International Journal on Emerging Technologies **4**(2): 42-46(2013)
ISSN No. (Print) : 0975-8364
ISSN No. (Online) : 2249-3255

Reduce Chiller and Server Power, and Removing Heat Transfer for Liquid Cooling with Minimal Coolant Volume

Mahima Sonakiya, Bhupendra Panchore and Nitesh Mishra
Department of Mechanical Engineering,
Shri Dadaji Institute of Technology & Science, Khandwa, (MP)

(Received 15 July, 2013 Accepted 17 August, 2013)

ABSTRACT: Uptime, operating cost and serviceability are the major concerns for data center cooling. This paper will describe an air and liquid cooling system that is optimized for data centers and does not involve major infrastructure changes. This system uses a minimum heat exchanger volume and coolant volume in order to simplify the installation at the server and data center level. The goal is to reduce chiller and server power, and removing heat from the CPU and the air in the server achieves that goal. Secondary goals include reducing the volume of liquid inside the server and making it easy to swap out racks, servers or components. The proposed system uses a normal base plate and fin heat sink, with a high performance liquid cooled heat exchanger built into it. This allows heat to be extracted from the CPU as well as the air that flows through the heat sink, removing up to 100% of the heat from the entire server with the liquid cooled heat sink which makes it a literal 'heat sink'.

Key words: Liquid Cooling, Heat Sink, Data Center, Server Cooling

I. INTRODUCTION

CPU power densities are returning to the levels they were at in the early 90s when the fastest computers ran power hungry ECL logic. At that time, liquid cooling systems were deployed by a number of mainframe manufacturers. Now CMOS power densities are at the same levels, approximately 100 watts/cm2. Liquid cooling is attractive again because of the low thermal resistance; low pumping power and high density that comes with liquid cooling. Spills are a concern with any liquid cooling system, but the risk can be eliminated by using negative pressure. Optimization of data center power for a given amount of computer services has been recognized as a way to save energy and reduce carbon output. Software improvements such as virtualization and putting excess hardware to sleep are currently being implemented. Efficiency improvements are underway for hard drives, power supplies and memory. For server CPUs the goal is to provide as much computing power as possible at the ~140-watt level that is the maximum that a typical air-cooled heat sink can dissipate. This leads to a power concentration that is attractive for liquid cooling. The diffuse heat generation in the remainder of the system lends itself to air-cooling, although in some cases the RAM may also benefit from liquid cooling. A liquid cooling system which runs in parallel with the existing air-cooled infrastructure was selected for uptime and serviceability considerations, as any liquid cooled system will be easier to service if it can be air cooled on the workbench. A CPU heat sink that is both liquid and air cooled leads to a few no obvious advantages.

II. COMBINED LIQUID AND AIR COOLED SYSTEMS

The liquid cooled architecture that we propose is an addition to an existing air-cooled heat exchanger, so that the air-cooling system still works as originally designed and the liquid cooling passages fit inside the existing heat sink, within the base or attached to the fins. Heat from the CPUs, which represents up to 50% of the overall server power, can

be removed directly to an outside cooling tower, with no chiller, room fans or air handlers required. The wet bulb temperature in the US never gets above 81oF (27oC), and the processor maximum case temperature is 145-167oF(63-75oC) so if there is a low thermal impedance liquid heat path from the CPU to a cooling tower, there is plenty of margin and no chiller is required for the liquid cooling system. Bypassing the A/C system saves the data center cooling power and adds redundancy to the cooling system. The cooler processors have lower leakage currents, and if the system can take advantage of the lower temperature, it can run faster and/or use a lower processor voltage and even less power. If the data center temperature is significantly higher than the wet bulb temperature, as it is most of the time, the fins on the heat exchanger will remove heat from the data center air as well as the CPU. This is illustrated in Figure 1. This system adds redundancy and reduces cooling loads. ASHRAE recommends a data center operating temperature of 80oF, (27oC), but many servers speed up the fans at this temperature and the CPU requires more power due to higher leakage currents, so the savings in HVAC power are not fully realized. If the CPUs are liquid cooled, the data center can run warmer with no issues. Often the data center ambient temperature is set based on operator comfort at the hottest location. If liquid cooling is used, the stack rise through the server is reduced, the hot aisle temperature remains reasonable and the room temperature setting can be increased. The liquid/air heat sink performance may be predicted based on the data center temperature, the coolant temperature, the CPU power and the thermal resistances.

