INVITED PAPER

# Cooling a Microprocessor Chip

*Heat spreading, choice of thermal interface materials, and proper heat-sink design can enhance cooling of microprocessor packages and systems.*

By Ravi Mahajan, *Senior Member IEEE*, Chia-pin Chiu, *Member IEEE*, and Greg Chrysler

**ABSTRACT** | Increasing microprocessor performance has historically been accompanied by increasing power and increasing on-chip power density, both of which present a cooling challenge. In this paper, the historical evolution of power is traced and the impact of power and power density on thermal solution designs is summarized. Industrial and academic researchers have correspondingly increased their focus on elucidating the problem and developing innovative solutions in devices, circuits, architectures, packaging and system level heatsinking. Examples of some of the current packaging and system thermal solutions are provided to illustrate the strategies used in their design. This is followed by a brief discussion of some of the future trends in demand and solution strategies that are being developed by academic and industrial researchers to meet these demands. Potential opportunities and limitations with these strategies are reviewed.

**KEYWORDS** | Density factor; heatsinks; heat spreaders; microprocessors; packaging; thermal interface materials; thermal management

## I. INTRODUCTION

The past few decades have seen a revolutionary increase in computing performance and computers have become increasingly pervasive in all aspects of modern life. Evolution of the microprocessor is one of the most visible and representative facets of the computing revolution. Following Moore's law [1], the semiconductor industry has successfully doubled transistor density every two years and the microprocessor has been the flagship product, successfully exploiting the increased performance with each new technology generation along Moore's law. One of the historical consequences of increasing microprocessor performance is an associated increase in power dissipation. This is not a new issue and was highlighted as early as 1965 [1]. This paper will elaborate on the issues and solutions associated with thermal management of the microprocessor. It will highlight the increased understanding of the thermal management challenges and show some of the innovations developed to meet these challenges.

The microprocessor typically requires thermal management in three distinct environments

1) Cooling is required during functional test to prevent transient temperature rises from causing false performance readouts or failures.
2) During burn-in where infant mortality fails are identified as a function of the temperature, requiring accurate temperature control.
3) During functioning in a user environment. Managing the thermal environment is essential to ensure reliable, long-term performance.

Item 1) requires transient thermal management and 2) and 3), to the first order, require steady state thermal management. The focus in this paper is on 3), i.e., cooling the microprocessor in a user environment. The paper begins with a description of the thermal problem and describes the impact of cooling hot spots on the die. This is followed by a description of the thermal design strategies and the constraints, under which the thermal problem needs to be solved. The paper concludes with some brief comments on the pros and cons of some of the technologies being developed by academic and industrial researchers to meet future challenges.

## II. THERMAL MANAGEMENT OF MICROPROCESSORS

### A. Cooling Demand

Fig. 1 shows the evolution of Thermal Design Power (TDP) as a function of frequency, which is one measure of performance. TDP is of primary interest to the thermal solution designer and it represents the maximum sustained power dissipated by the microprocessor, across a set of realistic applications. It is possible for there to be brief bursts of activity where power dissipated by the microprocessor is larger that TDP, but if the bursts are

Manuscript received February 8, 2005; revised February 3, 2006.
The authors are with the Assembly Technology Development, Intel Corporation, Chandler, AZ 85226 USA (e-mail: greg.chrysler@intel.com).

Digital Object Identifier: 10.1109/JPROC.2006.879800

 0018-9219/$20.00 ©2006 IEEE

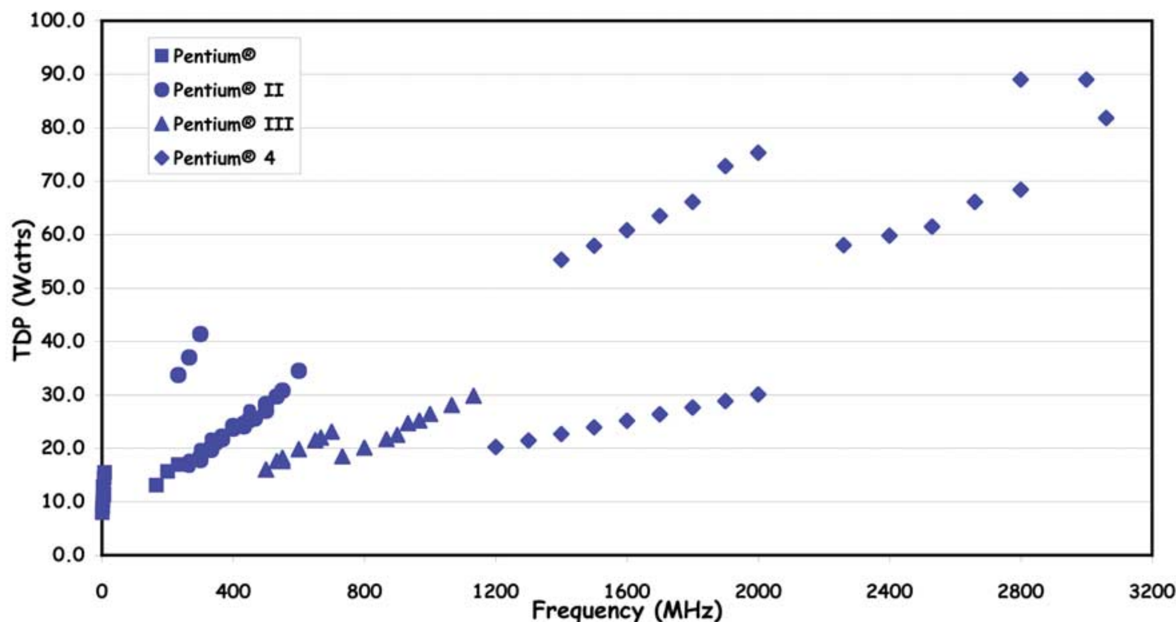


**Fig. 1.** *Historical trend chart of microprocessor TDP versus frequency. Data is based on Intel microprocessors and different symbols in the graph represent different classes of microprocessors. The TDP trend for each class is expected to level off with the transition to multicore processors.*

within the thermal time constant and do not violate the thermal specifications of the microprocessor, they are not of interest to the thermal designer [2]. Modern microprocessors also include advanced thermal monitors and fail-safe mechanisms that prevent catastrophic failures by automatically shutting down the microprocessor if the temperature exceeds a predefined limit [3]. Thus, designing for TDP is adequate to ensure reliable long-term performance.

