DATA CENTER ENERGY SAVINGS: TOTAL LIQUID COOLING VS. INDIRECT COOLING

Electronics COOLING

JUNE 2015
electronics-cooling.com

FEATURE

Application of Low Melt Alloys as Compliant Thermal Interface Materials:

A Study of Performance and Degradation under Thermal Duress

BACK COVER

INSIDE

COMPARISON OF HPC/TELECOM DATA CENTER COOLING METHODS BY OPERATING AND CAPITAL EXPENSE

CONTENTS

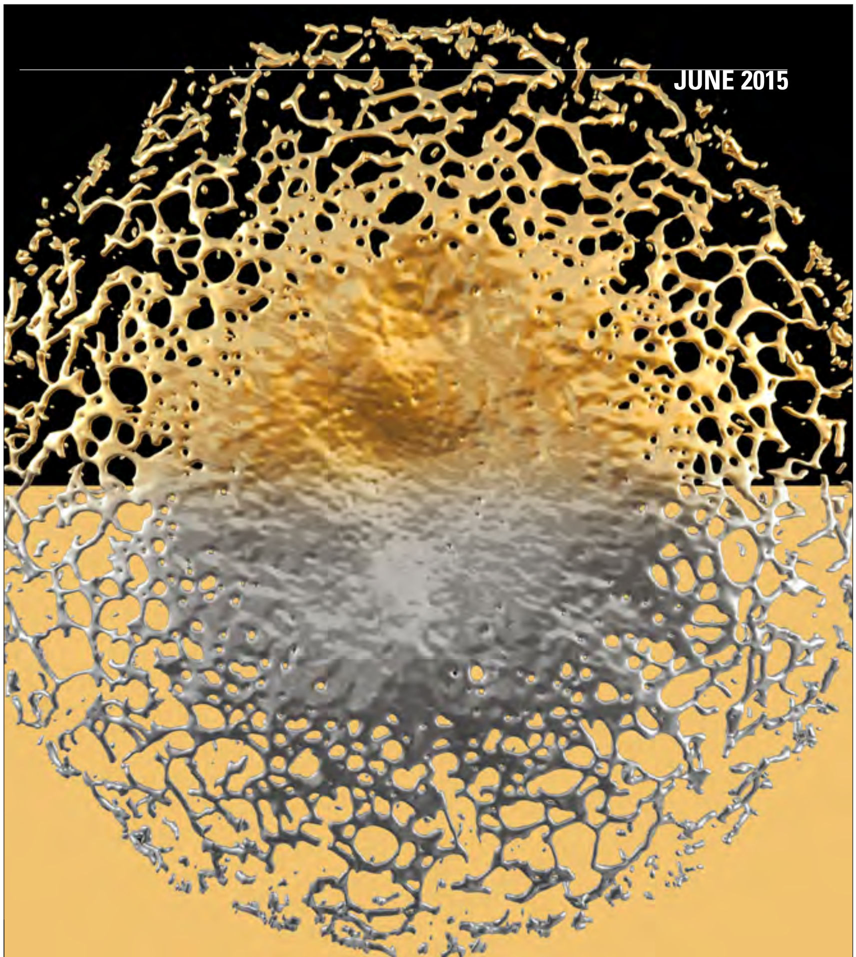

FEATURE ARTICLES

Editorial

The Joy of Engineering

Bruce Guenin, Acting Editor-in-Chief, June 2015

THE RECENT FINANCIAL crisis and the subsequent difficulties encountered by many recent college graduates in finding satisfying jobs, particularly for those in the humanities, has led to many questioning the real value of a college education.

In contrast, graduates in the STEM disciplines, namely, science, technology, engineering, and math have fared much better. Consequently, out of concern for the future career success of elementary and secondary school students, parents and educators have been encouraging them to think seriously about engineering.

Of course, one applauds the good intentions of these adults in attempting to guide our youth into these lucrative and societally critical careers. However, I'm sorry to say that the message our youth often hears is that training for these fields is very challenging and takes a lot of effort, but in the end, you'll be rewarded with steady employment and financial security. One must really wonder how effective these "eat your spinach" exhortations will really be in motivating our young folks to persevere for the many years involved in becoming an engineer.

This editorial is based on a different premise, namely that the act of doing engineering can be an intrinsically rewarding and, indeed, an even joyful experience. As we practicing engineers know so well, engineering involves many different activities. They include understanding, inventing, designing, calculating, making things, and problem solving, among others. Also, in our era, engineering is a very social activity. Most engineering challenges are complex and require the efforts of diverse teams. Communication and social skills are as important as the technical ones. This social aspect of the profession can further enhance the intrinsic enjoyment that comes with our profession.

I would say that our biggest challenge in motivating the young to become the engineers of the future is how to give them an early taste of the joy of engineering. In our day, there are fewer examples of ordinary adults doing engineering simply by repairing their own cars and homes as in earlier times.

Furthermore, the lives of many kids tend to be highly programmed, with most of their time devoted to school and other activities, leaving them little time (or parental permission) to freely explore their environment and improvise their own solutions as they encounter various challenges in their physical world.

What can we do about it? As engineers, we can create opportunities to share "roll up your sleeve" moments with young folks and to give them a feel for how to approach a technical problem, by defining a plan to solve it and patiently (we hope) seeing the project through to completion. This, in itself can be quite satisfying, but the real thrill occurs when, after struggling with a problem, an "ah ha" moment occurs and the way forward suddenly reveals itself.

We can certainly do this within our families and even look to do so in our wider society by volunteering in our schools and youth organizations, such as the scouts. Thus we can introduce our young folks to the joy of engineering and help to create the engineers of the future one child at a time.

Electronics COOLING

www.electronics-cooling.com

ASSOCIATE TECHNICAL EDITORS

Bruce Guenin, Ph.D.
Principal Hardware Engineer, Oracle
bruce.guenin@oracle.com

Peter Rodgers, Ph.D.
Professor, The Petroleum Institute
prodgers@pi.ac.ae

Jim Wilson, Ph.D., P.E.
Engineering Fellow, Raytheon Company
jsw@raytheon.com

PUBLISHED BY
ITEM Media
1000 Germantown Pike, F-2
Plymouth Meeting, PA 19462 USA
Phone: +1 484-688-0300; Fax:+1 484-688-0303
info@electronics-cooling.com; electronics-cooling.com

CONTENT MANAGER
Belinda Stasiukiewicz
bstas@item-media.net

MARKETING MANAGER
Dawn Hoffman
dhoffman@item-media.net

GRAPHIC DESIGNER
Erica Osting
eosting@item-media.net

EDITORIAL ASSISTANT
Allison Titus
atitus@item-media.net

PRESIDENT
Graham Kilshaw
gkilshaw@item-media.net

REPRINTS
Reprints are available on a custom basis at reasonable prices in quantities of 500 or more. Please call +1 484-688-0300.

SUBSCRIPTIONS
Subscriptions are free. Subscribe online at www.electronics-cooling.com. For subscription changes email info@electronics-cooling.com.

ENZOTECH™
WWW.ENZOTECHNOLOGY.COM

Pure Copper , Ultra Low Profile

Enzotechnology has designed a new ultra low profile copper heat sink cooler series, providing a perfect solution for any compact systems.

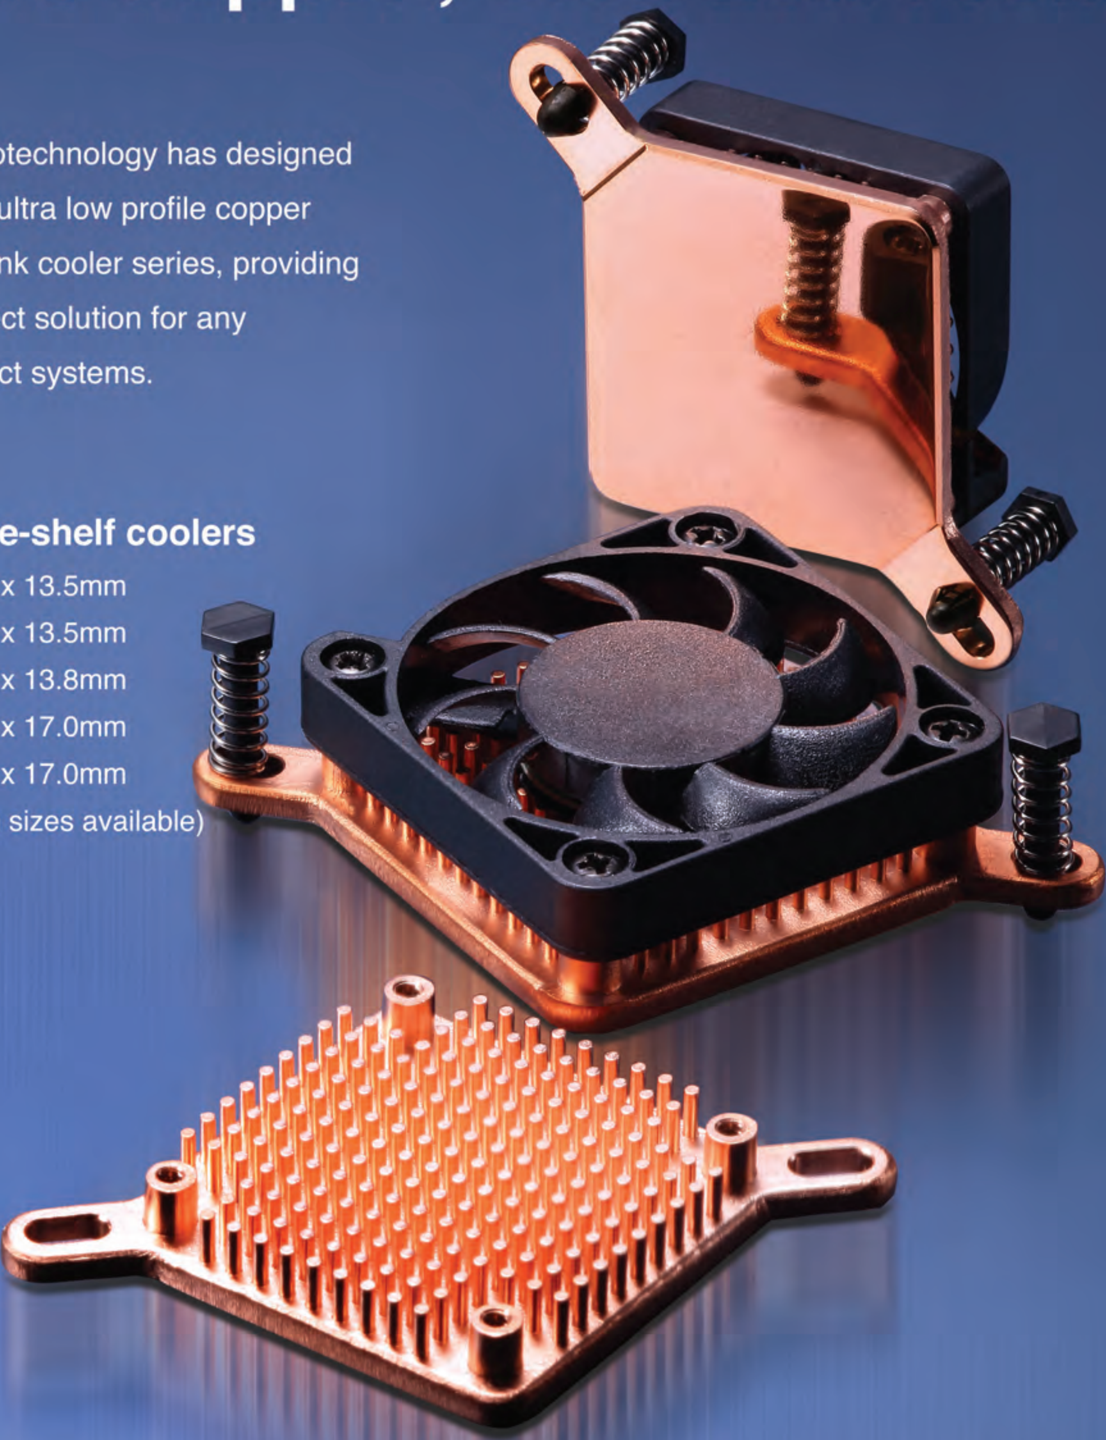

Off-the-shelf coolers

32 x 32 x 13.5mm
37 x 37 x 13.5mm
40 x 40 x 13.8mm
50 x 50 x 17.0mm
60 x 60 x 17.0mm
(custom sizes available)

ENZOTECHNOLOGY CORP.

Address : 14776 Yorba Ct. Chino CA 91710 USA
Tel : 909-993-5140
Fax : 909-993-5141
E-mail : sandra@enzotechnology.com
Website : www.enzotechnology.com

www.enzotechnology.com

Cooling Matters

News of thermal management technologies

COOLING SYSTEM FOR COMP. PROCESSORS COULD SAVE CONSUMERS BILLIONS

5/5/15 - A patented passive cooling system for computer processors could save U.S. consumers more than 6.3 billion dollars in electricity costs per year. The system is currently being manufactured at the University of Alabama in Huntsville.

The cooling system uses "convection to circulate 3M's Fluorinert FC-72 liquid through channels in a computer's processor and then into a heat sink that serves as an external radiator." Fluorinert FC-72 is a colorless, odorless, nonflammable, chemically stable and biologically inert, dielectric liquid with a boiling point of 56°C.

In the system's convection cycle, heat from the processor vaporizes the FC-72 liquid, which then moves into a heat exchanger. The heat is released into the environment and is condensed into a heavier liquid which is moved to a holding tank. From the holding tank, liquid then travels again to the processor and the entire process starts over.

The system eliminates the need for fans and associated wiring. If adopted globally, the system could also save computer manufacturers more than $540 million dollars annually in manufacturing costs.

Source: Science Daily

NEW DATA CENTER DESIGN REDUCES SPACE; INCREASES COOLING PRODUCTIVITY

3/24/15 - One company hopes to create an entirely new design for data centers for the industry. Vapor IO introduces a new "Data Center Aisle" concept. This vapor chamber design encompasses a room filled with round black cylinders that are nine feet in diameter. Vapor IO claims this chamber design is cheaper to produce and operate than the traditional data center designs. This design also reduces the amount of space needed for cooling due to its shape.

The concept was developed for edge data centers that are in or close to major population centers, storing data and content that needs to be delivered to users who live there.

The vapor chamber is ideal for situations where physical space is limited - it's designed to provide large capacity within a small footprint. A single chamber can accommodate up to 150 kW across six 42RU racks, according to Vapor IO. The chamber also offers rectifiers, PDUs, backup batteries, and fire detection and suppression.

Source: Data Center Knowledge

COMPANY BRINGS LIQUID COOLING TECHNOLOGY TO SMARTPHONES

4/27/15 - Fujitsu introduces a 1mm thick liquid cooling heat pipe for compact electronic devices such as smartphones and tablets.

Consumers constantly feel their smartphones and tablets become hot after continuous use and Fujitsu believes its new heat pipe will solve this problem. The heat pipe can transfer more than five times the heat than current heat pipes in the industry.

"The miniature heat pipe is a closed system comprising of an evaporator and a condenser. The evaporator will sit closer to a hot spot (CPU or GPU) and the condenser would be located at a relatively cooler section of the phone. Tiny pipes will connect the two to form a closed loop. The evaporator has six perforated sheets of copper, each about 0.1mm thick. The heat from the hotspot causes the liquid to vaporise and the vapor line will carry these vapors to the condenser where the lower temperature causes the vapours to condense back to liquid state and release heat energy in the process. The liquid moves in the system through capillary action which makes the orientation of the smartphone irrelevant," according to Fujitsu.

Source: Crazyengineers.com

Datebook

JUNE 9-12

Guangzhou International Lighting Exhibit
GUANGZHOU, CHINA
www.lightstrade.com

JUNE 16-20

THERMPROCESS 2015
DÜSSELDORF, GERMANY
www.mdna.com/trade-shows/details/thermprocess

AUGUST 5-6

Advancements in Thermal Mgmt. 2015
DENVER, COLO. U.S.
www.thermalnews.com/conferences

SEPTEMBER 16-19

LED China 2015
SHANGHAI, CHINA
www.chinaexhibition.com

VIRUS USED TO SLOW DOWN WATER BOILING; IMPROVES COOLING OF ELECTRONICS

4/14/15 - Matthew McCarthy, a scientist from Drexel University, has discovered the process of boiling water can be sped up with aid of a virus.

McCarthy focused his research on the point of contact between the surface and the water. He noticed that when liquid boils, it creates an insulating vapor between the water and the boiling surface, which slows down the boiling process.

In order to quickly distribute the vapor and to keep the boiling surface as wet as possible, McCarthy used a virus found on tobacco leaves, the Tobacco mosaic virus. This virus is surrounded by protein strands; when the boiling surface is coated with the virus, the protein strands, after coated in nickel, increase surface area expels vapor faster.

McCarthy's research will improve efficiency and safety of boiling water in many industries, including electronics cooling.

Source: Electronics Weekly

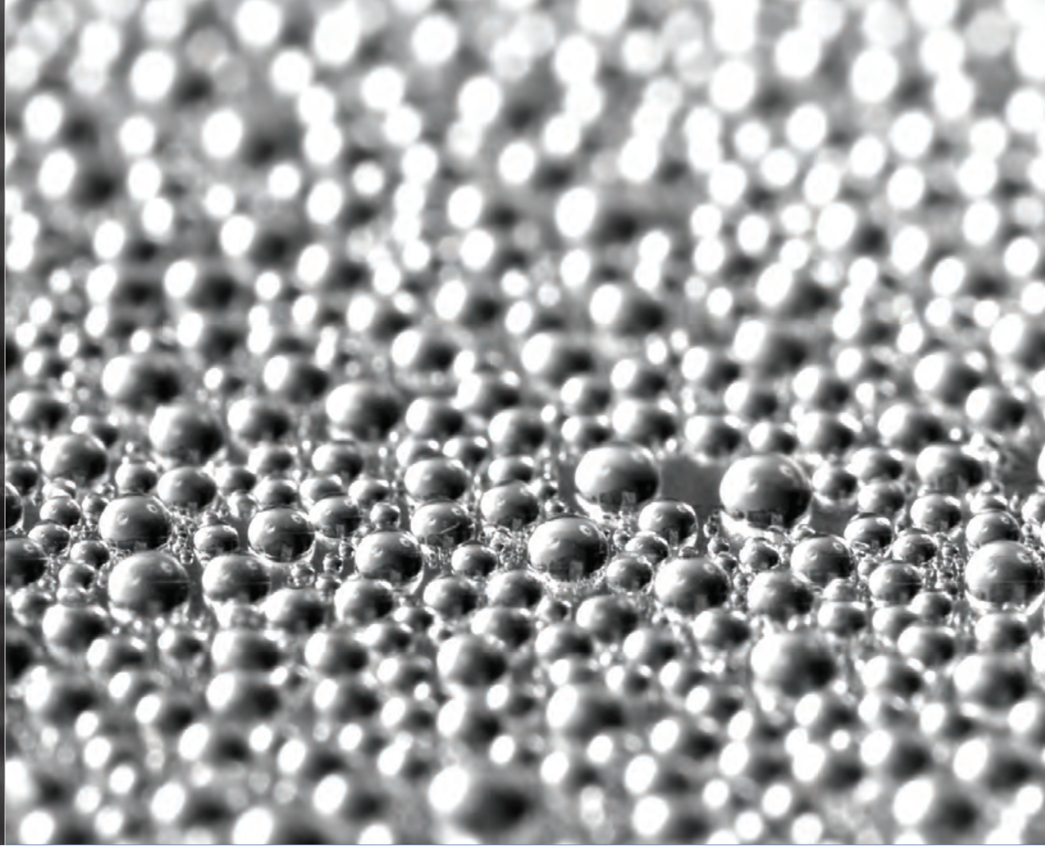

OVERHEATING ISSUES STALL RELEASE OF SONY'S XPERIA Z4

3/24/15 - The release of Sony's Xperia Z4 has been delayed due to heat dissipation issues in the Snapdragon 810 processor in the device. The Snapdragon 810 chip reportedly has overheating issues and Sony is working to fix the problems.

