

An advanced multichip module (MCM) for high-performance UNIX servers

by J. U. Knickerbocker
F. L. Pompeo
A. F. Tai
D. L. Thomas
R. D. Weekly
M. G. Nealon
H. C. Hamel
A. Haridass
J. N. Humenik
R. A. Shelleman
S. N. Reddy
K. M. Prettyman
B. V. Fasano
S. K. Ray
T. E. Lombardi
K. C. Marston
P. A. Coico
P. J. Brofman
L. S. Goldmann
D. L. Edwards
J. A. Zitz
S. Iruvanti
S. L. Shinde
H. P. Longworth

In 2001, IBM delivered to the marketplace a high-performance UNIX®-class eServer based on a four-chip multichip module (MCM) code named Regatta. This MCM supports four POWER4 chips, each with 170 million transistors, which utilize the IBM advanced copper back-end interconnect technology. Each chip is attached to the MCM through 7018 flip-chip solder connections. The MCM, fabricated using the IBM high-performance glass-ceramic technology, features 1.7 million internal copper vias and high-density top-surface contact pad arrays with 100- μm pads on 200- μm centers. Interconnections between chips on the MCM and interconnections to the board for power distribution and MCM-to-MCM communication are provided by 190 meters of co-sintered copper wiring. Additionally, the 5100 off-module connections on the bottom side of the MCM are fabricated at a 1-mm pitch and connected to the board through the use of a novel land grid array technology, thus enabling a compact 85-mm \times 85-mm module footprint that enables 8- to

32-way systems with processors operating at 1.1 GHz or 1.3 GHz. The MCM also incorporates advanced thermal solutions that enable 156 W of cooling per chip. This paper presents a detailed overview of the fabrication, assembly, testing, and reliability qualification of this advanced MCM technology.

1. Introduction

From the very early stages of the physical definition of the Regatta multichip module, or MCM, the module design approach was driven by the goal of achieving aggressive system objectives and providing the optimum and most cost-effective module performance. This required the development of new module attributes and, for each new attribute, the associated plans and actions for its introduction into manufacturing and scale-up production. The new attributes developed include the following:

1. The high-performance glass-ceramic (HPGC) MCM, which required new surface-finishing methods and structures for interconnection to both chip and capacitor pads at feature sizes as small as 200- μm

©Copyright 2002 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/02/\$5.00 © 2002 IBM




Figure 1

Cutaway assembly view of the multichip module (MCM).

pad-to-pad pitch for controlled-collapse chip connection (C4) solder contacts.

2. An advanced thermal cooling solution which could support chips at 150 to 160 W each.
3. A new land grid array (LGA) spring connector (beyond that available in the industry at the time) for more than 5100 module input/output (I/O) connections at an aggressive 1-mm pitch.

The advanced MCM which resulted interconnects four POWER4 ultrahigh-performance CMOS chips which operate at up to 1.3 GHz in the IBM high-end p690 UNIX** class eServer. The POWER4 four-chip MCM provides copper interconnection for 12 unidirectional, 7-byte-wide address/control buses that operate at half the processor clock rate and support a communication bandwidth of 41.6 GB/s between POWER4 chips. The MCM is the fundamental processor building block; it provides 28072 C4 solder connections, utilizing 200- μ m C4-to-C4 pitch for the four copper chips, 190 meters of copper wiring within the package for interconnection between chips, capacitors, and the 5184-module I/O for connection to the board.

The module also provides a mechanical platform to support the adhesive thermal interface (ATI), with high-thermal-conductivity silicon carbide (SiC) heat spreaders which are placed between the chips and the copper cooling hat to provide a thermal path to cool the 624-W module. A thin adhesive bond line provides the thermal connection from chips to SiC, and the IBM advanced thermal compound (ATC) 3.8 is used between each SiC heat spreader and the copper top plate, or hat. **Figure 1** shows a cutaway view of the IBM multichip module showing the copper hat, the SiC heat spreaders, the IBM

high-performance glass-ceramic substrate, the LGA spring connector, and the IBM advanced board and precision alignment pins for assembly. The expanded view of the MCM, the LGA spring interconnection, and the advanced board shows the level of interconnection needed to support this system solution. The module, LGA connector, and board are rigidly held together with nearly half a ton of force using a simple mechanical structure to ensure the integrity of the 5184 individual spring LGA connections for reliable field operation.

This paper begins with a brief system overview and then provides more insight into the development, fabrication, assembly, and reliability qualification of this leadership module, which is currently in volume production.

2. System overview

The POWER4 four-chip MCM is the fundamental processor building block of the IBM high-end p690 UNIX-class eServer. This system was announced in 2001 and has been generally available since December 2001. Systems based upon single central electronic complex (CEC) drawers are available from 8-way up to 32-way systems with processors operating at 1.1 GHz or 1.3 GHz. These CEC drawers are designed to sit within a 24-inch rack. The rack contains a bulk power adapter, I/O, direct-access storage device (DASD), battery backup, and/or switch drawers connected to the CEC I/O ports via cables to provide greater capacity and accessibility. The drawers can be interconnected further to create large computer complexes encompassing 512 or more processors.

A backplane and power distribution board bisect the CEC drawer so that the power regulators and system I/O are available from the rear of the rack. The blowers and memory cards are available from the front of the rack. This allows easy access to the $N + 1$ dc adapters and blowers for maintenance, upgrade, and servicing while the unit remains online. Although not accessible for concurrent maintenance, the MCM and L3 cache modules mounted directly on the backplane can be accessed, removed, and replaced without removal of the backplane from the drawer or the drawer from the rack.

The logical organization of the system begins with IBM CMOS chips, which feature two POWER4 processors per physical chip. These chips, implemented in the IBM 0.18- μ m, seven-layer copper metallization technology, contain more than 170 million transistors. Each POWER4 processor is an out-of-order superscalar design with eight execution units: two fixed-point, two floating-point, two load/store, a branch unit, and an execution unit to perform logical operations on the condition register. Instructions can be issued to each execution unit every cycle, though the maximum instruction retirement rate is five per cycle. Associated with each processor is a 64Kb Level 1 (L1) instruction cache and a dual-ported store-

through 32Kb L1 data cache. Four such chips are mounted per MCM, and up to four MCMs onto a backplane [1].

The processors are interconnected within the CEC by a distributed crossbar switch providing low-latency interconnections. By distributing the crossbar function over the chips housing the processors, the total chip count (and hence cost) is kept to a minimum, and overall reliability is enhanced. There are three kinds of processor crossbar switch interconnects: 1) Within a chip, the two processors are connected to an on-chip L2 cache and fabric controller; 2) from chip to chip on the MCM, each fabric controller interconnects each chip with each of the other three chips; and 3) from each chip on the MCM to an equivalent chip on each of the other MCMs, a logical connection is made by using one of four unidirectional physical loops through the board.

The chip-to-chip communication pattern on the Regatta MCM is schematically illustrated in **Figure 2**. The four chips on the MCM are interconnected by 12 unidirectional, source-synchronous, 7-byte-wide address/control buses that operate at half the processor clock (p-clk) rate. The data transfers among these chips are accomplished on 12 unidirectional, source-synchronous, 16-byte-wide data buses which are physically configured to connect orthogonal adjacent chips on the MCM. Of the 12 data buses, eight provide logical communication paths in both directions between the orthogonal adjacent chips, and four provide a physical ring structure for transfer of data between chips which are placed diagonally with respect to each other on the MCM. At the 1.3-GHz p-clk rate, there is 41.6 GB/s of bandwidth between POWER4 chips on the MCM.

The high chip-to-chip data transfer rate achieved in these connections on a given MCM is in large part due to the exceptional clock performance obtained in this design. A central clock card provides low-skew, independent differential-load-terminated clock signals to each of the chips on all MCMs. Phase-lock loops (PLLs) in each dual-processor chip multiply the clock signal from the clock card to derive the on-chip p-clks. Independent-card-mounted LC filters, coupled with RC filters distributed between on-MCM capacitors and on-chip resistors and capacitors, provide low-noise power sources to these PLLs so that jitter is minimized. Care was taken in the card and module designs to make sure that ground connections to the filter capacitors were implemented so that no package delta-I noise shows up across the supplies to the PLLs at the chip circuits.

All of the functional off-MCM buses, including the MCM-to-MCM, the MCM-to-L3/memory, and the MCM-to-I/O-system paths, are implemented using synchronous wave pipeline elastic interfaces. All signals on these buses, other than the clocks, are unidirectional, single-ended, and source-terminated. Each receiver is capable of inserting a

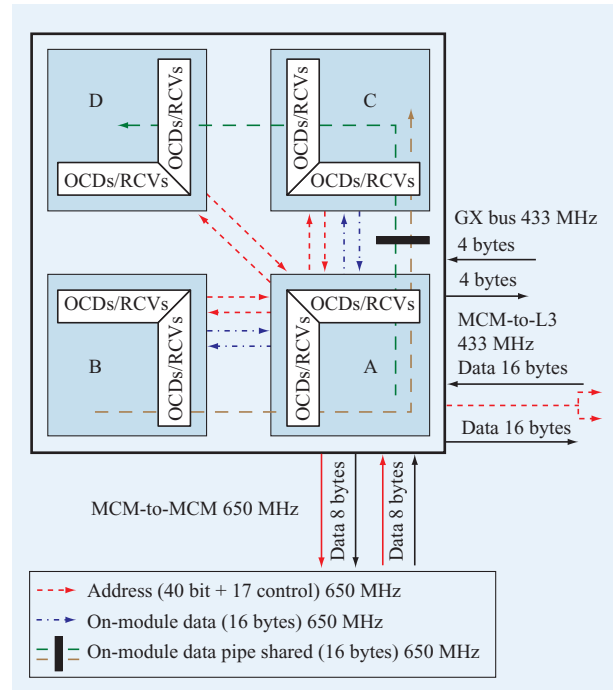


Figure 2

Regatta MCM chip communication pattern. Reprinted with permission from [2]; ©2001 IEEE.

delay into the bitstream so that all data-valid windows of bits in a bus are aligned for clocking into latches. The clocks are unidirectional, differential, and differentially terminated. Careful attention was paid in design to building a 50-Ω path from the source end of each net to the load. Typical module substrate technologies, which yield a characteristic impedance close to 30–40 Ω for slower, more general-purpose applications, have been optimized to closely match the rest of the 50-Ω path of the cards and connectors. Continuity of reference planes and proper return path decoupling are ensured at the connector boundaries. In those cases where signal references differ between source and load, capacitors placed within module packages provide decoupling required for signal return current paths.

