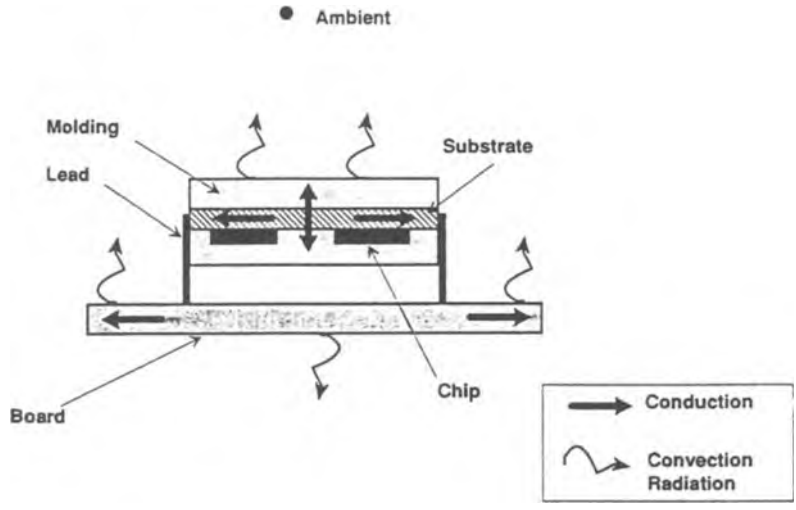

Figure 12-3 Heat transfer in a low performance multichip module.

potentially multiple heat sources in thermal communication with each other via the substrate and its ambient.

The heat generated at the chips seeks the most conductive path to reach the sinks. The sinks to which the heat is eventually transferred are the cooling fluid and the boards. The paths available for heat flow are through the substrate then the molding material and the leads. The flow of heat is impeded by each material, regardless of its thickness, as it travels from the sources to the sink. As the heat reaches the leads, part of it is conducted to the board, and the rest is either radiated or convected to the ambient.

The flow of heat spreads, seeking the greatest ratio of A to L by Equation 12-2 to maximize the conductive heat transfer rate. As a side effect of this heat spreading, each chip has a higher temperature due to the presence of its neighbors. The closer the chips are spaced, the greater will be this effect.

However, heat spreading also can be beneficial because of the increase in heat transfer rate. In plastic SCMs, this is done by inserting an aluminum plate, called a "heat spreader," into the package body. In some MCMs this is done with copper plates. As long as these copper plates conduct much more heat to the heatsinks than to each other, the net result is beneficial.

A similar process occurs as the heat reaches the physical boundaries of the component. Figure 12-3 gives a schematic depiction of different modes of heat

transfer and heat flow process in an MCM. Combination of multiple heat sources and different possible avenues for heat flow have created a rather complex and nonuniform temperature field.

12.3.3 The Concept of Thermal Resistance

The concept of thermal resistance is associated with impeding the flow of heat through a medium. Analogous to electrical resistance, if resistance is decreased, less voltage is required to pass the current through the wire. Current is similar to heat flow, and voltage to temperature. Hence, if thermal resistance is decreased, smaller temperature differences across a medium result. Thus, it becomes intuitive that if thermal resistance is minimized, internally and externally, the junction temperature is reduced.

The package thermal response can be viewed by two resistances, external and internal. The internal resistance, Θ_{jc} (junction to case resistance), addresses heat transfer within the package, from the chip to the surface of the package. The external resistance, Θ_{ca} (case to ambient resistance), is a measure of thermal transport occurring between the package surface and the ambient.

Θ_{jc} and Θ_{ca} are defined by the following equations,

$$\Theta_{jc} = \frac{T_j - T_c}{P} \tag{12-5}$$

and

$$\Theta_{ca} = \frac{T_c - T_a}{P}. \tag{12-6}$$

In Equations 12-5 and 12-6, subscripts a, c and j refer to ambient, case and junction, respectively. The denominator, P, is the total power dissipation in the component. The units of these resistances are °C/W. Figure 12-4 shows a schematic representation of these resistances.

There are three problems, however, with this use of the concept of thermal resistance. First, it is difficult to calculate a single number for thermal resistance with the presence of multiple heat paths. Second, even if you did calculate it, you could not reuse the resistance in other designs. Consider what would happen to the heat flow pictured in Figure 12-4 if another MCM were placed on the bottom of the board. Less heat would flow in this direction due to the smaller ΔT. The total heat flow would change and the thermal resistance would increase. These problems are common to MCMs and SCMs [13]. In fact, a common error in SCM design is to assume that the data sheet value for thermal resistance

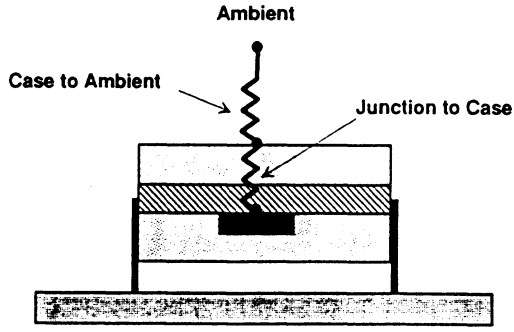

Figure 12-4 Thermal resistance representation in an electronic component.

applies to a crowded two-sided board. It often has to be increased by 50% or more to compensate for the ideal conditions under which the manufacturer measured it. The third problem is unique with MCMs. With multiple chips, each with a different power, P, each chip will have a different temperature, T_j, and thus the choice of values to be used in Equations 12-5 and 12-6 are arbitrary and somewhat meaningless.

Thermal resistance is a widely used concept. The preceding discussion has shown the weaknesses associated with using Equations 12-5 and 12-6 even for SCMs. These equations contain even higher error levels when applied to low performance MCMs. In the case of high performance MCMs, thermal resistance in this form is completely useless. Although thermal resistances for these MCMs still are reported, it typically is for the chip and not the entire module.

Nevertheless, thermal resistance is a very valuable concept for qualitatively understanding the thermal effects of different materials and cooling approaches. Equations 12-2 and 12-3 show conduction and convection heat transfer. Conductive and convective thermal resistances are defined as

$$R_k = \frac{L}{kA} \tag{12-7}$$

$$R_h = \frac{1}{hA} \tag{12-8}$$

which refer to internal and external resistances. By definition, Equations 12-7 and 12-8 are identical to equations 12-5 and 12-6. We see that R_k is inversely proportional to k and A (area normal to the direction of heat flow). R_h is inversely proportional to h and A (convective surface area). To reduce these resistances, the denominator would have to increase.

Consider an MCM where the chips are epoxied to the substrate. The thermal resistance between the chip and the substrate (where the heat is conducted away) is the one imposed by the epoxy. The area, A, associated with the epoxy is constrained to the size of the chip. Often times in the application of epoxy, air gaps are created adding to overall chip to substrate resistance. The resistance is reduced by removing the air gaps or using an epoxy with a higher thermal conductivity or another bonding technique. The Hitachi Silicon Carbide (SiC) RAM [14] is an example of such practice where 52 solder bumps are used for heat transfer purposes only (Figure 12-5). Of course, as we change process or epoxy, we have to be concerned with material compatibility to avoid uneven expansion. Otherwise, stresses induced as the result of uneven expansion may

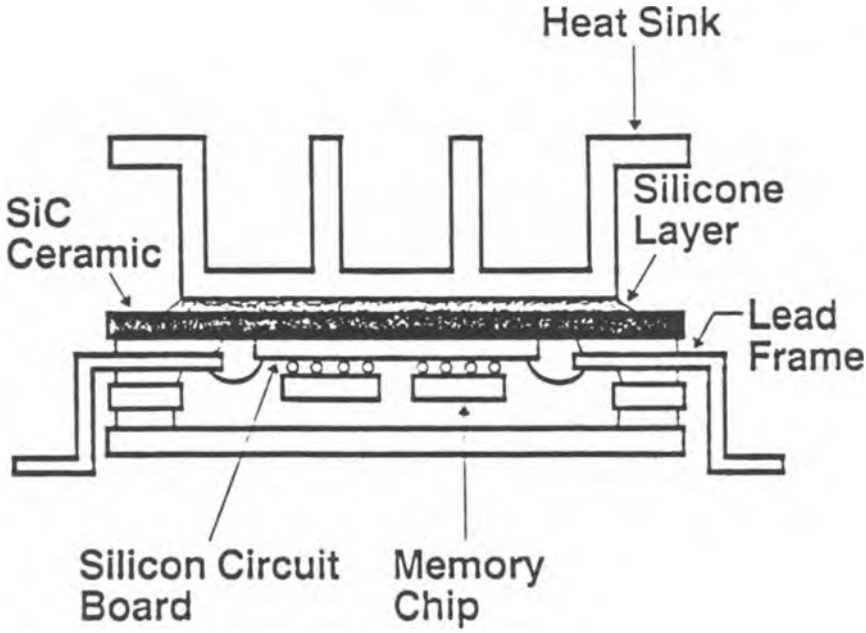

Figure 12-5 Cross sectional view of Hitachi air-cooled multichip module.

result in component failure. For this reason, thicker epoxies typically are used in laminate MCMs.

Similarly, if we look at R_h, it can also be reduced by increasing h or A. The heat transfer coefficient, h, is increased by going to higher velocity flows (fans or jet impingement) or changing the fluid (gas to liquid). The surface area is increased by adding heatsinks to the component.

Based on the above discussion, thermal resistance plays a pivotal role in the magnitude of junction temperature. Reducing thermal resistance either internal or external to the MCM package impacts thermal control positively. In the design of MCMs, it is important to locate the heat sources and the thermal resistances on their paths. The resistances should be reduced so that minimum spreading takes place, and the path from the chip to the sink has the least thermal resistance.

12.3.4 Heat Transfer On a Board

For thermal design purposes, each component cannot be considered in isolation. The heat being produced by one component is transferred amongst the others. This is referred to as thermal coupling and is discussed here at the board level, and in the next section at the system level. As discussed above, thermal coupling can be very strong within an MCM.

Consider the forced air cooled board shown in Figure 12-6. There are two mechanisms leading to thermal coupling. First, the MCMs and SCMs heat the air as it passes over them. The downstream parts will experience a hotter fluid

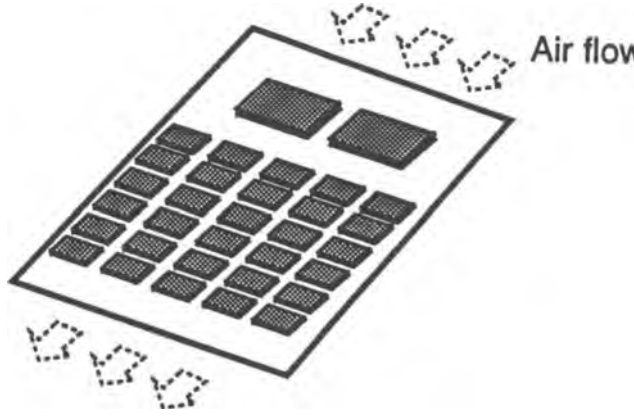

Figure 12-6 Flow over a circuit board containing SCMs and MCMs.

temperature, T_f and, according to Equation 12-3, the component is hotter, by the same amount, to dissipate the generated heat. Second, some of the heat produced by each component is passed to the board through leads, conduction and radiation across the air gap. Though part of this heat is convected away from the board, part of it is conducted to the other components on the board, raising their temperature. While the board dielectric is a poor conductor of heat, the copper layers within it are excellent heat conductors. An eight layer board exhibits a stronger conductive thermal coupling between components than does a four layer board.

The net effect of this is that the chip with the highest T_j might not be the chip that produces the most heat. The critical chip is the one that has a T_j closest to or over the specified limit. For example in Figure 12-6, the high power CPU core MCM is placed near the fan while the low power main memories are placed at the other end. Often this arrangement requires that special memories be purchased with an aluminum heat spreader within them.

12.3.5 Thermal Coupling in Electronic Enclosures

To appreciate the impact of the system (enclosure) on the thermal performance of MCMs, it is necessary to review the thermal phenomenon in an enclosure. (See Figure 12-7 for an example of an enclosure.)

The shelf or card holder (cage) is where a circuit board resides in the system. Boards are normally inserted into the shelves through card guides. Except in some specialized cases where a latching mechanism is used to rigidly attach the board to the shelf, the boards are loosely fitted inside the shelf. Therefore, the necessary contact to facilitate conduction heat transfer from the board to the shelf does not usually exist.

The backplane, or motherboard, in a PC is another avenue for the heat to be transported to the ambient or the shelf. If the thermal conductivity of the board is very large, that is multilayered boards with several layers of copper, conduction heat transfer through the backplane can be significant. However, the thermal coupling by convection and radiation heat transfers is significantly larger than conduction heat transfer.

Frames or enclosures that house single or multiple shelves generally are designed to be isolated from the shelves. Thus, the heat generated within the system normally is convected through the vent holes. Although this constitutes the bulk of heat flow, there exists significant thermal coupling between the boards (and shelves) and the frame. The thermal coupling, in the order of significance, is by radiation, convection and conduction heat transfer. Since the frame is in contact with the system ambient, it acts as a sink and source of heat for the system.

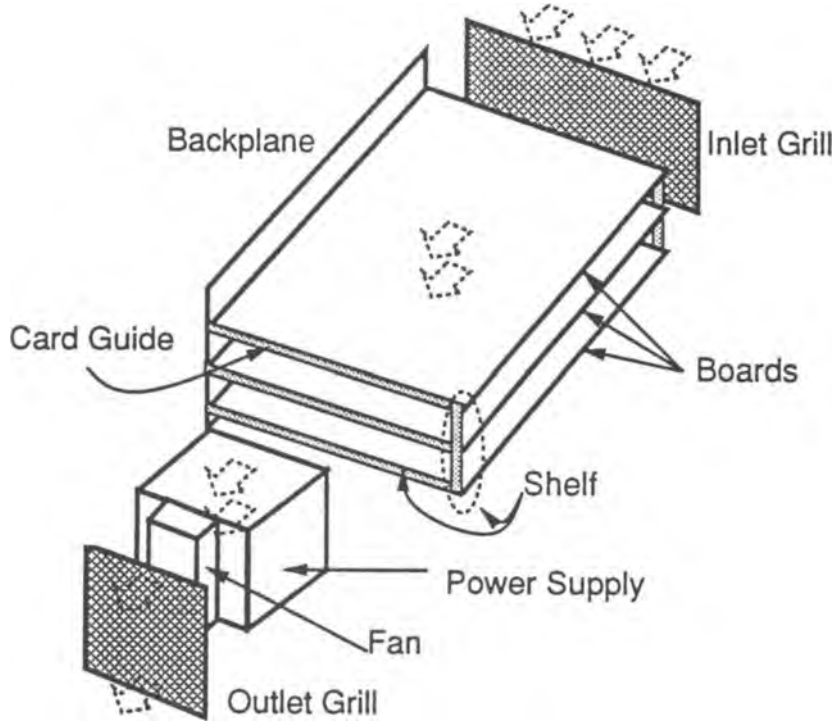

Figure 12-7 Schematic view of electronic system configuration.

The magnitude of conduction heat transfer is very system dependent. The radiation heat transfer, however, is generally the predominant mode of thermal coupling between the shelves and the frame. The radiation heat transfer tends to be even more significant if the system is cooled by natural convection.

The frame is coupled to the surrounding ambient via radiation and convection heat transfers. The system ambient also can act as a source and a sink. The magnitude of these heat transfers varies significantly with the changes in the system surroundings.

The thermal transport process in electronic systems is quite involved and can become very complex. Because of many different thermal processes and strong coupling at various system levels, thermal bookkeeping is necessary for accurate design. In addition, it should be clear that we cannot only focus on a component (module) without considering the system, environment and other parameters affecting thermal design.

12.4 THERMAL MANAGEMENT OF MCMs

Power dissipation levels exceeding 4 W/cm^2 and specific system performance requirements have forced design of highly customized cooling systems for some MCMs. MCMs typically have higher power dissipations than SCMs and generally are placed on circuit boards or substrates that contain other potentially high power components. The propensity for thermal spreading within the circuit board or the substrate has sometimes led the designers to build individual cooling elements for every chip on the substrate. This combination at times has created a challenging problem for thermal management of MCMs. The challenge has embraced packaging of the MCM itself and integration of its cooling system into the overall frame.

In low performance systems, thermal management of the MCM is typically an after thought and is constrained by the application. In high performance systems, the cooling system design is an integral part of the design cycle. Because of the customized nature of cooling systems, their spatial restrictions in terms of system compactness or electrical distance between high speed components has been an added challenge to thermal management.

The constraints imposed by temperature limits and system compactness have produced innovative packaging and thermal control methods. In this section, some of these techniques, with emphasis on cooling method, are presented. The objectives of thermal management, with respect to design and manufacturing constraints, are discussed. Then, the cooling approaches and order of application are reviewed.

12.4.1 Alternate Thermal Control Methods for MCMs

There are alternate thermal control methods for both the internal and external paths. The choice of the primary internal path is closely related to the choice of chip attach. The primary internal path alternatives are (Figure 12-8) through the substrate, through the substrate with thermal vias or thermal cutouts, and chip backside.

With through the substrate cooling, the main heat dissipation surface in contact with the fluid is beneath the chip. The primary internal heat path is through the substrate. There might be considerable heat spreading and the thermal resistance might be high, particularly as the thermal conductivity of the most common substrate materials is low. It can be reduced by using different materials, such as aluminum nitride instead of alumina. However, highly thermally conductive alternatives to polyimide and laminate MCM materials are not available.

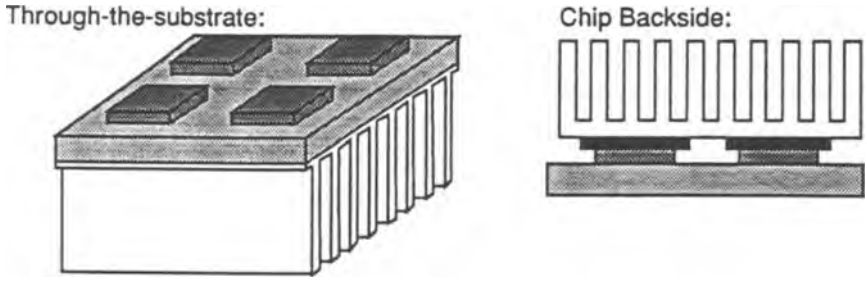

Figure 12-8 Through-the-substrate and chip backside technology for MCMs.

If the material properties cannot be improved, then the heat spreading and the thermal resistance can be reduced, at the expense of wiring capacity, by placing copper vias in the substrate or by sinking the chip into the substrate. Either of these techniques can be used with any chip connection technique. If wire bond or TAB leads are used, the chip can be attached with a thermal epoxy. Part of the heat is transferred through the thermal epoxy and part through the metal leads. With solder bump attachment, the heat is transferred through the bumps and the air gap. Often extra bumps are added to reduce thermal resistance.

With chip backside cooling, the main heat dissipation surface is above the chip and is attached with a metal mount, a high thermal conductivity solder, a thermal epoxy or a thermal grease. Only flip techniques (flip TAB or solder bumps) can be used because of potential damage to the surface of the chips. Chip backside cooling generally has the smallest internal thermal resistance.

In either thermal path, careful attention must be given to the interfaces. For example, if the epoxy chip interface has many air bubbles in it, the thermal resistance increases substantially.

In general, the external thermal control methods can be categorized as follows:

- Natural convection
- Forced convection
- Conduction or radiation cooling
- Liquid immersion
- Phase change (boiling)

Table 12-2 Thermal Control Methods.

Primary Cooling Mechanism	Typical HTC (W/m^2K)	Relative Effectiveness	Achievable Density	Complexity
Natural convection (air)	10	0.1	Low	Very low
Forced convection (air)	100	1.0	Medium	Low
Natural convection (liquid)	100	1.0	Medium	Medium
Forced convection (liquid)	1000	10.0	High	High
Phase change (liquid)	5000	50.0	High	High

Natural convection cooling is the case when no fluid movers are used in circulating the fluid in the system. Cooling by forced convection utilizes a fluid mover to circulate the fluid. Conduction or radiation cooling is when a cold plate [15] or radiation plate is used to remove heat. The application of the latter is seen in military and space electronics. Liquid immersion is when a component or system is immersed in a liquid. The liquid can be fluoroinerts or others such as liquid nitrogen. With boiling, the fluid boils at the MCM contact surface. General features of the convection-based cooling modes are given in Table 12-2 [16]. In that Table, HTC is the heat transfer coefficient h (see Equation 12-3).

A designer often is confronted with the decision of selecting a cooling method, keeping in mind manufacturing issues and end use application constraints. Selecting a thermal control method is a function of component temperature rating and the heat removal capacity of a specific design. The coolant fluid, gas or liquid, typically sets the capability of these cooling methods apart. This is made more clear by looking at Figure 12-9 [17]. The figure merges the thermal control methods with power dissipation (heat flux) and temperature rise. It shows a representation of the expected temperature rise over ambient for different cooling methods. It also hints to a potential thermal control method as a function of component power density.

Figure 12-9 does not suggest or provide an absolute case for a thermal control method. It suggests the expected range or type necessary for heat removal to ensure the junction temperature meets its constraint. For example, for an MCM with 10 W/cm^2 and a temperature rating, T_j, of 85°C, Figure 12-9 suggests some sort of liquid cooling. However, the same component can be

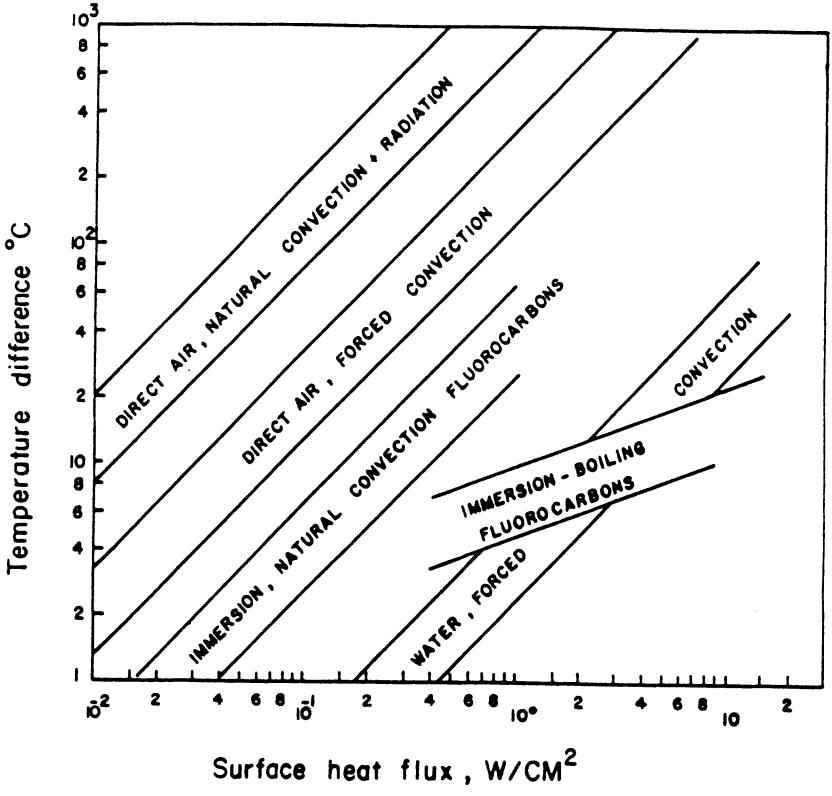

Figure 12-9 Temperature differences attainable as a function of heat flux for various heat transfer modes and various coolant fluids.

effectively cooled with a high level air system (jet impingement, where the air is blown directly onto each heatsink [18]) and may not require exotic cooling. The high end air-based thermal control methods tend to be not as expensive as liquid based ones, yet they yield junction temperatures within the rated limits. This has sparked much interest in the community to explore and expand air cooling limits.

To provide an overall view of cooling techniques practiced in the industry, it is worthwhile to review some of the designs and highlight their salient thermal management features. References [19] and [20] provide an excellent overview of this subject and excerpts from these and other references are used in the forthcoming discussion. Figure 12-5 shows Hitachi's SiC RAM representing low performing systems. The module has six 1 W chips that provide 1 kbit of

memory. The cooling system is an 8 mm high 4 fin simple heatsink with air as the coolant fluid. Using a finned heatsink increases the area in contact with the coolant and thus reduces T_j. The chips contain 77 solder bumps; 52 of them are there purely for thermal reasons. The bumps act as thermal paths, carrying the dissipated heat from the chips to the substrate. The substrate's thermal conductivity is approximately 14 times that of alumina, thus it is very effective for spreading the heat [20]-[21].

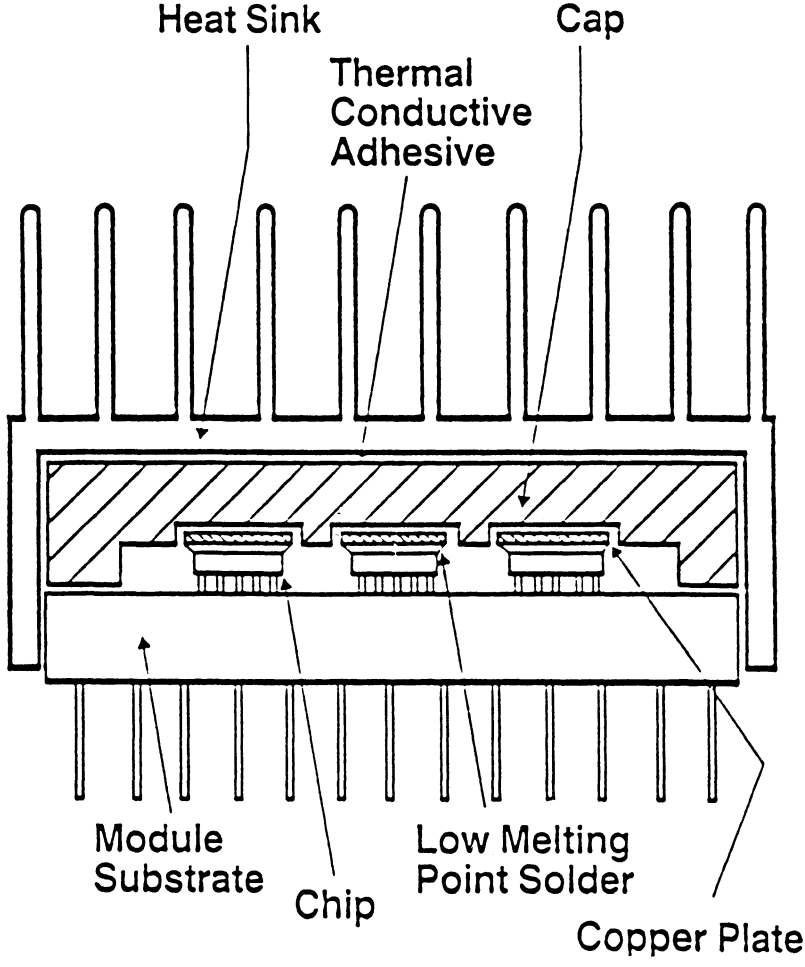

Figure 12-10 Cross sectional view of Mitsubishi air-cooled high thermal conduction module.

Figure 12-10 shows another example of air cooled modules designed by Mitsubishi known as the High Thermal Conduction Module. The component's total power dissipation is 36 W generated uniformly in nine 3K gate ECL chips. The internal thermal paths consist of both the top and bottom of the chips. The chips are placed on bumps to accommodate heat transfer from the bottom side to the substrate. The combination of a copper plate heat spreader placed on the top of the chips (Figure 12-10), and a cap heatsink assembly provide thermal paths from the top side of the chips. The heat generated by the chips is dissipated through the heatsink and convected away by air with a velocity of 6 m/s.

References [19] and [22] discuss the details of heat transfer analysis and subsequent thermal resistances for this module. A note worthy point that relates to our earlier discussion is the variations in heat flow paths and thermal resistance. Reference [22] reports that in the absence of heat transfer from the pins to the air, the central chip conducts 13% of its heat through the solder bumps versus 18% for the peripheral chips. The central chip has 3°C/W and peripheral chips have 2.5°C/W chip to heatsink thermal resistance, resulting in 0.5°C/W difference in thermal resistance for these chips. The combination of internal multiple heat paths (stemming from heat spreading) and the nature of the air flow through the heatsink accounts for this difference. The analysis presented in the noted references has ignored the heat path to the PCB via the module's pins. Inclusion of this heat path may have led to an even higher difference between these chips.

Although the chip power dissipation is uniform, the boundaries of the peripheral chips are different from the central one. The central chip is surrounded, on all sides, by 4 W chips. The peripheral ones have at least one side facing the periphery of the module. Therefore, as a result of thermal coupling through the substrate and the surrounding gas, the central chip is expected to operate at a higher temperature. As a side note, this example typifies the difficulty of reporting a single thermal resistance value for MCMs.

The IBM 4381 module, Figure 12-11 [19], is an example of a high performance air cooled MCM that uses higher level air cooling (impingement). The module is dimensioned 64 mm × 64 mm × 40 mm and contains 36 chip sites. The chips are solder bumped on a multilayered ceramic (MLC) substrate and separated from the ceramic cap by a layer of thermal paste designed by IBM. The power dissipation per module varies from 36 W - 90 W resulting in a typical circuit board power dissipation of 1.3 kW. Reliable system operation requires that the chip temperature not exceed 90°C.

The chips are closely spaced to minimize delay. The power dissipation per chip is high (3.8 W maximum). For this power dissipation level and temperature rating constraint, each chip requires individual thermal management and control.

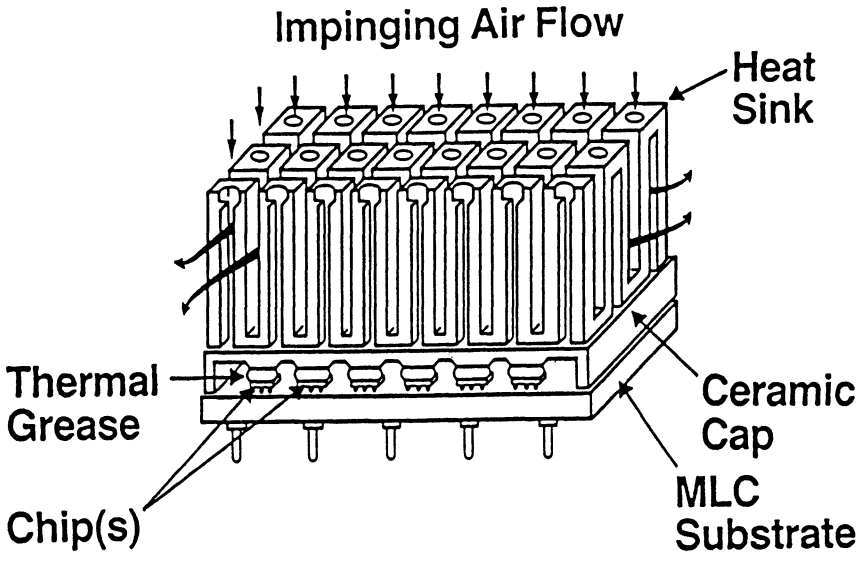

Figure 12-11 Impingement air-cooled MCM used in IBM 4381 processor.

The heatsink, air impingement combination provides an adequate level of thermal control at the chip level to ensure chip temperature below 90°C.

Perhaps the most talked about example of liquid cooling is the IBM Thermal Conduction Module (TCM), used in the 3081 processor, Figure 12-12 [19]. The design requirement specified an 85°C temperature limit for the chip and achieved 69°C. The system performance required that up to nine TCMs be mounted on a single PCB. Since each TCM can dissipate up to 300 W, the total power on the PCB is 2700 W. Stringent chip temperature limit and high system performance was the driving force behind development of this MCM. The thermal design objective was reached by removing heat from the chips as directly as possible and minimizing thermal spreading. The heat dissipated by each chip is conducted via the spring loaded piston in a helium atmosphere to the water-cooled heat exchanger, Figure 12-12.

Fujitsu's FACOM M-780 is a water cooled MCM that departs markedly from the IBM and other similar water cooled modules (Figure 12-13) [19] and [23]. In this design, the thermal control unit consists of bellows and water jets packaged in a closed system. The tip of the bellows is in contact with chip surface through a compliant material to ensure adequate thermal contact.

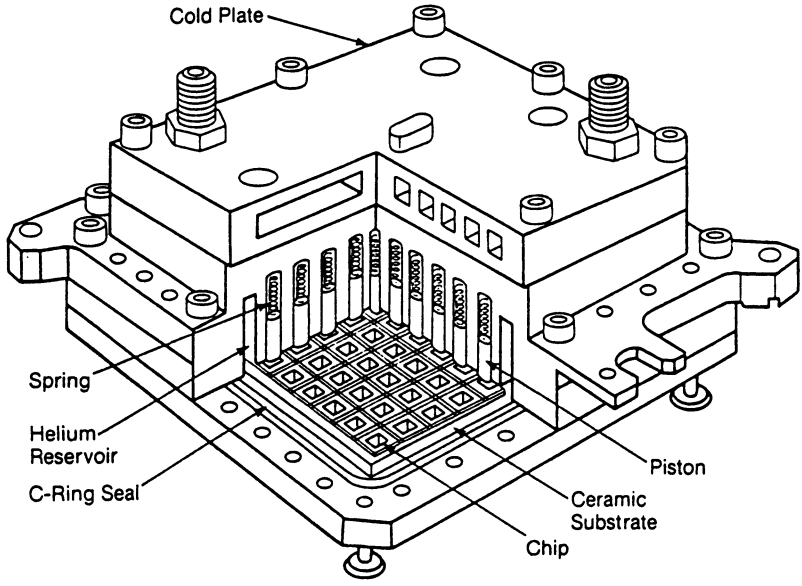

Figure 12-12 IBM thermal conduction module (TCM) with water cooled cold plate.

The FACOM M-780 has 336 single chip modules mounted on both sides of 540 mm x 488 mm PCB. The maximum chip power is 9.5 W and the board dissipates 3,000 W [19]. The cold plate is introduced to the section of the PCB containing the single chip modules. The cold plate is factory assembled and cannot be separated for field repairs.

The next level of cooling that a few computer companies have gravitated toward is liquid immersion. The same criteria drive the selection of the thermal control method: system performance and temperature limit. One of the advantages of immersion cooling is eliminating interface resistances seen in the cold plate thermal control methods. By immersing the MCM or the entire circuit board in a Fluorocarbon (FC-72, FC-77), immersion cooling is attained. The most noted forced immersion cooled system is the CRAY-2 supercomputer.

The immersion cooled portion of the CRAY-2 consists of SCMs mounted on eight PCBs, dissipating a total of 600 - 700 W for a heat density of 0.21 W/cm^2. Though this is within air cooling limits, the large air flow rate required would have been impractical to design. Hence the CRAY-2 uses FC-77 fluorocarbon cooling forced horizontally over the surface of the PCBs with 2.5 cm/sec. velocity.

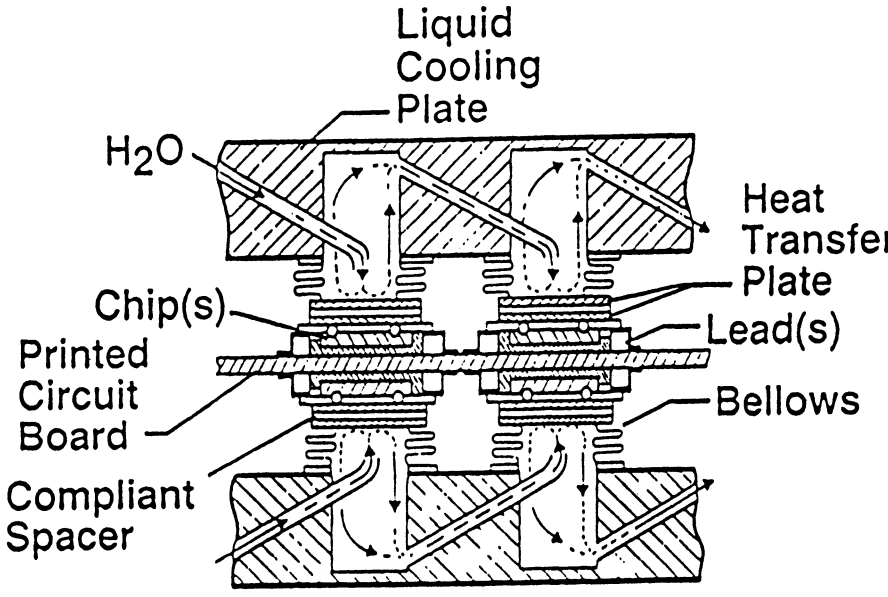

Figure 12-13 Cross sectional view of Fujitsu water-cooled bellows cold plate cooling system.

Before closing this section, it is important to revisit a few points regarding the thermal control techniques used for MCMs. We looked at many examples from the simple application of a flat finned heatsink to immersion cooling of MCMs. Two issues should be evident by now. First, thermal control techniques, beyond the use of a heatsink, are system dependent, or customized. This dependency stems from system level packaging and performance requirements. For example, CRAY-2 designers determined that with 0.21 W/cm^2 heat flux, air cooling was possible. Volumetric air flow requirements, however, were not practical and a liquid cooling system was designed. Therefore, customization of thermal control methods suggests that it is not safe to define general design rules.

The second point is thermal resistance minimization. In all the examples reviewed here, and in many more, the designers have gone through much effort to ensure that the internal and external resistances are as small as possible. An example to be cited is the IBM TCM, in which helium is used in place of air to improve conduction heat transfer through the gas. By replacing air with helium inside the TCM, the internal thermal resistance was reduced from 25°C/W to 8.08°C/W [24]. The importance of this point cannot be emphasized enough. Use of different chip mounting technologies and materials (discussed in other

chapters) for internal chip design can have a major impact in thermal performance of the MCM. Hence, examination of alternatives, with regard to thermal tradeoffs, should be a routine exercise for an MCM designer.

12.4.2 Cooling Methods - Cost Impact of Thermal Management Techniques

The electronics industry can be segmented into four categories:

- Computer
- Military and Space
- Telecommunications
- Consumer products

MCMs have been and will continue to be used in products produced by these industries. The thermal control technique is a function of system application and, therefore, varies between each industry. The high end computer industry has led the way in cooling system design and can afford to use exotic, although not desirable, cooling methods. The military has close tolerance requirements, resulting in unique and system specific cooling systems. Telecommunications tends to gravitate toward lower level cooling methods. Consumer electronics seeks passive cooling techniques because of their application. With the increase in processing speed, all these industries seek higher order cooling methods. This trend will continue until significant changes in the packaging of electronics components are introduced.

The power density, system packaging and junction temperature specifications set the foundation for selection of a cooling method. Additionally, this selection is constrained by manufacturability and cost. The hierarchy of the cooling methods are:

- Air in natural convection, with or without heatsink
- Air in forced convection, with or without heatsink and other flow enhancement methods
- Air jet impingement
- Radiation
- Liquid cooling:
  - Natural convection
  - Forced convection
  - Jet impingement
  - Immersion
- Immersion by boiling
- Cryogenics

The above list is ordered in the ease of implementation, but in a reverse order of cooling capacity (see Figure 12-9). Natural convection, with air as the working fluid, is most desirable since no cooling system design is required. This does not imply that thermal analysis is not needed. It suggests that no cooling system (fluid movers, heat exchanger etc.) is required for thermal management of the system. Forced air convection is the next most desirable mode. Use of heatsinks with air cooling can further increase the heat removal capacity of forced air convection. Some industry segments (telecommunications and consumer products) tend to shy away from fans (fluid movers) because of reliability issues and noise. Fluid movers are unavoidable if power dissipation and temperature specifications are at such a level that fans are required.

From jet impingement to higher cooling methods, heat removal capability goes up significantly. But implementation in a system from cost, manufacturing and user impression becomes complex. For example, jet impingement requires a compressor and placement of jet nozzles throughout the circuit board, creating a reliability and physical design dilemma. Additionally, there is a whistling noise as the air expands and leaves the nozzle. Imagine your PC or workstation whistling continuously - fan noise is uncomfortable enough.

Liquid cooling is a very attractive proposition for high powered components and systems. But implementation, cost, maintenance and use is difficult and liquid cooling only is seen in a small number of systems. These difficulties have invigorated the search for air cooling methods for high powered systems. Reference [25] gives an excellent review of liquid cooling in microelectronics components. Reference [26] provides a general review of thermal management of electronic equipment.

12.4.3 Parameters Impacting MCM Thermal Performance

We have discussed thermal phenomena at length in MCM and SCMs. The intent of this section is not to revisit them, but to give a list of parameters impacting junction temperature.

External

External parameters are outside the package (module). These include the following:

- Environment where the system resides
- Method of thermal control
- Coolant fluid temperature, velocity and flow regime - laminar, transition or turbulent

- Neighboring circuit board power, locational proximity and surface emissivity
- Circuit board or substrate thermal conductivity
- Neighboring component's power dissipation, size and locational proximity to the MCM
- The size of the gap between the component and the circuit board. In the case of military applications, this gap is typically filled with a conductive material
- Method of lead attachment to the circuit board

Internal

The internal parameters are specific to inside the package (module). These include the following:

- Dimensions - exact dimensions of component including leads, wire bond, chips, via type and density, etc.
- Method of chip attachment to the substrate (bonded, epoxy) with material dimension and thermal properties
- Method of lead attachment to the substrate
- Molding compound material property and dimension
- Power dissipation of each chip on the substrate
- Metallization - material and dimensions
- Material and dimensions of the substrate
- Property and the dimension of interface material used between the chip and the cooling tower

The response to the above list provides the information necessary for thermal analysis or experimental simulation of MCMs. In the case of analytical simulation, the impact of various parameters can easily be highlighted. For example, one can calculate the junction temperature as a function of chip

attachment method and molding compound thermal conductivity. Various what-if cases allow for selection of more appropriate material to help reduce junction temperature. Similarly, if a system level analytical model is developed, the what-if cases provide the designer with insight into alternative options in system design and cooling method implementation.

12.5 TOOLS FOR THERMAL DESIGN

12.5.1 Overview of Design Analysis Tools

The tools used for thermal analysis do not discriminate between SCMs and MCMs. The tools are of generic nature and apply to any thermal problem. In this section, the use and the domain of applicability of these tools are presented. The discussion is followed by two examples to demonstrate the application and utility of these tools.

The principal tools in thermal design are categorized (Table 12-3) into three areas:

- Integral (Analytical)
- Numerical (Computational)
- Experimental

The integral method is the first and most essential method for forming the solution. Because of its analytical nature, it typically is referred to as a first order solution. The numerical solution is a second order solution with limited domain of application, mainly component and circuit board. The experimental method is the highest order solution utilized when the other two methods are not suitable for the specific problem. The important point is that the solution techniques are interdependent. The integral method often forms the foundation or the starting point of the numerical solution procedure.

The requirements of the two higher order solutions show an interdependency with the integral method. The numerical method needs boundary or initial conditions to initiate a solution. The conditions are obtained either by direct measurement or by application of an integral method. When experimentation is deemed necessary, the physical trend and pertinent parameters should be identified before measurement is done. These and the premises for experimentation are obtained by forming the solution by the integral method. Since analytical methods play such a pivotal role in the analysis and design of electronic components and systems, it is worth the time to further develop this base.

Table 12-3 Comparison and Application of Solution Techniques in Electronics Cooling.

	INTEGRAL	NUMERICAL	EXPERIMENTAL
Order of Solution	First	Second	Third
Accuracy	Moderate	High	Best
Effort	Small	Small-High	Difficult
Domain of Application	All problem domains	Component and circuit pack	All program domains
Cost	Small	Medium-High	Very high
Expertise	Introductory	Introductory-Specialized	Specialized
Tools	PC	High speed computers and software	Laboratory facility

12.5.2 Analysis Tools

Thermal problems are a synthesis of heat transfer, fluid mechanics and thermodynamics. The analytical tools are based on the conservation laws and equation of state. These laws are:

- Conservation of mass
- Conservation of momentum
- Conservation of energy
- Equation of state

Thermal problems can be formulated by the general application of these laws to a finite or infinitesimal region of the problem domain, called a control volume (CV). To accurately formulate the problem, it is necessary to understand the utility and the domain of the application of each law. Therefore, we start by defining each equation and briefly discuss its utility when applied to a thermal design problem.

Conservation of mass simply states that the mass in a given thermal process is conserved. The equation is formed from the cross sectional area, velocity and

density in two regions of the problem. The equation can be expressed as follows:

Rate of mass flow into CV - Rate of mass flow out of CV =
Rate of accumulation of mass inside the CV, (12-9)

where CV denotes control volume. The conservation of mass is a very useful tool, and its utility or application is often overlooked in thermal analysis problems.

The flow field in thermal problems is resolved by the application of the conservation of momentum (Newton's second law). This equation relates the forces that govern the flow to the rate of the fluid's change of momentum. Once applied to a fluid enclosed by a CV, the equation can be expressed by the following:

Rate of momentum accumulating in the CV
= Rate of momentum into the CV
- Rate of momentum out of the CV
+ Sum of the forces acting on the CV. (12-10)

Application of the conservation of momentum yields information regarding velocity and pressure distribution within the system. The equation is nonlinear in nature and inherently unstable. Hence, its solution can become quite involved and may require higher order solution techniques such as numerical computation or experimentation.

The last equation in the conservation laws is the energy equation. Conservation of energy, or the first law of thermodynamics, simply states that within a thermal process, the energy is conserved. In general, the energy can have many facets, from viscous heating to heat released by a chemical reaction. Therefore, to include all forms of energy, let us state the conservation of energy as applied to a CV by the following:

Rate of energy entering the CV
+ Rate of generation of energy in CV
- Rate of energy leaving CV
= Rate of accumulation of energy in CV. (12-11)

The solution of the conservation of energy yields the temperature and heat transfer of the thermal process. If fluid motion also is associated with the problem, then the conservation of energy and momentum equations are coupled. The energy equation, in this case, includes the velocity terms. Their magnitude

must be known before a solution can be obtained. Furthermore, if the fluid properties are temperature dependent the conservation of energy and momentum equations also must be solved simultaneously. Otherwise, each equation can be solved independently, significantly simplifying the solution process.

All of these laws contain thermodynamic properties describing the fluid at a given state. If these properties become non-constant and vary from state to state (for example at different temperatures), their values must be known. Therefore, an equation of state, such as the ideal gas law, is used for determination of the properties.

12.5.3 Solution Procedure

Complexities in electronic cooling problems stem from the intercoupling that occurs among various elements of thermal transport. To develop a solution, it is necessary to establish a methodical approach or engineering bookkeeping. Hence, the general approach to solution of thermal problems is first outlined. We then focus more specifically on the electronic cooling.

The general approach for problem solving, independent of electronic cooling, is shown in a flow chart in Figure 12-14 [27]. The approach consists of four parts: problem preparation, solution preparation, solution procedure and solution verification. The approach provides a method for accurately making assumptions and keeping track of the solution process. Additionally, the critical issue of what method is necessary for the solution procedure can be addressed early on. This also provides a process for generating the necessary information to select a solution technique, for example, boundary conditions for numerical simulation. The general procedure for electronic cooling consists of fluid flow and heat transfer analyses, with the outcome focused on the reliability predictions of the components (Figure 12-15).

12.5.4 Analytical Modeling - Integral Method

The thermal model of an electronic system is developed by applying Equations 12-9 through 12-11 to a CV embracing the region of interest. The model yields an analytical expression describing the desired variable in terms of other parameters governing the problem. The process of modeling, however, is not confined to a specific problem domain. The procedure is best described by the following example:

Example 1: Governing equations for a single component residing in a circuit board channel

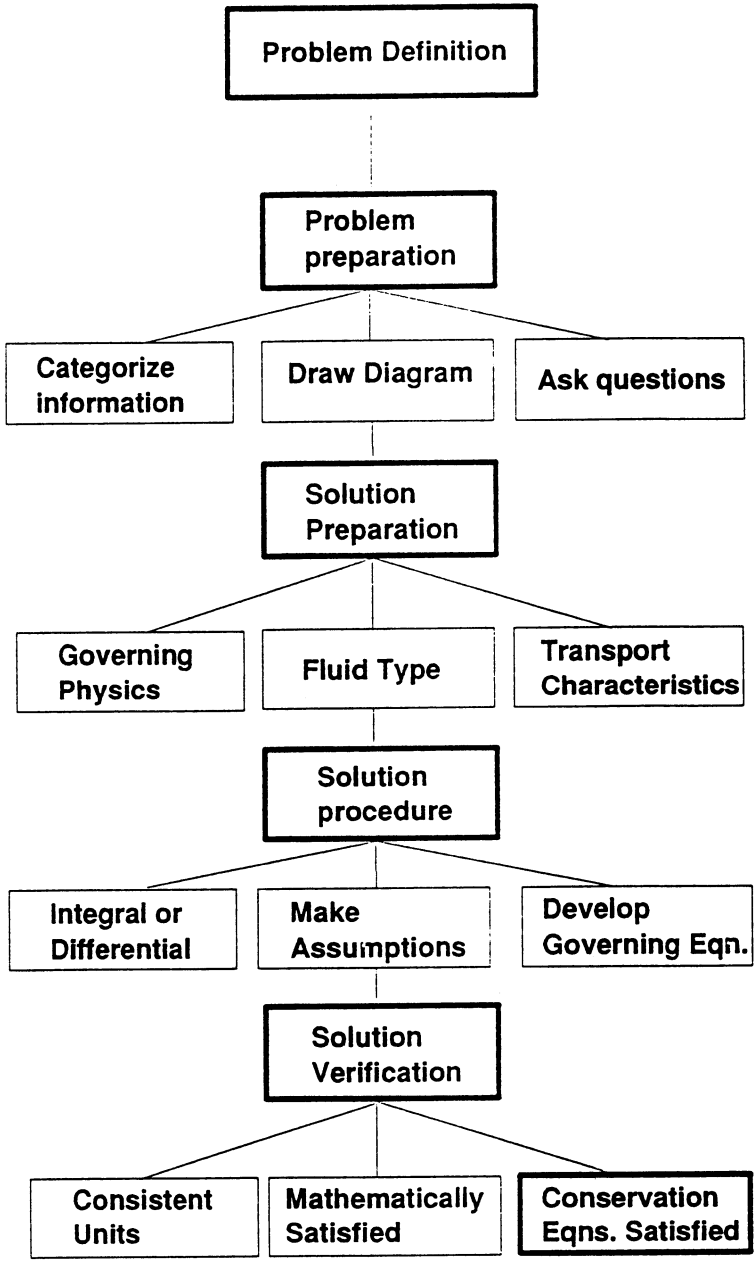

Figure 12-14 Flow chart for general problem solving approach.

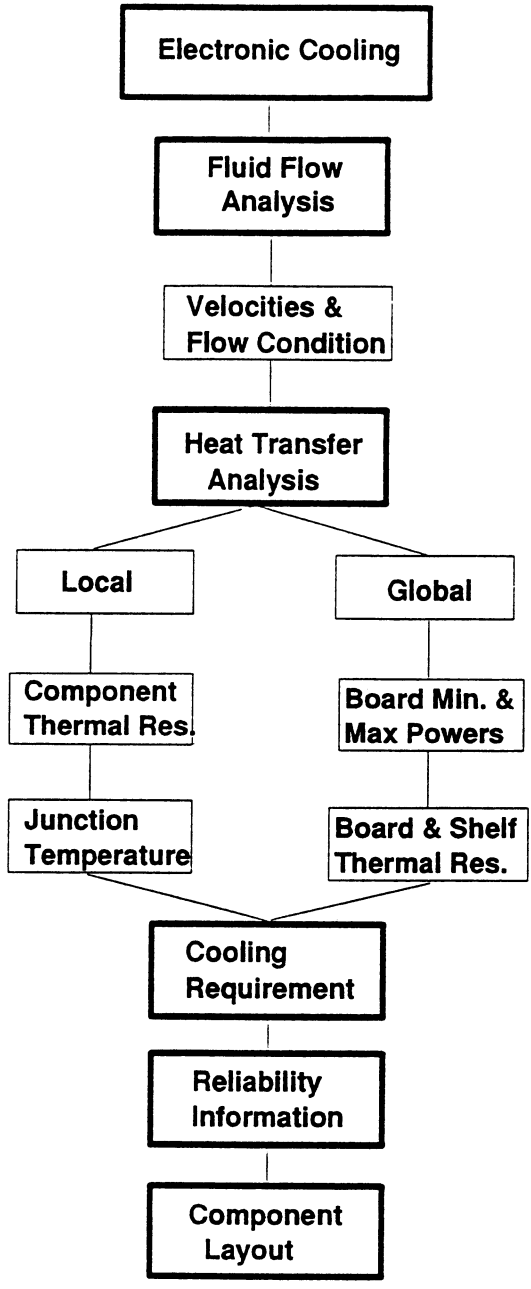

Figure 12-15 Flow chart for problem solving approach in electronic cooling.

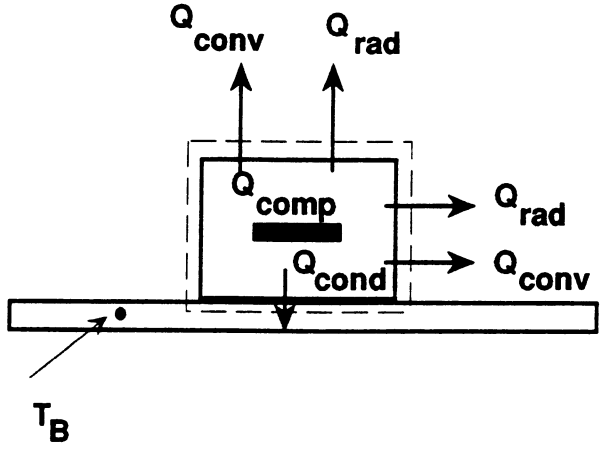

Figure 12-16 Single SCM with control volume around component.

Consider an air cooled SCM residing on a board. The objective is to obtain the junction temperature of the component as a function of SCM board and coolant parameters. This will be done by constructing the equations applicable to each CV.

The following assumptions are made within the CV:

- The problem is at steady state.
- The fluid behaves like an ideal gas, and is incompressible.
- In the designated CVs, the fluid temperature is uniform. Note that since the first order solution method is selected, the temperatures and velocity are averaged over an area of interest, for example, an area perpendicular to the direction of the fluid flow.
- The board temperature is uniform over the component's footprint. This area is equal to twice the planar area of the component.
- The convective heat transfer coefficient is also uniform within the CV.

Component

The information generally available for the component is the geometry and an estimate of the power dissipation. However, on may occasions, the power dissipation is not known exactly. This is not an obstacle in the analytical model since the power dissipation can be treated as a variable. Applying Equation 12-11 to the CV shown in Figure 12-16 we get:

$$Q_{comp} = Q_{cond} + (Q_{conv} + Q_{rad})_{top} + (Q_{conv} + Q_{rad})_{bottom} \cdot \tag{12-12}$$

Replacing the heat flows, Q, in Equations 12-2 through 12-4 with their respective temperature definitions and linearizing the radiation heat transfer terms, we find:

$$Q_{comp} = \frac{kA}{L}(T_1 - T_B) + hA_{top}(T_1 - T_f) + 4\sigma\varepsilon F_{1,top} A_{top} T_m^3 (T_1 - T_{board,\,top}) + hA_{side}(T_1 - T_f) + 4\sigma\varepsilon F_{1,side} A_{side} T_m^3 (T_1 - T_{board,\,top})$$

(12-13)

where T_1 is the component case temperature and T_B is the board temperature in the footprint of the component. A, L and k are the lead total cross section, length and thermal conductivity, respectively. A_{top} is the area of the top of the component, and A_{side} is the total area of the component sides. T_f is the mean fluid temperature $T_f = (T_{f,i} + T_{f,o})/2$, where $T_{f,i}$ and $T_{f,o}$ are the fluid inlet and outlet temperatures, respectively. $T_{board,\,top}$ is the temperature of the board top away from the component. T_m is the mean board temperature.

Board

The energy balance for the board with heat flows shown on Figure 12-17 becomes:

$$Q_{cond} = Q_{cond,B} + (Q_{rad} + Q_{conv})_{top} + (Q_{rad} + Q_{conv})_{bottom} \cdot \quad (12\text{-}14)$$

Substituting the temperatures from Equation 12-2 through 12-4 and linearizing the radiation heat transfer:

$$\frac{kA}{L}(T_1 - T_B) = k_B \frac{t_B W_B}{L_B}(T_B - T_{amb}) + 4\sigma\varepsilon A_{f,p} T_m^3 (T_B - T_{board,\,top}) + h(A_{f,p} - A_{top})(T_B - T_f) + 4\sigma\varepsilon A_{f,p} T_m^3 (T_B - T_{board,\,bottom}) + hA_{f,p}(T_B - T_{f,\,bottom})$$

(12-15)

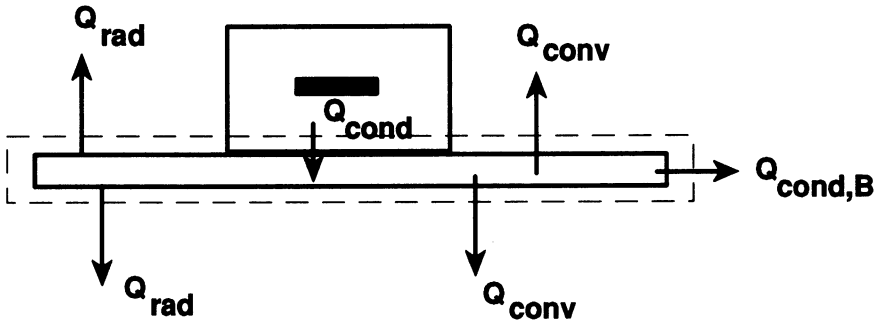

Figure 12-17 Single SCM with control volume around the circuit board.

where k_B is the thermal conductivity of the board and t_B, W_B and L_B are board thickness, width and the distance from the component to the thermal connection with the frame, respectively. T_{amb} is the ambient temperature of the frame. $A_{f,p} = 2A_{top}$ is the area of the board assumed to be at temperature T_B (component footprint), $(A_{f,p} - A_{top})$ is the area not covered by the component, $T_{board,\ bottom}$ is the temperature of the board away from the component and $T_{f,bottom}$ is the mean fluid temperature on the bottom of the board. The radiation view factor is ignored since it approximately is equal to 1.

Fluid

Conservation of energy for the fluid (Figure 12-18) yields the following:

$$\dot{m} C_P (T_{f,o} - T_{f,i}) = h(A_{top} + A_{side})(T_1 - T_f) + h(A_{f,p} - A_{top})(T_B - T_f), \quad (12\text{-}16)$$

where $\dot{m}$ is the mass flow rate of the coolant and C_P is its specific heat. Conservation of mass becomes

$$V_1 A_1 = V_2 A_2, \quad (12\text{-}17)$$

where A_1 is the inlet temperature, A_2 is the outlet temperature, V_1 is the inlet velocity and V_2 is the outlet velocity. If V_1 is given, then V_2 and $V = (V_1 + V_2)/2$ can be calculated and h determined.

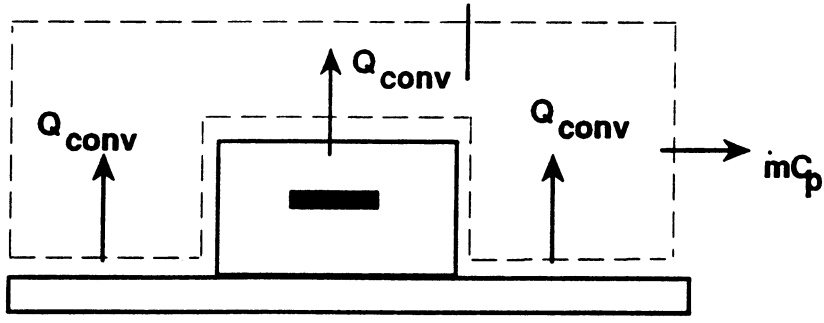

Figure 12-18 Single SCM with control volume around fluid.

Component Interior

The equation applicable to the CV inside the component body is

$$Q_{comp} = \frac{k_{eff} A_{eff}}{L_{eff}} (T_j - T_1) \qquad (12\text{-}18)$$

where k_{eff} is the effective conductivity, A_{eff} is the effective area and L_{eff} the effective length of the SCM internal path.

Thus, if Q_{comp}, $T_{f,i}$, V_1, $T_{f,bottom}$, $\dot{m}$, $T_{board,top}$, $T_{board,bottom}$ and T_{amb} are given, then the designer can solve for T_j, T_1, T_B and $T_{f,o}$.

12.5.5 Computer Based Tools - Numerical Method

The numerical approach to the solution of thermal design problems is based on the differential form of the conservation laws [28]. The differential model yields the distribution of the desired parameters: velocity and temperature. The solution can be obtained by either finite difference method (FDM), finite element method (FEM) or spectral method (SM). Since the differential model tends to be highly nonlinear and inherently unstable, the numerical approach to these problems must be done with much care.

The numerical solution can be categorized into three areas: solid, fluid or a combination of the two. For solids (component or board), the modeler must have a good insight for fluid to solid (heat transfer coefficient) and solid to solid (board where component resides) boundary conditions. The fluid case yields velocity and pressure distributions in the flow field. If the boundary conditions are not available at the interface of the solid and fluid, the solution becomes a simultaneous simulation of the two.

The accuracy of the numerical simulation is a function of boundary or initial conditions. In addition, the solution is dependent on the mesh or grid size, resulting in an iterative process to verify the dependency of the solution on the numerical mesh. Since the differential models are developed on the basis of infinitesimal CVs, their application is suitable indeed for localized analyses. This implies that higher accuracy with reasonable effort can be obtained when the domain of analysis is comparatively small, such as with components and circuit boards. Therefore, to obtain the highest accuracy, some rules of thumb can be developed when using numerical techniques.

1. Do not model the entire system or the shelf unless you are seeking some very qualitative data.

2. Increase the mesh or gird density at the points of concentration. These points can be categorized as areas such as heat sources and points of separation in the flow field. Many tools do this automatically.

3. Avoid idealization of the boundary conditions. That is, if the boundary condition is a gradient, do not assume it is uniform unless there is a sound engineering reason for making such as assumption.

4. Do not automatically assume that the answer obtained from the numerical simulation is correct. The answer from numerical simulation must always be verified via integral methods or experimentation.

5. Do not to analyze data by just looking at the color graphics output. In most electronic cooling problems, the geometry is small. Therefore, spatial variation can be significant. Color graphics outputs, even with 16 colors, do not have the necessary resolution to highlight the spatial variations. The color graphics, therefore, can be very misleading and should be used only for obtaining insight into the trend of the problem. The XY plots or numeric outputs are the best way to see the variations in the area of interest.

There are three numerical methods for obtaining the solution to thermal design problems, FDM, FEM and SM. They might be used in two dimensional form (usually FDM) for obtaining quick low order solutions or in three dimensional form for high order solutions. The two dimensional forms typically use analytic expressions (such as the ones above) to handle the effects of the third dimension. There are many good references that discuss these methods in detail [29]. It is beyond the scope of this text to discuss these methods in detail. Instead this section attempts to answer a common question that confronts a practitioner - "which method is most suited for my application?"

Let us start by first giving a general description of each method. The finite difference method uses a Taylor series expansion to discretize the entire problem domain and converts the differential equation into a series of analytic expressions. The finite element method uses approximate functions as a local solution and also converts the differential equation into a set of analytical expressions. The spectral method can be categorized as an extension of finite element methods in which the spatial frequency domain also is utilized. From this simple description of each method, one can conclude that the FDM is perhaps the most simple to model and the SM the most complex.

To determine which method is most suited for thermal design of MCMs, or other electronic components, let us divide the domain of application into the fluid and solid solvers. The FDM and FEM are most suited as solid solvers, that is, component and board. The suitability for solid modeling is basically a function of geometry. The FDM, which converges much faster than FEM, is ideal for regular geometry, such as rectangular. Irregular geometries also are handled effectively if the FDM utilizes the Body Fitted (or fixed) Coordinate (BFC) methodology. The FEM addresses the irregular geometry much more accurately. The SM method appears to be an over description, with a very slow rate of convergence, as a solid modeler. Therefore, FDM and FEM are the techniques of choice for obtaining temperature and stress distribution in a solid.

Four points are of interest in selecting the fluid solver: geometry, condition of the boundaries, transient/separated flows, and rate of convergence. The FDM method is most suited for regular geometry and the FEM for irregular geometry. If the boundary conditions are non-stationary, FEM is more suitable for modeling than FDM. If the flow field contains separation or transience, that is, flow past a component, SM appears to be most capable to address the physics associated with the flow regime. One should add that all these methods are capable of handling transient problems. But much more effort may be required to obtain a solution with one method versus another. The last point of concern is the rate of convergence. Computational fluid dynamics tend to be very computer time intensive. The author's comparison of all three methods applied to the same problem on similar computers has resulted in the following order (where least

computer time is first): FDM, FEM, and SM. It is important to point out that the numerical methods are not black boxes. These are tools that must fit the problem. Therefore, the selection of a technique should be based on the type of problem and the capabilities of the code.

To see a typical application of the computational fluid dynamics method, the following example is presented.

Example 2: Temperature and air flow distribution in a horizontal channel containing an MCM.

Consider a single MCM placed inside a horizontal channel. The total power dissipation of the module is 5 W. The center chip has 4 W and the side chips have 0.5 W of power dissipation each. The MCM also is surrounded by four SCMs each dissipating 0.5 W. The components are placed on a four layer board and cooled with natural convection. We have two objectives in this case. First, we would like to see how the module functions in natural convection - does it meet the 125°C temperature specification? Second, what if a flow of 200 ft/min (forced convection) is introduced from left to right - what is T_j in this situation?

A commercially available computational fluid dynamics (CFD) software, FLOTHERM™ [30], was used for the simulation. Figures 12-19 and 12-20 show the results of such analyses. These results show the temperature distribution in the circuit board and the velocity distribution in vectorial form, with arrow length proportional to air flow velocity. The results show that natural convection is not adequate for cooling of this MCM since junction temperature is very close to the design limit, Figure 12-19. The margin is not sufficient to approve this design for natural convection cooling.

In the case of forced convection (Figure 12-20), although the fins are misaligned with the flow direction, we see that the junction temperature is considerably lower. The temperature would be lower yet if the fins were aligned with the air flow. This maximizes the average flow rate. Fin misalignment is a common error (and done here solely for illustrative purposes). The temperature with forced air cooling provides a sufficient margin in T_j making it the recommended solution. Note the thermal coupling evident in Figure 12-20. The component downstream of the MCM resides in a slower moving fluid at an elevated temperature. Thus, its T_j is hotter.

12.5.6 Experimentation - Why, When and How

The highest order solution in a thermal design process is experimentation. As the ranking of the solution techniques show, it is last on the list. Although we tend to believe experimental data over analytical ones, and rightfully so, it is

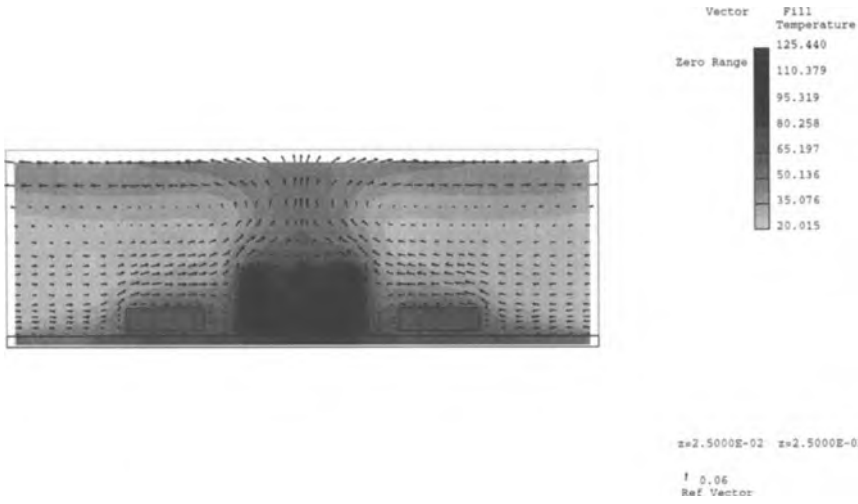

Figure 12-19 Flow distribution around a 5W MCM cooled by natural convection.

normally avoided because of its cost and duration. Therefore, this section highlights when an experiment is necessary and what parameters should be measured. It is certainly not the scope of this section to show methodologies for fluid flow and heat transfer measurements. For a detailed description of these types of measurements, references [31] and [32] are recommended.

As discussed earlier, because of thermal intercoupling that exists inside the electronic circuit boards, analytical and numerical predictions are difficult. The level or magnitude of intercoupling indicates when an experiment should be conducted. When the boundary conditions are not clearly defined, an experiment closely simulating the actual case must be conducted. Another case when experimentation is deemed necessary is when the problem is not phenomenologically understood. These are cases when the problem at hand is new and has not been investigated before. These cases typically are encountered in new technologies or processes. A good experiment is one that closely resembles the actual process. An experimenter should understand the details of the problem and design the experiment accordingly. Often because of the nature of the problem, such as chip size or limitation in the test facility, scaling is required [33]. In this case, the dimensionless fluid dynamics numbers such as Reynolds (Re), Grashof (Gr) and Prandtl (Pr) [11], must be identical for the model and actual process. Another important aspect of a good experimental design is the uncertainty analysis [34].

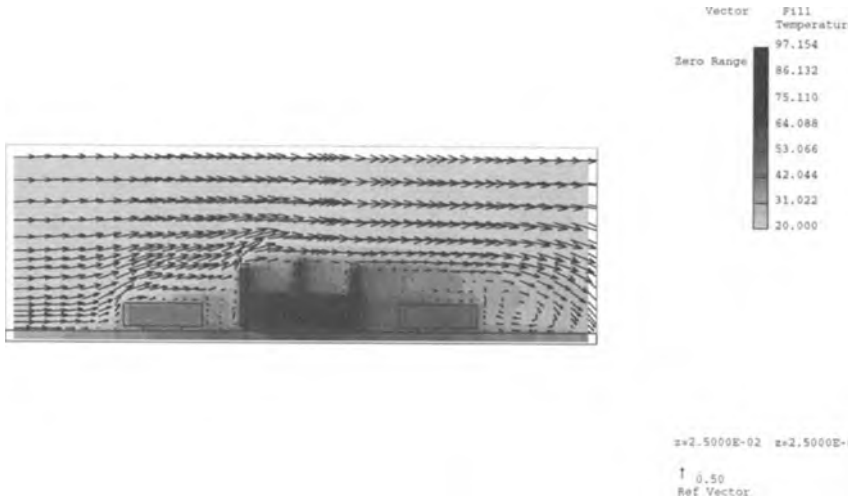

Figure 12-20 Flow distribution around a 5W MCM cooled by moderate forced convection.

The selection of parameters that best describe a process is always a difficult task. This difficulty is even more profound in electronic cooling since the problem can be a very dynamic one. For example, if the case temperature is to be measured, where should the temperature sensor be placed? What follows are some general guidelines for measurement in electronic cooling.

Let us start by stating that pressure (velocity) and temperature are the two parameters that readily can be measured in a laboratory. The parameters and points of measurements should be the ones that best describe the process. Two parameters, component temperature and temperature rise in the system, can be quantified as the most important for temperature measurement. The component temperatures are junction, case on top of the hottest chip, and case at the edge. The two case temperatures provide the limits, minimum and maximum, reached by the case under test conditions.

Use of thermal resistance in MCMs should be avoided. If one insists on reporting Θ, chip junction to ambient should be used. Here, the ambient is defined as the reference temperature or the fluid temperature in the vicinity of the component. For conventionally cooled MCMs (low performance systems), the latter is calculated by an energy balance performed on the components, preceding the component of interest. For liquid cooled MCMs, fluid temperature at the source is used as the reference.

Temperature rise in the system is defined as the increase of cooling fluid's temperature over an established reference. The reference temperature in the case of air cooled systems is the environment's temperature, and for liquid cooled systems, it is the supply temperature.

Likewise, fluid velocity and pressure drop would have to be measured to accurately quantify the component thermal performance and the system cooling mechanism. Fluid velocity and temperature measurements at the local level are labor intensive and do not yield much practical information. For air cooled systems, the measurements are at the boundaries, that is, inlet and outlet.

12.6 SUMMARY

In this chapter a description of thermal design and management of MCMs is presented. The chapter describes thermal phenomena in MCMs, thermal control methods and different methodologies for analysis of MCMs. There are some salient points that merit reiteration. These are:

- Principles of thermal management and design have not changed for MCMs from what was practiced with SCMs.

- The challenge or complexity in thermal control of MCMs versus SCMs stems from multiple chips on the substrate. Air cooling of MCMs is the most desirable method of cooling. This is based on the ease of use, abundance of air and lesser manufacturing and maintenance costs.

- A logical progression for thermal control of MCMs is to first exhaust all air cooling options before considering any form of liquid cooling.

- Liquid cooling tends to be highly customized for most MCM designs. Prepackaged cooling systems with liquid cooling capability typically are not found.

MCM thermal performance is increased by minimizing internal and external resistances. The internal resistance can be reduced by selecting:

- Materials with higher thermal conductivity
- Minimizing thermal spreading near the multiple chips and
- Considering alternate chip mounting techniques.

The external resistance is reduced for:

- Air cooled MCMs by delivering air directly to the MCM and minimizing flow resistances.
- Liquid cooled MCMs (excluding immersion) by reducing interfacial resistances between the cold plate and the chip.

An important point to note for low performance MCMs is the impact of board layout on thermal response. Judicious placement of components on the board significantly improves component thermal response. Often, a combination of thermally appropriate board layout and minimization of internal resistance can avoid higher order cooling methods.

MCM thermal analysis for either characterization or verification are done by three methods: integral, computational and experimental. Each method's domain of utility is specified. An important point is that the integral method forms the foundation for any analysis initiative. Computational and experimental studies, although very useful, are boundary condition dependent. A noteworthy point is the utility of integral and computational methods for parametric studies. The what-if scenarios performed early in the design cycle play a pivotal role in the robust design of MCMs. Additionally, the design duration and its associated development cost can be reduced if these analyses are performed at the onset of system conceptualization.

Acknowledgments

The author would like to thank the Editors for their very significant technical and organizational help in defining the structure of the chapter. He also would like to thank the reviewers for their concrete and substantial suggestions.

References

1 D. J. Klinger, Y. Nakada, M. A. Menendez, *AT&T Reliability Manual*, New York: Van Nostrand Reinhold, 1990.

2 M. Bonnifiat, M. Cader, "Thermal Simulation for Electronic Components using Finite Elements and Nodal Networks," *ASME HTD*, vol. 57, pp. 183-188, 1986.

3 H. Hardisty, J. Abboud, "Thermal Analysis of a Dual-in-Line Package Using Finite Element Method," *IEE Proc.*, (UK), vol. 134, Part 1, pp. 23-31, 1987.

4 M. Aghazadeh, D. Mallik, "Thermal Characteristics of Single and Multi-Layer High Performance PQFP Packages," *Proc. 6th Annual Semi. Thermal and Temperature Measurement Symp.*, (Phoenix AZ), pp. 33-39, 1990.

