Journal of Electronic Materials, Vol. 16, No. 5, 1987

A High-Performance Thermal Module for Computer Packaging

DEEPAK NAYAK, LIH-TYNG HWANG, IWONA TURLIK,[1] AND ARNOLD REISMAN[2]

Microelectronics Center of North Carolina (MCNC)
Research Triangle Park
P.O. Box 12889, NC 27709

[1]MCNC Resident Professional
Bell Northern Research
Research Triangle Park
P.O. Box 13478, NC 27709

[2]Also with the Dept. of Electrical and Computer Engineering
North Carolina State University
P.O. Box 5275, Raleigh, NC 27650

A thermal module was designed to transfer heat efficiently from high power dissipation chips to a liquid coolant via forced convection. Turbulent and laminar flow regimes were investigated. Channel geometries for deep channels (1000 μm deep, and used for turbulent flow), and shallow channels (100 μm deep, and used for laminar flow) were optimized for high heat transfer coefficient, ease of fabrication, and better structural rigidity of the module. A 4″ × 4″ module, made out of Cu, was tested using a 4″ Si "thermal" wafer as a heat generating source as well as a temperature sensor. Wafer scale integration and high energy ion implantation were employed to obtain nine 1 × 1 cm heat sources, and temperature sensing diodes embedded within the thermal wafer. For the deep channel design, the maximum device temperature rise on the module was 18° C for a power dissipation of 42 W/chip, and a flow rate of 126 cc/sec. For the shallow channel design, the temperature rise was 19° C for a flow rate of 19 cc/sec, and a power dissipation level of 42 W/chip. With all nine chips on the thermal module powered to 42 W/chip, the maximum chip to chip temperature variations were found to be 2 and 8° C for deep and shallow channel designs, respectively.

Key words: Packaging, heat dissipation, multi-chip carrier.

INTRODUCTION

Semiconductor devices are being scaled to submicron dimensions to achieve high-density, high-speed integrated circuits. For a scaling factor $\alpha(>1)$,[1] the density of circuits/unit-chip-area increases by a factor of α^2, and the power dissipation/device decreases by a factor $1/\alpha^2$. This would result in constant power dissipation/unit-chip-area as dimensions are scaled. In practice, however, the power supply voltage, V, is not scaled[2] to V/α, but is either kept constant (5V for NMOS & CMOS), or is scaled by a factor κ where $\alpha/\kappa > 1$. This introduces an undesirable increase in power dissipation/unit-chip-area. Furthermore, advances in processing technology have lead to an increase in chip size with acceptable yield. Therefore, power dissipation/chip is increasing dramatically.[3] At high levels of integration, power dissipation of NMOS circuits can exceed 2 W/chip. This is above the level at which convective air cooling can be employed effectively to cool a chip. With CMOS technology, power dissipation/circuit increases with an increase in switching frequency (for example, in the case of CMOS, it is CV^2f, where C is the circuit capacitance, and f is the switching frequency). At the GHz level, CMOS and NMOS dissipation levels can approximate each other. In any event, the power dissipation is a much more severe problem in case of bipolar technology. The power dissipation of a large computer employing high power dissipation chips can be large. For example, IBM's 3081 processor complex,[4] employing TTL circuitry, dissipates 23 kW, and the CRAY-2 computer,[5] employing ECL circuitry, dissipates 194 kW.

As the circuit density on a chip increases, it is necessary to enhance the device-circuit reliability. Many of the physical and chemical processes that can cause chip failure are accelerated by elevated temperature.[6] At the system level, for example, electromigration, corrosion and interfacial diffusion mechanisms increase with increasing temperature leading to premature failures in metallization and bonded interfaces. The mean time to failure is decreased by a factor of two with approximately a 10° C increase in operating temperature.[2] Another factor which affects performance is chip to chip temperature variation, as well as absolute temperature. The former impacts the forward characteristic spread of diodes while the latter impacts leakage characteristics. It is, therefore, desirable to maintain a low device operating temperature, and to minimize chip to chip temperature variations.

A thermal module design is described which exhibits a maximum chip operating temperature of

(Received February 5, 1987)

0361-5235/1987/1401-357$5.00©87 AIME

40° C when the inlet water temperature is at 20° C, and all the chips on the thermal module are powered to 42 W/chip. With all the chips on the module powered to 42 W/chip, the maximum chip to chip temperature variation was found to be 2 and 8° C for deep and shallow channel designs, respectively. The worst case chip to chip temperature variation occurs when all but one chip at the inlet end of the thermal module are powered to 42 W/chip. In this case, the maximum chip to chip temperature variations are 18° C for the deep channel design (flow rate 126 cc/sec), and 19° C for the shallow channel design (flow rate 19 cc/sec).

Natural convection as well as forced air cooling are not adequate to meet the demand of high performance packages.[2,3] Tuckerman and Pease[7,8] have demonstrated that heat removal from a chip can be enhanced by forced water cooling using microscopic channels in a Si substrate. This design enables removal of 790 W/cm^2 for a temperature rise of 71° C for a single chip. Their technique requires sophisticated anisotropic etching of 50 μm wide grooves on 100 μm centers to a depth of 300 μm (out of a 400 μm thick Si wafer). It also requires a special bonding process to seal those microscopic channels in this thermal module. The pressure drop across the channel is 31 psi for a channel length of 1.4 cm. This would translate into an unacceptable pressure drop of 200 psi (assuming the same mean flow velocity) for a 9 × 9 cm multi-chip thermal module. In addition, the caloric heating of water would be high due to the small flow rate possible. In a similar microchannel approach, Mahalingam[2] was able to remove 1000 W from a 3.8 × 3.8 cm 'super chip' on a 5 × 5 cm Si substrate while allowing a temperature rise of 30° C.