3. Packaging overview

IBM-manufactured high-performance glass-ceramic (HPGC) MCMs are used to achieve the high bandwidth and low latency required by the p690 system. **Figure 3** shows an expanded cutaway view of the module. Several of the IBM mainframe Enterprise systems have utilized glass-ceramic technology in past years [3, 4]. Like these earlier systems, this advanced MCM was based on glass-ceramic technology to take advantage of its low dielectric constant




Figure 3

MCM package construction. Reprinted with permission from [2]; ©2001 IEEE.




Figure 4

Regatta MCM cross section. Reprinted with permission from [2]; ©2001 IEEE.

of 5.1 compared to that of alumina, which is about 9.8. However, these past systems had more than 20 chips per module and also used MCM-D technology in which sintered copper wiring and copper/polyimide multilayer thin-film technology was used to provide the chip-to-chip interconnection [3, 5, 6]. The current MCM, on the contrary, accomplishes all of the chip-to-chip interconnections in the glass-ceramic with the use of a unique 4-mil pad size on 8-mil pitch, area array C4 footprint. When placed on the MCM, the chips are rotated 90° relative to one another. This allows a separation of the horizontal and diagonal communication between the chips and optimizes performance. This special pattern was made possible because of the capability of the glass-ceramic technology to permit a deep stacked-via structure. **Figure 4** shows the via density and via depth in a cross-section view.




Figure 5

MCM trace pattern on upper layer.

The MCM wiring traces that enable the chip communication pattern for Chip A in **Figure 2** are depicted in **Figure 5**. The wiring in **Figure 5** shows interconnections between chips and contains wiring parallel to the mesh (in the x, y directions) and diagonal to the mesh. To control the impedance and crosstalk, the reference structures are implemented as mesh planes with a 200- μm pitch in the direction of the parallel interchip traces. The diagonal wiring is referenced to a specially constructed mesh plane structure, as shown in **Figure 6**. This feature keeps the impedance level of the diagonal wires matched to that of the parallel wiring.

In addition to the breakthrough substrate design described above, the MCM uses the innovative land grid array (LGA) instead of the traditional pin grid array (PGA) for module I/O. With the smaller I/O pad pitch, 1.00 mm instead of 2.54 mm interstitial, higher-I/O-count modules can be achieved in a much smaller substrate size, thus lowering package costs.

4. Glass-ceramic substrate technology

Ground rules and substrate process overview

The 85-mm² MCM substrates used in the Regatta system are made from patented IBM cordierite glass-ceramic dielectric and contain copper internal circuitry. Substrates are fabricated to the feature ground rules given in **Table 1** using the multilayer ceramic (MLC) substrate technology pioneered by IBM in the early 1970s [7]. The process, outlined in **Table 2**, begins with the combining of alumin-

silicate glass particles with organic materials to form a slurry which is cast into sheets using a doctor blade technique. A blanking operation then forms individual greensheets from the cast material. Each greensheet is punched to form small circular holes, called vias, that will eventually become the connections between wiring paths in different layers. In the screening operation, a Cu metal-organic paste or ink is forced through a stencil to simultaneously form a circuit pattern on the greensheet and fill the vias in it. Generally, each of the 70 layers that make up the MCM substrate has a unique pattern. After all of the greensheet layers are precisely stacked, moderate heat and pressure are used to coalesce the organic matrix into a laminate. The laminates are subjected to the sintering operation, in which the organic materials are burned off and the Cu-based circuitry and glass particles are co-densified. The IBM sintering process maximizes the positional accuracy of the top-surface and bottom-surface features. Once densification is reached during sintering, additional thermal processing crystallizes the glass-copper composite into a cordierite-based multilayered ceramic with co-fired copper interconnects. After sintering, the substrates are sized to their proper dimensions, the surface features are plated, and an electrical testing protocol is carried out in order to ensure the electrical functionality and reliability of each unit.

For the current MCM substrate, enhancements were made to the punching, screening, stacking, and sintering sectors to facilitate the formation of the smaller circuit features, control the dimensions of the vias making up the 200- μm reduced-pitch chip footprint, and control the coplanarity of the substrate I/O pads required for the LGA application. The absence of thin-film structures on the top and bottom of the substrate also required the implementation and qualification of new, top-side (TSM) and bottom-side (BSM) metal structures compatible with the co-sintered MLC technology.

Dimensional control

Controlling the dimensional tolerances of the MCM substrate presented some unique challenges compared to glass-ceramic/Cu substrates used in previous IBM servers. In order to reduce the cost of the substrate, planarization to control top- and bottom-surface flatness and thin-film I/O pads to optimize via positional accuracy were not employed in this application. Specifications for TSM flatness and via locations, along with those for BSM I/O pad true position and coplanarity, had to be met by the nonplanarized, as-sintered surfaces. Dimensional control in the glass-ceramic/Cu system was achieved in sintering and required other processing variables, including greensheet properties, paste properties, and lamination conditions, to be controlled prior to sintering [3, 8].




Figure 6

Diagonal mesh for impedance control.

Table 1 Circuit feature dimensions.

Feature	Dimensions
Line width	60–66 μm
Line spacing	134 μm
Wiring length	190 m
Smallest via diameter	89 μm
Largest via diameter	102 μm
Via spacing	100 μm
Vias per layer	~25 K
Vias per substrate	~1.7 M

Table 2 MLC process flow.

Preparation of raw materials
Greensheet casting
Greensheet blanking and curing
Screening
Via punching
Metal paste preparation
Inspection
Stacking
Lamination
Substrate sizing
Sintering
Plating
Electrical test

Dimensional control in sintered substrates requires that all surface features (vias and I/O pads) be located within specified limits around their designed positions. In MLC

Table 3 Dimensional tolerances.

<i>Dimensional property</i>	<i>Regatta glass-ceramic substrate</i>
TSM camber	<75 μm
I/O pad coplanarity	<40 μm
I/O pad true position	<0.20 mm

substrates, deviation of these features from their ideal location is caused by a radial expansion or contraction of the feature pattern from its center. Therefore, the larger the pattern area, the more deviated from their design position the outside features may be. The amount of expansion or contraction is also influenced by the amount of metal loading on the greensheets and the interconnect pattern density. The MCM substrate consists of four approximately 20-mm \times 20-mm chip sites with 90- μm vias on the TSM. The BSM pattern comprises a 76-mm \times 76-mm array of 5184 I/O pads of 0.7-mm diameter. This required much tighter expansion/contraction control on the BSM than the TSM because of the larger BSM footprint area. Furthermore, to meet the required true position specifications for the I/O pads, the substrates were sized after sintering.

Besides the feature location specifications, the MCM substrate also had very challenging tolerances with respect to TSM flatness and BSM pad coplanarity. The TSM flatness requirement was governed by the tolerances of the thermal solution, while the BSM pad coplanarity was determined by the LGA connection requirements. Normally, either of these requirements would be very difficult to meet on an unplanarized MLC substrate; however, enhancements to the IBM patented conformal sintering process [9] for the glass-ceramic enabled these tight requirements to be met. **Table 3** shows the typical dimensional control obtained on the MCM substrate used in the current system. As is seen later, these tolerances were adequate to meet the thermal, chip-attachment, and LGA needs.

TSM and BSM structures

The chips are attached to the TSM structure by solder bumps and hence require a metal surface that is wettable by liquid solder. To maintain high reliability, there should be no voids in the solder joints. The I/O pads on the BSM side have to meet the requirements for the second-level connector (LGA). The copper-based conductors on the TSM and BSM of the ceramic substrates have a higher coefficient of thermal expansion (CTE) than the ceramic. To mitigate any effects stemming from the CTE mismatch between the metal surface features and the glass-ceramic, novel TSM and BSM structures were used with the MCM

substrate. These structures replace the thin-film metal features and intermediate polymer layers used on previous substrates.

5. C4 flip-chip joining, rework, and underfill

Flip-chip devices have been used in IBM high-end systems since the early 1980s [7]. When first introduced, the bipolar devices typically contained a total of 120 to 200 C4 solder ball connections, with a C4 pitch of 250 μm . These early low-I/O-count devices were attached to a multilayer ceramic substrate made of alumina. Since the introduction of ES/9000* in 1990, IBM has been using glass-ceramic MCM modules with copper internal wiring for high-end servers [3, 4]. In addition, the glass-ceramic MCMs used in previous high-end servers have utilized thin-film-deposited wiring on the top surface of the substrate. The introduction of the Regatta system posed significant challenges in chip-join, chip-rework, chip-encapsulation, and module-sealing technologies because of the following new requirements:

1. Joining a very large die with 7018 solder bumps on a 200- μm C4 pitch to a glass-ceramic substrate without planar thin-film-deposited top-surface metallurgy (TSM).
2. Developing a new plating metallurgy to cap the sintered copper vias that ensures good solder wettability without joining defects.
3. Improving chip placement accuracy and cleaning to enable high yield joining in bond and assembly (BA).
4. Developing a rework process for chip removal and replacement.
5. Underfilling the large die with a newly developed encapsulant material.