5 M. A. Zimmerman, K. Azar, C. D. Mandrone, "Thermal Characteristics of Molded

Multi-Chip Modules," *ASME Winter Annual Mtg.*, (Dallas TX), 90-WA/EEP-21, 1990

6 W. M. Kays, M. E. Crawford, "Convective Heat and Mass Transfer," 2nd Edition, New York: McGraw Hill, 1980, pp. 139.

7 R. A. Wirtz, W. McAuliffe, "Experimental Modeling of Convective Downstream from an Electronic Package Row," *ASME J. of Electr. Packaging*, vol. 11, pp. 207-212, 1989.

8 R. A. Wirtz, P. Dykshoorn, "Heat Transfer from Arrays of Flat Packs in a Channel Flow," *Proc. Inst. of Electr. Packaging Symp.*, pp. 318-326, 1984.

9 E. M. Sparrow, J. E. Neithammer, A. Chaboki, "Heat Transfer and Pressure Drop Characteristics of Arrays of Tectangular Modules Encountered in Electronic Equipment," *Internat. J. of Heat and Mass Transfer*, vol. 25, no. 7, pp. 961-973, 1982.

10 S. Sridhar, *et al.*, "Heat Transfer Behavior Including Thermal Wake Effects in Forced Air Cooling of Rectangular Blocks," *ASME HTD*, WAM, (Dallas TX), vol. 153, pp. 15-26, 1990.

11 F. P. Incropera, D. P. Dewitt, *Introduction to Heat Transfer*, New York: Wiley and Sons, 1985.

12 F. M. White, *Heat Transfer*, New York: Addison-Wesley Publishing Company, 1984.

13 V. P. Manno, K. Azar, "The Effect of Neighboring Components on Thermal Performance of Air-Cooled Circuit Packs," *ASME J. of Electr. Packaging*, vol. 113, pp. 50-57, March 1992.

14 K. Okutani, *et al.*, "Packaging Design of SiC Ceramic Multi-Chip RAM Module," *Proc. of Internat. Electr. Packaging Soc.*, pp. 299-304, 1984.

15 K. Azar, C. D. Mandrone, "Effect of Pin Fin Density on Thermal Performance of Un-Shrouded Pin Fin Heat Sinks," *Proc. of Electro/92 Conf.*, (Boston MA), 1992.

16 A. D. Kraus, A. Bar-Cohen, *Thermal Analysis and Control of Electronic Equipment*, Hemisphere Publishing, 1983, pp. 593-596.

17 R. Hannemann, *Thermal Control for Mini and Microcomputers: The Limits of Air Cooling," Heat Transfer in Electronic and Microelectronic Equipment*, Hemisphere Publishing, 1990, pp. 61-84.

18 K. Azar, J. R. Benson, V. P. Manno, "An Experimental Investigation of Microjet Impingement Cooling," *27th Nat. Heat Transfer Conf.*, July 1991.

19 A. Bar-Cohen, "Thermal Management of Air and Liquid-Cooled Multichip Modules," *IEEE Trans. CHMT*, vol. CHMT-10, no. 2, pp. 159-175, 1987.

20 R. C. Chu, R. E. Simons, "Recent Developments in Thermal Technology For Electronics Packaging," *Electr. Packaging Forum*, vol 2, New York: Van Nostrand Reinhold, 1991.

21 K. Otsuka, *et al.*, "Considerations of VLSI Chip Interconnection Methods," *IEEE Computer Society Spring Workshop*, (Palm Desert CA), 1985.

22 M. Kohara, *et al.*, "High Thermal Conduction Package For Flip Chip Devices," *IEEE Trans. of Components, Hybrids, Manuf. Techn.*, vol. CHMT-6, pp. 267-271, 1983.

23 H. Yamaoto, Y. Udagawa, M. Suzuki, "Cooling System for FACOM M-780 Large Scale Computer," *Proc. of Internat. Symp. on Cooling Techn. for Electr. Equipment*, (Hawaii), pp. 110-125, 1987.

24 Personal communication with R.E. Simons of IBM.
25 A. E. Bergles, A. Bar-Cohen "Direct Liquid Cooling of Microelectronic Components," *Advances in Thermal Modeling of Electr. Components and System*, ASME Press, vol. 2, pp. 233-342, 1990.
26 W. Nakayama, "Thermal Management of Electronic Equipment: A Review of Technology and Research Topics," *Appl. Mechanics Review*, vol. 39, no 12, 1986.
27 K. Azar, "Problem Solving Techniques in Electronic Cooling," *Proc. 5th Annual Semi. Thermal and Temperature Measurement Symp.*, (San Diego CA), pp. 37-44, 1989.
28 B. R. Bird, W. E. Stewart, E. N. Lightfoot, "Transport Phenomenon," New York: John Wiley and Sons, 1976, pp. 83-85.
29 C. A. J. Fletcher, "Computational Techniques for Fluid Dynamics," Springer Verlag, 1988.
30 FLOTHERM, Flomerics Corp., Westborough, MA 01581.
31 R. J. Goldstein, "Fluid Mechanics Measurements," Hemisphere Publishing, 1984.
32 E. R. Eckert, R. J. Goldstein, *Measurements in Heat Transfer*, Hemisphere Publishing, 2nd ed., 1976.
33 G. Murphy, *Similitude in Engineering*, Ronald Press Co. 1950.
34 H. W. Coleman, W. G. Steele, "Experimentation and Uncertainty Analysis for Engineers," New York: John Wiley and Sons, 1989.

13

ELECTRICAL TESTING OF MULTICHIP MODULES

Thomas C. Russell and Yenting Wen

13.1 INTRODUCTION

Electrical testing is used throughout the multichip module fabrication process to verify the quality of each processing step and component which goes into the module. Once module assembly is complete, a final test is performed to make sure that the module functions to its specifications. It is desirable to locate any faults as early in the manufacturing process as possible since this results in a lower overall cost. Testing can be divided into three basic areas: substrate test, integrated circuit (IC) test and module test. Substrates are tested during fabrication and prior to component attach. ICs and other components are tested prior to mounting on the substrate. The assembled module is tested prior to final sealing to permit repair of faulty components.

Multichip modules (MCMs) provide an increased level of performance over printed circuit boards by virtue of their dense packaging. It is this same dense packaging that complicates the testing process. The small size of the substrate and component pads makes access by electrical probes difficult and creates a need for test techniques which are unique to MCMs. The ICs mounted on the substrate are unpackaged and typically run at high speeds. A bare IC is often called a die and requires special handling and test to meet MCM needs. Finally, the completed module itself is a complex electrical system which must be tested

as a single unit. Simulation and specialized diagnostic techniques are required to test fully the operation and to locate any faults.

MCMs provide considerable testing challenges which will be addressed in this chapter. While not all of the problems associated with MCMs test have been solved as of the writing of this book, each of the testing issues is presented here along with the current state of the art. Substrate testing is discussed first since this forms the foundation for the module. Testing of ICs and other components is then presented, followed by the simulation and testing of assembled modules. Figure 13-1 presents a flow chart for the entire MCM testing process.

13.2 SUBSTRATE TEST

13.2.1 Introduction

Electrical testing of the MCM substrate is an integral part of the overall MCM fabrication process. The substrate contains the wiring used to interconnect all of the components on the completed MCM. Testing is used to verify the connectivity of the substrate and to monitor the fabrication process for quality control. A final electrical test is performed typically at or near the end of the substrate fabrication. In some cases, electrical testing is performed in-process to monitor the yield after various steps.

MCM substrates contain a number of electrical networks (nets) which are terminated in pads on the surface. These pads are used to bring the signals in and out of the substrate. Some pads connect to the chips which are mounted on the substrate while others connect to off module components. Electrical testing is performed by bringing one or more probes into contact with the pads on the substrate and making measurements. Each net on the substrate must be tested to ensure that there are no open circuits within the net and also that there are no short circuits to any other net. Nets also can be tested to make sure that the resistance from one pad to another meets design criteria. Additional testing can include measurements of impedance, propagation delay, crosstalk and high voltage leakage. Typically, every substrate receives testing for opens and shorts. Additional measurements are made as required to meet specifications and to maintain quality control.

MCMs fabricated using thin film technology on silicon or ceramic substrates are built up in successive layers using lithographic techniques. Pads occur on only one side of the substrate. These substrates are characterized by pad size $\leq$ 100 μm, pad spacing $\leq$ 200 μm and pad densities $>$ 100 pads per square cm. The layers are defined using a lithographic process which results in very accurate

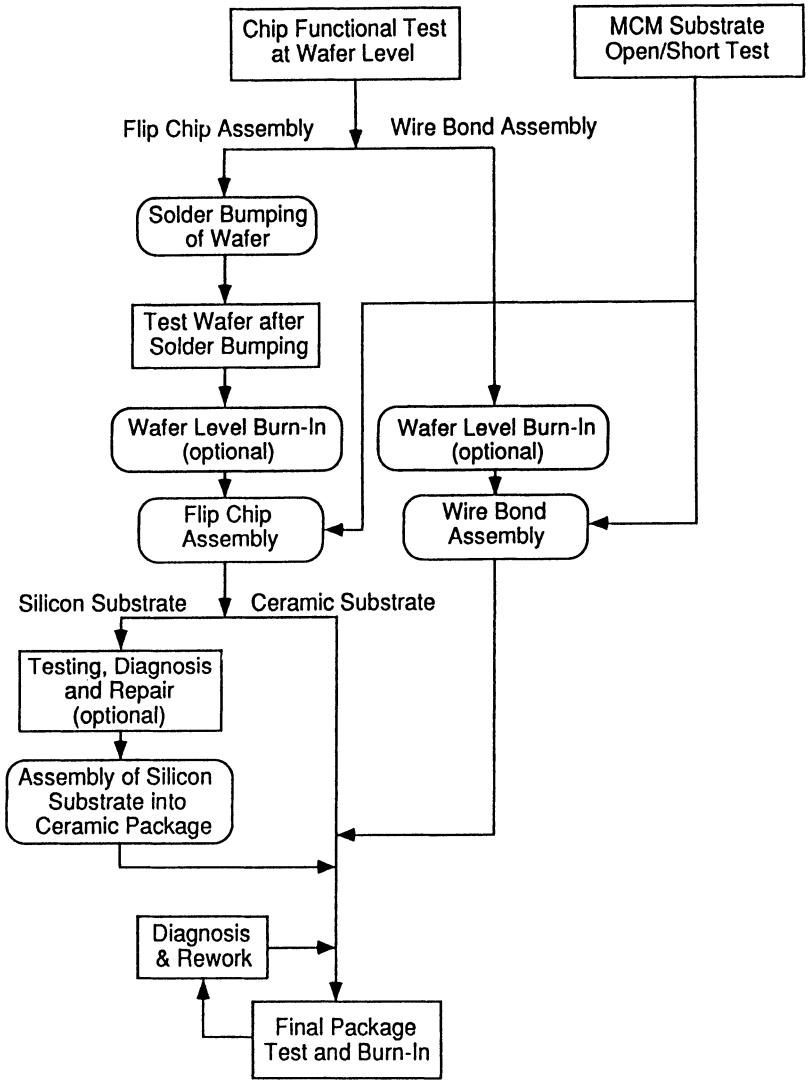

Figure 13-1 Flow chart of testing procedure for multichip modules.

placement of the pads across the entire substrate. Pads are aluminum, nickel, copper or gold providing good contact for probes, but they can be damaged by excessive probe pressure. A single substrate may contain several modules which range in size from 10 - 15 cm across, with larger sizes anticipated in the future. The thermal mismatch between the substrate and the thin film layers can cause camber (warpage) of the wafer, as discussed in Chapter 7. This camber does not usually cause a problem during interconnect testing since the substrate is held flat on a vacuum chuck. Automatic handling equipment is readily available and the modules are cut apart after interconnect testing.

MCMs fabricated using cofired ceramic technology contain layers of ceramic material which are fired at high temperature. The firing process causes the substrate to shrink, introducing distortion in the placement of pads across the substrate. This distortion complicates the alignment of the probes used with testing systems. The pad size is typically ≥ 100 μm, the spacing ≥ 200 μm and the densities typically lower than those found on thin film substrates. Cofired ceramic MCM substrates often contain pads on both sides and testing must include verification of the interconnection from top to bottom side. In addition to shrinkage, cofired ceramic parts are subject to camber resulting in variations in the height of the test pads and require adjustments in the probing level. The metallization is tungsten followed by layers of nickel and gold. Interconnect testing of cofired parts is done after firing, when the part is at or near its final size. In some instances, the circuitry can come very close to the edge of the part which complicates the fixturing. Automatic handling of ceramic parts tends to be more complex than that for thin film parts since the cofired part is formed in a variety of sizes and shapes.

MCMs fabricated using an overlay technology as described in Chapter 7 start with the bare die mounted on a ceramic substrate and deposit interconnect layers over the die to complete the circuit. Since there is no interconnect substrate distinct from the completed module, testing is performed as part of the module test described in Section 13.4.

No matter what the fabrication technology, MCMs are essentially miniature printed wiring boards (PWBs). Some testing techniques from PWBs can be employed in testing MCMs, but the small size and high density of pads creates a formidable testing challenge. Access to the individual pads is more difficult than on a PWB, so there is a temptation to wait and test the functionality of the assembled MCM from its input/output (I/O) pins. Unfortunately, this is not often viable economically since repairing the module at that point would require removal of all of the ICs and then mounting them on a fresh substrate or simply discarding the entire assembly. The rule of tens holds in this case - it is ten times as expensive to diagnose and correct a defect at the next level of integration. It is therefore wise to locate the fault as early in the process as

possible. Adding this fact to the substantial benefit gained by using electrical measurements to monitor process quality provides an overwhelming motivation for full electrical test of the substrate.

The test methods described in this section are used to test for the most common types of electrical failure modes: opens, high net resistance and shorts. The selection of any test method will depend on the level of testing required and the volume of parts to be tested. A testing program should be developed as an integral part of the design and manufacture process. Proper design of the substrate can greatly reduce the testing time and decrease the testing cost. Design for testability applies equally to the substrate as well as to the whole system. It no longer is possible to create designs without considering the test issues. Careful consideration of testability results in a final product that is not only less expensive to produce, but also higher in quality.

13.2.2 Fixed Probe Array Testing

A fixed probe array tester (often called a bed-of-nails tester) uses a head with many probes arrayed to contact all of the pads on the substrate at the same time. Each probe is connected to one channel of an analog multiplexer. Each channel of the multiplexer can connect to either side of the measurement unit (typically an ohmmeter). This allows the meter to be connected between any pair of pins or between two groups of pins. A schematic diagram of a 2 × n switch matrix is given in Figure 13-2. In actual practice, the switch channels are often configured as a 4 × n matrix to permit 4-wire Kelvin type measurements to be made which negate the effect of the resistance in the matrix and the wiring.

Opens and high resistance within a net can be located by measuring the resistance between all the pads on the net. The total number of tests required for this is calculated using:

$$\text{number of open tests} = \sum_{i=1}^{n} (p_i - 1) \qquad (13\text{-}1)$$

where n = number of nets and p_i = number of pads in net i. Shorts between nets are located by measuring from each net to every other net. This is a simple combination of n objects taken two at a time; the number of tests required is figured as:

$$\text{number of short tests} = \frac{n(n-1)}{2}. \qquad (13\text{-}2)$$

For a circuit with 100 nets and an average of four pads per net, the total number of tests is:

$$(100)(4-1) + \frac{(100)(99)}{2} = 5{,}250 \text{ tests.} \qquad (13\text{-}3)$$

An analog multiplexer uses solid state relays and can make 2400 measurements per second [1] so the total test time is still less than 3 seconds.

If the number of nets increases to 1000, then the number of tests required, and consequently the test time, increases dramatically:

$$(1000)(4-1) + \frac{(1000)(999)}{2} = 502{,}500 \text{ tests.} \qquad (13\text{-}4)$$

Fortunately, if the multiplexer allows more than one probe to be connected to the measuring unit at the same time, the number of tests required to locate shorts can be reduced radically. One method is to check continuity from each net in turn to all the other nets which have been tied together. If the part is good, then only n measurements will be needed. For each shorted net identified, additional (n-1) measurements will be needed to isolate the short. Other methods employing binary search techniques can reduce the number of measurements still further.

Bed-of-nails testers are commonly employed to test PWBs. Test heads for PWBs are typically a regular array of probes since many boards are designed on fixed centers (either 100 or 50 mils). Since MCMs contain bare die, with no pads on a grid, a custom probe array is required for each new substrate design. Various types of probes are employed to build the test head including spring loaded pogo probes, buckling beam probes and cantilever probes. When the head is contacted to the board, each probe aligns with a single pad on the board. As the size and density of the pads increase, the probe head becomes increasingly difficult to fabricate. While there is no absolute limit to the complexity of the test head, the viability of the method diminishes due to the high per pin cost of the test head and the fact that a head is useable for only one design. In addition, the reliability of individual probes tends to go down with their size. A single bent or nonfunctional probe will disable the entire tester.

For silicon substrates, the pad density is often too high to use a fixed probe array approach. If the pad density is low except for one or two small areas of

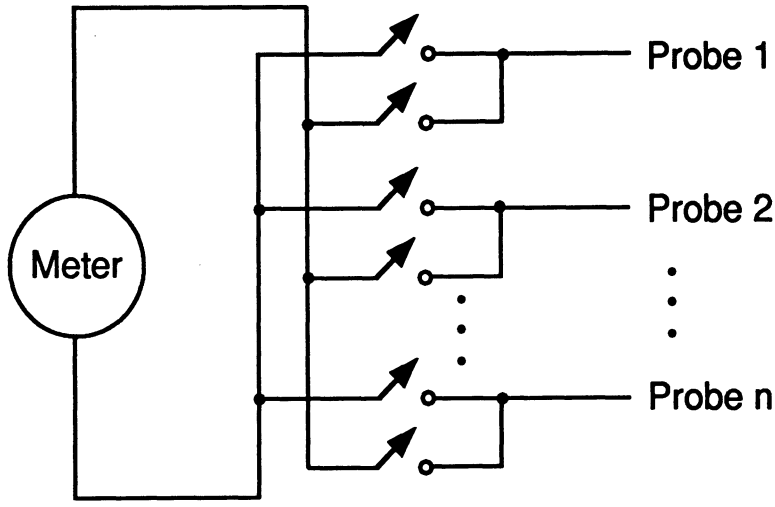

Figure 13-2 Schematic diagram of fixed probe array tester.

high density, it is often possible to build up a test head out of varying probe technologies. Usually, cofired ceramic MCM parts cannot be tested using a fixed probe array due to shrinkage variations across the substrate and from part to part. Cofired parts typically have a size specification of ±0.5 percent across the entire substrate [2]. For a part measuring 4" on a side, this translates into a pad placement error of ±20 mil, which is often larger than the pad to be probed. A fixed probe array only can be used if the pads are larger than the errors in pad placement. One way to reduce the pad placement error is to apply the top layer metallization after the part is fired. This layer contacts pads on the fired layer below and provides probing pads with much less distortion across the substrate.

Fixed probe array testing provides very high throughput since all the pads are contacted at once and only electronic switching is required to complete the test. It is employed most favorably for high volume testing since the tooling cost is high. It also adds considerable lead time to the test setup. Although it cannot be used for cofired parts with small pads or with very high density parts, it can be used with many lower density designs. For this reason, it may be prudent to design the substrate to accommodate this type of testing if high volume is anticipated.

13.2.3 Single Probe Testing

Single probe testers build on an idea originally patented by Honeywell [3] and made popular by Teledyne/TAC [4]. Each net in an MCM substrate is isolated

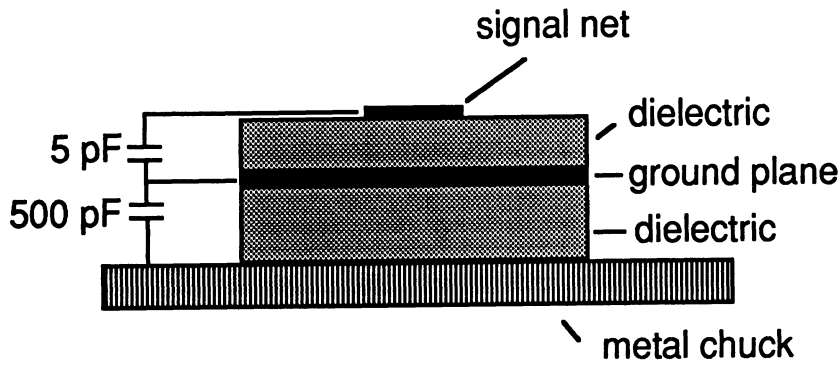

Figure 13-3 Capacitance relationship of a signal net to the ground plane on a substrate.

from all other nets by a dielectric material (polymer, ceramic etc.). Due to this relationship, there is a certain capacitance between each net and a nearby ground plane based solely on the geometry of conductors in the net, the separation between the net and the ground plane and the dielectric constant. The ground plane itself can be an internal plane inside of the part or it can be the chuck that the part rests on during the measurement process. The tester operates by connecting one side of a capacitance meter to the ground plane and connecting the other side to a single probe which is brought into contact with each pad on the substrate. While the theoretical capacitance of the net can be determined from its design, measurement of this capacitance on a known good part (also called a golden part) to create a reference is more common. If a golden part is not available, capacitance is measured on several parts and the median value for each net is computed to create a statistically derived reference. During testing, if the capacitance of any net has increased above the reference value, then its size must have increased, and it is likely therefore that it is shorted to another net. If the capacitance has decreased, then the net probably contains an open circuit.

The chuck used as the ground plane often has a thin dielectric coating over it to prevent shorting of any exposed circuitry on the substrate. Use of an external ground plane is the most common method since only one probe needs to contact the substrate, simplifying the mechanical system of the tester. A ground plane inside the substrate also can be used, but this requires the addition of a fixed probe to contact a pad connected to the plane. If an external plane is

used, the measured capacitance of the net is dominated by the capacitance between the net and any internal ground plane. This is illustrated in the Figure 13-3. Assume the capacitance between the signal net and the ground plane is 5 pF and the capacitance between the ground plane and the external ground of the chuck is 500 pF. Since these two capacitances are effectively in series, the total capacitance is calculated using:

$$\frac{1}{C_T} = \frac{1}{C_1} + \frac{1}{C_2} = \frac{1}{5} + \frac{1}{500} \qquad (13\text{-}5)$$

$$C_T = 4.95\,\text{pF}$$

If the capacitance between the internal ground plane and the chuck varies due to changes in substrate thickness or to mounting on the chuck, the effect on the measured capacitance is very small.

Figure 13-4 shows two nets labeled N1 and N2 on a substrate. If net N1 has a nominal capacitance of 5 pF and net N2 has a nominal value of 8 pF, then a short between the two would result in a capacitance of approximately 13 pF. The total might be slightly more or less than the sum of the two net capacitances due to the additional metal area of the short itself and the possibility that there are some overlapping areas in the nets. A fault condition would be indicated if the capacitance of either net were significantly greater than its nominal value. The tolerance setting is somewhat dependent on the substrate technology and on the degree of feature size control, but a typical value is 25%.

Using this method, a substrate with n nets will require exactly n measurements to verify that there are no shorts. It should be pointed out that not all shorting nets can be identified specifically using this method. If a 20 pF net shorts to a 1 pF net, then the total capacitance will be about 21 pF. This is much larger than the expected value of 1 pF, so that net will be flagged as a fault. However, since it is too close to the expected value for the 20 pF net, no fault will be indicated on that net. For this reason, the capacitance method is used primarily as a go/no go test. Further testing of the part is required usually to isolate the fault.

Opens within a net are indicated by a reduction in capacitance. Figure 13-5 shows a typical net with two possible break sites labeled X1 and X2. If the capacitance test is made at pad A, then a break at X1 will cause a loss of about half of the capacitance. A break at X2 will cause only a very small change in capacitance since only one pad was lost from the net and might not be detected as a fault. If, however, an additional capacitance measurement were made at pad D, then the capacitance measured would be much lower than the expected value since pad D is very near break X2. In order to make sure that opens are

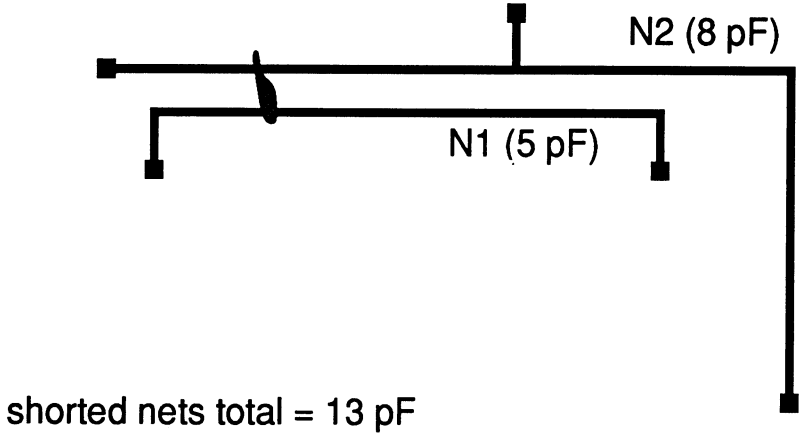

Figure 13-4 Capacitance of two nets shorted together will be approximately equal to the sum of the capacitance of each net.

identified regardless of their location in the net, a capacitance measurement is required for each pad on the net. The relative capacitance values measured at each pad on the faulty net can be used to help locate the fault within the net.

Single probe testing can be used to provide opens and shorts testing on a wide variety of MCM substrates. There is no inherent limitation on pad size or density since only a single probe is involved. As pad sizes decrease, the probe can be made smaller and the motion system more accurate. Since there is no fixed array of probes, distortion from one part to the next, as is common for co-fired ceramic parts, can be handled by an alignment procedure which maps out the distortion. The simple mechanical system results in a small equipment footprint with no need for the bulky and complex switch matrix required by a bed-of-nails tester. New parts can be set up in a matter of hours since there is no fixture required. The pad locations for the new part are simply programmed and the capacitance values learned from a sample of parts.

Single probe testing does have its limitations. The capacitance method can locate opens and shorts but it cannot find high resistance within a net. To the capacitance meter, any resistance less than a few megohms looks like a short circuit. This may not be adequate for MCM parts, particularly in the early prototyping stage, where high resistance faults due to line width and thickness variations may be common. Since the capacitance reference values are learned

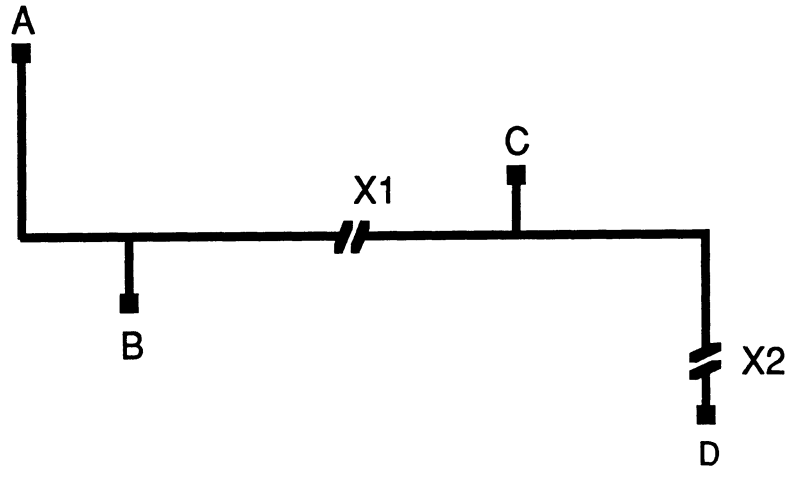

Figure 13-5 Net showing two possible break locations, X1 and X2.

from actual parts, a design error may go undetected. If the substrate has circuitry on both sides, then it must be turned over to complete the test. The testing speed of this tester is on the order of five measurements per second and is a strong function of the distance from one pad to the next. While the throughput is much lower than that for the bed-of-nails tester, the setup cost for a new product is much less and there is an inherent improvement in reliability realized by the use of a single probe instead of thousands of probes.

13.2.4 Two Probe Testing

Two probe testers provide the benefits of single probe testers with the added ability to make resistance measurements. In a two probe tester, the substrate is held stationary on a vacuum chuck or other fixture and the two probes move over the substrate making measurements at the required pad pairs. The measurement can be the resistance from one pad to another within a net or the leakage current between two nets. The simplest two probe testers have only resistance capability. These can perform measurements similar to those of a bed-of-nails tester by moving the probes from pad to pad. This is adequate for opens and high resistance testing within a net, but too slow for shorts testing. As was discussed in Section 13.2.2, the number of tests required to detect all shorts

increases geometrically with the number of nets. To locate shorts on a substrate with 1000 nets would require over 500,000 measurements. A high performance two probe tester can make about ten measurements per second so it would take nearly 14 hours to test a single substrate. This could be reduced somewhat by eliminating tests between nets that could not possibly be shorted due to their routing, but in many substrates this would not be adequate.

Two probe testers typically add capacitance test capability to their resistance testing. In this way, they can perform with all of the advantages of a single probe tester for shorts identification but with the added ability to isolate shorts to specific nets. A resistance measurement is made between any net with excessive capacitance and all other nets with similar capacitance to determine which pair or pairs of nets is shorted. Integri-Test® and Bath Scientific Limited (BSL) both offer testers for the MCM market utilizing this combined measurement capability. Integri-Test® uses a patented method to combine the tests for resistance and capacitance and thus provide a higher throughput [5].

By making use of the resistance measurement feature of a two probe tester, a design verification of a new product can be performed. A design verification checks that the net list of the physical substrate matches the design. During verification, each net is checked for opens and then each net is checked to every other net for shorts. The shorts verification can be made faster by using measurements of capacitance to limit the number of tests, checking the resistance only between nets with similar capacitance. Also, two probe testers can make net to net leakage measurements using a high voltage source and an electrometer to measure the current.

Dual sided two probe testers, such as the BSL PRECISIONPROBE™ pictured in Figure 13-6, can test both sides of an MCM. The part is mounted in a fixture so as to expose both sides at the same time. The motion system is duplicated above and below the part to permit probing of connections on each side of the part and also from one side to the other. When a part is mounted for probing of both sides, there is no external ground plane provided by a mounting chuck. In this case, one of the probes is brought in contact with a pad connected to an internal power or ground plane to provide the electrical reference needed for capacitance measurement. If the part does not have internal metal planes, then it cannot be tested using the capacitance method unless it is mounted on a conductive chuck to provide the plane. Unfortunately, this restricts access to one side of the part at a time. An MCM with pads on both sides should include an internal metal plane in its design if the part is to be tested on a dual sided two probe tester.

The testing speed of both single and two probe testers is limited by the mechanical motion of the probes. Increases in the positioning speed of the probes translate directly into decreases in test time. The positioning speed

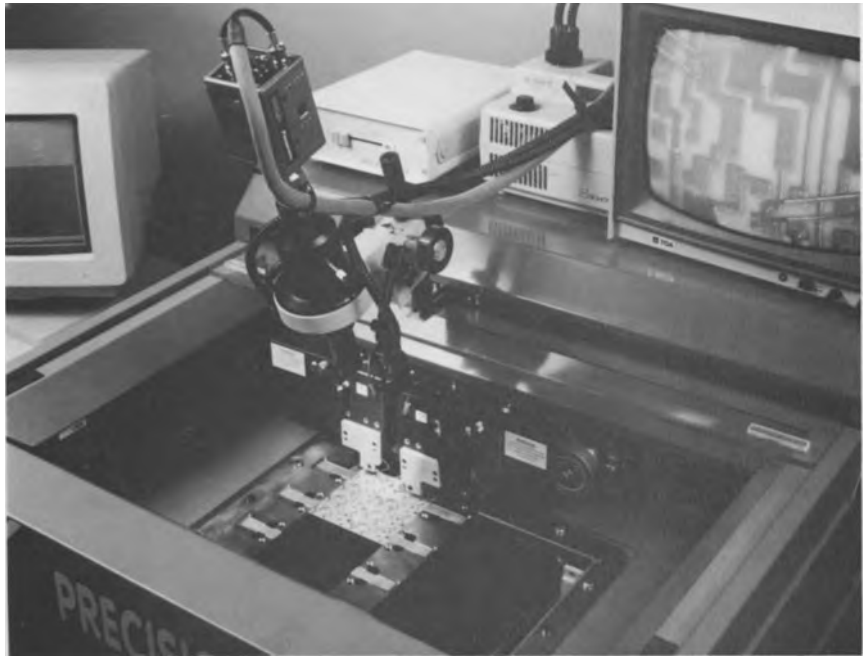

Figure 13-6 BSL PRECISIONPROBE™ two sided tester: interior view (courtesy of Bath Scientific Limited).

improves as the distance covered from one measurement to the next decreases, so moving probe testers operate faster on smaller substrates. A proprietary two probe tester designed and built by the author at AT&T Bell Laboratories averaged 12 measurements per second while testing thin film on silicon substrates.

One method to improve the testing speed is to reduce the total distance that the probes need to travel to make all measurements. For single probe testers, this is equivalent to a popular mathematical problem known as the Traveling Salesman Problem. This problem postulates a salesman who needs to visit a

number of cities and would like to determine the best order in which to make the visits covering the least possible total distance. Solutions of this problem fall into a class of problems which are intractable; that is as the number of cities gets larger, the calculation time required to solve the problem grows explosively [6]-[7]. Exact solutions to the problem are feasible for a few hundred cities at most. There are many solutions to the problem which yield a good but not optimum route. A single probe moving around an MCM tester can be considered as a kind of traveling salesman.

Two probe testing somewhat complicates the problem since now we can consider the probes as two salesman which must visit different cities at the same time. Integri-Test® uses a combination of serpentine motion of one probe with the other probe moving as required to contact the remaining points on the net [8]. The two probe tester developed at AT&T Bell Labs used a scheduling algorithm which considered all measurements for all nets at once and then sorted them to obtain an optimal route [9]. The original sorting method used at Bell Labs was a Greedy Algorithm solution to the traveling salesman problem modified to handle the notion of two probes. In the greedy algorithm approach, the next probe position selected is the one which requires the least possible movement from the current position. Improvements in the route are achieved using more sophisticated algorithms and additional computation [10].

Two probe testers provide great flexibility in MCM substrate testing. Opens, shorts and high resistance faults can all be identified. Other two terminal measurements can be made by simply connecting an external meter to the probes. Two probe testers can accommodate parts with pad sizes down to 2 mils and even smaller with improvements in the positioning systems. Substrate distortion found on cofired MCMs can be handled using distortion mapping just like in single probe testers. Parts with die cavities present no problem by using a mechanism with programmable probe actuators. Probing from one side of a part to the other side also is possible using a tester configured for this mode of test. The primary disadvantage of two probe testing is the testing speed. Typical speeds are three to six tests per second on commercial machines, but this is improving with new models. As the testing speed increases however, it becomes increasingly difficult to bring the probes to a vibration free stop due to the high acceleration and deceleration required.

13.2.5 Electron Beam Probing

All of the probing methods described above use a physical probe which is brought into contact with the test pads and which can leave a mark on the pad.

In a properly designed test system, this mark usually does not cause any problems with wire bonding or with flip chip attachment. There is a practical minimum size to the pads which can be physically probed. Mechanical probing can usually handle any interconnect pad on the surface of the substrate since these must be large enough to connect to a component such as a die. In some instances, however, it may be desirable to probe a substrate during the fabrication process. Thin film MCMs are built up one layer at a time and the interconnect pads may not exist until the final layer is applied. Before application of the final layer, only lines and vias may be available for probing with probing areas well below 20 μm across. In this case, automated mechanical probing is not feasible and electron beam probing provides an alternative.

Electron beam probing operates on the same principle as a scanning electron microscope. An electron beam is directed toward test points on the surface of the substrate. The beam charges the selected test point to a preset voltage. A secondary electron energy analyzer is used then to locate other places on the substrate which also have been charged. Opens can be detected by charging one pad on a net and then checking for charge on every other pad. Shorts can be detected by charging one net and then checking that no other net has been charged. An electron beam flood gun is used to discharge or charge all nets at once. Electron beam testers are not commercially available yet. Significant development efforts have been made both at the IBM Research Center [11] and at Microelectronics and Computer Technology Corporation (MCC) [12]. The system developed at IBM has the additional capability of an electron beam flood gun mounted below the substrate. This is useful for testing continuity through a substrate by flooding the bottom surface with electrons and then locating all pads on the top surface which become charged.

The electron beam is deflected across the substrate circuit in the same way that a television paints a picture on the screen. Since there are no moving parts, the testing speed can be very fast. Speeds of 100 tests per second have been achieved and greater speeds are possible [12]. Electron beam testers are not without their disadvantages, however. They require that the substrate be placed in a high vacuum chamber that takes several minutes to pump down. Systems can be equipped with load air locks to speed this process, but there will always be a throughput penalty. There is also a tradeoff between maximum substrate size and resolution. As the substrate becomes larger, the resolution of the beam placement decreases. Errors on the order of 100 ppm are achievable, [12] so for a 10 cm substrate the placement accuracy can be held to 10 μm. While this is adequate to probe pads on the substrate, it may not be enough to probe vias or lines. As substrates grow larger, mechanical movement of the substrate under the beam may be required which would diminish the speed advantages. Electron beam testers also cannot check for high net resistance. The net charging is

Table 13-1 Comparison of MCM Substrate Test Methods.

Test Method	Typical Tests/sec	Fixturing Required	Minimum Feature Size	Distortion Compen-sation	Double Sided
Fixed Probe Array	2400	yes	4 mils	no	yes
Single Probe	3 - 10	no	2 mils	yes	no
Two Probes	3 - 10	no	2 mils	yes	yes
Electron Beam	100+	no	<1 mil	yes	limited

similar to the capacitance method used by a single probe tester and is insensitive to high resistance within the net. Resistances of less than several megohms all appear as shorts to an electron beam system. Finally, electron beam testers are inherently expensive due to the vacuum chamber and sophisticated electronics required for their operation. Despite these limitations, electron beam probing may provide high throughput production capability for MCM testing in the future as the technology is developed further.

13.2.6 Developing a Substrate Test Strategy

The development of a substrate test strategy requires careful consideration of a number of factors including volume of parts, technology of substrate, failure modes, level of testing required and where in the process testing will occur. For high volume parts and low pad density, a fixed probe array tester probably provides the best solution. If the pad density is high and the process is well developed, then single probe testing may be adequate. If the resistance of each net must be tested, then a two probe tester is a good choice. If probe marks cannot be tolerated on the finished component, then testing the substrate prior to final metallization will be required. Above all, consideration of these factors should feed back to the design area to provide a design which is easy to test. Table 13-1 summarizes some of the features of each of the available test methods.

13.2.7 Summary

Electrical testing of MCM substrates is essential to ensure a quality product and

commercial equipment is readily available. Fixed probe array testers provide high speed but are limited in the range of pad sizes and densities that can be handled. Single probe testers provide a simple go/no-go test for a variety of substrates at limited speed. Two probe testers provide excellent test flexibility with limited speed, but may provide a cost advantage compared to fixed probe testing when used for low volume production runs [8]. Electron beam probers provide very high speed and can handle a wide range of MCM substrates but are not yet commercially viable. For many MCM substrates, including high density interconnects on silicon and ceramic, the only testing method which can provide a range of test capabilities is the two probe tester. This technology is fairly mature, but further increased in testing speed can be expected as the motion systems are improved. While all of the substrate test technologies described above have been demonstrated for pilot line volumes, further development will be required to handle the large volume of MCM production expected in the coming years. Substrate test will be a significant factor in the determination of the cost effectiveness of MCMs.

13.3 IC TEST

13.3.1 Requirements for MCM Modules

By definition, MCMs contain two or more bare ICs mounted to a common substrate. These ICs can belong to any logic family (CMOS, ECL, etc.) and may be analog as well as digital. The dividing line between hybrids and MCMs is indistinct, but hybrids usually contain a variety of components including lower lead count analog and digital ICs, while MCMs primarily contain high lead count application specific integrated circuits (ASICs), microprocessors and memory ICs. Since the vast majority of current MCM designs do not contain any analog ICs, we will focus our attention on the testing of digital components. While IC testing is a mature technology, MCMs present some unique challenges. Most ICs today are not tested fully and burned-in until they have been mounted into their packages since it is much easier to perform the final testing of the chip in its packaged form. ICs which are destined for MCM use must be tested fully as bare die since it may be difficult to perform a complete test once the die are mounted to the MCM substrate.

MCMs are utilized typically for high speed applications and require ICs and substrates capable of operating at high speed. Flip chip bonding is likely to be the method of choice for chip-to-substrate connection as MCM technology matures since it provides the shortest electrical path from the die to substrate and permits the die to be packed very close together. Some ICs designed for MCM

use may have very small output drivers since they will only need to connect to other devices on the same MCM through low capacitance paths. Testing of these devices will require special electrical test equipment designed to present a very small load on the device.

The quality of die used in an MCM is a key factor to producing a finished part with high yield. To illustrate this point, consider an MCM with 10 die. Even if there is no yield loss from the assembly process, but 5% of the die fail during module burn-in, then the yield will be $(0.95)^{10}$ or only 60%. This means that 40% of the MCMs will require some sort of rework to obtain a fully functional module. This rework is usually costly since the die are bonded to the substrate and not placed in sockets like on a PWB. From this example, it is clear that a die which is to be used on an MCM should be as free of defects as possible prior to installation on the substrate.

While it very desirable to obtain known good die, there are still some technical and infrastructure obstacles which need to be overcome. It is difficult to perform high speed functional testing of ICs in wafer form due to parasitics introduced by the wafer probes. There is also no satisfactory method for burning in ICs either in wafer form or as individual die. From a business standpoint, most companies are not interested in selling bare die since part of their profit is derived from the package and it is difficult to guarantee the quality of bare die due to the technical problems mentioned above. In addition, IC manufacturers often shrink the size of their die without telling the customers since the die are fully packaged anyway. The success of MCMs hinges on the availability of known good die. It is for this reason that many companies are now participating in programs to create standards for bare die [13].

13.3.2 Introduction to IC Test

A complete test of a digital IC includes testing of the functionality of the device as well as its quality. IC testing is broken down into two basic areas: functional and parametrics. In functional testing, the IC is subjected to a sequence of input states known as vectors, and is monitored for the correct sequence of output states. Parametric testing provides a measure of how quickly the IC responds and the electrical specifications of the inputs and outputs. As the complexity of ICs grows, it becomes increasingly difficult to test completely all the functions from the I/O pads. This has given rise to a built-in self test (BIST) method which uses a set of structures designed into the IC to exercise some of the internal functioning of the IC. The self test process is controlled from the I/O pads and usually complements the traditional functional test. In addition to BIST, boundary and internal scan have become important design for testability features. Boundary scan provides access to all of the I/Os on a IC through four

special purpose pads added to the IC. Internal scan provides a similar function to access internal nodes of the IC. Boundary scan is of great benefit to MCM test since access to all of the device I/O pads may be impossible.

13.3.3 Test Generation

Most of the ICs incorporated into MCMs consist of microprocessors, memory controllers, digital signal processors, memories and ASICs. Functional testing for all of these IC types is similar. Prior to any testing, the IC is modeled on a digital simulator and the functionality of the IC represented by software models. In some simulations, each gate on the device is modeled, while in others, the functionality of groups of gates or cells is modeled. Once the model has been created, a set of bit patterns is applied to the model inputs and the corresponding outputs are determined by simulation.

These inputs and outputs form the basis for a functional test of the IC. Creating the model is a fairly straightforward task since it is a representation of known physical elements. Generating the set of inputs to exercise the model is quite another matter. A thorough knowledge of the expected use of the device is required to create the input patterns. Creating the input patterns for a simple combinatorial logic circuit is easy since the device does not have any internal memory. This is not true for a microprocessor where the output state depends not only on the current input state, but also on one or more previous states. For this reason, it can be very difficult to create a set of input patterns to generate completely all possible output states.

In addition to covering all possible output states, the input patterns should be able to identify all possible faults inside of the IC. Faults inside the IC can take many forms, but the most common are stuck-at faults and shorts and opens. Stuck-at faults are gates which will not change state regardless of their inputs. Shorts and opens are caused by defects resulting in improper interconnection of gates. An internal fault can be detected only if the application of input patterns causes an unexpected set of output patterns. The ability of a set of input patterns to reveal a fault is measured as the test or fault coverage. For complex circuits, it is usually not practical to provide 100% fault coverage. The fault coverage typically increases with the length of the test pattern, so coverage usually can be increased by making the test longer. Microprocessors often require millions of input test patterns to obtain adequate fault coverage. Clearly, patterns of this length cannot be generated by hand. One method to create the test patterns for a microprocessor is to write a program using the appropriate assembler language which makes use of all of the microprocessor functions. The compiled test program is captured in its binary form and applied as inputs to the simulation program of the microprocessor. Input test patterns for memory devices can be

created by simple algorithms such as walking bit and checkerboard. In ASICs, the generation of test patterns is aided by creating the device from a set of standard cells.

13.3.4 Boundary Scan and Built-In Self Test

A new standard has emerged recently to improve the testability of both ICs and completed MCMs. The standard, IEEE 1149.1 [14], commonly referred to as boundary scan, provides for the addition of a shift register at each of the IC I/Os. All of the shift registers are connected together in a serial chain with each end of the chain connected to an added I/O pad. Two additional pads are added to the IC to provide clocking and control of the modes of operation of this boundary chain. The boundary scan controller or test access port (TAP) is accessible through these four new pads. A detailed description of the TAP modes can be found in the data sheets for the boundary scan equipped IC or the IEEE 1149.1 standard. A single boundary scan register is shown in Figure 13-7. When the IC operates normally, the internal logic is connected to the I/O pads. In this mode, the boundary is transparent, although there may be some delay penalty due to the addition of the extra gates. In the test mode, the IC logic is disconnected first from the I/O pads. A sequence of bits then is shifted into the

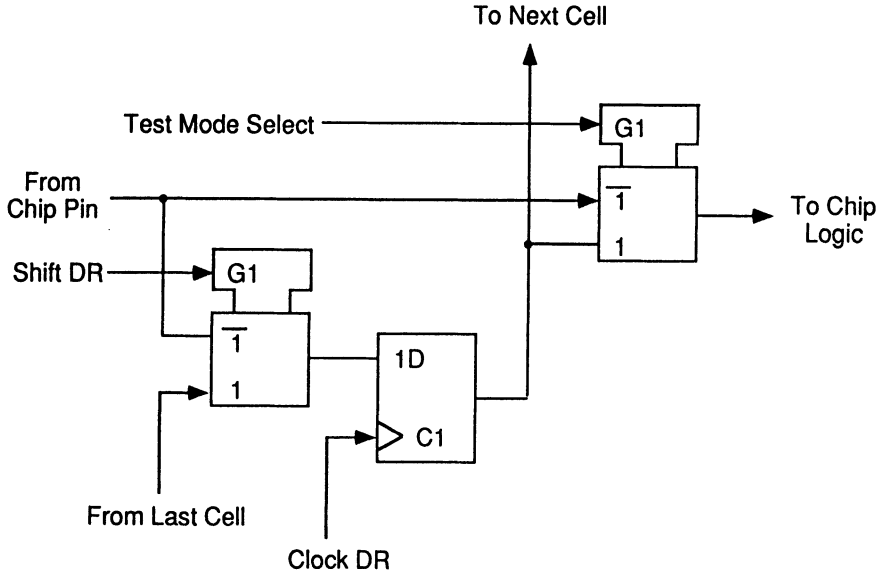

Figure 13-7 Basic boundary scan cell configured for an IC input pin.

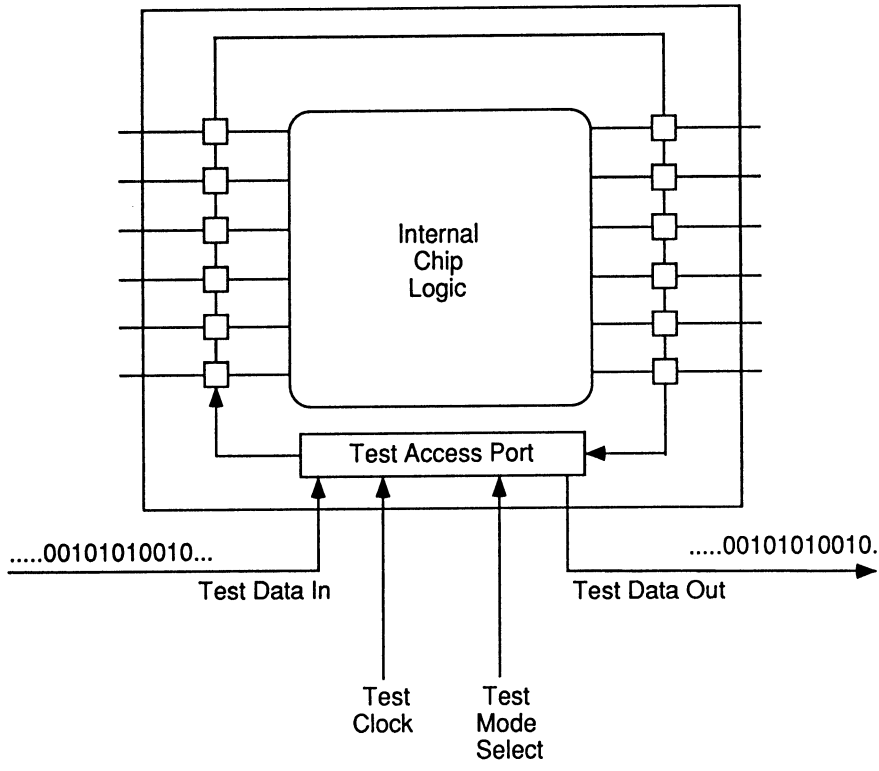

Figure 13-8 IC showing boundary scan cells connected together in a serial chain with TAP controller.

boundary so that each register contains the required logic level for each input on the device. The boundary control then connects the internal logic to the boundary registers and clocks the data on the boundary into the IC logic. Next, the state of the outputs is transferred to their boundary registers. The IC logic is disconnected from the boundary and all of the boundary registers are shifted out to read the results. Figure 13-8 shows an IC with boundary scan registers (or cells) and a TAP controller.

By using boundary scan, it is possible to access any I/O on the IC through connection to the TAP only. Internal scan builds on the boundary scan concept by adding additional serial chain paths inside of the IC as illustrated in Figure 13-9. These paths allow internal nodes of the IC to be accessed and thus tested.

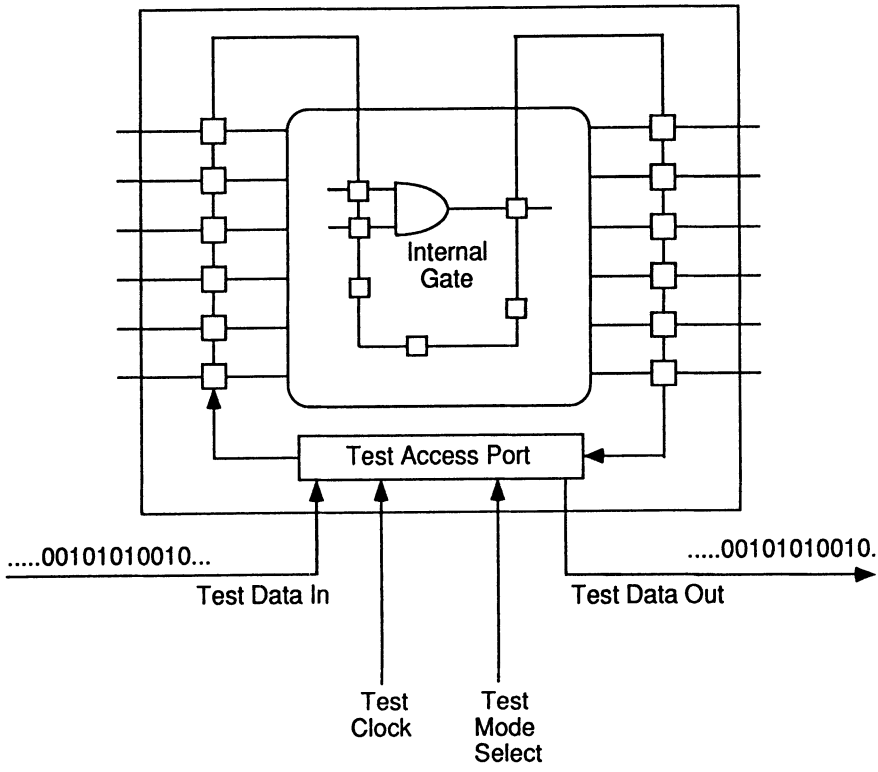

Figure 13-9 IC showing boundary scan and internal scan cells connected together in a serial chain with TAP controller.

In fact, if all gates in the device were accessible through internal scan paths, it would be possible to test fully the device by individually testing all of its gates and interconnections without an understanding of its function at all. Of course, this would not be practical as it would add greatly to the size of the IC and also would decrease its speed.

Another form of advanced IC testing is BIST. The concept of BIST is to provide additional logic on the IC itself which is designed to exercise some of the functions of the IC. There are many forms which BIST can take, but in general, an external tester provides clocking and control signals through IC pads to the BIST structure. After the test runs, the BIST reports its results to the external tester. This result is often in the form of a result code or check sum which is then compared to the expected results in the external tester. High speed

ICs can benefit greatly from BIST by qualifying the operation of the device at-speed [15] without the need for a costly tester with very high speed drivers and receivers. BIST devices require only that the tester provide a high speed clock signal and then wait for the final results. The use of BIST greatly simplifies the testing of complex ICs, but adds to the number of gates and, as a result, increases the die size. BIST can also be initiated from boundary or internal scan paths to reduce the need for extra I/O pads. Several articles and books have been published describing the use of BIST [16]-[18].

The most important function of boundary scan is to simplify testing of ICs once they have been mounted on a PWB or MCM. It is possible to connect the boundary scan paths of each IC together to form one long chain. This provides access to all of the chip I/Os through a single TAP. This feature is important particularly to MCMs and will be discussed in more detail in the section on module test. Boundary scan can be used in IC testing to perform functional testing without connection of all of the I/O pads. Devices with a large number of I/Os can be tested by converting all of the functional test vectors into serial chains and then loading them through the TAP. This does not provide a complete test, however, since there is no way to be sure that the IC I/O pads are connected to the IC logic since no signal is applied to the pads. In addition, an IC cannot be tested at full operating speed since all of the inputs and outputs must be accessed by loading in a serial fashion instead of in parallel as in normal operation.

13.3.5 Parametric Testing

Testing of ICs does not end with simply verifying the functionality of the device. Parametric testing is performed also to make sure that the device meets its performance specifications. Parametrics include both AC and DC measurements. Some of the more commonly measured AC parameters include propagation delay from an input to an output (t_{pd}), rise and fall time of output signals (t_r and t_f), the time that an input must be present before a clock transition (set up time, t_{su}), and the time the input must be maintained after a clock transition (hold time, t_{hd}). AC parametric testing is especially critical in ICs destined for MCMs due to their high operating speeds. DC parametrics include the voltage levels associated with an input high and low (V_{ih} and V_{il}), the voltage levels produced by an output high and low (V_{oh} and V_{ol}) and the quiescent and dynamic current required by the device (I_{DDQ} and I_{DD}) [19]. The specifications are determined by the designer in order to meet the requirements for devices in that logic family. An IC may be operating perfectly according to its functional test, but may be unable to meet its parametric specifications. In this case, the IC will be unable to interact properly with other components in the circuit.

13.3.6 Wafer Probing

During the fabrication of the wafers each of which contain many ICs, several tests are performed to maintain the quality of the process. The first functional tests of the ICs are made when the wafer is complete and the devices have not been separated into individual die. ICs which fail this test are discarded and the rest are mounted into packages. Typically a more complete test is made once the package has been sealed.

Functional and parametric testing is performed using automatic test equipment (ATE). A functional tester contains a number of driver/receiver channels which are wired to a fixture which makes contact with the pads of the device under test (DUT). Each signal pad of the DUT is connected to one channel of the tester. Power is also supplied to the DUT from power supplies built into the tester. The tester presents each of the input patterns which have been stored in a special pattern memory to the DUT in sequence and collects the results from the output pads. Discrepancies between the collected and expected results are reported as failures. Functional testers can operate over a wide range of clock speeds and provide the levels needed to drive a variety of logic families. The same testers usually include parametric measurement units to permit these measurements on the selected pads.

The wafer form provides an easy vehicle for handling, and permits testing of many parts without operator intervention. Functional testers configured for wafer testing use a wafer probing attachment equipped with a probe card as shown in Figure 13-10. The wafer is translated under the probe card so that each IC can be aligned and contacted by the probes. A low speed functional test typically is performed on each IC in turn. While a high speed test is possible at this point, it is complicated by the parasitic capacitance and inductance inherent in the probe card. High speed functional testing can be performed through the use of specially designed probe cards and connections to the functional tester [20]-[21]. The selection of tests to be performed at wafer level is determined by trading off the increased cost of additional testing and the expense of packaging a defective die. In the case of ICs to be used in MCMs, it is very desirable to test the part as completely as possible in wafer form since it will not be mounted in an individual package prior to installation on the MCM.

13.3.7 Die Carriers and Packages

Once the wafers have been probed, they are sawn apart to produce the individual die which are mounted then in single chip or multichip packages. In single chip

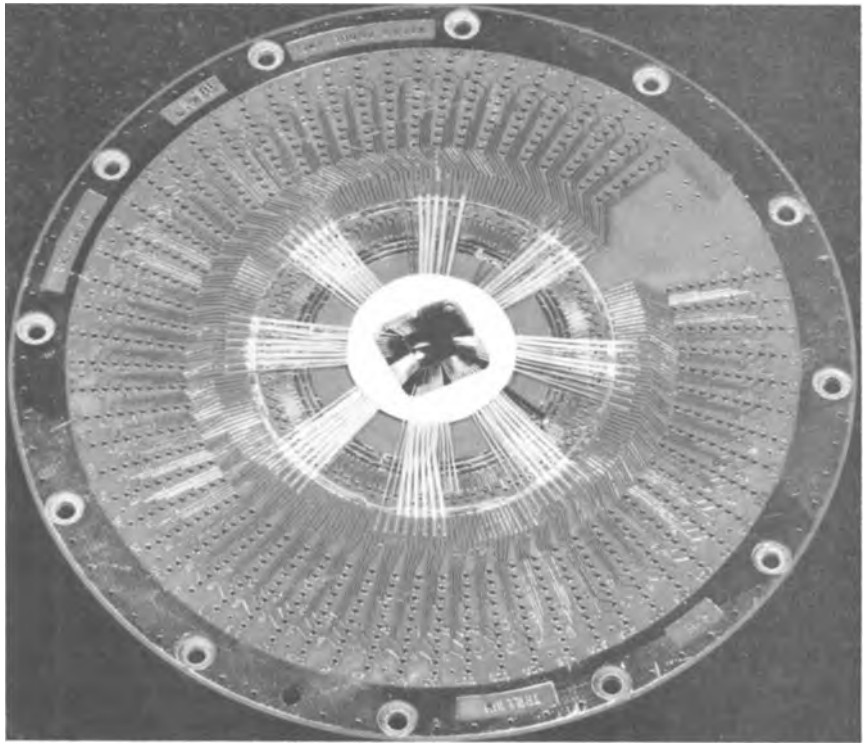

Figure 13-10 Typical probe card for use in testing IC wafers. (Courtesy of Micro-Probe, Inc.)

package applications, the IC undergoes a final test and burn-in once it is sealed inside the package. In the case of die to be used in MCMs, usually it is desirable to perform a complete test and burn-in prior to mounting on the MCM, as will be discussed later. Unfortunately, it is very difficult to handle bare die and there are no commercially available sockets which can hold a single die. One solution to this problem is the use of a carrier to hold the die during test and burn-in. The die may be removed from the carrier after test or it can be mounted directly on the MCM with the carrier.

One common carrier form is tape automated bonding (TAB) [22]. TAB is suitable for very high production since the carrier takes the form of a continuous tape onto which are mounted one die after another. Each TAB carrier consists of a lead fanout which is connected electrically to the pads on the die. There are many types of TAB carriers, but to be useful for electrical test and burn-in, the

lead fanout must be supported with a nonconductive tape using a two or three layer construction. It is possible to test the ICs while still in tape form or to cut the tape up into individual carriers and test each part in much the same way as packaged parts are tested. Once the IC has been tested and burned-in, the TAB lead frame is trimmed to its final size and the IC is mounted directly to the package.

There are many issues surrounding the use of TAB carriers. The custom TAB design required for each new die can be costly and add to the manufacturing lead time. The testing socket must provide adequate electrical and thermal performance without harming the carrier. Handling can be difficult due to the fine lead pitch. For very high speed ICs, additional parasitic capacitance and inductance in the untrimmed TAB leads may make it impossible to test the part at full speed. Ideally, for MCM applications, an IC manufacturer would supply die mounted on TAB which have been tested fully and burned-in. While TAB is popular in Japan, it is not in widespread use in the United States, making it difficult to obtain die in this form.

Another carrier form is a simplified plastic or ceramic sacrificial package. This carrier can take many forms including a simple fanout pattern on a substrate or an actual pin grid array (PGA) package. The die is mounted to the carrier and removed after testing is complete. If the die is wire bonded to the carrier, then these bonds must be broken to remove the die. This can be a difficult and costly process which has its own yield loss. If removal of the die from the carrier is too difficult, then the die with carrier can be mounted on the MCM [23]. In this case, the carrier is designed to be as small as possible and to have minimal impact on the electrical performance. If the die is designed for flip chip solder attachment, then a carrier can be designed to permit easy mounting and removal of the die with no damage. The carrier can be designed so that the die can be separated mechanically or reheated to melt the solder bump pads. After separation, the die is cleaned and additional solder added as required.

13.3.8 Final Test

Final testing is usually performed on the IC mounted into its package or carrier. Fixturing of the part to the ATE depends on the type of package. For those parts such as PGAs which are designed to be inserted into a socket, a matching easy insertion socket is provided on a special board which mounts on the tester. This performance board is customized to each IC design, connecting signal and power to the proper channels on the tester. A special socket is required to hold surface mount packages or carriers during testing without lead damage. The electronics and wiring of the tester is designed to mimic the actual operating conditions of the IC and thus determine its true performance. A photo of a typical functional

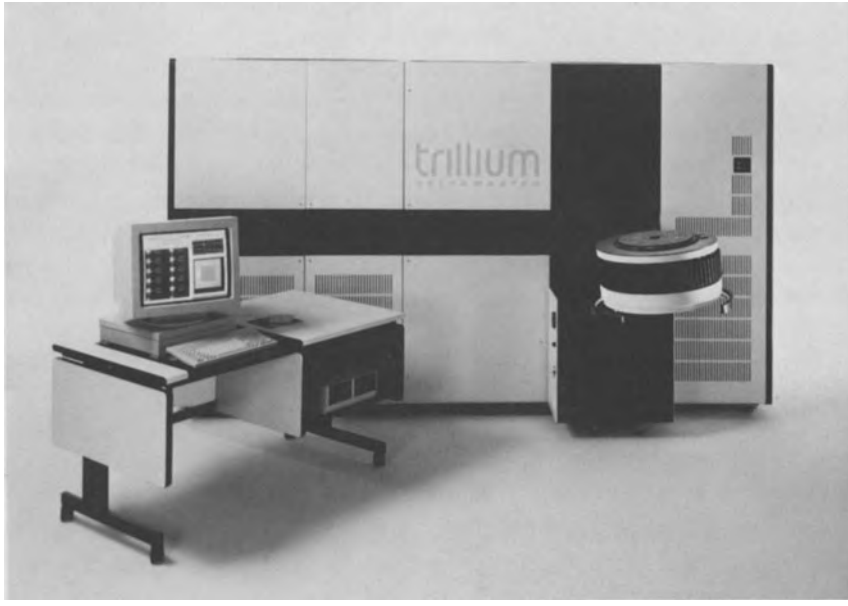

Figure 13-11 LTX/Trillium® Deltamaster™ functional tester capable of 256 pins with data rates of 160 MHz and pattern depth up to 128 million words (Courtesy of LTX Corp.)

tester configured to test packaged parts is shown in Figure 13-11. If no carrier or package is to be used, then the final testing must be performed on the IC in wafer form.

Final testing includes a full complement of tests. The part is first tested functionally at its full rated clock speed. If it does not pass this test, the clock speed is typically reduced and the test repeated. The parts are sorted according to their maximum clock speed to provide varying levels of performance. A full parametric test is performed then, and again parts may be sorted according to the results. Rejected parts often go on to diagnostic stations to determine the exact cause of the defect. The defect data can be transferred back to the manufacturing and design area to help improve the yield.

13.3.9 Burn-In

Burn-in is the process where the ICs are subject to accelerated aging conditions for many hours. A certain percentage of the ICs will fail during burn-in due to

infant mortality (a term used in testing for early operational failure). These failures result from manufacturing defects, including gate oxide pinholes in CMOS parts, for example or photoresist or etching defects resulting in poor geometry and contamination [24]. The need for burn-in is related to the maturity of the IC fabrication process. New IC designs which utilize smaller design rules or new process steps are more likely to fail during burn-in. As the process is refined, the failure modes are analyzed and process steps corrected, reducing the number of failures at burn-in. In some cases, the failures at burn-in can drop low enough to warrant eliminating the burn-in step. In fact, one recent study has shown that in some cases burn-in creates more failures than the actual defects that it detects [25].

Burn-in is performed inside of an environmental chamber. The ambient temperature in the chamber is elevated to provide accelerated life testing. The IC can be operated in either a static or dynamic mode. In static burn-in, the IC is powered and static loads placed on the outputs. After a period typically ranging from 48 hours for some commercial parts to hundreds of hours to meet military specifications [26], the IC is removed from the chamber and again subjected to a full functional and parametric test. The power supply current to each IC is monitored during burn-in and if it changes radically, the IC has failed and no further testing is needed. In dynamic burn-in, the IC is exercised by supplying a limited number of input vectors and monitoring of the output states. If the outputs or current consumption changes from the expected results, the IC has failed. In both static and dynamic burn-in, the power supply voltage may be increased above nominal to accelerate the aging of the device.

Burn-in systems resemble large ovens with racks for several burn-in boards. Each burn-in board holds sockets for a number of individual IC packages. These sockets are similar to the sockets used in a functional tester. They must be rugged enough to tolerate the high temperature environment in the chamber and permit easy insertion and removal of the part without pin damage. Burn-in boards are customized for each product type since the pin-out and burn-in requirements vary. Depending on the burn-in test to be performed, the ICs are wired in parallel with some monitoring lines from each IC brought out to the edge of the burn-in board. Since many parts are burned-in at one time, the burn-in system has large power supplies and fuses to isolate each burn-in board or individual IC. If one device shorts during burn-in, its fuse will blow and the rest of the parts can continue to operate.

Burn-in is perhaps the most troublesome problem for ICs designed for use in MCMs. While it is possible to perform a full functional test on the IC in wafer form if required, wafer level burn-in or the equivalent is not yet a reality. Instead, if burn-in of the IC is required, it must be done in the bare die form

using a carrier. In prototype MCM designs, the solution has been simply to eliminate the burn-in step for the die and substitute a burn-in (and possible subsequent repair) of the assembled module. While this may be acceptable for prototyping, it is unlikely to be viable economically for full scale production. While the subject of bare die versus module burn-in is still being debated, there are several arguments that suggest bare die burn-in will be required.

Perhaps the most persuasive argument for bare die burn-in relates to the nature of functional testing to be performed on the MCM. In many MCM designs, the final testing will verify only that it has been assembled correctly and operates at its rated speed. It is unlikely that a complete test of each die will be performed after MCM assembly since the access to each device I/O is limited. For comparison, in PWBs the most common final test is an in-circuit test which confirms that each IC is present and connected correctly. A functional test may also be performed from the edge connector of the board, but this test does not exercise each IC fully. It is extremely difficult to obtain high fault coverage for every IC from the edge of the module. Since the burn-in process stresses the internal logic of the ICs, it is essential that a full functional test be performed on each IC after burn-in (and before installation on the module) to ensure that no faults have developed.

In addition to the limitation imposed by functional test requirements, MCMs, by their very nature, integrate a range of components. An MCM may contain some CMOS devices as well as some bipolar or ECL devices. Each of these technologies will have it own burn-in requirements. Burn-in of the full module may not meet the needs of all of the die on the module. Even within a single logic family, individual part types may have differing burn-in needs related to the maturity of the product. A complex microprocessor may need a longer burn-in than a simpler part with more relaxed design rules. Finally, the ICs that find their way into MCMs tend to be the newest and highest performance products. These devices often use new technologies that usually have a higher fallout rate than for more mature technologies. As a result, burn-in is even more critical for these ICs than for lower performance devices found in single chip packages

The most promising solutions to bare die burn-in are wafer level reliability and wafer level burn-in. Wafer level reliability subjects the die and test structures on the wafer to a range of tests designed to provide stress aimed at likely failure modes. Typical tests include elevated temperature, voltage and output currents. Increasing the temperature and current will cause shorts or opens in improperly formed structures. High voltage will cause electromigration and breakdown in gate oxides and junctions. Wafer level reliability is really an extension of standard quality control procedures, stressing the ICs in the most revealing way possible. Wafer level burn-in, on the other hand, seeks to move the burn-in process back to the wafer. The entire wafer is burned-in, eliminating

the need to handle and to socket bare die. The individual ICs are connected together on the wafer using additional layers of metallization. The wafer is then contacted using some sort of socket or probe card and placed in the burn-in system. After burn-in this metallization is either stripped off or disconnected in some way. This is not yet a commercially available process and those IC manufacturers pursuing wafer burn-in employ their own proprietary methodology.

Despite the technical problems, burned-in bare die are available from some IC manufacturers. The methods used are proprietary, but utilize wafer level reliability techniques and in some cases, wafer burn-in. Burned-in bare die are also available from vendors who mount the die on carriers to perform the burn-in. The die is supplied either on the carrier or it is unmounted if required. This is done most commonly for the military market where the cost of the finished MCM is high and the volume is small. At the current time, the supply of burned-in bare die is very limited and the price is much higher than that of fully packaged die. As MCM volume increases, it will become essential that known good bare die be available at a reasonable cost. It is likely that this goal will be met only by a combination of quality control and wafer level testing which can guarantee defect-free die without the need for further test or burn-in.

13.3.10 Summary

MCMs present some additional requirements to the standard methods of test and acquisition of ICs. The use of bare die on the MCM necessitates that full functional and parametric tests be performed on the die in unpackaged form. While this is not the current practice, it is possible technically and only needs the demand from MCM fabricators to make it happen. Burn-in, however, presents some additional difficulties. There is currently no widely accepted method of performing burn-in on bare die, and this will continue to have a significant depressing effect on the introduction of multichip modules. Once a technology is in place to produce known good bare die at a reasonable cost, the yield of MCMs will increase and their ultimate cost will be reduced. In the interim, die carriers can be used to permit fixturing of the die during burn-in and thus enable the MCM fabricator to obtain reliable functional parts. Users of ICs for MCMs also are more likely to require test facilities such as boundary scan and BIST which make it easier to test completed modules. These points will be discussed in the next section.

13.4 ASSEMBLED MODULE TEST

13.4.1 Introduction

After the die and other components have been mounted onto the MCM substrate, the completed module must be tested fully to ensure that it meets the design

specifications. This testing includes both functional testing and measurement of AC and DC parametrics, similar to that performed on each individual die mounted onto the substrate. The module may also be subjected to a burn-in procedure to locate any defects in the assembly process or in the substrate itself.

MCMs can be viewed as either large, complex components or as small, high performance subsystems. The testing process used is influenced strongly by the way that the MCM is considered. Since most MCMs are destined to be plugged in or mounted to a PWB, the system designer will view the MCM as a component. As a component, the defect rate must be as low as that of other semiconductor components on the board. On the other hand, the MCM contains a number of individual components and may be repaired by diagnosing the fault and replacing the offending component. MCMs thus have a *split personality*: while they are being fabricated, it is possible and perhaps even essential that individual components be replaced as needed. Once the MCM is complete and sealed, the end user has no way to repair it. This leads to a testing methodology which fits into the dual nature of the MCM.

Fabrication of MCMs requires known good components. As previously mentioned, diagnosing faults as early in a process as possible invariably leads to the least cost. The assembled MCM, while still at the manufacturer, needs to be treated as a very small and complex PWB. Test equipment, diagnostic techniques and repair methods will be based on this requirement and focus on functionality. Prior to shipment to a system house where the MCM will be used, it will require testing to verify its quality as a component. This will include high speed functional testing, parametrics and burn-in.

This section reviews the methods that can be employed to test assembled MCMs along with a guideline to design for testability. The module test deployment sequence is covered from test generation through bring-up to production test. A reference covering design for test, boundary scan, built-in self test, modeling and fault simulation is available [18]. While this reference does not cover MCMs specifically, its concepts provide an excellent introduction.

13.4.2 Testing Strategies

There are a number of different test strategies that can be employed favorably to test the assembled MCM. The selection of the testing method depends on a number of factors, including quality of the individual die and substrate, acceptable failure rate at the customer's site, ease of fault diagnosis and cost of repair. The most aggressive testing program subjects the module to a full test of the internal logic of each individual die on the module. A greatly reduced

level of testing just verifies that the components have been assembled properly onto the substrate. A review of these testing options will enable the MCM fabricator to select the most desirable methodology.

Full Functional Testing

From a fault coverage standpoint, the ideal module test would combine all of the tests for the individual components into one large test. Running this test would guarantee the functionality and performance of all of the internal logic in the module. The problem with this testing concept is that it may be extremely difficult and time consuming to generate a test which can exercise fully all of the internal logic from the module I/O pins. The test itself could have an enormous number of test vectors which would lengthen the test time and add to the testing cost. Despite these problems, there are some MCM designs that lend themselves well to this testing methodology. An MCM consisting of a microprocessor and SRAM in a bus arrangement can be tested by combining the tests of the microprocessor and SRAM together as long as the SRAM is accessible from the module I/O pins. If the SRAM is connected only to an internal bus, then the microprocessor can be used to test the SRAM, although this requires creating a new combined test. A full functional test of all components in a module has the advantage that it can detect logic defects that might have been created during the assembly or burn-in process.

Limited Functional Testing

If the assumption is made that the primary defects that show up during module test are assembly errors, (an assumption often made in the PWB industry) then a simpler testing method can be employed. Simulation of the MCM can be used to generate a test that provides a high degree of fault coverage for faults related to the individual chip I/Os. This limited functional test is much reduced in scope from the full functional test, but require much less effort to create it. In some cases, adequate coverage can be obtained by modeling the way the ICs interact without considering their internal behavior or gate level logic. Prior to limited functional test, passive testing can be used to make sure that no module I/O pins are shorted together.

