

2017

Electronics Thermal Management in Information and Communications Technologies: Challenges and Future Directions

S. V. Garimella

Purdue University, sureshg@purdue.edu

T. Persoons

J. A. Weibel

Purdue University, jaweibel@purdue.edu

V. Gektin

Huawei

Follow this and additional works at: <http://docs.lib.purdue.edu/coolingpubs>

Garimella, S. V.; Persoons, T.; Weibel, J. A.; and Gektin, V., "Electronics Thermal Management in Information and Communications Technologies: Challenges and Future Directions" (2017). *CTRC Research Publications*. Paper 321.

<http://dx.doi.org/10.1109/TCPMT.2016.2603600>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

Electronics Thermal Management in Information and Communications Technologies: Challenges and Future Directions

Suresh V. Garimella, Tim Persoons, Justin A. Weibel, and Vadim Gektin

Abstract — This paper reviews thermal management challenges encountered in a wide range of electronics cooling applications from large-scale (data center and telecommunication) to small-scale systems (personal, portable/wearable and automotive). The paper identifies drivers for progress, and immediate and future challenges, based on discussions at the 3rd Workshop on Thermal Management in Telecommunication Systems and Data Centers held in Redwood City, California, on November 4-5, 2015. Participants in this workshop represented industry and academia, with backgrounds ranging from data center thermal management and energy efficiency to high-performance computing and liquid cooling, thermal management in wearable and mobile devices, and acoustic noise management. By considering a wide range of electronics cooling applications with different length and time scales, the paper identifies both common themes and diverging views in the thermal management community.

Index Terms — Information and communication networks; Energy management; Thermal management of electronics.

I. BACKGROUND

The content of this paper is informed by presentations and discussions during the 3rd Workshop on Thermal Management in Telecommunication Systems and Data Centers held in Redwood City, California, November 4-5, 2015. The workshop drew together representatives from academia¹ and a range of industries² spanning large-scale electronics and

telecommunications systems, to small-scale personal electronics. While the underlying thermal management challenges are similar in many of these applications, each is inevitably characterized by different drivers and constraints. This paper is inspired by discussions and presentations at this Workshop, and organized as follows: (i) observations related to the background for ICT thermal management, (ii) the business and technological drivers for progress, and (iii) a review of the wide range of thermal management challenges in large- and small-scale systems, using the workshop participants' opinions as a starting point. (iv) Finally, an overview of the implementation of thermal solutions is presented, broken down into different technologies and heat transfer modes.

A number of recent publications have addressed the thermal management challenges in large-scale electronics systems, including those by the authors [1,2] that were based on prior workshops in this series organized by the Cooling Technologies Research Center, a National Science Foundation Industry/University Cooperative Research Center at Purdue University. Because of the growing energy consumption by large-scale electronic systems and the data center sector in particular [1], thermal management challenges in this sector have received a lot of attention in recent research and review papers [3,4]. The current paper acknowledges the significant challenges and opportunities in large-scale electronic systems; yet, its aim is to provide a more comprehensive context by also considering the needs and restrictions in small-scale electronic systems. The increasing degree of interconnectivity, continuing growth in data volumes at the level of personal electronic devices, and proliferation of embedded electronics and sensors in transportation systems will inevitably have repercussions on the energy usage by large-scale systems providing services such as cloud-based data storage and computing.

A widely adopted energy efficiency metric encountered in data center thermal management, and one which therefore

¹Purdue University, Trinity College Dublin, Kyushu University, University of Houston, and Stanford University

²Amazon Lab126, CoolIT Systems Inc, EXA Corporation, Fujitsu Limited (Japan), Hewlett Packard Labs, Huawei Technologies (US, Sweden and China), Intel Corporation, Qualcomm, Samsung Electronics, and Toyota Research Institute of North America

appears frequently in this report, is the power utilization effectiveness (PUE), defined as:

$$PUE = \frac{P_{tot}}{P_{IT}} = \frac{P_{IT} + P_{th} + P_e}{P_{IT}} = 1 + \underbrace{\frac{P_{th} + P_e}{P_{IT}}}_{\varepsilon} \quad (1)$$

in which P_{IT} , P_{th} , P_e are the power consumption of the ICT equipment, thermal management and electrical power delivery systems, respectively.

The amount of annual data generated globally in the last decade has rapidly increased from 0.1 ZB in 2005, to 1.2 ZB in 2010, and 8.5 ZB in 2015, with a projection of 40 ZB for 2020 [5]. Amazon, Google, and Uber are just a few examples of market-disrupting businesses that rely on real-time processing of ‘Big Data’. The handling of vast amounts of scientific data, such as human DNA sequences, would benefit enormously from cloud-based data storage and computing [6]. In terms of the energy footprint of public cloud computing, a survey by Uptime Institute has revealed that the average data center’s power usage effectiveness (PUE) has been only slightly reduced from 1.89 in 2011 to 1.70 in 2014 [7], still leaving ample room for improvement compared to the ideal value of unity.

In addition to the traditional air-cooled servers in a data center, ‘hybrid’ thermal management solutions employing a combination of air cooling and liquid cooling at the server level are being gradually introduced for high-end computing facilities. In a hybrid air/liquid cooled server, only the components with the highest heat loads such as processors are liquid cooled. More conventional liquid-based heat spreading techniques (*e.g.*, heat pipes, vapor chambers) are not taken into consideration here. One key benefit of the hybrid cooling approach over traditional air cooling is the increased potential for recuperation of waste heat from the liquid coolant stream. Different system configurations are being considered to make use of this waste heat. Some of these involve complex thermodynamic cycles for maximizing energy efficiency and computational performance [8]. A hybrid liquid cooling solution (60% by water, 40% by air [8]) is cheaper than a fully liquid-cooled analog, and yet achieves a PUE as low as 1.3 [9]. With a tendency towards density-driven designs, liquid cooling (even if only used on high-power components such as CPUs) allows for about three times denser CPU packaging for the same footprint area, and a similar increase in volumetric power density (W/m^3) [9]. Indeed, performance is increasingly being expressed in terms of volumetric performance ($Gflops/W$ and $Gflops/m^3$) [9]. This will require an even closer integration of ICT equipment and thermal management approaches in future designs.

Extreme heat fluxes beyond those encountered in typical data centers must be dissipated in radar, power electronics, and high-performance computing (HPC) systems. Evaporative cooling strategies in direct contact with the semiconductor device are therefore the focus of the current DARPA Intrachip/Interchip Enhanced Cooling (ICECool) program

[10]. Targeted heat density levels are $> 1 \text{ kW/cm}^2$ and $> 1 \text{ kW/cm}^3$ with a chip temperature rise below 30°C and maximum temperature difference across the chip footprint below 10°C .

Thermal design for portable electronics requires approaches that are quite different from the traditional paradigms in large-scale electronics cooling. Here, the focus is on customer-centered performance, and thermal targets are driven by ergonomic considerations such as limits on the skin temperature. However with increasing performance and decreasing form factors, this end of the design spectrum also relies on concurrent interaction between electrical and thermal engineering teams. Electrothermal co-design emerges as a common theme for next-generation systems and devices, whether at large or small scale.

II. DRIVERS FOR PROGRESS

A. Business drivers

Current investment strategies in large-scale ICT facilities seem to be driven more by operational expenditure (OpEx) than by capital expenditure (CapEx), as evidenced by the increasing number of large-scale bare-bones systems relying on free cooling, designed for resilience rather than redundancy [1]. One might wonder whether relying solely on air cooling could result in a reluctance to invest in the development of new technology. Without research and development investment in higher risk cooling technologies including liquid cooling, higher CapEx costs may prevent these technologies from reaching a high enough penetration to achieve sector-wide savings in OpEx by a reduction in energy use [1]. However, changes in policy could induce a change in the CapEx balance between air and liquid cooling. Acoustic noise regulations are set by standards and thus subject to policy changes. An increase or decrease in noise emission thresholds would change the level of detail needed from aeroacoustics simulations in the development of servers. Along with other drivers including enhanced reliability and increased energy efficiency, this may push up the cost of air-cooled systems.

In other applications such as hybrid electric vehicles (HEVs), the power density of inverters is such that air cooling would require high fan loads resulting in excessive acoustic noise levels. This has driven decisions towards liquid-cooled cold plates in personal vehicles. A similar shift may well arise with the advent of 3D integrated circuits and denser packages in HPC and volume servers.

Market drivers for liquid cooling also vary by region. Computational performance using overclocked CPUs is the main driver in North America, whereas energy efficiency and density are the main drivers for Europe and Japan due to the higher electricity cost.

The market demand for mobile phones with added features, functionality and performance inevitably increases thermal susceptibility. To keep up with performance demands, the average CPU power density per core has been steadily

increasing with phone generations. Top market drivers include price, battery life, functionality, as well as sleekness of the product which complicates electro-thermal design. Phone makers have a paradoxical choice concerning sleekness versus battery life – although sleekness is not a top priority based on customer surveys, manufacturer vie for dominance in the race for ever-thinner phones.

Comparing maximum power dissipation and the approximate cost of the thermal solution across different form factors, (i) a laptop is limited to 15-20 W at a cooling cost of \$1.50/W, while the budget is (ii) 7-10 W and \$2.50/W for a tablet and (iii) 2-3 W and \$3.50/W for a mobile phone [9].

Other possible practical business constraints to a further penetration of energy-saving cooling methods could include the need to accommodate legacy air-cooled products, and being able to guarantee customers an established level of reliability.

B. Technological drivers

1. Computing versus communication costs

A key technological driver is the balance between the cost for computation versus the cost for communication [1]. This balance exists on the micro level, *e.g.*, driving the evolution towards integration of logic and memory chips in a single package, as well as on the macro level in the balance between (on-device) local computing versus (off-device) cloud computing in data centers.

It is worth asking if there may be a shift back from cloud computing toward local computing if cheap, high-density options for local computing became available (such as IBM's concept for a sugar cube-size supercomputer [11]). Or would this lead to data centers simply adopting and multiplying this enhanced technology on a larger scale?

It remains impossible to combine all functionality in a single device; therefore, the distributed heterogeneous model is the only viable solution. Intel's Knights Landing is a 'many integrated core' coprocessor chip [12,13] which is a practical example of this model.

From a provider's point of view, it is more economically viable to carry out as much local computing on the device as possible, since these energy costs are at the user's expense, whereas server energy consumption is at the provider's expense.

As transmitted data volumes increase, photonics and combined electronics/photonics modules for optical communication will gain in importance. The thermal management requirements are quite different in this case, with laser-based photonics requiring very precise temperature control to prevent laser wavelength drift. However, the same cooling technologies can be used for both electronics and photonics packages; for photonics, actively controlled thermoelectric coolers are considered in addition [14,15].