The Xperia Z4 is expected to be an extremely thin smartphone; thus the need for successful heat dissipation is important. The Xperia Z4 tablet is also powered by the Snapdragon 810; however, there is more surface area for heat dissipation on this 10-inch tablet than the small smartphone.

Source: Ubergizmo

IMMERSION COOLING EMERGES AS NEW COOLING TECHNOLOGY

3/9/15 - Immersion cooling, in which electronics are submerged in liquid coolants, is gaining popularity in the cooling industry. Companies have become aware of the benefits of immersion cooling in extreme environments, such as oil rigs or in the desert. The U.S. military is also considering it to save energy in tropical camps.

Liquid immersion cooling has proved more successful than other cooling methods, since air cooling and other methods still require fans or air conditioners. Immersion cooling also saves 20% on costs, 40% on power and 60% on space.

"Liquid cooling will grow at about 16 per cent per year through 2019. The military is expected to drive modular designs because it operates in remote locations and requires security and mobility," according to TechNavio.

Source: University of Michigan

NUTELLA USED AS UNUSUAL THERMAL PASTE

2/20/15 - On February 5, which was World Nutella Day, Cool Master celebrated the hazelnut-chocolate spread by using Nutella as a thermal paste.

The firm heated the Nutella first and cleaned the CPU before applying the spread. Surprisingly, the Nutella kept the test processor running at a maximum of 50 degrees Celsius for a while. The firm urges not to try this at home, since Nutella would eventually dry out and damage the processor.

Source: PCR Online

SEPTEMBER 22-24

Advanced Technology Workshop and Tabletop Exhibit on Thermal Management
LOS GATOS, CALIF., U.S.
www.imaps.org/thermal

SEPTEMBER 29-31

SAE 2015 Thermal Management Systems Symposium
TROY, MICH., U.S.
http://saeevents.org

OCTOBER 6-8

THERMAL LIVE 2015
ONLINE
http://thermallive2015.com

DECEMBER 7-10

Gartner Data Center, Infrastructure & Operations Mgmt. Conference
Las Vagas, NV, U.S.
www.gartner.com/events

A Simple Method to Understand Trade-Offs in Data Center Cooling

Dustin W. Demetriou, Ph.D.
IBM

COOLING AND THERMAL management are critical to data center reliability. Many organizations see cooling as a differentiating factor in the lifecycle cost of their data center. Recent industry guidelines [1] for data center cooling have suggested energy savings in air-cooled data centers by increasing the temperature of the cooling air. This increase enables two trends: reduced refrigeration power at higher refrigeration system evaporator temperatures and an increased number of hours available for using free-cooling. Free-cooling uses ambient air to provide cooling to the data center and reduce or eliminate the need for mechanical refrigeration. However, often overlooked is the significant power used in data centers for moving the cooling air. This article will use a simple thermodynamic and heat transfer analysis to highlight how trade-offs in Information Technology (IT) system power, computer room air conditioner (CRAC) fan power, and refrigeration power must be balanced to optimally operate the data center. Prior studies [2, 3, 4] have illustrated how simple models can be very effective at elucidating meaningful results.

NOMENCLATURE

COP	Coefficient of Performance
c_p	Specific heat
K_c	CRAC pressure coefficient
$\dot{m}$	Mass flow rate
N_c	Number of CRAC units
P	Power
T	Temperature
T_C	CRAC supply temperature
$\dot{Q}_c$	CRAC heat transfer
$\dot{V}_c$	CRAC air volume flow rate

GREEK SYMBOLS

ΔT	Temperature rise
η	Efficiency
λ	Leakage air flow as a fraction of active cooling air
ρ	Density
φ	Hot air recirculation fraction

SUPERSCRIPTS

i	Inlet
o	Outlet
$*$	Prescribed value

SUBSCRIPTS

c	CRAC
f	Fan
L	Leakage
r	Rack
t	Total cooling requirement
ref	Refrigeration
x	Data center exhaust

ESTIMATING THE AIR FLOW IN A DATA CENTER

Referring to Figure 1, which shows a diagram of a simple model of an air-cooled data center, the cold air emanating from the CRAC (at temperature T_C) is divided into an active cooling portion $\dot{m}_t$ and a leakage component $\lambda\dot{m}_t$, which bypasses the cold aisle and blends with the IT rack's exhaust air. The leakage could represent air that escapes through cable cutouts or through gaps in the raised-floor. The IT rack consumes a given power P_r, which is a combination of the power required by the electronics and the power of the IT system's cooling fans. The cooling air that enters the IT

rack is a combination of air supplied by the CRAC to the cold-aisle and a fraction ϕ of hot exhaust air that is recirculated back to the front on the IT equipment at the temperature of the CRAC return air T_x.

A mass and energy balance at the inlet of the IT rack shows that the rack inlet temperature must satisfy,

$$T_r^i = (1-\varphi)T_c + \varphi T_x \tag{1}$$

The inlet air, after being heated by the electronics, exits the rack at a temperature,

$$T_r^o = T_r^i + \frac{P_r}{\dot{m}_r c_p} \tag{2}$$

where, c_p is the specific heat of the air. The rack's exhaust air mixes with the leakage air to form the data center exhaust air. The data center exhaust temperature, which is both the recirculated air and the CRAC inlet air temperature, is given by an energy balance on the exhaust node as,

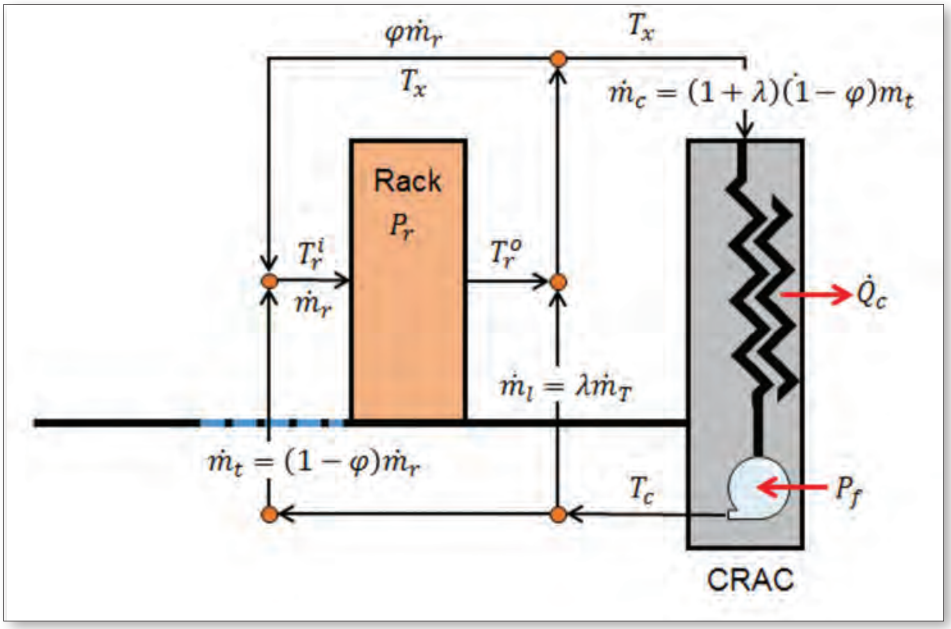

FIGURE 1 - Schematic of a Simple Model of an Air-Cooled Data Center.

$$T_x = \frac{T_r^o + \lambda(1-\varphi)T_c}{1+\lambda(1-\varphi)} \tag{3}$$

Equation 1 can be combined with Equation 3 to show that the data center's exhaust temperature must also satisfy,

$$T_x = \frac{T_r^o + \lambda T_r^i}{1+\lambda} \tag{4}$$

Lastly, an energy balance of the CRAC shows that the CRAC supply temperature must satisfy,

$$T_c = T_x - \frac{P_r}{(1+\lambda)(1-\varphi)c_p\dot{m}_r} \tag{5}$$

DATA CENTER COOLING POWER CONSUMPTION

The data center CRAC cooling power P_c is comprised of two components: a refrigeration component that is expect to be higher with lower CRAC supply temperature and an air moving component that is expected to vary with the air flow rate as,

$$P_f = \frac{N_c K_c}{\eta_f}\left(\frac{\dot{V}_c}{N_c}\right)^3 = \alpha\left[\frac{(1+\lambda)(1-\varphi)}{N_c}\dot{m}_r\right]^3 \tag{6}$$

where, N_c is the number of identical CRAC units in the data center, K_c is the CRAC pressure coefficient, η_f is the overall fan efficiency, assumed to be driven by variable speed motors, and $\dot{V}_c$ is the total CRAC volume flow rate. The refrigeration power consumption can be computed from,

$$P_{ref} = \frac{\dot{Q}_c}{COP} = \frac{P_r + P_f}{COP} \tag{7}$$

where, $\dot{Q}_c$ is the CRAC's cooling requirement and COP is the refrigeration system's Coefficient of Performance. It can be obtained from an energy balance of the data center. Therefore, the total cooling can be computed as, $P_c = P_{ref} + P_f$.

OPTIMIZED COOLING FOR AN ENCLOSED AISLE DATA CENTER

As an example, let's assume that the data center has implemented cold-aisle containment [2, 5]. Theoretically, in an enclosed aisle data center, the air provided through the perforated tiles must equal the air required by the racks (i.e., ϕ = 0.0). The above set of equations can be used to parametrically study the optimal operating condition, given values of λ, K_c, η_f ,the variations of rack power and rack flow rate, and the variation in CRAC COP. Table 1 provides the parameters used in the analysis.

The variation in rack flow rate and power is not straightforward and the design varies from manufacturer-to-manufacturer. This example uses the relationships provided by ASHRAE [1], which have been reproduced in Figure 2, which shows that, as a function of inlet temperature, the required power and air flow rate can increase by 20% and 250%, respectively, as the inlet temperature is increased from 15 to 35°C. The figures also highlight the variation in IT equipment design by expressing the relationships as a band, with high performance and highly utilized servers as the upper limit (ASHRAE HIGH) and lower utilized, lightly configured servers as the lower limit (ASHRAE LOW). To

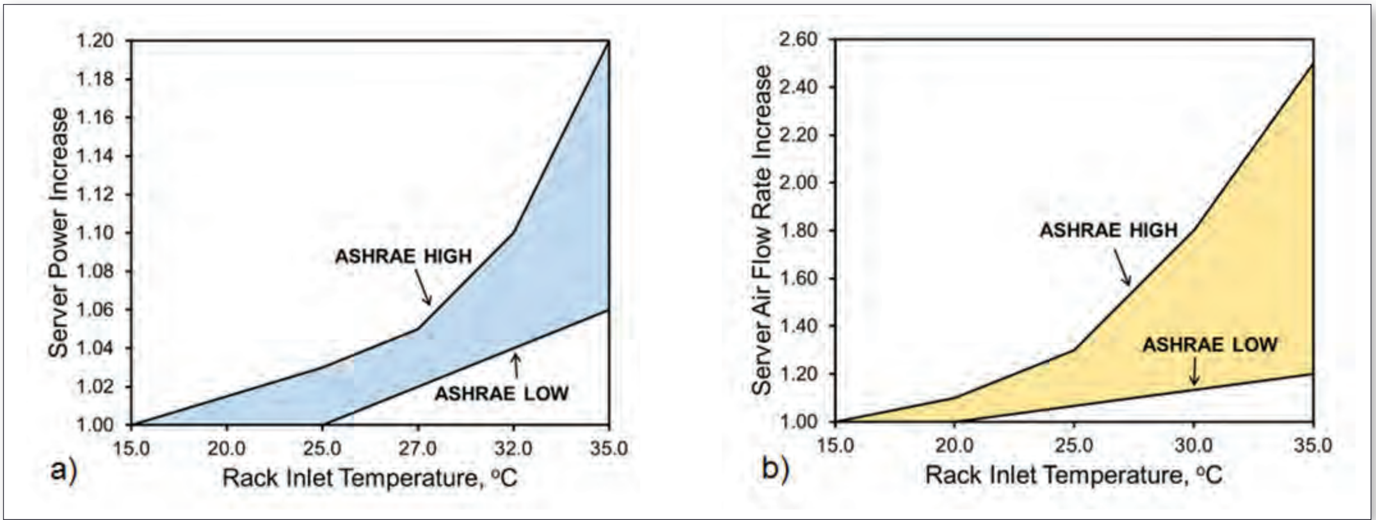

FIGURE 2 - a) ASHRAE Volume Server Power Increase with Inlet Temperature and b) ASHRAE Volume Server Flow Rate Increase with Inlet Temperature. Reproduced from [1].

use these figures, we will prescribe the rack power P_r^* at 15°C and the server's temperature rise ΔT_r^* at 15°C. The required rack flow rate at 15°C can then be computed from Equation 2. With the IT equipment characteristics defined at 15°C, the trends in Figure 2 can be used to find the power and flow requirements at higher inlet temperatures. For the CRAC performance, we will use the simple relationship described in [6], which expresses the CRAC COP as a monotonically increasing function of the CRAC supply temperature.

P_r^* (kW)	1000.0
ΔT_r^* (°C)	15.0
λ	0.20
K_c (Pa/(m^6/s^2))	20.0
N_c	12
η_f	0.65
ρ (kg/m^3)	1.225
c_p (kJ/kg-K)	1.012

TABLE 1 - Parameters Used in Optimized Cooling Example Analysis for an Enclosed Aisle Data Center

Using Equations 1 – 7, the relationships in Figure 2 and the parameters defined in Table 1, the data center's cooling power requirement can be computed as a function of rack inlet temperature. Figure 3 shows the cooling power as the rack inlet temperature is varied from 15 to 30°C. It shows that the optimum operating point (i.e., the point with lowest cooling power consumption) occurs around 23°C. The figure shows the trade-off that data center operators should consider. Increasing the inlet temperature increases the amount of power and airflow required by the IT equipment and thus the amount of air that must be delivered by the CRACs.

Even though the increased temperature enables a reduction in refrigeration power, this reduction is not sufficient to overcome the increased power requirement of the CRAC fans. Exclusive focus on reducing refrigeration power consumption without regard to the power consumed in moving the cooling air would lead to sub-optimum, possibly misleading results.

Figure 3 highlighted the trend in cooling power for ASHRAE low power and lightly utilized IT equipment. Using the trends from ASHRAE in Figure 2, Figure 4 shows how the design of the server's cooling algorithm, utilization and power consumption can impact the data center cooling optimization. High performance and/or highly utilized servers exhibit a clear data center minimum power consumption point; whereas, lower utilized and/or lightly configured servers, which do not have as significant air flow requirements, exhibit their minimum data center power consumption point at higher inlet temperatures. Interestingly, for this class of servers, the reduction in power consumption beyond 24°C is relatively small, compared to the savings going from 15°C to 24°C. Since uptime and reliability are paramount to data center design, a designer may consider operating at a lower temperature, and sacrifice small energy savings, to provide more resiliency in the event of a cooling failure.

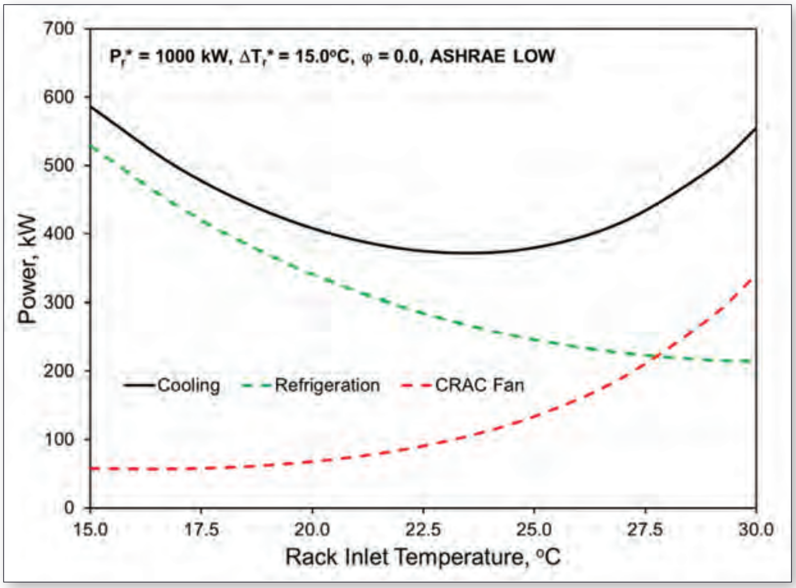

FIGURE 3 - Predicted Data Center Cooling Power Consumption Breakdown based on Equations 1 – 7 and Table 1 Parameters.

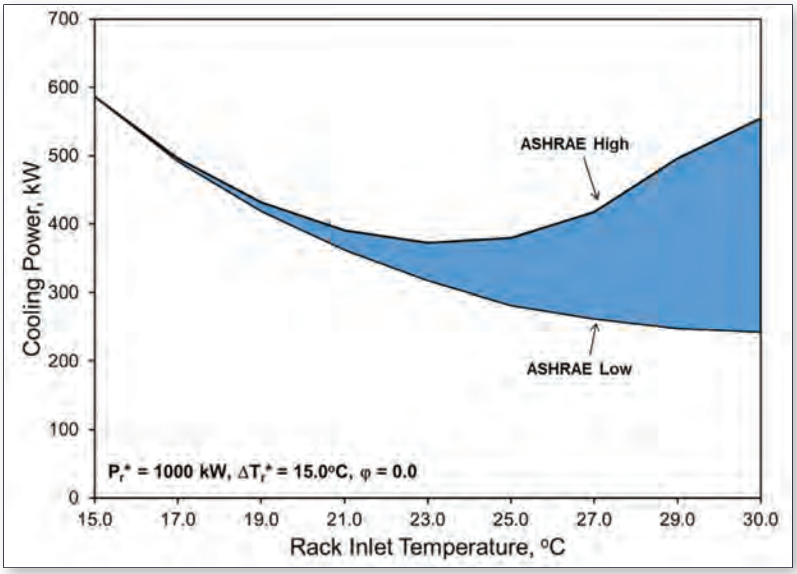

FIGURE 4 - Predicted Volume Server Impact on Data Center Cooling Optimization based on Equations 1 – 7 and Table 1 Parameters.

CONCLUSIONS

This article showed how a simple thermodynamic and heat transfer model can be effective to understand the most efficient manner for cooling IT servers in a data center. It highlighted the trade-off between refrigeration and air moving device power that data center operators need to consider in order to optimize their data center's cooling system. Hopefully, the article described to the reader a methodology that can be considered at the early design stage of any thermal-fluid system to fully understand the problem at hand, before considering more detailed and time-consuming techniques.