Kishimoto and Ohsaki[9] of NTT, Japan, have used 800 μm wide, 400 μm deep microchannels pitched at 2.54 mm in their thermal module. This cooling system allowed a dissipation of 16 W/chip for a temperature rise of 50° C. The cooling capability was limited by the low thermal conductivity of an alumina substrate used in this package. Another high performance package is IBM's Thermal Conduction Module (TCM).[10,11] A 9 × 9 cm multi-layer ceramic TCM supporting about 100 chips can dissipate up to 300 W. The device junction temperature can be controlled between 40 and 85° C. The maximum thermal resistance from a device on the chip to the liquid coolant was measured as 91° C/W. This high value of thermal resistance was due to the internal resistance from the device to the cold plate via a He ambient atmosphere, conductive pistons, springs, and assembly hardware. The maximum heat dissipation/chip for any single chip (0.46 × 0.46 cm) on this module was only 4 W.

In the present report, we have used a 3 × 3 array of 1 × 1 cm "thermal" chips constructed within a 5.4 × 5.4 cm substrate to test the heat removing capability of an analytically designed 25 chip thermal module prototype (9 × 9 cm).[12] The channel geometries (deep and shallow) on the 9 chip module are the same as on the 25 chip "theoretical" modules, but the water flow rate is scaled accordingly to keep the linear water velocity the same in the deep and shallow channels. The module design is intended to enable removal of 40 W/chip with a maximum operating temperature of 40° C (with an inlet water temperature of 20° C). Channel design was optimized not only in terms of high heat transfer coefficient, but also for ease of fabrication. Turbulent flow conditions were obtained in the deep channel design, and provided a high heat transfer coefficient coupled with a smaller pressure drop in comparison to shallow channel design. The spacing between channels was large enough to obtain easy and reliable bonding of the top and bottom plates. Bonding techniques such as soldering, anodic bonding,[13] welding or brazing,[14] and SiO_2-SiO_2 bonding (in case of a Si-Si package)[15] can be used. The channels can be made by machining, etching, molding, sawing, etc. Different substrate materials such as SiC, AlN, Mo, Si, Pyrex 7740 and Cu can be employed to obtain optimum combinations of thermal coefficient of expansion and thermal conductivity. Furthermore, the maximum temperature rise of the water due to caloric heating is small (for a power dissipation of 42 W/chip, for example, the water temperature rise is less than 1° C for a water flow rate of 126 cc/sec).

It has been found experimentally that the maximum temperature rise of a chip (all nine chips powered to 42 W/chip) is 18° C for the deep channel design employed (turbulent flow at 126 cc/sec). The pressure drop in this case was 8 psi (0.55 × 10^5 nt/m^2). For shallow channels (using the same lateral dimensions as in the deep channel design, but with reduced channel depth), and for laminar flow (flow rate of 19 cc/sec), it has been found experimentally that the maximum temperature rise on a chip was 19° C. The pressure drop across the channel was 41 psi (2.83 × 10^5 nt/m^2) for a flow rate of 19 cc/sec.

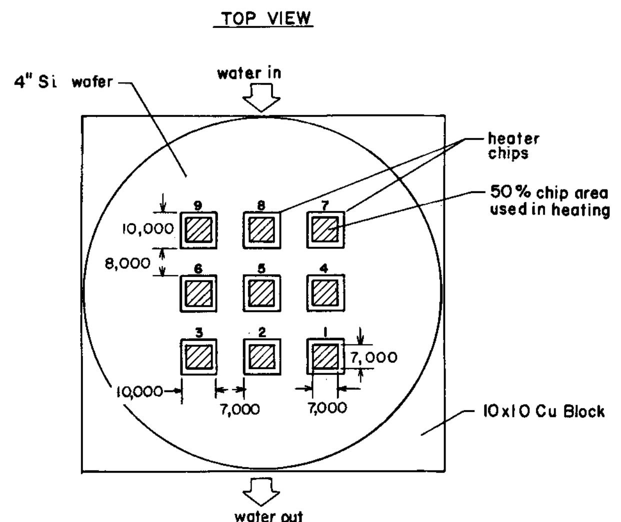

Fig. 1(a) — Top view of the thermal module showing the layout of 9 heat sources on a 4" Si wafer. Only 50% of each chip area is used in heating.

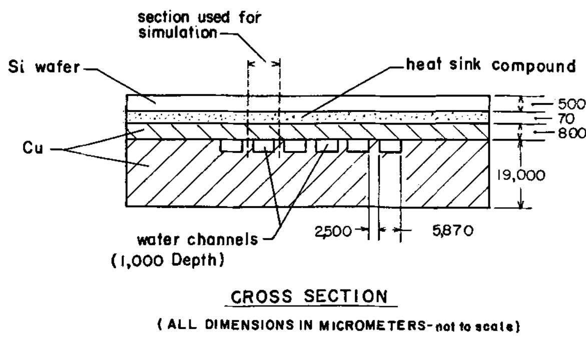

Fig. 1(b) — Cross sectional view of the thermal module.

