

AD-A278 730



WL-TR-94-7023

①

AN OVERVIEW OF MULTICHIP MODULES

Mickie A. Phipps

DTIC
ELECTE
APR 29 1994
S F D

January 1994

408 **94-13041**

FINAL REPORT FOR PERIOD 1 NOVEMBER 1993 - 28 JANUARY 1994

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

WRIGHT LABORATORY, ARMAMENT DIRECTORATE
AIR FORCE MATERIEL COMMAND UNITED STATES AIR FORCE EGLIN AIR FORCE BASE

DTIC QUALITY INSPECTED 3

94 4 28 14 2

NOTICE

When Government drawings, specifications, or other data are used for any purpose other than in connection with a definitely Government-related procurement, the United States Government incurs no responsibility or any obligation whatsoever. The fact that the Government may have formulated or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication, or otherwise as in any manner construed, as licensing the holder, or any other person or corporation; or as conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

This report has been reviewed and is approved for publication.

FOR THE DIRECTOR



MICHAEL L. DELORENZO, Lt Col, USAF
Chief, Advanced Guidance Division

Even though this report may contain special release rights held by the controlling office, please do not request copies from the Air Force Wright Laboratory's Armament Directorate. If you qualify as a recipient, release approval will be obtained from the originating activity by DTIC. Address your request for additional copies to:

Defense Technical Information Center
Cameron Station
Alexandria VA 22304-6145

If your address has changed, if you wish to be removed from our mailing list, or if your organization no longer employs the addressee, please notify WL/MNSN, 101 Eglin Blvd Ste 483, Eglin AFB FL 32542-6810, to help us maintain a current mailing list.

Do not return copies of this report unless contractual obligations or notice on a specific document requires that it be returned.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE 28 Jan 94	3. REPORT TYPE AND DATES COVERED Final 1 Nov 93 - 28 Jan 94		
4. TITLE AND SUBTITLE An Overview of Multichip Modules		5. FUNDING NUMBERS JON: 9993G800		
6. AUTHOR(S) Mickie A. Phipps				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Assessment and Instrumentation Division Armament Directorate Air Force Wright Laboratory (WL/MNS) Eglin AFB FL 32542-6810		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Assessment and Instrumentation Division Armament Directorate Air Force Wright Laboratory (WL/MNS) Eglin AFB FL 32542-6810		10. SPONSORING/MONITORING AGENCY REPORT NUMBER WL-TR-94-7023		
11. SUPPLEMENTARY NOTES No distribution limitations. Availability of this report is specified on verso of front cover.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release: distribution unlimited.			12b. DISTRIBUTION CODE A	
13. ABSTRACT (Maximum 200 words) This report provides a high-level tutorial on the electronics packaging technology known as multichip modules (MCMs). The terminology associated with MCM development is explained in words which can be understood by personnel who do not have a background in electronics packaging techniques. The technology is introduced with a brief synopsis of the MCM history and market, proceeding into an in-depth discussion of interconnect substrates, known-good-die issues, chip attachment technologies, and testing/rework methods. The report concludes with a look at the fledgling industry infrastructure that supports MCM development, the high-speed chips that are used in MCM designs, the software support required, and the future of MCMs. Although by no means a comprehensive treatment of these subjects, this report provides the basics essential to an understanding of current state-of-the-art high-density electronics packaging. This information should be particularly beneficial to personnel involved in the development of miniaturized signal and data processors.				
14. SUBJECT TERMS Electronics packaging, multichip modules (MCMs), high density electronics, signal and data processing.			15. NUMBER OF PAGES 40	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR	

PREFACE

This report did not result from a specific Air Force project but is a general discussion of the state-of-the-art of the electronics packaging technologies involved in the development of multichip modules. The information presented in this report was extracted from a wide variety of sources, primarily electronics trade magazines and vendor marketing brochures. This information is intended to be used as an overview of the multichip module technology for personnel involved in the development of miniaturized signal and data processors.

Accession For	
NTIS	CRA&I ☒
DTIC	TAB ☐
Unannounced	☐
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

TABLE OF CONTENTS

Section		Page
I	INTRODUCTION	1
	1.1 The Multichip Module Technology	1
	1.2 History	2
	1.3 The Market	3
II	MCM DESIGN	5
	2.1 High-Speed Design Issues	5
	2.2 Consortia	5
	2.3 Computer Aided Design (CAD)	6
	2.4 Prototyping	6
III	MCM INTERCONNECT SUBSTRATES	8
	3.1 Trade-Offs	8
	3.2 MCM-L	8
	3.3 MCM-C	9
	3.4 MCM-D	9
	3.5 Other Interconnect Methods	10
IV	AVAILABILITY OF KNOWN-GOOD-DIE	12
V	CHIP ATTACHMENT TECHNOLOGIES	13
	5.1 Trade-Offs	13
	5.2 Wire Bonding	13
	5.3 Tape Automated Bonding (TAB)	13
	5.4 Flip-TAB	14
	5.5 Flip-Chip	15
	5.6 High Density Interconnect (HDI)	16
VI	TESTING TECHNIQUES	17
VII	REWORK METHODS	19
VIII	HEAT REMOVAL	20
IX	PRODUCIBILITY	22
X	INFRASTRUCTURE	23

XI	DSP, FFT, AND MEMORY CHIPS	24
	11.1 The Four Primary DSP/FFT Players	24
	11.2 Other DSP/FFT Vendors	25
	11.3 CISC vs. RISC	25
	11.4 Memory	25
	11.5 The Future of ICs	26
XII	SOFTWARE SUPPORT	27
XIII	THE FUTURE	29
	BIBLIOGRAPHY	30

LIST OF FIGURES

FIGURE	TITLE	PAGE
1	WIREBOND	13
2	TAPE AUTOMATED BONDING (TAB)	14
3	FLIP-TAB	15
4	FLIP-CHIP	16
5	HIGH-DENSITY INTERCONNECT (HDI)	16

LIST OF TABLES

TABLE	TITLE	PAGE
1	INTERCONNECT SUBSTRATE COMPARISONS	10

SECTION I

INTRODUCTION

1.1 THE MULTICHIP MODULE TECHNOLOGY

This report provides a high-level tutorial on the electronics packaging technology known as multichip modules (MCMs). Additional information can be found in the documents listed in the Bibliography.

High density electronics packaging is gaining rapid momentum in both research and in application; in fact, MCMs are the hottest topic in the semiconductor industry since the advent of very large-scale integration (Ref 1). Every once in a while, a key technology emerges that bridges some chasm, or blasts through a heretofore unpassable technological barrier, or initiates the buildup of a new and significant industry. MCMs may actually be that next great brick in the building of the electronics industry. But the implications of the widespread use of miniaturized electronics go far beyond better, faster, smaller technology. The entry of MCMs in the marketplace will force those involved in the procurement of military or civilian systems to change their entire framework of electronics specification, acquisition, testing and maintenance practices (Ref 2).

An MCM includes several integrated circuit (IC) die on one substrate that acts like a miniature printed-circuit board (PCB). The ICs are electrically connected to the substrate, which is connected to the module's input/output (I/O) pins (Ref 3). At first glance, MCMs appear to be an extension of application specific IC (ASIC), PCB, and hybrid technologies -- particularly of hybrids, which have been around for more than 30 years. But a closer look reveals that MCMs use different substrate and interconnect materials and introduce difficulties in package design, thermal management, signal integrity, testability, and ultimately cost, which are far more complicated than the problems raised by a traditional hybrid device (Ref 4).

