

MARCH 1992

WRL Research Report 92/1



Packaging a 150 W Bipolar ECL Microprocessor

*William R. Hamburgen
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, 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.

Packaging a 150 W Bipolar ECL Microprocessor

William R. Hamburg

John S. Fitch

March, 1992



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

Abstract

Recent developments in computer-aided design have enabled highly automated layout of custom ECL circuits. These layouts have a much higher circuit and power density than gate array designs. It is now possible to place an entire ECL microprocessor, including floating point unit and cache memory, on one large die. To demonstrate the capability of supporting such a die, we built and tested low-cost, air-cooled single-chip packaging for a 12.6 mm x 15.4 mm die. Our PPGA package supplied the required current and maintained junction temperatures at less than 100° C while dissipating 150 W. This required innovation in five areas: die metalization, bondwire layout, PPGA package design, die attach, and cooling by a thermosiphon.

This is a preprint of a paper that will be presented at the
42nd Electronic Components and Technology Conference,
San Diego, California, May, 18-20, 1992.

Copyright © 1992 IEEE

Table of Contents

Nomenclature	vii
1. Introduction	1
2. Die Metalization	4
3. Bondwire Layout	8
3.1. Bondwire Pitch	9
3.2. Power Bondwire Resistance	10
3.3. Power Bondwire Fusing	10
3.4. Signal Bondwires	16
4. PPGA (Plastic Pin Grid Array)	18
5. Die Attach	20
6. Thermosiphon Cooling	22
7. Conclusions	25
8. Acknowledgements	25
References	27

List of Figures

Figure 1:	Package for 150W bipolar ECL microprocessor die.	2
Figure 2:	Bottom view of package assembly.	3
Figure 3:	Test die in cavity of PPGA package.	5
Figure 4:	40 μm thick gold test pattern with 101.6 μm wide lines. End view @100\times (<i>top</i>), 25.4μm drawn space @1000\times (<i>bottom</i>).	6
Figure 5:	Defects rate for gold plating process with various line spacings and plating thicknesses.	7
Figure 6:	Two-row non-staggered gold ball bonding @15\times.	9
Figure 7:	Film coefficient vs wire diameter.	11
Figure 8:	Fusing currents for single wires and arrays of wires. Dashed lines are the short and long wire models.	13
Figure 9:	SEM photo showing bondwire loop clearances @100\times.	17
Figure 10:	Exploded view of package assembly.	20
Figure 11:	Section view of the condenser tube sealed on the boiler by the shape memory ring. The land feature is 0.3 mm high.	23
Figure 12:	Land penetration into the condenser tube versus tube hardness. This land was made of nickel-plated OFHC copper.	23
Figure 13:	Junction-to-air thermal resistance as a function of die power, with copper-finned condenser and 7 liter/s airflow.	24

Nomenclature

$A_{surface}$	wire surface area	R_{sheet}	lumped resistance of the sheet of power bondwires on the long edge of the die
C_1	long wire fusing current coefficient	R_{wire}	resistance of a single bondwire
C_2	short wire fusing current coefficient	t	die metal layer thickness
D	bondwire diameter	ΔT	surface to ambient temperature difference
h	convection or film coefficient	ΔT_{melt}	wire melting point to ambient temperature difference
$I_{fuse,sheet}$	fusing current for the sheet of power bondwires on the long edge of the die	V_{cc}	-5.2V supply voltage
I_{fuse}	fusing current for a single bondwire	V_{ee}	ground voltage
I_{wire}	current in a single power bondwire	V_{tt}	-2V input terminating voltage
k	thermal conductivity	ΔV	voltage drop
L_{wire}	bondwire length	W_{stripe}	busbar or stripe width
L_{die}	die length (long dimension)	W_{die}	die width (narrow dimension)
Nu	Nuselt number (dimensionless convection coefficient)		
P	busbar or stripe pitch		
P_{ideal}	ideal pitch (“perfect bonder” model)		
P_{real}	realistic pitch (“real bonder” model)		
Q_{in}	power or heat generated in a wire		
Q_{out}	power or heat transferred out of a wire		

Greek Symbols

α_{ideal}	ideal P/D ratio
α_{real}	realistic P/D ratio
Ψ	die current demand per unit area
ρ	resistivity

1. Introduction

In traditional multichip processor implementations, the inter-chip signaling delay increases the basic machine cycle time. While multichip modules can partially mitigate this effect, the additional complexity tends to increase costs and time-to-market. As cycle times continue to drop, this off-chip signaling penalty will become still more onerous.

Over the last twenty years, microprocessors have evolved to include more and more system functions on a single die. As device feature sizes shrank to 1.0 μm , and die sizes increased to over 1 cm^2 , it became possible to put an entire processor, including floating point unit, memory management, and cache memory, on a single CMOS (complementary metal-oxide semiconductor) die [14]. By using this *fully integrated* microprocessor design approach, the performance penalty due to off-chip signaling latency is small; off-chip delays affect execution speed only on occasions when the processor is stalled by an on-chip cache miss. Examples of the fully integrated approach include most of today's fastest microprocessors, including Intel's i486 [11] and Digital's recently announced 21064 RISC microprocessor [5]. This trend toward increased integration has sounded the death knell for multichip processor implementations. Shortly even mainframe computers - the last bastion of the older, multichip design style - will be built using fully integrated microprocessors [7].

Today mainframes are almost universally built with emitter coupled logic (ECL). ECL is a very fast logic family with superior load driving capabilities. However, in its usual gate array or standard cell form, ECL has a lower layout density than custom CMOS. This lower density and ECL's high static power dissipation have made it uninteresting for microprocessor applications thus far.

A custom ECL approach offers the potential for extremely high performance microprocessors, but such die would dissipate well over 100 W. Even single-chip packaging presents a considerable challenge, particularly in power distribution and cooling. ECL's small 600 mV signal swings (compared to 3.3V for modern CMOS) call for correspondingly small supply voltage variations. The size and high watt density of these die, preclude conventional power distribution strategies. Similarly, acoustic noise considerations preclude using a conventional solid metal heatsink; either a phase-change or pumped-liquid cooling system would be required.

To demonstrate the capability of supporting such a die, we built and tested low-cost, air-cooled, single-chip packaging for a 12.6 mm x 15.4 mm ECL microprocessor. (Figures 1 and 2). The system uniformly distributed 30 A and maintained junction temperatures of less than 100° C while dissipating 150 W. It also provided 348 signal pins in a field-replaceable package. This required innovations in five areas:

- die metalization
- bondwire layout
- plastic pin grid array (PPGA) package design
- die attach
- cooling with a thermosiphon

With the exception of the die attach and thermosiphon filling and sealing, conventional manufacturing processes were used. Cost of the complete package assembly, exclusive of the die, is less than \$300 in small quantities, and would be less than half that in volume production.

While our work was directed toward meeting the demands of ECL designs, some of these packaging techniques may eventually prove useful even for CMOS. Today's fastest and most power-hungry CMOS microprocessors dissipate only 30 W, [5] but history suggests that higher dissipations are inevitable.




Figure 1: Package for 150W bipolar ECL microprocessor die.




Figure 2: Bottom view of package assembly.

2. Die Metalization

For a large, high-powered die, the design of the packaging must begin with consideration of the local distribution of power to the logic cells. Our CAD methodology uses polysilicon for the shortest connections within cells, with metal 1 (M1) primarily used for the remaining intracellular routing. Metal 2 (M2) and most of metal 3 (M3) form an x,y signal layer pair used for connections between cells. The remainder of the M3 layer and all of metal 4 (M4) form an x,y power layer pair.