Photonics thermal management challenges may become more extreme for optical power transfer applications [16]. The introduction of more and more sensors in automotive systems has led to a significant increase in the amount of copper wiring. Significant weight savings would be achieved if a sensor could be connected with a single optical fiber for communication and power delivery. There are currently about 100 sensors in a typical car, but this number will be significantly higher for autonomous vehicles [17,18]. Vehicle-to-network communication is already in place (to exchange traffic information, for example) but the future will see more vehicle-to-vehicle communication as well.

2. Proliferation of different computing systems platforms

At present we have a good grasp of the thermal behavior of components; the computing revolution continues, however, with new and as yet unimagined applications in the 'Internet of Things' (IoT) [19]. As new applications emerge, the focus has shifted to systems rather than components. Computing will evolve along three vectors: (i) small, flexible, light-weight interconnected devices; (ii) increased computing performance at manageable cooling costs in data centers and other server markets; and (iii) increased proliferation of embedded computing (*e.g.*, in automotive, healthcare, buildings, space, and other commercial/consumer markets).

The number of devices has increased from 2 billion in 2006 to 15 billion in 2015, and is projected to approach 20-50 billion by 2020 [9,20,21]. This evolution will see a rise in background computing in data centers. Today, the data center business includes enterprise ICT, cloud service providers, telecommunication service providers and scientific computing. From 2014 to 2018, a 15% growth in data center business is expected [9]. The main growth drivers are (i) the cloud market offering low-cost quick access, (ii) high performance computing (HPC), and (iii) Big Data manipulation.

From the 1990s to 2015, devices have evolved from clunky to sleek systems with more functionality, becoming slimmer and more portable. The next step seems to be even thinner, flexible devices which could be integrated in new ways, *e.g.*, as part of clothing in the form of wearables. This broad spectrum of devices, packages and form factors will see a growing set of communication protocols. If different devices should be able to talk to one another, they will have to evolve an ability to support multiple protocols.

Across the spectrum from portable devices to data centers, a key focus is energy. Intel's Shekhar Borkar anticipates an exascale data center with a power consumption below 20 MW to be realizable by 2020 [22]. This achievement would require a hypothetical processor with 4,096 cores on a die capable of 16 teraflops at double precision for below 100 W, or 200 Gflops/W. This exascale machine could be realized by combining 62,500 of these processors with four threads per core.

3. Thermal management for mobile phones

Experts on mobile phone thermal design see two clear trends: (i) towards thinner phones, requiring thinner electronics packages and boards, with package thickness reducing from 1.4 mm to 1 mm, and further to 0.5 mm in the near future; and (ii) customer demand for longer battery life that results in an increase in battery size [9]. However, because the printed circuit board (PCB) size has not changed, these trends lead to a consideration of different internal layouts. The PCB takes up 30-40% of the phone by area. Typical layouts have the battery below and PCB on top, or a C-shaped PCB with an elongated battery alongside. The board has very dense routing and this leads to a risk of hot spots on the outer skin surface of the phone within the limited package size.

Portable consumer electronics such as mobile phones have seen significant evolution in performance and to a lesser extent also in battery capacity. Between 2010 and 2012, the processor performance for Samsung Galaxy models has increased 5.9 times while the battery capacity increased 1.4 times [23]. The trends are similar for the Apple iPhone. So while performance and heat dissipation are sharply increasing, battery capacity is not following at the same rate.

The trend towards smaller form factors combined with higher performance means that hot spots become more prevalent and problematic to resolve. It is becoming more important to balance computing performance and power consumption. Key factors for innovative, active thermal solutions include heat radiating capacity, noiseless and vibration-free operation, scalability to smaller form factors, reliability, and low cost.

For wearables, low-power operation is crucial both from the perspective of battery life and ergonomic constraints related to skin temperature.

4. How to extend the life of air cooling?

At a previous Workshop in this series in 2012, the following list of required developments was proposed to extend the life of air cooling in data centers [1]:

1. At the system level: More efficient air movers (coefficient of performance, COP > 20) and a reduction in fan noise to below 60 dB(A).
2. At the board level: Optimized heat sinks with low pressure drop and thermal resistance, R_{th} (e.g., < 0.1°C/W for a 140 CFM 1U volume server) and alternative air movers to decouple pumping power and heat transfer, such as piezo fans, synthetic jets, and electro-aerodynamics.
3. At the component level: Limit the component power dissipation and reduce R_{th} for component and thermal interface material (TIM).

Today there still seems to be general agreement on these targets, with particular emphasis on the importance of component-level developments. The system- and board-level developments are determined by economic drivers but the

same board-level targets were deemed valid at the 2015 Workshop as well [9].

Regarding acoustic noise in data centers, there are two limitations to be considered: (i) annoyance and ergonomics, and (ii) regulatory safety limits for maintenance workers. While there is debate over the necessity of the former in data center environments, the latter will always remain because data centers cannot afford to shut down to carry out maintenance in quiet conditions. In residential areas, the annoyance limitation may yet prove important.

A significant contribution to the noise problem in air cooling comes from excessive fan backpressure resulting from poorly designed grilles and backplane connectors. Even in a notebook computer, the pressure drop is split approximately equally between the heat sink and the grilles and ducting. Improvements are possible by simply arranging connectors differently. Part of the solution would be to better educate other design teams on thermal-fluid engineering and aeroacoustics. Fans are already being customized for certain applications, with increased dialogue between ICT systems manufacturers and fan suppliers. This tailored design approach should include the placement of the fan and all components making up the air path.

A sometimes overlooked contributor to fan load is the need for air filtration, especially in uncontrolled building environments located in emerging markets with often high levels of airborne pollutants. The pressure drop across the filters may cause the fans to operate closer to stall conditions, thereby further increasing power consumption and acoustic noise levels [24,25,26].

For certain high-density board layouts, air cooling may become so challenging that liquid-cooled cold plates are the only viable option. The three main drivers for liquid cooling are performance, efficiency and density. Most practical implementations would not be 100% liquid-cooled, but instead involve a hybrid combination of liquid cooling for high-power components and air cooling for low-power components. This may shift the backpressure balance in favor of either fans or alternative air movers, although many of the latter (e.g., piezo fans and synthetic jets) can overcome only a small pumping head and require an integrated design approach to provide sufficient steady-state cooling performance [27,28].

In mobile devices, widespread introduction of forced air cooling is unlikely. However in an IoT context, there could be a need to plug in a mobile device into docking stations which could provide active cooling to unlock increased performance in the device. A large contact area with low thermal resistance would be required to accomplish this or some other form of thermal connection. At least for now, handheld devices do not seem to have reached a plateau where active cooling has become essential. Power-hungry functionality such as video recording is already commonplace, and other limits such as battery life are often hit before thermal limits.

III. IMMEDIATE AND FUTURE CHALLENGES

A. Diversity of perspectives

A summary of perspectives offered by Workshop participants from diverse sectors, facilitated by a polling device with real-time feedback, is provided below. For the results discussed in sections III.A.1 and III.A.2, responses were collected from twenty six participants. The top challenges discussed in section III.A.3 were based on responses from nine senior industry experts within this group.

1. Appropriateness of thermal management design metrics

In the case of large-scale ICT facilities such as data centers, participants believed by a 3:1 margin (Fig. 1a) that reducing the absolute energy consumption P_{tot} is preferable as a thermal management design target to reducing the PUE value (Eq. (1)) by minimizing the relative fraction ε through means such as running barebones servers in hot conditions and utilizing free cooling. This requires a collaborative effort between electrical and thermal design teams which could bring about joint reductions in P_{IT} and the ancillary $P_{th} + P_e$ can be accomplished in parallel.





Figure 1. Workshop participants' perspectives on (a) the required thermal management design target for large scale systems (A: reducing PUE or B: reducing total power consumption) and (b) the predicted penetration of convective liquid cooling techniques in the next five years on a scale of 1 to 9 (1 = air only, 9 = liquid only)

As a metric, PUE values can be manipulated depending on how the power consumption is categorized. Its simplicity may have contributed to the proliferation of its use in the media. Indeed, it has focused the agenda on the issue of energy consumption and efficiency in these systems, and reducing thermal management energy consumption in certain facilities with high PUE values ($PUE > 2.5$) remains a meaningful objective. Apart from PUE-related metrics, other metrics have been considered such as a coefficient of performance where the leakage, idling and rack-internal fan power are subtracted from the ICT equipment power and added to the (rack-external) infrastructure power consumption, *i.e.*, thermal management and electrical power delivery [29].

2. Penetration of liquid cooling in future large-scale ICT facilities

Workshop participants predicted a 39% penetration-level of liquid-cooling techniques into large-scale ICT facilities (*e.g.*, data centers, HPC and telecommunication systems) in five years' time, relative to a penetration level of 5-10% in 2015 [30]. This current level was pointed out to the participants prior to collecting their responses. Only *convective* liquid cooling techniques are considered, and not the more ubiquitous fully sealed heat spreaders such as heat pipes or vapor chambers. On a scale of 1 to 9 (1 being air-only and 9 being liquid-only), the histogram of responses in Fig. 1b shows a mean response of 4.1. The responses were distributed bimodally with a highest peak at 25% and smaller peak at 63% penetration levels.

Liquid-cooling proponents gave reasons such as an increase in energy efficiency and performance. High-performance computing is the primary market, with compound annual growth rate (CAGR) in excess of 15%. Some companies expected liquid cooling to be first introduced in their HPC facility, after which it may be adopted in other areas.

There has long been a moving horizon claiming an imminent transition to liquid cooling 'less than 5 years away', as well as a history of irrational fear of the introduction of liquid coolants. This apprehension seems to have abated somewhat in recent years as more evidence reaches the media [31]. Depending on a company's culture, some will take longer to even start considering liquid cooling as an option.

A move from air to liquid cooling could be instigated by a demonstrated paradigm shift in performance, economics, or by regulatory mandates. The latter could become important as governments may decide to regulate energy usage. This could be done in the form of incentives or other mandates, which may tip the economic balance (through a shift in the total cost of ownership, TCO) toward liquid cooling. While European regulations do not apply to the US market, multinational corporations may tend to favor solutions that meet more than one market.

A transition to liquid cooling could take two forms: (i) a complete system redesign, or (ii) a gradual introduction of

component-level liquid cooling solutions, *e.g.*, replacing fan-based CPU coolers with cold plates and a closed circulation loop with a secondary heat exchanger. These approaches will likely yield different gains in performance. In the personal computing market, liquid cooling has long been used in the gaming community, but it is unlikely to spread to a wider field any time soon.

