.

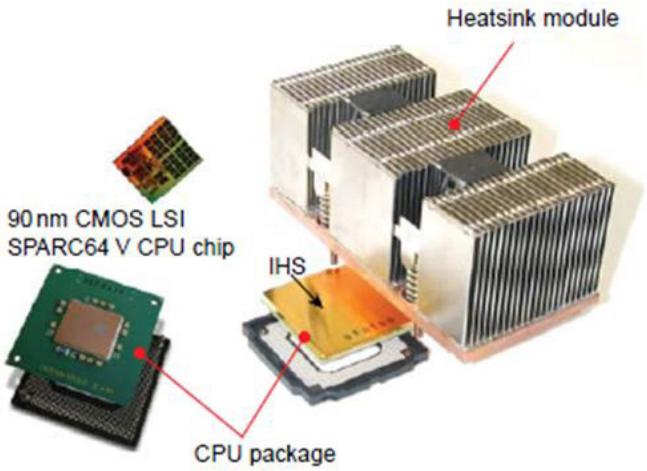

Figure 5. PRIMEPOWER 2500 CPU Package and Heat Sink Module, [1]

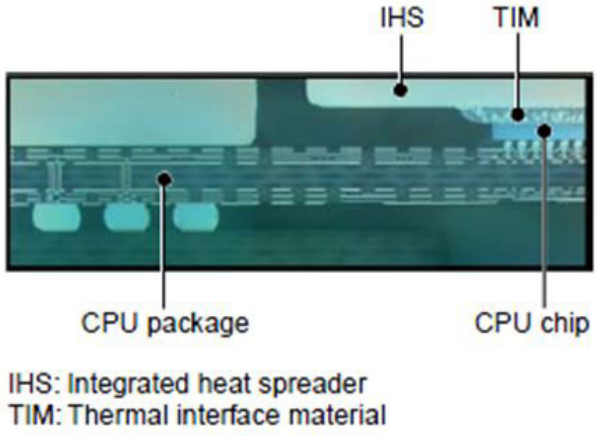

Figure 6. CPU Package. [1]

The TIM used in between the CPU chip and the IHS are crucial to the CPU's operation. It has two key functions: to conduct heat from the chip to the IHS and to reduce the CPU chip stress caused by the mismatch of the coefficient of thermal expansion (CTE) between the CPU chip and IHS. Fujitsu developed a TIM made of In-Ag composite solder for the above application. The In-Ag composite has a low melting point and a high thermal conductivity. It is relatively soft, which is good for absorbing thermal stress between the chip and the IHS.

Wei [2] also investigated the impact of thermal conductivity on heat spread performance. He found a diamond composite IHS (k=600 W/(mK)) would result in a lower temperature gradient across the chip and low temperature hot spots, compared with aluminum nitride (k=200 W/(mK)) and copper (k=400 W/(mK)). The simulation results are shown in Figure 7.

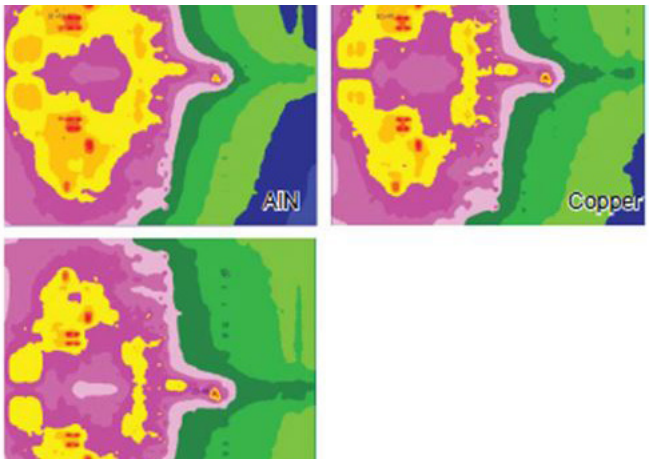

Figure 7. Heat Spreader Material Comparison [2]

In high-performance servers like the PRIMEPOWER 2500, the thermal performance gains by optimizing the TIM and the IHS are small, because they compose only a small portion of the total thermal resistance. Heat sinks dissipate heat from the CPU to air and have an important role in the thermal management of the server. In a server application, the heat sink needs to meet not

only the mechanical and thermal requirements, but also the weight and volume restraints. Hence, heat pipes, vapor chambers, and composite materials are widely used in place of high-performance heat sinks.

Koide et al [1] compared the thermal performance and weight of different heat sinks for server application. The results are shown in Figure 8. They used the Cu-base/AL-fin heat sink as a benchmark. Compared with the Cu-base/AL-fin heat sink, the Cu-base/Cu-fin heat sink is 50% heavier and gains only 8% performance.

If a heat pipe is used in the base, the heat sink weight can be reduced by 15% and the thermal performance increases by 10%. If the vapor chamber is embedded in the heat sink base, it reduces the heat sink weight by 20% and increases the heat sink performance by 20%.

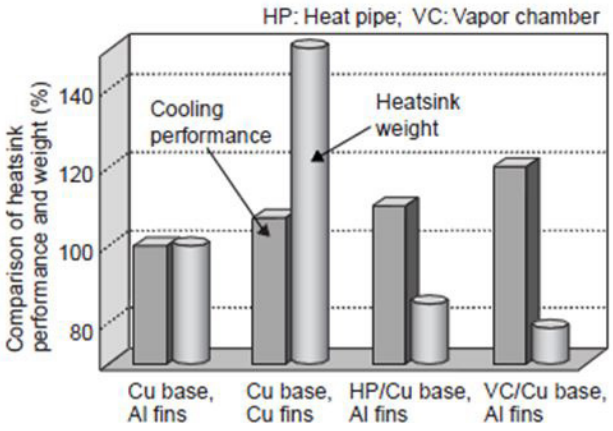

Figure 8. Thermal Performance and Weight Comparison of Different Heat Sinks.

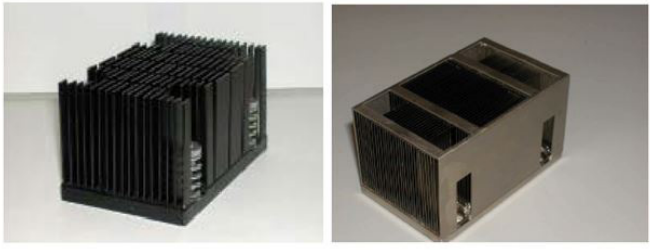

Figure 9. (a) USIII Heat Sink for Sun Fire 15K Server, (b) USIV Heat Sink for Sun Fire 25K. [3]

Sun Microsystems' high-performance Sun Fire 15K Server uses the USIII heat sink to cool its 72 UltraSparc III (USIII) processors. In Sun Fire 25K Server, the CPUs are upgraded to UltraSparc IV (USIV), which has a maximum power of 108 W. To cool the USIV processor, Xu and Follmer [3] designed a new USIV heat sink with copper base/copper fin, see Figure 9. The old USIII heat sink has 17 forged aluminum fins, the USIV heat sink has 33 copper fins. Both heat sinks have the same base dimensions and height.