In the sections below, details of the process and materials improvements developed to meet these challenges are discussed.

Chip join

At approximately 20 mm on a side, the Regatta chip area of 427 mm² is more than ten times greater than the area of the first bipolar chip used on glass-ceramic packages with thin-film wiring in 1990. Similarly, at 7018 solder bumps, it has more than ten times the number of solder balls per chip. A micrograph of the Regatta substrate surface is shown in **Figure 7**. The high-density region of the C4 array is confined to an L-shaped section that is 15 rows deep from the diced chip edge. This region, which contains most of the signal nets, is on a 200- μm pitch compared to the balance of the array, which is on a 400- μm interstitial pitch. (The closest C4 pitch on a diagonal is 284 μm .) The chips are assembled in a rotated fashion

such that the L-shaped regions face each other in order to provide the shortest signal nets between devices. The optics on the chip-placement tool were improved from 20- μm to 12- μm accuracy to enable joining of these fine-pitch devices. Mounted around each chip site are 28 C4-terminated decoupling capacitors (decaps), which use standard 250- μm -pitch solder balls [10]. Chips and decaps are joined to the module simultaneously in a furnace reflow cycle and go through a post-reflow flux cleaning.

Figure 8 provides a visual perspective of the higher density of the fine-pitch C4 array used in Regatta. It compares C4 arrays at three different pitches: 250 μm , 225 μm , and 200 μm . On Regatta, the 100- μm -diameter C4 ball has a nominal solder volume of $70 \times 10^4 \mu\text{m}^3$; this is 44% lower than that of a 125- μm -diameter C4 ball on a 250- μm pitch, having a volume of $125 \times 10^4 \mu\text{m}^3$. The reduced volume results in a significantly lower solder joint height after the device is attached to the substrate. Typical 125- μm -diameter solder balls have a joined height between 95 and 110 μm , while 100- μm solder balls have joined heights only between 70 and 85 μm . There is a greater than 50% increase in connections/ cm^2 for a 200- μm -pitch C4 array compared to a 250- μm -pitch array: 2500 solder bumps per cm^2 compared to 1600 bumps. One consequence of tighter pitch is difficulty in flux cleaning after chip joining. The lower joined height of about 75 μm , along with a large increase in chip size, required flux cleaning optimization in the manufacturing line to ensure removal of flux residues after the reflow operation.

Figure 9 shows a process flow for the bond, assembly, and test operations for the MCM.

The Regatta glass-ceramic substrate has a nonplanar thick-film top- and bottom-surface metallurgy which is different from those used in previous glass-ceramic substrates. The previous thin-film top-surface metallurgy consisted of Cr/Cu/Ni/Au which, in the first generation of glass-ceramic technology, was deposited by electron-beam evaporation [3, 4]. More recent glass-ceramic modules have used an electroplated Cu/Ni/Au TSM with a sputtered Cr/Cu seed layer [11, 12]. The glass-ceramic substrate in Regatta, in contrast, has a Ni/Au metallization which is plated directly on sintered Cu/Ni C4 vias or microsockets.

The plating, which is done by an electroless process, deposits 3 to 5 μm of Ni, followed by an 80–120-nm layer of Au. This metallurgy is different from both the thin-film metallurgy previously used on glass-ceramic and the Ni-B/Au electroless plated metallurgy used in past IBM alumina substrates. Significant development work was done to optimize the plating processes and post-plating Ni/Au diffusion to ensure good wettability of the low-tin solder (97 wt.% Pb, 3 wt.% Sn) to this metallurgy.




Figure 7

Top surface of MCM substrate.




Figure 8

Images of C4 solder balls on pitches of (a) 250 μm , (b) 225 μm , and (c) 200 μm at constant magnification.

With respect to the process flow in Figure 9, removal of a defective chip is routinely used to achieve high multichip module yield. Typically, the module is qualified to 13 chip reflow cycles at a peak temperature between 345°C and 375°C to allow an initial chip join, followed by up to four individual chip rework cycles.

Excellent C4 solder wettability has been demonstrated for the Ni/Au metallurgy capping the Cu/Ni vias co-sintered in glass-ceramic substrates. Wettability is verified by tensile pull removal of chips after joining to the MCM substrates and performing fracture surface inspection of the separated parts. Both a minimum pull strength and an allowed fracture mode are specified. The fracture should preferably occur in the C4 solder balls with no interface

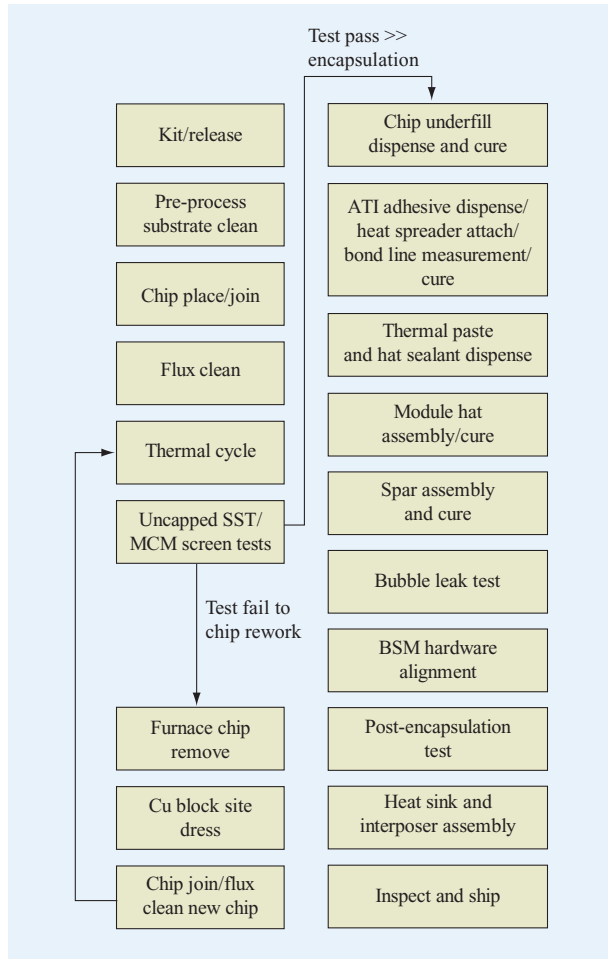


Figure 9

MCM assembly, test, and rework process flow.

delaminations or solder voids. Typical wettability pull strength data after initial join, multiple reflows, and chip rework is shown in **Table 4**. The wettability data is equal to or better than that seen on thin-film metallurgy used in previous glass-ceramic MCMs. The dense structure of sintered Cu/Ni vias, along with the generally defect-free quality of the plated Ni/Au layer, limited the formation of both large voids and clustered small voids in the C4 joints [13]. Large voids covering 50% or more of the via of the ball-limiting metallurgy (BLM) on the Si chip can form because of the trapping of flux during solder wetting of the TSM pad.

Finally, one of the concerns associated with using a Ni metallization in the C4 microsocket is the formation of a Ni intermetallic phase after prolonged exposure at elevated temperatures. The Ni intermetallic compound observed presents a non-wettable surface, and can lead to localized “non-wets” in tensile tests used to assess

Table 4 Regatta typical chip wettability data.

<i>Wettability data, following:</i>	<i>Pull strength (lb) (Min. spec. = 238 lb)</i>
Initial chip join (IJ)	332
IJ + 5 reflows	311
IJ + 12 reflows	306
1X rework*	324
4X rework*	335

*Each rework cycle consists of three furnace reflows.

solder wettability. The wettability data after 13× solder reflows (confirmed by microstructure analysis) showed that there was no significant formation of Ni intermetallic compound in the plated metallurgy.

Chip rework process

Since the introduction of multichip module technologies in the early 1980s, IBM has used individual chip rework (ICR) processes to enable individual removal of a defective chip and replacement with a new device. A chip may have to be removed because it has a joining defect such as signal I/O opens or shorts between power and ground or between I/O and power C4 balls. Typically, the rework driven by device joining defects is quite low (less than 1%). Most chip rework is done to optimize module function when module-level testing reveals non-optimum timing between chips. All CMOS chips used in server MCMs are burned-in and tested at operational speed in a single-chip carrier using a process known as temporary chip attach carrier (TCA).

The ICR process has evolved over time and has been described in detail elsewhere [14, 15]. The three basic steps of ICR are individual chip removal, dressing of the TSM chip site with a sintered porous Cu block, and joining a new chip using the same furnace reflow process used for initial chip join (Figure 9). Because the chip size has grown from about 40 mm² used in the first generation of MCMs to more than 400 mm² in CMOS-based MCMs, the infrared method used to locally heat the back of the chip and remove the chip with molten C4s using a vacuum tip was found to be inadequate. The current removal process, known as the “hot-spider” process, utilizes a bimetallic disc which is attached to a gripper fixture. The gripper fixture is assembled so that one or more defective chips can be removed simultaneously by subjecting the MCM to a furnace reflow cycle. The bimetallic disc is designed so that it exerts a lifting force on the gripper and removes the chip from the module when the C4 solders are molten above the liquidus temperature of about 320°C. An appropriate gripper fixture and bimetallic discs were designed for the 427-mm² Regatta die with 7018 C4 balls. It was found that these devices could easily be removed in a furnace reflow using the modified hot-spider process.