COOLANT CHANNEL DESIGN

The worst-case expected theoretical temperature rise of the chips was calculated using a thermal distribution model,[12] in which the heat is removed by water only, i.e., natural convection due to air, and heat spread via conductive metal paths were ignored. In the thermal distribution model employed, isolated heat sources and sinks (the water channels) were assumed, and the temperature distribution was calculated by solving the conduction equation with appropriate boundary conditions. Because of the symmetrical nature of the chip-channel configuration, we could obtain the temperature distribution on the module by calculations on one section of the module, Figs. 1 and 2. Because of the structural symmetry of the channels, the temperature gradients at the left and right sides of each section are zero. Therefore, it was assumed that each section is thermally isolated from adjacent sections on either side of it. Since it has been further assumed that the heat is only removed by the coolant, the top and bottom surfaces must also be thought of as thermally isolated. The forced convective heat transfer coefficients of the water channels were calculated using the methods described in Refs. 16 and 17. Temperature distribution within the analyzed section was solved by converting the 2-D conduction equations to finite difference equations for every node

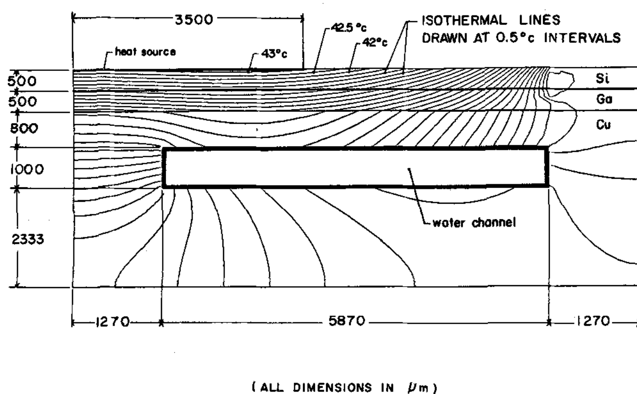

Fig. 2(a) — A plot of the temperature distribution on a cross section of the thermal module for deep channels (1000 μm deep). The section is taken through the center of a chip located at the outlet end. Inlet water temperature 20° C; Power dissipation/chip 42.4 W; Heat transfer coefficient 2.98 W/cm^{2}° C; Isothermal lines drawn at 0.5° C interval.

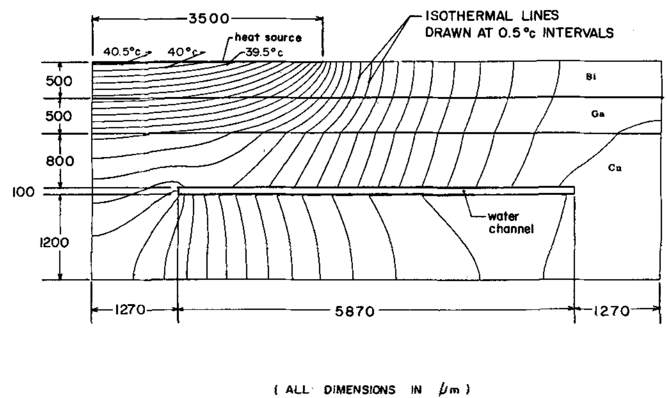

Fig. 2(b) — A plot of isothermal lines for shallow channels (100 μm deep) for same conditions as in Fig. 2(a). Heat transfer coefficient is 2.47 W/cm^{2}° C.

within, and at the boundary of the section, and solving it by matrix inversion techniques.[18,19] Since the theoretical model predicts the worst case temperature distribution, the calculated temperature rise is higher than that expected experimentally.

Figure 2(*a*) and (*b*) show calculated isothermal plots for the deep and shallow channel designs, respectively. The top layer is the structure containing "the thermal chips." In our case the top layer and the chips were assumed to be made of Si. The second layer (heat sink compound material) is the Ga seal. Typical thickness of the seal is in the range 70 to 100 μm. For analysis, it was assumed that the seal was 100 μm thick. If the temperature distribution for the seal were to be calculated, the grid size of the simulation has to be much smaller than 100 μm. In this case, the number of nodes would be too large (20,000 for the entire section if a grid size of 50 μm were used) to implement in a computer. Because of this, the thickness of Ga was artificially enlarged to 500 μm, and the thermal conductivity of the Ga was also scaled up by a factor 5 to preserve the temperature difference across the Ga layer. Under this situation, the grid size could be made larger (166 μm, for example, in our simulation) to reduce the number of nodes analyzed. On Figs. 2(a) and (b), the temperature difference between adjacent isothermal lines is 0.5° C. Theoretical results for deep and shallow channel designs, Figs. 2(a) and 2(b), respectively, are shown in Table-I.

The theoretically required water pressure can be calculated as shown in Refs. 16 and 20, by assuming no additional pressure drops due to contraction, expansion, and right-angle turns. For a given average linear stream velocity, the calculated turbulent flow pressure is less than the calculated laminar flow pressure because of boundary layer thickness differences (the boundary layer is thicker in the laminar flow case). For example, for an average linear stream velocity of 5.4 m/sec, the calculated pressure drop is 2 psi in the deep channels (189 cc/sec) vs 44 psi in the shallow channels (19 cc/sec). The experimental pressure drops were found to be 16 psi for the deep channels (189 cc/sec), and 41 psi for shallow channels (19 cc/sec). In case of deep chan-

Table I. Comparison of theoretical and experimental results for maximum temperature rise on a chip at the outlet end for different power dissipation/chip.

POWER / CHIP (WATTS)	FLOW RATE cc/sec	MAXIMUM TEMP RISE ON A CHIP (AT THE OUTLET END) (ΔT, °C)		
		DISTRIBUTION MODEL	EXPERIMENT	
42.4	189	23.4	14.75 *	deep channel
32.5	189	18.0	11.0 *	deep channel
22.8	189	12.5	9.0	deep channel
11.19	189	6.0	5.0	deep channel
42.4	19	25.1	18.7	shallow channel
32.5	19	19.34	15.0	shallow channel
22.8	19	13.65	9.0	shallow channel
11.19	19	6.5	4.7	shallow channel

* EXTRAPOLATED VALUES FROM FIGURE 5.

nels, the reason for the large descrepancy between the theoretical and the experimental pressure drop is not clear, but is believed to be due mainly to the pressure drop associated with the change of momentum of a large quantity (compared to a shallow channel) of water flowing through the right angle turns of the connecting pipes. It is to be noted that the flow of water (between the two pressure gauges in the experiment) changes its direction four times at the right angle turns.

