AUGUST 1991

WRL
Research Report 91/9

Interleaved Fin Thermal Connectors for Multichip Modules

William R. Hamburgen

digital
Western Research Laboratory 250 University Avenue Palo Alto, California 94301 USA

The Western Research Laboratory (WRL) is a computer systems research group that was founded by Digital Equipment Corporation in 1982. Our focus is computer science research relevant to the design and application of high performance scientific computers. We test our ideas by designing, building, and using real systems. The systems we build are research prototypes; they are not intended to become products.

There is a second research laboratory located in Palo Alto, the Systems Research Center (SRC). Other Digital research groups are located in Paris (PRL) and in Cambridge, Massachusetts (CRL).

Our research is directed towards mainstream high-performance computer systems. Our prototypes are intended to foreshadow the future computing environments used by many Digital customers. The long-term goal of WRL is to aid and accelerate the development of high-performance uni- and multi-processors. The research projects within WRL will address various aspects of high-performance computing.

We believe that significant advances in computer systems do not come from any single technological advance. Technologies, both hardware and software, do not all advance at the same pace. System design is the art of composing systems which use each level of technology in an appropriate balance. A major advance in overall system performance will require reexamination of all aspects of the system.

We do work in the design, fabrication and packaging of hardware; language processing and scaling issues in system software design; and the exploration of new applications areas that are opening up with the advent of higher performance systems. Researchers at WRL cooperate closely and move freely among the various levels of system design. This allows us to explore a wide range of tradeoffs to meet system goals.

We publish the results of our work in a variety of journals, conferences, research reports, and technical notes. This document is a research report. Research reports are normally accounts of completed research and may include material from earlier technical notes. We use technical notes for rapid distribution of technical material; usually this represents research in progress.

Research reports and technical notes may be ordered from us. You may mail your order to:

Technical Report Distribution
DEC Western Research Laboratory, WRL-2
250 University Avenue
Palo Alto, California 94301 USA

Reports and notes may also be ordered by electronic mail. Use one of the following addresses:

Digital E-net:	`DECWRL::WRL-TECHREPORTS`
Internet:	`WRL-Techreports@decwrl.dec.com`
UUCP:	`decwrl!wrl-techreports`

To obtain more details on ordering by electronic mail, send a message to one of these addresses with the word ‘‘`help`’’ in the Subject line; you will receive detailed instructions.

Interleaved Fin Thermal Connectors for Multichip Modules

William R. Hamburgen

August, 1991

digital **Western Research Laboratory** 250 University Avenue Palo Alto, California 94301 USA

Abstract

Packaging VLSI devices in multichip modules provides a means of building high performance mainframe computers. The very best substrate wiring densities and power distribution are obtained when flipped-chip mounting is used. With flipped-chips, however, heat removal through the substrate is seldom adequate; a means is needed to effectively couple each die to a common lid or coldplate. This is complicated by the need to be able to disassemble the module for repair. The usual approach to this problem applies thermal conduction from the back of the die across an approximately planar gap to a cap assembly designed to accommodate manufacturing tolerances, shock and vibration, and thermal expansion. Typically, a thermally conductive material enhances heat transfer across the gap. Conductive gas such as helium complicates module sealing, and liquids or greases can present difficulties in assembly and rework.

An alternative approach is to greatly increase the surface area for heat transfer across the chip-to-cap interface with arrays of interleaved fins attached to the mating surfaces. This reduces or eliminates the need for a highly thermally conductive interstitial material. In this paper, some of the many proposed configurations of such a thermal connector will be reviewed. It will be shown that for a given fin and interstitial material, the thermal resistance of an optimal design is proportional to the inter-fin gap. For example, at a 70 micron gap, an optimal copper connector with air as the interstitial fluid attains 1 $cm^2\ °C/W$. Particular attention will be paid to connectors with the separable interface at the gap between the mating fin sets.

This is a preprint based on a paper that will be presented at the
1991 IEPS International Electronic Packaging Conference,
San Diego, California, September 15-19, 1991.

Table of Contents

List of Figures

1. Introduction

The recent introduction of new mainframe computers by IBM [16] and Hitachi [10] marks the coming of age of interleaved fin thermal connectors (figure 1). Designers have long grappled with the difficult problem of efficiently removing heat from the backside of a flipped chip in a multichip module, without also transferring damaging forces from a coldplate assembly to the die and its substrate. A further complication is the need for a separable interface to facilitate module repair. The interleaved fin thermal connector, in its many varieties, offers the potential to transfer heat with little or no mechanical coupling between source and sink. It is the objective of this paper to review the motivation for this style of thermal connector and several proposed configurations, and to explore some fundamentals of its design and optimization.