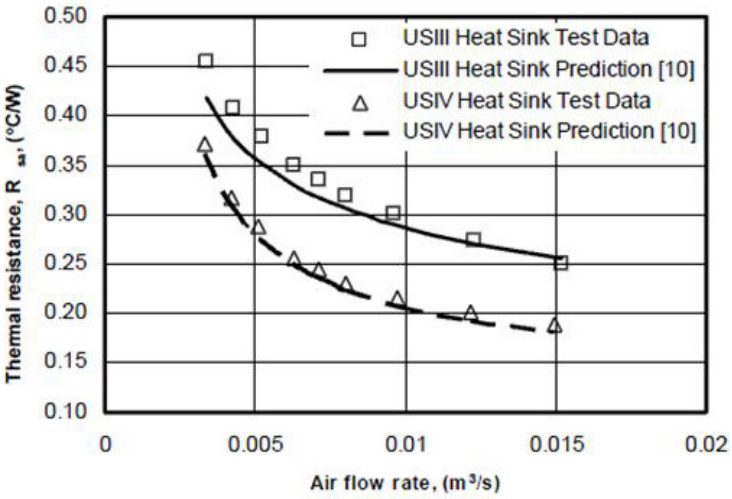

Figure 10. Thermal Resistance Comparison between USIII Heat Sink and USIV Heat Sink.

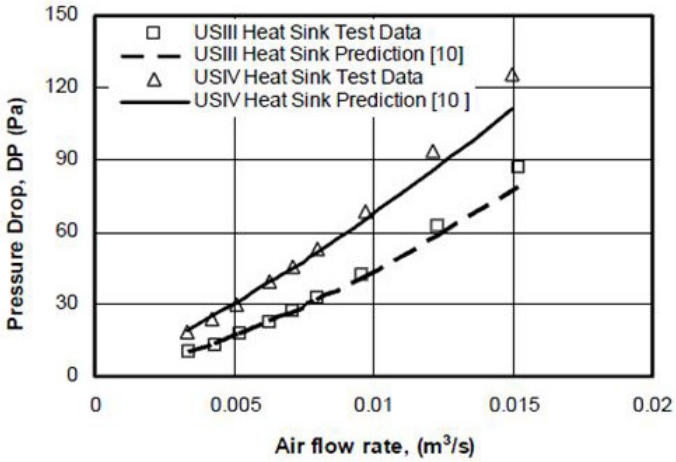

Figure 11. Pressure Drop Comparison between USIII Heat Sink and USIV Heat Sink. [3]

Figure 10 shows the thermal resistance comparison between the USIII heat sink and the USIV heat sink. The thermal resistance of the USIV heat sink is almost 0.1°C/W lower

than that of the USIII heat sink at medium and high flow rates, which is a huge gain in thermal performance. The thermal performance improvement of the USIV heat sink is not without penalty.

Figure 11 shows the pressure drop comparison between the USIII heat sink and the USIV heat sink. For the same air flow rate, the pressure drop of the USIV heat sink is higher than that of the USIII heat sink. That means the Sun Fire 25K Server needs stronger fans and better flow arrangements to ensure the USIV heat sinks have adequate cooling flow.

The design of the cooling method in high-performance servers follows the same methodology used in the design cooling solution of other electronic devices, but at an elevated scale. The main focus is to identify the hottest components, which in most cases is CPUs. Due to the extreme high power of CPUs, memory modules, heat spreaders, TIMs, and heat sinks are needed to achieve desired cooling in the server. The goal of thermal management is to find cost-effective ways to maintain the junction temperature of the CPU lower than specifications and ensure the continuous operation of the server. Wei [1] has proved a 40 kW server can be cooled by forced air cooling.

However, it requires highly integrated design and a huge amount of air flow that the 54 fans inside PRIMEPOWER 2500 can generate. In the near future, it would be very difficult for a forced air-cooling method to cool cabinets with more than 60 kW power. It would require bigger fan trays to deliver huge amounts of air flow and large size heat sinks to transfer heat from the CPUs to air, which makes it impossible to design a reliable, compact and cost-effective cooling system for the server.

We have to find alternative ways to deal with this problem. Other cooling methods, such as air impinging jets, liquid cooling and refrigeration cooling systems, have the potential to dissipate more heat. But it will require intuitive packaging to integrate them into the server system.

References

1. Wen, J., Thermal Management of Fujitsu's High-performance Servers, source: http://www.fujitsu.com/downloads/MAG/vol43-1/paper14.pdf.

2. Koide, M.; Fukuzono, K.; Yoshimura, H.; Sato, T.; Abe, K.; Fujisaki, H.; High-Performance Flip-Chip BGA Technology Based on Thin-Core and Coreless Package Substrates, Proceedings of 56th ECTC, San Diego, CA, USA, 2006, pp.1869-1873.

3. Xu, G; Follmer, L.; Thermal Solution Development for High-end System, Proceedings of 21st IEEE SEMI-THERM Symposium, San Jose, CA, USA, 2005, pp. 109-115.

Application of Air Jet Impingement in a 1-U System

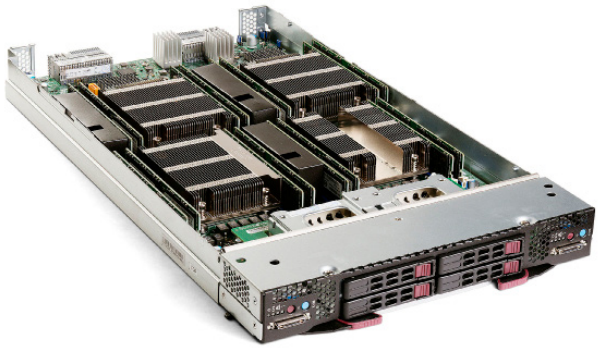

Cooling high-powered 1U servers requires new thermal management techniques, including advanced air jet impingement. (Wikimedia Commons)

As their speeds increase, the heat dissipation from high performance processors requires more innovative cooling techniques for heat removal. One such technique, jet impingement, provides one of the highest heat transfer coefficients among cooling methods. This property of jet impingement has been used by Advanced Thermal Solutions, Inc. (ATS) in a 1-U server application. Jet impingement has also been applied to ATCA chassis.

