

THERMAL CONTROL OF ELECTRONICS: PERSPECTIVES AND PROSPECTS

Dr. Robert Hannemann
Charlespoint Group
Boston, MA

ABSTRACT

One of the most prominent industrial applications of heat transfer science and engineering has been electronics thermal control. Driven by the relentless increase in spatial density of microelectronic devices, integrated circuit chip powers have risen by a factor of 100 over the past twenty years, with a somewhat smaller increase in heat flux. The traditional approaches using natural convection and forced-air cooling are becoming less viable as power levels increase. This paper provides a high-level overview of the thermal management problem from the perspective of a practitioner, as well as speculation on the prospects for electronics thermal engineering in years to come.

1 INTRODUCTION AND HISTORICAL PERSPECTIVE

The thermal control of electronic devices and systems has been an integral part of the progress of what is now a dominant component of the world economy. Current research efforts, such as the EDIFICE project described by Amon [1], represent a new wave of cooling technologies driven by the need to remove significant amounts of heat at high flux, with compact and highly reliable devices. This paper provides both a perspective on the challenges of electronics cooling and an overview of future directions and research needs. In the interest of conciseness, the focus is computer-centric; a wider-ranging overview is presented in a recent industry roadmap document [2].

Figure 1 is a useful starting point for considering the problem. While absolute power levels in microelectronic devices are relatively modest (a few tens to a few hundred watts), heat fluxes can be significant ($\sim 50 \text{ W/cm}^2$ in current electronic chips; up to 2000 W/cm^2 in semiconductor lasers). Further, temperatures at the chip surface must be kept to rather low levels ($\sim 100 \text{ C}$) for both performance and reliability reasons.

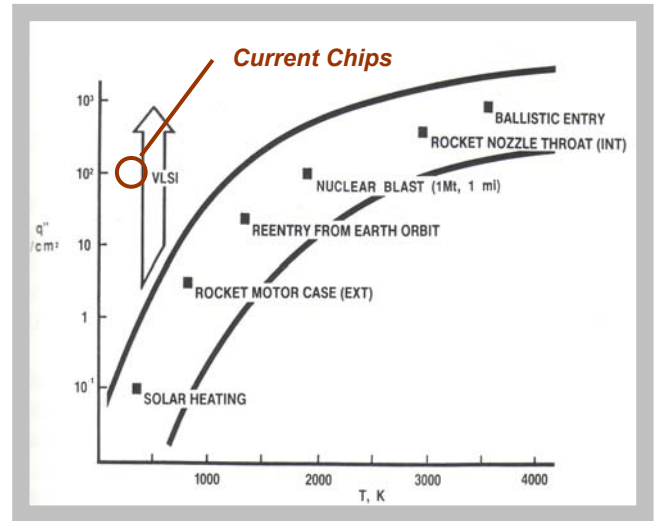


Figure 1: The electronics cooling challenge [3]

The challenge in electronics cooling may be viewed in terms of three different but non-separable problems. First, the chip temperature must be maintained at a relatively low level despite high local heat density. Second, high heat loads must be handled at the assembly or module level. Finally, the thermal environment in the computer machine room, office space, or telecommunications central office must be controlled and the overall rack heat load dealt with. In Table 1 below, the changing thermal control landscape is shown vs. arbitrary generations of products; green indicates no problem, yellow indicates that careful design is required, and red implies a serious challenge.

and experimental optimization of the multi-chip extended surface.

TIMEFRAME	GENERATION	CHIP POWER	MODULE POWER DENSITY	RACK POWER
1945-1955	HISTORIC	●	●	●
1955-1965	TRANSISTOR	●	●	●
1965-1975	SSI	●	●	●
1975-1985	MSI	●	●	●
1985-1990	LSI	●	●	●
1990-2000	VLSI	●	●	●
2000-	ULSI	●	●	●

Table 1: Technology generations

During the period 1975 through 1990, electronics cooling was an intense focus of R&D organizations in industry (IBM, HP, Digital, Fujitsu, etc.) and considerable progress was made in developing design approaches and cooling technology to ensure that thermal management did not become a limit to progress. For reasons of cost and practicality, “exotic” cooling techniques (like liquid cooling and immersion) were limited to mainframes and supercomputers. In other product spaces, the main challenge was to extend the use of forced-air cooling while simultaneously maintaining an acceptable acoustic environment. Work was also being done in anticipation of the major problems currently being addressed – high local heat density and rack-level heat load management. Figures 2, 3, and 4 are examples of this work that are very familiar to this author, but they are representative of a very large body of R&D carried out during that period.

