Review of Cooling Technologies for Computer Products

Richard C. Chu, Robert E. Simons, Michael J. Ellsworth, Roger R. Schmidt, and Vincent Cozzolino

Invited Paper

***Abstract*—This paper provides a broad review of the cooling technologies for computer products from desktop computers to large servers. For many years cooling technology has played a key role in enabling and facilitating the packaging and performance improvements in each new generation of computers. The role of internal and external thermal resistance in module level cooling is discussed in terms of heat removal from chips and module and examples are cited. The use of air-cooled heat sinks and liquid-cooled cold plates to improve module cooling is addressed. Immersion cooling as a scheme to accommodate high heat flux at the chip level is also discussed. Cooling at the system level is discussed in terms of air, hybrid, liquid, and refrigeration-cooled systems. The growing problem of data center thermal management is also considered. The paper concludes with a discussion of future challenges related to computer cooling technology.**

***Index Terms*—Air cooling, data center cooling, flow boiling, heat sink, immersion cooling, impingement cooling, liquid cooling, pool boiling, refrigeration cooling, system cooling, thermal, thermal management, water cooling.**

I. Introduction

ELECTRONIC devices and equipment now permeate virtually every aspect of our daily life. Among the most ubiquitous of these is the electronic computer varying in size from the handheld personal digital assistant to large scale mainframes or servers. In many instances a computer is imbedded within some other device controlling its function and is not even recognizable as such. The applications of computers vary from games for entertainment to highly complex systems supporting vital health, economic, scientific, and military activities. In a growing number of applications computer failure results in a major disruption of vital services and can even have life-threatening consequences. As a result, efforts to improve the reliability of electronic computers are as important as efforts to improve their speed and storage capacity.

Since the development of the first electronic digital computers in the 1940s, the effective removal of heat has played a key role in insuring the reliable operation of successive generations of computers. The Electrical Numerical Integrator and Computer (ENIAC), dedicated in 1946, has been described as a "30 ton, boxcar-sized machine requiring an array of industrial cooling fans to remove the 140 kW dissipated from its 18 000 vacuum tubes" [1]. Following ENIAC, most early digital computers used vacuum-tube electronics and were cooled with forced air.

The invention of the transistor by Bardeen, Brattain, and Shockley at Bell Laboratories in 1947 [2] foreshadowed the development of generations of computers yet to come. As a replacement for vacuum tubes, the miniature transistor generated less heat, was much more reliable, and promised lower production costs. For a while it was thought that the use of transistors would greatly reduce if not totally eliminate cooling concerns. This thought was short-lived as packaging engineers worked to improve computer speed and storage capacity by packaging more and more transistors on printed circuit boards, and then on ceramic substrates.

The trend toward higher packaging densities dramatically gained momentum with the invention of the integrated circuit separately by Kilby at Texas Instruments and Noyce at Fairchild Semiconductor in 1959 [2]. During the 1960s, small scale and then medium scale integration (SSI and MSI) led from one device per chip to hundreds of devices per chip. The trend continued through the 1970s with the development of large scale integration (LSI) technologies offering hundreds to thousands of devices per chip, and then through the 1980s with the development of very large scale (VLSI) technologies offering thousands to tens of thousands of devices per chip. This trend continued with the introduction of the microprocessor and continues to this day with chip makers projecting that a microprocessor chip with a billion or more transistors will be a reality before 2010.

In many instances the trend toward higher circuit packaging density has been accompanied by increased power dissipation per circuit to provide reductions in circuit delay (i.e., increased speed). The need to further increase packaging density and reduce signal delay between communicating circuits led to the development of multichip modules beginning in the late 1970s and is continuing today. An example of the effect that these trends have had on module heat flux in high-end computers is shown in Fig. 1. As can be seen heat flux associated with Bipolar circuit technologies steadily increased from the very beginning and really took off in the 1980s. There was a brief respite with the transition to CMOS circuit technologies in the 1990s; but, the demand for increased packaging density and performance reasserted itself and heat flux is again increasing at a challenging

Manuscript received August 30, 2004.
The authors are with the IBM Corporation, Poughkeepsie, NY 12601 USA (e-mail: rcchu@us.ibm.com).

