CHARACTERIZATION OF SERVER THERMAL MASS

Electronics COOLING

DECEMBER 2013
electronics-cooling.com

FEATURES

COOLING THE CLOUD:

ENERGY-EFFICIENT WARM-WATER COOLING OF SERVERS

THERMAL MANAGEMENT OF MANY-CORE PROCESSORS USING POWER MULTIPLEXING

BACK COVER

It's Lights Out Again.

Do You Really Want To Take Your Brand Integrity Down This Road?

Bergquist Thermal Clad® keeps your LEDs intense, bright and reliable wherever they go.

Built for long-term reliability.

Choosing the right IMS (insulated metal substrate) can make the difference of a successful product or not. Knowing the long term performance and reliability of the materials in your design, will give your customers a quality product and protect your reputation. This includes:

- *Thermal impedance, not just thermal conductivity*
- *Long-term dielectric strength, not just breakdown strength*
- *U.L. Listed*
- *Long-term temperature testing*

A balance of innovative thinking and expertise.

Not all substrate materials have the Bergquist company's 25+ years of testing and proven field performance. Our testing is geared to the application, not just numbers, which give a true performance picture you can count on in your design.

Bergquist delivers cooler performance for Power LEDs.

Dielectrics: Bergquist offers a wide selection of dielectric and thickness choices to meet the demanding specs of high performance applications (Bergquist HPL and HT versions), as well as, lower performance requirements (Bergquist MP, HR T30.20 versions).

Base Plates: Bergquist offers a wide range of thicknesses, aluminum for packaged components and copper base for high power density and COB applications.

Circuit Copper: From 35µm to 350µm (1/2 oz. to 10 oz.)

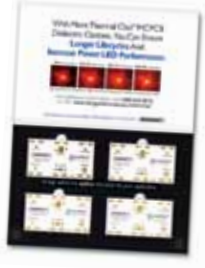

Call or visit us to qualify for a FREE T-Clad Samples Kit:
www.bergquistcompany.com/coolkit

www.bergquistcompany.com **1.800.347.4572**

952.835.2322 fax 952.835.0430

18930 West 78th Street • Chanhassen, Minnesota 55317

THE BERGQUIST COMPANY

Thermal Materials • **Thermal Substrates** • **Fans and Blowers**

CONTENTS

FEATURED ARTICLE

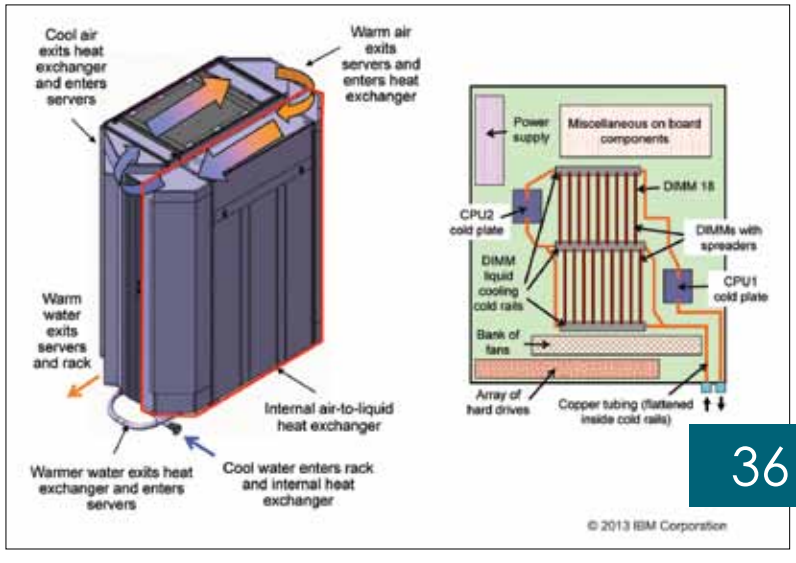

Editorial

Madhusudan Iyengar, Editor-in-Chief, December 2013

RECENTLY I READ AN ARTICLE on cutting-edge research demonstrating the feasibility of data storage on media for a million years! A few years ago this might not have resonated with me as a thermal engineer working to cool electronics, but now it does. Let me explain why.

Research & Development on IT equipment thermal management over the past few decades has justifiably focused on microprocessor (CPU) based systems. This may be explained by increased CPU transistor density, as per Moore's law, hence rising heat fluxes. However, in parallel, the growth rate of storage hardware that complements the compute engine (i.e., microprocessors) has been dramatic in recent years, not only in data centers that support internet based products and services (e.g. social media, photo/video hosting, utility/cloud computing, e-commerce), but also in the governmental, medical, and industrial sectors. While microprocessors come and go with new electronics hardware, the expectation for how long the data need to be stored for can easily range from decades to a lifetime to potentially centuries (e.g., in the case of historic and scientific data). As humans with intricately embedded electronics, we generate and use data at an ever increasing rate. This should not come as a surprise to any of us, engineers or otherwise, based on our own data storage needs at work or at home. On a personal note, I own several storage devices ranging from CDs and consumer disk drive devices to DVDs. The number of storage devices has continuously grown rapidly over the last few years, especially with the addition of so many photographs of my daughter since her birth.

So, why should this matter for a thermal engineer in the electronics cooling field? In addition to cooling traditional electronic devices and systems, some of us are responsible for reliably cooling storage devices and systems (and data centers). Storage devices could be made of magnetic media (tapes or rotationally spinning disk drives) or silicon (flash) or something else. As an example, a common form factor for hard disk drive with spinning magnetic media is 4" by 6" by 1", with a power in the range of 10 W and maximum case temperature of approximately 60°C. By itself, this thermal management problem is a piece of cake for a thermal engineer, but consider hundreds or thousands or millions of such drives closely packed with each other, with fan-based air cooling doing the job. Now consider that some of the devices store "hot" or active data, and others store "cold" or inactive data. Each of these storage tiers have different cooling needs over their lifetime. I anticipate that this area of thermal engineering will increase in importance in the coming years, and I thus wanted to use my editorial capital for this issue to discuss this topic with the ElectronicsCooling audience.

For this edition of the magazine, I am excited to offer the readers several interesting topics such as justification for embedded liquid cooling for very high performance microprocessors, numerical analyses to investigate the concept of intentional transient power manipulations in microprocessor cores, transient thermal characterization of servers, and highly energy-efficient liquid cooling systems for data center facilities. In addition, we have our traditional favorites, which are the calculation corner and the thermal facts and fairy tales sections. I hope that you enjoy this stimulating issue of *Electronics Cooling* and would like to offer you all my warmest wishes for a great holiday season. Stay warm or cool based on your preference, but I hope it is thermally energy efficient!

Electronics COOLING

www.electronics-cooling.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

BUSINESS DEVELOPMENT MANAGER
Casey Goodwin
cgoodwin@item-media.net

GRAPHIC DESIGNER
Evan Schmidt
eschmidt@item-media.net

EDITORIAL ASSISTANT
Aliza Becker
abecker@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.

MALICO INC

Special line attachment maker

Malico is not an ordinary company,
It's unique and extra ordinary product driven not profit driven.
Our products are so carefully made and almost zero defect.
We take extreme care of every part in our products.
So, Our products are so reliable and fool prove.
You can rely on us 100% for using our products and care free.
Let's grow together in the industry and be prosperous.

Contact offices:
Boston , MA Tel: 978-771-9285
Chino , CA Tel: 909-598-3738
Dallas , TX Tel: 214-514-9836
14776 Yorba Ct. , Chino , CA 91710

www.malico.com

Cooling Matters

Applications of thermal management technologies

RESEARCHERS: CLAY IS KEY INGREDIENT IN HIGH-TEMP. SUPERCAPACITORS

Clay, an abundant natural material, may hold the answer to developing super high-temperature supercapacitors for powering devices in extreme environments, says new research from Rice University.

According to a report, the supercapacitor created by researchers at Rice University using naturally-occurring clay and room temperature ionic liquids (RTIL) is reliable at temperatures of up to 200°C (392°F), and possibly higher. With such a tolerance, the new supercapacitor design could help raise temperature restrictions on devices used in oil drilling, military and space applications.

The Rice University team created a paste comprised of a room-temperature ionic liquid (RTIL) first developed in 2009 with natural Bentonite clay. The paste was then sandwiched between layers of reduced graphene oxide and two current collectors to form the supercapacitor. Tests and subsequent electron microscope images showed very little change in the paste material after it was heated up to 300°C. Researchers also observed that despite a slight drop in capacity observed in the initial charge/discharge cycles, the supercapacitor remained stable through 10,000 test cycles.

Source: Scientific Reports

'ELECTRONIC BLOOD' POWERS AND COOLS BRAIN-INSPIRED COMPUTER

IBM has unveiled a prototype of a new brain-inspired computer powered and cooled by what the company is calling "electronic blood." The company says it is learning from nature in an effort to develop small, highly efficient computing system.

"We want to fit a supercomputer inside a sugar cube," IBM researcher Bruno Michel, Ph.D., told the BBC. "To do that, we need a paradigm shift in electronics—we need to be motivated by our brain."

According to Michel, the human brain is able to consume energy at a low rate "because it uses only one—extremely efficient—network of capillaries and blood vessels to transport heat and energy—all at the same time." The new IBM prototype computer emulates this concept using "redox flow" system, which pumps an electrolyte "blood" through the computer. The liquid is charged via electrodes and pushed into the computer, where it discharges energy to the processor chips and then carries heat away.

Ultimately, IBM hopes to shrink a one petaflop computer—which would fill half a football field today—down to the size of a standard desktop computer by 2060, and one day use its "electronic blood" concept to achieve zettascale computing.

Source: BBC

MIT RESEARCH DISCOVERS NEW SURFACES EASE THE DISSIPATION OF EXTREME HEAT

A new method of cooling extremely hot surfaces more effectively from researchers at MIT could benefit industrial equipment and electronic devices.

Liquid cooling is commonly used in a number of applications. MIT researchers developed a way to keep the water droplets from bouncing as part of a method to cool hot surfaces more effectively.

The method involves adding tiny microscale silicon posts to a smooth silicon surface to create a textured surface.

Source: MIT

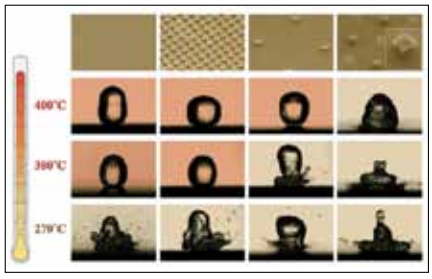

▲ A micrograph showing water droplets landing on specially designed silicon surfaces at different temperatures. At higher temperatures, the droplets begin to exhibit a new behavior: instead of boiling, they bounce on a layer of vapor, never really wetting the surface.■

Datebook

DECEMBER 11-13

Electronics Packaging Technology Conference
SINGAPORE, TAIWAN
www.clocate.com/conference/15th-Electronics-Packaging-Technology-Conference-EPTC-2013/28329/

JANUARY 7-10

Consumer Electronics Show 2014
LAS VEGAS, NEV., US
www.cesweb.org

FEBRUARY 5-6

9th European Advanced Technology Workshop on Micropackaging & Thermal Management
LA ROCHELLE, FRANCE
www.france.imapseurope.org

MARCH 9-13

SEMI-THERM
SAN JOSE, CALIF., USA
www.semi-therm.org

SEMI-THERM THIRTY 30

NEW HEAT EXCHANGER RESEARCH COULD LEAD TO BETTER LIQUID COOLING

A New Mexico State University assistant professor is investigating methods of increasing heat exchanger efficiency that could enable space missions to remain in orbit for longer periods of time and lead to better automotive, defense, data center and power generation thermal management systems.

Currently, the length of space missions is severely limited because of the gradual loss of cryogenic propellants due to boil-off caused by solar heating. Thus far, the team has found that under certain operating conditions, cryogenic heat exchangers could be reduced to at least one-third the size of current models, with more than 10-15 percent improvement in thermal performance. In cooperation with Brian Motil, chief of the Fluid Physics and Transport Branch of NASA GRC, they are working to integrate their findings into cryocooler systems.

Source: The El Paso Times

HIGH TEMPERATURE CAPACITORS REMOVE NEED FOR EV COOLING SYSTEMS

A new lead-free, high temperature ceramic capacitor developed by scientists at the National Physical Laboratory UK could improve the efficiency and reliability of electric and hybrid vehicles by enabling the removal of heavy complex cooling systems.

NPL researchers claim their new material, called HITECA can hold more energy and function at much higher temperatures—over 200°C—than conventional capacitors. HITECA is made from a ceramic paste with a granular stucture comprising a bismuth ferrite ($BiFeO_3$) compound doped with strontium-titanate ($SrTiO_3$). The NPL research team is now looking to test the capacitor technology in an electric vehicle power electronics system.

Source: The Engineer

SONY RELEASES SOURCE CODE FOR THERMAL MANAGEMENT SOLUTION

Sony has released the source code for Thermanager, a thermal management solution developed for and used in the company's Android Open Source Project (AOSP) for Xperia projects. The source code is available through the Sony channel on GitHub.

"Thermal management for these projects has been handled by binaries released through Developer World. While this solves the immediate problem of devices overheating, it does not provide developers a way to customize, improve or analyze the behavior of this important part of the system," the company said in a recent statement.

Source: SONY

NEW COOLING TECHNOLOGY FOR MRI MACHINES OFFERS HELIUM ALTERNATIVE

A new way to cool magnets to the extremely low temperatures needed for MRI machines may hold the answer to overcoming the effects of the increasing global helium shortage.

Well-known for its use in party balloons, helium is also used in the cooling of superconducting magnets, with the main commercial application being in MRI scanners and particle accelerators.

Now, a company called Cryogenic has developed a technique to cool magnets that requires only a small fixed amount of helium—approximately half a liquid liter—and mechanical coolers which run using electrical power and cooling water.

Source: *Wired*

APRIL 7-9

IEEE Euro-SimE
GHENT, BELGIUM
http://www.eurosime.org

APRIL 27-MAY 1

International Thermal Conductivity Conf. and International Thermal Expansion Symposium
WEST LAFAYETTE, IND., US
www.thermalconductivity.org

MAY 7-9

Smart Lighting 2014
BARCELONA, SPAIN
www.smartlighting.org

JUNE 2-5

Thermal Imaging Conference 2014
LAS VEGAS, NEV., US
www.thermalimaging conference.com

Evolving the Role of the Thermal Engineer
from Analyst to Architect

Scott Johnson, Raytheon Space and Airborne Systems
Brendon Holt, Raytheon Missile Systems

EVOLVING TRENDS within the aerospace industry are creating a new set of challenges that directly influence thermal management choices and planning that should be addressed in the early phases of the product development process. Within the functional side of the aerospace industry, the term "system architect" has been created to describe individuals with expertise in subsystem design, acquisition and integration. Hence, it is the position of the authors that the role of the thermal engineer must also evolve to have a similar systems perspective.

Based on past and recent experiences of the authors, there are four levels of interaction between thermal engineers and the product development process. These levels of thermal interaction fall into four categories:

1. *Thermal Analysis.* Interaction at this level is common to almost all product development efforts. The thermal engineer performs equation based and/or detailed thermal analysis using computational techniques and reports requirements compliance at the component level.
2. *Thermal Design.* This level of interaction is also common to most product development efforts. The thermal engineer works in parallel with product design team to provide balanced design solutions that enable compliance with thermal design requirements.
3. *Thermal Systems Engineering.* Interaction at this level is less common. The thermal engineer creates system level thermal performance models, uses these to derive and document thermal design requirements and then manages execution by product design teams across all program phases (design, build, qualification, production).
4. *Thermal Architect.* Interaction between a thermal engineer and a product development activity at the architect level is desirable but infrequent. The thermal engineer engages during the pursuit and competition/capture phases of the program to determine the best over all thermal systems approach that satisfies program objectives. The au thors refer to this as Mission Level Engagement. Ideally, the thermal systems engineer and the thermal architect are the same individual.

The impact of where these roles are applied in the product development process is illustrated in the left side of Figure 1. The highest perspective is the mission level and the lowest the component level. Ideally, the thermal architect would be engaged at the mission level (or as close as possible) early in the program pursuit and capture phases. The thermal architect would then transition to the role of thermal systems engineer, deriving and documenting thermal requirements at the system level and then managing thermal design and analysis activities at the subsystem, subassembly and component level. In general, thermal management and design decisions made at higher levels afford the benefits of more flexibility, more opportunities for design optimization and synergy, and less risk. These benefits are illustrated by the

Scott T. Johnson is an engineering fellow with more than 30 years of experience at Raytheon Space and Airborne Systems in El Segundo Calif., where he is part of the Mechanical and Optical Engineering Center (MOEC). Johnson specializes in system level thermal and thermodynamics management for military airborne and ground based sensors. Scott also specializes in the development of advanced thermal technologies for very high heat flux component design and holds multiple patents in this area. Johnson received his BSME degree from the University of Florida in 1981.

Brendon R. Holt, Ph.D. is currently a senior principle engineer at Raytheon Missile System in Tucson Az where where he is a thermal section head. He leads a group that is responsible for providing thermal engineering analysis, design and testing support for military land, sea, air and space systems. He has been at Raytheon and legacy Texas Instruments Defence Segment since 1995. He received his Ph.D. from the University of Texas at Arlington (1995) in Fluid Dynamics and Heat transfer with phase change.

