

JULY 1990

WRL Research Report 90/6



A One-Dimensional Thermal Model for the VAX 9000™ Multi Chip Units

John S. Fitch

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, UCO-4
100 Hamilton 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
DARPA Internet:	WRL-Techreports@decwrl.dec.com
CSnet:	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.

Abstract

A thermal resistance network is used to predict the performance of the Multi Chip Units (MCUs) in the VAX 9000™ computer. This branched network is comprised of resistors defined by analytical, numerical and experimental techniques. Effects of thermal conduction, contact resistance and convection are included. A comparison is made between the model's temperature predictions and test data. Agreement within 15% is achieved, demonstrating that the chips in the MCU will operate well below their specification limit of 85° C.

This is a preprint of a paper that will be presented at the
American Society of Mechanical Engineers,
Winter Annual Meeting, Dallas, Texas, 1990.

Copyright © 1990 ASME

Table of Contents

Nomenclature	1
1. Introduction	2
2. Hardware Description	3
2.1. Silicon Chips	4
2.2. Epoxy	4
2.3. Baseplate	4
2.4. Heat Sink	4
3. Thermal Model	5
3.1. Silicon Thin Film Heater Chips	6
3.2. Epoxy	6
3.3. Baseplate	6
3.4. Interface	8
3.5. Heat Sink	9
4. Test Results and Model Comparisons	11
5. Conclusions	13
Acknowledgments	15
References	17

List of Figures

Figure 1:	An exploded side view of the MCU.	3
Figure 2:	Thermal resistance network for the MCU with up to n various chips.	5
Figure 3:	Baseplate temperature profiles in a plane located 1.3mm below the top of the baseplate. T_{ref} is a temperature at the bottom of the baseplate. $L = 5\text{ mm} =$ one half of this heater chip's width.	7
Figure 4:	Thermal contact resistance between the baseplate and the heat sink.	10
Figure 5:	Thermal resistance of a heat sink with aluminum fins.	11
Figure 6:	Average surface temperatures of silicon heater chips.	12

Nomenclature

A	area (m^2)
B	coefficient
C	coefficient
H	contact microhardness (MPa)
h	heat conductance (W/m^2-K)
k	thermal conductivity ($W/m-K$)
k_s	harmonic mean thermal conductivity ($W/m-K$)
L	half the width of a chip (m)
m	mean absolute asperity slope
M	gas parameter
Nu	Nusselt number
P	contact pressure (MPa)
Pr	Prandtl number
Q	heat flow rate (W)
R	thermal resistance (C/W)
r	radial distance from bolt center (m)
Re	Reynolds number based on fin diameter
T	temperature (C)
T_j	average chip junction temperature (C)
V	air flow rate (m^3/s)
x	distance (m)
Y	distance between contacting planes (m)
σ	effective RMS surface roughness (m)

Subscripts

i	chip or circuit branch i
j	chip junction
bp	baseplate
$chip$	chip
$epoxy$	epoxy
hs	heat sink
int	baseplate-heat sink interface
air	ambient air

1. Introduction

The VAX 9000™ computer extends Digital's style of computing into the main frame domain. It is a highly reliable computer with five times the performance of previous VAX™ systems. This is achieved, in part, by packaging several Multi Chip Units (MCUs) in close proximity to each other. These MCUs generate large amounts of heat, yet are cooled efficiently with air. Designing MCUs which maintain low chip temperatures, leading to higher system reliability, was a major goal of the program. A thermal model which aided in the design process was essential.

Modeling Multi Chip Units is generally more complicated than modeling single chip packages. This is due to three major differences: 1) Multiple heat sources mean that many of the heat paths are in parallel, leading to more complicated thermal networks. 2) A chip can influence its neighbor, particularly if some are operating intermittantly. 3) Within one computer there are many MCUs, each with different chips in different patterns. A thermal model must be robust enough to account for these complexities and variations. In many cases, the odd geometries and boundary conditions in MCUs are too complex to model analytically. Layout variations and complexities may make numerical grids large and solutions impractical, particularly when design changes are frequent, requiring repeated solutions.