3. Top thermal management challenges

Workshop participants, representing a range of industries from large-scale ICT hardware to automotive and handheld device markets, identified their top thermal management challenges. These were identified by the respondents without a list of suggested challenges being provided by the workshop organizers in advance. As illustrated in Fig. 2, the top priorities were different according to the sector: (1) **low cost** operation is a top priority for representatives from personal electronics companies, whereas (2) **heat spreading, IC integration, co-design** and **collaboration** are priorities for representatives from large-scale ICT system companies.




Figure 2. Distribution of top challenges according to workshop participants: — all participants, — participants from companies involved in large-scale systems (*e.g.*, data centers, HPC), — participants from companies involved in personal electronics (*e.g.*, handheld devices, automotive)

The eight most common themes in the responses are plotted in Fig. 2, and together represent 75% of all responses. The blue line represents the responses from all participants, while the red and green lines indicate the responses from participants representing companies active in large-scale (red) versus small-scale systems (green). The values for each category in Fig. 2 correspond to the sum of occurrences in each respondent's top three challenges, weighted by factors of 3, 2 and 1 from the first to third response. Responses from nine

experts were taken into account for Fig. 2; five from industries involved in large-scale electronics and four with activities in personal electronics and automotive. The results should be seen as indicative and not definitive given the limited size of the companies polled.

B. Thermal and electrical challenges

From the early 1990s to the mid-2000s, processor chips based on CMOS technology were becoming increasingly thermally constrained, from the Intel 486 processor generating only 4 W, to the Pentium initially designed for 15 W but exceeding 100 W by 2004. Hot spots became an issue with the introduction of the latter, as local zones of high heat dissipation were in close proximity to local memory. With thermal design focusing on controlling junction temperature and hot spots, heat spreading became very important.

During this period, the thermal community responded with several initiatives that improved TIMs [32,33,34,35], thermal metrology [36,37], modelling [38,39] and various building block technologies such as microchannel heatsinks [40,41,42,43,44], ionic wind [45], heat pipes [46,47], and thermoelectric coolers [48,49]. It also saw the emergence of major research efforts funded in the US by DARPA programs such as HERETIC (starting 1992), followed by nTIM, MACE, TGP, NJTT, and ACM (2008), and most recently by the Intra/Interchip Enhanced Cooling (ICECool) program (2012).

As continued increases in clock frequency on a single core would not be sustainable, the mid-2000s saw a shift to a new paradigm with the introduction of multicore architectures. This provided a way to constrain power dissipation without constraining performance. Multicore architectures eased the cooling challenges and reduced hot spots, leading to more uniform power maps. The rate at which the power dissipation increased also fell. For instance, at an approximately constant dissipation of 130 W, several versions of the Intel Xeon processor with increasing number of cores from 2 to 10 were able to increase their performance at 75% per year between 2007 and 2012 [50].

A number of technological thermal management challenges identified by Intel are:

- For high performance computing (HPC) systems: combining improved multichip packaged (MCP) components with high-performance, rack-level liquid cooling solutions (30-100 kW/rack). Multichip packaging is important to minimize the power spent in transmitting signals by means of shorter optimized interconnects. [51]
- For client systems and cloud servers: lower thermal resistance for MCP processors needed with at least a 2.5-fold improvement over the current junction-to-case resistances of $\sim 0.1\text{-}0.2\text{ K cm}^2/\text{W}$, and TIMs with an increased compliance to accommodate die height variations in MCP processors.
- For wearables and portable devices: flexible thermal

interface materials (TIMs), and more options for reducing skin temperature other than throttling performance.

In the face of these challenges, the past decades have seen good collaboration between academia and industry, and cooling demands have hitherto been met in all cases. As the semiconductor industry evolves to cater to a wide gamut of devices, thermal management will be needed for a wide range of form factors, powers and power densities. The challenge is for the thermal management community to develop a portfolio of building block solutions in anticipation of this demand.

1. Large-scale electronics systems

A significant fraction (about 90%) of the total energy used in computing facilities goes to moving data around. In the absence of a cheap and low-energy medium for storing large amounts of data, current systems are forced to use a series of computing and memory storage devices ranging from CPU cache, random access memory, and local hard drives to network storage. This hierarchy is currently unavoidable but needs a lot of power. Flattening this hierarchy, and removing the multitude of copper interconnects will be a significant challenge [52].

To cope with future demands related to Big Data, the main characteristics of next-generation memory are that it be non-volatile, high density, very low power, and fast. Thermal management of memory encompasses a wide range of time and length scales down to the molecular level, requiring molecular dynamics-based thermal modelling in order to correctly simulate the device's thermal characteristics [53]. Electro-thermal co-design of server architectures is being used to optimize layouts of a high number of dual in-line memory modules (DIMMs). DIMMs are traditionally not considered the thermal bottleneck in servers but this situation is now changing. For future 3D stacked designs, although the heat dissipation in memory devices is small (~ 100 mW) the thermal resistance they add between the processor and the outer case is important. For a 100 W processor, the maximum junction temperature can be reached with a single added memory die on top; thus embedded cooling within the package will likely be required for 3D chips [54].

Future high-speed data movement could see the integration of CMOS and photonic devices, with the added challenges of temperature-sensitivity of the latter. HP identifies the opportunities with the highest potential return as being at the nanoscale [9]: saving energy at the nanometer-scale offers the greatest return for the overall system.

2. Automotive applications

Hybrid electric vehicle (HEV) powertrains contain a large number of power electronics operating in harsh environments with strong accelerations, fouling and electro-magnetic fields. While allowable die temperatures are generally higher compared to logic devices, a growing power density

requirement makes power electronics a strong emerging market for advanced thermal management solutions. Whereas traditional power electronics use narrow band gap materials (e.g., silicon, gallium arsenide), future devices will use wide band gap materials (e.g., silicon carbide, gallium nitride) which are more efficient and yet operate at higher temperatures ($\sim 300^\circ\text{C}$) [55]. Required heat transfer coefficients for these new technologies exceed $10,000$ $\text{W}/(\text{m}^2\text{K})$, beyond the limits of single-phase water-cooled cold plates [56]. Alternatives being considered are single- or two-phase jet impingement and flow boiling. Solutions that are insensitive to accelerations such as jet impingement are preferred in vehicular applications.

3. Small-scale personal electronics

a) Challenges related to skin temperature

The constraint of greatest interest for mobile phone thermal designers is not the junction temperature but the skin temperature at the outer surface of the device. Designs target a maximum skin temperature of $40\text{-}45^\circ\text{C}$ and the highest possible temperature uniformity (i.e., an average-to-maximum temperature ratio above 0.8), although there are no standards yet for tolerable heat loads on the human body in terms of temperature, contact time, or location. Portable electronics pose particular challenges for thermal design, given the extremely limited physical space available for inserting thermal solutions and the limits on allowable surface temperature. This currently limits power dissipation in tablet computers to $7\text{-}10$ W and smart phones to $2\text{-}3$ W. Emerging devices such as smart watches amplify the challenge due to their continuous contact with the human body [9].

b) Challenges related to system design

In the design of thermal management controls for mobile electronics, complete system models are used to correctly represent the thermal behavior from the transistor level (microsecond response times) to the device level (response times up to minutes). Different devices have different ergonomic limits and usage times, from seconds for phones to minutes for tablets. These thermally aware systems contain several sensors for predictive temperature control, e.g., by throttling performance [57]. Internal heat conduction and capacitance is optimized using TIM layers, but the layout remains a challenge for architectures with memory on top of a CPU, resulting in a restricted thermal pathway to the outside. In mobile phones, a system-on-a-chip (SoC) package contains a processor, memory, mixed-signal and radio-frequency functions in a single package. Some challenges include the doubled heat flux in two-die stacks and the need to satisfy different thermal constraints in a single package by stacking memory on top of logic processors.

C. Acoustic noise emission challenges

1. Acoustic noise sources in electronics systems

In some application areas including personal electronics,

acoustic noise emissions are growing with the use of smaller, higher-speed fans producing strong tonal noise. Best practices are not always followed when designing air flow pathways. Active noise cancellation is impractical in most instances due to the distributed nature of aero-acoustic noise sources in turbulent boundary and mixing layers. Scaling laws for rotary fans can express the characteristic flow rate Q , pressure head Δp and acoustic noise emission L_w in terms of fan speed N and diameter D , as $Q \propto ND^3$, $\Delta p \propto N^2D^2$, and $L_w \propto 10 \log_{10}(N^n D^8)$ [58]. The exponent n depends on the fan type but generally $n = 6$ [9]. Given the current trend of downsizing, to maintain the same flow rate the fan speed should increase as $N \propto 1/D^3$, causing an increased noise level $L_w \propto 10 \log_{10}(1/D^{10})$. Furthermore, the fan blade pass frequency (= number of blades $\times N$) also increases. If this highly tonal noise reaches 1-5 kHz, the fan will be perceived as annoying. For portable electronics, end-user surveys show that sharp tonality (even at lower dB levels) is perceived as very to extremely annoying [9].

Whereas rigid health and safety guidelines exist for noise emissions levels [59,60], these are only expressed as an overall averaged dB value and do not account for variations in the frequency spectrum. These criteria are appropriate for the classical situation of a person standing in the vicinity of a stationary device with a constant noise emission. By contrast, the psychoacoustic or ergonomic noise limitations for portable electronic devices are not static but evolving in time. The perception is the combined result of ‘sensation’ and ‘expectation’, where the expectation cost is not well understood. What the end user expects of a device, such as how hot they expect a phone to feel compared to a cup of coffee, plays an important role. The environment also plays a role: in environments with high background noise such as an airplane, users would tolerate more noise from a laptop fan. Based on an internal study at Intel, some challenges in ergonomics have not seen much change, while others are rapidly evolving. For instance, preconceptions about phones are moving quickly, and users are ready to accept different ideas, but that is not true in all fields. For design engineers, this means that a certain degree of fluidity should be allowed in requirements, which will require a change in mindset [9].

2. Aero-acoustic modeling and validation

Developing quieter air cooling methods requires insight into aeroacoustics in confined geometries. Flow-induced noise simulations can be used as a design tool but need experimental validation. For noise measurements, near-field acoustical holography (NAH) [61] or acoustical beamforming [62] relies on a microphone array to measure the spatially and phase-resolved sound pressure field. Numerical aeroacoustics can take several approaches: brute force direct numerical simulation (DNS) is impractical for routine electronics cooling design purposes. Several hybrid methods exist to separately solve the core flow with unsteady computational fluid dynamics (CFD) and the acoustic field on a larger domain, such as the method based on the Lighthill acoustic analogy

and the Ffowcs-Williams and Hawkings equation [63]. For electronics cooling configurations, the latter can be further simplified with negligible loss of accuracy. To improve the near-wall treatment of eddies which are important for the noise source modelling, the large eddy simulation (LES) and unsteady Reynolds-averaged Navier-Stokes (RANS) approaches can be combined as detached eddy simulation (DES) [64]. Other techniques such as the lattice Boltzmann method (LBM) [65] could be used as well.