Air Jet Impingement

Cooling high-powered 1U servers requires new thermal management techniques, including advanced air jet impingement. (Wikimedia Commons)

A Ujet-1000™ 1-U chassis made by ATS was assessed in the company's thermal test lab. Four identical heat sinks were tested under conventional parallel flow and ATS' proprietary Therm-Jett™ impingement flow. The results showed a 20-40 percent improvement in the thermal performance of the heat sinks.

The Ujet-1000™ is a 1 to 2 KW 1-U chassis system (depending on component case and ambient temperature) designed for the most demanding, telecom and server applications. Lab tests demonstrated that four heat sinks on four simulated chips located on a PCB achieved 0.16 to 0.18°C/W thermal resistance. The power dissipation of each simulated component was maintained at 200 W. On the other hand, the thermal resistance of the same heat sinks with parallel flow, using the same fans, was almost 20 to 40% worse.

The new ATS Therm-Jett™ technology uses a specially made duct with an impingement plate beneath it to create jet impingement on top of the components and heat sinks. The tremendous increase in heat transfer coefficient leads to significant reduction of thermal resistance compared to the other conventional 1-U systems. A Therm-Jett™ system can be built for any specific configuration. The impingement duct is less than 5 mm thick and is located in the chassis on top of the motherboard. In addition to high heat transfer coefficient, fresh air is distributed between all heat sinks at inlet temperature.

In contrast, in conventional cooling systems, the upstream heat sinks and components receive air at inlet temperatures which are cold and gradually warm up as the air moves downstream. The increase of air temperature effectively reduces the cooling effect of the air downstream.

The other advantage of Therm-Jett™ is that there is no need to make a special duct for each heat sink, thus freeing the motherboard for other components. Even by adding ducts, other components such as memory cards, resistors and capacitors located upstream of the heat sinks on the PWB would deprive the heat sink of the flow at its most critical point, which is close to the base.

Figure 1 shows a CAD drawing of the real system under test. The cooling is provided by eight 40 mm high capacity double fans located in the midsection of the 1-U chassis. The power to the heat sinks was provided by four heaters attached below and dissipating 200 W each. The

heat dissipation of the power supply was simulated by attaching a rectangular heater strip under the power supply which dissipates about 100 W. A “U” shape frame made of aluminum was located under the hard drives. The power of four hard drives was simulated by placing a rectangular heater under the “U” shape frame which dissipates about 80 W.

Four thermocouples were placed in holes at the center of the base of the heat sink downstream. The holes were filled with thermal grease to minimize the interfacial resistance. Three thermocouples were attached to the aluminum “U” frame, and their average temperature was recorded as an approximate temperature of a real hard drive. One thermocouple was also attached to the base of the power supply to measure its approximate temperature. All temperature measurements were taken using J type thermocouples.

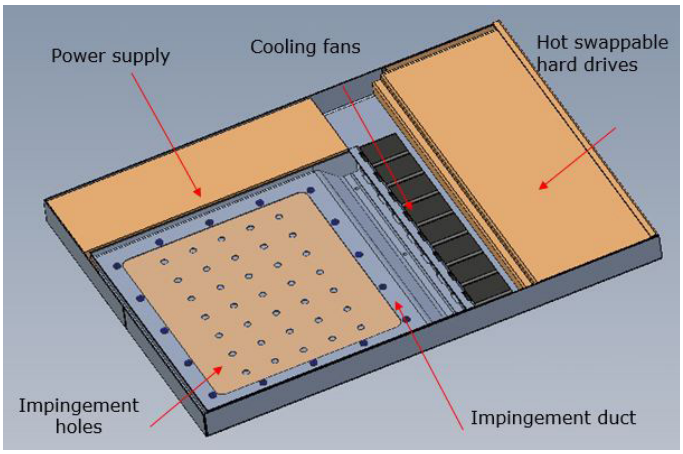

Figure 1. Schematic of the Ujet-1000™ and the Therm-JETT™ cooling duct. (Advanced Thermal Solutions, Inc.)

Figure 2 shows the exploded view of the Ujet-1000™ chassis.

Figure 3 shows a conventional cooling system for a 1-U system. In this system, the air flow from the fans is parallel to the heat sinks.

Results

Figure 4 shows the schematic set up of the conventional cooling system. In this configuration, eight (8) blowers move the air in parallel to the heat sink fins.

Figure 5 shows the implementation of an ATS Therm-Jett™, which provides jet impingement on the same four heat sinks, and the location of impingement holes with respect to the heat sinks.

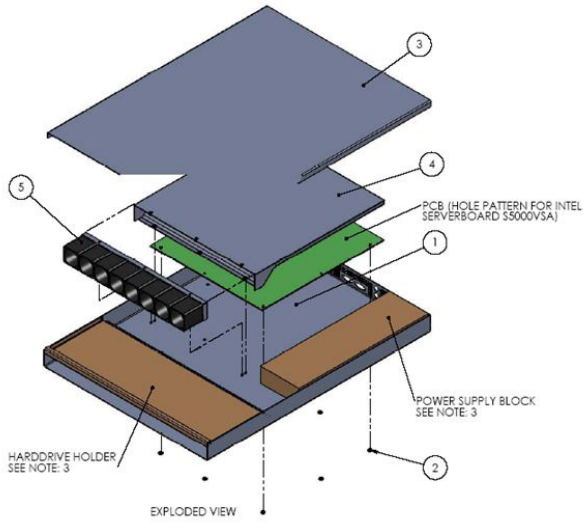

Figure 2. Exploded view of the Ujet-1000™ and the Therm-JETT™ cooling duct. (ATS)

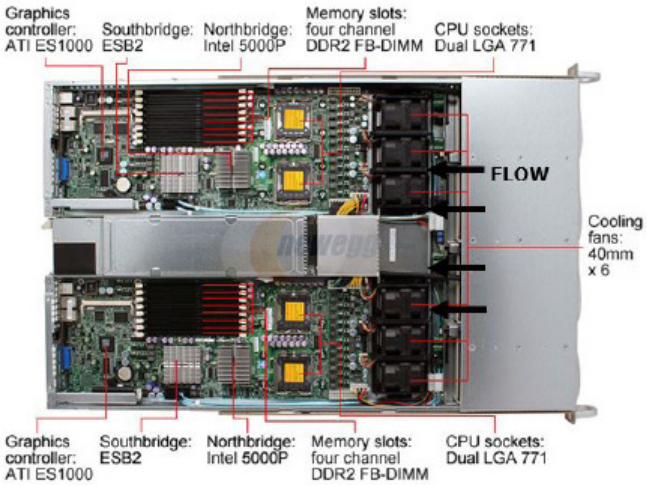

Figure 3. Conventional Cooling System in a 1-U Application. (ATS)

Tables 1 and 2 show the experimental values obtained within a 1-U chassis made by ATS. The two sets of tests were done for both 12 and 6 volts to the fans. The thermal resistance data of all four heat sinks, hard drives and the power supply were obtained for both conventional cooling and jet impingement cases. The acoustic noise for each case was also recorded for comparison.