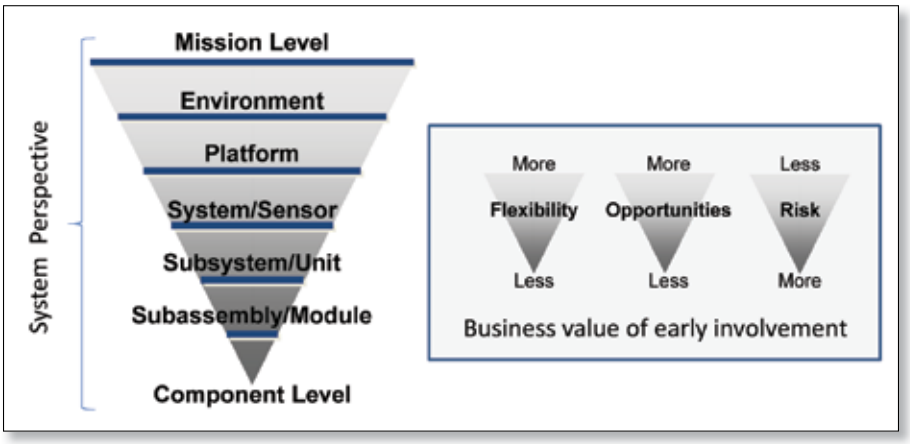

FIGURE 1: Levels of Thermal Management Interaction with Product Development.

smaller inverted pyramids on the right side of the figure.

At present, mission level engagement of the thermal architect is something of an idealization. Many product development activities simply are not structured to benefit from the engagement of a thermal engineer at the mission level. Also, mission level engagement of the thermal engineer would require something of a culture change within the industry and these are slow to occur. So in order to precipitate change in the culture the thermal engineer's skill and tool set will have to evolve beyond the traditional thermal analysis and thermal design roles into the role of the thermal systems engineer and thermal architect. When it comes to modeling, the attributes most critical to the thermal architect and systems engineer's effectiveness are speed and scalability. This means modeling only what is necessary and doing it quickly and with the needed accuracy and flexibility. This also means not being seduced by the siren's call of the "rainbow plots" created using modern computational fluid dynamics (CFD) and simulation codes too early in the program. The level of detail required for these object based models is typically not available during the early phases of the program when thermal management decisions have the most impact.

In 2011, one of the author's participated in a technology study derived from the Air Force Research Laboratories (AFRL) INVENT (Integrated Vehicle Energy Technology) initiative [1]. As part of the effort, a complete equation based transient model was created for a critical pod mounted sensor and integrated to run in MATLAB Simulink. This was combined with a platform (aircraft) level environmental cooling system (ECS) Simulink model and used to fly virtual missions. One of the objectives of the INVENT initiative was to extend this approach to integrate all of the platform critical energy subsystems. These would include energy "demand" functions like critical sensors and energy "supply" functions such as the engine, fuel delivery and electrical power generation subsystems. The benefit of this approach is true concurrent design at the mission level. The model evolves with the program and is used to actively manage the product development process. This is a real example of how a thermal engineer would function as a thermal architect engaging near the top of the System Perspective illustrated in Figure 1.

So what is the moral of this "facts and fairy tales" story. The fairy tale is that the role of the thermal engineer is defined entirely by thermal analysis. The fact is that acting as thermal architects, we have the opportunity to influence design choices at the earliest program phases in a way that maximizes product performance while minimizing technical and business risk.

REFERENCES

[1] http://acdl.mit.edu/mdo/mdo_10/INVENT%20M&S%20-%20Wolff.pdf

Thermal Packaging
From Problem Solver to Performance Multiplier

Avram Bar-Cohen
Microsystems Technology Office, Defense Advanced Research Projects Agency (DARPA)

THE INCREASED integration density of electronic components and subsystems, including the nascent commercialization of 3D chip stack technology, has exacerbated the thermal management challenges facing electronic system developers. The sequential conductive and interfacial thermal resistances associated with the prevailing "remote cooling" paradigm in which heat must diffuse from the active regions on the chip to the displaced coolant, have resulted in only limited improvements in the overall junction-to-ambient thermal resistance of high-performance electronic systems during the past decade. These limitations of Commercial Off-The-Shelf (COTS) thermal packaging are undermining the cadence of Moore's Law and leading to a growing number of products that fail to realize the inherent capability of their continuously improving materials and architecture. Moreover, the uncritical application of this "remote cooling" paradigm, has resulted in electronic systems in which the thermal management hardware accounts for a large fraction of the system volume, weight, and cost and undermines efforts to transfer emerging components to small form-factor applications.

To overcome these limitations and remove a significant barrier to continued Moore's Law progression in electronic components and systems, it is essential to implement aggressive thermal management techniques that directly cool the heat generation sites in the chip, substrate, and/or package. The development and implementation of such "Gen-3" embedded thermal management technology, following on the Gen-1 air-conditioning approaches of the early years and the decades-long commitment to the Gen-2 "remote cooling" paradigm, is the focus of the current DARPA Intra/Inter Chip Enhanced Cooling (ICECool) thermal packaging program. Launched in 2013, ICECool aims to develop and demonstrate "embedded cooling" techniques capable of removing kW/cm^2 chip heat fluxes and kW/cm^3 chip stack heat densities, while suppressing the temperature rise of multi-kW/cm^2 sub-mm hot spots [1,2]. The ICECool program is composed of two thrusts: a 3-year ICECool Fundamentals effort, involving several university teams which are developing embedded cooling building blocks and modeling tools, and a 2.5-year ICECool Applications effort, led by several aero-

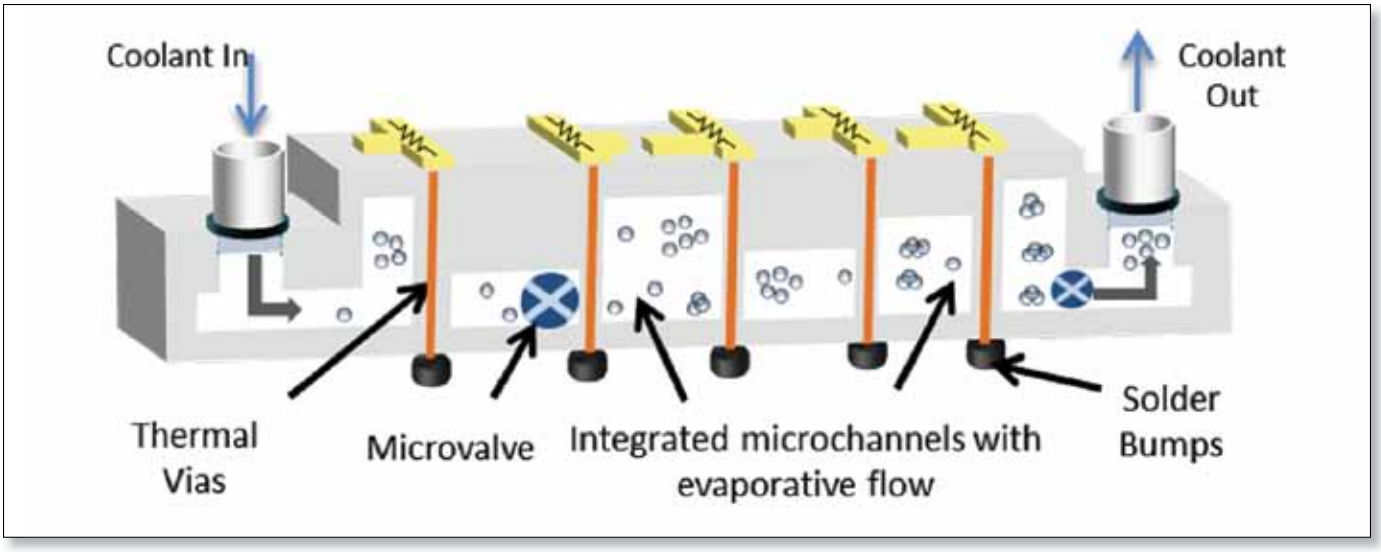

FIGURE 1: A Cross-Sectional Conceptual Schematic of an Embedded Cooling, Gen-3 (ICECool) device.

sapa:

Shaping the future

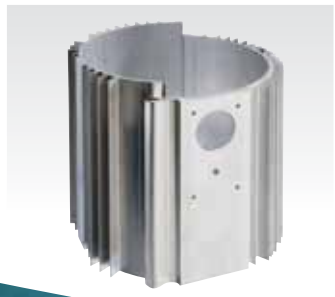

COOLER. FASTER.

ALUMINUM EXTRUSIONS MAKE THE DIFFERENCE

And Sapa is the largest heat sink extruder in North America, and the world!

Sapa has been providing heat sink extrusions for over 30 years, with 16 manufacturing locations in North America and Thermal Management Centers of Excellence on both coasts that extrude and fabricated finished heat sinks.

The strategic location of production facilities on both coasts means reduced shipping time and cost for thermal management customers, as well as providing centralized expertise for more complex thermal management products through Sapa's North American Technical Center—a research, development, and design center within the company.

In addition, Sapa has developed a new method of manufacturing high ratio air cooled heat sinks. Utilizing our experience in aluminum extrusion combined with Friction Stir Welding technology we're now able to achieve heat sink ratios in excess of 40:1. Our new technique utilizes a modular concept allowing for maximum flexibility in fin geometry, fin ratio, and heat sink footprint. It also allows for the production of large scale heat sinks up to 20" width.

Maximize your product designs with highly efficient thermal management solutions from Sapa Extrusions North America – it's what we like to call **Aluminology**™, or the art and science of making aluminum extrusion technology work for you.

MAKING THE CASE

- ✓ Keep your products cooler longer – extruded heat sinks are up to 53% more efficient than aluminum castings.
- ✓ Incredible design flexibility with fin ratios up to 40:1.
- ✓ Nationwide network of production facilities for rapid response to your needs.

Sapa Extrusions North America

877-710-7272 | www.sapagroup.com/NA

NorthAmerica.Sales@sapagroup.com

space performers and culminating in functional electronic demonstration modules. ICECool performers are pursuing the creation of a rich micro/nano grid of thermal interconnects, using high thermal conductivity, as well as thermoelectric, materials to link on-chip hot spots to microfluidically-cooled microchannels. Such intra/inter chip enhanced cooling approaches are required to be compatible with the materials, fabrication procedures, and thermal management needs of homogeneous and heterogeneous integration in 3D chip stacks, 2.5D constructs, and planar arrays. A conceptual ICECool device is shown in Figure 1.

An intrachip approach would involve fabricating micropores and microchannels directly into the chip [3,4] while an interchip approach would involve utilizing the microgap between chips in three-dimensional stacks [5,6], as the cooling channel. In addition to the inclusion of an appropriate grid of passive and/or active thermal interconnects, it is expected that a combination of intrachip and interchip approaches, linked with thru-silicon and/or "blind" micropores, will confer added thermal management functionality. These microchannels and/or micropores will be integrated into a fluid distribution network, delivering chilled fluid to the chip or package and extracting a mixture of heated liquid and vapor to be transported to the ambiently-cooled radiator.

Some 30 years of thermofluid and microfabrication R&D, driven initially by the publication of the Tuckerman & Pease microchannel cooler paper in 1981 [1] and more recently by compact heat exchanger and biofluidic applications [4,8], has created the scientific and engineering foundation for the aggressive implementation of the "embedded cooling" paradigm. Nevertheless, substantial development and modeling challenges must be overcome if Gen-3 techniques are to supplant the current "remote cooling" paradigm. Successful completion of the DARPA ICECool program requires overcoming multiple microfabrication, thermofluid, and design challenges, including:

- Subtractive and additive microfabrication in silicon, silicon carbide, and synthetic diamond of high aspect ratio, thin-walled microchannels and high aspect ratio micropores; low thermal boundary resistance, high thermal conductivity thermal interconnect grids; on-chip, high power factor, high COP thin-film thermoelectric coolers; and hermetic attachment of liquid supply and liquid/vapor removal tubes.

- Convective and evaporative thermofluid transport in microchannels and micropores– removal of 1 kW/cm^2 chip heat fluxes with 2-5 kW/cm^2 sub-millimeter "hot spots"; low pumping power liquid-vapor manifolds with Coefficients-of-Performance (CoP) between 20 and 30; high-exit-quality, greater than 90%, evaporative flows without flow instabilities and/or local dryout; and high fidelity thermofluid models for single- and two-phase flow in microchannels, microgaps, and micropores.

- Thermal and electrical co-design which moves progressively from passive, thermally-informed designs, which recognize the impact of temperature on functional performance, to active thermal co-design which places functional paths and blocks in the most favorable locations on the chip, to fully-integrated co-design which deals with the impact of

microfluidic channels and thermal interconnects on the electrical design and placement of electrical devices and cells, to mature designs that interactively balance the use of resources to optimize layout for energy consumption and functional performance.

- Physics of Failure models that address the failure mechanisms and reliability of the Gen-3 thermal management components, including erosion and corrosion in microchannels, microgaps, and micropores; failure modes induced in the electrically active areas of the chip and/or substrate; and the impact of microfabrication and embedded cooling operation on the structural integrity and stress profile of the microchanneled substrate (intrachip) and/or the chip-to-chip bonding (interchip).

Successful development and implementation of this Gen-3 thermal packaging paradigm would place thermal management on an equal footing with functional design and power delivery, transforming electronic system architecture and unleashing the power of nanofeatured device technology, while overcoming the SWaP (size, weight, and power consumption) bottleneck encountered by many advanced electronic systems. After decades of mere "problem solving" with Gen-1 (HVAC) and Gen-2 (spreaders, heat sinks, and TIMs) thermal management technology, it is expected that widespread adoption of Gen-3 "embedded cooling" techniques will provide a significant performance multiplier for advanced electronic components.

REFERENCES

[1] Intrachip/Interchip Enhanced Cooling Fundamentals (ICECool Fun), DARPA BAA 12-50, Posted 7 June 2012.

[2] ICECool Applications (ICECool Apps), DARPA BAA 13-21, Posted 6 February 2013.

[3] D.B. Tuckerman and R.F.W. Pease. "High-performance heat sinking for VLSI." IEEE Electron Device Letters, vol. 2, no. 5, pp. 126–129, 1981.

[4] L.P. Yarin, A. Mosyak and G. Hetsroni. Fluid Flow, Heat Transfer, and Boiling in Micro-Channels. Springer, Berlin, 2009.

[5] A. Bar-Cohen and K.J.L. Geisler. "Cooling the Electronic Brain." Mechanical Engineering, pp. 38–41, April 2011

[6] A. Bar-Cohen, J. Sheehan and E. Rahim. "Two-Phase Thermal Transport in Microgap Channels - Theory, Experimental Results, and Predictive Relations." Microgravity Science and Technology, pp. 1–15, Sept. 2011.

[7] S.Garimella, T. Harirchian, Microchannel Heat Sinks for Electronics Cooling, Volume 1, Encyclopedia of Thermal Packaging, 2013, WSPC, Singapore

Disclaimer: *The views, opinions, and/or findings contained in this article/presentation are those of the author/presenter and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense.*

Distribution Statement A, Approved for Public Release, Distribution Unlimited.

Use of JEDEC Thermal Metrics in Calculating Chip Temperatures (without Attached Heat Sinks)

Bruce Guenin
Assoc. Technical Editor

INTRODUCTION

JEDEC SINGLE-CHIP package thermal metrics are widely used as a means of characterizing the thermal performance of semiconductor packages. They correlate the peak temperature of a uniformly-heated semiconductor chip (the junction temperature, T_J) with the temperature of a specified region along the heat flow path. The values of these metrics are determined from temperature measurements under standardized conditions, which specify the test method, test board, and thermal environment in rigorous detail [1].

These metrics take the form of an Ohms law resistance calculation. Thermal resistances are calculated using the equation:

$$\Theta_{JX} = \frac{T_J - T_X}{P} \qquad (1)$$

where T_X is the temperature of the region to which the heat is flowing and P is the total power dissipated in the device.

Figure 1 indicates the primary heat flow paths in the four JEDEC-standard test environments: 1) natural convection [2] and forced-convection [3], and the conduction-cooled test environments : 2) junction-to-case [4] and 3) junction-to-board [5].

The junction-to-case and junction-to-board test environments force nearly 100% of the heat to flow along the indicated path.