IV. IMPLEMENTATION OF ADVANCED THERMAL MANAGEMENT SOLUTIONS

A. Heat conduction, spreading and storage

1. Developments in high conductivity substrate materials

For high heat flux applications (*e.g.*, high-end servers, transportation, and defense), work is ongoing to develop near-junction high thermal conductivity materials and interfacing techniques to facilitate targeted heat conduction and spreading. The heat fluxes range from ~ 1 MW/cm² for nanoscale IC heat sources at ~ 1 nm length scales to ~ 1 W/cm² for IC packages at ~ 1 mm length scales [9]. For heat spreading applications, diamond has an exceptionally high conductivity. Breakthroughs in the deposition of synthetic diamond on silicon were made by Mercedes in the late 1990s [66], which led to an uptick in follow-up publications since the mid-2000s [67]. Carbon-based materials are being considered for production of solutions with a high conductivity and low density [68]. However graphene and carbon nanotubes only feature a high conductivity in one or two directions, whereas diamond offers high 3D conductivity.

Two examples of extreme limits for electronics cooling are (i) 3D integrated packages with stacked memory and logic chips for HPC, and (ii) high-electron-mobility wide bandgap transistor electronics (*e.g.*, using GaN, GaAs, AlGaAs, InGaAs) [69] for radar and telecommunication applications. These wide bandgap electronics are facing extremely high (> 10 kW/cm²) local fluxes or > 1 kW/cm² averaged over the chip. Techniques are being developed for co-fabrication of GaN with a 30 nm-thick intermediate amorphous layer of SiN or SiC bonded onto a diamond substrate [70]. There is an evolution towards substrates of lower thermal resistance, using different material combinations to different extents (*e.g.*, GaN on silicon, GaN on SiC, GaN on diamond).

2. Heat spreading and storage in mobile devices

Mobile phone thermal management approaches typically include both software- and hardware-based solutions. On the software side, static or adaptive algorithms are used to control power dissipation by changing both the processor clock frequency and voltage, thereby throttling performance as needed. An adaptive system tries to predict if more power will be needed, and adjusts the load accordingly [9]. On the hardware side, mobile phone manufacturers require efficient

heat spreading within a highly constrained thickness to mitigate hot spots and ensure surface temperature uniformity on the outer skin of the device. Internal air gaps can be used strategically to limit heat flow to the outer surface and thereby avoid surges in the skin temperatures. The thermal inertia of the battery and frame are also used as a buffer.

Compact thermal models can be used for design and operational control, but the models must account for the thermal coupling between the battery and processor [71]. Power consumption in a mobile phone processor comprises a static component (due to leakage and standby current) and a dynamic component (due to capacitance current). To handle the dynamic response, energy storage is typically used in the form of thermal mass of the battery and mechanical components, but phase-change materials (PCM) have also been proposed for use. However, PCMs typically have low thermal diffusivity making it difficult to remove heat at high fluxes. Graphite sheets used for heat spreading are only 10-100 μm thick; replacement of their volume with PCMs (in small quantities) could only be part of the solution, because it would only take a few seconds to heat up and melt.

Thermal management solutions for mobile phones have evolved in steps: when the Apple iPhone was introduced in the mid-2000s, performance throttling and a TIM on a metal surface were sufficient. Since then, a thermal management unit with multiple sensors and graphite sheet heat spreaders has become necessary [72]. Now, these thermal management units have become more advanced, with multiple embedded sensors. While it is unlikely that active liquid cooling will appear in mobile phones, heat pipes are now being adopted to provide more efficient spreading (*e.g.*, SonyXperiaZ5, Samsung Galaxy S7, and others), which was unimaginable only a short time ago. These solutions typically place a heat pipe near the edge of the phone [73], likely due to severe thickness limitations in the middle of the frame (less than ~ 0.4 mm); much thinner vapor chamber technologies are required to place these solutions closer to the heat source.

While vapor chambers may offer a viable solution for passive spreading within mobile devices, conventional design approaches and performance metrics have focused on increasing the maximum achievable heat load and power density. The requirements for mobile thermal management are starkly different, and call for heat spreading efficiency to meet ergonomic constraints at ultra-thin form factors. The key heat transfer mechanisms governing the performance of vapor chambers with shrinking thickness have been recently identified in order to delineate the performance thresholds within which the effectiveness of a vapor chamber, as a function of geometry and heat input, is greater than that of a comparable solid heat spreader [74]. For ultra-thin vapor chambers operating at a low power, the heat spreading resistance is dictated by the vapor-phase behavior; design approaches and working fluid selection criteria must be revised with these considerations.

Existing ultra-thin heat spreader characterization approaches

do not reflect the mobile thermal management application environment or prioritize the ergonomic considerations. A new characterization approach has been recently developed to assess the behavior of heat spreaders in a simulated scenario of heat rejection to the ambient via natural convection [75]. The surface temperature distribution is measured directly, and performance metrics have been proposed to characterize heat spreader performance in terms of the surface temperature uniformity. The testing methodology can be used as a tool for the assessment of vapor chambers and heat spreaders intended for use in portable electronics platforms.

B. Convective air cooling

Recent reviews of current server cooling techniques for large-scale electronics systems and data centers are presented by Kheirabadi and Groulx [76] and Ebrahimi *et al.* [77], among others. As noted in a previous review by the authors [1], current air cooling technology can still be further optimized. As discussed in connection with the acoustic noise emission challenges in Section III.C, although fans may be well-designed components, the air ducting is not always implemented according to best practices. Closer collaboration of thermal and aeroacoustics experts with system architecture designers is crucial to extend the life of air cooling in servers.

As discussed in Section II.B.4, there are various limitations to air cooling, indicating that a gradual transition to liquid cooling may be imminent (see Section III.A.2). Some companies involved in large-scale electronics systems remain reluctant to make this transition. Moreover the gamut of electronics cooling applications continues to proliferate towards many different smaller platforms (see Section II.B.2), for which liquid cooling will remain impractical. Many of these smaller platforms currently rely on radiation and natural convection as the ultimate heat sink, yet the higher-end devices may soon resort to forced air convection.

Fan-based air cooling has proven very resilient in the face of growing thermal demands. Advances in alternative air-moving techniques (such as the ones described below) may gain importance in the coming years for the following situations: (i) to enhance natural convection without resorting to rotary fans for small-scale platforms, or (ii) to supplement fans in the cooling of server components with moderate heat fluxes *e.g.*, alongside liquid cooled CPUs in hybrid-cooled servers.

- Piezoelectric fans [27,78,79,80] have proven a viable cooling solution in controlled environments. Their main advantages are low cost, low power consumption, low noise and good reliability. However in harsh environments, dust fouling and accumulation on the blade changes the resonance frequency of the oscillating fan. In this case vibration frequency tracking can provide a partial solution, changing the actuation frequency based on a measurement of the free resonance frequency [81]. Good thermal performance is observed in the moderate heat flux range of 0.1-1 W/cm^2 , making these most suitable for low-power platforms (below 40 W) and thereby bridging the overlap

zone (15-40 W) between natural and forced convection heat sinks.

- Synthetic jets could also provide solutions in this moderate heat flux range between natural and forced convection, although with properly integrated design of the actuator and ducting, these devices can achieve significant local cooling comparable to steady impinging jets [82] in the heat flux range $0.2\text{-}5\text{ W/cm}^2$. Since both synthetic jets and piezo fans are essentially agitators, the overall air path should be carefully designed to avoid recirculating hot air. A synthetic jet can be quieter than an axial fan for the same or better thermal performance [83]. Furthermore, a pair of phase-controlled adjacent jets can be used to control the angle of the jet flow [84], thus providing opportunities for active cooling control. Synthetic jets (and piezo fans) can be used in combination with low speed fans to decouple local heat transfer from thermal advection due to bulk air flow. With the introduction of compact liquid cooled heat sinks on server CPUs instead of high pressure drop air cooled heatsinks, the overall pumping power budget may shift dramatically to enable synthetic jets and piezo fans to obtain a foothold in server cooling. In data centers this could reduce maintenance cost and increase energy efficiency.
- Other alternative air movers are being studied as well; from bladeless fans to avoid the blade passage-related tonal noise to electro-aerodynamics for small platforms ($< 3\text{ mm}$) for which bearings are unsuitable. The latter could use low-voltage piezoelectric transformers to generate a corona discharge in air [85].

Other opportunities to enhance air cooling include (i) advanced flow simulation and optimization tools, and (ii) combining numerical shape optimization techniques with additive manufacturing of validation prototypes [86,87] such as selective laser sintering of aluminum alloys (*e.g.*, AlSi₁₂).

C. Convective liquid cooling

Since liquids (and water in particular) have a much greater volumetric heat capacity than air, component-level liquid cooling remains an active research area. Recent research focuses on targeted liquid chip cooling and heat transfer enhancement [88,89,90]. This section does not provide an exhaustive research overview, but highlights selected topics related to the implementation of single- and two-phase liquid cooling.

1. Single-phase liquid cooling

Current end users of single phase liquid cooling make up a diverse group including gaming enthusiasts and HPC facility operators, *e.g.*, in universities with limited floor space. For instance, rack-level liquid cooling was installed in an academic supercomputing center in Poznan, Poland [91] using Huawei E9000 chassis with CH121 nodes (30 kW/rack, 910 Gflops/node). CPU cold plates achieve a thermal resistance of 0.03 K/W, with secondary cold plates cooling the DIMMs. Hard plastic corrugated tubing is used to facilitate tight

bending radii. The fittings are cold-inserted into the tubing which provides a robust seal. Custom-designed quick disconnect dry-break plugs are used at the manifold [92]. These are reliable but expensive, and the dry-break technology requires coolant filtration to prevent fouling on the quick disconnect couplings. Interestingly, the system filtration demands are not determined by the microchannel cold plates but by these quick disconnect couplings. Per rack, more than 40 kW of heat can be dissipated at a pressure drop of about 14 kPa at 500 mL/min and 40°C inlet water temperature [9].

More expensive dielectric fluids may be necessary in case of direct contact between the coolant and electronic circuits (*e.g.*, immersion server cooling [93]). Usually however, water or glycol water mixtures are preferred due to their low cost, excellent heat transfer characteristics and established use in other engineering applications [2]. A typical coolant is a propylene glycol water mixture with additives specially designed to prevent corrosion and biological growth. With the right choice of inhibitors, the coolant can be used without problems in combination with dissimilar metals including copper, aluminum, stainless steel, etc. [92].