TEST WAFER DESIGN

There is no standard test-chip to test the thermal characteristics of a ceramic or plastic single-chip package.[21] Infrared radiometry could be used to measure directly the temperature of a surface. The most popular and accurate technique, however, is to use diffused resistors either in series or in parallel combinations to generate heat, and to employ constant current driven diodes to measure the junction temperature.[21,22,23] In our multi-chip module testing, the resistors and diodes were formed using high energy ion implantation on a 4″ Si wafer. The layout of nine 1 × 1 cm "test-chips" is shown in Fig. 1(a). Fig. 1(b) shows the cross sectional view of the thermal module. In an actual chip, only 40–50% of the chip area is utilized for fabricating active devices. Therefore, on each thermal chip (1 × 1 cm) only 50% of the area was used as heating elements, Fig. 1(a). All the transistors on a chip do not switch at all times. Depending on the function being executed, only a portion of the circuits on a chip are switching. This introduces some non-uniformity in heat dissipation on the entire chip. In order to account for this non-uniform heating, each heater element was configured with several heater stripes (210 μm × 7000 μm each), Fig. 3.

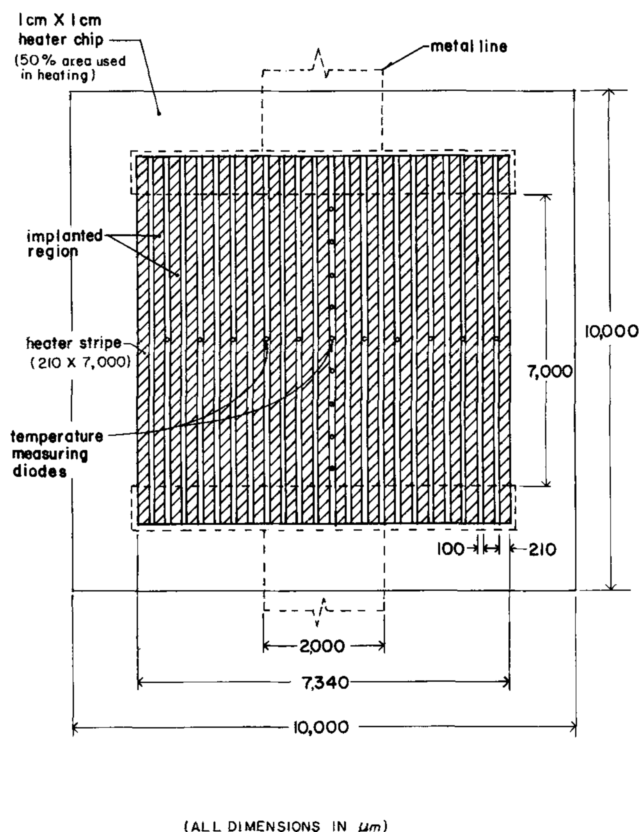

Fig. 3 — A detailed layout of chip-5 showig the heater doping and diode positions. Each heat source consists of parallelly connected 24 heater stripes (210 × 7000 μm each).

Each "thermal" chip was designed to dissipate a maximum of 40 W. The applied voltage to a chip, and the current in a chip to achieve a maximum power dissipation of 40 W were 40 V and 1 A, respectively. This required that the resistor diodes exhibit a reverse junction breakdown voltage greater than 40 V. The reverse breakdown voltage is a function of substrate doping and junction depth. For a p-type, 6 Ω-cm Si wafer as the substrate, and a 4 μm deep junction (a greater depth would require a longer implant energy and a higher drive-in time), the reverse breakdown voltage for an abrupt junction was estimated[24] to be 80 V (for a spherical junction geometry). A 4 μm deep junction was obtained using a 150 keV Phosphorus ion implant (dose = 2 × $10^{15}/cm^2$) followed by a 160 minutes drive-in at 1150° C. The sheet resistivity was measured to be 36 Ω/□, and the junction was reverse biased up to 42 V without observed breakdown.

In addition to resistor diodes, temperature measuring diodes were built within each chip. There was one diode at the center of each chip. In addition, there were 12 more diodes at different locations on chip #5, Fig. 3, to enable detailed temperature distribution measurements on a single chip. The diode area was 10 × 10 μm, and the distance between diode doping and heater doping was 40 μm. The diodes were operated at a constant current of 100 μA. The forward voltage drop across a diode is a function of temperature[25] at that point on the substrate. Typ-

ically, the change in diode voltage with temperature was found experimentally to be −2.5 mV/°C, and each diode was individually calibrated.

PROCESSING OF TEST WAFER

As mentioned, a p-type, 6 Ω-cm, ⟨100⟩, 4″ Si wafer was used in a 3-mask process. First, a 700 nm field oxide was grown at 950° C. Using mask-1, selected areas on the oxide were wet etched to expose Si for doping of both heater and diode regions. Phosphorus ions were implanted at 150 keV for a total dose of $2 \times 10^{15}/cm^2$. This was followed by a drive-in step at 1150° C, for 10 min in dry O_2, and then 150 min in Ar. This high temperature and long drive-in time was required to obtain a deep junction (4 μm). The process was simulated by using PREDICT.[26] The oxide was wet etched to open contact holes using mask-2. A 1000 nm Al-2%Si layer was evaporated. The metal interconnections were patterned using mask-3. The wafers were then annealed in H_2 for 20 min at 400° C. The back sides of the wafers were metallized by evaporating 100 nm of Cr followed by 1000 nm of Cu. This was necessary to obtain good soldering between the test wafer and the Cu thermal module. Finally, wafers were again annealed in H_2 for 20 min at 400° C.