Non voltage-compensated logic is the fastest and densest ECL design style, but it demands that power and ground voltages be held within a 15 mV range over the entire chip. To minimize voltage variation, the M4 layer must be thick and consist of an array of stripes on a pitch P , that cross the narrow dimension of the die W_{die} (Figure 3).

Each stripe has width W_{stripe} . The stripes alternate between supply ($V_{ee} = -5.2V$) and return ($V_{cc} = \text{ground}$) to minimize supply inductance and to bring power and ground close to every cell in the die. Thus, each stripe supplies current to a region $2P$ wide. Each end of each stripe has a power bondwire, supplying half the stripe; the current crossing the middle of a stripe is zero. Assuming the die has a uniform current demand per unit area, Ψ , the current per bondwire, I_{wire} , is determined by:

$$I_{wire} = \Psi W_{die} P \quad (1)$$

We obtain the M4 voltage drop ΔV , by integrating dV from the end of the stripe to the center. For M4 with resistivity ρ and thickness t , we have,

$$\Delta V = \frac{1}{4} \frac{P}{W_{stripe}} \frac{\rho}{t} \Psi W_{die}^2 \quad (2)$$

The W_{die}^2 term reflects the strong incentive to run the M4 stripes across the narrow die dimension. Note that W_{stripe}/P is the percent utilization of M4, and ρ/t is the M4 sheet resistance (Ω/square).

Our die will require 30 A, for an average flux $\Psi = 15.5 \text{ A/cm}^2$. We'll assume the M4 aluminum is $2.5 \mu\text{m}$ thick and has resistivity $\rho = 4.2 \times 10^{-8} \Omega\text{-m}$ @ 100°C (sheet resistance = $17 \text{ m}\Omega/\text{square}$). If the utilization $W_{stripe}/P = 0.95$, ΔV for our 12.6 mm wide die will be 110 mV. This greatly exceeds the 15 mV budget.

Another issue is electromigration. Even with this unusually thick aluminum, the average current flux at the end of a stripe ($\Psi W_{die} P / t W_{stripe}$) is $0.82 \text{ mA}/\mu\text{m}$. This average value is already dangerously close to exceeding the 1-2 $\text{mA}/\mu\text{m}$ electromigration limits typical for aluminum thin-films, and current crowding at the bondpads and non-uniform current demand from stripe-to-stripe will push the peak value even higher. To avoid these potentially crippling problems, a better metal system would be needed for M4.




Figure 3: Test die in cavity of PPGA package.

To meet this need, we used a process developed for tape automated bonding (TAB) bumps to plate thick gold busbars that entirely covered all M4 aluminum. To test an existing gold plating process which had originally been developed to plate 25 μm thick bumps, we made a mask with a gold pattern intended to be identical to our final die.

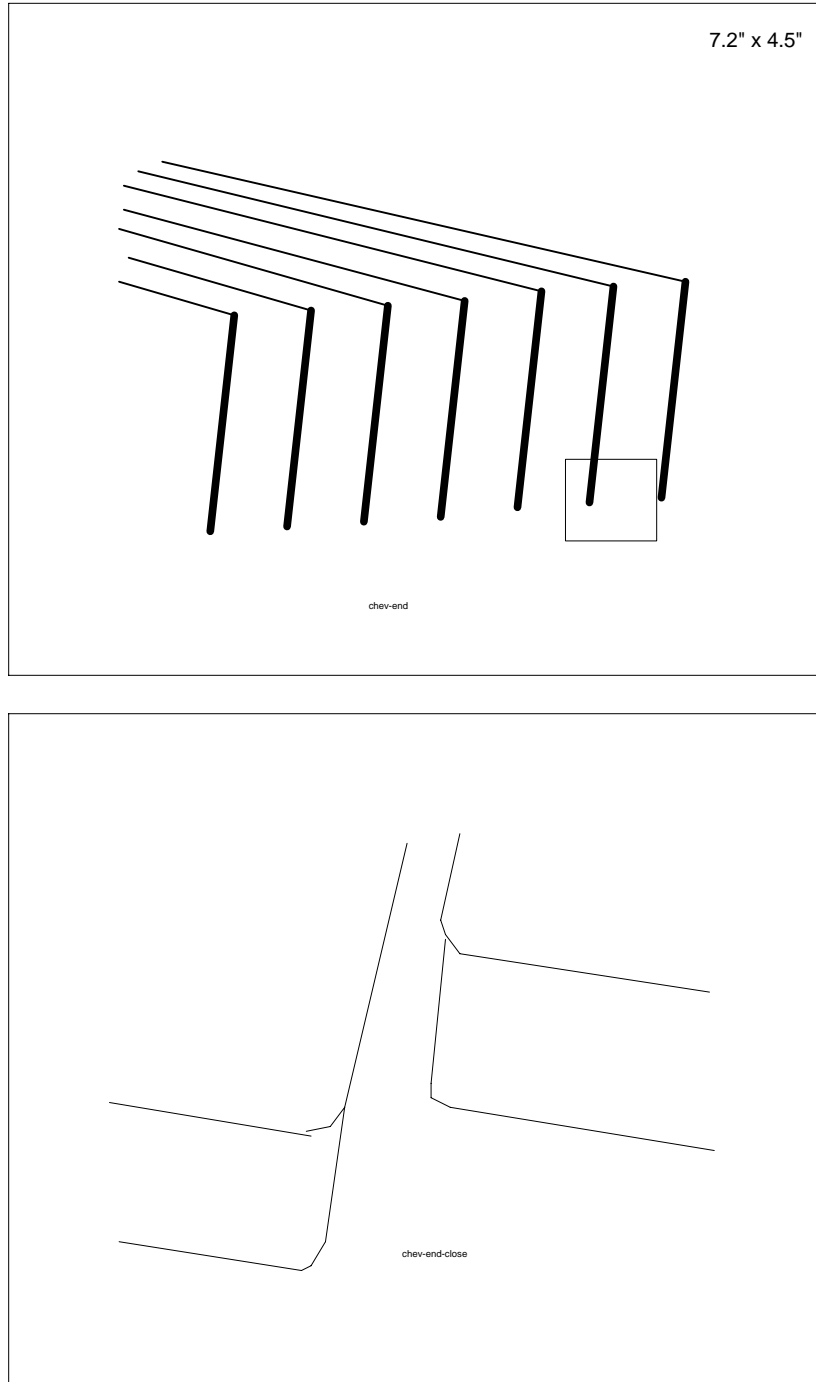


Figure 4: 40 μm thick gold test pattern with 101.6 μm wide lines.
End view @100 \times (*top*), 25.4 μm drawn space @1000 \times (*bottom*).

This mock die had 81 busbars, each 127 μm wide, on a 177.8 μm pitch. Included were four chevron-shaped resolution test patterns (Figure 4), each with a different linewidth (101.6, 127, 152.4 and 177.8 μm mask dimensions). Within each test pattern, a range of spaces were tested (25.4, 38.1, 50.8, 63.5, and 76.2 μm mask dimensions). Four different gold thicknesses were tested (10, 20, 30, and 40 μm).