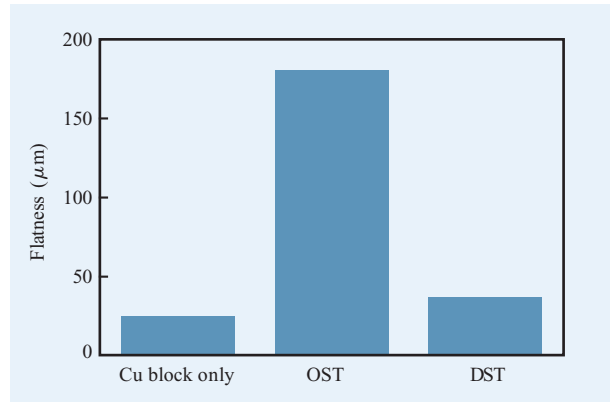


Figure 10

Effect of tinning on copper blocks used for chip site dress.

It was determined that the porous Cu block had to be optimized to properly dress the chip site on the module. Site dressing is required to remove residual solder on the TSM pads after a chip is removed. Without site dressing, solder can build up after one or more chip rework cycles, and can cause increased and uneven solder volume in the C4 joints. Increased solder volume may result in shorts between adjacent solder balls, especially for a tight pitch of 200 μm. Uneven solder height can lead to chip tilting, which can interfere with achieving uniform thermal contact between the heat spreader and the cooling hat, as described in a later section. Cu blocks used in site dressing typically have the side facing the chip site tinned (one-side tinned, or OST). The tinned surface soaks up the residual solder during furnace Cu block dressing and prevents the transfer of Cu particles from the block to the dressed site. One disadvantage of the OST blocks is the warpage associated with tinning one side only. Cu blocks used for the Regatta chip site are 21 mm × 21 mm and utilize an enhanced double-side tinning (DST) process after sintering to prevent warping of the block.

Data showing the flatness of Cu blocks with one-side and two-side tinning, along with as-sintered and coined Cu blocks, is shown in **Figure 10**. It is clear that for a Cu block of such large area, OST leads to unacceptable warpage (greater than 150 μm). The residual solder height after Cu block dress should be less than about 35 μm. OST Cu blocks leave the middle of the large chip site generally undressed because of their warped shape. DST Cu blocks have flatness comparable to as-sintered Cu blocks, and do an excellent job of removing residual solder from the reworked chip site. Typical solder heights, after multiple chip rework and site dress cycles, were measured to be in the range of 5 to 20 μm. Micrographs of an area of the microsockets from a chip site are shown

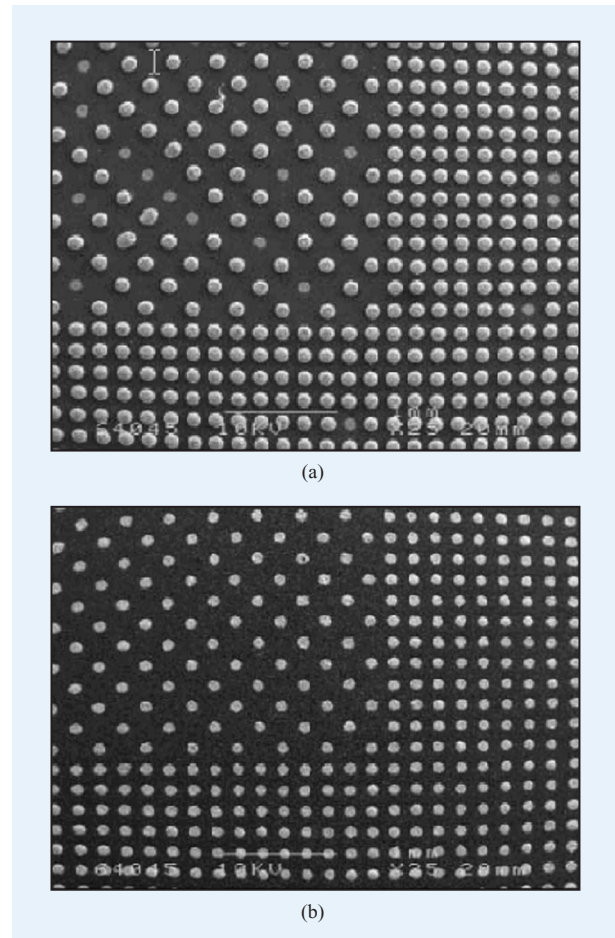


Figure 11

Micrographs of a chip site (a) after chip removal; (b) after furnace Cu block dress.

after chip removal [**Figure 11(a)**] and after Cu block site dress [**Figure 11(b)**]. As discussed before, the wettability of a new chip joined to a dressed site remains excellent after multiple chip rework cycles. Chip join and up to four chip rework processes have been successfully qualified for the large Regatta die using the enhancements described above.

Underfill process

After the four chips and the capacitors have been joined to the Regatta module, an organic encapsulant material is dispensed in a line along the edges of each chip. The underfill material consists of a thermoset organic resin with inorganic filler dispersed in the resin. The encapsulant fills the gap between each chip and the glass-ceramic substrate, and fully encapsulates the C4 solder balls. After the underfill has been dispensed, the modules are placed in a box oven to cure the underfill encapsulant




Figure 12

(a) Micrograph of a cross section of the MCM at the interface between the chip and the glass-ceramic module. (b) Magnified image of MCM cross section showing solder connections to Si die.

prior to attachment of the heat spreader and other subsequent process steps, as shown in the process flow in Figure 9.

The Regatta module posed several challenges with respect to underfill because of the large chip size (>20 mm on an edge), very high C4 count, and small chip-to-substrate gap height (approximately 70–85 μm). The underfill was optimized to address these challenges. In particular, the underfill developed for this application was able to meet the requirements for flow characteristics and adhesion to both the glass-ceramic substrate and the polyimide passivation on the chip. The new material exhibits excellent wetting and flow characteristics and viscosity low enough to be able to completely fill the small gaps between the large, high-I/O chips and the MCM substrate.

Acoustic microscopic images of the Regatta modules showed that the new, improved underfill exhibited no

Table 5 Typical properties of underfill materials.

<i>Cured properties</i>	
CTE	25–30 ppm/ $^{\circ}\text{C}$
Glass transition temp. (T_g)	$\geq 130^{\circ}\text{C}$
Elastic modulus	9 GPa @ 25 $^{\circ}\text{C}$
Water absorption	<0.5% by wt.

chip-to-underfill or underfill-to-substrate delaminations after the modules were stressed. This was further confirmed by cross-sectional analysis. After Group B stressing (1200 g mechanical shock and 20 g sinusoidal vibration), the underfill showed no delaminations, demonstrating excellent adhesion.

Figure 12(a) is a low-magnification cross section of an assembled Regatta module. The figure shows lead–tin C4 solder balls joining the silicon die to the copper vias in the high-performance glass-ceramic substrate. The underfill material fills the gap between the chip and substrate, completely encapsulating the solder balls. The SiC heat spreader, which is attached to the chip by means of a conductive epoxy, is also shown. **Figure 12(b)** is at higher magnification, showing three encapsulated C4 balls. The filler particles can be seen. The gap between the chip and substrate in the photo is approximately 75 μm .

In the past, IBM produced hermetic packages with small die sizes (<10 mm), and no underfill was needed to ensure acceptable C4 reliability [16]. As nonhermetic packages with larger chip sizes were designed, underfill materials were utilized to enable structural coupling of the chip and substrate, effectively decreasing the shear stress sustained by the C4 balls and thus lowering the applied strain on the solder joints [17]. This strain was due to the mismatch in thermal expansion between the silicon chips [coefficient of thermal expansion (CTE) of 2.6 ppm/ $^{\circ}\text{C}$] and the substrates (alumina substrate materials have a CTE of 6–7 ppm/ $^{\circ}\text{C}$, while organic packages have CTEs in the 17-ppm/ $^{\circ}\text{C}$ range).

In the case of the Regatta module, a high-performance glass-ceramic substrate with a thermal expansion coefficient which is essentially matched to that of the silicon die (CTE of HPGC substrate = 3.0 ppm/ $^{\circ}\text{C}$) is used. Hence, there is actually minimal strain on the solder balls due to CTE mismatch. The underfill in Regatta serves primarily to provide environmental protection of the C4s, as well as providing overall mechanical rigidity of the structure and electrical isolation.

Table 5 shows several typical properties of underfill materials [18]. Key parameters of underfills are low viscosity, high glass transition temperature (T_g), high modulus, moderately low thermal expansion coefficient, inorganic filler of proper particle size that is adequately

dispersed in the underfill resin, and superior adhesion to the chip and substrate surfaces.

Figure 13 shows a comparison of several properties of the new underfill used in Regatta and the standard underfill material that was in use previously. The new underfill showed faster flow time (under 17-mm test chips), shorter cure time, and improved adhesion strength after exposure to stringent JEDEC¹ Level 1 stressing of 168 hr at 85°C/85% relative humidity, as measured by an IBM-developed adhesion testing method [19]. Standard alumina ceramic test vehicle substrates using the new underfill have gone through 10000 deep thermal cycles of -55°C to 125°C with no electrical fails observed.

The bond and assembly advances implemented for handling the fine-pitch, high-I/O Regatta devices in the areas of chip placement, joining, and flux cleaning, as well as the improved underfill encapsulant material used, make possible the very high reliability achieved in the C4 interconnections.

6. Module sealing

After the application and cure of the underfill material, the next step in the Regatta MCM assembly process is the attachment of individual silicon carbide heat spreaders to each of the four processor chips on the MCM. The thermal and mechanical joint between the chip and heat spreader is accomplished with a thermally conductive adhesive. A key aspect of this thermal joint is a thin and consistent bond line, which is guaranteed through several manufacturing process controls and metrology. Foremost among these are a measurement process developed to ensure an acceptable bond line before the heat spreaders are cured into place and other controls designed to ensure the cleanliness of the chips and spreaders during the assembly, which is performed in a stringently controlled assembly room.