As seen from Fig. 1, TDP has in the past increased steadily with increasing microprocessor performance, requiring increased focus on thermal management. The transition to multicore microprocessors should alleviate the growth in TDP with increased performance, and the expectation is that there will be a power cap for each product segment. However, in addition to the TDP, thermal design engineers need to account for thermal nonuniformity (typically referred to as hot spots, where power densities of 300+ W/cm$^2$ are possible) caused by nonuniform distribution of power on die (Fig. 2).

The thermal impact of nonuniform power distribution is schematically illustrated in Fig. 3.

The thermal management problem is one of transporting the TDP from the die surface, where the hot spot temperature is maintained at or below a certain temperature specification (typically referred to as the junction temperature $T_j$) to ensure reliable performance, to the ambient air at temperature $T_a$. The transfer of the TDP is by conduction through the many solid layers and interfaces of the package into the heatsink, conduction through the base and into the fins of heatsink, and finally convection from the fins of the heatsink to the cooling air stream. Using a simple thermal resistance model, the cooling demand may be represented as

$$\text{Required thermal resistance} = (T_j - T_a)/\text{TDP}. \quad (1)$$

In general terms, the temperature difference $(T_j - T_a)$ is expected to slowly reduce over time, since $T_j$ can be forced lower by reliability and performance expectations, and $T_a$ can be forced higher due to heating of the inside box air caused by increased integration and shrinking box sizes. The thermal problem can thus become increasingly difficult either due to increases in TDP, reductions in $(T_j - T_a)$ or a combination of both. The thermal solution designer is faced with the challenge of developing a thermal solution that has a thermal resistance at or below the required thermal resistance.

### B. Cooling Solution Design and Development Strategies

The general strategy for thermal management focuses on:

1) minimizing impact of local hot spots by improving heat spreading;
2) increasing the power-dissipation capability of the thermal solutions;
3) expanding the thermal envelopes of systems;

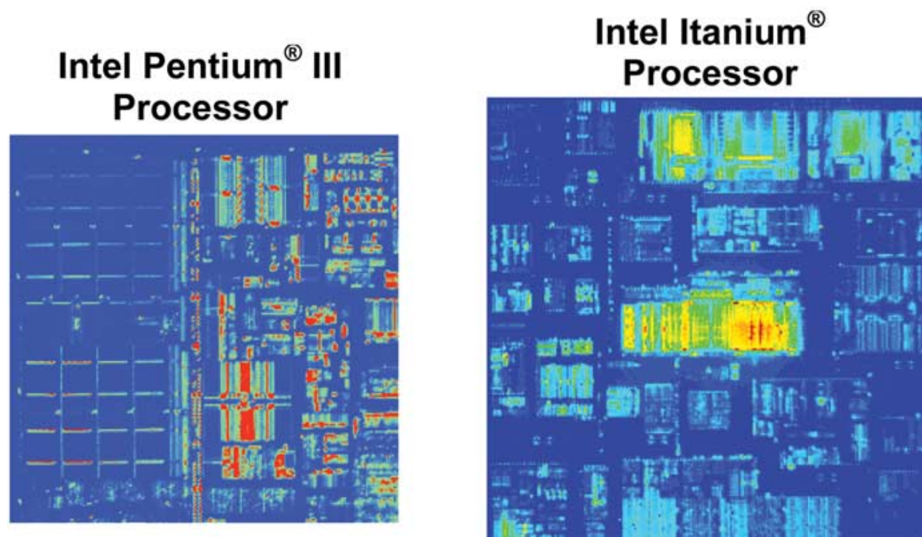


**Fig. 2.** *Illustrative IR images of two typical microprocessors showing on-die hot spots due to nonuniform power distributions on die. Die sizes are not to scale.*

4) developing thermal solutions that meet cost constraints imposed by business considerations;
5) developing solutions that fit within form factor considerations of the chassis.

Thermal management is thus a technical and economic challenge. Cost is one of the most important considerations in the selection of a thermal solution and is often the reason why a new technology may find barriers to introduction especially if it cannot displace an existing technology on a cost performance basis. Fit considerations are also especially important when increasingly higher power components need to be accommodated in the same chassis to prolong use of the same form factor.

Most of today's high-performance microprocessors use an area array, flip chip interconnect scheme to connect the active (circuit) side of the die to an organic or ceramic package substrate (see Fig. 4 for a schematic illustration). The package substrate is either soldered to the computer motherboard through a grid array of solder joints, or has pins that are inserted into a socket that is soldered to the motherboard (another alternate socket is the land grid array socket where socket fingers contact pads on the surface of the package). In all cases, when dealing with high cooling demand, and in attempting to establish cooling envelopes, a reasonable first order assumption is that the bulk of the heat will have to be removed from the inactive

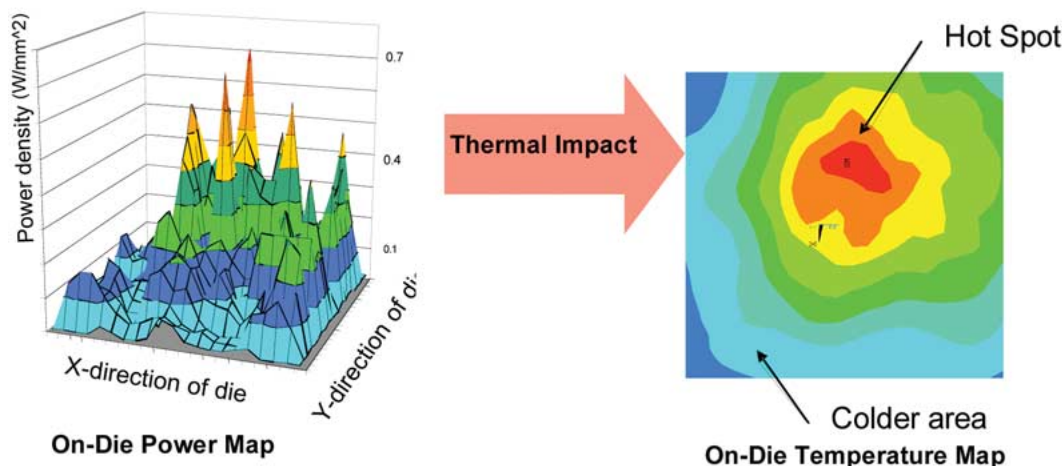


**Fig. 3.** *Schematic illustrating typical die power map and the hot spots on the corresponding die temperature map. The red region represents the highest temperature spot.*