Before MCMs can become a mainstream packaging technology, many factors must come together. Among other things, the packaging industry must address issues such as the availability of known-good-die (KGD), repairability, testing, producibility, vendor infrastructure, software support, and availability of sophisticated design tools (Ref 5). In spite of these challenges, the future looks very bright for MCMs. Today's investment towards solving those problems promises to make MCMs a viable alternative for applications requiring high throughput and significant memory in small packages.

1.2 HISTORY

This promising forecast for MCMs has a solid foundation which is based on the past decade's phenomenal growth in electronics. In 1980, the state-of-

the-art in IC technology included 3.5 μm feature sizes, single metal, n-channel metal oxide-semiconductor (NMOS) circuitry, a maximum 300 mil² die size, a maximum 28-40 I/O pins, 8-10 MHz clock rates, and 200-400 nanosecond instruction cycles. By 1993, device designers were investigating 0.1 μm feature sizes using complementary metal-oxide-semiconductor (CMOS); chips had grown to greater than 450 mil² with 240 pins; clocks were running at speeds greater than 80 MHz, computation units handled 32-bit data, and instruction cycles were less than 25 nanoseconds. Device speed alone accounts for a 5-10 fold improvement in digital signal processor (DSP) throughput simply because of shorter instruction cycles (Ref 6).

The desire to translate those higher operating speeds achieved at chip level into higher operating speeds at the subsystem level will certainly increase the acceptance of MCMs as a desirable technology. MCMs can handle the increased speed because the interconnects represent lower capacitance and controlled-resistance metallic lines (Ref 5).

1.3 THE MARKET

To date, the primary customer for MCMs has been the military. Due to the compact size, the harsh operational, environmental, and launch conditions, and the severe power constraints, the requirements levied on the onboard components of tactical and strategic weapon systems are very stringent. MCMs have been developed to satisfy that demand, providing miniaturized, high throughput, programmable processors for a variety of sensor and guidance processing applications.

For MCMs to flourish, the civilian market must also demand the functional advantages they provide in terms of speed, size, and weight for high-volume, low-cost applications. A major catalyst that may pull this technology into the mainstream is the multimedia industry (Ref 5), which is on the verge of revolutionizing the communications industry. Multimedia has the interactive capability of the personal computer, the communication power of the publishing press, and the entertainment factor of the broadcast industry all rolled into one. With its demand for full motion video and sound as well as the power of a supercomputer, multimedia is pushing the use of high throughput DSPs to perform computationally intensive image processing and real-time image compression and decompression. The Interactive Multimedia Association is discussing a cross-platform, cross-DSP application programming interface that would allow the multimedia industry to march forward without crowning another set of chip and operating system kings (Ref 7).

Today's typical full-complex, multimedia system is housed in a desktop-sized, conventional computer system, but MCMs provide the enabling technology that will allow multimedia users to become mobile. That mobility feature is the key to widespread use of these revolutionary communications technologies. As the integration of multimedia systems into the fabric of day-to-

day life continues, the MCM industry is likely to get a critical boost from the civilian market.

MCMs are also being investigated for uses in medical imaging systems, diagnostic equipment, and cellular communication systems. The discussion below focuses on the MCM development process -- from the module design to selection and fabrication of the multi-layer substrate; from selection and testing of the ICs to die attachment methods; from final testing of MCM to MCM applications.

SECTION II

MCM DESIGN

2.1 HIGH-SPEED DESIGN ISSUES

The latest surge of activity in MCM development is occurring primarily because significant hunks of technology, missing in earlier generations of electronics packaging, are now in place. Additionally, during the late 1980's academia recognized the growth potential of these technologies and began to include training in DSP use and MCM application development in the engineering curriculum (Ref 8).

Design engineers are beginning to understand that the success of a high-speed design is pegged on three issues -- architecture, component selection, and physical layout. Architecture is the first concern. The designer must be able to control what happens at each clock cycle -- what system elements need to be closely tied to the system clock and what elements can function asynchronously. It may be necessary to apply specialized high-speed components to handle the clock distribution and the interfaces between processor and cache memory, and between cache and main memory. Very high clock rates also make component selection a very important issue. The designer must choose between systems that utilize high-speed CMOS and those that use a bipolar technology. The third issue is physical layout and the availability of design tools to help visualize the analog effects of the physical layout. Tools are needed to model the impedance effects of IC lead lengths and interconnect line widths. The effects of impedance mismatch -- the invisible monsters of noise, ringing, crosstalk, and ground bounce -- will destroy a high-speed design. Tools that allow visualization of these transmission line effects on a workstation screen will become increasingly valuable (Ref 9).

According to the market research firm Dataquest, the worldwide market for MCM design tools is forecast as the most rapidly growing sector of the physical design market -- from revenues of \$20 million in 1991 to \$30 million by 1995 (Ref 2).

2.2 CONSORTIA

And yet even with this burst of development activity, one of the greatest gaps in the MCM technologies is still the lack of integrated MCM-design systems. Hopefully this area will see some changes with the introduction of such consortia as the Microelectronics and Computer Technology Center (MCC) in Austin, Texas, the worldwide Advanced Packaging Consortium in Huntsville, Alabama, and the Technical Alliance for Multichip Engineering (TAME) (sponsored by Harris Electronic Design Automation Division) in Fishers, New York. These groups and others like them bring MCM manufacturers,

consultants, suppliers, and end users together to advance the use and application of advanced packaging techniques, providing system designers with a means to successfully implement MCM technologies.

MCC is a cooperative enterprise whose mission is to strengthen and sustain the competitiveness of member companies by addressing R&D issues in evolving technologies, such as MCMs. Its members include a diverse group of companies from the aerospace, computer, electronic design automation (EDA), electronics, and manufacturing industries (Ref 4).

Through TAME, members can create cooperative marketing activities and benefit from each other's services and expertise. Membership is designed to provide wide access to design software, technology consulting, design service bureaus, university research organizations, fabrication-tooling companies, MCM and ceramic foundries, assembly services, semiconductor suppliers, and test-services companies (Ref 10).

2.3 COMPUTER AIDED DESIGN (CAD)

To develop high-integrity MCMs, intelligent computer aided design (CAD) tools must be available to handle the requirements for rerouting, precision placement, and thermal/electrical simulations. Designers must have validated Spice models (Simulation Program, Integrated Circuit Emphasis, a widely used circuit simulator) to design and analyze high-performance electrical interconnections. By the end of the 1990s, engineers will use software that accepts system specification in a high-level language, recommends partitions into various packages automatically, and synthesizes individual modules. Advanced autorouting and design-rule-checking programs will intelligently consider multiple design rules simultaneously (Ref 11).

Design kits for high density interconnect (HDI) technology, a chips-first approach from the MCM foundry at Texas Instruments (TI), and for laminated technology from IBM Microelectronics, have been introduced by Harris EDA for its Finesse MCM Design System. The TI kit contains a comprehensive set of data and system parameters and supports TI's Open Foundry concept for its HDI process. Users of the IBM kit can exploit the advanced capabilities of IBM's Endicott foundry. Harris EDA has also developed a universal design kit for MCMs that enables designers to select from low-temperature cofired-ceramic, silicon, or laminate chip-on-board technologies (Ref 12).

Harris EDA's Finesse MCM system is a state-of-the-art CAD system designed to handle the advanced packaging of MCM and ceramic hybrids. Finesse users can specify rules for via construction, feature size, metal plane construction, layer build order, routing, chip attach technology, and other process dependencies (Ref 13).

2.4 PROTOTYPING

Flexible, inexpensive prototyping capabilities must also be developed to provide design verification (Ref 11). Verification of the electrical performance of the design is critical and can only be completely accomplished through a fully-functional, electrically equivalent prototype. This type of prototype allows the measurement and post-test analysis of electrical characteristics of the MCM, to include power dissipation, voltage fluctuations, operational speed, and timing synchronization requirements.