EXPERIMENTAL PROCEDURE

The experimental set up is shown in Fig.4. The grooves for the channels and the manifolds were machined out of a 1.9 cm thick Cu block. This block was soldered to a thin 0.08 cm Cu plate to form the water channels, Fig.1(b). Soldering was done by using a 60Sn/40Pb solder preform sheet in a N_2 ambient atmosphere. The thermal contact between the Si test wafer and the Cu plate was obtained using Ga (thermal conductivity, k = 0.281 W/cm°C). Electrical contacts to the test wafer were made via spring loaded probes. The temperature (T_i), pressure (P_i), and flow rate F at the inlet, and at the outlet (T_0,P_0) were recorded. For deep channels, the flow rate was varied from 13 cc/sec (average linear stream velocity = 0.4 m/sec) to 189 cc/sec (average linear stream velocity = 5.4 m/sec). For shallow channels, measurements were made at 13 cc/sec (average linear stream velocity = 3.7 m/sec) and 19 cc/sec (average linear stream velocity = 5.4 m/sec). The module (Cu block, Si wafer, and brass connectors at the water inlet and outlet ends) were thermally insulated using a ceramic fiber insulator (50% Al_2O_3 − 49% SiO_2) to prevent heat loss due to natural convection and radiation. Prior to use, each individual diode on the module was calibrated to determine its voltage-temperature relation in a closed oven at a constant current of 100 μA. The accuracy of the temperature measurement was ±0.3° C.

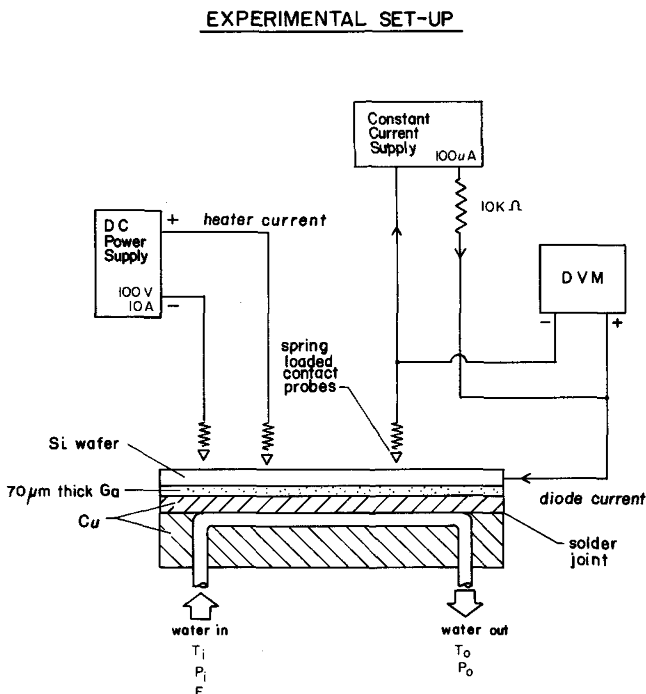

Fig. 4 — Experimental set up used to test the thermal module. Spring loaded probes make electrical contacts with the test wafer.

EXPERIMENTAL RESULTS

For deep channels, Fig. 5 shows the temperature on a chip at the outlet end for different power dissipations and flow rate values with an input water

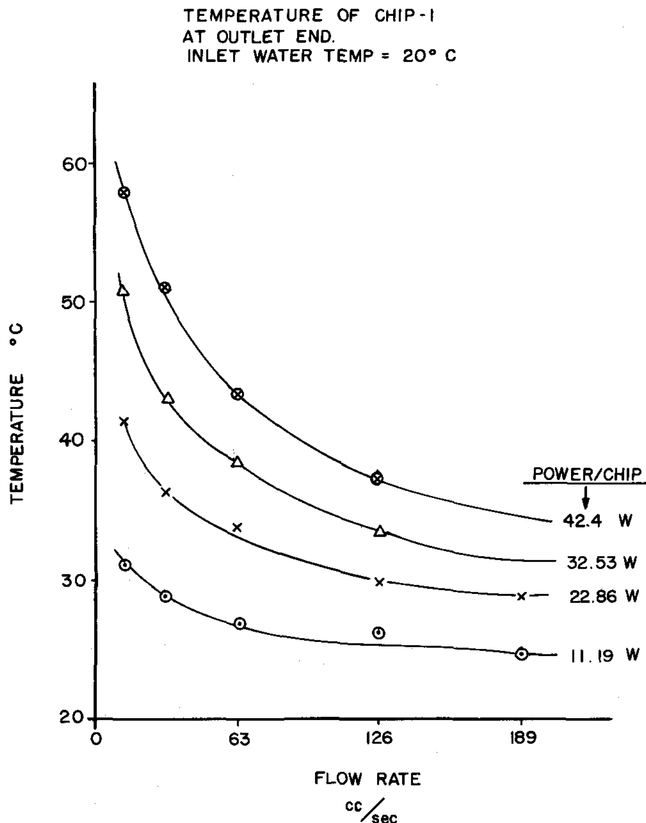

Fig. 5 — A plot of maximum temperatures (for deep channels) on chip-1 located at the outlet end as a function of flow rate and power dissipation/chip. Inlet water temperature 20° C.