The thermal model must serve the needs of packaging engineers concerned with electrical and mechanical performance, reliability engineers concerned with maintaining a cool, reliable system, and manufacturing inspectors who must verify the thermal performance of manufactured units. When the model must be used repeatedly by a large community, a familiar and easily used format, such as a spread sheet, is desirable.

In many cases applying a one-dimensional network of thermal resistances to the MCU meets these needs. Each component in the thermal path can be assigned a resistance. Accurately defining that resistance then becomes the challenging part of the model. In some cases, direct analytical relationships can be used. Numerical results that are easily scaled to changes in chip types or powers may also be useful. Finally, experimental results lead to empirical definitions which fine tune the resistance values.

The VAX 9000 MCU thermal model is comprised of resistances determined using all of these techniques. The chip and epoxy are analytically described using Fourier's law. Numerical results are combined with experimental results to describe the baseplate and contact resistance between it and the heat sink. Test results define the performance of the heat sink empirically. The model is easy to use and helped make the VAX 9000 a very fast computer with chips operating well below the specified operating limit of 85°C.

This paper explains how the VAX 9000 MCU thermal model was created. Examples of calculations and test results are given with sufficient detail that the approach can be understood and applied to other situations.

2. Hardware Description

The VAX 9000 MCUs can accommodate various configurations of gate array, RAM and custom chips. The chips are surrounded by the High Density Signal Carrier (HDSC), a complex polyimide and copper structure which supplies signals and power. Connections from the chips to this carrier are via Tape Automated Bonding (TAB). The outer leads of the TAB are soldered to the HDSC. Both the HDSC and the chips are supported by a copper baseplate. The HDSC, which generates no significant heat, is simply laminated to the copper baseplate. Each chip is epoxied to the baseplate through cutouts in the HDSC. The baseplate serves as the assembly foundation for the electrical components. An air-cooled heat sink is attached to the opposite side of the baseplate with a regular array of nine cap screws. See figure 1.




Figure 1: An exploded side view of the MCU.

2.1. Silicon Chips

There are four sizes of silicon chips in the VAX 9000. The gate array chips are $9.8 \times 9.8\text{mm}$. The custom clock chip is $6.4 \times 6.4\text{mm}$. Some MCUs require RAM chips that are either $3.6 \times 4.9\text{mm}$ or $4.2 \times 6.3\text{mm}$. All are roughly 0.5mm thick. These emitter-coupled logic (ECL) devices have high power densities: A gate array may consume 30 watts (30 W/cm^2), while a small RAM consumes up to 2 watts (11 W/cm^2).

2.2. Epoxy

The bottom sides of the chips are epoxied to the top of the copper baseplate. Several attributes are required of this epoxy. Electrically, it must be an insulator. Thermally, it must be an excellent conductor. The epoxy must be resilient because the thermal expansion coefficient of the copper baseplate is much higher than that of the silicon. Therefore, a high percentage of elongation is needed in the epoxy. Finally, it must perform consistently in the manufacturing process, yielding thin, strong, continuous bond lines.

The resin system used has a very high percentage of elongation. To obtain a high thermal conductivity and an electrically insulating joint, fine diamond particles are added to the resin. By keeping the bond line thickness between 0.025 and 0.05mm , the thermal performance is acceptable, and the structural integrity of the joint is assured.

2.3. Baseplate

The structural foundation for the MCU is the baseplate. The HDSC is laminated to this flat, rigid, $10 \times 10 \times 0.9\text{cm}$ thick plate. The chips are bonded to the plate through the cutouts in the HDSC.

The baseplate must have a high thermal conductivity so that the heat introduced at the epoxy joint can spread effectively through its thickness. Nine blind tapped holes on 3 cm centers in the bottom of the baseplate allow the heat sink to be attached with #10-32 cap screws. Chrome-copper alloy (CDA 182) was found to meet these requirements. To reduce scratching during manufacturing, an electroless nickel plating is applied.