TThe natural convection and forced-convection environments bear a closer resemblance to the majority of the end-

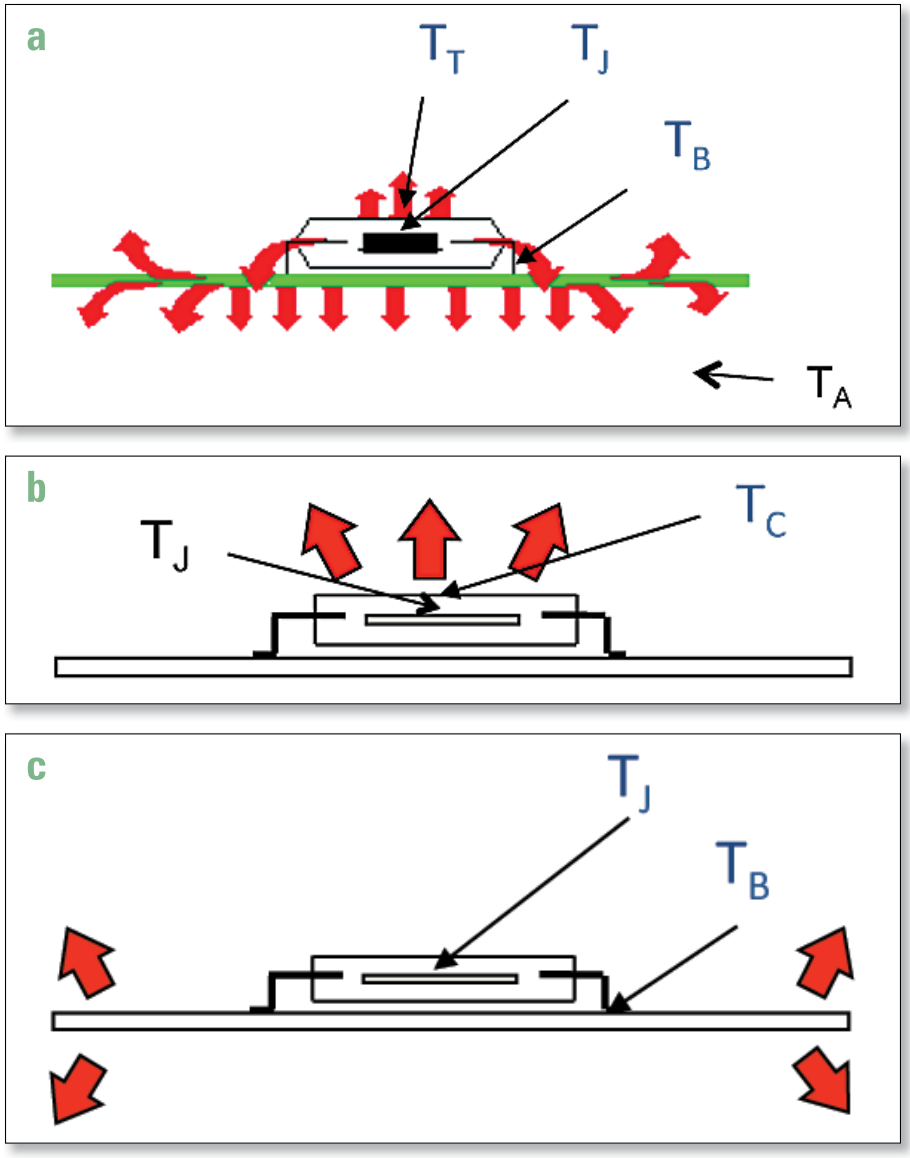

FIGURE 1: Diagrams representing typical heat flow paths for a package mounted to a test board for the four JEDEC standard test environments: **(a)** natural and forced convection, **(b)** junction-to-case, and **(c)** junction-to-board. The indicated temperature measurement locations are those specified by the respective test standard for each environment.

use environments for electronics in comparison with the other two. In these convection-cooled environments, the packages lose a portion of the dissipated heat to the air out of the top surface and the remainder to the board. To a reasonable approximation, for most applications, the two paths account for nearly all of the heat flow out of the package.

There is another class of JEDEC thermal performance metrics that are very useful in calculating junction temperatures for an application environment. They are called thermal characterization parameters and are represented by the Greek letter Ψ (pronounced "psi") [1]. They are calculated in the same fashion as the theta metrics, as follows:

$$\Psi_{JX} = \frac{T_J - T_X}{P} \quad \textbf{(2)}$$

A key difference between a thermal characterization parameter and a thermal resistance is that, in the case of the former, only part of heat flows to the region represented by the temperature, T_X. They are useful in estimating T_J, when T_X and P are known. The metrics that are most relevant to our purposes here are Ψ_{JT} and Ψ_{JB}, where T_T represents the temperature at the top center of the package and T_B is the board temperature, measured on a surface trace to which the package is soldered within a 1 mm distance from the edge of the package. The trace should contact the package at the middle of one of its sides.

A simple model can be exploited to gain additional insight into the thermal performance of the package/board assembly. Figure 2a represents the two heat flow paths using a simple thermal resistance network. The division of the total heat between these two paths is ultimately determined by the total thermal resistance along each path, which are equal to $\Theta_{JC} + \Theta_{CA}$ and $\Theta_{JB} + \Theta_{BA}$, respectively. The resistances Θ_{CA} and Θ_{BA} represent the efficiency of heat transport from the top of the package and the board, respectively, to the ambient air and are each a function of the air velocity. Unless a heat sink is attached to the top of the package, Θ_{BA} will be considerably less than Θ_{CA} due to the much larger area of the board surfaces for exchanging heat to the air compared to that of the top of the package. An inspection of this circuit shows that changes in air velocity will modify the values of Θ_{CA} and Θ_{BA} , leading to changes in the power flowing along the two paths and subsequent changes in T_T and T_B and in Ψ_{JT} and Ψ_{JB}.

The diagram in Figure 2b illustrates the interrelationship between Ψ_{JT} and Ψ_{JB} as the power flowing out the top of the package (P_{TOP}) changes. It shows how at small values of P_{TOP}, nearly all the power is flowing to the board and Ψ_{JB} approxi-

mates the value of Θ_{JB}. Conversely, if a heat sink is attached to the top of the package, such that P_{TOP} nearly equals the total power, then Ψ_{JT} approximates the value of Θ_{JC}.

This article explores the relationship of these various metrics and how they can be used 1) with an existing electronic system to estimate the junction temperature or 2) to predict the junction temperature for a system during the design process.

BEHAVIOR OF JEDEC THERMAL METRICS IN FORCED AIR TEST ENVIRONMENT

Figure 3a contains a graph of several metrics measured on a 35 x 35 mm, 388 ball, PBGA package with a 4-layer laminate containing two copper planes. The package was mounted to a 100 mm sq. JEDEC-standard board with two copper internal planes. The two upper curves are Θ_{JA} and Ψ_{JB}. A third curve, Ψ_{BA} was obtained by subtracting Ψ_{JB} from Θ_{JA}.

The Θ_{JA} curve displays a significant dependence on the air velocity, primarily due to the role of the board as a fin as it exchanges heat with the convective air flow over its entire area.

In contrast, Ψ_{JB} is nearly independent of air velocity. This results from the fact that most of the heat flow from the junction flows to the board by way of an internal conduction process. This suggests that Ψ_{JB} is relatively robust and is an effective measure of the heat conduction efficiency of the package with a relatively small contribution of the convective environment.

The Ψ_{BA} curve manifests the greatest sensitivity to the air velocity since, by definition, its magnitude is a direct result of the coupled conduction-convection cooling mechanism of the board.

Figure 3b plots Ψ_{JT} versus air velocity. As expected, it is a sensitive function of air velocity, since the heat flow from the die through the overmolded plastic cap is proportional to the heat transfer coefficient associated with a certain air flow velocity. One notes also that at air velocities up to 2.5 m/s, Ψ_{JT} is less than 1 ˚C/W. The implication of this is that for modest power levels, the temperature of the top center of the package is only slightly cooler than the junction. This serves to reduce the error in calculating T_J since the value of ΔT_{JT} calculated using Eqn. 2 would then be a very small correction to add to the much larger measured value of T_T.

This analysis leads to the conclusion that the main value of Ψ_{JB} is that it is moderately robust and can be used to predict the temperature difference between the junction and a board, when the dissipated power is known. Conversely Ψ_{JT} is more useful if one has an operating electronic system in front of him and wants to know the die temperature in a plastic package by a measurement of the temperature at the top of the package.

It is not always the case that Ψ_{JT} is provided by the package vendor. In this situation, it is relatively easy to calculate it.

In a previous column, the following equation was derived [6]:

$$\Psi_{JT} = \frac{h\,\Theta_{JA}\,t_{EMC}}{\kappa_{EMC}} \qquad (3)$$

where h is a value of heat transfer coefficient calculated at the air velocity of interest, and t_{EMC} and κ_{EMC} are the thickness

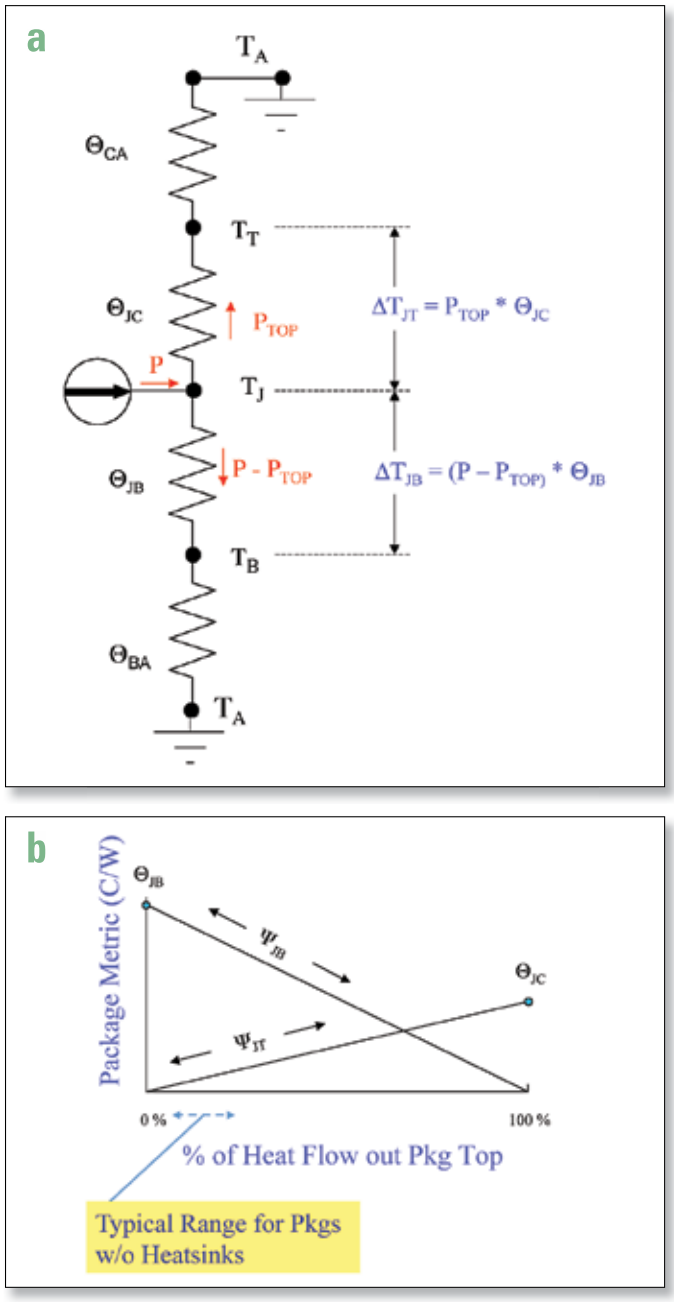

FIGURE 2: (a) Thermal resistor network representing junction-to-board-to-air path and junction-to-case-to-air path. **(b)** Diagram illustrating the relationship between Θ and equivalent Ψ values as a function of heat flow out of the top of the package to air.

of the epoxy mold compound above the die and its thermal conductivity. In the present example, t_{EMC} and κ_{EMC} are 0.86 mm and 0.7 W/mK, respectively.

The calculated values of Ψ_{JT} are plotted in Figure 3b and are also listed in the Table along with the remaining parameters input into Eqn. 3. The differences between the measured and calculated values of Ψ_{JT} are small in absolute terms.

USING Ψ_{JB} IS TO PREDICT JUNCTION TEMPERATURE IN A DESIGN

In order to successfully use the Ψ_{JB} metric in predicting the junction temperature, one needs to know T_B and the dissipated power and air flow conditions in the application.

There are numerous methods available for getting an estimate of T_B. Among the options is to build a load board, having the overall dimensions and conductivity of the actual board and applying the expected heat loads using heaters and also an appropriate cooling method [7]. Alternatively, one could create either an appropriate computational fluid dynamics (CFD) or finite element or finite difference conduction model of the board and accounting for the heat loads and cooling. Lastly, if a quick and moderately accurate solution is needed, it is possible to estimate the board temperature by including a single-package board in the model. The appropriate board size would be determined by taking the area of the actual board and multiplying by the ratio of the power of the device of interest to the total power due to all the devices on the board [8].

The usefulness of Ψ_{JB}, which is generated in a standard test environment, to the prediction of junction temperatures in an application results from the fact that in both environments, there is a convective boundary condition applied to the top of the package, which, ideally, would be at the same air velocity.

In general, a single resistance link from the die to the top of the package is not consistently accurate for an arbitrary boundary condition. To have robust boundary condition independence requires a compact thermal model, for example [9].

The initial work devoted to the development of the Ψ_{JB} standard showed similar values for a given package whether it was tested on JEDEC-standard multilayer boards or on an application board [10].

CONCLUSIONS

The JEDEC-standard thermal characterization parameters, Ψ_{JB} and Ψ_{JT}, can be valuable in facilitating the calculation of the junction temperature in a semiconductor package both for a live system and also in a predictive calculation. These metrics, determined by measurement, complement a purely thermal simulation approach by providing convenient experimental validation of simulated junction temperatures.

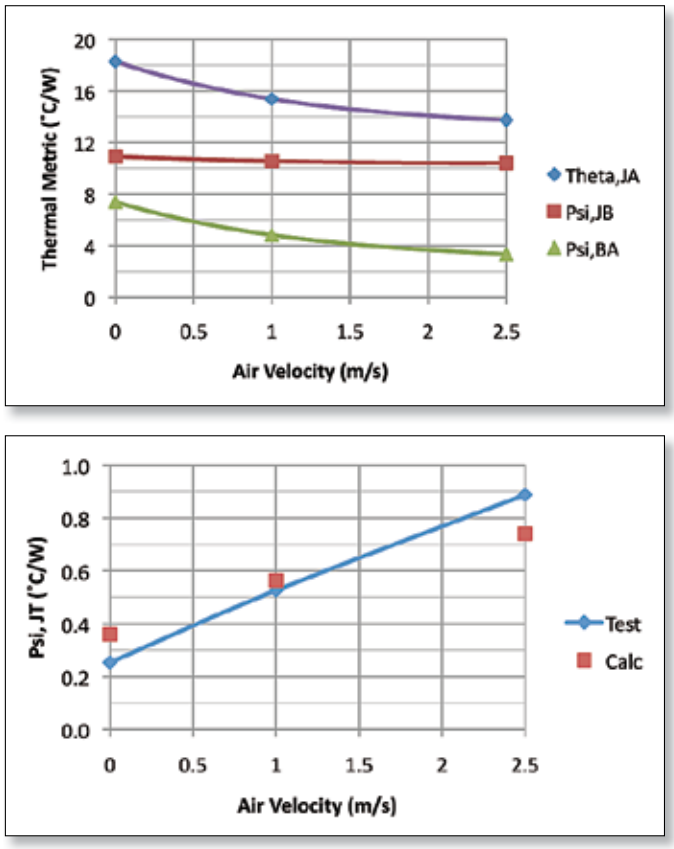

FIGURE 3: Thermal test results versus air velocity for a 35 x 35 mm, 388 ball, PBGA package having a 4-layer laminate attached to a 4-layer JEDEC-standard board.

TABLE: COMPARISON OF CALCULATED AND MEASURED VALUES OF Ψ_{JT}

35mm, 388 ball PBGA pkg, 4-Layer laminate on 4-Layer board				
Vair	h	Θ_{JA} Test	Ψ_{JT} Calc	Ψ_{JT} Test
m/s	W/m^2-K	C/W	C/W	C/W
0	16	18.3	0.36	0.25
1	30	15.3	0.56	0.53
2.5	44	13.7	0.74	0.89

REFERENCES

[1] JEDEC Standard JESD51-12, "Guidelines for Reporting and Using Electronic Package Thermal Information." All JEDEC standards are available for download at no charge from www.jedec.org.

[2] JEDEC Standard JESD51-2A, "Integrated Circuits Thermal Test Method Environmental Conditions - Natural Convection (Still Air)."

[3] JEDEC Standard JESD51-6, "Integrated Circuit Thermal Test Method Environmental Conditions - Forced Convection (Moving Air)."

[4] JEDEC Standard JESD51-14, "Transient Dual Interface Test Method for the Measurement of the Thermal Resistance Junction to Case of Semiconductor Devices with Heat Flow Through a Single Path."

[5] JEDEC Standard JESD51-8, "Integrated Circuit Thermal Test Method Environmental Conditions — Junction-to-Board."

[6] B. Guenin, "Determining the Junction Temperature in a Plastic Semiconductor Package, Part I," ElectronicsCooling, Vol. 5, No. 2, May, 1999.

[7] B. Siegal, "Thermal Load Board Design Considerations," ElectronicsCooling, Vol. 18, No. 3, September, 2009.

[8] B. Guenin, "Characterizing a Package on a Populated Printed Circuit Board," ElectronicsCooling, Vol. 7, No. 2, May, 2001.

[9] S. Shidore, "Compact Thermal Modeling in Electronics Design" ElectronicsCooling, Vol. 13, No. 2, May, 2007.

[10] B. Joiner "Use of Junction-to-Board Thermal Resistance in Predictive Engineering," Electronics Cooling, Vol. 5, No. 1, January, 1999.

Keep Your Optics Clear, Gap Filler 1500LV Offers Over 10x Reduction In Outgassing.

Headlight assembly with a standard thermal interface material

Headlight assembly with new Bergquist Gap Filler 1500LV

Bergquist Gap Filler 1500LV supports greater clarity inside sensitive electronic applications where fogging of lenses or optics from outgassing is a consideration.

Control outgassing, improve thermal performance for a clearer environment.