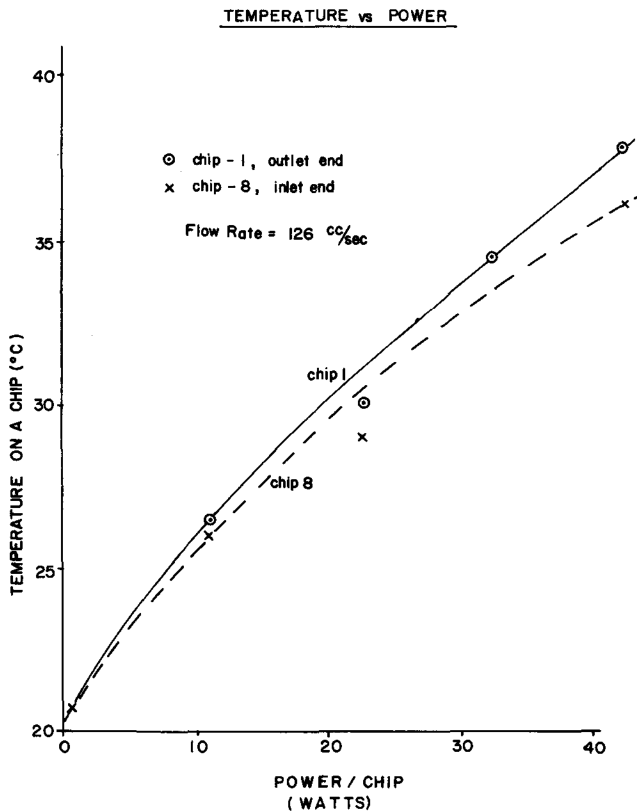

Fig. 6 — A plot of maximum temperature (for deep channels) on chip-1 (at the outlet end) and on chip-8 (at inlet end) as a function of power dissipation/chip. Flow rate 126 cc/sec; Inlet water temperature 20° C.

temperature of 20° C. At a flow rate of 126 cc/sec, and a pressure drop of 8 psi (0.55 × 10^5 nt/m^2), the temperature rise was 18° C for a power dissipation of 42.4 W/chip in each of the 9 chips. There is a minimum water flow rate for maximizing heat transfer from the thermal module to the water. Thus, the temperature on a chip reaches a plateau value for flow rates above 126 cc/sec, Fig. 5. For the same chip dissipating 42.4 W, the temperature rise is 38° C for a flow rate of 13 cc/sec. It was found from calculation that the flow was laminar for flow rates below 45 cc/sec (Reynolds number, Re = 2180, and average linear stream velocity = 1.3 m/sec). The temperature variation between two chips one located at the inlet and the other at the outlet end is shown in Fig. 6. The variation is less than 2° C for a power dissipation of 42.4 W/chip, and a flow rate of 126 cc/sec (in a deep channel design). Fig. 7 shows the temperature profile on chip #5 (see Fig. 1) at a power dissipation of 22.9 W/chip. The maximum temperature differential between the chip center and its edge is less than 2° C. All the above results were obtained by using Ga (thermal conductivity, k = 0.281 W/cm°C) as the contact medium between the module and test wafer.

Experimental results for shallow channels were also obtained by using Ga as the thermal contact medium between the module and the test wafer. Figure 8 shows the temperature variation on a chip at the outlet end for different flow rates and power

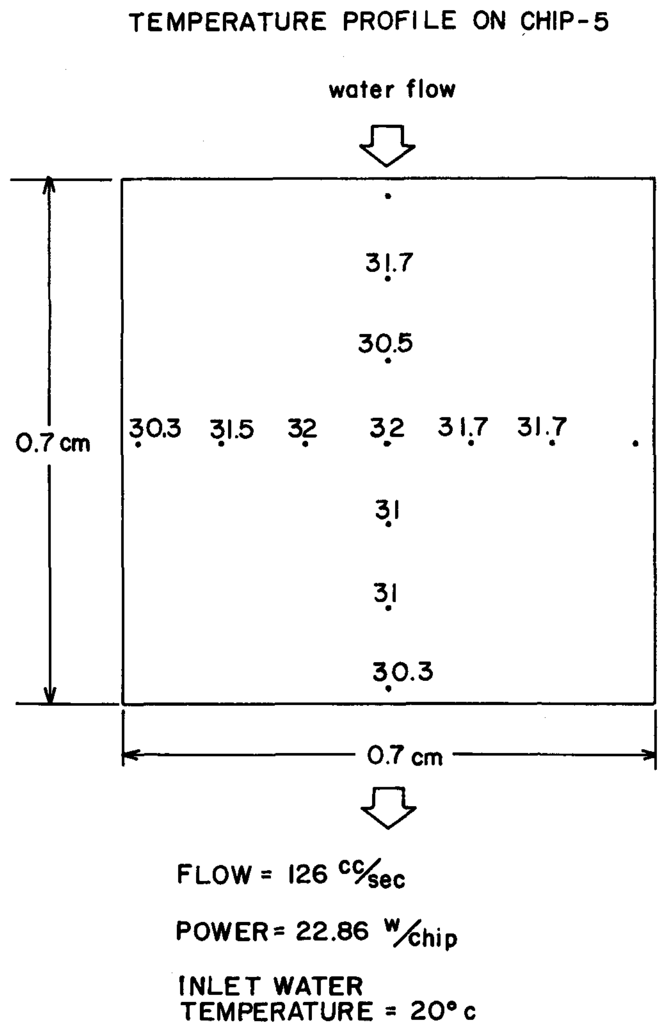

Fig. 7 — A temperature (°C) distribution profile on chip-5 for deep channels. Flow rate 126 cc/sec; Power dissipation/chip 22.86 W; Inlet water temperature 20° C.