While the photos show apparently good lithography with 40 μm gold, Figure 5 shows that at this thickness, shorts were abundant at all but the largest spaces. This was primarily due to the incomplete resist removal and barrier metal etchback*. With 10 μm and 20 μm thick gold, no shorts were found, even at the smallest line spacings. The 960 mm of line spaces on each mock die site bore out these test pattern results. The 50.8 μm spaces between the busbars had substantial defects only with the 40 μm thick plating, and the 10 and 20 μm thick samples were defect free. On this basis, we selected a 25 μm gold thickness for subsequent work. The thick gold busbars on top of the M4, and gold's low resistivity (at 100° C, typically $3.5 \times 10^{-8} \Omega\text{-m}$ for plated gold vs $4.2 \times 10^{-8} \Omega\text{-m}$ for aluminum thin-film) reduced edge-to-center voltage drop from 110 mV to 11 mV, comfortably within the 15 mV specification.

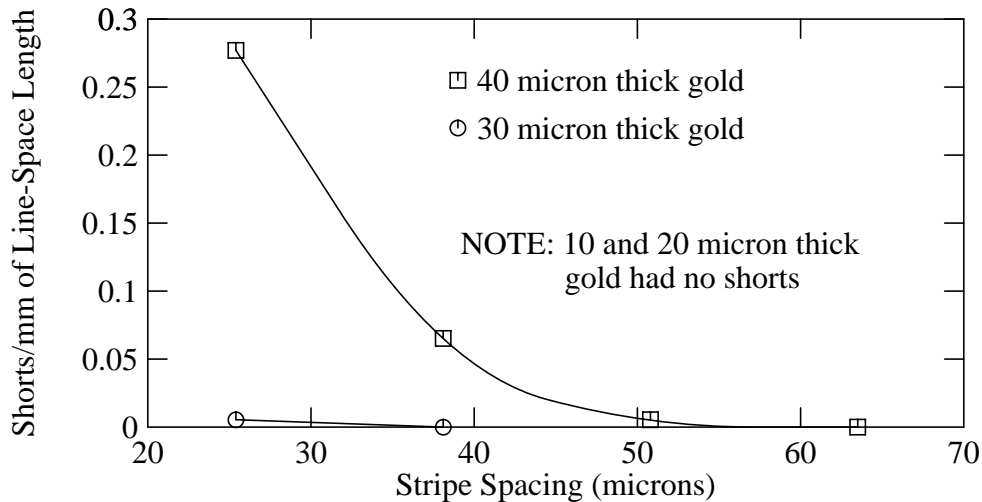


Figure 5: Defects rate for gold plating process with various line spacings and plating thicknesses.

*We believe minor process improvements would reduce or eliminate these defects

3. Bondwire Layout

Engineering the bondwires was one of the principal challenges in packaging our die. Except in transfer-molded packages, most VLSI assembly is done with 33 μm diameter aluminum wires with wedge bonding. Aluminum wires are easy to bond at room temperature, and offer good compatibility with the usual aluminum bondpad metallurgy. However, aluminum oxidizes and has a higher resistivity and a lower melting point than gold (660° C vs 1064 ° C). Because of this, usual practice limits the current per wire to 200 mA. Just providing the 30 A to our die would require 300 of these aluminum power wires! We also have an additional 348 signals and another 100 or so I/O reference wires. Considering that 33 μm wires are seldom bonded on pitches less than 100 μm , it was clear that on a die this size, conventional wire bonding practice would not suffice. A higher wiring density was needed. We briefly considered TAB, but the complexity, lead time, and cost made it unattractive for low volume research prototypes.

To increase bondwire density, other designers have tried staggering the pads around the edge of the die [23]. While this can accommodate wider bonding tools and bonds, it increases wire-to-wire shorting risk. This is particularly true with multi-tier packages and their poor tier-to-tier registration. With wedge bonding, the shallow wire departure angle at the first bond means that to alleviate this potential for shorting, the spacing between rows on the die must be increased, making wires longer and consuming die real estate. Consequently, wedge bonding development has focused on single row bonding. Pitches of 40 μm have been demonstrated with 10 μm diameter gold wires [19]. This specialized technique would not offer adequate current capacity for our die. Gold ball bonds, however, have a vertical departure angle, so they do not incur a row spacing penalty. In fact, looping heights can be controlled sufficiently that even staggering the pads is unnecessary. This two-row, non-staggered, gold ball bonding is the approach we selected and developed (Figure 6).

Our design analysis begins by examining bondwire selection to minimize voltage drop, and follows with an analysis of fusing current, including experimental data for arrays of closely-spaced wires. We'll then look at how the signal wires are accommodated, and show how we used a test die to verify the current capacity of our bondwire array.




Figure 6: Two-row non-staggered gold ball bonding @15 \times .

3.1. Bondwire Pitch

Design rules usually specify some minimum centerline separation between any two bondwires; an *ideal* wire-bonder would be able to bond wires onto a substrate at this pitch. If, as is frequently the case, the design rule specifies the minimum wire spacing in terms of wire diameter, we can define this ideal pitch as:

$$P_{ideal} = \alpha_{ideal} D \quad (3)$$

Any *real* bonder, however, has limitations that may require choosing a larger than ideal pitch:

- **Positioning errors** - Real bonders have finite positioning errors.
- **Non-ideal bond or bond tool size** - The ball bond may be larger than ideal or lopsided, requiring a larger pitch to avoid shorts. Also, the capillary tip diameter may limit how close an adjacent bond can be placed without capillary contact with the neighboring wire.
- **Pad size** - Pad sizes and potential for damage to overlapping passivation are important when the pad metal is lower than the surrounding materials. This is not the case bonding to our thick gold.

Using K&S model 1484 bonders to ball bond 33 μm gold wires, Shu [22] reported 30 ppm ball shorting rates at a 109 μm pitch, and sub part-per-billion rates at a 127 μm pitch. Thus, this bonding process has virtually no shorting defects at $P/D=3.85$. We can idealize this real process by equation 4, and assume that it's valid locally over a range of wire diameters, with $\alpha_{real} = 3.85$.

$$P_{real} = \alpha_{real} D \quad (4)$$

3.2. Power Bondwire Resistance

Like die metalization, bondwires contribute voltage drop. However, if the die demands equal current from all bondwires, they will all have the same voltage drop. This can be offset by increasing the supply voltage. This differs from die metalization, where voltage differences occur even with a uniform current demand. The degree to which bondwires contribute to voltage nonuniformity across the die depends on the variation of flux demand across the die.

Wire resistance is inversely proportional to the diameter squared,

$$R_{wire} = 4\rho L_{wire}/\pi D^2 \quad (5)$$

where ρ = wire resistivity. On a die edge with length L_{die} , we can fit L_{die}/P_{real} wires. By Equation 4, for the entire edge, the resistance of the sheet of power bondwires, R_{sheet} is inversely proportional to the wire diameter.

$$R_{sheet} = 4\rho L_{wire} \alpha_{real} / \pi L_{die} D \quad (6)$$

Large wires keep the resistance low, and the corresponding large pitch makes wafer probing easier. However, there are a number of reasons why using the largest possible wire diameter may not be desirable:

- **Economics** - There are economic benefits to using conventionally configured VLSI bonders, which have limited wire diameter capacity. Ideally, the same bonder could be used both for power and signal wires.
- **Die area** - As wire diameter grows, the bondpad region consumes an increasing fraction of the die area. For example, using the real bonder model (Equation 4), and assuming the bondpad region width equals the pad pitch, 100 μm diameter wires imply a 385 μm wide bondpad region on each of the long edges of the die. On our 12.6 mm wide die, these oversized power bondpads alone would consume 6% of the die area!
- **M3 voltage drop** - The bondwire and hence the M4 stripe pitch, directly affects voltage drops on the M3 power straps. The more current each M3 strap must carry, the wider it must become to avoid excessive current density or voltage drop. Since these M3 power straps consume precious signal routing channels, it may be desirable to reduce the bondwire diameter and pitch to obtain adequate signal routability on M3.
- **Inductance** - For a given L_{die} , L_{wire} , and α , 2-dimensional modeling shows that the total inductance of the bondwire sheet is proportional to the bondwire diameter [20]. This is not of great importance with ECL's constant currents, but a low supply inductance is critical for squeezing the highest speeds from CMOS logic.