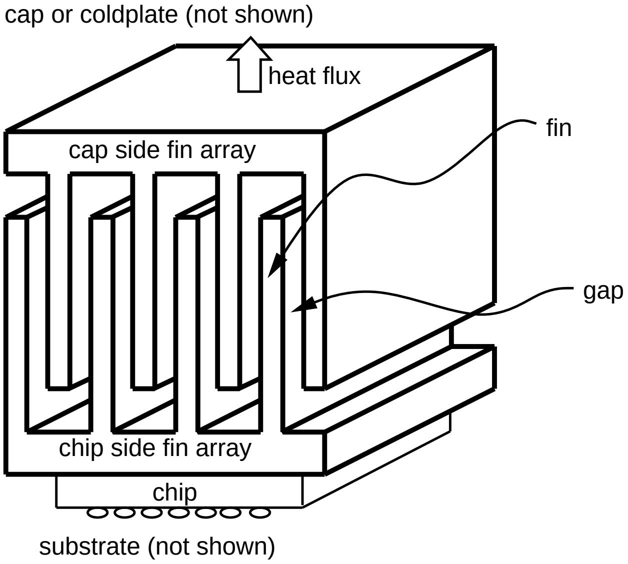

Figure 1: Generic interleaved fin thermal connector

The designer of a multichip module faces a choice between two basic module configurations: conventional and flipped-chip. In conventional hybrid construction, the back of the die is bonded to a thermally conductive substrate, and TABed or wire bonded to electrically connect it to the substrate wiring. Heat is removed by conduction through the substrate. The fundamental limitations to this approach are that the substrate design must trade off thermal and electrical characteristics, and that the chips cannot be positioned close together due to the space needed to wire the chips to the interconnect. In flipped-chip construction, the active face of the die is attached to the substrate via bumps or flipped-TAB. In selecting the flipped-chip approach, the designer can no longer remove heat through the substrate; even with a heavily bumped chip, heat transfer through the bumps to the substrate is inadequate for all but the lowest power levels. The advantages of flipped-chip construction are that the substrate can be optimized for its electrical characteristics and that the chips can be mounted closely together. The heat can then be removed either by direct immersion as in the Cray 3 [4] or indirectly by conduction through the back of the chip and into some cooling structure. While direct immersion in a dielectric fluid appears to be an attractive option, the high cost of the preferred perfluorinated fluids, and long term material compatibility issues have prevented its wider adoption. A fundamental difficulty with indirect conduction cooling of a flipped-chip module is obtaining efficient heat removal without transferring large forces to the die. If, for example, the chips are all bonded directly to a coldplate, then external forces such as shock and vibration, and internally generated forces such as differential expansion due to mismatched materials or thermal gradients, could apply large forces to the die and its electrical connections to the substrate. A further complication is the need to disassemble the module for repair. This is a requirement for most multichip modules, since fresh lot assembly yields fall as the number of chips increases.

Designers have taken a multitude of approaches to address these problems. There have been proposals for soldering the backs of the chips directly to the coldplate, utilizing creep in the indium solder to absorb cyclic thermal stresses [6]. While it was suggested that the indium creep provides adequate decoupling, most manufacturers prefer greater compliance. The IBM 4381 used grease to conduct approximately 45% of the heat directly from the top of each chip into a cap/heatsink, 30% was conducted into the cap/heatsink via the chip bumps and substrate, and the remaining 25% was removed through the substrate pins into the motherboard [2]. The Honeywell SLIC micropackage used a 0.2 mm thick conformal copper diaphragm held against the chips by the pressure of cooling water on the opposite side [18]. These approaches worked because the chips were relatively low powered.

The first successful indirectly cooled flipped-chip module with high powered chips was the IBM TCM, first used in the IBM 3081 series computers [3], [13]. Spring loaded aluminum pistons transferred heat from die to cap assembly, and a helium fill enhanced conduction across the various gaps. To accommodate tilt of the bumped chips, the piston face contacting the chip had a spherical shape. This point-contact arrangement was adequate for the small chips in the IBM 3081, but would prove challenging to scale to larger chips in later machines. NEC used a similar arrangement for their SX supercomputer series; but the pistons were screw-locked into position during initial assembly, and grease on the piston tip accommodated small scale strains [17]. An alternative approach used a coldplate with deep v-grooves and a trigonal prismatic piston at each die site [12]. While the bottom of the piston was flat to contact the die, the two faces mating with the coldplate v-groove had 200 mm radius cylindrical sector surfaces to accommodate one axis of chip tilt. The piston was also split along its axis of symmetry, and the two halves pressed apart by a helical compression spring; this accommodated the other axis of chip tilt and chip height variations. Each half of the piston was in constant line contact with its side of the coldplate v-groove. All these designs share a common limitation; they attempt to conduct the heat through a small area and across an approximately planar gap. The limited available area forces the use of tiny gaps and gap conductivity enhancing materials. Controlling manufacturing tolerances to maintain small clearances and confining gap conductivity enhancing materials can be problematical.