REFERENCES

[1] ASHRAE, 2012, Thermal Guidelines for Data Processing Environments, 3rd Edition, Atlanta.

[2] Demetriou, D.W., Khalifa, H.E., 2012, "Optimization of Enclosed Aisle Data Centers Using Bypass Recirculation," Journal of Electronics Packaging, Vol. 134, 020904.

[3] Khalifa, H.E., Demetriou, D.W., 2010, "Energy Optimization of Air-Cooled Data Centers," Journal of Thermal Sciences and Engineering Applications, Vol. 2, pp, 041005-1:13.

[4] Simons, R.E, 2007, "Using a Simple Air Recirculation Model to Explore Computer Rack Cooling," Electronics Cooling. http://www.electronics-cooling.com/2007/02/using-a-simple-air-recirculation-model-to-explore-computer-rack-cooling/, last accessed April 25, 2015.

[5] Muralidharan, B., Shrivastava, S. Ibrahim, M., Alkharabsheh, S., Sammakia, B., 2013, "Impact of Cold Aisle Containment on Thermal Performance of Data Center," Proceedings ASME InterPACK, Vol. 2, Burlingame, CA, July 16 – 18.

[6] Moore, J., Chase, J., Ranganathan, P., Sharma, R., 2005, "Making Scheduling "cool": Temperature-aware Resource Assignment in Data Centers," in 2005 Usenix Annual Technical Conference, April 2005.

INTRODUCING COOLSPAN® TECA

thermally & electrically conductive adhesive

Rogers can help by being your reliable conductive adhesive film source

Get the heat out of those high-power PCBs. COOLSPAN® Thermally & Electrically Conductive Adhesive (TECA) Films are ideal for dissipating heat in high-frequency circuits. COOLSPAN adhesives feature outstanding thermal conductivity (6 W/m/K) and reliable thermal stability. Keep things cool, with Rogers and COOLSPAN TECA film.

ROGERS CORPORATION

www.rogerscorp.com

MEET YOUR COOLSPAN® TECA FILM SUPPORT TEAM

Leading the way in...
• Support • Service • Knowledge • Reputation

SCAN THE CODE TO GET OUR CONTACT INFO.

Greg Bull
Applications Development Manager
Northern Territory (U.S.) & Canada

Dale Doyle
Applications Development Manager
Western Territory (U.S.)

John Dobrick
Applications Development Manager
Southern Territory (U.S.) & South America

Scott Kennedy
Applications Development Manager
Eastern Territory (U.S.)

John Hendricks
Regional Sales Manager
Europe

Kent Yeung
Regional Sales Director
Asia

If you are unable to scan a VR code please visit our Support Team website at www.rogerscorp.com/coolspan

The Holy Books of Heat Transfer: Facts or Fairy Tales?

"There must be an ideal world
A sort of mathematicians' paradise
Where everything happens
As it does in textbooks."

- Bertrand Russell

Clemens J.M. Lasance
Guest Editor

THE ANSWER TO the question stated in the title is: both! It all depends on whether or not your situation is truly similar to the situation described in the book (physically congruent, fluid mechanically and thermally similar). Very few heat transfer books provide detailed dimensions and flow conditions for the experimentally observed data they contain– more often they simply describe the overall situation briefly and idealistically. In other words, it all depends if you believe in the axioms or postulates that are, more often than not, hidden in the introductory chapters. I found it enlightening to see the parallelism with some other Holy Books, and hence this column is about why people believe in books or trust gurus (I do not belong to this category of course). It is all about the acceptance of axioms or assumptions. In mathematics, an axiom is any mathematical statement that serves as a starting point from which other statements are logically derived. In common language, an axiom is a premise so evident as to be accepted as true without controversy. However, progress in science has learned otherwise. It is not correct to say that axioms are propositions that are regarded as true without proof. Rather, it is better to look at axioms as a set of constraints. If any given system (or book) satisfies these constraints, then one is in a position to instantly know a great deal of extra information about this system (or book). Applied to heat transfer this means that if and only if the physical situation conforms to the assumptions does the experimenter have the right to assume that the predicted results will be obtained. That means much more than simply matching Nusselt (Nu) and Reynolds (Re) numbers. Specifically, all analytical (and numerical) studies assume uniform flow velocity with a specified turbulence (often zero), uniform temperature and the origins of the velocity boundary layer and the thermal boundary layer on the surface.

Let me quote from my column of April 2010 [1]: "Having read myriads of papers/articles/books/reports on thermal management I feel there is a lot of misunderstanding about what really should drive a sound approach of how to tackle the thermal problems that land on the desk of thermal designers." The problem in a nutshell is that the handbooks are inherently based upon a set of constraints that are not satisfied in real-life. Even more important, this set is rarely described. When you believe in the following axioms, then the Holy Books of Heat Transfer are consistent and comprise a wealth of information, very useful for a basic understanding of the physics.

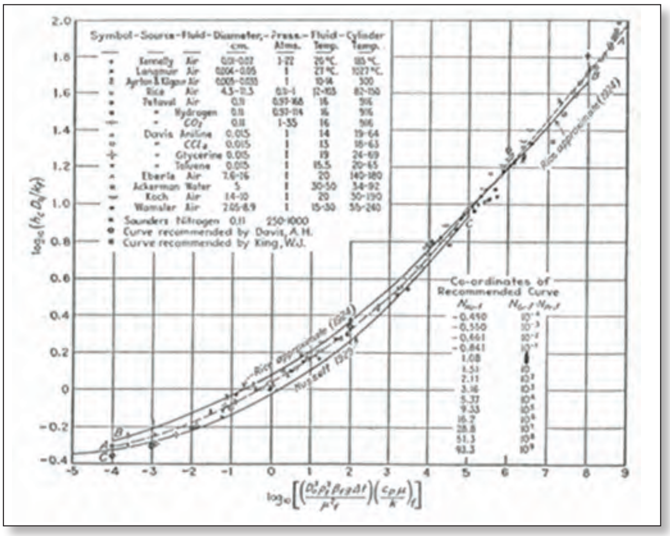

Figure 1 - Nu-Gr correlation for horizontal cylinders [6].

- Uniform boundary conditions, either constant temperature or flux
- Uniform approach flow with a degree of turbulence as close as possible to zero (that's why research type wind tunnels are huge).
- (Very) simple geometries: smooth, flat and thin plates, parallel plate channels, pipes
- Single source, especially for natural convection
- Constant properties
- Fan dynamics based on air flow chamber testing
- Extended surfaces based on Murray-Gardner assumptions
- 'Complex shapes' means there exists an analytical solution
- Heat spreading limited to one-layer, one-sided heat transfer
- Radiation diffuse and grey

For many fields of heat transfer, such as turbulence, boiling, heat exchangers, channel flow, etc., these axioms form a sound base. Not so for thermal management of electronics. When your situation is consistent with all of the assumptions described above (geometric, fluid mechanic and thermal), then the presented results can be immediately useful. Otherwise, they must be taken only as "data taken along an idealized path" through the real world that you are working in. It is then up to you to decide how much to raise or lower the provided values of h.

There is another aspect of heat transfer theory that deserves attention, with their own axioms, and this is correlation theory that tends to pollute every heat transfer book or paper. A long time ago I coined the term 'Correligion', because that is exactly how it feels to someone like me who is not trained as a mechanical engineer but as a physicist: a religion, based on axioms that are subject to serious doubt in practice. Quoting Dr. Bob Moffat: "Years of poorly controlled and inadequately described experiments have filled the literature with data that appear to be 'comparable' but are not. It is the result of standard practice: all data 'must' be made dimensionless and collapsed on log-log paper. The average user seems to accept the fact that there will be always ± 20-50% scatter in heat transfer data.Reynolds number scaling is much more subtle than many heat transfer researchers think." Why this scatter you may wonder that would be unacceptable in other fields of science? Because it seems forgotten that correlations only make sense when they obey the following three axioms:

- Congruency
- Similarity
- The boundary layer approximation

Ref [2] shows that in most practical cases especially the lack of congruency causes the problems. The axiom of congruency implies that all dimensions, velocities, times and forces should scale. Dr. Moffat and his students have published some interesting results, for example that the accuracy of your correlation increases significantly when you scale also your wind tunnel. The consequence in practice? Suppose you want to study flow between two plates with a variable space. When you change this dimension, formally you should also scale the plate thickness, the plate area and your wind tunnel!

In ref [3], I point to another phenomenon that puzzled me for a long time: the fact that Nu (and hence the heat transfer coefficient) is constant for fully-developed channel flow, independent of the velocity. As a physicist, this does not make sense to me. I showed that if you rewrite the local h in an adiabatic form, the results do make sense.

Another flaw, I would not call it an error, is the curious fact that by using dimensionless equations there is a danger of imposing a slope, caused by confounding of those parameters that are common to both sides of the equation by making the equation dimensionless, first discussed by Wilkie [4]. I discussed this flaw extensively in refs.[2] and [5], but I would like to repeat here an interesting interpretation of one of the famous and often cited graphs in one of the earliest Holy Heat Transfer Books [6], which is reproduced in fig. 1. Especially when a large data set is being used, for example when one of the common variables has been varied over a wide range, an imposed slope becomes clearly visible. In fig. 1, we see a correlation for natural convection from horizontal cylinders spanning more than 13 decades in Grashof (Gr) number. Note that the upper-half slope, which represents turbulence, is ⅓, which is exactly the exponent for the imposed slope. The conclusion that turbulent heat transfer scales with the same slope is not warranted at all. The point is that you get, by definition, the ⅓ slope if you randomly vary the heat transfer coefficient and put the results on double-log paper. Hence, is it a fact or a fairy tale that the turbulent heat transfer coefficient scales with the velocity to the power ⅓? You may not conclude this from the graph alone. What one should do, is to plot convective heat transfer coefficient (h) versus velocity (v), with the diameter as the parameter.

I quote here only one Holy Book. I could quote many others, but it suffices to say that I have read them all, including many that are specifically devoted to thermal management of electronics. None of them devote a chapter or even a paragraph to the problems designers are facing who have to deal with the complex geometries and the non-uniform boundary conditions, and the consequences thereof for applying the plethora of equations derived for simple geometries.

Disclaimer: I am not discussing the tremendous value of these books for other fields of heat transfer!

I challenge all readers to show me and the other readers that I am wrong. This, I hope, would start at last a discussion on a very important topic in electronics cooling: the huge gap what is required in real-life and what is found in holy books: "What did you learn in school today, dear little boy of mine?" "About little boxes made of ticky tacky, little boxes all the same."

Note: the original version continued here with a discussion on how perfectly rational the arguments of believers in the real Holy Books are, provided the postulates are taken for granted. It continues with an overview of modern fiction-based religions such as the Tolkien Religion. This later discussion is not included here since it is considered to not fit within the framework of this magazine. The complete version may be obtained by contacting the author at lasance@onsnet.nu.

ACKNOWLEDGMENT

Many thanks to Dr. Bob Moffat who commented upon an early version. I gratefully used his valuable critic. Here is his recommendation when you still would like to use that impressive correlation, despite all the foregoing:

Be sure you know the following six things about the situation it describes:

- The Velocity Field (Reynolds number)
- The Fluid Properties (Prandtl number)
- The Shape of the Test Rig and the Specimens
- The Temperature Boundary Conditions
- The Temperature Initial Condition
- The Reference Temperature

REFERENCES

[1] Lasance C., Most of us live neither in wind tunnels nor in the world of Nusselt, Electronics Cooling, April 2010.

[2] Lasance, C., "Sense and Nonsense of Heat Transfer Correlations Applied to Electronics Cooling", Proc. EUROSIME 2005, pp. 8-16.

[3] Lasance C., Fully-developed channel flow: why is Nu constant? Electronics Cooling, September 2010.

[4] Wilkie D., 1985, The Correlation of Engineering Data Reconsidered, Int. J. Heat & Fluid Flow, vol.6, pp. 99-103.

[5] Lasance C., Does Your Correlation Have an Imposed Slope? Electronics Cooling, September 2011.

[6] McAdams W., 1954, Heat Transmission, 3d Edition, McGraw-Hill, p. 176.

An All-Silicon Field Programmable Thermal Emulator for Integrated Circuits

Wen Yueh, Khondker Z. Ahmed, and Saibal Mukhopadhyay,
School of ECE, Georgia Institute of Technology

INTRODUCTION

THERMAL TEST CHIPS with integrated thin-film heaters via post-processing are often used as the test structure to evaluate thermal integrity of the packages and cooling solutions [1]. In advanced silicon technologies (sub-45nm nodes), the chip's thermal design point (TDP) is often a determining factor for the performance and energy-efficiency in multi-core microprocessors [2]. However, time-varying workloads in processors can lead to time-varying temperature, and impact the integrated circuit (IC's) performance and power. Therefore, it is crucial for the future test vehicles to characterize the interactions among application dependent power patterns, chip/package's thermal properties, cooling technologies, and circuit functionalities. The thermal test-chip should be designed and fabricated in the standard silicon process and may be embedded in the packaging and cooling technology candidates. The chip's programmability to emulate the time-varying power patterns of real applications is required for on-chip characterization of spatiotemporal temperature variations and its impact on the circuit's performance, power, and reliability; instead of only ensuring thermal integrity at the peak power density. An all-silicon structure with on-line programmable power and on-chip characterization sensors integrates well with the IC design flow, and is a preferable solution for holistic power, temperature, and circuit analysis. This article introduces a thermal test structure, referred to as the field programmable thermal emulator (FPTE), to achieve this goal.

FIELD PROGRAMMABLE THERMAL EMULATOR

Parallel to the field programmable gate array (FPGA) in functional emulation, FPTE emulates the thermal field using programmable on-die CMOS based heater array and on-die sensors for temperature and circuit properties. On die programmable heaters are controlled with the integrated registers, to emulate time-varying power patterns and generate time-varying temperature pattern. Multiple digitally programmable FPTE cores are integrated on-chip to characterize the thermal effects in multi-core processing including core-to-core thermal coupling. The FPTE cores are augmented with the analog temperature sensors to record the temperature patterns and digital circuits to characterize the changes in the electrical characteristics. The conceptual diagram in Figure 1 shows an FPTE chip integrated in a board. An on-board microcontroller can be used to drive test applications and record the sensor outputs for automated measurement and characterization. The

Wen Yueh received his B.S. and M.S. degrees in Electrical & Computer Engineering from Rutgers University, New Jersey, in 2009. He is currently pursuing his Ph.D. degree at Georgia Institute of Technology, Atlanta. His research interest includes system level multi-physics simulator, self-adaptive circuit design for many-core processor thermal management, and energy-aware low power memory architecture.

Khondker Zakir Ahmed received BSc. in Electrical and Electronic Engineering from Bangladesh University of Engineering and Technology (BUET) in 2004. He was an Analog IC Designer in Power IC Ltd., Dhaka, Bangladesh from 2005 to 2010. Since 2011, he has been pursuing his PhD in Electrical and Computer Engineering at Georgia Institute of Technology, in Atlanta, Georgia. Mr. Ahmed was a summer intern in 2013 and 2014 at Intel's Circuit Research Lab (CRL) in Hillsboro, Oregon; where he worked on efficiency improvement techniques for integrated voltage regulators. His research focus at Georgia Tech includes power management circuits for low voltage energy harvesting systems and on-chip integration of voltage regulator with digital loads.

Prof. Saibal Mukhopadhyay is an Associate Professor at School of ECE, Georgia Institute of Technology. He has received a Ph.D. from Purdue University in 2006. His research interests include power and thermal management, low-power systems, and nonconventional computing. Dr.. Mukhopadhyay has received multiple best paper awards including the IEEE Transactions on Component and Packaging Best paper Award in 2014. Dr. Mukhopadhyay was a recipient of the Office of Naval Research Young Investigator Award in 2012, the National Science Foundation CAREER Award in 2011, and the IBM Faculty Partnership Award in 2009 and 2010. He has co-authored over 150 papers in refereed journals and conferences, and holds five U.S. patents.

Powerful Solutions

FOR ALL THERMAL MANAGEMENT NEEDS

AAVID THERMALLOY

Aavid Thermalloy offers the widest array of thermal management solutions and thermal engineering services **IN THE WORLD!**

With a standard catalog of thousands of cooling products as well as manufacturing and design centers around the globe, Aavid Thermalloy is the world's largest provider of innovative thermal solutions.

For over 50 years we have been developing advanced cooling technology to conquer every thermal challenge that has emerged in the marketplace. Aavid understands that the right cooling solution plays an integral part in the performance of all electronics across every industry and in the ultimate success of your design.

Our engineers are constantly working to expand our catalog and improve the performance of our standard products from stamped and extruded heat sinks to innovative liquid cold plates utilizing Aavid's HiContact™ process. We also have dedicated engineering, design and prototype teams focused on creating modified or fully custom solutions for any thermal challenge from the board to the facility level.

1.855.32.AAVID

WWW.AAVID.COM

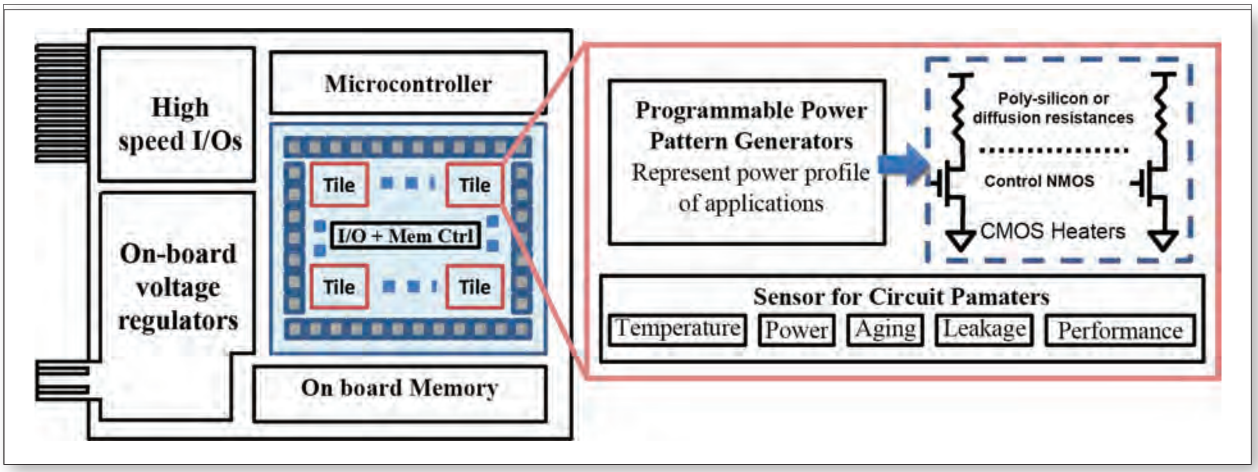

Figure 1 - A conceptual diagram of a field programmable thermal emulator (FPTE) integrated in an instrumentation board for thermal characterization.

cartoon highlights an on-line field programmable thermal characterization framework on board.