3.3. Power Bondwire Fusing

Wire resistance causes self heating, and at the fusing current, the wire melts and disconnects. Bondwire heat is removed by a combination of conduction to the ends, radiation and convection to the surrounding fluid and by conduction to the molding compound in an over-molded package. In logic families with small static and large dynamic current demands, such as CMOS, understanding the transient thermal response of bondwires is necessary. Recent research has emphasized such transient thermal analysis. Coxon et al. [4] gave closed-form solutions for

transient and steady state heat transfer under a variety of assumptions and boundary conditions. These solutions were compared with experimental data by A. N. Peddy^{**}. The work highlighted the importance of accounting for property variations - in particular, the temperature coefficient of electrical resistivity for the wire. Hill et al. used this analysis by Coxon et al., to calculate the transient temperature rise of gold wires, and made infra-red measurements to confirm the models [8]. They obtained good correlation and noted that temperature rise was highly sensitive to the actual diameter of the wire used. However, ECL is a constant current logic family, and our fusing analysis need only consider the constant current case.

For long wires, convection dominates, and the fusing current is independent of wire length. Conduction out the ends cannot help cool the middle of a long wire. Due to gold's extraordinarily low emissivity, radiation contributes little, < 1% for 25 μm diameter wires at the 1064° C melting point [4].

Free convection over these thin wires occurs at very low Rayleigh numbers. For this regime, Holman suggests a constant Nusselt number independent of wire diameter [10].

$$Nu = hD/k_{air} = 0.4 \quad (7)$$

Therefore, the convection or film coefficient h varies inversely with the wire diameter,

$$h = 0.4k_{air}/D \quad (8)$$

Air's thermal conductivity midway between room temperature and gold's 1064° C melting point is $k_{air} = 0.059 \text{ W/m-K}$. Figure 7 shows that this idealized relationship correlates well with values suggested by Coxon et al. [4].




Figure 7: Film coefficient vs wire diameter.

^{**}These experiments were described in a technical report which we were unable to obtain: "Fusing bond wire study," Hughes IDC 5687.0733B/014, Aug 7, 1978. Awkward et al. [1] cited this same report as an "Inter-departmental Correspondence, Hughes Aircraft Corp."

By combining the wire surface area $A_{surface} = \pi L_{wire} D$, with Equation 8, we may write an expression for the power, or heat, transferred out of a long wire, which is independent of the wire diameter!

$$\begin{aligned} Q_{out} &= h A_{surface} \Delta T \\ &= 0.4 \pi k_{air} L_{wire} \Delta T \end{aligned} \quad (9)$$

For a current I_{wire} , the power generated in the wire is,

$$Q_{in} = I_{wire}^2 R_{wire} = 4 I_{wire}^2 \rho L_{wire} / \pi D^2 \quad (10)$$

by Equation 5. Since $Q_{in} = Q_{out}$, we can combine Equations 9 and 10, and evaluate at the melting temperature to determine the fusing current of a long wire in air:

$$I_{fuse} = \pi D \sqrt{0.1 k_{air} \Delta T_{melt} / \rho} = C_1 D \quad (11)$$

At gold's 1064° C melting point, $\rho = 13.7 \times 10^{-8} \Omega\text{-m}$, therefore $C_1 = 21,000 \text{ A/m}$. By equation 4, the number of bondwires that fit on the die edge varies inversely with wire diameter, so for long wires, the maximum current that can be delivered to the edge of the die is *independent* of the wire diameter.

$$I_{fuse, sheet} = L_{die} C_1 / \alpha_{real} \quad (12)$$

Thus, for $\alpha_{real} = 3.85$, sheets of long gold wires will fuse at 5.45 A per mm of die edge length, regardless of the diameter chosen.

As a wire becomes shorter, conduction to the ends begins to reduce the wire midpoint temperature. With very short wires, the contribution of convection becomes negligible. By ignoring convection cooling entirely, we can approximately model wire behavior by one-dimensional conduction with uniform internal heat generation, in a material with constant properties. Substituting our wire geometry into a solution from Holman [10], gives:

$$\begin{aligned} I_{fuse} &= \pi \sqrt{k_{wire} \Delta T_{melt} / 2 \rho} D^2 / L_{wire} \\ &= C_2 D^2 / L_{wire} \end{aligned} \quad (13)$$

At the fusing current, the temperature along the wire varies from ambient temperature at the wire ends, to the melting point at the center. Temperature dependent properties vary accordingly, and numerical methods are necessary to accurately account for such property variation [1]. From the one-dimensional wire temperature profile, we know the average wire temperature rise $\Delta T_{ave} = \frac{2}{3} \Delta T_{melt}$. With the wire ends assumed to be at 100° C, (the maximum specified die temperature), $\Delta T_{ave} = 743^\circ \text{ C}$. By interpolation of handbook data for gold [21], we find average properties $k_{wire} = 270 \text{ W/m-K}$ and $\rho = 9.05 \times 10^{-8} \Omega\text{-m}$. Thus, for short gold wires with 100° C ends, $C_2 = 3.77 \times 10^6 \text{ A/m}$.

By equating Equations 11 and 13 we find that the crossover point between the long-wire and short-wire models occurs at $L/D = C_2/C_1$. For gold wires, this is $L/D = 180$. Since gold wire bonding rules usually limit L/D to less than this, we'd expect the short wire models to apply to most practical bonding problems.

Again, by equation 4, the number of bondwires that fit on the die edge varies inversely with wire diameter. For short wires, by Equation 13, the fusing current varies as D^2 . Thus, the max-

imum current that can be supplied to the edge of a die by short bondwires is *proportional* to the wire diameter.

$$I_{fuse, sheet} = \frac{\pi L_{die}}{\alpha_{real} L_{wire}} \sqrt{k_{wire} \Delta T_{melt} / 2 \rho D} \quad (14)$$

This dependence is similar to the scaling law for wire resistance (Equation 6). For sheets of short bondwires, both fusing current and resistance improve in direct proportion to wire diameter. Which constraint will be binding depends on the supply current and voltage uniformity specifications for a particular design.

Our fusing analysis is thus far based on models for isolated wires. When an entire sheet of wires is powered, it is unclear to what extent adjacent wires heat each other, reducing fusing current. To understand if this was a significant issue, we performed a series of wire fusing experiments (Figure 8). Arrays of gold wires were bonded into PPGAs, and tested both singly and with adjacent wires powered. Shorter wires were bonded orthogonally in eight-wire groups from the first package tier to the package slug. These were connected in parallel and powered by a single supply. Longer wires were bonded in three-wire groups and connected in series to assure that each wire carried the same current. In both cases, the end wires do not run as hot as those toward the center, so the fusing currents measured would be somewhat lower than in a continuous sheet of power wires. Furthermore, the limited cooling allowed wire end temperatures to rise, significantly in the case high test currents. An extreme case was the 8-wire group tests with the shortest 50.8 μm wires; at currents over 20 A/group, the package resin would begin to smoke.