A fundamentally different approach is to increase the area of the joint. We are all familiar with the idea of adding fins to a heat sink to increase surface area and improve the heat transfer to a passing fluid stream. The interleaved fin thermal connector applies the same thinking to transferring heat across a pair of mating surfaces; both surfaces are extended to form an array of interleaved fins, the convoluted interface having a much larger surface area than the original planar mating surfaces. If the surface area is greatly enhanced, efficient heat transfer can be obtained without tiny gaps or a high conductivity interstitial material. Also, this interface arrangement provides 3 degrees of freedom of motion between the mating surfaces (two in translation, one in rotation). If the gap is sufficiently large, additional degrees of freedom can be provided with a limited range of motion. It is even possible to design the mating fin sets so that they can be repeatedly engaged and disengaged, thus providing a separable interface for module repair.

Analytical studies have investigated the idea of increasing surface area to enhance contact conduction heat transfer. Research into the effects of surface roughness on conduction across planar pressure contacts led to an analysis of saw-tooth, sine-wave, and square-wave profiles on mating surfaces [15]. The system of interest was the interface between a hard, high thermal conductivity material, and a softer, low conductivity material. It was found that low aspect-ratio (ratio of feature height to feature pitch) surface roughening treatments improved thermal performance, but there were diminishing returns with higher aspect-ratios. A subsequent analysis extended these results [14]. Neither study explicitly considered varying gaps between the mating surfaces or the possibilities of relative movement between the surfaces.

An early disclosure of a hardware implementation of the interleaved fin concept had one set of fins as an integral part of the coldplate [7]. The second set consisted of individual fins which floated in the spaces between the coldplate side fins. These individual fins were caged so they couldn’t fall out and were independently spring biased against the heated chip surface. While this accommodated a tilted chip, the line contact between the individual fins and the heated surface limited performance. A later implementation envisioned using an interleaved fin connector as an inseparable modular interposer [11]. This connector had flat top and bottom surfaces to contact the chip and coldplate. In between were two fixed sets of interleaved fins, spring biased to provide pressure contact at the chip and coldplate interfaces. This structure was designed to control the average inter-fin gap to 1/2 the fin thickness. The controlled gap and the flexibility of the fins could accommodate a tilted chip. While this addressed many of the system requirements, the two planar pressure contacts at the chip and coldplate interfaces limited overall performance.

The interleaved fin connector used in the Hitachi M-880 [10] was first described in an earlier patent [5]. One set of fins was an integral part of the module cap. Each chip site had a second, monolithic fin array with a planar or cylindrical-sector surface which was spring biased against the chip. Helium gas enhanced conduction across the small gap. The fin gap and/or the cylindrical surface against the chip were used to accommodate lateral chip tilt. Numerous variations of fin shape and biasing mechanisms were described. A subsequent paper described analysis of, and experiments on, the connector [1]. It included the effect on performance of machining tolerances, roughness, waviness and warp of the fins, and chip tilt. Models included the effect on apparent contact resistance when the fin gap approached the mean free path of the interstitial gas. Local design optimization was also discussed.

IBM S390/ES9000 Model 900 mainframe computers may also utilize an interleaved fin connector. A recent conference paper had an illustration depicting such a connector [16]. However in the accompanying oral presentation, it was stated that because of manufacturing costs and capabilities, IBM is currently shipping the bell piston, a refinement of the original TCM design. In either case, an oil fill has replaced the helium formerly used to enhance thermal conduction across gaps. Further details are promised in an upcoming article [19]. A recent IBM patent described many related interleaved fin connector designs [9]. While otherwise similar to the Hitachi approach, the module cap side of IBM’s preferred connector has thin and flexible flat spring fins which are biased against the mating rigid fins of the chip side fin array. Each slot in a chip side fin array contains a pair of the flexible module cap side fins pressing in opposite directions. Thus each spring fin can maintain intimate contact with its corresponding rigid fin, and still accommodate chip tilt and displacement. For a particular connector design for 10 mm square chips, an overall thermal resistance of $0.5°C/W$ was predicted, using a poly(alphaolefin) oil fill.

While both the IBM and Hitachi designs have a separable planar interface between the chip and the chip side fin array, it is possible to eliminate all such interfaces. This is done by bonding both the cap side and the chip side fin arrays to their respective mounting surfaces, and separating the assembly at the gap between the mating fin sets [8]. To avoid mechanical interference in the mated fin assemblies, generous gaps and/or matched assembly techniques may be used. A hinge or dowel aligns the cap with the module base, and tapered leading edges on the fins ease initial interpenetration of the fin assemblies. A particular implementation of this approach is described in **Design Example** section.