2. Two-phase liquid cooling

Two-phase liquid cooling remains an active area in both academic and industrial research and development. This section covers a selection of ongoing work.

An example of industrial research and development in this area is by Toyota Research Institute of North America for thermal management of power electronics in HEVs. Unlike other two-phase cooling approaches, jet impingement is insensitive to the orientation of gravity (or strong accelerations) in automotive environments.

A two-phase cooling test facility was developed [94] for single and multiple jets impinging on a copper heat spreader, using a fluorocarbon working fluid (3MTM NOVECTM 7100). Both smooth and finned impingement surfaces were tested. For a system pressure drop below 1 kPa, the finned surface gives the best heat transfer performance. For two-phase jets impinging on a finned surface, the averaged heat transfer coefficient is 5.5 times higher with 50% lower pressure drop compared to single-phase operation. Alternative porous surfaces [95] featuring open or closed tunnel-like passages [96] and pin-fin designs [97] dissipated heat fluxes of up to 100 W/cm^2 with heat transfer coefficients of up to $70,000\text{ W/(m}^2\text{K)}$. Using R245fa as working fluid, a small-scale two-phase system is capable of handling heat fluxes up to 200 W/cm^2 and heat transfer coefficients up to $100,000\text{ W/(m}^2\text{K)}$ for pin-fin surfaces.

In single-phase cooling systems, an increase to such heat transfer coefficients would require a significant reduction of the hydraulic diameter of the cooling channels. The use of such small structures would drive up the cost significantly and increase the risk of clogging. Two-phase operation can provide a solution without the need for such small-scale flow

channels, although the overall system cost could be higher because of the added complexity of a two-phase flow loop.

As part of the DARPA ICECool Fundamentals program, a research team at Purdue University is dissipating high heat fluxes by feeding an array of high-aspect-ratio microchannel heat sinks in parallel with HFE-7100 via a hierarchical manifold for fluid distribution. Two performance-enabling aspects of this design are the (1) fabrication of suitably high aspect ratio microchannels etched into the silicon test chip to provide the necessary surface area enhancement to dissipate the target heat fluxes at an allowable surface temperature rise, and (2) parallelization of flow across an array of short-flow-length microchannel heat sink elements to minimize the pressure drop.

A hierarchical manifold microchannel heat sink array test vehicle, with all flow distribution components heterogeneously integrated, has been fabricated to demonstrate thermal and hydraulic performance of this technology [98]. High-aspect-ratio microchannels are etched into a silicon chip (with heaters and sensors embedded directly on the back side) and thermo-compression bonded to a silicon plenum plate. The heated chip area is cooled by a 3×3 array of microchannel heat sinks that are fed with HFE-7100 through the plenum using a hierarchical manifold. Background heat fluxes of up to 500 W/cm^2 and hotspot fluxes of greater than 2500 W/cm^2 are simultaneously dissipated at an average chip temperature $\sim 30 \text{ }^\circ\text{C}$ above the fluid inlet temperature [98]; background heat fluxes alone of up to 800 W/cm^2 have been dissipated at $\sim 60 \text{ }^\circ\text{C}$ [9].

Flow boiling instabilities that induce maldistribution are a challenge for implementation of two-phase microchannel heat sinks [99,100]. Several geometric enhancements have been proposed to suppress instabilities, such as inlet restrictions [101], channel tapering [102] and the use of surface roughness [103]. Recent work at Stanford has focused on vapor/liquid phase separation, *i.e.*, vapor extraction from a microchannel flow using permeable membrane surfaces [104] or heat pipe-inspired separation using phase-separator coatings on a porous medium which keeps liquid contained within structure by surface forces, yet allows evaporation at the boundaries. One remaining challenge with membrane vapor separation is membrane clogging. Recently published work has demonstrated capillary-fed porous copper structures capable of dissipating over 1200 W/cm^2 in boiling with water, at a low superheat not exceeding 10 K at maximum dissipation rates [105].

Nucleate boiling is an efficient mode of heat transfer with applications throughout the industrial and power generation sector. Its main limitations are susceptibility to a larger wall temperature excursion prior to bubble incipience and at critical heat flux (CHF). The goals for boiling heat transfer enhancement are (i) to promote onset of nucleate boiling at low wall superheat, (ii) to increase the slope of the boiling curve, *i.e.*, increase the boiling heat transfer coefficient, and (iii) to increase CHF.

Surface wettability can be used to affect nucleate boiling, but with different effects: A hydrophobic surface promotes nucleation whereas a hydrophilic surface promotes rewetting and thus increases CHF. Hybrid surface treatments which include both hydrophobic and hydrophilic areas can be manufactured by chemical or physical means [106,107]. Results are promising but fabrication is complicated, and the wetting distribution is fixed while boiling is a highly transient phenomenon.

One way to actively control boiling relies on electro-wetting (EW), in which the surface wettability is controlled by applying an electric field. Most literature on EW is on the control of droplets [108,109] yet its effect on bubbly flows is not yet fully understood [110,111,112,113]. Preliminary experiments at the University of Houston have demonstrated that the bubble dynamics can be effectively controlled by EW and that nucleate boiling heat transfer can be favorably improved over the entire range of boiling regimes. For a pool boiling experiment with a hydrophobic heated surface, a $2\times$ increase in CHF and a 50% increase in boiling heat transfer coefficient were observed in the presence of the applied electric field, compared to the case without EW actuation [9]. Electrowetting in combination with hydrophobic surface treatment could be considered for other applications in high heat flux thermal management, for instance to control two-phase instabilities in parallel microchannel heatsinks [114].

D. Waste energy recuperation

The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) Technical Committee 9.9 issued guidelines in 2004 and 2008 for the operation of data centers. These early guidelines had quite a narrow recommended room temperature range ($18\text{-}27^\circ\text{C}$) and an allowable maximum temperature of only 35°C [115]. The 2011 ASHRAE regulations [24] have widened the operating temperature envelope to a maximum allowable room temperature of 45°C . The greater temperature difference with the outside air facilitates so-called free cooling using external ambient air as heat sink without the need for computer room air conditioning (CRAC) units. Energy reuse has become an important aspect, and various technologies are being considered to upgrade low-grade waste heat from large-scale electronic systems, such as organic Rankine cycles (ORC), fuel cells, and absorption chillers [1,9]. Depending on the exhaust water temperature, different technologies can be used as heat sink. For water at 50°C , free cooling seems the only viable option. To avoid mechanical chillers, evaporative-assisted dry coolers are preferred. For water above 65°C , room heating can be considered. Electricity generation only becomes viable for water above 80°C [9].

An example of a multi-stage waste energy recuperation strategy for liquid-cooled ICT equipment was presented at the 2015 IMAPS Thermal Workshop [116], comprising fuel cells, adsorption chillers, and evaporative-assisted air-side heat exchangers. This was a demonstration of how to recuperate

useful energy from an ICT coolant outlet temperature of only 55°C using a combination of the aforementioned technologies.




Figure 3. Word cloud visualization of the key areas discussed in each topic during the Q&A discussions following presentations at the 3rd Workshop on Thermal Management in Telecommunication Systems and Data Centers, held in Redwood City, California on November 4-5, 2015 (image generated by wordclouds.com).

V. SUMMARY AND OUTLOOK

This paper summarizes discussions among thermal management experts from a range of industries spanning large-scale electronic systems (*e.g.*, data centers, HPC) to small-scale systems and devices (*e.g.*, automotive, personal handhelds).

Thermal design for portable electronics, with its strong emphasis on customer ergonomics, requires a quite different from the traditional approaches in large-scale, higher heat flux electronics cooling. With trends pointing to increased performance and smaller form factors, design for portables increasingly relies on – indeed demands – a closer integration of electrical and thermal engineering teams.

In large-scale system design, next generation cooling technologies are sought to increase computational performance and decrease energy usage, while increasing reliability and reducing acoustic noise emission. While liquid

cooling could provide reliable solutions for many of these challenges, air cooling remains resilient and the first choice. Research and development teams have been resourceful in developing novel solutions (*e.g.*, piezo fans, synthetic jets, bladeless fans, electro-aerodynamics) to augment or replace traditional rotary fan-based architectures.

With high-end portable electronics such as tablet computers pushing the limits of natural convection and radiation as ultimate heat sinks, some of these novel air-moving techniques could find acceptance in a range of emerging platforms.

A closer integration of multidisciplinary design teams seems to emerge as a common theme for the thermal management of next generation systems and devices, whether at large or small scale [117]. This holds true across the transistor level (material scientists, electro-thermal design engineers), package level (thermo-mechanical and thermo-fluid design), server board and rack level (aeroacoustics and fluid dynamics) and system integration level (thermodynamics, energy policy and techno-economics). Cross-communication between engineers to bridge these notional interfaces is critical to keep pace with the foreseeable thermal management challenges in both large- and small-scale electronics systems.

VI. ACKNOWLEDGEMENTS

The authors gratefully acknowledge the contributions of the participants and speakers at the 3rd Workshop on Thermal Management in Telecommunication Systems and Data Centers held in Redwood City, California on November 4-5, 2015.