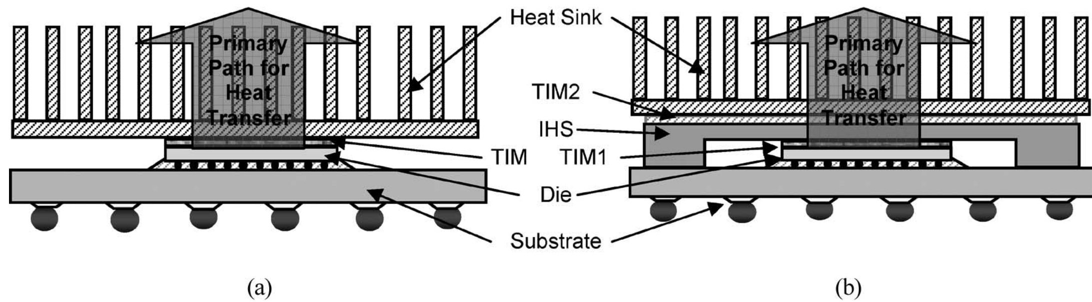


**Fig. 4.** *Schematic of the two basic thermal architectures, illustrating their primary heat transfer path. (a) Architecture I. (b) Architecture II.*

side that is farther away from the motherboard. Given the limited airflow and the presence of significant amounts of lower thermal conductivity organic material on the active side, this is a reasonable first assumption. As discussed in [4] and illustrated in Fig. 4, there are two thermal design architectures. Architecture I is one where a bare die interfaces to the heatsink solution through a thermal interface material (TIM) and Architecture II is one where an integrated heat spreader (IHS) is attached to the die through the use of a TIM and the heatsink interfaces to the IHS through a second TIM. Architecture I has a lower profile compared to Architecture II, and often used for microprocessors in mobile and handheld computers. Architecture II is typically used for microprocessors in desktop and server applications. There are a number of technical and business considerations, beyond the scope of this paper, that dictate selection of a particular thermal architecture for an application.

*1) Power Management at the Chip:* A holistic thermal design strategy would have to consider all aspects of the thermal hierarchy. One of the key aspects is innovation in the development of devices and microarchitectures (including multicore architectures) targeted at reducing TDP [5], [6]. These are beyond the scope of this paper and will not be discussed, however strategies implemented at the device and microarchitectural level have a significant impact on thermal demand and will modulate the scope of the thermal problem going forward. The first aspect discussed here is the definition of TDP and the on-chip power distribution. Over the past few years there have been increasing interactions between chip designers and thermal solution designers to more accurately define and validate the TDP demand and on-die power maps. Based on the power envelope and maps, thermal designers are able to compute hot spot temperatures and help chip designers optimize power maps to be more thermally friendly.

*2) Package Level Cooling:* The goal of package level cooling in Architecture II is to use the IHS to spread the heat while transporting it from the die to the heatsink. The heatsink in turn dissipates heat to the local environment (see Fig. 5 for a pictorial representation of this process). In Architecture I, the base of the heatsink serves the function of an IHS in terms of spreading the heat. Since Architecture II serves to better illustrate the cooling strategy, we will use it in most of the discussion in this section. The TIM between die and IHS is referred to as TIM1 and the TIM2 is the interface material between IHS and the heatsink.

As discussed in Section II-A, there is a need to cool the hot spot at or below $T_j$. As on-die nonuniformity increases, for a constant $(T_j - T_a)$, the overall cooling capability of a thermal solution decreases. This is best illustrated by an example calculation of the impact of power nonuniformity for a typical power map in Fig. 6.

Fig. 6 shows that thermal management solutions will perform better if the source of heat being cooled is uniformly spread. To help quantify the nonuniform power effects, a density factor (DF) that is independent of the power profile on the die has been proposed [7]. The DF is simply the ratio of the actual package thermal resistance at the hottest spot to the die-area-normalized uniform power resistance or thermal impedance, and has the units of inverse area $A^{-1}$. Two very different power profiles could actually result in the same DF, and would therefore result in equal thermal management challenges. On the other