The data shows an improvement of thermal resistance of 22% to 42% for the heat sinks from jet impingement as compared to conven-

tional cooling. The power supply shows a 10% improvement in the thermal resistance.

The hard drive, though, shows a 20% degradation. This is due to the fact that, with jet impingement, the pressure drop on the fans increases, consequently decreasing the flow through the system. However in an actual system the percentage will be smaller. That's because the heat generated will be more volumetric compared to the current setup where heat is generated on the surface of the "U" shape aluminum piece located at the bottom of hard drives.

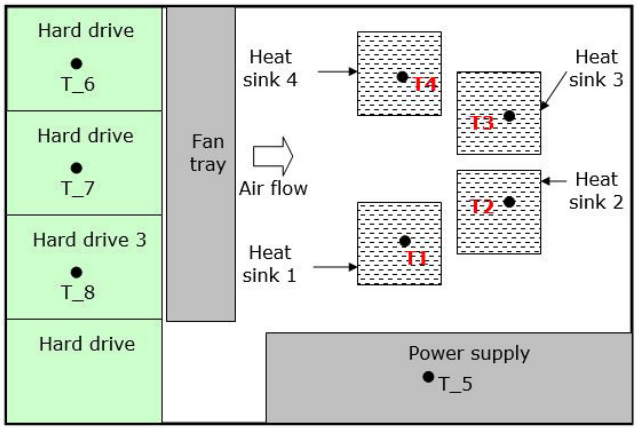

Figure 4. Configuration of conventional layout in a 1U system for temperature measurement. (ATS)

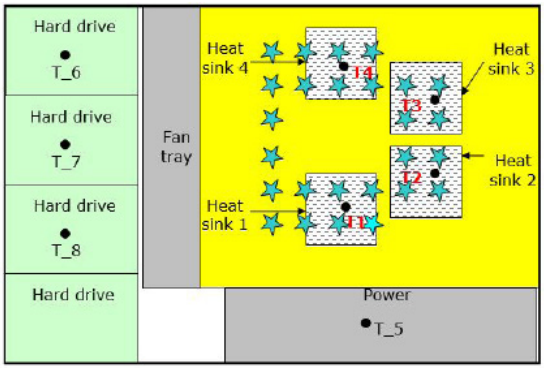

Figure 5. Configuration of an ATS Therm-JETT™ application in a 1U system for temperature measurement. (ATS)

In that case, the decrease of flow through the system will have less impact. Additionally, the increase in hard drive temperature is less than 2°C in this experiment, which is generally not large enough to be a concern.

The question might be raised as to whether the performance of the heat sinks could be improved if we removed the impingement duct,

Layout	Fan Voltage	R_{S-1} (°C/W)	R_{S-2} (°C/W)	R_{S-3} (°C/W)	R_{S-4} (°C/W)	R_{HD} (°C/W)	R_{PS} (°C/W)	Noise (db)- 0 m	Noise-(db)-1 m
Conventional	12	0.25	0.26	0.25	0.22	0.13	0.17	88	79
Therm-Jet	12	0.17	0.15	0.17	0.17	0.16	0.15	79	79
Therm-Jet Improvement % Over Conventional		40%	40%	33%	21%	-20%	10%		

Table 1. Experimental test results comparing conventional and Therm-JETT™ results with 12 volts to the fans.

Layout	Fan Voltage	R_{S-1} (°C/W)	R_{S-2} (°C/W)	R_{S-3} (°C/W)	R_{S-4} (°C/W)	R_{HD} (°C/W)	R_{PS} (°C/W)	Noise (db)- 0 m	Noise-(db)-1 m
Conventional	6	0.37	0.38	0.36	0.30	0.17	0.17	80	68
Therm-Jet	6	0.23	0.22	0.23	0.24	0.22	0.16	80	68
Therm-Jet Improvement % Over Conventional		37%	42%	35%	20%	-26%	9%		

Table 2. Experimental test results comparing conventional and Therm-JETT™ results with six volts to the fans.

increased the heat sink height by the height of impingement duct and ducted the flow.

We analyzed this situation and found that the improvement would be at most 5%, if we assume that the heat sink is ducted, and the pressure drop is the same in both short and tall versions. To study this problem in detail one must consider the fan curves instead of using a fixed volumetric flow rate. Interested readers will find an article in a previous issue of the ATS Qpedia Thermal eMagazine [1] with more information about this topic.

Table 3 shows the temperatures of the four heat sinks, the hard drive and the power supply. As we mentioned earlier, the heat sinks were mounted on 200 W devices, the power supply was dissipating 100 W and the hard drives were dissipating 80 W. The results are shown for jet impingement and conventional parallel air flow over 23.5 mm tall heat sinks, and ducted flow over heat sinks with 28.5 mm tall heat sinks. It can be seen that heat sink temperatures are significantly lower for jet impingement even compared with a taller heat sink with ducted flow.

Table 4 shows the improvement in temperature

of the four processors between the jet impingement and the two cases of 23.5 mm heat sink and the ducted 28.5 mm heat sinks.

By comparing the results for 6 and 12 volts to the fans, it can be seen that at lower voltage the jet impingement temperature difference is even more than with higher voltage to the fans. This implies low pressure drop fans can significantly benefit from the application of jet impingement.

Flow Type	Fan Voltage	HS #1 (°C)	HS #2 (°C)	HS #3 (°C)	HS #4 (°C)	Hard Drive (°C)	Power Supply (°C)
Impingement with 23.5 mm Heat Sink	12	59.8	55.6	58.4	59.2	36.8	39.7
23.5 mm Heat Sink-Parallel Flow	12	75	77	75.8	68.8	34.7	41.4
Ducted with 28.5 mm Heat Sink	12	71.9	73.8	72.6	66.1	34.7	41.4
Impingement with 23.5 mm Heat Sink	6	70.8	68	70.6	73.4	41.9	40.2
23.5 mm Heat Sink-Parallel Flow	6	99	100.2	96.4	85.6	38.1	41.8
Ducted with 28.5 mm Heat Sink	6	95.6	96.7	93.1	82.8	38.1	41.8

Table 3. Comparison of temperatures for jet impingement and parallel flow with 23.5 mm heat sinks and ducted flow with 28.5 mm tall heat sinks.