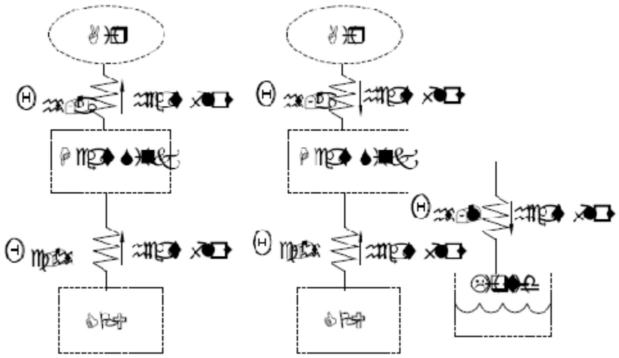

Fig. 1. Heat flow reversal with air vs. air and liquid cooled heat sink.

The heat flow may be estimated using the simplified thermal model in Figure 1 or a more detailed one using CFD and FEA. It can be modeled as a simple voltage divider using an electrical analog between the data center air, the heat sink and the liquid coolant with the processor modeled as a current source. The various thermal resistances can be adjusted via heat sink design to minimize overall data center power consumption. For example, if the heat sink has a thermal resistance of .04oC/w from the water to the Heat sink and .2oC/W to the air, and the CPU power is 100 watts, then if the water temperature is 59oF (15oC) and the air temperature is 77F (25oC), the heat sink will remove 33 watts from the air in addition to the 100 watts of CPU power. If the same heat sink is used with a CPU power of 50 watts and an 86oF(30oC) data center temperature, then. 105 watts of heat will be extracted from the data center air. This means that if the data center is allowed to warm up during periods of low occupancy, the computers are at low power and the wet bulb temperature is
low then the HVAC system will stay off and the PUE drops
to near 1 without having to build a custom building. The precise power savings for a given data center over a year can be calculated using the method above, combined with available data on hourly wet bulb temperature vs. time from the NWS and data center power consumption vs. time.

III. SERVICEABILITY AND RELIABILITY

For serviceability, the volume of coolant in the system should be minimized. In the event that water is used for cooling, a minimal amount of water could spill and potentially damage the electronics. For other fluids, this leads to a minimum amount of fluid loss. In addition to minimizing the coolant volume, a method to remove the coolant from the server or blades is desirable, so that any fluid remaining will not attract dirt or spill. If the liquid heat exchanger flow path has a minimum cross sectional area, then it is possible to remove all the coolant by purging with air when a component is removed. For maximum system reliability, the coolant is run at negative pressure throughout the server chassis and the rack so that any leaks are of air into the cooling system instead of coolant out of the system. This limits the delta pressure through the heat exchanger based on the difference between the local atmospheric pressure and the vapor pressure of the coolant at its warmest point. This provides an available delta pressure of about 5 psi in a worst-case scenario of high temperature, altitude and humidity (Denver during a thunderstorm). This is sufficient delta pressure to extract a few hundred watts from any processor. Any dissolved air in the coolant will come out of solution at the hottest point as the solubility of air in water decreases with temperature. In addition air will enter the system during removal or installation of components so the pump and plumbing must be designed to reject air in the return line. One more requirement for servers is that server replacement should not require the entire cooling system be shut down. This has been done with double shutoff connectors which seal when disconnected, but the subsystem is still full of coolant which must be removed before component replacement or device storage. A better way to solve the problem is to use a connector which purges the server of coolant when it is disconnected, as shown schematically in Figure 2. . Such a connector will allow air to enter the server as it is disconnected. If the server has a minimum volume of coolant, this process occurs in less than one second. Then, if component repair is required inside the server, plumbing can be disconnected without loss of coolant.