Figure 2 is a sketch describing the use of an optimized heat sinking arrangement for LSI chips, 50 to 100 of which were being used in the design of a high performance minicomputer [4]. The large number of chips being used militated against the use of large individual heat sinks with ducted air paths such as are used in today’s microprocessor-based systems, but the design constraints were similar. Air velocities of 3 – 8 m/s were used, and many engineer-hours of work went into the analytical




Figure 2: Heroic measures to extend air cooling

The rise of so-called multi-chip packaging during this era also provided major problems – heat spreading and the minimization of chip-package thermal resistance. Figure 3 is a patented approach that integrated a two-phase cooling approach into the basic multichip package – an integral heat pipe structure. This approach was never commercialized, but was prototyped [6] and anticipated the so-called vapor-space heat spreaders that are beginning to be used in certain applications.




Figure 3: Integral heat pipe package [5]

The issue of high heat loads at the rack level was also addressed. Figure 4 is a block diagram of a two-phase pumped liquid thermosyphon that was developed for removal of heat

from racks and frames, originally researched as an approach for balancing heat loads on large spacecraft [7].

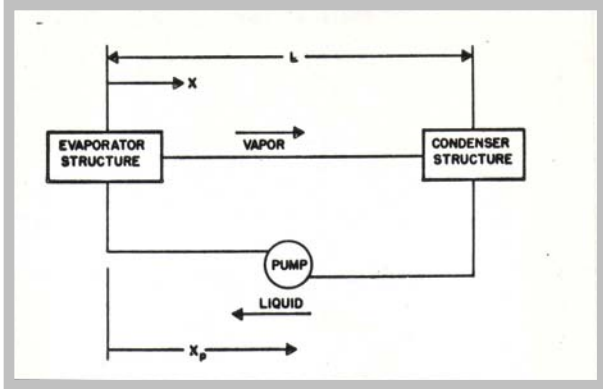


Figure 4: Pumped-liquid thermosyphon

Many of these advanced concepts became unnecessary to deploy with the advent of microprocessor-based systems for almost all applications. While chip power continued to rise, the number of such highly powered chips in each system was greatly reduced. The optimization of micro-heat sinks with channeled forced-air cooling [8] was sufficient for another three generations of products.

2 CURRENT TRENDS AND CHALLENGES

2.1 Product and Technology Requirements

A new generation of problem-solvers is working to develop the advanced thermal management approaches that are appropriate across the product spectrum for the ULSI era. They are employing new tools and technologies (e.g., MEMS devices), but the constraints and design requirements have not changed greatly. Thermal management cannot be the driving force behind new product designs. It must be subservient to other requirements and constraints. The major design goals for any thermal management solution are *performance*, *cost*, *physical size*, and *reliability*. An elegant technical approach will not see the light of day if it is not cost-effective and provide the appropriate levels of reliability and maintainability.

As will be shown in upcoming sections, the use of technologies beyond forced-air cooling will be required for the next major generation of mainstream products. At the same time, the cost of cooling is beginning to be recognized as a major issue (c.f. Figure 5). Once again, a major research and development effort is called for in electronics cooling.

2.2 Driving Trends

Microprocessors are the key element of modern computer design. A quick examination of the power dissipation history of

Intel X86 processors shows that the year 2000 represents a watershed in the requirements for thermal management.




Figure 5: Cost of cooling a microprocessor

Figure 6 shows the power dissipation trends for Intel processors. In general, for CMOS device technology, power dissipation is proportional to the capacitance of the logic elements, the square of the operating voltage swing, and the operating frequency:

$$\text{Power} \sim NCV^2f$$

While the logic element capacitance declines with feature size, and while operating voltages have been significantly reduced, the increase in the number of devices per chip (N) and the operating frequency have driven power levels for the next generation of microprocessors to very high levels (70 – 200W).