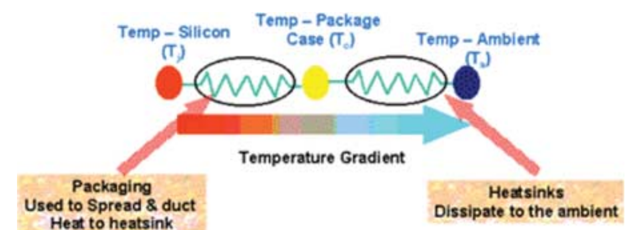


**Fig. 5.** *Temperatures and temperature gradients along the heat transfer path.*

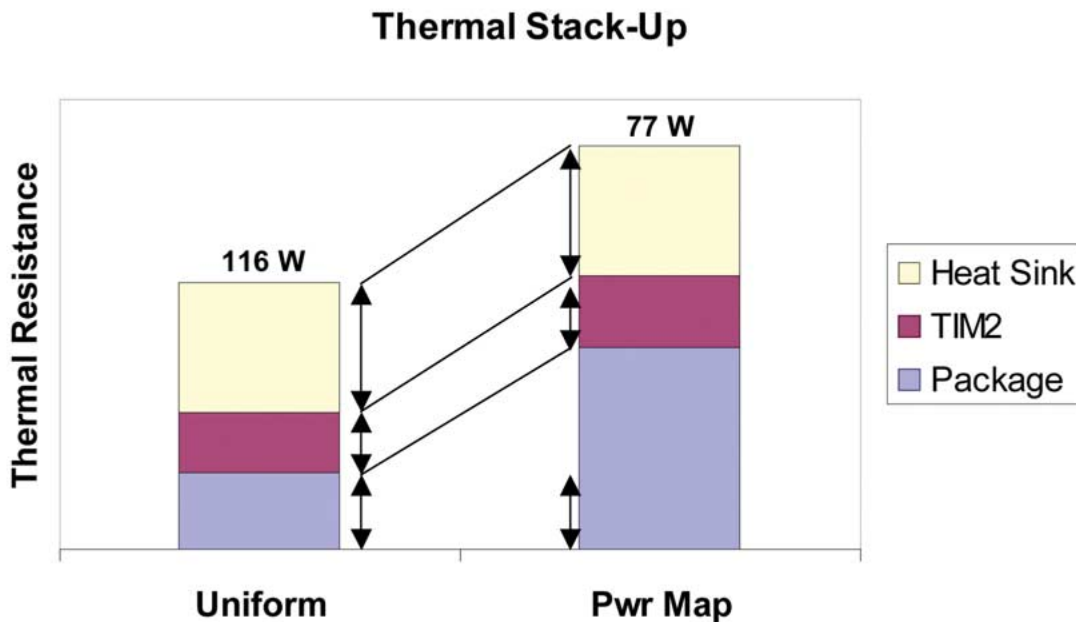


**Fig. 6.** *Sample calculation illustrating the reduction in cooling capability due to nonuniform power map when compared to a uniform on-die power distribution. Vertical arrows represent the magnitude of the thermal resistance of each component under uniform power conditions.*

hand, slight changes in power profiles could result in very different DFs, including very different challenges in thermal management. This factor can help in assessing different cooling schemes without the need to understand specifically the on-die power distribution. An example of the use of this factor is graphically illustrated in Fig. 7, which shows a plot of the overall cooling capability for a die in a fixed package, but with two different heatsinks, as a function of the DF. It can be seen that cooling capability can be increased by reducing DF.

Consider the upper curve in Fig. 7. A 1-cm$^2$ die would result in a DF of 1.0 cm$^{-2}$ if the die power were uniform. For this case, a die dissipating 110 W could be cooled by the package and heatsink combination. If, however, power were concentrated into a small area on the die, then the DF would be greater than 1 cm$^{-2}$; the total power that could be cooled in the package–heatsink combination would be reduced. Furthermore, the value of the DF will also depend on where the concentrated power lies. As way of example, if 50% of the die power were concentrated in only 36% of the die area, then the DF would be 2.7 cm$^{-2}$ if the "hot spot" were located in the center of the die, 3.6 cm$^{-2}$ if located along one edge, and 4.7 cm$^{-2}$ if located in a corner of the die. The corresponding power that could be cooled drops from 110 to 59, 48, and 39 W, for the center, edge, and corner hotspots, respectively. These points are shown as open symbols in Fig. 7.

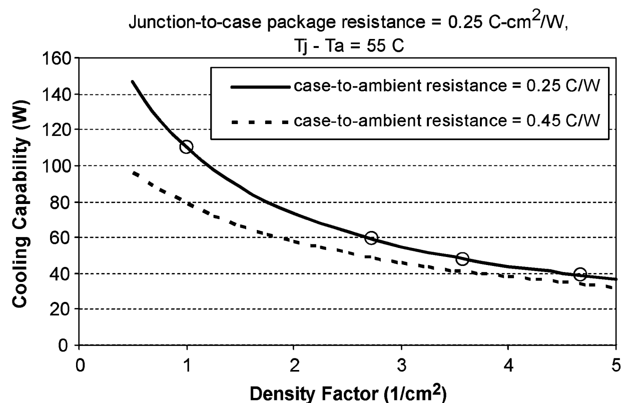


**Fig. 7.** *Impact of die power nonuniformity on cooling capability—increasing DF reduced cooling capability.*

To increase cooling capability, the strategy is to even out the temperature profiles due to nonuniform power distributions, as close to the source as possible, by spreading out the heat. Heat generated at the device is mostly conducted through the thickness in TIM1 with a minimal amount of spreading. The focus in optimizing TIM1 thermal performance is to minimize thermal resistance, resulting in a lower temperature drop across the thermal interface. A lower temperature drop here allows for higher temperature drops across other components without affecting the overall thermal budget. This is accomplished by managing three parameters: 1) the intrinsic thermal conductivity of the TIM; 2) the thermal contact resistance of the die/TIM1 and TIM1/IHS interfaces; and 3) the thermal-interface thickness, also referred to as the bond-line thickness. Considerable advances have

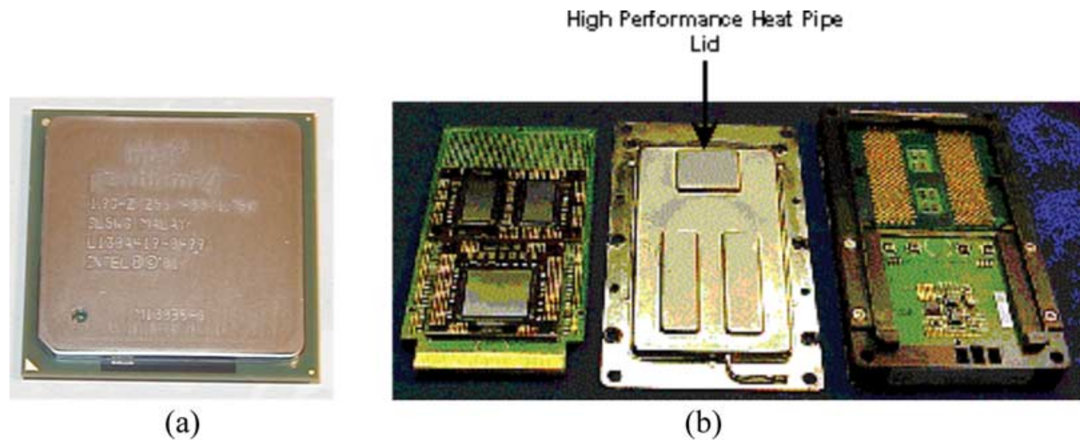


**Fig. 8.** ***(a) Use of an IHS heat spreader for Intel Pentium 4 processor (flip-chip package with Architecture II). (b) Use of a low-weight heat-pipe lid that has high lateral spreading for the Intel Itanium processor.***

been made in recent years in the development and testing on TIMs [8]–[13].

At the IHS level, the heat spreads and some of the peaks in the power profile are smoothed out. The considerations in designing the IHS, are to optimize two factors: the thermal conductivity of the IHS material and the thickness of the IHS while ensuring that the weight to the package is within acceptable limits. A high thermal conductivity and thicker IHS will enhance heat spreading. Some examples of actual applications are shown in Fig. 8.