AN EXAMPLE FPTE TEST-CHIP

A simplified FPTE test-chip to demonstrate the design concept is presented in Figure 2 [3]. The die photo of the designed test chip shows five digitally controllable heaters, five digital sensors, and serial-to-parallel-interface (SPI) registers. The on-chip SPI programs heater registers for heating and records data from the sensor registers [3]. The FPTE hotspots at the corners are 600 μm apart horizontally and 750 μm apart vertically. An additional hotspot is placed at the center of the chip. The core of the FPTE block is the programmable CMOS heater that is made of n-well resistor to generate heat within the silicon at the junction. An NMOS transistor with binary ('high' and 'low') states controls the current through the resistor. The resistor bank is arranged in groups of resistors with maximum equivalent resistance as 250 ohm. Each resistor bank is form by 4 x 4 sub-bank tiles. Each tiles are 60 μm x 50 μm in size and shares the control signal to improve intra-bank uniformity in generated heat. Each resistor bank is designed to generate 16 current levels with maximum 165 mA at 3.3 V, resulting in the maximum instantaneous heater power of 544.5 mW at the granularity of 34 mW. The total area of the heater hotspot is 250 μm x 150 μm i.e. 0.0375 mm2. Each FPTE cell dissipates ~9 uW of standby power. The maximum power of each bank may be controlled using an external voltage source or on-chip voltage regulators. The internal registers can then tune the individual heater output at a fine-grain level. The latch based registers for heaters and sensors have a footprint of 250μm x 150μm. The fill-factor for a single FPTE is designed to be 50% and the density may further improve if SRAM is used for the register cells. The heater drivers are designed to sink the full swing current with rise/fall time at ~20 ns.

The conventional bipolar junction transistor (BJT) based analog temperature sensors are included in the design. The outputs of the analog sensors are quantized using external analog-to-digital converters (ADCs). To characterize the interaction between delay and temperature patterns, digital ring oscillators (performance sensors) are integrated within the FPTE blocks. The delay based sensors store temperature history in 8 bit word and the first-in-first-out (FIFO) buffer holds up to 16 entries in the register.

MEASUREMENT RESULTS

The test chip was designed and fabricated in 130 nm CMOS process. Figure 3a shows the measurement setup to verify the operation of the FPTE. A program stored on the microcontroller controlled the heaters and obtained the sensor readings. The automated measurement demonstrates the feasibility of a high-speed test and characterization using the platform. Figure 3b shows the experimental result where a time-varying arbitrary power pattern is generated in the chip and recording from the delay based digital temperature sensors. The test case demonstrated programmable time-varying power pattern in the heaters. The FPTE sensors identified the coupled time-varying temperature. The measured reading from the delay sensors showed that due to the power-thermal-performance correlation, the power dissipation in the chip modulates a circuit's performance. The arbitrary power profile may be driven by realistic power trace from measurements of current processor or from architecture simulators for predictive architectures. The measurements can be repeated for different packaging conditions as well as using different external cooling technologies to understand their transient properties.

APPLICATIONS

The FPTE uses on-chip digital heaters and sensors to emulate time-varying power patterns on-chip, generate the corresponding spatiotemporally varying temperature pattern, and characterize the resulting variations in circuit properties. The

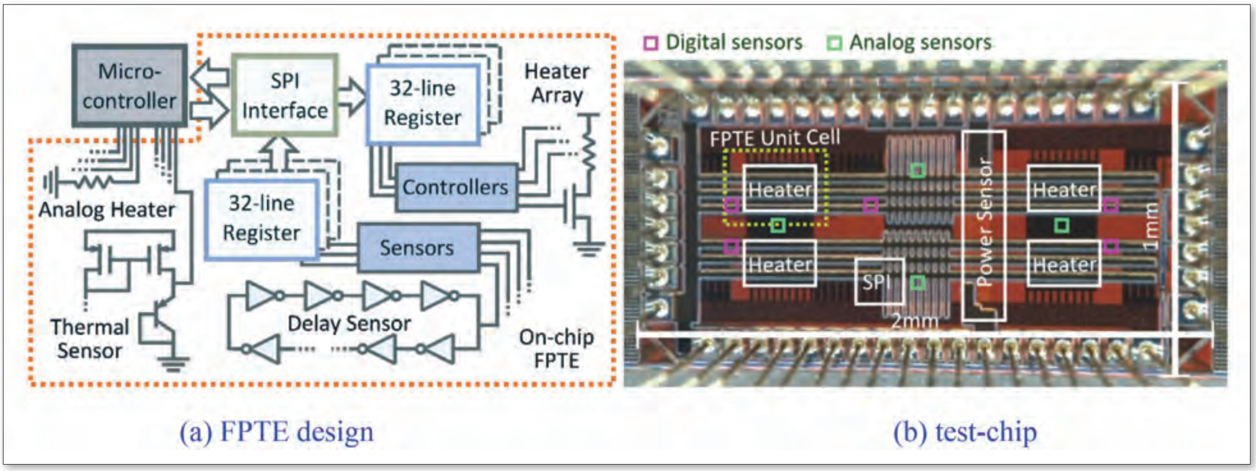

Figure 2 - . An example implementation of an FPTE design and corresponding test-chip in 130nm CMOS [3].

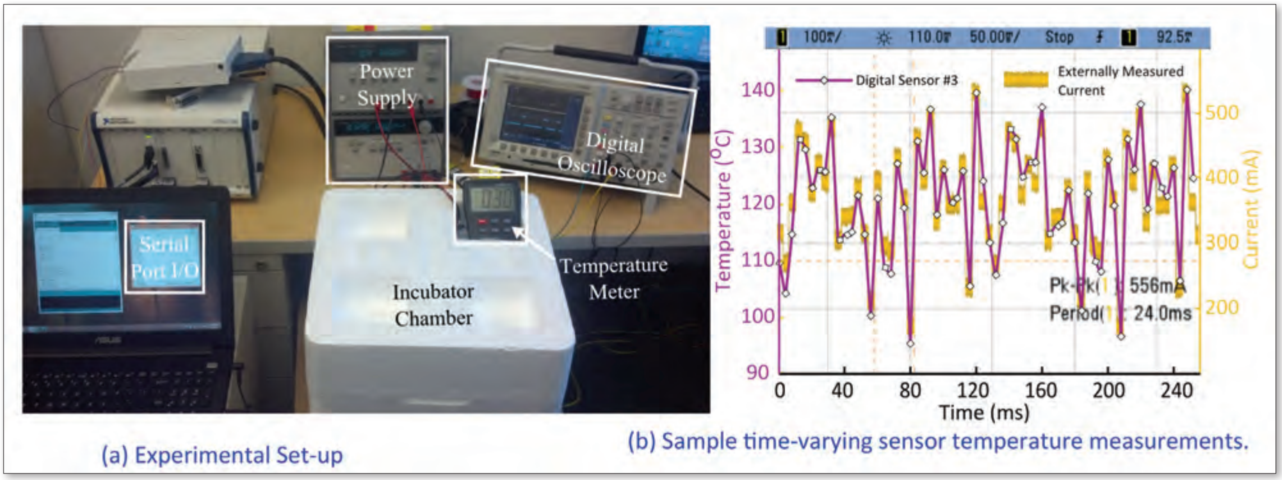

Figure 3 - A snapshot of the measurement environment and sample results.

application domains of FPTE include a thermal test-vehicle as well as an on-line thermal test-structure.

As a thermal test-vehicle, FPTE has the advantage over existing thin-film heater based approaches due to its compatibility with standard CMOS process, ability to generate controllable and time-varying power patterns, and directly characterize the effect of temperature patterns on device characteristics. Application of FPTE as a test-vehicle to evaluate advanced microfluidic cooling has been presented by Wan *et. al.* [4]. In the experiment, the silicon interposer with etched micro pin-fins was attached directly with the FPTE die. The FPTE was used to directly characterize the effect of fluidic cooling on the circuit properties, for example, the trade-off between flow rate (cooling power) and the potential leakage power saving.

As an on-chip test structure, an FPTE blocks may be embedded within a microprocessor core with minimal overhead, because its low standby power and small area. A built-in self-test routine may apply test power patterns to validate/ensure thermal fidelity of a specific packaged IC considering process variations as well as the time-dependent degradation in the thermal properties [5]. The thermal characteristics of packaged chips may be extracted in the form of frequency domain thermal filters as presented by Kung et. al. [6]. The extracted filter was used to accurately predict transient temperature pattern for spatiotemporally varying power patterns. The filter also estimate power pattern from measured temperature accurately, which implies FPTE can facilitate on-line real-time temperature/power prediction in an IC.

PLANS FOR FUTURE DEVELOPMENT

The FPTE test-chip presented in Figure 2 can be extended to enrich the characterization framework. The increasing the density of the heater cells will help emulate higher hotspot density. Currently, the design is capable of programing arbitrary power pattern over time; a potential extension is to improve the spatial controllability to generate arbitrary power pattern in space to enable true spatiotemporal controllability. The presented design depends on externally available power pattern to emulate a workload. In future, it will be important to consider built-in 'power pattern generators' that consists of fundamental logic and memory blocks with controllable activity patterns. Finally, it will be interesting to pursue embedding of FPTE as an Intellectual Property (IP) block within a processor design.

CONCLUSIONS

In integrated circuits, especially in high performance microprocessors, the run-time thermal management is critical to ensure performance, power, and reliability of the IC. A CMOS based design is effective to emulate thermal effects of any time-varying power patterns utilizing standard digital I/Os. The CMOS-only test vehicles, such as the FPTE, reduce test cost, improve testability, and extend observability. The FPTE structure presented in this article is an important step towards this direction and opens opportunities for future studies.

ACKNOWLEDGMENTS

This work is supported in part by Semiconductor Research Corp (#2084.001) and Sandia National Laboratory.

REFERENCES

[1] B. Siegal and J. Galloway, "Thermal Test Chip Design and Performance Considerations," in IEEE SEMI-THERM 2008, pp. 59-62, March 16-20, March 2008.

[2] J. Kong, S. W. Chung, and K. Skadron., "Recent Thermal Management Techniques for Microprocessors," ACM Computing Surveys vol. 44(3), pp. 13-42, June 2012.

[3] W. Yueh, K. Z. Ahmed, and S. Mukhopadhyay, "Field Programmable Thermal Emulator (FPTE): An All-silicon Test Structure for Thermal Characterization of Integrated Circuits," in IEEE SEMI-THERM 2014, pp. 66-71, March 9-13, 2014.

[4] Z. Wan, W. Yueh, Y. Joshi, and S. Mukhopadhyay, "Enhancement in CMOS Chip Performance through Microfluidic Cooling ," in THERMINIC 2014, pp. 24-26, September 2014.

[5] M. Cho, W. Song, S. Yalamanchili, and S. Mukhopadhyay, "Thermal system identification (TSI): A Methodology for Post-silicon Characterization and Prediction of the Transient Thermal Field in Multicore Chips," in IEEE SEMI-THERM 2012, pp.118-124, March 18-22, 2012.

[6] J. Kung, W. Yueh, S. Yalamanchili, and S. Mukhopadhyay, "Post-silicon Estimation of Spatiotemporal Temperature Variations Using MIMO Thermal Filters," IEEE Trans. Compon. Packag. Manuf. Technol. 2014, in press.

Data Center Energy Savings: Total Liquid Cooling Versus Indirect Liquid Cooling

Yong Quiang Chi[1], Peter Hopton[1], Keith Deakin[1], Jonathan Summers[2], Alan Real[2], Nik Kapur[2] and Harvey Thompson[2],
[1]Iceotope Ltd., [2]University of Leeds

Yong Quiang Chi is a PhD Student at the University of Leeds sponsored by Iceocope Ltd and researching into the natural-convection based thermodynamic performance of liquid-cooled data centre systems.

Peter Hopton is a green tech entrepreneur, he provides innovative solutions to help improve energy efficiency in IT and solve the real environmental problems within the industry. CVO and Founder of Iceotope, Peter invented and patented the original Iceotope server cooling technology.

Keith Deakin is Design and Development Director at Iceotope. Over the last 6 years his passion and determination has driven him to develop an innovative concept into an efficient and commercialised product; and a band of 3 into a tight team of 33 skilled and dedicated people.

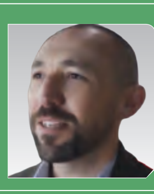

Jon Summers leads a group of researchers looking at thermal management of information systems from chip to data centre. Having developed and operated High Performance Computing (HPC) systems, Jon chairs the HPC user group at the University of Leeds.

Alan Real heads Advanced Research Computing at the University of Leeds, is Technical Director of the N8 regional HPC consortium and is the current Chair of the UK High Performance Special interest group (HPC-SIG).

Nik Kapur holds a chair in Applied Fluid Mechanics and works extensively at the academic/industry interface to bring fundamental science to bear on real problems. His interests are in flows on all length-scales – from microfluidics through to industrial scale and his research highly interdisciplinary in nature.

Dr. Harvey Thompson is Professor of Computational Fluid Dynamics and Director of Research & Innovation in the School of Mechanical Engineering at the University of Leeds, UK. His research interests are in heat transfer, industrial flow and process optimization and he has published more than 120 papers in these areas

INTRODUCTION

IMPROVEMENTS IN ENERGY efficiency and performance of data centers are possible when liquid is supplied to the racks [1]. A solution which has become popular for dense racks is the rear-door water-cooled heat exchangers (RDHx) [2]. Moving liquid to the rack reduces the travelling distance of conditioned air, however RDHx solutions rely on air-cooled design datacom components and is classified as indirect liquid cooling (ILC).

There are a growing number of suppliers that provide liquid cooled solutions, where the liquid directly cools the datacom equipment, classified as direct liquid cooling (DLC) [3]. A subset of these does not require any form of supplementary cooling via air and are classified as total liquid cooling (TLC), which rely on the use of non-conducting (dielectric) liquids that are either pumped [4], rely on phase change [5] or make use of natural liquid convection [6].

METHOD OF COMPARING ILC AND TLC

The configuration and experimental data for the comparison of the two cooling approaches is acquired from two operational systems and the collected data is applied to construct two design systems for comparison of ILC and TLC, where the IT is exactly the same. A single 2U rack chassis node with a single motherboard, converted from one of the fully immersed liquid-cooled systems was employed for comparison of the two cooling methodologies. The

only physical difference between the 2 systems is in the rack level cooling methods, that being ILC and TLC.

The enclosed, immersed, TLC system was deployed in a laboratory environment as shown in Figure 1(a), where the TLC modules (1), rack (2) and radiators heated by return water from the TLC modules (3) are indicated. Datacom performance, PDU power measurements and temperatures obtained from this system are used for the analysis of the rack level performance and power consumption. The cooling configuration of the ILC High Performance Computing (HPC) system (see Figure 1(b)) uses fan-assisted rear door liquid-loop heat exchangers connected to 500kW of external chillers, which has provided power and cooling data for the comparison.

The IT, rack and data center power and thermal data acquired from both real systems is used to construct the design systems with identical IT hardware, but comparing the performance of the two different cooling strategies. The same server level motherboard and CPU arrangement (namely SuperMicro X9D with Intel dual 2670 Xeons) and the Airedale Ultima Compact FreeCool chillers for the final heat removal step are included to complete the total energy calculation. Power consumption measurements were performed with (i) the rack of TLC modules, and (ii) a single node of the total liquid-cooled system was removed from its enclosure and run in a single air-cooled rack configuration to conduct the same power consumption measurements, but cooled by air and fans rather than the dielectric liquid. The results from these 2 tests provide the basic thermal and power consumption data to build up the design based total liquid-cooled and indirect liquid-cooled data centers.

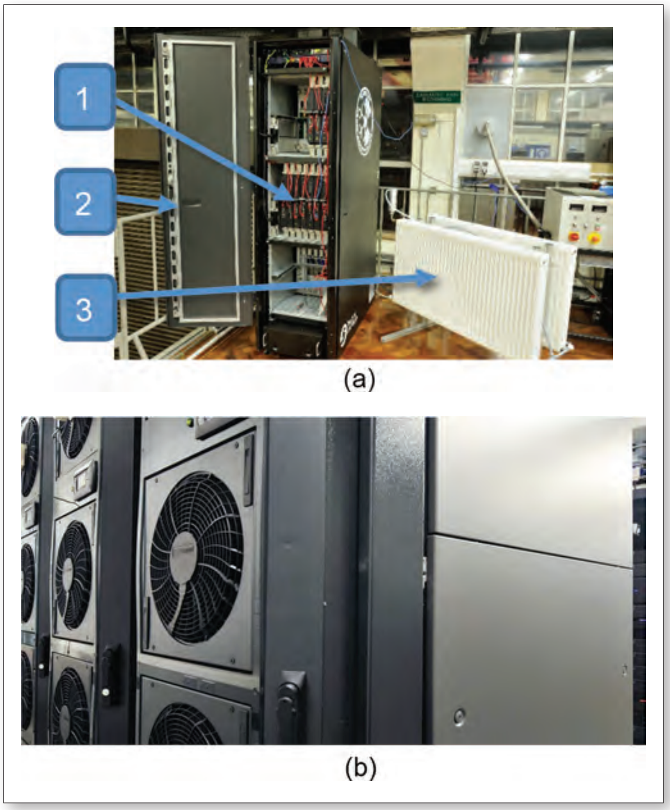

Figure 1 - (a) Rack with integrated TLC modules with (1) TLC modules, (2) rack, and (3) radiators [7]; (b) part of the data center HPC system showing the fan-assisted rear door liquid-loop heat exchangers.

DETERMINE RACK LEVEL TEMPERATURES AND DATA CENTER RESIDUAL HEAT.

Constructing a data center design based on TLC requires experimental data on the supply and return water temperatures to the racks and the residual heat emitted to the data center. In [9] calculations for the small TLC system in Figure 1(a) demonstrated that for a consumption of 2128W, 1367W of thermal power would enter the surroundings from the rack at a supply and return water temperature of 32°C and 38°C respectively. Baying racks together would reduce the exposed rack surface area so the thermal losses are expected to be smaller than the value obtained from the experimental system, but for a fully loaded 48 server rack it is assumed that

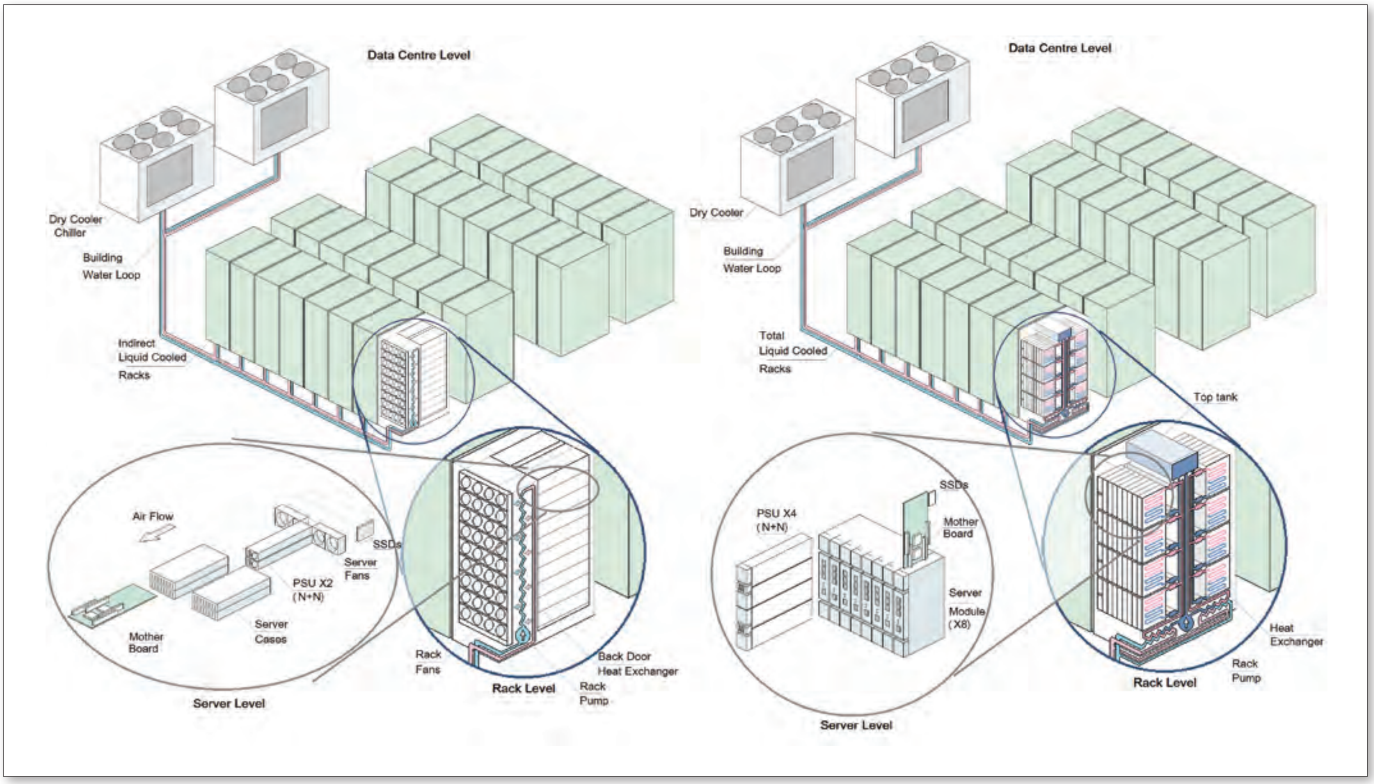

Figure 2 - Schematics showing ILC (left) and TLC (right) system levels of the data center.