VII. GLOSSARY OF ABBREVIATIONS AND ACRONYMS

ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers
CAGR	compound annual growth rate
CapEx	capital expenditure
CDU	coolant distribution unit
CHF	critical heat flux
COP	coefficient of performance
CPU	central processing unit
CRAC	computer room air conditioning
DARPA Agency	Defense Advanced Research Projects Agency
DIMM	dual in-line memory module
EW	electro-wetting
flops	floating point operations per second
HEV	hybrid electric vehicle
IC	integrated circuit

ICECool (DARPA program)	Intrachip/Interchip Enhanced Cooling
ICT	information and communications technology
IoT	Internet of Things
MCP	multichip packaged (components)
OpEx	operational expenditure
PCB	printed circuit board
PCM	phase-change material
PUE	power utilization effectiveness
RRH	remote radio head
SoC	system-on-a-chip
TCO	total cost of ownership
TDTR	time-domain thermo-reflectance
TIM	thermal interface material
TSV	through-silicon via
TTR	transient thermo-reflectance
VOF	volume of fluid

VIII. REFERENCES

- 1 S. V. Garimella, T. Persoons, J. Weibel, and L.-T. Yeh, "Technological drivers in data centers and telecom systems: Multiscale thermal, electrical, and energy management," *Appl. Energy*, vol. 107, pp. 66-80, 2013.
- 2 S. V. Garimella, L.-T. Yeh, and T. Persoons, "Thermal Management Challenges in Telecommunication Systems and Data Centers," *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 2, pp. 1307-1316, 2012.
- 3 H. N. Zhang, S. Q. Shao, H. B. Xu, H. M. Zou, and C. Q. Tian, "Integrated system of mechanical refrigeration and thermosyphon for free cooling of data centers," *Appl. Therm. Eng.*, vol. 75, pp. 185-192, 2015.
- 4 S. W. Ham, M. H. Kim, B. N. Choi, and J. W. Jeong, "Energy saving potential of various air-side economizers in a modular data center," *Appl. Energy*, vol. 138, pp. 258-275, 2015.
- 5 J. Gantz, and D. Reinsel, *The Digital Universe in 2020: Big Data, Bigger Digital Shadows, and Biggest Growth in the Far East*, International Data Corporation (IDC), sponsored by EMC, Dec. 2012.
- 6 L. D. Stein, B. M. Knoppers, P. Campbell, G. Getz, and J. O. Korbel, "Data analysis: Create a cloud commons," *Nature*, vol. 523, no. 7559, pp. 149-151, 2015.
- 7 Uptime Institute, *2014 Data Center Industry Survey*, Retrieved from <http://journal.uptimeinstitute.com/2014-data-center-industry-survey/> on Dec. 9, 2015.
- 8 Y. Sekiguchi, "The Plant Technology of the K Computer," *J. Heat Transfer Soc. Japan*, vol. 52, no. 220, pp. 27-32, 2013 (in Japanese).
- 9 Communicated in technical presentations and discussions of participants at the 3rd *Workshop on Thermal Management in Telecommunication Systems and Data Centers*, Redwood City, California, Nov. 4-5, 2015.
- 10 *Intrachip/Interchip Enhanced Cooling Fundamentals*, Solicitation No.: DARPA-BAA-12-50, 2012.
- 11 J. Morgan, *IBM unveils computer fed by 'electronic blood'*, BBC, Retrieved from <http://www.bbc.com/news/science-environment-24571219> on Dec. 12, 2015.
- 12 T. P. Morgan, *Intel Stacks Knights Landing Chips Next To Xeons*, thenextplatform.com, Retrieved from <http://www.nextplatform.com/2015/08/28/intel-stacks-knights-landing-chips-next-to-xeons/> on Dec 7, 2015.
- 13 E. Gardner, *What public disclosures has Intel made about Knights Landing?*, Intel Developer Zone, Retrieved from <http://software.intel.com/en-us/articles/what-disclosures-has-intel-made-about-knights-landing> on Dec 7, 2015.
- 14 D. Hernon, and R. Enright, "Modeling: Photonics cooling meets high-speed-communications energy demands through simulation," *LaserFocusWorld*, Nov. 4, 2014, Retrieved from <http://www.laserfocusworld.com/articles/print/volume-50/issue-11/features/modeling-photonics-cooling-meets-high-speed-communications-energy-demands-through-simulation.html> on Dec. 12, 2015.
- 15 R. Enright, S. Lei, K. Nolan, I. Mathews, A. Shen, G. Levaufre, R. Frizzell, G.-H. Duan, and D. Hernon, "A Vision for Thermally Integrated Photonics Systems," *Bell Labs Tech. J.*, vol. 19, pp. 31-45, 2014, Retrieved from <http://www.tips2020.eu/uploads/Enright-et-al.-2014-BLTJ-TIPS-vision1.pdf> on Dec. 12, 2015.
- 16 J. Mukherjee, S. Jarvis, M. Perren, and S. Sweeney, "Efficiency limits of laser power converters for optical power transfer applications," *J. Phys. D: Appl. Phys.*, vol. 46, no. 26, art. 264006, 2013.

- 17 InformationWeek, *Big Data Drives The Smart Car*, Retrieved from <http://www.informationweek.com/big-data/big-data-analytics/big-data-drives-the-smart-car/d/d-id/1127767> on Dec. 12, 2015.
- 18 J. Healey, *If cars could talk, accidents might be avoidable*, TED Talk, Intel, Retrieved from http://www.ted.com/talks/jennifer_healey_if_cars_could_talk_accidents_might_be_avoidable on Dec. 12, 2015.
- 19 P. Moynagh, YouTube, Retrieved from <http://www.youtube.com/watch?v=C8NMRpZb6Ss>, accessed on Dec 7, 2015.
- 20 *Press release*, Gartner, Stamford, Connecticut, March 19, 2014, Retrieved from <http://www.gartner.com/newsroom/id/2684616> on Dec 7, 2015
- 21 H. Bauer, M. Patel, and J. Veira, *The Internet of Things: Sizing up the opportunity*, McKinsey & Company, Frankfurt, Germany, Retrieved from http://www.mckinsey.com/insights/high_tech_telecoms_internet/the_internet_of_things_sizing_up_the_opportunity on Dec 7, 2015.
- 22 T. P. Morgan, *Future Systems: Intel Fellow Conjures The Perfect Exascale Machine*, thenextplatform.com, Retrieved from <http://www.nextplatform.com/2015/08/12/future-systems-intel-fellow-conjures-the-perfect-exascale-machine/> on Dec 7, 2015.
- 23 K. Sekar, "Power and thermal challenges in mobile devices," 19th *Annu. Int. Conf. Mobile Comput. Network.* (MobiCom), Miami, Florida, Sep. 30-Oct. 4, 2013.
- 24 ASHRAE, *Whitepaper: 2011 Thermal guidelines for data processing environments – expanded data center classes and usage guidance*, ASHRAE Technical Committee 9.9, 2011, Retrieved from http://ecoinfo.cnrs.fr/IMG/pdf/ashrae_2011_thermal_guidelines_data_center.pdf on Dec. 13, 2015.
- 25 R. Ungar, and M. Thiry, *Overcoming hidden costs in the data center*, ABB White Paper, Jan. 2013, Retrieved from <http://abb-decathlon-dcim.tnbpowersolutions.com/overcoming-hidden-costs/> on Aug. 18, 2016.
- 26 P. Lin, J. Niemann, and L. Long, *APC White Paper 215: Choosing between direct and indirect air economization for data centers*, Schneider Electric, Aug. 2015, Retrieved from <http://www.apc.com/wp?wp=215> on Aug. 18, 2016.
- 27 M. Kimber, K. Suzuki, N. Kitsunai, K. Seki, and S. V. Garimella, "Pressure and flow rate performance of piezoelectric fans," *IEEE Trans. Compon. Packag. Technol.*, vol. 32, no. 4, pp. 766-775, 2009.
- 28 P. de Bock, A. Mann, and B. Whalen, "Evaluation and opportunities for use of thin form factor synthetic jets to low profile electronics cooling applications," *Electronics Cooling*, December 2012.
- 29 A. J. Shah, C. E. Bash, N. Kumari, T. Cader, T. J. Breen, E. Walsh, and J. Punch, "On the need for energy efficiency metrics that span integrated IT-facility infrastructures," *Proc. ASME 2011 Pacific Rim Tech. Conf. Exhibition Packag. Integration Electron. Photon. Syst. (InterPACK)*, IPACK2011-52140, 2011.
- 30 R. Miller, *Rise of Direct Liquid Cooling in Data Centers Likely Inevitable*, Data Center Knowledge, Retrieved from <http://www.datacenterknowledge.com/archives/2014/12/09/rise-direct-liquid-cooling-data-centers-likely-inevitable/> on Dec. 11, 2015.
- 31 E. Pfanner, *Liquid-Cooled Supercomputers, to Trim the Power Bill*, The New York Times, Feb. 11, 2014, Retrieved from http://www.nytimes.com/2014/02/12/business/international/improving-energy-efficiency-in-supercomputers.html?_r=0 on Apr. 14, 2016.
- 32 R. S. Prasher, "Surface chemistry and characteristics based model for the thermal contact resistance of fluidic interstitial thermal interface materials," *J. Heat Trans.*, vol. 123, no. 5, pp. 969-975, 2001.
- 33 R. S. Prasher, J. Shipley, S. Prstic, P. Koning, and J. L. Wang, "Thermal resistance of particle laden polymeric thermal interface materials," *J. Heat Trans.*, vol. 125, no. 6, pp. 1170-1177, 2003.
- 34 J. Xu, and T. S. Fisher, "Enhancement of thermal interface materials with carbon nanotube arrays," *Int. J. Heat Mass Trans.*, vol. 49, no. 9-10, pp. 1658-1666, 2006.
- 35 M. A. Panzer, G. Zhang, D. Mann, X. Hu, E. Pop, H. Dai, and K. E. Goodson, "Thermal properties of metal-coated vertically aligned single-wall nanotube arrays," *J. Heat Trans.*, vol. 130, no. 5, art. 052401, 2008.
- 36 D. G. Cahill, K. E. Goodson, and A. Majumdar, "Thermometry and thermal transport in micro/nanoscale solid-state devices and structures," *J. Heat Trans.*, vol. 124, no. 2, pp. 223-241, 2002.