*3) Heatsink Design:* Heat is next transported between the IHS and the heatsink through another TIM (TIM2). TIM2 is designed with considerations similar to TIM1 and with an additional requirement that it be reworkable. Due to the heat spreading on the heatsink base, the heatsink fins see a more uniform gradient as compared to the IHS, and the heatsink has the primary function of ducting heat to the environment. In typical desktop computer applications, natural or forced convection of air through fins on the heatsink is used to transfer heat to the local ambient. Historically, heatsinks were usually made of aluminum; chosen for its price/thermal-performance ratio and weight advantages. However, with increases in TDP, higher conductivity materials such as copper are being increasingly used.

A properly designed heatsink can help the processor run more reliably and minimize the acoustic noise levels generated from the system by allowing the cooling fan to run at lower speeds. Heatsinks can be categorized into two types i.e., passive heatsinks which depend on airflow from system fans or natural convection and active heatsinks which incorporate a fan to produce direct air impingement for efficient heat removal. These heatsink designs are mostly used in the desktop, workstation, and server systems.

A typical notebook cooling solution is more sophisticated than that for a desktop system. With limited space, and varying notebook design, layout, and processor location, notebook cooling solutions vary greatly between notebooks from different manufacturers. In all notebooks, however, the processor uses one or both of the two cooling methods i.e., passive and active cooling. A remote heat exchanger (RHE) based design offers more flexibility, because the actual heatsink and fan can be placed far from the processor. Fig. 9 shows the RHE design concept. Heat is transferred from the processor to an attachment block, through which runs a heat pipe to an RHE. Localized airflow at the RHE then evacuates the heat to the outside air. Heatsink design considerations include: 1) thermal performance that is optimized given the airflow and pressure drop available and 2) ensuring the "fit" within the computer system in terms of volume and weight.

*4) System Level Cooling:* It is important to optimize effective local airflow through the heatsink to enhance heat transfer from the heatsink to the chassis air flow and to ensure adequate venting in the system. System airflow is typically determined by chassis design, chassis size, location of chassis air intake and exhaust vents, power supply fan capacity and venting, location of the processor, and placement of add-in cards/cables/memory card/other tall components. System integrators must ensure adequate airflow through the system to allow the heatsink to work effectively. Proper attention to airflow when selecting subassemblies and building systems is important for good thermal management and reliable system operation. A typical list of guidelines to be used when integrating a system includes:

- ensuring chassis vents are functional and not excessive in quantity;
- they are properly located;

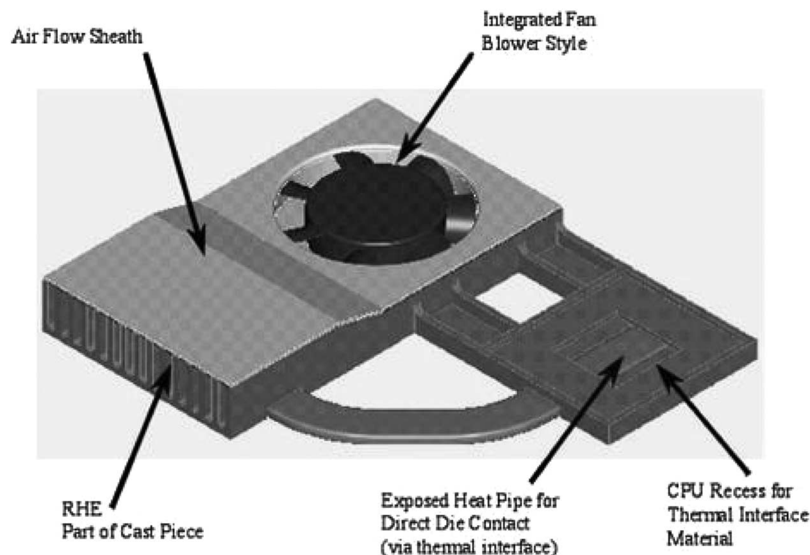


**Fig. 9.** *Typical RHE design for a mobile application.*

- choose a power supply that provides sufficient venting and a fan that sufficiently exhausts air in the proper direction;
- utilizing system fans to improve airflow within the system;
- ensuring that the system fan draws air in the same direction as the overall system airflow;
- avoiding hot spots in the chassis by optimizing placement of exhaust fans.

## III. FUTURE TRENDS

There some recent trends that need careful focus from a thermal management perspective.

1) The recent transition from single-core microprocessor architectures to multicore architectures [14] is a significant trend that somewhat changes the landscape for future thermal management demand. The primary impact is that due to a focus on performance per watt, microprocessor powers will not increase according to historical trends. The focus of the thermal engineer needs to shift towards quantifying thermal power and power density demands and on developing technologies to increase the thermal envelopes.
2) Additionally, there is more interest in microsystems where heterogeneous technology integration at various levels is being actively pursued to enhance performance [15].
3) There is also interest in developing thermal solutions that enable smaller and sleeker form factors.

These trends indicate that a holistic system systems approach to thermal management is needed to ensure adequate problem scoping and solution development. At the device level and microarchitectural level, engineers and designers are more power aware and look for better means of creating thermally friendly designs, while in parallel reducing the power dissipated by the device. The packaging community continues to invest in materials and process technologies to reduce the thermal path resistance. Research in improved heatsinking technologies is also receiving attention. Considering the spectrum of thermal technologies under evaluation, one finds focus on improving the thermal resistance of TIMs and heat spreaders on the one end [16] and more exotic technologies including nanofluid-enhanced liquid cooling [17]–[19] and solid-state refrigeration on the other [20].

In this section, some of the future technologies under active investigation at the package and heatsink level are discussed. Given that a variety of potential directions are being pursued for enhancing component and system computing performance, thermal management will continue to be an important area. Investments in cooling technologies and resolution of integration challenges to fit system requirements will enable increased performance. The key area of focus will have to be on understanding the capability envelopes of new technologies and their cost. The cost performance of new technologies and their ability to be integrated in computing systems will eventually dictate their adoption.

### A. Technologies Under Evaluation

There are quite a few technologies under evaluation by industrial and academic researchers. A few key ones are summarized in Table 1 and discussed in additional detail in the following paragraphs in this section.