	Type	Load	Number	ILC	TLC
Server/ IT Level	Motherboard	305W*	672	**205kW**	**205kW**
	Cooling Fans	23.4W	672	15.7kW	N/A
	Storage	0.85W	672	0.57kW	0.57kW
	PSU Fan	15.8W	168	2.65kW	N/A
	PSU Loss	7% Loss	168	15.7kW	14.4kW
*Under maximum performance and maximum load			Performance	223TFLOPS	223TFLOPS
			Total Load	**240kW**	**220kW**
			GFLOPS/W	930	1014
Rack Level	Cabinet Fan	161W	14	2.25kW	N/A
	Pump	45W	14	N/A	0.63kW
	Telco	33.4W	28	0.94kW	0.94kW
	PDU	3.5% Loss	28	8.38kW	7.7kW
			Component	11.6kW	9.3kW
			Total Load	**252kW**	**229kW**
Data Center Level	UPS	4% Loss	2	10.0kW	9.2kW
	Chiller	EER=3.03 Chiller On	2	82.9kW	N/A
		EER=19.3 Chiller Off	2	N/A	11.9kW
	Ventilation	880W	10	8.8kW	8.8kW
			Component	102kW	29.9kW
			Total Load	**354kW**	**259kW**
			pPUE	1.48	1.18
			MFLOPS/W	630	861
			Total Cooling	**112kW**	**21.3kW**
			Power Saving		**95kW**

Table 1 - Overall power data of an ILC based data centre compared to one with TLC server technology.

less than 10% of thermal energy will be lost to the data center, which is based on the fact that the system power consumption will increase but the surface area for heat loss remains constant.

For the ILC system, temperatures and residual heat can be obtained from a single air-cooled computer node. An experiment was set up to monitor the inlet and outlet temperatures, together with the air flow rate to obtain typical temperature increases over the inlet temperature for the main server components. To capture a realistic outflow temperature, the outlet air was funnelled into an insulated duct with an exhaust diameter of 75mm and a length of 400mm. This was found sufficient to characterize the outlet temperature from a single measurement. The mean velocity was calculated by taking a series of measurements across the pipe diameter, and thus calculating the mean airflow.

	ILC Based System			TLC Based System		
	medium	Temperature		medium	Temperature	
		In	Out		In	Out
Ambient	Air		25°C	Air		25°C
Chiller	R407c/water	25°C	20°C	Water	38°C	32°C
Building Water	Water	20°C	22°C	Water	32°C	38°C
Ventilation	Air / Water		24°C	N / A		
Rack	Air	24°C	36°C	Water	33°C	39°C
Computer Node	Air		36°C	Water	33°C	39°C
				Dielectric Coolant		53°C
CPU	Air		70°C	Dielectric Coolant		70°C
Max Delta Temperature	CPU to Chiller			CPU to Ambient		
	50°C			45°C		

Table 2 - Measured temperatures of the components for the comparison of ILC and TLC based data centers.

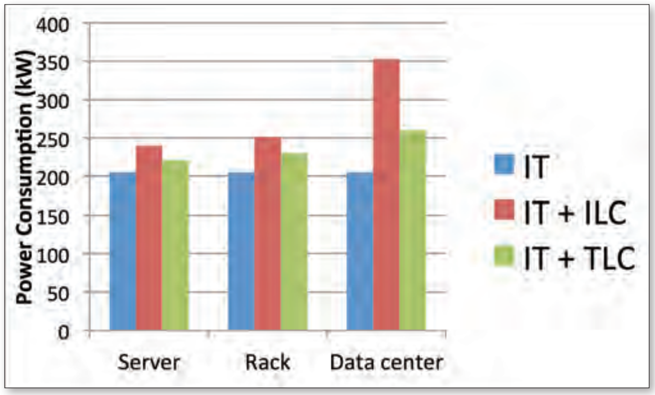

Figure 3 - Histogram of power consumed by ILC and TLC at the various levels within the data center.

The server level tests made use of two software tools to produce a near 100% load on the server, namely StressLinux [10] and High-Performance Linpack [11]. This resulted in a total energy load reading of 305W, which does not include the server fans as these are powered externally.

Temperature and velocity readings were taken at a range of points over the diameter of the exhaust pipe and all readings became repeatedly stable after 2 hours of continuous running. After several runs the temperature delta from the inlet to the outlet was 12.54°C. Under the measured exhaust conditions of 37°C, the air density, ρ, is approximately 1.138kg/m^3 and the mass flow rate, $\dot{m}$, of the outlet can be calculated as,

$$\dot{m} = \rho \times A \times \bar{v}$$

where A is the surface area of the ducted outlet, 0.004418m^2, and $\bar{v}$ is the averaged air velocity of the outlet, 4.45 m/s, which results in the mass flow rate is $\dot{m}$=0.02237kg/s.

The average specific heat capacity (SHC) of air between inlet at 25°C and outlet at 37°C is approximately 1005 J/(kg·°C), and since the temperature difference from the inlet to outlet, ΔT, was measured to be 12.5°C, the total heat flux Q that is released into the air would be,

$$Q = \dot{m} \times SHC \times \Delta T$$

yielding a value of Q of 282W, which is close to the system power of 305W, where the difference can be accounted for the heat losses through the server chassis.

In [2], the fan assistant RDHx units were found to be able to remove almost 94% of the heat, but when combined with the thermal losses in the server this would drop to 87% of the heat being removed. Therefore, as with the TLC solutions based on liquid immersed servers, there will be the requirement for CRAH units to handle a maximum of 10% of both TLC and ILC system thermal loads.

OPERATIONS OF THE DESIGN DATA CENTER

The design data center using the two rack level cooling configurations are shown in the Figure 2. The full sets of calculated results are detailed in Table 1 for a target of 250kW of total power consumption for both ILC and TLC based solutions.

From Table 1 it can be seen that the TLC based solution is able to achieve a partial Power Usage Effectiveness (pPUE) of approximately 1.18 at full load, which means that 0.18kW of power is required to cool 1kW of operating IT. This is compared to an equivalent ILC based system, which has a pPUE of 1.48. However this is based on the assumption that the computer systems are fully loaded and that the outside ambient air temperature was 25°C, as shown in Table 2, then for the ILC systems require mechanical cooling, which is a worst case cooling scenario for ILC.

Chillers usually have an Energy Efficiency Ratio (EER) of around 3.0, indicating that 3.0kW of cooling can be provided by 1kW of power, under fully loaded conditions, which results in the pPUE of the data center that is at least 1.33 when the chiller is operating at 100% mechanical load.

It is also possible to directly compare the data center level Performance Per Watt (PPW) of these 2 hypothetical systems since they contain identical computing hardware. From Table 1, the ILC system is rated to 632 MFLOPS/W, while the TLC system is rated at 858 MFLOPS/W.

From Table 2 it can also be seen that the ILC solution has a supply and return temperature that is lower than ambient whereas the TLC solution has supply and return temperatures greater than ambient, which is the reason why the chillers would be operating in mechanical cooling mode to deal with the additional cooling requirements.

The key results of Table 1 can be seen quite easily in Figure 3, which demonstrates the larger cooling overhead for ILC compared to TLC at the different levels, which are depicted in Figure 2.

CONCLUSION

A data center based on ILC systems is calculated to operate with a partial PUE of 1.48. This is compared with a TLC based data centre, which is calculated to operate at a partial PUE of 1.18 – 63% more effective at heat removal. Comparing performance yields a value for the ILC system of 630MFLOPS/W, which can be compared to the TLC based system of 861MFLOPS/W – 37% better performance. Finally the data centre based on TLC systems, under 25°C ambient air conditions, consumes 95kW less power, which saves 27% of the total power and more than 81% of cooling power over the ILC based system.

ACKNOWLEDGMENTS

The authors are grateful for the financial support from Iceotope Ltd and the University of Leeds' Digital Innovation Research Hub.

REFERENCES

[1] Iyengar, Madhusudan, *et al.* "Server liquid cooling with chiller-less data center design to enable significant energy savings." Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2012 28th Annual IEEE. IEEE, 2012.

[2] Almoli, Ali, *et al.* "Computational fluid dynamic investigation of liquid rack cooling in data centres." Applied Energy 89.1 pp150-155, 2012.

[3] 451 Research. "Liquid-Cooled IT: A flood of disruption for the datacentre", July 2014.

[4] Attlesey, Chad Daniel. "Case and rack system for liquid submersion cooling of electronic devices connected in an array." U.S. Patent No. 7,905,106. 15 Mar. 2011.

[5] Tuma, Phillip E. "The merits of open bath immersion cooling of datacom equipment." Semiconductor Thermal Measurement and Management Symposium, 2010. SEMI-THERM 2010. 26th Annual IEEE. IEEE, 2010.

[6] Hopton, Peter, and Jonathan Summers. "Enclosed Liquid Natural Convection as a Means of Transferring Heat from Microelectronics to Cold Plates." Semiconductor Thermal Measurement and Management Symposium, 2013. SEMI-THERM 2013. 29th Annual IEEE. IEEE, 2013.

[7] http://www.leeds.ac.uk/news/article/3374/wet_computer_server_could_cut_internet_waste

[8] Airedale, University of Leeds Case Study , High density IT cooling, Airedale, 2012.

Available from http://www.airedale.com/web/file?uuid=0fd59a99-917c-4363-9926-10a314de648c&owner=de3aa83e-cfaf-4872-8849-d5a40ddcd334&contentid=2027&name=University+of+Leeds+Case+Study.pdf.

[9] Yong Quiang Chi, *et al.* "Case study if a data centre using enclosed, immersed, direct liquid-cooled servers." Semiconductor Thermal Measurement and Management Symposium, 2014. SEMI-THERM 2014. 30th Annual IEEE. IEEE, 2014.

[10] http://www.stresslinux.org/sl/

[11] http://www.netlib.org/benchmark/hpl/

Application of Low Melt Alloys as Compliant Thermal Interface Materials: A Study of Performance and Degradation Under Thermal Duress

Chandan K. Roy[1], Sushil Bhavnani[1], Michael C. Hamilton[2], R. Wayne Johnson[3], Roy W. Knight[1], & Daniel K. Harris[1]

[1]Department of Mechanical Engineering, Auburn University, Auburn, AL

[2]Department of Electrical and Computer Engineering, Auburn University, Auburn, AL

[3]Department of Electrical and Computer Engineering, Tennessee Tech University, Cookeville, TN

INTRODUCTION

AS THE POWER DENSITY of microelectronics steadily increases, thermal management continues to play a key role in the continuous march towards miniaturization, performance improvements, and higher reliability demands [1]. One crucial thermal management consideration in electronics packaging involves reducing the thermal interfacial resistance between devices and their adjoining heat sink or substrate through improvements in thermal interface materials (TIMs). In a high power package, the thermal resistance of the TIM can account for as much as 50% of the overall thermal resistance [2]. Any desirable TIM candidate material would need to have a high thermal conductivity, low thermal resistance at a thin bond line thickness (BLT), excellent wetting properties, and all at a low cost while also being environmental and health friendly [3]. In addition to thermal considerations, mechanical considerations of TIMs are equally important to the performance and health of high power electronics. TIMs can be either compliant (e.g. greases and gels) or non-compliant (e.g. solders and adhesives), where the former bears no interstitial mechanical stress due to thermally induced strains. Non-compliant TIMs must be able to withstand the mechanical stresses from coefficient of thermal expansion (CTE) mismatches between the adjoining materials (e.g. silicon - copper). If the CTE strain overwhelms the mechanical properties of the TIM, the joint will ultimately fail. Therefore, high performing compliant TIMs are a desirable design option for better thermal performance and improved reliability.

Commonly used commercially available TIMs include greases, phase change materials (PCM), gels, and pads, which are polymer based materials loaded with conductive particles (metal or ceramic) to enhance the thermal conductivity. Greases are the most widely used TIMs. They offer a moderate level of thermal performance [3-4]. Carbon-based materials such as carbon nanotubes [5] and graphene [6] have been investigated by many researchers for use in TIMs. The thermal performance of different commercial and emerging TIMs based on previous research works is presented in Figure 1.

Recently, several researchers [7-9] have encouraged using liquid metal alloys (LMAs) as compliant, high performing TIMs. LMAs possess the requisite high thermal conductivity (an order of magnitude higher compared to the traditional TIMs) while also

Chandan K. Roy is currently pursuing Ph.D. in Mechanical Engineering at Auburn University, Auburn, AL. He has been working as a research assistant at the Auburn Thermal Management Laboratory, AUTherMML since fall 2011. Mr. Roy received Master of Science and Bachelor of Science degree in Mechanical Engineering from The University of Texas at El Paso (UTEP) and Bangladesh University of Engineering and Technology (BUET) in 2011 and 2007 respectively. His research interests include thermal interface materials, thermal management of electronics, heat transfer, thin film analysis and thermal barrier coatings.

Daniel K. Harris has been on faculty at Auburn University's Department of Mechanical Engineering since 1997. He holds a M.S. in Applied Mathematics from John Hopkins Applied Physics Laboratory and a Ph.D. in Mechanical Engineering from Purdue University, and has over 48 published journal articles and conference proceedings, as well as several book chapters on Lunar Engineering. He is the director of Auburn's Thermal Management Lab, AUTherMML, housed in the Enhanced Heat Transfer Laboratory. Dr. Harris' research has included several technology developments including post fabrication bendable copper-water heat pipes and micro-heat pipes using liquid metal working fluids. His current research efforts include the use of nano-particles in thermal composites for use as interfacial thermal materials. In addition, Dr. Harris has 10 years of industry experience in the thermal design and analysis of satellites and airborne avionics.

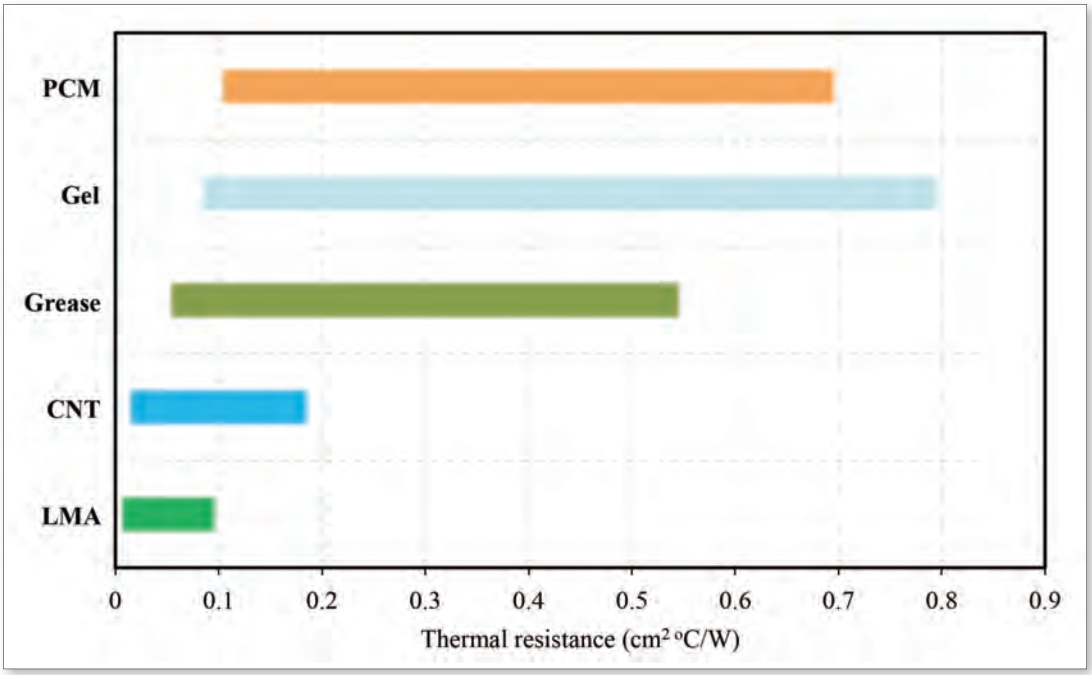

Figure 1 - Thermal performance comparison of variety of commercial TIMs with emerging TIMs such as CNTs and LMAs [4,5,7,10,12,13].

offering adequately low (<0.1 cm^{2}°C/W, Figure 1) thermal resistances at small contact pressures. LMAs can be used either as a thin foil or in combination with a substrate [8-10] or as a filler material in composites [11]. Alloys of indium, bismuth, gallium, and tin are preferable to use as TIMs [8-10]. Mercury, lead and cadmium based alloys are usually avoided due to their toxicity and environmental issues. Hamdan *et al.* [7] reported the thermal resistance of liquid mercury micro droplets as low as 0.0025 cm^{2}°C/W. However, mercury should be avoided due to its toxicity. Webb and Gwinn [8] observed that cycling 20°C above the melting point of the alloy (51In/32.5Bi/16.5Sn) resulted in a significant increase in the thermal resistance. However, Hill and Strader [9] did not find any notable performance degradation after cycling 90°C above the melting point of the same alloy (51In/32.5Bi/16.5Sn). Thus, a good understanding of the reliability of LMAs as TIMs is needed for today's designers considering their use.