- 37 P. Kim, L. Shi, A. Majumdar, and P. L. McEuen, "Thermal transport measurements of individual multiwalled nanotubes," *Phys. Rev. Lett.*, vol. 87, no. 21, art. 215502, 2001.
- 38 D. Liu, and S. V. Garimella, "Analysis and optimization of the thermal performance of microchannel heat sinks," *Int. J. Num. Methods Heat Fluid Flow*, vol. 15, no. 1, pp. 7-26, 2005.
- 39 S. Bhopte, D. Agonafer, R. Schmidt, and B. Sammakia, "Optimization of data center room layout to minimize rack inlet air temperature," *J. Electron. Packag.*, vol. 128, no. 4, pp. 380-387, 2006.
- 40 I. Mudawar, "Assessment of high-heat-flux thermal management schemes," *IEEE Trans. Compon. Packag. Technol.*, vol. 24, no. 2, pp. 122-141, 2001.
- 41 S. G. Kandlikar, and W. J. Grande, "Evolution of microchannel flow passages - Thermohydraulic performance and fabrication technology," *Heat Trans. Eng.*, vol. 24, no. 1, pp. 3-17, 2003.
- 42 P. S. Lee, and S. V. Garimella, "Investigation of heat transfer in rectangular microchannels," *Int. J. Heat Mass Trans.*, vol. 48, no. 9, pp. 1688-1704, 2005.
- 43 A. Bar-Cohen, and E. Rahim, "Modeling and prediction of two-phase microgap channel heat transfer characteristics," *Heat Trans. Eng.*, vol. 30, no. 8, pp. 601-625, 2009.
- 44 L. A. Zhang, E. N. Wang, K. E. Goodson, and T. W. Kenny, "Phase change phenomena in silicon microchannels," *Int. J. Heat Mass Trans.*, vol. 48, no. 8, pp. 1572-1582, 2005.
- 45 D. B. Go, S. V. Garimella, T. S. Fisher, and R. K. Mongia, "Ionic winds for locally enhanced cooling," *J. Appl. Phys.*, vol. 102, no. 5, art. 053302, 2007.
- 46 P. C. Stephan, and C. A. Busse, "Analysis of the Heat-Transfer Coefficient of Grooved Heat Pipe Evaporator Walls," *Int. J. Heat Mass Trans.*, vol. 35, no. 2, pp. 383-391, 1992.
- 47 S. V. Garimella, "Advances in mesoscale thermal management technologies for microelectronics," *Microelectron. J.*, vol. 37, no. 11, pp. 1165-1185, 2006.
- 48 P. E. Phelan, V. A. Chiriack, and T. Y. T. Lee, "Current and future miniature refrigeration cooling technologies for high power microelectronics," *IEEE Trans. Compon. Packag. Technol.*, vol. 25, no. 3, pp. 356-365, 2002.
- 49 R. Chein, G. M. Huang, "Thermoelectric cooler application in electronic cooling," *Appl. Therm. Eng.*, vol. 24, no. 14-15, pp. 2207-2217, 2004.
- 50 Intel Brief, March 2012, Retrieved from <http://www.intel.ie/content/dam/www/public/us/en/documents/performance-briefs/workstation-design-throughput-xeon-e5-2600-brief.pdf> on Apr. 6, 2016.
- 51 S. Chittor, "Memory requirement trends and challenges: Servers to devices," 2015 *IEEE Int. Solid-State Circuits Conf.*, San Francisco, California, Feb. 22-26, 2015.
- 52 HP, *The Machine: A New Kind of Computer*, Hewlett Packard Labs, Retrieved from <http://www.labs.hpe.com/research/themachine> on Apr. 14, 2016.
- 53 S. E. Savel'ev, A. S. Alexandrov, A. M. Bratkovsky, and R. Stanley Williams, "Molecular dynamics simulations of oxide memristors: thermal effects," *Appl. Phys. A*, vol. 102, no. 4, pp. 891-895, 2011.
- 54 N. Kumari, R. Shih, S. Escobar-Vargas, T. Cader, A. Govyadinov, S. Anthony, and C. Bash, "Air Cooling Limits of 3D Stacked Logic Processor and Memory Dies," 2014 *IEEE Intersoc. Conf. Therm. Thermomech. Phenom. Electron. Syst. (ITHERM)*, Orlando, Florida, May 27-30, 2014.
- 55 DOE Advanced Manufacturing Office, *Wide Bandgap Semiconductors: Pursuing the Promise* (DOE/EE-0910), United States Department of Energy, April 2013, Retrieved from http://www1.eere.energy.gov/manufacturing/rd/pdfs/wide_bandgap_semiconductors_factsheet.pdf on Feb. 19, 2016.
- 56 C. J. M. Lasance, "Advances In High-Performance Cooling For Electronics," *Electronics Cooling*, Nov. 1, 2005.
- 57 F. Paterna, J. Zanotelli, and T. S. Rosing, "Ambient variation-tolerant and inter components aware thermal management for mobile system on chips," *Design, Automation & Test in Europe Conference and Exhibition (DATE)*, Dresden, Germany, Mar. 24-28, 2014.
- 58 C. Harris, *Handbook of Acoustical Measurements and Noise Control*, Acoustical Society of America, 1991.
- 59 *OHSA Regulation 1910.95. Occupational Noise Exposure*. Occupational Safety and Health Administration, Dec. 2008.
- 60 *Council Directive 89/391/EEC on the introduction of*

- measures to encourage improvements in the safety and health of workers at work, European Agency for Safety and Health at Work, June 1989.
- 61 E. G. Williams, "Regularization methods for near-field acoustical holography," *J. Acoust. Soc. Am.*, vol. 110, no. 4, pp. 1976-1988, 2001.
- 62 J. Antoni, "A Bayesian approach to sound source reconstruction: Optimal basis, regularization, and focusing," *J. Acoust. Soc. Am.*, vol. 131, no. 4, pp. 2873-2890, 2012.
- 63 K. S. Brentner, and F. Farassat, "Analytical comparison of the acoustic analogy and Kirchhoff formulation for moving surfaces," *AIAA J.*, vol. 36, no. 8, pp. 1379-1386, 1998.
- 64 J. A. Ekaterinaris, "High-order accurate, low numerical diffusion methods for aerodynamics," *Progr. Aero. Sci.*, vol. 41, no. 3-4, pp. 192-300, 2005.
- 65 X. W. Shan, X. F. Yuan, and H. D. Chen, "Kinetic theory representation of hydrodynamics: a way beyond the Navier-Stokes equation," *J. Fluid Mech.*, vol. 550, pp. 413-441, 2006.
- 66 A. Flöter, H. Güttler, G. Schulz, D. Steinbach, C. Lutz-Elsner, R. Zachai, A. Bergmaier, and G. Dollinger, "The nucleation and growth of large area, highly oriented diamond films on silicon substrates," *Diamond & Related Mat.*, vol. 7, no. 2-5, pp. 283-288, 1998.
- 67 Data from ISI Web of Knowledge (<http://apps.webofknowledge.com>), retrieved on Dec. 4, 2015.
- 68 A. M. Marconnet, M. A. Panzer, and K. E. Goodson, "Thermal conduction phenomena in carbon nanotubes and related nanostructured materials," *Rev. Mod. Phys.*, vol. 85, pp. 1295-1326, 2013.
- 69 T. Mimura, "The Early History of the High Electron Mobility Transistor (HEMT)," *IEEE Trans. Microwave Theory Techn.*, vol. 50, no. 3, pp. 780-782, 2002.
- 70 J. Cho., Z. Li, E. Bozorg-Gayeli, T. Kodama, D. Francis, F. Ejeckam, F. Faili, M. Asheghi, and K. E. Goodson, "Improved Thermal Interfaces of GaN-Diamond Composite Substrates for HEMT Applications," *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 3, pp. 79-84, 2013.
- 71 Q. Xie, J. Kim, Y. Z. Wang, D. Shin, N. Chang, and M. Pedram, "Dynamic Thermal Management in Mobile Devices Considering the Thermal Coupling between Battery and Application Processor," 2013 *IEEE/ACM Int. Conf. Computer-Aided Design (ICCAD)*, San Jose, California, pp. 242-247, Nov. 18-21, 2013.
- 72 K.-C. Chen, S.-Y. Lin, and A.-Y. Wu, "Design of thermal management unit with vertical throttling scheme for proactive thermal-aware 3D NoC systems," 2013 *Int. Symp. VLSI Design, Automation, and Test (VLSI-DAT)*, Hsinchu, Taiwan, Apr. 22-24, 2013.
- 73 D. Luis, *The Xperia Z2 has liquid heat-pipe cooling technology*, phonearena.com, Apr. 9, 2014, Retrieved from http://www.phonearena.com/news/The-Xperia-Z2-has-liquid-heat-pipe-cooling-technology_id54935 on Dec. 13, 2015.
- 74 Y. Yadavalli, J. A. Weibel, and S. V. Garimella, "Performance-Governing Transport Mechanisms for Heat Pipes at Ultra-thin Form Factors," *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 5, pp. 1618-1627, 2015.
- 75 G. Patankar, S. Mancin, J. A. Weibel, S. V. Garimella, and M. A. MacDonald, "A Method for Thermal Performance Characterization of Ultra-Thin Vapor Chambers Cooled by Natural Convection," *J. Electron. Packag.*, vol. 138, art. 011003, 2016.
- 76 A. C. Kheirabadi, and D. Groulx, "Cooling of Server Electronics: A Design Review of Existing Technology," *Appl. Therm. Eng.* (in press), doi: 10.1016/j.applthermaleng.2016.03.056.
- 77 K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities," *Renew. Sustain. Energy Rev.*, vol. 31, pp. 622-638, 2014.
- 78 M. L. Kimber, and S. V. Garimella, "Cooling performance of arrays of vibrating cantilevers," *J. Heat Transf.*, vol. 131, no. 11, art. 111401, 2009.
- 79 *Press Release*, GE Global Research, Retrieved from <http://www.geglobalresearch.com/innovation/ges-dual-piezoelectric-cooling-jets-dcj-cool-quiet> on Dec. 13, 2015.
- 80 H. P. de Bock, T. Persoons, and K. Bodla, "Particle image velocimetry study on dual cooling jet flows," 2016 *Intersoc. Conf. Therm. Thermomech. Phenom. Electron. Syst. (ITHERM)*, Las Vegas, Nevada, May 31-Jun. 3, 2016.
- 81 A. Raman, and S. V. Garimella, *Piezoelectric device with feedback sensor*, US Patent 6,713,942 B2, 2004.
- 82 T. Persoons, A. McGuinn, and D. B. Murray, "A