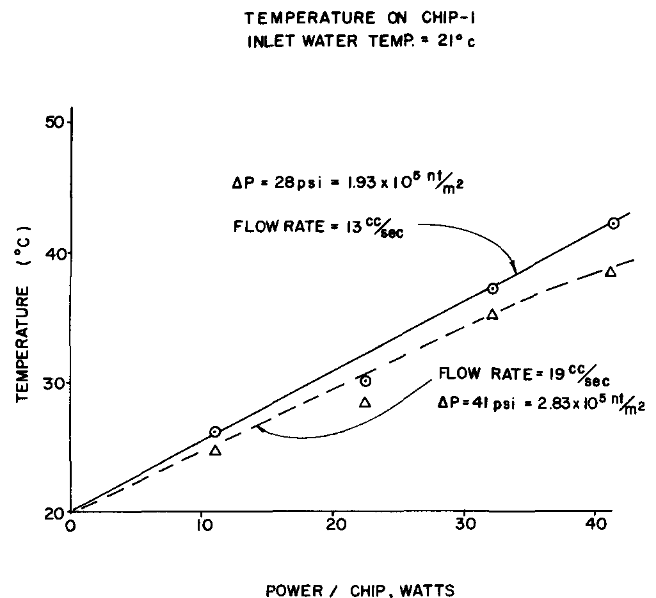

Fig. 8 — A plot of maximum temperature on chip-1 (for shallow channels) located at the outlet end as a function of flow rate and power dissipation/chip. Inlet water temperature 21° C.

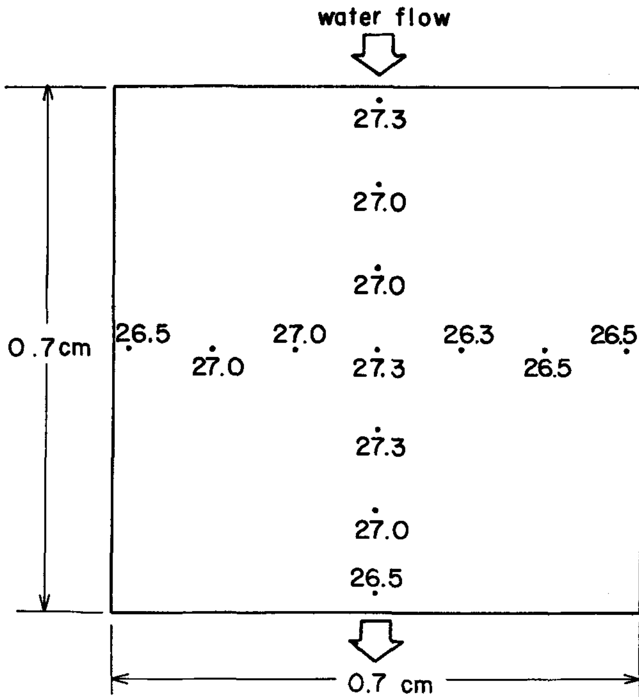

Fig. 9 — A temperature (°C) distribution profile on chip-5 for shallow channels. Flow rate 13 cc/sec; Power dissipation/chip 22.57 W; Inlet water temperature 21° C.

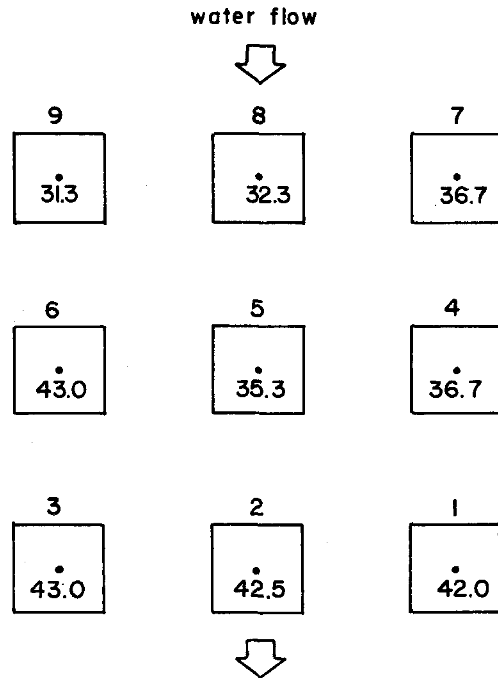

Fig. 10 — A temperature (°C) distribution profile on the thermal module for shallow channels. Flow rate 13 cc/sec; Power dissipation/chip 41.6 W; Inlet water temperature 21° C.

dissipations. At a power dissipation of 41.6 W/chip, the maximum temperature rise on the module is 19° C at a flow of 19 cc/sec and is 22° C for a flow of 13 cc/sec. At 19 cc/sec, ΔP is 41 psi (2.83×10^5 nt/m^2). Figures 9, and 10 show the temperature variation on chip #5, and on the entire module, respectively.

DISCUSSION

Table-I compares the results obtained from the theoretical analyses and experiments. The lateral spreading of heat depends upon the channel geometry, the thermal conductivity, and the distance between the channels and the heat source. The theoretical temperature rise is larger, than the experimentally determined data because we intentionally neglected the effects of thermal conduction by wires, air and the Cu block.

Equivalent thermal performance can be achieved by using either deep channels or shallow channels, see Table-I. The flow rate in deep channels is 10 times higher than that used for shallow channels for the same average linear stream velocity. The experimental pressure drop is, however, lower for deep channels (16 psi at 189 cc/sec for deep channels vs. 41 psi at 19 cc/sec for shallow channels, with the same average linear stream velocity 5.4 m/sec). The stresses introduced by water pressure for both the deep and shallow channel designs are much less than the rupture strength of both Cu and the solder bond. However, using a 4″ Si wafer (500 μm thick) as the cooling substrate instead of a Cu block, a pressure drop of about 50 psi ruptured the Si wafer at the inlet manifold. From structural rigidity, material options, and reliability considerations, a low pressure drop is therefore desirable. Furthermore, the water caloric temperature differential on the module for a small flow rate is significant (4.8° C at 42 W/chip for a flow rate of 19 cc/sec). In addition, an optimal flow rate can be determined for a given channel design as the heat transfer coefficient does

not increase linearly with flow rate. In our example, the heat transfer coefficient increases very slowly above a flow rate of 126 cc/sec, Fig. 5.