Boundary Scan and BIST

If the chips on the module incorporate boundary scan features, then it is possible to create a test that detects all opens and shorts at the chip I/Os without the need to understand the functionality of the ICs [24]. Boundary scan is useful particularly in MCMs since it permits access to the chip I/Os with the addition of just four pins on the module. Figure 13-12 shows the boundary interconnection of several ICs on a module. Opens and shorts testing can be

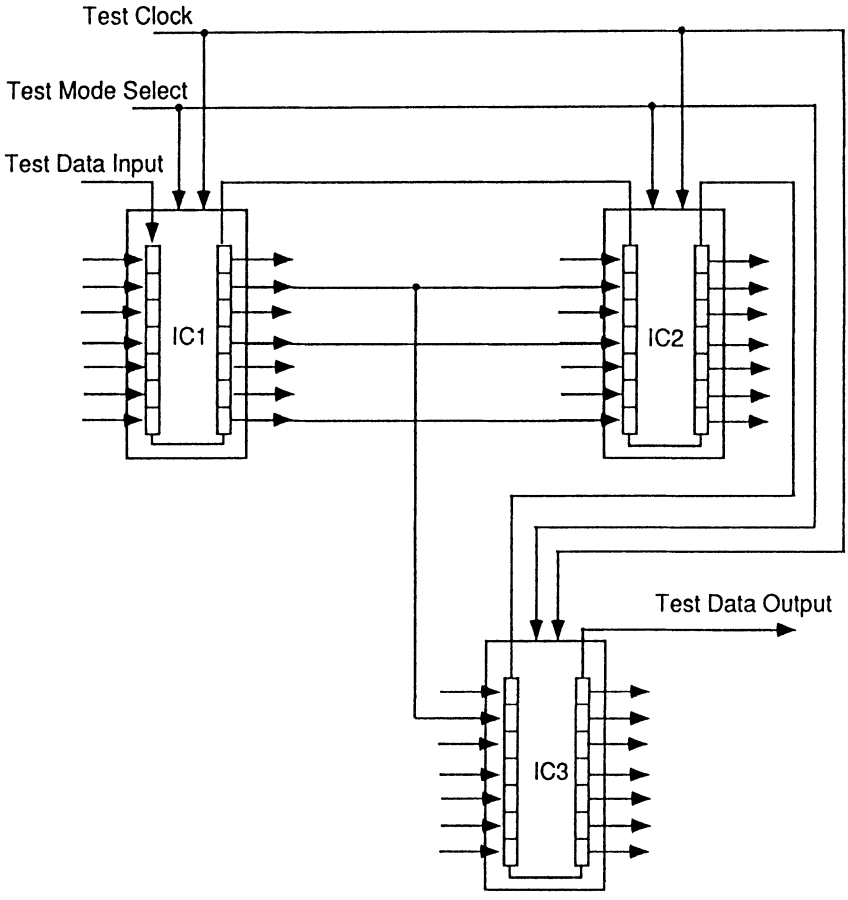

Figure 13-12 Boundary scan interconnection of several ICs on a module.

accomplished by applying an input test pattern to the TAP and by using the boundary scan chain to shift a known state to a single chip output [29]. The state at the output is transferred to all of the inputs on the same net through the substrate interconnect. The boundary scan chain is shifted out then, and examined to see if the pattern matches the expected result. This process is repeated for each chip output on the MCM, creating a walking bit type test. This test can be quite lengthy since the number of input bits = $(n_o)(n_t)$, where n_o is the number of chip outputs and n_t is the total length of the boundary scan chain.

The test time can be reduced by utilizing algorithms which use an adaptive technique to reduce the number of tests [30]. The connections from the module pins to the chip I/Os can be tested by applying vectors at the module pins, clocking the states that appear at the chip I/Os into the boundary scan registers, and then serial shifting the results out the TAP. In the case where an MCM contains some ICs with boundary scan and some without, a combination of boundary scan and limited functional test is required to test the interconnects fully.

The boundary cells are located near the edge of the die so the cells themselves are particularly vulnerable to mishandling of the die. Just the fact that the boundary scan path is intact provides some measure of the quality of the assembly. Boundary scan can be used also in a module to perform an IC level functional test. The test vectors which have been designed for parallel application are converted to serial form, shifted through the boundary to the correct chip and then clocked into the internal logic of the IC. Then, results from the chip outputs are shifted out in the same manner. This is useful for detecting any failures in the IC level logic which may have occurred as a result of assembly errors. Internal testing of the logic can be enhanced if the chips contain BIST features. In this case, the boundary scan path can be used to initiate and further to read the results of BIST [31].

Final Performance Test

MCMs are often considered for those applications requiring very high speed operation due to the short distances between die and the low substrate parasitics. If the substrate has been fully tested and characterized and the MCM assembled from die which meet the required high performance specifications, then no at speed testing of the MCM should be required. Unfortunately, there are many factors which can contribute to an MCM not running at its rated speed. Timing skew between chips, ground bounce, crosstalk between lines, variations in loading on the IC outputs and errors in the assembly process can combine to reduce the performance of the MCM. It is for this reason that a final test of the module at rated speed is required. The AC and DC parametrics are tested also to ensure that the module performs up to its expectations.

A final performance test of the MCM is similar to the limited functional test used to diagnose assembly defects. In the performance test, the vectors selected are designed to make sure that all ICs on the substrate can communicate together properly and that the MCM can communicate at its rated speed with external components. An MCM may pass a limited functional test and fail its at-speed or parametric test. Similarly, the at-speed test may not provide full coverage of all open, short and stuck-at faults for the internal chip I/Os.

The proper choice of module tests will depend strongly on the quality of both the incoming components and the assembly yield. In keeping with the dual nature of the MCM, a limited functional test or boundary scan test will be performed after assembly. Any defects in assembly are repaired and then the module will undergo a high speed functional and parametric test as a final qualification. While these two tests could be combined into one single test, it may be more appropriate to utilize a PWB type tester for the assembly test and an IC type functional tester for the final test. Board testers have many features designed to aid in isolating a fault to an individual component, as will be discussed later. Table 13-2 presents a comparison of the testing strategies available for module test.

13.4.3 Design for Testability

It is important to consider the testability of any product designed for volume production. This is particularly true of MCMs due to their complexity and small size which limits access to their internal components. Changes in design can have significant impact on the ease of testing. The designer should incorporate as many testability features as possible while still providing the required performance. A list of design for test features for MCMs is given below in the approximate order of desirability:

1. Design the circuit with modularity in mind. Breaking the circuit down into functional units at design time will lead to more modular test programs and simpler fault isolation.

2. Bring as many internal chip I/Os as possible to the module pins, even if not needed to operate the module. This reduces the problem of limited access since the internal nodes become external nodes. The selection of which I/Os to bring out is made by evaluating any improvements in fault coverage, and the ease in test generation and fault diagnosis. This simple concept is overlooked often in the rush to make the MCM design as elegant as possible, although many designs can tolerate the additional trace length and extra module pins.

3. Use ICs that have outputs which can be placed in a high impedance state. The use of ICs with tristate outputs permits one chip to be isolated from all the others by simply turning off the other ICs. This is a particularly good design for test feature when combined with Item 2 above, since it makes it possible to test a chip with its standard IC test vectors isolated from the other ICs.

Table 13-2 Comparison of Test Strategies for MCMs.

Test type	Test IC Logic	Test Interconnects	Test At-speed	Test Module Pins	AC Parametric	Fault Diagnosis Level	Physical Access
Full Functional	yes	yes	yes	yes	yes	logic level	module pins and test pads
Limited Functional	some	yes	yes	yes	no	die level	module pins
Final Performance	some	no	yes	yes	yes	die level	module pins
Boundary Scan Interconnect	no	within module	no	no	no	interconnect	module TAP
Boundary Scan + Module Pins	no	yes	no	yes	no	interconnect	module pins
Boundary Scan + Internal Scan	some	within module	no	no	no	partial logic level	module TAP
Boundary Scan + Chip BIST	some	within module	within chip	no	no	partial logic level	module TAP
Module Level BIST	some	some	within module	some	no	die level	module pins

4. Provide test access pads where possible for internal chip I/Os. This is similar to Item 2 above, but requires only that accessible test pads be added to the substrate instead of using more module I/O pins. The test pads can be accessed using a separate probe card or a manual probe. The test pads can be used for diagnostics or they can be part of the standard testing sequence. Once the module is complete, the test pads may be hidden by the package lid or heat sink, if desired.

5. Use ICs equipped with boundary scan and BIST. Boundary scan is the next best thing to having physical access to the internal nodes. Assembly defects which cause opens and shorts can be detected easily and the boundary can be used for diagnostic probing of an internal node in the event of a failure in the functional test. BIST allows high performance testing of IC level logic and can be used to make sure that the chip still operates at its rated speed.

6. Incorporate a boundary scan controller on the module where appropriate [32]. Additional chips can be placed on the MCM to control testing and to aid in fault diagnosis. In addition to running boundary scan tests, a module level BIST controller can perform a variety of test functions. If a microprocessor is present on the MCM it may take the place of a dedicated BIST controller.

13.4.4 Test Generation

Simulation

The first step in generating a test for a MCM is to create a computer simulation of the circuit. The simulation is used to model the behavior of the circuit in response to input patterns. The simulator calculates the output states of the MCM which form the basis for a functional test. AC and DC parametric tests are generated by combining the specifications for each IC used in the module to calculate expected values such as propagation delay through the MCM or quiescent current. It should be noted that the test generation process is accomplished best by the designer of the MCM who possesses an intimate working knowledge of its functions and its intended uses. The MCM fabricator needs to rely on the designer to lead the test generation process. The only exception to this is the generation of an assembly test for MCM parts designed with ICs utilizing boundary scan. These MCMs can be tested for interconnect opens and shorts without any knowledge of the IC functionality. Of course, a final test of the MCM functionality will still require design knowledge.

Simulation programs typically run on workstations and may require hours of computation time to provide an acceptable level of coverage. Common simulators include VERILOG™ by Cadence Design Systems, Inc. and QUICKSIM™ by Mentor Graphics, Inc. These same simulators are used to model the performance of individual ICs. Simulators typically represent an IC by modeling either its logic elements (AND, OR etc.), its overall behavior including timing or just its bus level interaction. In a logic level model, the response of each gate on the IC is modeled. This provides the most detailed simulation, but it is not required usually for MCM design. A behavioral model calculates the response of the IC to input patterns using an algorithm appropriate for the device. For example, it is quite straightforward to write a small program to model the behavior of a memory IC since it acts the same way as an array in a high level computer language. Models for more complex devices can be much more difficult, but they still only require the data book for the device, since this book fully describes the device behavior. The simplest models are known as bus level models. These models are created for ICs which operate on a bus and deal only with the communication to and from the IC. This can be quite useful for MCM test generation if the primary goal is to make sure that the ICs can run at full speed. It does not, however, provide very good fault coverage for opens and shorts at the chip I/Os. Both behavioral and bus models are offered by Logic Modeling Corp., for a wide range of ICs. In some cases, the IC manufacturer may be willing to supply the needed models.

Hardware Modelers

If a software model is not available, a hardware modeler can be used in its place. Instead of relying on software, the hardware modeler uses an actual IC to determine the next state. Most simulators have interfacing available to connect to a hardware modeler. Input patterns from the simulator are sent to the hardware modeler which contains a packaged IC mounted on a board inside of the unit. The modeler can hold several ICs at the same time, each on a separate board in the unit. The IC processes the input vectors and returns to the simulator the output vectors while operating at the clock speed of the modeler. Hardware modelers are an excellent method of running simulations when the software model is not available yet. One disadvantage to using this method is that the response returned from the device is not the typical response, but is specific to that copy of the part. One could create a simulation based on hardware modeling and then find that it does not work correctly due to manufacturing tolerances of the ICs. In addition, some ICs designed specifically for MCM use may not be available in a single chip package.

Timing

Timing information for the ICs must be included also in the simulation. This is available from the data book for each IC or may be part of a purchased hardware or software model. Since MCMs are designed for high performance, the timing information is critical to obtaining a reliable simulation. Timing information allows the simulation to check that input and output waveforms are present when expected. The propagation delay of a signal traveling through the substrate interconnect can be added to the model to provide a more accurate timing simulation.

Input Vectors

Perhaps the most difficult simulation task for the MCM designer is the creation of a set of input patterns for the MCM model. Just creating a model of the MCM does not provide test vectors. A set of input conditions is required also. The goal is to create a set of input patterns that will exercise the ICs on the MCM to obtain adequate fault coverage. The determination of adequate coverage is made by the system designer, but corresponds to the most likely set of failure modes. In many instances, the MCM contains a microprocessor, and responds to programmed instructions. In this case, input vectors can be created by writing programs designed for the MCM and capturing the binary form of the compiled program. If the MCM performs a simple function such as static memory or digital switching, an algorithmic approach can be used to create a program for generation of the inputs. A complex MCM will require the combination of several different approaches to create the input vectors. A sound philosophy to this process is to break up the functions of the MCM and then generate vectors designed to exercise each of these functions. This process is facilitated greatly by employing design for test strategies as outlined in the previous section.

Fault Simulation and Coverage

The test coverage level provided by the input patterns is obtained by observing the operation of the simulation in the presence of faults introduced into the model. For a fault to be observable, it must be propagated to an MCM output pin where it changes the expected output state. Functional tests designed to uncover assembly defects use a set of inputs that causes each internal node on the MCM to change state and propagate this change to the MCM pins during the test. An at-speed test might have less node coverage, instead emphasizing the interchip communications. A full functional test that tests the logic of each individual chip will require enough input vectors to observe any internal fault from the MCM pins. This can be extremely difficult and it is not expected to be used as a general method.

The level of fault coverage can be determined during the simulation by inducing faults into the internal nodes of the MCM and then by observing if the output vectors have been changed. Opens, shorts and stuck-at faults can be added easily to the simulation, and the effect of each noted. Fault simulation can, however, be an extremely time-consuming process. As an example, if an MCM has 500 nets, then the simulation will need to be run 500 times to introduce a single stuck-at fault into each net. If multiple faults are to be simulated, the number of simulations required increases dramatically. To simulate a stuck-at fault in just two nets will require (500 (499) / 2 = 124,750 simulations. The computation time for this would be excessive. The designer must rely on knowledge of the likely failure modes to choose an acceptable level of fault simulation.

In addition to determining the level of fault coverage, a fault simulation also can be used to create a fault dictionary. The changes in the output patterns are recorded as a function of the induced fault. This dictionary can be used by the tester to isolate a fault by comparing the fault dictionary to the actual vectors measured on the MCM under test. Unfortunately, if an MCM has more than one fault, the test vectors collected may not match any single entry in the fault dictionary. In this case, some testers employ a best fit approach to selecting the most likely cause of failure based on the test results.

Once the simulation has been run and fault coverage determined to be adequate, a test program specific to the test equipment must be generated. This can be done with the assistance of translator programs which convert the output of popular simulators into the specific format needed by a wide range of test equipment.

13.4.5 Test Equipment

Selection of test equipment to be used for the module test is dependent on the nature of the test to be performed. In the past, a variety of test equipment has been used to perform the functional test. Functional testers can be divided into two broad categories: board testers and IC testers. Board testers have high pin counts and diagnostic features including the ability to handle fault dictionaries and provide a trace back method for locating defects on a PWB using a manual probe. IC testers are configured to test single devices and can run at much higher speed with greater timing accuracy, but their per channel cost is much higher than a board tester. IC testers usually can handle test programs with long sets of test vectors, often one million vectors or more. Since performance has been a prime motivation in MCMs, IC testers have been used often for functional testing. In addition, early MCM designs often provided access to all I/Os of all chips through pins on the module, permitting each chip to be tested the same way as was done for single chip packaging. As the industry develops, it is

expected that MCM test will require a two step approach. A low cost board type tester will be used to locate assembly defects using either limited functional testing or boundary scan. An IC type tester will be used then to check the at-speed performance and measure the parametrics.

Fixturing of the MCM to the ATE can be more difficult than for a single chip package or PWB. MCMs tend to have high pin counts and unusual lead arrangements. If the MCM is designed for solder installation, then a special test socket will be needed to hold the MCM without damage to the leads. The high speed of many MCMs also presents some challenges to fixturing. It may be necessary to bring controlled impedance lines to the MCM to terminate its I/O pins properly. This is practiced already for ultra high speed ICs and is adaptable readily to MCMs. As MCMs gain in popularity, standard packages will emerge making fixturing much easier. For thin film MCMs, it may be desirable to test the module after die attach, but before assembly into the final package. In this case, the modules will need to be placed in carriers and tested using a wafer probe as is done for ICs.

13.4.6 Bring-Up

Bring-up is the process whereby a new MCM design is tested and debugged. The test program generated from the design simulation is installed in the ATE and an appropriate socket wired on the tester adapter board. The tester is then configured for the part by setting up the mapping of test channels to package pins, by setting of the timing parameters and by configuring the power and ground. As with any complex system, there is a good chance that the module will not pass the test program the first time. Problems can include wrong channel assignments in the tester, wrong timing for the module, errors in the simulation, errors in the models and finally errors in the design itself. In addition, the module may contain assembly errors or defective die presenting a very difficult diagnostic problem. Once again we turn to design for testability as the savior of this situation. Strict adherence to modular design, along with features designed to assist in the diagnostic process, can provide a great deal of assistance. Steps to diagnosing a failed module at bring-up include:

1. Make sure that the module pins are mapped correctly to the tester channel pins and that the continuity at the pins of the module is as expected.

2. Run the boundary scan test to make sure that the parts are installed and wired as expected. This can go a long way to revealing errors since both errors in the substrate interconnect and in improper descriptions of the boundary will be revealed from this test.

3. Run a limited functional test which has been organized by the functionality of the module. The first section of the test might just reset all the ICs and check the output states. Diagnose any timing errors at this time.

4. If all seems well in the functional test, but the module does not produce the expected response, go back to the simulation with information on which vectors have failed and locate any errors in simulation.

5. Re-check the performance of relevant software models and make sure the hardware modeler is producing the expected responses.

6. If all else fails, go back to the original substrate and circuit design and make sure that the module will function as intended.

One simple method to verify that an MCM is functioning properly is to plug it into its intended final circuit and see if it works. This is a useful procedure in the bring-up stage, but does not provide adequate fault coverage in production testing. In addition, if the MCM does not work, it is difficult to diagnose the fault. An advantage to this type of test during bring-up is that once the MCM is found to be functioning in a real circuit, it is easier to debug the test program without the uncertainty of a faulty part.

If there is no test program available for the MCM because no simulation was performed, it is possible to capture the test vectors from the MCM operating in a real circuit. The clocking between the circuit and the tester must be synchronized and then the vectors can be captured. This method is not recommended. It is imprudent to design and fabricate an MCM without performing a simulation. The functional tester is operated in a vector capture mode and connected to the circuit in parallel with the MCM. This method has an additional disadvantage. It provides no way to diagnose faults since the intent of each vector in the test is unknown.

13.4.7 Production Test

Production testing of MCMs, like virtually any other manufactured item, will always be a tradeoff of cost versus yield. The more time spent on the final test, the more reliable the final product will be and the higher the yield will be after it is shipped. Testing does have a cost in time and equipment, so the desired quality level must be determined before any production test program can be implemented. A production test will include an initial test for assembly defects, a test of MCM performance as a component and an optional burn-in step.

Since each MCM product will be different in the quality of its components, the level of its performance and the desired final quality, it is useful to provide an ideal case as an example. Assume that the substrate has been tested fully for opens and shorts and that the ICs utilize boundary scan and have been tested fully at-speed and burned-in, if required. After assembly, the module will be tested on a board type tester by testing first for shorts with no power applied. Power is applied and a DC parametric test is performed. A boundary scan test is performed then to confirm that the die are installed correctly and that the substrate interconnect is correct. If the MCM fails at this stage, the test results are used to guide the repair or replacement of a die or in the case of overlay interconnect technology, the replacement of the interconnect itself. This cycle is repeated until the MCM has passed its assembly test. A high speed tester then is employed to perform a limited functional test of the MCM at its rated speed. The AC and DC parametrics also are measured at this time. Failures are diagnosed with the aid of a modular test program and the use of a fault dictionary generated during the simulation. Die are replaced as needed and the module starts at the beginning of the test process again.

Once the module is tested fully, it is then subject to a burn-in if required. Burn-in of a module uses the same equipment and methodology as is employed for single chip packages. A module will fail at burn-in due to a number of possible causes including assembly, interconnect, substrate and die defects. Module burn-in is recommended highly for new MCM products which do not have adequate production experience. If the die on the module have been burned-in and the level of defects in assembly and substrate are small, this step may be eliminated. As with IC burn-in, some failures will be detected on the burn-in unit itself by monitoring selected output pins and power supply currents. If no failures are detected on the burn-in unit, the module is retested to make sure no subtle problem has developed. Modules that fail after burn-in can be repaired and retested. In some cases, burned-in parts will not be repaired since it would require sending the entire module through burn-in again which might shorten the life of some of the components.

13.4.8 Summary

The testing of the assembled MCM presents many challenges to the world of semiconductor testing. Design for testability plays a key role in the success of bringing up a new design and in isolating the faults to a single component. Extensive simulation of the circuit is required to assure adequate fault coverage. Boundary scan will be an essential component in nearly every MCM design in the future since it provides the ability to probe internal nodes of the device without the use of extra test pads or pins. Finally, the MCM must be treated as

both a large component and a small system. The small system nature assumed by an MCM during fabrication requires diagnosis and repair of defects, while the large component nature requires a very high level of quality for the part as shipped to the customer.

13.5 CONCLUSION

Testing of MCMs requires a broad knowledge of technologies. The substrates must be tested using special equipment which can probe the very small pads. There is a tradeoff in testing speed, flexibility and cost, but commercial equipment is available to achieve the desired goals. A successful MCM program requires a source of known good die. At the current time, fully tested bare die are in very limited supply, but as the demand increases, the supply will expand. Burn-in of these die remains a roadblock to the widespread availability of known good die. In the interim period, MCMs can be fabricated using die placed on carriers which ease the test and burn-in problems. Module test is mostly a case of employing a proper design methodology from the conception of the product through to the creation of the test. Good design for test methods can greatly simplify first time bring-up of the device and isolation of faults. Boundary scan and BIST are new technologies that are here to stay in MCM fabrication and their use cannot be emphasized too much. Simulation of the module as a whole is essential also. It is no longer possible to create a schematic and wire a circuit without the benefit of an extensive simulation. While MCM testing is not yet a mature technology, it offers the excitement of a new area of endeavor and the rewards of ensuring the quality of the finished product.

The testing of MCMs presents one of the greatest challenges to their widespread use. Significant improvements are still required in substrate testing speed, quality of bare die, module level test program generation and fault diagnosis. Once these problems have been brought under control, MCMs will move closer to offering cost and performance improvements over single chip packaging in many applications.

References

1 Kryterion 4000 Specifications, *Everett/Charles Test Equipment, Inc.*, Pomona, CA, Document Number 8803-10M, 1988.

2 "Cofired Ceramic Multichip Module Design Guide," *Alcoa Electronic Packaging, Inc.*, San Diego, CA, Rev. 91-08, Aug. 1991.

3 L. J. Webb, "Non-Contact Coupling Plate for Circuit Board Tester," U.S. Patent Number 3,975,680, Issued Aug. 17, 1976.

4 C. F. Spence, "Capacitive Continuity Testing of Multilayer Hybrid Substrates: Its Capabilities and Limitations," *Proc. of the 1985 Intl. Symp. on Microelect.*, pp. 174-182, 1985.

5 R. P. Burr, *et al.*, "Method and Apparatus for Testing of Electrical Interconnection Networks," U.S. Patent Number 4,565,966, Issued Jan. 21, 1986.

6 S. Lin, "Computer Solutions of the Traveling Salesman Problem," *Bell System Technical J.*, vol. 44, pp. 2245-2269, 1965.

7 Michael Garey, David Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, New York: Freeman, 1979.

8 J. A. Conti and P. Hohne-Kranz, "Advantages of Moving Probe Testers Compared to Fixtured Systems," *IPC 31st Annual Meeting*, (Lincolnwood, IL), Technical Paper IPC-TP-679, 1988.

9 J. A. Moran, T. C. Russell, "Testing Process for Electronic Devices," U.S. Patent Number 5,023,557, Issued Jun. 11, 1991.

10 S. Z. Yen, *et al.*, "A Multi-Chip Module Substrate Testing Algorithm," *Proc. of the Fourth Annual IEEE Intl. ASIC Conf. and Exhibit*, pp. 9-4.1-4, Sept. 1991.

11 S. D. Golladay, *et al.*, "Electron-Beam Technology for Open/Short Testing of Multi-Chip Substrates," *IBM J. of Res. and Dev.*, vol. 34, no. 2/3, pp. 250-259, Mar./May 1990.

12 O. C. Woodard Sr., *et al.*, "Voltage Contrast Electron Beam Tester for Testing Unpopulated Multichip Modules," *Proc. of the 1990 Intl. Symp. on Microelect.*, pp. 370-377, 1990.

13 Microelectronics and Computer Technology Corp., 12100-A Technology Blvd., Austin, TX, a consortium of high technology companies, has hosted several workshops on the acquisition of known good die.

14 IEEE Standard 1149.1-1990, IEEE Standard Test Access Port and Boundary-Scan Architecture, New York: IEEE Standards Board, May, 1990.

15 L. Whetsel, "Event Qualification: A Gateway to At-Speed System Testing," *Proc. of the 1990 Intl. Test Conf.*, pp. 135-141, 1990.

16 B. Koenemann, G. Zwiehoff, and R. Bosch, "Built-In Testing: State-of-the-Art," *Proc. of the 1983 Curriculum for Test Tech. Conf.*, pp. 83-89, 1983.

17 E. J. McCluskey, "Built-In-Self-Test Techniques," and "Built-In-Self-Test Structures," *IEEE Design & Test*, pp. 21-36, Apr. 1985.

18 M. Abramovici, M. A. Breuer, and A. D. Friedman, *Digital Systems Testing and Testable Design*, New York: W. H. Freeman and Company, 1990.

19 Parametric measurements are defined in most digital data books. See: *Advanced CMOS Logic Data Book*, Texas Instruments, 1990, pp. 1.6-9 as an example.

20 T. Tada, *et al.*, "A Fine Pitch Probe Technology for VLSI Wafer Testing," *Proc. of the 1990 Intl. Test Conf.*, pp. 900-906, 1990.

21 S. P. Atahn, D. C. Keezer and J. McKinley, "High Frequency Wafer Probing and Power Supply Resonance Effects," *Proc. of the 1991 Intl. Test Conf.*, pp. 1069-1072. 1991.

22 T. C. Chung, D. A. Gibson and P. B. Wesling, "TAB Testing and Burn-in," *Handbook of Tape Automated Bonding*, J. H. Lau, ed., New York: Van Nostrand Reinhold, 1992, pp. 243-304.

23 G. Coors and R. Buck, "Pad Array Carriers—The Surface Mount Alternative," *Proc. of the 1991 Intl. Elect. Packaging Conf.*, pp. 373-389, 1991.
24 D. J. Klinger, Y. Nakada and M. A. Menendez, *AT&T Reliability Manual*, New York: Van Nostrand Reinhold, 1990, p. 44.
25 D. Romanchik, "Burn-In: Still a Hot Topic," *Test & Measurement World*, pp. 51-54, Jan. 1992.
26 MIL-STD-883C—Test Methods and Procedures for Microelectronics, Alexandria, VA: National Technical Information Service, 1983.
27 J. Reedholm and T. Turner, "Wafer-Level Reliability," *Microelect. Manuf. Tech.*, pp. 28-32, Apr. 1991.
28 P. T. Wagner, "Interconnect Testing with Boundary Scan," *Proc. of the 1987 Intl. Test Conf.*, pp. 52-57, 1987.
29 C. M. Maunder and R. E. Tulloss, eds., *The Test Access Port and Boundary Scan Architecture*, Los Alamitos, CA: IEEE Computer Society Press, 1990, pp. 11-21.
30 N. Jarwala, C. W. Yau, "A New Framework for Analyzing Test Generation and Diagnosis Algorithms for Board Interconnects," *Proc. of the 1990 Intl. Test Conf.*, pp. 565-571, 1990.
31 P. Hansen, "Taking Advantage of Boundary-Scan in Loaded-Board Testing," *The Test Access Port and Boundary Scan Architecture*, C. M. Maunder and R. E. Tulloss, eds., Los Alamitos, CA: IEEE Computer Society Press, 1990, pp. 81-96.
32 C. W. Yau and N. Jarwala, "The Boundary-Scan Master: Target Applications and Functional Requirements," *Proc. of the 1990 Intl. Test Conf.*, pp. 311-315, 1990.

Part C—Case Studies

Alice went timidly up to the door, and knocked.
"There's no sort of use in knocking," said the Footman,
"and that for two reasons:
First, because I'm on the same side of the door as you are:
secondly, because they're making such a noise inside,
no one could possibly hear you."

"How am I to get in?" Alice asked again in a louder tone.
"Are you to get in at all?" said the Footman.
That's the first question, you know."
"I shall sit here," he said, "on and off, for days and days."

"Oh, there's no use in talking to him," said Alice desperately,
"He's perfectly idiotic!"
And she opened the door and went in.

Alice in Wonderland
by Lewis Carroll

How are all the elements pulled together to create MCM-based products? The intent of the following chapters is to provide reports from some companies with successful multichip products. Some of these products are commercial in that they compete with non-MCM-based products in applications sensitive to cost. In each case, the MCM program was developed internally from existing expertise. In these case studies, the authors deal with issues that created real and imaginary barriers to considerations of MCM technology.

For the first time, some of the leading companies that have produced MCM-based products tell their **own story.** The list of companies included is not meant to be complete; *it is meant to be a sample.* In particular, the excellent MCM-based products developed at AT&T, Fujitsu, Hitachi, IBM, Mitsubishi, NEC, Siemens, to name a few, are missing from this Part. Case studies also could have been written by MCM vendor companies (that make bare and assembled MCM substrates), such as Alcoa, Cypress, Kyocera, Pacific Hybrid Microelectronics, NTK and TI. At the time of writing there are over

30 MCM vendor companies which also offer design services to help their customers realize their first MCM product. An even larger number of companies are gaining experience by building MCM prototypes.

In this part of the book, several leading companies describe how they designed and implemented complete MCM products. From these chapters, managers will gain the insight needed to make crucial technology decisions and to organize programs for an MCM-based product or series of products. System engineers will learn what the critical product decisions are and how to make them. Design and manufacturing engineers will see how the concepts discussed in the previous two parts of the book are put into practice. All of these people will be able to understand how concurrent engineering is applied to MCM product development. Marketing personnel will use these case studies to relate packaging alternatives to the needs of their customers.

For the first time, this part provides the reader with:

- Detailed descriptions indicating how some systems houses have made critical MCM-related packaging decisions which best meet their cost and performance goals. For example, Chapter 14 discusses how Unisys made critical technology decisions for a number of its computer products. Chapter 17 shows how Digital Equipment Corporation decided on the details of the MCM technology that was used in the VAX-9000 computer.

- Complete presentations of the end-to-end process of conceptualizing, designing and manufacturing MCM-based products. These descriptions put many of the issues, such as yield and testability, that concern potential MCM product developers, in their proper perspective. Authors also discuss unexpected problems that arose and the solutions they found to them.

- An analysis of the packaging needs of products for aerospace and military applications and how MCMs can be applied as a solution.

- An overview of the most popular thin film MCM technology, silicon on silicon, with a description of the tradeoffs and the range of applications for which this silicon-based technology has been used.

- A full presentation of the manufacturing process for a high end, thin film MCM product, tying together many of the MCM manufacturing issues raised in previous chapters.

The main message conveyed by these chapters is that ***MCMs are real*** and that products can be made. The complete range of potential MCM types are covered in these chapters, from a simple ceramic package to perhaps the most advanced thin film MCM found. These chapters show the reader that not only is the technology accessible and beneficial to many electronics product manufacturers, but also that it is possible to use the high end of the technology to gain significant systems advantage.

14

THE DEVELOPMENT OF UNISYS MULTICHIP MODULES

John A. Nelson and Randy D. Rhodes

14.1 INTRODUCTION

Do mainframe computers really need multichip modules (MCMs)? Won't advancements in silicon technology alone be sufficient for next generation computers? What are the challenges in designing MCMs? Can we design future computers with current packaging technology? These are some of the questions this chapter will answer.

One way to begin to answer these questions is to look at trends in computer packaging at Unisys. A mainframe computer consists of many components. The major elements are the instruction processor, the input/output (I/O) subsystem and the mass storage subsystem. All three of these elements utilize semiconductor technology. However, the instruction processor is the brain of the computer where the most demanding packaging challenges originate.

The largest Unisys instruction processor of the late 1970s contained nearly 25,000 semiconductor packages residing on 240 printed wiring boards (PWBs). In addition to long electrical signal paths between computing elements, 400,000 solder joints were required. Major advances in silicon technology permitted nearly a seven to one reduction in the packages and boards for the mid-1980s Unisys computer. Specifically, 3,900 packages were contained on 29 PWBs. Silicon was faster and electrical paths shorter, resulting in significantly increased

performance. Reliability also increased because the number of solder joints was reduced to 117,000. The 1991 instruction processor in this family utilized silicon advances along with MCMs to achieve an even larger gain. Only 187 packages on four PWBs were required. Performance was over twice that of the previous generation. Big gains were made in reliability because solder joints were cut to approximately 35,000. The latest Unisys instruction processor, introduced in late 1991, utilized double sided multichip packages and leading edge PWB technology to get down to a single board processor.

How can you do better than that? The whole instruction processor on one PWB! What about putting two to four processors on one PWB? The drive to improve performance, reliability and cost continues. This chapter will preview the type of packaging technology that will permit two to four mainframe instruction processors per single PWB.

Multichip technology doesn't come easily! Electrical requirements, high power densities, complex assembly challenges and lack of an industry infrastructure are all issues to be addressed. This chapter will also provide some insight into the decision making process necessary to resolve these issues.

Finally, this chapter will address the significance of these packaging trends for mainframe computer manufacturers. Mainframes are a modest growth segment of the computing market. The most optimistic predictions result in growths of less than 10% per year. Workstation supporters come up with less than that. Modest growth and drastic parts count reductions result in underutilized factories. It is an industry problem. During 1991 we have seen competing companies willing to sell their excess packaging capacity to their competitors. Are silicon and packaging technologists putting themselves out of business? An article in the *Harvard Business Review* suggests that this may be the case [1]-[2]. However, there is a large market segment out there waiting for the right time to apply MCMs to their products. Exciting work in this area will face packaging engineers as this time arrives.

14.2 THE DRIVING FORCES BEHIND MULTICHIP PACKAGING

The development and application of multichip packaging is primarily driven by the interrelated requirements of performance and density [3]. As silicon performance has increased, the delay associated with the propagation of a signal from one silicon die to another has moved from the status of insignificant to become a limiting factor in system performance. In the past decade, this packaging delay moved from being a very small percentage of the clock cycle

time to requiring nearly 100% of the clock cycle. The frequency at which this occurs depends on the number of devices, fanout required, etc.

The term packaging delay as used in this chapter warrants clarification. This delay consists of three major parts. The first is the delay of the output buffer that is required to drive the signal leaving the die. This buffer performs no logic function, therefore, its delay is strictly related to the buffering of the signal as required by the packaging environment.

The second part of the packaging delay is the delay that results from the total interconnect path from the output pad of the source die to the input pad of the destination die. This includes the interconnection of the die to the package, the package interconnect, the next level of interconnect (usually a PWB), the delay of the receiving package interconnection and finally the delay of the package to die connection for the receiving die. Note that this is the description of the delay for a single chip packaging environment.

The third part of this delay is that of the input buffer required to receive the incoming signal. Like the output buffer, this input buffer does not perform any logic function. It is necessary to terminate the incoming signal and isolate the internal nets of the receiving die.

In the case where this packaging delay reaches nearly 100% of the clock cycle, it is not possible to perform any logic function and a die crossing in one clock cycle. An entire cycle must be used to propagate the signal from a register in the source die to a register in the receiving die. This entire cycle is used to move information from one die to another. This cycle is a total loss from a system performance point of view.

14.3 ADVANCES IN SILICON DELAY MULTICHIP MODULES

It must be noted that the advances in silicon density also have helped to limit the impact of the packaging delay on system performance. Even though the packaging delay as a percentage of clock cycle had increased, the number of die crossings (or packaging delays) per logic function have been reduced. This is a direct result of the increase in density, allowing more functionality on each die. Today the increased silicon density allows entire mainframe processors or at least entire processor functions, to reside on one die. In the early 1980s, thousands of die were required. This decrease results not only in a reduced number of die crossings but also reduces the delay itself by shortening the length of interconnect required. This increase in density along with architectural innovations has limited the impact of the packaging delay and delayed the widespread use of MCMs.

14.4 SINGLE CHIP PACKAGES HAVE LIMITATIONS

This same increase in density does result in another issue which favors the use of MCMs. While there is always debate on the relationship of gate count to pin count, the pin count has continued to rise as the gate count increases. As the silicon gets increasingly faster and denser, multiple paths on and off the die will be required, even with MCMs, to provide the data fast enough to maximize the die performance. This increase in pin count per die will make single chip packaging impractical. Single chip packaging gets unattractive above approximately 600 pins, and highly impractical as pin counts near 1000. Array I/O technology will provide relief for the I/Os per die limit. However, the issues with package size, attach to the next level interconnect and the cost impact on the next level of interconnect all contribute to the impracticality of single chip packaging.

The requirement for greater packaging density is the second major driver for MCMs. Small size may be driven by sheer space limitations or by cost considerations for the overall system. Some applications have very strong limits on available space. The most recognized example might be a laptop computer. Actually, all levels of computers have some size restrictions. The allowed size for systems such as personal computers or workstations is basically defined by the application environment where these systems are used. The drive to make them more and more powerful within the same size limitations makes packaging density an issue. Another example is the desire to place powerful digital signal processing capability in satellites, missiles and mobile artillery. Using MCMs would eliminate the necessity to transmit data back to large systems for processing and significantly increase performance of such systems.

Even large mainframes are density driven. Floor space costs for a computer room are quite significant. It is recognized that mainframes have more latitude in dealing with thermal management challenges than do desktop personal computers and workstations. However, reductions in mainframe floor space and power requirements translate directly into cost savings for the user.

The drive to increase density is often related to performance. As packaging density is increased, the physical distances across which signals must be propagated are reduced [4]. This can be translated into higher performance at all levels of the system.

14.5 BASIC PACKAGING GOALS

In any discussion of multichip packaging, one should always review the basic goals. The first goal relates to the drive to increase performance. ***The goal is***

to make the die crossing delay equal to the on die delay. If this goal is achieved then a signal would cross from one die to another with the same delay and electrical requirements as going from one gate to another on the same die. This means that input and output buffers would be eliminated, making packaging a nonfactor in system performance. The system architect could ignore die partitioning in developing the system. The partitioning would then be driven by other cost factors such as optimum die size, pin count, etc.

The next goal relates to the driving requirement of increased density. The ideal packaging design would not add to the size of the die. ***The goal is to have a package to chip area ratio of one to one***. This means that the area required by the package is no more than the die itself. Achievement of this goal would benefit both performance and cost at the system level.

Reliability is always a concern with any packaging and must be considered in any MCM. ***The goal is to improve the reliability of MCMs over single chip packaging.*** The packaging reliability should be insignificant when compared to the silicon itself.

Any packaging must have a cost goal. ***The goal for MCMs is to cost less than the closest equivalent single chip packaging.***

In attempting to achieve the above goals, requirements for thermal performance must be addressed. Thermal performance is a major consideration which requires attention from the very start. Any MCM must be designed to provide the appropriate thermal management in a cost effective manner. ***The goal is to provide the required junction temperatures for the silicon while meeting cost and density goals.*** The junction temperature requirement is dictated by the reliability goals for the silicon and/or the performance goals. (In CMOS applications, performance goals usually dictate the maximum junction temperature.) Thermal management is quite often the limiting factor in MCM development.

14.6 MULTICHIP MODULE DEVELOPMENT PROCESS

14.6.1 Designing Through Technology Change

This section reviews the major decisions to be made when selecting technology for an MCM program. Techniques for evaluating choices are reviewed. Some experiences are shared to illustrate how unanticipated tasks were dealt with on Unisys MCM programs.

Traditionally, for a given silicon technology, transistor content doubles approximately every three years. As a result, the impact of silicon advances on MCM development programs is more significant than for single chip programs.

Single chip development programs can be very complex but seldom as lengthy as a multichip program. This insures that a custom single chip package has a reasonable life before it is made obsolete by advancements in the silicon family for which it was designed.

14.6.2 Develop a Multichip Packaging Strategy First

All multichip programs are designed for a specific silicon family of devices as well. In order to have a useful life, the multichip package should be developed and available well before the replacement silicon family arrives and reduces its cost/performance effectiveness. Generally, the new silicon family does not fit the prior generation MCM because the increased number of transistors are often in a larger die, consuming more power and containing more I/O.

On the front end of a MCM program, it is important to deal with this and decide on a strategy. One strategy is to accept that with each new silicon family a new multichip package development is required. Development time is reduced since extendability is designed in to the new strategy. No time is required to predict the impact of growth on the package features.

The second strategy is to plan for growth and design the multichip package to be able to accommodate larger die, more pins and more power when it becomes available. This approach results in a more complex development program which will require more resources and more time. The benefits are a longer product life, lower cost, fully characterized reliability and greater production efficiencies. The IBM thermal conduction module used in the 3081, 3090 and S390 machines is an example of this approach, where extendability was designed in to permit several generations of silicon to be accommodated [5]-[6]. The recently announced ES9000 system has a new MCM which leverages many of the design features of the earlier TCMs [7]. These systems illustrate the benefits of an extendable technology by accommodating at least four silicon families spanning more than a decade.

The Unisys Single Chip A-series Mainframe Processor (SCAMP), memory array module (MAM) and multiple random access memory (MRAM) MCMs described in this chapter are an example of the first approach. They were designed to serve one generation of silicon. Their development times were relatively short as they were extensions of mature technology.

14.6.3 Technical Decision Making in an MCM Program

There are less than ten major decisions to be made in the development of an MCM [8]. The following ones will occupy a major part of the program.

- How to connect the die to the package: tape automated bonding (TAB), wire bond or flip chip?
- What type of substrate is to be selected: printed wiring board, metallized ceramic, thin film in combination with a base substrate or silicon?
- Is the module to be hermetic or nonhermetic?
- How is heat to be removed: air, liquid or other?
- What is the strategy for test and repair?
- How will the module be connected to the next level of packaging: leaded or leadless, array or peripheral, socket or no socket, other?
- What type of CAD tools are required and how will they be integrated?

There are other decisions but these are more than can be covered in this chapter.

14.6.4 Making Multichip Module Tradeoffs in a Disciplined Way

To develop a multichip package, start by identifying the core technologies available within the industry. A summary of the most common ones is presented in Tables 14-1 through 14-3 titled Multichip Packaging Core Technology. Unisys has applied MCMs in products ranging from desk top to high end mainframes. As a result, Tables 14-1 through 14-3 have three categories of cost and performance for the differing product needs. Tradeoff studies were completed by considering which of the technologies we already had in-house. Other important factors are the performance attributes of the technologies as well as cost, development time, qualification time and risk factor. The amount of support from industry was also heavily considered. It is effective to prepare lists of performance characteristics and limitations to formalize the tradeoff process. Table 14-4 summarizes the limitations of the packaging core technologies from our perspective. These attributes are considered in narrowing the choices.

As a part of the final selection process, a risk analysis needs to be completed. The amount of development risk you are willing to accept is an important factor when opting for new materials and processes in an MCM program. Table 14-5 illustrates a way to look at the hermetic versus nonhermetic decision. During a development process, analysis is performed during the study phase of a new design. Since most MCM packaging concepts are quite complex,

Table 14-1 Multichip Packaging Core Technology Lowest Cost for Performance.

Chip to substrate connection	Wire bond, TAB or COB, low die I/O
Substrate technology	Low technology PWBs with fanout at OLB
Substrate to PWB connection	Solder module directly to PCB, low I/O count <400
Module sealing concept	Nonhermetic plastic encapsulation, usually at chip level
Module cooling concept	Still air or small fan, <5 Watts
Module manufacturing methods	Highest degree of automation due to volume

Table 14-2 Multichip Packaging Core Technology Medium Cost for Performance.

Chip to substrate connection	Wire bond, TAB or COB, higher die I/O counts
Substrate technology	High technology PWB or ceramic or composite, some fanout may be required
Substrate to PWB connection	Solder module directly to PCB or socket up to 700 I/O
Module sealing concept	Nonhermetic plastic encapsulation, usually at chip level
Module cooling concept	Forced air impinge air, 10 Watts or more
Module manufacturing methods	Fairly high degree of automation, modest volumes

Note:

Most of the above building blocks have limitations that are restrictive and limit extendability to higher levels of integration with increasing performance goals. These limits are summarized in Table 14-4.

Table 14-3 Multichip Packaging Core Technology Highest Cost for Performance.

Chip to substrate connection	TAB or flip chip, die I/O too high for repair with wire bond
Substrate technology	Multilayer ceramic of composites, OLB fanout is not required
Substrate to PWB connection	Sockets most common up 2,000 I/O
Module sealing concept	Usually sealed at module level, not always hermetic
Module cooling concept	Liquid impinged air immersion, 100 Watts or more
Module manufacturing methods	Some automation but volumes are not high

Note:
Most of the above building blocks have limitations that are restrictive and limit extendability to higher levels of integration with increasing performance goals. These limits are summarized in Table 14-4.

it is impossible to anticipate all of the surprises one experiences during the implementation phase. A redesign or design tweak is considered a cycle of learning. Note that cycle of learning times are shorter for hermetic designs due to the maturity of industry experience. Furthermore, the number of cycles required is less, as is the sample size in testing programs. The nonhermetic MCMs does not have as much industry data available to reinforce a decision. Therefore, the development program will have to accommodate a larger number of cycles in the schedule.

We are not implying that one is preferred over the other. Rather, the example is designed to show that when you make a decision like this, do it in a disciplined way. Also, it should be noted that there is much interest and work in the area of nonhermetic modules now underway. The IEEE Gel Task Force driven by Jack Balde was an example of industry cooperation to complete evaluations of nonhermetic packages in a shorter time and with less resources overall. Another example is the MCM Substrate Size Task Force, where the goal is to encourage infrastructure participation by standardizing. As stated previously, it is very important to leverage the work done in the industry as a whole.

Table 14-4 Multichip Core Technology Limitations.

Core Technology	Density Limitation	Performance Limitation	Thermal Limitation	Manufacturing Issues
Wire Bonding	I/O count due to peripheral lead/spacing	Inductance of long wire and power distribution	Must remove heat through package	Repair of high I/O not feasible
TAB	Similar to wire bonding	Similar to wire bonding	Not an issue, can be done several ways	Unique lead frame and thermode for each die
Flip Chip	I/O not an issue	Inductance and power distribution not an issue	Not an issue	Major development effort on CAD and die process and assembly
Low Technology PWBs	Lines/cm feature geometries	Fanout at OLB reduces device density	Not an issue	Boards cannot handle temperatures needed for chip assembly
Multilayer Ceramic	Lines/cm feature geometries	Fanout at OLB reduces device density, high dielectric constant	Not an issue	Has upper size limits, long lead time, costly tooling
High Technology PWBs	To be studied	To be studied	Not an issue	To be studied
Composite Substrates of Ceramic and Thin Film	Not an issue, lines/cm approaches silicon	Not an issue	Not an issue	Has upper size limits, long lead time. costly tooling

Table 14-4 Multichip Core Technology Limitations (continued).

Core Technology	Density Limitation	Performance Limitation	Thermal Limitation	Manufacturing Issues
Composite Substrates of Thin Film on Silicon	Not as good as ceramic, lines/cm approaches silicon but layer counts limited by strength	It can only provide peripheral I/O connections, thus it is limited	Must be put in another package that reduces thermal performance	Has upper size limits
Module to Board Connection (hard or soft connect)	Pins per unit of area peripheral or array are only options	Connectors add inductance	Not usually an issue	Soldering high pin count modules into a PWB is a significant challenge
Air Cooling, No Fan	Chips have to be spaced apart	Longer signal runs, more chip crossings	< 2 Watts per device, SCAMP @ 4 Watts used a fan	Lowest cost thermal system
Air Cooling, Fan	Chips can be closer	Some improvement, but still medium performance A-15 class	< 8 Watts per device, A-15 was 5 - 6 Watts	Next lowest cost
Air Cooling, Impingement	Chips almost can be brickwalled	Medium to high performance	Up to 25 Watts per device	Getting to be costly
Liquid Cooling	Chips can be brickwalled	High performance	35 - 40 Watts per device	Very costly
Immersion Cooling	Chips can be brickwalled in three dimensions	Cray, as example	High power density	Very costly

Table 14-5 Package Reliability Assessment Planning Table.

ITEM	HERMETIC		NONHERMETIC	
	Conventional Ceramic Package	**Advanced MCM**	**Conventional Plastic**	**Advanced MCM**
Cycle of learning	9 weeks	15 weeks	12 weeks	24 weeks
Number of cycles of learning required	2	5 - 7	3 - 4	5 - 10
Sample size	50	100	100	100 plus
Residual risk	Very low	High	Medium	High

14.6.5 Temperature Hierarchy Management

The construction of any MCM is driven by the temperature hierarchy of the various processes used in sequence to produce it. Temperature hierarchy refers to the processing temperature at each step in the manufacture of an MCM. This hierarchy is managed to insure that processes down the line don't damage earlier process results. Many potentially useful processes are not feasible because they do not fit well with the other processes. We found it very useful to prepare a formal document as reflected on Table 14-6, titled *Package Temperature Hierarchy - Single and Multichip Formats.* Every process used to make and assemble the module along with the temperature and times involved is listed. It is important to list the processes where you have the potential to reduce reliability so they stand out for frequent review. As this list changes with time, as the MCM development program proceeds, it needs to be kept up to date. The importance of temperature hierarchy management cannot be overemphasized. During Unisys MCM program design reviews, the temperature hierarchy chart was a key program management tool. Final process selections and the required reliability verification tests were developed through using it.

14.6.6 Problems Encountered and Solutions Developed

No matter how well you plan your program, there are bound to be some surprises that result in additional tasks and longer development times. Some are worthy of note as they resulted in some significant achievements for us.

Table 14-6 Process Temperature Hierarchy.

PROCESS	PEAK TEMP (°C)	TIME @ PEAK TEMP (minutes)	REMELT/DAMAGE TEMPERATURE (°C)	REMARKS
Package Processing				
Fire Packages	1600		> 1500	
Braze Pins and Heatriser	830		> 780	
Plate Packages	25		> 300 (time dep)	Nickel migration
Precondition Vacuum Bake	150	1440		
Logic Die Assembly				
Die Attach (Ag/Glass)	440	2	> 450 (time dep)	Nickel oxide formation
Plasma Etch	25	5		
Wire Bond (Au Thermosonic)	150	10	> 300 (time dep)	Intermetallics at die
H_2 Reduction	335	5		
Lid Seal (Au/Sn Solder, N_2)	335	5	> 280	
Stabilization Bake	150	1440		
RAM Die Assembly				
Die Attach (Ag/Epoxy)	180	60	>180 (time dep)	Solids loss up to 200°C acceptable
Plasma Etch	25	5		
Wire Bond (Au Thermosonic)	150	10	> 300 (time dep)	Intermetallics at die
RAM Repair (if required)				
Repair	150 (substrate)	5		
	250 (die)	4		
Die Attach (Ag/Epoxy)	180	60	> 180 (time dep)	
Plasma Etch	25	5		
W/B (Au Thermosonic)	150	10	>300 (time dep)	
Lid Seal (Au/Sn, N_2)	335	5	>280	
Temperature Cycle	-40/+105	10 cycles		Package level temperature cycling
Heatsink Attach				
Attach AlN (Sn/Ag Solder)	250	2	> 221	
Burn-in	170	1×10^4		
Package-to-Board Attach	170	5	> 140	
Package Removal	185	0.5		
Stress Testing	-40/+100	20 cycles		Board level thermal cycling (5°C/min)

Conventional leak testing of hermetic packages involves subjecting the part to pressures from 15 psi to as high as 60 psi. A lid with nearly 4 square inches of surface area will sustain destructive loads under these pressures. Some companies have dealt with this problem by putting a pillar in the cavity to support the lid during pressure bombing. However, we did not have room for this so we had to seek an alternate method. A computer monitored, deflection leak detection system utilizing very low bombing pressures was successfully developed. On the front end of a new program, it is easy for a design team to assume that standard assembly and test processes will continue to work.

Another significant challenge was module sealing process development. Initially, we intended to use a conventional seam welder to hermetically seal our packages. This was consistent with temperature hierarchy management to minimize the temperature rise during sealing. However, with a cavity down design, seam welding processes existing at the time required that some pins be left off to provide room for the electrodes. System designers do not like to lose I/O pins. The solution was to introduce laser welding as a replacement for seam welding. The laser limited cavity internal temperatures to slightly over 100°C, and did not require any pins to be deleted. Implementation proved quite challenging as little laser welding experience was available from industry at large. Laser welding proved to be a very reliable, high yielding sealing method, once established.

Other development tasks which consumed more time and resources than planned were custom multichip package test methods, test hardware, burn-in sockets and accessories for internal handling and shipment of the product.

Burn-in strategy and the burn-in hardware is a significant development task for a MCM. High power dissipation and constraints on temperature hierarchy have to be carefully considered early in the program since they can influence numerous design and process decisions.

As an MCM developer, you need to be sure you understand all of the environments the MCM will see in its assembly and during its lifetime of use. One challenging environment is a board level stress test called Environmental Stress Screening or ESS testing. A typical test could consist of 20 temperature cycles between -40°C and 100°C to break marginal solder joints permitting in factory repair rather than a customer field failure. An environment like this may require additional package strength or special pin metallurgy to provide the reliability margin required to survive the test. Identify these items early in the program when they can be accommodated in the design.

14.6.7 Good News for the MCM Developer

One of the concerns frequently discussed in industry circles is final yield for an MCM. The concern stems from the belief that bare die testing and assembly

process complexity combine to impair yields. In practice, bare die testing is facilitated by chip design and by very effective bare die test methods. Also, an MCM can be designed so that with proper assembly process selections, it is possible to remove a defective die and replace it. There is considerable technical information available from published articles, consultants and suppliers to the assembly industry. Those of you about to design your first MCM should know that by paying attention to detail, it is possible to achieve final assembly yields in the high 90s. Continuous improvement programs are a very effective way to insure this.

The following items were confirmed during our programs: package suppliers can build very complex packages very competently; the industry infrastructure exists to supply many of the needed materials and piece parts; published works on others' experiences provide much of the guidance needed, and finally, helpful consultants exist if you need outside support.

14.6.8 Verify the Reliability of the Multichip Module

Upon completion of the design phase, every multichip program must go through a rigorous testing and qualification phase. During this time, the performance properties are measured to verify that system goals are met. Environmental testing confirms that the design is rugged enough to function reliably over the range of conditions anticipated during its lifetime. MCMs are more complex and more costly than an individual single chip package. Therefore, the testing phase of the program requires considerable planning so that costs for test specimens and special test methods are reasonable.

14.6.9 Summary of Development Phase Challenges

In summary, an MCM development program needs careful planning and implementation using proven development methods. Frequent design reviews are an excellent tool, along with test vehicles, to get early feedback on the design features which are new to the development team. A strong team approach with all critical skills represented is a must as well. Use the wealth of industry information available on MCM technology. Attend several of the MCM conferences held every year. Join several of the prominent technical societies and participate. These should include the IEEE, ISHM and IEPS. Follow industry standards activities like JEDEC to stay in the mainstream if you are not a large vertically integrated company.

14.7 UNISYS MULTICHIP MODULE IMPLEMENTATIONS

Historically, Unisys computer systems heavily utilize high performance memory devices for critical functions such as caches, control stores and register files. Industry SRAMs are used for these wherever possible to achieve the best cost/performance advantage. This results in a critical function that contains two die crossings/packaging delays. The first is from an application specific integrated circuit (ASIC) device that generates the address, writes data and controls signals to the static random access memory (SRAM). The second is the data from the SRAM to an ASIC which utilizes the data. It was the need to improve the performance and density of this critical area that drove the first Unisys development and application of MCMs.

14.7.1 SCAMP

The first Unisys MCM was driven by both performance and density requirements. This module, known as the single chip A-series mainframe processor (SCAMP), was developed for the MICRO A system [9]-[10]. For this module, a single ASIC was developed that implemented a mainframe processor. The module contained this ASIC and the required SRAM die for the control store. The features of this MCM, shown in Figure 14-1, are:

- 1 ASIC CMOS processor die
- 10 - 32k × 8 CMOS SRAMs
- Cavity up ceramic PGA
- 5.00 cm × 5.84 cm
- 255 pins, 2.54 mm (100 mil) grid
- 10 ceramic layers
- 4 watts total power

The ASIC die required 346 bond wires at 0.13 mm pitch and utilized aluminum wedge bonding. All the die were attached with a silver filled epoxy. The SRAMS were mounted on the same package layer as the package bond pads, requiring the wires to be down bonded with gold ball bonding. (The package wire bond pads were on the same surface as the die attach pads.) The epoxy die attach allowed the SRAM die to be replaced up to twice per site. The ASIC die was not replaceable due to the I/O pitch on the package. The package bond pads, due to the fine pitch, are not large enough to allow for new bond wires after a die replacement.

The temperature hierarchy required the use of a laser lid seal on a large lid of 4.6 cm × 5.2 cm. The epoxy die attach could not easily withstand the

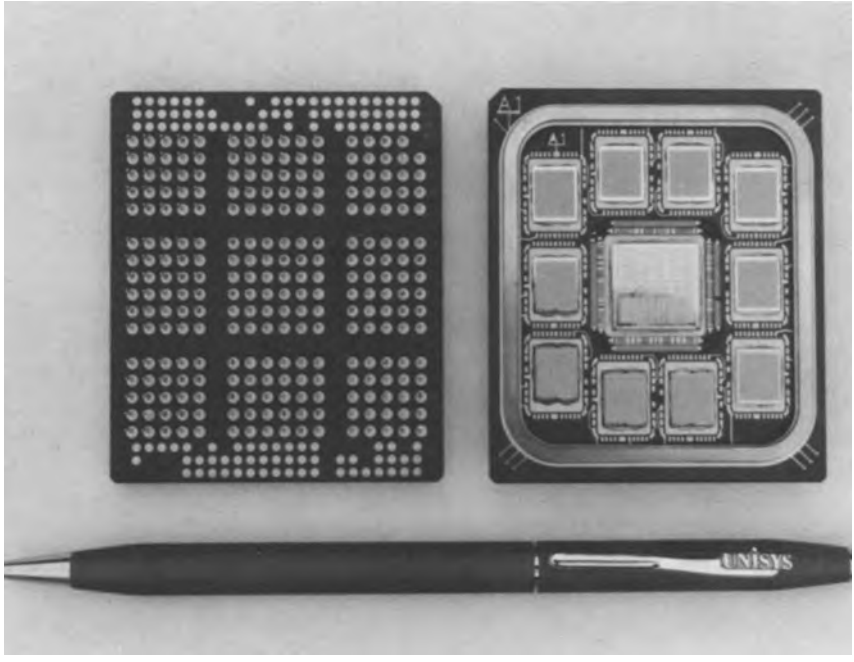

Figure 14-1 Single chip A-series mainframe (SCAMP) multichip module.

temperature of a solder seal operation. The laser allowed sealing without exceeding the temperature, which would damage the epoxy die attach.

The relatively low power of this CMOS module allowed the design to be Cavity up. Cavity up is a term which means the package cavity is opposite the pin grid array pin field. The resulting thermal performance achieved junction temperatures of 65°C at 0.27 meters per second (MPS) of air flow.

This module achieved the packaging delay reduction necessary for the system performance. The packaging delay for the two chip crossings (to and from the SRAM) would have been 10 ns in the single chip implementation. The MCM reduced this delay to 2 ns.

The MICRO A system required the placement of a mainframe as a coprocessor in a personal computer. This imposed severe space limitations since a single board must contain the processor, its control store, memory control, local main memory and an interface into the personal computer environment. The

ratio of package to silicon area for the single chip implementation would have been 12:1. The SCAMP module reduced this ratio to 4:1. This significant improvement allowed the development of the desired system.

14.7.2 A16/A19

The next Unisys module was developed originally for the A16 system which first shipped in June 1990. This large mainframe system required 51 ASICs and 216 SRAMs per processor [10]. For this system the goal was a single board processor.

The MCM developed for the caches, control stores and register files of this system was the result of a highly interactive effort between the system architects and the packaging technologists. This effort resulted in a single ceramic package design used for all the multichip functions in the system by changing only the ASIC die. The resulting processor card contained 27 MCMs and 24 single chip packages.

The basic features of this module, shown in Figure 14-2, are as follows:

- 1 ASIC 10k gate array - emitter coupled logic (ECL)
- 8 - 1k × 4 SRAMs - 5 ns ECL
- 28 watts total power
- Cavity down ceramic PGA
- 4.6 cm × 4.6 cm
- Copper/tungsten die attach slug
  - for thermal performance
- 155 pins, 2.54 mm (100 mil) grid
- 14 ceramic layers
- Copper/Invar/copper convoluted heatsink

The 28 watts of total power made the thermal performance the main challenge in this module. This dictated the cavity down design with the thermal slug that extends through to the back of the package. Cavity down means the cavity is on the same side as the pin grid array pin field. This was necessary to provide as direct a thermal path to the heatsink as possible. The ASIC die was attached directly to this slug with a silver glass die attach. The convoluted heatsink was then soldered to the top of this slug. The result was a junction to ambient for the ASIC of 40°C and 44°C for the SRAMs with 2.9 CFM of impingement air.

The A16 system required only a single cabinet. The floor space required was 80% less than its predecessor. The package to die ratio for the single chip implementation would have been 23:1. The MCM used achieved a ratio of 12:1.

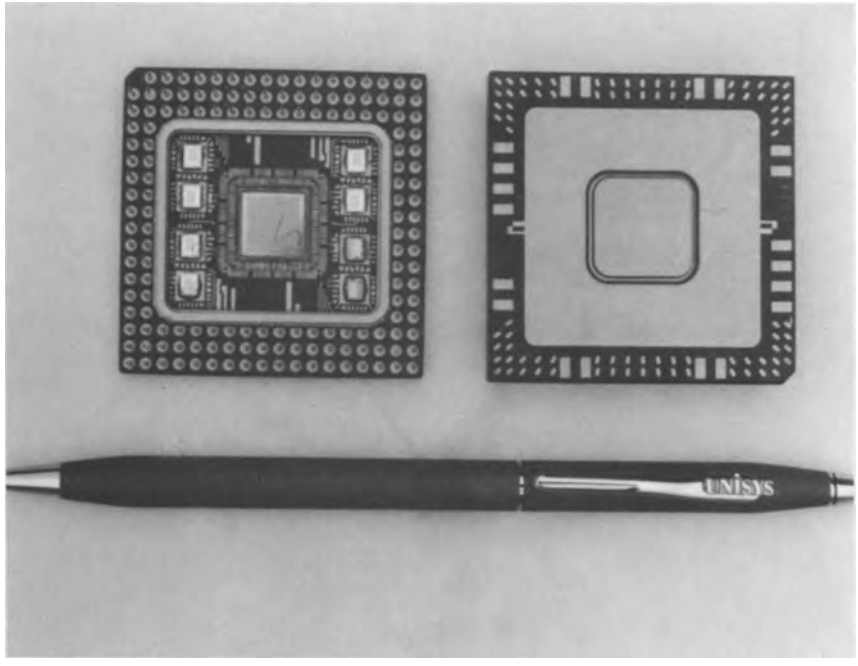

Figure 14-2 A16/A19 memory array module (MAM).

The packaging delay for the SRAM access would have been 6 ns for the single chip implementation. The MCM achieved a 1 ns packaging delay for this critical path. It is interesting to note that the packaging delay at 6 ns for the single chip implementation would have been greater than the 5 ns access time for the SRAM.

The same module is also being used in an even higher performance mainframe, the A19, which was announced in the first quarter of 1991. This system uses a four card processor with a flexible superscalar architecture. A single processor gives a 51 MIPS performance while a 6 processor system yields 240 MIPS performance. The packaging allows this system to require only 186 square feet of floor space, use only 66 KVA of power and use air cooling.

14.7.3 2200/900 Double Sided Multichip Module

The third MCM was developed for a mainframe that required 197 ASICs and

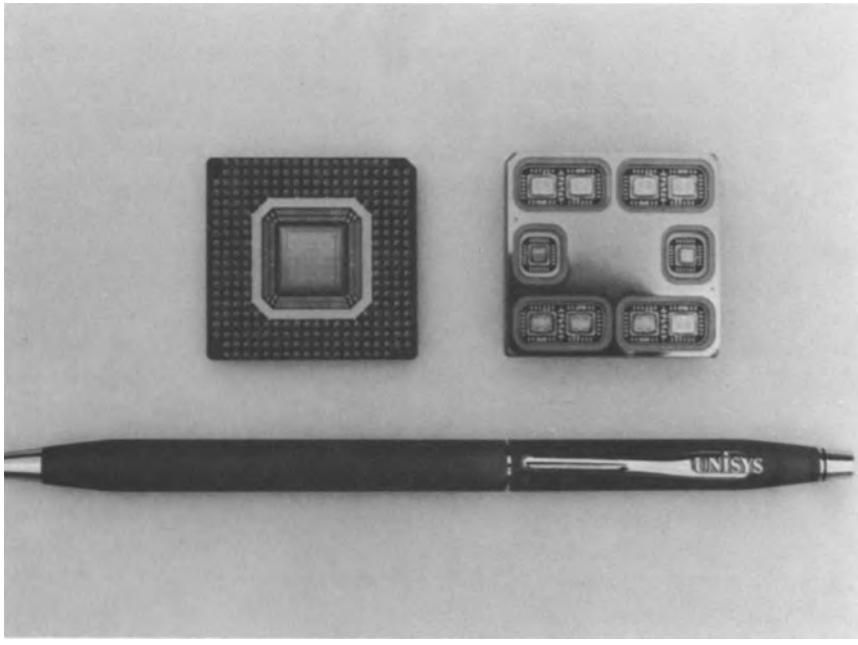

Figure 14-3 2200/900 double sided multichip module.

404 SRAMs per processor [10]. As with the A16, the goal was for a single board processor, but the initial estimates indicated that four cards would be required.

An innovative double sided package was developed to meet the requirements. This package used a cavity down approach for the ASIC and mounted the SRAMs in cavities on the top side above the pins. The basic features of this module, shown in Figure 14-3, are:

- 1 ASIC 9k gate array - ECL
- 10 - 1k by 4 SRAMS - 3 ns ECL
- 34 watts total power
- ASIC cavity down, SRAMs cavities up
- 3.0 cm × 3.2 cm
- Copper/tungsten slug
- Aluminum nitride heat spreader
- 252 pins, 1.52 mm (60 mil) grid
- 14 layers

There were challenges in every area for this package design. The thermal slug extended through the package and covered the entire top side. It contained cavities for the SRAM die that were sealed with individual lids. The aluminum nitride heat spreader was then soldered to the top. An individual liquid cooling jacket was soft attached to each package. The resulting junction to liquid temperature rise for the ASIC is only 18°C and 25°C for the SRAMs.

The single chip package to silicon density ratio for this would have been 10:1. This MCM achieved a ratio of 5:1. The performance improvement was again from a 6 ns packaging delay for the single chip implementation to 1 ns for this package. In this case, the packaging delay would have been twice the access time of the SRAM if single chip packaging were used.

14.7.4 Limits of Cofired MCMs

Each of these first three packages was based on cofired ceramic PGA technology. Each was limited to a single high pin count ASIC for several reasons. The first reason was repairability. The tight bond pad pitch required on the ceramic for the ASIC could not allow room for new wire bonds to be placed reliably if a die were replaced. The spacing on the bonds for the SRAMs allowed for large pads which made die replacement possible. It was found to be acceptable from a cost standpoint to have one die that could not be replaced, rather than multiple ones. This can be understood when one considers that even with a 99% yield on each die, a ten die module would result in 10% of the modules being bad. If the die yield goes to 95%, then over 40% of the modules would be defective without repair.

The high dielectric constant (approximately 10) of the ceramic is also a significant performance limit. For the short lines connecting from the single ASIC to the SRAMs this was not a major factor, but it would be if the general interconnect required for multiple ASICs were implemented. The ability to interconnect multiple ASICs would also be limited severely by the line density limit of approximately 20 lines per centimeter per layer of the cofired ceramic. For modules containing die with I/O counts in the 300 or more range, line densities of 200 to 400 total per centimeter are often required. This would require 10 to 20 ceramic routing layers or more.

The shrinkage factor for cofired ceramic also eliminates effective use of TAB. The final ceramic has a total dimensional uncertainty of up to 1%. For fine pitch TAB, this uncertainty would cause the leads to miss the package pads. Solder bump technology would increase the ability to use such tolerances at low pin counts, but the same problem would occur for high pin count die (along with other stress issues).

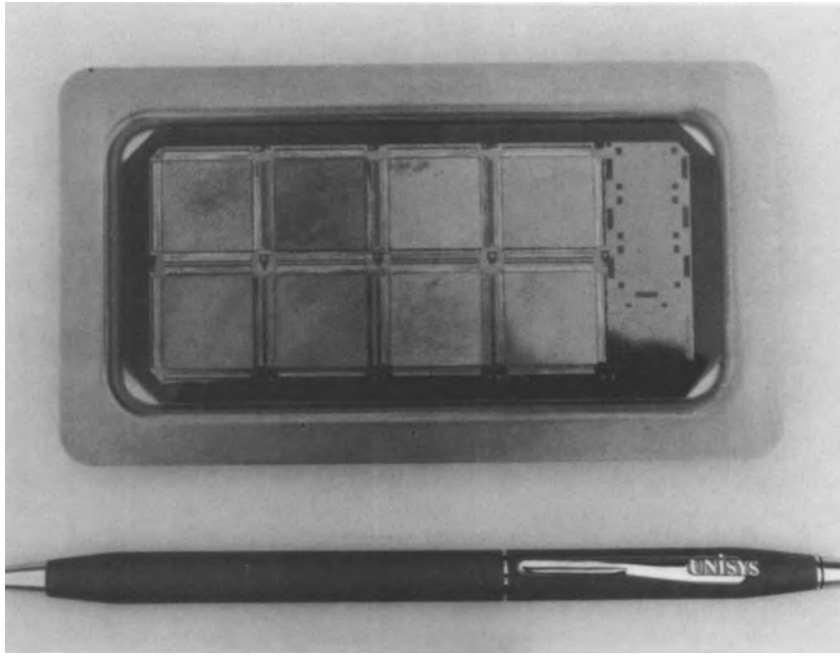

Figure 14-4 Unisys thin film multichip module.

14.7.5 Thin Film MCMs

UNISYS Thin Film

The next module developed by Unisys overcomes these limits by utilizing thin film on ceramic. This technology provides 200 lines per centimeter per layer of copper lines with a polyimide dielectric (dielectric constant of 3.5) on a ceramic base. These packages offer higher performance, higher density, higher reliability, full repairability and lower costs. One such module is shown in Figure 14-4.

Many versions of thin film MCMs are being discussed at this time. Some are done on silicon, some on metal and some on ceramic. The interconnect density and packaging features sizes vary significantly. The approach taken by Unisys offers significant advantages when compared to many of these.

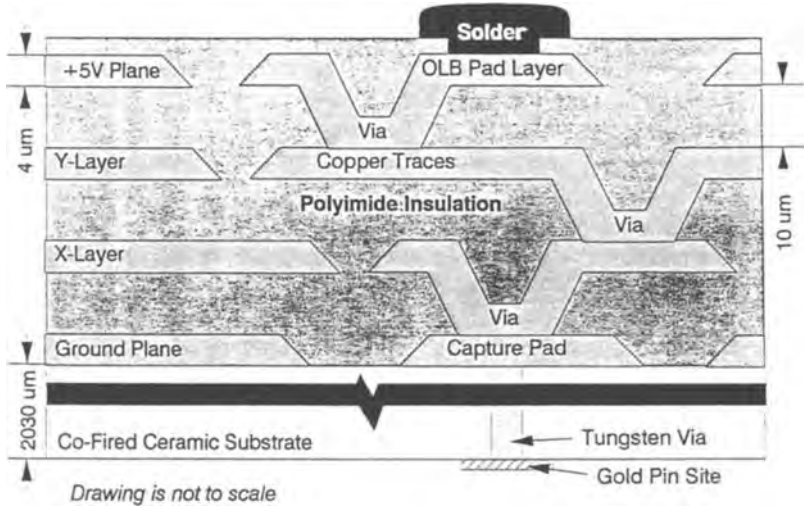

Figure 14-5 Thin film module cross section.

Interconnect Density

The first area of advantage is the interconnect density. The 200 lines per centimeter per layer allows the use of only two signal layers. A total of only four thin film layers are used. A cross section of this is shown in Figure 14-5. The first layer is a combination capture pad and ground plane. The capture pads connect to the vias in the ceramic and are sized to insure connection over the shrinkage range of the ceramic base. The two layers are the x and y signal layers. The top layer is a combination of the attach pads for the die and a five volt plane. This provides a stripline structure for the signals with excellent impedance control. The typical module is designed with the following characteristics:

- Line width = 12 μm,
- Line pitch = 50 μm
- Line impedance = 48 Ω
- Line capacitance = 1.3 pF/cm
- Line resistance = 4.0 Ω/cm

This significantly reduces cost over approaches, that require up to 10 or 12 layers to provide the same level of interconnect. The interconnect is also copper for lower line resistance. This is a major advantage in both AC performance and in the management of DC drops when compared to the aluminum used by many today.

Ceramic Base

The use of a cofired multilayer ceramic base offers several major advantages. First, the high dielectric constant of the ceramic is used to advantage by placing alternating power and ground planes in it. This provides excellent capacitive decoupling. If other substrates are used then this must be accomplished by using additional thin film layers, which is costly, or filter capacitors, which do not provide as good filtering and take up valuable area. The ceramic also provides a strong hermetic base for the module. If silicon is used, it then must be placed in some type of package at additional material and assembly cost.

The ceramic provides for optimum power, ground and signal pin location in a PGA type format. The pins are located under the die and thus provide the best signal and noise performance possible. There are several types of thin film substrates using silicon as the base. There is aluminim/polyimide on silicon, copper/polyimide on silicon and aluminum or copper using silicon oxide as the dielectric. The silicon base does not provide through vias to support an array type I/O design. This requires all power, ground and signals to exit the substrate along an edge. Thus, all interconnection paths to the next level are longer, resulting in a reduction of performance. The location of the pins under the die also results in smaller modules. The periphery connections of the other approaches require area both on the thin film substrate and the next level of packaging. This lowers both performance and increases cost when compared to the ceramic base approach.

Flip TAB

The use of flip TAB (and/or flip chip) is a significant advantage over the other alternatives. No lead forming or die attach is required. Full routing channels are also available under the die, which is not inherent in some of the other approaches. The result is the ability to place the die closer to each other improving the performance (by reducing signal delay) and reducing module size. Again, this results in lower cost and better performance than other methods.

It is recognized that thermal management is always a challenge, regardless of die orientation. With the die mounted to the substrate, the substrate has to be chosen to provide a low resistance thermal path. This usually results in a compromise in reduced routing channels or a high dielectric substrate. With the die backside up, as it is with flip tab, clever heat exchangers have to be

employed. The IBM modules previously referenced provide one example as to how this is done. Having accepted this task, the substrate material can now be chosen to optimize electrical properties.

ILB = OLB

The ability to do outer lead bond (OLB) pitch (the attach of the TAB lead to the package) equal to inner lead bond (ILB) pitch (the attach of the TAB lead to the die) is another major performance and density advantage. Improvements of four to one or more over approaches that require the TAB leads to be fanned out to a greater pitch are realized. Once fanout is required the addition of a support ring is often required, which dictates the leads to be even longer. This is illustrated by Figure 14-6.