There are also some very important reasons to pursue the development of non-functional MCM prototypes. Some believe that the day is not far off when engineers will complete a CAD model, select an icon on the computer, and create a part right there in the office -- a part which can be picked up from a peripheral just as a printout is picked up from a printer. Although this rapid prototyping capability will not allow electrical performance verification, it will provide solid models of the MCM that could be useful in selling concepts, optimizing designs, checking fit and tolerance, and even creating production tooling. Rapid prototyping parts are created via CAD, with the image "sliced" by software every 5-10 thousandths of an inch, something like a CAT scan. Each computer slice is a map for equipment that physically builds up the model, layer by layer, in paper, plastic, wax, metal, and even ceramics, using technologies such as stereolithography, laser sintering, and laminated object modeling (Ref 14).

SECTION III

MCM INTERCONNECT SUBSTRATES

3.1 TRADE-OFFS

Although there are other substrate technologies undergoing current research, such as diamond and laminated film, there are basically three types of MCM multi-layer interconnect substrates -- MCM-L (L for laminate), MCM-C (C for ceramic), and MCM-D (D for dielectric). The substrate serves as the chip-to-chip interconnection mechanism as well as the means to connect the chips' I/O to the outside world. See Table 1 for a comparison of substrate characteristics.

Each of the existing MCM substrate technologies comes with cost and performance trade-offs that make choosing a substrate a frustrating process. Without a doubt, there are times when the cost of the substrate is in line with the budget, but either the substrate can't support the routing density the circuit requires or its performance isn't up to the signal-propagation demands (Ref 15). Those are typical conflicts which continue to spur the development of new substrate technologies to expand the search for the ideal cost and performance balance.

3.2 MCM-L

MCM-L uses laminated PCB substrates, usually with copper conductors and vias (the connections between layers), with conductive patterns imaged by subtractive or additive deposition techniques. This method has been in existence the longest and most resembles the printed wiring board (PWB). At \$1/in²/conductive layer (Ref 15), cost weighs heavily in favor of MCM-L, but performance is significantly degraded when compared to that of MCM-C and MCM-D. Only in the area of speed can the MCM-L perform better than MCM-C technology because the trace resistance and capacitance can be lower. Thermal-expansion mismatches are generally more severe, and it's limited in terms of dimensional stability and power dissipation; therefore, MCM-L is usually limited to low-power, low-reliability applications. 12-mil vias are standard on MCM-L substrates, though 8-mil vias are possible. The maximum wiring density is 300 cm/cm², with minimum line widths of 60-100 μ m, and line spacings of 625-2250 μ m. The future for MCM-L includes research into improved performance at temperatures above 150°C, reduced dielectric constants into the 3-4 range, and reduced variation in physical properties such as thickness, flatness, electrical resistance, and homogeneity (Ref 5).

3.3 MCM-C

MCM-C consists of insulating materials with dielectric constants higher than five stacked between signal planes or signal and ground planes; conductors can be metals such as tungsten or molybdenum or screen-printable thick-film conductors such as gold, silver, or copper. The sequential, additive nature of the process means that the substrate can't be fully tested from an electrical standpoint until it's completed (Ref 15). MCM-C's advantage over MCM-L is the greater line density that it offers, and it also has the edge in thermal performance (Ref 3). At \$3/in²/conductive layer, MCM-C costs three times as much as MCM-L. The materials, such as green tape, are expensive, as is the gold metalization that's typically applied to the successive layers. Substrate shrinkage during sintering, a high heat process used to weld the components together without melting, sometimes causes problems as well. In addition, the capital equipment required to create ceramic substrates is costly (Ref 15). 4-mil vias are achievable with good yield, thus required wiring density can be obtained with fewer layers. Once 2-mil vias, spaces, and lines can be mass produced with high-yield processes, the electrical performance (speed) of MCM-C and MCM-L will be comparable. Current research is being conducted to lower the dielectric constant for MCM-Cs as well as to reduce layer thickness and via/conductor feature size (Ref 5).

3.4 MCM-D

MCM-D uses insulating materials with a dielectric constant lower than five (Ref 16). It is a thin-film technology, and its interconnect wafers are fabricated by depositing alternating layers of metal and dielectric, such as aluminum and polyimide, on a silicon-wafer base. Lithography is the most important determinant of electrical performance in MCM-D substrates and is an area of continued research. Silicon is used as the carrier for two reasons -- its expansion coefficient matches that of the ICs, and the thin-film manufacturing equipment used to fabricate the multi-layer substrate assembly is the same IC fabrication equipment used to process silicon wafers -- hence the term wafer-scale integration. This process also brings with it all of the attendant yield problems associated with IC manufacturing. When spinning on liquid polyimides, which can cost as much as \$400/gram, a large percentage of the material is wasted in the process, and the whole procedure must be carried out in clean rooms, which drives up the cost even further. The sequential, additive nature of the process means that the substrate can't be fully tested from an electrical standpoint until it's completed (Ref 15).

Although work at MCC suggests that MCM-D technology can surpass MCM-C in both cost and performance in some applications (Ref 5), it currently carries a cost of \$15/in²/conductive layer (Ref 15). Its greatest advantage is its line density, which is over 500 lines/inch, similar to a state-of-the-art VLSI chip (Ref 17).

Although dielectrics are not as critical as some of the other factors in establishing MCMs as a mainstream technology, the future of MCM-D technology rests on current research into the dielectrics. The goal is to reduce dielectric constants to less than 2 to achieve higher performance and to reduce on-module delays and maximize signal bandwidth. Reducing the moisture absorption of polyimide, which presents a problem for hermetic sealing requirements, is also a focus of much of the work in dielectric development. In addition, industry is attempting to simplify the required dielectric manufacturing process in order to reduce the cost. MCM-D manufacturers put cost reduction at the top of the list, because dielectrics typically make up 20% of a module's cost. Reduction of the environmental impact associated with dielectric use is another research area. Dielectrics currently employed in MCMs include epoxy-glass, alumina, silicon dioxide (SiO₂), benzocyclobutene (BCB), and polyimides, some of which are either toxic compounds themselves or involve the use of hazardous materials in the development process (Ref 5).

Another trend emerging from efforts to increase the functional density of MCMs is to employ thinner and thinner silicon substrates for MCM-Ds. In some instances, the silicon becomes so thin that it is transparent (advantageous for certain sensor applications). In addition to occupying less volume, thinned silicon also provides better thermal characteristics (Ref 16).

TABLE 1. INTERCONNECT SUBSTRATE COMPARISONS (Ref 5)

	MCM-L	MCM-C	MCM-D
Wiring Density	300 cm/cm ²	800 cm/cm ²	2000 cm/cm ²
Line Widths	60-100 μm	75-100 μm	15 μm
Line Spacings	625-2250 μm	125-450 μm	25-50 μm
Via Diameters	200 - 300 μm	50 - 100 μm	18-25 μm
Cost	\$1/in ² /conductive layer	\$3/in ² /conductive layer	\$15/in ² /conductive layer
Advantages	Low cost; established infrastructure	Greater density; better thermal performance	Greatest density; higher speeds achievable
Disadvantages	Limited to low-power, low-reliability applications	Difficult to test; expensive processing equipment	Higher cost; difficult to test; lacks vendor support

3.5 OTHER INTERCONNECT METHODS

Investigation continues into other interconnect methods, including a laminated-film (MCM-LF) substrate technology. MCM-LF, developed by Packard-Hughes Interconnect, of Irvine, California, reportedly approaches the speed and density potential of MCM-C and MCM-D while keeping costs closer to the MCM-L end of the scale. MCM-LF substrates consist of layers of a modified polyimide film (3.4 dielectric constant). Because this polyimide can be laminated at half the temperature required by earlier polyimides, it results in less thermal stress. The layers, which are only about 18 μm thick, are processed individually and then laminated together in a final step. This

process offers the advantage of being able to electrically test each layer before committing it to the final assembly. MCM-LF substrates permit laser drilling of vias of less than 5 mils. Materials and capital equipment required for MCM-LF fabrication are typically far less costly than for either MCM-D or MCM-C technologies. MCM-LF costs about \$3/in²/conductive layer (Ref 15).