Improving convection could be as simple as increasing the fan speed to provide additional air flow over the fins of

Table 1 Summary of Key Technologies Being Studied to Enhance Thermal Performance

| Technology | Pros | Cons |
|---|---|---|
| Enhanced Convection<br>Improve heat transfer coefficient over extended heatsink surfaces | • Simple to implement<br>• Relatively inexpensive | • Improvement is potentially limited<br>• Acoustics will be impacted |
| Spreader Materials<br>High conductivity, carbon fiber, graphite, composites, vapor chambers, heat pipes | • Potentially improved performance can be realized<br>• Potential for weight reduction | • Difficult to compete with copper in terms of cost<br>• Difficult to manufacture compared to copper |
| Closed Loop Liquid Cooling<br>Pumped liquid, single or two phase, closed loop systems for enhanced heatsinking | • Heat Exchanger can be distant from microprocessor with possibly increased surface area for improved cooling | • Chassis size/shape can limit achievable enhancements<br>• Cost<br>• Reliability<br>• Resistance limited to > 0 °C/W |
| Refrigeration<br>Vapor compression, gas compression refrigeration | • Heat Exchanger can be distant from microprocessor with possibly increased surface area for improved cooling surface areas are possible<br>• Resistance < 0 °C/W are possible | • Limited space within chassis<br>• Cost<br>• Reliability |
| Solid State Refrigeration (thermo-electric/thermionic)<br>Targeted hot spot cooling or overall cooling | • Active cooling with no moving parts<br>• Cooling capable of generating thermal resistance < °0 C/W | • Current material sets don't provide efficient cooling |

a heatsink. This would be a fast and easy way to increase heat transfer capability. There is typically little cost increase for a fan motor that can run at a higher speed. However, there are some drawbacks. Increased fan speed comes with a cost of increased acoustic emission and reduced fan reliability. In addition, the improvement achieved is somewhat limited. The convective resistance of a heatsink is only part of the total resistance from chip junction to air. There are conduction resistances which are unaffected by the increase in air velocity. For a high-performance processor package increasing the heat transfer coefficient by 50% over the fins of the heatsink will reduce the heatsink resistance by only 10%. And since the heatsink may be less than half of the total junction to air resistance, this translates to a reduction of less than 5% overall [21].

Another area under active evaluation for improved thermal performance is spreader materials. This wide ranging group includes nonconventional solids, like graphite, carbon fibers, and other composites, and diamond. It also includes vapor chambers and all sorts of heat pipes. In every case the material's effective thermal conductivity is greater than the more common solids like aluminum or copper. Since all microprocessor packages must provide a spreading of the heat from the (small) microprocessor to a (larger) convective area, improvements in the spreader material are important. However, the improvement in thermal resistance is not directly proportional to the thermal conductivity. Doubling the thermal conductivity, from 400 to 800 W/m-K, will not reduce the package resistance to half the original value, since the spreader is not the only thermal resistance in the stack-up. The thermal resistance across TIM1 is independent of the thermal conductivity of the spreader and is a constant term in the stack-up. Even if the thermal conductivity of the spreader were made infinite, the package resistance would not be zero, but rather the sum of the silicon die and the TIM resistances. The major problems with these new materials are cost and handling/forming. Graphite, carbon composites, or diamond materials are typically more expensive than copper. Thus, the cost performance tradeoffs become important in selecting new spreader materials. Heat pipes, vapor chambers, thermosiphons, etc., all provide a spreading capacity and can serve as efficient devices to transport heat [22].

Closed-loop liquid cooling systems are receiving a lot of attention of late. The idea is quite simple: use a liquid to transfer the energy dissipated from the microprocessor to a

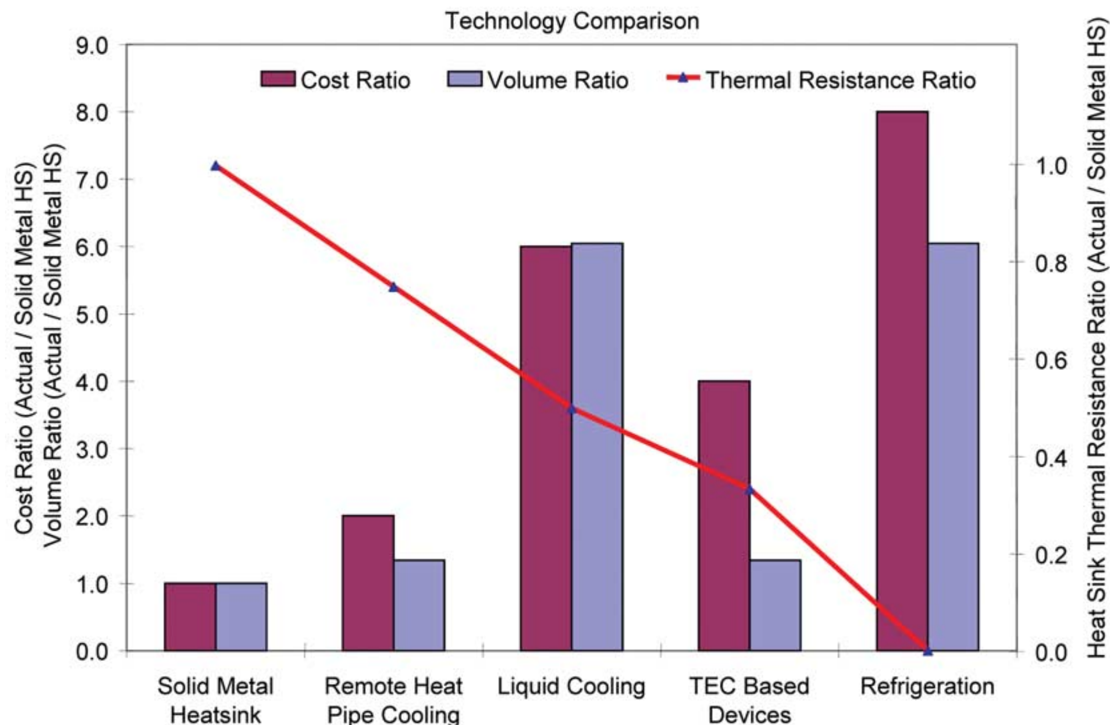


**Fig. 10.** ***Comparison of various advanced thermal solutions.***

remote convective surface to expel it to the air. Since the heat generating and heat dissipating components can now be separated by a distance, the convective surface area is not constrained by the area of the microprocessor, and hence a larger heat exchange surface can be used. The heat exchanger can be placed at any convenient position within the chassis. Still, there are practical limits to the size of the heat exchanger. The pumped liquid system has many components: lines, fittings, pump, cold plate on processor package, and heat exchange. The system can be prone to leaks. The ability of pumps to provide both the pressure head and flow necessary for microprocessor cooling, while fitting within the chassis is a challenge along with their reliability.