Flow Type	Fan Voltage	Heat Sink #1 (°C)	Heat Sink #2 (°C)	Heat Sink #3 (°C)	Heat Sink #4 (°C)
Jet Impingement Improvement Compared to 23.5 mm Heat Sink with Parallel Flow	12	15.2	21.4	17.4	9.6
Jet Impingement Improvement Compared to 28.5 mm Ducted Heat Sink	12	12.1	18.2	14.2	6.9
Jet Impingement Improvement Compared to 23.5 mm Heat Sink with Parallel Flow	6	28.2	32.2	25.8	12.2
Jet Impingement Improvement Compared to 28.5 mm Ducted Heat Sink	6	24.8	28.7	22.5	9.4

Table 4. Comparison of temperature improvement for jet impingement and parallel flow with 23.5 mm heat sinks and ducted flow for 28.5 mm tall heat sinks.

Figure 6 is a graphical representation of Table 4. The figure shows the significant temperature increases in the case of parallel flow and ducted flow for the heat sinks compared to jet impingement technology.

Components (heat sinks) 2 and 3 are hotter than components 1 and 2 because they are downstream, and the approach air temperature is higher for a ducted flow. In the jet impingement mode, the impingement flow is at upstream temperature and therefore much cooler than the air received in ducted flow.

In impingement mode, there is another flow coming axially toward the components, called cross flow. It is the interaction of cross flow and impingement that causes the cooling of the component (heat sink).

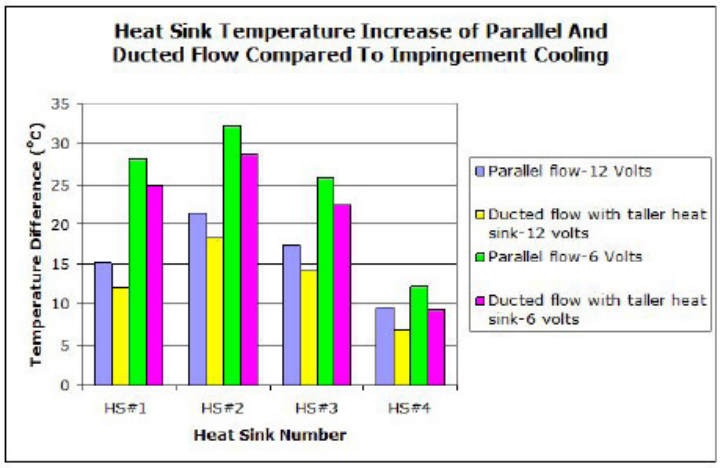

Figure 6. Heat sink temperature increase of parallel and ducted flow compared to jet impingement cooling. (ATS)

It should be noted that the above experiment was done for a heat source that is the same size as the heat sink base; hence, the spreading resistance is zero because it is almost independent of the heat transfer coefficient.

The spreading resistance can be added to the above numbers for other sizes of heat sources. The same concept of jet impingement has been applied to simulated components in ATCA chassis. The data improvement is promising.

Even though conventional air-cooling technology is fast approaching its thermodynamic limit, there are still numerous potentials for air cooling which will enable this technology to be used in the years to come.

Reference

1. “Heat Sink Thermal Resistance as a Function of Height-Ducted Flow with Fan Curve,” Qpedia Thermal eMagazine, Advanced Thermal Solutions, Inc., January 2009.

Thermal Performance of Heat Sinks with Heat Pipes or Vapor Chambers for Servers

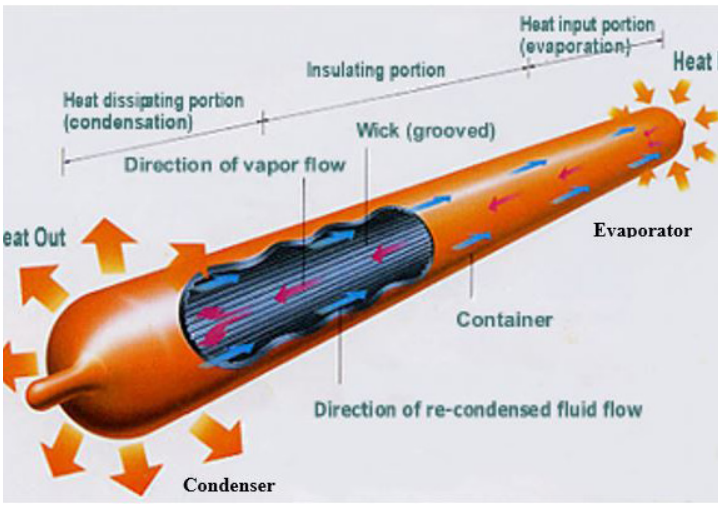

Figure 1. Typical heat pipe. [4]

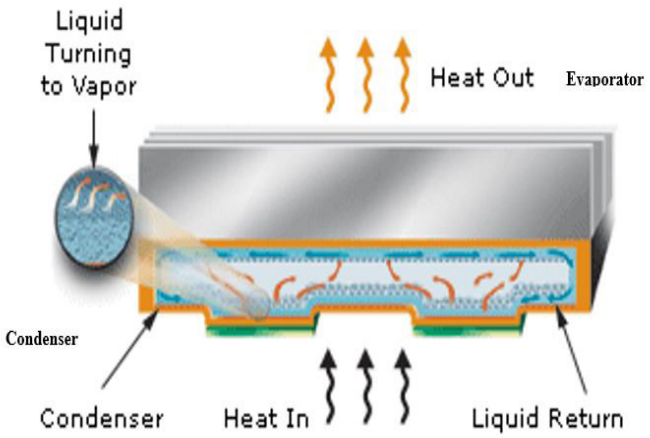

Figure 2. Vapor chamber. [5]

Most blade servers for data and telecommunication systems use air to cool the high-power chips inside. As the power levels of these chips keep increasing, the pressure is on thermal engineers to design ever higher performance air-cooled heat sinks. In recent years, advancements in manufacturing of thinner heat pipes and vapor chambers have enabled engineers to integrate the heat pipes and vapor chambers into the blade server heat sinks.

A heat pipe is a device with high thermal conductance that can transport large amounts of heat with a small temperature difference between its hot and cold ends. The idea of a heat pipe was first proposed by Gaugler [1]. However, only after its invention by Grover [2, 3] in the early 1960s, were the remarkable properties of heat pipes realized by scientists and engineers. It is now widely used to transport heat from one location to another location or to smooth the temperature distribution on a solid surface.