The connector works by venting the supply line, and shutting off the liquid supply which allows air to be pulled into the heat sink, displacing the coolant. After the liquid is purged, the return line is shut and only then can the connector be removed. A small amount of air may be pulled into the system, but it will be automatically be pumped out.

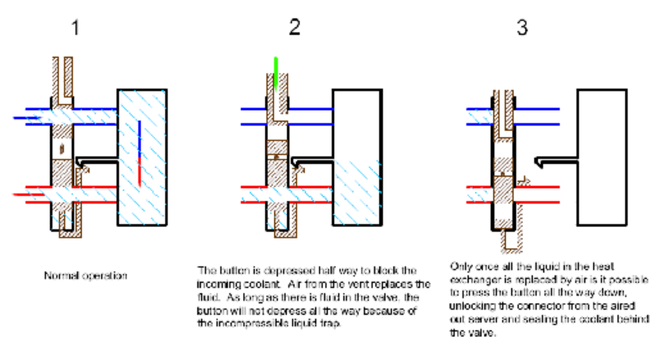

Fig. 2. No Drip Hot Swap Connector Operation Schematic.

The CPU is the most heat sensitive and has the highest power density so typical server cooling systems are generally designed for it, and the other components run colder than necessary. Furthermore, the server is generally designed to work with a fan failure; so redundant fans are added to make up for the flow resistance or leakage through a failed fan. This leads to a system where the hard drives in particular end up running too cool, as Pinheiro et al found a study4 where they found that reliability was slightly improved at the higher end of the hard drive temperature range. Often the hard drives are located at the front of the server, exposed to the coldest air. Therefore, as liquid cooling is added to a server, the fans can be slowed, allowing the hard drives to warm up, saving power and increasing reliability for fans, CPU and hard drives.

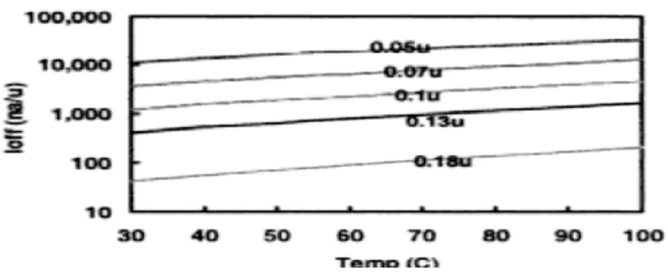

Fig. 3. Leakage Current Increases with Temperature. (From Narendra3).

that cooler CMOS uses less power due to leakage current reductions, and as the number of transistors goes up, the size of the features gets smaller, the leakage current adds up to more of the total current making the power sensitivity to temperature greater. The sensitivity of leakage to temperature is shown in Fig. 3.

A. Heat exchanger design

The design starts with the existing air-cooled system. In order to provide the best cooling with minimum volume and input power, a spiral cooling channel with a Reynolds number just above the laminar limit is used. This provides the best cooling with a reasonably sized channel that can pass contamination.

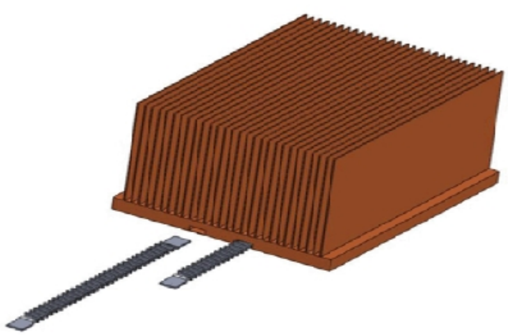

Fig. 4. Turbulators added to air-cooled heat sink.