Refrigeration systems, vapor compression or gas compression, have one major advantage over the other technologies discussed thus far. They are capable of generating subzero effective thermal resistances. By the definition of the thermal resistance, all that is required is that the evaporator (the cold plate attached to the processor package) has a contact temperature less than the cooling air temperature—something a refrigerator can provide. This feature could substantially increase the heat dissipation possible from a processor. However the low thermal resistance comes at a price. Refrigeration systems require input power to operate and they generate noise. Compressors tend to be large and bulky, from a microelectronics point of view. Compact vapor compression refrigeration sized for electronics cooling has been made [23], however, there is no body of long-term reliability data on compressors sized for electronics cooling.

Solid-state refrigeration (primarily based on thermoelectric and thermionic components) can also provide thermal resistances approaching or even below 0 °C/W [24] and, unlike compression-based refrigeration systems, solid-state refrigeration systems have no moving parts—hence improved reliability. However, there are other concerns dealing with the variations in CTE of components and the resulting stresses as the thermoelectric cooler heats up and cools down. Thermoelectric coolers must be sandwiched between the heat source (the die or package) and the heatsink. The operational power dissipated by the cooler must then be dissipated by the heatsink. This extra energy dissipation may result in higher ambient temperatures at the heatsink and may impact the performance of components downstream of the processor. This concern can be minimized by developing thermoelectric coolers with increased efficiencies.

When compared to the standard solid metal heatsink, there are advantages and disadvantages of each of these newer cooling technologies. Typically the newer systems will require a larger volume and cost more. These are disadvantages. But they do provide lower thermal resistances so that greater powers can be dissipated at reasonable temperatures. A summary of these tradeoffs is presented in Fig. 10.

## IV. CONCLUSION

Thermal management of microelectronics components, microsystems, and systems has been of increasing importance in the past few decades and will continue to be important in the near future with a continued push for performance. The microprocessor, which has most visibly exploited the benefits of Moore's law evolution, has helped define thermal management issues and solution strategies. It is interesting to note that the thermal management is not an entirely new subject and was discussed as a rhetorical question in the paper that defined Moore's Law: "Will it be possible to remove heat generated by tens of thousands of components in a single silicon chip?" [1]. Today it is possible to cool a microprocessor with 1.7 billion transistors [25], an accomplishment that is the result of significant innovations in thermal management. Historically, the key technical issue dominating microprocessor thermal management has been the need to cool a significant TDP, nonuniformly distributed over the die, while maintaining a temperature difference defined by the die hot spot temperature on one end and system ambient on the other. The key nontechnical issues have been meeting cost and form factor constraints. Effective interactions between product designers, silicon process technologists, and packaging and system technologists at the design, development, and research levels have led to innovations that continue to optimize the tradeoffs between performance and cost, and which have led to our success to date. Moving forward, the challenge will be ensuring continued availability of thermal solution technologies so that thermal management is not a limiter to performance. It is important to understand the capability envelopes and limitations of the thermal management technologies currently being developed in industry and academia. These must fit within a holistic management strategy and must meet cost and performance envelopes to find acceptance. ■

## Acknowledgment

The authors would like to thank J. Barletta, M. Rausch, M. Garner, G. Choksi, D. Brown, M. Chen, B. Sankman, V. Bissessur, R. Willoner, and E. Pope from Intel for their help in reviewing this paper.

## REFERENCES

[1] G. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, pp. 114–117, Apr. 19, 1965.

[2] R. Viswanath, V. Wakharkar, A. Watwe, and V. Lebonheur. (2000). Thermal performance challenges from silicon to systems. *Intel Technol. J.*, 3rd quarter. [Online]. Available: http://www.intel.com/technology/itj/q32000/articles/art_4.htm

[3] Intel pentium processor data sheet. [Online]. Available: ftp://download.intel.com/design/Pentium4/datashts/303128.pdf

[4] V. Atluri, R. Mahajan, P. Patel, D. Mallik, J. Tang, V. Wakharkar, G. Chrysler, C. Chiu, G. Choksi, and R. Viswanath, "Critical aspects of high-performance microprocessor packaging," *MRS Bull. Microelectron. Packag. Integr.*, vol. 28, no. 1, pp. 21–34, Jan. 2003.

[5] S. Chou, "Innovation and integration in the nanoelectronics era (keynote address)," in *Dig. Tech. Papers Int. Solid-State Circuits Conf.*, 2005, pp. 36–41.

[6] P. Bai, C. Auth, S. Balakrishnan, M. Bost, R. Brain, V. Chikarmane, R. Heussner, M. Hussein, J. Hwang, D. Ingerly, R. K. James, J. Jeong, C. Kenyon, E. Lee, S.-H. Lee, N. Lindert, M. Liu, Z. Ma, T. Marieb, A. Murthy, R. Nagisetty, S. Natarajan, J. Nierynck, A. Ott, C. Parker, J. Sebastian, R. Shaheed, S. Sivakumar, J. Stegerwald, S. Tyagi, C. Weber, B. Woolery, A. Yeoh, K. Zhang, and M. Bohr, "A 65 nm logic technology featuring 35 nm gate lengths, enhanced channel strain, 8 cu interconnect layers, low-k ILD and 0.57 $\mu m^2$ SRAM cell," in *IEDM Tech. Dig.*, 2004, pp. 657–660.

[7] J. Torresola, C. Chiu, G. Chrysler, D. Grannes, R. Mahajan, and R. Prasher, "Density factor approach to representing impact of die power maps on thermal management," *IEEE Trans. Adv. Packag.*, vol. 28, no. 4, pp. 659–664, Nov. 2005.

[8] A. Russell and C. Chiu, "A testing apparatus for thermal interface materials," in Proc. 1998 Int. Symp. Microelectronics, 1998, pp. 1036–1041.

[9] C. Chiu, G. Solbrekken, and T. Young, "Thermal modeling and experimental validation of thermal interface performance between non-flat surfaces," in *Proc. 7th Intersociety Conf. Thermal and Thermo-Mechanical Phenomena in Electronic Systems (ITHERM 2000)*, vol. 1, pp. 52–62.