Gap Filler 1500LV is a low volatility, two-component, liquid-dispensable thermal interface material. This ultra conformable elastomer provides a thermal conductivity of 1.8 W/m-K. It offers the high temperature resistance and low modulus of a silicone material, with significantly lower levels of outgassing. At higher temperatures, volatiles from outgassing can create issues in enclosed fixtures such as lighting where the clarity is crucial. Gap Filler 1500LV is ideal for use in these lighting applications or other situations where fogging of lenses or optics is a potential consideration.

Low assembly stress on electronic components.

Gap Filler 1500LV is thixotropic and although it will remain in place after dispensing, the material will flow easily under minimal pressure resulting in little to no stress on fragile components during assembly.

Gap Filler 1500LV cures at room temperature – and the curing process can be accelerated with the addition of heat. When cured, this material provides a soft, thermally conductive, form-in-place elastomer that is ideal for filling unique and intricate air voids and gaps.

Another innovative thermal solution for optimized dispensing.

Gap Filler 1500LV is just one of an expanding line of liquid dispensed materials. Unlike precured gap filling materials, liquid dispensed materials offer infinite thickness options and eliminate the need for specific pad thicknesses or die-cut shapes for individual applications. Applying precise amounts of material directly to the target surface results in an effective use of material with minimal waste.

To qualify for your FREE Gap Filler 1500LV Kit, call or visit our website: www.bergquistcompany.com/nofog

For manual or automated dispensing.

www.bergquistcompany.com 1.800.347.4572

952.835.2322 fax 952.835.0430

18930 West 78th Street • Chanhassen, Minnesota 55317

THE BERGQUIST COMPANY

Thermal Materials • Thermal Substrates • Fans and Blowers

Thermal Management of Many-Core Processors using Power Multiplexing

Man Prakash Gupta, Satish Kumar
G. W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology

INTRODUCTION

DUE TO THE growing demands of higher performance and faster computing, the number of cores in a microprocessor chip has been increasing consistently. The transition from single core to multi-core technology has already been observed in the past few years and with the strong potential of parallel computing, the transition from multi-core to many-core is also imminent where the number of cores on a single chip is expected to reach in hundreds or even thousands per single processor die. Such large-scale integration and very high power densities on chip will bring a significant challenge of heat dissipation. The traditional air-cooling methods begin to reach their flow and acoustic limits for very high power density (~1.5 W/mm^2) apart from being inefficient from economic point of view when applied to many-core technology [1, 2]. Moreover, the uneven workload on the cores leads to spatiotemporal non-uniformity in the thermal field on chip which can be detrimental to its performance and reliability [3]. The leakage power also increases exponentially with temperature resulting in higher power dissipation, and cooling costs [4, 5].

Another way to obtain a uniform on-chip temperature distribution and lower peak temperature is efficient redistribution of heat within the chip which can help to improve the energy efficiency and coefficient of performance (~compute/cooling power). This brings new opportunities for the dynamic thermal management (DTM) techniques, and their role to address the challenges of power dissipation in many-core processors becomes very important. Many DTM techniques have been explored such as clock gating, dynamic voltage and frequency scaling, and thread migration for single and multi-core processors [6-9]. All these reactive methods can have power and performance overhead apart from the hardware and software implications.

Power multiplexing which is a proactive method can be utilized as a supplementary approach to the reactive methods for effective thermal management of many-core processors [10, 11]. Power multiplexing technique involves redistribution (or migration) of the workload of the cores in the chip at regular time intervals to control the thermal profile on the chip. This approach is different from the reactive DTM techniques which wait for the temperature to increase beyond a certain threshold value. The idea is to improve the thermal profile by using idle or underutilized cores efficiently. The guiding rule which governs the redistribution of workload is called migration policy. The time interval at which this migration takes place is referred to as timeslice. A smaller timeslice corresponds to faster multiplexing. The value for the timeslice is typically chosen such that it is smaller than the characteristics thermal time constant (τ) of the system. In the present case, this time constant

Man Prakash is a Ph.D. candidate in the department of Mechanical Engineering at Georgia Tech. He received his B.Tech. - M.Tech. dual degree in mechanical engineering from Indian Institute of Technology Kanpur, India in May 2009. His current research work is related with thermal management of microelectronics devices and electro-thermal transport in carbon nanotube based composites.

Dr. Satish Kumar joined Georgia Tech in 2009 as an Assistant Professor in the George W. Woodruff School of Mechanical Engineering. Prior to joining Georgia Tech, he worked at IBM Corporation where he was responsible for the thermal management of electronic devices. Kumar received his Ph.D. in mechanical engineering and M.S. degree in electrical and computer Engineering from Purdue University, West Lafayette in 2007. He received his M.S. degree in mechanical engineering from Louisiana State University, Baton Rouge in 2003 and B.Tech. degree in mechanical engineering from the Indian Institute of Technology, Guwahati in 2001. His research interests are in thermal management, thermo-electric coolers, electro-thermal transport in carbon nanotube, graphene and 2D materials based flexible-electronics. He is author or co-author of over 65 journal or conference publications.

Find solutions at interferencetechnology.com

ITEM™ interference technology
THE INTERNATIONAL JOURNAL OF ELECTROMAGNETIC COMPATIBILITY

τ is defined as the time for the chip peak temperature to reach 63% of the steady-state after turning on the power under flow conditions described below. The value of τ is estimated to be 0.1 s. This criterion for the timeslice selection is based on the requirement that the 2D effects of power multiplexing need to be realized faster than the 3D thermal diffusion in order to get full advantage of multiplexing. A tile-type homogeneous 256-core processor is considered where the cores are arranged in a 16x16 2D array [12]. The power dissipation value has been selected based on the prediction by International Technology Roadmap of Semiconductor (ITRS) for 16 nm node technology. The model considers 2 W of power dissipation in each active core which is reasonable for cores with 16 nm node technology running at 3 GHz. The total power dissipation on the chip is considered to be 128 W, i.e., at one instant only 25% cores (~64 cores) are active.

Three migration policies namely, random, cyclic and global policies are explored here (Figure 1). Random policy involves random redistribution of all active cores at each timeslice. In cyclic policy active cores are assigned in a checkerboard configuration and shifted in a circular fashion at each timeslice maintaining checkerboard configuration. Global policy involves the swapping of workload between hottest and coolest cores at regular time intervals.

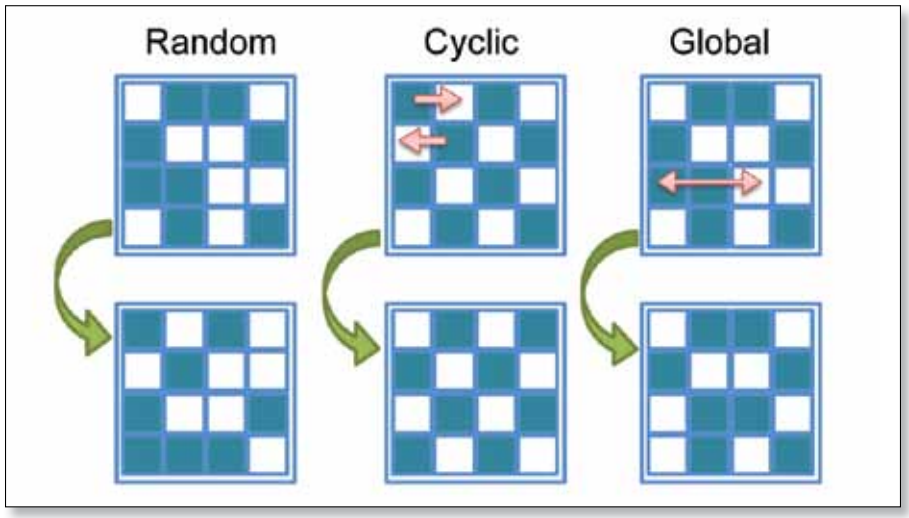

FIGURE 1: Illustration of the migration policies for power multiplexing on many-core processors. Random multiplexing involves arbitrary exchange of workloads among all cores at regular time intervals. Cyclic multiplexing policy preserves checkerboard configuration during multiplexing. Global policy involves exchange of workload between hottest and coolest cores.

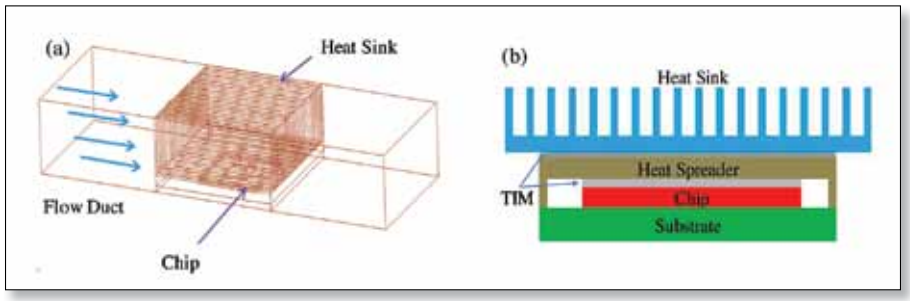

FIGURE 2: (a) Flow duct with a heat sink and an electronic package used for the thermal modeling. **(b)** Schematic of the heat sink and electronic package of the many-core processor which include heat spreader, thermal interface material (TIM), chip and substrate (view along the direction of inlet flow).

TABLE 1: MATERIAL PROPERTIES OF THE COMPONENTS OF CHIP PACKAGE

Component	Material	ρ (kg/m³)	Cp (J/kgK)	Value (W/mK)
Heat Sink	copper	8978	381	387.6
Heat Spreader	aluminum	2719	871	202.4
TIM	grease	2550	700	4
Chip	silicon	2330	712	141.2

METHODOLOGY AND RESULTS

Using computational fluid dynamics (CFD), a detailed heat transfer analysis of the electronic package is performed. The computational domain is comprised of a flow duct, a heat sink, a heat spreader, the thermal interface material (TIM), a chip and a substrate (Figure 2) [12]. The properties of the various components of the system are considered to be constant and are listed in Table 1. It should be noted that temperature dependent thermal conductivity of the components does not cause any significant change in the results since the temperature variation is between 300 and 330 K only. The dimensions of chip are 12 mm x 0.5 mm x 12 mm and the typical size of a grid cell inside chip is 0.375 mm x 0.1mm x 0.375mm. A uniform velocity profile at the inlet of the air flow tunnel is considered with constant velocity of 5 m/s. An outflow boundary condition is imposed at the outlet of the tunnel and no-slip boundary condition is imposed at the walls of the tunnel and outer surfaces of the electronic package (Figure 2 (a)). The flow inside the tunnel is turbulent as Reynolds number based on the inlet flow rate and duct hydraulic diameter width is 20,000. As accurate turbulent flow computations are not critical in

the present study, Spalart-Allmaras turbulence model [13] was used, which is a simple one-equation model and appropriate for applications involving wall-bounded flows and for avoiding fine meshing near the wall. We consider SIMPLE scheme for pressure-velocity coupling, implicit scheme for transient formulation and second order upwind scheme for the discretization of all governing equations [14].

Three cases (slow, fast or no multiplexing) are investigated corresponding to each migration policy to examine the effect of timeslice variation. For random power multiplexing, results suggest faster multiplexing (at timeslice = 0.0033 τ which equals to 10^6 clock cycles) provides 10°C reduction in the peak temperature (T_{max}) and 15°C reduction in the maximum spatial temperature difference (T_{max}-T_{min}) [12]. A graphic comparison of the thermal

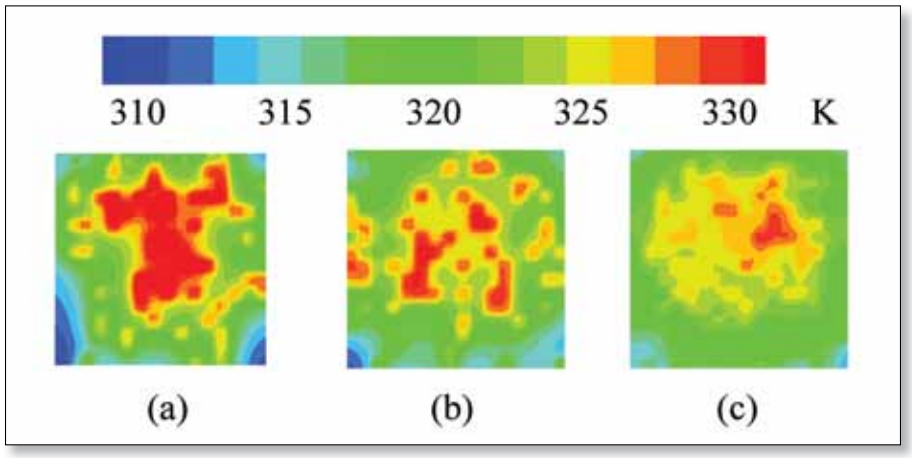

FIGURE 3: Thermal profile on 256 core chip at time instant, t/τ = 6.6, for **(a)** no multiplexing, **(b)** slow multiplexing with timeslice = 0.033τ (10^7 clock cycles), and **(c)** fast multiplexing with timeslice = 0.0033τ (10^6 clock cycles) with random core migration policy. 25% active cores with total power = 128W.

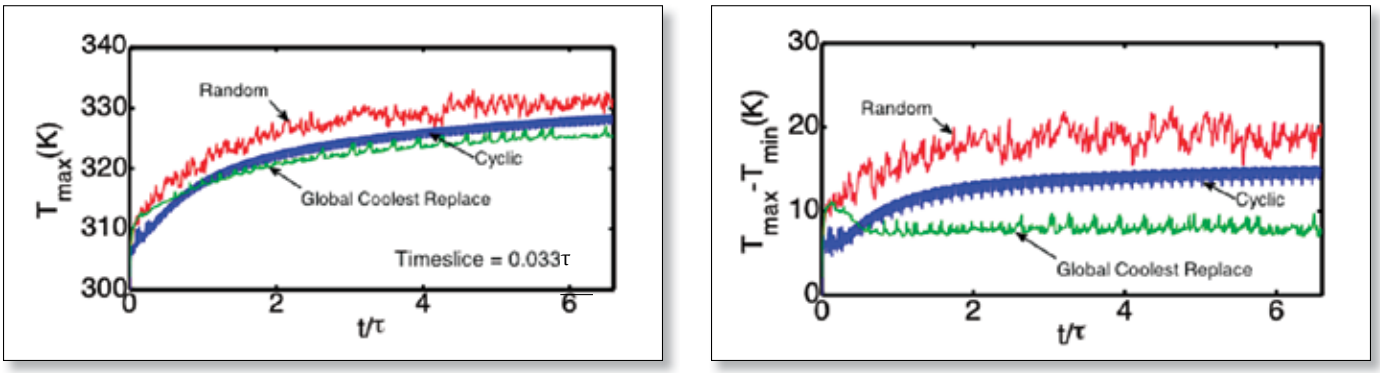

FIGURE 4: Comparison of the effect of different migration policies on **(a)** peak temperature, **(b)** spatial temperature difference. Timeslice is kept as 0.033τ during power multiplexing. 25% cores are considered to be active with total power = 128W.

profile on the chip at time instant, t = 6.6 τ, is shown in Figure 3.

For cyclic policy, results indicate that it reduces the peak temperature by only 3°C even for vary fast multiplexing. This small reduction can be attributed to the pre-existing checkerboard configuration of active cores. The maximum spatial temperature difference across the chip is however significantly reduced (by 7°C).

Global policy is intrinsically different from the previous two policies as it takes decisions based on the instantaneous chip temperature and also, fewer cores are involved in the multiplexing. To begin with, only a pair of cores is considered for the global multiplexing, *i.e.*, the workload is swapped between the hottest and the coolest core at each timeslice. It is found that global policy shows significant improvement in thermal profile even for very slow multiplexing. Analysis of the power map at each migration step, finds that the global coolest policy ingeniously places the active cores away from the center of the chip such that it not only reduces peak temperature by a significant amount but also reduces thermal non-uniformity. By comparison of the three policies, it is found that cyclic policy shows better performance compared to random policy but global policy outperforms the other two in terms of higher peak temperature reduction and better thermal uniformity on the chip (Figure 4). A graphic comparison of the thermal profile on chip can be seen in Figure 5. It should be noted that increasing the number of cores involved in the swapping of workload during the global policy does not bring any significant improvement in the thermal profile of the chip [12]. Thus, the results advocate the strength of global policy as it requires swapping of workload on only a pair of cores and even slow multiplexing can get higher reduction in the peak temperature and uniformity in the temperature profile.

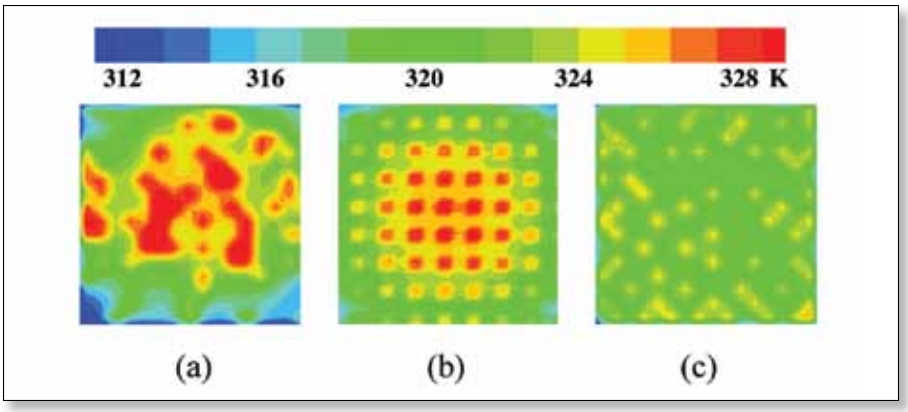

FIGURE 5: Thermal profile on a chip at t/τ = 6.6 for **(a)** random, **(b)** cyclic, and **(c)** global coolest replace policies. Timeslice is taken as 0.033τ. Very high spatial thermal uniformity can be seen for the global multiplexing. 25% active cores with total power = 128W.