For example, if a 140 watt CPU is to be cooled with water, and we can accept an 18oF (10^{o}C) temperature rise, then we need a flow rate of 220 cc/minute based on the heat capacity and mass flow rate of water. Now we take a page from rocket nozzle cooling system design2, which is often done with an array of tubes that cool the nozzle and preheat the fuel on the way to the combustion chamber. The goal is to adjust the length and diameter of the parallel tube array to get the optimum cooling for a given flow rate. In our case, we want the water outlet temperature and the heat sink temperature to be within 1oC. So we select a fluid path with a Reynolds number slightly higher than 2100, so that the flow is turbulent, but the pressure drop is reasonable. In this case we use two helical flow passages, .055 inch (1.4 mm) in diameter. This system has been analyzed using empirical heat transfer equations for flow in a tube, modeled in CFD, and tested with a Xeon processor running stress software. The thermal resistance heat sink to water, based on the temperature of the water into the heat sink, was .04 watt/oC with 230 cc/minute flow rate per CPU. A modified heat sink design with coolant passages in the base is shown in Figure 4. The analysis of the heat flow is shown in Figure 5, and a prototype is shown installed in a HP Server in Figure 6.

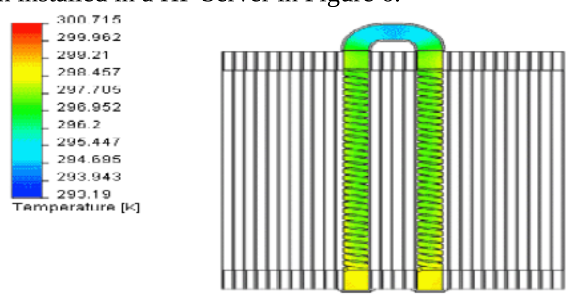

Fig. 5. CFD analysis of Turbulators.

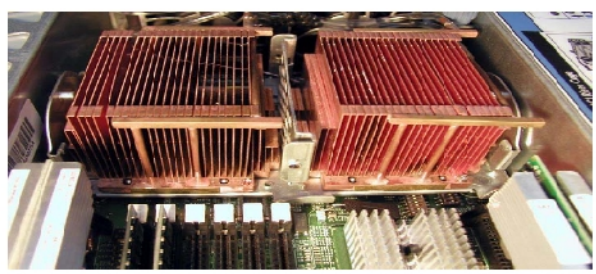

Fig. 6. Air Liquid Cooled Prototypes Installed.

This heat sink worked exactly as predicted, but when the flow was increased, it was discovered that it could remove heat from the entire system. A stack of three DL380 servers was run at idle power levels in an insulated box and the heat sinks were able to remove *all* the heat (700 watts) from the computers. In this case the ambient air was 107oF (42^{o}C) and the coolant inlet was at 76oF (22^{o}C).

IV. SERVER TEST RESULTS

A test was conducted with a 2 Kw rack of servers in an office environment at 75**o**F (24**o**C) ambient. The servers were either air-cooled or water-cooled using an outdoor miniature cooling tower with water at 65**o**F (18**o**C). The temperature data is shown in Figure 7. Note that when the liquid cooling was turned off the HVAC system was not able to keep up, so the door was opened slightly to keep the temperature relatively constant. A set of 7 Servers (3x HP Proliant DL380 G4 2x3.4 GHz and 4 Verari 2U 2x Opteron 245) consumed 2 Kw using air-cooling while running a processor stress test program (2 instances of BurnK7). With liquid cooling, and slowed down fans, the power was reduced to 1.8 Kw with 1 Kw of heat extracted using the liquid cooling system. In addition the average processor temperature decreased 25 F (14 C). The hard drives warmed slightly with liquid cooling due to the reduced airflow, but as mentioned before, they last longer at warmer temperatures. The RAM temperatures were lower with liquid cooling because the RAM chips were located downstream of the heat exchanger. Assuming a typical data center power distribution1 of 56% Servers, 30% HVAC, 5% UPS and 6% other, the total power required for the original air cooled system would be 3.6 kw (Server Power/.56). Using liquid cooling allows 1 Kw to bypass the HVAC system and go directly outdoors, saving HVAC power. It also saves 10% of the server input power due to lower fan power and because the processors require less power at lower temperatures The liquid pump and cooling tower fan use only 50 watts.