[10] C. Chiu, J. Maveety, and Q. Tran, "Characterization of solder interfaces using laser flash metrology," *Microelectron. Reliab.* vol. 42, no. 1, pp. 93–100, Jan. 2002.

[11] C. Chiu, B. Chandran, M. Mello, and K. Kelley, "An accelerated reliability test method to predict thermal grease pump-out in flip-chip applications," in *Proc. 51st Electronic Components and Technology Conf.*, 2001, pp. 91–97.

[12] G. Solbrekken, C. Chiu, B. Byers, and D. Reichebbacher, "The development of a tool to predict package level thermal interface material performance," in *Proc. 7th Intersociety Conf. Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM 2000)*, pp. 48–54.

[13] C. Chiu, G. Solbrekken, and Y. Chung, "Thermal modeling of grease-type interface material in PPGA application," in *Proc. 13th Annu. IEEE Semiconductor Thermal Measurement and Management Symp. (SEMI-THERM)*, 1997, pp. 57–63.

[14] Intel multi-core processors: Leading the next digital revolution (white paper). [Online]. Available: ftp://download.intel.com/technology/computing/multi-core/multi-core-revolution.pdf

[15] R. Tummala, "SOP: What is it and why? A new microsystem-integration technology paradigm-Moore's Law for system integration of miniaturized convergent systems of the next-decade," *IEEE Trans. Adv. Packag.*, vol. 27, no. 2, pp. 241–249, May 2004.

[16] I. Sauciuc, G. Chrysler, R. Mahajan, and R. Prasher, "Spreading in the heat sink base: Phase change systems or solid metals," *Thermal Challenges in Next Generation Electronic Systems, Thermes*, pp. 211–220, 2002.

[17] R. Prasher, P. Bhattacharya, and P. Phelan, "Thermal conductivity of nanoscale colloidal solutions (nanofluids)," *Phys. Rev. Lett.*, vol. 2, no. 94, p. 025901, Jan. 1–4, 2005.

[18] G. Upadhya, P. Zhou, J. Hom, K. Goodson, and M. Munch, "Electro-kinetic microchannel cooling system for servers," in *PROC. 2004 Intersoc. Conf. Thermal Phenomena*, pp. 367–371.

[19] S. Garimella and C. Sobhan, "Transport in micro-channels—A critical review," *Annu. Rev. Heat Transf.*, vol. 13, pp. 1–50, 2003.

[20] R. Venkatasubramanian, E. Silvola, T. Colpitts, and B. O'Quinn, "Thin-film thermo-electric devices with high room-temperature figures of merit," *Nature*, vol. 413, pp. 597–602, Oct. 2001.

[21] M. Saini and R. Webb, "Heat rejection limits of air cooled plane fin heat sinks for computer cooling," in *PROC. 8th Intersoc. Conf. Thermal and Thermo-mechanical Phenomena in Electronic Systems, ITherm 2002*, pp. 1–8.

[22] Y. Joshi, "Microfabrication alliance for innovative cooling of microelectronics," presented at the DARPA HERETIC Principle Investigator's Meeting, Heat Removal by Thermo-Intergrated Circuits, Atlanta, GA, May 23–25, 2001.

[23] J. Peeples, "Vapor compression cooling for high performance applications," *Electron. Cooling*, vol. 7, no. 3, pp. 16–24, 2001.

[24] I. Sauciuc, H. Erturk, G. Chrysler, V. Bala, and R. Mahajan, "Thermal devices integrated with thermoelectric modules with applications to CPU cooling," presented at the ASME InterPACK'05, San Francisco, CA.

[25] S. Naffziger, B. Stackhouse, and T. Grutkowski, "The implementation of a 2-core multi-threaded Itanium family processor," in *Dig. Tech. Papers IEEE Int. Solid-State Circuits Conf. 2005 (ISSCC)*, pp. 182–183, 592.

## ABOUT THE AUTHORS

**Ravi Mahajan** (Senior Member, IEEE) received the B.S. degree from the University of Bombay, India, in 1985, the M.S. degree from the University of Houston, Houston, TX, in 1987, and the Ph.D. degree in mechanical engineering from Lehigh University, Bethlehem, PA, in 1992.

He is a Principal Engineer in the Pathfinding Group within Intel's Assembly Technology Development, Chandler, AZ. In this capacity, he is responsible for setting technology directions for future high density packaging and associated assembly processes. He is also responsible for technical direction for Intel and consortia funded research in assembly and packaging. He specialized in fracture mechanics during his work towards his M.S. and Ph.D. degrees. He has authored several technical papers in the areas of experimental and analytical stress analysis and thermal management. He holds several patents in the area of microelectronic packaging and has been editor and one of the founding members for the section on microelectronics for the Society of Experimental Mechanics. He is also one of the founding editors for the *Intel Assembly and Test Technology Journal*, an Intel internal journal that documents challenges and current progress in the area of assembly and packaging.

Dr. Mahajan currently serves as an Associate Editor of the IEEE TRANSACTIONS ON ADVANCED PACKAGING.

**Chia-pin Chiu** (Member, IEEE) received the B.S. degree in mechanical engineering from the National Taiwan University in 1984 and the M.S. and Ph.D. degrees in mechanical engineering from the University of Minnesota, Minneapolis, in 1992.

After graduation, he worked as a postdoc for nine months and then joined the Assembly Technology Development group, Intel Corp., Chandler, AZ. Among his significant accomplishments are thermal designs for the Pentium, Pentium II, and Pentium III microprocessors. He is a Principal Engineer and manages a Thermal Core Competency Group at Intel, where he is responsible for thermal technology development and product support. His major research includes thermal interface materials, thermal characterization metrology, and the development of new cooling solutions for both CPU and non-CPU products. He holds 26 patents, has authored over 40 technical papers, and is an active member of the JEDEC JC15 committee.

**Greg Chrysler** received the Ph.D. degree in mechanical engineering, specializing in thermal sciences, from the University of Minnesota, Minneapolis, in 1984.

He is a Principal Engineer in the Pathfinding Group within Intel's Assembly Technology Development, Chandler, AZ. His major activities include identification of new thermal management and packaging technologies. He has authored several technical papers, was an Associate Editor of the *ASME Journal of Heat Transfer*, and holds over 60 patents in packaging and cooling of electronics.