The recently announced MCM-L/O (laminate overlay) process from TI is also a marriage of laminate and deposited MCM technologies. MCM-L/O combines a laminate substrate with an overlay of high-density thin-film interconnect. 1-oz copper traces provide the power and ground connections in the laminate substrate while the HDI process forms all signal paths. In high-volume production, the cost of MCM-L/O will be the same as that of conventional surface mount technology assemblies (Ref 16).

The search for materials with thermal characteristics that let MCMs handle the next high-powered generation of chips may have ended for some applications with the use of synthetic diamond as a substrate. Diamond's thermal conductivity is five times that of copper (Ref 18), and 50 times that of aluminum oxide, meaning that the die can rest directly on the substrate (Ref 19). Norton Diamond Film of Northboro, Massachusetts, as well as Sandia National Laboratories of Albuquerque, New Mexico, are investigating synthetic diamond substrates. Low producibility potential is still the greatest obstacle in making diamond the material of choice for MCM substrates.

SECTION IV

AVAILABILITY OF KNOWN-GOOD-DIE (KGD)

The primary problem that MCM designers face is the lack of availability of known-good-die (KGD). They need chips that have been tested at full speed with properly prepared bonding locations (bonding techniques are discussed in depth in following sections). Unfortunately, availability of KGD is not just a question of developing the right testing technology; it is also a problem that is surrounded by logistic influences and business conflicts. The industry infrastructure simply does not exist to support a small quantity, military-dominated demand. Until the civilian market also demands more bare die with KGD guarantees, vendors will be unlikely to support the development of the necessary test equipment. The reduced guarantee that comes with today's unpackaged die means added risk to the module user, usually in terms of an undefined MCM yield (Ref 5).

In addition to the need to obtain KGD, economic MCM production will rely on not over-testing the bare die. For example, National Semiconductor offers MCM makers the flexibility to choose between several testing flows to meet their manufacturing requirements. These testing flows range from complete ac/dc tri-temp and burn-in, to Mil-Std-883 requirements, to simple dc-only probing. This points out one of the problems faced by KGD users -- not all KGD are created equally. It's possible to get KGD which have not been tested at their specified operating speed. Testing at speed can be expensive and often isn't done, but if the correlation between the tested speed and the operating speed are known with a high level of confidence, testing at slower speeds generally will not cause a problem and will cost considerably less than high-speed testing (Ref 16).

After the KGD problem is solved, the most dominant fault type will be structural, in the form of bad interconnects and faulty bonds.

SECTION V


CHIP ATTACHMENT TECHNOLOGIES

5.1 TRADE-OFFS

The dominant chip-attachment technology will likely be determined by how the KGD issue is solved (Ref 5). If KGD are not available, then chip-attachment methods will be significantly affected by the need to test individual die before committing them to the MCM and by the inevitable requirement to provide rework capabilities. There are currently four primary methods used to attach the ICs to the MCM interconnect substrate: wire bonding, tape automated bonding (TAB), flip-chip, and the chips-first method used in the High Density Interconnect (HDI) process.

5.2 WIRE BONDING

Wire bonding (see Figure 1) is a mature, widely accepted chip attachment method that has flexibility to accommodate any design (Ref 4). It has the advantage of an established infrastructure because of its long history. However, with the increased number of chip I/O and the increased demand for smaller/faster technology, the labor intensive wire bonding technique is no longer considered practical. Moreover, rework is extremely difficult for chips with high pin counts (Ref 5).



5.3 TAPE AUTOMATED BONDING (TAB)

In TAB (see Figure 2), a chip with metal film conductors is laid down on multi-layer polymeric tape, similar to a sprocketed 35-mm photographic film. Customized metal conductors which have been patterned on the tape are bonded to the pads on the chip in the same place where wires would be bonded

in the conventional process. The tape, with bare chips attached, can be fed as reels, as strips, or as individual frames to test and assembly machinery. TAB offers significant test and handling advantages and is capable of achieving both very high assembly rates and packing densities (Ref 20). The disadvantage is the level of detail and attention required to design the TAB tape. Designers must specify the material for the tape along with precise dimensions, clearances, stiffness, pull strength, and test sockets. These requirements present a series of complex, often interrelated choices to the TAB designer.





FIGURE 2. TAPE AUTOMATED BONDING (TAB)

There are still plenty of challenges ahead in the TAB business, including such things as wafer bumping, precision alignment methods, rework, effects of lead inductance, vendor support, and the need for better, faster design tools. The wafer bumping process will very likely prove to be the most difficult part of the TAB development, but many of the problems can be avoided if the passivation layer of the interconnect wafer is designed specifically for bumping. Precision alignment and placement of the TAB'd die is critical, requiring a significant investment in sophisticated tooling (Ref 23). And each tape has to be custom-designed for each type of IC and must be redesigned even when trivial circuit changes are made, resulting in high non-recurring engineering (NRE) costs (Ref 20). TAB can also present severe lead inductance problems. And even though TAB provides the critical capability to test and burn-in individual die, industry has spawned few sources for TAB and even fewer cooperative IC vendors, providing perhaps the most important contributor to the slow growth of TAB applications.

Recognizing the mounting interest in MCMs, Chip Supply, of Orlando, Florida, is forging close alliances with customers, chip makers, and a specialized packaging house to deliver fully tested die in TAB form. They serve as the broker between the chip makers and the MCM manufacturers. Through their industry alliances, Chip Supply obtains the undiced wafers (which chip makers are reluctant to deliver to potential competitors), solder-bumps and dices the wafers, custom designs the TAB lead frames for the chips, and then tests the chips. One recent achievement, partly as a result of such alliances -- TAB tape standards have been established (Ref 1).

5.4 FLIP-TAB

Flip-TAB (see Figure 3), where the IC is actually mounted face down on the substrate, provides a smaller footprint with leads that come straight off of the die with no fan-out. And the leads don't have to be shaped or formed before outer lead bond (attachment of the outer end of the lead to the interconnect substrate).



It's tempting to say that TAB will continue to be a niche solution, but that's probably a US perception, based on the slow growth of the TAB infrastructure. The size of the Japanese TAB industry can't be ignored. They've solved the KGD problem, but in the US it's considered to be too costly and to take too long. The Japanese have shown that TAB can be cost-effective in large volumes. However, TAB may become obsolete as bare-die yield becomes acceptable enough for flip chip (Ref 5).