For example, for an ILB of 0.2 mm, fanning out to 0.5 mm requires leads of 12 mm. With an OLB pitch equal to the ILB pitch the length is less than 1.0 mm. This results in switching noise improvement by a factor of 20 over the 12 mm long leads.

The fine pitch outer lead bonding processes to support this method are challenging. Fortunately, the industry is working in this area. Laser bonding, single point bonding and various combinations of these are emerging as processes to address the challenge.

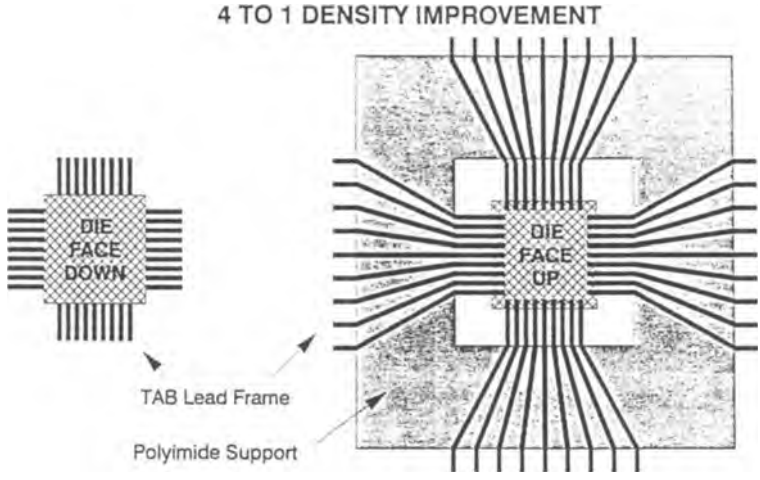

Figure 14-6 TAB lead fanout.

Reliability

All these modules contributed to improved system reliability in their respective applications. The thin film module will offer the greatest gain since it offers the greatest density (a ratio of two for the package to silicon) and the greatest reduction in connections. All examples offered reductions in:

- Interconnects
- PWBs
- Connectors
- Cables
- Backplanes
- Cabinets

These reductions all contribute to improved reliability.

Cost

The thin film MCMs also offer significant cost improvements. First, the packaging material itself is cheaper than the equivalent single chip ceramic PGA if approximately four or more die are placed in a module. The scrap cost also is reduced. Since the thin film module is completely repairable, a defective die is simply replaced. There is no package scrap due to a defective die as is true with ceramic PGAs. The cost of assembly is reduced. Once the die are mounted in the module only a single physical part is handled for test, burn-in, and insertion into the next level of packaging.

While all the above listed cost savings are significant, the greatest savings is at the system level. The savings are the result of:

- Lower PWB cost
  - Most of the interconnect is in the MCM
  - Smaller boards required
- Fewer PWBs
- Fewer connectors
- Fewer and smaller cabinets

14.8 MULTICHIP MODULE INDUSTRY ISSUES

The development and application of multichip packaging has been slower on an industry wide basis than many have predicted [11]. There are many challenges which have faced and continue to face multichip packaging.

Die availability is one of the major issues. Many suppliers of semiconductors are reluctant to sell bare die and some totally refuse to do so. This reluctance exists for several reasons. The bare die does require more careful treatment in handling and storage than a packaged part. Many potential customers are not equipped to meet the requirements of such handling and storage. The supplier is concerned that blame will be misplaced if parts fail in the field due to inadequate handling by the multichip assembler. Ability to burn-in bare die is a complex issue. While TAB claims to facilitate bare die burn-in, it can be a complicated and expensive process for high I/O die. Some suppliers lack the TAB and/or solder bump capability that is desirable for MCMs. While this does not make it impossible for them to supply die, it can be a significant handicap.

The management of die quality is another concern for both supplier and customer. Most suppliers do not have the capability to fully test a die prior to packaging [12]. (TAB, if available, does help this problem in that better testing can be done by more suppliers.) Both the supplier and the customer know how to do business when they both can fully test the component being exchanged (with no potentially damaging processing required by the customer). The issues of how to price, determine initial yields, and determine parts for return are some of the issues.

These are valid concerns that must be addressed. The customers need to develop the expertise to properly handle and store die. Even packaged parts can be damaged by mishandling, but customers long ago developed the expertise to handle them. It will take time for the same to happen for die. Additionally, suppliers will develop better die test methods and also more robust die. Time and the requirements will drive this development.

Semiconductor suppliers eventually will become MCM developers. This will eliminate some of the issues since both the die supplier and customer would then be the same company. However, that same fact will be a limitation. Most modules will require die from multiple suppliers. Semiconductor suppliers will not, in most cases, be able to negotiate for other suppliers' die.

A profitable price structure for the die business must also be developed. Semiconductor suppliers are structured such that packaging is a significant portion of their value added income. Their business is based on the profits generated not only by manufacturing the die, but also from packaging it. Their initial reaction is to view selling die as limiting their profits. This is another serious issue but one which time and volumes should solve.

Another factor which has slowed development is testing. In order to generate a good module test, the simulation models for the die in the module are required. Often suppliers of standard VLSI devices are reluctant to make these simulation models available. The gate level model is basically the design and

is considered proprietary by these suppliers. This problem exists even with single chip devices since they also must be tested after board assembly. Application of board level methods and the use of higher level models will solve this problem.

Whereas bare chips for MCM use have not been readily available for the reasons stated above, an interesting phenomenon is occurring in the multichip package/substrate area. Figure 14-7 is a trend plot of the parts count and mainframe growth history discussed in the introduction to this chapter. Vertically integrated mainframe developers have seen parts counts drop, growth almost level and development complexity climb rapidly for each new machine. This trend is predicted to result in excess factory capacities and excess inventory.

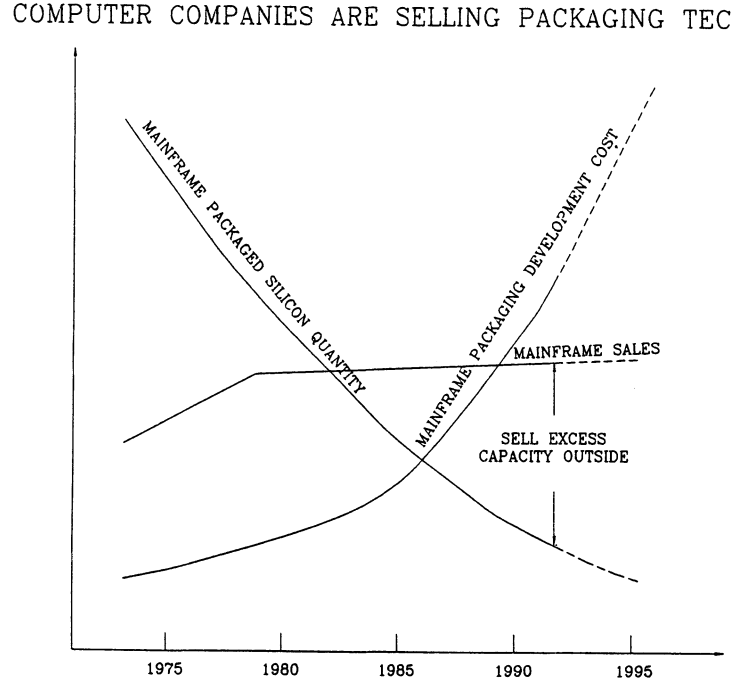

Figure 14-7 Mainframe component trends.

Furthermore, these factories are operated with expensive and complex equipment mandated by the complexity of the products required by mainframes. Although companies will attempt to fill these voids with products from faster growing lines such as workstations, differences in technology will not always make this practical.

One way to deal with this excess capacity and/or inventory is to sell packaging technology/plant capacity to others, including competitors. A number of the Fortune 500 companies publicly announced during 1991 that their packaging technology could be purchased [13], including thin film substrates, advanced ceramic modules, assembly services and so on. This trend should have a positive effect on developing the infrastructure required for smaller companies to develop MCMs. It should also lower the cost of developing modules for the entrants into the marketplace.

MCMs will make it! The development is following a chicken and egg scenario. Today, everyone would use them if the infrastructure were there, and the infrastructure would be there if the volumes warranted it [14]. This is the normal situation for a new technology. Surface mount technology followed a similar, though less complex, development.

Initial uses will be in areas where performance and/or density dictate that no other reasonable alternative exists. These applications will drive the development of the infrastructure required. They will also slowly increase the volume and drive the cost down such that more applications will be addressed.

Acknowledgment

The authors wish to acknowledge the contributions of a large number of people in Mission Viejo, Rancho Bernardo, Roseville and Tredyffrin, who contributed to the developments covered in this chapter.

References

1 A. S. Rappaport and S. Halevi, "The Computerless Computer Company," *Harvard Business Review*, vol. 52, pp. 69-80, July/Aug. 1991.

2 "Should the U.S. Abandon Computer Manufacturing," *Harvard Business Review*, vol. 52, pp. 140-160, Sept./Oct. 1991.

3 E. Davidson, "The Coming of Age For MCM Packaging Technology," *IEEE Computer Packaging Workshop*, (Lake Tahoe, CA), May 1991. No published proceedings.

4 H. W. Markstein, ed., "Multichip Modules Pursue Wafer Scale Performance," *Electronic Packaging & Production*, vol. 31, pp. 40-44, Oct. 1991.

5 A. J. Blodgett and D. R. Barbour, "Thermal Conduction Module: A High Performance Multilayer Ceramic Package," *IBM J. Res. and Devel.*, vol. 26, pp. 30-36, 1982.

6 J. U. Knickerbocker, *et al.*, "IBM System/390 Air-Cooled Alumina Thermal Conduction Module," *IBM J. Res. and Devel.*, vol. 35, pp. 330-348, May 1991.

7 V. L. Gani, *et al.*, "IBM Enterprise System/9000 Type 9121 Model 320 Air Cooled Processor Technology," *IBM J. Res. and Devel.*, vol. 35, pp. 342-351, May 1991.

8 B. Blood, "ASIC Design Methodology For Multichip Modules," *Hybrid Circuit Technology*, vol. 8, pp. 21-27, Dec. 1991.

9 J. Nelson and C. Cheves, "Multichip Module For A Cost Driven Mainframe," *Proc. Internat. Electronic Packaging Soc. Conf.*, IEPS, (San Diego CA), pp. 230-240, Sept. 1989.

10 R. D. Rhodes, "Multichip Packages at Unisys - Cofired Ceramic to Thin Film," *NEPCON Proceedings*, (Anaheim CA), vol. 1, pp. 477-485, Feb. 1991.

11 "Emerging Technologies, What The Experts Say," *Inside ISHM*, vol. 18, pp. 5-12, Nov./Dec. 1991.

12 J. D. Mosley, ed.,"Multichip Modules, Lack of Standards Impedes Design Issues," *EDN*, vol. 21, pp. 35-42, Jan. 1992.

13 K. O'Brien and T. Williams, eds., "A Changing Hybrid Marketplace," *Hybrid Circuit Technology*, vol. 8, pp. 14-18, Dec. 1991.

14 "The MCM Market: It's Anybody's Guess," *Hybrid Circuit Technology*, vol. 9, pp. 6-7, Jan. 1992.

15

HIGH PERFORMANCE AEROSPACE MULTICHIP MODULE TECHNOLOGY DEVELOPMENT AT HUGHES

Robert L. Bone, Dennis F. Elwell, James J. Licari,
Robert S. Miles, Kevin P. Shambrook, Stanley M. Stuhlbarg,
Philip A. Trask, David F. Zarnow

Technology development occurs only when needs are recognized and financial support is applied to focus the necessary personnel and materials. Much recent progress in the development of multichip modules (MCMs) can be attributed to efforts supported by aerospace industries. These industries were among the first to recognize that semiconductor device complexity was beginning to outstrip an engineer's ability to ensure effective intercommunication between chips in a system and among chip clusters.

This chapter discusses the MCM evolution as it parallels the development of military and commercial systems over the past ten years. It stresses the changing role of packaging as a driving force in system definition, describes several key technological developments that are spearheading this packaging drive and outlines the need for a comprehensive computer aided design, manufacturing and testing operation to bring about successful systems implementation.

15.1 MCMs MEET THE NEEDS OF SYSTEMS EVOLUTION

An increase in sophistication and performance levels of military electronic systems, made possible by the rapid development of VHSIC and VLSI chips,

Weapon Comparison

	ADCAP	RDT
Length	20.0 Ft.	9.16 Ft.
Diameter	21.0 In.	6.5 In.
Weight	4000 Lbs.	173 Lbs.

Electronic Unit Comparison

	ADCAP	RDT
Weight	460 Lbs.	22.4 Lbs.
Volume	5.8 Cu. Ft.	0.23 Cu. Ft.
Complexity	1.0	1.0

Figure 15-1 The trend for torpedoes is toward higher performance and smaller electronics. Although the new reduced diameter torpedo has approximately 4% of the volume of the ADCAP torpedo, its electronic unit is of similar complexity.

must be achieved simultaneously with requirements for lighter, more compact microelectronic packaging. This twin challenge - higher performance, but smaller packages - has driven the military electronics industry to seek radical new approaches to the fabrication of MCMs.

The increased complexity of military electronics throughout the aerospace industry parallels a rapid rate of change in the commercial market. In particular, computer manufacturers have developed novel packaging techniques that provide speed and density performance to match the rapid evolution in the speed of memory and logic chips. In military packaging, there are additional requirements of ruggedness and reliability that must be met by manufacturers. Commercial products also are becoming more reliable.

15.1.1 Examples of Trends in Weapons

Examples of military systems that require order-of-magnitude packaging improvements over previous generations can be found in the seas, in the air, and in space. Increasingly sophisticated weapon systems are required to counter improvements in a potential enemy's weapons. Even if the apparent end of the Cold War reduces the need for tactical weapons, a national need is likely to

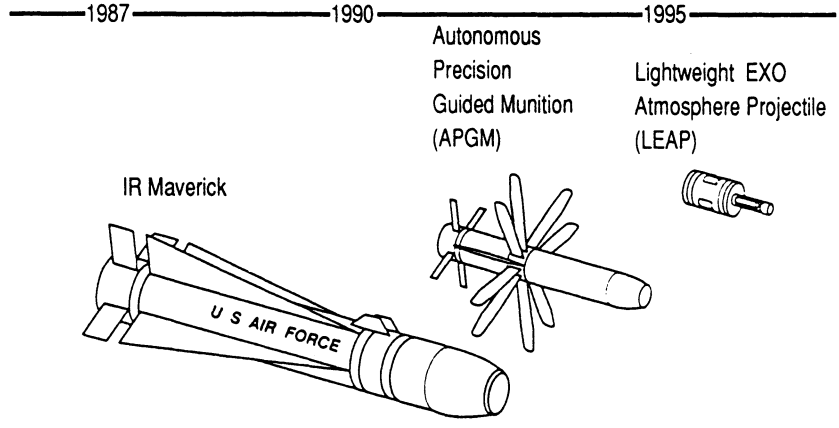

Weapon Comparison

	Maverick	APGM	Leap
Volume	6.4 Cu. Ft.	0.34 Cu. Ft.	0.14 Cu. Ft.
Weight	675 Lbs.	125 Lbs.	5.5 Lbs.

Electronic Unit Comparison

	Maverick	APGM	Leap
Weight	100 Lbs.	1.5 Lbs.	0.52 Lbs.
Volume	1.64 Cu. Ft.	0.02 Cu. Ft.	0.025 Cu. Ft.
Thru Point	4.0 MIPs	20 MIPs	4.2 MIPs

Figure 15-2 Higher performance and smaller electronics is also the trend for missiles. The developmental APGM and LEAP have similar or faster data processing capability than the MAVERICK, but are tiny in comparison.

remain for the most modern surveillance systems and electronic aids to counter terrorism and drug operations. Such equipment generally processes data from various sensors and, in many cases, evaluates the input and provides an automatic response.

A "smarter system" is one with more powerful computers and intelligent sensors. Smarter and smaller weapons lead to an increase in performance and a decrease in cost.

In the Seas

Figure 15-1 shows a size comparison between the Department of Defense (DoD)/Hughes' advanced capability torpedo (ADCAP), developed in the mid to late 1980s, and the projected reduced diameter torpedo, which represents the next generation. The weight of the new torpedo will be approximately 5% of the ADCAP weight. At the same time, the electronic sophistication has increased.

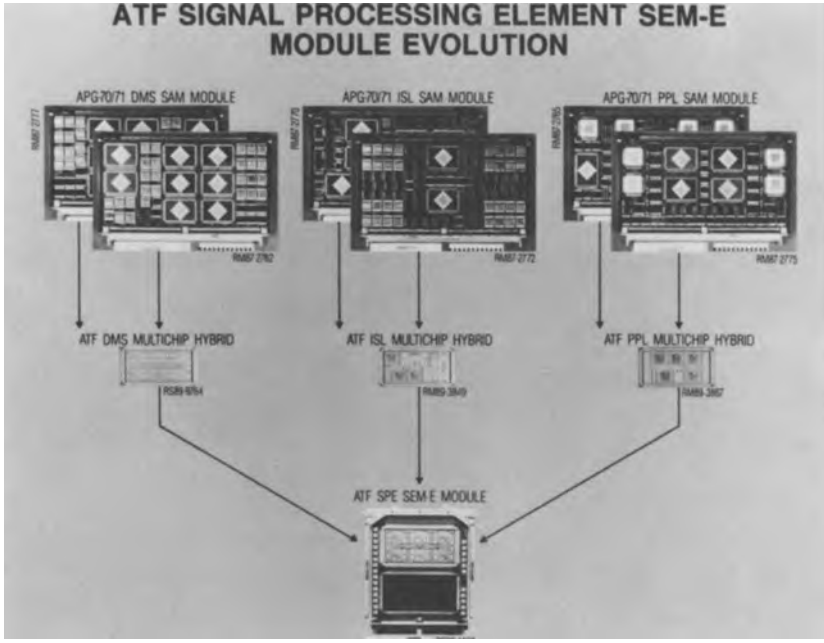

Figure 15-3 Evolution of the ATF signal processing element, SEM-E, module.

In the Air

Successive generations of missiles and projectiles are compared in Figure 15-2. The MAVERICK missile, currently in production, is shown on the same scale as the Autonomous Precision Guided Munition (APGM) and the Lightweight Exo-Atmospheric Projectile (LEAP) missiles. Even though modern, highly integrated semiconductor devices and advanced hybrid technologies are used to control the MAVERICK missile, the electronic units of the APGM and LEAP must be less than 2% and 1%, respectively, of the weight and volume of the MAVERICK electronics unit.

The Advanced Tactical Fighter (ATF) Signal Processing Element (SPE) is undergoing significant packaging evolution. As shown in Figure 15-3, 2" × 4" MCMs have become key elements in this evolutionary change. The modules can be mounted, two per side, on Standard Electronic Modules (SEMs), or three per side on Standard Avionic Modules (SAMs), shown in Figure 15-4.

Figure 15-5 shows the 6" diameter LEAP Guidance Unit electronics implemented in the High Density Multichip Interconnect (HDMI) packaging technology, discussed later in this chapter.

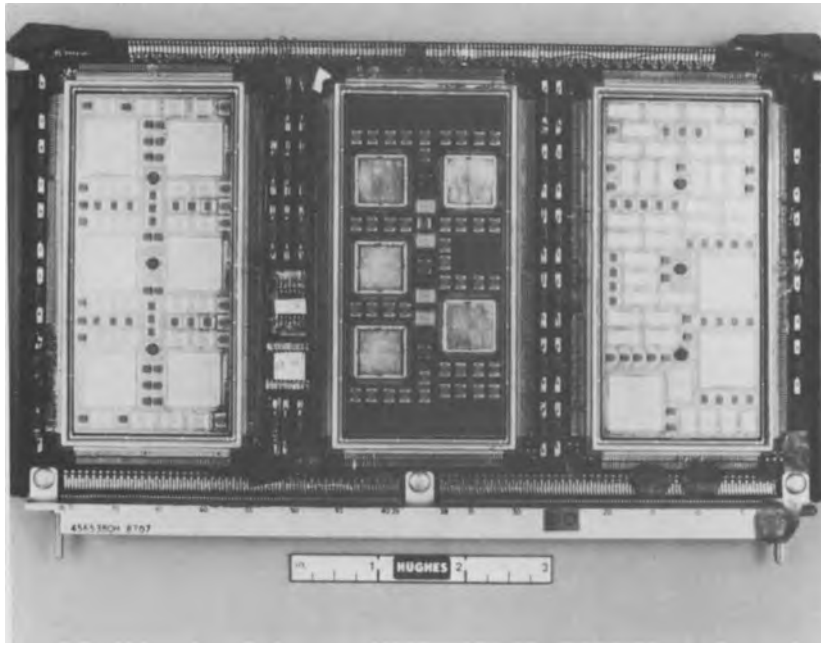

Figure 15-4 MCMs mounted on SAM board.

In Space

The system requirements of the Strategic Defense Initiative (SDI) are making radically new demands on packaging technology. Figure 15-6 is an artist's impression of the Hughes' Boost Surveillance and Tracking System (BSTS), one of the surveillance units to be deployed as part of SDI. According to an article in *Defense Electronics* [1], the BSTS will require analog to digital (A/D) converters with a greater dynamic range than required in any previous system. In this case, military packaging technology is likely to be ahead of that needed in other leading commercial technologies.

The DoD designated microelectronics and packaging as "one of the eight most critical technologies... required to ensure long term qualitative superiority of United States weapons systems" in its 1991 "Critical Technologies Plan" [2]. While there is no ranking order within those eight technologies, improved microelectronics and packaging are the only hardware elements designated as essential to all of the DoD top 12 mission goals. DoD succinctly and

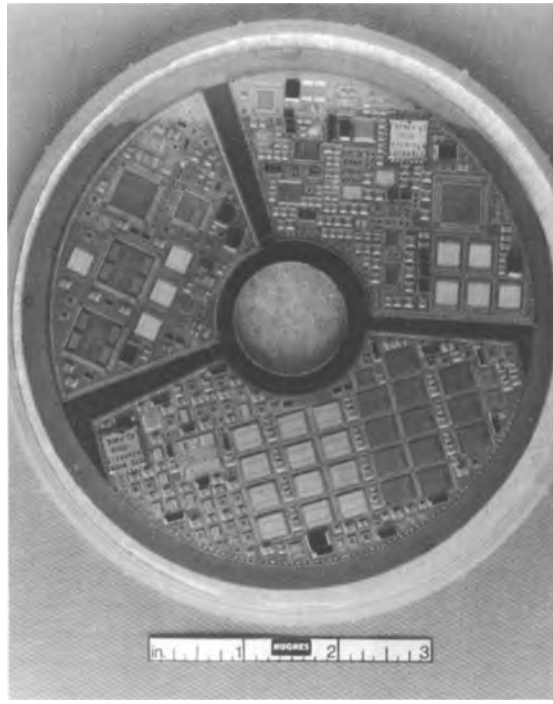

Figure 15-5 Use of HDMI MCM technology in the SDI LEAP program results in a single card electronic guidance unit which weighs only 140 grams.

emphatically states, "... microelectronics technology has a pervasive effect on virtually every U.S. weapon system, current or future" [2]. Order-of-magnitude advancements in microelectronic packaging provides unexcelled leverage in weapon system improvements and are, consequently, essential.

Microelectronic packaging advances are crucial for continued U.S. weapons superiority and essential to the success of the U.S. electronics industry in international commercial competition. Great cost savings can be made by using new miniature, high speed microelectronics to upgrade performance and reliability of existing weapons systems. Order-of-magnitude miniaturization levels and accompanying performance improvements now are needed to create radically new weapons that will fulfill the long term critical mission requirements of the DoD.

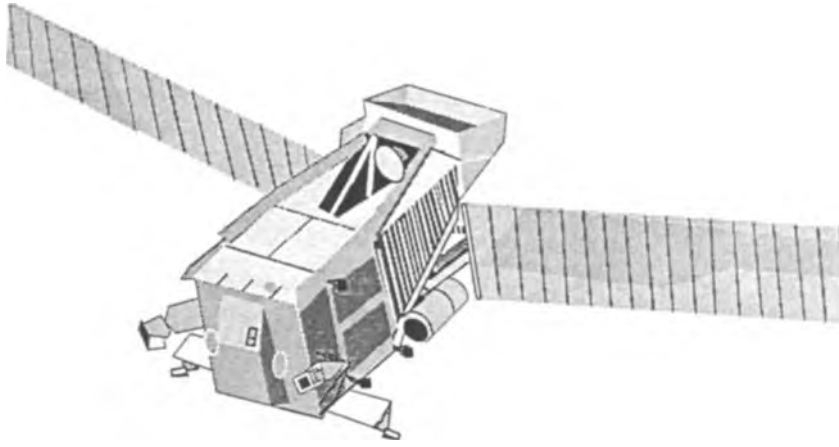

Figure 15-6 An artist's impression of the Boost Surveillance and Tracking System (BSTS), which will be deployed as part of the Strategic Defense Initiative. The BSTS will advance the state of the art for A/D converters.

15.1.2 Measuring the Trends

Recent years have seen progressive increases in the number of input/output (I/O) leads, chip area, power dissipation and speed of semiconductor chips. Future trends in these chip parameters are shown in Table 15-1.

Input/Output Leads

Increase in the complexity of individual semiconductor devices generally requires more I/O pads. The difficulty of packaging chips with large numbers of I/Os is one of the factors that led to MCM technology. One approach toward minimizing I/O connections is wafer scale integration (WSI), in which a 6", or even 8", silicon wafer is filled with functioning, interconnected ICs, including integrated component parts such as resistors and capacitors. Successful WSI usage remains elusive, at least for the present. Multichip microelectronic packaging as a field has maintained a healthy growth rate, because, however complex individual chips have become, circuit designers usually need to mount two or more onto the same substrate. A large number of I/Os on the chips requires a high line density.

Table 15-1 Semiconductor Research Corporation (SRC) Packaging Roadmap - IC Chip Complexity Trends.

	1994	2001
CMOS		
(Gate Array and MPU)		
• Size	1.7 (2.0)	2.5
• I/O	600 (1000)	2000
• Power (watts)	15 (50)	60
• Rise Time (ps)	700 (200)	150
Switching Lines		
• Voltage	3.3 (2.1)	1.6
• Number Lines (MPU)	128 (256)	400
Bipolar		
(Gate Array)		
• Size (cm)	2.2	2.5
• I/O	600 (1000)	2000
• Power (watts)	60 (100)	200
• Rise Time (ps)	100 (75)	40

Note:
Average anticipated values are shown.
Maximum value for small number of devices are shown in parentheses.

Line Density

The width and spacing of conductors used to interconnect circuit devices have been used extensively as a measure of packaging efficiency. Figure 15-7 illustrates the relative rapidity of transition from 0.025" wide conductor lines and 0.025" spacing in 1980 to the present 25 μm line widths with 75 μm spacing. In special cases, these figures can be reduced to 15 μm and 35 μm, respectively. Within the past decade (the 1980s), a notable measure of packaging trends has been the transition from mils to microns as the unit used to describe conductor line widths and spacings. As lines get narrower, they also get shorter due to closer spacing of die, reducing conductor area and capacitance, leading to greater signal throughput.

Signal Throughput

One overall measure of packaging efficiency for military guidance unit processors is signal throughput per unit volume. In Figure 15-8, increased throughput of military signal processors is depicted, using as units millions of

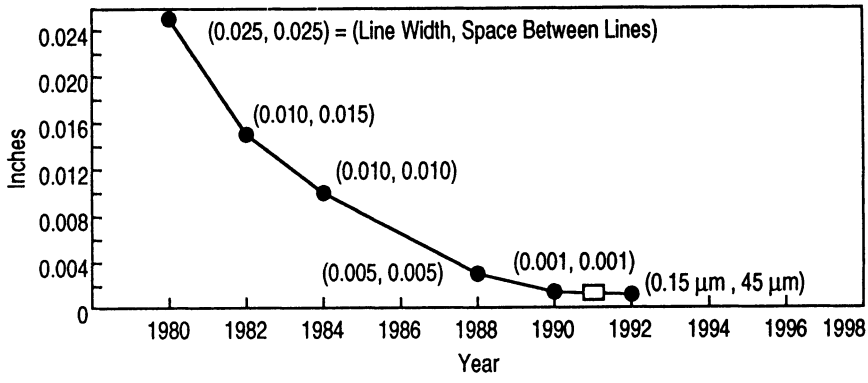

Figure 15-7 Conductor line widths and spaces for typical Hughes interconnect structures, 1980-1992.

instructions per second per cubic inch. From a level of 100 in 1990, this number is expected to increase sharply over the next few years. Higher signal throughput and higher packaging densities lead to more heat dissipation.

Power Density and Power Dissipation

An additional measure of packaging efficiency is power density (watts dissipated per unit volume of package), as shown in Figure 15-9. This value increased from 400 mW per cubic inch for the Hughes' PHOENIX missile in 1980 to 1.2 watts per cubic inch in the currently produced AMRAAM missile. In the LEAP program, this figure will increase to 7.3 watts per cubic inch. As this trend continues, power dissipation will continue to be a challenging factor. A similar problem has been experienced in mainframe computers, where cooling air or water is once again a requirement as it was in the 1950s when computers used vacuum tubes.

Volume Fraction

A common parameter used as a measure of packaging efficiency is the fraction of package surface area occupied by active semiconductor devices. Extending this to the *volume fraction* occupied by active devices is a particularly sensitive parameter and a more useful measure. In Figure 15-10, values of this parameter are plotted for well established technologies and for some developmental and projected approaches. Traditional hybrid microcircuits mounted onto epoxy based printed wiring boards (PWBs), with semiconductor devices mounted on one side

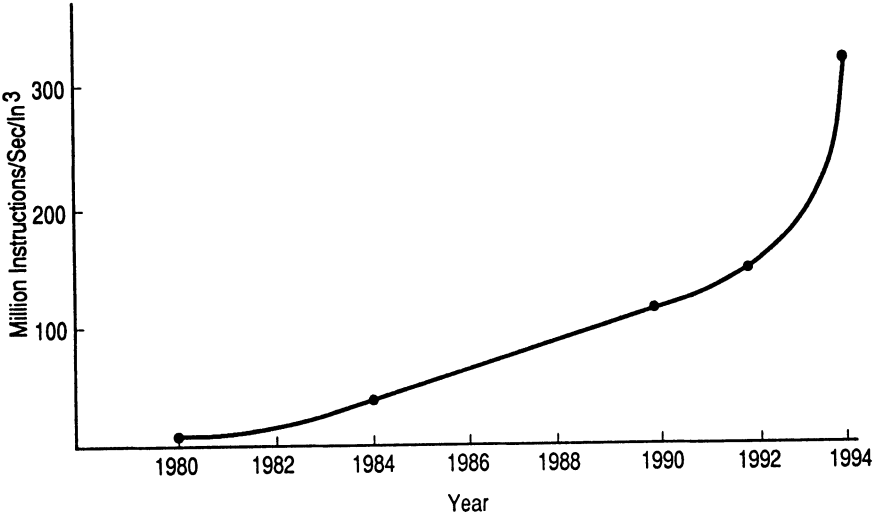

Figure 15-8 Processor throughput per unit volume.

along with typical surface mount passive component parts, have volume packaging efficiencies less than 0.5%. With thick film multilayer ceramic wiring boards (CWBs), developed in the early 1980s, this figure has increased to around 1%.

The earliest attempts at high density MCMs have led to an improvement in volume fraction by a factor of two, while the developmental Hughes' LEAP packaging technology has a volume efficiency around 3.5%. Volumetric efficiency advancements in three dimensional MCM packaging with integrated passive devices is projected to reach efficiencies up to 25%. For monolithic WSI, as yet unrealized on large scale systems, this efficiency figure would become about 40%. Projections for other near term developments are also shown in Figure 15-10. One of the key evolving steps in military packaging is the use of flip chip attachment of the semiconductor die, avoiding leads which extend laterally and occupy considerable surface area. The use of a hermetic sealing coating, eliminating the necessity of a lidded package, would also have a great impact on packaging efficiency.

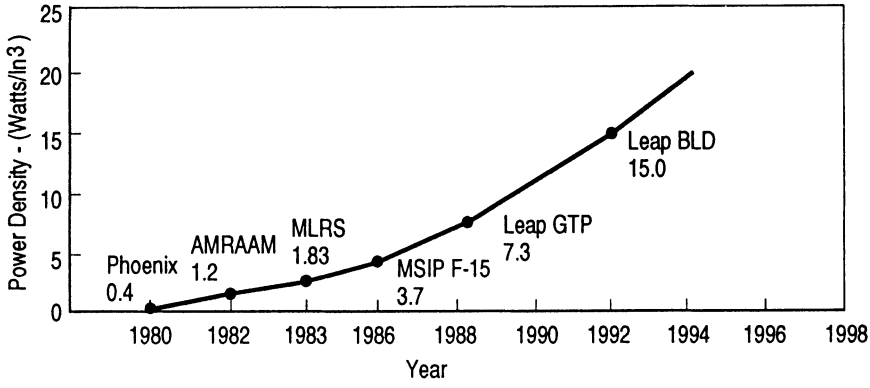

Figure 15-9 Military electronic power density trends, including projections through 1994.

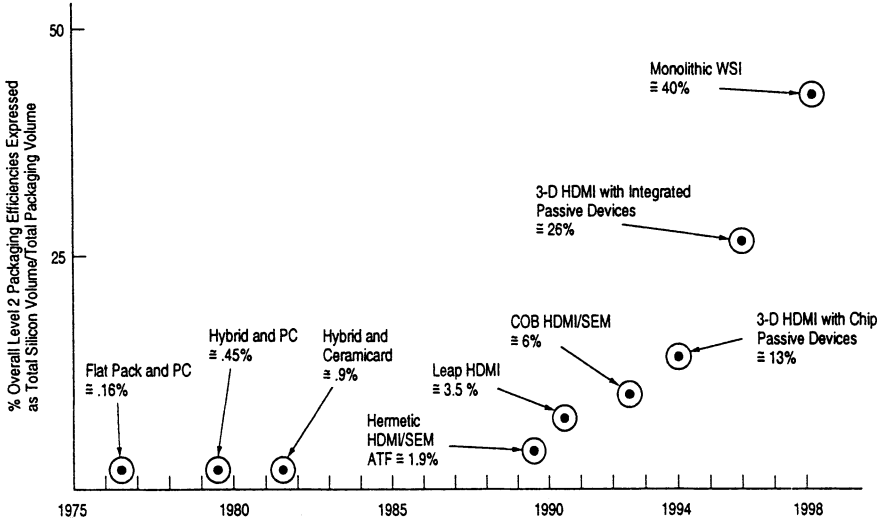

Figure 15-10 While major increases in volumetric packaging efficiency have been made with high . density planar MCMs, significant future improvements require three dimensional packaging and chip-on-board approaches.

15.2 NEW PACKAGING ROLES

Within the aerospace industry, progress in the past has been determined mainly by advances in system technology. During the 1970s and 1980s, advances in packaging technology began to accelerate as semiconductors incorporated increasing numbers of electronic functions per unit area of silicon. Various active and passive component parts were assembled within smaller volumes, as building blocks progressed from PWBs with parts assembled through holes, to hybrid microcircuits on CWBs. This acceleration has increased continually to the point where it has impacted the classical systems supported by the packaging scheme described above. Now and in the future, the progress of electronic systems will be determined more and more by advances in semiconductor and packaging technology. Packaging is becoming an increasingly dominant technology.

Two major packaging technologies, described in Section 15.3, are multilayer low temperature cofired ceramic (LTCC) circuits and multilayer thin film substrates. Hughes developed LTCC in the early 1980s for use in microwave and high frequency (over 300 MHz) digital packaging. By the late 1980s, it became clear that multilayer thin film substrates also were required. A new in-house packaging program was set up to address the commercial and aerospace markets.

15.2.1 A High Density Multichip Interconnection Program

In 1987, the High Density Multichip Interconnect (HDMI) program was created to develop advanced MCM packaging required for the DoD/Hughes Advanced Tactical Fighter (ATF), BSTS, LHX, LEAP and ADCAP weapon programs. Progress in the HDMI program led to the creation of an aluminum/polyimide HDMI substrate fabrication facility in Newport Beach CA, with comprehensive high density MCM design, manufacturing and test capabilities. This rapid progress was instrumental in obtaining several key DoD electronic packaging technology contracts, including two Air Force Smart Skins contracts and the Naval Ocean Systems Center (NOSC) VLSIC Packaging Technology (VPT) program [3].

15.2.2 A Company Wide Electronics Packaging Program

Based upon success of the HDMI program and the need to provide timely and unified responses to DoD customers, the Hughes Electronics Packaging Program (HEPP) was established in late 1988. The scope of HEPP was increased to include all high density electronic packaging technologies. Additionally,

managers were assigned to major customer electronic packaging categories and to interface with the Microelectronic and Computer Technology Corporation (MCC). To date, the most valuable work of HEPP has been the timely development of high density multichip packaging technologies used to demonstrate packaging viability for programs such as the DoD/Hughes ATF, BSTS, LEAP, LHX, MMIC, and ADCAP and for commercial products.

As known, the increased dominance of packaging technology in overall systems design has modified the classical relationship of the packaging function supporting systems engineering in achieving developmental and production goals. The packaging function now plays a partnership role with systems engineering - an example of concurrent engineering in action.

15.3 DEVELOPING TWO KEY MCM TECHNOLOGIES

MCM development has been driven by system needs, packaging improvements and the increasing complexity and speed of semiconductors, particularly CMOS. Semiconductor interconnection density has been increasing more rapidly than that of PWBs (Figure 15-11), leading to a widening gap in interconnect capability. The first approach toward increasing connectivity was to increase the line density and number of layers in PWBs, leading to development of the MCM-L

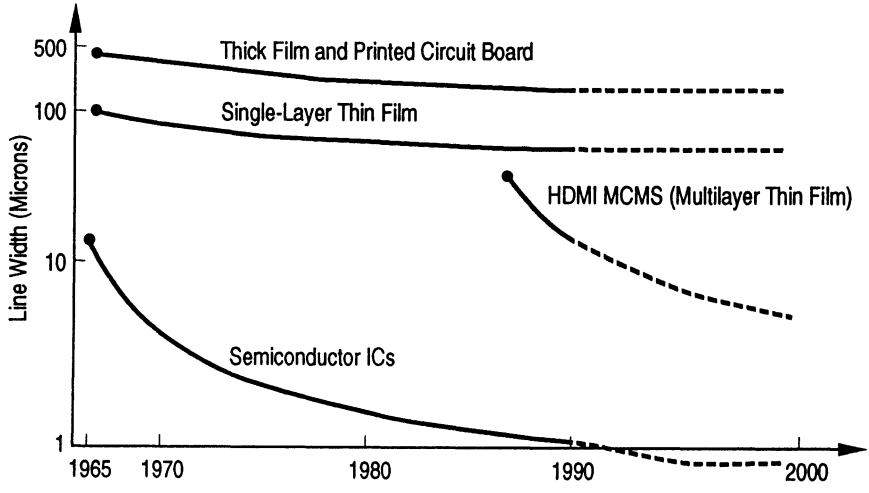

Figure 15-11 HDMI technology reduces this IC versus packaging gap to improve packaging density.

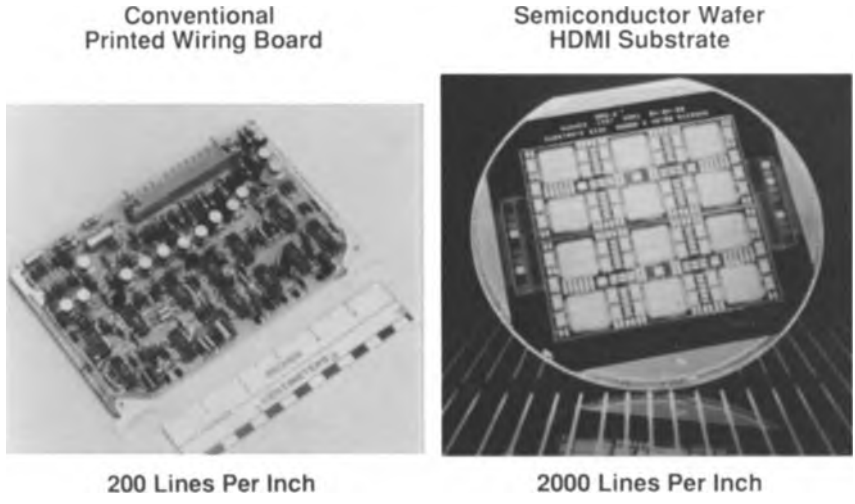

Figure 15-12 A 10:1 improvement in wiring density closes the interconnectivity gap.

technology (see Chapter 5). Other principal approaches have been developed to close this gap. One is a multilayer ceramic MCM (MCM-C), which uses LTCC technology, described in Section 15.3.1 and further described in Chapter 6. The other approach is a multilayer thin film (MCM-D), described in 5.3.2 and in Chapter 7. (See Figure 15-12.)

15.3.1 An Evolutionary Approach

An early and important example of MCM-C was the IBM Thermal Conduction Module (TCM), which has die flip chip bonded on a high temperature cofired ceramic (HTCC) interconnect (more than 30 layers). The TCM uses an area array die I/O pattern for two reasons. The desired I/O density would require an extremely fine pitch if it were peripheral I/O. This I/O density cannot be matched by the pad pitch in the ceramic. Making the die larger to achieve a larger I/O pitch creates unacceptable stress problems when using flip chip bonding, as it also reduces die yield. Multilayer ceramic interconnect matches semiconductor area arrays and flip chip bonding very well, but standard die do not have area array interconnects. Aerospace companies generally use relatively small quantities of chips and, therefore, cannot economically justify special area array die configurations. Consequently, their usage has not become a generic aerospace packaging approach. Now, because of the increasing number of I/O

contacts required and the development of MCMs, area array patterns are becoming more common.

The traditional MCM was HTCC, which entailed high set up and change costs and excluded the use of gold, silver or copper metals or buried resistors/capacitors. Recent developments have introduced LTCC based on cordierite or glass filled thick film-type materials, which fire around 900°C. IBM is using this LTCC technology as a successor to its TCM. It has been developed at Hughes for both aerospace and commercial applications.

The LTCC approach to MCMs does not address directly the problem of die that achieve high I/O by a fine pitch, about 6 mils in 1990 and reduced to 4 mils in production by 1992. In the early 1980s, Hughes decided it needed a fine line interconnection technology to use most advantageously the fine pitch microprocessor and configurable gate array die projected for the 1990s. Research was done on how best to adapt the technology developed for semiconductors to the die interconnection problem. Traditional PWBs and thick film networks use conductor patterns which are about two orders-of-magnitude greater than the line width of 1 μm being achieved by semiconductors and are limited by their manufacturing process (Figure 15-11). By adopting a photolithographic approach, the line width of the base network is no longer limited by the process but may be optimized for resistance and crosstalk, two parameters which also might limit width and spacing. To match the fine pitch on the die, multilayer fine pitch MCM-D technology was developed. The approach developed is described in detail in Sections 15.2.1 and 15.3.3. The various technologies are compared in the following section.

15.3.2 Technology Comparisons

There are a number of technical choices required in developing an MCM-D technology. Alternative approaches have been adopted by different companies; these are summarized in Table 15-2. Each company optimized different parameters. (Some of these approaches are described in the other case studies.) This is intended not to exhaustively list companies and processes, but to represent the diversity of technology. A primary distinction among the processes is the method of forming vias. Hughes chose a simple approach to achieve high yields. Joining the MCC allowed access to alternative technologies.

Microelectronics and Computer Technology Corporation

The MCC, a consortium of companies, has developed a copper plated thin film technology (Figure 15-13) that plates up the posts that become z-direction interconnects. The polyimide dielectric is then spun on and cured. Before the next plating process, the top surface must be lapped to provide planarization.

Table 15-2 Alternate HDMI Approaches.

FOUNDRY	UNIQUE FEATURE
• MCC	Via: Plated post
• Hughes	Via: Etched, staggered
• ATT*	Via: Filled, stacked
• GE	Overlay: Laser process
• nCHIP	SiO_2 dielectric

This provides a sturdy interconnect using copper (good conductivity) and the number of layers is not limited by planarity. This type of interconnect is also excellent for thermal conduction, which is important since the polyimide dielectric typically used is a thermal barrier.

AT&T

AT&T developed an Advanced VLSI Packaging (AVP) system, in which the via was sputtered and plated in copper and then filled with a nickel plug (see the discussion in Chapter 16 and also Figure 16-9). This allows the stacking of vias as in the MCC approach. AT&T has focused on flip chip bonding and, thus, needed a good thermal path under the solder bumps.

nCHIP

nCHIP departed from the norm by using silicon dioxide as the dielectric rather than polyimide. (See Figure 16-20 as well as the discussions in Chapter 7 pertaining to silicon dioxide as a dielectric material.) Silicon dioxide traditionally has been used in semiconductors for very thin insulation layers, but such thin layers create unacceptable capacitance when extended to the larger length and width of interconnect. nCHIP has developed techniques to create thicker layers of silicon dioxide without cracking. This has certain advantages, since silicon dioxide has good thermal conductivity.

General Electric

The GE overlay process (as described in Figure 7-17) is an innovative approach, requiring that the die be placed in cavities, followed by the interconnect layer on top of the die. This is done by applying a sheet of dielectric material. Vias are created with a laser by aligning the laser to the die pads. A key advantage of this approach is that the interconnect makes a sputtered connection to the die pad, eliminating the need for wire bonding. Another advantage is that the die

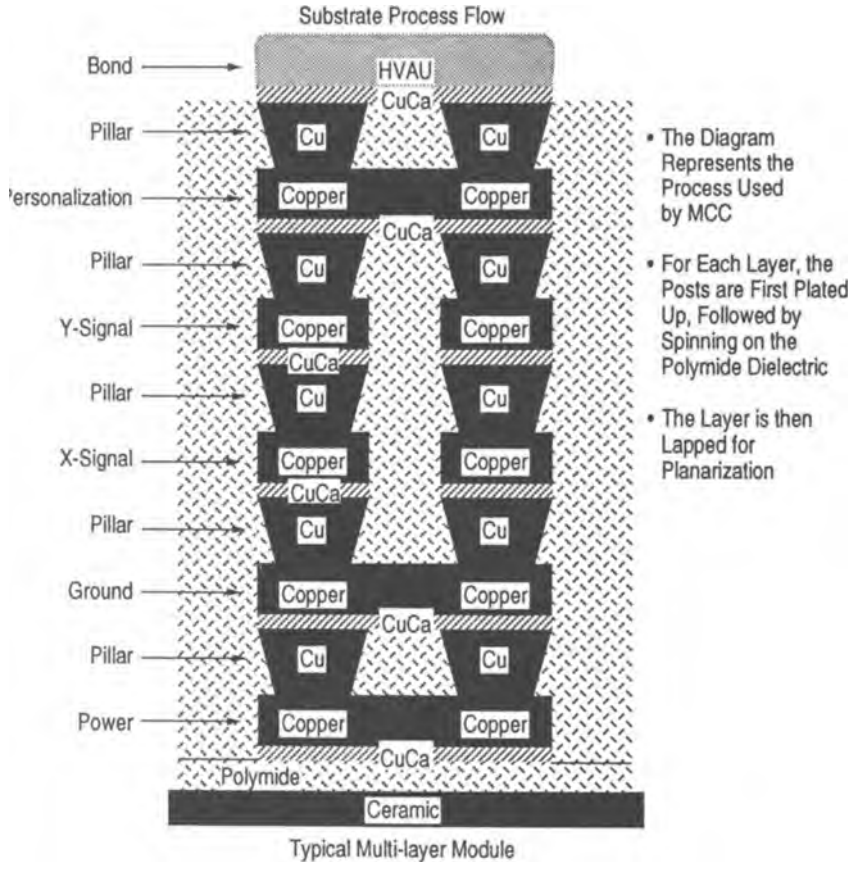

Figure 15-13 The MCC "Copper Post" process.

are placed directly on the base substrate without any intervening multilayer polyimide. The disadvantage is, that to replace a failed die, the entire interconnect layer must be stripped off and rebuilt.

15.3.3 HDMI Technology

High Density Multichip Interconnect (HDMI), is the term used by Hughes for a high density packaging approach based on techniques that were originally developed to produce semiconductor devices. The base technology involves the

use of a spin on polyimide dielectric, with alternate aluminum metallization layers on a variety of substrates, such as silicon, alumina, metal matrix or aluminum nitride [4]-[5].

Many options were considered when confronting the requirements of high interconnect density imposed both by system size and weight on one hand and increasing chip complexity and speed on the other. Two main thrusts were pursued: HDMI for computer-type interconnects, and LTCC for memory, microwave and low cost interconnects. Of course, HDMI can be used for microwave, and LTCC can be used for computers. Even though the technologies are radically different, they both relate to MCM technology.

The most pressing need for HDMI arose from the ATF program at Hughes (Figure 15-3). Initial ATF signal processor speed requirements were below 25 MHz, although evolutionary increases above that figure are expected. The technology has been characterized as high as 350 MHz. It was felt that alumina ceramic substrates could be used initially, and the resulting MCMs could be hermetically sealed within large area alumina ceramic packages. The use of silicon substrates offered the possibility of increased thermal conductivity and processing flatness, plus the possibility of incorporating active devices within the surface area. Thermal mismatch between silicon substrates and alumina ceramic packages created a potential problem, which could be solved through use of an aluminum nitride ceramic package. A family of such packages is being developed in the NOSC program described later in this section.

HDMI-1 describes a five layer baseline process developed primarily for low to moderate digital clock speeds. Substrates with additional layers (HDMI-1A, -1B, etc.) have been developed to extend the range of circuit applications.

Process Description of HDMI

The simplest way to describe this MCM-D process is to look at a cross section of a substrate (Figure 15-14). The interconnection is built on a base substrate, typically ceramic, by depositing a layer of metallization, followed by a layer of polyimide dielectric. Vias are created by etching through the dielectric and sputtering them to the next interconnection layer in the manner used for IC processing. The vias are typically staggered (not placed directly on top of a previous via) to avoid build up of nonplanarity, particularly in that they are not filled. These vias also can be stacked as discussed in Chapter 7.

HDMI technology uses multilayer thin film polyimide and aluminum applied in a Class 100 cleanroom. The substrate is a 150 mm diameter silicon or ceramic wafer. Silicon is available in well polished flat wafers for the semiconductor industry, but this application required new standards for flatness in the ceramic substrate. Fine ceramic dust particles are a hazard, since they may be introduced by the ceramic wafer into an otherwise ultraclean chamber.

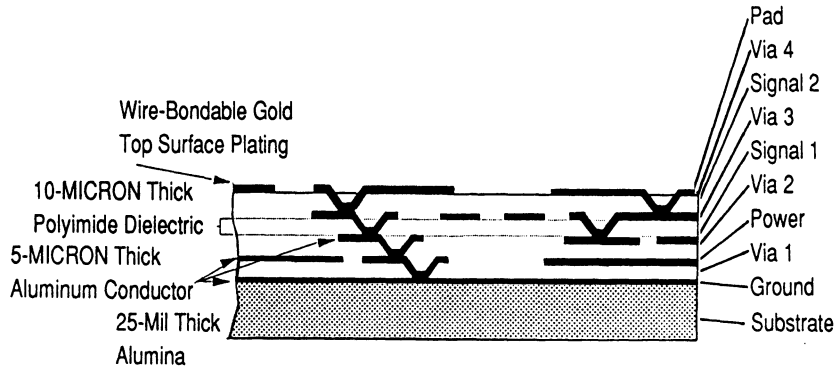

Figure 15-14 Cross section of Hughes HDMI-1 substrate with five aluminum/polyimide interconnection layers.

Polyimide is spun onto the wafers to a typical thickness of 5 μm. It is cured and then metallized by sputtering aluminum to form the ground plane. A second layer of polyimide is then applied and contacts to the ground plane are formed at selected locations, as described above. Via formation occurs by plasma etching after standard lithography is used to define the pattern. The vias are normally filled by aluminum, although alternatives are possible. Sputtering of the next aluminum layer then contacts the vias to form interconnects normal to the substrate. Photolithography and etching are then used to pattern the first level of signal interconnects with 100 μm spacings. (See Figure 15-15 for HDMI-1 baseline process flow.)

This polyimide spin on, via formation and conductor deposition process is repeated to form additional interconnection layers to complete fabrication of the substrate. The VLSI chips are then attached to the surface, with capacitors and other surface mounted component parts as necessary. The high density signal layers yield circuits of enormous complexity with just five conductor layers, one layer being the power plane. The HDMI base process was developed rapidly, because the basic components, using advanced sputtering and photolithographic techniques, were already available from MOS technology. High processing yields are achievable with this base process.

The key issue in production is yield. The substrate is best seen as a large (4" × 4") IC. One defect per wafer in a 150 mm process is 0.005 defects per square centimeter! This low defect density in a simplistic model yields 4" × 4" substrates at 0%; 2" × 4" substrates at 50%: and 2" × 2" substrates at 75%. Note that defect size is different from that of semiconductors: the typical defect is 10 μm or more - a boulder in normal semiconductor terms. However, the defect sources are numerous: the polyimide, the substrate and, especially, the

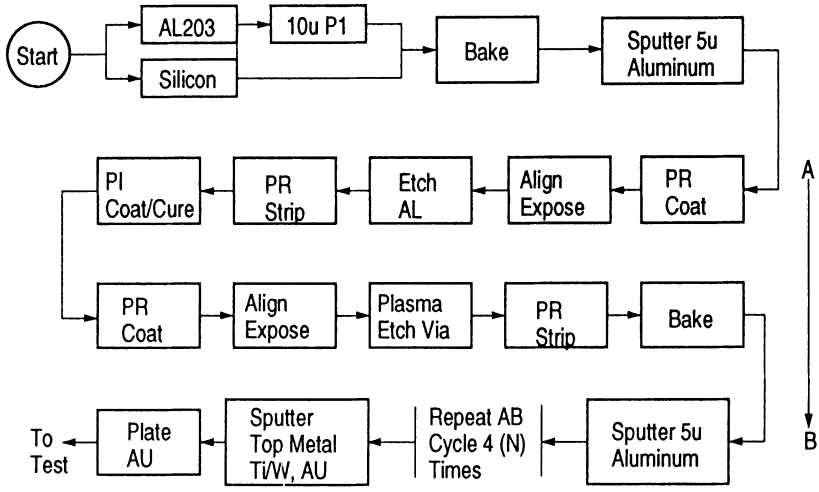

Figure 15-15 HDMI-1 substrate fabrication process flow.

equipment. This is one reason for using standard semiconductor equipment designed especially to eliminate production of any particles. It also relates to the problems of developing new equipment especially suited to substrates. Yield has two components: line yield - the percent of wafers that successfully complete all the process steps within specification, dependent on process maturity and control, and electrical yield - the percent of completed substrates that are good, dependent on defect density and substrate size.

As a process reaches maturity in the semiconductor industry, improvement and modification are needed to meet new demands. One demand was the addition of analog circuitry requiring three additional metal layers. Additional circuit layers may be cofired onto the alumina substrate prior to the first spun on polyimide layer or added as polyimide/conductor layers. This is indicated in Figure 15-16, which also shows the integration of capacitance into the silicon substrates by oxidization forming MOS capacitors. The need for some surface mounted devices is removed. The HDMI MCM performance can be extended to very high speed digital clock rates (2 - 5 GHz), by designing a configuration with signal and power planes closely matched in impedance.

In addition to the above developments, the pattern of history shows the need for even finer lines and spaces. It is possible to achieve 50 μm geometry, with 15 μm wide lines separated by 35 μm spaces. The resulting increase in the

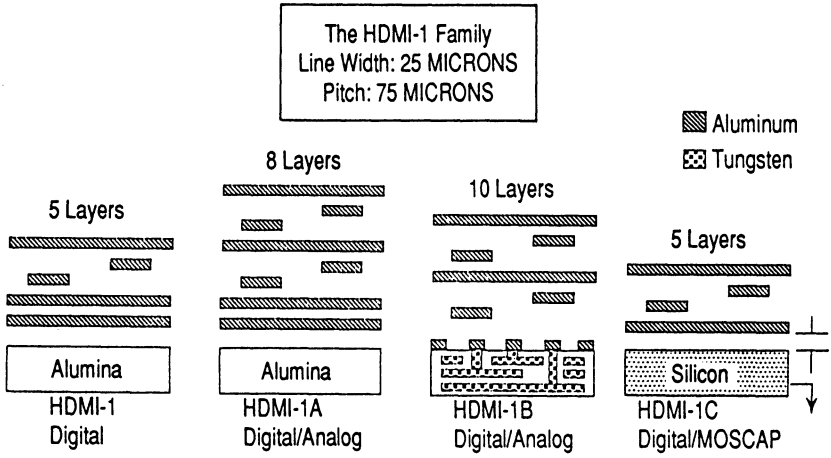

Figure 15-16 Schematic cross section views of progressively more complex substrate configurations. (Beginning with the baseline five layer digital system and extending to integration of capacitive functions for high speed digital processing.

conductor line resistance can be compensated by switching from aluminum to copper, since the conductivity of copper is about 60% higher. Copper does present new challenges; typically it does not adhere well to polyimide and also normally requires a migration barrier. Titanium/tungsten and chrome, respectively, are used as adhesion and diffusion barrier layers.

Table 15-3 shows key physical parameters and design features of the HDMI-1 substrates. This is a production process. Other processes are being developed continuously, but it is essential for production to standardize so that yields can be optimized. Table 15-4 shows electrical performance values for baseline the HDMI-1 substrates. Note that the two signal layers are not symmetrical (SIG1 is near the ground plane) and, hence, have different impedance values. This is not a problem, since most long interconnects contain several SIG1 and SIG2 segments, as illustrated in Figure 15-14, and the average effect predominates. For short interconnects, the impedance characteristics are not critical.

Related MCM-D Work

Hughes' participation in the NOSC VLSIC Packaging Technology (VPT) program included development of a 4.23" × 4.23" aluminum nitride package [6]. Successful performance under this program makes it possible to proceed with HDMI development on a dual substrate basis: alumina in an alumina ceramic package, and silicon in an aluminum nitride ceramic package. An example of a NOSC VPT HDMI test circuit module is shown in Figure 15-17.

Table 15-3 Some Design Features of Baseline HDMI-1 Substrates.

	STANDARD	SPECIAL
Substrate		
• Material	Silicon or alumina	Aluminum nitride
• Wafer size	6" (150 mm) diameter	—
• Substrate sizes	1.8" × 1.8" and 1.8" × 3.8"	3.8" × 3.8"
Conductors		
• Material	Aluminum	Copper
• Sheet resistance	6 mΩ/□	3 mΩ/□
• Width	25 μm	15 μm
• Pitch	75 - 100 μm	50 - 60 μm
• Number of layers	5 layers	7 layers
• Top metal	Titanium-tungsten/gold Titanium-tungsten/nickel/gold	Aluminum (for rad hard applications)
Dielectric		
• Material	Polyimide	—
• Thickness	5 - 10 μm	—
• Dielectric constant	2.9 at 1 MHz	—
• Dissipation factor	0.03% at 1 MHz	—
• CTE	3 ppm/°C	—
• Thermal stability	to 500°C	—
• Outgassing (1000 hours at 150°C)	< 5,000 ppm water	—
• Via diameter	35 μm	15 μm

The initial goals in the program applications of MCM-C and MCM-D technologies were to advance MCM-C and MCM-D packaging to achieve an order of magnitude improvement in all three DoD weapon system categories:

- Near term weapon programs
- Performance upgrades of existing weapons
- Radically superior future weapons

While much work remains, initial goals have been met. The greatest value of this project, to date, has been the development and building of high efficiency HDMI MCM-D and LTCC MCM-C modules that demonstrate packaging viabilities and that will be used to achieve performance improvements and costs savings in all three weapon system categories.

Table 15-4 Electrical Performance Values for Baseline Hughes HDMI-1 Substrates.

CHARACTERISTIC	SIG1	SIG2
• Capacitance (pF/in)	3.7	2.6
• Inductance (mH/in)	0.0085	0.0105
• Characteristic Impedance (Ω)	50	64
• Resistance (Ω/in)	9.0	8.1
• Crosstalk, near end (%) (50 μm space, 35 ps T_R)	2.6	5.0
• Crosstalk, far end (%)	2.3	3.8

The F-22 (formerly ATF) is one near term new program for which HDMI MCM-D packaging contributed to operational success. Weight and volume savings (96% and 93% respectively) were achieved for equivalent circuit functions when compared to prior F-15/F-15 radar packaging. By using very large scale integrated (VLSI) devices in conjunction with HDMI MCM-D packaging, the F/A-18 was provided with superior target resolution and higher land mapping resolution. This was a threefold packaging efficiency improvement and a 100% increase in radar processor throughput with a 50% cost reduction. These upgrades, brought about by packaging innovation, enable the F/A-18 to perform new reconnaissance missions which previously required both a reconnaissance plane and an escort.

HDMI packaging of LEAP electronics demonstrates two order weight and volume advancements and establishes credibility of lightweight processors which will be required in theater defense weapons: Brilliant Pebbles, Brilliant Eyes, Ground-Based Interceptor (GBI), advanced lightweight satellites and the Advanced Air to Air Missile (AAAM). LTCC MCM-C packaging technology will be used initially to build prototype transmit/receive (T/R) modules for the MMIC program. This will support active array radar systems developments proposed for upgrade applications in the F-15, F-15, and F/A-18 programs, the F-22 and Multirole Fighter (MRF) programs and in ground-based radar systems. Advancements in 3-D MCM chip on board (COB) packaging are continuing, and progressively more advanced versions of planar MCM packaging are being developed to achieve progressively greater performance reliability and cost advantages in military systems. Application of these advanced packaging

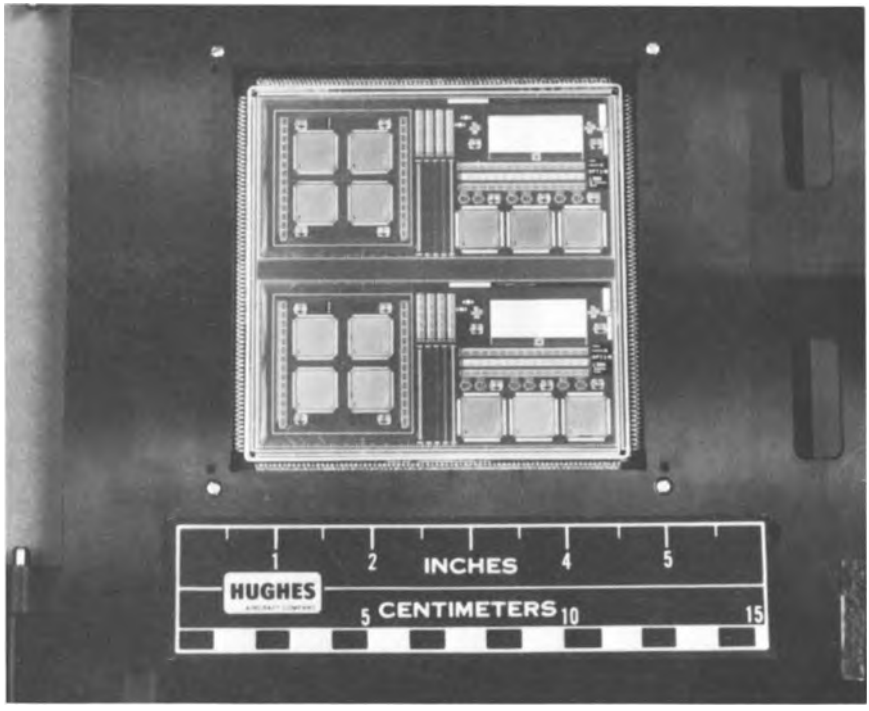

Figure 15-17 Photograph of 3.8" × 3.8" eight layer silicon-based hermetic aluminum nitride package developed under the NOSC VPT contract by Hughes, Coors Ceramics and W. R. Grace Company.

concepts to high end workstations, medical electronics, parallel processors and automotive electronics is a new focus in the commercial sector.

15.3.4 Low Temperature Cofired Ceramic (LTCC) Technology

One alternative method being developed at Hughes involves using a "green" ceramic tape to fabricate multilayer packaging. A schematic of a multilayer structure, together with an LTCC structure, shown in Figure 15-18, illustrates the substrate as an integral part of the package.

LTCC has two major advantages in meeting the challenges of both military and commercial packaging requirements. Although it uses thick film technology that does not have fine line definition, LTCC is effective in the use of the third

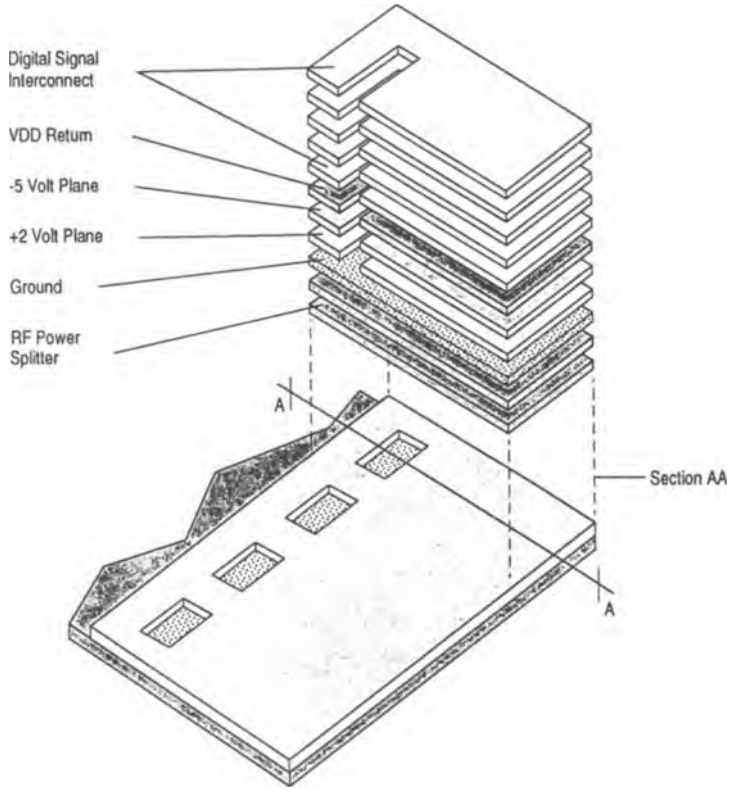

Figure 15-18 Artist's conceptual illustration of key features of LTCC interconnect package structure showing ability to install devices within cavities. Substrate integrated passive devices are also common LTCC structures.

or z-dimension. While polyimide-based packaging is currently limited to eight to ten conductor layers because of the stress build up, LTCC packages may have 20 or more layers. Prototype packages of 40 layers have been fabricated.

A second advantage of LTCC is its potential for low cost. In describing present trends in military packaging, the issue of cost was not stressed, but government program managers increasingly emphasize the importance of this issue as well as those of performance and reliability. LTCC has an advantage over conventional thick film technology because it does not require a firing stage after each conductor printing; this saves in processing time. Also, the use of ceramic tape lends itself to reel to reel processing through the lamination step. (LTCC technology is discussed further in Chapter 6.)

An additional advantage of LTCC is that of versatility. Large area packages can be made, limited only by the size of the tape and the maturity of the processing steps, with respect to preserving flatness. The composition of tape materials can be modified for expansion coefficient match with alumina, silicon or gallium arsenide. The dielectric constant may be quite low (4 - 5), but tapes with dielectric constants up to 100 or more can be obtained so that capacitors can be buried within the structure. Similarly, resistors and any other screen printable component parts may be totally embedded within the package, provided that excessive local stress is avoided. Similarly, cavities can be fabricated within the structure to accommodate embedded semiconductor chips as an alternative to surface mounting. The fired LTCC structures are hermetic. Semiconductor devices can be located within a sealed structure without the use of a separate package. Arrangements can be made for I/Os to be brought out to the edges of the LTCC structure for connection to external circuits.

LTCC has been found to be suitable particularly for microwave packages. The requirement here is for impedance matching, achievable through dimensional control of conductor lines and dielectric thickness. Gold is the metallurgy of choice for microwave applications because of its stability and relatively high conductivity. Although most tapes contain glass, the dielectric loss is acceptable for many military applications.

The low thermal conductivity of LTCC normally requires some arrangement to remove heat. The versatility of the structure does permit the insertion of conductive material blocks bondable into cavities formed below the ICs. Arrays of thermal vias also are used to transport heat effectively through the ceramic. The number of vias can be chosen according to the power to be dissipated.

LTCC tapes are now available from DuPont, ElectroScience Labs, Ferro and Rohm and Haas. Of these, DuPont has been most aggressive in developing a tape with low variability; their Green Tape™ Type 851AT has become the industry standard. A detailed review of this tape, with emphasis on economic analysis, has been given by Bender *et al.* [7]. Some early LTCC work at Hughes with this tape has also been reported [8]-[9].

The tape is formed by mixing glass, a ceramic or crystallizable glass powder and a suitable organic vehicle and placing this mixture onto Mylar. In forming a package, each layer of tape is first cut to size with registration marks. Then vias are formed in an appropriate pattern. The tape is patterned using thick film screen printing, including via fill. The individual layers are laminated together by the application of pressure and a small amount of heat to form the green, multilayer package. This then must be fired, with care taken to allow complete burnout of the organic material prior to densification of the glass-ceramic body. The important advance in LTCC over HTCC technology is that the maximum firing temperature is 850°C - 900°C, allowing the internal and external use of

gold, silver or copper metallization. High temperature cofiring is restricted to refractory internal conductors such as tungsten, with a resistivity three to four times that of silver.

DuPont has done valuable work in seeking to provide a wide range of conductor and resistor inks and dielectric tapes compatible with their LTCC tape. For military packaging, gold is normally required. Silver is available. A careful study conducted by Gaspar *et al.* [10] demonstrated that silver embedded within an LTCC structure can stand up well to military tests for mechanical integrity and electrical reliability under low DC voltages. Silver can cause warpage if present in high concentrations within a given layer, such as in a ground plane. Similarly, copper metallization is still developmental and suffers from the additional problem of oxidation susceptibility. Firing under a nitrogen atmosphere to prevent this increases production costs. An additional disadvantage is that organics must be evaporated rather than burned out. The concern is that densification will occur before all the organic material has been removed, potentially causing blistering as firing proceeds. Development of a trouble free, low cost silver or copper metallization is an urgent priority for LTCC technology in view of the high cost of gold. Table 15-5 shows key physical properties, dimensional tolerance and electrical properties of LTCC. Processing of LTCC tapes is discussed further in Chapter 6.

15.4 COMPREHENSIVE DESIGN/MANUFACTURING SERVICES

The technology and relationships discussed in Sections 15.1, 15.2 and 15.3 are important to applying successfully MCMs to complex, high performance military and commercial systems. Comprehensive design, fabrication and testing operations also are required to achieve MCM benefits successfully. To minimize cost, scheduling and technological changes in advanced programs, it is vital that comprehensive packaging systems be completed prior to program applications.

Figure 15-19 summarizes essential elements of a comprehensive packaging system. Many process steps are required and a broad scope of manufacturing and testing technologies are applied. Figure 15-20 shows a high performance 2" × 4" HDMI module to be used on the F/A-18 and F/22.

MCM technology is more difficult to apply than other more common PWB techniques. Because many chips are grouped into sealed units, problems associated with dense circuitry, such as testing and thermal dissipation, are amplified. MCM users need a complete design to assembly concurrent engineering process to apply this technology successfully. Such a process is described through a step by step review of Figure 15-19.

Table 15-5 Properties of LTCC.

	LTCC Package	Thick Film on 96% Alumina	High Temperature Cofired Package (92% Alumina)
Physical Properties			
• CTE at 300°C (ppm/°C)	7.9	6.4	6.0
• Density (g/cm^3)	2.9	3.7	3.6
• Camber (mils/in)	1-4[a]	1-2	1-4[a]
• Surface smoothness (microinches)	8.7	14.5	20.0
• Thermal conductivity (W/m-K)	16-20[b]	20[c]	14-18
• Flexural strength (kpsi)	22	40[c]	46
• Thickness/layer after firing (mils)	3.5-10.0	0.5-1.0[d]	5-20
Dimensional Thickness			
• Length and width	±0.2%	N/A	±1.0%
• Thickness	±0.5%	N/A	±5.0%
Electrical Properties			
• Insulation resistance (ohms @ 100 VCS)	$> 10^{12}$	$> 10^{12}$	$> 10^{12}$
• Breakdown field (volts/mil)	> 1,000	> 1,000	> 700
• Dielectric constant (1MHz)	7.1	9.3	8.9
• Dissipation factor (%)	0.3	0.3	0.03

[a] Function of firing setter and part design
[b] With thermal vias
[c] Property of substrate
[d] Property of printed dielectric

15.4.1 Design

The first requisite of a comprehensive MCM packaging system is that it be developed fully prior to systems applications. It includes design rules, design methods, standards design programs, standard design platforms and training of design personnel. Complete design analysis tools, characterizing electrical performance, mechanical performance, thermal performance, environmental withstanding capability, reliability and producibility must also be available prior to MCM applications in production systems. To support automated low cost MCM manufacture, computer aided testing (CAT) and computer aided manufacturing (CAM) databases, data extraction programs and formatting procedures also are required. New design capabilities had to be created, since MCMs are far more complex in comparison to more conventional hybrid microcircuits.

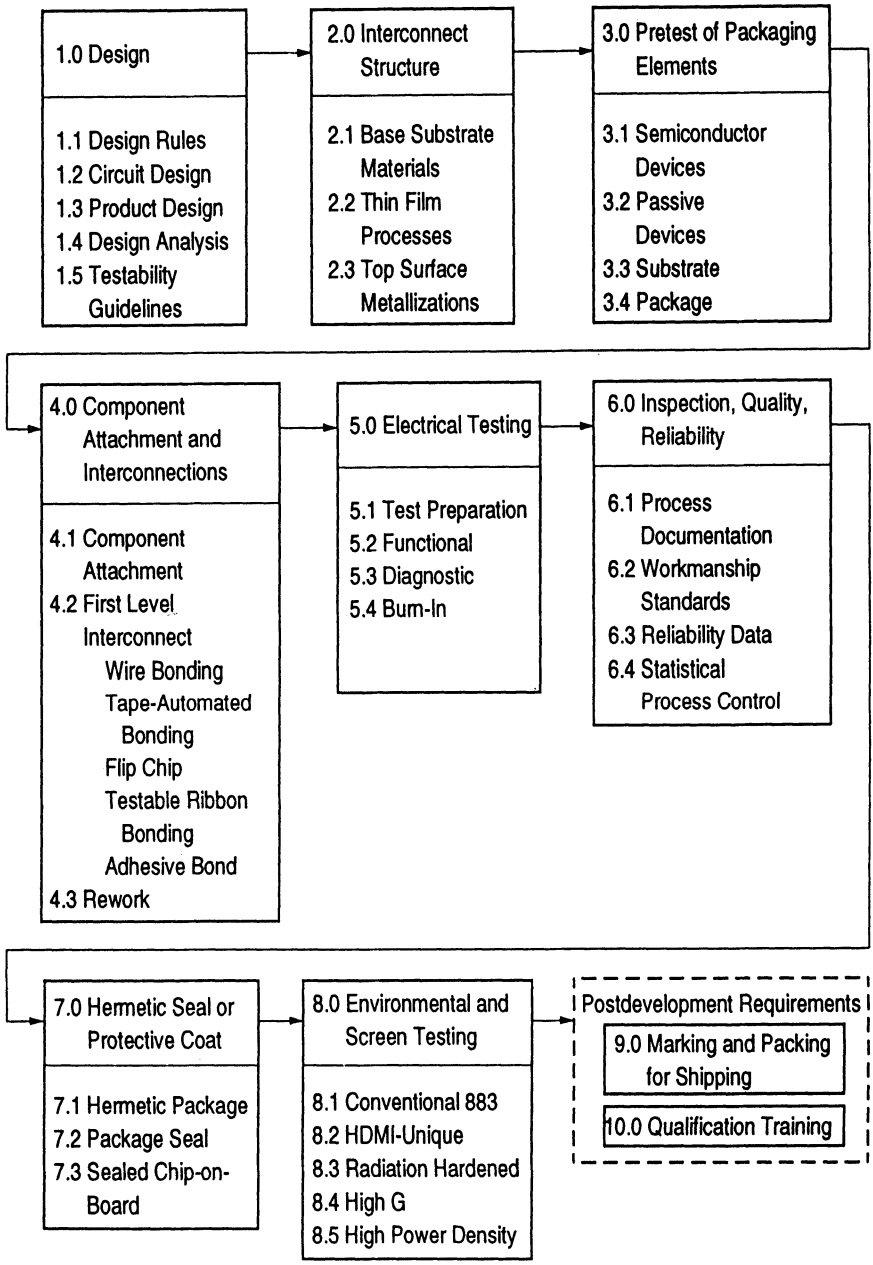

Figure 15-19 Essential elements of a comprehensive packaging system.