In the present study, Cu (thermal conductivity, k = 4.01 W/cm°C) was used as the substrate material only to verify the concept of turbulent flow and deep channel design. Using materials having comparable thermal conductivity as Cu, the temperature rise does not change significantly. It has been found from simulation that using SiC (k = 2.7 W/cm°C), and Si (k = 1.48 W/cm°C), the temperature rise on a chip (for a heat dissipation of 40 W/chip) will be about 1° C, and 3° C higher, respectively, than that obtained for Cu. For Al_2O_3 (k = 0.37 W/cm°C), the temperature rise will be 12° C higher than that obtained for Cu. While Cu is an easily fabricated material, a concern using it as a packaging material is the large coefficient of thermal expansion mismatch between Cu (16.7×10^{-6}/°C) and Si (2.5×10^{-6}/°C). If the temperature rise on the module can be maintained at ≤20° C, thermal mismatch would not be a major problem. The material selection for such a package will be largely driven by the electrical performance and the total packaging cost.

Even if the flow transport phenomena are optimized, heat removal can be degraded significantly if the interfacial contact resistance between the chips and the module is high. For proper performance of the thermal module, a good thermal contacting technique between the chips and the module to achieve a uniform thermal interface layer is necessary. Ga was used in this experiment because of its good thermal (k = 0.281 W/cm°C) and wetting properties, but using 60Sn/40Pb solder (k = 0.505 W/cm°C) instead would improve further the thermal performance of the module. The uniformity of the thermal interface layer and the presence of voids in this layer also affects the local temperature variation on the module. Ga was applied manually in the described experiments and the nonuniformity effects could be seen from temperature profile measurements. For example, chip #6 in Fig. 10 is at higher temperature than chips #4 and #5.

ACKNOWLEDGMENT

We would like to thank Alton Caviness of MCNC for his help in fabricating the thermal module and Neil Poley of IBM for valuable suggestions on bonding materials. We would also like to thank Drs. Carl Osburn of MCNC and NCSU, Dinesh Sharma and Scott Goodwin-Johansson of MCNC for valuable discussions. The help of the MCNC fabrication facility staff in processing the test wafer is gratefully acknowledged.

REFERENCES

1. A. Reisman, Proc. of IEEE, *71,* 550 (1983).
2. M. Mahalingam, Proc. IEEE *73,* 1396 (1985).
3. S. Oktay, R. Hannemann, and A. Bar-Cohen, Mech. Eng. 36 (1986).
4. M. S. Pittler, D. M. Powers, and D. L. Schnabel, IBM J. Res. Develop., *26,* 2 (1982).
5. R. D. Danielson, N. Krajewski, and J. Brost, Electron. Packaging and Production 44 (1986).
6. S. M. Sze, editor, VLSI Technology, New York, McGraw-Hill Book Co., 1983, pp. 624–627.
7. D. B. Tuckerman and R. F. W. Pease, IEEE Electron Device Lett. *EDL-2,* 126 (1981).
8. D. B. Tuckerman, Heat-Transfer Microstructure for Integrated Circuits, PhD Thesis, Dept. of Electrical Engineering, Stanford University, Feb. 1984.
9. T. Kishimoto and T. Ohsaki, Proc. of 36th IEEE Electronics Component Conference (ECC), pp. 595–601, May 1986.
10. A. J. Blodget and R. Barbour, IBM J. Res. Develop. *26,* 30 (1982).
11. S. Oktay and H. C. Kammerer, IBM J. Res. Develop. *26,* 55 (1982).
12. L. Hwang, I. Turlik, and A. Reisman, to be published.
13. G. Wallis and D. I. Pomerantz, J. Appl. Phys. *40,* 3946 (1969).
14. T. J. Moore, "Welding and Brazing of Silicon Carbide," NASA Tech. Briefs, pp. 118–119, July/Aug. 1986.
15. J. B. Lasky, Appl. Phys. Lett., *48,* 78 (1986).
16. R. K. Shah, and A. L. London, Laminar Flow Forced Convection in Ducts, New York, Academic Press, 1978.
17. B. S. Petukhov, "Heat Transfer and Friction in Turbulent Pipe Flow with Variable Physical Properties," Advances in Heat Transfer, vol. 6, Academic Press, 1970.
18. B. Gebhart, Heat Transfer, New York, McGraw-Hill, 1971, pp. 48–56.
19. J. J. Dongarra, C. B. Moler, J. R. Bunch, and G. W. Stewart, Linpack User's Guide, Philadelphia, Siam, 1979.
20. Fox and McDonald, Introduction to Fluid Mechanics, 3rd edition, New York, Wiley, 1985.
21. F. F. Oettinger, Solid State Technol., pp. 169–179, Feb. 1984.
22. A. Reisman, M. Berkenblit, C. J. Merz, and A. K. Ray, J. Electron. Mater. *11,* 391 (1982).
23. B. J. Shanker and R. T. Lambertson, "Sources of Error in Package Thermal Resistance Measurement," Proc. 36th IEEE Electronics Component Conference (ECC), pp. 161–167, May 1986.
24. S. M. Sze, Physics of Semiconductor Devices, New York, Wiley-Interscience Publication, 1981, pp. 32–33, 96–108.
25. E. S. Schlig, IEEE J. Solid-State Circuits *SC-3,* 271 (1968).
26. PREDICT (Process Estimator for the Design of Integrated Circuit Technologies) is a trademark of the Microelectronics Center of North Carolina, P.O. Box 12889, Research Triangle Park, NC 27709.