A heat pipe is a self-driven two-phase heat transfer device. A schematic view of a heat pipe is shown in Figure 1. At the hot section (evaporator), the liquid evaporates and turns to vapor. The vapor flows to the cold section (condenser) and liquefies there. The liquid is driven back from the cold section to the hot section by a capillary force induced by the wick structure. By using boiling and condensation, the heat pipes can transfer and spread the heat from one side to another side with a minimum temperature gradient.

Vapor chambers are flat heat pipes with very high thermal conductance. They have flat surfaces on the top and bottom sides, see Figure 2, which can be easily attached to a heat source and a heat sink.

Just like heat pipes, vapor chambers use both boiling and condensation to maximize their heat transfer ability. A vapor chamber generally has a solid metal enclosure with a wick structure lining the inside walls. The inside of the enclosure is saturated with a working fluid at reduced pressure. As heat is applied at one side of the vapor chamber, the fluid at locations close to the heat source reaches its boiling temperature and vaporizes. The vapor then travels to the other side of the vapor chamber and condenses into liquid. The condensed fluid returns to the hot side via the gravity or capillary action, ready to vaporize again and repeat the cycle.

In electronics cooling, heat pipes are generally used to move the heat from electronics to heat dissipation devices. For example, in a desktop computer, multiple heat pipes are used to transfer heat from a CPU to an array of cooling fins, which dissipate the heat to ambient environment through convection. Vapor chambers are generally used to spread heat from a small size device to a larger size heat sink, as it is shown in Figure 2. If used in server heat sinks,

the heat pipes and vapor chambers are both used to spread the heat due to the low profile and large footprint of the heat sinks.

Compared to copper heat spreaders, heat pipes and vapor chambers have the following merits.

First, they have a much higher effective thermal conductivity. The pure copper has a thermal conductivity of 401 W/m°C and the best conductive material of diamond has a thermal conductivity of 1000-2000 W/m°C. The effective thermal conductivity of a well-designed heat pipe and vapor chamber can exceed 5000 W/m°C, which is an order of magnitude higher than that of pure copper. Second, the density of the heat pipe and vapor chamber is much lower than that of copper. Due to its hollow structure, the heat spreaders made by vapor chambers are much lighter than those made of copper. These properties make them the ideal candidate for high heat flux and weight sensitive heat spreading applications.

Dynatron Corporation is an electronic cooling provider specializing in heat sinks for servers. This article compares the thermal performance of its server heat sinks, some of which have integrated vapor chambers. Figure 3 shows the photos of two Dynatron 1U passive server heat sinks for Intel's Sandy Bridge EP/EX Processors. The R12 is made of pure copper with skived fins. The R19 has a vapor chamber base and stacked copper fins. The heat sink specification is listed in Table 1. The R19 is 150g lighter than the R12.

Figures 4 and 5 show the thermal performance of R12 and R19 at different flow rates. At 10CFM, both heat sinks have a thermal resistance of 0.298ºC/W. When the flow rate increases to 20CFM, the R19's thermal resistance is 0.218ºC/W, which is 0.020ºC/W lower than that of R12.

Figure 6 shows the photos of two Dynatron 1U active server heat sinks for Intel's Sandy Bridge EP/EX Processors. The R18 is made of copper with skived fins. The R16 has a vapor chamber base and stacked copper fins. Both heat sinks use the same blower. The heat sink specification is listed in Table 2. The R16 is 90g

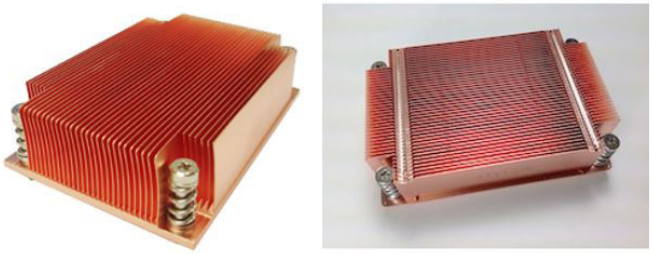

Figure 3. Dynatron passive heat sinks R12 (left) and R19 (right). [6]

	R12	R19
Material	Copper 1100 Skiving Fins	Vapor Chamber Base with Copper Stacked Fin
Overall Dimension	106.0 x 70.0 x 25.5 mm	106.0 x 70.0 x 27.0 mm
Overall Weight	550g	400g

Table 1. Dynatron passive heat sink specification.

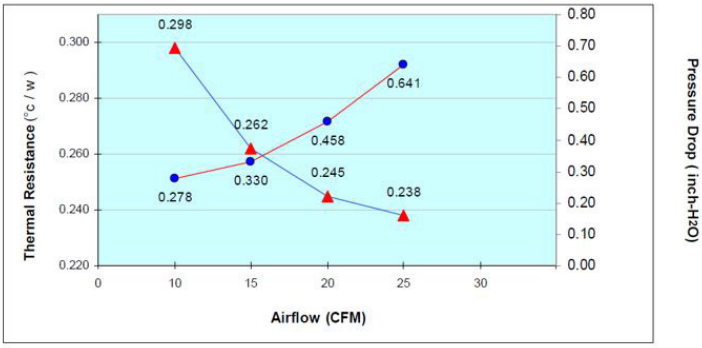

Figure 4. Dynatron R12 heat sink performance. [6]

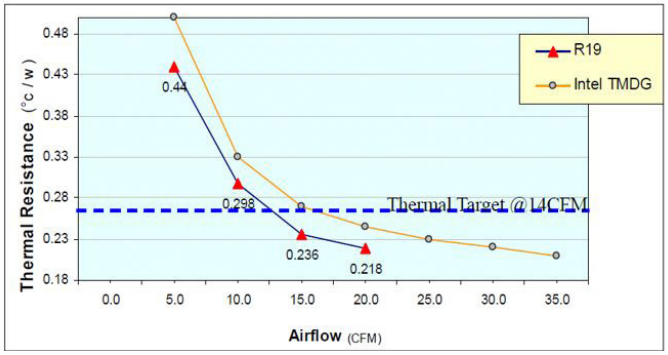

Figure 5. Dynatron R19 heat sink performance. [6]

Figure 6. Dynatron active heat sinks R18 (left) and R16 (right). [6]

lighter than the R18.

Figures 7 and 8 show the thermal performance of R18 and R16 at different blower speeds. At 3000RPM, the R18 and R16 heat sinks have a thermal resistance of 0.437ºC/W and 0.393ºC/W, respectively. When the blower speed increases to 6000RPM, the R18's thermal resistance is 0.268ºC/W and the R16's thermal resistance is 0.241 ºC/W. The R16 is constantly able to outperform the R18 at different blower speeds and its thermal resistance is 10% lower than R18.