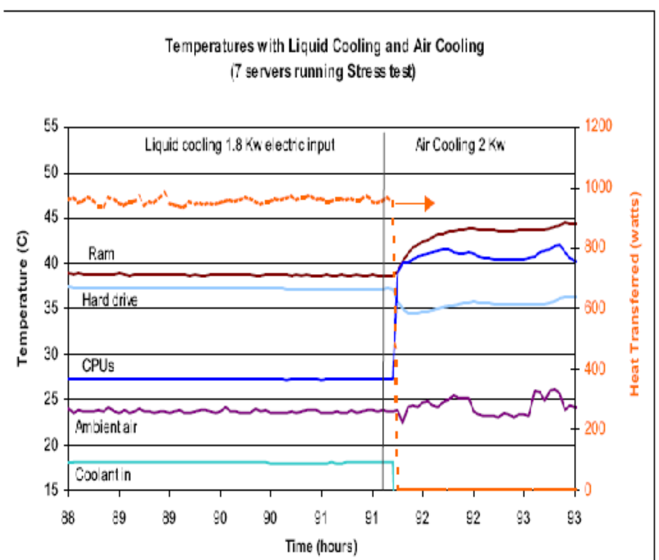

Fig. 7. Test data for 7 servers under stress test.

This reduces the overall power consumption based on typical data center power distribution to 2.9 Kw, a power reduction of ~20%. The power reduction is diagrammed in Fig. 8. This experiment was done using a miniature cooling tower that was only 52% efficient which lowered the water temperature down to (65°F) 18°C in a (50°F) 11°C wet bulb environment. A commercial grade cooling tower with 75% efficiency would be able to reduce the temperature of the cooling water to (59°F) 15°C. Assuming that the heat removed is proportional to the difference between the ambient and the cooling water inlet, the more efficient cooling tower would boost the heat removal by 50%, leading to a predicted power savings of 25% overall.

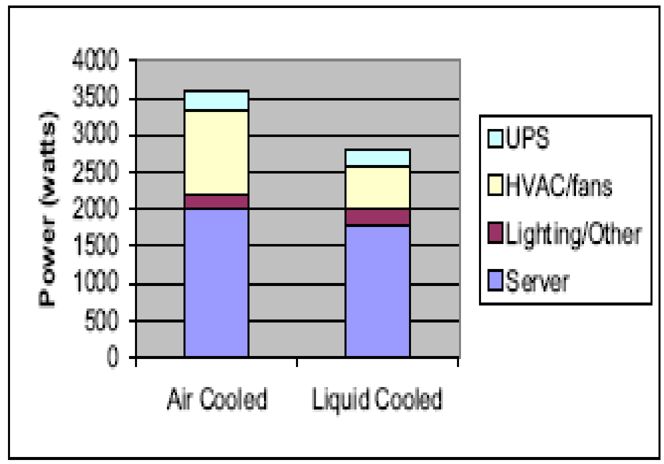

Fig. 8. Power Distribution Based on Test Data.

V. CONCLUSIONS

The combination of an air cooled heat sink modified for redundant liquid cooling, a negative pressure system to prevent leaks, and a connector that automatically purges the coolant adds up to a system that offers a path from the current air-cooled technology to the liquid cooled data center of the future, without having to modify the building. The key to obtaining a respectable ROI for this system is the water and air cooled heat sink that reverses the system thermodynamics so that the heat sink removes heat from the CPUs and the server interior and the data center in general in order to reduce the HVAC loads and fan power by a large margin.

REFERENCES

[1]. Pelley, S.; Meisner, D.; Wenisch, T. F., and VanGilder, J. W. 2009. "Understanding and abstracting total data center power." *In Proceedings of the Workshop on Energy Efficient Design (WEED'09).*

[2]. Dieter K. Huzel and David H. Huang Modern Engineering for Design of Liquid-Propellant Rocket Engines (Progress in Astronautics and Aeronautics) AIAA 1992.
[3]. Narendra, Siva G. and Chandrakasan, Anantha P. “Leakage in Nanometer CMOS Technologies” (Integrated Circuits and Systems) Springer 2010.
[4]. Pinheiro, Eduardo; Weber, Wolf-Dietrich: and Andr´e Barroso, Luiz “Failure Trends in a Large Disk Drive Population” Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST’07), February 2007.