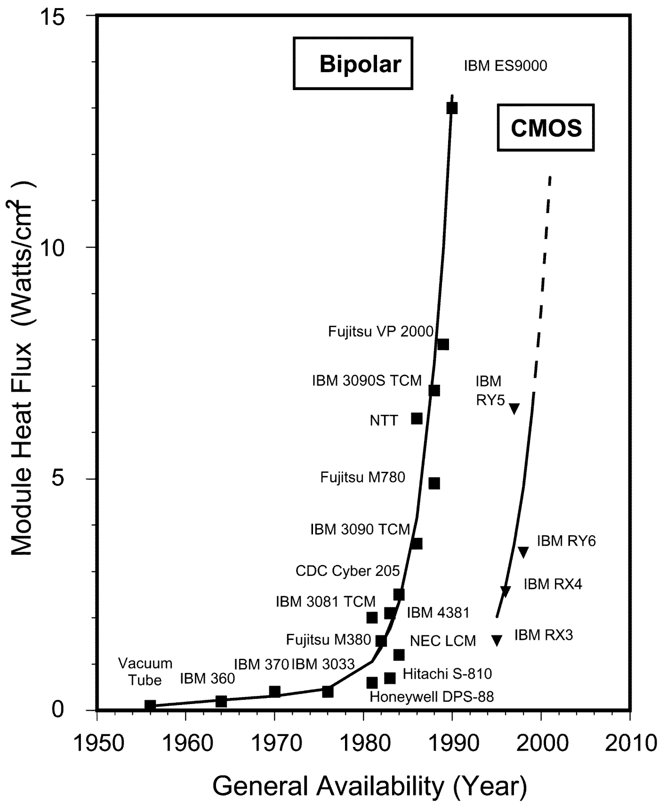

Fig. 1. Evolution of module level heat flux in high-end computers.

Throughout the past 50 years, cooling and thermal management have played a key role in accommodating increases in power while maintaining component temperatures at satisfactory levels to satisfy performance and reliability objectives. Sections II–V of this paper will discuss the various techniques that have been used to provide temperature control in computers in the past and present, as well as some of the methods being explored for the future.

II. Module-Level Cooling

Processor module cooling is typically characterized in two ways: cooling internal and external to the module package and applies to both single and multichip modules. Fig. 2 illustrates the distinction between the two cooling regimes in the context of a single-chip module.

A. Internal Module Cooling

The primary mode of heat transfer internal to the module is by conduction. The internal thermal resistance is therefore dictated by the module's physical construction and material properties. The objective is to effectively transfer the heat from the electronics circuits to an outer surface of the module where the heat will be removed by external means which will be discussed in the following section.

In the case of large multichip modules (MCMs) where variation in the location and height of chips had to be considered, an approach (Figs. 3 and 4) was adopted that employed a spring-loaded mechanical cylindrical piston touching each chip

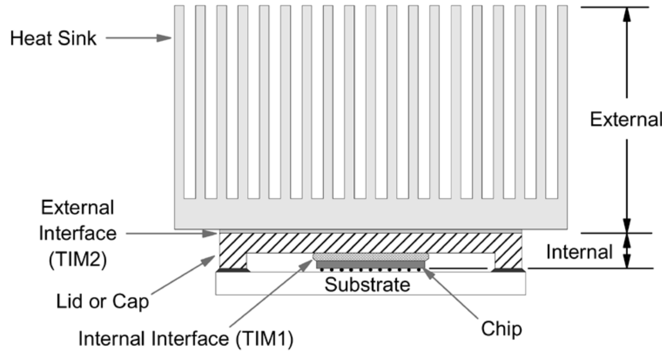

Fig. 2. Cross-section of a typical module denoting internal cooling region and external cooling region.