SUMMARY

Power multiplexing approach has been presented as a prospective thermal management technique for many-core processors. The global power multiplexing has been found to be the most effective among the three policies discussed in this article. The peak temperature reduction of 10°C and the maximum spatial temperature difference reduction of 20°C have been observed on a 256-core chip using global policy based power multiplexing. This can be attributed to its inherent approach to optimize the proximity of active cores on a finite size chip by automatically considering the effect of geometrical and thermal properties of the 3D system through the temperature distribution at each migration step. The work presented in this article may be considered as a first order analysis of migration policies as simple policies are applied in case of the homogeneous many-core processors. More evolved policies can be formulated to handle thermal management of heterogeneous many-core processors.

REFERENCES

[1] J. Michael J. Ellsworth, "High Powered Chip Cooling - Air and Beyond," Electronics Cooling, 2005.

[2] P. Zhou, J. Hom, G. Upadhya, K. Goodson, and M. Munch, "Electro-kinetic microchannel cooling system for desktop computers," Twentieth Annual IEEE Semiconductor Thermal Measurement and Management Symposium, Proceedings 2004, pp. 26-29, 2004.

[3] R. Mukherjee and S. O. Memik, "Physical Aware Frequency Selection for Dynamic Thermal Management in Multi-Core Systems," in IEEE/ACM International Conference on Computer-Aided Design (ICCAD). , 2006, pp. 547-552.

[4] M. Janicki, J. H. Collet, A. Louri, and A. Napieralski, "Hot spots and core-to-core thermal coupling in future multi-core architectures," in 26th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM). , 2010, pp. 205-210.

[5] E. Kursun and C. Chen-Yong, "Temperature Variation Characterization and Thermal Management of Multicore Architectures," IEEE Micro., vol. 29, pp. 116-126, 2009.

[6] D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," in The Seventh International Symposium on High-Performance Computer Architecture (HPCA). , 2001, pp. 171-182.

[7] P. Chaparro, J. Gonzalez, G. Magklis, Q. Cai, and A. Gonzalez, "Understanding the Thermal Implications of Multi-Core Architectures," IEEE Transactions on Parallel and Distributed Systems, vol. 18, pp. 1055-1065, 2007.

[8] V. Hanumaiah, S. Vrudhula, and K. S. Chatha, "Maximizing performance of thermally constrained multi-core processors by dynamic voltage and frequency control," in IEEE/ACM International Conference on Computer-Aided Design - Digest of Technical Papers (ICCAD). , 2009, pp. 310-313.

[9] R. Rao, S. Vrudhula, and C. Chakrabarti, "Throughput of multi-core processors under thermal constraints," in ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED). , 2007, pp. 201-206.

[10] C. Y. Cher and E. Kursun, "Exploring the Effects of On-Chip Thermal Variation on High-Performance Multicore Architectures," ACM Transactions on Architecture and Code Optimization, vol. 8, Apr 2011.

[11] M. K. Cho, C. Kersey, M. P. Gupta, N. Sathe, S. Kumar, S. Yalamanchili, and S. Mukhopadhyay, "Power Multiplexing for Thermal Field Management in Many-Core Processors," IEEE Transactions on Components Packaging and Manufacturing Technology, vol. 3, pp. 94-104, Jan 2013.

[12] M. P. Gupta, M. K. Cho, S. Mukhopadhyay, and S. Kumar, "Thermal Investigation Into Power Multiplexing for Homogeneous Many-Core Processors," Journal of Heat Transfer-Transactions of the ASME, vol. 134, Jun 2012.

[13] P. Spalart and S. Allmaras, "A one-equation turbulence model for aerodynamic flows," American Institute of Aeronautics and Astronautics, vol. Technical Report AIAA-92-0439 1992.

[14] S. V. Patankar, Numerical heat transfer and fluid flow / Suhas V. Patankar. Washington : New York :: Hemisphere Pub. Corp. ; McGraw-Hill, 1980.

Characterization of Server Thermal Mass

Mahmoud Ibrahim
Panduit

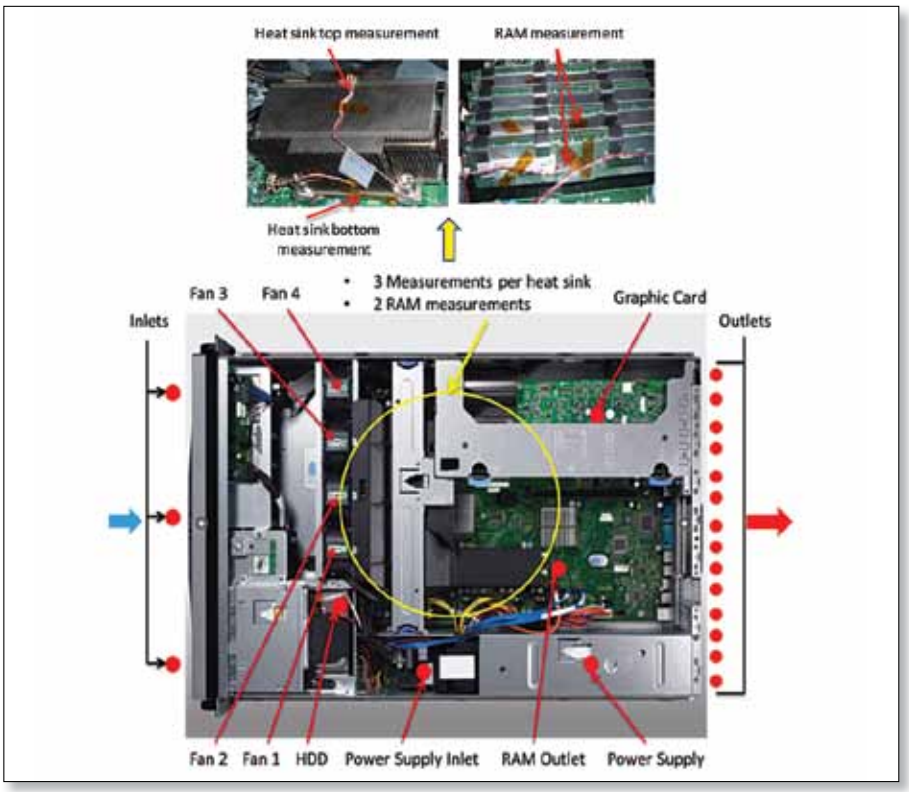

FIGURE 1: Server temperature measurement, with thermocouple locations highlighted by the red dots.

Mahmoud Ibrahim is a senior analysis engineer at Panduit Corporation, working on advanced energy-efficient cooling controls and technologies for data centers. He received his PhD in Mechanical Engineering from Binghamton University in 2012. He is a member of ASME, and ASHRAE. He has authored and co-authored over 20 technical publications in journals, conferences, and book chapters.

THE CONTINUOUS growth of internet business and social media is necessitating the construction of new data centers to support the rising volume of internet traffic. This is in addition to the increasing number of colocation facilities being built to support different business sizes with different functions. Generally, data centers are considered mission critical facilities that house a large number of electronic equipment, typically servers, switches, and routers that demand high levels of electrical power which eventually dissipates as heat. These facilities require a tightly controlled environment to ensure the reliable operation of the IT equipment. In many cases, with the IT equipment being a core player in the customer-business relationship regardless of the sole function of the business (healthcare, finance, social media, etc.), facility managers and data center operators tend to unnecessarily overcool their data centers fearing the presence of hotspots within the facility. This explains why in the 2007 EPA report to Congress it was estimated that U.S. data centers used 61 billion kWh of electricity in 2006, representing 1.5% of all U.S. electricity consumption and double the amount consumed in 2000 [1]. Consequently, data center cooling technology has gained considerable research focus over the past few years in an effort to lower these high levels of energy consumption.

Dynamic cooling has been proposed as one approach for enhancing the energy efficiency of data center facilities. It involves using sensors to monitor continuously the data center environment and making real time decisions on how to allocate the cooling resources based on the location of hotspots and concentration of workloads. In order to effectively implement this approach, it is vital to know the transient thermal response of the various systems comprising the data center, which is a function of thermal mass.

Another crucial requirement for mission critical facilities, in addition to providing adequate cooling to IT equipment, is securing constant electrical power supply. A survey conducted by Ponemon Institute on data center outages [2] stated that 88% out of 453 surveyed data center operators experienced a loss of primary utility power and hence the power outage of their data centers, with an average of 5.12 complete data center outages every two years. When an outage occurs, it is standard practice to continue to furnish power to the servers using uninterruptable power supplies (UPSs) for a period of time. This allows the fans in the server to continue to operate while allowing sufficient time for the server to attempt a controlled shutdown. However, without active cooling the ambient air temperature within the data center will rise. How quickly the data center environment reaches the critical limit of the IT equipment is once again a function of thermal mass. A number of studies have addressed this issue and developed simple energy based models to predict the data center room temperature rise during power failure [3-4]. While such models may provide useful information, they are based on combining all parameters in the data center into one, so as to provide an overall sense of the thermal response. These models however are not capable of providing a detailed view of the data center whether during dynamic cooling control or during power failure. In order to do so, the thermal mass of the various objects present in the data center must be determined.

This article introduces an experimental technique used to extract the thermal mass of servers, which can further be used as a compact model embedded in Computational Fluid Dynamics (CFD) simulations for data center level analyses, or analytical based thermal models. All the experimental testing and results presented are conducted on a typical 2U server used in the market.

APPROACH

In order to obtain a thermal mass of the server, the following energy balance equation is applied:

$$\dot{Q} = (\dot{m}\, c_p)_{air}\, \Delta T_{air} + (pVc_p)_{unit} \frac{\partial T_{unit}}{\partial T} \quad \textbf{(1)}$$

where $\dot{Q}$ is the total energy generated by the server per unit time, $\dot{m}$ is the air mass flow rate within the server, and $(pVc_p)_{unit}$ is the thermal mass of the server. The first term on the right-hand side of Equation 1 represents forced convection cooling by the server fans, with the second term accounting for the server thermal storage. Measuring the air velocity inside the server, the power consumed, the air temperature rise (ΔT_{air}) across the server and a representative temperature of the server with time, one can extract the thermal mass. Although it appears straightforward since all the variables are measured, obtaining a server unit temperature T_{unit} poses difficulties that require further assumptions.

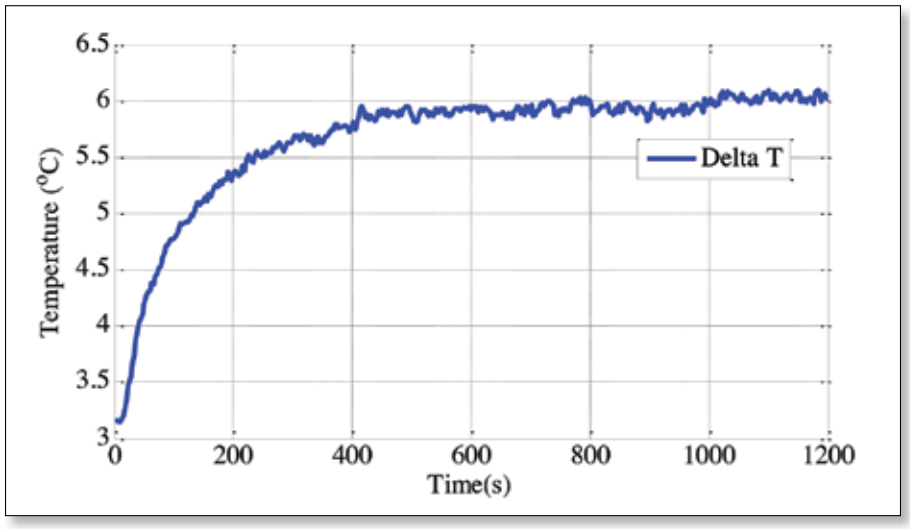

FIGURE 2: Measured air temperature rise (ΔT) across the server at $\dot{Q}$ of 350 W.

Given that the server is comprised of numerous components, obtaining an appropriate average temperature of the unit as a whole requires a weighted allocation of the individual components. The major components that may have an effect on the average unit temperature are CPU/heat sink assemblies, hard disk drives, graphic card, memory modules, fans, power supply unit, and the chassis including the motherboard. The primary materials used in the production of these various components are aluminum, steel, copper, and silicon. Knowing that the specific heat capacity of the above materials within the range of 0.5-0.9 KJ/kg.K, and that the range of heat capacity for condensed matter is 1-4 J/m^3.K [5], one can assume that the contribution of each component to the unit temperature will be solely determined by the ratio of the component weight to the overall server weight. By obtaining the weights of the various server components, the percentage contribution of each component to the overall unit temperature can be computed. The unit temperature is then calculated by multiplying each component's percentage by its measured temperature as represented by Equation (2). The equation coefficients represent the relative mass of each component.

$$T_{unit} = 0.03\, T_{fans} + 0.05\, T_{GC} + 0.07\, T_{HDD} + 0.04\, T_{RAM} + 0.08\, T_{PS} + 0.11\, T_{CPU/HS} + 0.61\, T_{Chassis} \quad \textbf{(2)}$$

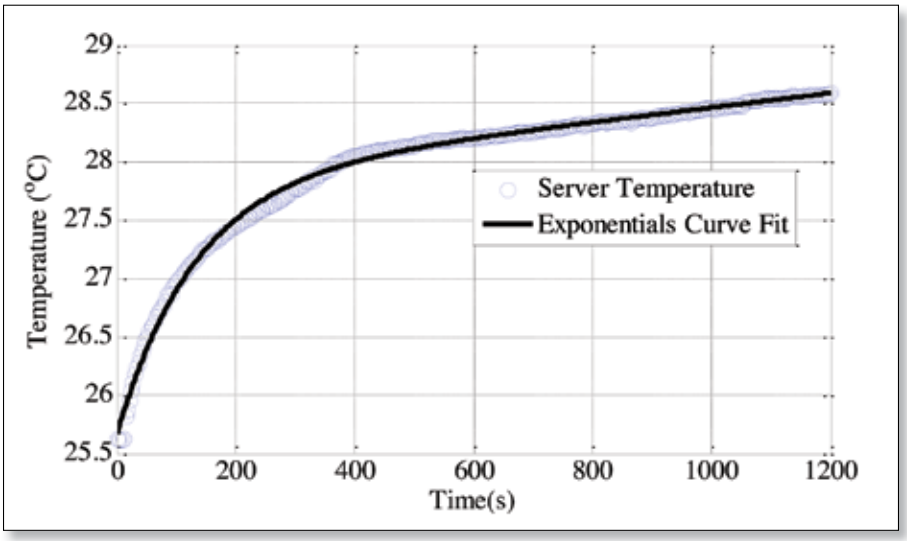

FIGURE 3: Curve fit of the calculated server transient temperature profile (T_{unit}) at $\dot{Q}$ of 350 W.

TEST SETUP

The server was located in a large room to insure a uniform inlet ambient air temperature, not affected by the server exhaust temperature. Type J thermocouples with a measurement uncertainty of ±1.1°C, and a thermal response time of less than 1 second were used for temperature measurement. The thermocouples were attached in various locations to measure temperatures of the different components within the server, as well as the air intake and exhaust temperatures (Figure 1). The thermocouples were attached using a thermal adhesive tape to obtain the best possible thermal contact. A data acquisition (DAQ) unit was used to

log the temperature data. A power meter was used to measure the power consumed by the server with an accuracy of ±1% and a response time less than 5 seconds.

In addition to the thermocouples, the server is equipped with sensors from the manufacturer, taking temperature measurements of different components, as well as measuring the four CPU fans' speed. Software was installed to log the reported fan speed measurements and temperature measurements by the various components in the server. Data was collected every three seconds, and it was synchronized between the three sources: power meter, DAQ unit, and server sensors using the clocks on each computer.

Fig. 1 shows a top view of the server, highlighting where the thermocouples were placed. Fourteen thermocouples were used to obtain an average server air outlet temperature, so as to account for the large variation in outlet temperature across the width of the server.

Through operating the server at a specific power level and measuring the corresponding temperatures, fan speeds, and power consumption, one can extract the server thermal mass using Equation (1). Various power levels were tested to ensure the validity and accuracy of the approach. All the tests were conducted in an identical manner, where the server was operated initially at idle state and the temperature measurements were used to determine a steady state operation. Steady state condition was determined by observing that the server air exhaust temperatures as well as the component surface temperatures remained constant (±0.5°C) for at least 10 minutes. A computational load is applied to the server to run it at a specific power level causing component surface temperatures and server exhaust temperature to increase. Loads applied to the server varied from 220W to 350W while the fan airflow was kept constant. Once a steady state is reached, computations are terminated allowing the server to go back to its idle state.

RESULTS

Figures 2 and 3 show one of the experimental runs used to determine the thermal mass of the server. The server power consumption ($\dot{Q}$) is constant at 350 W, while the fans are operating at maximum speed of 12,000 RPM. The flow rate ($\dot{m}$) at this speed is about 6.61E-2 kg/s. In Figure 2, the measured air temperature rise across the server (ΔT_{air}) as a function of time is shown. The cor-

responding increase in the average server temperature (T_{unit}) is shown in Figure 3, which is curve fitted using an exponential function. The component temperatures ranged from 27 °C for the chassis to 75 °C for the CPU. By finding the derivative of the function with respect to time, along with the power, flow rate, and temperature difference across the server, the server thermal mass is obtained using the method of least squares. For this particular run, the computed thermal mass was 11.3 kJ/K.