The comparison of the Dynatron heat sinks shows that heat sinks with vapor chambers have a slight thermal edge vis-a-vis its copper counterparts even though they are light. This is true for both passive and active heat sinks.

Glover et al., for Cisco, have tested different heat sinks either with embedded heat pipes or vapor chambers for their servers and published their findings [7]. They tested five different heat sinks from different vendors, who utilized different manufacturing technologies to fabricate the heat sinks. The five heat sinks are similar in size: 152.4 x 101.6 x 12.7mm. Table 3 summarizes the physical attributes of these five heat sinks.

Figures 9-11 show the three vapor chamber heat sinks with different vapor chamber structures and fin designs. Heat sink A-1 is an extruded aluminum heat sink with a vapor chamber strip. The 40 mm wide vapor chamber strip is embedded in the center of the base. It is the lightest one among the five tested heat sinks. Heat sink B-1 and C-1 have full base size vapor chambers and aluminum zipper fins.

Figures 12-13 show the two heat sinks with embedded heat pipes. Heat sink C-2 has heat pipes embedded inside its aluminum base. It uses zipper fins and has a copper slug in the middle of the base. Heat sink D-1 has three flat heat pipes embedded in its base. It has a copper plate as a base.

Glover et al. tested the five heat sinks at different mounting orientations and air velocities. Table 4 presents the summary results of the heat sinks at 3m/s approach air velocity. The

	R18	R16
Material	Copper 1100 Skiving Fins	Vapor Chamber Base with Copper Stacked Fin
Overall Dimension	90.0 x 90.0 x 27.7mm	90.0 x 90.0 x 28.0 mm
Overall Weight	480g	390g
Blower Model	DB128015BU-PWMG	DB128015BU-PWMG
Blower Dimension	80 x 80 x 15mm	80 x 80 x 15mm

Table 2. Dynatron active heat sink specification. [6]

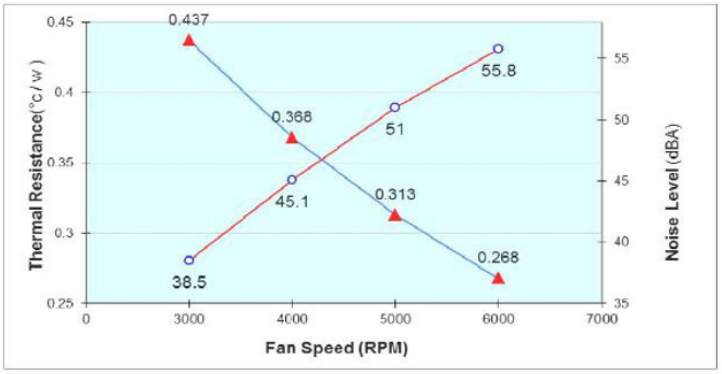

Figure 7. Dynatron R18 heat sink performance. [6]

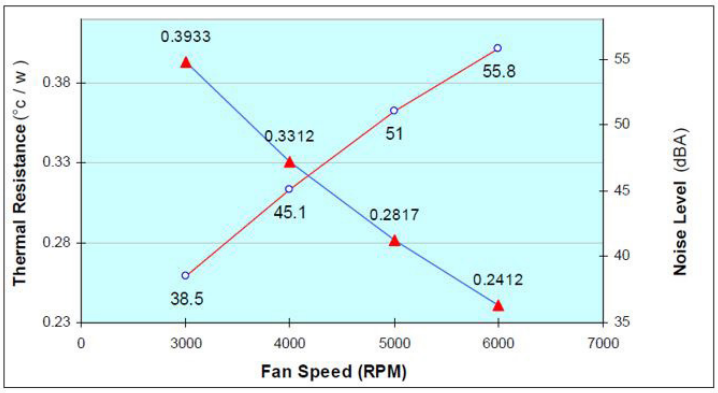

Figure 8. Dynatron R16 heat sink performance. [6]

Heat Sink #	Base Thickness (mm)	Weight (g)	Description
A-1	3	270	Extruded heat sink with 2 mm thick embedded sintered powder VC strip in base, thicker fins.
B-1	3	373	Wire-mesh VC base with Al zipper fins.
C-1	3	318	VC base wick structure: Heating side sintered powder, cooling side-wire-mesh; Al zipper fins.
C-2	3	246	Al base with 3 embedded heat pipes in parallel, inserted 1"x1" Cu block in center of base; Al zipper fins.
D-1	4	320	3 heat pipes cantilevered off the Cu center block base with Al stacked fins.

Table 3. Cisco tested heat sink specification. [7]

tested heat sinks were mounted horizontally with heater sources underneath the heat sink bases.

The C-1 heat sink has the lowest thermal resistance; thus, its values are used as the benchmark for other heat sinks. The performance of heat sinks is purely design dependent. For vapor chamber heat sinks, the thermal resistance value varies from 0.19 to 0.23°C/W for

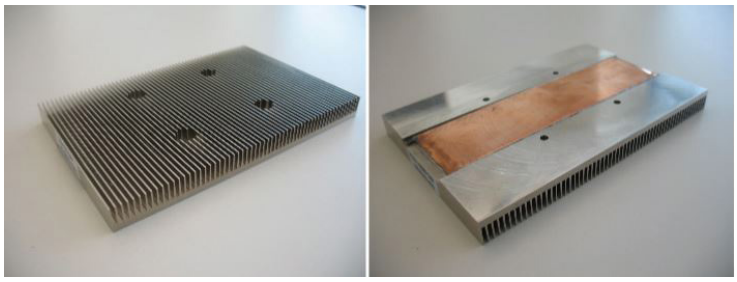

Figure 9. Heat sink with vapor chamber A-1. [7]

Figure 10. Heat sink with vapor chamber B-1. [7]

Figure 11. Heat sink with vapor chamber C-1. [7]

30 W of power. For heat sinks with heat pipes, the C-2 heat sink has a thermal resistance of 0.23°C/W, which matched with that of A-1 and B-1.

The D-1 heat sink has the highest thermal resistance, which is the result of inferior design and manufacturing. However, the D-1 heat sink still has relatively low thermal resistance when it is compared to a regular heat sink without a heat pipe and vapor chamber.

Figure 14 shows the thermal resistance of the five heat sinks for 60W of input power at different air velocities. The C-1 heat sink performs best for all velocities and the D-1 heat sink's performance is the worst.