2. Theory

A simple lumped parameter model of the interleaved fin connector provides a useful introduction to its design and optimization. We begin by assuming that the connector consists of a square array of a large number of rectangular fins, that the gaps are the same on each side of each fin, and that the fins are fully intermeshed (figure 2). By assuming a large number of fins, we can safely ignore edge effects of the fins at the ends of the stack. By symmetry, each gap is served by half the thickness of each of the two adjacent fins. Thus, we choose fin centerlines to bound our elemental control volume consisting of two half-fins, and one gap. Looking at two limiting cases suggests appropriate values of lumped gap resistance and lumped fin resistance. If the fins were constructed of a material with infinite thermal conductivity ($k_f = \infty$), then each fin would be isothermal and at its respective base temperature, and each gap would have a thermal resistance of $\theta_g = s/(k_g l d)$. Similarly, if the gaps were vanishingly small or if the gap material had an infinite thermal conductivity across the gap, then adjacent half-fins would have the same temperature at adjacent points. The half-fins would behave thermally as if joined to form a single fin. This suggests a thermal resistance for a single fin (a pair of half-fins) of $\theta_f = l/(k_f b d)$. Summing a lumped fin and gap resistance, and dividing by the total number of fins (or gaps) provides an estimate of the thermal resistance across the active region of the entire connector.

$$\theta_{approx} = \frac{s+b}{d}\left(\frac{l}{k_f\, b\, d} + \frac{s}{k_g\, l\, d}\right) \tag{1}$$

$$= \frac{1}{\text{number of fins}}\ (\text{ resistance of one fin } + \text{ resistance of one gap })$$

Note that by assuming that the gaps are the same on each side of each fin, we obtain an upper bound on thermal resistance; as the gaps become increasingly asymmetrical, the gap resistance term will vanish, and the total thermal resistance will asymptotically approach that of the lumped fin resistances.

We notice that the active fin length l in equation 1 occurs in the numerator of the plate resistance expression and the denominator of the fin resistance expression. Thus there is some op-

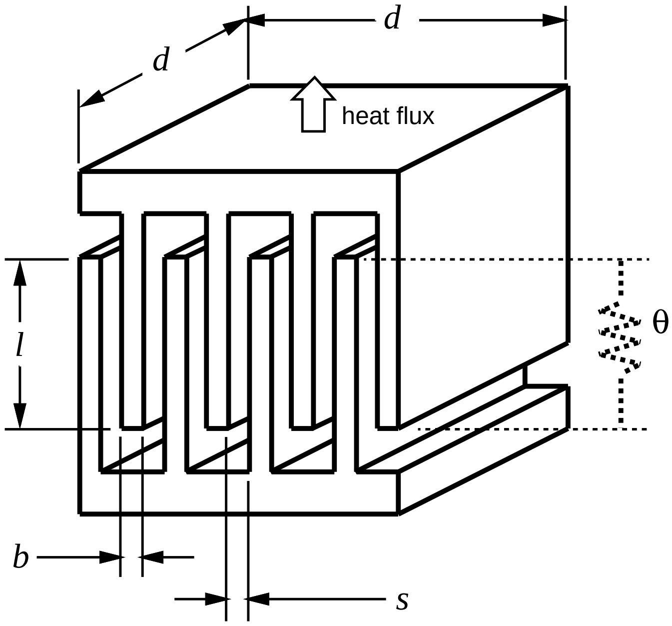

Figure 2: Nomenclature for interleaved fin thermal connectors

timum value of l which minimizes the sum of the two resistances, and hence the total thermal resistance θ_{approx}. This happens when the two terms are equal. Equating and solving for l gives the approximate optimal fin length $l_{approx-optimal}$.

$$l_{approx-optimal} = \sqrt{s\,b\,k_f / k_g} \tag{2}$$

The fin thickness b occurs in both the numerator and denominator of equation 1. By inspection, there is also an optimum value of b which minimizes the total thermal resistance. Differentiating θ_{approx} with respect to b and setting equal to zero, we solve for this approximate optimal fin thickness $b_{approx-optimal}$.

$$b_{approx-optimal} = l\sqrt{k_g / k_f} \tag{3}$$

Substituting the approximate optimal fin length and thickness back into equation 1, we obtain an approximate expression for the total thermal resistance of an optimal design.

$$\theta_{approx-optimal} = \frac{4s}{d^2\sqrt{k_f k_g}} \tag{4}$$

Once the connector footprint and the fin and gap materials have been chosen, the total thermal resistance of an optimal design depends only on the gap. Since the total thermal resistance is directly proportional to the gap, it is apparent that the gap is a variable of fundamental importance, and that we wish to design with the smallest possible gaps.

We can make a few additional observations based on this simple model. Substituting equation 2 into equation 3 and solving for the fin thickness b reveals that at the optima, the fin thickness equals the gap. This implies than in addition to tiny gaps, we'll want to design our connectors with very thin fins. Substituting $s = b$ back into equation 2 shows that the optimal fin aspect ratio (l/b), equals $\sqrt{k_f/k_g}$. When the fin and gap thermal conductivities are similar, the optimal design has short, stubby fins, and when the thermal conductivities differ greatly, the result is tall, thin fins.