2.4. Heat Sink

A goal of the VAX 9000 design was to use quiet, low velocity air to cool the MCUs. This was accomplished by utilizing a pin fin heat sink for each MCU.

This heat sink is made by pressing cylindrical pins into an aluminum base. The 600 staggered pins increase the heat transfer surface area by 8:1 over that of the flat base alone. Various designs of the pin fin heat sink were considered, including copper and aluminum pins of different diameters and pressing the pins into the base versus dip brazing or casting them.

Air is supplied perpendicularly to the center of the heat sink, from a 5cm diameter nozzle. The flow behaves like uniformly-approaching flow or wedge flow, since the nozzle is large compared to the heat sink. As the air enters the center of the fin array and hits the base of the heat sink, it axisymmetrically turns and passes across the fins in a radial direction.

3. Thermal Model

The thermal network for the MCU consists of conduction and convection paths. Heat generated on the top surface of the chips, passes through the silicon and then through the epoxy. From the epoxy, the heat spreads into the baseplate, through the bolted joint and into the heat sink where it is liberated to the passing air. A very small amount of heat passes out of the chip through the TAB leads and is not included in the model. The thermal network for the MCU may be depicted as in figure 2.




Figure 2: Thermal resistance network for the MCU with up to n various chips.

Note that portions of the network are in parallel. Since an MCU includes chips of various powers and cross-sectional areas, parallel paths are needed to adequately describe the system. The temperature at the bottom of the baseplate is assumed to be isothermal. Experimental results support this assumption. Thermal coupling from chip to chip, within the baseplate, is not handled explicitly. This would require that the model be multi-dimensional and possibly solved iteratively. This coupling is discussed in the *Baseplate* section below. From the bottom of the baseplate to the air, the network is based on the total MCU area and power. A boundary condition on the system is the ambient air temperature. Note that the heat fluxes are constant, but not equal, through the various branches of the network. Each resistance in this network is discussed below.

3.1. Silicon Thin Film Heater Chips

In order to estimate the thermal performance of the MCUs early in the program, silicon thin film resistance heaters were used to simulate the real chips. These heater chips were produced to the same dimensions as the real chips. A thin film of NiCr across the top of the silicon provided a uniform resistance. These heater chips allowed us to examine various MCU configurations and verify the thermal model. The following results and discussions are based on experiments with such heater chips.

A simple model describes the temperature drop through a chip by Fourier's law in one dimension. This description requires that the heat flux be assumed uniform. The expression below provides the average thermal resistance through a chip with a given thickness, conductivity and area.

$$R_{chip} = \frac{\Delta x_{chip}}{k_{chip} A_{chip}} \quad (1)$$

3.2. Epoxy

Fourier's law is also used to describe the resistance through the epoxy as

$$R_{epoxy} = \frac{\Delta x_{epoxy}}{k_{epoxy} A_{epoxy}} \quad (2)$$

The thermal conductivity of this new diamond-filled epoxy was not known. Tests were designed to determine this essential property in geometries similar to the real epoxy joint. Experiments produced a value of $k_{epoxy} = 1.4 \text{ W/m-K}$ with an experimental uncertainty of $\pm 10\%$.

3.3. Baseplate

Analytical, numerical and experimental techniques were used to understand the thermal performance of the baseplate. Since this was the element that coupled all the chips to the heat sink, it received a great deal of attention. Analytical techniques were initially attempted using the results of Kennedy. [3] Bolt holes in the baseplate and a non-uniform contact pressure distribution made this solution difficult to use, and results were found to be unreliable.

Numerical models were used to study one symmetrical quadrant of a chip and the components between it and the cooling air. Both axisymmetrical (SINDA) and cartesian (ANSYS®) grid geometries were investigated. This method proved very useful, as many baseplate thicknesses, materials and bolt hole sizes could be evaluated to determine optimum combinations. The boundary conditions at the bottom of the baseplate required particular attention, and are discussed in the *Interface* section.