5.5 FLIP-CHIP

Flip-chip (Figure 4), which provides the greatest density, has become more of a necessity with increasing chip I/O count. With the flip-chip attachment method, each chip has solder-bumps placed on each I/O pad. The chip is then flipped face down and precisely aligned to matching bumps on the interconnect substrate. Flip-chip is electrically and thermally sound, offers high reliability, is reworkable, makes the most efficient use of substrate area, and allows high I/O counts. There is reason to believe that the problem of expansion mismatch between chip and substrate for large die with high pin counts will be solved in near future. Flip-chip, though it's a complicated and expensive procedure which is still in its infancy, is likely to be the dominant chip attachment technology in coming years; however, the transition from wire bonding may take 5-10 years. MCM vendors must first show that they have manufacturable flip-chip processes that don't require specially designed or specially processed ICs. Bumped chips are scarce, and the technology is in dire need of standardized tests; however, the infrastructure for solder-bumped chips

is definitely growing. Flip-chip has been used in IBM computers for years, using their proprietary C4 attachment method. They have demonstrated a phenomenal reliability record -- zero failures in 12 years (Ref 5).




FIGURE 4. FLIP-CHIP

5.6 HIGH DENSITY INTERCONNECT (HDI)

High Density Interconnect (HDI) (Figure 5) is the process wherein complex bare chips are placed face-up into cavities on a base substrate, then interconnected by conductive and polyimide insulating layers that are deposited on top of the chips and substrate. Via holes are opened in the insulating layers by lasers, which also photo pattern metal connections (Ref 21). The interconnect layers can be stripped away to replace defective die. For HDI to be a viable alternative for high volume applications, the KGD issue must be resolved.




FIGURE 5. HIGH-DENSITY INTERCONNECT (HDI)

SECTION VI

TESTING TECHNIQUES

Testing may be the greatest challenge facing MCM developers. Better bare die testing techniques must be made available before there can be any hope of providing KGD for widescale MCM use. A host of other issues surrounding MCM testing also demands immediate attention, including temperature sorting and testing of ICs; speed or access-time sorting; burn-in; interconnect wafer test; test and diagnostics of assembled modules; isolating bad die for rework; developing high MCM fault-coverage test vectors; and developing tools for evaluating MCM fault coverage. Reduced accessibility of internal nodes in the MCM simply means that designers are finally forced to do what they've known they should do all along -- that means real consideration of test issues early in the design cycle, more disciplined use of boundary scan and other design-for-test techniques, and more interplay between design and test personnel throughout development (Ref 5).

One of the key questions about bare die testing is "who does the testing?" Should it be the chip manufacturer or the module assembler? In the past, simply finding an IC vendor who would provide bare die was a major challenge, and it was primarily the responsibility of the module assembler to perform any IC tests. However, today many IC manufacturers are beginning to accept the challenge of providing well-tested die for MCM use. Membrane-probe technology is advancing rapidly for at-speed bare-die testing, which should allow drive circuits to be placed electrically close to the die under test. And temporary interconnect schemes, including TAB, chicklet substrates, and sacrificial substrates may be used for complete test and burn-in of individual die. Even with TAB, testing at full speed is difficult, and at very high speeds, chips may have to be selected to work together as a group to perform tests (Ref 5).

To be fair to silicon vendors, it's important to note that testing bare dies is a tricky proposition. Most silicon manufacturers are set up to do a low-speed test on wafers, categorize them according to specification and then do a full-speed test on the packaged die. Testing bare die is slow, difficult, and potentially harmful to the silicon, and most vendors are reluctant to invest the resources necessary to provide tested die or test information because the market is still so small (Ref 4).

With multiple chips involved, as in MCMs, fault isolation is often a time-consuming and costly process. Regardless of the choice of test equipment, or the ultimate promise of KGD, there is a near-complete consensus among MCM manufacturers and test system developers alike on the need for widespread implementation of JTAG (IEEE 1149.1) boundary scan techniques. Although JTAG won't find missing resistors, it can give 100% fault coverage for shorts and opens in boundary scan circuitry. However, since it is a static test

technique, JTAG doesn't directly aid performance testing, such as finding timing faults (Ref 2). In the traditional view, the benefits of IEEE-1149.1 boundary scan are aimed primarily at board manufacturing. But 1149.1 can aid product design as well by streamlining design tasks, reducing development costs, and improving time to market. Boundary scan circuitry typically takes up only 1-10% of the real estate in a VLSI chip of one million transistors. The 1149.1 test access port (TAP) also offers a "non-invasive" way to access a system while it's running instead of taking the system down or even taking it apart (Ref 22).

Algorithms that drive testability issues will have to be included to automatically add the appropriate level of boundary-scan circuit for optimal access to critical nodes (Ref 11).

Probe tests at the wafer level may need to include extra fault coverage, elevated voltage levels to induce faults in weak oxides, and temperature-controlled chucks to test over wide temperature ranges (Ref 5). Though expensive, device testers provide good critical path timing and full-performance testing. The drawback comes in testing modules with multiple complex die, where the development of functional test vectors for use on a chip tester can quickly become overwhelming (Ref 2).

Although MCM testing isn't fundamentally different than testing a dense surface mount technology (SMT) board, MCMs provide an interconnection medium with very low capacitive loading. The test equipment now used in chip manufacturing can't test devices in such low-load conditions (Ref 5).

SECTION VII

REWORK METHODS

Repairability during the MCM manufacturing process is a critical cost control issue that adds many additional requirements to MCM test systems, especially in the area of diagnostics (Ref 2). To help reduce that cost risk, repairability must be stressed from the beginning of the interconnect wafer design. For example, pad patterns on the wafers might be designed to handle 2-3 repair bonds; adhesives used for die attach should be adequately re-workable; and additional test ports should be designed on the wafer for probe access (Ref 23).

Sintered ceramics such as alumina, aluminum nitride, or beryllia may offer an answer to the nagging problem of reworking for MCM thin-film applications. Conventional silicon-dioxide substrates are so fragile that attempts at reworking the MCM -- removing a defective chip, for example -- generally ruin the MCM. The Via/Plane substrate technology developed by Micro Substrates Corp., Tempe, Arizona, reportedly solves this problem by making both sides of the substrate available for mounting chips (Ref 18).

The study of MCM failure mechanisms is in its infancy and the accompanying questions of field-return repairability have really yet to be tackled at all (Ref 2).

SECTION VIII

HEAT REMOVAL

The best way to handle heat is to avoid generating it in the first place, but high speed, dense circuits -- like MCMs -- are inherently more power hungry than a single chip solution, often creating a thermal management nightmare. But recent innovations in device design, such as smaller feature sizes and lower voltage requirements, promise a 20-30% reduction in power dissipation (Ref 5). Smaller features reduce parasitic switching losses, thus reducing the power consumption (if clock frequencies don't go any higher than they are today). And as devices get smaller, existing circuits also shrink in size -- as each block's area shrinks, the block can either run faster or consume less power (Ref 24). Lower power chips (1.5-3.3V) are already on the market.

CMOS gates implemented with 0.5 μm features dissipate about 2 $\mu\text{W}/\text{MHz}$ when powered by a 3.3V supply. So if just 4% of a 20 million transistor chip (4 transistors = 1 gate) switches at top speed (100 MHz), 200,000 gates will be simultaneously consuming power. Multiplying those numbers out, just the high-speed gates can consume about 40 W. Add in the slower switching elements and the chip could end up dissipating 50 W or more. But at 40-50 MHz, bipolar-CMOS circuits consume less power than CMOS circuits. The crossover occurs because biCMOS-circuit power consumption increases more slowly than does the power of standard CMOS circuits. By judiciously using biCMOS for speed-critical paths and mostly CMOS circuitry for slower circuits, designers can keep the power manageable (Ref 24).

MCM designers are also attempting to reduce both the capacitive loading in interconnections and the voltage change during switching to reduce power consumption in high-speed circuits (Ref 5).