Before attachment of the heat spreader, each module with a serial number is placed in an optical measurement tool that records the height of each of the four processor chips. This data is stored for comparison later in the process. Next, both the spreaders and chip surfaces are inspected under oblique lighting to check for particulate contamination. Any contamination found is removed using a clean nitrogen blow-off technique. The epoxy material is then applied to the individual chip surfaces using an automated dispensing tool to accurately control the dispensed pattern and quantity of material. Specialized fixtures are used to locate and hold the spreaders against the chips during curing in an oven. Before the curing operation is begun, the pre-cured assembly is returned to the measurement tool to record the height of each heat spreader, with the data being compared to the previous chip height measurements. Any measurement anomalies

¹ JEDEC: Joint Electron Device Engineering Council.

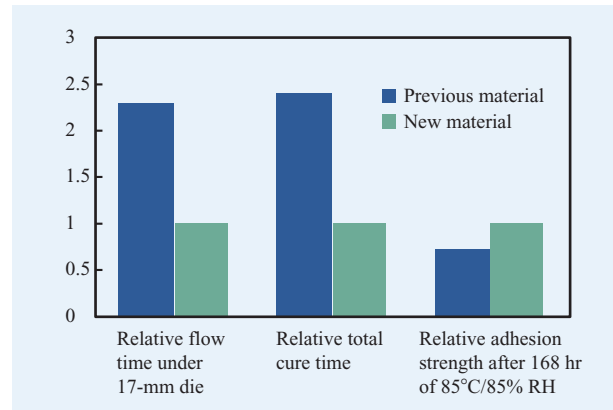


Figure 13

Relative properties of new underfill used in Regatta MCM compared to previous-generation material.

(height, tilt, etc.) are called out to the manufacturing operator, who can perform a process to remove the pre-cured assembly from the fixture, clean off the epoxy, and repeat the heat-spreader attachment process. While fallout at this step is rare, this process is essential to prevent the loss of an expensive module for a heat-spreader bond line which is out of specification.

Concurrent with the attachment and curing of the heat spreaders is the attachment of a small circular metal bearing plate to the center of the substrate. This bearing plate is used in conjunction with the adjustable center post, or “spar,” described later in the encapsulation process. The spar is used to provide added mechanical stability to the assembly and to limit substrate flexure under LGA loading, described in the next section. The spar is put in place after the copper cooling hat has been attached to the carrier. The curing profile has been optimized to provide good bond strength for both the heat spreaders and the bearing plate. The concurrent joining process also reduces the manufacturing cycle time.

The remaining elements of the encapsulation process serve to attach the large copper lid, or “hat,” to the assembly above. To complete the thermal path from the heat spreaders to the hat, which completely covers the four heat spreaders and is joined to the outer top-side edge of the substrate, IBM advanced thermal compound (ATC) is used. The ATC is applied to the inside cavity of the hat using either an automated dispensing tool or a screening template, with the weight applied to each module hat monitored to ensure complete coverage between the heat spreaders and the hat. The silicone adhesive which joins the hat to the substrate is then applied using another automated dispensing tool to form the wide seal band around the perimeter of the hat. The




Figure 14

Components associated with LGA connection of MCM to board.




Figure 15

Schematic of system CEC assembly (shown configured with one MCM and four L3 modules in an eight-way configuration).

hat with both the ATC and the silicone adhesive is placed in a specialized alignment and compression fixture with the module and heat spreader subassembly. This fixture is then assembled to guarantee a thin adhesive bond line between the substrate and the hat. After the assembly is completed, the adhesive is cured. Once the seal has been cured, the spar set screw is inserted into the threaded hole at the center of the hat, adhesive is applied, and the spar is tightened to specifications. The assembly is then heat-cured to properly seal the spar. Finally, after the module is cooled from the curing step, it is leak-tested according to ASTM Test Method 883E for fluid tightness.

From here the module proceeds to LGA interposer alignment. This step involves mounting the alignment hardware on the cap and making adjustments needed to bring the interposer into position, in effect lining up the connector buttons with the I/O pads on the bottom side of the ceramic carrier. The LGA interposer is then attached to the module. This also prepares the assembly for mounting on the board, which is the next level of system packaging. To complete the assembly, a heat sink is mounted on the cap, and the module is prepared for LGA actuation/loading.

7. Second-level attachment using land grid array (LGA) technology

The decision to employ LGA technology to achieve module-to-board interconnection for the Regatta-H CEC imposed considerable challenges and constraints on the packaging technology selected for the MCM used in this system. It provides a useful example of how system performance strategy affects the choice of the components that make up the system, and how packaging technology is required to solve system-driven challenges.

LGA overview and performance metrics

LGA is a mechanical, and therefore readily separable, interconnect, and is one of several means for achieving second-level interconnect to provide the electrical pathway for signal and power that connects a module to a card or board. Alternative technologies include hard connections such as solder, ball grid array (BGA), and column grid array (CGA), and other mechanical connections such as pin grid array (PGA). All of these are “area array” connections, for which the I/O density is determined by the I/O spacing or pitch [20].

As shown schematically in **Figure 14**, the LGA connection is achieved by an interposer with embedded conductive elements, typically supported by an insulating carrier that is sandwiched between matching metallized I/O pads on the module and the card. An actuation hardware system must be employed to apply a normal force that holds this sandwich together, since it is only through this normal force that a low-resistance, high-reliability connection is maintained.

In order to provide a competitive server for the targeted market, system designers in the IBM Enterprise Server Group selected a flexible design strategy that would allow the same basic system frame to cover a range of performance levels from an entry-level eight-way parallel processing system to a maximum 32-way system, each operating at 1.3 GHz. In addition, the system had to provide expandable cache from one to sixteen units of 32 Mb each, and expandable system memory from 8 Gb to 256 Gb. This flexibility was not to be simply

a factory option, but was required to be field-upgradable so that a system could grow with a customer's specific needs.

To achieve this performance strategy, each CEC must allow the installation of one to four of the eight-way MCM processor modules discussed above, each requiring 5184 I/O connections, and one to sixteen 32Mb cache modules, each requiring 1247 I/O connections. Since these modules must be field-replaceable units (FRUs), the CEC requires approximately 40000 field-pluggable connections on a single large, complex backplane card. Density and module size limitations demand a 1-mm I/O pitch, and 1.3Gb performance demands high-performance glass-ceramic for the MCM. With an 85-mm MCM module, the only technology potentially able to meet this requirement is LGA. However, at 85 mm and 5184 I/O, this is the largest LGA array ever produced. At 40000 pluggable LGA connections, the backplane requirement is also a technology first, and the reliability per I/O requirement exceeds any previously established for an LGA. The layout of the CEC, with four cache modules and one MCM, is shown schematically in **Figure 15**. In addition, as discussed in Section 8, the 1.3-GHz speed requires high module power, which in turn requires innovation in thermal management that has further affected the LGA requirements.

Specific module requirements

Several key module requirements stem directly from the need to use LGA technology for the MCM, most significantly strength, pad integrity, and module flatness.

The requirement of considerable constant load on each LGA connection imposes a significant strength requirement for the module and the actuation system. Although reliable contact is achieved with loads of 50–100 g per I/O, the total module load, when multiplied by 5000 I/O, becomes quite high. Consequences of these large loads are discussed later in this section.

Mechanical interconnection is achieved by intimate metal-to-metal contact. Although most LGA connectors provide some element of “wipe,” whereby the mating surfaces are scrubbed by the connector elements, particulate or film contamination can interfere with the formation of reliable contact. In addition, since the connection is demountable, the pad metallurgy must be capable of withstanding a number of actuations without undue wear, and must be robust to environmental factors such as moisture and dust in the field.

Finally, because of the large active area and the requirement for uniform loading, the module must provide

a flat surface or high level of coplanarity between the metallized I/O pads. A significant amount of effort was focused on achieving tight control of flatness across a large-area sintered ceramic module, as described earlier. The consequences of an out-of-flatness situation, as observed during the learning stages of the program, are shown dramatically in **Figure 16**. In Figure 16(a), a contour map of the module topography reveals the presence of troughs between the quadrants of I/O pads. Early substrates showed recessed areas of the order of 100 μm . Imprints on pressure-sensitive film applied between the LGA interposer and the same module illustrate the resultant nonuniform loading of the individual conductive elements, as seen in Figure 16(b). Contact resistance measurements reveal high-resistance tails associated with the reduced contact pressure [Figure 16(c)]. When improvements in control of coplanarity were achieved, uniformity in pressure-sensitive imprints was seen, and concomitantly the contact resistance distribution was tightly controlled to within specifications.

The application of LGA technology to the Regatta-H system required additional significant enhancements in technology that will be discussed in future publications. Some of these are briefly mentioned here. The connector technology chosen for this application utilized a fuzz button contact, known for excellent reliability but periodically subject to button retention problems in assembly. An IBM patented process was developed that achieved a 10 \times improvement in button retention. The actuation hardware for applying high, consistent, and uniform loading at very high tolerances required a unique design based on extensive modeling and analysis. Finally, the card technology, card pad metallization, and ability to withstand loading required the development of unique card and card assembly technology.

The LGA connector required extensive reliability testing utilizing test vehicles that accurately simulated the card and module designs used in the real product. The final design exceeded all test requirements. A sample from one such test, thermal cycle stressing, shown in **Figure 17**, demonstrates the excellent durability of the LGA connector scheme. In Figure 17, the values on the y axis correspond to the percentage of the sample population having a resistance change less than the resistance change plotted on the x axis.