Using a one-dimensional fin conduction model, we can more accurately calculate the fin temperature distribution and local heat transfer across the gap. Solving the one-dimensional fin equation with fixed base temperature and adiabatic fin tip boundary conditions yields:

$$\theta = \frac{l\,(s+b)}{k_f\, b\, d^2}\left(1 + \frac{2(\,1+\cosh\lambda)}{\lambda\,\sinh\lambda}\right) \qquad \text{where: } \lambda = 2l\sqrt{k_g/(k_f\, s\, b)} \tag{5}$$

A parametric comparison of this solution with the earlier approximation shows that the simple lumped-parameter model (equation 1) underestimates thermal resistance by up to 20%. This maximum error occurs near the optima, and the two models converge as deviation from optimality becomes large. This suggests that the simple lumped parameter equations are adequate for understanding tradeoffs and for preliminary design, and that more precise solutions are necessary only for accurate refinement of well characterized designs.

At the optima, the thermal resistance predicted by the two models differs by a fixed ratio of 1.138, so we can rewrite equation 4 as:

$$\theta_{optimal} = \frac{4.55\,s}{d^2\sqrt{k_f k_g}} \tag{6}$$

By comparison of the two models it can be shown that the lumped parameter model correctly predicts the optimal fin thickness $b_{optimal}$ to be equal to the gap s. However, the lumped parameter model overestimates the optimal fin length l by a factor of 1.14. Thus equation 2 can be rewritten as:

$$l_{optimal} = 0.88\, s\,\sqrt{k_f / k_g} \tag{7}$$

An important advantage of choosing a design point near the optima is that the connector will exhibit minimal sensitivity to changes of the design parameters. For example, beginning with an optimal design and then making the fins twice their optimal thickness results in only a 10% increase in thermal resistance. Similarly, there is approximately a 20% loss if the fins are half or twice their optimal length. Unfortunately, doubling the gap on an optimal design results in a more than 2x degradation of performance; maintaining a small gap is fundamental to interleaved fin connector performance. By assuming uniform fin gaps across the entire assembly, we've determined an upper bound on thermal resistance. Any real assembly will do better. A lower bound can be constructed by setting the gap resistance term in equation 1 to zero.

3. Design Example

The IBM and Hitachi interleaved fin connectors described earlier used thermal conductivity enhancing interstitial materials (oil and helium respectively) and had separable planar interfaces at the back surface of the chip. Here we'll look at an alternative implementation that that uses copper fins, air as the interstitial fluid, and is designed for separation at the gap between the mating fin sets. Assuming that a 15.6 mm square area is available for the fin array, and that an 84 micron fin gap is desired, equation 6 predicts the thermal resistance θ of an optimal connector will be 0.50 °*C*/*W*. Depending of course on the performance of the other cooling system elements, this connector would be adequate for chips dissipating 30-40 watts each. The optimal fin thickness b is 84 microns and by equation 7, the optimal fin length l is 9.0 mm. For ease of fabrication, we might choose a larger blade thickness, and perhaps shrink the gap slightly to compensate.

Figure 3: Prototype connector designed for separation at gap between the mating fin sets

Figure 3 shows a prototype of just such a connector. Fins are etched in strip format from 0.2 mm thick CDA151 zirconium-copper, a common lead-frame material, and separated from the strip with a punch tool that also forms a beveled lead-in at the tip of each fin. Arrays of the 9.0 mm high by 15.6 mm wide fins are soldered to 17.5 mm square OFHC copper bases. Mating fin arrays are manufactured as matched sets. The blades for a fin array set are sequential stacked into a fixture in alternating orientations and interwoven with a tape separator of Kapton® film[*]to

[*]trademark of E. I. du Pont de Nemours & Co.

maintain the desired gap. A thin layer of solder paste is then applied to one side of each of a pair of copper bases, and the bases clamped to opposite sides of the fin stack and reflowed into place. After cooling, the clamp and fixture are removed. The mating halves can then be broken loose and pulled apart slightly to increase the effective height of the assembly. Friction between the tape and the fins maintains the relative positions of the mating halves. A group of such assemblies is placed between the module cap and the flipped chips on their substrate. Aligned by dowels or a hinge, the cap and substrate are pressed together and the fin assemblies are reflow soldered or adhesive bonded into place. Then the cap and substrate are separated, the Kapton® tapes pulled loose, and the assembly cleaned. When the module is subsequently reassembled, the tapes are no longer present, and the resulting gaps structurally decouple cap and substrate. If a chip needs to be replaced, that die site receives a new matched fin assembly by a similar process. Test assemblies have been built with tape thicknesses, and hence gaps, ranging from 8 to 125 microns, though tapes less than 18 microns thick tear easily during assembly and the stiffness of tapes greater than 75 microns thick impedes winding while stacking the blades into the fixture.