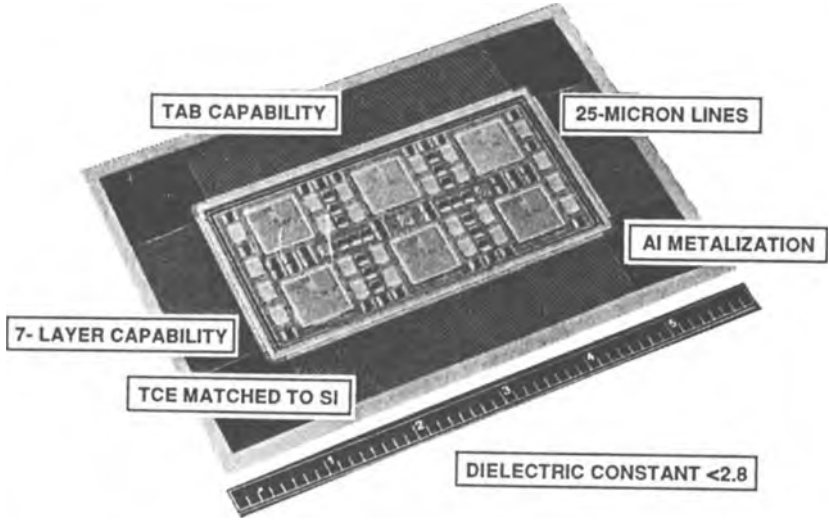

Figure 15-20 Typical 2" x 4" HDMI MCM-D used in F/A-18 and F-22.

Design Automation

Design automation is critical to placing MCM technology into the hands of board designers. A complete MCM CAD strategy must assist at every step, from design entry through verification and manufacturing. Design systems for hybrids usually have been modifications of IC or PWB tools, since hybrids are used much less frequently. The advent of commercial MCM use is changing this situation. A significant amount of commercial software specifically has been created to assist in MCM design.

At Hughes, CAD personnel initiate each customer project with up front strategic advisement. The first step is to identify areas where MCMs can best improve a particular project. Assisting designers to choose the correct MCM technology and set up a development plan is only the beginning. The design and testing of multichip packages requires a system level approach, forcing the integration of system, IC and packaging engineering. Work with customers focuses on understanding and integrating the various disciplines necessary to produce manufacturable, testable MCMs. The ability to integrate all these disciplines in a timely manner may provide the key competitive edge for

designers in the next decade. Complete production capability is offered for MCM users. Capabilities exist which include semiconductor processing, hybrid assembly, and experience in bare die purchasing. A need for new materials also can exist, as well as a need for the same processes, discipline and cleanroom facilities used to manufacture semiconductors.

Workstation Role

The most difficult task for MCM users is the design and verification of the circuit. Design automation can simplify much of the task. Mentor Graphics design automation tools were used because of the range of high density MCM technologies on engineering workstations they offer [10].

CAD support at Hughes began in proprietary efforts. Early in the development of its CAD program, it was found that most design automation tools supported either IC or PWB design, but not hybrids or MCMs. Therefore, new hybrid design software was developed on internal DEC VAX computers. This CAD system was based on a variation of an existing PWB router for its thick film processes.

As the company developed the CAD technology for internal HDMI efforts, it became apparent that new resources would be required to keep up with the rapid evolution of the CAD industry. An engineering effort therefore began to evaluate commercial vendors, seeking to adapt existing board or IC design tools to its nascent HDMI design process. To evaluate the capabilities of different vendors, the company used a benchmark circuit of a real processor design used in the ATF project. The benchmark circuit contained six 368 pad ASIC chips and more than a dozen memory chips. Several commercial tools were evaluated for their ability to place and route the circuit, given the special constraints of HDMI technology, such as blind and staggered vias.

Based on this benchmark, the Mentor Graphics program called "Hybrid Station™" was chosen to develop a commercial design capability. In an informal joint effort, Hughes and Mentor Graphics attacked the difficulties of MCM design, including via placement, test point extraction and development of GDS II output for the Hughes verification software. The end result of this effort was a design kit that, when used with the Mentor Hybrid Station™, allows systems designers to create high density digital MCM designs on their engineering workstations. In addition, there is in-house standardization of the internal MCM CAD, from design entry through verification, on Hybrid Station™.

Hybrid Station™ offers design capture, placement, routing and verification for hybrids and MCMs. It also can be used in conjunction with the Mentor Graphics tools for simulation, timing verification, electrical analysis and thermal analysis. The design kit contains the process rules and guidelines for design, placement and routing. With input from up front consulting efforts, the board

designer can fully specify an MCM. There is in-house testing of the final design and verification steps (adding test structures) that validate particularly complicated MCM aspects for manufacturing and assembly.

Product Design/Simulation

Having determined the design, testing technology and manufacturing parameters, the designer can begin to design the MCM circuit. At this point, assumed timing delays can suffice for first pass verification. The designer can also implement and verify testability of the circuit and begin to develop test vectors for the MCM.

Given a verified netlist for the design, the designer has the choice of performing physical design on Hybrid Station™ or proceeding directly to complete the MCM. Using Hybrid Station™, the designer can interactively or automatically place components depending on the timing constraints of the design. A critical task for layout is creation of the die outlines; they generally are nonstandard and not normally published by semiconductor manufacturers. The designer must enter the die outlines in parts libraries; Hughes provides template pad outlines in its design kit. With a placement that satisfies density and thermal requirements, the designer can begin routing with preplacement of critical nets. For designs with hundreds or thousands of nets, most of the routing is done automatically. The design kit provides process rules for the automatic router in Hybrid Station™. For very critical nets, the designer can reestimate parasitic effects based on actual routing and resimulate the design.

Design Analysis

Other analyses, such as those conducted on transmission lines, are available through Hybrid Station™ for MCMs, but generally are not necessary for moderate speed digital circuits (up to 100 MHz). The fundamental five layer HDMI modules have been characterized electrically with acceptable performance at clock rates up to 350 MHz. Most designs use simulation with nominal timing information or critical path analysis. After routing is completed, designers can verify that, based upon actual line lengths, specified electrical performance parameters have not been exceeded.

Given these design capabilities, a design process is in place for the commercial development of MCMs. The first step is joint in-house evaluation of the project. This initiates the concurrent engineering process necessary to meet all the criteria of for MCM manufacture, assembly, and test.

Thermal analysis of MCMs is particularly difficult. Most thermal dissipation occurs through conduction through the substrate. PWB thermal analysis tools, however, evaluate thermal convection from component parts to ambient air flowing across the PWB. Designers can use a spreadsheet to program basic

thermal equations for accuracy within about 10%. For thorough analysis, system designers can create three dimensional models using general purpose analysis tools such as ANSYS™ (a standard finite element analysis program), or using Mentor Graphics' Autotherm™.

Future Enhancements of MCM Design

The first enhancement will be further characterization of the interconnect systems to provide more parametric data. In addition to capacitance and resistance, inductance will be provided for estimating the impedance of interconnects. This step will become more important as microwave and very high frequency signals are used.

Another evolutionary step will help designers detail specific timing constraints up front. Worst case timing delays and relative timing sensitivity could be noted on schematics. Placement and routing tools then would use this information to preplace certain component parts, preroute critical nets and limit the overall length of timing sensitive lines. This timing driven layout will help to minimize the number of verification and design iterations needed to meet performance goals.

As design automation matures, larger production volumes and varieties of MCM products will be encouraged, tending to decrease manufacturing costs. These convergent trends will help to make MCM technology more accessible for commercial projects that require maximum functional density and performance.

Concurrent engineering of MCM design and manufacturing can identify areas of opportunity in the case of custom devices. For example, I/O pins on die can be arranged to simplify wire bonding and I/O drivers may be downsized according to the parasitic load of the substrate. Similarly, smaller die and I/O drivers can help to reduce the amount of bypass capacitance that may be necessary. Such tradeoffs can improve the operating frequency and yield of the MCM.

15.4.2 Interconnection Structures

The second essential element of a comprehensive MCM packaging system, as indicated in Figure 15-19, is the interconnection structure or substrate. LTCC and HDMI substrates are two possible interconnection structures.

15.4.3 Pretesting of Packaging System Elements

It is particularly important to achieve high initial manufacturing yield, minimizing MCM costs, because high performance MCMs are inherently complex. In early stages of production programs, when the circuit designs and

ASIC semiconductor devices are immature and not fully proven, it is usually advantageous to pretest all critical packaging system elements (Figure 15-19) including substrates, semiconductor devices, passive devices and high pin count hermetic packages. To minimize loss of modules caused by defects during the manufacturing cycle, substrates are tested for both continuity and shorts prior to and after temperature cycling.

The principal source of electrical test failures in the manufacture of complex MCMs is defective semiconductor devices. It is essential that only known good bare semiconductor devices be installed in MCMs. Semiconductor device manufacturers do not usually perform tests that are thorough enough to determine that the bare die satisfy circuit performance speeds and temperatures required of the MCMs. In addition, probe methods used by semiconductor manufacturers frequently cause bare die I/O pad damage, which precludes the multiple probing required for complete die pretesting.

To overcome limitations of conventional current bare die testing, alternative configurations were created, including beam leaded devices and tape automated bonding (TAB) formats. Still, a high proportion of MCMs employ wire bonding to form electrical connections from bare semiconductor device I/O pads to corresponding pads on the interconnection substrates. This involves use of bare die. Consequently, it was necessary to develop methods for performing complete dynamic pretesting of bare die at high, ambient and low temperatures to assure adequate yields of wire bond interconnected MCMs.

Probe testing of bare die minimizes pad damage. It was first used successfully on the LEAP program. This approach was based upon prior work conducted by Hewlett-Packard Corp. and IBM Corp. It utilized carefully balanced probe contacts with soft contact materials, in place of conventional hard tungsten probe tips. This approach was termed the "soft touch" probe card configuration. The soft touch probe cards were highly instrumental in achieving high initial test yield in the manufacture of LEAP MCM sector card assemblies. The extension of conventional probe card technology to high pin count advanced ASIC and microprocessor devices is difficult and costly. Therefore an even less damaging membrane contactor probe system was developed, to pretest in increasingly more complex devices with increasingly fine I/O pitches. Lithographically produced, patented membrane probe contactors show great promise for bare die and semiconductor wafer testing at clock rates as high as 5 GHz (Figure 15-21).

A further issue is pretesting hermetic packages in which valuable MCM assemblies are installed. The hermetic MCM package, shown in Figure 15-22, contains 368 leads on 0.030" centers. The procedure for pretesting this complex package is to install it into the electrical test fixture and perform continuity and short testing in conjunction with temperature cycling.

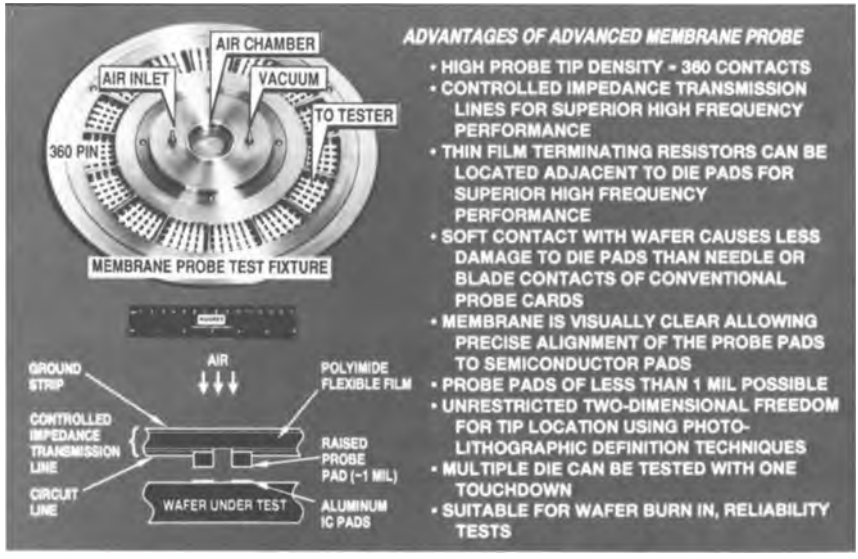

Figure 15-21 Membrane probe technology developed for die pretest.

15.4.4 Component Attachment and Connections

As indicated in Figure 15-19, the next principal packaging system element is attachment of component parts and formation of level 1 (chip level) connections. Manual assembly, while used extensively in less complex hybrid manufacture, is not well suited for production of MCMs. Precision and uniformity required in advanced MCM manufacture require high quality automated assembly equipment. Automated die attach systems are available commercially for high speed, low cost assembly of complex MCMs.

The evaluation and selection of die attach materials is quite demanding for high yield MCM production. For low power MCMs, organic adhesives are used typically for semiconductor and passive device attachment. Extensive tests of both thermoset and thermoplastic adhesives were conducted before adhesives for MCM component attachment were chosen. For MCMs with high power densities, solder attachment of devices to molytab heat spreaders, with subsequent solder attachment of molytabs to metallized pads on the MCM substrate, will result in low thermal resistance paths.

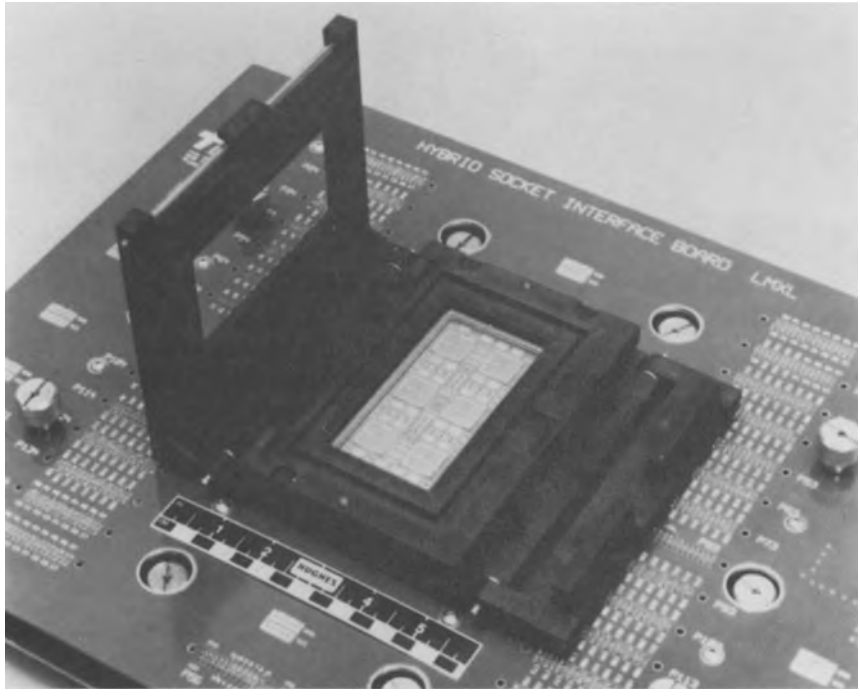

Figure 15-22 HDMI module in carrier and test socket on DUT board (fine pitch, 368 lead electrical test fixture).

First Level Connections

The baseline HDMI-1 and LTCC-1 packaging systems use thermosonic gold wire bonding to form electrical connections from bare semiconductor device I/O pads to corresponding pads on substrate top layer metallized pads. With recent decreases in I/O pad spacings on complex semiconductor devices below 0.004", new automated wire bond software and greater positioning precision had to be developed. To accommodate the great variety of MCM types and provide die pretest capability, JEDEC standard TAB technology was developed. To achieve even greater packaging efficiency, high density TAB technology was developed, which reduced the TAB footprint area by 75%. TAB nonrecurring tape and tooling costs are high, however, and TAB tape development cycle times are too great for use in early program development stages. A new connection method, software controlled and using testable ribbons was developed (Figure 15-23). This approach enables electrical test and burn in testing of bare die prior to

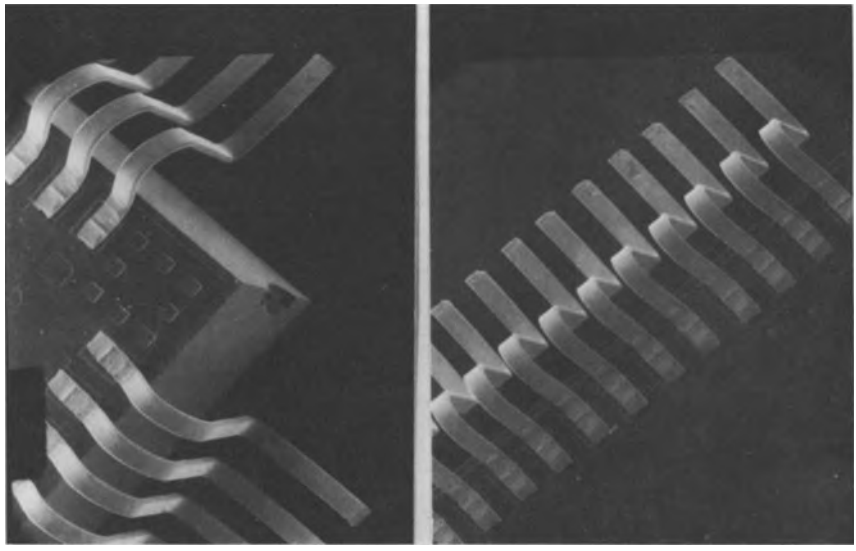

Figure 15-23 Testable ribbon bond (TRB) first level connection. This novel approach provides TAB-like pretest advantages while minimizing nonrecurring cost and cycle times. a) Face up version on left side, b) Flip TAB on right side.

MCM assembly. Moderate pitch flip chip technology is being used to further increase packaging efficiency and to reduce MCM manufacturing cost.

Rework

Rework of defective MCMs with many large, closely spaced devices has been a significant concern. An automated rework station was developed which can remove adjacent semiconductor devices spaced as closely as 0.020" with no resulting damage to neighboring devices or the substrate.

15.4.5 Electrical Testing

Even with extensive use of advanced semiconductor devices and MCM built in test features, functional and diagnostic testing of complex high pin count MCMs frequently exceeds the capabilities of commercially available automated test equipment. MCMs with as many as 1,000 I/O pins are now becoming common. For these complex high performance modules, very precise test fixtures are

required, as illustrated in Figure 15-22. It may be necessary to partition tests to account for I/O limitations of commercially available test equipment.

15.4.6 Inspection, Quality and Reliability

Because visual inspection methods are not adequate for large, complex MCMs with more than 100 devices and very small pattern features, automated optical inspection systems are now becoming standard in complex MCM manufacture. In addition, it is necessary to tailor some segments of existing military MCM specifications to accommodate increased package and substrate sizes. Similarly, new reliability models must now be generated to accurately predict overall systems reliability values.

15.4.7 Encapsulation

Hermetic alumina packages, illustrated in Figure 15-20, were initially developed for use in high performance airborne radar processor systems. To provide higher thermal performance and a close coefficient of thermal expansion match with silicon substrates, the team of Hughes, Coors Ceramics and W. R. Grace company developed a family of hermetic aluminum nitride packages under the NOSC VPT program. A 3.8" × 3.8" thick aluminum nitride package with 612 I/O contacts is shown in Figure 15-17.

The weight of a hermetic package into which the bare MCM assembly is installed often is approximately 50 times greater than the weight of the bare MCM assembly. These large hermetic packages also represent a significant portion of overall MCM cost. Accordingly, there are intensive incentives to develop passivation and protective coatings that obviate the need for hermetic packages. Consequently, there is now significant work within the electronics industry to develop and qualify suitable passivation and protective coatings. Several promising passivation and protective coatings, including silicon nitride, epoxy and silicone materials, are being developed and evaluated currently. Coated functional modules have passed Highly Accelerated Screen Tests (HAST). In addition, Hughes is a participant in the Air Force-funded Reliability without Hermeticity (RWoH) program to develop, evaluate and qualify SCOB coating.

Candidate materials are usually silicon dioxide or silicon nitride, alone or in combination [11]. The deposition temperature must be low to avoid damage to the circuitry. Plasma or ion assisted depositions are alternatives to purely chemical deposition methods being evaluated. Diamond-like carbon, a relatively new coating material, is of interest as a possible hermetic coating and will be included in future tests.

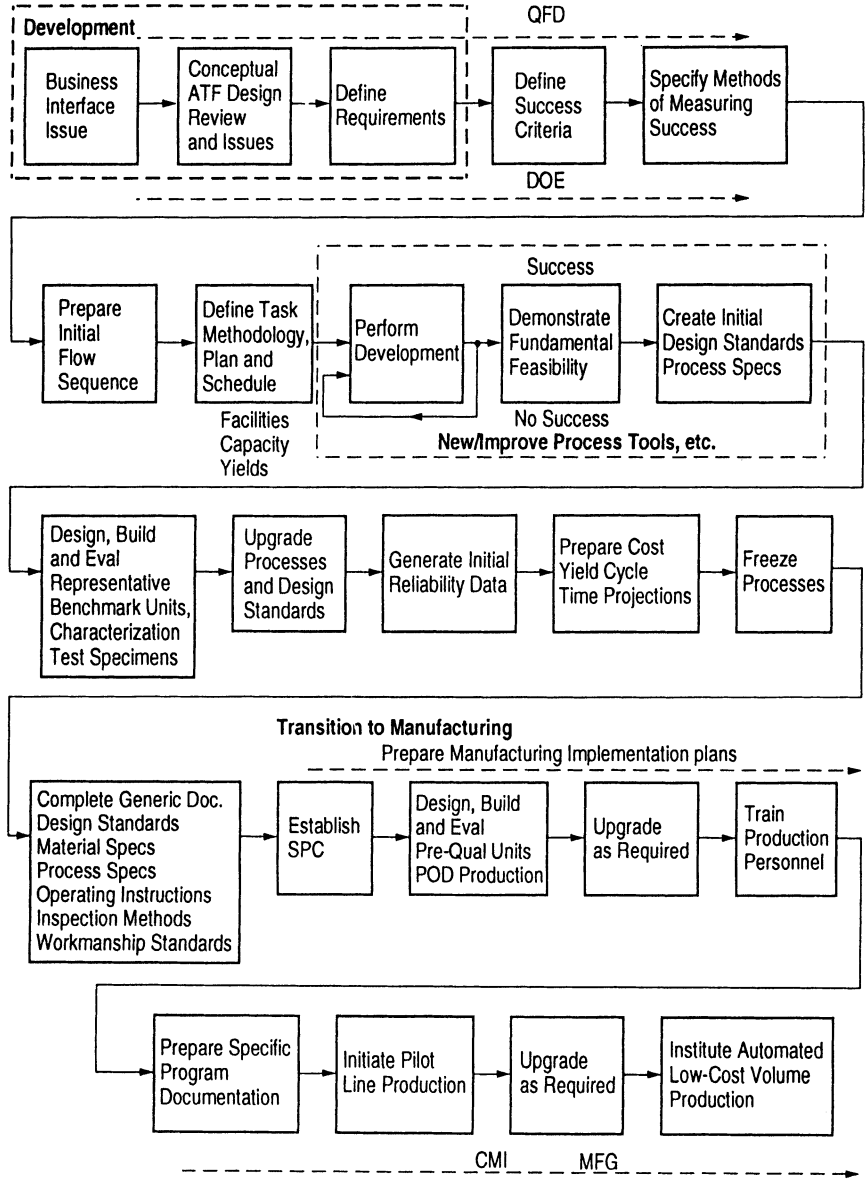

Figure 15-24 Packaging development steps showing transition to manufacturing.

15.5 SUMMARY

Like many aerospace manufacturers, Hughes Aircraft Company has recognized the significance of MCMs. As weapon systems technology becomes more complex, signal speeds and the need for high volumetric efficiency increase. Such changes have shifted interrelationships among systems and packaging engineers. Packaging knowledge now is becoming a strong factor in determining the extent and rate of overall technological process to be expected.

At Hughes, progress in two particular areas has determined the degree of overall packaging related advances occurring over the past several years. These two areas, HDMI and LTCC, have been described in this chapter as the core of a foundry which includes a full range of concurrent capabilities from computer aided design through manufacturing and testing. These capabilities are being applied to a broad range of applications. A possible sequence of development and transition to manufacturing is indicated in Figure 15-24.

References

1 F. C. Painter, "Building a Strategic Defense System," *Defense Electronics*, vol 21, p. 58, April 1989.

2 "Critical Technologies Plan for the Committees on Armed Services," *Department of Defense*, United States Congress: Defense Logistics Agency, Defense Technical Information Center, p. ES-1, March 1990.

3 J. J. Licari, "Fabrication and Packaging of High Density Multichip Interconnect Substrates," *Proc. NAECON*, (Dayton, OH), May 1989, pp. 1682-1688.

4 K. P. Shambrook, P. A. Trask, "High Density Multichip Interconnect," *Proc. 39th Electr. Components Conf.*, (Houston, TX), May 1989, pp. 656-662.

5 J. J. Licari, "Hybrid Thin-Film Processing Enters a New Era," *Elect. Packaging & Prod.*, vol. 29, no. 2, pp. 58-63, Sept. 1989.

6 L. E. Dolhert, *et al.*, "Performance and Reliability of Metallized Aluminum Nitride for Multichip Module Applications," *The Internat. J. for Hybrid Microelect.*, vol. 14, no. 4, Dec. 1991.

7 M. F. Bender, *et al.*, "Low-Temperature Cofired Tape Circuits vs. Thick Film: An Economic Analysis," *Proc. ISHM* , (Seattle WA), pp. 12-24, 1988.

8 W. A. Vitriol, J. I. Steinberg, "Development of a Low Temperature Cofired Ceramic Technology," *Proc. ISHM*, (Philadelphia PA), pp. 593-598, 1983.

9 R. G. Pond and W. A. Vitriol, "Custom Packaging in a Thick Film House Using Low Temperature Cofired Multilayer Ceramic Technology," *Proc. ISHM*, (Dallas TX), pp. 268-271, 1984.

10 K. P. Shambrook, E. C. Shi, and J. Isaac, "Bringing MCM Technology to Board Designers," *Printed Circuit Design*, vol. 7, no. 10, pp. 40-49, Oct. 1990.

11 G. Chandra, "Surface Protected Electronic Circuits," *Government Microcircuit Appl. Conf., GOMAC 1987 Digest,* (Orlando FL), vol. 13, pp. 86-88, Oct. 1987.

16

Silicon-Based Multichip Modules

R. Wayne Johnson

16.1 INTRODUCTION

Advances in silicon semiconductor processing and manufacturing technologies are allowing submicrometer scaling of semiconductor features, increasing both the transistor density and switching speed. The need to interconnect ever increasing numbers of transistors on the semiconductor chip has led to significant research and development in thin film multilayer technology. Current semiconductor chips may contain as many as four layers of metal interconnect. The issues addressed in the development of chip level multilayer interconnections include process compatibility, planarization, step coverage, build-in stress, and reliability [1]-[2]. Aluminum and aluminum alloys are used for interconnection metallization. A variety of glasses deposited by chemical vapor deposition (CVD) and spin-on techniques, and polyimides have been used as interlayer dielectrics. Complex processes have been developed to address planarization and step coverage issues.

As transistor switching speeds and semiconductor integration levels increase, the interconnection between chips becomes the limiting factor in realizing high speed system performance. Buschbom [3] has estimated that 55% of the cycle time of a high performance CPU is related to packaging and interconnection delays. Extending semiconductor technology to the level of wafer scale

integration (WSI) has been proposed to reduce these delays and to improve performance [4]. While successfully applied to highly repetitive circuits, WSI has been plagued by low yields and the need for elaborate redundancy or discretionary wiring schemes in typical logic systems.

Multichip thin film modules fabricated on silicon substrates provide an alternative to WSI in achieving high density, system level interconnections. The hybrid approach permits pretesting of the substrate and devices prior to assembly, repair of defective modules, flexibility in mixing device technologies and shorter development cycles. The concept of using a silicon substrate is not new. IBM [5] disclosed the concept of an active silicon chip carrier as early as 1972 where field effect transistors were mounted on the surface of a silicon substrate containing monolithic bipolar devices to produce mixed technology circuits.

The fabrication of multilayer thin film substrates is a natural extension of the on chip multilayer interconnection technology. The silicon MCM described by Speilberger [6] in 1984 utilized conventional IC aluminum glass multilayer processing to produce a silicon interconnect substrate. Hitachi also used this technology for a memory module reported in 1985 [7]. The requirements for multichip module (MCM) interconnections are different than for on chip interconnections. Key differences include the average interconnect length (resistive losses, signal attenuation), the module versus chip size (yield) and the need for controlled impedance lines (thicker dielectric layers). Evolution in the materials and processes has occurred in MCM fabrication. Current materials and processes will be discussed in the following sections along with a presentation of multichip applications. The fundamentals of these process steps can be found in Chapter 7. Issues in the selection of dielectric materials are covered in Chapter 8.

16.2 MATERIALS

16.2.1 Silicon Substrates

A number of substrate materials have been used for thin film MCM construction. Table 16-1 presents the properties for the more common substrate materials. In addition, MCMs have been fabricated on metal substrates to improve thermal performance in high power applications [8].

Silicon offers several benefits as a MCM substrate material including:

- Silicon is extremely smooth and flat and dimensionally stable, allowing fine line lithography. Current cleanroom facilities and lithography are used to fabricate semiconductor devices with < 1 μm features at an

Table 16-1 Material Properties for Common Substrate Materials.

	Si	Al_2O_3 99%	Al_2O_3 96%	SiC	AlN	Glass-Ceramic
ELECTRICAL PROPERTIES						
Volume Resistivity (Ω-cm)	2.3×10^5 (10^{-4} to 10^5 doped)	$\geq 10^{14}$	$\geq 10^{14}$	$\geq 10^{11}$	$\geq 10^{14}$	$\geq 10^{13}$
Dielectric Constant at 1 MHz	11.9	9.9	8.9-10.2	40	8.8	5-8
Dielectric Loss at 1 MHz	n/a	≤ 0.0004	0.001	0.05	≤ 0.001	0.002
Dielectric Strength (kV/mm)	n/a	25	23	1.0	14	1.5
THERMAL PROPERTIES						
Thermal Conductivity (W/m-K)	150	25	20	70-270	140-220	1-4
CTE from 25°C-400°C (ppm/°C)	3.5	6.5	7.1	3.8	4.1	3-8
Specific Heat (J/g-K)	0.71	0.88	0.88	0.8	0.7	
MECHANICAL PROPERTIES						
Density (g/cm^3)	2.33	3.89	3.75	3.2	3.26	2.9
Grain Size (mm)	Single crystal	1-5	5	2.5	5-10	1-5
Compressive Strength (kpsi)	700	375	340	560	300	
Tensile Strength (kpsi)	12-18	30	25	26	28	10-15
Bending Strength (kg/mm^2)	10-50	25-35	25-35	45	40-50	10-20
Young's Modulus (Mpsi)	23.5	40	40	59	40	15-20
Poisson's Ratio	0.28	0.22	0.22	0.15	0.25	0.26
Sintering Temperature (°C)	1412 (mp)	1300	1300	2250	1900	850
Maximum Use Temperature (°C)	1400	1500	1500	1900	1800	500

acceptable yield over areas approaching 1 cm^2. Scaling of multilayer thin film substrate features to 10 - 25 μm should allow acceptable wafer scale interconnection yields.

- Silicon is compatible with thin film multilayer processing.

- The use of a silicon substrate with silicon chips eliminates the coefficient of thermal expansion (CTE) mismatch, between the chip and substrate improving die attach reliability for large area VLSI and VHSIC devices. This is a major concern with flip chip connections since the stress level increases radially outward from the center of the chip. With an alumina substrate, the expansion mismatch is approximately two to one and the maximum chip size is limited for reliable flip chip solder connections. These effects also were discussed in Section 9.6 and are illustrated in Figure 9-42 [9]. Aluminum nitride (AlN) and the cofired glass/ceramic substrate recently developed by IBM for flip chip connections [10] have CTEs which closely match that of silicon. See Chapter 9 for additional discussion of the flip chip solder bump connection technology.

- The thermal conductivity of silicon is six to eight times higher than alumina and is comparable to silicon carbide and aluminum nitride. Since MCMs have higher packaging densities and higher power densities, this is an important consideration if the heat is to be removed through the substrate.

- Active and passive devices can be fabricated in the silicon substrate. Gigabit Logic [11] has fabricated diffused termination resistors and a reversed biased decoupling capacitor in a silicon carrier for a single chip GaAs package. The package is illustrated in Figure 16-1. Potential active circuits include memory arrays, voltage regulators for chips requiring 3.3 volts or other voltages, and off-chip line drivers to reduce chip size and on-chip power dissipation. To maximize substrate yield, the active circuits should be highly repetitive, allowing selective wiring of tested good circuits. MCM examples of active (telecom subscriber line card) and passive (conductive ground plane and decoupling capacitor) circuit elements using the silicon substrate will be discussed in a later section.

- Optical wave guides can be formed on silicon substrates with detectors fabricated in the silicon for electro-optic and photonic systems [12]-[16].

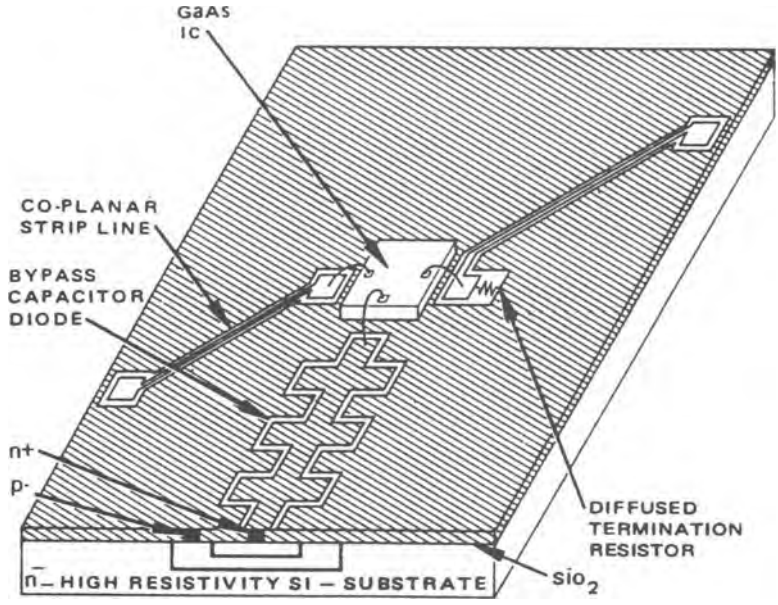

Figure 16-1 Illustration of silicon substrate carrier for single chip GaAs packaging [11].

- An established base of highly automated fabrication facilities for 4", 5" and 6" wafers exists. Semiconductor manufacturers are integrated vertically to produce silicon MCMs to meet specific application requirements. This will be particularly beneficial in producing MCMs with ASICs.

- Silicon can be micromachined to fabricate sensors [17]-[19] and advanced cooling structures [20]-[21] and to align optical fibers. In sensor applications, it often is necessary to place the signal conditioning electronics as close as possible to the sensor due to the typically low signal level of the sensor. Research is being directed to integrating the sensor and the electronics monolithically in the same piece of silicon. However, the process integration has been difficult, reducing yields and compromising electrical performance. MCM technology would allow interconnection of separately fabricated electronic devices onto a silicon substrate containing the micromachined sensor structure.

Microchannels, re-entrant cavities [20]-[21] and pyramids to initiate nucleated boiling have been fabricated in silicon to improve thermal performance. V-grooves etched in silicon provide a method to align optical fibers for off substrate optical communication.

Silicon does have its limitations as a substrate. Wafer bowing occurs during processing due to stresses arising from the CTE mismatch between the thin film dielectrics and the substrate. The stress in the thin films may be calculated using:

$$\Delta\sigma_f = \frac{E_f}{(1-\nu_f)}(\alpha_s - \alpha_f)\Delta T \qquad (16\text{-}1)$$

where: σ_f = film stress,
E_f = Young's modulus of the film
ν_f = Poisson's ratio of the film,
α_s = substrate CTE,
α_f = film CTE and
T = temperature.

The substrate radius of curvature, r, is given by the following equation:

$$r = \frac{E_s D^2}{6\sigma_f t(1-\nu)} \qquad (16\text{-}2)$$

where: E_s = Young's modulus of the substrate,
D = thickness of the substrate,
t = the film thickness and
ν = Poisson's ratio.

Young's modulus for silicon is approximately 50% of that for alumina and aluminum nitride. Thus, for equal stress levels and film and substrate thicknesses, the silicon substrate will bow more than the ceramic substrates. Substrate bowing adversely affects automated vacuum handling systems and fine line imaging. Lower stress dielectric films are being developed to address this issue. Thicker wafers also can be used to reduce bowing. Hagge [38] has estimated silicon wafers can be deflected to 0.010 in./in. bow without fracture and can withstand 10^8 bending cycles if the strain range is kept below 0.0005 in./in. For a sample subjected to bending, the strain is calculated as the change in length of the tensile (compressive) surface divided by the length. The strain range is the sum of the compressive and tensile strain.

A second limitation is the need for additional packaging of the silicon substrate for mechanical protection and to provide input/outputs (I/Os) for the next level interface. Various packaging options currently being used will be discussed in a later section. In all cases, the packaging schemes use perimeter I/Os. Internal output signals must travel to the substrate edge before exiting the module, increasing delay time. Power distribution to the substrate interior is also more difficult with perimeter interfaces than with multilayer cofired ceramic substrates using a pin grid array. The number of I/Os is limited by the substrate perimeter and the spatial fanout requirements of the next packaging level. At least one nonperimeter approach has been proposed by Little, *et al.* [22]. In this approach, aluminum is diffused through the silicon substrate to provide an area array of contacts on the substrate backside.

The CTE mismatch between the silicon substrate and the package base must be considered in package selection. The two options are to nearly match the CTE of the silicon substrate (AlN, SiC, mullite packages) or to use a compliant adhesive (silicones, modified epoxies) to attach the substrate into the package. If heat removal from the die is through the substrate and the package base, the thermal conductivity of the substrate attachment material and the package base material must also factor into the decision.

Silicon substrates currently are used in approximately 44% of the thin film MCMs produced worldwide [23]. The choice of substrate material depends on the specific application requirements and the fabrication facilities used. The ideal substrate does not exist. Therefore, substrate choice is based on optimizing the most critical parameters for a particular application. The preceding paragraphs have highlighted the advantages and disadvantages of silicon substrates and have presented comparative properties for other substrate materials. This can serve as a basis for making initial decisions on substrate materials.

16.2.2 Conductors

Aluminum, copper and gold are the three most common interconnection metallizations. Their properties are summarized in Table 16-2.

Early silicon MCMs used aluminum metallization as an extension of semiconductor technology and it is still widely used. The key advantage of aluminum is its ease of processing using sputtering and wet chemical etching. No adhesion or barrier layers are required typically. As operating speeds increase and line widths decrease, the resistivity of aluminum limits its performance. Copper metallization systems have been developed for this reason. Copper requires an adhesion layer typically chromium (Cr) or titanium (Ti). It can be sputter deposited, but is difficult to etch by wet chemical methods.

Table 16-2 Properties of Common Thin Film Metallizations.

	ALUMINUM	COPPER	GOLD
Resistivity	2.66 μΩ·cm	1.67 μΩ·cm	2.35 μΩ·cm
Adhesion layer used	Good adhesion layer	Ti or Cr adhesion used	Cr, TaN or NiCr
Deposition	Sputtering	Sputtering, plating	Sputtering, plating
Corrosion	Corrodes	Corrodes without barrier layer	No

Note: Copper reacts with polyamic acid; requires barrier layer (Ni or Cr).

Liftoff processes have been developed for fine line patterning of sputtered copper [24]. Copper also can be pattern plated to achieve fine lines. In a typical plating process, an adhesion layer followed by a thin copper layer is sputter deposited. Then photoresist is applied and imaged. After electroplating the thicker copper pattern, the photoresist is removed and the background copper and adhesion layer are etched away. When using polyimide dielectrics, the polyamic acid precursor reacts with the copper during the polyimide cure affecting local electrical properties and the polyimide to copper adhesion. Barrier layers such as nickel and chromium are used to protect the copper. These may be deposited by plating or sputtered and patterned by etching techniques. A number of variations exist for depositing and patterning copper metallization but the details will not be discussed here.

Gold metallization also has been used in MCM fabrication [8]. The corrosion resistance of gold is its primary attraction. Gold requires an adhesion layer, but does not react with polyamic acid and no barrier layer is needed. It is typically pattern plated.

Other interconnection systems including superconductors and optical wave guides have been discussed for MCMs and may be used in the future [15]-[16], [25].

16.2.3 Dielectrics

Thin film dielectrics used on silicon substrates include silicon dioxide, polyimide, modified polyimides, benzocyclobutene, and polyphenylquinoxaline. A summary of typical properties is presented in Table 16-3. While this table can be used for the general comparison of dielectrics, caution is advised since the data presented

Table 16-3 Summary of Interlayer Dielectric Properties.

	Standard Polyimide	Fluorinated Polyimides	Silicone Polyimides	Acetylene Terminated Polyimides	Low Stress Polyimides	BCBs	PPQs	SiO_2
ELECTRICAL PROPERTIES								
Dielectric Constant (1 kHz)	3.4-3.8	2.7-3.0	3.0-3.5	2.8-3.2	2.9 (z) 3.4 (xy)	2.7	2.7	3.5
Dissipation Factor (1 kHz)	0.002	0.002		0.002	0.002	0.0006	0.0005	0.0001
THERMAL PROPERTIES								
Decomposition Temperature (°C)	520-550	> 470	450	500-520	620-650	430	500	> 1000
Glass Transition Temperature (°C)	300-320	< 300	< 300	215-230	>400	350-360	365	
CTE (ppm)	20-40			38	3-6	45-70	55	0.5
Moisture Absorption (% wt.)	1.1-3.5	0.7	0.9	0.8-1.3	0.4-0.5	0.3-0.5	0.9	< 0.1
MECHANICAL PROPERTIES								
Young's Modulus (kpsi)	400-1000		50-250	300-450	1280	340		10,000
Stress (10^8 dynes/cm^2)	3	5.3	3	4	0.4-0.5	3.7		±20
PLANARIZATION								
Degree of Planarization (%)	24-34	25-30	22	91-93	10-28	91		0-5
Percent Solids	14-17	15-19	26	35	10-13.5	35-62	35-47	n/a
Viscosity (poise)	11-19	15-25	12	15	20-110	0.15-1.5	11-15	n/a
PROCESSING								
Deposition method	Spin coat, spray	Spin coat, spray	Spin coat, spray	Spin coat, spray	Spin coat, spray	Spin coat, spray	Spin coat, spray	CVD
Curing temperature (°C)	350-400		300-400	300-350	350-400	230-250	400-450	300-400

BCB = Benzocyclobutene PPQ = Polyphenylquinoxaline SiO_2 = Silicon Dioxide

were collected from various references. The dielectric processing and test methods vary and the individual references should be reviewed for specific details. See Chapter 8 and 7 for a more detailed discussion of dielectrics and processing with dielectrics, respectively.

Silicon dioxide, used in early modules, has found limited acceptance. Thick silicon dioxide layers, necessary for controlled impedance transmission lines, have been difficult to deposit crack-free. Horton [39] recently has reported the deposition of silicon dioxide layers greater than 10 μm thick with controllable stress levels by plasma chemical vapor deposition. Examples of this technology will be presented in a later section. Silicon dioxide has the advantage of negligible water uptake, chemical inertness and higher thermal conductivity than organic dielectrics; its dielectric constant is higher than organic dielectrics and it is conformal. The resistance of the dielectric to moisture penetration is a significant advantage for MCMs in nonhermetic packages.

Polyimides are the most widely used interlayer dielectrics. The dielectric typically is applied as a solution containing polyamic acid which is converted by a condensation reaction at high temperature (> 350°C). Standard polyimides are stable at high temperature and compatible with thin film processing and module assembly. Limitations of standard polyimides include high moisture uptake, only moderate planarization and reactivity during cure with copper. Moisture absorption by the cured film effects its electrical properties with the dielectric constant increasing with increasing moisture content. During processing, dehydration bakes are necessary before metal deposition to avoid blistering of the metal layer during subsequent high temperature processing steps. Solid metal power and ground planes must be avoided. Meshed power and ground planes should be used to allow moisture in the polyimide layers below to outdiffuse without blistering the metal planes. The low degree of planarization (20 - 35%) achieved with polyimides results in a loss of planarity as the number of metal and dielectric layers increases. Planarity is important for fine line lithography, step coverage and uniform characteristic impedance. Multiple dielectric coatings improve planarization. Techniques including polishing have been developed to maintain planarity.

Polyimides have been modified to produce materials with lower dielectric constants, lower moisture uptake, lower CTE (low stress) or which are photoimageable. The modifications are typically compromises with other film properties and specific materials should be evaluated against actual dielectric requirements. Compared to standard polyimides, the fluorinated and silicone polyimides have a somewhat lower moisture uptake and the fluorinated polyimides have a slightly lower dielectric constant. Both materials, however, have found limited application to date. The siloxane modified polyimides typically are preimidized and dissolved in suitable solvents, requiring only

solvent evaporation during the cure process. Since the material is preimidized, no polyamic acid is present and no water of condensation is released during cure, reducing its reactivity with copper. The acetylene terminated polyimides also are preimidized and crosslink via the acetylene groups located at the ends of the oligomer without releasing volatiles. Fluorinated, acetylene terminated polyimides have been formulated to improve the solvent resistance as compared to earlier fluorinated polyimides.

There is significant interest in the photoimageable materials as a means of simplifying processing for via formation [30]. By exposing the polyimide and developing it by wet chemistry, the process steps of depositing and patterning an etch mask, plasma etching the polyimide and stripping the etch mask are eliminated. Film shrinkage and via distortion during cure (after imaging) have limited the acceptance of these materials in the United States, however there is considerable usage in Japan and development work is continuing.

Benzocyclobutene (BCB) is a highly crosslinked polymer derived by the thermal rearrangement (> 200°C) of bisbenzocyclobutene monomers. No reaction byproducts are produced in the curing process. The prepolymer is B-staged and dissolved in a solvent for application to substrates. BCB has a low dielectric constant, has low moisture absorption, is planarizing, cures at lower temperatures and does not react with copper during curing. Its thermal stability limit is approximately 350°C, which is compatible with most processing and assembly operations, but may be insufficient for certain high temperature metal anneals and soldering operations.

Polyphenylquinoxaline (PPQ) recently has been explored as an interlayer dielectric [26]. This material is supplied as a fully cyclized polymer dissolved in a solvent. Only solvent evaporation is required to cure the dielectric film. The material has a low dielectric constant and a high thermal stability. Additional characterization of this material is required.

Research is continuing to improve existing materials and to develop new ones. Other materials currently being investigated include fluoropolymers and polyquinoline.

16.3 MCM EXAMPLES

Extensive research and development programs in silicon MCMs exist at companies and universities around the world. Many have produced modules. A sampling of these will be presented in this section to highlight the various processing, material and packaging options.

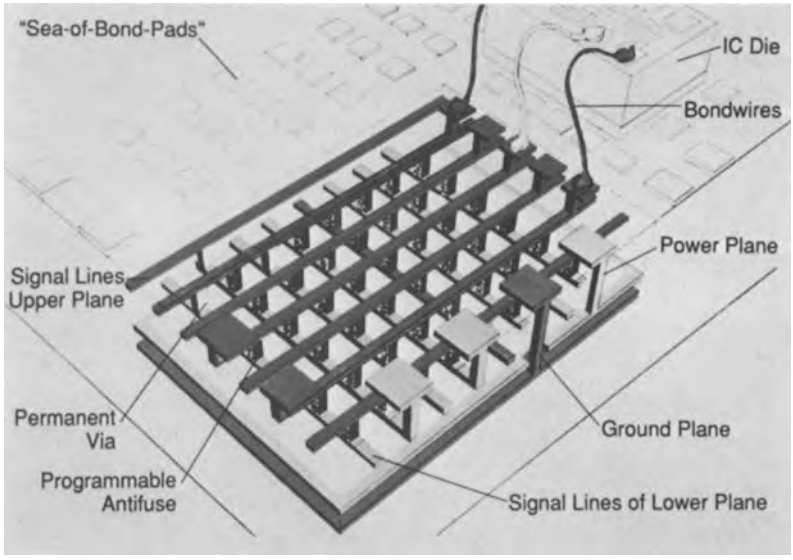

Figure 16-2 Illustration of Mosaic System's programmable silicon substrate. (Courtesy of ERIM.)

16.3.1 Prototype Silicon Substrates

As with any new technology, the development costs for prototypes are high and may be difficult to justify for small volume applications. Development times typically are also long. Mosaic Systems developed an electrically programmable silicon substrate technology to provide rapid prototyping capability [31]. A proprietary antifuse technology, based upon amorphous silicon, is used to form programmable connections between an x and y grid of aluminum lines as shown in Figure 16-2. All x and y lines extend to the perimeter of the substrate for programming. By applying a threshold voltage (approximately 20 V) between appropriate x and y lines, the amorphous silicon crystallizes, forming a conductive path between the two lines. Substrates can be mass produced and programmed directly from a CAD file. Integrated circuit chips are epoxy attached to the silicon substrate and electrically connected by ultrasonic aluminum wire bonding.

The Environmental Research Institute of Michigan (ERIM) has used the Mosaic Systems technology to fabricate the SEM-E memory module shown in

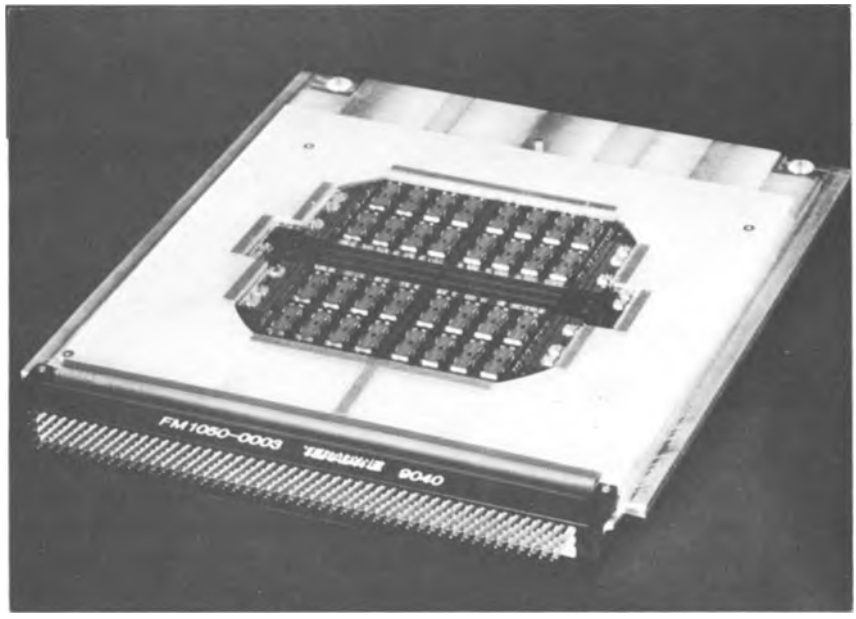

Figure 16-3 SEM-E Memory module. (Courtesy of ERIM).

Figure 16-3 [32]. The module contains 32 32k × 8-bit memory chips. The silicon substrate is mounted on a metal matrix composite (Al graphite) frame. A second example from ERIM is the Symbolics MacIvory processor module shown in Figure 16-4. In this implementation the circuit is partitioned into four smaller silicon substrates which are mounted subsequently on a programmable 4" diameter silicon substrate. The single MCM replaces two 12.5" × 3.9" printed wiring boards (PWBs) and contains 112 components. The thin silicon dioxide layers and the resistance (typically 4 Ω) of the crystallized via connections currently limits this technology to applications below approximately 40 MHz.

A prototype substrate technology based on copper conductors and polyimide dielectrics has been developed by Oerlikon-Contraves AG [33]. In their approach, a grid of x and y conductor lines are separated by a polyimide layer. At crossover points the dielectric is removed between the two conductor traces by plasma etching to form an air bridge. The resulting structure is shown in Figure 16-5. The air gap supports approximately 300 V. To program the substrate, appropriate x and y lines are connected at crossover points by using

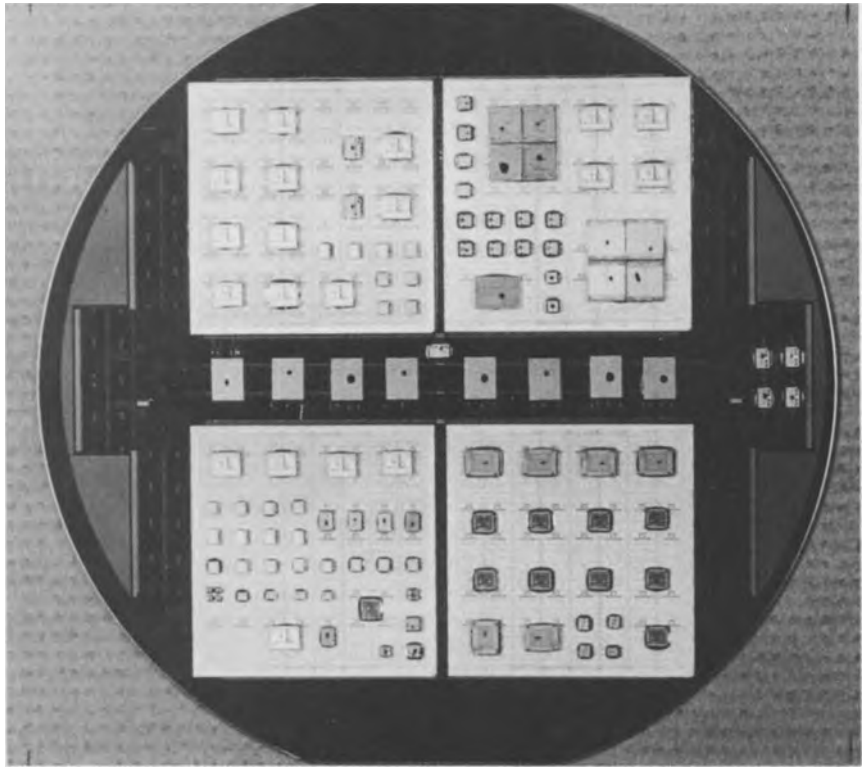

Figure 16-4 Symbolics MacIvory processor module. (Courtesy of ERIM.)

a modified ultrasonic wire bonder to weld the two conductors. A welded contact is shown in Figure 16-6. Laser welding techniques are under development. The laser is being used also to sever unused line segments from the node. These stubs could effect electrical performance at high frequency. Figure 16-7 illustrates a completed assembly.

Another technique, which is similar in philosophy to a gate array design, has been used by MCC (Microelectronics and Computer Technology Corp., Austin TX) for thin film MCM on ceramic substrates, but is also applicable to silicon substrate modules [34]. In their process, the substrate is fabricated using copper and polyimide to produce a standard interconnection pattern with power and

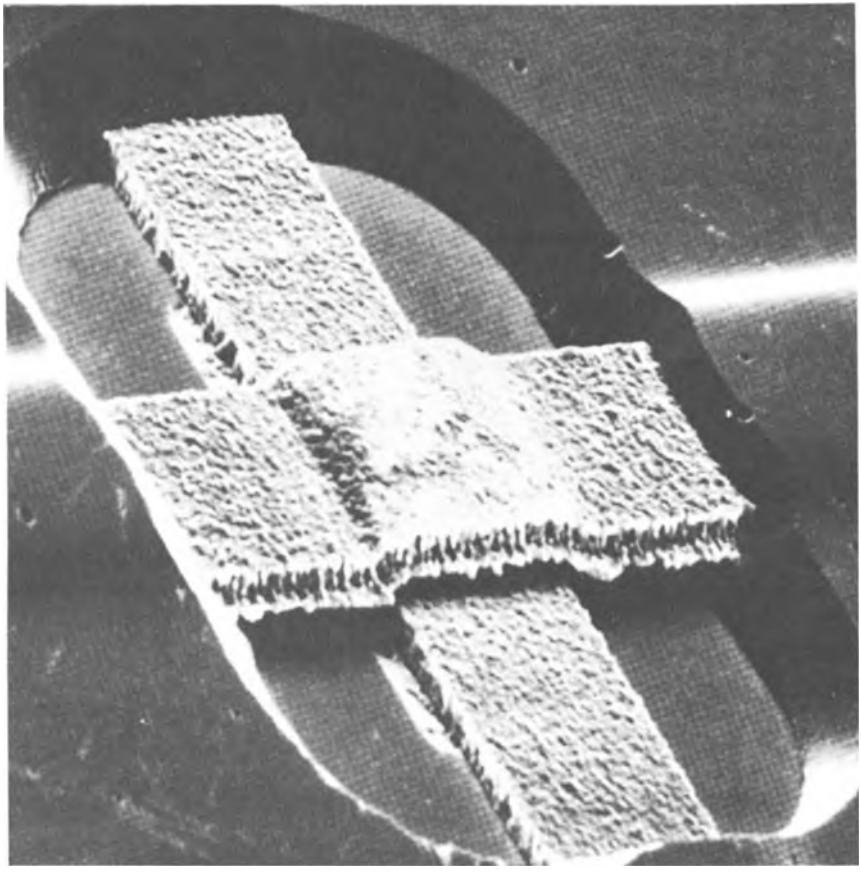

Figure 16-5 Air bridge structure for programmable x-y connections. (Courtesy of Oerlikon-Contraves AG.)

ground planes and short x and y signal line segments. These substrates are placed in inventory. The final interconnection layers are added to connect the appropriate signal line segments providing the necessary signal routing to customize the module.

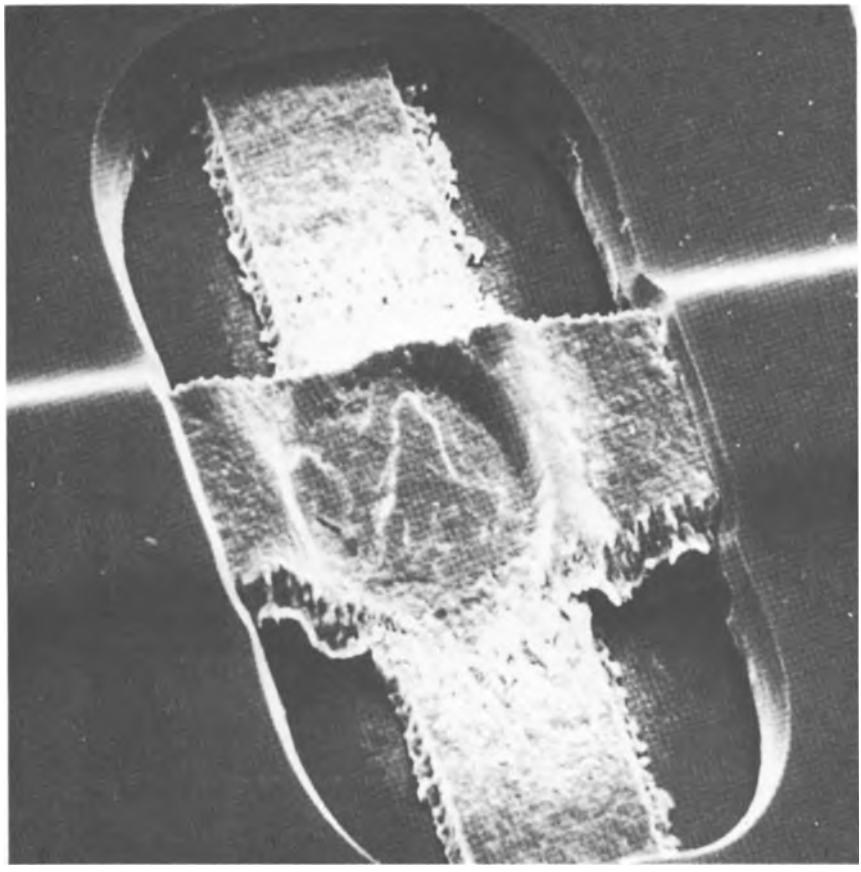

Figure 16-6 Ultrasonically welded connection between x- and y- signal lines. (Courtesy of Oerlikon-Contraves AG.)

16.3.2 Custom Silicon Substrates

As with custom semiconductor chips, custom silicon substrates provide wider latitude in module design. The silicon substrate can be used as an electrically active component of the module. In the case of AT&T, the substrate was heavily

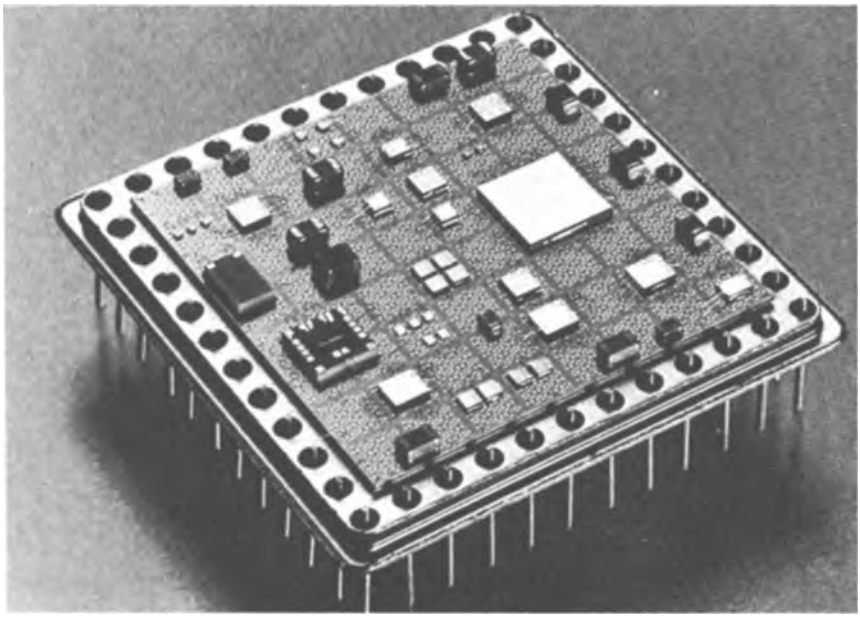

Figure 16-7 12-bit D/A converter. (Courtesy of Oerlikon-Contraves AG.)

doped to create a ground plane. The cross section of their MCM technology is illustrated in Figure 16-8 [35]. A thermally grown oxide and chemical vapor deposited silicon nitride dual dielectric layer forms a decoupling capacitor (25 nF/cm^2) between the power and ground planes. Two micrometer thick copper conductors were electroplated onto a sputter deposited Ti/Cu or Ti/Cr/Cu adhesion layer. Minimum conductor line widths are 10μm wide (line resistance approximately 10 Ω/cm). Characteristic impedances of 50 - 70 Ω for 10 μm wide/20 μm pitch signal lines are achieved with a bottom polyimide layer thickness of 10 μm and a second level polyimide thickness of 6 μm. Vias are filled with electroless Ni to insure electrical continuity through the vias and to allow via stacking. IC chips are bonded to the substrate using flip chip soldering. Figure 16-9 shows a three chip (WE 32100 central processor, memory management unit and math accelerator unit) module fabricated by AT&T. The silicon substrate is bonded to an aluminum heatsink using a thick silicone gel material to provide a compliant bond. The heatsink is assembled to a four layer bismaleimide triazine (BT) epoxy PWB. The PWB contains 160 I/O

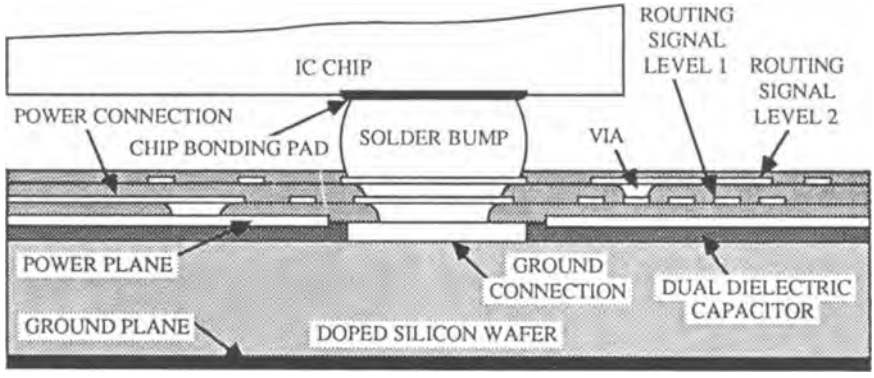

Figure 16-8 Cross section illustrating AT&T MCM technology. (Courtesy of AT&T.)

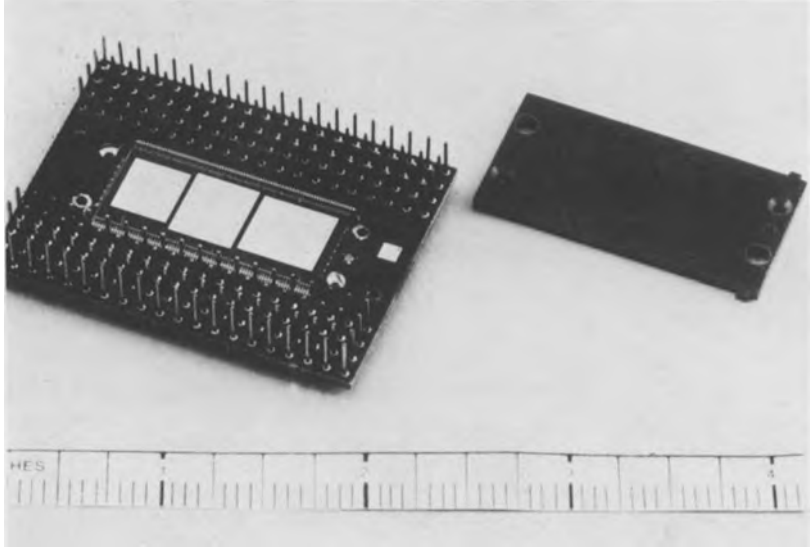

Figure 16-9 Three chip (WE 32100 central processor, memory management unit and math accelerator unit) module. (Courtesy of AT&T.)

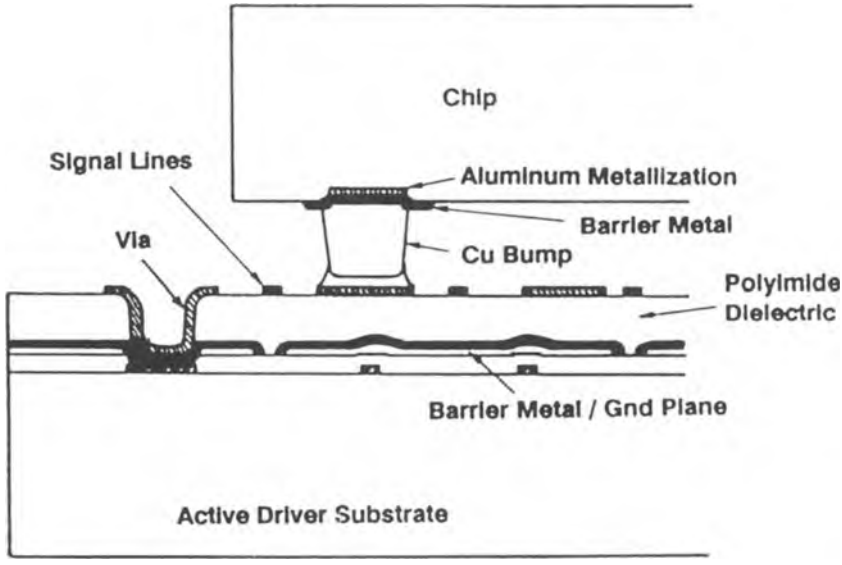

Figure 16-10 Cross section illustrating Texas Instruments' active silicon substrate module. (Courtesy of Texas Instruments.)

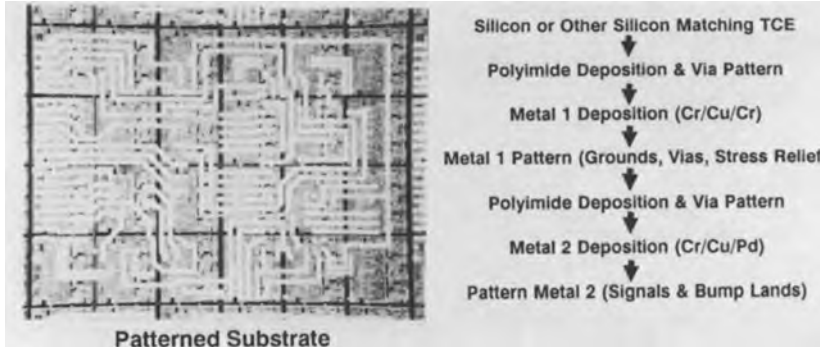

Figure 16-11 Active substrate containing DS3680 line drivers chips. (Courtesy of Texas Instruments.)

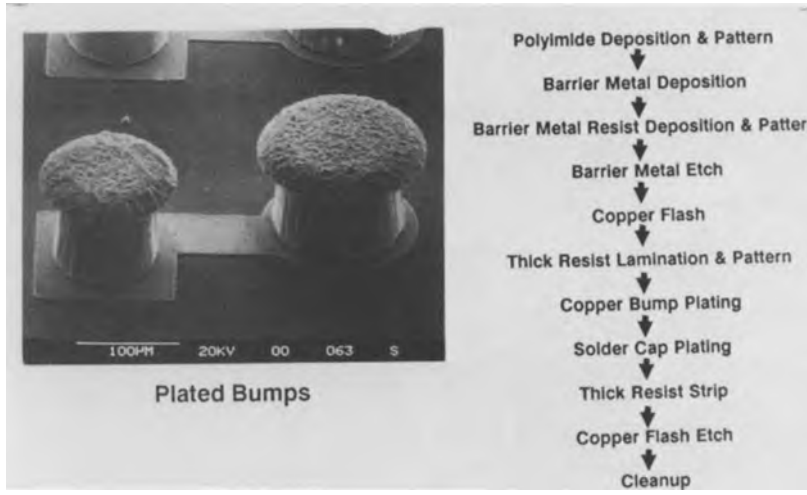

Figure 16-12 Solder capped copper bumps. (Courtesy of Texas Instruments.)

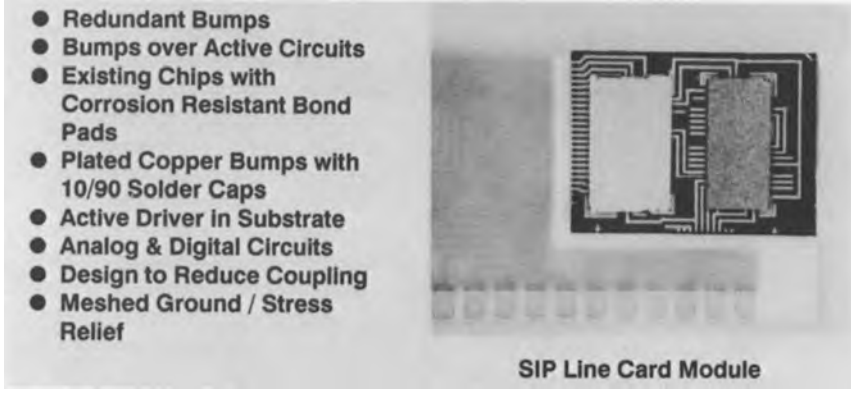

Figure 16-13 Telecom subscriber line card module. (Courtesy of Texas Instruments.)

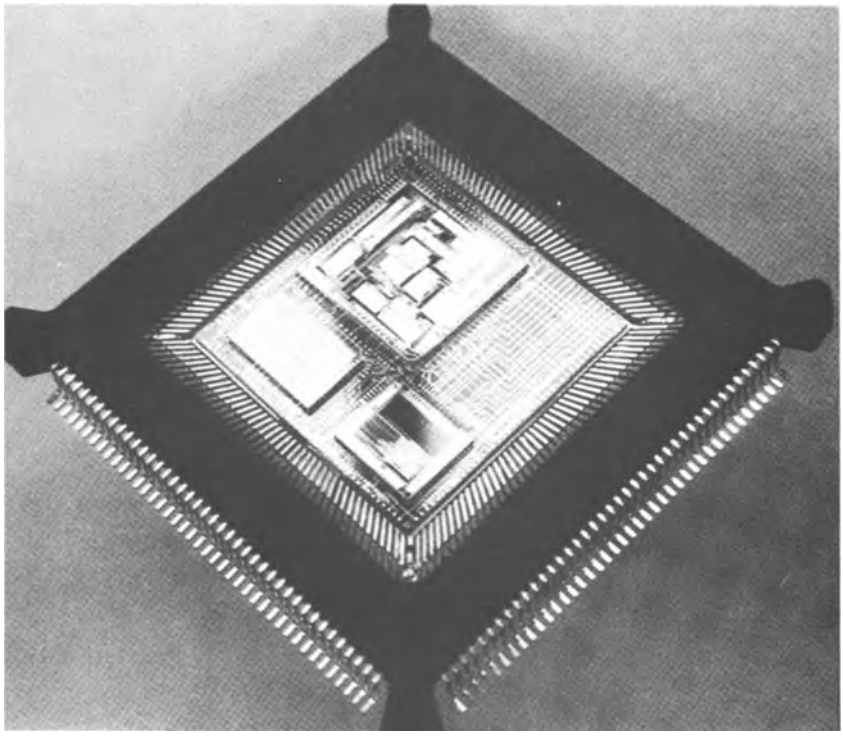

Figure 16-14 Prototype automotive MCM module. (Courtesy of Texas Instruments.)

pins in a grid array pattern. Thermosonic gold wire bonding is used to electrically connect the silicon substrate and the PWB. The BT resin system is chosen for the PWB to accommodate the thermosonic wire bonding temperature of 150°C. A plastic cover is attached to provide mechanical protection of the wire bonds.