An experiment was run to determine how well the numerical models described the temperature field within the baseplate. Temperatures within the baseplate were measured by probing 0.5mm diameter holes with thermocouples. This provided plots of temperature versus radial distance from the center of the chip, in various planes within the thickness of the baseplate. A comparison of two numerical results and the temperature measurements within a baseplate with

nine 30 watt chips on it, is shown in figure 3. The reasonable agreement between the experimental and numerical results permitted us to use the numerical model to determine the temperature distribution within the baseplate.

The network in figure 2 does not account for thermal coupling from chip to chip within the baseplate. At the cost of much greater complexity, the accuracy of this model could be improved by replacing the one-dimensional baseplate resistances with a mesh of inter-connected resistances. However, our simpler model may be justified by noting that already at $r/L = 3$ (1.5 chip widths away from the center of the chip) $\frac{dT_{bp}}{dr}$ approaches zero. This suggests the spacing at which the heat flux from one chip begins to encroach upon that of its neighbor. All gate arrays are spaced at least three chip widths apart, reducing their thermal influence on each other. RAM chips are spaced somewhat closer together, but their power per unit area of baseplate is lower.

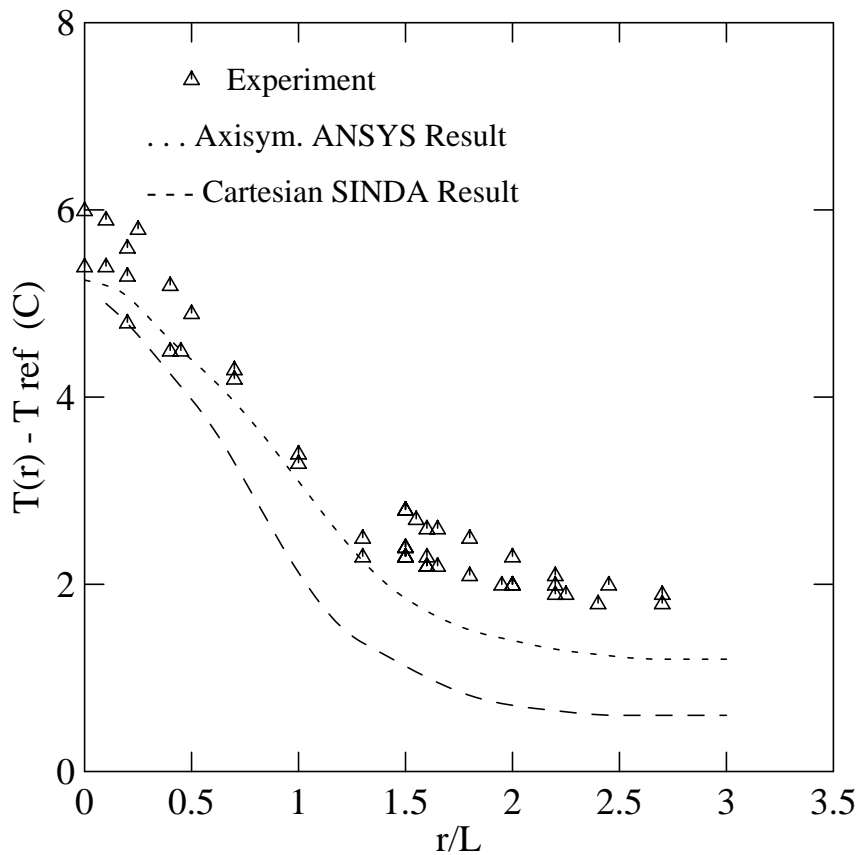


Figure 3: Baseplate temperature profiles in a plane located 1.3mm below the top of the baseplate. T_{ref} is a temperature at the bottom of the baseplate. $L = 5\text{ mm}$ = one half of this heater chip's width.