Limited thermal testing has shown the performance of these connectors to be within the range predicted by the bounding models. While even connectors with 8 micron gaps engage and disengage easily, fins on the samples built to date are sufficiently warped and wavy that there is less lateral compliance than the designed-in gap would imply. For thermally optimal designs, the connectors can be shown to have large lateral spring rates beyond the free-play range provided by the fin gap.

4. Conclusions

Interleaved fin thermal connectors have become a mainstream cooling technology for high powered multichip modules. The principal challenge in designing such connectors is minimizing the gap between the fins while maintaining the required compliance. Straightforward analytical techniques are available to guide optimization. By the use of matched assembly techniques, high performance connectors have been demonstrated with the separable interface at the gap between the mating fin sets. Interleaved fin thermal connectors are not yet a mature technology, and we can expect considerable refinement and additional applications over the coming years.

5. Acknowledgements

The author wishes to thank Jack Berenholz and Boris Mirman for their assistance with portions of the analysis and for their constant encouragement throughout the project, and John Fitch, Susan Brown, and Wade McGillis for their many useful comments on an early draft of the paper.

6. References

[1] N. Ashiwake, T. Daikoku, K. Kawamula, and S. Zushi.
A Flexible Thermal Contactor for the Cooling of Electronic Components.
In *Proceedings of the 67th JSME Conference*, pages 264-266. March, 1990.

[2] R. G. Biskeborn, J. L. Horvath, E. B. Hultmark.
Integral Heat Cap Heat Sink Assembly for the IBM 4381 Processor.
In *Proceedings of 4th International Electronics Packaging Conference*, pages 468-474. IEPS, Baltimore, Maryland, October, 1984.

[3] R. C. Chu, U. P. Hwang, and R. E. Simons.
Conduction Cooling for an LSI Package: A One-Dimensional Approach.
IBM Journel Res Dev 26(1):45-54, 1982.

[4] Seymour Cray.
Keynote address at Supercomputing ’88 Conference.
November 15, 1988.

[5] T. Daikoku.
Cooling Device of Semiconductor Chips.
U. S. Patent #4,770,242.
September 13, 1988.

[6] R. Darveaux, L. Hwang, A. Reisman, and Iwona Turlik.
Thermal/Stress Analysis of a Multichip Package.
In *Proceedings of 39th Electronic Components Conference*, pages 668-671. IEEE/EIA, Houston, Texas, May, 1989.

[7] E. L. Dombroski and U. P. Hwang.
Thermal Conduction Stud.
IBM Technical Disclosure Bulletin 19(12):4683-4685, May, 1977.

[8] W. R. Hamburgen.
Apparatus and Method for Removal of Heat from Packaged Element.
U. S. Patent #4,800,956.
January 31, 1989.

[9] J. L. Horvath.
High Conduction Flexible Fin Cooling Module.
U. S. Patent #5,014,117.
May 7, 1991.

[10] Fumiyuki Kobayashi, et al.
Hardware Technology for HITACHI M-880 Processor Group.
In *Proceedings of 41st Electronic Components and Technology Conference*, pages 693-703. IEEE/EIA, Atlanta, Georgia, May, 1991.

[11] L. W. Lipschutz.
Flexible Thermal Conduction Element for Cooling Semiconductor Devices.
U. S. Patent #4,498,530.
February 12, 1985.

[12] A. Morihara, Y. Naganuma, H. Yokoyama.
High-Performance Line Contact Cooling Module for Multichip Packaging.
In *Proceedings of 2nd InterSociety Conference on Thermal Phenomenon in Electronic Systems - I-Therm*, pages 468-474. Las Vegas, Nevada, May, 1990.

[13] S. Oktay, H. C. Kammerer.
A Conduction Cooled Module for High-Performance LSI Devices.
IBM Journel Res Dev 26(1):55-66, 1982.

[14] E. J. Pinto and B. B. Mikic.
A Novel Design Concept for Reduction of Thermal Contact Resistance.
In *Proceedings of Int'l Symp on Cooling Technology for Electronic Equipment*, pages 461-473. Honolulu, Hawaii, March, 1987.

[15] B. J. Rozon, P. F. Galpin, and M. M. Yovanovich.
The Effect of Geometry on the Contact Conductance of Contigous Interfaces.
J. Spacecraft and Rockets 23(3):225-230, 1978.

[16] R. R. Tummala, H. R. Potts, and Shakil Ahmed.
Packaging Technology for IBM's Latest Mainframe Computers (S/390/ES9000).
In *Proceedings of 41st Electronic Components and Technology Conference*, pages 682-688. IEEE/EIA, Atlanta, Georgia, May, 1991.