Texas Instruments [36] has demonstrated a telecom subscriber line card module using an active silicon substrate. A cross section of the assembly is illustrated in Figure 16-10. The circuit uses three chips: a DS3680 line driver, a TCM 4209 SLCC (supervisor control, hybrid and attenuation) and a TCM 2914 COMBO (CODEC and filter). The DS3680 line drivers chips are fabricated monolithically in the MCM substrate as shown in Figure 16-11. The substrate interconnection pattern overlays several DS3680 chips providing redundancy to improve the active substrate yield. The interconnection metallization used is

Cr/Cu/Cr and the interlayer dielectric is polyimide. Palladium metallization is used on the substrate as the solderable pads for the flip chip assembly. Palladium is readily wet by solder, has limited solubility in solder, and is amenable to thin film processing. Chrome is not solderable and copper is chemically reactive and tends to oxidize and corrode prior to soldering. Rather than using solder bumps on the TCM 2914 COMBO and the TCM 4209 SLCC chips, copper bumps with solder caps are plated on the chips. Examples of the copper bumps are shown in Figure 16-12. Compared to solder bumps, the copper bumps provide a lower thermal resistance path from the die to the substrate. Two copper bumps are plated for each I/O on the die to improve thermal performance and also provide redundancy. An assembled module is shown in Figure 16-13.

A surface mount, plastic cavity package for MCMs developed by Texas Instruments is shown in Figure 16-14. The circuit is a prototype for an automotive application. The interconnection substrate has aluminum metallization and polyimide dielectric. The plastic package provides a low-cost packaging option.

Rockwell has designed a number of silicon MCMs [37]-[38]. The avionics processor module shown in Figure 16-15 contains 53 ICs and 40 discrete devices on a 2.2" × 2.2" silicon substrate with aluminum polyimide interconnect. The module is a flight computer system with 64 kbytes of static data memory, 64 kbytes of nonvolatile program storage, system timer, multiple interrupt inputs, discrete output ports and an advanced 16-bit microprocessor operating at 30 MHz. The module is packaged in a Kovar package (2.4" × 2.4") with 180 I/O.

The Rockwell module in Figure 16-16 performs interface and signal processing functions in a Global Positioning System (GPS) receiver. The 1.2" × 1.2" module operates at 88 MHz and uses meshed power and ground planes in a stripline configuration around two signal layers. The aluminum benzocyclobutene multilayer structure has 25 μm wide signal lines.

The Avionics Memory Module shown in Figure 16-17 is packaged in a hermetic 2.2" × 2.2" cofired aluminum nitride package. The module is for storage of flight and operating condition data and contains 32 ICs and 64 chip capacitors and resistors on two silicon substrates. The post in the center of the package is to support the metal lid during the pressure variations encountered in avionics applications. The package contains 136 fine stranded copper leads to provide for ultracompliant surface mount assembly to the next interconnect level. This is necessary to minimize cyclic stress on the surface mount solder joints due to the CTE mismatch between the AlN package and the PWB.

The module in Figure 16-18 contains two 1750A processors with on board memory and selftest fabricated on a 2.5" octagonal substrate. The aluminum

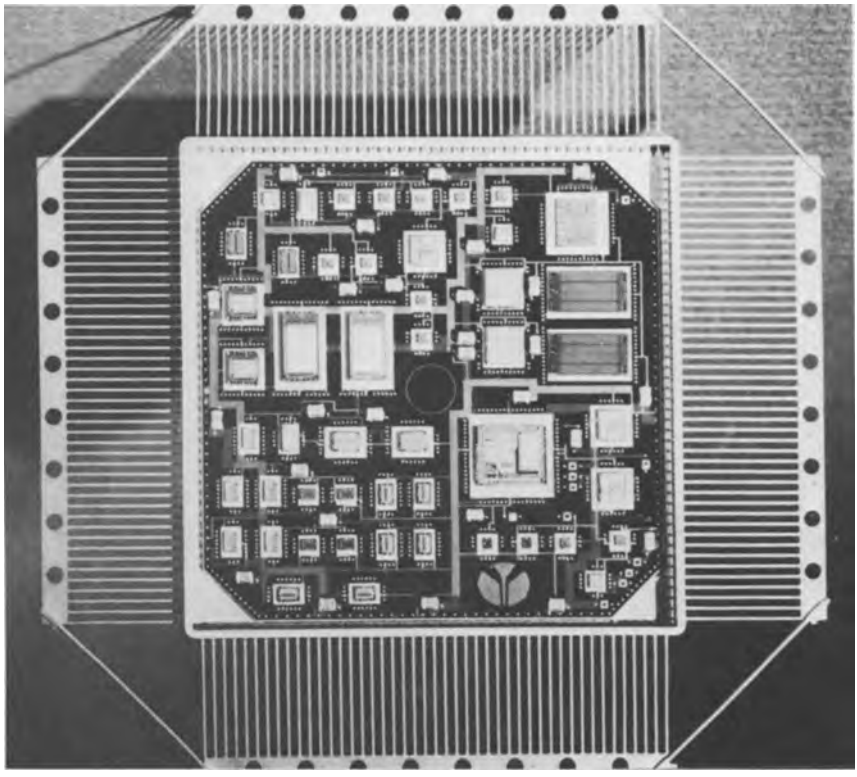

Figure 16-15 Avionics processor module. (Courtesy of Rockwell International).

benzocyclobutene substrate has four conductor layers and interconnects 134 components including 33 ICs. The module operates at 40 MHz and the combined throughput is 5 MIPs. Both hermetic (machineable aluminum nitride) and nonhermetic (graphite aluminum) packages containing an integral plug in connector have been fabricated for the module.

nChip's MCM technology utilizes thick silicon dioxide dielectric layers [39]-[41]. The dielectric is plasma enhanced chemical vapor deposited and undoped. The stress is controllable over the range $\pm 2 \times 10^8$ N/m^2. The three substrate variations illustrated in Figure 16-19 provide increasing performance (lower interconnect resistance, lower power/ground plane resistance, dual power planes, termination resistors and increased decoupling capacitance) with increasing substrate complexity (cost). Target applications for the nC-1000 Series (Figure

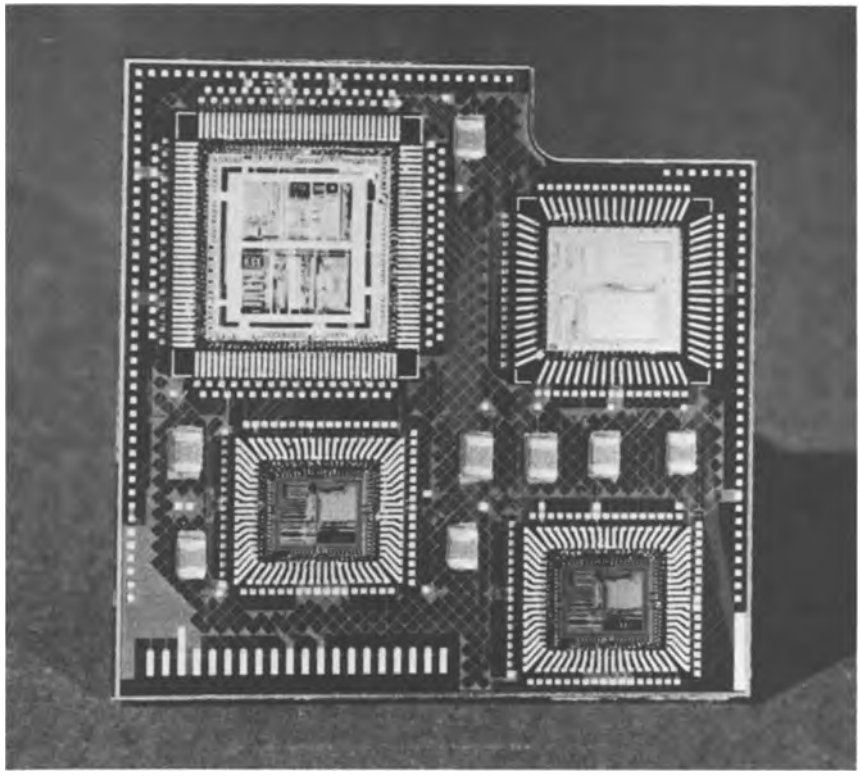

Figure 16-16 Global positioning system (GPS) receiver. (Courtesy of Rockwell International.)

16-19a) are 40 - 80 MHz CMOS circuits, for the nC-2000 (Figure 16-19b) are > 60 MHz CMOS/BiCMOS circuits, and for the nC-3000 Series (Figure 16-19c) are > 100 MHz GaAs and BiCMOS/ECL circuits. Figure 16-20 shows a five chip, SPARC Central Processor Module packaged in a 224 pin ceramic quad flat pack. The substrate size is 0.9" × 1.2" and replaces a 3.3" × 7.25" PWB.

Three different RISC processor modules fabricated by nChip are shown in Figure 16-21. The five chip module is the SPARC module previously discussed. The other two modules are packaged in pin grid arrays. The eight chip module is CMOS and implemented with nC-1000 substrate technology. The three chip module is ECL and uses nC-2000 substrate technology. This module dissipates 60 watts using conventional air cooling.

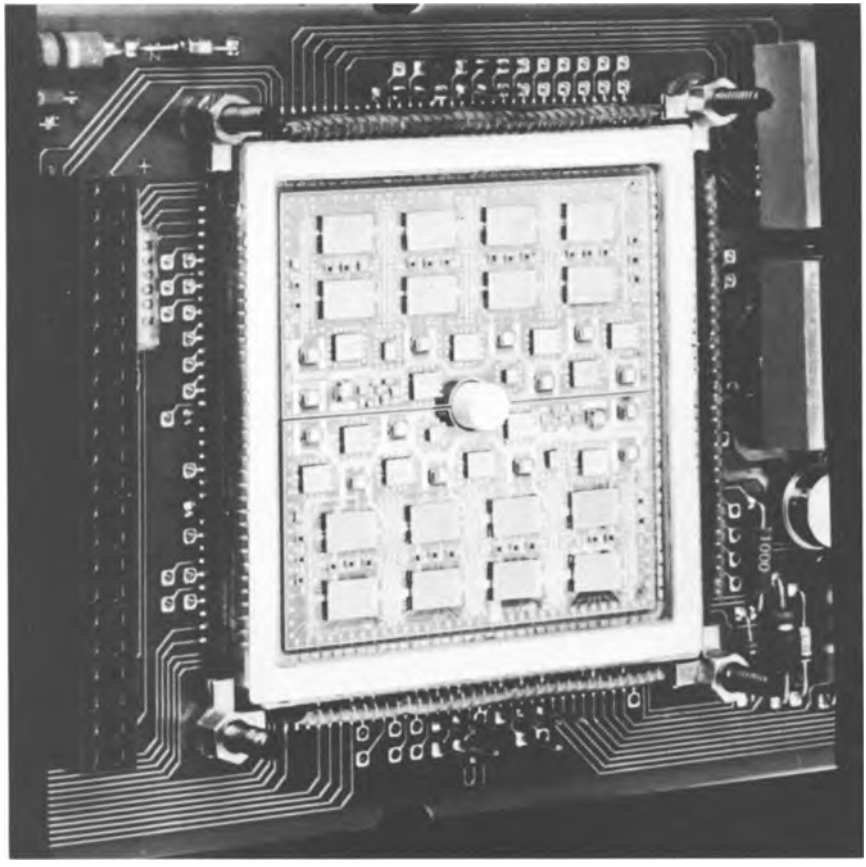

Figure 16-17 Avionics memory module. (Courtesy of Rockwell International.)

There continues to be significant development in the use of silicon as a substrate for MCMs and new modules are being announced on a regular basis. Examples of other silicon MCMs may be found in the literature.

16.4 SUMMARY

Continued advances in semiconductor technology, increases in system performance requirements, shorter design windows and short product life cycles will all contribute to the widespread application of MCM technology. Silicon

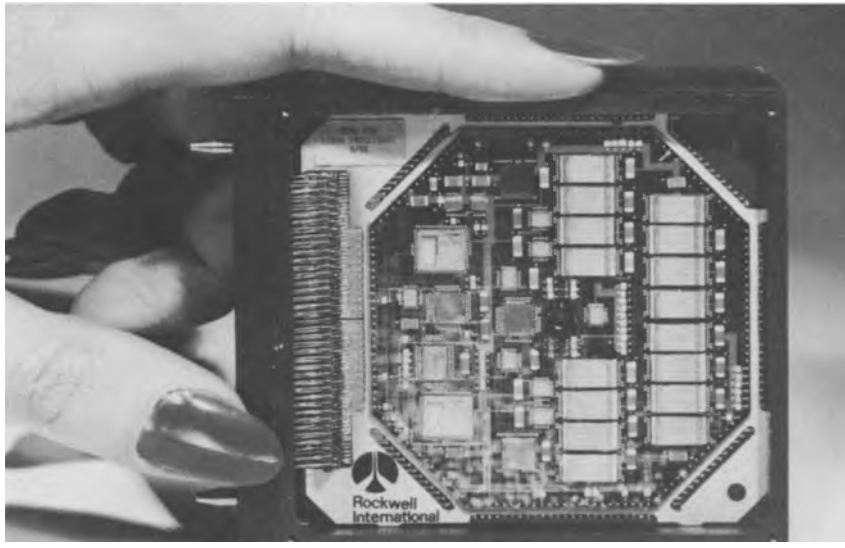

Figure 16-18 Dual 1750A Processors with on-board memory and self test. (Courtesy of Rockwell International.)

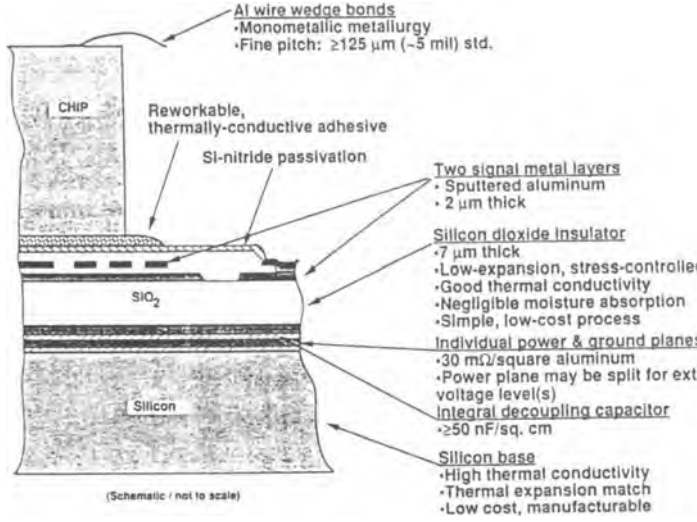

Figure 16-19a Cross section illustration of nC-1000 series silicon circuit board.

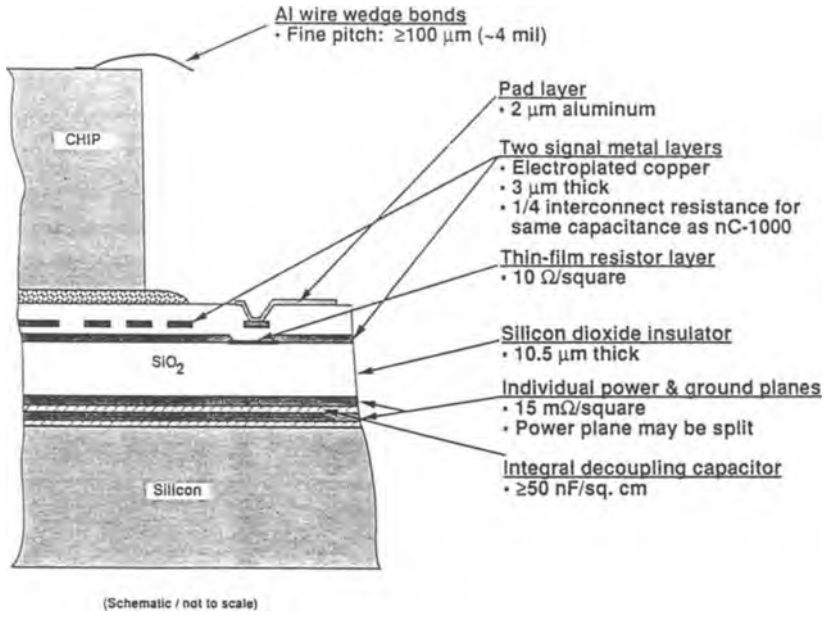

Figure 16-19b Cross section illustration of nC-2000 series silicon circuit board.

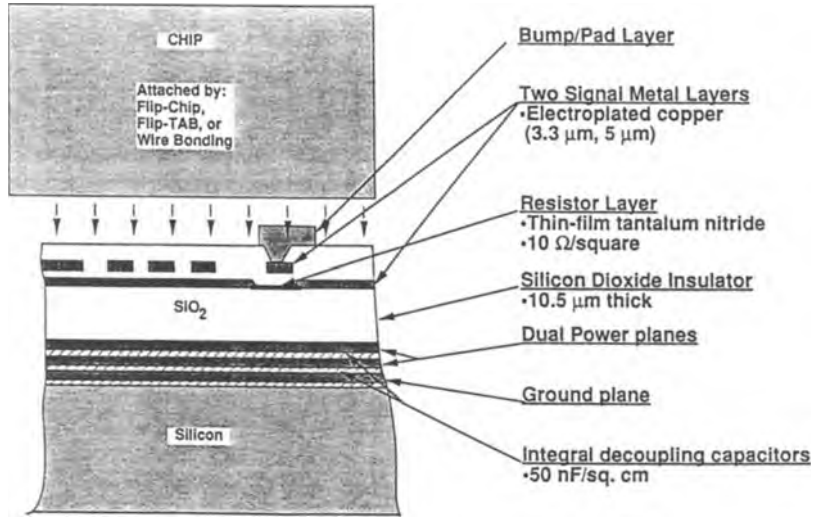

Figure 16-19c Cross section illustration of nC-3000 series silicon circuit board. (Courtesy of nChip.)

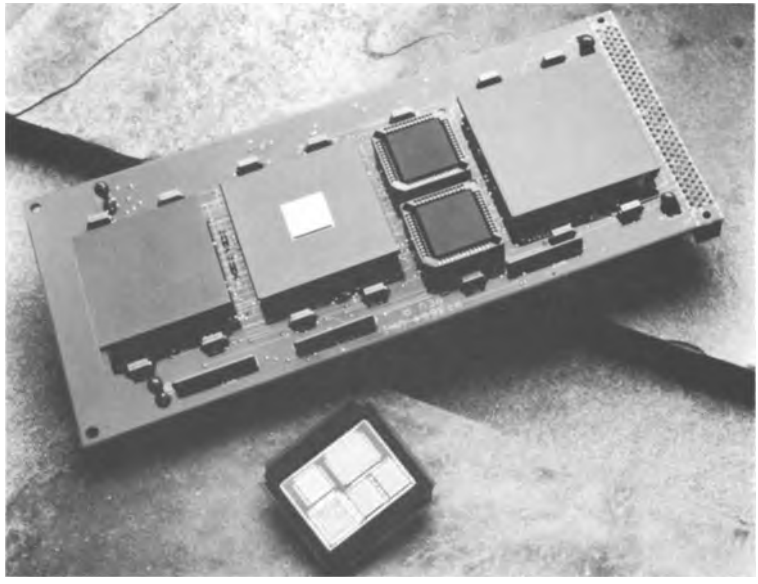

Figure 16-20a PWB and silicon MCM versions of the Ross Technology SPARC central processor chip set. (Courtesy of nChip.)

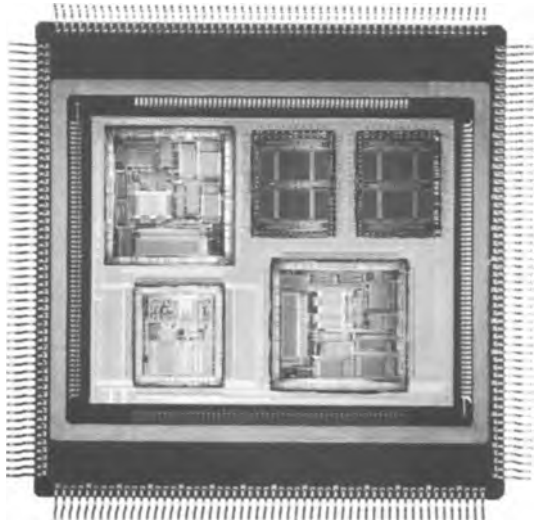

Figure 16-20b Five chip SPARC central processor module. (Courtesy of nChip.)

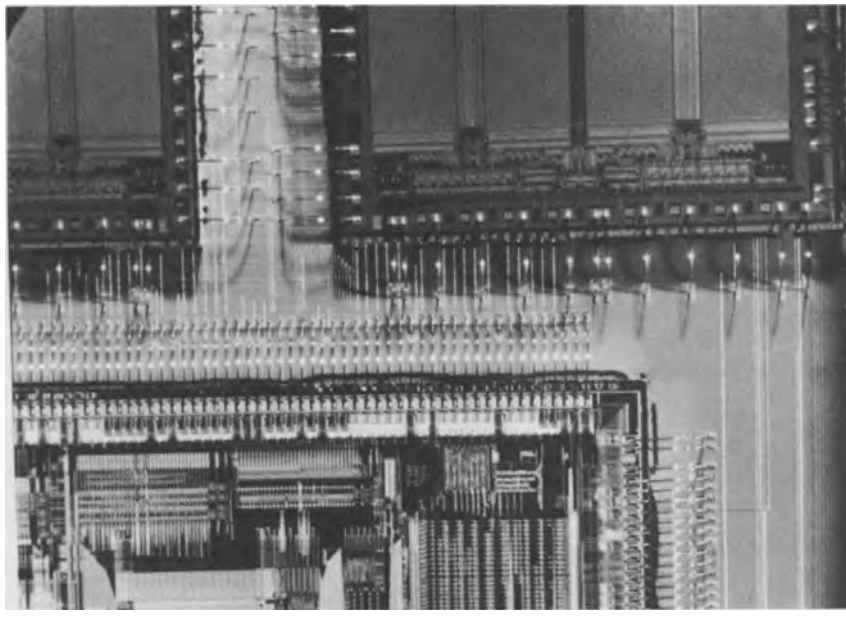

Figure 16-20c Close up of SPARC module illustrating the fine wire bond pitch. (Courtesy of nChip.)

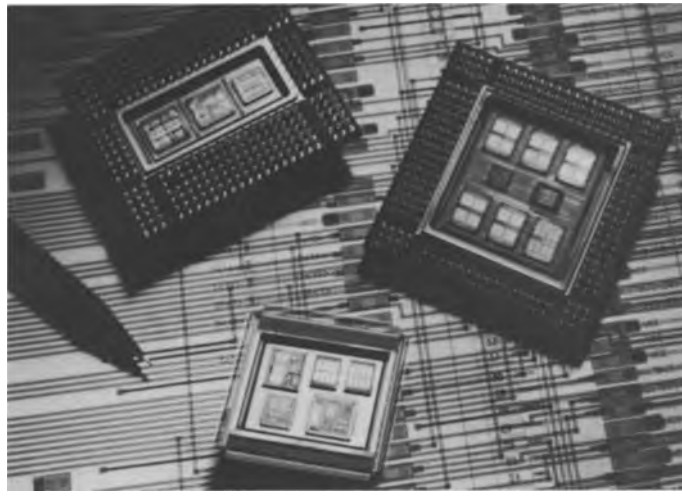

Figure 16-21 Three RISC processor modules. (Courtesy of nChip.)

provides a number of tradeoffs as a MCM substrate. It will not be universally used; silicon offers unique capabilities which can be exploited in a number of applications.

References

1 E. Ong, H. Chu, and S. Chen, "Metal Planarization with an Excimer Laser," *Solid State Technol.*, vol.34, no.8, pp. 63-68, Aug. 1991.

2 A. Nagy and J. Helbert, "Planarized Inorganic Interlevel Dielectric for Multilevel Metallization—Part I," *Solid State Technol.*, vol. 34, no. 1, pp. 53, Jan. 1991.
A. Nagy and J. Helbert, "Planarized Inorganic Interlevel Dielectric for Multilevel Metallization—Part II," *Solid State Technol.*, vol. 34. no. 3, pp. 77-80, March 1991.

3 M. L. Buschbom, "Modelling Issues for VLSI Package Design," *Proc. of the Techn. Prog.: NEPCON West*, (Anaheim CA), pp. 686-697, March 1989.

4 S. K. Tewksbury and L. A. Hornak, "Wafer Level System Integration: A Review," *IEEE Circuits and Devices Mag.*, vol. 5, no. 5, pp. 22-30, Sept. 1989.

5 D. J. Bodendorf, K. T. Olsen, J. F. Trinko, J. R. Winnard, "Active Silicon Chip Carrier," *IBM Technical Disclosure Bulletin*, vol. 15, no. 2, pp. 656-657, July 1972.

6 R. K. Spielberger *et al.*, "Silicon-on-Silicon Packaging," *IEEE Trans. on Components, Hybrids, and Manufacturing Techn.*, vol. CHMT-7, no. 2, pp. 193-196, June 1984.

7 T. Yamada *et al.*, "Low Stress Design of Flip Chip Technology for Si on Si Multichip Modules," *Proc. of the Fifth Annual Internat. Elect. Packaging Conf.*, (Orlando FL), pp. 551-557, Oct. 1985.

8 R. L. Hubbard and G. Lehman-Lamer, "Very High Speed Multilayer Interconnect Using Photosensitive Polyimides," *Proc of the 1988 Internat. Symp. on Microelect.*, (Seattle WA), pp. 374-376, Oct. 1988.

9 S. E. Greer, "Low Expansivity Organic Substrate for Flip-Chip Bonding," *Proc. of the 28th Elect. Components Conf.*, (Anaheim CA), pp. 166-171, April 1978.

10 R. R. Tummala, H. R. Potts, and S. Ahmed, "Packaging Technology for IBM's Latest Mainframe Computers (S/390/ES9000)," *Proc. of the 41st Elect. Components and Techn. Conf.*, (Atlanta GA), pp. 682-687, May 1991.

11 T. R. Gheewala, "Packages for Ultra-High Speed GaAs IC's," *Technical Digest of the 1984 GaAs IC Symp.*, (Boston MA), pp. 67-70, Oct. 1984.

12 S. Sriram, "Novel V-Groove Structures in Silicon," *SPIE Vol. 578, Integrated Optical Circuit Engineering II*, pp. 88-94, 1985.

13 R. A. Soref and J. P. Lorenzo, "Silicon Guided-Wave Optics," *Solid State Techn.*, vol. 31, no.11, pp. 895-898, Nov. 1988.

14 J. T. Boyd *et al.*, "Guided Wave Optical Structures Utilizing Silicon," *Optical Engineering*, vol. 24, no. 2, pp. 230-234, March/April 1985.

15 S. E. Schacham *et al.*, "Waveguides as Interconnections for High Performance Packaging," *Proc. of the 9th Annual Internat. Elect. Packaging Conf.*, (San Diego CA), pp. 1003-1013, Sept. 1989.

16 R. Selvaraj, H. T. Lin, and J. F. McDonald, "Integrated Optical Waveguides in Polyimide for Wafer Scale Integration," *J. of Lightwave Techn.*, vol. 6, pp. 1034-1044, 1988.

17 J. B. Angell, S. C. Terry, and P. W. Barth, "Silicon Micromechanical Devices," *Scientific Amer.*, vol. 248, no. 4, pp. 44-55, April 1983.

18 K. E. Bean, "Anisotropic Etching of Silicon," *IEEE Trans. on Electron Devices*, vol. ED-25, no. 10, pp. 1185-1193, Oct. 1978.

19 J. H. Atherton, "Sensor Signal Conditioning: An IC Designer's Perspective," *SENSORS*, vol. 8, no.11, pp. 23-30, Nov. 1991.

20 A. Goyal *et al.*, "Re-entrant Cavity Heat Sink Fabricated by Anisotropic Etching and Silicon Direct Wafer Bonding," *Proc. of the Eighth Annual Semi. Thermal Measurement and Management Symp.*, (Austin TX), pp. 25-29, Feb. 1992.

21 N. K. Phadke *et al.*, "Re-entrant Cavity Surface Enhancements for Immersion Cooling of Silicon Multichip Packages," *Proc. of the Third IEEE/ASME Intersociety Conf. on Thermal Phenomena in Elect. Sys.*, (Austin TX), pp. 59-65, Feb. 1992.

22 M. J. Little *et al.*, "A Three-Dimensional Computer for Image and Signal Processing," *Proc. of the IEEE 1985 Custom Integrated Circuits Conf.*, (Portland OR), pp. 119-123, May 1985.

23 M. W. Hartnett and E. J. Vardaman, "Worldwide MCM Status and Trends: Material Choices," *Proc. of the Technical Program: NEPCON West*, (Anaheim CA), pp. 1111-1120, Feb. 1991.

24 A. Schiltz *et al.*, "Lift Off Techniques of Packaging/ Processing Technologies for Multichip Modules," *Proc. of the Technical Prog.: NEPCON West*, (Anaheim CA), pp. 975-983, March 1990.

25 H. Kroger *et al.*, "Applications of Superconductors to Hybrid Electronics," *Proc. of the 1990 Internat. Symp. on Microelect.*, (Chicago IL), pp. 570-572, Oct. 1990.

26 L. Verdet, J. Reche, and G. Rabilloud, "The P's and Q's of PPQ," *Elect. Packaging and Prod.*, vol. 31, no. 1, pp. 58-61, Jan. 1991.

27 T. G. Tessier, G. M. Adema, and I. Turlik, "Polymer Dielectric Options for Thin Film Packaging Applications," *Proc. of the 39th Elect. Components Conf.*, (Houston TX), pp. 127-134, May 1989.

28 B. T. Merriman *et al.*, "New Low Coefficient of thermal Expansion Polyimide for Inorganic Substrates," *Proc. of the 39th Elect. Components Conf.*, (Houston TX), pp. 155-159, May 1989.

29 D. Burdeaux *et al.*, "Benzocyclobutene (BCB) Dielectrics for the Fabrication of High Density, Thin Film Multichip Modules," *J. of Elect. Materials*, vol. 19, no. 12, pp. 1357-1366, Dec. 1990.

30 T. Ohsaki *et al.*, "A Fine-Line Multilayer Substrate with Photo-Sensitive Polyimide Dielectric and Electroless Copper Plated Conductors," *Proc. of the IEEE Internat. Elect. Manufac. Techn. Symp.*, (Anaheim CA), pp. 178-183, 1987.

31 A. A. Bogdan, "An Electrically Programmable Silicon Circuit Board," *Proc. of BUSCON '87*, pp. 9-14, April 1987.

32 K. L. Drake *et al.*, "Equivalent Hermetic Packaging of High Density Multichip Modules Interconnect Substrate," *Proc. of the Technical Prog.: NEPCON West*, (Anaheim CA), p. 840, Feb. 1991.

33 E. Sutcliffe, "Multichip Modules for High Performance," *Proc. of the Technical Prog.: NEPCON West*, (Anaheim CA), pp. 824-831, Feb. 1991.

34 D. Carey, "A Program to Provide Quick Turnaround Delivery of Multichip

Modules," *Proc. of the Technical Prog.: NEPCON West*, (Anaheim CA), pp. 507-514, Feb. 1991.

35 C. J. Bartlett, J. M. Segelken, and N. A Teneketges, "Multichip Packaging Design for VLSI-Based Systems," *IEEE Trans. on Components, Hybrids, and Manufac. Techn.*, vol. CHMT-12, no. 4, pp. 647-653, Dec. 1987.

36 K. G. Heinen *et al.*, "Multichip Assembly with Flipped Integrated Circuits," *Proc. of the 39th Elect. Components Conf.*, (Houston TX), pp. 672-680, May 1989.

37 J. K. Hagge, "Ultra Reliable HWSI with Aluminum Nitride Packaging," *Proc. of the Technical Prog.: NEPCON West*, (Anaheim CA), pp. 1271-1283, Feb. 1989.

38 J. K. Hagge, "State-of-the-Art Multichip Modules for Avionics," *Internat. Symp. on Advances in Interconnects and Packaging*, (Boston MA), SPIE/OPTCON-90, pp. 175-194, Nov. 1990.

J. K. Hagge, "Ultra-Reliable Pachaging for Silicon-on-Silicon WSI," *Proc. 38th Electr. Comp. Conf.*, (Los Angeles CA), pp. 282-292, May 1988.

39 T. Horton, "MCM Driving Forces, Applications, and Future Directions," *Proc. of the Technical Prog.: NEPCON West*, (Anaheim CA), pp. 487-494, Feb. 1991.

40 B. McWilliams, "Comparison of Multichip Interconnection Technologies," *Proc. of the 11th Annual Internat. Elect. Packaging Conf.*, (San Diego CA), pp. 63-68, Aug. 1991.

41 J. C. Demmin, "Thermal Modeling of Multi-Chip Modules," *Proc. of the Technical Prog.: NEPCON West*, (Anaheim CA), pp. 1145-1154, Feb. 1991.

17

THE TECHNOLOGY AND MANUFACTURE OF THE VAX-9000 MULTICHIP UNIT

Salma Y. Abbasi, Arun Malhotra, John D. Marshall,
Douglas N. Modlin, Ralph Platz, Hongbee Teoh,
Scott Westbrook, Jimmy Jun-Min Yang

17.1 INTRODUCTION

A multichip module (MCM) is a miniaturized electronic subsystem based on a high density interconnect (HDI) structure. The MCM connects unpackaged components (ICs and/or passive devices) and integrates them with an interconnection system.

Corresponding to the interconnection technology that it employs, the industry has defined MCM-D as having a thin film interconnection structure deposited on a suitable substrate (aluminum, silicon, ceramic). The MCM-D provides the highest density and performance but requires the largest investment in manufacturing technology. This chapter discusses Digital's approach to the development of MCM-D technology and manufacturing processes.

The information presented in Section 17.2 provides a historical perspective on the motivation for the development of MCM technology. Section 17.3 describes development of the Digital's VAX-9000 MCM-D technology as embodied in the multichip unit (MCU). The up front system performance goals are presented, along with the thought process that resulted in the choice of a packaging solution based on MCM-D technology. The VAX-9000 MCU engineering strategy also is described. In addition, the quality strategy used to obtain the very high level of reliability required for a mainframe class machine is presented.

In Section 17.4, the MCU technology is described in detail. Section 17.5 provides a description of the VAX-9000 MCU manufacturing process. Finally, Section 17.6 summarizes the applicability of Digital's MCM-D technology to the development of future products. Examples of applications are presented, along with a strategy for achieving high quality and reliability in the cost sensitive environment in which the industry now finds itself.

17.2 MCM SYSTEM PERSPECTIVE

In order for system designs to reach increasingly higher performance capabilities, designers must continue to develop more effective methods of integrating larger numbers of logic gates and driving them at increasingly faster clock speeds. As clock speeds increase, the intergate signal propagation delays must correspondingly decrease, putting more stringent demands on the design of the physical package and on the physical gate interconnect structures. Shortening the interconnect distances to the minimum by packing the logic into the closest possible physical proximity, optimizing the capacitive, resistive and inductive characteristics of the interconnects, distributing cleanly decoupled power to the circuits and removing the heat are all essential to achieving high clocking speeds. To meet these requirements, two approaches can be used concurrently by the designer: packing the gates densely into ICs and packing the ICs densely onto MCMs.

Integrated circuit technology has been and remains the first order method for achieving dense logic and high performance because constant improvements in chip technology continue to enable higher gate counts, higher gate densities and higher clock frequencies to be realized. However, as the experience of wafer scale integration demonstrates, when the demand for physical density increases by very large factors, numerous problems become exceedingly difficult to resolve. These problems include managing higher thermal dissipation, implementing the denser and more complex physical interconnect routes, managing clock skews, ensuring signal integrity, providing all the required I/O and power connections and keeping the die and silicon feature sizes within the boundaries where cost effective manufacturing yields are achievable. In addition, some of the performance gain achieved in very large scale ICs is offset by the relatively high inductance and capacitance present in high pin count packages.

To address these issues, designers are turning to MCM technology since it provides alternative methods of achieving performance and density. With MCM technology the system designer can minimize the number of distinct packages and thereby reduce significantly the distances between the ICs. Figure 17-1 illustrates the dramatic size reduction benefits of MCM technology. The

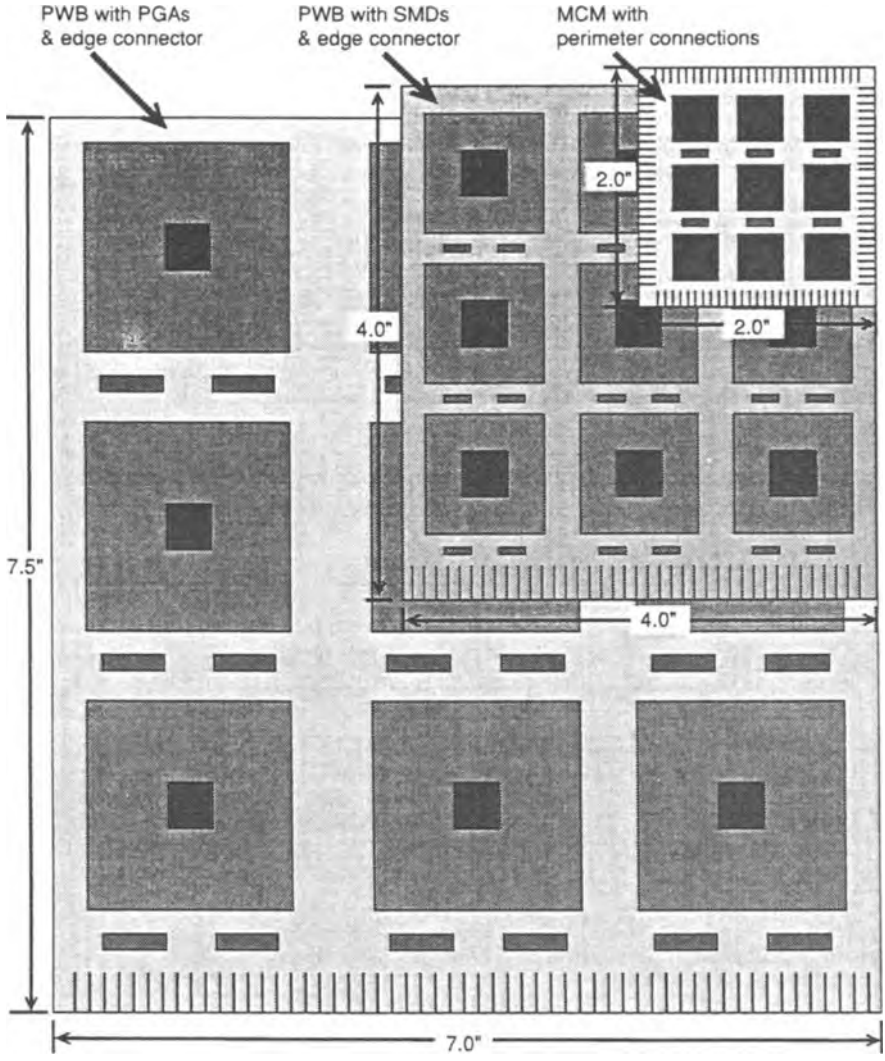

Figure 17-1 Scale perspective of one application implemented in standard through-hole PWB technology, in SMT and in MCM technology.

resultant reduction in interconnect distances, interconnect capacitance and interconnect inductance can reduce signal propagation delays by as much as 300%. Furthermore, MCM technology gives the designer additional options for

partitioning the logic in ways that improve its performance, cost or manufacturability.

17.3 VAX-9000 MCU DEVELOPMENT

17.3.1 VAX-9000 Performance Goals

In 1983, the designers of the VAX-9000 set out to design a top of the line high performance mainframe system. The performance goal for the VAX-9000 central processor unit (CPU) was set at 30 VAX units of performance (VPUs). This was to be achieved not only by implementing an innovative CPU architecture that would overlap the execution of CPU instructions and minimize the number of CPU cycles per instruction, but also by deploying a physical technology that would allow the design to be implemented with the fastest available logic gates driven at the highest possible clock rate.

Although system performance was the driving force behind the VAX-9000 MCU technology, aggressive goals were set also for system uptime. High availability was achieved through improvements in the physical machine design as well as through the inherently higher reliability of an MCM package. Specifically, since an entire level of interconnect is eliminated by directly connecting the integrated circuits to the MCM substrate, the projected MCM failure rate is lower than the failure rate projected for the same integrated circuits in conventional packages on a printed wiring board (PWB).

17.3.2 Analysis of Alternative Physical Technologies

To meet such demanding performance requirements, a number of different physical technologies were considered in the beginning of the development program (1983-1984). The evaluation required not only that the technical capabilities of the alternatives be compared, but also that projections be made regarding the feasibility of technically viable options being made available for manufacturing production in the time period of interest (early 1990s).

Against a reference baseline of the VAX-8000 series CPUs (which utilized PGA packaged Motorola MCA-I emitter coupled logic (ECL) gate arrays on large PWBs), the following options were projected and evaluated:

- Chilled CMOS.
- Motorola ECL Mosaic-III gate arrays (MCA-III) PGA packaged on PWBs.
- Motorola ECL Mosaic-III gate arrays (MCA-III) packaged on MCMs.

- GaAs gate arrays packaged on MCMs.
- Wafer scale integration.

To understand the performance differences between these options, a critical path analysis was undertaken for a typical logic path. For each of the several technologies, the signal propagation delay was evaluated through a path that included the arithmetic logic unit (ALU), the translation buffer and the cache. The normalized results of that analysis are illustrated in Figure 17-2.

The results of the evaluation projected the following:

1. Chilled CMOS in 1989 would achieve the same on-chip logic speeds as the earlier ECL logic.

2. Because of the increased level of integration, chilled CMOS would have significantly shorter off-chip delays.

3. The on-chip speeds for Mosaic-III 10,000 ECL gate arrays would be approximately twice as fast as for CMOS.

4. If packaged in PGAs, the performance of the Mosaic-III ECL gate arrays would be dominated by their off-chip delays.

5. The off-chip delays of the Mosaic-III circuits could be reduced by 50% by changing from the PGA package to uncased TAB packaging.

The estimated performance gains from GaAs and wafer scale integration did not appear sufficient to warrant their increased development and time to market risks.

17.3.3 VAX-9000 MCU Strategy

Based on the above analysis, a strategy was established for partitioning the VAX-9000 CPU (greater than 1,000,000 equivalent logic gates and greater than 3,200,000 bits of high speed local storage) into a set of 20 high speed modules which could be implemented with high power (up to 30 watts each), TAB packaged, MCA-III ECL 10,000 gate arrays driven at a clock cycle time of 16 nanoseconds [1]-[3]. To meet these performance demands the set of gate arrays for each module would be interconnected, powered and cooled through use of the proprietary MCM-D technology described in Section 17.4. The modular MCM-D units based on this design were called MCUs. The MCU is comprised of signal and power connectors, integrated circuits and the MCM interconnect is known as the high density signal carrier (HDSC).

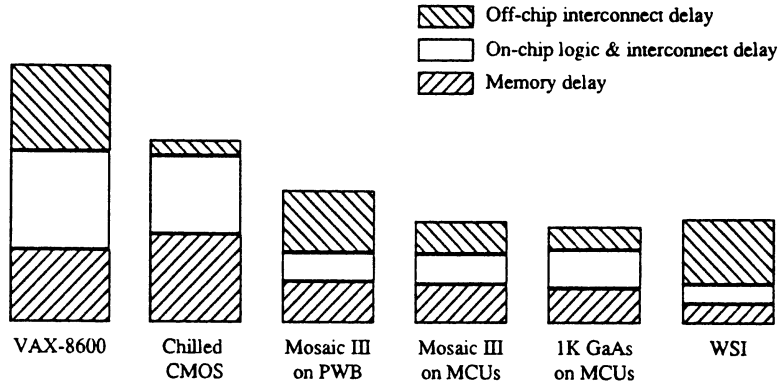

Figure 17-2 Normalized delay of a critical path in an ALU, translation buffer and cache for several technologies.

17.3.4 MCU Manufacturability Goals

To ensure the manufacturability and field serviceability of the VAX-9000 CPU, the MCU had to be designed for testability and diagnosibility. In manufacturing, electrical testing had to be executed progressively through the process at the earliest points of value added. Testability structures and mechanisms had to be established at the IC level, at the interconnect level and at the MCU module level. In the field, the test structures on each MCU had to be directly accessible from the VAX-9000 system console.

To manufacture MCUs in a cost competitive manner, high yielding base process technologies had to be developed and established. These processes had to support the production of 20 distinct MCU design options. In order to establish reliable and cost efficient processes and to meet time-to-market requirements, the processes were designed to use conventional semiconductor and PWB equipment wherever practical. The development of custom equipment was restricted to a selected set of processes (such as HDSC fabrication fixtures, HDSC test, IC TAB excise and lead form, die placement and lead bonding, MCU test) where the payback was high or where there was no other alternative.

17.3.5 Concurrent MCU Engineering and Manufacturing

Because of the long lead time associated with the scale up of an MCM-D process, it was necessary to begin operating the MCU manufacturing line long

in advance of the time when the ICs would be available. The concurrent nature of the VAX-9000 engineering and manufacturing efforts necessitated frequent and in depth communication between the respective engineering and manufacturing teams [4]. This was particularly challenging since the engineering group was located in Marlborough MA, and the manufacturing site was located in Cupertino CA. To minimize the impact of distance between the engineering and manufacturing groups, a general purpose networked database system was established in order to allow rapid transfer of design and manufacturing data.

Initial production start up was accomplished by running test vehicles in a small scale process development line. The process was then scaled up in a pilot line and finally transferred to volume production.

17.3.6 MCU Quality Engineering

At the start of the development program, specific MCU and HDSC reliability goals were established to drive the design:

- MCU dead on arrival (DOA) failures - 300 PPM
- MCU mean time between failures (MTBF) - 700,000 hours
- HDSC mean time between failures (MTBF) - 10,000,000 hours

A strong design for reliability (DFR) philosophy then guided the subsequent development of both the HDSC and MCU. The DFR focused specifically on defining the detailed reliability requirements for the design, the material and the manufacturing processes and on establishing methods to verify that the requirements were met. The qualification program encompassed not only internal development but also suppliers' contributions toward product reliability. Working jointly with the suppliers, clear definitions of requirements and responsibilities were negotiated in advance, auditing methods were established and the suppliers' components, subassemblies and subsystems were integrated into the overall verification and qualification processes.

For the HDSC and MCU, the reliability characterizations evaluated the critical areas:

- Process materials
- Contamination
- Solderability
- Die attach integrity
- Encapsulant integrity
- Outer lead bond (OLB) and flex lead bond (FLB) solder joint integrity
- Corrosion

The test and verification methodology was based on applicable military and industry standards:

- Thermal shock, -35°C - 125°C
- Temperature cycling, -65°C - 150°C
- Biased temperature/humidity, 85°C / 85% RH
- High temperature aging, 125°C
- Corrosive gas tests

Preliminary test results showed that some of the early, marginal MCU and HDSC engineering prototype samples had unexpected failures in PTH integrity and some corrosion was observed on these samples after biased 85/85 tests. Appropriate process changes were then implemented and as a result, no subsequent failures were observed.

The formal MCU/HDSC Design Verification Test (DVT) process was then used to confirm the integrity of the design and to verify that the product met its functional requirements and the requirements imposed by its manufacturing processes, its storage and transport environments and its operating conditions.

Ongoing reliability testing (ORT) was established to ensure that the HDSCs and MCUs produced in manufacturing continued to meet their specifications for infancy and steady state reliability. ORT sample units were randomly drawn from production lots and were stressed thermally and environmentally. The ORT results have reconfirmed the robustness of the HDSC substrate and have demonstrated the high yield and quality of the HDSC manufacturing process.

Finally, field return data is continuously analyzed in order to ensure an accurate empirical understanding of product reliability performance, of failure modes and of failure mechanisms so that corrective actions can be promptly implemented when required. Based on the population of HDSCs in the field over a period of 18 months (since the VAX-9000 new product introduction), an extensive set of data from system installations, customer system run time clocking and field returns has been tracked and analyzed. This data shows only four confirmed HDSC failures to date, with an accumulated total system run time of 37,769,784 hours. This translates to an impressive MTBF of 9,442,446 hours for the HDSC at the point in time at which this is written and this is 94% toward full demonstration of its design goal!

17.4 VAX-9000 MCU TECHNOLOGY

17.4.1 MCU Overview

The MCU serves several functions. It allows for fine pitch bonding so that high density chip placement can be achieved. The characteristic impedance of the

Table 17-1 Summary of MCU Specifications.

Maximum power dissipation	270 watts (air cooled)
Maximum IC junction temperature	85°C @ 25°C room temperature
Maximum number of VLSI chips	72
Minimum chip lead pitch	200 μm
Size: in plane	14.2 cm x 13.21 cm
height	5.44 cm
Minimum pitch on planar module	14.38 cm x 13.46 cm
Weight	1.59 kg (with heatsink)
Clock input frequency	320 MHz - 580 MHz
Signal I/O per MCU	800
Signal rise time	600 ps
Voltage levels	2 plus ground
Maximum current	40 A per voltage level

signal paths between the chips and off the MCU is controlled to within ±5% of the nominal impedance. At the same time, crosstalk is held to less than 5% (see Table 17-3). High pin out connectors increase the system architecture options. The MCU can provide up to 270 W of power with less than 20 mV of resistive voltage drop (IR) between the power supplies and ground. At the same time, the thermal path from semiconductor junction to heatsink is optimized so that no more than a 20°C thermal drop occurs. It is for these reasons that the MCU technology was chosen for packaging the VAX-9000 ECL circuits.

A summary of MCU specifications is given in Table 17-1. To the external world, this device presents four signal connectors capable of connecting to 201 signals each (804 total) and two power connectors capable of distributing two supplies at 40 amps each. Internally, the package can accommodate up to 72 ICs (a maximum of 45 in the actual designs) connected together by a special interconnect known as an HDSC. Variations of the MCU may contain different numbers and designs of ICs and different HDSC designs.

The basic elements of the MCU are shown in Figure 17-3. The lid serves primarily to prevent large objects from inappropriately contacting the HDSC and chip areas. The signal connectors and power connectors are mounted on the housing assembly which supports the HDSC. The pin fin heatsink is bolted to the back of the HDSC.

IC placement on the HDSC is defined by a 3 x 3 matrix of chip placement sites. The middle site is always used by a special clock distribution chip. Any of the other eight sites can be occupied either by one large logic circuit (MCA-III) or by up to nine smaller special RAM circuits.

The VAX-9000 system uses 20 different MCU designs, each utilizing a unique HDSC design. The 20 different MCU options use a total of 79 different IC designs. The common design base for all the MCU options allows for economy in the manufacturing of the device and enables the same manufacturing process to be used for all variations.

The VAX-9000 MCU solution is unique when compared to other solutions for this class of problems such as the IBM multilayer ceramic thermal conduction module (TCM) with solder bumped flipped ICs. There were several differentiating reasons. The investment and effort to follow a similar approach was considered prohibitive. Face up chip attach approaches were considered easier to develop and easier to test, debug and cool. The desire to have the MCU as a field replaceable unit (FRU) drove the power and signal connector to be separable. The signal connector is folded back above the MCU to minimize footprint on the planar board and allow cooling on the opposite side. The separate signal and power cores laminated together was seen to minimize the topography of a single sequential core and the drilled holes replaced the extra via layers of a single sequential core. In summary, the MCU was the point solution to a long series of intersecting constraints.

17.4.2 MCU Integrated Circuits

The MCU uses exclusively ECL based integrated circuits. All of these circuits operate from a -5.2 volt supply, while most also use a -3.4 volt supply. The -3.4 volt supply is used primarily to limit the power dissipation of the off-chip output drivers. There are three basic classes of ECL circuits used: gate array and custom logic circuits, a clock distribution circuit and two types of special RAM devices called "self timed random access memory" (STRAMs). The characteristics of these ICs are summarized in Table 17-2 and Table 17-3.

Gate Array and Custom Logic Circuits

The logic circuit chips all use the Motorola Macrocell Array III (MCA-III) technology. The gate array master slice has 10,000 available gates. The custom circuits were designed for logic functions that could not achieve high efficiency in a gate array implementation. These ICs feature less than 600 ps output rise times.

STRAM

Many of the MCU types require high speed local storage, such as caches, registers and control store. This is provided by 4K × 1 and 4K × 4 STRAMs. The STRAMs are proprietary designs with address latches and internal write pulse generation.

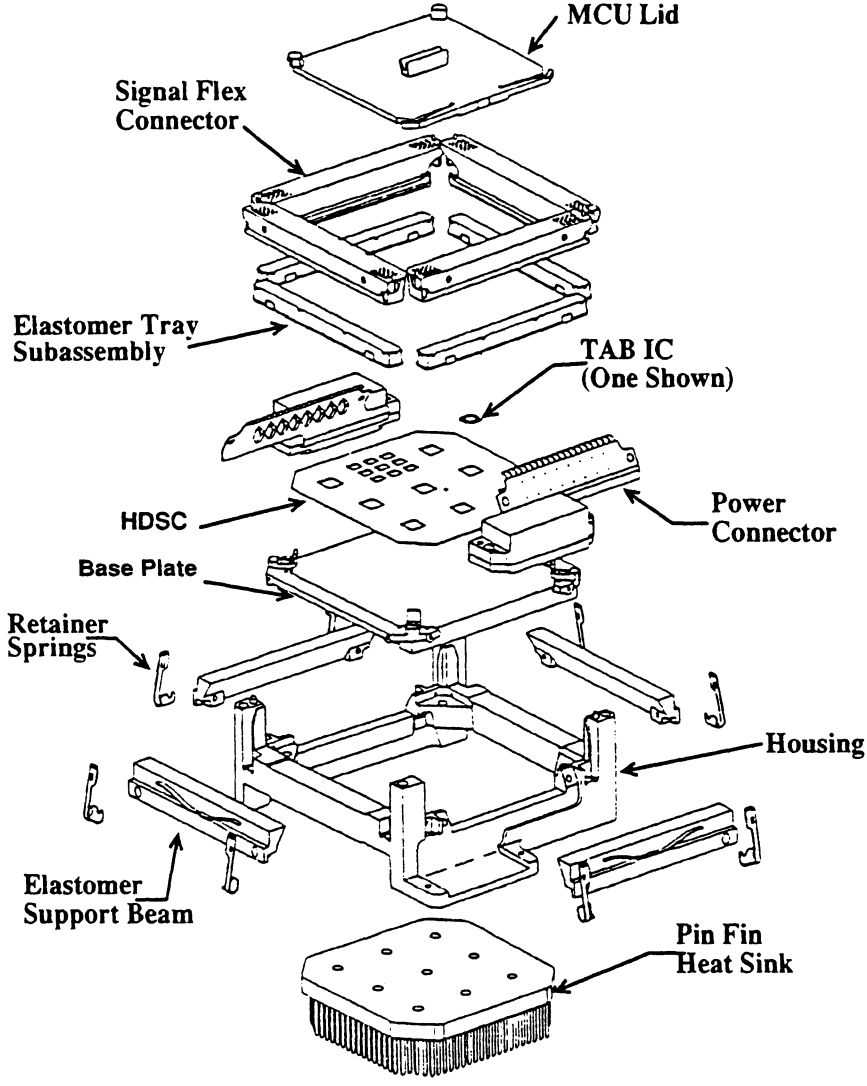

Figure 17-3 An exploded view of the MCU highlighting the individual components.

Table 17-2 Characteristics of MCU Chips.

	MCA-III Gate Array	CDXX (Clock Distribution)	4 K × 1 STRAM	1 K × 1 STRAM
Die Size (mm)	9.8 × 9.8	6.2 × 6.2	6.4 × 4.2	4.9 × 3.6
Signal Pins	256	170	35	33
Total Pins	360	272	48	48
Transistor Count	40.1K	7.2K	103.0K	28.0K
RAM (bits)			16384	4096
Power (watts)	30.0	13.9	2.4	2.4

Clock Distribution Circuit
Central to every MCU is a clock distribution chip (CDXX). This chip has master clock and reference clock inputs, from which de-skewed clocks are generated for every chip within an MCU. This chip also provides scan distribution and scan control for test and diagnostic purposes.

17.4.3 HDSC Interconnect Technology

The heart of MCU technology is the HDSC [5]. The HDSC is built with nine copper conducting layers separated by polyimide dielectric layers. Four metal layers are used for power distribution. Of the remaining five copper layers, two are used for signal interconnects in the x- and y- directions respectively, two act as reference planes for the signal layers and one layer provides surface pads for electrical connections. The HDSC is capable of supporting three to five times the routing channels per unit area available in conventional printed wiring boards. A cross section of the HDSC is shown in Figure 17-4 (not to scale).

Signal Interconnect
The performance of the signal interconnect is determined by the dielectric constant of the insulating medium and the resistance of the metallic interconnect traces. A fairly low dielectric constant (3.5) is achieved by using polyimide as an insulator in the HDSC. The physical dimensions of the interconnects are set so that the line resistance is 1 Ω/cm. The two signal layers in the HDSC are

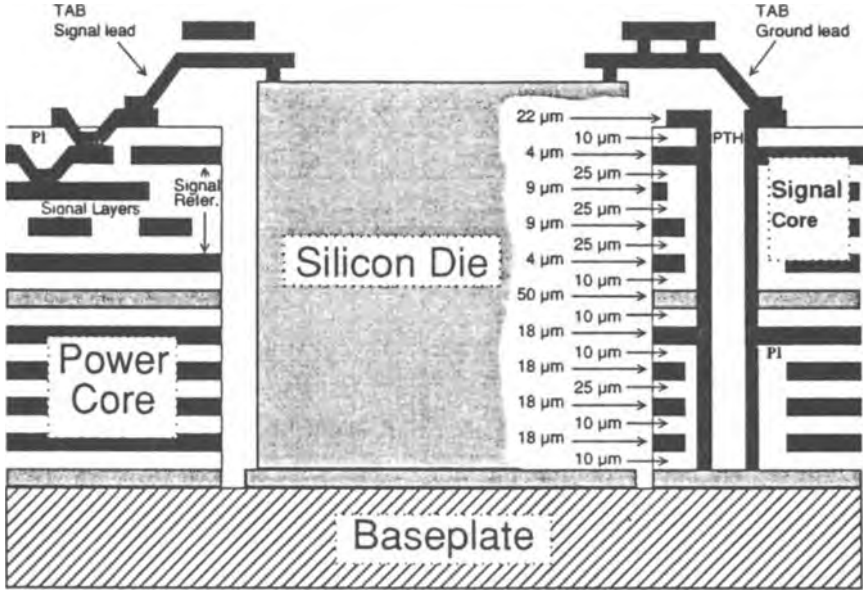

Figure 17-4 HDSC cross section with metal and dielectric thicknesses indicated.

enclosed by upper and lower reference planes to form a stripline. The cross section of this structure is symmetric so both signal layers have the same nominal impedance value of 60 Ω. Crosstalk is held to a minimum by controlling the line resistance and holding the minimum line pitch to 75 µm line pitch. Plasma etched vias create paths for connections between the signal layers and the top metal layer. The pitch of these vias allow for a minimum diagonal placement so that a true 75 µm line pitch can be achieved. Electrical connection to the lower reference plane is made with drilled and plated through-hole technology.

The master and reference clock inputs to the clock distribution chip require 50 Ω impedance, instead of the 60 Ω impedance of the signal interconnects. In this case, the signal traces are made wider (33 µm as opposed to 18 µm) to achieve the desired impedance.

Power Distribution

The HDSC power distribution is designed to minimize IR drops and provide the highest possible level of signal decoupling. Four power planes are used. Two power planes conduct ground or V_{CC} currents, one conducts V_{EE1} (-5.2 V)

Table 17-3 MCA-III Chip Parameters.

Chip Dimensions	385 mils square
Number of Signal Leads	256
Number of Power Leads	104
Total Number of Leads	360
Inner Lead Pitch	4 mils
Output Signal Rise Time	600 ps
Maximum Chip Power	30 watts

currents and the last one conducts V_{EE2} (-3.4 V) currents. Each of the supply planes is paired with a V_{CC} plane separated by polyimide thickness of 10 μm; this produces a distributed capacitance of 0.31 nF/cm^2. The effective impedance across a single power plane is on the order of 20 mΩ. The resistance of each is determined by the low sheet resistance of only 1 mΩ/□.

HDSC Assembly

The HDSC is made up of a top layer of signal interconnect (signal core) laminated to a lower layer of power interconnect (power core) as shown in Figure 17-4. This laminated structure is drilled and plated to connect the power interconnect to the top surface. The top layer of metallization, which later will be bonded to the ICs and the MCU connectors, is also formed in this plating step. Cut outs are excised at the sites where the chips are to be attached and bonded and, finally, the assembly is laminated to a heat spreading and mechanically supporting baseplate.

17.4.4 MCU Structure

The MCU is built in two stages. First the ICs are bonded to the HDSC to constitute the MCU subassembly. Then the MCU subassembly is joined to its housing, connectors, lid and heatsink to form the final MCU assembly.

MCU Subassembly

Prior to subassembly, the integrated circuits have copper clad tape bonded to the gold plated pads on the die [6] (see Figure 17-5). At this point the die are tested at speed and screened before shipment to the MCU assembly plant. The gate arrays and custom logic ICs use a special two metal tape that dedicates one metal layer to a reference ground in order to control impedance and crosstalk. The inner lead pitch is 100 μm for the logic and the clock chips and 450 μm for the STRAMs.

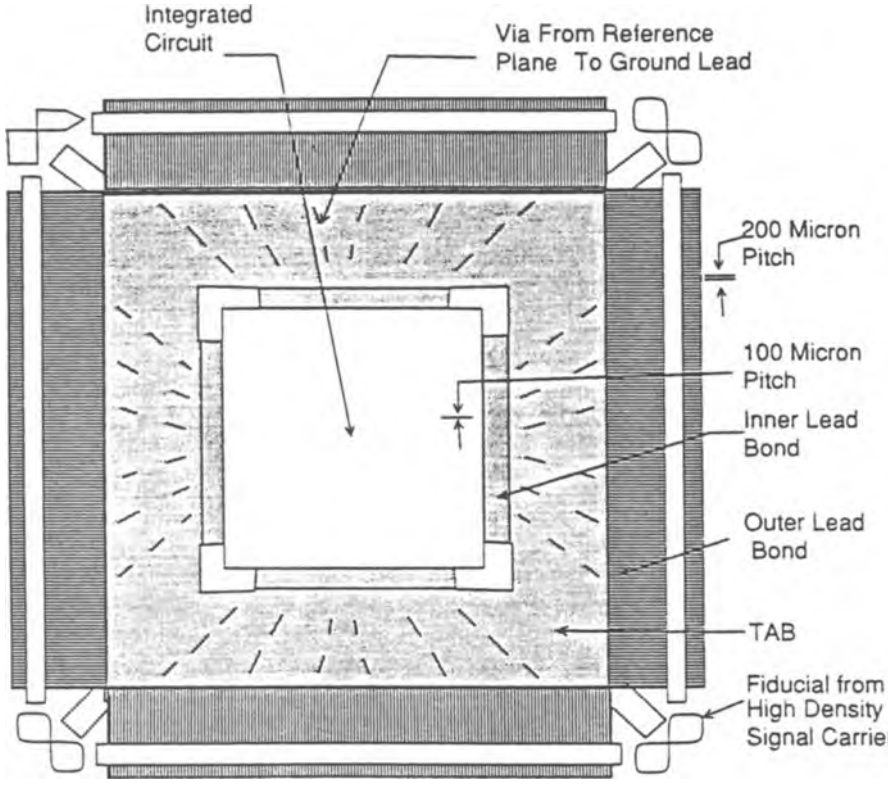

Figure 17-5 MCA-III on TAB frame. Top view of the TABed gate array.

The die are attached to the HDSC baseplate with a diamond filled epoxy. This provides a low thermal resistance path (0.3°C/W) that is electrically isolating (see Section 17.4.6). The electrical isolation is required because the back side of the die is at V_{EE1} potential while the baseplate is required to be at ground potential. During die placement, the outer copper leads of the ICs copper clad tape are formed into the shape required for their subsequent bonding to the solder coated pads on the HDSC. The outer lead pitch of these leads is 200 μm for the logic devices and 450 μm for the memory devices.

MCU Final Assembly

The MCU assembly is complete when it is ready for mounting onto a system's planar board. The electrical connectors are designed to make contact with the

planar board through static pressure (no soldering), so that the MCU can be quickly installed or removed, a fundamental requirement for a FRU. The key parts of the final assembly are the subassembly (the HDSC complete with bonded chips), the power connectors and the signal connectors. There are two identical power connectors bonded to the HDSC on opposite sides and there are four identical signal connectors bonded to the four sides of the HDSC.

MCU Power Connector
Each power connector consists of a semi-flexible tape of three isolated metal planes, eight discrete decoupling capacitors (0.22 μF each) and a spring connector through which it contacts the power bus bars on the planar board. This structure continues the low impedance design implemented in the power planes of the HDSC. Connection to the HDSC is made by soldering 20 metal tabs, each 3.3 mm wide, on a 3.6 mm pitch. Each tab makes connection to the power core of the HDSC through six plated through-holes. The tabs alternate between V_{CC} and either V_{EE1} or V_{EE2}.

MCU Signal Connector
The signal connector is also a flexible tape structure. Its leads are on a 300 μm pitch and are bonded to the HDSC pads through solder reflow. While each connector has 201 signals, there is also a ground return for every three signals, resulting in a total of 267 leads connected to the HDSC. On the side of the signal connector that mates to the planar board, an array of gold bumps on its flexible circuit makes a wiping contact as pressure is applied. Between the gold bumps and the soldered leads the length of the signal conductors is approximately 2.5". The signal conductors reside on one side of the tape and a solid ground plane on the other side. This microstrip structure guarantees controlled impedance signal lines with minimal crosstalk.

17.4.5 MCU Electrical Characteristics

Performance of the MCU is defined by the speed of the internal ICs, by the fidelity of signal transmission through the signal connectors and by suppression of noise in the power distribution system. Overall, these elements are controlled sufficiently to allow for a system cycle time of 16 nsec for the complex instruction set computer (CISC) CPU and for distribution of clocks of up to at least 750 MHz. At the same time, MCU power distribution noise is held within the noise budget of 20 mV.

Signal Integrity
All components of the MCU interconnect - the two layer TAB tape, the HDSC and the signal connector - are designed primarily for a characteristic impedance

of 60 Ω and this characteristic impedance is guaranteed by manufacturing tolerances of less than ±5%. In addition, the crosstalk for signals in the HDSC is controlled to less than 5.1%. The resistance of the HDSC signal lines is 1 Ω/cm, enough to help attenuate reflections without adversely affecting delay. Termination of the HDSC transmission lines is accomplished by the use of series terminating resistors of 30 Ω within the ICs.

The gate array ICs required tape to fanout from a 100 μm inner lead pitch to a 200 μm outer lead pitch. The distance required to achieve this fanout is more than 4.0 mm. The two metal tape provides a ground plane that controls the impedance and the inductance. With the plane, the impedance is 60 - 70 Ω and without it, the impedance will exceed 100 Ω.

Several design techniques were used for further improvements in signal integrity. All clocks are delivered in differential form and the clock traces in the HDSC are routed as fixed length, minimum pitch pairs. Restrictions are placed on standard signal traces to limit the length of adjacent traces and crossovers so that crosstalk is minimized beyond the structure specification. Series terminated ECL (STECL) outputs were used to reduce simultaneous switching noise by more than 50% compared to the parallel terminated option.

Power Connections and Distribution

The MCU DC circuit for ground current begins at the pad of the IC, flows through the inner lead bond, through the tape lead to the outer lead bond, through the conductors on the top metal layer of the HDSC and through the top vias to the top reference plane. The top reference plane then spreads the current through all V_{CC} plated through-holes to the power core. The power core shares the ground current between the two ground planes and conducts it to the HDSC perimeter, where additional plated through-holes conduct the current to the power connector bond pads. From the bond pads the current flows through the flexible tape to the spring contacts of the power connector. Similar paths are used for the two power supplies as well.

The resistance, as measured from the outer lead bond pad to the spring contacts of the power connector, is less than 3 mΩ for V_{CC} and less than 30 mΩ for either of the two power supplies. The total distributed capacitance between V_{CC} and either power supply is 30 nF in the HDSC. The distributed nature (ultra low inductance, less than 10 pH/□) of this capacitance provides very effective suppression of high frequency noise [7]-[8]. The supply noise is measured to be less than 30 mV on this power distribution.

17.4.6 MCU Thermal Management

The VAX-9000 was designed to provide an efficient cooling path for high power ECL circuits. This was achieved by conducting the heat generated by the silicon

chips through the die attach material, into the baseplate and, finally, into an air stream via a pin fin heatsink [9]-[11].

Since a single ECL chip of the type used in the VAX-9000 system can generate up to 30 watts, a typical MCU generates between 145 - 270 watts of heat. However, because the overall thermal resistance for each MCU chip site is less than 2.0°C/W, the maximum junction temperature will be 85°C for an ambient temperature of 25°C as specified in the design requirements. Achieving this low junction temperature was critical to meeting the reliability goals of the MCU. Consequently, extensive thermal modeling took place in the early design phase to test and verify the feasibility of such a system. A detailed description of the MCU thermal model has been published by Fitch [12].

A significant component of the overall MCU thermal management system is the unique die attach material that was developed specifically for this program. The die attach material ensures maximum heat transfer and high flexibility to accommodate the differing thermal coefficients of expansion between the copper chromium baseplate and the silicon chips. This material was heavily loaded with microscopic (25 μm) diamond particles to provide the highest thermal conductivity possible and yet remain electrically insulating. The design requirement of electrical insulation prohibited the use of conventional die attach materials such as silver filled epoxies. Optimization of the epoxy formulation took several years. In addition, it was necessary to develop concurrently a robust die attach process that would safely allow multiple chip replacements.

The thermal transfer process can be understood in relatively simple terms (see Figure 17-6). The principle is based on the conduction of the heat from the silicon chips through the diamond particles embedded in the epoxy material and into the copper chromium baseplate, from where it is conducted across a dry interface to the base of the aluminum pin fin heatsink. The heatsink is described as "pin fin" because it contains 606 aluminum pins, each 0.078" in diameter, pressed into its base with 0.75" protruding into the air stream. The heat is then removed by air plenums in the VAX-9000 cabinets, which direct approximately 14.6 *l*/sec of air into each of the MCU heatsinks. In this manner, the IC junction temperature is controlled in the VAX-9000 system.

17.4.7 MCU Test Technology

Before an MCU is installed in a VAX-9000 computer system, each module is tested at least 11 times at various stages in the manufacturing process. The sequence of test operations performed on each MCU is referred to as the test flow. The test flow for the VAX-9000 is shown schematically in Figure 17-7.

The three major test processes performed in the manufacture of each MCU

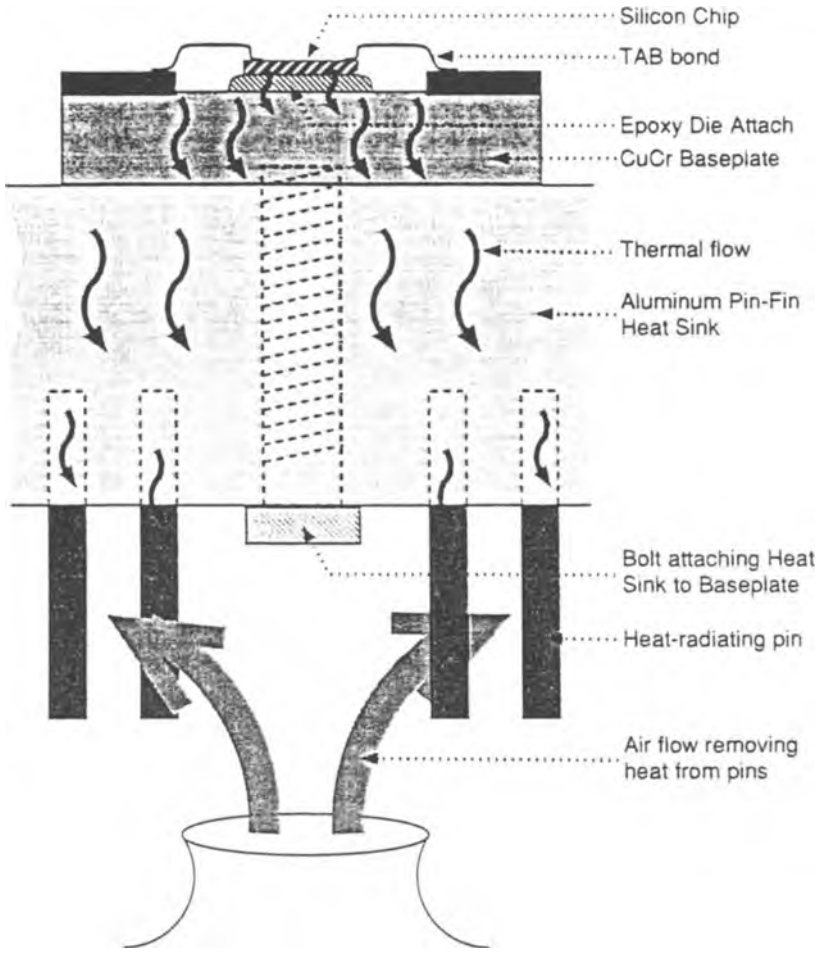

Figure 17-6 The heat conduction and convection paths for the MCU.

are in-process HDSC test, final HDSC test and final MCU test. Each of these three processes can be broken down further into sub-processes and each sub-process can be described as a sequence of test operations. For example, in-process signal core test can be broken down into three successive test operations: surface leakage resistance, via resistance and line resistance measurements. The

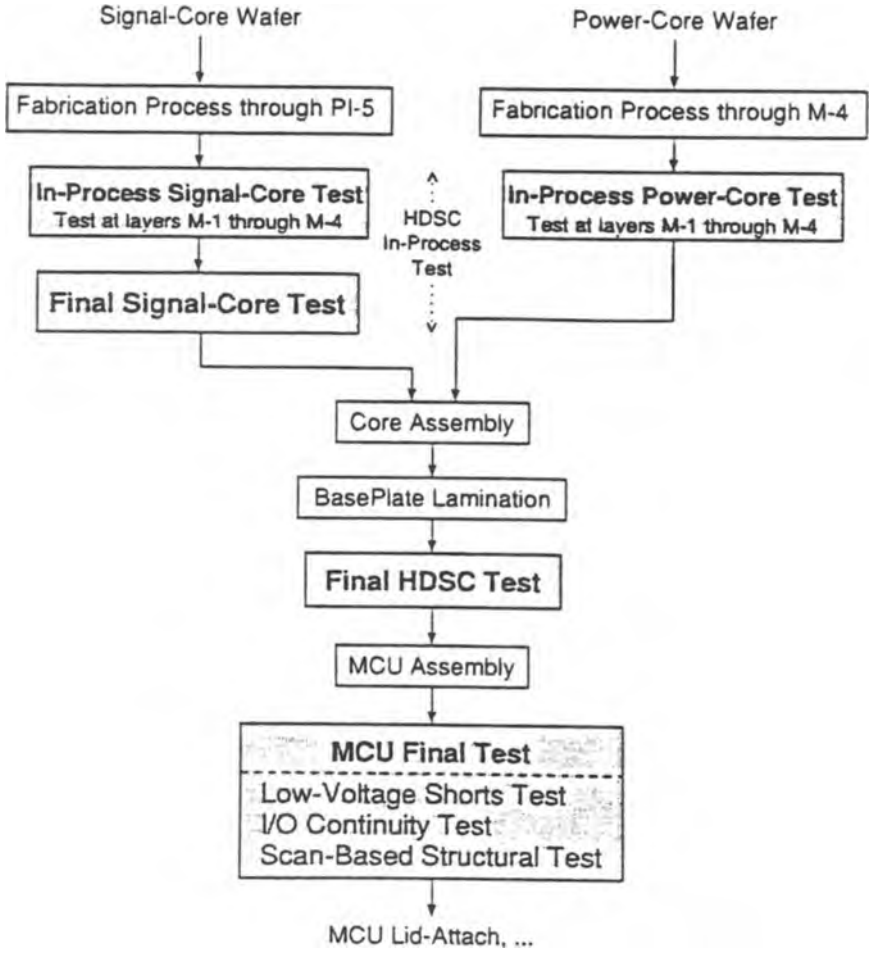

Figure 17-7 This test flow shows that the interconnect is tested several times before MCU assembly and final test.

type of operations and the sequence in which they are performed are chosen to provide visibility into important process parameters that must be monitored and controlled to optimize yield.