8. Modeling, test, and validation of MCM–LGA assembly integrity

Background

The Regatta module is the largest ever developed using the LGA-to-board attachment technology. The socket,

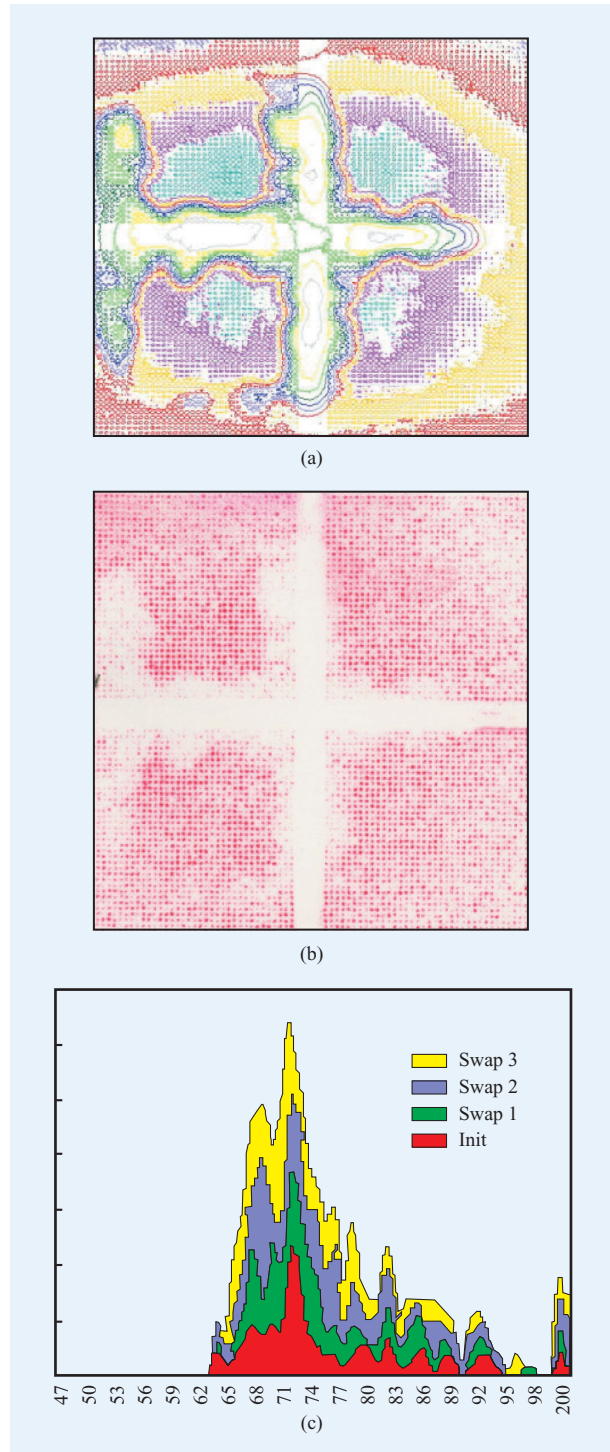


Figure 16

Effect of module flatness on LGA pressure and contact resistance uniformity: (a) Contour map of module (pad coplanarity of 92 μm); (b) pressure film imprint of LGA connector for same module; (c) contact resistance distribution (mΩ). The high-resistance tails of the distributions are correlated with the topography.

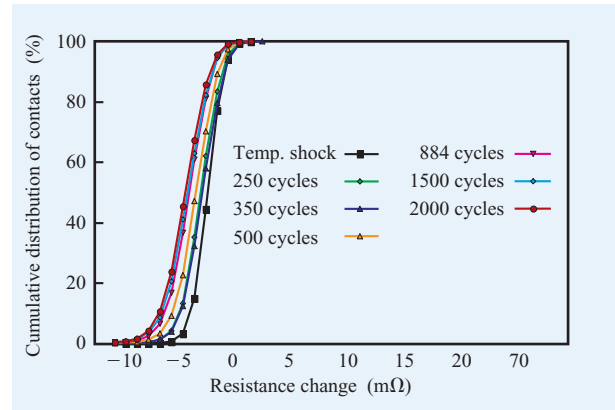


Figure 17

MCM-LGA temperature cycling stress test.




Figure 18

MCM socket interposer.

shown in **Figure 18**, comprises four separate 36 × 36 arrays of enmeshed wire buttons, each compressed when the module is socketed, resulting in a nominal retention force approaching 454 kg.

The retention force is transmitted through the module cover to the substrate along the peripheral seal band and is balanced by the total I/O reaction force distributed along the bottom expanse of the substrate surface. A simplified two-dimensional schematic diagram of substrate loading is shown in **Figure 19**. The noncollinearity of load and reaction forces tends to flex the substrate inward, giving rise to a number of potential concerns, including thermal paste squeeze-out, contact force reduction on the




Figure 19

Schematic diagram of mechanical loading on substrate from LGA attachment.

inner contacts, and substrate flexural stress. The last of these, which presented the greatest challenge, is the primary focus of this section. Stress was found to be further exacerbated by the underfilled chips and attached SiC heat spreaders which locally stiffen the substrate in flexure.

To accommodate the large size of the module while still limiting substrate flexure and its attendant concerns, a support has been introduced which couples the copper lid to the geographic center of the substrate. This post, or “spar,” reduces flexure and lowers stress by about 35%. To allow for product variability, the post is tailored to each module by threading it through a hole in the lid against the already-sealed substrate. To prevent the post from scratching or creating a point of high local stress on the substrate surface, a stainless steel bearing disc is first bonded to the substrate by an assembly process that was described in Section 6. The crown on the bottom of the threaded post, applied torque, bearing disc dimensions, and adhesive thicknesses have been optimized in order to minimize substrate stress, taking into account available substrate real estate and manufacturing tolerances. A photograph of the lid-and-post assembly is shown in **Figure 20**.

Stress modeling and strength evaluation

Stress under LGA loading was evaluated by a three-dimensional finite element model incorporating one quadrant of the module and retention hardware (**Figure 21**). Stress contours of a substrate without the module stabilizing post [**Figure 22(a)**] showed maximum stress on the top (chip) side just outside the stiffened regions, adjacent to the off-diagonal chip corners. With the post introduced to limit flexure, the computed stress [**Figure 22(b)**] was reduced by 35%. This model was used extensively, not only in the design process, but also to evaluate a multitude of hypothetical changes in design, material, tolerances, and loading conditions.




Figure 20

Photograph of cooling hat showing central support post and recessed “picture frame” sealing surface.




Figure 21

Finite element model showing one quadrant of module and retention hardware. For clarity, grid lines are not shown; die, flip-chip solder bumps, and underfill are recessed under heat spreader. (Modeling software from ANSYS, Inc.)

In parallel with the model, substrate strength was measured on a large representative sample of product substrates, using a biaxial ring-on-ring fixture. Pertinence of the numerical results was ensured by 1) testing actual hardware and 2) designing the support and load rings so that the region under tensile stress encompassed the entire region identified by the finite element model as having high stress in the LGA–MCM assembly.

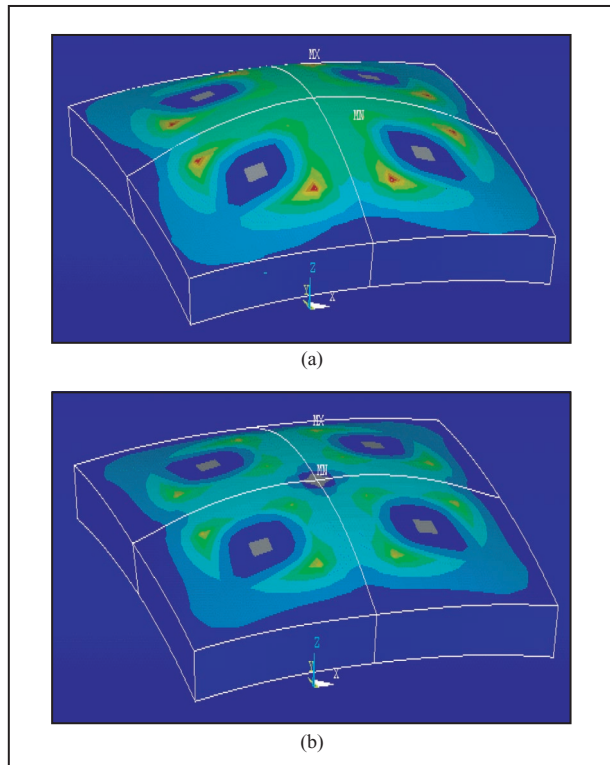


Figure 22

Stress contours on the top substrate surface: (a) Without module stabilizing post; (b) with post. Highest stress is indicated by red, lowest by blue. The same scale is used in (a) and (b). Distortions are not to scale. Stress peaks are just outside the off-diagonal chip corners. (Modeling software from ANSYS, Inc.)

Experimental validation of models

At this point in the evaluation, a critical experiment was carried out to test both the model and its comparison with experimentally measured strength. The approach employed was to assemble completed modules in actual retention hardware to the nominal actuation force, and then to continue tightening the retention screw to the eventual point of failure, thereby allowing comparison of measured overload capability to that predicted by the model.

For this experiment, modules were assembled in actual Regatta retention hardware. Force was measured by the number of screw turns, supported by instrumented posts and by pressure-sensitive measurement. Fracture was detected audibly and confirmed by dye penetrant testing of the lid that had been sheared off.

Modules were tested, with spar hardware and assembly, and without the central supporting post, or spar. Comparison to model prediction, summarized in **Table 6**, shows excellent agreement.

Long-term load ramifications

To complement the reliability assessment described in a later section, the analytical approach was expanded to better quantify resistance to stress cracking in the server operating environment. Two vital factors had to be considered: the impact of long-term load retention, and statistical variability and uncertainty.