Temperature profiles such as those in figure 3 are different for chips with various sizes and powers. Obtaining these specific curves for each unique chip layout would require hundreds of numerical models. In order to have an empirical model that easily fits the spread sheet paradigm, the temperature drop through the baseplate was converted into an apparent heat conductance for that portion of the baseplate.

$$h_{bp,i} = \frac{Q_i}{\Delta T_{bp,i} A_{bp,i}} \quad (3)$$

where $A_{bp,i}$ is each chip's allotted portion of the baseplate. For example, for a symmetrical layout of nine chips on an MCU, $A_{bp,i} = \frac{1}{9}$ of the total baseplate area. Or for a more complex case, if an MCU were made up of one tight cluster of six RAMS and seven separate custom chips, then the entire RAM array might be assigned $\frac{1}{8}$ of the baseplate's total area. Each of the six RAMs could then be allotted $\frac{1}{6}$ of the RAM array area or $\frac{1}{48}$ of the entire baseplate area. This approximation becomes less valid for highly asymmetrical chip layouts. $A_{bp,i}$ is graphically determined from MCU drawings. The $\Delta T_{bp,i}$ used in equation 3 was computed and measured for a particular, but representative, chip layout. The thermal resistance through the baseplate may now be approximated by assuming that the local heat conductance is the same throughout the baseplate. Under any particular chip,

$$h_{bp} = h_{bp,i} \quad (4)$$

which yields the following:

$$R_{bp,i} = \frac{1}{h_{bp} A_{bp,i}} \quad (5)$$

While it is clear that these approximations can limit the accuracy of the model, particularly for unusual chip layouts, the simplicity makes the model accessible to a large user community.

3.4. Interface

A contact resistance exists between the baseplate and the heat sink. This resistance can be very significant depending on materials, the condition of the surfaces and the contact pressure. A detailed understanding of these parameters was needed to design an effective, reliable joint.

Manufacturing and performance requirements constrained the material selection for the baseplate (Cr-Cu) and heat sink (Al). The nickel plating on the copper baseplate increases the thermal resistance of the joint, since it increases the surface hardness, however the benefit of scratch resistance outweighs the thermal cost. Experimentally it was found that a uniform surface finish between 0.2 - 0.4 micron RMS minimized the contact resistance. Finishes outside this range increased the resistance.

The joint contact pressure can be controlled by choosing an appropriate torque for the cap screws that attach the heat sink. However, the pressure distribution in the interface is not uniform, and depends on variables such as the bolt locations and sizes, plate materials and thicknesses.

To properly estimate the contact resistance in this joint, the pressure distribution was investigated. Mikic and Gould have described a method to determine the contact area between two thin plates that are bolted together at the center. [5] Analysts at Digital built a similar numerical model. This numerical scheme was used to predict the interface pressure versus distance, r , from the bolt. It is interesting to note that the model suggests that the contact pressure drops off to zero at roughly three bolt radii from the center of the bolt. The local pressure was then used in a contact conductance expression suggested by Yovanovich. [6] This expression defines the con-

ductance through the solid-to-solid contacts and the air gaps, which are much less significant in our case.

$$h(r) = 1.25 \frac{m k_s}{\sigma} \left(\frac{P(r)}{H} \right)^{0.95} + \frac{k_{air}}{Y + \sigma M} \quad (6)$$

$$m = \sqrt{m_{bp}^2 + m_{hs}^2} \quad (7)$$

$$k_s = \frac{2 k_{bp} k_{hs}}{k_{bp} + k_{hs}} \quad (8)$$

$$\sigma = \sqrt{\sigma_{bp}^2 + \sigma_{hs}^2} \quad (9)$$

$$Y = 1.53 \sigma \left(\frac{P(r)}{H} \right)^{-0.097} \quad (10)$$

where m is the mean absolute asperity slope, k_s is the harmonic mean thermal conductivity of the two contacting solids, σ is the effective RMS surface roughness, H is the contact microhardness of the softer material, Y is the distance between the mean planes of the contacting rough surfaces [1], and M is a gas parameter which is dependent upon surface characteristics and thermodynamic properties. [6] Equation 6 yields a local heat conductance, which was used in the baseplate numerical model, or may be integrated around the bolt for an average value.