Since some heat generation is inevitable, critical thermal management has become a spin-off industry of its own right. There are several techniques available to remove heat from MCMs. The easiest method to implement is also one of the most effective -- the inclusion of high thermal-conductivity substrate materials and heat sinks. But there are drawbacks. High thermal-conductivity substrate materials are usually very expensive and present severe handling and storage problems, and heat sinks add unwanted weight to the system. Low-thermal resistance paths from the chips to a heat sink can be provided through silicon substrates, then through brazed-in thermal slugs in ceramic hermetic packages. Or integral heat spreaders molded into the packages can also be combined with external heat sinks and forced-air cooling to carry the heat away from the MCM (Ref 5).

Heat removal can also be achieved with through-substrate cooling. When combined with low-noise ducted-air cooling, through-substrate cooling can handle chip-power densities of up to about 10-20 W/cm^2 , and module-power densities of up to about 2-4 W/cm^2 . Heat can be removed from the front side of

wire-bonded or TAB'd chips in a cavity-up configuration by using a machined lid and various thermally conducting, electrically insulating compounds between the chips and the lid. Above-substrate cooling will probably continue to dominate high-end systems, strongly competing with direct immersion of the chip or package in inert liquids such as fluorocarbons. Phase change techniques are another option for certain applications, as are liquid cooling techniques, but both have the potential for creating acoustical noise problems. Thermoelectric coolers are also attractive solutions for certain applications (Ref 5).

SECTION IX

PRODUCIBILITY

Three things that will help improve producibility are: low cost substrate fabrication (a must); metalization techniques for wide data paths, and lower die counts on the MCM.

One of the most pressing needs that must be met before MCMs will move into the mainstream is the need for low cost production techniques. The MCM must cost less than or the same as conventional packaging methods for it to become a viable alternative (Ref 5).

Because of noise and signal-integrity issues, connecting 32-, 64-, and 128-bit wide data paths to the external world is inviting trouble. To make wide on-chip buses practical, new metal systems together with improved planarization schemes must be developed to keep wiring resistance low and to permit four or more levels of interconnections (Ref 24).

The trend toward higher integration is making high die counts unnecessary, and the chances of assembling a good module are significantly increased if the die are KGD (Ref 24).

SECTION X

INFRASTRUCTURE

For MCM solutions to be readily available through vendors, there has to be an established, stable infrastructure to carry the product from inception to delivery, from substrate material choice to chip attachment to test. There has to be a spectrum of understood, available technologies that span the cost-performance domain between MCM-L and MCM-D (Ref 5).

Technical and business interfaces must be established. At the present time, only vertically integrated companies can become serious players in the MCM business because they can control both the IC production line as well as the substrate development. Vertically integrated contract-assembly houses have the potential to manufacture MCMs in high volume as inexpensively and quickly as single-chip packages (Ref 5).

In 1991, the US semiconductor industry produced less than 5% of the world's total production of ICs in packaged form. Without a strong IC "back-end" manufacturing base in the US, there may not be a significant MCM industry here either (Ref 5).

Major IC and CAD companies have not been leaders (except IBM) in stepping out and developing the necessary infrastructure and capabilities either. Until they do, there won't be a commitment to major-volume applications. The technology for direct chip-attach interconnections hasn't grown to a point where most manufacturers respond to the demand (Ref 5).

There is also an immediate need for standards in areas such as package footprints, die-procurement specifications, and data interchange formats. The lack of such standards makes each company's products unique with no opportunity for second sourcing (Ref 5).

Alliances and consortia are already a major force in the advancement of MCM technologies, and we are likely to see even more of them in the future. We may also see more MCM "brokers," such as Chip Supply, acting as the interface between the chip manufacturer and the user. The broker would deal with the problems of assembly and test and could help to overcome the reluctance of chip makers to supply tested wafers to other companies who are potential competitors (Ref 5).

Designing, fabricating, assembling, and using MCMs must be made easier. Off-the-shelf (OTS) assembly equipment will have to be available to eliminate the current practice of customizing complex hardware for each application and to delete the relatively high non-recurring engineering (NRE) costs. Lead times for MCM prototypes are very long compared to ASICs, making them less attractive for immediate applications. And the learning curve for the MCM user is extremely long (Ref 5).

SECTION XI

DSP, FFT, and MEMORY CHIPS

11.1 THE FOUR PRIMARY DSP/FFT PLAYERS

Any discussion of the MCM packaging technologies must also include some of the MCM components -- particularly the digital signal processors (DSPs), Fast Fourier Transform (FFT) chips, and memory chips. Since this paper primarily focuses on electronics packaging, this discussion will be brief, but that doesn't mean that this subject is of any less importance. There is a direct, dependent relationship between the major MCM components and the total MCM design/performance.

In the mid 1980s, four fixed-point processor architectures were introduced: DSP16 from AT&T Microelectronics, ADSP-2100 from Analog Devices, DSP56000 from Motorola, and TMS320C25 from Texas Instruments. Each of these chips formed the basis of an architecture upon which each company built DSP product families. Those vendors are still the top four players in DSP development today. Later derivatives of their chips offered improvements in speed and capability, but the fundamental architectures limited performance. Constrained by the technology that existed in the '80s, DSP vendors were forced to trade off features and performance. Low pin count packages meant limited external address space, and limited on-chip memory meant reduced performance and added system cost. Limitations in integration meant limitations in op-code length and, consequently, difficulties in programming. Because they wanted to maintain backward code compatibility, DSP vendors have perpetuated these limitations in later derivatives (Ref 6).

TI's TMS320C40, which is a throwback to the transputer, provides 50 MFLOPS throughput, 40 nanosecond instruction cycles, a pair of 32-bit external memory ports, and six high-speed byte-wide parallel 20 MByte/s I/O ports with independent DMA. The transputer had four 10Mbit/s serial ports, giving an aggregate I/O bandwidth of 5 MByte/s (Ref 25).

Motorola's DSP6002 provides two 32-bit, non-multiplexed address/data memory buses (Ref 6).

ADI's recently released ADSP-21060, the Super Harvard Architecture Computer (SHARC), is marketed as being an order of magnitude faster than any other floating-point DSP chip on the market. It provides 100 MFLOPS throughput, 128Kx32 SRAM, a host interface port, a DMS controller, and multiprocessing support. It was designed to permit construction of a massively parallel system with essentially no glue logic. It uses 3.3 V, 0.6 μ m CMOS technology to minimize power dissipation. And Ada software available for 21020 will be optimized and upward compatible to the 21060 (Ref 26).

The SHARC performs 1K FFTs in 0.46 milliseconds at 40 MHz. While that's an impressive FFT, it's a single benchmark. And in an overall system-level performance requirement like that of a seeker of a missile, or a navigation computer, or sonar system processing, an FFT is a small percentage of the overall system requirement. Benchmarks must be more system level oriented to carry much weight (Ref 26).

11.2 OTHER DSP/FFT VENDORS

Other major chip vendors include Zoran, Digital Electronic Corp (DEC), AT&T, Intel, and IBM.

DEC's Alpha AXP chip is a 64-bit/200 MHz RISC processor providing 400 MIPS. It is flexible enough to work with UNIX, Windows NT, Open VMS, and DECelx. It has a full set of comprehensive development support tools including debuggers, compilers, assemblers, and design tools (Ref 27). Second sourcing is available through Mitsubishi. The fastest Windows NT server available -- well beyond 200 MHz -- will be based on an Alpha-based MCM mounted on a diamond substrate (Ref 28).

AT&T's DSP32C provides a 25-Mbit/s serial port and a 16-bit, 12 Mbyte/s parallel I/O port with independent DMA (Ref 25).