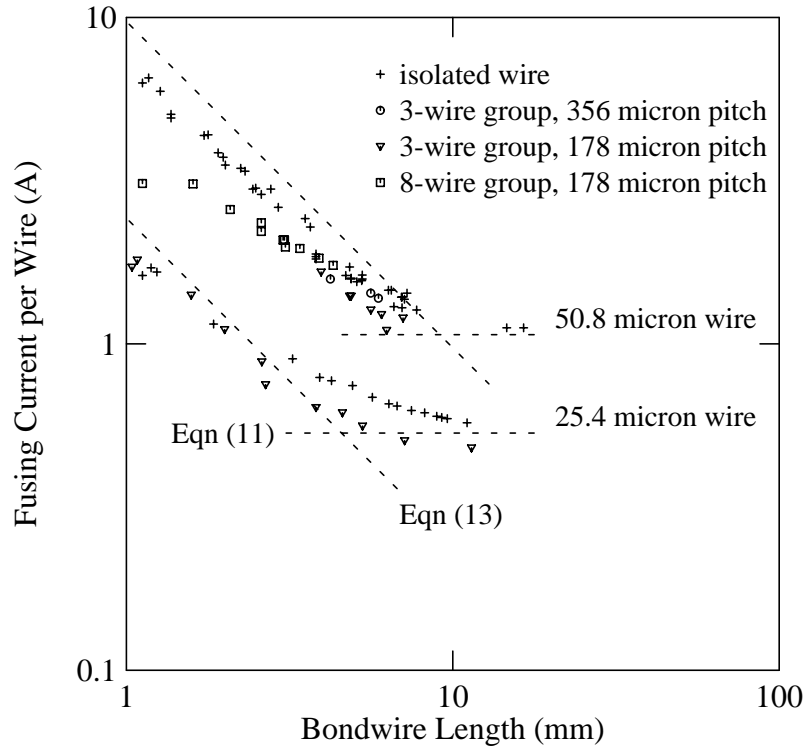


Figure 8: Fusing currents for single wires and arrays of wires. Dashed lines are the short and long wire models.

As the dashed lines in Figure 8 show, test data from long isolated wires was in good agreement with Equation 11. The short isolated wires fused at about 75% of the current predicted by the Equation 13. This suggests that the wire ends were reaching temperatures higher than the 100° C that we had assumed.

Our data shows that heating by adjacent wires is of little importance for the shortest wires, but reduces fusing current by up to 30% in the transition to the long wire regime.

The remaining question is the selection of a fusing current safety factor. For gold bondwires enveloped by organic encapsulants, the encapsulant's thermal expansion and decomposition characteristics affect the value chosen [8]. But for our free-standing gold wires, the choice is less obvious.

During our testing, we made several observations relevant to the choice of a safety factor. A slight reddish glow could be seen on gold wires just below the fusing current. A wire could be operated for hours at this current level without any apparent change, but would blow instantly when the current was increased. This suggests that the sort of thermally activated failure mechanisms that are important in aluminum, are less so in gold. Gold has no native oxides, so as long as physical loads are modest, and there is no interdiffusion of other species, gold wires might perform quite adequately at just below the melting point.^{***} It has been shown that at high temperatures, gold wires anneal or recrystallize, and grains grow to form a characteristic "bamboo" structure, with grain sizes larger than the wire diameter [9]. While this leads to low strength and high elongation limits, and causes early failure in molded packages, it is unclear whether such microstructural changes preclude safe, high temperature operation for a wire constrained only at its ends.

Under 40× magnification, bondwires jumped visibly due to thermal expansion when suddenly powered to 1/3 of their fusing current. This was observed for a wide range of wire lengths. Staying below this threshold assures that wires will not move as current is applied, possibly causing shorts. While this 3× safety factor is not a particularly satisfying criteria, in the absence of other data, we chose it for specifying the maximum design current. Clearly, further work is needed to understand and quantify the risks of operating free-standing gold wires at high temperatures and currents.

^{***}As an extreme example, fusing wires would occasionally repair themselves! As we approached the fusing current, stepping up the current in small increments, a wire would repeatedly collapse, shorten, and resolidify, forming a callus at the break site.

As a final check of our fusing current models, we took a “mock” die created during our die metalization evaluations, and wirebonded it into our PPGA. The measured diameter of the power wires was 47 μm . The V_{ee} wires - the longer power wires - measured 3.0 mm long. Figure 8 shows that at this wire length and our 178 μm pitch, adjacent wire heating effects are negligible. Using Equation 13, we predicted that the eighty V_{ee} wires would fuse at 224 A.

The predicted total current capacity was verified by shorting the on-chip busbars together with a silver-filled epoxy and supplying current to the package pins through a socket. External cooling kept the die under 150° C. The V_{ee} bondwires failed at 164 A, noticeably scorching the package in the neighborhood of the bondfingers. This current was 73% of the value predicted by Equation 13, and in good agreement with our earlier fusing experiments. Since a 150 W die requires only 29 A @ 5.2 V, it would appear that our margins are ample and wires have reserve capacity for non-uniform power demands.

3.4. Signal Bondwires

Our bondwire analysis has thus far assumed that the long sides of the die were used entirely for supplying power to the die core, relegating all the signal leads to the two short ends of the die. This is not enough for I/O-hungry microprocessors. A second row of pads is needed around the entire edge of the die. While putting the power pads in the outer row might appear attractive, owing to the shorter resultant bondwires, doing so would demand necking the busbars down to fit between the inner row signal pads. The large spaces required for gold busbar plating would lead to a layout that is little denser than a single row. To avoid this, we put the busbar bondpads and their 162 power wires on the inner row. The 348 signals and 108 I/O reference wires occupied the entire outer row plus the inner row on the narrow ends of the die (Figure 3). The inner row had a $177.8\ \mu\text{m}$ pitch to accommodate large power wires; by Equation 4, we'd expect to be able to bond at least $46\ \mu\text{m}$ diameter wires on this pitch. The outer row had a $152.4\ \mu\text{m}$ pitch, generous considering we planned to use $33\ \mu\text{m}$ wire. By Equation 4, we'd expect to be able to bond at least $40\ \mu\text{m}$ diameter wires on this pitch.

The wires were bonded in two completely separate arrays. Outer row wires went only to the first package tier and inner row wires went only to the second tier. Bondwires did not cross from outer row on die to second tier on package, and vice versa. This simplified inspection and repair after the first tier had been bonded. Provided there is sufficient control of wire looping, our methodology also allowed the pitches of the two rows to be completely independent. Figure 9 shows how the clearance between the rows of bondwires allowed for package tier misregistration and for arbitrary bondpad pitch on the die. While the two pitches are independent, choosing an appropriate ratio between them simplifies routing die signals to the outer row of pads.

Between the outer pad row and the scribe street, we provide terminating resistors for the 203 input signals. This is the lowest power-density portion of die.

Wafer probing strategy depends on bondpad layout. It would be very difficult to place a probe tip on every pad in both rows during wafer testing. Instead we probe only a single row around each diesite: the power connections on the long sides, and half the signals on the ends. With these probes, the die could be fully powered, data transferred by use of boundary scan, and the die exercised with clock and control signals.