HDSC In-Process Test

Both in-process and final signal core tests are performed with a test system consisting of a KLA 1007 wafer prober with auto-alignment and cassette to cassette wafer handling capabilities which is coupled to a semi-custom instrumentation package (Reedholm Instruments Corp. (RI-15a.). An 80386 IBM compatible PC acts as the system controller.

Parametric tests, such as 2-point and 4-point resistance and leakage tests, are carried out with the Reedholm Instruments DC measurement modules. Figure 17-8 illustrates the basic principles underlying the resistance test methods used in in-process testing.

Signal Core and HDSC Final Test

To optimize yield, each net on the signal core must be tested for opens and shorts before lamination to the power core. To optimize cost, test time must be minimized. The most complete test approach would directly test each net segment for DC continuity, measure the DC isolation of each net to every other net and evaluate the performance metric (impedance, crosstalk, propagation delay) of each net. The time and cost of this approach is prohibitive and unnecessary. Parallel probing (such as bed-of-nails) is not practiced due to the small pad size and pitch and high pad count (≈ 4000). Probing each net with one or two probes is acceptable for DC continuity testing but the number of probings required to do the DC isolation test grows as the square of number of nets. To achieve the best optimum of quality assurance and test cost, a capacitance based test approach can be used to verify net continuity and isolation. Since impedance is directly related to the capacitance, the performance of the signal core can be assured. This approach only requires a single probe test setup and only one probing per pad.

Open and shorts testing is carried out by measuring the capacitance at the ends of each net in the interconnect device using a capacitance meter such as the HP 4278A. In Figure 17-9a, it can be seen that for a good net the capacitance measured at either end of the net will have the same value. For an open net, as shown in Figure 17-9b, the capacitance at either end of the net will be lower than the expected net capacitance and the sum of the values at the ends of the net will be equal to the expected value. In the case of a short, as in Figure 17-9c, the capacitance measured at any pad of any shorted net will be equal to the sum of the expected net capacitances for the individual nets.

In practice, subtle processing defects create more complex interconnect faults than the simple cases of Figure 17-9. For example, Figure 17-10a illustrates a resistive open, that is, where a resistive path exists rather than a complete open. This situation, as well as the capacitive open illustrated in Figure 17-10b, can be caused by defective vias between two metal layers.

a) Two-Point Resistance Test (Force current, measure voltage)

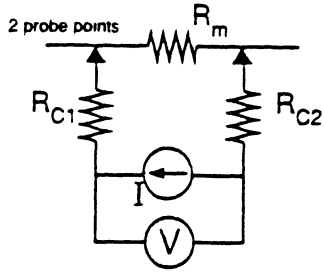

$$R_m + R_{C1} + R_{C2} = \frac{V}{I}$$

Assuming $R_{C1} + R_{C2} \ll R_m$

then $R_m \cong \frac{V}{I}$

b) Leakage Resistance Test (Force voltage, measure current)

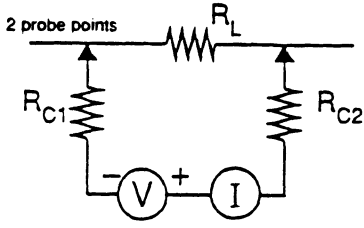

Assuming $R_{C1} \approx R_{C2} \ll R_L$

then $R_L = \frac{V}{I}$

c) 4-Point Resistance Test (Force current, measure voltage)

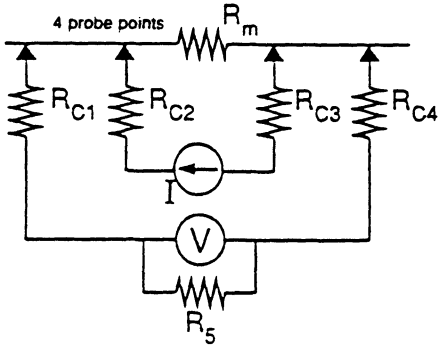

Assuming

$R_{C1} \approx R_{C2} \approx R_{C3} \approx R_{C4} \ll R_5$

then $R_m = \frac{V}{I}$

Figure 17-8 Basic resistance test configuration. The probe and meter connections for achieving three types of resistance tests.

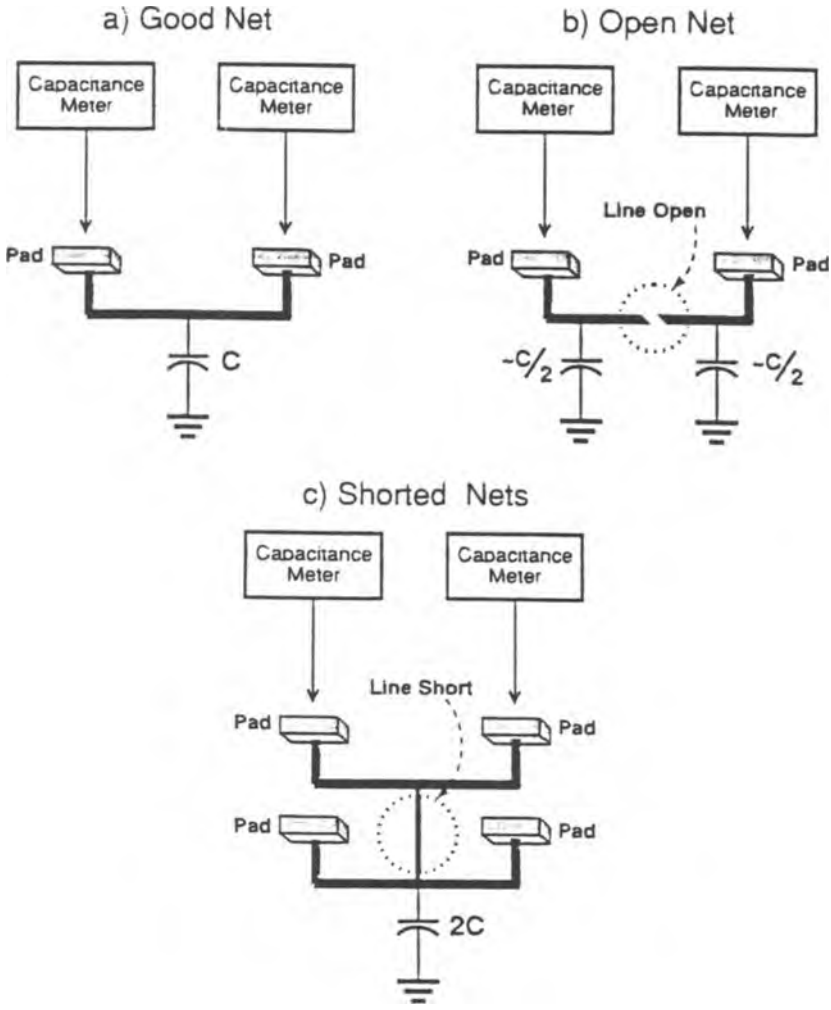

Figure 17-9 Ideal capacitance results for good, open and shorted net conditions.

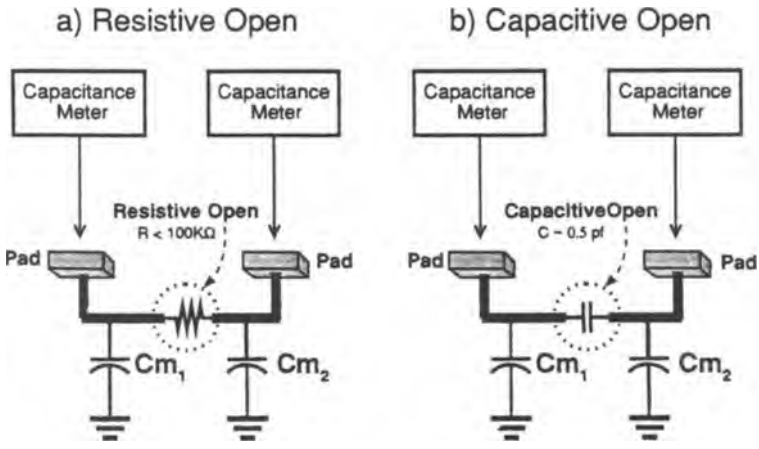

Figure 17-10 Effects of "soft" opens - resistive and capacitive.

In practice, after the capacitance measurements are taken by the tester, the results must be further analyzed to minimize the chance of either false negatives or false positives. False positives can result if a resistive or capacitive open is misinterpreted as a good net. False negatives can be caused by equipment malfunctions or by contamination on either the probe tip or pad.

The expected net capacitance values can be either calculated from design data or derived from measurements of known good devices (commonly referred to as the "learned mode" method). When the expected values are calculated from design data, there are many factors which must be taken into account in order to obtain accurate results. It is straightforward to calculate the capacitance of a parallel plate capacitor if the dielectric constant, the distance between the plates and the area of the plates are known. However, in controlled impedance interconnect networks, the geometries of the conductors are such that fringe fields become a very important factor in the actual capacitance value. Since more than 20 MCU design options were used in the VAX-9000 system, it proved to be worthwhile to determine the expected capacitance values directly from the design data. There were two compelling reasons for this: automated test file generation capability saved a great deal of time in the debugging phase, when design changes occurred regularly and it was possible to detect and correct masking errors and other sources of repeating defects.

An additive capacitance model was developed to estimate the capacitance of each net. The model used design dependent data such as pad type and pad count, line length and crossover count for input data. Capacitance parameters were determined by two dimensional capacitance modeling and a 10% tolerance band was applied to most of the parameters to generate limits that would screen opens and shorts. This tolerance band was also tight enough to screen out parts that would be outside the impedance specification.

Prior to chip connection and TAB bonding, the net capacitance test is repeated to insure that no damage was done during power core and baseplate lamination. In addition, the PTH and power planes are tested.

The PTH are tested using the force current, measure voltage technique described earlier and Figure 17-8c. This measurement allows weeding out parts with suspiciously high PTH resistance. It has been found that PTHs with resistances above several mΩ are subject to reliability problems in the event that the part undergoes extensive thermal cycling. Fortunately, the PTH yield has been found to be very high.

The dielectric integrity of power planes is tested using the force voltage, measure current method described earlier and in Figure 17-8b. A constant voltage (50 volts) is applied to the power planes for one minute, and then the current is measured. In this test, the reject threshold is 100 MΩ.

The final HDSC test suite, which includes resistance and leakage measurements, is performed on a flying probe tester (the STAR tester), designed by Digital especially for the VAX-9000 program. Multichip module substrate testing is covered in more detail in Section 13.2.

MCU Final Test

Because of Digital's experience in testing modular computer systems, the difficulties in testing the VAX-9000 MCU were recognized early in the design process. Given that the module would operate with a clock frequency of 500 MHz and would have over 800 signal inputs and outputs, it was clear that no available commercial tester could provide adequate test capabilities. It was also clear that a viable test strategy and methodology had to be developed (other than using the completed computer system as a tester), since 20 distinct MCU types are required for a VAX-9000 system even in its uniprocessor configuration. It was therefore necessary to undertake the development of a special tester for MCUs; it would be a scan based tester, capable of testing each MCU type individually and independently from the assembled computer system.

One of the important challenges that arises when scan based techniques are used to test a modular system is that even though a module may pass the scan test successfully, it can still fail to function correctly in the final system if the scan test coverage is less than 100%. In practice, test coverage is almost always

less than 100%, particularly when a design is newly released to production. Generally a period of maturation is required, during which the set of test vectors for each module is enhanced based on actual experience. Subtle and even not so subtle fault mechanisms are discovered as the number of modules manufactured increases. For this reason, the assembled set of modules must always be tested as a complete system prior to shipment.

Another way of describing the implications of scan based testing is to say that scan based testing provides a validation of an MCU's logical structure but it will not necessarily guarantee its functionality. Scan testing will not detect a logic design error if generated from design data. This error will only appear at the system level.

As was mentioned previously, since no suitable commercial tester was available, it was necessary for Digital to develop its own MCU tester. Digital's MCU tester has now been fully developed and deployed in a production environment for several years. The actual sequence of tests performed are the following:

1. Low voltage pin to pin shorts test (LVS)
2. Input and output pin contact check
3. Scan based structural test

These tests are described in more detail below.

In the LVS test, a small DC voltage (below the level to turn on a PN junction) is applied to the power supply pins, input pins and output pins while the current flow is observed. If all of the currents are below a preset threshold, then the unit passes. If any pins fail, then the MCU is sent for diagnosis and repair.

In the input/output pin contact check, the voltage on each pin is set to a level just high enough to turn on the substrate diode and thus cause a small current to flow. If current flow is sensed at all pins, then the unit passes. If the unit fails, then the problem may be in the connector or the MCU. If the unit does not pass after being re-seated in the test fixture, then the MCU is sent for diagnosis and repair.

Finally, in the scan-based test, an automatically generated set of test vectors (unique to each MCU design option) is scanned into the module through a dedicated scan port. The module is clocked and the data is scanned back out and interpreted by a computer program. If a fault is detected, then the module is sent for diagnosis.

17.5 VAX-9000 MCU Manufacturing

17.5.1 Overview of Manufacturing Process

MCU manufacturing has five major processes: Signal Core Fabrication, Power Core Fabrication, HDSC Assembly, MCU Subassembly and MCU Assembly (Figure 17-11). The signal core and power core processes are similar to IC processes. The HDSC and MCU subassembly processes are like PWB fabrication and IC packaging processes respectively. The core processes are performed in a class 100 cleanroom environment, whereas the assembly processes are executed in a class 10K environment.

The philosophy guiding HDSC/MCU manufacturing has been that of problem prevention and reduction of variability. Up front product specifications were solidified and process capabilities were assessed and improved to ensure a high C_{pk}. Equipment preventive maintenance (PM) and calibration methodologies were developed and implemented. Training modules were developed; it was made mandatory that operators pass certification tests before being allowed to work in the fabrication areas. The fabrication facility was certified and its key parameters were measured and kept in control. Multifunctional process improvement teams were formed and underwent thorough statistical process control (SPC) training. A shop floor control software system was used to manage the work in process (WIP) and to collect and analyze the process data. Carefully designed experiments were used to characterize the process and to expedite problem resolution. In process electrical tests and inspections were developed to ensure the quality of the WIP and to provide quick process feedback. As the processes matured and became more stable, sample sizes for test and inspection were reduced to streamline the process without jeopardizing the product quality. An on going Pareto analysis of rejects was maintained and failure analysis tools were used to find and fix the root causes of rejects.

Great effort was put into ensuring the quality of incoming materials by partnering with key suppliers. Ship to stock programs were developed and followed rigorously. Die were fully tested at speed in their TAB frames by the vendors before shipment to the MCU assembly operation. Vendors' process capabilities and SPC programs were assessed to ensure that product specifications would be met consistently and several subassemblies were redesigned in order to match the supplier's capabilities.

In the following sections the core and assembly processes are discussed in detail.

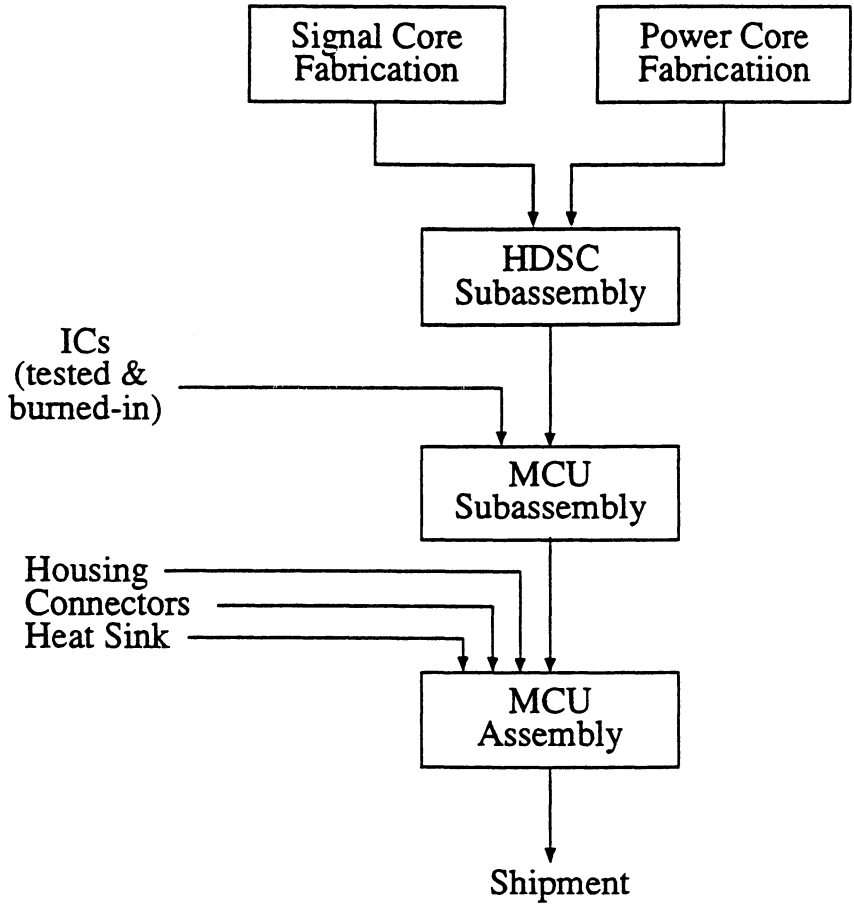

Figure 17-11 High level flow for the complete MCU manufacturing process.

17.5.2 HDSC Signal Core Fabrication Process

The signal core process follows the process flow shown in Figure 17-12. An aluminum wafer of 6" in diameter and 50 mils in thickness is the starting material. The wafer is first solvent cleaned and then is coated on its back side with polyimide to a thickness of 20 μm; this coating isolates the aluminum from the subsequent chemical processing. Photoresist spinners are used for the coating and the soft baking of the polyimide (PI) films. To achieve films of such thicknesses, multiple coatings are used. The polyimide is then cured in a diffusion furnace. Precautions are taken to minimize oxidation in the furnace by controlling the leak back of ambient air and also by cooling the wafers before removing them from the furnace. Another layer of 10 μm thick polyimide is then spun on and cured on the front side of the wafer; this layer protects the first metal layer when the aluminum substrate is later removed.

A Cr/Cu/Cr layer then is sputtered on in a vacuum system. A photoresist layer is spun on and exposed through a mask to define the reference plane. A chemical wet etch process is used to etch off the excess metal. A third layer of polyimide is spun on and cured to separate the reference plane from the signal lines.

A Cr/Cu seed layer is sputtered on and a signal line pattern is defined in the photoresist. Copper is electroplated to form the signal lines. To ensure the uniformity of the thickness, care is taken in pre-cleaning steps and in bath maintenance. The mask layers are thieved to create a uniform pattern and a special wafer plating holder is used. This holder allows for multiple perimeter contacts to the wafer but is sealed from exposure to the bath; the seal prevents plating to the contacts and to the wafer perimeter and ensures repeated uniform contact.

The resist is then stripped, another layer of thicker resist (≈40 μm) is spun on and holes are opened where vias between the two signal layers are required. The holes are then filled with copper by an electroplating process. Because there is minimal redundancy, this is a very critical step; all the holes have to be filled with copper - any small residue left in the hole could prevent or slow down the plating and lead to open vias.

A seed layer is then chemically etched and another thick polyimide layer is spun on and cured. A via mask is then patterned in the thick resist and is plasma hardened so that it can withstand the plasma during the subsequent SF_6 + O_2 plasma etch process that opens the vias. After the via etch and resist strip steps, another Cr/Cu seed layer is sputtered and the process for defining the signal layer and the vias is repeated. A Cr/Cu/Cr layer is sputtered and the second reference plane is defined and etched.

A final polyimide layer deposition and a via etch process complete the signal core fabrication. Figure 17-13 shows the cross section of a finished signal core at various steps of the process.

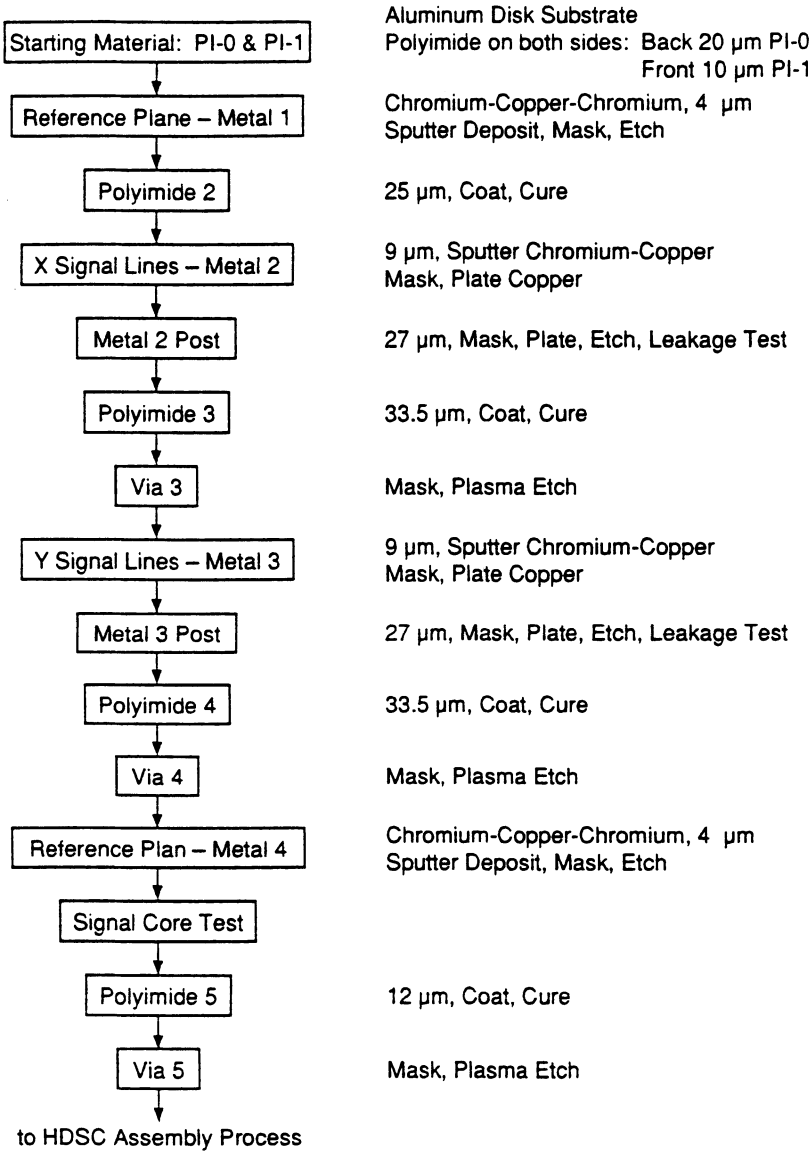

Figure 17-12 Major steps of the HSDC signal core process flow.

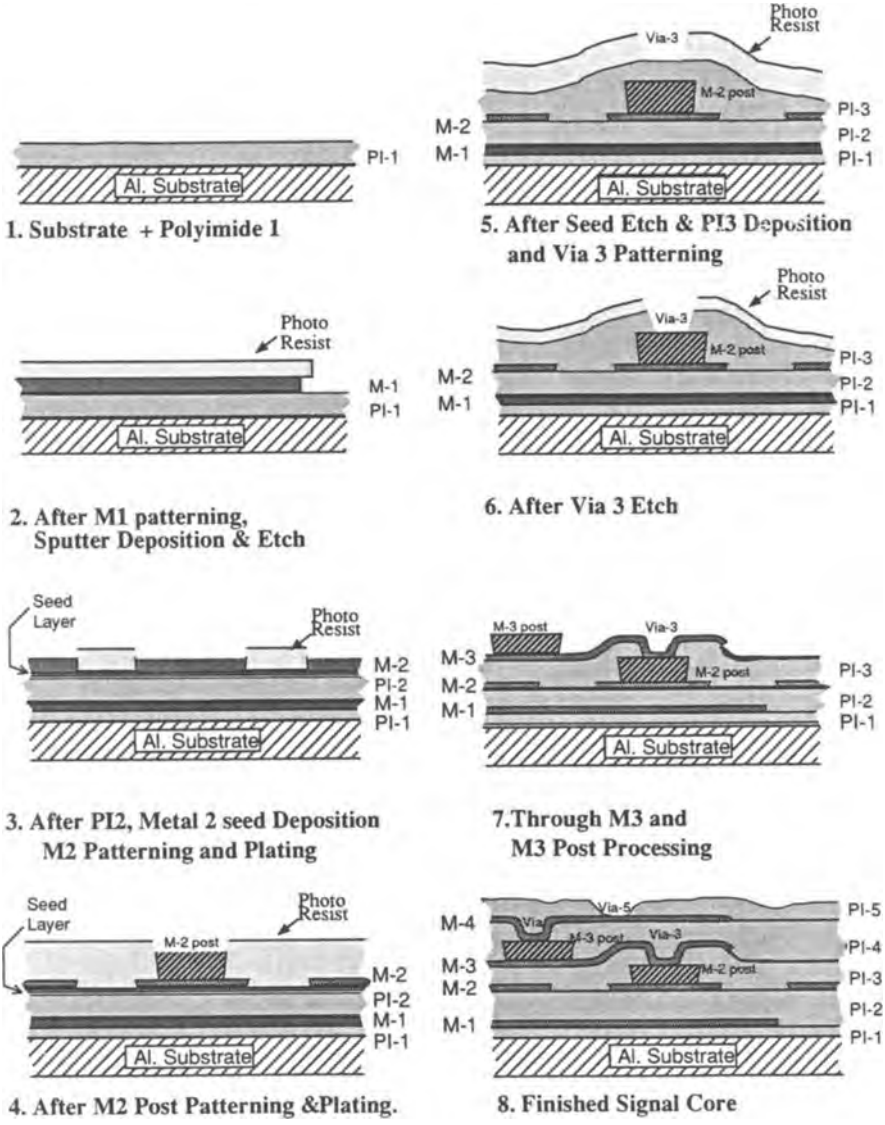

Figure 17-13 Cross sections of the signal core after major steps of the process.

17.5.3 HDSC Power Core Fabrication Process

The starting material for the power core is also a 6" aluminum wafer. The back side and the front side are coated with polyimide through processing steps similar to those used in the signal core process. After the front side polyimide coating, a seed layer of Cr/Cu is sputtered and then built up by use of electroplating. A thin layer of chromium is then sputtered on to ensure good adhesion between the power planes and the polyimide.

Using a photolithographic and a wet etch process, the first power plane is defined. Both interlayer and intralayer tests are performed on all the wafers. Rejected wafers are sent to failure analysis and good wafers are moved forward through the process.

The polyimide coating, metal definition and test steps are then repeated three more times. A top polyimide coating layer then completes the fabrication of the power core. A cross section of a completed power core is diagrammed in Figure 17-14. The detailed process flow is illustrated in Figure 17-15. As can be seen in the figures, there are no vias in the power core; the connections between the power planes and the outside world are made using plated through-holes built during the HDSC assembly process.

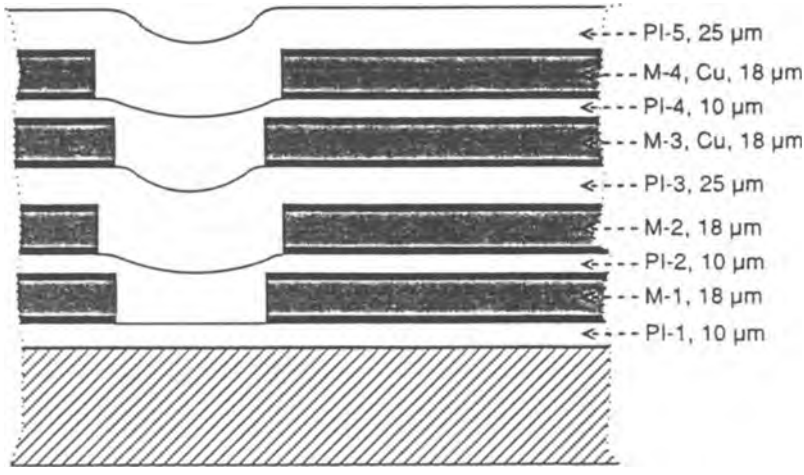

Figure 17-14 Power core cross section with layer thicknesses shown.

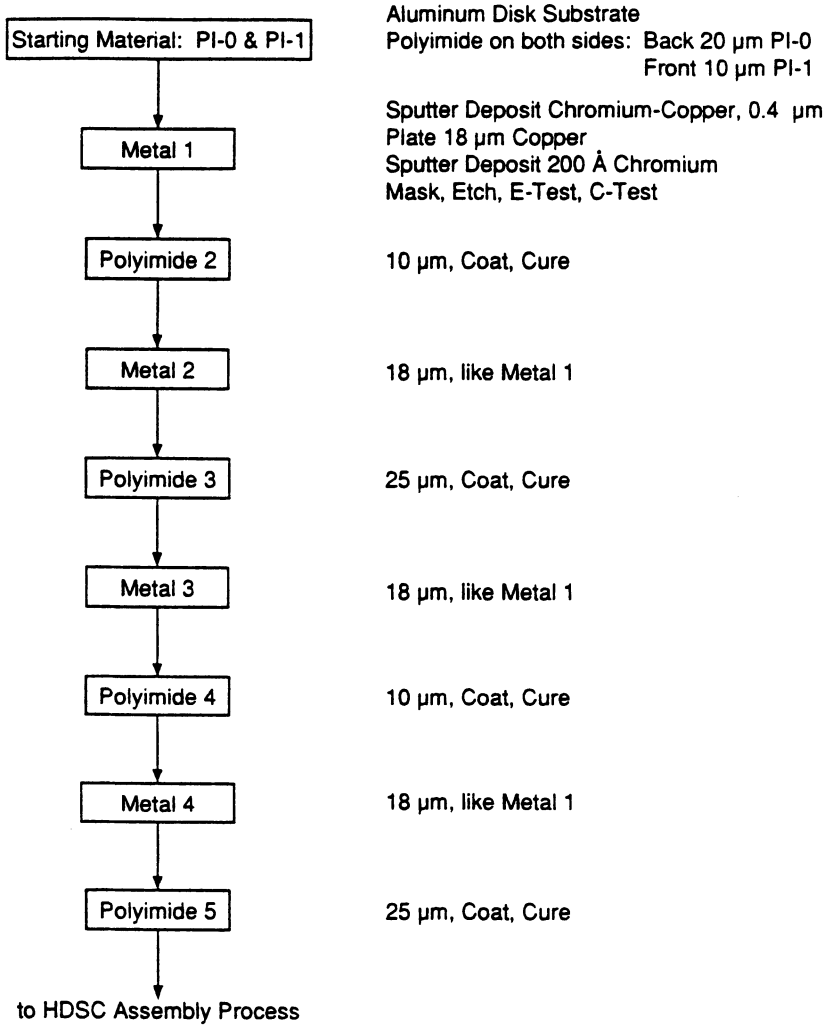

Figure 17-15 Major steps of the power core process flow.

17.5.4 HDSC Assembly Process

In the HDSC assembly process, a signal core and a power core are joined together and then attached to a copper chromium baseplate. The assembly process flow is described in Figure 17-16.

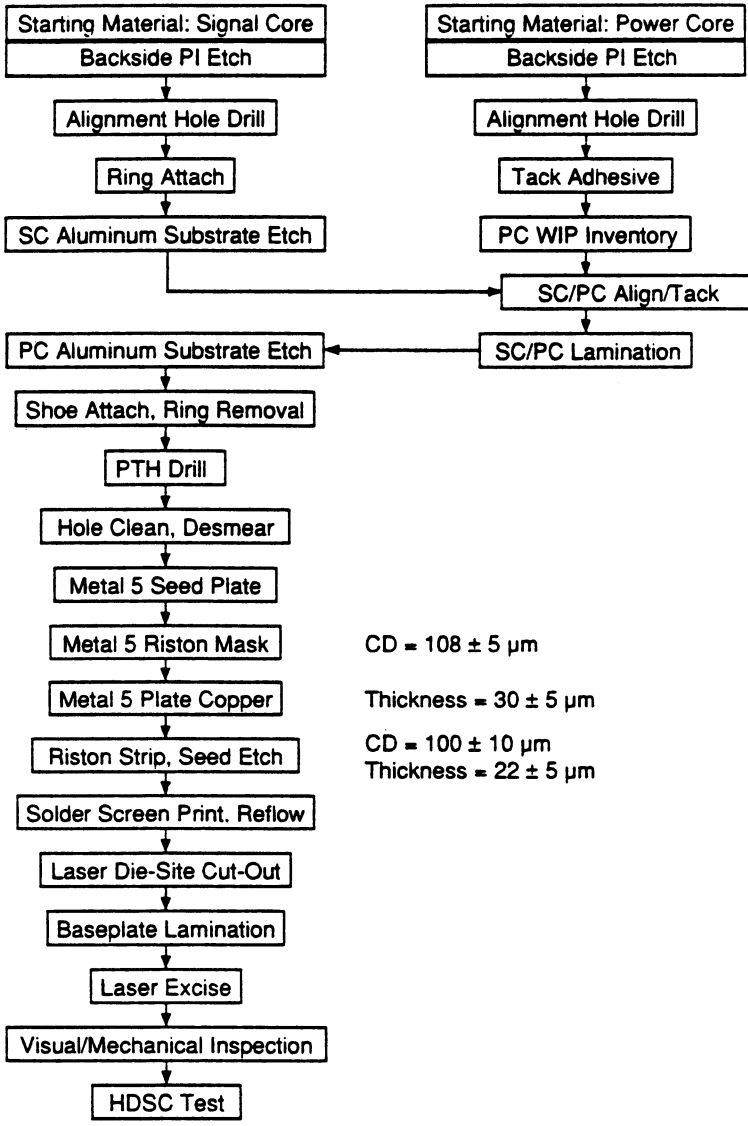

Figure 17-16 Major steps of the HDSC assembly process flow.

The back side polyimide is first plasma etched from both the signal core and the power core. After attaching a support ring to the signal core, its aluminum wafer is etched away, leaving its thin film interconnect on the ring.

An adhesive is tacked onto the power core. The signal core is then aligned to the power core and is tacked to the adhesive on the power core. The whole assembly is then put in a hydrostatic lamination press and the adhesive is cured under pressure.

The aluminum wafer is then etched from the power core to form a composite interconnect decal.

Using a mechanical drill, holes are drilled where connections are required between the power planes and the reference planes. The holes are then cleaned to remove debris and smear, if any, from the metal planes. The debris is removed with an aqueous shear stress clean; smear is removed by plasma etching. This is a very critical step and is controlled by monitoring the amount of polyimide removed from a test wafer. All traces of smear must be removed before any subsequent metal deposition in order to ensure the reliability of the PTH.

A Cr/Cu seed layer is then sputtered and built up using electroless and electroplating processes.

A dry photoresist film is laminated, exposed and developed to define the exterior footprint layer. Copper is electroplated, resist is stripped and the seed layer is wet etched to leave the copper footprint layer on the HDSC surface.

For the subsequent bonding of the TABed ICs, solder is required on the HDSC bonding pads. This is deposited by screening solder paste onto the bonding pads and then reflowing.

Next, the polyimide is excised with a laser to form openings in the decal where the die will be placed so that they can make contact with the baseplate. The decal is then attached to the baseplate with an adhesive. Figure 17-17 shows the cross section of the HDSC after particular process steps.

The finished HDSC then undergoes a complete electrical test (Section 17.4.7) and a thorough visual/mechanical inspection (VMI) and is then transferred to the MCU manufacturing area.

17.5.5 MCU Subassembly Process

An MCU subassembly is built by attaching and bonding the required ICs to the HDSC assembly. See steps 1 through 7 in Figure 17-18.

1. A plasma clean process is used to prepare the surface of the HDSC assembly for the die attach and lead bonding operations.

2. Through a computer controlled system that understands the footprint personality of each distinct MCU type, controls the motion of the HDSC on an x-and y- table and controls the rate of epoxy dispensing, the specified pattern and quantity of epoxy is automatically placed into the die site cut outs on the surface of the HDSC. Unique patterns have been developed for each IC type in order to allow the maximum thermal transfer to occur.

3. Through a similar computer controlled system, a low activation flux is automatically dispensed over those bonding pads on the surface of the HDSC where the TAB tape of the ICs is to be soldered. This activates the surface of the solder so that clean intermetallic bonds occur during the bonding process.

4. With automated bonding equipment, the MCU's ICs, one after another, are excised from their carrier frames and their leads are formed into the required gull wing shape. Each IC is precisely aligned by a vision system to its designated position on the surface of the HDSC. With computer managed control of the time, pressure and temperature parameters, a gang bonding thermode is then applied to the leads of its TAB tape. The solder reflow bonds the ICs leads to its corresponding surface pads on the HDSC.

5. When all ICs have been attached and bonded to the HDSC, the epoxy is cured in a nitrogen purged belt furnace. The bond line thickness, (the thickness of the epoxy material between the die and the HDSC baseplate) is critical for optimum thermal transfer. It is controlled during this process by mechanically applying pressure to the ICs during the curing.

6. To ensure that all soldered leads are bonded reliably, the lead bonds are inspected for possible shorts, opens, and/or weak bonds. Shorts and misalignments are detected by an automated vision system (MRSI 781) that calls failing and marginal conditions to the operator's attention. Opens and weak bonds are detected by an automated system (Vanzetti 6215) which strikes each bond with a pulse of laser energy and measures its thermal decay profile; this system also calls failing and marginal conditions to the attention of the operator. The operator then determines whether touch up corrective action is needed.

7. When touch up is necessary, it consists either of localized removal of shorts or of single point bond reflow.

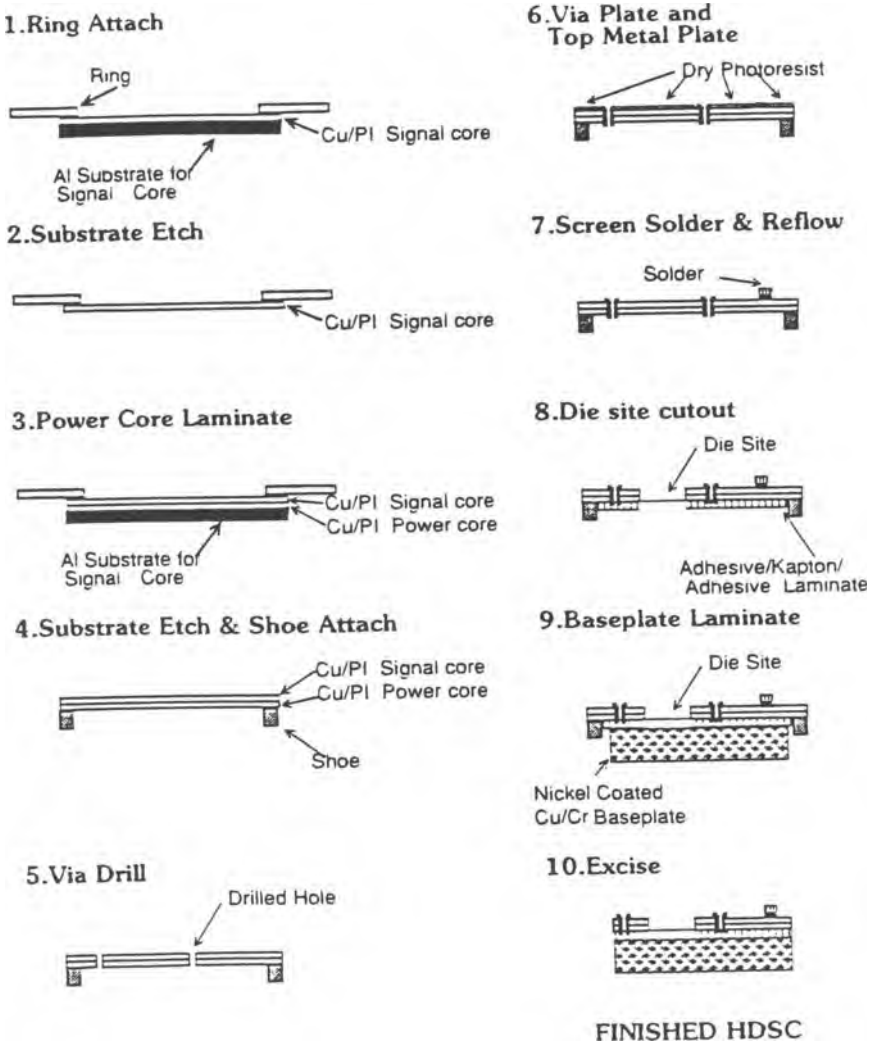

Figure 17-17 HDSC cross sections after major steps of the process.

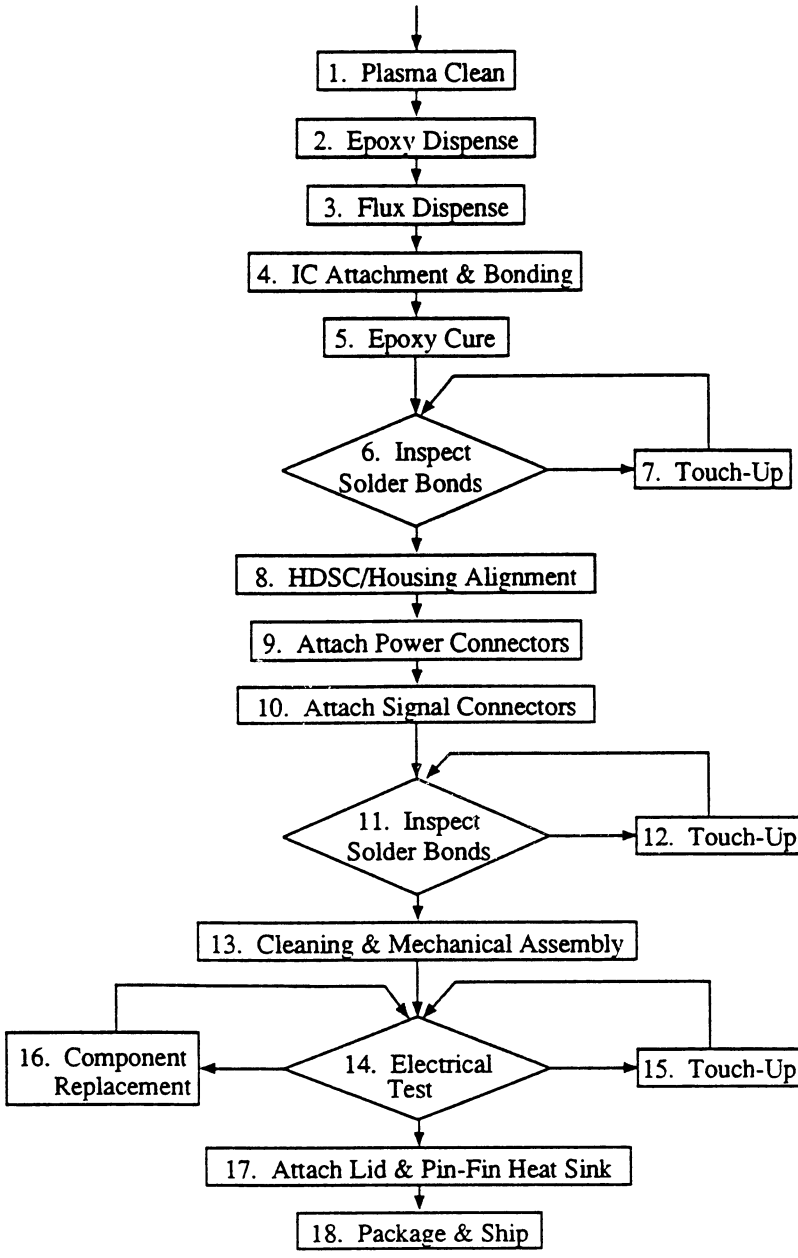

Figure 17-18 Major steps of the MCU assembly process flow.

17.5.6 MCU Assembly Process

In the second phase of assembly, the MCU subassembly is put into a mechanical housing and its remaining hardware (power connectors, signal connectors, lid, pin fin heatsink etc.) is attached to it. See steps 8 through 18 in Figure 17-18.

1. The MCU subassembly (HDSC + ICs) is put into the rigid frame of a mechanical housing and, through use of specialized equipment, is precisely aligned and locked into position with bolts. This alignment ensures that the connector systems, when mounted onto the housing, can readily be aligned with their corresponding bonding pads on the surface of the HDSC.

2. The MCU's two power connectors are bolted onto the housing, their bonding pads are prepared with flux and electrical connections to the HDSC surface pads are established by solder reflow with a single point bonder.

3. The MCU's four flexible signal connectors are attached to the housing, precisely aligned and clamped into position. The bonding pads are prepared with flux and electrical connections to the HDSC surface pads are established by solder reflow effected by series of multi-lead bonder operations.

4. With the same types of systems as mentioned above in Item 6 of Section 17.5.5, the solder bonds of the connectors are examined for shorts, misalignments, opens and weak bonds. The operator is notified if failing or marginal conditions are detected. The operator then determines whether touch up corrective action is needed.

5. When touch up is necessary, it consists either of localized removal of shorts or of single point bond reflow.

6. All flux residues are removed from the MCU through a solvent based cleaning process and solvent residues are driven from the HDSC surface with a baking process. The signal flex connectors are then folded into their final configurations and the final mechanical assembly operations are completed.

7. The completed MCU undergoes a full electrical test. If necessary, appropriate touch up operations are performed or an IC is replaced.

8. When the MCU has satisfactorily completed electrical test, a lid is placed on it to protect the ICs and the pin fin heatsink is bolted to the back side of the HDSC baseplate. The MCU is then ready for packaging and shipment to the VAX-9000 system plant.

17.6 FUTURE EVOLUTION OF MCM TECHNOLOGY

MCMs based on MCM-D technology were originally developed to meet the high performance needs of mainframe computers and military applications since performance, not cost, was the primary driving factor. However, as the evolution of MCM technology has progressed, the potential applications for MCM-D technology have become more numerous and diverse. MCM-D solutions which were once exclusive to high end systems are now being incorporated into a wide range of products in the telecommunications, aerospace, instrumentation and computer workstation market areas.

Lassen [13] has proposed a rationale under which the progression of MCM-D technology from mainframe to lower cost applications can be understood easily. In essence, he argues that as the density of interconnection pads per unit area increases, the cost of a conventional PWB solution begins to climb dramatically due to the need for more and more layers, finer line pitch and smaller PTH. As the pad density moves beyond 250 - 300 pads per square inch, MCM-D technology becomes cost competitive with PWB technology, assuming both must meet the same performance requirements. As the pad density increases beyond 300 - 400 pads per square inch, MCM-D becomes the more cost effective solution.

A case in point is the evolution of the MCM-D technology at Digital Equipment Corporation. The high speed performance of Digital's VAX-9000 mainframe CPU is based primarily on the MCM-D technology developed for and embodied in the VAX-9000 MCU. The MCM-D learning experience that has been gained over the past eight years in the design and manufacture of MCUs is now being extended to the development of more general purpose, low power dissipation MicroModules. When the interconnect pad density exceeds about 300 pads per square inch, Digital's MCM-D can meet the performance/cost requirements of workstations, low end systems, peripheral controllers and similar types of high speed electronic devices.

The two main attributes of MCM packaging from which performance improvements come are the high density device packaging and the elimination of package parasitics. The high density packaging results in time of flight and load capacitance improvements. The elimination of package parasitics improves load capacitance and inductance. The reduction in load capacitance and time of

flight reduces signal propagation delay and improves rise and fall times. The reduction of load capacitance also reduces power dissipation. Furthermore, reduced inductance in the signal path and power distribution system reduces noise and ensures reliable high frequency operation. The system speed range at which these attributes will begin to make a difference is from 50 - 100 MHz.

One of the key objectives in the migration from MCU to MicroModule technology is reduced cost. The value added performance of the MCU technology is expected to remain while the base cost is dramatically reduced, allowing for lower cost high performance systems to become practical. In the initial MicroModule applications, cost parity with conventional technology is the goal; the additional performance gained results in an improved price to performance ratio. As systems designers learn to take advantage of this technology, much greater cost benefits will arise. Specifically, if large single chip silicon designs are replaced with equal or better performing multiple chip designs, the expected cost improvement will be dramatic due to the power law governing silicon die yield. Once the potential of low cost MCMs begins to be realized through optimized strategies for design, test and manufacturing, many new application areas are expected to become apparent.

Figure 17-19 illustrates the product benefits of MicroModule technology (reliability, cost and performance) in terms of the characteristic features of the technology. Here, it can be seen how the superior heat removal, high packing density, thin film interconnect and superior power decoupling provided by MicroModule technology are interrelated and can be applied to produce products with superior cost/performance ratios.

For low power applications, it is possible to simplify significantly the processes used for the VAX-9000 MCU and yet retain essentially all of the attractive features of the MCM-D technology. This is so because low power CMOS devices require a much simpler power distribution scheme (V_{CC} and V_{DD} only) and systems with contemporary CMOS ICs consume less power and operate at considerably slower edge rates than the ECL-based system of the VAX-9000.

17.6.1 Strategy for Migration of MCM Technology from Mainframe to Low Power Applications

A successful technology migration strategy or evolutionary approach seeks to take advantage of the inherent strengths of a technology while correcting its perceived weaknesses. One strength of Digital's mainframe MCU technology is its maturity. The computer aided design (CAD) systems are ready, the module test strategies have been developed and the manufacturing processes are stable

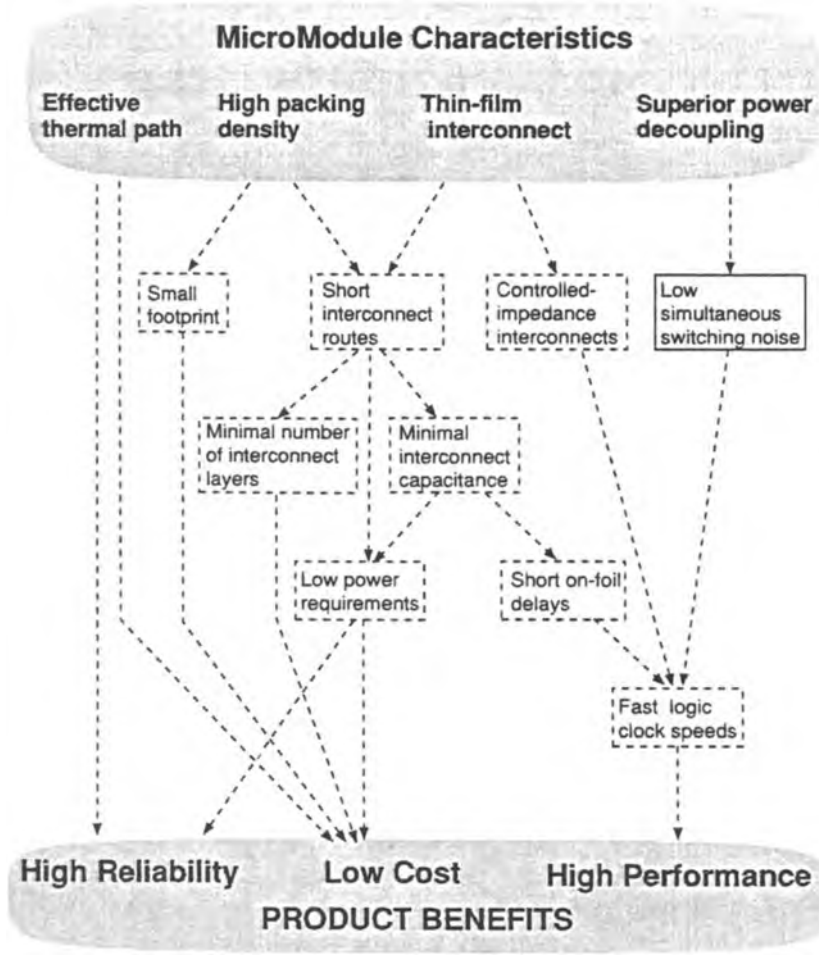

Figure 17-19 How MicroModules create product benefits.

and qualified. The weakness of most concern is cost, a consequence of designing to support high power applications. The key to lowering cost is the exploitation of the lower power requirements of CMOS-based MicroModules.

The MCU technology was developed to handle up to 270 watts. For the applications expected to be addressed with MicroModule technology, the power

requirements are expected to be well under 100 watts. Furthermore the greater signal swings of CMOS technologies will tolerate greater noise and/or DC drops. The impact on the HDI (HDSC) design and process is significant. A separate power core is no longer required. Eliminating the power core not only deletes its material contribution to the cost but also eliminates the need to remove the signal core and power core interconnects from their aluminum substrates, the need for any subsequent lamination and the need for the drill and PTH processes. The lower thermal dissipation per IC eliminates the need for the copper baseplate and the excising of die site cut outs and allows for direct die placement. Special copper/polyimide structures can be designed to operate as thermal vias, thereby creating a low thermal resistance path for the hotter chips. In as much as an electrical connection from the top metal to the aluminum substrate is made by a stairway of vias and metal pads, a thermal connection is made as well. Arrays of such structures or thermal vias can achieve less than 0.5°C cm^2/W.

Since it is no longer necessary to remove the interconnect from the substrate, the substrate itself now can be used for one of the power distribution layers. This is true whether the substrate is aluminum or highly doped silicon. While aluminum substrate does not offer the thermal expansion match that a silicon substrate does, it does offer a superior low resistance (< 0.05 mΩ/□); 500 times better than silicon. Aluminum can also serve as a mechanical element of the final package, whereas silicon would require mounting to a plate or a supporting package. Due to the slower CMOS edge rates and the fact that the entire system is contained on one MicroModule, the stripline structure of the MCU (with top and bottom ground planes) can be replaced with a simpler microstrip configuration. Because of the lower power requirement, the ability to take advantage of the substrate for power distribution and the need for only a microstrip configuration, the resulting interconnect needs only four added metal layers as compared with the nine layers of the MCU technology. A cross section of such a low power MicroModule is shown in Figure 17-20.

There are several additional approaches to lowering the cost of the MicroModule. One is to adopt standards so that multiple applications can use the same basic package. For example, a package can be developed to fit a current application but be designed with broader capabilities in mind for the future; an example of such a package is a pin grid array package (PGA). Another approach is to develop a package according to an industry defined standard, such as the 40 mm quad flatpack package (QFP). Note that for the QFP the aluminum substrate not only is electrically active but also serves as a structural element of the package. Other opportunities to reduce cost stem from the fact that the higher levels of integration in CMOS ICs are expected to result in fewer chips per system and therefore to consume less substrate area. This will allow multiple MicroModule substrates to be built on a single wafer.

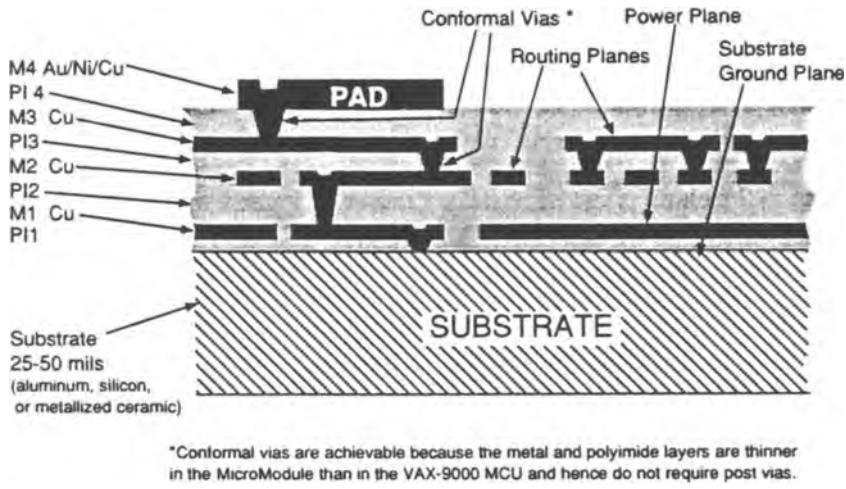

Figure 17-20 Cross section of a low cost microstrip MicroModule interconnection.

Another opportunity to lower system cost is to consider the tradeoffs of partitioning a single large (expensive) chip into smaller, higher yielding (less expensive) chips. The performance capability of the MicroModule substrate is high enough that this can be done not only without sacrificing performance, but also with potential overall system performance improvements. This can lead to greater design flexibility in that silicon performance, for example, can be optimized separately for memory if logic no longer has to be part of the same IC design. More memory could be added to the system since the constraint of the size of single yieldable chip and the size of the stepper field have been removed.

Given the need to accommodate a variety of possible die bonding formats from multiple IC vendors, it is advantageous to provide for maximum flexibility in the assembly process. Although the VAX-9000 MCU assembly was restricted to TAB, the MicroModule process currently supports both TAB and wire bond. In the future, flip chip solder connections will also be supported at the same time, on the same substrate. This flexibility (wire bond, TAB or flip chip capability) minimizes the cost of tooling and shortens the time-to-market.

A key enabler for the wide spread acceptance of MCMs is the availability of so called "known good die." The die requirements for MCM modules are discussed in more detail in Section 13.3. For the MicroModule test flow, known good die are obtained and assembled onto known good substrates.

After assembly, it is necessary to screen out assembly induced defects prior to functional test in order to simplify the diagnostic process on failed units. A suite of *in situ* tests are usually performed; these are tailored to the specifics of the MCM under assembly. For example, open wire bonds or solder joints can be found with visual inspection, boundary scan or other methods. In addition, die attach quality can be checked with ultrasonic, X-ray or thermal imaging. Once the assembly induced defects have been eliminated, the module can be functionally tested with minimal time spent on diagnosis and repair.

It is expected that, in some cases, MicroModules will require burn-in and a second functional test. Burn-in must be carried out prior to lid attach in order to allow for repair, since the lid is not removable in most low cost nonhermetic packaging schemes. After lid attach, the module is given a quick functional test prior to shipment. The issues of module test yield, fault isolation and repair are fundamental to the broad acceptance of MCM technology. This holds true not only for the MCM-D approach outlined, but also for MCM-C and MCM-L technologies. The costs of MCM are dominated by the silicon cost. Therefore the expense of scrapping good silicon due to an unrepairable or unisolated fault is severe.

Acknowledgments

The authors wish to thank the management team of Digital's MicroModules Systems Business Unit for their encouragement and help in our writing of this chapter. We are particularly grateful for the support of Chung Ho, Sharon McAfee-Hunter, James McElroy and Frank Swiatowiec. The authors also thank Jeffrey Slaney for his untiring efforts in researching the material used in this chapter.

References

1 D. E. Marshall, J. B. McElroy Sr., "VAX-9000 Packaging—The Multi Chip Unit," *COMPCON Spring '90: 35th IEEE Comp. Soc.*, (Los Alamitos CA), 1990.

2 P. B. Dunbeck, *et al.*, "HDSC and Multichip Unit Design and Manufacture," *Digital Technical Journal*, vol. 2, no. 4, Fall 1990.

3 T. Fossum, D. B. Fite, Jr., "Designing a VAX for High Performance", *COMPCON Spring 90: 35th Comp. Soc.*, (Los Alamitos CA), 1990.

4 S. G. Baust, R. J. Dischler, S. Westbrook, "Implementing a Packaging Strategy for High-Performance Computers", *High Performance Systems*, pp. 28-31, Jan. 1990, .

5 U. Deshpande, S. Shamouilian, G. Howell, "High Density Interconnect Technology for the VAX-9000 System," *Proc. IEPS Conf.*, (Marlborough MA), Sept. 1990.

6 J. P. Joseph, M. A. Kniffin, "Designing TAB Interconnect for the VAX 9000 Computer," *Proc. IEPS Conf.*, (Marlborough MA), pp. 683-693, Sept. 1990.

7 H. Hashemi, P. Sandborn, "Design and Analysis for a Reduced Parasitic Power Distribution System," unpublished proceedings of the *IEEE VLSI & GaAs Chip Packaging Workshop*, (Westford MA), Sept. 1990.

8 F. C. Chong, *et al.*, "High Density Multichip Module," *WESCON '85*, (San Francisco CA), Nov. 1985.

9 J. M. McPhee, T. S. O'Toole, M. Yedvabny, "Cooling the VAX 9000," *Proc. ELECTRO '90 Elect. Conf.*, (Boston MA), May 1990.

10 S. Heng, J Pei, "Air Impingement Cooled Pin Fin Heatsink for Multi-Chip Unit," *NEPCON West '91 Conf.*, (Anaheim CA), pp. 1155-1166, Feb. 1991.

11 J. Pei, S. F. Heng, "Cooling Components Used in the VAX 9000 Family of Computers," *Proc. IEPS Conf.*, (Marlborough MA), pp. 587-601, Sept. 1990.

12 J. S. Fitch, "A One-Dimensional Thermal Model for the VAX 9000 Multi Chip Units," *Proc. ASME Annual Mtg.*, (Dallas, TX), pp. 64-72, Nov. 1990.

13 C. L. Lassen, "Integrating Multichip Modules into Electronic Equipment," *Proc. IEPS Conf.*, (Marlboro, MA), pp. 3-15, Sept. 1990.

Part D—Closing the Loop

"You're thinking about something my dear,
and that makes you forget to talk.
I can't tell you now what the ***moral*** *of that is,*
but I shall remember it in a bit."

"Perhaps it hasn't one," Alice ventured to remark.

"Tut, tut child!" said the Duchess.
"Everything's got a moral,
if only you can find it."

Alice in Wonderland
by Lewis Carroll

Over the last 17 chapters we have travelled far and wide in the alternative worlds of packaging! At the start of the book we established a framework, or basis, for understanding the material that followed. This included definitions of the relevant terms and technologies. The main goal in **Part A—Making Decisions: The Big Picture** was to provide a *global picture* of the issues involved in package selection: packaging technologies, materials, manufacturing, systems performance and costs. The chapters in **Part B—The Basics** provided an understanding of the principal technology elements that make up a multichip module and the design sciences that are specific to multichip modules. The **Case Studies** in **Part C** showed examples of how some companies develop MCM products. Now it is time to **Close the Loop,** to tie specific new understanding back to the global picture described in the framework.

The following chapter serves two purposes.

First, it provides ***emphasis***. In **Part B—The Basics**, the intent was to provide and describe the set of fundamental technologies

needed for MCM package selection. Here the goal is to emphasize which of those technologies are the most important. This emphasis will be mainly given from a systems perspective.

Secondly, it provides a ***forward view***. It describes what are the open challenges that need to be solved if future systems are not to be *overconstrained* in performance and cost, due to limitations in the packaging technologies.

These purposes are accomplished through a review of the complementing technologies influential in the package selection process for small and large systems.

18

COMPLEMENTING TECHNOLOGIES FOR MCM SUCCESS

Ronald W. Gedney and
Donald P. Seraphim

18.1 INTRODUCTION

In previous discussions in this book, a number of key issues emerge critical to the future success of MCM-based products and to the future growth of packaging technology in general. By success, we mean volume manufacturing of MCM-based products by a range of both large and small companies. By future growth, we mean those technology elements whose further development are critical if packaging technology growth is to match the systems requirements and silicon chip technology density and performance improvements. Often success and future growth issues overlap.

A major conclusion that the reader has derived by now is that MCMs are generally more expensive than the same die packaged in single chip packages. Thus, the use of MCMs must be justified [1]-[3] either by cost savings at higher levels of the packaging hierarchy (such as printed wiring boards (PWBs), connectors, etc.) or by performance gains which result in system level cost/performance improvements. Both of these tie into several of the issues to be discussed here.

Substantial differences in the complementing technologies used to assemble large and small systems are influential in the package selection process. The relevance of these differences to package selection is summarized in the

following list and then are considered in more detail in corresponding sections of this chapter.

- **Separable connectors and the packaging hierarchy**. The connectors between different levels of packaging are the single most important constraint on system wide interconnect capacity.

- **The effect of semiconductor type on package design**. Until recently, most MCMs used bipolar parts. The growing use of CMOS devices creates some unique challenges to success.

- **System performance level (cycle time)**. This section reinforces the main advantage gained by using MCMs.

- **Level of chip integration, function size, and chip I/O**. These are the main chip related factors that force the need for high interconnection capacities.

- **Chip joining interface, burn-in, testability and repair**. This is one of the two reasons why using MCMs increases cost. (The other is substrate cost, covered in Chapter 4 and not to be discussed here.) It is a major limitation to success and a major opportunity for future growth.

- **Power density and cooling methods**. MCMs increase power density creating a unique challenge for their use. More care must be taken with cooling on the MCM than is usual in the single chip package.

- **Timeliness of introduction**. It is important that the use of MCMs does not increase time-to-market as discussed in Chapter 3.

Finally, we discuss the advances and limitations of single chip alternatives and the importance of considering the low chip count MCM alternative before concluding.

18.2 SEPARABLE CONNECTORS AND THE PACKAGING HIERARCHY

The previous chapters have shown that performance in the largest systems depends on the footprint and the areal density of the die and circuit functions, as

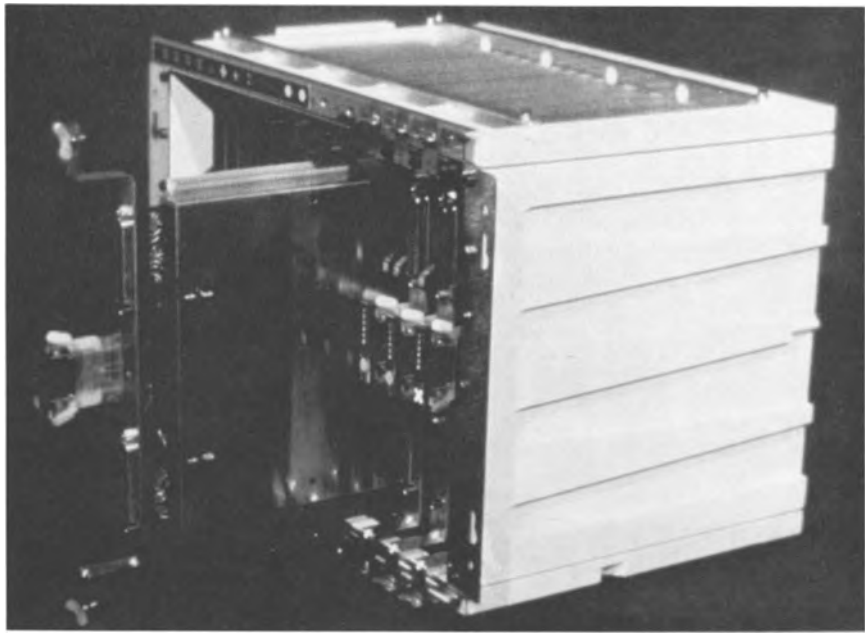

Figure 18-1 A three dimensional packaging system. A zero insertion force separable connector system along top and bottom edges of the PWBs supports connection of 500 contacts per PWB to two back panel PWBs. Tolerances are in the range of 0.2 mm.

well as on the ability to contain all of those functions within a cycle time. The opportunity does exist to use a variety of packaging architectures to achieve these capabilities. Until the late 70s the dominant approach was to use a three dimensional-type architecture with PWBs plugged edgewise into printed wiring back panels with separable contact arrays (see Figure 18-1). The back panels could interconnect 20 or more PWBs. For larger systems containing hundreds of boards, back panels were interconnected with cables into system configurations. Essentially all packaging of semiconductor main memory still is achieved in this way at extremely low circuit costs and continually improving performance [4]. There appeared to be no real boundary for three dimensional packaging density except for one, all important, limiting factor. This factor, separable contact density, was the key that locked the door on the use of three dimensional packaging for logic in the highest performance systems. Separable contact density is, however, more than adequate for low I/O connection demand of memory and acceptable for small systems and workstations.

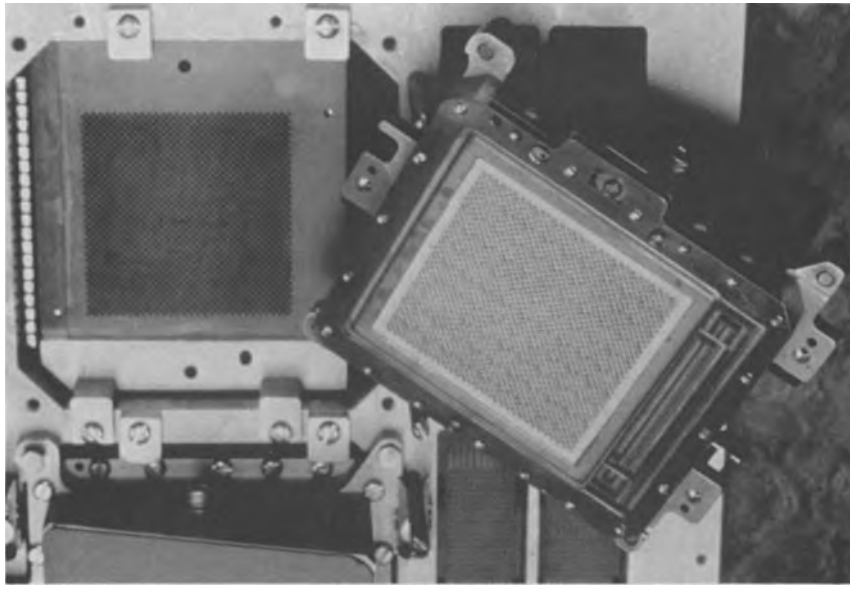

Figure 18-2 An area array separable connection system for MCMs. This system is extendable from the 2772 I/O shown here to the range of 8000 I/O.

One significance of the MCM lies in its areal array contacting capability (Figure 18-2) that has grown from 400 I/O (signal and voltage) contacts in the late 70s [3] to the range of 8,000 contacts today [5]-[6]. Even today, there is no qualified competitive approach for contacting large arrays of 2,000 or more separable contacts on the edge of PWBs for back panel connections.

For low end systems (PCs, workstations, servers), and even intermediate systems, a steady but slow stream of innovations in making conventional mechanical pin-in-spring separable contacts allows the continual growth of functional I/O capability of PWBs (Figure 18-3). One approach [7] is to use two opposite edges of the board as shown in Figure 18-1. This separable contact system supplies some 500 contacts. A precision frame with tight tolerances is required with a design to guarantee contact forces in the range of 100 grams per contact, and with a wiping action of the contact of several thousandths of an inch. This is a very reliable system. However, the cost per contact is higher

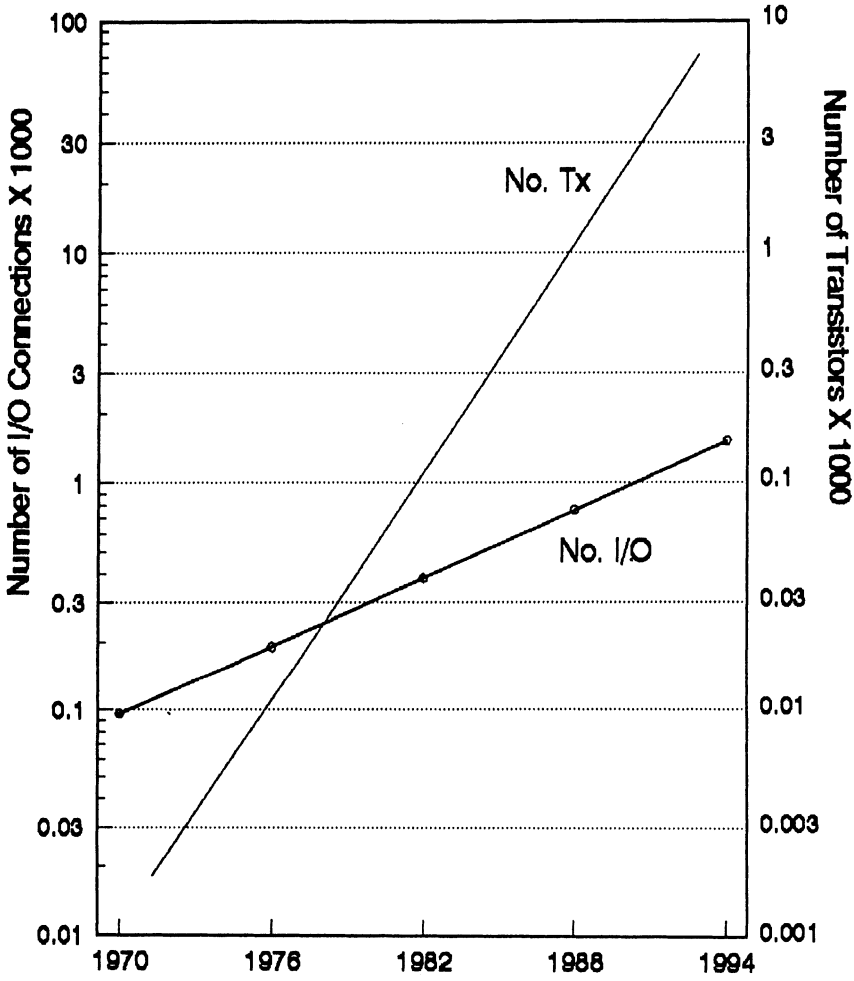

Figure 18-3 Projection and history of semiconductor integration and separable contacts on the edge of PWBs.

than desired by low end system managers. Precision frames are too expensive unless they are amortized over a large number of PWB connectors as applied in intermediate range and larger systems.

The largest board edge (three dimensional) connector systems available, utilizing one edge without precision frame contacts, are now in the range of 800

to 1,400 contacts. One method of extending the contact availability is to extend the edge of the board to an unwieldy, 10 - 15", dimension. Another method is a flex printed circuit to turn the corner into the connector or pin housing from the board edge (see Figure 18-4). This allows sufficient density to place multiple rows of springs beyond the two rows normally used on each side of the PWB. This transforms an edge array on the PWB to an area array at the next package level (PWB backpanel). An MCM pin grid array design accomplished this easily.

An alternative to the flex connection system is elastomeric connections [8]. Either of these technologies are qualified in large arrays. However, there is considerable development activity and it is possible that these two technologies will be developed sufficiently for more general use in the next three to five years. One concern is that the fundamental science base for contacts of these types has not been generated, even though substantial promising empirical data exists. It is essential to provide this science base in the next few years. Clearly, developments in this design space are strategic if PWB alternatives are to be competitive with MCMs and if lower cost connection techniques for MCMs are to be derived from these technologies.