- general correlation for the stagnation point Nusselt number of an axisymmetric impinging synthetic jet,” *Int. J. Heat Mass Trans.*, vol. 54, pp. 3900-3908, 2011.
- 83 R. Mahalingam, S. Heffington, L. Jones, and R. Williams, “Synthetic jets for forced air cooling of electronics,” *Electron. Cooling*, vol. 13, pp. 2-5, 2007.
- 84 E. Fanning, T. Persoons, and D. B. Murray, “Heat transfer and flow characteristics of a pair of adjacent impinging synthetic jets,” *Int. J. Heat Fluid Flow*, vol. 54, pp. 153-166, 2015.
- 85 M. J. Johnson, and D. B. Go, “Piezoelectric transformers for low-voltage generation of gas discharges and ionic winds in atmospheric air,” *J. Appl. Phys.*, vol. 118, art. 243304, 2015.
- 86 E. M. Dede, S. N. Joshi, and F. Zhou, “Topology Optimization, Additive Layer Manufacturing, and Experimental Testing of an Air-Cooled Heat Sink,” *J. Mech. Des.*, vol. 137, no. 11, art. 111403, 2015.
- 87 T. Nomura, E. M. Dede, J. Lee, S. Yamasaki, T. Matsumori, A. Kawamoto, and N. Kikuchi, “General topology optimization method with continuous and discrete orientation design using isoparametric projection,” *Int. J. Num. Methods Eng.*, vol. 101, no. 8, pp. 571-605, 2015.
- 88 C. S. Sharma, M. K. Tiwari, S. Zimmermann, T. Brunschwiler, G. Schlottig, B. Michel, and D. Poulikakos, “Energy efficient hotspot-targeted embedded liquid cooling of electronics,” *Appl. Energy*, vol. 138, pp. 414-422, 2015.
- 89 A. Renfer, M. K. Tiwari, R. Tiwari, F. Alfieri, T. Brunschwiler, B. Michel, and D. Poulikakos, “Microvortex-enhanced heat transfer in 3D-integrated liquid cooling of electronic chip stacks,” *Int. J. Heat Mass Trans.*, vol. 65, pp. 33-43, 2013.
- 90 Y. Y. Wang, and Y. Peles, “An Experimental Study of Passive and Active Heat Transfer Enhancement in Microchannels,” *J. Heat Trans.*, vol. 136, art. 031901, 2014.
- 91 Electronics Cooling, *Second Deployment of Liquid Cooling Solution Scheduled for Supercomputing Center*, Nov. 16, 2015, Retrieved from <http://www.electronics-cooling.com/2015/11/second-deployment-of-liquid-cooling-solution-scheduled-for-supercomputing-center/> on Dec. 13, 2015.
- 92 CoolIT Systems Inc., *Product Guide*, Retrieved from <http://www.coolitsystems.com/index.php/data-center/technology-q-a.html> on Dec. 13, 2015.
- 93 Press release, 3M, St. Paul, Minnesota, Oct. 26, 2015, Retrieved from <http://investors.3m.com/news/press-release-details/2015/3M-Novoc-7100-Fluid-Used-in-Worlds-Largest-Two-Phase-Immersion-Cooling-Project/default.aspx> on Aug. 18, 2016.
- 94 S. N. Joshi, M. J. Rau, E. M. Dede, and S. V. Garimella, “An experimental study of a multi-device jet impingement cooler with phase change using HFE-7100,” 2013 ASME Summer Heat Trans. Conf. (HT2013), Minneapolis, Minnesota, Jul. 14-19, 2013.
- 95 3M, *Microporous Metallic Boiling Enhancement Coating (BEC) L-20227*, Development material, 2009, Retrieved from <http://multimedia.3m.com/mws/media/5635660/3mtm-microporous-metallic-boiling-enhancement-coating-l-20227.pdf>, on Feb. 19, 2016.
- 96 S. N. Joshi, and E. M. Dede, “Thermal Management of Future WBG Devices using Two-phase Cooling,” *Int. Exhib. Conf. Power Electron. Intel. Motion Renew. Energy & Energy Manag. (PCIM)*, Nuremberg, Germany, May 10-12, 2016.
- 97 S. N. Joshi, and E. M. Dede, “Effect of sub-cooling on performance of a multi-jet two phase cooler with multi-scale porous surfaces,” *Int. J. Therm. Sci.*, vol. 87, pp. 110-120, 2015.
- 98 K. P. Drummond, J. A. Weibel, S. V. Garimella, D. Back, D. B. Janes, M. D. Sinanis, and D. Peroulis, “Evaporative Intrachip Hotspot Cooling with a Hierarchical Manifold Microchannel Heat Sink Array,” 2016 Intersoc. Conf. Therm. Thermomech. Phenom. Electron. Syst. (ITHERM), Las Vegas, Nevada, May 31 – Jun. 3, 2016.
- 99 S. G. Kandlikar, “Fundamental issues related to flow boiling in minichannels and microchannels,” *Exp. Therm. Fluid Sci.*, vol. 26, pp. 389-407, 2002.
- 100 L. C. Ruspini, *Experimental and numerical investigation on two-phase flow instabilities*, Doctoral thesis: Norwegian University of Science and Technology, Trondheim, Norway, 2013.
- 101 S. Szczukiewicz, N. Borhani, and J. R. Thome, “Two-phase flow operational maps for multi-microchannel evaporators,” *Int. J. Heat Fluid Flow*, vol. 42, pp. 176-189, 2013.
- 102 S. G. Kandlikar, T. Widger, A. Kalani, and V. Mejjia, “Enhanced Flow Boiling Over Open Microchannels With Uniform and Tapered Gap Manifolds,” *J. Heat Trans.*, vol. 135, no. 6, art. 061401, 2013.

- 103 L. Zhang, E. N. Wang, K. E. Goodson, and T. W. Kenny, "Phase Change Phenomena in Silicon Microchannels," *Int. J. Heat Mass Trans.*, vol. 48, no. 8, pp. 1572-1582, 2005.
- 104 M. P. David, J. Miler, J. E. Steinbrenner, Y. Yang, M. Touzelbaev, and K. E. Goodson, "Hydraulic and thermal characteristics of a vapor venting two-phase microchannel heat exchanger," *Int. J. Heat Mass Trans.*, vol. 54, no. 25-26, pp. 5504-5516, 2011.
- 105 J. W. Palko, C. Zhang, J. D. Wilbur, T. J. Dusseault, M. Asheghi, K. E. Goodson, and J. G. Santiago, "Approaching the limits of two-phase boiling heat transfer: High heat flux and low superheat," *Appl. Phys. Lett.*, vol. 107, art. 253903, 2015.
- 106 A.R. Betz, J. Xu, H. Qiu, and D. Attinger, "Do surfaces with mixed hydrophilic and hydrophobic areas enhance pool boiling?," *Appl. Phys. Lett.*, vol. 97, art. 141909, 2010.
- 107 H. Jo, H. S. Ahn, S. Kang, and M. H. Kim, "A study of nucleate boiling heat transfer on hydrophilic, hydrophobic and heterogeneous wetting surfaces," *Int. J. Heat Mass Transf.*, vol. 54, no. 25-26, pp. 5643-5652, 2011.
- 108 H. Oprins, J. Danneels, B. Van Ham, B. Vandeveld, M. Baelmans, "Convection heat transfer in electrostatic actuated liquid droplets for electronics cooling," *Microelectron. J.*, vol. 39, no. 7, pp. 966-974, 2008.
- 109 E. Baird, and K. Mohseni, K., "Digitized heat transfer: A new paradigm for thermal management of compact micro systems," *IEEE Trans. Compon. Packag. Technol.*, vol. 31, no. 1, pp. 143-151, 2008.
- 110 S. K. Chung, K. Rhee, and S. K. Cho, "Bubble actuation by electrowetting-on-dielectric (EWOD) and its applications: A review," *Int. J. Precis. Eng. Manuf.*, vol. 11, no. 6, pp. 991-1006, 2010.
- 111 Y. Zhao, and S. K. Cho, "Micro air bubble manipulation by electrowetting on dielectric (EWOD): transporting, splitting, merging and eliminating of bubbles," *Lab Chip*, vol. 7, no. 2, pp. 273-280, 2007.
- 112 S. H. Ko, S. J. Lee, and K. H. Kang, "A synthetic jet produced by electrowetting-driven bubble oscillations in aqueous solution," *Appl. Phys. Lett.*, vol. 94, art. 194102, 2009.
- 113 A. Sur, Y. Lu, C. Pascente, P. Ruchhoeft, and D. Liu, "Bubble ebullition on a hydrophilic surface," *J. Heat Transf.*, vol. 137, no. 2, art. 020905, 2015.
- 114 L. Yu, A. Sur, and D. Liu, "Flow Boiling Heat Transfer and Two-Phase Flow Instability of Nanofluids in a Minichannel," *J. Heat Trans.*, vol. 137, no. 5, art. 051502, 2015.
- 115 C. Kelley, H. Singh, and V. Smith, "Data center efficiency and IT equipment reliability at wider operating temperature and humidity ranges," (S. Strutt, ed.) The Green Grid, White Paper #50, Oct. 23, 2012, Retrieved from <http://www.thegreengrid.org/~media/WhitePapers/WP50-Data%20Center%20Efficiency%20and%20IT%20Equipment%20Reliability%20at%20Wider%20Operating%20Temperature%20and%20Humidity%20Ranges.pdf?lang=en> on Jun. 14, 2016
- 116 T. Cader, "The Evolving Data Center: How Power, Cooling, and Standards are Driving Change," *Adv. Technol. Workshop Tabletop Exhib. Therm. Manag., Int. Microelectron. Packag. Soc.*, Los Gatos, California, Sep. 22-24, 2015.
- 117 S. V. Garimella, "The Challenge of Thermal Management," 2016 Intersoc. Conf. Therm. Thermomech. Phenom. Electron. Syst. (ITHERM), Las Vegas, Nevada, May 31- Jun. 3, 2016.



Suresh V. Garimella is Purdue University's Executive Vice President for Research and Partnerships, and the Goodson Distinguished Professor of Mechanical Engineering. He is also Director of the NSF Cooling Technologies Research Center. Dr. Garimella's research group has made seminal contributions in micro- and nano-scale thermal and fluidic engineering, novel materials for thermal management, materials processing and manufacturing, and renewable energy. He has supervised over 90 PhD and MS students, and co-authored 500 refereed journal and conference publications and 21 patents and patent applications. Nineteen alumni from his research group are now faculty members in prestigious universities around the world. Dr. Garimella serves in editorial roles with several leading journals. He is Fellow of the AAAS and the ASME. His contributions to thermal management were recognized with the 2016 IThERM Achievement Award.



Tim Persoons is Assistant Professor in the Department of Mechanical and Manufacturing Engineering at Trinity College Dublin. He is a visiting Faculty member in the NSF Cooling Technologies Research Center at Purdue University. He received his Doctorate in Engineering from KU Leuven in 2006. Dr. Persoons' research activities include multi-scale convective heat transfer in electronics thermal management using unsteady flows, active flow control for sustainable energy devices, and developing experimental thermo-fluid measurement techniques. He has authored over 100 refereed journal and conference publications, and was co-recipient of the 2013 Hartnett-Irvine Award from the International Centre for Heat and Mass Transfer.



Justin A. Weibel is a Research Assistant Professor in the School of Mechanical Engineering at Purdue University and serves as the Associate Director of the NSF Cooling Technologies Research Center (CTRC). He received his PhD in 2012 and BSME in 2007, both from Purdue University. Dr. Weibel's research explores methodologies for prediction and control of phase-change and heat transport across interfaces to enhance the performance and efficiency of thermal management technologies, energy transfer processes, and other multiphase and psychrometric thermal systems. He recently investigated the development of ultra-thin vapor chamber devices for high heat flux generating applications as part of the DARPA Thermal Ground Plane program, and received the 2011 ASME Electronic & Photonic Packaging Division (EPPD) Student Member of the Year Award.



Vadim Gektin is a Principal Engineer with Futurewei Technologies (Huawei USA). He received MS and Ph.D. degrees from the University of Minnesota, Minneapolis, MN, both in Mechanical Engineering, in 1993 and 1997, respectively. Prior to Futurewei Technologies, Dr. Gektin was a Principal Hardware Engineer with Oracle in Santa Clara, CA. His research interests include thermal management of telecommunications, computing, and storage systems and electronic packaging. Dr. Gektin is a licensed Professional Engineer in the state of California.