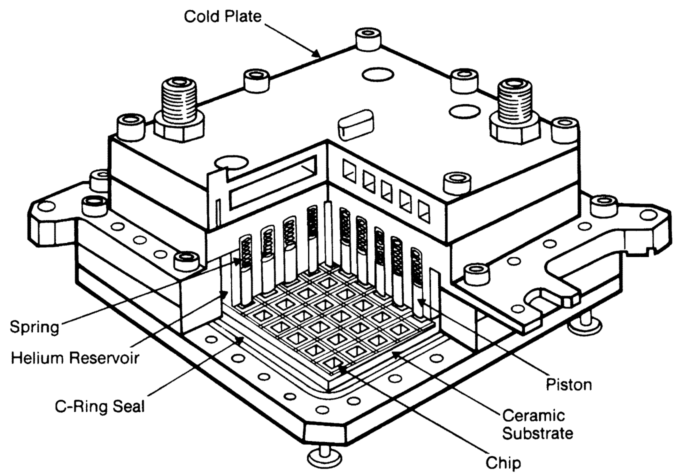

Fig. 3. Isometric cutaway view of an IBM TCM module with a water-cooled cold plate.

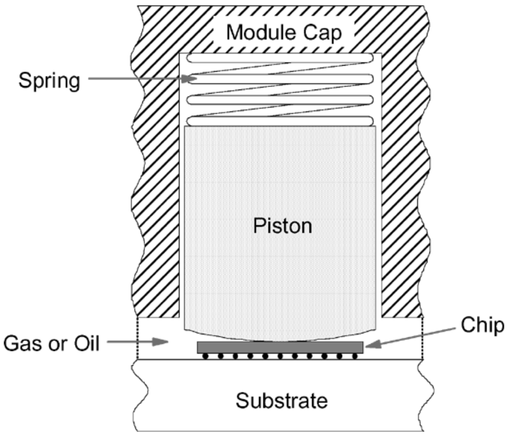

Fig. 4. Cross-sectional view of an IBM TCM module on an individual chip site basis.

The volume within the module was filled with helium gas to minimize the thermal resistance across the gaps and achieve an acceptable internal thermal resistance. The total module cooling assembly was patented as a gas-encapsulated module [4] and later named a thermal conduction module (TCM). TCM cooling technology evolved through three generations of IBM mainframes: system 3081, ES/3090, and ES/9000, with about

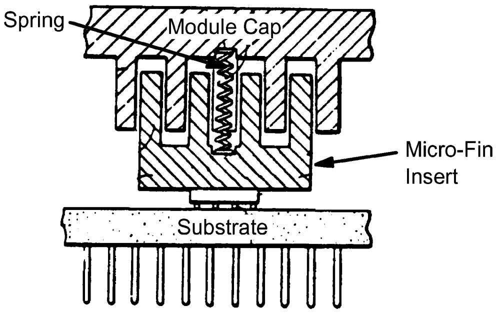

Fig. 5. Cross-sectional view of a Hitachi M-880 module on an individual chip site basis.

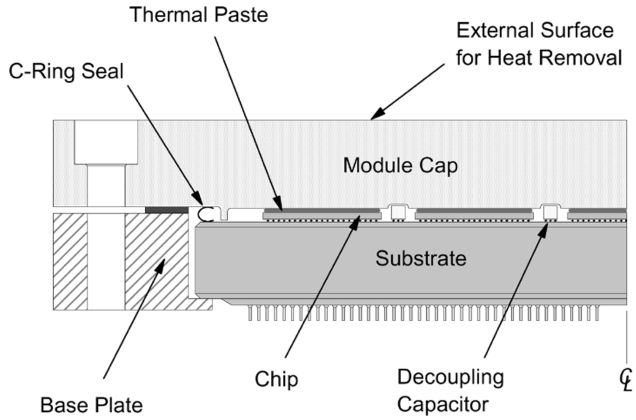

Fig. 6. Cross-sectional view of central processor module package with thermal paste path to module cap [9].