The summary of all runs is shown in Fig. 4 with an average thermal mass value of 11.1 kJ/K. The corresponding average percent error for each run compared to the average thermal mass value is ±7%. The range of server power tested was 220 W to 350 W, which is a typical range for a 2U server, with the server power at idle been 150 W. The server thermal mass is expected to be similar for a wider range of powers as it is a characteristic of the server.

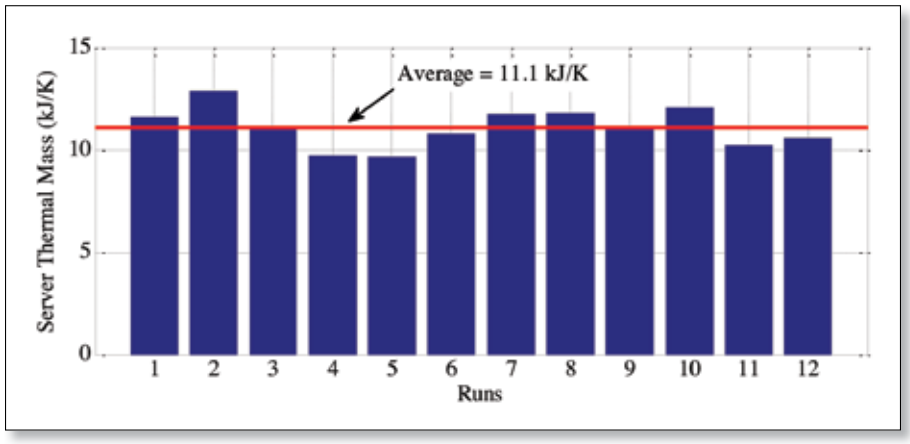

FIGURE 4: Calculated server thermal mass obtained from 12 experimental runs.

With the obtained server thermal mass, the server can be treated as a black box in which for a given scenario with specified inlet air flow rate and temperature, and server power consumption, the server outlet air temperature can be obtained with time. For data center operators this is primarily the information they would require, rather than being concerned with the details of the temperature distribution within the server. This model can also be used to construct a CFD model that represents a server behaving similarly to the 2U server tested here. The CFD model would approximately account for the thermal mass of the server, allowing a transient investigation. A room level CFD model can be populated with these simple server models and the transient thermal response of the data center can be investigated for different scenarios. Although not evaluated yet, cases such as power shutdown or cooling equipment failure can be modeled in CFD to account for the transient response of the IT equipment.

REFERENCES

[1] U.S. Environmental Protection Agency ENERGY STAR Program, 2007, "Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431," http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.

[2] Ponemon Institute, 2010, "National Survey on Data Center Outages," Retrieved from http://www.ponemon.org/local/upload/file/Data_Center_Outages.pdf.

[3] Khankari, 2010, K., "Thermal Mass Availability for Cooling Data Centers During Power Failure Shutdown," ASHRAE Transactions, 116(2), 205-217.

[4] Zhang, X., and VanGilder, J., 2011, "Real-Time Data Center Transient Analysis", ASME 2011 Pacific Rim Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Systems, MEMS and NEMS: Volume 2, pp. 471-477.

[5] Wilson, J., 2007, "Thermal Diffusivity," ElectronicsCooling Magazine, 13(3), http://www.electronics-cooling.com/2007/08/thermal-diffusivity/.

FloTHERM XT–
BRIDGING EDA AND MCAD DESIGN DOMAINS

FROM CONCEPT TO VERIFICATON | Mentor Graphics latest thermal simulation software, FloTHERM XT, is designed for the complex mechanical design environments and processes in today's electronics companies. FloTHERM XT enables supply-chain integration and fits anywhere in the mechanical design flow, from design concept to final product verification. Check out this next generation thermal simulation software at: www.mentor.com/flotherm-xt

Mentor Graphics®
THE EDA TECHNOLOGY LEADER

30th Anniversary
The 30th Annual Thermal Measurement, Modeling & Management Symposium

SEMI-THERM 30 THIRTY

SEMI-THERM®

THERMAL INNOVATIONS THAT MAKE THE WORLD'S TECHNOLOGY COOL...

ADVANCE PROGRAM

SEMI-THERM 30 is coming soon! The symposium will feature critical information for thermal and non-thermal professionals. Learn about and share advances in thermal management, measurement and modeling.

Register Now
www.SEMI-THERM.org

SEMI-THERM 30
March 9-13, 2014
San Jose, California

For program details, registration, exhibition and hotel information visit www.semi-therm.org today!

Doubletree by Hilton
San Jose, CA USA

- **Pre-conference short-courses** from world-class thermal experts
- **Networking opportunities** with industry leaders & innovators
- **Over 55 peer-reviewed papers** presented by the best and brightest thermal professionals and educators.
- **Technical Sessions** include 3-D Electronics, Consumer Electronics, Computation Fluid Dynamics, Data Centers, Harsh Environments, High Heat Flux Cooling, LEDs, Measurements and Characterization, New Materials, Thermo-mechanical modeling/characterization and more…
- **New for 2014!**
  - Welcome reception
  - Product teardowns
  - Problem solving sessions
  - "How-to" courses
- Program includes **embedded tutorial**, **evening events** and **special luncheon speakers**
- **THERMI** and **Harvey Rosten Award** presentations
- Two-day SEMI-THERM Exhibition with **over 40 vendors and vendor workshops** presenting the most recent technical information
- **Complimentary receptions** Monday and Wednesday evenings, take advantage of these opportunities to network with thermal experts
- Convenient location **in the heart of Silicon Valley**, minutes from San Jose International Airport

Platinum Sponsors:

Mentor Graphics
Mechanical Analysis Division

future facilities

Media Sponsor:

Electronics COOLING

Check out the information on the next few pages and visit **www.semi-therm.org** for registration and program details and to see other events coming to SEMI-THERM 2014!

THANK YOU TO OUR SEMI-THERM 30 PROGRAM COMMITTEE

General Chair	*Genevieve Martin*, Philips Research
Program Chair	*Peter Rodgers, Ph.D.*, The Petroleum Institute
Vice-Program Chair	*Rahima Mohammed, Ph.D.*, Intel

SEMI-THERM 30 SCHEDULE

SUNDAY
MARCH 9, 2014

- Short Courses

MONDAY
MARCH 10, 2014

- Short Courses
- Welcome Reception

TUESDAY
MARCH 11, 2014

- Opening Ceremony
- Keynote
- Technical Sessions
- Luncheon Speaker
- Exhibition
- Vendor Workshops
- Problem Solving Sessions

WEDNESDAY
MARCH 12, 2014

- Embedded Tutorial
- Technical Sessions
- Exhibition
- Vendor Workshops
- How to Courses
- Product Teardowns
- Networking Reception

THURSDAY
MARCH 13, 2014

- Technical Sessions
- Award Luncheon

EXHIBITION

Don't miss the SEMI-THERM 30 exhibits!

Entrance is free.

Exhibit space is available.

semi-therm.org

SEMI-THERM TECHNICAL SESSIONS

SEMI-THERM 30 presents 3 days of technical papers on the latest advances in thermal management. Sessions run 8AM to 12:30PM on Tuesday, Wednesday and Thursday.
Registration is open now. **Visit www.semi-therm.org**

♦ 3-D Electronics

3D electronics is gaining considerable momentum to be the next generation solution for the electronics industry due to smaller foot print and shorter interconnects between multiple devices. In addition, partitioning a large chip into multiple smaller chips with 3D stacking can improve yield and the fabrication cost. Due to the complexity of 3D electronics, the thermal characterization of its packaging, for one thing, becomes more challenging than the standard single or multi-chip packages. This session will consider such challenges for both thermal measurement and modeling of 3D electronics.

♦ Consumer Electronics

This session will explore the challenges in thermal management of consumer electronics, i.e. electronic equipment used everyday in entertainment, communications and productivity, such as audio equipment, TVs, game consoles, mobile phones, laptops, tablets and wearable electronics. Topics of interest include system level thermal design and cooling; integrated thermal solutions combining thermal and non-thermal design functions; specific consumer electronics concerns such as extreme thin and slim product shapes (e.g., displays), touch temperatures of portable and wearable electronic devices, and thermal management of chargeable batteries.

♦ Computation Fluid Dynamics

This session will highlight the role of CFD software in the design and thermal management of electronics, from advanced component to system level applications. Emphasis will be placed on understanding the fundamentals of CFD simulation, and best practices in modeling.

♦ Data Centers

This session will present best practices in data center thermal management, including hot/cold aisle containment, VFDs, and floor layouts. The use of thermal performance and energy efficiency metrics to evaluate data center performance will be highlighted based on practical case studies. Industry trends such as innovative/non-conventional cooling schemes, electrical distribution, energy sources, and intelligent workload scheduling will be considered.

♦ Harsh Environments

This session will explore the thermal management challenges of electronics operated in harsh environments. Applications include defense and aerospace equipment, under-hood automotive and other transportation systems, and outdoor equipment. Topics include system level cooling in high ambient temperatures and/or low pressures; integrated designs to address thermal management as well as severe environmental effects such as corrosive environments, extreme pressures, and high shock/vibration.

♦ High Heat Flux Cooling

In this session, advanced thermal management concepts and heat transfer aspects of electronic systems including two-phase heat transfer will be featured.

♦ Light Emitting Diodes (LEDs)

Billions of LEDs are being made every year. This number is increasing rapidly and soon most new lighting or light equipped devices (such as those found in healthcare and displays) will be made from LEDs. Thermal management is critical to the lifetime and performance of LED systems. This session will address the many challenges related to design aspects, characterization and reliability. The session will present novel (noiseless) cooling technologies, prediction of LED junction temperature, thermal management of active LED applications, implementation of multi-domain (optical-thermal) LED model, reliability of LEDs, measurement methods.

♦ Measurements and Characterization

This session addresses the thermal metrological needs of the microelectronics field, such as the characterization of the thermal properties of thin-film materials and the temperature of microscale features, both accessible and embedded.

♦ New Materials

This session provides a forum to address how advancements in materials research can significantly enhance electronics thermal performance. A variety of packaging materials used in the design of semiconductor packaging across the spectrum of electronics applications will be highlighted. Materials include organic compounds used as thermal interface materials, metal alloys used as package stiffeners and heat spreaders, and solders for mechanical joining and electrical connection.

♦ Thermo-Mechanical Modeling/Characterization

This session will focus on the thermo-mechanical modeling and characterization various semiconductor applications and packaging technologies, including silicon based microprocessors, memory devices, Gallium based components, power components, and FETs.

♦ And more

SEMI-THERM TECHNICAL SHORT COURSES

Highly informative short courses are being offered **Sunday March 9 - Monday March 10, 2014**. These reasonably priced and outstanding technical courses are focused on the thermal sciences and are presented by leaders in the thermal field.
Registration is open now! Visit **www.semi-therm.org**

SHORT COURSE 1 SUNDAY MARCH 9, 2014, 1PM-5:30PM

Turn On the Power of Your Desktop for Complicated Thermal Transient Problems Using Microsoft Excel
David Billings, P.E., ON Semiconductor Corporate R&D

Course Description
The course will teach how to use some of the basic tools in Excel to do complex transient thermal analysis and Design of Experiments response analysis. Having a simplified spreadsheet tool that can be shared with colleagues or customers allows them to play with some of the variables. This provides better intuition of the mechanics of the problem which can optimize or improve the processes that are involved with product development.

The attendees will be instructed in:

- The process of converting transient thermal data to useful Foster or Cauer compact thermal networks
- How to combine multiple time varying heat sources using linear superposition
- How to fit multi-variable Design of Experiment data to a response surface equation
- Using Microsoft® Excel engineering tools to solve and analyze the above problems
- Making calculators to be shared with colleagues or customers

About the Instructor

David Billings has been a packaging engineer and thermal/mechanical analyst for Motorola for 13 years and ON Semiconductor for 15 years. David has over 2000 Excel spreadsheet analysis tools that are used by internal and external customers to help them solve steady state or transient thermal problems.

SHORT COURSE 2 SUNDAY, MARCH 9, 2014 1PM-5:30PM

Thermal and Electrical Modeling of Three-Dimensional Integrated Circuits (3D ICs)
Ankur Jain, Assistant Professor, Mechanical and Aerospace Engineering Department, University of Texas at Arlington
Ioannis Savidis, Assistant Professor, Electrical and Computer Engineering Department, Drexel University

Course Description
While much work has been carried out in implementation of 3D IC technology in future microelectronic devices, thermal management of 3D ICs continues to remain a critical technological challenge. It is important to view thermal management requirements in conjunction with electrical modeling of 3D ICs since thermal and electrical design of 3D ICs are closely coupled with each other, and a concurrent co-design will help maximize possible benefits of this technology. This tutorial will present an overview of 3D ICs and discuss specific thermal management and electrical modeling issues.

About the Instructors

Dr. Ankur Jain is an Assistant Professor in the Mechanical and Aerospace Engineering Department at the University of Texas, where he directs the Microscale Thermophysics Laboratory. Research includes thermal characterization of 3D ICs, microscale thermal transport, and bioheat transfer. Jain has held R&D positions in leading semiconductor firms including AMD and Freescale.

Dr. Ioannis Savidis is an Assistant Professor in the Electrical and Computer Engineering Department at Drexel University. His research interests include analysis, modeling, and design methodologies for high performance digital and mixed-signal ICs, with an emphasis on electrical and thermal modeling and characterization, signal and power integrity, and power and clock delivery for 3-D IC technologies.

SHORT COURSE 3 MONDAY, MARCH 10, 2014, 8AM-12PM

Noise Emission and Control: Fundamentals
David Nelson, Principal Consultant, Nelson Acoustics

Course Description
Increasing power density of fan-cooled equipment, coupled with decreasing public tolerance for noise, make it increasingly important to understand the role that noise plays in product acceptance. Important noise criteria are being added, refined, and/or tightened. This half-day course will focus on a comprehensive introduction to acoustics and noise control terminology and concepts. It will also provide insight into testing methods, and basic troubleshooting.

SHORT COURSE 4 MONDAY, MARCH 10, 2014, 1PM-5:30PM

Noise of Cooling Fans: Quiet by Design
David Nelson, Principal Consultant, Nelson Acoustics

Course Description
This course is aimed at thermal engineers and design engineers involved in the design, development, optimization, and testing of electronic products. Those involved in development of any fan-cooled equipment, from small desktop pieces to large telecom servers, will benefit from the insight provided.

The course assumes familiarity with basic acoustical concepts and terminology such as that obtained in the companion seminar "Noise Emission and Control : Fundamentals".

About the Instructor

David Nelson is Principal Consultant of Nelson Acoustics, an Austin, TX firm that specializes in noise control using quiet-by-design methods. David studied acoustical engineering at MIT (BS '81) and the University of Texas, Austin (MSE '84). He is Board Certified, Institute of Noise Control Engineering, a professional engineer in Oregon and Texas and a member of Acoustical Society of America. Mr. Nelson has presented more than 30 training sessions to corporate, university, and government audiences on subjects related to acoustics and noise control.

SHORT COURSE 5 MONDAY, MARCH 10, 2014 8AM-5:30PM

Fundamentals of Thermal Systems Design for Electronic Products
Wendy Luiten, Philips Research, the Netherlands

Course Description
At system level, thermal functioning is a result of electrical design determining heat production, mechanical design determining cooling capabilities, and performance requirements ultimately determining allowed temperatures. This 1-day short course will focus on early system level thermal design and trouble shooting, answering the typical early-on thermal questions:

- Is this box large enough to handle x watts?
- Do we need a fan ?
- Will this sink do?
- Will component A heat component B ?
- Do we have a product?

About the Instructor

Wendy Luiten is a thermal specialist at Philips Research in Eindhoven, the Netherlands. She has over 25 years experience in the thermal and mechanical field, in product R & D, especially in Consumer Electronics and LED products. She is a lecturer at Electronics Cooling courses at Philips and the High Tech Institute Eindhoven. Wendy has authored and co-authored 20 papers and holds 5 patents. Luiten received a MSc in mechanical engineering (heat & fluid flow) from Twente Technical University, the Netherlands.

SEMI-THERM HOW TO COURSES

How to courses are a series of free courses developed to introduce practical knowledge of thermal issues to marketing or technical personnel who are new to the thermal management field. The courses will be presented from 4PM to 6PM in two parallel sessions on **Wednesday, March 12, 2014**, during the SEMI-THERM 30 Exhibition. Each course will last about 50 minutes. The courses are open to anyone attending the symposium and/or the exhibition. Seating will be limited so attendees should plan on coming early.

How To Select a Semiconductor Package

Presented by Tom Tarter, Package Science Services LLC

Introduction to design, materials, assembly and cost attributes to consider when selecting a package for your semiconductor device. The course provides a basic understanding of packaging and processes to help engineers in all disciplines make sound packaging decisions early in the design cycle of the chip or system

About the Instructor

This course is being taught by Tom Tarter who is the proprietor and principal engineer of Package Science Services. He was a working professional in the area of thermal management and electrical characterization of packaging structures. He spent over 16 years at Advanced Micro Devices (Sunnyvale, CA) in package characterization and left as a Senior Member of the Technical Staff. Tom has authored or co-authored over 20 published papers and numerous short courses and lectures on thermal and electrical phenomenon in microelectronic packaging and most recently in optoelectronic packaging. He has presented short courses and technical papers at conferences and technical meetings around the world. An invited lecturer and author, he has also lectured at graduate level short courses on micro- and opto-electronic packaging at UC Santa Cruz extension and San Jose State University.