A defining feature of module-to-board attachment using the LGA is that the retention force is sustained over the lifetime of the server. Thus, not only must the substrate withstand the flexural stress inherent in LGA attachment, but it must do so for many years. Ceramics share with most other brittle materials a slow crack-growth phenomenon known as stress corrosion cracking, created, as its name suggests, by both the applied stress field and the environment. This phenomenon is well-documented and described in the literature [21, 22]. If the slow-growing crack reaches a critical size, the part will fracture. A structural element susceptible to stress corrosion cracking and exposed to a sustained load must thus be designed to withstand an operating target stress which is lower than if the stress were applied only briefly or intermittently. An equation describing the more stringent long-term sustained stress target [20] is

$$S = S_0 [T_0 / T(n + 1)]^{1/n} \quad (1)$$

Here S is the modified stress target for the longer load retention time T ; S_0 is the sustainable stress (i.e., fracture strength, or “modulus of rupture”) when measured in ramp loading over time T_0 (for example, in the ring-on-ring test); and n is the exponent relating the velocity of slow crack growth to stress intensity. For ceramics, n is of the order of 30–40 and thus is the dominant factor in the equation. For example, if $n = 31.3$, $T = 7$ years, and $T_0 = 3$ minutes, the allowable stress level is reduced by 43%.

One common method of determining the exponent n is to measure strength at two or more test speeds, the theory being that at slower test speeds, slow crack growth will proceed just a little more, and the apparent strength will be slightly reduced. For sufficient differentiation, rates should encompass a range of at least two orders of magnitude. It may easily be shown that on a log-log plot of strength versus test speed, the slope is $(n + 1)^{-1}$. Employing this technique with our ring-on-ring test, we arrived at a best estimate of $n = 31.3$.

As a result, the predicted breaking force for a seven-year sustained load can be expressed as a percentage of the module rupture strength given in Table 6. Taking into account the actual test time and projected module lifetime, and employing Equation (1), the estimated sustained-load capability is nearly twice the nominal retention force required for acceptable LGA engagement and low contact resistance.

Statistical variability

Finally, it was necessary to consider various sources of variability and uncertainty and to incorporate them into the predictive process. In lieu of a simple safety factor, we used a more robust approach, by isolating, quantifying, and then combining the various sources of uncertainty and variability in a statistical algorithm. Product specifications on substrate thickness, out-of-flatness, and other dimensional parameters were taken into account, utilizing the stress model and translating the extreme values into 3σ limits. Variability in substrate strength and nominal socket force was also factored in. The stress increment from machine operation was added to the mix, as was an adder reflecting possible model inaccuracies. The final result was a prediction of one part per million product fallout due to substrate cracking over the projected operating lifetime of the Regatta FRU, a number that was well below target.

9. Module cooling challenges and solutions

Chip operating temperatures are a function of the efficiency of the total module cooling solution, which includes internal cooling (inside the module) and external cooling (how the heat is removed from the back of the module). Although liquid cooling is more efficient than air cooling, air cooling is less expensive, requires less infrastructure within the frame, and does not require the customer to supply chilled water. For the present server application, air cooling was chosen. This forced the internal cooling solution to be more efficient in order to maintain operating temperature limits that affect performance, functionality, and reliability.

The internal cooling of multichip modules (MCMs) is more challenging than that of single-chip modules (SCMs). In the past, IBM has used either spring-loaded pistons (as in the original TCM) or thermal compounds [7]. As chip size and power density have grown, thermal pastes have been used more extensively. Thermal paste cooling is effective for a range of chip powers and sizes, but as power density increases, on-chip thermal gradients may become significant. **Figure 23** shows how chip power density has increased over time, first for bipolar technology and more recently for CMOS technology chips.

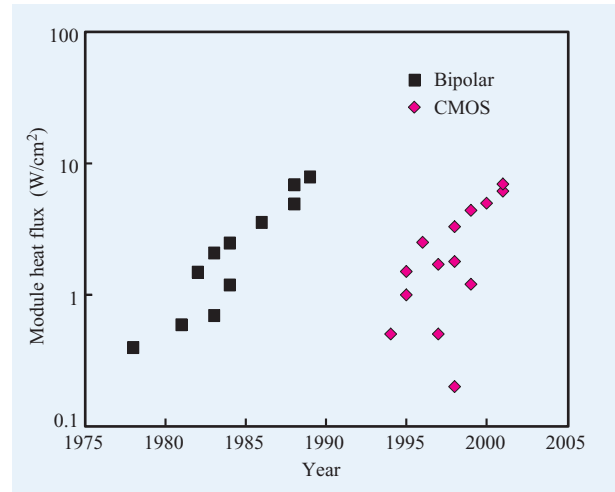


Figure 23

Heat flux explosion, comparing bipolar and CMOS chip technologies.

The Regatta MCM utilizes four high-power chips, each with two processors, and an integrated L2 cache. The unique chip floorplan (design) produces a chip power distribution that is highly nonuniform, as shown in **Figure 24**, with regions exceeding 100 W/cm^2 . To address the high power density, it is desirable to use high-thermal-conductivity heat spreaders that are thermally coupled with the chips in order to minimize the on-chip thermal gradients. A novel adhesive thermal interface (ATI) solution to this problem has been developed.

ATI involves bonding individual heat spreaders directly to the chips with an adhesive, as shown in **Figure 25**. The thermal resistance of the bond line is minimized by utilizing a very thin layer of a thermally conductive adhesive. In this application the effective thermal conductivity of the epoxy was found to be about $1.23 \text{ W/m}\cdot\text{K}$. A high-grade form of SiC was selected as the spreader material for its unique combination of low CTE and high thermal conductivity. The CTE of the SiC heat spreader closely matches the CTE of the chip, avoiding stress problems when the module heats up during use. The

Table 6 Comparison of measured and calculated breaking forces.

	Predicted breaking force from model (lb) (mean $\pm 1\sigma$)	Measured breaking force (lb) (mean $\pm 1\sigma$)
With support post	3264 \pm 414	3043 \pm 335
Without support post	2057 \pm 261	2072 \pm 124




Figure 24

Finite element model of chip power distribution. (Modeling software from ANSYS, Inc.)




Figure 25

MCM cross section showing adhesive thermal interface (ATI) cooling.

thermal conductivity of the SiC spreader is approximately 275 W/m·K, and it is nearly isotropic. The thickness of the spreader was optimized on the basis of extensive thermal modeling. To further enhance the effectiveness of the heat spreading within the module, the chip spacing was optimized to improve heat flow. To minimize the thermal resistance from each chip to its heat spreader, the heat spreaders were attached independently, so that thin bond lines were maintained and the effects of varying chip height and tilt were minimized. The thermal benefit of using ATI with adhesive thermal compound (ATC) instead of a solution with only ATC is that the thermal resistance of the ATC layer is significantly reduced by having the heat spread out over the spreader, which has more than

twice as much area as the chip. The combined thermal resistance of the ATI structure with the ATC layer is less than the thermal resistance of an ATC layer alone. An additional benefit of having the heat spreader thermally coupled to the back of the chip is that the thermal gradients in the chip are significantly reduced as well:

Thermal resistance =

$$\frac{\text{distance (m)}}{\text{thermal conductivity (W/m}\cdot\text{K)} \times \text{area (m}^2\text{)}} \quad (2)$$

The outer shell of the module is a thick copper hat that cools and protects the module. The top outer surface of the hat is mechanically and thermally connected to an impingement heat sink that receives ducted airflow. The inside of the hat encloses the heat spreaders, the chips, and the top of the substrate. Because the heights and tilts of the heat spreaders vary, a thermally compliant layer is needed to complete the heat path from the spreaders to the inside of the hat. A thin layer of IBM ATC 3.8 is used to fill this gap. This material has a thermal conductivity of 3.8 W/m·K. The copper hat is sealed to the substrate with Sylgard** adhesive to form an enclosure, which provides support to meet application requirements.

10. Regatta MCM reliability/qualification

The Regatta MCM is one of the first IBM products that incorporates many new technology elements: HPGC, tighter-pitched C4s (100 μm on 200-μm centers) on large chips (up to 427 mm²), and a high-density LGA connection. The qualification of this package included reliability evaluation of these new elements and other aspects of package technology under the new processes and operating conditions.

Reliability evaluation is an essential part of the IBM product qualification program. Several of the Regatta MCM technology components were thoroughly evaluated over the last five years through multiple phases from development stage to prototyping to early user hardware and the final product form factor. Reliability evaluations were performed using a combination of specially designed test-vehicle (TV) and product-form-factor modules.

Table 7 is a summary of the stress plan for the final qualification.

Test-vehicle description

The test-vehicle module is a four-chip MCM with the same dimensions, number of layers, and ground rules as the Regatta product, but with special wiring to access C4/LGA test features as well as heaters/thermal sense resistors on the chips. The test module is mounted on a test card, also with similar mechanical support (stiffeners, spring plate, interposer, etc.) to allow powering of the modules, monitoring of temperature,

Table 7 MCM qualification stress plan.