Figure 9: SEM photo showing bondwire loop clearances @100x.

4. PPGA (Plastic Pin Grid Array)

We chose a plastic package rather than conventional cofired alumina ceramic primarily for the high performance of the copper metallization. Plastic pin grid array (PPGA) packaging is gradually gaining acceptance and offers many advantages over conventional ceramic pin grid array (CPGA) construction [18]. Our 504-pin PPGA is built from copper-clad bismaleimide-triazine (BT) laminates, plated with soft gold to facilitate wirebonding, and pinned in a 25 x 25 array on a 100 mil grid to fit a standard ZIF (zero insertion force) socket. The BT's 175° C glass transition temperature makes it possible to wirebond at the high temperatures which are important for good gold-to-gold bonds. This all-gold interconnect system - die pad, bondwire, bondfinger and pin - is advantageous from a galvanic corrosion viewpoint.

While alumina CPGAs are strong and provide a hermetic environmental seal, PPGAs offer several advantages over ceramic:

- **Smaller bond pad pitch** - PPGAs offer fine patterning; 50 μm minimum lines and spaces are permitted in the package bondpad (bondfinger) area, so a PPGA can accommodate as many wires in two tiers as a CPGA can in four. This means shorter wires and a more robust bonding process. While CPGAs have recently become available with a thin film signal layer on the outermost tier, this doesn't help us; the fine pitch is needed on the *first* tier, which handles the signal bondwires from the outer rows on the die.
- **Lower sheet resistance package planes** - The sheet resistance of each copper power plane in our PPGA is about 1 $\text{m}\Omega$ per square. Ceramic planes are typically an order of magnitude higher. While CPGAs offer smaller via obstructions and a low incremental cost for adding supplemental planes, it is difficult to overcome the limitations of the high-resistivity refractory metal conductors.
- **Wrap-around die cavity plating** - Each layer pair in a PPGA can have plated-through holes (PTH) formed prior to lamination. This PTH metal allows wrapping a connection from one side of a power plane layer pair, around the die cavity wall, to the other power plane. This creates a low inductance and resistance connection from the bottom plane to the bondfingers, and allows one PPGA tier to replace two on a CPGA. It also offers a convenient way to distribute the input terminating voltage ($V_{\text{tt}} = -2.0 \text{ V}$) to the first tier.
- **Compatible with copper slug** - Our thermosiphon technology requires that the metal die attach "slug" at the bottom of the package cavity also function as the thermosiphon boiler. Thermal conductivity, thermosiphon sealing and material compatibility considerations require the boiler to be made of copper. Attaching a copper boiler directly to a CPGA is difficult, and using an intermediate package slug would adversely affect cooling.
- **Improved signal integrity** - The PPGA's lower dielectric constant and line resistance reduces parasitics. In fact, quite unlike CPGAs, the challenge with PPGAs is getting a *low* enough line impedance! An etch-after-plating process produces stub-free signal lines. Plate-through-the-pins and electroless plating processes are becoming available to offer stub-free lines for CPGAs.
- **Lower tooling and part costs** - For comparable complexity, PPGAs can be up to an order of magnitude less expensive than CPGAs to tool, and finished part costs are substantially lower as well.

For maximum fusing current, the PPGA cavity must be laid out to keep the power bondwires short. Minimizing bondwire length requires careful attention to process limitations and capabilities: clearances, tolerances, and so on. Unfortunately, there are no big terms to engineer into insignificance, each of many small ones must be painstakingly minimized.

Even with fine patterning of the PPGA metal, laying out the bondfingers for orthogonal bonding (having the same pad pitch on package and die) was difficult. This was particularly true for the 152.4 μm pitch pad on the outer row of the die. Radial bonding (having a larger pitch on the package than on the die) would not have helped much due to the shortness of our bondwires and the size of the die [12]. Instead, we staggered the bondfingers on the package. While this made the first package tier wider, we felt that the benefit of adequate bondfinger width outweighed the liability of lengthening the bondwires to the second tier. Radial bonding on the second tier was rejected since varying power wire lengths would have caused varying fusing currents and increased voltage variations on the die.

Pads for adhesive bonding 504-size decoupling capacitors were provided on the second bond tier immediately adjacent the bondwires, for the lowest possible series inductance (Figure 3). We also provided mounting holes at the edge of the package for subminiature microwave connectors for diagnostic signals.

The 156 power pins are dispersed evenly through the pin field, both to spread their heat into the motherboard, and to improve signal integrity. There are 83 V_{cc} (ground) pins, 61 V_{ee} pins and 12 V_{tt} (terminating voltage) pins. The V_{ee} pins are most heavily loaded, but with a 30 A supply, each pin averages only 0.49 A. At these high current levels, however, the motherboard must feed power symmetrically to the pin field to ensure uniform voltages across the die.

5. Die Attach

An important aspect of our thermal design philosophy was to eliminate all the thermal interfaces between the die and the thermosiphon, except the die attach. We achieved this by substituting our copper thermosiphon boiler for the usual copper PPGA slug (Figure 10). This approach reduced the number of thermal interfaces, but it required that the thermosiphon condenser be assembled after the die was attached and wire bonded into the PPGA. Had a conventional package slug been used, we could have attached a completed thermosiphon or heatpipe to it, at the cost of higher die temperatures and developing an additional assembly process.




Figure 10: Exploded view of package assembly.

To obtain good spreading of heat under the die, we used copper for the boiler. Since silicon has a much smaller thermal expansion coefficient than copper, we attached the die to the boiler with a flexible epoxy, Ablestik™ ECF563. This unsupported silver-filled epoxy preform of-

ferred high thermal conductivity, a low elastic modulus and adequate adhesion. Using preforms eliminated the need to develop paste dispensing programs. The bondline could be easily controlled by varying the pressure applied during cure. Final bondlines were typically 10 to 20% less than the original 75 μm preform thickness. Acoustic and infrared studies of the epoxy joint indicated void- and delamination-free joints after repeated liquid thermal shocks from -65°C to 150°C [6]. At an estimated 1.9 W/m-K, the epoxy thermal conductivity was high enough to keep the die attach thermal resistance below 0.15°C/W .

Since the epoxy was electrically conductive, the entire thermosiphon was connected to the same -5.2 volt DC potential as the backside of the ECL die. The thermosiphon would have to be covered by a plastic shroud to protect service personnel. Dielectric materials, such as ceramics, were considered for the boiler or for isolation layers on the boiler / epoxy interface. This would allow the condenser to be grounded, but at the cost of higher die operating temperatures.

The die and PPGA were simultaneously bonded to the boiler, using a single epoxy preform and bonding fixture (see Figure 10). The following die attach process was used:

- Electroless nickel plated boiler placed in fixture with flat side facing upward.
- Epoxy preform placed on boiler.
- PPGA placed onto fixture locating features and lowered onto epoxy.
- Die placed in cavity on exposed epoxy, located by tabs that are part of the PPGA.
- Pressure distribution pad placed on die.
- Spring mechanism pressed on pad, squeezing die into epoxy. A second spring pressed PPGA into epoxy.
- Fixtured assembly heated to 130°C for two hours.