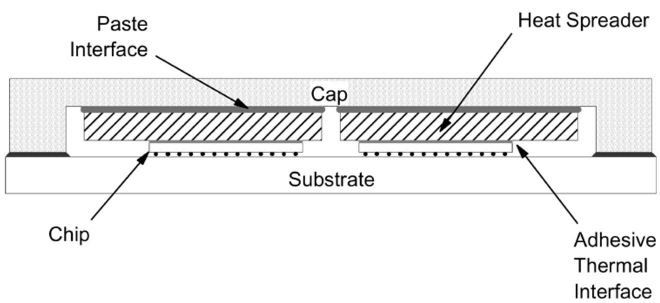

Fig. 7. MCM cross-section showing heat spreader adhesively attached to chip (adapted from [10]).

[5]. The last generation TCM incorporated a copper piston (the original piston was aluminum) with a cylindrical center section and a slight taper on each end to minimize the gap between piston and cap while retaining intimate contact between the piston face and the chip [6]. Additionally, the volume inside the module was filled with a PAO (polyalphaolefin) oil instead of helium to reduce the piston-to-cap and chip-to-piston thermal resistances. Hitachi packaged a similar conduction scheme in their M-880 [7] and MP5800 [8] processors. Instead of a cylindrical piston Hitachi utilized an interdigitated microfin structure (Fig. 5).

In the 1990s when IBM made the switch from bipolar to CMOS circuit technology [10] the conduction cooling approach was simplified and reduced in cost by adopting a "flat plate" conduction approach as shown in Fig. 6. The thermal path from chip to cap is provided by a controlled thickness (e.g., 0.10 mm to 0.18 mm) of a thermally conductive paste. This was possible largely due to improved planarity of the substrate, better control of dimensional tolerances and enhanced thermal conductivity of the paste.

As time went on, chip power levels continued to increase. In addition, concentrated areas of high heat flux 2 to 3 times the average chip heat flux referred to as hot spots emerged. To meet internal thermal resistance requirements, in 2001 IBM chose to attach a high-grade silicon carbide (SiC) spreader to the chip with an adhesive thermal interface (ATI) and then use a more conventional thermal paste between the spreader and the cap [10]. This configuration is shown in Fig. 7.

The adhesive thermal interface (ATI), while not as thermally conductive as the thermal paste, could be applied much thinner resulting in a lower thermal resistance. SiC was chosen for the spreader material for its unique combination of high thermal conductivity and low coefficient of thermal expansion (CTE). The CTE of the SiC closely matches that of the silicon chip thus avoiding stress fracturing the interface when the module heats up during use. The thermal resistance of this package arrangement is lower than just using thermal paste between chip and cap because of the use of the lower thermal resistance ATI on the smaller chip area. The thermal paste thermal resistance is mitigated by applying it over a much larger area.

B. External Module Cooling

the system environment. This is accomplished primarily by attaching a heat sink to the module. Traditionally, and preferably, the system environment of choice has been air because of its ease of implementation, low cost, and transparency to the end user or customer. This section, therefore, will focus on air-cooled heat sinks. Liquid-cooled heat sinks typically referred to as cold plates will also be discussed.

1) Air-Cooled Heat Sinks: A typical air-cooled heat sink is shown in Fig. 8. The heat sink is constructed of a base region that is in contact with the module to be cooled. Fins protruding from the base serve to extend surface area for heat transfer to the air. Heat is conducted through the base, up into the fins and then transferred to the air flowing in the spaces between the fins by convection. The spacing between fins can run continuously in one direction in the case of a straight fin heat sink or they can run in two directions in the case of a pin fin heat sink (Fig. 9). Air flow can either be through the heat sink laterally (in cross flow) or can impinge from the top as seen in Fig. 10.