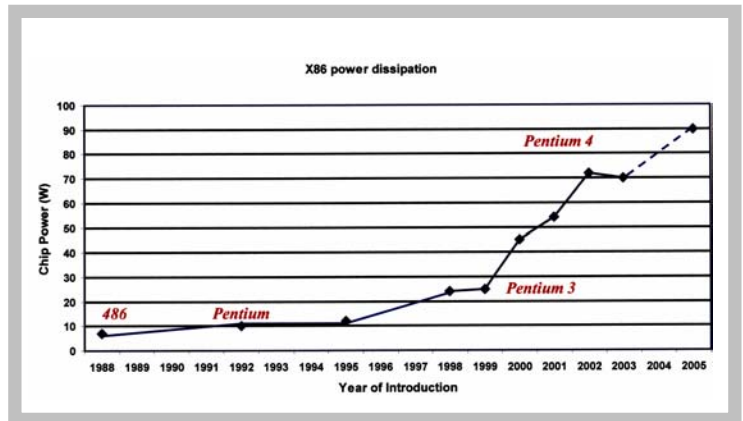


Figure 6: Intel microprocessor power trends

In general, the industry is in agreement as to the general trends for power dissipation. Shown in Figure 7 are the predictions in 1993, 1999, and 2000 for the heat generated by high-end devices – microprocessors and microcontrollers [9,10,11]. Two observations are important. First, it is clear that we will reach well above 100W per chip by the middle of the decade. Second, it is also clear that expectations for power dissipation are *rising* as time goes on.

Figure 8 presents similar data for chip heat flux. This data is also indicative of the rising importance of micro-package level cooling solutions.




Figure 7: Industry power dissipation forecast

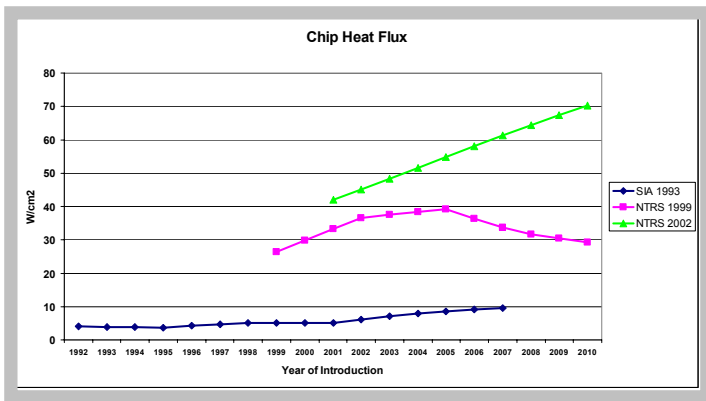


Figure 8: Chip heat flux forecast [9,10,11]

For this generation of products, the two main areas of concern are (1) the thermal control of high-power microprocessors and/or similar ULSI components, given the very high local heat density, and (2) the thermal management of the high overall heat loads generated by equipment in racks and machine rooms. Figure 9 presents the historical and forecast data for the latter issue, for a variety of particular rack environments.

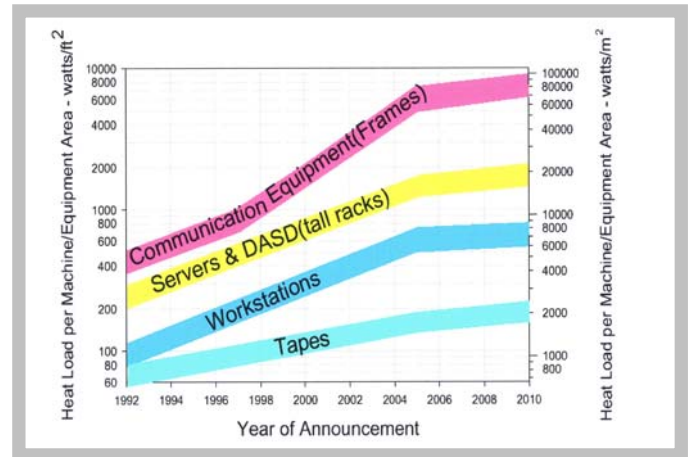


Figure 9: Rack / equipment room heat load forecast [12]

Telecommunications equipment rooms, whether central offices or smaller remote huts, have reached their limit in terms of power density. Computer facilities have lower overall heat loads, but concentrations of power within these rooms are a serious issue. Office areas (typified by the “workstations” data) have rising heat loads that within an open office environment pose a very significant HVAC problem. These problems are beginning to be addressed with liquid-cooled equipment frames (allowing heat to be transported to a remote location) and spot cooling units.

2.3 Large System Challenges

For purposes of this paper “large systems” includes mainframe computer systems, storage arrays, server farms, and network offices. For these systems, localized thermal control for very high power chips is a major issue as chip powers increase beyond 100W. Ducted forced air cooling and active heat sinks (heat sinks with small, dedicated air movers) will continue to be used, but miniaturized liquid cooling loops and in some cases MEMS-based coolers such as EDIFICE require development. These devices, however, must meet stringent cost and reliability requirements – no more than \$100 or so as a cost target, and perhaps 1000 – 5000 FITs (failures in 10⁹ hours) for reliability performance.