To distribute the pressure evenly across the large die, we used an elastomer pad with a KaptonTM polyimide film backing. This pad was bonded to a steel backing plate that was guided by two of the PPGA pins. This accurately located the pad to the die. A spring mechanism pushed on this backing plate through a steel ball bonded to the exact center of the plate. The ball assured that during cure, the load remained accurately centered and perpendicular to the die. By changing springs, the pressure on either the die or the PPGA-to-boiler joint could be independently adjusted about the 400 kPa nominal value. Under these high cure pressures, the locating tabs in the PPGA cavity proved important in preventing the die from squirming out of position.

The Kapton polyimide layer ensured that the gold bond pads on the die would not be contaminated during die attach. Several metal foils were also evaluated for this purpose; all marred the die surface to some extent. Based on surface examination with laser microprobe mass spectrometry (LIMS), tantalum and Kapton did not transfer contaminants to the gold during a simulated curing cycle. The Ni foil was marginal, and Ti, Co, V, and W all substantially contaminated the gold. Kapton was selected for its low cost and ease of handling.

A temperature controller and cartridge heaters in the bonding fixture brought the assembly to the desired cure temperature. To reduce oxidation and void formation, we experimented with maintaining a vacuum on the fixture as the epoxy heated and began to flow. This proved unnecessary since atmospheric cure conditions gave reliable and void-free bonds.

6. Thermosiphon Cooling

A thermosiphon is a phase-change heatsink consisting of a sealed vessel containing a volatile liquid. When heat is supplied to the bottom end - the boiler or evaporator - the liquid vaporizes and moves to the cooled portion - the condenser. There it condenses, giving up its energy to the air passing over the external fins. The condensate then returns to the heated end where it repeats the cycle (Figure 10). A thermosiphon is essentially a wickless heatpipe [13]. It relies on gravity, rather than the capillary forces of a wick, to return condensate to the boiler. The wick can limit the performance of a heatpipe [2]. We chose a thermosiphon instead of a heatpipe because of its simplicity and the large heat fluxes it could handle. We used a mercury tip switch to interrupt power to prevent overheating should the thermosiphon inadvertently be tipped upside down. A thermosiphon is more expensive than a solid metal heatsink, but in VLSI packaging applications it becomes attractive when die power exceeds 50 to 100 W. At these power levels, a solid heatsink would have to be very large and would require significant air flow. Because it has no apparent moving parts, the casual observer views the thermosiphon as nothing more than an efficient air-cooled heatsink.

We built our thermosiphon by combining the package boiler with spirally finned copper tubing as the condenser. The nickel-plated copper boiler that holds the die had fins cut into one side for the boiling surface of the thermosiphon. The fin geometry was chosen after boiling studies of water at subatmospheric pressures [15]. We built condensers of 2.85 cm diameter copper tubing with two types of fins. Copper fins spirally wrapped and brazed onto the tube provided the best thermal performance, while tension-wrapped aluminum fins provided the lightest condenser. All fins had an outside diameter of 5.4 cm and a pitch of 4 fins per cm. The condenser was 12 cm long to accommodate two adjacent 60 mm diameter fans and to allow the assembly to fit inside a standard desk-top box. The two fans, in parallel, quietly supplied 7 liters/s (≈ 17 c.f.m.) of air at a static pressure drop of 18 Pa (0.07 inches of H_2O).

A safety cap plugged the condenser end of the thermosiphon. This nickel-plated copper cap had a hole drilled through it which was plugged with a low temperature solder. Should the thermosiphon be exposed to a fire or some other high temperature fault condition, the solder would melt and release the internal pressure before vessel rupture.

The choice of operating fluid significantly affects thermosiphon performance. Water is a desirable working fluid for a thermosiphon since it has such a high heat of vaporization, high thermal conductivity, and is nontoxic and nonflammable. However, from our earlier studies, we found that at subatmospheric pressures and at our heat fluxes, the boiling behavior of water became sporadic. [16] By adding a small amount of alcohol to the water, the surface tension and thermodynamic conditions changed significantly enough to provide steady boiling. In fact, at certain mixture ratios, the critical heat flux could be improved above that of pure water by 100%! [17] A mixture of 22% propanol in water, by volume, was chosen for our thermosiphon because of performance characteristics, its non-toxicity and because at this concentration, the U.S. Department of Transportation does not classify it as a flammable or combustible liquid.

Cleanliness of assembled parts was essential because non-condensable gases and other contaminants degrade the performance of a thermosiphon. Scale on the hot regions of the thermosiphon could hinder boiling performance, ionics could cause corrosion, and noncondensable gases could gather in the cooled end, blanketing condenser surface. In order to avoid trapping

solder fluxes or additional materials in the thermosiphon, an intrinsically clean sealing technique was used. The copper tube of the condenser was hermetically sealed onto the boiler and cap by a shape memory ring supplied by Raychem Corp. When electrically heated to 180° C, this nickel-titanium alloy ring shrinks by roughly 5% in diameter. This squeezed the tube onto a small positive feature, or sealing land, on the outside diameter of the boiler and cap [3] (Figure 11). A




Figure 11: Section view of the condenser tube sealed on the boiler by the shape memory ring. The land feature is 0.3 mm high.

soft tube material and hard sealing land were essential to achieve a plastically deformed seal with good land penetration into the tube (Figure 12). A helium leak detector verified that with annealed copper tubing and a precipitation-hardened chrome-copper boiler and cap, reliable seals were obtained, even after twenty -65° C to 150° C thermal shock cycles.




Figure 12: Land penetration into the condenser tube versus tube hardness. This land was made of nickel-plated OFHC copper.

The following assembly procedure was used for the thermosiphon:

- Boiler, with PPGA and wire-bonded die, attached to condenser using shape memory ring.
- Liquid mixture added to condenser.
- Chip powered to boil and degas liquid. Boiling continues to purge gases from condenser.
- At desired fill volume (~20%) and desired alcohol-water mixture (22% alcohol by volume), safety cap sealed in place with a second shape memory ring.

Sealing the condenser while the liquid was boiling ensured a saturated state inside the thermosiphon. This minimized the internal operating pressure, and hence operating temperature.

To check cooling system performance, we used a test die with TaN thin-film resistor patterns between and underneath the gold busbars. The resistor layer was patterned so that when the test die was wirebonded into a package, it duplicated the current draw and power dissipation of a live part. We coated a thin layer of filled silicone to the surface of the die, to increase its emissivity for infrared temperature measurements. Figure 13 shows the performance of the entire cooling system at various die powers and constant air flow. The nonlinear behavior was due to the changing internal pressure, which caused boiling variations as the heat flux changed. The flat portion of the curve corresponds to steady nucleate boiling. In this power range, the thermosiphon's thermal resistance, from the boiler to the cooling air, was below 0.32°C/W . The flat portion of the performance curve in Figure 13 may be shifted significantly by changing the surface area of the boiler.




Figure 13: Junction-to-air thermal resistance as a function of die power, with copper-finned condenser and 7 liter/s airflow.

7. Conclusions

While much work remains to formally qualify this packaging approach, it appears to offer a low-cost solution for powering, cooling, and connecting high wattage microprocessor die. The necessary technologies include:

- Thick gold busbars across the surface of the die to uniformly distribute large supply currents.
- Two-row, non-staggered ball bonding around the edge of the die, for good signal and power distribution.
- PPGA package with low sheet resistance copper planes for distributing power, and fine patterning capabilities for dense bondfinger pitch.
- Low modulus silver-filled epoxy die attach directly to the copper heatsink.
- Thermosiphon with alcohol-water fill for quiet air-cooling.