The thermal performance of the heat sink is a function of many variables. Geometric variables include the thickness and plan area of the base plus the fin thickness, height, and spacing. The principal material variable is thermal conductivity. Also factored in is volumetric air flow and pressure drop. Many optimization studies have been conducted to minimize the external thermal resistance for a particular set of application conditions [11]–[13]. However, over time, as greater and greater thermal performance has been required, fin heights and fin number have

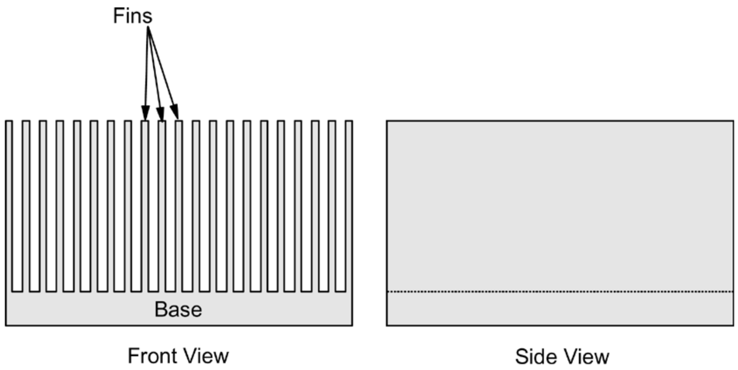

Fig. 8. Typical air-cooled heat sink.

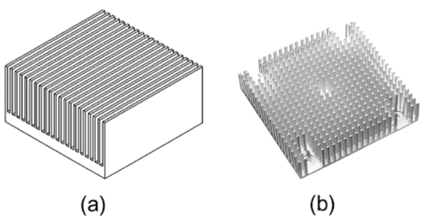

Fig. 9. Typical (a) straight fin heat sink and (b) pin fin heat sink.

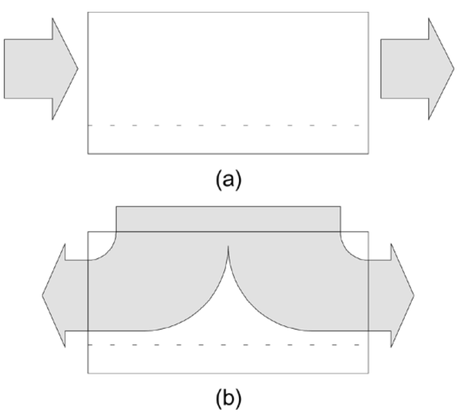

Fig. 10. Air flow path through a heat sink: (a) cross flow or (b) impingement.

(with thermal conductivity ranging from 150–200 W/mK) to aluminum fins on copper bases (with thermal conductivity ranging from 350–390 W/mK) to all copper. In certain cases heat pipes have been embedded into heat sinks to more effectively spread the heat [14]–[16].

Heat sink attachment to the module also plays a role in the external thermal performance of a module. The method of attachment and the material at the interface must be considered. The material at the interface is important because when two surfaces are brought together seemingly in contact with one another, surface irregularities such as surface flatness and surface roughness result in just a fraction of the surfaces actually contacting one another. The majority of the heat is therefore transferred through mechanical means using screws or a clamping mechanism. Air has traditionally existed at the interface but more recently oils or even phase change materials (PCMs) have been used [18] to reduce the thermal resistance at the interface. Another method of attachment has been adhesively with an elastomer or epoxy. This method has worked well on smaller single-chip modules where heat sinks do not have to be removed from the module.