Rearranging the average heat conductance yields an average thermal resistance for the interface.

$$R_{int} = \frac{1}{h A_{int}} \quad (11)$$

Equation 6 was used to help make decisions about various design options. Once a design was chosen, experiments were performed to evaluate the interface performance. Figure 4 shows results that were easily reproduced, even when parts were mixed and matched. At low contact pressures (low bolt torques) the resistance becomes erratic and increases toward a limit described by conduction through an air gap. At high contact pressures all significant plastic deformations in the surface asperities have taken place, maximizing the solid-to-solid contact area, thereby minimizing the thermal resistance. Based on thread strength in the baseplate, a bolt torque of 4 N-m was specified for the MCU assembly. $R_{int} = 0.01^\circ\text{C}/\text{W}$ was chosen for use in the model since it is a reasonable upper limit given normal process variations.

3.5. Heat Sink

Analytical modeling of the heat sink proved to be very difficult. Air flow approaching perpendicularly to a flat heat sink without fins can be described by wedge flow solutions [2], but the addition of pin fins in the flow field disqualifies direct use of this type of solution. At best, closed form solutions reveal trends which aid in choosing fin characteristics. Empirical expressions for the heat sink performance were the only reasonable methods to use in our model. The design of the heat sink was strongly driven by manufacturability and cost considerations. Once a design was chosen, tests were run on the heat sink to examine nozzle designs, pressure drops, thermal resistances and noise at various flow rates. These data were then fitted to an approximate theoretical correlation.

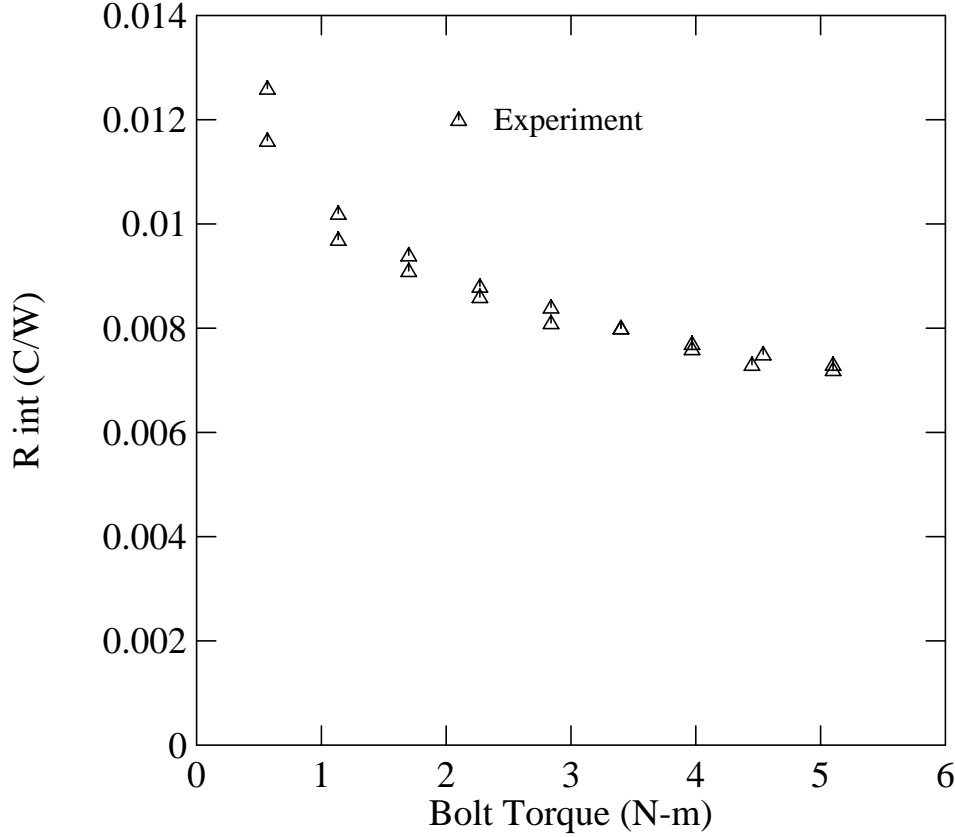


Figure 4: Thermal contact resistance between the baseplate and the heat sink.