Zoran Corp developed the ZR38000, which is optimized to perform 4-cycle radix-2 FFT butterfly operations and execute 25 MIPS. It provides 32-bit wide instructions, 20-bit-wide data, and a 20x20 bit multiplier. A 1024-point complex FFT requires just 0.88 ms, about one third the time required by Motorola's DSP56002 and TI's C51 (Ref 29).

11.3 CISC vs. RISC

CISC and RISC processor architectures will become indistinguishable from each other by the next century, and on-chip memory will increase dramatically. Already, the latest CISC-style CPUs are exploiting architectural tricks employed by today's current crop of RISC CPUs, such as superscalar operation and superpipelining. Today's high-performance RISC CPUs or DSPs typically employ less than 400,000 transistors each to form the actual central-processor core and math-ccprocessor blocks. In the future, CPUs will pack close to 50 million transistors, offer 64- and 128-bit-wide buses, and use half to two-thirds of their transistors for large amounts of on-chip cache memory (assuming 6-transistor SRAM storage cells are still the preferred approach, that would equate to over half a Mbyte of SRAM for cache) (Ref 24).

11.4 MEMORY

Memory designers expect to be producing 256-Mbit DRAMs by the end of the decade (32 Mbytes on one chip). By then, 1-Gbit DRAMs will be in the prototyping stage (Ref 24).

Using chip-stacking technology, both Texas Instruments and Irvine Sensors have developed 3-D memory packages. Texas Instruments took planar memory chips, routed all of the I/O off of a single side of each chip, stacked them with alternating layers of dielectric, turned the stack so that the memory chips stood in a vertical orientation, then solder bumped the completed stack to the MCM substrate. Irvine Sensors' component is a thin package of horizontally stacked memory chips that's physically interchangeable with a single-chip package. The individual memory chips in the stack are linked at the top by a "cap chip," which allows various conventional attachment techniques to be used, including wire bonding, TAB, or solder bumps (Ref 30).

11.5 THE FUTURE OF ICs

The future of ICs holds great change ahead in terms of physical construction and the level of functional complexity. Reduction of feature sizes will continue. Today's most advanced volume-produced chips typically have minimum drawn feature sizes of about 0.6 μm . That's less than one-tenth the size of the 6-10 μm features employed about two decades ago. It seems that feature size decreases about 50% per decade. If that trend continues, the scenario goes as follows: The early 90's 0.6 μm features will give way to 0.5 μm in 1994, and current research is looking at 0.4-0.35 μm features for 1995-96 and 0.25-0.3 μm by 1997/98. Exploratory work with 0.1 μm features has already started, with production targeted for early next decade (Ref 24).

Chip sizes are getting larger. Today's largest volume-producible monolithic circuits are about 600 mils on a side. By the end of this decade, chips that measure about 800 mils on a side will be manufacturable, almost doubling the available area for circuits on the chip (Ref 24).

An area increase alone may not be enough to satisfy the need to integrate multiple functions or large amounts of memory on one chip. Silicon designers are already looking at the vertical, or third dimension -- digging trenches into the silicon and/or creating multiple layers of active elements above the substrate surface (Ref 24).

As features shrink and insulating-oxide layers go below 100 angstroms, the electrical field stresses imposed across regions such as between the gate and the substrate must also be scaled or the stress will cause punch-through (short circuits). The stress induces a hole to open in the dielectric layer that separates the gate from the channel and allows a direct connection between the gate and channel (Ref 24).

Optical lithography is getting a new lease on life in the submicron range thanks to a new technology that allows the printing of 0.35 μm feature sizes on semiconductor wafers. It can print the equivalent of two 64 Mbit DRAM circuits in one exposure with a resolution of less than 0.45 μm (Ref 31).

SECTION XII

SOFTWARE SUPPORT

DSPs, a major component of MCM designs, are the Rubix-Cubes of the programming world. Only 10% of programming time is spent developing the DSP algorithms; 90% of the time is spent trying to optimize the assembly code to make the algorithm run properly on the chip (Ref 6).

Engineers working on designs having multiple DSPs for an embedded system have a wide variety of development tools at their disposal -- tools that provide assistance throughout the many stages of the design effort. Special operating systems tailored to the needs of DSP are becoming available. In general, DSP tools address three main design stages: algorithm development, software coding, and debugging (Ref 31).

The heart of a DSP system's functionality is its algorithms. They describe the mathematical procedures the DSP uses to operate on the input to generate an output. To design and evaluate an algorithm, simulation of the DSP portion of the embedded system is required. It should provide a flexible graphics interface, access to standard DSP functions (FFTs, digital filters, etc.), the ability to represent the rest of the system as a macro or other high-level description, and the ability to accept user-supplied data as input (Ref 31).

Responding to the needs of their customers, DSP chip manufacturers now usually provide a reasonable code development environment for their own products. Today's DSP chips, with their large address space and rich set of registers, make it much easier to create compilers with decent efficiencies. So high-order-language (HOL) compilers for DSP-executable code are much more common now than they were in the early days of DSP. Segments that need optimizing can always be rewritten in assembly language (Ref 31).

Some source-level debugging tools are also available for symbolic debug of HOL DSP code. There are in-circuit DSP emulators and DSP-chip simulators, too. A simulator can provide results similar to those of an emulator, but it runs significantly slower and, by its nature, precludes the use of real-time data (Ref 31).

Debugging a multiprocessor DSP system brings a whole new set of problems beyond those found with a single-processor system. Controlling and monitoring the state of each processor can be difficult (Ref 31). Most debuggers let programmers view the state of only one processor at a time. Only a handful (such as the Signal Processor Applications Development Environment (SPADE) developed by Rockwell for their Signal Processor Packaging Design (SPPD) family) let the programmer observe the state of multiple DSPs. Ideally, a multiprocessor debugger should provide a separate window for each DSP. Each window should supply information about all of the tasks running on that DSP, together with all of the resources that the DSP shares with other DSPs (Ref 33). Because embedded DSP systems are often designed to deal in real time with a

continuous stream of input data, setting breakpoints and single-stepping through code must be done with care (Ref 32).

Recent advances in DSP design and performance are generating an interest in providing operating system features such as multitasking, memory management, event scheduling, I/O and communications facilities, timer services, interrupt handling, and device-driver support within embedded DSP systems (Ref 32).

Spectron Microsystem's SPOX, is probably the first operating system that was developed specifically for use on DSP chips (Ref 32). SPOX, 3L's Parallel C, and Perihelion's Helios provide real-time, pre-emptive multitasking and memory management, and handle synchronization and interprocess communication mechanisms (mailboxes, variable and fixed size FIFOs). SPOX and Parallel C support virtual channels, which let the programmer implement interprocess and intertask communication without specifying physical connections in the application program. The actual communication mechanism is determined by a driver selected via a configuration table; this allows the programmer to change the system configuration without changing the source code (Ref 25).

Comdisco's Multi-Prox supports Motorola's DSP96002 and C40-based multiprocessor platforms. It allows the application to be described by interconnecting blocks that represent functions such as FFTs, filters, and I/O, then drawing bounding blocks around portions of each diagram to be implemented on a particular DSP. It automatically compiles each block and establishes the communication, insulating the user from the details of the underlying DSP hardware (Ref 25).

Although tools such as compilers, debuggers, libraries, and real-time operating systems give designers the tools needed to develop DSP applications, attracting mainstream designers with minimal DSP experience requires a high-level approach to application development. Real-time DSP designers have always resisted using higher-level tools that separate them from the bits and bytes of their DSPs; however, mainstream designers will insist on being separated as much as possible from the underlying details of the DSP hardware (Ref 33).