[17] Toshihiko Watari and Hiroshi Murano.
Packaging Technology for the NEC SX Supercomputer.
In *Proceedings of 35th Electronic Components Conference*, pages 192-198. IEEE/EIA, May, 1985.

[18] E. A. Wilson.
Cooling Modern Mainframes - A Liquid Approach.
Computer Design :219-225, May, 1983.

[19] S. Young.
Module Cooling Technology for IBM's System 390/ES9000 Model 900 Mainframe Computers.
IBM Journel Res Dev , to be published in September, 1991.

WRL Research Reports

‘‘Titan System Manual.’’
Michael J. K. Nielsen.
WRL Research Report 86/1, September 1986.

‘‘Global Register Allocation at Link Time.’’
David W. Wall.
WRL Research Report 86/3, October 1986.

‘‘Optimal Finned Heat Sinks.’’
William R. Hamburgen.
WRL Research Report 86/4, October 1986.

‘‘The Mahler Experience: Using an Intermediate Language as the Machine Description.’’
David W. Wall and Michael L. Powell.
WRL Research Report 87/1, August 1987.

‘‘The Packet Filter: An Efficient Mechanism for User-level Network Code.’’
Jeffrey C. Mogul, Richard F. Rashid, Michael J. Accetta.
WRL Research Report 87/2, November 1987.

‘‘Fragmentation Considered Harmful.’’
Christopher A. Kent, Jeffrey C. Mogul.
WRL Research Report 87/3, December 1987.

‘‘Cache Coherence in Distributed Systems.’’
Christopher A. Kent.
WRL Research Report 87/4, December 1987.

‘‘Register Windows vs. Register Allocation.’’
David W. Wall.
WRL Research Report 87/5, December 1987.

‘‘Editing Graphical Objects Using Procedural Representations.’’
Paul J. Asente.
WRL Research Report 87/6, November 1987.

‘‘The USENET Cookbook: an Experiment in Electronic Publication.’’
Brian K. Reid.
WRL Research Report 87/7, December 1987.

‘‘MultiTitan: Four Architecture Papers.’’
Norman P. Jouppi, Jeremy Dion, David Boggs, Michael J. K. Nielsen.
WRL Research Report 87/8, April 1988.

‘‘Fast Printed Circuit Board Routing.’’
Jeremy Dion.
WRL Research Report 88/1, March 1988.

‘‘Compacting Garbage Collection with Ambiguous Roots.’’
Joel F. Bartlett.
WRL Research Report 88/2, February 1988.

‘‘The Experimental Literature of The Internet: An Annotated Bibliography.’’
Jeffrey C. Mogul.
WRL Research Report 88/3, August 1988.

‘‘Measured Capacity of an Ethernet: Myths and Reality.’’
David R. Boggs, Jeffrey C. Mogul, Christopher A. Kent.
WRL Research Report 88/4, September 1988.

‘‘Visa Protocols for Controlling Inter-Organizational Datagram Flow: Extended Description.’’
Deborah Estrin, Jeffrey C. Mogul, Gene Tsudik, Kamaljit Anand.
WRL Research Report 88/5, December 1988.

‘‘SCHEME->C A Portable Scheme-to-C Compiler.’’
Joel F. Bartlett.
WRL Research Report 89/1, January 1989.

‘‘Optimal Group Distribution in Carry-Skip Adders.’’
Silvio Turrini.
WRL Research Report 89/2, February 1989.

‘‘Precise Robotic Paste Dot Dispensing.’’
William R. Hamburgen.
WRL Research Report 89/3, February 1989.

‘‘Simple and Flexible Datagram Access Controls for Unix-based Gateways.’’
Jeffrey C. Mogul.
WRL Research Report 89/4, March 1989.

‘‘Spritely NFS: Implementation and Performance of Cache-Consistency Protocols.’’
V. Srinivasan and Jeffrey C. Mogul.
WRL Research Report 89/5, May 1989.

‘‘Available Instruction-Level Parallelism for Superscalar and Superpipelined Machines.’’
Norman P. Jouppi and David W. Wall.
WRL Research Report 89/7, July 1989.

‘‘A Unified Vector/Scalar Floating-Point Architecture.’’
Norman P. Jouppi, Jonathan Bertoni, and David W. Wall.
WRL Research Report 89/8, July 1989.

‘‘Architectural and Organizational Tradeoffs in the Design of the MultiTitan CPU.’’
Norman P. Jouppi.
WRL Research Report 89/9, July 1989.

‘‘Integration and Packaging Plateaus of Processor Performance.’’
Norman P. Jouppi.
WRL Research Report 89/10, July 1989.