2) Water-Cooled Cold Plates: For situations where air cooling could not meet requirements, such as was the case in IBM's 3081, ES/3090, and ES/9000 systems in the 1980s and early 1990s, and the case in Hitachi's M-880 and MP5800 in the 1990s, heat was removed from the modules via water-cooled cold plates. Compared to air, water cooling can provide almost an order of magnitude reduction in thermal resistance principally due to the higher thermal conductivity of water. In addition, because of the higher density and specific heat of water, its ability to absorb heat in terms of the temperature rise across the coolant stream is approximately 3500 times that of air. Cold plates function very similarly to air-cooled heat sinks. For example, the ES/9000 cold plate is an internal finned structure made of tellurium copper [19]. As with the air-cooled heat sinks, changes in material properties and geometry were made to improve performance. A higher thermal conductivity tellurium copper was chosen over beryllium copper used in previous generation cold plates. Additionally, fin heights were increased and channel widths (analogous to fin spacings) were decreased. The ES/9000 module also marked the first time IBM used a PAO oil at the interface between the module cap and cold plate to reduce the thermal interface resistance.

In an effort to significantly extend the cooling capability of liquid-cooled cold plates, researchers continue to work on microchannel cooling structures. The concept was originally demonstrated over 20 years ago by Tuckerman and Pease [20]. They chemically etched 50 μm-wide by 300-μm-deep channels into a 1 cm $\times$ 1 cm silicon chip. By directing water through these microchannels they were able to remove 790 W with a temperature difference of 71 °C. More recently, aluminum nitride heat sinks fabricated using laser machining and adhesively attached to the die have been used to cool a high-powered MCM and achieve a junction to ambient unit thermal resistance below 0.6 K-cm^2/W [21]. The challenge continues to be to provide a practical chip or module cooling structure and flow

C. Immersion Cooling

Immersion cooling has been of interest as a possible method to cool high heat flux components for many years. Unlike the water-cooled cold plate approaches which utilize physical walls to separate the coolant from the chips, immersion cooling brings the coolant in direct physical contact with the chips. As a result, most of the contributors to internal thermal resistance are eliminated, except for the thermal conduction resistance from the device junctions to the surface of the chip in contact with the liquid.

Direct liquid immersion cooling offers a high heat transfer coefficient which reduces the temperature rise of the heated chip surface above the liquid coolant temperature. The magnitude of the heat transfer coefficient depends upon the thermophysical properties of the coolant and the mode of convective heat transfer employed. The modes of heat transfer associated with liquid immersion cooling are generally classified as natural convection, forced convection, and boiling. Forced convection includes liquid jet impingement in the single phase regime and boiling (including pool boiling, flow boiling, and spray cooling) in the two-phase regime. An example of the broad range of heat flux that can be accommodated with the different modes and forms of direct liquid immersion cooling is shown in Fig. 11 [22].

Selection of a liquid for direct immersion cooling cannot be made on the basis of heat transfer characteristics alone. Chemical compatibility of the coolant with the chips and other packaging materials exposed to the liquid is an essential consideration. There may be several coolants that can provide adequate cooling, but only a few will be chemically compatible. Water is an example of a liquid which has very desirable heat transfer properties, but which is generally undesirable for direct immersion cooling because of its chemical and electrical characteristics. Alternatively, fluorocarbon liquids (e.g., FC-72, FC-86, FC-77, etc.) are generally considered to be the most suitable liquids for direct immersion cooling, in spite of their poorer thermophysical properties [22], [23].

1) Natural and Forced Liquid Convection: As in the case of air cooling, liquid natural convection is a heat transfer process in which mixing and fluid motion is induced by differences in coolant density caused by heat transferred to the coolant. As shown in Fig. 11, this mode of heat transfer offers the lowest heat flux or cooling capability for a given wall superheat or surface-to-liquid temperature difference. Nonetheless, the heat transfer rates attainable with liquid natural convection can exceed those attainable with forced convection of air.

Higher heat transfer rates may be attained by utilizing a pump to provide forced circulation of the liquid coolant over the chip or module surfaces. This process is termed forced convection and the allowable heat flux for a given surface-to-liquid temperature difference can be increased by increasing the velocity of the liquid over the heated surface. The price to be paid for the increased cooling performance will be a higher pressure drop. This can mean a larger pump and higher system operating pressures. Although forced convection requires the use of a pump and the associated piping, it offers the opportunity to remove heat from high power chips and modules in a confined space.