How to Select Air Movers

Presented by Guy Wagner, Electronic Cooling Solutions

This course will focus on the basics of fan/blower selection and their interaction with system airflow. Topics include basic types of air movers, choosing the proper air mover for the system, determining the most efficient operating point, the effect on noise and power draw, and other applications-oriented issues.

About the Instructor

The course is being taught by Guy Wagner who is currently a director at Electronic Cooling Solutions specializing in thermal design of electronics, thermoelectric power generation and thermal design of photovoltaic (solar cell) systems. Guy holds 27 patents - most of them are for electronic cooling technology. He was Chief Scientist at Hewlett-Packard and Agilent Technologies and an engineering director at Cornice, Inc. specializing in micro-hard drives for consumer electronics applications.

How to make Temperature Measurements

Presented by Prof. Robert Moffat, Stanford University

Prof. Moffat will discuss the best practices for measuring the temperatures in electronic systems. This will include measuring air stream temperatures, component temperatures and the use of various sensors for reading the temperatures. The pros and cons of various methods will be discussed.

About the Instructor

This course is being taught by Professor Bob Moffat. Dr. Robert J. Moffat is a Professor of Mechanical Engineering (Emeritus) at Stanford University and President of Moffat Thermosciences, Inc. Prof. Moffat started his professional career in the Gas Turbine Laboratory at General Motors Research Laboratories upon graduation from the University of Michigan. After 10 years at GMR, he left for graduate studies at Stanford University, completing the requirements for Master of Science and Ph.D. in Mechanical Engineering. He was appointed Acting Associate Professor, 1966, Associate Professor, 1967, and Professor of Mechanical Engineering, 1972. He served as the Director of the Thermosciences Industrial Affiliates Program from 1967 to 1986 and as Chairman of the Thermosciences Division from 1973 to 1986.

How to select the best TIM

Presented by Dave Saums, DS&A LLC

This course will focus on the how to select the best thermal interface materials for a wide range of applications. The number of materials in the market place makes the selection of the correct material a daunting task. Dave will walk you thought the steps in narrow down the choices to the appropriate material.

About the Instructor

This course is being taught by Dave Saums who is a principle at DS&A Associates. DS&A is a thermal market strategy consulting firm focused on electronics thermal management materials, thermal components, and thermal systems. Dave has been studying new electronic materials for over 20 years and has become well-regarded as an expert in the application of new materials including TIM materials in the field of electronics packaging.

SEMI-THERM PROBLEM SOLVING SESSION

Experts address participant-generated thermal problems in real time!

This session will be held during the evening of **Tuesday, March 11, 2014.**

SEMI-THERM PRODUCT TEARDOWN SESSIONS

Product teardown sessions are open to anyone attending the symposium and/or the exhibition.

These sessions will be held during the evening of **Wednesday, March 12, 2014.**

Handheld Devices Teardown

Handheld devices are increasingly capable of running applications that used to require laptop and desktop computers. The requirement that these devices provide good performance with a smaller form factor or size presents significant challenges, especially with the limitations of passive cooling.

This session will feature a teardown of popular, commercially available tablets to illustrate the techniques that have been used to package the internal components and provide the thermal management that is necessary to keep both the external skin temperature and the component temperature within the maximum acceptable limits. Factors affecting the maximum possible power dissipation are discussed. The effects of the selection of the outer shell materials, thermal interface materials, heat spreaders and air gaps are presented using thermal models derived from the testing and disassemble of the tablets.

SEMI-THERM®

THERMAL INNOVATIONS THAT MAKE THE WORLD'S TECHNOLOGY COOL...

Register now at www.semi-therm.org

Cooling the Cloud:
Energy-Efficient Warm-Water Cooling of Servers

Milnes P. David, Robert Simons, David Graybill, Vinod Kamath, Bejoy Kochuparambil, Robert Rainey, Roger Schmidt,
IBM Systems & Technology Group, Poughkeepsie N.Y. and Raleigh N.C.
Pritish Parida, Mark Schultz, Michael Gaynes, Timothy Chainer,
IBM T. J. Watson Research Center, Yorktown Heights N.Y.

INTRODUCTION

DATA CENTERS NOW consume a significant fraction of US and worldwide energy (approximately 2% and 1.3% respectively); a trend which continues to increase, particularly in the cloud computing space [1,2]. Energy efficiency in data centers is thus a key issue, both from a business perspective, to save on operating costs and to increase capacity, and from an environmental perspective. Traditional chiller-based data centers, which dominate the market, typically have a PUE ~ 2 where PUE is Power Utilization Effectiveness, defined as the ratio of the total data center power consumption to the power consumed by the IT equipment alone. Ideally energy consumed outside of the IT equipment, such as by the electrical and cooling infrastructure, would be kept to a minimum leading to a PUE close to one. Traditionally only half the total data center energy is used at the equipment with typically about 25%-40% of the total data center energy consumed by the cooling infrastructure. Most of this is consumed by the site chiller plant, used to provide chilled water to the data center, and by computer room air conditioners (CRAC) and air handlers (CRAH), used to cool the computer room. Several earlier publications have shown that significant energy savings can be obtained by reducing the annual operating hours of, or even completely eliminating, the chiller plant and CRAH/CRACs units [3-5]. A project, jointly funded by the US Department of Energy (DOE) and IBM, was conducted at the IBM site in Poughkeepsie, NY, to build and demonstrate a data center test facility combining direct warm-water cooling at the server level and water-side economization at the facility level, to maximize energy efficiency and reduce the cooling power to less than 10% of the IT power (or <5% of the total data center power). Key elements of the design and results are presented in this article, with additional details available in references [6-11].

Dr. Milnes P. David is a development engineer in the Advanced Thermal Energy Efficiency Lab at IBM in Poughkeepsie, N.Y. His work focuses on a variety of areas related to improving the thermal management and energy use of computing systems and data centers, such as liquid cooling of servers, holistic thermal/energy analyses for data centers, energy-efficient controls, and energy recovery and renewables. David received his undergraduate degree from Georgia Tech in 2004 and his doctoral degree from Stanford University in 2011, both in mechanical engineering. At Stanford he studied and developed novel two-phase microfluidic heat sinks for high performance microprocessors under the guidance of Prof. Kenneth E. Goodson. David is an active member of ASME, IEEE and ASHRAE and is an industrial mentor in the NSF I/UCRC program.

CLOSED RACK WATER COOLING OF SERVERS

Direct liquid cooling of the IT equipment has been shown to offer both performance enhancements and energy savings [12-14]. It was further chosen for this program to enable warmer-water cooling and subsequently expand the use of free-air cooling over the year and to a wider geography. The IT equipment rack, shown in Figure 1, is comprised of 40 1U (1U = 44.45mm) 2-socket volume servers (often used for cloud applications), dissipating approximately 350 W each, with a total rack heat load of 14 kW. The processors are conduction cooled using cold-plates, and the memory modules by heat spreaders attached to cold rails. Measurements made comparing the device temperatures in an air-cooled server with the hybrid air & water-cooled version showed that the CPU junction temperatures were 30°C cooler when using 25°C water (in the air & water-cooled server) versus 22°C air (in the air-cooled server) [8]. Similarly, memory module temperatures were on average 18°C cooler and generally more uniform in the water-cooled server versus the air-cooled server. Such reductions in device temperatures can yield improvements in server performance and reliability and enables the use of warmer-water cooling. For example, at 45°C water temperature, the processor temperatures were below that of their

Puzzled about Environmental Testing or Design?

› Learn more about the interrelationship between electronic products, systems, devices and their environments.

etd | environmental test & design

envirotestdesign.com

counterparts cooled with 22°C air and the memory temperatures were nearly the same. Approximately 60%-70% of the heat dissipated at the server level is directly absorbed by the water circulating in the server level cooling hardware. Since less air flow is required at the server, three of the six server fans were removed, reducing IT power consumption. Heat that is not directly absorbed by the water-cooled server hardware is absorbed by the water flowing through the internal rack-enclosed air-to-liquid heat exchanger. Air is circulated through this internal heat exchanger by the remaining server fans and then returned to the servers in a closed loop, with no air exchange with the computer room. This fully enclosed rack design results in almost complete heat removal (~99%) to circulating water, protection from room air contamination and a very quiet rack. This combination of server-level direct water cooling and rack-level air-to-water cooling enables expanded warm-water (and thus free-air) cooling, elimination of energy-intensive and expensive computer room air conditioning and reduction in the number of server level fans and their associated power consumption and noise.

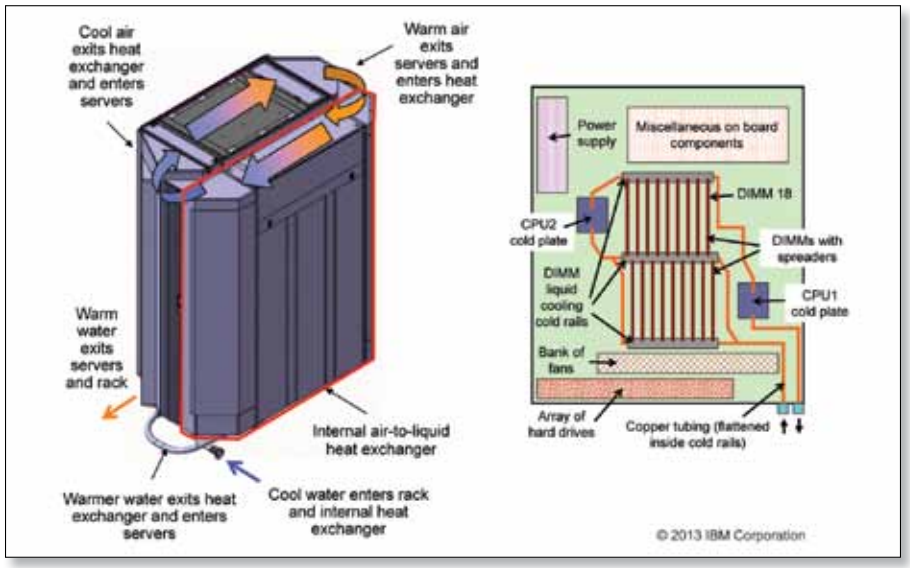

FIGURE 1: (Left) Model of modified IT rack with enclosed air-to-liquid heat exchanger to cool the hot air exhausting from the servers. No air leaves the IT rack into the computer room. (Right) Schematic of the volume server showing the direct water-cooled components.

WATER-SIDE ECONOMIZED DATA CENTER

Water or air-side economization utilizes ambient external conditions directly to offset some or all the necessary cooling, thereby reducing or eliminating intermediate refrigeration. Air-side economization utilizes outside air to provide the necessary cooling [15,16]. Filtration and humidification/dehumidification of the incoming air is necessary to lessen any risk associated with condensation, arcing, fouling or corrosion [17]. Water-side economization uses an intermediate economizer liquid-to-liquid heat exchanger to connect the facility water loop to the building or system water loop. This heat exchanger either supplements or completely replaces the chiller plant. In our data center test facility (shown in Figure 2), a 30kW capacity dry cooler, with 5 controllable fans, a facility coolant pump and a recirculation valve for use in the winter, is directly plumbed to a liquid-liquid buffer heat exchanger unit. No intermediate chiller is installed, resulting in significant energy and capital cost savings. The facility water supply temperature is allowed to vary with the ambient dry bulb temperature. Though a wet cooling tower can expand use of this data center architecture over most locations in the US and worldwide, dry coolers are attractive where water is expensive or difficult to obtain. As shown in Fig. 2, a dual loop data center design was used, which isolates the facility side coolant loop, containing a 50% water-propylene glycol mix, from the internal system loop containing distilled water. The propylene glycol mix was

chosen based on the historic minimum of -34°C in Poughkeepsie, New York. Though a single coolant loop arrangement is also possible, simulations found that the use of the propylene glycol mix, which has poorer thermal properties and is more viscous, in the full loop, reduces the potential thermal and pumping power improvements [9]. Separating the loops also has the advantage of needing to maintain a high coolant quality (anti-corrosives, fungicides, bactericides, filtration, etc.) only on the system side.

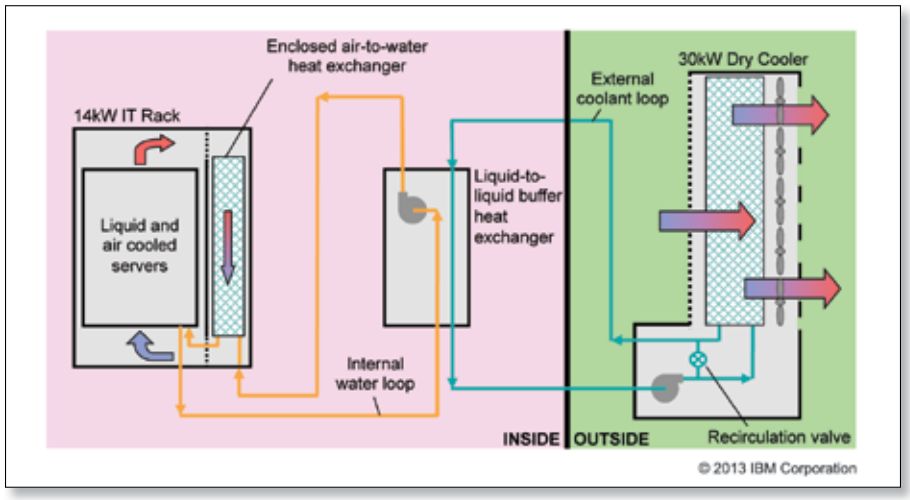

FIGURE 2: Schematic of the dual-loop water-side economized data center test facility that was constructed in Poughkeepsie, NY.

COOLING PERFORMANCE CHARACTERISTICS

The cooling characteristics of the rack level hardware, buffer heat exchanger and dry cooler as well as the cooling equipment power consumption were modeled using empirical functions derived from a series of characterization experiments, discussed in detail in reference [6]. The thermal resistances relationships obtained for pure water in both the internal and external loops are described by equations 1-3.*

Figure 3 shows the predicted CPU coldplate to outside dry bulb thermal resistance (determined from Eqs. 1-3) and the corresponding cooling equipment power for an internal loop flow rate of 0.25 l/s (4 GPM) and as external loop flow rate and dry cooler fan speed are varied. The results show that the total thermal resistance is dominated by the coldplate to rack inlet water thermal resistance. If the internal loop flow is increased to 0.5 l/s (8 GPM) this portion reduces significantly. The buffer thermal resistance reduces while the dry cooler resistance increases as external loop flow is increased. The cooler resistance reduces with increasing fan speed but the improvement is small after 8.3 Hz (500 RPM). Figure 3 (bottom) shows the total power consumption where P_{fans} is the power consumed by the dry cooler fans, P_{ext} is the power consumed by the external loop pump and P_{int} is the power consumed by the pump in the buffer unit that drives flow in the internal loop. The total power consumption doesn't change significantly until fan speeds of 8.3 Hz (500 RPM) but then dramatically increases up to the maximum 24.2 Hz (1450 RPM). Reference [6] further discusses the impact of adding propylene glycol to the external loop. These results highlight the importance of characterization to help determine the most energy efficient operating ranges for the different pieces of cooling equipment and in turn for developing the most ideal control algorithms to provide additional energy savings.

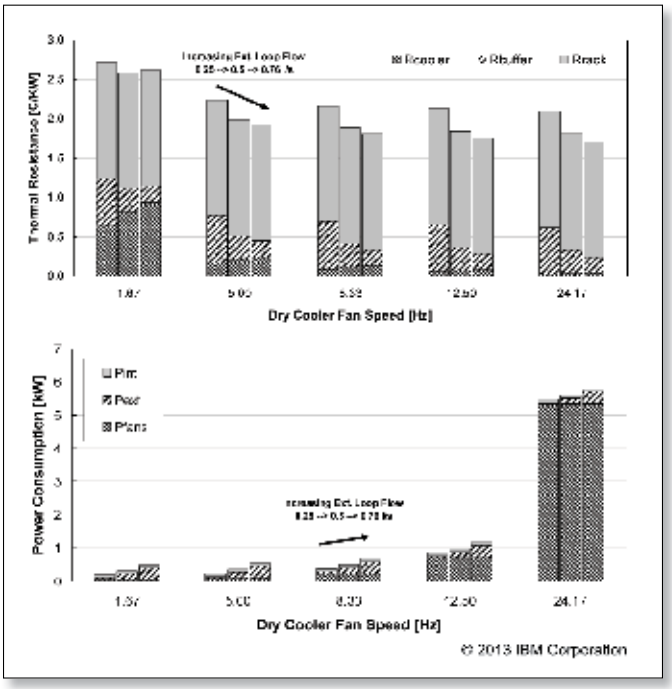

FIGURE 3: (Top) Thermal resistance from CPU coldplate to outside dry-bulb for 0.25 l/s (4 GPM) internal loop flow and varying external loop flow and fan speed. Rrack, Rbuffer and Rcooler are defined by Eqns. 1-3. (Bottom) Cooling equipment power consumption for same conditions. P_{fans} represent the dry cooler fans, P_{ext} the external loop pump and P_{int} the internal loop pump.