‘‘A 20-MIPS Sustained 32-bit CMOS Microprocessor with High Ratio of Sustained to Peak Performance.’’
Norman P. Jouppi and Jeffrey Y. F. Tang.
WRL Research Report 89/11, July 1989.

‘‘The Distribution of Instruction-Level and Machine Parallelism and Its Effect on Performance.’’
Norman P. Jouppi.
WRL Research Report 89/13, July 1989.

‘‘Long Address Traces from RISC Machines: Generation and Analysis.’’
Anita Borg, R.E.Kessler, Georgia Lazana, and David W. Wall.
WRL Research Report 89/14, September 1989.

‘‘Link-Time Code Modification.’’
David W. Wall.
WRL Research Report 89/17, September 1989.

‘‘Noise Issues in the ECL Circuit Family.’’
Jeffrey Y.F. Tang and J. Leon Yang.
WRL Research Report 90/1, January 1990.

‘‘Efficient Generation of Test Patterns Using Boolean Satisfiablilty.’’
Tracy Larrabee.
WRL Research Report 90/2, February 1990.

‘‘Two Papers on Test Pattern Generation.’’
Tracy Larrabee.
WRL Research Report 90/3, March 1990.

‘‘Virtual Memory vs. The File System.’’
Michael N. Nelson.
WRL Research Report 90/4, March 1990.

‘‘Efficient Use of Workstations for Passive Monitoring of Local Area Networks.’’
Jeffrey C. Mogul.
WRL Research Report 90/5, July 1990.

‘‘A One-Dimensional Thermal Model for the VAX 9000 Multi Chip Units.’’
John S. Fitch.
WRL Research Report 90/6, July 1990.

‘‘1990 DECWRL/Livermore Magic Release.’’
Robert N. Mayo, Michael H. Arnold, Walter S. Scott, Don Stark, Gordon T. Hamachi.
WRL Research Report 90/7, September 1990.

‘‘Pool Boiling Enhancement Techniques for Water at Low Pressure.’’
Wade R. McGillis, John S. Fitch, William R. Hamburgen, Van P. Carey.
WRL Research Report 90/9, December 1990.

‘‘Writing Fast X Servers for Dumb Color Frame Buffers.’’
Joel McCormack.
WRL Research Report 91/1, February 1991.

‘‘Analysis of Power Supply Networks in VLSI Circuits.’’
Don Stark.
WRL Research Report 91/3, April 1991.

‘‘TurboChannel T1 Adapter.’’
David Boggs.
WRL Research Report 91/4, April 1991.

‘‘Procedure Merging with Instruction Caches.’’
Scott McFarling.
WRL Research Report 91/5, March 1991.

‘‘Don’t Fidget with Widgets, Draw!.’’
Joel Bartlett.
WRL Research Report 91/6, May 1991.

‘‘Pool Boiling on Small Heat Dissipating Elements in Water at Subatmospheric Pressure.’’
Wade R. McGillis, John S. Fitch, William R. Hamburgen, Van P. Carey.
WRL Research Report 91/7, June 1991.

‘‘Incremental, Generational Mostly-Copying Garbage Collection in Uncooperative Environments.’’
G. May Yip.
WRL Research Report 91/8, June 1991.

‘‘Interleaved Fin Thermal Connectors for Multichip Modules.’’
William R. Hamburgen.
WRL Research Report 91/9, August 1991.

WRL Technical Notes

‘‘TCP/IP PrintServer: Print Server Protocol.’’
Brian K. Reid and Christopher A. Kent.
WRL Technical Note TN-4, September 1988.

‘‘TCP/IP PrintServer: Server Architecture and Implementation.’’
Christopher A. Kent.
WRL Technical Note TN-7, November 1988.

‘‘Smart Code, Stupid Memory: A Fast X Server for a Dumb Color Frame Buffer.’’
Joel McCormack.
WRL Technical Note TN-9, September 1989.

‘‘Why Aren’t Operating Systems Getting Faster As Fast As Hardware?’’
John Ousterhout.
WRL Technical Note TN-11, October 1989.

‘‘Mostly-Copying Garbage Collection Picks Up Generations and C++.’’
Joel F. Bartlett.
WRL Technical Note TN-12, October 1989.

‘‘Limits of Instruction-Level Parallelism.’’
David W. Wall.
WRL Technical Note TN-15, December 1990.

‘‘The Effect of Context Switches on Cache Performance.’’
Jeffrey C. Mogul and Anita Borg.
WRL Technical Note TN-16, December 1990.

‘‘MTOOL: A Method For Detecting Memory Bottlenecks.’’
Aaron Goldberg and John Hennessy.
WRL Technical Note TN-17, December 1990.

‘‘Predicting Program Behavior Using Real or Estimated Profiles.’’
David W. Wall.
WRL Technical Note TN-18, December 1990.

‘‘Systems for Late Code Modification.’’
David W. Wall.
WRL Technical Note TN-19, June 1991.