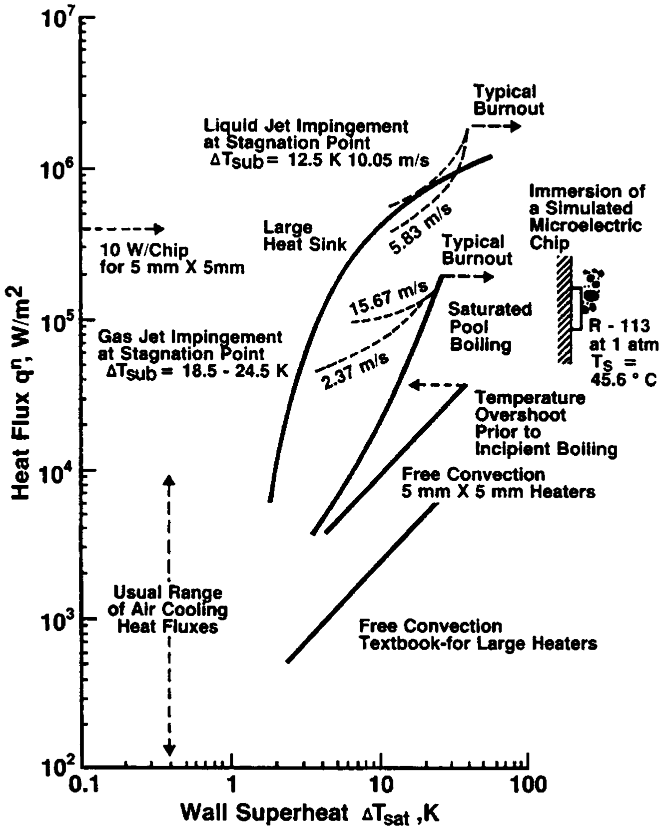

Fig. 11. Heat flux ranges for direct liquid immersion cooling of microelectronic chips [22].

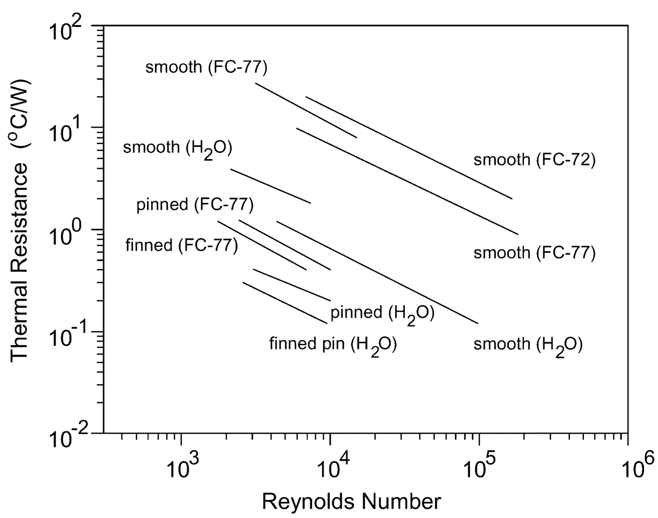

Fig. 12. Forced convection thermal resistance results for simulated 12.7 mm × 12.7 mm microelectronic chips (adapted from [24]).

Experimental studies were conducted by Incropera and Ramadhyani [24] to study liquid forced convection heat transfer from simulated microelectronic chips. Tests were performed with water and dielectric liquids (FC-77 and FC-72) flowing over bare heat sources and heat sources with pin-fin and finned pin extended surface enhancement. It can be seen in Fig. 12 that, depending upon surface and flow conditions (i.e., Reynolds number), thermal resistance values obtained for the

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | **sales@docketalarm.com** | **1-866-77-FASTCASE**

fastcase®
Smarter legal research.