Theoretically, flow over banks of cylinders may be described by the following relationship,

$$Nu = B Pr^{.33} Re^C \quad (12)$$

where B and C are constants based upon the cylinder size and row and column layout. [4] The VAX 9000 heat sink may be described as a bank of cylinders with air passing through the field radially. For our geometry, values for the constants might be expected to be in the neighborhood of $B=0.5$ and $C=0.57$. There are two major difficulties in directly applying this kind of model. One is that since the flow of air is spreading out radially, the Reynolds number decreases with distance from the center of the heat sink. The second is that our fin layout is staggered differently from the available tube bank correlations. In spite of this, the experimental data of figure 5 had a good fit to the correlation in equation 12.

Noting that

$$R \sim \frac{1}{Nu} \quad (13)$$

and the volumetric flow rate is proportional to the Reynolds number,

$$V \sim Re \quad (14)$$

our data fits the following empirical expression.

$$R_{hs} = 0.012 Pr^{-.33} V^{-.55} \quad (15)$$

Equation 15 provides an adequate prediction of R_{hs} for any flow rate of interest in this application. The change in the enthalpy of the air is included in the resistance of the heat sink. A single 5cm diameter nozzle per MCU produced the best performance and was used for the data of figure 5. Other designs tested consisted of single or multiple smaller nozzles.




Figure 5: Thermal resistance of a heat sink with aluminum fins.

4. Test Results and Model Comparisons

By assembling the resistances described above, a model can be built for any particular MCU. The model can handle any of the four chip types in various layouts. A spread sheet was used to handle this task, making all of the calculations quick and easy for a large community, including those without thermal engineering experience. For any particular chip, i , the predicted junction temperature is

$$T_{j,i} = (R_{chip,i} + R_{epoxy,i} + R_{bp,i}) Q_i + (R_{int} + R_{hs}) Q + T_{air} \quad (16)$$

with an uncertainty of $\pm 15\%$.

To verify the model's accuracy, the silicon heater chips were assembled to baseplates in place of real chips. Several MCU patterns with all four heater chip types were built using normal manufacturing equipment and processes. While these MCUs were powered, an infra-red camera

was used to determine the average temperatures on the surface of the chips. The uncertainty for the temperature measurement was ± 4 C. These measured temperatures were then compared to those predicted by the thermal model. Figure 6 shows good agreement between the model and