The limitations of the board edge connection systems for application in low end systems are evident in the above discussion. First, the connectors are unwieldy in size requiring board sizes much larger than desired for the component functions and possibly much larger than desired for desktop or portable applications. Second, the conventional mechanical contact systems are reliable only at large forces per contact (60 - 100 grams) and generally with guaranteed wipe. These design constraints impose the use of precision housing and frame tolerances expensive to assemble.

Nevertheless, systems applications growth continues to double in function approximately every three years (see Figure 18-3). The relationship between the growth of circuit functions and increase in I/O is Rent-like, increasing I/Os proportionally to the square root of the number of circuits. This leads to the doubling of I/O and bus widths approximately every six years in the application of microprocessors, as well as in the high end MCM applications [9].

At this juncture, we look for either a new connection technology for PWB edges or for the extended use of MCM arrays and array contact systems to allow functional growth in low end processors.

We conclude that low end systems, with the persistent growth of functions in logic and their persistent pervasive application into all systems use territories, are reaching the same limiting factors faced by high end systems in the late 70s. The question is whether or not the cost learning curves of the high end packaging technologies have been strong enough to penetrate into the low end market. Furthermore, major innovations in the packaging technology have

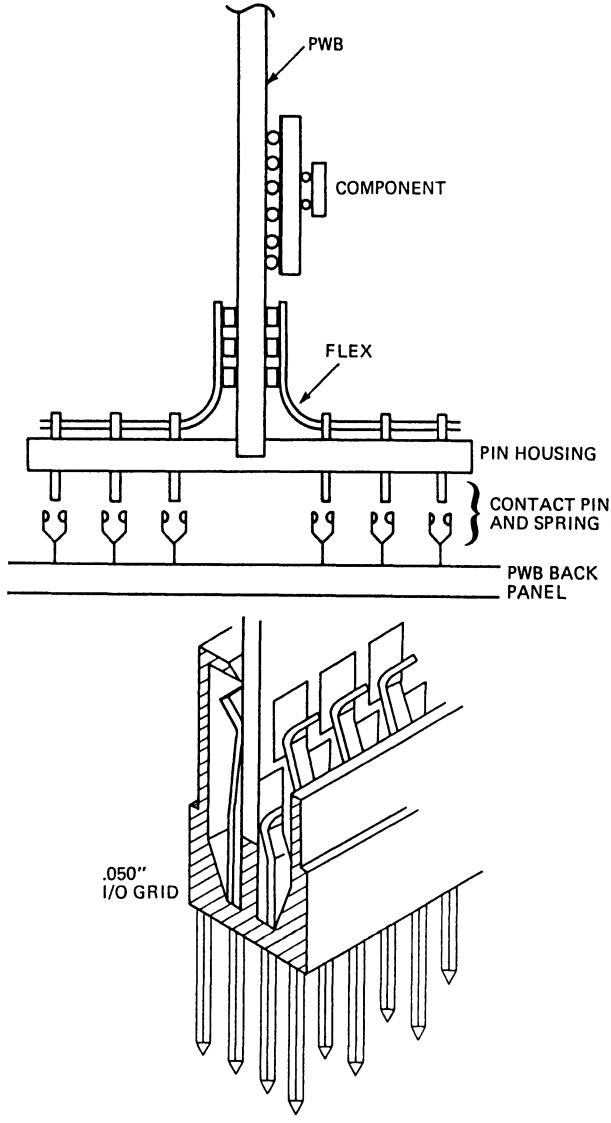

Figure 18-4 Interconnection techniques from PWB edges to PWB back panels. Flexible circuitry is used to extend the number of rows of contacts.

occurred and are developing that may allow alternatives to be applied at the low end. A discussion of the influential technology differences is continued in the sections to follow.

18.3 EFFECT OF SEMICONDUCTOR TYPE ON PACKAGE DESIGN

The discussion on I/O capability should not distract attention from the significance of packaging the die close together. This is an attribute of the MCM. It is desirable to purchase components from a variety of sources at the lowest possible costs. In small systems, there is a tendency for the PWB to contain a variety of components, such as surface mount plastic leaded chip carriers (PLCCs), surface mount small outline package (SOPs), memory boards with edge contacts, dual in line package (DIPs) as well as pin grid arrays (PGAs) and discrete componentry all on the same board surface (see Figure 18-5). At the same time, component sizes are decreasing steadily while device functions are increasing. Capabilities exist to build the processors and the complement of adapters on a single PWB (see Figure 18-6). MCMs, if used at all, are used

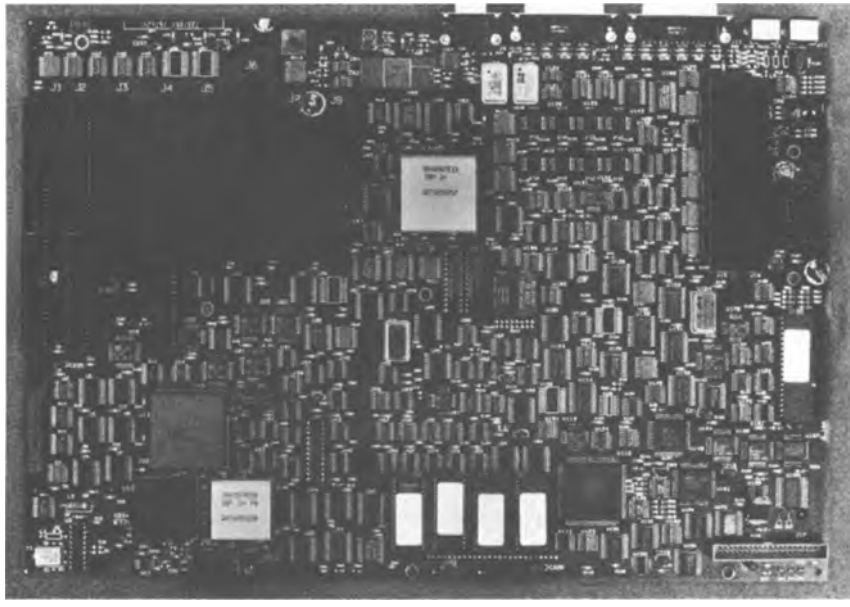

Figure 18-5 A typical component menu for a Personal System PWB.

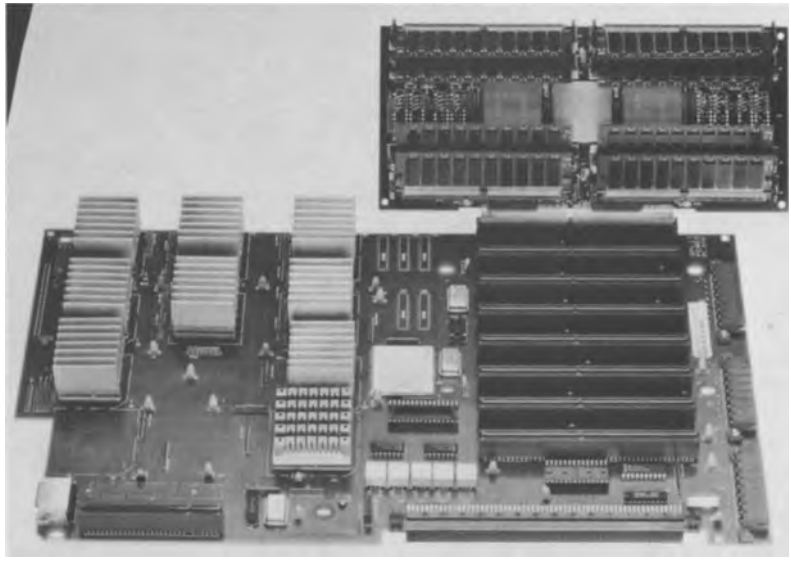

Figure 18-6 A workstation processor board with three dimensional memory board and large SCMs.

sparingly to increase board performance capabilities as a "complement" to the single chip capabilities and to shrink the size of the chip function, allowing it to be packaged on one board.

A substantially larger system with multiple processor logic devices and adapters is more likely to use an MCM. The cost savings, if any, might be minimal. However the initiative to use MCMs in a system of this type is for a cost/performance increase, or to avoid the use of a second PWB and its connections. Another reason to choose MCMs is that printed wiring density, optimized to bring adapters close together for performance increase, can cause the PWB to be very expensive. An MCM allows the use of a much lower cost PWB. The combination of the MCM with a lower cost board meets permissive system costs while providing an important increase in performance.

Logic components used in the low end systems have been primarily CMOS. Substantive improvements have been made in these semiconductors in recent years. However, the very rapid growth of integration, shrink of feature size,

decreasing thickness of gate oxide and other factors tend to widen the parametric distribution of CMOS devices. This means that some CMOS parts need to be burned-in and sorted for speed. Once the infant mortality (specification) failures are removed through burn-in, an excellent MCM yield and good thermal and environmental-related reliability can be expected. However, the need for burn-in of logic devices is still a topic of debate, corresponding to the difficulty of accomplishing burn-in on bare die.

The higher level of integration of CMOS in comparison to bipolar requires an order of magnitude fewer chips to make an equivalent function. This is a very positive factor in MCM assembly yields and reduces the number of chips needing to be burned-in to 10 chips per MCM.

It took IBM almost 10 years to develop the manufacturing capability for custom bipolar MCMs. A particular design for a low end multichip processor would require 24 to 36 months to qualify. This is assuming that an appropriate strategy and capability is available for test and burn-in of unpackaged chips, and that the manufacturing capability is in place.

A feature of single chip modules (such as TABs, PGAs, PLCCs, etc.) is the ability to test the function, including its performance. If the chip is faulty, the entire component can be scrapped without concern for its neighbors. For MCM applications, repair of one chip interacts with the other devices and the substrate on which they are mounted. Semiconductor manufacturers must provide good bare die to MCM assemblers as the probability is high that the silicon requirements encompass the product line of more than one manufacturer (see Chapter 1).

18.4 SYSTEM PERFORMANCE LEVEL (CYCLE TIME)

The decision to use an MCM usually is made by a system manager able to assess various alternatives at the complete box level. Since it is very difficult for any MCM to be cheaper than the sum of the single chip carriers it replaces, cost reductions resulting from the use of MCMs are apparent only when the entire system is considered. The system manager also can assess the value of increased performance gained by using a MCM.

Davidson showed that an MCM had the potential for putting 16 or more devices within a 1 ns cycle [10]. If the system cycle time is 20 ns or more, a savings of 1 ns is 5% or less of the cycle; other solutions to meet performance objectives can probably be found. However, when system cycle times are 10 ns or less (100 MHz), a 10% or better improvement can be garnered from packaging alone. Shiao and Nguyen showed that a high performance cache system could see cycle time improvements of up to 25% through use of

Cu/polyimide MCM modules versus single die on board packages when both package and I/O buffer delays were taken into account [11]. Neugebauer *et al.* showed that for a 4.36 MGate array processor system, moving from SCMs to simple 10 chip MCMs would result in a total package footprint reduction of 69%. If 72% of the interconnections are moved from the PWB to the MCM, the power multiplied by the delay product in the average signal path is diminshed by over 75% [12].

Another example of the space and performance advantages of MCMs is the IBM Risc/6000™ system. With slight advances in silicon technology, nine complex single chip modules were condensed into a seven chip MCM, providing a 12× reduction in physical size on the planar performance and a 2× improvement in system performance (see Figure 18-7). IBM also found that the MCM provided a cost reduction at the system level. Krusius *et al.* developed

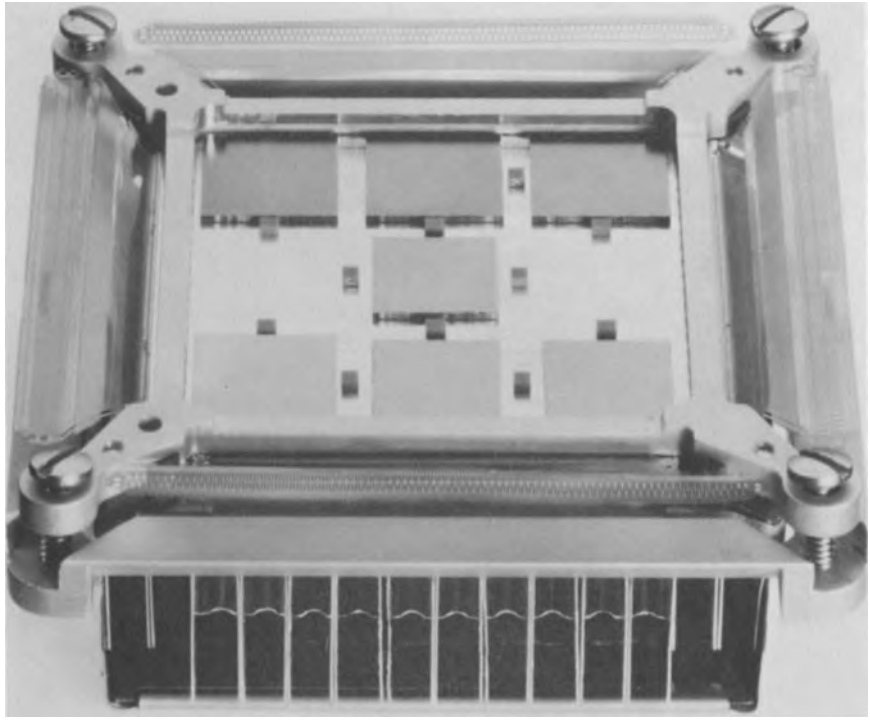

Figure 18-7 A MCM (SOS) replacement for the processor and cache (nine large single chip modules) shown in Figure 18-6.

computer models for evaluating package performance [13]. Various MCM alternatives to single chip on PWB designs have been implemented and substantial performance gains for MCMs have been demonstrated.

It is reasonable to assume that most systems with performances greater than 50 MHz can see measurable cycle time reductions through use of MCMs, particularly those of the thin film, Cu/polyimide design. While clever use of high performance single chip carriers, ordered wiring [14] and high performance surface mount PWBs can achieve the desired system performance, well designed MCMs may result in better cost/performance tradeoffs at the system level.

18.5 LEVEL OF CHIP INTEGRATION, FUNCTION SIZE AND CHIP I/O

The number of CMOS chips on an MCM has to be small and the quality high to reduce failures on MCMs during burn-in. Otherwise the test, rework and scrap costs charged to the MCM will be excessive. Furthermore, a degree of sophistication in the user's system of assembly and test is required beyond that necessary for single chip components.

18.5.1 Memory MCMs

An illustration of the impact of test, scrap and rework cost is IBMs original work with memory MCMs. During the 70s, IBM consistently used multichip memory modules. These modules, on a 24 mm substrate, would have four DRAM die. Two substrates could be stacked (pins joined to pinheads with solder) to provide eight die per module. With the introduction of the 1 Mb DRAM, IBM went back to a single die module.

Referring to the single chip carrier case in Table 18-1, functional test can represent as much as 30% of the total package cost. Given the cost of a single chip small outline integrated circuit (SOIC), it is not cost effective to rework such a package. Typically, one can expect a complex (not burned-in) CMOS device (such as a 1 Mb DRAM) to have a chip join to final test yield of 90% to 94%. Assuming a median yield of 92%, approximately 28% of the four chip modules require repair. If this were a true eight chip substrate, the first pass yield would be only 51%, but since the individual four chip decks are tested good before stacking, a very high yield is achieved on the stacked substrates. Usually there is only one bad die on a substrate, but occasionally there are two or more. In a production environment, economics would dictate that more than two bad die would result in scrapping the entire substrate.

Table 18-1 Module Process Steps.

SCM	MCM	
1 Die	4 Die	8 Die
1) chip place and join 2) encapsulate 3) burn-in (optional) 4) functional test 5) ship	1) chip place and join 2) functional test 3) repair and rework 4) functional test 5) encapsulate 6) functional test 7) ship	1) chip place and join 2) functional test 3) repair and rework 4) functional test 5) stack substrates 6) functional test 7) repair and rework 8) functonal test 9) encapsulate 10) functional test 11) ship

Table 18-2 Cost Considerations for MCMs and SCMs.

Assume:	SCM cost (SOIC) = $X
	Unreworked MCM cost = $X/chip
	(MCM = SCM cost x Number of SCMs
	Rework + functional test = $2X
	(28% of the modules are reworked.)
	Cost of additional test = $0.3X
Then:	Added module cost over production lot is 0.28 x 2X = 0.56X
	Added cost per packaged die = (0.3X + 0.56X)/4 = 0.22X

Independent of the rework, which is time consuming and expensive, consider the cost of testing. In the case of the four chip module, two complete functional tests on 100% good parts and three on repaired parts have been completed. Keep in mind that, as each die is individually tested, the four chip test is four times as long as the test for a single chip package. (Handling times are usually insignificant compared to test times.) The cost of one of the more expensive parameters in building a module has been doubled. This, coupled with the additional cost of rework and repair (and another test), on 28% of the product means a premium of 22% per chip for the MCM over the SCM cost (see Table 18-2).

This added cost per packaged die might be affordable if cost savings could be made elsewhere in the system (smaller or less dense PWB). However, in the case of memory, the package is almost the same size as the die. With double sided surface mount and stacked SOICs, these savings are not available. This is an example of multichip packaging that did not have any payback at the system level.

We project from this discussion that the number of CMOS chips on an MCM needs to be relatively small and the quality needs to be very high to reduce, to an inconsequential few percent, failures on the MCMs during burn-in. Otherwise the test, rework and scrap costs charged to the MCM will be excessive. To achieve this level of confidence, semiconductor suppliers need to offer tested, burned-in bare die for MCM assemblers.

18.6 CHIP JOINING INTERFACE, BURN-IN, TESTABILITY AND REPAIR

For MCMs to be practical, test and rework processes have to be developed (see Chapter 13). This calls for a chip removal process, a site preparation process where the chip is replaced, a joining process that does not risk the neighboring chips, precision placement capability [15] and a functional test capability with analytical ability built into the design [16] to isolate the fault to a particular chip. All of these factors must be considered during the design stage.

18.6.1 Known Good Die

The next logical question to ask is, "Why not burn-in and test the chip prior to module mounting?" If producers of MCMs could work with known good die, then yields would be far superior and rework and repair minimized, if not eliminated. The problem, of course, is handling, testing and burn-in of devices individually or at the wafer level. Test probes to provide functional test of individual die or at the wafer level are both complex and expensive as discussed in Chapter 13. Connections to bare die for burn-in are also quite difficult, although under active investigation. Our experience to date is that test and burn-in at the die or wafer level leads to device costs equal to or higher than the same device in a single chip carrier. Although we believe this will change in the future as new manufacturing concepts and techniques are devised, there are always some added costs. Given that simple packages cost $0.05 to $0.15 per I/O, it is doubtful that known good die will ever allow MCMs to be cost equivalent at the chip carrier level. The exception to this is the case where, early in a program, a die is committed to an expensive, high I/O, hermetic, multilayer ceramic package and the cost of the MCM interconnection is amortized over a large number of chips.

18.6.2 Chip Joining Interface

Still another MCM design consideration is the physical interface to the semiconductor user. This has been predominantly wire bond technology. Unfortunately, wire bonding is a difficult technology for MCM production. One detriment is the ability to remove a bad die, clean the bond sites and put another die in its place as discussed in Chapter 9, Section 9.3. Furthermore, the inductance and capacitance of the wires are unfavorable (for high performance applications) compared to some of the alternative methods.

We believe that either an assembly process like TAB [17], flip TAB or a process like flip chip (C4), not commonly used by semiconductor manufacturers, has to be introduced as an industry standard for the semiconductor producer and the user of MCMs.

18.6.3 Flip TAB

The flip TAB technology [18] appeared first in the high end systems of NEC [5] where it was used to join chips to a very small multilayer single chip carrier (see discussion in Chapter 9, Section 4). The carrier was interconnected to an MCM with an area array of gold tin bumps. The area array of bumps of the single chip component was on a grid similar to that of the wiring interconnection vias of the MCM. This eliminated the need for fanout interconnection layers, simplifying the MCM design and moving the finest line interconnections to a component package to be tested individually and scrapped, if necessary. This decreased the probability that the MCM had to be reworked after chips were placed and joined because of the simpler, less dense interface and an adequate pretest of the units joined. An advantage of the NEC approach is the versatility in achieving a dense area array grid joined to the surface of the MCM. In fact, the NEC approach is very similar to placing a die directly on a PWB, but it does not eliminate a level of packaging. As pointed out earlier, the dense fanout layers for direct chip joining on an MCM can be equivalent to the fanout interconnections on the single chip component level of packaging.

The single chip carrier may be eliminated in the future by using the flip TAB interconnect directly to the MCM or a printed circuit card. A flip TAB array of this type, recently developed by IBM (see Figure 18-8), could be used to mount chips directly on ceramic or organic packages, eliminating one assembly operation.

18.6.4 Flip Chip

As semiconductor densities increase, the number of I/O, even for low end devices, soon will exceed 600 and push the 1,000 to 1,200 range. Flip TAB on large (> 14 mm) die will handle 600 I/O and might extend to the 800 I/O range.

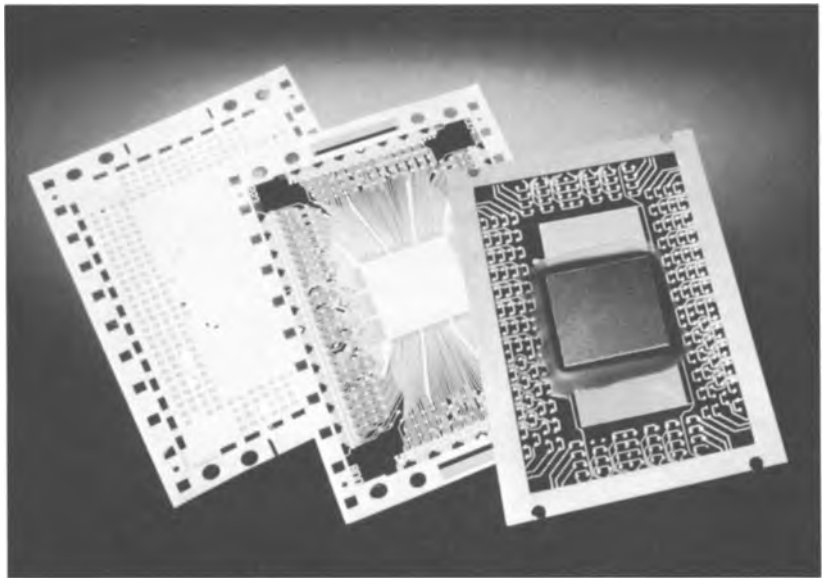

Figure 18-8 A solder ball array TAB. The ATAB has two metal layers (ground and signal).

After that, flip chip technology (see Chapter 9) seems to be the only fully qualified solution.

Putting solder bumps on the die necessitates additional wafer processing. However, since C4 joints can be placed over active areas (unless there is alpha particle sensitivity), the die does not need a pad for wire bond or TAB mounting outside the active silicon area. Thus, die that are 10% to 30% smaller can be designed with comparable increases in wafer productivity, depending on whether the semiconductor design is circuit or I/O limited. Also, if the flip die bond is designed properly, the via opening in the chip passivation layer is covered and the die is effectively sealed, eliminating the need for hermetic packages.

18.7 POWER DENSITY AND COOLING METHODS

Functions on MCMs for mid-range range systems using bipolar die have run at about 90 watts (Figure 18-9) requiring impingement air cooling on very large heatsink areas [19] and thermally conducting grease between the heatsink and the back of the chip. Power density is about 20 watts per square inch, with air flow up to 1,000 feet per minute. This is not a low end technology. With the advent

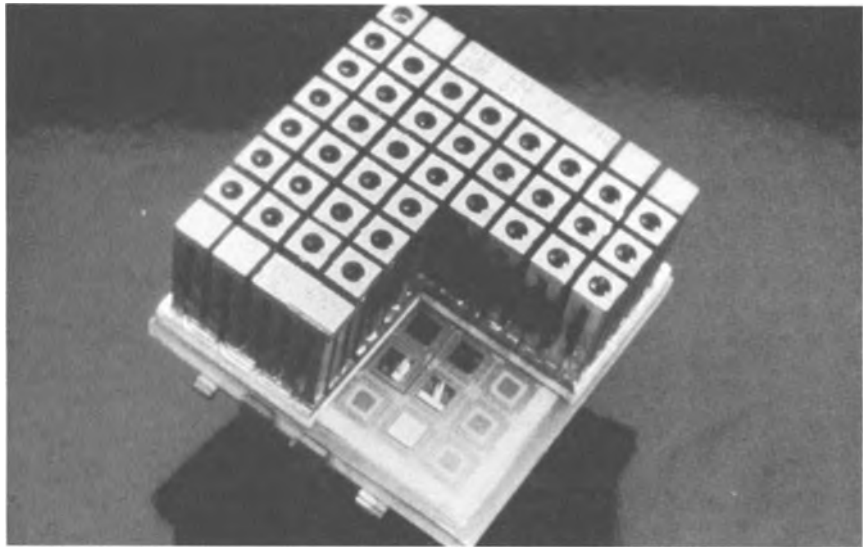

Figure 18-9 Impingement air cooled 64 mm MCM (about 90 watts).

of CMOS processors, which run in the range of 3 - 10 watts per chip, a six to nine chip MCM that dissipates about 30 watts is reasonable. This can be cooled with a small fan providing up to 500 feet per minute in a small system environment. For higher powers, it is essential to find efficient methods for heat removal, while decreasing the power levels per circuit. In this case, the cooling cap provides a hermetic or semihermetic seal. This seal eliminates concerns for ionic migration with the high density of interconnection and the resulting high electric field on the top surface of the MCM. CMOS technology is helping by going to lower voltages with resultant lower powers. Devices currently are being announced with voltages of 3.0 - 3.5 volts and expectations that devices with 2.5 volts or less are coming.

Aluminum nitride (AlN) ceramic substrates can be used in future designs since its heat spreading capability is nearly an order of magnitude better than that of alumina. A single C4 joint has a thermal resistance of about 550°C per watt. For high I/O chips, a reasonably good thermal path is provided between the chip and the substrate. (Extra C4 pads can easily be added to assist thermal management.) AlN allows efficient spread of heat through the ceramic without needing a complex process and design.

In the near future (three to five years), we expect 20 - 40 watt CMOS chips still will be used. In this case, the cooling technology will be complex, similar

to that discussed for bipolar die and MCMs for low end systems will have powers ranging up to 100 watts.

18.8 TIMELINESS OF INTRODUCTION

Assuming that the complement of assembly processes, interface designs and system cooling configuration capability exists, MCMs can be designed and built with a turnaround time equivalent to the semiconductor. The prerequisites to the achievement of this goal are:

- Freezing device pad assignments allowing module design to start early (concurrent engineering)
- Good information and design computer aided design systems
- Rapid prototype facilities
- Good simulation systems for electrical and thermal modeling

These techniques all exist today, able to meet requirements for rapid design iterations and performance increases expected in small systems. Following the semiconductor evolution, it is reasonable to expect small systems to double in performance every three years.

Other alternatives should be considered. One approach is for a consortia of companies to develop functional building blocks utilized by its members. Another alternative is for semiconductor houses to develop and sell functional building blocks. Systems houses currently buy the processor device from one supplier and complementary devices and memory from other suppliers. If a semiconductor house supplied MCMs with the whole function, they could supply more of their devices to the system purchaser.

18.9 ALTERNATIVES TO MCMs

Consider a system consisting of nine large chips each with 600 I/Os (400 signal I/Os, 200 power and ground) connected to a large memory array and some other low density parts. First, consider how the system can be partitioned by interconnection needs. Unless the system performance is particularly critical to small improvements in delay to the memory array chips, there is little benefit gained by packaging the memory chips in an MCM. The memory and low density part of the system is slightly larger, slightly slower and, certainly, much more cheaply packaged with plastic single chip packages on a four layer (two signal, power and ground) board.

On the other hand, if the nine chips in single chip packages are spaced 2" apart, the required wiring capacity necessary is about 1,800" per component or 450" per square inch. This could be provided by a PWB with 9 - 10 signal layers, as long as the average line pitch was better than 1 line per 20 mils. This is a credible state of the art for leading PWB suppliers. The critical technology for this board would be providing through hole vias on a 40 -60 mil grid. This will always be a critical technology element for future high performance PWBs, or similar laminate MCMs, the other critical element being coefficient of thermal expansion (CTE) match. Progress in this technology was discussed in Chapters 2 and 5. Despite this, their cost is significantly cheaper than the equivalent MCM.

However, consider the edge connector of this board. If they are all to be on one edge of the PC board (6" in length), then the edge contact density is approximately $\left(\sqrt{9} \times 400\right) = 1200$ or 200 per inch, using a Rent's rule relationship. This is equivalent to 10 rows of contacts on a 50 mil grid or 5 rows on each side of the edge. A PWB edge connection system of this type has not yet been qualified. However, a peripheral leaded MCM flex demonstrates the contact technology and was released in the DEC 9000 computer series (see Chapter 17).

For comparison purposes, consider an equivalent nine chip MCM. With modest device spacing, the same nine chips could be mounted on a 2.5" square substrate. Adequate wireability [21] and I/O are achievable with either a multilayer ceramic thick film (MCM-C) or thin film (MCM-D) design with I/O on a staggered 0.100" grid (area array). Conventional spring connectors are available with this spacing. Also, existing designs [19] cool up to 90 watts in air for these size modules.

The single chip PWB approach results in a 6" × 6" (36 square inch) package and a formidable connector challenge. The MCM alternatives result in a 2.5" square package (6.25 square inch) with a readily achievable interconnect solution. One readily can discuss performance differences in the range of 15 - 20% based on the usual rule of thumb that performance is one-third chip, one-third circuit and one-third package delay.

The other issue is cost. The PWB has two advantages: it is low on the manufacturing learning curve, and many pieces can be cut from the large panel processed. The MCM, in comparison, can be six times as expensive per unit area at the substrate level.

Finally, examine the significance of improvements in the cost/performance ratio with the use of MCMs. For example, if mounting the chips on an MCM substrate doubles the performance and also doubles the cost compared with single chip mounting, the cost/performance ratio is unchanged at the mounting level. This may well be sufficient to sell additional applications. However, at

a system level, the cost/performance ratio may be improved significantly. Consider a case where the system level cost split is 25% to semiconductor devices, 50% to the package and 25% to assemble and test. Then, assuming that chip mounting is approximately 50% of the total package cost:

Cost/performance for single chips = [.25 + (.25+.25) + .25] / 1 = 1/1

Cost/performance for MCMs = [.25 + (.25+.50) + .25] / 2 = 1.25/2

Thus the improvement in the cost/performance ratio at the system level is 1/0.6125 = 1.63 which is substantial.

In conclusion, as performance continues to increase in low end systems, and, especially as multiprocessor applications increase, MCMs will capture a significant fraction of the electronic packaging market. Estimates by Sage and Hartley are in the range of 4% to 12% of an $11.5 million unit total market for just work stations, high end PCs and portable computers by 1995 [14].

18.9.1 Future Growth in Single Chip Packaging

PLCCs, with I/O counts greater than 300, are now available and sample quantities of 500+ I/O are offered. For our conceptual package we chose 2" spacings for the components. Even with 0.3 mm OLB spacings, the achievement of 600 peripheral leads on a flat pack will not be contained in the space allotted. Also, it is recognized in industry that peripheral lead packages with 0.5 mm or smaller spacings are difficult to join at high yield and reasonable cost to the printed circuit carrier. It is barely possible to achieve a 3σ joining yield. A goal of 6σ joining does not appear practical without new processes and paradigms being invented. On the other hand, surface mount array (SMA) packages, with area array I/O, can achieve reasonable sizes at 600 I/O (Figure 18-10) and, at a 0.050" grid, can be joined readily at the 6σ level with current processes. Both pin grid array and land (or pad) grid array ceramic packages are readily available for surface mounting (Figure 18-11).

Although advanced area array packages provide higher packaging density, they also drive higher PWB wiring capability (density of interconnection) [23], because the same number of wires have to be contained in the smaller area allowed by the array. However, the wires are shorter proportionally to the pitch or spacing of the components, if the component spacing on the PWB is decreased proportionally to the component size.

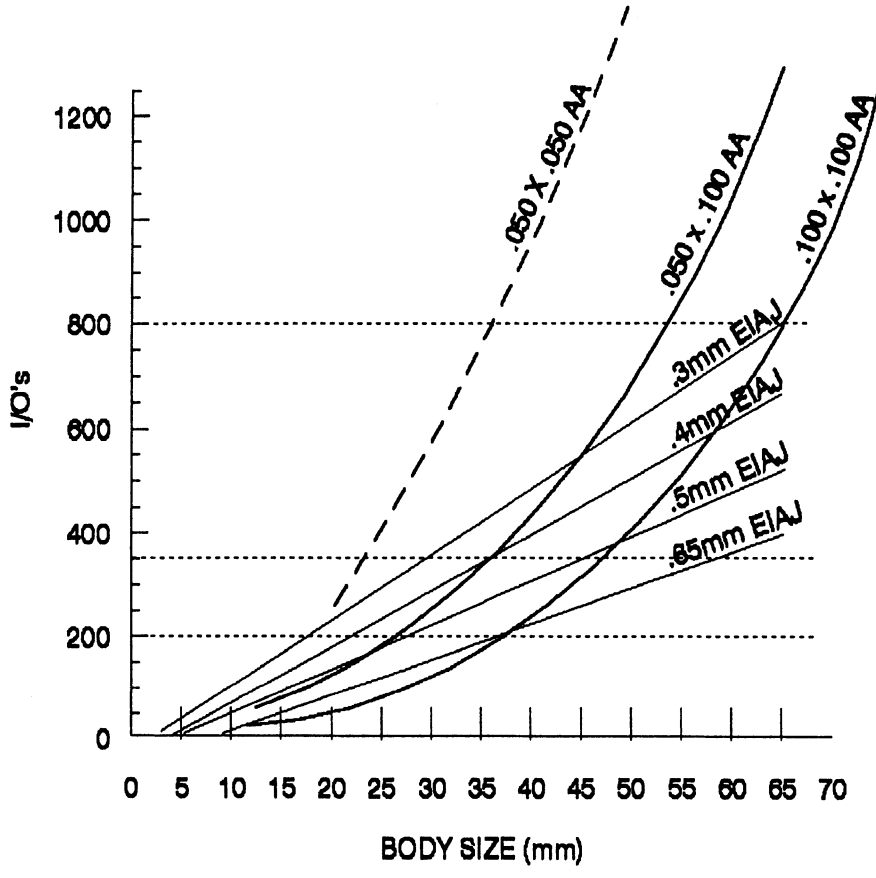

Figure 18-10 Comparison of I/O capability versus size for area array and peripheral lead packages.

PLCC packages increase in size for a given lead pitch as they provide higher and higher I/O. The same number of wires (as compared to area arrays) are contained in a larger area. The wires in this case are longer proportionally to the pitch of the components. If the larger areas are used effectively, some relief in wiring density and area for air cooling is available with the larger components. However, the PWB size is increased.

The system designer has several decisions to make. If the electronic package size is not particularly important (such as a desktop PC), low cost PLCCs can be spread out on a surface mount PWB. The wiring demand on the

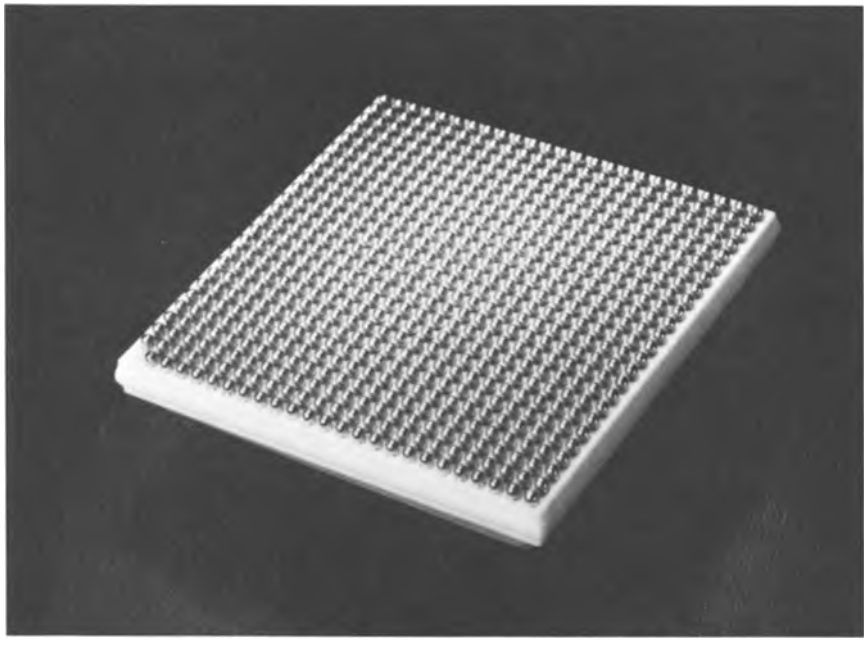

Figure 18-11 A surface mount array MLC MCM with 600+ I/O on 0.050" grid.

PWB is not very high; spreading the components out assists in thermal management. If size of the electronics is important, one alternative to an MCM is the SMA package. This very dense package reduces the PWB area required at the next level, at the expense of wiring demand. The wiring demand probably doubles from a 0.5 mm peripheral lead package to a 50 mil grid SMA package, but the required area is reduced substantially. A dense SMA package on a dense PWB should permit up to 100 MHz clock times without resorting to MCMs.

18.10 INEXPENSIVE, LOW COMPLEXITY MCMs

So far our discussion has dealt mainly with the system processor or central electronic complex. However, there is another class of MCMs, generally with only a few chips [22], that are successful. In this breed of MCM, there only may be one to three semiconductor devices and some passive components (resistors, capacitors, etc.). These modules are often associated with unique functions such as analog circuits, sense amplifiers, optical driver/receivers or semipower devices such as hammer drivers. The many passive devices accommodated on a PWB can be placed on an MCM as cheaply. In many cases,

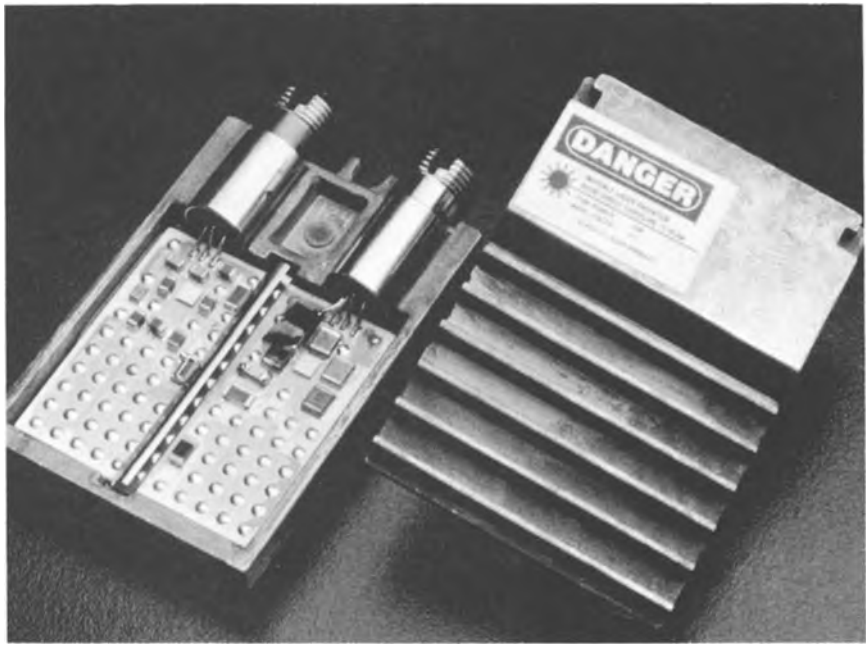

Figure 18-12 A 200 MHz optical driver/receiver MCM.

the circuit performance improvements of the MCM are considerable. In the case of very small signal amplifier circuits, the ability to actively trim resistors and tune the circuits for performance is vital to success of the function.

An optical driver/receiver MCM substrate produced by IBM is shown in Figure 18-12. This MCM design provides both driver and receiver circuits for a 200 Mb/sec optical data communications module and provides interconnection of the fiber optic cable in a simple 30 mm package. A more sophisticated module has been investigated by GTE [23]. This MCM (28 die plus assorted passive components) packaged many functions of a digital signal processor (DSP) board on one MCM. The MCM resulted in a 60% space saving and substantial performance gain.

It is conceivable that the MCM industry will grow, not from a "trickle down" of technology from the very high end, but be driven up by the growing demands for density and performance increases of single chip packaging technology [9].

The other thing that the MCM developer has to keep in mind is that silicon density keeps increasing with little sign of slowing growth. A function that requires an MCM today will, in all likelihood, be designed into a single chip

tomorrow. It is feasible that an MCM with three to ten chips will be integrated into a single die three years after the MCM goes into production. For the careful planner, this provides an opportunity. As a single die with the same function requires the same number of I/O, a functional MCM design can be migrated to a single chip design at a lower cost without disturbing the next level of package.

Trends indicate that high end PCs, work stations and large systems will embrace multiway designs to enhance performance in the near future. Assuming that a one chip processor can be achieved, a four way design will require four of these chips plus associated cache. An MCM design can handle the dense interconnections required now; a single die with four processors can replace the MCM when the technology becomes available.

18.11 SUMMARY

The decision to use MCMs in a system, whether small or large, is driven by the need for performance and competitive cost. MCMs have two main performance advantages: their ability to space chips closely, minimizing interconnect delays, and their ability to provide very large amounts of wiring, within the MCM, from the chip to the MCM and through the connector to the next level of packaging. Cost benefits will be achieved, and should be optimized, at the system level. However, consideration needs to be given to a number of potential factors that might inhibit MCM use. In particular, the need for burn-in and testing of complex CMOS systems needs to be considered, along with the higher heat density and the extra design challenge. One solution is to use small, less than 10 chip MCMs. This fortunately is becoming more and more feasible with the growth in semiconductor density.

Looking toward the future, a number of challenges are evident that require new solutions. First and foremost is the need for high I/O count, low insertion force, separable connectors for both PWBs and MCMs. Next, better solutions to the test and burn-in issues are needed if high chip count CMOS MCMs are to be a success. This is partly an infrastructure issue, as the chip manufacturers have to provide qualified bare die for this to happen. Finally, there is the challenge of designing a high speed, high power MCM with first pass success in the manufacturing environment.

We believe these challenges will be met and that the MCM technology will find widespread use in the electronics industry.

References

1 D. P. Seraphim, I. Feinberg, "Electronic Packaging Evolution in IBM," *IBM J. of Res. & Devel.*, Anniversary Issue, vol. 22, no. 5, p. 617, Sept. 1981.

2 R. Johnson, R. K. F. Teng, J. W. Balde, *Multichip Modules: System Advantages, Major Constructions, and Materials Technologies*, New York: IEEE Press, 1991:
J. W. Balde, "New Packaging Strategy to Reduce System costs," pp. 7-10
W. H. Knausenberger, L. W. Schaper "Interconnection Costs of Various Substrates—The Myth of Cheap Wire," pp. 11-13.
M. G. Sage, "Future of Multichip Modules in Electronics," pp. 14-20.
G. Messner, "Cost Density Analysis of Interconnections," pp. 21-29.

3 Albert J. Blodgett, Jr., "A Multilayer Ceramic Multichip Module," *IEEE CHMT Trans.*, vol CHMT-3, no. 4, pp. 634-637, Dec, 1980. Also *Scientific American*, vol. 249, pp. 86-96, July 1983.

4 L. E. Thomas, F. W. Haining, J. S. Macek, "Second Level Packaging for the IBM ES/9000 Processors", *Internat. Electr. Packaging Soc. Conf.*, (San Diego CA), pp. 1049-1055, Sept. 1991.

5 D. Akihiro, W. Toshihiko, N. Hideki, "Packaging Technology for the NEC SX-3/SX-X Supercomputer," *Proc. of the IEEE 40th Electr. Components and Tech. Conf.*, (Las Vegas, NV), pp. 522-533, May 1990.

6 A. Kaneko, K. Kuwabara, S. Kikuchi, T. Kano, "Hardware Technology for Fujitsu VP2000 Series," *Fujitsu Scientific and Tech. J.*, vol. 2, no. 27, pp. 158-168, June 1991.

7 T. Chikazawa, *et al.*, "Sensor Technology for ZIF Connectors," *ASME J. of Electr. Packaging*, vol. 112, pp. 187-191, Sept. 1990.

8 *Proc. NEPCON West*, (Anaheim CA), 1990.
L. S. Buchoff, "Solving High Density Electronic Problems with Elastomeric Connectors," pp. 3-7.
K. Gilleo, "A New Multilayer Circuit Process Based on Anisotropicity," pp. 8-31.
J. A. Fulton, *et al.*, "The Use of Anisotropically Conductive Polymer Composites for High Density Interconnections," pp. 32-46.
R. C. Moore, *et al.*, "Testing Bare Boards with an Anisotropically Conductive Elastomer," pp. 47-56.
H. Yonekura, "Oriented Wire Through Connectors for High Density Contacts," pp. 57-72.

9 "Computer Development and Parallel Processing," *BPA COMPASS Update*, no. 1, 1992.

10 E. Davidson, "The Performance Choice: A Study of Chip and Package Densities," *Proc. of the 40th Electr. Components and Tech. Conf.*, (Las Vegas NV), pp. 147-150, 1990.

11 J. Shaio, D. Nguyen, "Performance Modeling of a Cache System with Three Interconnect Technologies: Cyanate Ester PCB, Chip-on-Board and CuPI MCM," *Proc. of the 1992 IEEE Multichip Module Conf.*, (Santa Cruz CA), pp. 134-137, March 1992.

12 C. Neugebauer, R. D. Carlson, R. A. Fillion, T. R. Holler, "Multi-chip Module Designs for High Performance Applications," *Proc. NEPCON West*, (Anaheim CA), pp. 401-413, vol. 1., 1989.

13 D. E. O'Brien, B. M. Hahne, J. P. Krusius, "Multichip Modules vs. High-density Printed Wiring Board: A Trade-off Study," *Proc. of the IEEE Electr. Components and Tech. Conf.*, (San Diego CA), pp. 31-41, May 1992.

14 M. G. Sage, N. Hartley, "System Issues for Multichip Packaging," *Proc. of the Internat. Soc. Optical Engr.*, (Boston, MA), pp. 302-310, vol. 1390, 1991.

15 A. H. Landzberg, *Microelectronics Manufacturing Diagnostic Handbook*, New York: Van Nostrand Reinhold, 1992, Chapter 8.

16 D. S. Karpenske, "Interconnect Verification of Multichip Modules Using Boundary Scan," *Proc. IEEE 1991 VLSI Test Symp.*, (Atlantic City NJ), pp. 85-91, April 1991.

17 H. Wessely, "Packaging System for High Performance Computer," *Proc. IEEE IEMT Symp.*, (Nara Japan), pp. 85-91, April 1989.

18 J. H. Lau, *Handbook of Tape Automated Bonding*, New York: Van Nostrand Reinhold, 1992.

19 S. Oktay, B. Dessauer, J. I. Horvath, "New Internal and External Cooling Enhancements for the Air Cooled IBM 4381 Module," *IEEE Internat. Conf. on Comp. Design*, (Portchester NY), pp. 2-5, Nov 1983.

20 U. Deshpande, G. Howell, S. Shamouilian, "High Density Interconnect Technology for VAX-9000 System," *Proc. of the SPIE*, (Boston MA), pp. 489-501, vol. 1390, pt. 2, 1990.

21 L. R. Buda, R. W. Gedney, T. F. Kelley, "High Density Packaging: Future Outlook," *Proc. of the IEEE Electr. Components and Tech. Conf.*, (San Diego CA), pp. 36-41, May 1992.

22 K. M. Nair, "Hybrids and Higher Level Integration," *Electronic Materials Handbook: Volume 1, Packaging*, City: ASM International, 1989, Section 3.

23 R. Chew, "A Study on the Feasibility of Using MCMs: A Designer's Perspective," *Proc. of the 1992 IEEE Multichip Module Conf.*, (Santa Cruz CA), pp. 127-129, March 1992.

EPILOGUE

THE CONCLUSION OF THE MATTER

(Ecclesiastes 12:9-14)

Not only was the Teacher wise,
but he also imparted knowledge to the people.
He pondered and searched out and set in order
many proverbs.
The Teacher searched to find just the right words,
and what he wrote was upright and true.

...BUT, BE WARNED, MY SON....

Of making many books there is no end,
and much study is a weariness of the flesh!

**NOW ALL HAS BEEN HEARD.
HERE IS THE CONCLUSION OF THE MATTER.**

Fear God and keep his commandments,
for this is the whole duty of man.
For God will bring every deed into judgment,
including every hidden thing,
whether it is good or evil!

ABOUT THE AUTHORS

Salma Y. Abbasi, Principal Engineer, Digital Equipment Corp., Cupertino CA. Ms. Abbasi has 10 years of experience in the field of advanced computer and semiconductor packaging and interconnect technologies, including high pin count, high performance packaging, thermal management of single chip and multichip packaging, tape automated bonding, die attach and process development and characterization in both multichip module and surface mount technologies. Additionally, she has extensive experience in TQM, DOE and DFM. She has worked on semiconductor packaging and interconnect related technology development for Advanced Micro Devices and Monolithic Memories for four and two years respectively, before joining Digital Equipment in 1988 as a senior process development engineer in the Multichip Unit Development Engineering Group. She has a BS from Westminster University in England, and is currently enrolled in an MS degree in Engineering Management at San Jose State University CA.

Kaveh Azar received his PhD in hydrodynamic stability from University of Connecticut in 1983. He joined AT&T Bell Laboratories in 1985 and has been involved in electronics cooling for the past six years. His activities include development of generic cooling techniques and thermal consulting. Dr. Azar also has been holding an adjunct position at Tufts and Northeastern Universities and

had developed graduate courses in electronics cooling and thermal modeling. He was the past general chairman of IEEE SEMITHERM conference, holds five patents in electronics cooling and has many publications.

Robert L. Bone received a BS degree in Industrial Technology, Electronics Option from California State University at Long Beach in 1968. He currently is Manager of Application and Product Development Engineering at Hughes Aircraft Company, Microelectronic Circuits Division, Newport Beach CA, responsible for engineering, program management, marketing and operations.

Steve J. Bezuk is the Manger of Chip Interconnect Technology, Unisys Corp., San Diego CA. He received the BS (Chemistry) from the University of Pittsburgh, the PhD (Chemistry) from the University of Minnesota and the MBA from the University of Redlands. His prior experience includes management of the Advanced Interconnect Technology Group, Sperry Corp. (now UNISYS), and a member of the Energy Systems Group, RCA David Sarnoff Research Center. His experience includes amorphous silicon research, CMOS processing, GaAs microwave processing, as well as microwave, thin film, wire bond, TAB and FCSB packaging. He has published papers on plasma diagnostics, $TaSiO_2$ - polysilicon gate materials for CMOS devices and laser photochemical and pyrolytic deposition of metals.

Allison Casey Dixon is Product Engineering Manager for Multichip Modules in Motorola's ASIC Division located in Chandler AZ. Ms. Dixon has championed the MCM development effort in the ASIC Division since 1988 and currently is responsible for bringing Motorola MCM products to market. Ms. Dixon has 10 years of experience at Motorola in a variety of engineering positions including wafer fabrication, product engineering and RF gallium arsenide development. She has a special interest in creating products from emerging technologies. Ms. Dixon graduated in 1982 from the M.I.T. with a BS in Chemistry. She is currently pursuing an MBA Administration from Arizona State University.

Dennis F. Elwell received a PhD in Physics from Sheffield University, Sheffield, England, and a Bachelor of Physics from Imperial College, London, England. He has served on the physics faculty of Portsmouth Polytechnic, England, researching electrical conduction in magnetic oxides and crystal growth from high-temperature solutions, has published over 100 technical papers and four books, and has received one patent. He currently is Manager of Technology and Senior Scientist/Engineer for the Process Technology Development Department, Microelectronic Circuits Division, Hughes (IEG), responsible for coordinating division internal technical research efforts in materials development processes

and analysis, and heads a team developing LTCC technology in support of pilot manufacturing. Earlier, he managed projects in atomic layer epitaxy, high-temperature superconductors, and nonlinear optical materials.

Claudius Feger, Research Staff Member, IBM T. J. Watson Research Center, Yorktown Heights NY, since 1984. His research concerns all aspects of polymeric electronics packaging materials, particularly polyimides. He holds three patents and has published over 50 technical papers and edited two books, He has organized several conferences including two conferences on polyimides. Before joining IBM he was working at the University of Massachusetts, Amherst, and at UFRGS, Porto Alegre, Brazil. He is currently adjunct professor at the Electrical Engineering Department of the University of Maine, Orono. He received the chemistry diploma from the Albert-Ludwigs University at Freiburg Germany in 1976 and from the same university, the PhD in polymer science in 1981.

Christine Feger, Staff Engineer, IBM Corp., East Fishkill NY. Ms. Feger has 15 years experience in all aspects of thin film process technology. She is presently responsible for polyimide process development and technology transfer in IBMs East Fishkill Packaging Laboratory. Prior to joining IBM in 1989, she worked for Rockwell International developing processes for thin film hybrid circuits. She was responsible for setting up the prototype multichip module fabrication facility for Rockwell in Anaheim CA. She received a BS in Chemistry from the University of California, Santa Barbara in 1977 and an MS in Electrical Engineering and Computer Science from UCSB in 1982.

Ronald W. Gedney, Program Manager, IBM Packaging Development Laboratory in Endicott NY. As engineer and manager, he has been involved in electronic packaging and development for some 30 years. He was responsible for the development of IBM's "metallized ceramic" chip carriers from 1974 to 1982, the most widely used single and multi-chip packaging technology in IBM. From 1982 to 1984, he was a consultant at IBM's plant in Havant, England, during the transfer of a substrate manufacturing line. In 1986-87, he established IBM's TAB/flex capabilities. Over the last three years, he developed IBM's low end chip carrier packaging strategy. A senior member of IEEE, he was President of the CHMT Society during 1990 and 1991. He holds a BS degree in Electrical Engineering from Tufts University.

Dean R. Haagenson is a Senior Engineer, Semiconductor Memory and Packaging Organization, Unisys Corp., San Diego CA. Mr. Haagenson received his BSME from the University of Minnesota, Twin Cities. He joined Sperry Corp., Semiconductor Operations (later UNISYS Corp.) in 1985, primarily

focusing on TAB assembly development as applied to multichip module programs. He transferred to his present location at the Unisys semiconductor facility in San Diego in 1987, where he joined the surface mount development group charged primarily with developing advanced SMT processes for manufacturing, including TAB bonding. Currently, Mr. Haagenson is working on advanced TAB and flip chip assembly development programs for Unisys. He has one TAB-related patent.

Eugene F. Heimbecher, II is a Senior Engineer, at Motorola Inc. in Chandler AZ. Mr. Heimbecher recently joined the Motorola ASIC Division Multichip Module Group as a system engineer for substrate analysis, signal integrity and design guidelines. Previously, he worked in Motorola Government Electronics for ten years. His most recent experience is three years as a system engineer with responsibility for design guidelines and noise management in hugh speed CMOS digital processing systems. He has written a CMOS system noise management handbook for use at Motorola. He also has seven years experience in high speed analog and frequency circuit design and analysis, primarily in the area of high performance low noise frequency synthesizers and radio receivers. He received his BSEE from the University of Wisconsin in 1981. He has served three years in the U.S. Army as an electronics technician.

Leo M. Higgins, III is Manager of Multichip Module and C4 Packaging and Technology Development, Motorola Inc., Austin TX. Dr. Higgins has 14 years of experience in microelectronic packaging in the areas of super minicomputer system packaging, ECL multichip module development, thermal management of very high power multichip modules and single chip PGAs, high performance/high power PGA and chip carrier development, inner and outer lead tape automated bonding, flip chip assembly, development and manufacture of ceramic packaging systems, etc. At Motorola, he manages a group responsible for SMT application engineering support and assembly technology development, TAB ILB/OLB engineering, and introduction of C4 assembly into Motorola. He is establishing a prototype/pilot assembly facility for C4/wire bond/TAB multichip modules constructed with MCM-L, MCM-D, and MCM-C substrates. Previously he worked on single and multichip ECL and CMOS packaging at Prime Computer (Framingham MA) for three years. Prior to this, he worked on cofired ceramic packaging technology at Cabot/Augat Technical Ceramics (North Attleboro MA) for five years as Manager of Engineering, and Vice President Research and Technology. Before this, he worked on a broad range of ceramic packaging for three years at Kyocera International (San Diego CA) as Senior Research Scientist, then Manager of Applied Technology. He has a BS, MS, and PhD in Ceramic Science and Engineering from Rutgers. He holds five patents, and has published more than 10 papers.

Ronald J. Jensen, Engineering Fellow, Honeywell Solid State Electronics Center, Plymouth MN. Dr. Jensen has more than ten years of experience in multichip module technology. In 1981, he joined Honeywell's Physical Sciences Center, where he developed the materials and processes for fabricating thin film multilayer (TFML) interconnecting substrates for MCMs. He also developed design and testing methods for MCM substrates, and demonstrated the TFML technology in a variety of test vehicles and functional MCMs for Honeywell divisions and other companies. In 1990, he transferred the TFML technology to a production facility at SSEC and implemented SPC and DOE methods for optimizing processes. He is now involved in other aspects of MCM technology, including bare die testing and MCM design and assembly. He received a BS in Chemical Engineering from Iowa State University and a PhD in Chemical Engineering from the University of California at Berkeley. He has published more than 20 papers and book chapters in areas related to MCM technology.

R. Wayne Johnson received the BE and MS degrees in 1979 and 1982 from Vanderbilt University, Nashville, TN, and the PhD degree in 1987 from Auburn University, Auburn, AL. All are in electrical engineering. He is an Associate Professor of Electrical Engineering at Auburn University. At Auburn, he has established teaching and research laboratories for thick and thin film hybrid technology. His thin film research efforts are focused on materials, processing and modeling for multichip modules. Research in power hybrids and high temperature electronics is being directed in the thick film laboratory. Dr. Johnson has published in technical journals and presented numerous papers at workshops and conferences. He also worked in hybrid microelectronics for DuPont, Eaton and Amperex. Dr. Johnson is a member of the IEEE CHMT Society, IEPS and was 1991 President of ISHM.

Alan D. Knight joined IBM after receiving his BS degree in Mechanical Engineering from Clarkson College of Technology in 1965. He has held various positions in Manufacturing Engineering, Industrial Engineering, Process Development Engineering, Product Engineering and currently is a manager of a connector development group. He has co-authored and presented numerous technical papers and holds several patents pertaining to design and fabrication of cables and connectors.

Erik N. Larson received a BS degree in Mechanical Engineering from San Diego State University. In 1987, he joined Unisys Corp. as a Mechanical and Process Development Engineer. His primary responsibilities at Unisys have been in the areas of thermal design and assembly process development for the corporation's Semiconductor, Memory and Packaging Operations division in San

Diego CA. He has been responsible for developing both wire bond and tape automated bonding processes for single and multichip VLSI ASIC packaging. He has also participated in the development of advanced liquid cooling systems for multichip packaging. Erik has received two corporate level achievement awards for enhancing wire bond production yields.

James J. Licari has a degree of Bachelor of Science in Chemistry from Fordham College and a PhD from Princeton University, where he was a DuPont Senior Fellow. At American Cyanimid, Remington-Rand Laboratories and Rockwell International, he has made extensive contributions in the development of new materials and in the study of fundamental materials properties for electronic applications. He is a Senior Scientist at Hughes Aircraft Company, Microelectronic Circuits Division, Newport Beach CA, where he serves as Program Manager for the Navy's VLSIC Packaging Technology (VPT) contract.

Arun Malhotra, Director of Thin Film Head Operations, IOMEGA Corp. He was manufacturing engineering manager for the VAX-9000 MCU Interconnect (High Density Signal Carrier - HDSC) at Digital from 1986 to 1991. Dr. Malhotra was a manufacturing engineering manager at Hewlett-Packard from 1984 to 1986. From 1976 to 1983 he worked on the NMOS III process development and copper-clad Teflon board development for Hewlett-Packard in Ft. Collins CO. He received the BS degree in Electrical Engineering from IIT in New Delhi in 1970, and the MS and PhD degrees in Electrical Engineering from Purdue in 1972 and 1974 respectively.

John D. Marshall, Product Engineer, Digital Equipment Corp., Cupertino CA. He was born in New Haven CT in 1954. He the BSEE degree from the University of Rochester in 1976, the MSEE degree from the California Institute of Technology in 1977 and the PhD degree in Electrical Engineering from Stanford University in 1988. His doctoral research was concerned with modeling field effect transistors for analog and digital circuits and focused on performance limits of silicon E/D MESFET integrated circuits. Since then he has been employed by Digital Equipment Corp. working on the design, modeling and characterization of high density, thin film circuit interconnect structures. Dr. Marshall is a member of Tau Beta Pi and Sigma Xi.

Robert S. Miles received both a BS degree in Electrical Engineering and a Masters degree in Business Administration from California State University, Los Angeles CA. In his 25 years at Hughes Aircraft Company, he has worked on the Maverick, Phoenix, AMRAAM and LEAP programs, and has held program management, CAD, process development, automated manufacture, and electrical

test positions at MSG, S&CG, and IEG. Mr. Miles currently works in the Hughes Microcircuits Division Process Technology Development Department, and manages the Hughes Corporate-wide Electronic Packaging Program (HEPP).

Douglas Modlin, Project Manager, Affymax Research Institute. Throughout his professional career, Dr. Modlin has specialized in the development of instrumentation containing integrated electronic, optical, mechanical and computational systems. From 1986 through November of 1991, Dr. Modlin was employed by Digital Equipment Corp. and led the team which developed the testers used to test the VAX-9000 Multichip Unit interconnect substrates. This included both in-process testers based on existing wafer probers and flying probe testers designed specially for testing the VAX-9000's fine pitch copper-polyimide interconnect substrates. Dr. Modlin is currently employed by Affymax Research Institute where he is developing optically based test equipment for medical diagnostics applications. He received his BS degree from California Polytechnical, San Luis Obispo in 1975 and his Masters and PhD degrees from Stanford University in 1978 and 1983, respectively.

Edward G. Myszka, Staff Engineer, Motorola's Corporate Manufacturing Research Center in Schaumburg, Illinois. He currently is a Project Manager responsible for the development of high density interconnect substrates for high performance MCM and PWB applications. He holds several patents and has published several papers on this subject. During his tenure at Motorola, he also was involved in the design of fiber optic transceiver modules and high speed, high power PAC packages. Prior to Motorola, Mr. Myszka worked for Lytel, Inc., where he was responsible for the design of component and subsystem packages for semiconductor lasers, light emitting diodes and photo diodes for telecommunications applications. Mr. Myszka received his Bachelor degree in Mechanical Engineering from the Polytechnic Institute of New York in 1984 and his Master's in Materials and Metallurgical Engineering from the Illinois Institute of Technology in 1990. He is currently pursuing an MBA degree from the Illinois Institute of Technology.

John A. Nelson Staff Engineer in the Product Technology group of the Unisys Semiconductor, Memory and Packaging Operation in Rancho Bernardo CA. His primary responsibility is in the area of advanced packaging. He holds several patents in packaging technology. At Burroughs/Unisys, he has held several management positions in packaging design, package test, ceramic package fabrication and package assembly. Prior to joining Burroughs in 1974, he had worked in avionics and computer systems packaging at Honeywell. He is active in the packaging activities of IEEE, ISHM and IEPS. Society activities have

included speaking, writing, chairing conference and session programs. In 1957, he received his BS in Aeronautical Engineering from the University of Minnesota. He also served as an officer in the U.S. Air Force.

Lee Hong Ng, Project Manager at IBIS Associates, Inc. in Wellesley MA. IBIS specializes in the competitive assessment of alternative technologies. Since joining IBIS, Dr. Ng has been involved with the analysis of alternative packaging schemes ranging from single chip packaging to multichip modules, and has published numerous papers on these topics. Dr. Ng also is involved in business planning and strategic analysis for companies in the electronics industry. Dr. Ng earned a PhD from the Materials Science and Engineering Department at M.I.T. She holds an MS degree from the same department at M.I.T. and a Bachelor's degree in Mechanical Engineering from the University of Texas at Austin.

Rajendra D. Pendse received the BS degree in Metallurgical Engineering from Indian Institute of Technology, Bombay, in 1980 and the PhD degree in Materials Science from the University of California, Berkeley, in 1985. In 1985, he joined National Semiconductor Corporation's Package Engineering Department, where he worked in the areas of thermal, stress and electrical modeling of packages, design of test chips for package reliability studies and several aspects of tape automated bonding. He joined Hewlett-Packard Company in 1989 in the Electronics Packaging Laboratory, where he was co-inventor of a proprietary new package called demountable TAB. He was involved in the development of designs, materials and processes for DTAB, including its subsequent transfer to manufacturing. He continues to work on advanced, high lead count and emerging packaging technologies.

Ralph Platz, Program Director, Maxtor Corp. From 1971 to 1992, he worked for Digital Equipment Corp. and managed the support of PDP-11 systems, the development of HSC50 Storage Subsystem, the development of a VAX-based parallel processor research system, and the development of the VAX-9000 MCU assembly process. Mr. Platz holds a BA degree in philosophy and an MA degree in Classics.

Karl J. Puttlitz, Sr., Senior Engineer, IBM Corp., Packaging Laboratory East Fishkill NY. He has been involved with the design, development and evaluation of interconnections for electronic packages since the early 1960s when IBM instituted its solder ball flip chip (SBFC) program. His early work included the development and implementation of SBFC replacement technologies practiced by IBM. More recently, he has been involved with hybrid interconnection applications and also substrate surface mount interconnections, such as the area-

array, solder-ball connection (SBC). He holds a BS, MS and PhD degrees in Metallurgy from Michigan State University, published numerous papers and credited with over 30 published inventions. His memberships include Sigma Xi, ASM International and New York Academy of Sciences.

Robert E. Rackerby, Principal Engineer, Unisys Corp., San Diego CA. Mr. Rackerby has been working in packaging and process engineering at Unisys for more than eight years. Early in his career at Unisys, he became involved with multichip packaging, working heavily in reliability, mechanical modeling and testing. Since 1986, he has been involved in the design and assembly of ceramic multichip modules. His primary responsibilities in MCM assembly have been to evaluate adhesives and strategies for die attach, select die attach materials and equipment, engineer attach processes and certify processes to production. He is currently involved with developing TAB and C4 interconnect processes. Mr. Rackerby has been awarded one patent and received his BSE from the departments of Chemical and Mechanical Engineering from California State University at Northridge in 1983.

Randy D. Rhodes, Department Manager, Applications Function of the Product Technology group of the Unisys Semiconductor, Memory and Packaging Operation in Rancho Bernardo CA. His responsibilities include packaging applications, electrical design, IC applications, memory engineering, product engineering and test. Mr. Rhodes has been involved in the packaging development at Burroughs/Unisys since joining Burroughs in 1983. Prior to that time, he worked on the first CMOS microprocessor and associated development systems at RCA. Mr. Rhodes received his BSEE from Georgia Tech in 1973.

Thomas C. Russell, Manager of the Design and Test Development Department, Alcoa Electronic Packaging, Inc., San Diego CA. Dr. Russell has been employed by Alcoa Electronic Packaging since 1989. His department is responsible for the design and test of advanced ceramic and silicon packages. During 1985 to 1989 he was a Member of the Technical Staff at AT&T Bell Laboratories in the Advanced VLSI Packaging Technology Development Department. While at Bell Laboratories, he was responsible for the development of a two probe, high speed robotic test system used in the electrical test of thin film substrates, ceramic hybrids and printed circuit boards. Dr. Russell holds a patent in a methodology for testing interconnect substrates. He received his ScB from Brown University in 1978, his MS in Mechanical Engineering from Stanford University in 1981 and his PhD in Electrical Engineering from Brown University in 1985.

Donald P. Seraphim, is a retired IBM Fellow and consultant. Dr. Seraphim has been developing advanced package technologies since 1965. He directed IBM's early programs on multi-chip bipolar packaging and during the 1970s directed the advanced printed circuit board program. Dr. Seraphim was a member of IBM's Corporate Technology Committee between 1977 and 1979, and was elected an IBM Fellow in 1981. He served as manager of materials science and engineering in IBM Endicott's Packaging Technology Laboratory from 1981 to 1986. He has written more than 50 articles, co-edited the book *Principles of Electronic Packaging*, and has over 30 patents and patent publications. Dr. Seraphim holds bachelor and masters degrees in applied science from the University of British Columbia, and a Dr. of Engineering in metallurgy from Yale University.

Kevin P. Shambrook has an AEA degree from Stanford University, a BSEE from Manchester University in England, an MSEE from University of Pittsburgh, and a PhD in Engineering and Economic Development from University of California, Los Angeles CA. He joined Hughes in 1983 after more than 20 years of experience in the management of engineering projects. At Hughes, he managed an Air Force-sponsored Technology Modernization program involving a paperless hybrid microcircuit manufacturing facility using a fault-tolerant computer for MRP II, and a Corporate-wide High Density Multichip Interconnect (HDMI) program. Currently, Dr. Shambrook is Director of the Center for Applied Competitive Technologies at Orange Coast College.

Stanley M. Stuhlbarg received a BS degree in Electrical Engineering from University of Cincinnati OH and a Management Certificate from Indiana University. At Hughes since 1974, he was responsible for engineering, marketing and strategic planning projects involving hybrid microcircuits, tape chip carrier development, and multichip module packaging. He initiated an Air Force technology modernization program in microcircuit manufacturing and a Mantech program on VLSIC packaging technology. Currently, Mr. Stuhlbarg is President of Engineering Technologies Associates, a consulting firm specializing in assistance to companies desiring expansion in Manufacturing Technology, Industrial Modernization Incentive Programs (IMIP) and Advanced Technology Projects (ATP).

Hongbee Teoh, Reliability Engineer, Digital Equipment Corp., MicroModule Systems, Cupertino CA. Dr. Teoh is currently working on the reliability testing and qualification of multichip modules technology for high performance applications. Since joining Digital in 1987, he has been working in the Quality and Reliability organization in Digital's Semiconductor Operations in

Hudson/Andover MA. He has been active in the area of single chip and multichip packaging, process technology and product qualifications, interconnect reliability and nondestructive testing of IC interconnects. Prior to joining Digital, Dr. Teoh was the Quality Assurance Manager for Kyocera America, Inc., a manufacturer of ceramic IC packages located in San Diego. Dr. Teoh received his BSME from the University of Rochester NY, an MS and PhD in Materials Science and Engineering from the University of California, Los Angeles CA (1984).

Philip A. Trask received a B.S. degree in Engineering from UCLA, an M.S. degree in Solid State Engineering from Arizona State University, and a Masters degree in Business Administration from Pepperdine University. He was responsible for R&D projects involving integration of passive component parts within silicon substrates, and development of a one-metabit MNOS memory using wafer scale integration. He has established a Class-100 fabrication line from High Density Multichip Interconnect (HDMI) Substrates, and has set up deep UV resist lithography in a prototype CMOS line. Currently, Mr. Trask is Manager of HDMI Substrate Engineering within the Hughes Semiconductor Products Center in Newport Beach, California.

William A. Vitriol, Senior Scientist, CTS Corp., Elkhart IN. Dr. Vitriol has more than 21 years experience in thick film ink development, hybrid microcircuit processing, and ceramic packaging technology. Since joining CTS, he has served as a technical resource for the Corp. in the area of electronic packaging, and, in particular, low temperature cofired ceramic technology. Prior to joining CTS, Dr. Vitriol worked for Hughes Aircraft Company, Newport Beach, CA for 10 years, where he was an Assistant Department Manager in the Technology Development Laboratory, for McDonnell Douglas Corp., and for Beckman Instruments. He has a BS and a PhD degree in ceramics from Rutgers University. He also holds six patents, plus two pending, in the area of thick film and ceramic packaging technology and has published over 15 technical papers. Dr. Vitriol is currently the ISHM Materials Division Chairman.

Yenting Wen, Staff Engineer, ALCOA Electronic Packaging Inc., San Diego, CA. Mr. Wen has been employed by Alcoa since January 1990. He has been involved in MCM development including design, test and high speed digital applications. He is also working on interconnect modeling and packaging characterization. Yenting Wen was a test engineer at AT&T Bell Laboratories from 1986 through 1990. From 1989 to 1990, he worked at Electrical Material Research Department and was responsible for MCM functional testing and electrical design and modeling. From 1986 to 1989, he worked in the Advanced VLSI Packaging Technology Department and was responsible for the functional

testing of thin-film multichip modules. He received his BS in Physics from National Central University, Taiwan in 1980 and MSEE from Rutgers, the State University of NJ in 1985.

Scott Westbrook, Member of Technical Staff, Design and Product Engineer Supervisor, Digital Equipment Corp., Cupertino CA. Since coming to California in 1980, he has worked in sequence for Hewlett-Packard, Trilogy Systems, and Digital Equipment Corp. During his career he has researched and developed MOS power devices, MOS memory, bipolar wafer scale technology and most recently thin film interconnect for multichip packaging. He was one of the developers of the copper/polyimide interconnect technology for multichip packaging that is used in the VAX-9000 mainframe/super computer. He is currently responsible for New MicroModule Products Engineering. Mr. Westbrook holds a BS (1978) and an MS (1980) in Electrical Engineering from M.I.T.

Chee C. Wong received his BS and PhD in Materials Science from M.I.T. in 1981 and 1986, respectively. Since then he has been with AT&T Bell Laboratories, where he is currently a Member of Technical Staff in the Electronics Packaging Research Department. His research activities include amorphous silicon devices, thin film processing and artificial neural network hardware implementations using amorphous silicon photoconductors. He was responsible for the development and transfer of solder bumping technology into manufacturing. More recently, he has been engaged in advanced interconnection technology for electronic/photonic multichip module applications.

Jimmy Jun-Min Yang, Principal Reliability Engineer, Digital Equipment Corp., MicroModule Systems, Cupertino, CA. Mr. Yang started at the Digital Taiwan Terminal Design and Manufacturing plant as Mechanical Engineer in 1980. He has previously been involved in VAX-9000 Manufacturing program office for Quality and Reliability planning, MCU Process Qualification as principal Quality Engineer in Marlborough MA, DEC disk and tape drives manufacturing plant as a Senior Component/Quality Engineer. He has a BS degree in Mechanical Engineering from Feng-Chia University, Taiwan, and a MS degree in Manufacturing Engineering from Pittsburgh State University. He is also a candidate for an MS degree in Manufacturing System Engineering from NTU, and recognized by the ASQC as both a Certified Quality Engineer, and a Certified Reliability Engineer. Mr. Yang has contributed to DEC internal Quality Planning Handbook and has written several Quality and Reliability technical papers. His current interests include System Quality and Reliability, Process Quality Modeling, System Reliability prediction and modeling techniques.

David F. Zarnow received a BS degree in Computer Engineering from the University of Illinois. He was associated with the Naval Avionics Center in Indianapolis IN, as a member of the Advanced Microelectronics Engineering staff, co-founding their hybrid microcircuit pilot production facility. Joining Hughes Microelectronic Circuits Division (MCD), Newport Beach CA in 1982, he became Manager of the Corporate MultiChip Module Computer Aided-Design (CAD) program, establishing standard design tools and company-wide foundry interfaces for multichip product designs. Currently, Mr. Zarnow heads product design and design information activities for the Hughes Microelectronic Circuits Division.

INDEX