Although LMAs offer a low thermal resistance, there are several concerns such as oxidation/corrosion, intermetallic

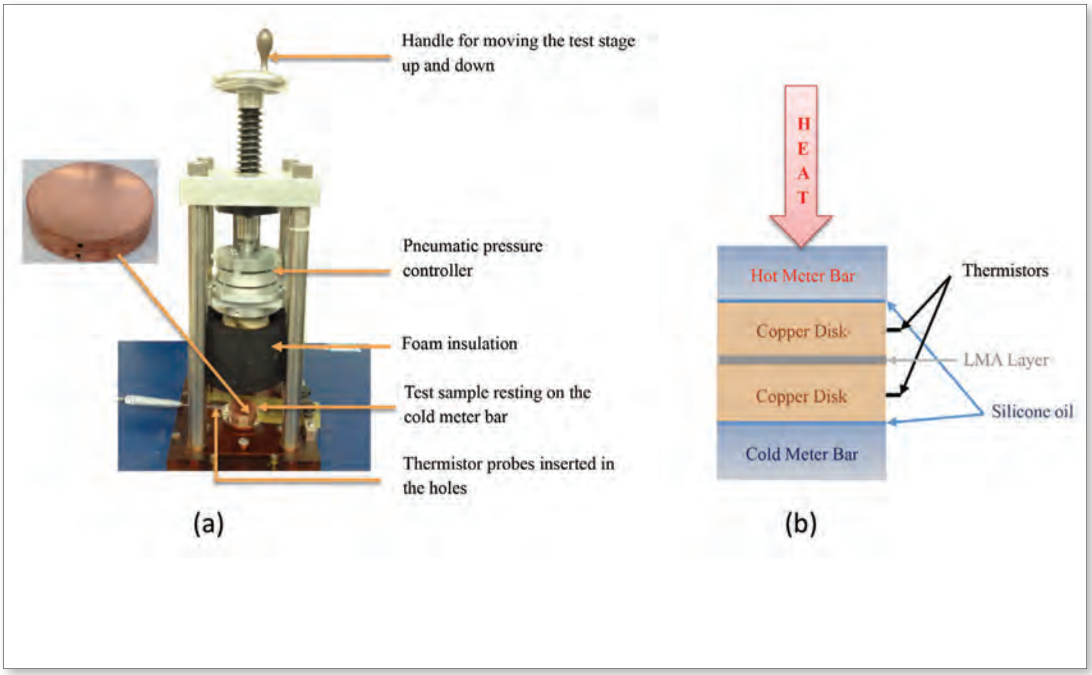

Figure 2 - (a) Testing of LMA TIM using modified test rig; the cooper disks assembly with the TIM at the interface was placed between the TIM tester surfaces (b) Schematic diagram of the testing methodology.

growth, dewetting and migration. Several investigators have offered different approaches to mitigating these concerns [8-10]. In this work, three alloys were chosen to test their performance as TIMs. The properties of these alloys are presented in Table 1. These alloys were selected for study because they have a wide range of MPs (from 16°C to 60°C) spanning the range of interests for most application in today's markets.

Alloy	Composition (% by mass)	Melting Point (°C)	Density (gm/cm³)
1	75.5 Ga/ 24.5 In	16	6.35
2	100.0 Ga	30	5.90
3	51 In/ 32.5 Bi/ 16.5 Sn	60	7.88

Table 1 - Properties of the LMAs with melting temperature ranges from 16°C to 60°C.

EXPERIMENTAL

Description of the Apparatus

The thermal performances of the LMAs reported here were generated using an ASTM D5470 standard TIM tester. The detailed specifications of the apparatus used can be found in reference [14]. To avoid any contamination of the TIM tester surfaces and to improve the accuracy of the test results, the LMAs were tested by placing them between copper disks (alloy 110) of thickness 3.2 mm and 33 mm in diameter. The disks are both sides polished with a surface flatness within 7-8 microns. This high level of surface finish was needed for accurate thermal performance measurement but it is not required in actual application. However, it should be noted that the quality of the mating surfaces directly affects interfacial thermal performance. The resulting disk assembly with the alloy at the interface was then placed under the tester, Figure 2. Silicone oil (Xiameter PMX-200, viscosity: 1000CS) was applied on the top and bottom surfaces of the adjoining copper disks to ensure reproducible contact between the test surfaces of the TIM tester and the copper disks. The temperature differential (ΔT) across the LMA bond line was measured using two high precision thermistor probes (1 mm diameter, accuracy 0.05°C) inserted in the middle of each copper disk. The hole was filled with thermal grease (laird tech. Tgrease 880) to ensure reproducibility of the measurements.

Sample Preparation and Diffusion Barriers

The oxide layers on the copper disk surfaces were first chemically removed using methanol, acetone, and a 5% hydrochloric acid solution prior to application of the LMAs. This method of surface preparation was sufficient for testing purpose, but not necessary for normal applications. It should be followed that the contact surfaces remain free from particulates or any other contaminants before applying LMAs. To enhance the wetting, the disk surfaces were mechanically scrubbed with the alloy using cotton swabs. In the case of Ga and In-Bi-Sn alloy, both the alloy and the disk were heated above the melting point of the alloy after which the molten metal was applied onto the heated disk's surface using a brisk mechanical rubbing application technique.

Since LMAs are known to react with and form intermetallics with Cu [8-10], surface treatments were used in order to retard the diffusion of the alloy components into the disks. Hence, a thin metallic barrier layer was applied on the copper disk surfaces by sputtering. As tungsten (W) provides superior protection against gallium (Ga) and Ga-based alloys, the copper disks were coated with 2 µm W. For better adhesion of the W, a 50 nm layer of titanium (Ti) was first applied.

RESULTS

In-Situ Thermal Resistance

The thermal resistances of six different substrate-alloy combinations is presented in Figure 3. Three samples of each combination were tested at 138 kPa. The calculated experimental uncertainties [15] are represented by the error bars. Each resistance value presented here is the average of three repeated measurements. The measured lowest resistance was as low as 0.005 cm^2°C/W with W/In-Bi-Sn/W (In-Bi-Sn alloy between W coated surfaces) and as high as 0.065 cm^2°C/W with Cu/In-Bi-Sn/Cu (alloy 3 between bare Cu surfaces). The variation in thermal resistance for the similar substrate-alloy combination results from the unique nature of each sample. Since LMAs are highly conductive and the joints are relatively thin (SEM cross-sectional analysis reveals that the BLT of Cu/Ga-In/Cu joint is about 37 µm), the contact resistances (mostly due to surface irregularities such as surface roughness and flatness) dominate the interfacial resistance. Even a small change in surface properties would result in observable changes in the overall thermal resistance. It was assumed that all the disks had the same degree of surface roughness and flatness, but this is not valid in reality. It should be noted that it was not possible to maintain the exact amount of LMAs at the interface for each pair of disks during testing. Another source of variation might appear from the mechanical rubbing (wetting) of LMAs onto the disk surfaces. During the wetting process, it was found that some samples were more easily wetted while others required hard scrubbing to induce wetting. If the LMAs do not wet the mating surfaces properly, small air pockets might be present at the interface, which in turn increases the thermal resistance. Other properties of the alloys such as viscosity, surface tension, and oxidation state during application may have caused some variation in the thermal resistance measurements. Considering all these factors, each disks pair is different and it is a challenge to reproduce the result even with similar substrate-alloy combination.

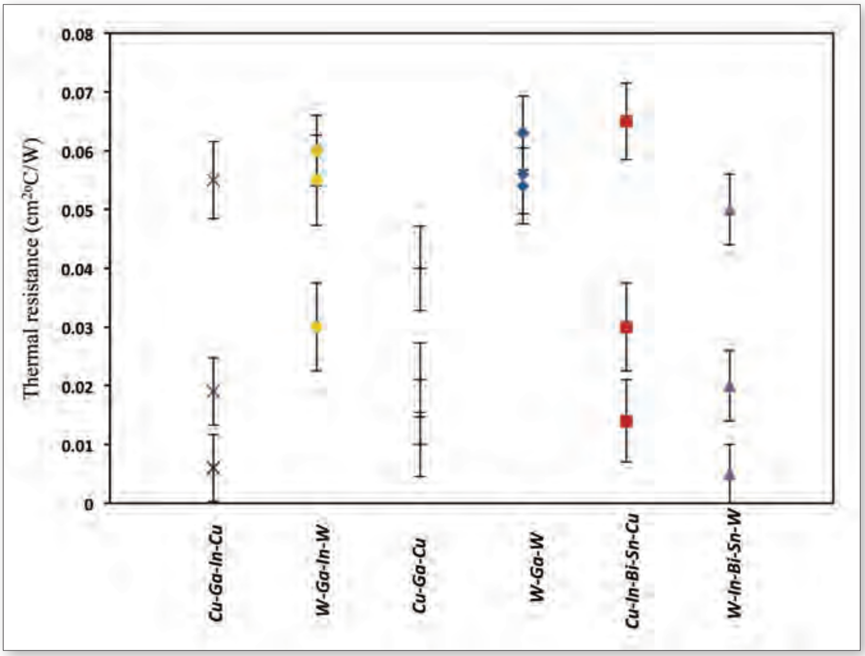

Figure 3 - In-situ thermal resistances of three alloys between Cu and W surfaces at 138 kPa.

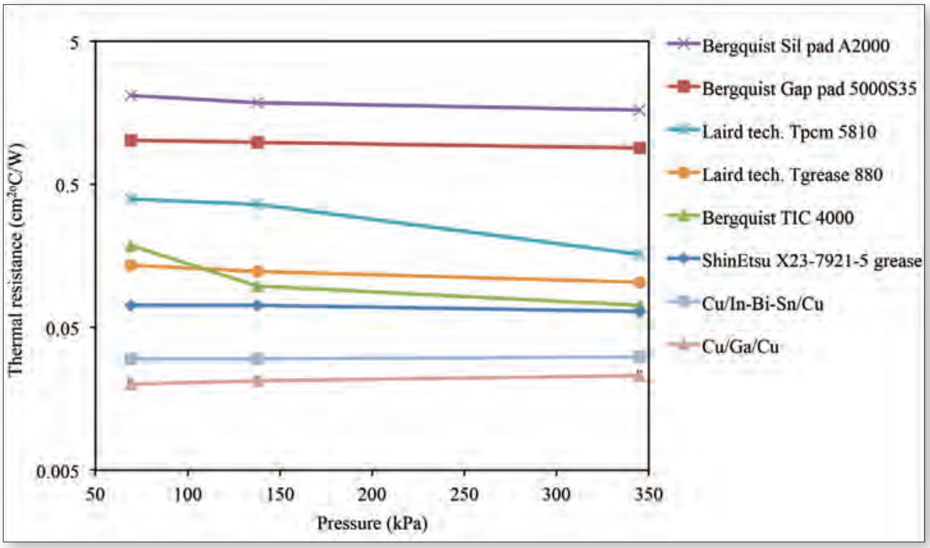

Figure 4 - Thermal resistances of variety of commercial TIMs and LMAs as a function of applied pressure.

The thermal performance of LMAs and some commonly used commercial TIMs are tested at different pressures using the same apparatus and similar methodology and measurements are compared (Figure 4). The results indicate that the thermal resistances of LMAs are independent of applied pressures in the range 69-345 kPa. This study concludes that LMAs offer excellent thermal bond even at a small pressure. None of the commercial TIMs were found to have a thermal joint resistance as low as the resistances of LMAs. Current high conductivity greases can offer resistances at the high end of what LMAs offer. The thermal resistances of the commercial TIMs tested were found to be as low as 0.065 $cm^{2o}C/W$ and as high as 2 $cm^{2o}C/W$ depending on the type of material being tested.

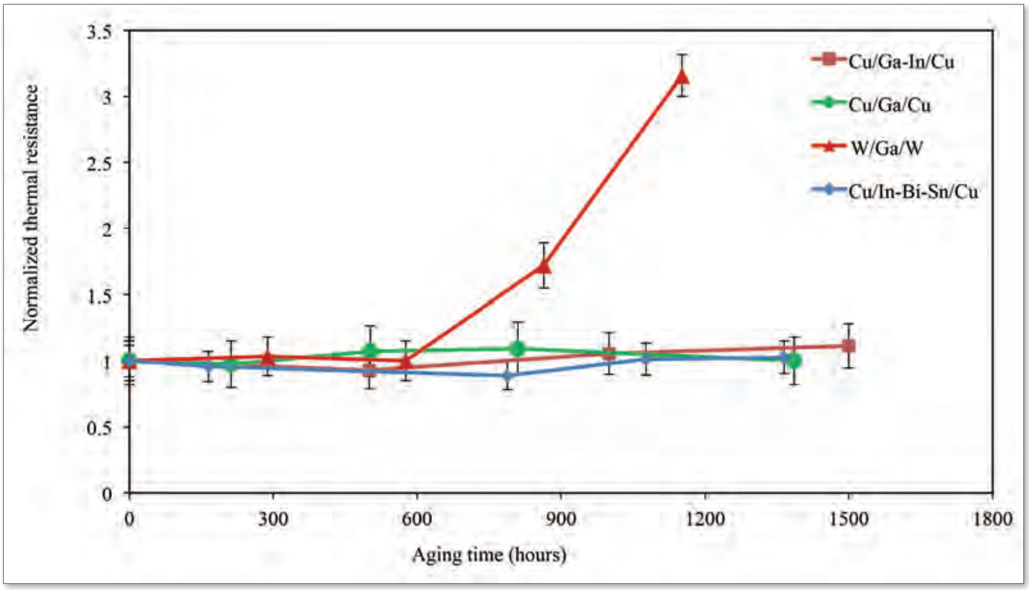

Figure 5 - .Isothermal aging (at 130^{o}C) of three alloys as a function of aging time.

Isothermal Aging

Accelerated aging was carried out by exposing the samples at an elevated temperature of 130^{o}C (followed DARPA's Thermal Management Technologies (TMT) program's reliability profile [16]) in an atmospheric furnace for extended periods of time without any additional pressure and/or constrain during aging. The thermal aging test results of Ga between different substrate surfaces (Cu and W) and Ga-In and In-Bi-Sn alloys between Cu surfaces are presented in Figure 5. The thermal resistances are normalized by the initial resistance values. It was found that the thermal resistance of all three alloys between bare Cu remained unchanged after long periods of exposure (1500 hours for Ga-In alloy), which indicated a stable thermal joint. Ga between W surfaces also showed a negligible change in thermal resistance up to 576 hours of aging. However, the resistance started to increase rapidly thereafter. The thermal resistance increased by about 220% after 1152 hours of aging. The superior aging performance of all three alloys between bare Cu surfaces is believed to be due to the enhanced wetting of the alloys with Cu and the diffusion of the alloy into the Cu substrate. Since the alloys contain Ga and In, they will most likely diffuse into Cu and form intermetallics [9,10]. Our results suggested that any diffusion and formation of these intermetallics that may have occurred did not negatively impact the thermal performance over time. Since the thin layer of W acted as a diffusion barrier, the alloy could not interact with the underlying Cu substrate, thereby, causing the thermal resistance to increase significantly upon aging.

Thermal Cycling

The purpose of these tests is to cycle each sample above and below the melting point of the interfacial alloy. The Cu disk assembly with alloys at the interface was placed in a thermal cycling chamber to cycle from -40^{o}C to 80^{o}C [16] without any additional pressure and/or constrain. The heating and cooling ramp rates were 3^{o}C per minute. The samples were soaked for 20 minutes at the two extreme temperature plateaus, ensuring that the interface reached the chamber temperature. The time needed to complete a thermal cycle was two hours. Figure 6 shows the thermal cycling results of all three alloys between bare Cu and W surfaces. The results showed that the thermal resistance of all three alloys between bare Cu did not change significantly even after 600 cycles. However, alloys placed between W coated surfaces, the resistance increased significantly just after 200 cycles. It can be noticed from the aging studies (Figure 5) that the thermal resistance of Ga between W surfaces started to increase after 576 hours. However, the resistance increased significantly just after 200 cycles (400 hours). This is because, during aging tests, the samples were always kept at a constant temperature, whereas during cycling, the samples experienced thermal expansion/contraction resulting from the heating and cooling of the joint. All these results suggest that the bare Cu-alloy surface combination makes a more reliable LMA thermal joint.

SUMMARY AND CONCLUSIONS

The application of LMAs as efficient TIMs has been investigated herein. Three alloys (Ga-In, Ga, In-Bi-Sn) with different melting points and various compositions were chosen to test the thermal performance. The thermal resistances of the alloys were tested by placing them between bare Cu and W coated Cu substrates using the ASTM D5470 standard methodology. The in-situ thermal performance shows that LMAs can offer extremely lower thermal resistances at small contact pressures compared to many commercial TIMs. Accelerated tests include isothermal aging at 130^{o}C and thermal cycling from -40^{o}C to 80^{o}C. It was observed that all the alloys between bare Cu surfaces survived both isothermal aging and thermal cycling. However, alloys applied between W coated surfaces failed to withstand both aging as well as cycling. The results indicate that any Cu-alloy interactions (which is primarily the diffusion and interfacial reaction) produce a more reliable thermal joint.

In actual processor-heat sink application, the LMAs are placed between two dissimilar materials, typically silicon (die) and copper (heat sink). In this work, LMAs are found to be effective when applied between bare Cu substrates. LMAs are thin liquids in molten form and they can accommodate any thermo-mechanical stress that results from the CTE mismatch of different materials. Therefore, it is presumed that they will continue to perform reliably between dissimilar material joints. LMAs are environmental and health friendly unless they contain any mercury, lead or cadmium compounds. There are a host of issues that would have to be addressed before LMAs are ready for volume production and reliable operation in the field. However, the level of thermal performance demonstrated here could be a powerful incentive for the industry to invest in the required development effort to enable their use in the most demanding applications.

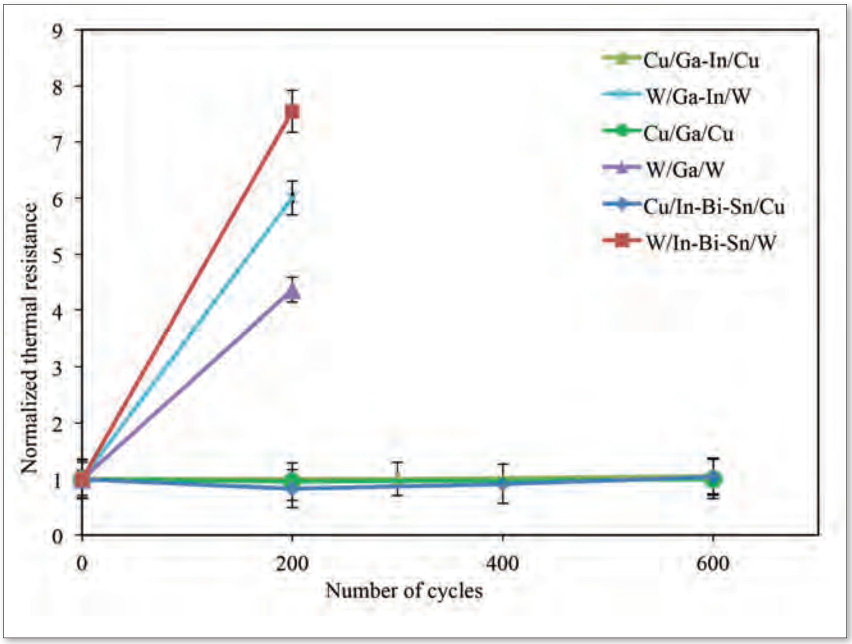

Figure 6 - Thermal cycling (-40°C to 80°C, 3°C/min, dwell: 20 minutes, 2 hours/cycle) of three alloys between Cu and W surfaces.