Also, rack-level cooling via liquid cooled frames and compact rack-mountable air conditioners (both meeting size, cost and reliability targets) are of interest. While these cooling approaches may seem straightforward, components such as compressors and extended surface arrays meeting the size, cost, and lifetime constraints are not yet available.

Finally, equipment rooms pose interesting HVAC design problems. Improved modeling tools and sensing and control devices are needed.

- Phase change materials for transient applications

2.4 Office System Challenges

Office systems include high performance workstations, printers, desktop computers, and office- and equipment-closet installed networking and telecommunications products.

Cost/performance microprocessors and networking chips will reach 80 – 100W within 2 years. The cooling of these chips will pose unique challenges, since some options open to large system designers will be unavailable. Cooling solutions must be extremely compact, quiet, and cost-effective. Proven-reliability micro heat pipes, breakthrough micro-sized liquid loops, and especially optimized active heat sinking systems may provide thermal management solutions.

Acoustic performance will remain very important in this segment.

2.5 Other Challenges

Telecommunications central offices are already stretching power availability and power handling limits. Photonic components used in these systems (as well as data networking systems) have unique and serious micro-cooling and materials challenges. For example, semiconductor lasers have chip-level heat fluxes on the order of 2000 W/cm². These devices also have very narrow thermal control limits, because performance is significantly impacted by operating temperature (e.g., wavelength). Historically, this segment has been a major applications space for thermoelectric cooling – and conventional, cost-effective TE coolers are reaching their limits in newer applications.

Other important electronic systems must operate in harsh environments, calling for ruggedized cooling solutions with some unique requirements (such as transient high heat loads) but no less pressure on cost and reliability. These applications include automotive, telecom, space, and military systems.

3 RESEARCH AREAS

A previous reference [2] contains a reasonably comprehensive discussion of relevant research areas. For this digest paper, a few of the major topics of interest are listed below.

In *materials*:

- Optimized fluids for liquid cooling
- Materials for package construction / thermal spreading with improved characteristics

Devices:

- High static pressure, low acoustic noise blowers
- Micro heat pipe structures
- Two-phase cooling approaches at small scale
- MEMS components for liquid / two phase cooling

Design:

- Small-footprint liquid cooling systems
- Package-scale jet impingement / spray cooling
- Advanced, integrated thermal design tools
- Frame and rack coolers
- Equipment room thermal design
- Distributed sensing systems

4 CONCLUSION

The importance of thermal management has waxed and waned through the technology and product generations of the past 50 years. Current integrated circuit and photonic technologies as well as the ubiquity of electronic system applications are once again providing a serious challenge to heat transfer engineers in terms of basic theory, tools, components, and innovative design. Breakthroughs are needed in advanced cooling and pragmatic design at all levels.

We are on the cusp of a renaissance in thermal engineering for electronics.

ACKNOWLEDGMENTS

Stimulating discussion and data were provided by Shlomo Novotny of Sun Microsystems, for which I am grateful.

REFERENCES

1. Amon, C., MEMS-based “Thermal Management of High Heat Flux Devices (EDIFICE)”, Rohsenow Symposium on Future Trends in Heat Transfer, Cambridge, MA 2003.
2. Chu, R. and Joshi, Y., eds. , “NEMI Technology Roadmap”, National Electronics Manufacturing Initiative, Herndon, VA 2002.

3. Hannemann, R., Bar-Cohen, A., and Oktay, S., High Heat from a Small Package, Mechanical Engineering, October 1986.
4. Hannemann, R., "Thermal Analysis of Heat Pin Packaging", internal memorandum, Digital Equipment Corporation, Maynard, MA 1979.
5. Saaski, E., Hannemann, R., and Fox, L., "Integral Heat Pipe Module", US Patent 4,833,567, 1989.
6. Hannemann, R., Fox, L., and Kromann, G., "two Phase Internal Cooling technique for Electronic Packages", ASME National Heat Transfer Conference, 1985.
7. Hannemann, R., "Externally Pumped Rankine Cycle Thermal Transport Devices", Progress in Aeronautics and Astronautics, vol. 70, pp. 293-306, 1979.
8. Hannemann, R., Fox, L., and Mahalingham, M., "Heat Transfer in Microelectronic Components", in *Cooling Techniques for Computers*, Hemisphere, New York, 1991.
9. ____, "National Semiconductor Roadmap", Semiconductor Industry Association, 1993.
10. ____, "International Technology Roadmap for Semiconductors", SEMATECH, 1999.
11. ____, "International Technology Roadmap for Semiconductors", SEMATECH, 2002.
12. Computer / telecommunications industry data, 2002.