The pressure drop across the heat sink at different air velocities was also measured and the results were plotted in Figure 15. The B-1, C-1 and C-2 heat sinks have similar fin structures. Therefore, their pressure drop is similar, too. The pressure drop of the A-1 and D-1 heat

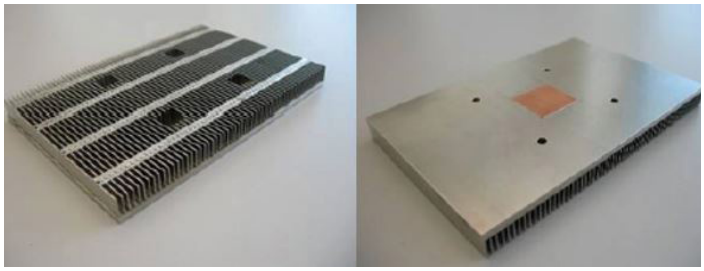

Figure 12. Heat sink with heat pipes C-2. [7]

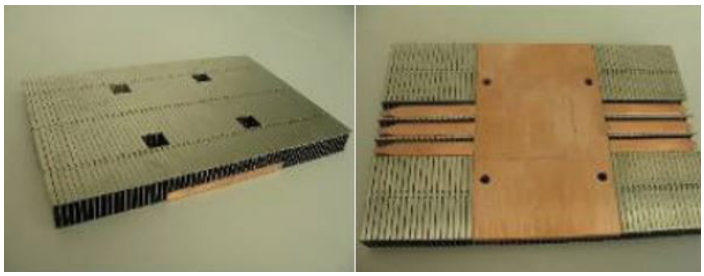

Figure 13. Heat sink with heat pipes D-1. [7]

Heat Sink	P=30W				P=60W			
	θ_{sa}		T_s	ΔT_{base}	θ_{sa}		T_s	ΔT_{base}
	(°C/W)	Diff	(°C)	(°C)	(°C/W)	Diff	(°C)	(°C)
A-1	0.23	21%	42	4.4	0.22	21%	48	7.9
B-1	0.23	20%	42	5.0	0.21	13%	47	7.6
C-1	0.19	0%	41	3.8	0.18	0%	46	6.1
C-2	0.23	21%	42	4.7	0.22	23%	48	9.1
D-1	0.26	37%	43	NA	0.25	39%	50	NA

Table 4. Heat Sink Performance at 3 m/s with horizontal mounting position and bottom heating. [7]

sinks are similar and higher than the other heat sinks. This is because the A-1 heat sink has thicker fins and the D-1 heat sink has a thicker base.

Because the heat pipes and vapor chambers use capillary force to drive liquid back from the condensation section to the evaporation section, their thermal performance is prone to orientation variation. Glover et al. also investigated the effects of the mounting orientation on the performance of the five heat sinks. They found the effect of the orientation is design dependent and is the result of both the wick structure and the entire heat sink assembly construct.

The heat sink specification from Dynatron Corporation and the test results from Cisco, show that the server heat sinks with embedded heat pipes or vapor chamber have a better thermal

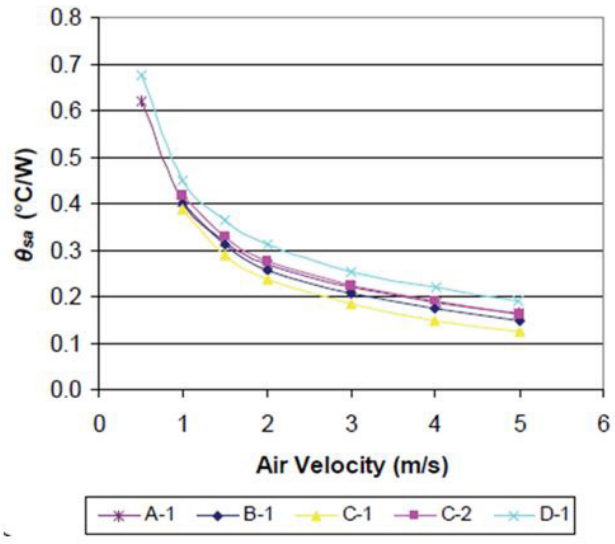

Figure 14. Heat sink thermal resistance at 60 W. [7]

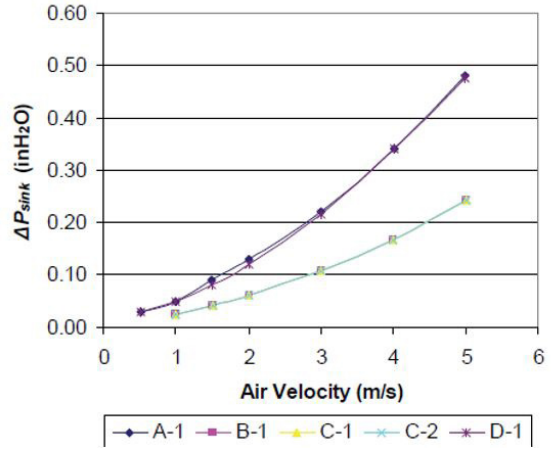

Figure 15. Heat sink pressure drop. [6]

performance than their copper counterparts. The heat sinks with embedded heat pipes or vapor chamber are also lighter than the pure copper heat sinks, which make them more suitable for applications which are weight sensitive.

If the cost of such heat sinks is justified, they are definitely good candidates for server cooling applications.

References

1. Gaugler, R. S., US Patent 2350348, Appl. 21 Dec, 1942. Published 6 Jun. 1944.
2. Grover, G. m., US Patent 3229759. Filed 1963.
3. Grover, G. M., Cotter, T. P., and Erickson, G. F., "Structure of Very High Thermal Conductance." J. App. Phys., Vol. 35, P. 1990, 1964.
4. http://www.lightstreamphotonics.com/technology.htm
5. http://www.thermacore.co.uk/vapour-chamber
6. http://http://www.dynatron-corp.com
7. Glover, G., Chen, Y., Luo, A., and Chu, H., "Thin Vapor Chamber Heat Sink and Embedded Heat Pipe Heat Sink Performance Evaluations," 25th IEEE SEMI-THERM Symposium, San Jose, CA, USA, 2009.

Written and published by the engineers of Advanced Thermal Solutions, Inc. (ATS), Qpedia is a technical resource to help the engineering community solve the most challenging thermal problems. Each issue provides practical approaches to thermal management challenges and explains the new technologies for successfully cooling today and tomorrow's hot-running electronics.

For additional Server Cooling Articles, as well as our complete Qpedia Library of articles, visit www.qats.com/Qpedia-Thermal-eMagazine.