Many variations are possible to meet differing requirements, but thick die metal, double-row wire bonding, a minimum number of thermal interfaces, and phase-change cooling are all essential ingredients.

Future work may include investigating the behavior of gold wires at sustained high current, obtaining die coatings with good adhesion and environmental protection, performing long-term thermosiphon tests, and developing enhancements to handle larger die and higher power levels.

8. Acknowledgements

We wish to acknowledge the vital contribution of Wade R. McGillis, who conducted the fundamental boiling studies that guided our thermosiphon fluid selection and boiler fin design. Norm Jouppi was our key resource in defining the project and guiding packaging design decisions. The authors also wish to thank Betsy Fitch and Mary Jo Doherty for reviewing the manuscript and providing their many critical comments and corrections.

References

- [1] K. A. Awkward and R. D. Lewis.
Model for Determining Thermal Profiles of Bond Wires Using "PSPICE" Analysis.
In *7th SEMI-THERM Proceedings 1991*, pages 86-90. SEMI-THERM Conference,
Phoenix, Arizona, February, 1991.
- [2] B. K. Bage and G. P. Peterson.
A Review of Entrainment Limits in Thermosyphons and Heat Pipes.
In *Porous Media, Mixtures, and Multiphase Heat Transfer*. Winter Annual Meeting,
ASME, San Francisco, December, 1989.
- [3] T. Borden.
Shape Memory Alloys Form a Tight Fit.
Mechanical Engineering 113, No. 10:66-72, 1991.
- [4] M. Coxon, C. Kershner and D. M. McEligot.
Transient Current Capacities of Bond Wires in Hybrid Microcircuits.
IEEE Trans. Components, Hybrids, and Manufacturing Technology
CHMT-9(3):279-285, September, 1986.
- [5] D. Dobberpuhl et al. (1992).
A 200MHz 64 Bit Dual Issue CMOS Microprocessor.
In *Proceedings of the 1992 International Solid-State Circuits Conference*. San Francisco,
California, February, 1992.
- [6] J. S. Fitch.
A Comparison of Acoustic and Infrared Inspection Techniques for Die Attach.
In *I-THERM III Proceedings 1992*, pages 159-164. I-THERM Conference, Austin,
Texas, February, 1992.
- [7] J. L. Hennesey and N. P. Jouppi.
Computer Technology and Architecture: An Evolving Interaction.
Computer 24(9):18-29, September, 1991.
- [8] G. J. Hill, B. C. Roberts and C. P. Strudwick.
Safe Operating Limits for Backdriving in ATE.
In *Proceedings of the 1987 International Test Conference*, pages 798-804. IEEE, 1987.
- [9] R. Hofbeck, K. Hausmann, B. Ilscher and H. U. Kunzi.
Fatigue of Very Thin Copper and Gold Wires.
Scripta Metallurgica 20:1601-1605, 1986.
- [10] J. P. Holman.
Heat Transfer.
McGraw-Hill, Washington , 1976 - 4th Ed.
ISBN 0-07-029598-0.
- [11] J. Schultz.
A CMOS 100Mhz Microprocessor.
In *Proceedings 1991 IEEE International Solid State Circuits Conference*, pages 90-91.
IEEE, San Francisco, February, 1991.

- [12] W. E. Jahsman.
Lead Frame and Wire Length Limitations to Bond Densification.
Journal of Electronic Packaging 111(4):289-293, December, 1989.
- [13] A. D. Kraus and A. Bar-Cohen.
Thermal Analysis and Control of Electronic Equipment.
Hemisphere, Washington , 1983.
ISBN 0-07-035416-2.
- [14] L. Kohn and S. Fu.
A 1,000,000 Transistor Microprocessor.
In *Proceedings 1989 IEEE International Solid State Circuits Conference*, pages 54-55.
IEEE, New York, February, 1989.
- [15] W. R. McGillis, V. P. Carey, J. S. Fitch and W. R. Hamburgren.
Pool Boiling Enhancement Techniques for Water at Low Pressure.
In *Proceedings of the Seventh IEEE Semiconductor Thermal Measurement and Management Symposium*, pages 64-72. 1991 IEEE Semiconductor Thermal Measurement and Management Symposium - SEMI-THERM, Phoenix, Arizona, February, 1991.
- [16] W. R. McGillis, V. P. Carey, J. S. Fitch and W. R. Hamburgren.
Pool Boiling on Small Dissipating Elements at Subatmospheric Pressure.
In *Proceedings on Phase Change Heat Transfer*, pages 27-36. 1991 ASME National Heat Transfer Conference, Minneapolis, Minnesota, July, 1991.
- [17] W. R. McGillis, J. S. Fitch, V. P. Carey and W. R. Hamburgren.
Boiling Binary Mixtures at Subatmospheric Pressure.
In *I-THERM III Proceedings 1992*, pages 127-136. I-THERM Conference, Austin, Texas, February, 1992.
- [18] T. Miwa, K. Otsuka, et al.
High Reliability and Low Cost Design in Plastic PGA Package with High Performance.
In *Proceedings 41st Electronic Components & Technology Conference*, pages 183-189.
IEEE, Atlanta, Georgia, May, 1991.
- [19] Y. Ohno, et al.
Development of Ultra Fine Wire for Fine Pitch Bonding.
In *Proceedings 41st Electronic Components & Technology Conference*, pages 519-523.
IEEE, Atlanta, Georgia, May, 1991.
- [20] D. Priore, Digital Equipment Corp., personal communication, September, 1990.
September, 1990.
- [21] E. M. Savitskii (ed).
Handbook of Precious Metals.
Hemisphere, New York, 1989.
ISBN 0-89116-709-9.
- [22] B. Shu.
Fine Pitch Wire Bonding Development Using a New Multipurpose, Multi-pad Pitch Test Die.
In *Proceedings 41st Electronic Components & Technology Conference*, pages 511-518.
IEEE, Atlanta, Georgia, May, 1991.

- [23] D. Tanksalvala, et. al.
A 90Mhz CMOS RISC CPU Designed for Sustained Performance.
In *Proceedings 1990 IEEE International Solid State Circuits Conference*, pages 52-53.
IEEE, San Francisco, February, 1990.

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. Hamburgren.

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. Hamburgren.

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.
- “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.

“A Simulation Based Study of TLB Performance.”

J. Bradley Chen, Anita Borg, Norman P. Jouppi.
WRL Research Report 91/2, November 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. Hamburgren, 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. Hamburgren.

WRL Research Report 91/9, August 1991.

“Experience with a Software-defined Machine Architecture.”

David W. Wall.

WRL Research Report 91/10, August 1991.

“Network Locality at the Scale of Processes.”

Jeffrey C. Mogul.

WRL Research Report 91/11, November 1991.

“Packaging a 150 W Bipolar ECL Microprocessor.”

William R. Hamburgren, John S. Fitch.

WRL Research Report 92/1, March 1992.

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.

“Unreachable Procedures in Object-oriented Programming.”

Amitabh Srivastava.
WRL Technical Note TN-21, November 1991.

“Cache Replacement with Dynamic Exclusion”

Scott McFarling.
WRL Technical Note TN-22, November 1991.

“Boiling Binary Mixtures at Subatmospheric Pressures”

Wade R. McGillis, John S. Fitch, William R. Hamburg, Van P. Carey.
WRL Technical Note TN-23, January 1992.

“A Comparison of Acoustic and Infrared Inspection Techniques for Die Attach”

John S. Fitch.
WRL Technical Note TN-24, January 1992.