SECTION XIII

THE FUTURE

The future looks bright for MCMs. The turn of the century will reveal the fruits of the intensive research going on today in such areas as optical interconnects; 3-D, stackable MCMs; KGD; and 0.1 μm (and smaller) feature sizes. We are likely to see a wide variety of DSP applications in the civilian market, accompanied by a host of sophisticated design and software development tools. The emergence of an industry-wide infrastructure that can support MCM production is inevitable.

Most of the basic trends in mainstream electronic systems work in favor of increased MCM usage -- it becomes more valuable as IC speeds and pin counts increase, as signal-integrity issues multiply, and as noise margins decrease (Ref 5). In fact, in future high-performance systems, even the metal interconnects between MCMs will be too slow and designers will be forced to employ optical buses. Rockwell International is already developing an MCM that includes four fiber optic I/Os to interconnect 128 64-bit DSPs all running in parallel. The resulting signal processing subsystem will be capable of 20 BOPS, but it will consume about 460 Watts on a SEM-E card (Ref 16).

In order to realize the high performance promised by hybrid optical-electrical systems and mixed digital-analog systems, available fabrication and CAD/CAE technologies for MCM applications are an absolute must. CAD/CAE and test technologies must suit the high speed MCMs being designed today -- ones that run at 2-5 GHz clock frequencies and 20-50 psec edge speeds (Ref 5).

1994 is likely to be the year that DSPs stake their claim to some space on that valuable square foot of PC real estate. While there are compelling technical arguments for DSP's being intimately integrated into the PC's basic hardware specification, that could never happen without a strong bridge to Windows. It now seems imminent that Microsoft will announce the availability of a software developer's kit to construct that bridge. It's also safe to bet on the SPOX DSP operating system from Spectron Microsystems Inc. as Microsoft's choice for handshaking with DOS and Windows. Digital-signal processing is just too embedded in the technologies for transforming sound, music, telephony, and video to be ignored., and there simply are not enough spare MIPS floating around in a 486 to handle real-time, multiple-interrupt functions such as telephony and music synchronization (Ref 34).

Research continues in the area of high-temperature superconductors for MCM interconnects. Superconducting interconnects between chips could be over 10 times smaller than conventional metal interconnects, allowing chips to be stacked very closely together, while still avoiding cross talk between chips (Ref 35).

All in all, it's probably safe to say that MCMs are here to stay and may be the key to the next generation of high-performance processors.

BIBLIOGRAPHY

1. **Military & Aerospace Electronics**, Sep 91, "Distributor Eases into Bare Die Supply for MCMs" by Lisa Burgess.
2. **Military & Aerospace Electronics**, 20 Sep 93, "In Search of a Military MCM Test Strategy" by John S. Haystead, p. 27-33
3. **IEEE Spectrum**, Nov 92, "Multichip Routing and Placement" by Wayne Mei-Ming Dai, p. 61-64
4. **Computer Design**, Aug 92, "MCMs Push Design and Test Tools to the Limit" by Mike Donlin, p. 59-65
5. **Electronic Design**, 25 Nov 92, Special Feature: 40 Years Technology -- Packaging & Components, "MCMs: 11 Experts Debate the Future" by David Maliniak, p. 157-171
6. **Computer Design**, Nov 93, "High-performance DSP Shuns Old Architecture" by Jeff Child, p. 46-48
7. **Electronic Design**, 14 Oct 93, "DSP Gets Ready for Multimedia" by Jack Shandle, p. 64
8. **Computer Design**, Aug 92, "DSP Boards Reach Performance Highs" by Warren Andrews, p. 69-79
9. **Computer Design**, Aug 92, "Making the 400-MHz Computer a Reality" by Stephan Ohr, p. 83-99
10. **Electronic Design**, Jul 92, Technology Newsletter, "Technical Alliance Furthers Cause of MCMs"
11. **Electronic Design**, 25 Nov 92, "The Future of Packaging Design" by Tony Mazzullo, p. 120
12. **Electronic Design**, 16 Dec 93, Technology Newsletter, "MCM Design Kits Expand System's Reach" p. 28
13. Harris EDA Promotional Brochure on FINESSE MCM Design System
14. **Industry Week**, 1 Feb 93, "Rapid Prototyping Moves to Desktop" by Tim Stevens, p. 38-44
15. **Electronic Design**, 11 Nov 93, "Laminated-Film MCMs Aim For Mass Market" by David Maliniak, p. 57-64
16. **Military & Aerospace Electronics**, 20 Sep 93, "ARPA Leads the MCM Charge" by Jeffrey D. Shepard, p. 31-33
17. **Inside ISHM**, Mar/Apr 91, Vol 18 No 2, "Multichip Module Cost: The Volume-Yield Relationship" by Dr. Lee H. Ng, p. 14
18. **Electronic Design**, 16 Dec 93, 1993 in Review, "Technology Rolled Up Impressive Gains in 1993" by Jack Shandle, p.53-64
19. **Electronic Design**, 18 Feb 93, Technology Newsletter, "CVD Diamond Makes Substrate Debut" p. 23
20. **IEEE Spectrum**, Jun 85, "The DIP May Take its Final Bows" by Reed Bowlby, p. 37-43
21. **Electronic Design**, 17 Dec 92, "Waferscale Integration Moves to Three Dimensions" by Dave Bursky, p. 30

22. Electronic Design, 7 Jan 93, "Boundary Scan will also Improve the Design Process" by Peter Hansen, p. 109
23. Paper presented at 1991 IEEE conference, "TAB Technology for Miniature Multiprocessor High Speed Interceptors" by Robert K. Scannell, Tactical Systems Division, Rockwell International Corporation, Anaheim, CA
24. Electronic Design, 25 Nov 92, "Digital ICs Become the System as Complexities Increase" by Dave Bursky, p. 63-70
25. Computer Design, Aug 92, "DSPs in Multiprocessing" by Tony Agnello, p. 25-27
26. ADSP-21060 SHARC marketing material from Analog Devices, Oct 93
27. Alpha AXP advertisement in Electronic Design, 11 Nov 93
28. Electronic Design, 27 May 93, Technology Newsletter, "Alpha-Based MCM Will Ride on Diamond", p. 25
29. Electronic Design, 11 Nov 93, "Improved DSP ICs Eye New Horizons" by Dave Bursky, p. 69-82
30. Electronic Design, Technology Newsletter, "Chip-Stacking Method Boosts Memory Density" by DM
31. Electronic Design, 7 Jan 93, Technology Newsletter, "Stepper Handles Small Features, Large Die" by JS, p.28
32. IEEE Spectrum, Nov 92, "Tools for Embedding DSP" by Robert W. Robison, p. 81-82
33. Electronic Design, Embedded Controllers Applications Issue, 1 Oct 93, "Multiprocessing Exploits DSP Performance Potential" by Tony Agnello, p. 50-55.
34. Electronic Design, 16 Dec 93, Technology Briefing, "The Signs are Right for DSPs and PCs" by Jack Shandle, p. 18.
35. Ballistic Missile Defense Organization Technology Applications Update, Oct 93, "Market for High Temperature Superconductivity Feels Warming Trend" by Leslie Aitcheson, p. 11

DISTRIBUTION LIST

DTIC/DDAC	1
Cameron Station	
Alexandria VA 22304-6145	
ARPA/SPO	4
Attr: Jim Silverthorn	
3701 N. Fairfax Drive	
Arlington VA 22203	
USSTRATCOM/J534	1
ATTN: Lt Col Swann	
901 SAC Blvd, Ste BB18	
Offutt AFB NE 68113-6200	
WL/MNSN	6
WL/MNG	2
WL/MNGA	2
WL/MNGI	1
WL/MNS	1
WL/MNA	1
WL/MNM	1