REFERENCES

[1] Bloschock Kristen P., Bar-Cohen, Avram, "Advanced Thermal Management Technologies for Defense Electronics, " Defense Transformation and Net-Centric Systems 2012, SPIE Defense, Security, and Sensing, 23-27 April 2012 in Baltimore, Maryland.

[2] Guenin, Bruce. "Use of the Monte Carlo method in packaging thermal calculations." Electronics Cooling Magazine, December 2014 issue.

[3] Gwinn, Joshua P., and R. L. Webb. "Performance and testing of thermal interface materials." Microelectronics Journal 34.3 (2003): 215-222.

[4] Prasher, Ravi. "Thermal interface materials: historical perspective, status, and future directions." Proceedings of the IEEE 94.8 (2006): 1571-1586.

[5] Cola, Baratunde A. "Carbon nanotubes as high performance thermal interface materials." Electron. Cooling Mag 16.1 (2010): 10-15.

[6] Shahil, Khan MF, and Alexander A. Balandin. "Graphene–multilayer graphene nanocomposites as highly efficient thermal interface materials." Nano letters12.2 (2012): 861-867.

[7] Hamdan, A., *et al.* "Characterization of a liquid–metal microdroplet thermal interface material." Experimental Thermal and Fluid Science 35.7 (2011): 1250-1254.

[8] Webb, R.L., Gwinn, J.P. "Low melting point thermal interface material." ITHERM. 2002. 671-676.

[9] Hill, R.F., and Strader, J.L. "Practical utilization of low melting alloy thermal interface materials." Semiconductor Thermal Measurement and Management Symposium, 2006 IEEE Twenty-Second Annual IEEE. Dallas, TX, 2006. 23-27. Macris, C.G. *et al.* "Performance, Reliability, and Approaches Using a Low Melt Alloy as a Thermal Interface Material." IMAPS, 37th Int. symposium on Microelectronics. Long Beach, Ca, 2004.

[10] Macris, C.G. *et al.* "Performance, Reliability, and Approaches Using a Low Melt Alloy as a Thermal Interface Material." IMAPS, 37th Int. symposium on Microelectronics. Long Beach, Ca, 2004.

[11] Carlberg, B. *et al.* "Nanostructured polymer-metal composite for thermal interface material application." 58th Electronics Components and Technology Conference (ECTC). 2008. 191-197.

[12] Wasniewski, Joseph R., *et al.* "Characterization of metallically bonded carbon nanotube-based thermal interface materials using a high accuracy 1D steady-state technique." Journal of Electronic Packaging 134.2 (2012): 020901.

[13] Blazej, D. "Thermal Interface Materials." Electronics Cooling. 9.4 (2003): 14-20.

[14] Analysis Tech. Material Thermal Testers: TIM Tester 1400 Specifications. 2012. 18 April 2014.

[15] Roy, Chandan K., *et al.* "Investigation into the application of low melting temperature alloys as wet thermal interface materials." International Journal of Heat and Mass Transfer 85 (2015): 996-1002.

[16] DeVoto, Douglas, *et al.* "Thermal performance and reliability characterization of bonded interface materials (BIMs)." Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2014 IEEE Intersociety Conference on. IEEE, 2014.

Comparison of HPC/Telecom Data Center Cooling Methods by Operating and Capital Expense

Dr. Alexander Yatskov
Thermal Form & Function Inc.

INTRODUCTION

CURRENT HIGH-PERFORMANCE computing (HPC) and Telecom trends have shown that the number of transistors per chip has continued to grow in recent years, and data center cabinets have already surpassed 30 kW per cabinet (or 40.4 kW/m2) [1]. It is not an unreasonable assumption to expect that, in accordance with Moore's Law [2], power could double within the next few years. However, while the capability of CPUs has steadily increased, the technology related to data center cooling systems has stagnated, and the average power per square meter in data centers has not been able to keep up with CPU advances because of cooling limitations. With cooling systems representing up to ~50% of the total electric power bill for data centers [3], growing power requirements for HPC and Telecom systems present a growing operating expense (OpEx). Brick and mortar and (especially) mobile, container based data centers cannot be physically expanded to compensate for the limitations of conventional air cooling methods.

In the near future, in order for data centers continue increasing in power density, alternative cooling methods, namely liquid cooling, must be implemented at the data center level in place of standard air cooling. Although microprocessor-level liquid cooling has seen recent innovation, cooling at the blade, cabinet, and datacenter level has emerged as a critical technical, economic, and environmental issue.

In this article, three cooling solutions summarized in Table 1 are assessed to provide cooling to a hypothetical, near-future computing cluster. Cooling Option 1 is an air-cooled system with large, high-efficiency, turbine-blade fans pushing air through finned, heat-pipe equipped, copper heat sinks on the blade. Rear-door air-to-liquid heat exchangers on each cabinet cool the exiting air back to room temperature so that no additional air conditioning strain is placed on facility air handlers. The water & propylene glycol (water/PG) mixture from the heat exchangers is pumped to the roof of the facility, where the heat absorbed from the CPUs is dissipated to the environment via a rooftop compressor-enabled chiller.

Cooling Option 2 uses water-based touch cooling on the CPUs via a copper coldplate loop on each board. These loops are connected to an in-rack manifold that feeds a water/PG mixture in and out of each blade. A coolant distribution unit (CDU) collects heated water via overhead manifolds from each cabinet, cools the water through an internal liquid-to-liquid brazed-plate heat exchanger, and pumps it back through the overhead manifolds to the cabinets for re-circulation. On the other side of the heat exchanger, a closed water loop runs from the CDU to a rooftop dry cooler, where the heat from the CPUs is ultimately dissipated into the atmosphere.

Cooling Option 3, which considers two approaches, removes water from the server cabinets and instead uses refrigerant (R134a) as the heat transfer fluid in the server room. This approach uses a coldplate and manifold system similar to the water-cooling approach,

Dr. Alexander Yatskov is Vice President of Engineering at Thermal Form & Function Inc., bringing extensive experience in high technology innovation and product development. He received his Doctor of Engineering degree from Moscow Technological University, and MSEM degree from Tufts University. Prior to joining Thermal Form & Function Inc., he worked for 10 years at Cray Inc., where he engineered cooling systems for Red Storm, XT-3, XT-4, Jaguar, Titan, and other supercomputer platforms. He authored sixteen articles and holds twenty five US patents in electronics cooling and packaging area.

SEMI-THERM®

Thermal Innovations for a Cool World
Semiconductor Thermal Measurement, Modeling and Management Symposium
T H I R T Y - S E C O N D   A N N U A L

CALL FOR PAPERS

March 13-17, 2016
Doubletree Hotel, San Jose, CA, USA

The conference

SEMI-THERM is an international forum dedicated to the thermal management and characterization of electronic components and systems. It provides knowledge covering all thermal length scales from IC to facility level. The symposium fosters the exchange of knowledge between thermal engineers, professionals and leading experts from industry as well as the exchange of information on the latest academic and industrial advances in electronics thermal management. Areas of interest include, but are not limited, to the following:

Topics:

- Component, Board- and System-Level Thermal Design Approaches
- Air mover technologies with low acoustics
- Thermal Integration in the Product Design Process
- Multi-Physics Based Reliability, including Accelerated Testing
- Multi-objective design and optimization, modeling and characterization
- Novel materials: Heat Spreaders, Thermal Vias and Thermal Interface Materials
- Nanotechnology: Thermal, Mechanics, Material and Process Related Issues In Nanostructures
- Micro-Fluidics
- Characterization and Standardization of Material Property Measurements
- Energy Harvesting materials, Thermo-electrics
- Novel and Advanced Cooling Techniques/Technology,
- Roadmaps, Specifications and Traditional Cooling Limits
- Mechanical Modeling, Simulations and Characterization
- Characterization and Modeling of Multi-Scale Heat Transfer Problems
- Computational Fluid Dynamics (CFD) Analysis and Validation
- Multi-physics Modeling and Characterization of Products and Processes
- Thermal Control Methodologies

Application areas:

- Processors, ICs, and Memory
- 3-D electronics
- Wireless, Network, Computing Systems
- Data Centers
- Portable and Consumer Electronics
- Power Electronics
- Harsh Environments
- Commercial, Defense, and Aerospace Systems
- Solid-State Lighting
- Solid State Energy Generation/Cooling
- Medical and Biomedical electronics
- Instrumentation and Controls
- Micro- and nano-scale devices
- MEMS and sensors
- Alternative and Renewable Energy
- Wearables and IoT

Conference highlights

- Dedicated sessions on various application areas, novel materials, thermo-mechanical modeling/characterization techniques
- Two days of short courses (March 13 and 14, 2016), embedded tutorials, technical workshops, "How To" courses
- Vendor exhibits: Cooling technologies, air movers, experimental characterization equipment and CFD simulation software

To submit an abstract	To submit a peer-reviewed paper
You are invited to submit an extended abstract describing the scope, contents, key results, findings and conclusions. This abstract of 2 to 5 pages is supported by figures, tables and references as appropriate. Abstracts must demonstrate high technical quality, originality, potential impacts. Upload your abstract electronically in RTF, DOC or PDF formats at www.semi-therm.org	You are invited to submit a full paper to be reviewed by the SEMI-THERM Technical Committee. Papers must include a one paragraph abstract and keywords and must demonstrate high technical quality, originality, potential impacts. Upload your paper electronically in RTF, DOC or PDF formats at www.semi-therm.org. Special consideration is given to papers based on academic/industry cooperation projects.
Abstract Submission Deadline: September 29, 2015	Paper Submission Deadline: May 31, 2015
Abstract Acceptance Notification: November 3, 2015	Paper Acceptance Notification: September 15, 2015
Photo-ready Full Manuscript Due: January 5, 2016	Photo-ready Full Manuscript Due: December 15, 2015

SEMI-THERM actively solicits student papers, students to co-chair sessions, and awards travel stipends and reduced conference fees.

For further information please contact the Program Chair:
Veerendra Mulay, Research Engineer, Facebook Inc., E-mail: vmulay@fb.com
Visit our website: http://www.semi-therm.org

but may or may not have a refrigerant distribution unit in the building. Cooling Option 3a, having pumps mounted in the cabinets, circulates the refrigerant through a water-cooled, brazed-plate heat exchanger in a refrigerant distribution unit (RDU) to condense the refrigerant after it absorbs the heat from the CPUs. In Cooling Option 3b, pumps move the refrigerant straight to the roof, where a rooftop condenser dissipates the heat to the environment (Table 1).

The goal of this article is to provide an "apples-to-apples" comparison of these three cooling systems by suggesting a hypothetical, high-power near-future data center specification for each method to cool. When the options are compared side-by side, differences in fluid dynamics and heat transfer will translate into differences in efficiency, and comparisons between various cooling methods become more easily visible. This article is not a complete guide to installing or selecting equipment for each of these cooling systems, but it is a general overview of what power usage advantages each system offers.

	Cooling Method		
	Option 1: Air Cooling	**Option 2: Water Cooling**	**Option 3: Pumped Refrigerant Cooling**
CPU Cooling	Copper-fin heat sinks with embedded heat pipes	Option 2a: Series-parallel copper cold plates Option 2b: Parallel copper cold plates	Series-parallel copper cold plates
Heat Exchanger Location	Air-to-liquid on rear door of cabinet	Liquid-to-liquid in Coolant Distribution Unit	Option 3a: Liquid-to-liquid in Refrigerant Distribution Unit Option 3b: Liquid-to-air in Rooftop Condenser
Pump Location	Rooftop Chiller	Coolant Distribution Unit	Cabinet Level Pumps
Rooftop Cooling	Compressor-Enabled Chiller	Water-to-Air Dry Cooler	Option 3a: Water-to-Air Dry Cooler Option 3b: R134a Vapor-to-Air Condenser

Note: External Ambient air temperature = 25 °C

Table 1 - Overview of cooling options

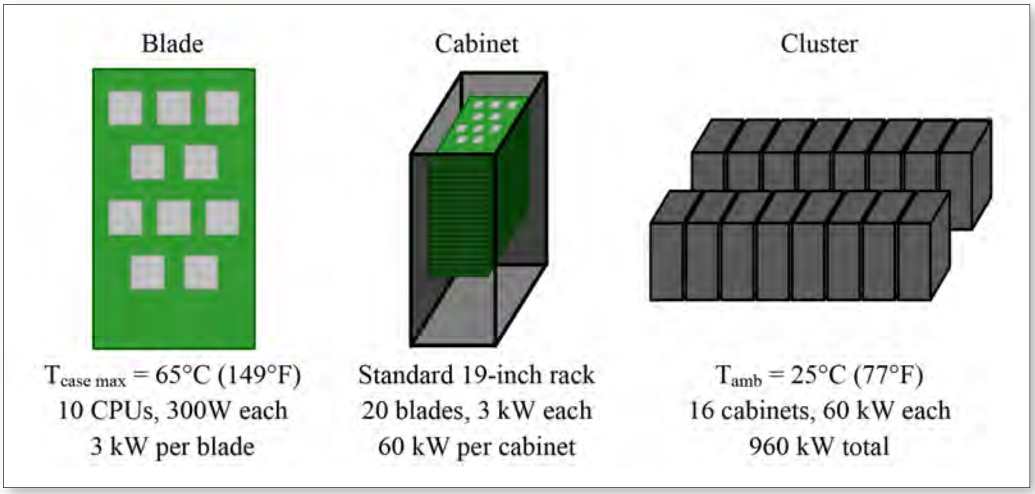

Figure 1 - Hypothetical data center module specifications

HYPOTHETICAL COMPUTER CLUSTER SPECIFICATION AND THERMAL ASSESSMENT

In order to bound the comparison in a meaningful way, a set of specifications was developed from extensive discussions with the Liquid Cooling (LC) forum of LinkedIn professional network [4], an association of motivated multidisciplinary professionals from HPC, Telecom and electronics cooling industries.

After several weeks of discussion, the LC forum agreed on the following system configuration and operation conditions for analysis: the hypothetical computer cluster under consideration should produce ~1MW of IT power. The distribution of power is shown in figure 1.

The cabinet architecture in this specification assumes horizontal card inserted from the front, with no cards inserted from the back of the cabinet. Alternative architectures exist, and can be cooled by a variety of means, but for this analysis, a simple, easily-relatable architecture was desired, so only front-facing, horizontal cards were considered.

A hypothetical data center could be equipped either with a dry cooler or compressor equipped chiller located on the roof of the building, at up to 60 m (≈200 feet) above computer HPC cluster level floor. To maintain the 65°C case temperature limit, the air-cooling method requires a compressor-equipped chiller, but a dry cooler is preferred for liquid cooling methods.

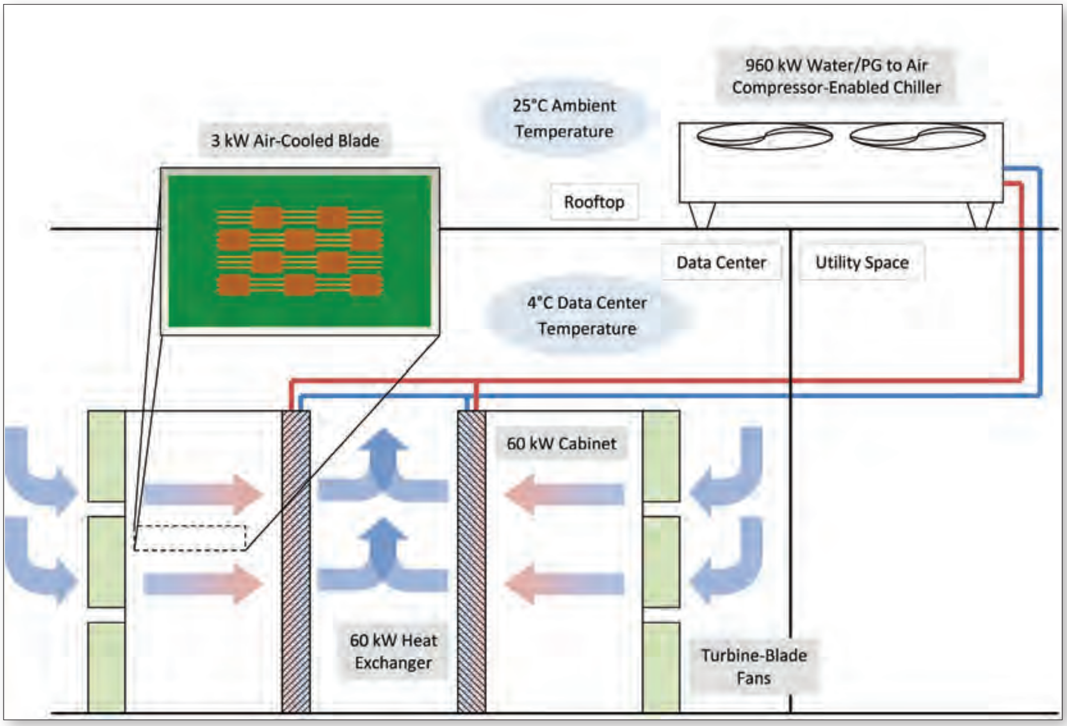

Figure 2 - Option 1: Air cooling with rear-door water/PG to air heat exchangers

One of the most important requirements for this analysis was its full transparency, so anybody could have a closer look, if desired. When the study results were presented to LC forum participants at a webinar in February, 2015, a number of forum participants requested and (later) participated in a "deeper dive", so a step-by-step process with a higher level of detail was provided on request. This "deeper dive" can be found on the Electronics Cooling website (**www.electronics-cooling.com**).

COOLING OPTIONS VIEWED VIA EXPLICIT COMPARISON

Option 1: Advanced Air Cooling

Cooling 60 kW in a single rack required a staggered, heatsinked 10-CPU layout, heat pipes, rear door cooling heat exchangers, and powerful, high-pressure turbine-blade fans. In addition, the air in the data center still needed to be significantly lower than room temperature (4°C) to achieve the desired case temperature. Even though air cooling may not be an economically feasible solution at this power density, and even though it would clearly not meet the NEBS GR-63 acoustic noise level standard [5], we still had to devise an air cooling option to compare to the LC options.

In this approach, each cabinet was supplied with a rear-door cold water/PG cooler with 3 "hurricane" turbine-blade fans (with operating point of ~3.7 m3/s at ~3.7 kPa pressure difference [6]). The cold water/PG solution circulates around the system to the roof chiller, where heat from the cabinets is dissipated to the ambient air. The system would be monitored with on-board temperature sensors and would either increase the fan speed or throttle CPU performance if the CPU approached the 65C case temperature threshold. This approach is illustrated in Figure 3.

Option 2: Water/PG Touch Cooling

In option 2a , (Fig. 3, top board layout) 10 CPUs per blade were arranged in two parallel groups of 5 serially connected cold plates (option 2a). 20 horizontal blades were plugged into vertical supply/return manifolds, and these cabinet manifolds were supplied with water/PG from the Coolant Distribution Unit (CDU) via overhead manifolds. Pump and water/PG-to-facility water heat exchanger was needed to reduce the pressure losses at cold plates. A separate loop brings the heat from the CDU to a rooftop dry cooler unit, where it is rejected into the ambient air. This method would include a control system in the CDU that would increase the water flow rate in the case of an increased load. Either passive or active flow regulators would also be placed at the inlet of each blade to ensure even flow distribution across the whole cabinet.