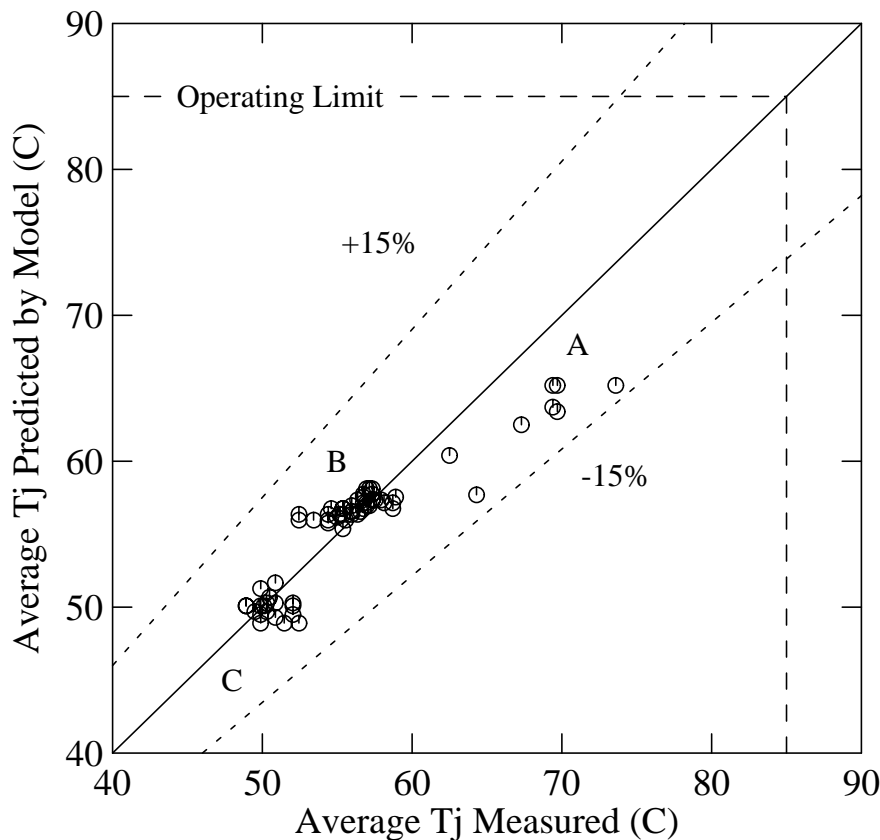


Figure 6: Average surface temperatures of silicon heater chips.

experiment. As expected, the hottest devices are the gate arrays (figure 6, group A), while the RAMS run cooler. The data in figure 6 are from many different MCUs with various chip patterns and total powers ranging from 134 to 220 watts. It is important to note that the same type of chip on two different MCUs may not run at the same temperature. Note that some low-heat flux RAMS (figure 6, group B) ran hotter than high heat flux RAMS (figure 6, group C). This occurred when the low heat flux RAMS were on an MCU with high total power, and vice versa. A particular chip's operating temperature is directly affected by the total MCU power, as seen in the thermal network, figure 2. Herein lies a complication with modeling a multi chip unit one-dimensionally.

It is also interesting to note the percentages of the total temperature rise that each component in the thermal network accounts for. For a typical MCU generating 200 watts, the chip and contact interface account for less than 5% each. The epoxy and baseplate each account for 15 to 20%. The heat sink has the largest impact on the chip temperature, contributing more than 50% of the total temperature rise.

5. Conclusions

For the VAX 9000, the many MCU chip layouts demanded a simple modeling approach. A thermal network composed of resistances representing major components in the MCU was developed. A combination of analytical, numerical and experimental approaches were used to derive the values of these thermal resistances. The measured temperatures of heater chips are within 15% of the model's predictions, indicating that chips on the VAX 9000 should operate well below the specified limit of 85°C.

Although this specific thermal model can not be used directly on other multi chip module designs, it suggests that this approach to modeling can have universal appeal to a large user community and can be done quite simply and effectively.

Acknowledgments

Much of the thermal hardware described here was originated by Jack Berenholz, Pat Gildea and many others. James Lindsay, Wade McGillis and Lou Palmieri provided experimental data. Karl Jacobson, Sheera Knecht, Boris Mirman and Fabio Piergentili performed numerical analyses. Bill Hamburg and Brian Reid reviewed the manuscript.

References

- [1] V. W. Antonetti.
Thermal Contact Resistance In Microelectronic Equipment.
The International Society for Hybrid Microelectronics 7(3):44-50, September, 1984.
- [2] W. M. Kays and M. E. Crawford.
Convective Heat and Mass Transfer.
McGraw-Hill, New York, 1980.
ISBN 0-07-033457-9. See in particular chapters 7 and 9.
- [3] D. P. Kennedy.
Spreading Resistance in Cylindrical Semiconductor Devices.
Applied Physics 31(8):1490, August, 1960.
- [4] F. Kreith and W. Z. Black.
Basic Heat Transfer.
Harper & Row, New York, 1980.
ISBN 0-700-22518-8. See in particular chapter 5.
- [5] Mikic and Gould.
MIT Heat Transfer Laboratory - Report no. DSR71821-68.
Technical Report, MIT, June, 1970.
- [6] M. M. Yovanovich.
New Contact and Gap Conductance Correlations for Conforming Rough Surfaces.
In *AIAA-81-1164*. AIAA 16th Thermophysics Conference, Palo Alto, CA, June, 1981.

ANSYS is a registered trademark of Swanson Analysis Systems Inc.

VAX and VAX 9000 are trademarks of Digital Equipment Corp.

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. Hamburguen.

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. Hamburguen.

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 Satisfiability.”

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.

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.