SYSTEM OPERATION AND CONTROLS

To study the actual energy consumption under operating conditions and to determine the cooling PUE (also referred to as mechanical PUE) several one day runs were carried out under simple table based control algorithms where the required equipment speeds are linearly interpolated for between user defined discrete settings. A long two month run was also carried out using a more sophisticated Proportional-Integral (P-I) control algorithm, details of which are discussed in references [10,11]. The control algorithms were implemented via a monitoring and control platform that was designed to allow both real-time tracking of system and facility environmental and power states as well as control of the dry cooler fans, internal and

external loop pumps and external loop recirculation valve. Summary results from the various one day runs as well as the two month run are shown in Table 1. Both the simple table based control and the P-I control implementation leads to similar cooling energy consumption of approximately 420-440W. However, the dynamic P-I control system was able to maintain the rack inlet water temperature, on average, within 0.5°C of the set-point of 35°C over the two month period. This helps maintain a more constant environment for the IT equipment and reduces reliability issues associated with rapid changes in operating environment. The cooling energy use of 420W for an IT load of 13kW translates to a cooling PUE of approximately 1.035 or a Cooling to IT power ratio of 3.5%. Compare this with a Cooling to IT power ratio of almost 50% for a traditional chiller based data center and this represents an over 90% reduction in data center cooling energy use. For a typical 1 MW data center this would represent a savings of roughly $90-$240k/year at an energy cost of $0.04 - $0.11 per kWh.

Date	4-Aug	19-Oct	5-Oct	11-May
Duration	1 Day	1 Day	1 Day	62 Day
Season	Summer	Fall	Fall	Spring - Summer
Weather	clear	rainy	clear	mixed
Ambient Temp [C]	24.0	14.9	10.8	21.6
Control Type	Table based	Table based	none	Model based PI Control
Fan Speed Setting, Hz [RPM]	2.8 - 8.3 (170 - 500)	2.8 - 8.3 (170 - 500)	fixed speed	2.8 - 8.3 (100 - 750)
Pump Speed Setting, Hz [RPM]	fixed flow	fixed flow	fixed flow	18.3 - 57.5 (1100-3450)
Ext Loop Flow, l/s [GPM]	0.45 (7.1)	0.45 (7.1)	0.24 (3.9)	0.33 (5.2)
Int Loop Flow, l/s [GPM]	0.45 (7.1)	0.46 (7.2)	0.25 (4.0)	0.38 (6.0)
Fan Speed, Hz [RPM]	3.3 (200)	2.8 (170)	2.8 (170)	3.2 (189)
Rack Inlet Liq Temp [C]	33.8	26.8	23.2	34.6
Server Inlet Air Temp [C]	36.5	30.5	29.6	36.6
IT Power [kW]	13.14	13.10	13.21	12.76
Cooling Power [kW]	0.44	0.43	0.21	0.42
Cooling / IT%	3.2	3.3	1.6	3.5

TABLE 1: Summary of operating conditions, temperatures, power consumption and efficiency for several day long runs and a for a long two month run.

CONCLUSION

The results of this project demonstrate the extreme energy efficiency attainable in a water-side economized data center with water-cooled volume servers, and support the development and deployment of other similar systems. One large-scale commercial implementation of a similar architecture is the 3PFlop SuperMUC HPC system unveiled in 2012 at the LRZ in Munich, Germany [18]. The SuperMUC system, which shares some technological features with this project, is currently ranked 9th on the Top500 list based on performance [19] and 15th greenest system (considering only 1MW+ systems) on the Green500 list based on a performance per watt of 846 MFlop/Watt [20].

ACKNOWLEDGEMENT

This project was supported in part by the U.S. Department of Energy's Industrial Technologies Program under the American Recovery and Reinvestment Act of 2009, award no. DE-EE0002894. We also thank the IBM Research project managers, James Whatley and Brenda Horton, DOE Project Officer Debo Aichbhaumik, DOE Project Monitors Darin Toronjo and Chap Sapp, and DOE HQ Contact Gideon Varga for their support throughout the project. We also thank George Manelski, Sal Rosato and the IBM Site & Facilities team for all their help.

$$R_{rack} = \frac{(T_{CP}\text{-}T_{rack\ inlet\ water})}{Q_{IT}} = 0.642\ (\text{int flow in 1/s})^{-0.602} \quad \textbf{(1)}$$

$$R_{buffer} = \frac{(T_{rack\ inlet\ water}\text{-}T_{buffer\ inlet\ water})}{Q_{buffer}} = 0.256 \cdot (\text{int flow in 1/s})^{0.37} \cdot (\text{ext flow in 1/s})^{-0.99} \quad \textbf{(2)}$$

$$R_{cooler} = \frac{(T_{buffer\ inlet\ water}\text{-}T_{ambient})}{Q_{cooler}} = 1.927 \cdot (\text{ext. flow in 1/s})^{0.35} \cdot (\text{fan speed in Hz})^{-1.22} \quad \textbf{(3)}$$

REFERENCES

[1] J.G. Koomey, "Growth in Data Center electricity use 2005 to 2010", Oakland, CA: Analytics Press., Aug 2011.

[2] R. Brown et al., "Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431," Lawrence Berkeley National Laboratory, 2008.

[3] M. Iyengar, R. Schmidt, J. Caricari, "Reducing Energy Usage in Data Centers Through Control of Room Air Conditioning Units," 2010 12th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), Jun 2-5, 2010, pp. 1-11.

[4] ASHRAE Datacom Series, "Best Practices for DataCom Facility Energy Efficiency," 2nd Edition, Altanta, 2009.

[5] Vision and Roadmap – Routing Telecom and Data Centers Toward Efficient Energy Use, May 13, 2009, www1.eere.energy.gov/industry/dataceters/pdfs/vision_and_roadmap.pdf.

[6] M. P. David et al. "Impact of operating conditions on a chiller-less data center test facility with liquid-cooled servers," 2012 13th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), San Diego, CA, May 30–Jun 1, 2012, pp. 562 – 573.

[7] M. K. Iyengar et al., "Server liquid cooling with chiller-less data center design to enable significant energy savings," 2012 28th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), Mar 18-22, 2012, pp. 212–223.

[8] P. R. Parida et al. "Experimental investigation of water-cooled server microprocessors and memory devices in an energy efficient chiller-less data center," 2012 28th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), Mar 18-22, 2012, pp. 224–231.

[9] P. R. Parida et al. "System-Level Design for Liquid-cooled Chiller-Less Data Center", ASME 2012 International Mechanical Engineering Conference and Exposition (IMECE), Nov 9-15, 2012, Paper No. IMECE2012-88957, pp. 1127-1139.

[10] M. P. David et al., "Dynamically Controlled Long Duration Operation of a Highly Energy Efficient Chiller-less Data Center Test Facility," 2013 12th ASME International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems (InterPACK2013), July 16-18, 2013, Paper# InterPACK2013-73183.

[11] P. R. Parida et al., "Cooling Energy Reduction During Dynamically Controlled Data Center Operation," 2013 12th ASME International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems (InterPACK2013), July 16-18, 2013, Paper# InterPACK2013-73208.

[12] M. J. Ellsworth et al., "An Overview of the IBM Power 775 Supercomputer Water Cooling System," J. Electron. Packag., 2012, 134(2):020906-020906-9.

[13] M. J. Ellsworth and M. Iyengar, "Energy Efficiency Analyses and Comparison of Air and Water-cooled High Performance Servers," Proceedings of ASME InterPACK 2009, July 19-23, 2009, Paper# IPACK2009-89248, pp. 907-914.

[14] J. Wei., "Hybrid Cooling Technology for Large-scale Computing Systems – From Back to the Future," ASME 2011 Pacific Rim Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Systems, MEMS and NEMS: Vol. 2, July 6-8, 2011, Paper# IPACK2011-52045, pp. 107-111.

[15] Timothy P. Morgan, "Yahoo! opens chicken coop data center," The Register, Sep. 2010, http://www.theregister.co.uk/2010/09/23/yahoo_compute_coop/

[16] C. Metz, "Welcome to Prineville, Oregon: Population, 800 Million". Wired, Dec. 2011, http://www.wired.com/wiredenterprise/2011/12/facebook-data-center/

[17] M. Iyengar et al., "Energy Efficient Economizer Based Data Centers with Air-cooled Servers," 2012 13th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), May 30 - Jun 1, 2012, pp. 367-376.

[18] H. Huber et al., "Case Study: LRZ Liquid Cooling, Energy Management, Contract Specialities," 2012 SC Companion: High Performance Computing, Networking, Storage and Analysis (SCC), Nov, 2012, pp.962-992.

[19] Top500 List – June 2013, accessed: Oct 28, 2013, http://www.top500.org/list/2013/06/

[20] The Green500 List – June 2013, accessed: Oct 28, 2013, http://www.green500.org/lists/green201306

Index of Advertisers

Electronics COOLING

Alpha's Next Generation Heat Sink

Custom or off-the-shelf.
Simple to complex.
Prototype to mass production.

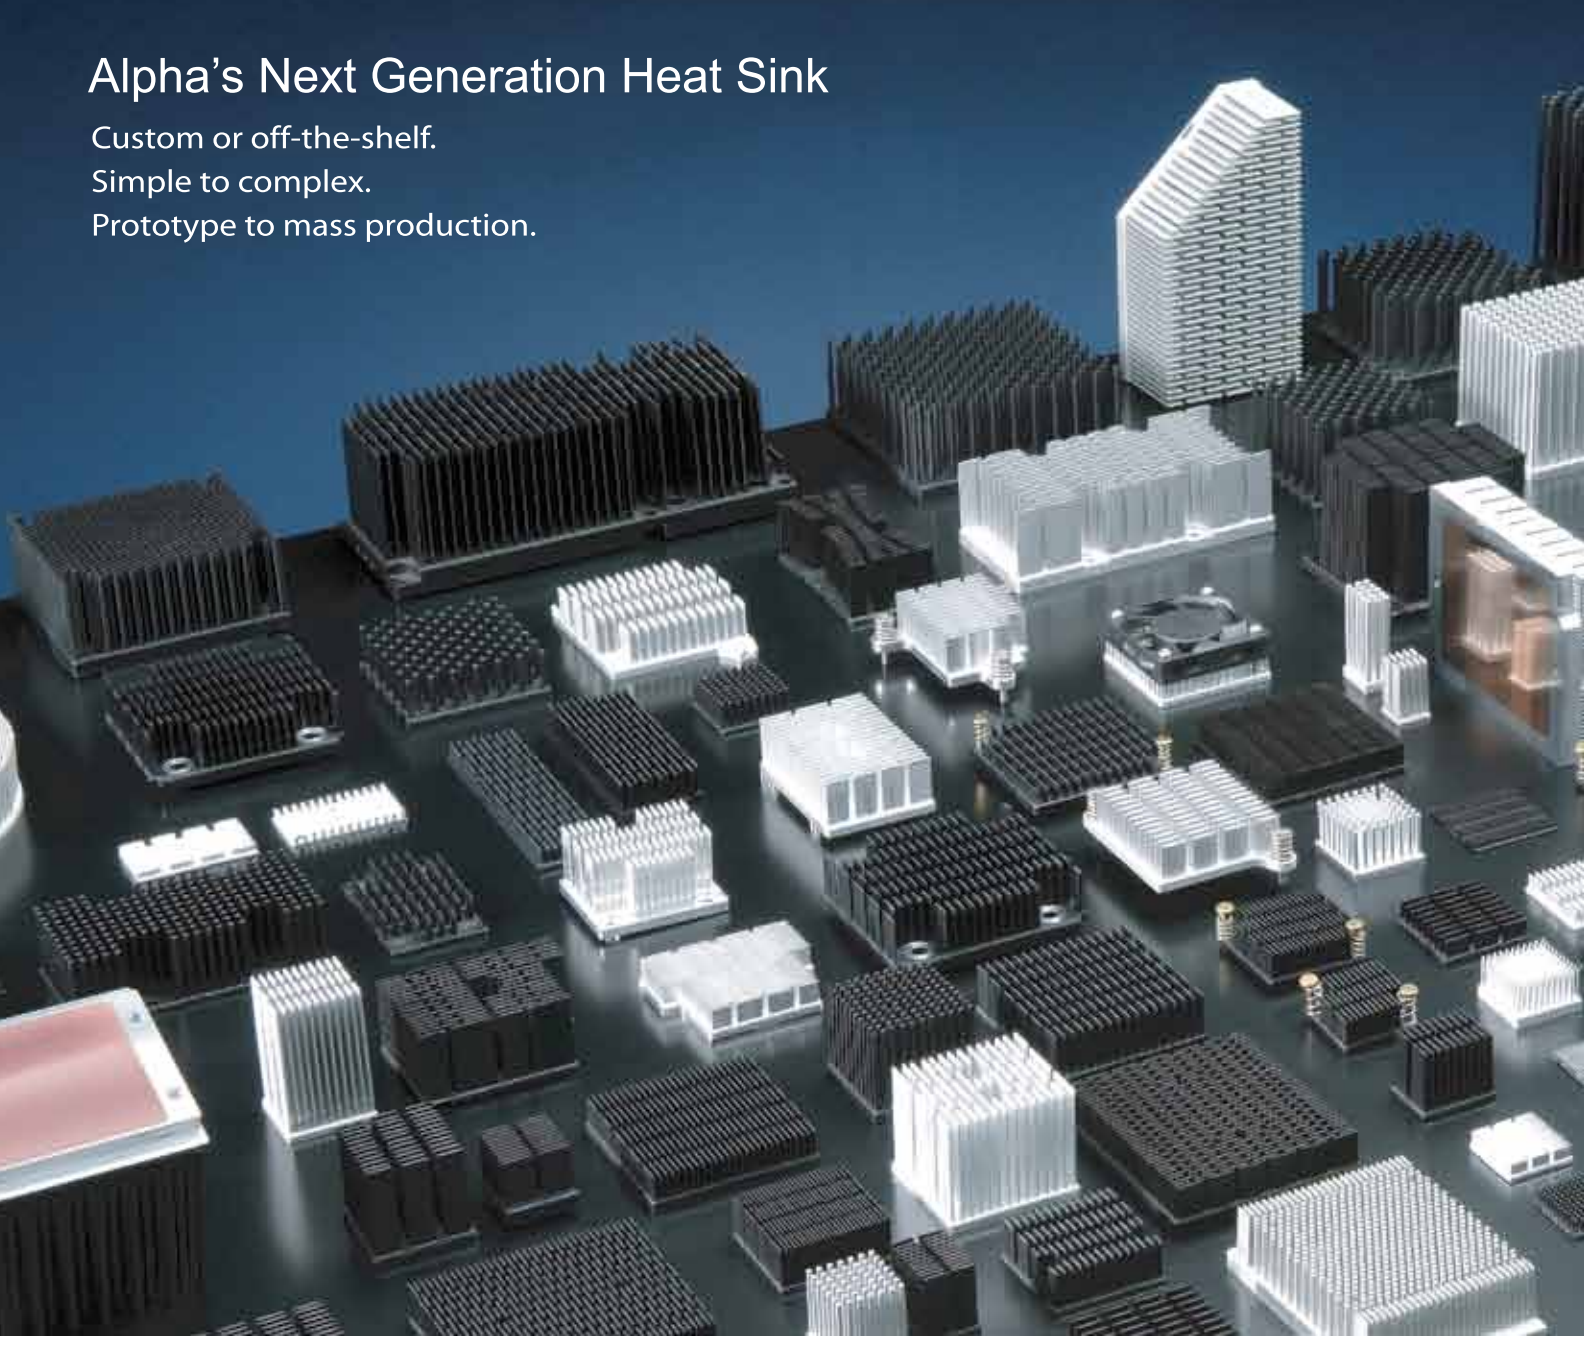

QSZ Clip

Minimum PCB Area

Pins only require 1.8mm diameter holes in the PCB.

Mounting Security

Easy-Install

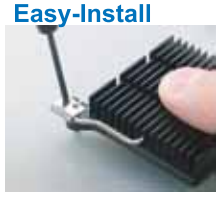

ALPHA SUPPORT

Variety of catalog parts : The most suitable heat sink, for many different environments, can typically be found from Alpha's standard catalog.

Data Library : Download CAD files, Flotherm data models and RoHS CofC from Alpha's website.

Engineering Support : Alpha provides thermal solutions with a focus on short lead times and low cost while providing outstanding customer support.

ALPHA

Your partner for thermal solutions

ALPHA Co., Ltd.
Head Office
www.micforg.co.jp
256-1 Ueda, Numazu City, Japan 410-0316
Tel: +81-55-966-0789 Fax: +81-55-966-9192
Email: alpha@micforg.co.jp

ALPHA NOVATECH, INC.
USA Subsidiary
www.alphanovatech.com
473 Sapena Ct. #12, Santa Clara, CA 95054 USA
Tel:+1-408-567-8082 Fax: +1-408-567-8053
Email: sales@alphanovatech.com

GREEN TECH

The highest performing fans on the market.

Whether for cooling server rooms or switch cabinets, complex machines or sensitive medical technology - our intelligent and controllable high performance S-Force fans can be suited for many different applications. Each series represents the benchmark in its class with respect to air performance and pressure build-up with optimum motor efficiency and long service life. These compact fans reach nominal speeds up to 14,000 rpm and operating values that had previously been attained only by large fans or blowers. To learn more about our IP rated fans, please visit **www.ebmpapst.us/sforce** or contact Application Engineering at 860-674-1515.

ebmpapst

The engineer's choice