After the first pass of Water/PG system simulation it was discovered that water/PG speeds in the blade exceeded allowable ASHRAE [1] velocity limits, an additional LC option (option 2b) was added - where all 10 cold plates were connected in parallel. This required additional onboard

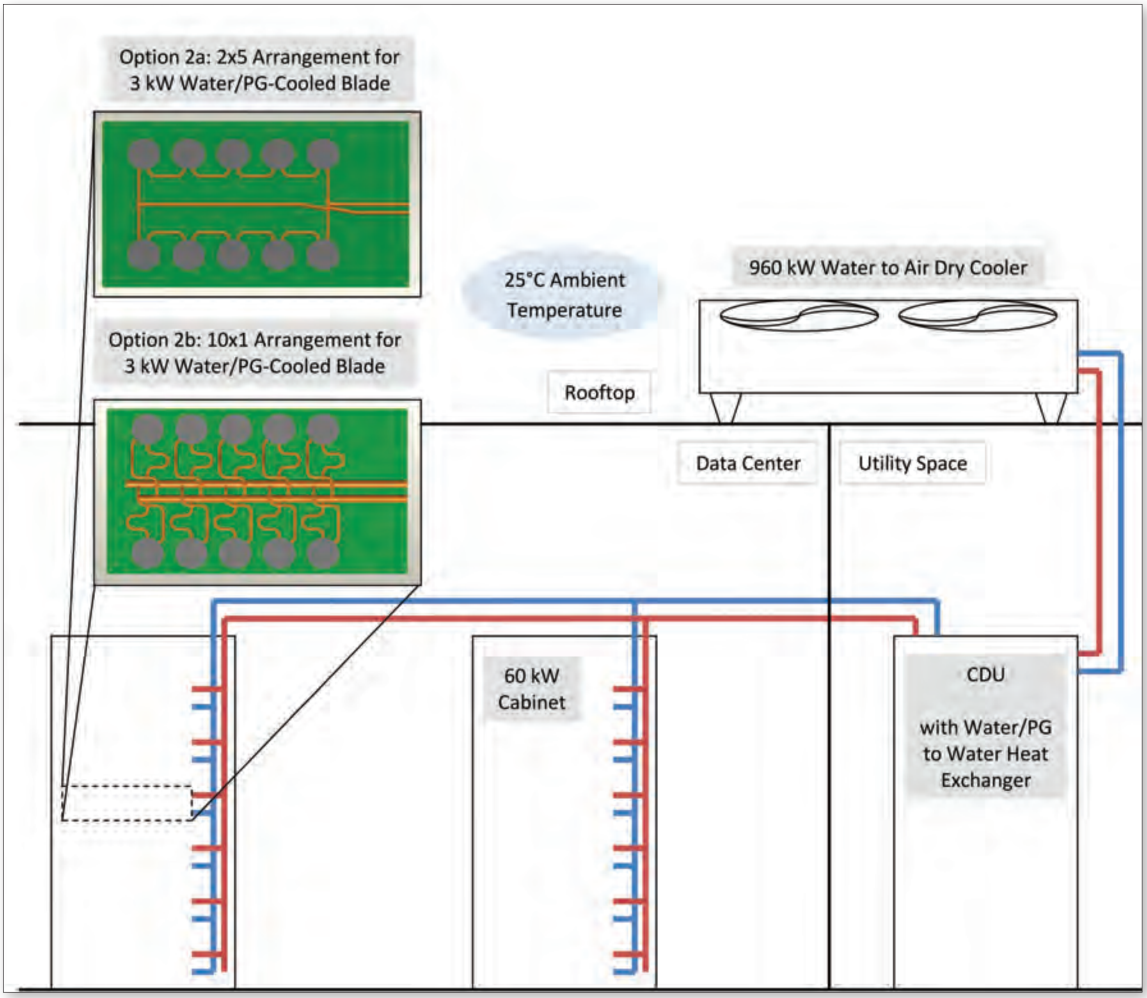

Figure 3 - Option 2: Water/PG touch cooling

manifolds with additional piping to allow for a compliant cold plate/CPU interface.

Option 3: R134a Touch Cooling

The high heat of vaporization for refrigerant enables refrigerant systems to use a flow rate that is approximately 5 times less than the required flow rate for a water system with the same power. Because of this, the cold plates in Option 3 do not need to be connected in parallel like in option 2b, and can have an arrangement similar to option 2a. As before, blades were connected to vertical supply/return (refrigerant) manifolds, but because of the lower flow rate, a refrigerant pump can be fit into each cabinet to pump refrigerant through the blades and manifolds.

In option 3a, the manifolds transport the refrigerant to a refrigerant distribution unit (RDU), where the heat is transferred to a closed water loop feeding into a rooftop dry cooler. In option 3b, the refrigerant is sent straight to the roof, where a rooftop condenser dissipates the heat to the atmosphere. The layouts of these two options are shown schematically in Figure 4.

The control system for a refrigerant cooling system includes a built-in headroom for the refrigerant capacity. The system is designed so that under a full load, the refrigerant quality (the fraction of refrigerant that is vapor, by mass) does not exceed 80%, so a 20% safety factor is already built into the system at the worst-case scenario. Flow regulators at the inlet of each blade ensure even flow distribution across the entire height of the cabinet. In the event of an increase in CPU power, either the RDU would increase the water flow rate or the rooftop condenser would increase its fan speed to fully condense the refrigerant.

Capital and Operating Expenditures

With identical performance specifications (maximum case temperature and environmental ambient temperature), the differences between cooling systems can be easily compared

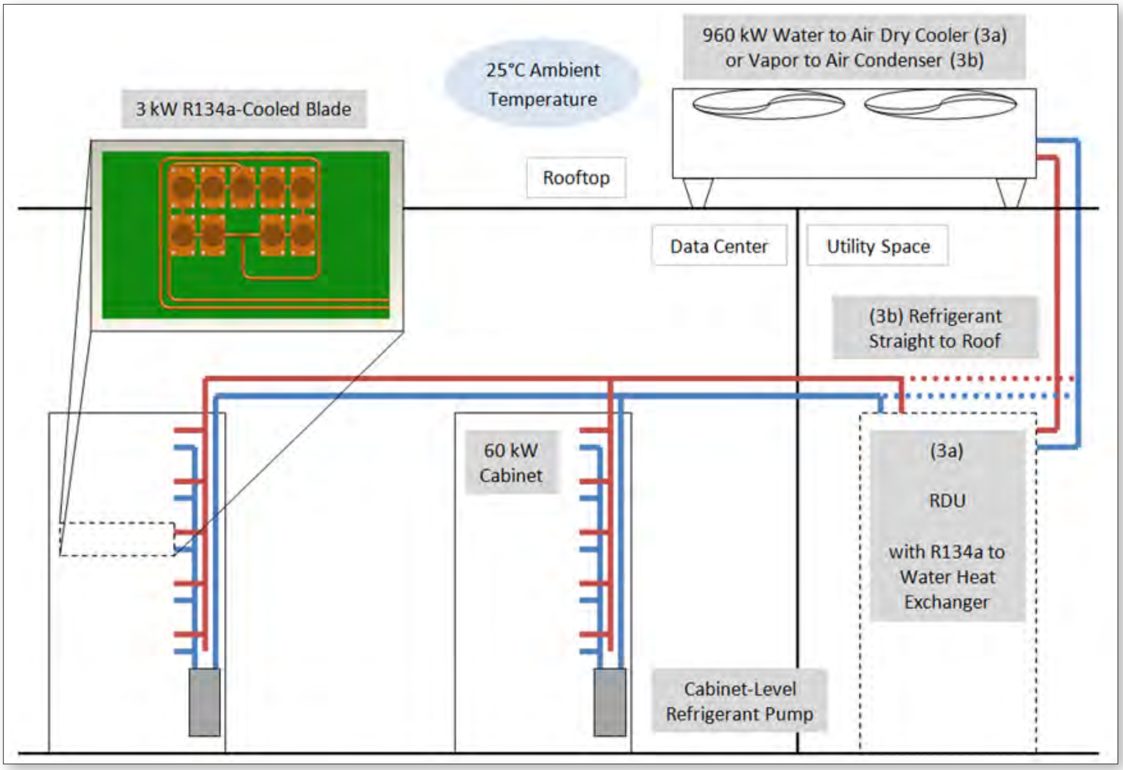

Figure 4 - Option 3: R134a touch cooling

in terms of capital and operating expenses (CapEx and OpEx, respectively).

It is important to mention that the presented first-pass analysis was not intended to produce entirely optimal designs for each cooling option. For equipment selection, computational fluid dynamics (CFD) analysis [7], flow network analysis [8], a two-phase analysis software suite, and vendors' product selection software [9, 10] were used to analyze pressure drops and heat transfer across different system components.

For each cooling option, capital expenses were determined by obtaining quotes of only the main components of cooling hardware (fans, heat exchangers, pumps, coldplates, refrigerant quick disconnects, etc.) from the manufacturer, and 10% of the cost was assessed for installation, piping, etc. The cost for the electric power supply, controls, hose and pipe fittings, UPS, etc. was not included. This cost is certainly an underestimate of the total cost of installation, but it would be representative of the main cost drivers associated with each cooling system.

One important note about the CapEx estimates presented in figure 5 is that the water and refrigerant coldplates are assumed to be of equal cost. In reality, water coldplates require higher flow rates and therefore larger tube diameters to cool the same power, but refrigerant coldplates must withstand higher pressures. The actual cost of manufacturing depends more on the manufacturing technique than it does on the fluid used, so the two coldplates were assumed to be of similar cost.

Operating expenses for all cases were calculated by determining the cost of electricity needed to pump the coolant around the loop and to run the fans. To do this, CFD and flow network analysis were used to calculate the pressure drop and flow rate of each fluid through the system, and then an average operational efficiency was used to determine the total power draw. This analysis used $0.10 / kW-hr, without demand charges, and operating hours per year were 8,760 for all methods.

Since the refrigerant-based option does not require periodic flushing and replacement, in our analysis the cost of R134a was only added to CapEx. With water cooling, in order to keep electro galvanic corrosion inhibitors and microbiological growth suppressants active, water/PG Coolant mixture requires regular (every 2-3 years) flushing, so the cost of water cooling additives was added to OpEx as well as the initial CapEx estimate.

Figure 6 shows that using a direct, rear-door air-cooling approach, low data center temperature, heat sinks with em-

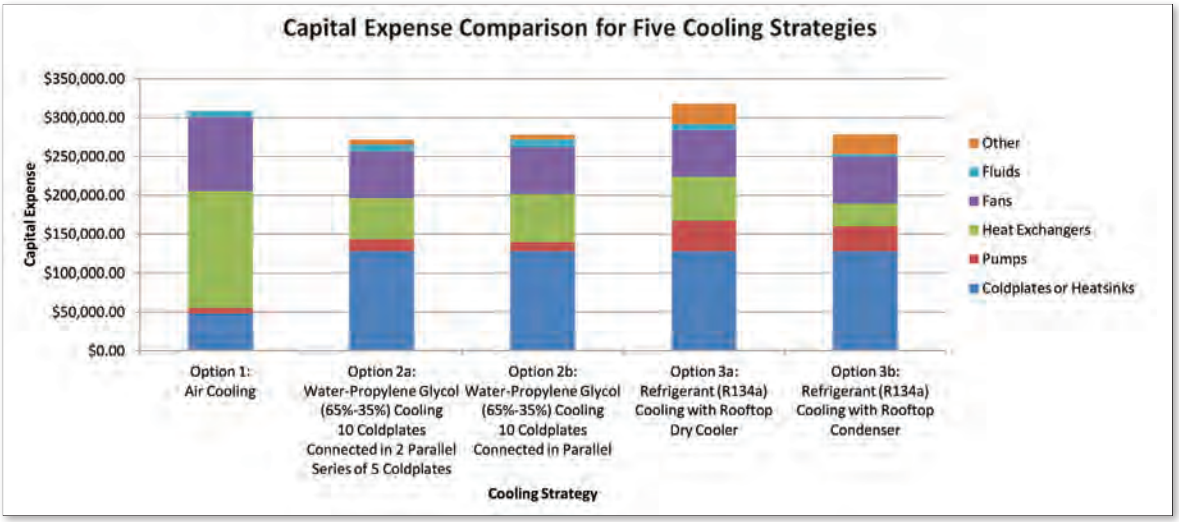

Figure 5 - Capital expenses for five cooling strategies

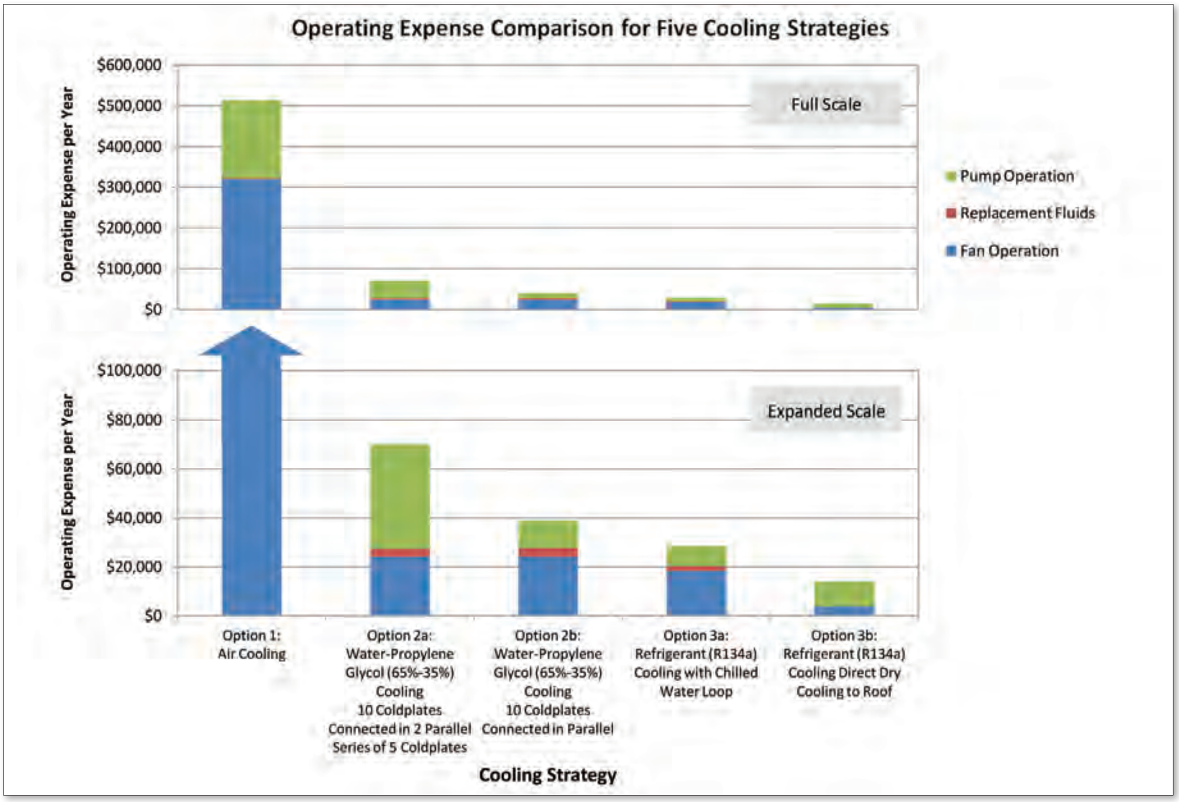

Figure 6 - Capital expenses for five cooling strategies

bedded heat pipes, and turbine-blade high-efficiency fans, air cooling will cost far more than a liquid-cooling option. Again, this best effort for air cooling option was presented only for the sake of comparison.

Figure 6 also indicates that that switching from air-cooling to any of the four liquid-cooling options will cut the operating expense (to run the cooling system, not to run the entire data center) by at least a factor of 5, In addition, the direct refrigerant cooling option (3b) shows the lowest operating cost of all the cooling options, with less than 1⁄30th of the cost of the air-cooled option. With this operating cost, an existing air-cooled data center (per option 1) would greatly benefit from switching to a refrigerant-cooled data center (option 3b), and, assuming no additional retrofit expenses, would recover the switching cost within the first year of operation.

CONCLUSION

Although the comparison in this paper is a preliminary, predictive analysis of several different cooling systems, the differences in power consumption revealed here show that a data center outfitted with liquid cooling provides a tremendous advantage over air cooling at the specified power level (60 kW per cabinet). As HPC and Telecom equipment continues toward higher power densities, the inevitable shift to liquid cooling will force designers to choose between water and refrigerant cooling. It is the author's belief that the industry will eventually choose direct refrigerant cooling because of the advantage it has over other cooling systems in operating cost (at least 2.5 times cheaper) with similar capital cost, the minimal space requirements on the board, the absence of microbiological growth, electro galvanic corrosion, and corresponding need to periodically flush the system, as well as the fact that conductive leaks are nonexistent.

ACKNOWLEDGMENTS

This study was inspired, promoted, and scrutinized by dedicated professionals representing all cooling approaches, from the Liquid Cooling forum at LinkedIn, and carried out by Thermal Form & Function, Inc.

REFERENCES

[1] ASHRAE TC 9.9 2011 "Thermal Guidelines for Liquid Cooled Data Processing Environments", 2011 American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

[2] Brock, David. Understanding Moore's Law: Four Decades of Innovation (PDF). Chemical Heritage Foundation. pp. 67–84. ISBN 0-941901-41-6. http://www.chemheritage.org/community/store/books-and-catalogs/understanding-moores-law.aspx, Retrieved March 15, 2015.

[3] Hannemann, R., and Chu, H., Analysis of Alternative Data Center Cooling Approaches, InterPACK '07 July 8-12, 2007 Vancouver, BC, Canada, Paper No. IPACK2007-33176, pp. 743-750.

[4] LinkedIn Liquid Cooling Forum. https://www.linkedin.com/grp/post/2265768-5924941519957037060, Retrieved May 11, 2015.

[5] Network Equipment-Building System (NEBS) Requirements: Physical Protection(A Module of LSSGR, FR-64; TSGR, FR-440; and NEBSFR, FR-2063). NEBS Requirements: Physical Protection GR–63–CORE, Issue 1, October 1995

[6] Xcelaero Corporation, "Hurricane 200," A457D600-200 datasheet, 2010. http://www.xcelaero.com/content/documents/hurricane-brochure_r5.indd.pdf, Retrieved May 11, 2015.

[7] ANSYS Icepak. Ver. 16. Computer Software. Ansys, Inc., 2015.

[8] AFT Fathom. Ver. 7.0. Computer Software. Applied Flow Technology, 2011.

[9] LuvataSelect. Ver. 2.00.48. Computer Software. Luvata, 2015. http://www.luvataselectna.com, Retrieved May 11, 2015.

[10] Enterprise Coil Selection Program. Ver. 3.1.6.0. Computer Software. Super Radiator Coils, 2015. http://www.srcoils.com/resources-support/coil-sizing-software/, Retrieved May 11, 2015.

Index of Advertisers

SUNON Active Cooling Solution for 400W High Power LED Lighting

Original Cooling Solution
(without fan)

Sunon Cooling Solution
(with fan)

Reduce size, weight up to **90%!!**

Product can meet IP68 & GR487 Rated

SUNON®

Sunonwealth Electric Machine Industry Co., Ltd

Headquarters (Taiwan) URL : www.sunon.com

Sunon Inc. (U.S.A.) E-mail : info@sunon.com | Tel : +1-714-255-0208