<i>Reliability concerns</i>	<i>Stress type</i>	<i>Stress modules</i>	<i>Stress conditions</i>
C4 integrity	Wettability	Test vehicles/products	Chip pull after initial join, TCA, 13× reflows, 4× rework
C4 fatigue	Power cycling Thermal cycling	Test vehicles Test vehicles	P/C 185 W, 25/102°C, 2 Kcyc 1. 20/100°C, 2 Kcyc 2. 20/100°C, 2 Kcyc 3. 40/120°C, 2 Kcyc 4. 20/120°C, 2 Kcyc 5. 0/120°C, 2 Kcyc 6. -55/125°C, 1 Kcyc
C4 migration/corrosion	Temperature/humidity/bias	Test vehicles	85°C/65%RH//15 V and 85°C/85%RH/15 V 1000 hr
LGA pad integrity	Multiple actuations/Group B	Test vehicles/products	1×, 5×, 10× actuations/Module Group B
LGA pad reliability	Temperature/humidity/bias	Test vehicles	85°C/65%RH//15 V, 1000 hr
Thermal performance and reliability	Thermal cycle	Test vehicles	1. 20/100°C, 2 Kcyc 3. 40/120°C, 2 Kcyc 4. 20/120°C, 2 Kcyc 5. 0/120°C, 2 Kcyc
	Thermal age/power cycle	Test vehicles	125°C 500 hr 25/102°C-185 W, 1 Kcyc
	Thermal age	Test vehicles	125°C, 1.5 Khr
	Temperature/humidity	Test vehicles	85°C/65%RH, 1000 hr
	Power cycle	Test vehicles	P/C 185 W 25/102°C, 2 Kcyc
	Thermal age-power cycle	Test vehicles	125°C, 500 cyc + P/C 185 W, 25/102°C, 1 Kcyc
Seal integrity	Power cycling Thermal cycling	Test vehicles/products	185 W, 25/102°C, 2 Kcyc 0/100°C, 3 Kcyc
Glass-ceramic substrate integrity	Thermal cycling Liq/liq thermal shock	Test vehicles/products	-25/125°C, 3 Kcyc -65/150°C, 1.5 Kcyc
Glass-ceramic substrate reliability	Temperature/humidity-thermal cycling under load	Products	T/H (90°C/85%RH, 90°C/70%RH, 65°C/70%RH, and T/H-T/C (-55/125°C, -25°C/125°C), 85g/100g/120g/166g per I/O load

and electrical/thermal measurements of all test features.

The test chip is similar in size to the product chip and contains 7194 C4s (100-μm balls on 200-μm spacing). This test chip is a combination C4/thermal test vehicle to provide information on C4 reliability, electromigration, metal migration/corrosion, and thermal characteristics and performance. There are four heaters on the chip, one in each quadrant. This test chip can be powered up to a 200-W level for extended stressing periods and can generate up to 500 W during readout. By individually controlling the heaters, various thermal gradients across the chip can be obtained. Multiple thermal sense resistors monitor temperature in various locations on the chip. A ring connection of signal C4s enables the testing of C4 failures at different distances to the neutral point (DNP) to monitor C4 fatigue behavior. Furthermore, the chip contains test sites designed for testing electromigration and metal migration/corrosion between adjacent C4s.

Product-form-factor design and process were used for the encapsulation of the modules and thermal assembly.

Stress preconditioning

Prior to stressing, all test modules were subjected to modified LGA Industry Standard EIA (Electronic Industries Alliance) Group B preconditioning to simulate the worst-case mechanical stresses the MCM module must withstand during module/system assembly, shipment, and installation. The Group B stresses include thermal shock (EIA 540B0AE-EIA364 Procedure 32), impact shock (EIA 540B0AE-IBM Modified EIA364 Procedure 27 Condition A), and vibration sweep (EIA 540B0AE-IBM Modified EIA364 Procedure 28 Condition V-B). The IBM modification to the impact shock is to perform three shocks for each of the three axes, for a total of nine, and the IBM modification to the vibration is a modification of the sweep spectral density.

Table 8 MCM stress corrosion results.

Load (grams/contact)	Temperature/humidity (T/H) condition (°C/%RH)	Cumulative T/H time (hr)	No. of fails/No. of samples	125-hr T/H followed by –55°C to 125°C temp. cycling		125-hr T/H followed by –25°C to 125°C temp. cycling	
				Cum. cycles	No. of fails/No. of samples	Cum. cycles	No. of fails/No. of samples
85	90/85	2000	0/2	928	0/2		
100	90/85	2000	0/2	928	1/2		
120	90/85	2000	0/2	1146	0/2	728	0/2
120	90/85	2000	0/2	1028 *300 hr T/H	0/1		
120	90/70	500	0/2	928	1/2		
120	65/70	500	0/2	1146	0/2		
166	90/85	2000	0/2	507	2/2	728	0/1

Glass-ceramic substrate integrity

Historically, the main cause of latent open defects in glass-ceramic substrates is contamination in vias or lines. The contaminants were typically burned off during sintering, leaving voids that could be filled with glass. The substrate burn-in thermal cycling test has proven to be very effective in screening out these latent defects by breaking them into detectable opens. Nonlinear conductance testing is employed after burn-in to further ensure that the products meet very stringent field reliability requirements. In spite of the fact that glass-ceramic has been used in many prior high-end products, extensive reliability stressing, as shown in Table 7, was performed to verify design robustness and ensure high-quality manufacturing processes.

Glass-ceramic substrate reliability

Since the Regatta MCM is the first application of IBM HPGC with LGA interconnection, a new reliability risk for glass-ceramic substrates has arisen. Of particular concern is the potential weakening of the substrates if they are not protected from extreme high temperature and humidity during shipping and storage. This concern is based on laboratory studies of the breaking strength of glass-ceramic under dynamic loading, which show strength degradation with increasing temperature and humidity (see Section 6). Extensive modeling showed that the addition of a center spar significantly reduced the risk of stress corrosion cracking (SCC). To verify the model, a test matrix was constructed as shown in Table 8, covering extreme ranges of temperature, humidity, and cycling conditions while subjecting modules to LGA loading forces from nominal to more than twice the nominal load levels. This matrix of varying stress conditions allows us to

evaluate the factors of temperature, humidity, and loading force contributing to the stress corrosion cracking mechanism. The data was used to predict safe usage, which includes shipping, storage, and both operating and non-operating conditions.

C4 integrity evaluation

As described in earlier sections, the devices used in Regatta MCM represent the highest-density IBM C4 products with regard to tighter C4 ground rules and large die size. The C4 solder balls are joined directly to the pads disposed on the glass-ceramic surface and not to thin-film metallurgy, as in previously introduced products. For this pad metallurgy, all aspects of C4 integrity, joinability, wettability, reflow limits, reworks, underfill integrity, corrosion, and fatigue were assessed as part of a rigorous reliability qualification.

Extensive evaluations were carried out through each stage of the bond, assembly, and test (BAT) processes. The integrity of C4s was determined through a combination of electrical measurement, physical characterization (cross-sectional analysis, examination of BLM interfaces, C4 size measurement, and Auger surface analysis), and wettability testing. Wettability data was collected by cold-chip pulling to evaluate the integrity of the C4 joining representing all BAT conditions (temporary chip attach, or TCA, initial chip join, up to 13× reflows, and 4× chip reworks). The wettability inspection data included chip pull strength, evidence of pulled pads, geodes, MSVs, non-wets, and planar fracture (i.e., cold-pull fractures occurring at interfaces rather than within the solder). The results of these studies were discussed in Section 5 and summarized in Table 4.

C4 reliability evaluation

The C4 failure rate projection for the Regatta MCM, based on the IBM C4 stress database, indicated that no fails would be expected in field operation. However, the combination of this high-density I/O die with several newly introduced technology elements dictated extensive stressing, as shown in Table 7, to evaluate different aspects of C4 reliability from thermal cycling fatigue susceptibility to metal migration or corrosion that may be induced through a combination of process variations (defects, contamination, etc.).

LGA pad integrity/reliability

The integrity of the pads is determined by visual inspection and multiple actuations up to 10× to represent the maximum number of resocketing cycles the module may experience. Reliability testing includes subjecting parts to multiple reflows, Group B and temperature/humidity/bias stressing with visual inspection, and electrical measurements. Since thermal and electrical measurements of test features in all stress cells were obtained with modules mounted on cards through LGA socket connections, good and repeatable measurements at stress intervals also reflect reliable LGA pads. In addition, in earlier phases of the evaluation of the HPGC with LGA, we had designed a special test vehicle with a stitch pattern of rings between the card and the substrate. The resistance of these rings was measured to verify the LGA interconnect reliability, with the failure criterion being greater than 10% increase of initial ring resistances.

Thermal integrity/reliability

Degradation of thermal paste translates to a temperature rise over the life of the module. This could lead to a significant increase in failure rate because of the exponential nature of many failure mechanisms. The thermal integrity and reliability were evaluated by both stressing and construction analyses. Thermal evaluation included 1) determination of initial module (time-zero) thermal performance; 2) determination of thermal paste void fraction and chip-to-hat gaps; 3) verification of the thermal model; and 4) determination of both internal and external thermal interface performance stability and reliability.

Stressing consisted of power cycling (P/C), thermal cycling (T/C), thermal aging (T/A) and temperature and humidity (T/H) exposures. Thermal cycling cells were treated at different conditions to study the effects of peak temperatures and the magnitude of thermal excursions on thermal performance. Combination stress tests with both T/A and P/C were also performed to verify the internal thermal performance with respect to the kinetics of interface separation. Measurements

were performed with more frequent readouts in the beginning (as often as every 50 cycles) to monitor thermal performance. Construction analyses included examining chip-to-hat gaps and measuring void fractions of thermal paste.

Seal (hermeticity) integrity/reliability

Regatta MCMs are nonhermetic modules, but they do need to remain fluid-tight. Degradation or loss of hermeticity to fluids has many adverse effects on module reliability, since oxidation and corrosion increase degradation of both C4 and thermal performance. An elastomeric adhesive has previously been used in many IBM MCM and SCM products, but primarily on alumina modules. Regatta MCMs are the first HPGC modules using this adhesive sealant and LGA interconnection to the card.

The seal evaluation assessed the following reliability concerns: 1) defective seals caused by handling damage and/or contamination which may occur during build, cap rework, chip rework, and thermal paste cleaning; 2) adhesion of Sylgard adhesive to cap and glass-ceramic substrates; 3) seal degradation induced by transient temperature differences and the thermal coefficient of expansion mismatch between the cap material and the glass-ceramic substrate; and 4) degradation of certain seal areas which may be introduced by the spar during thermal excursions with the module under LGA load. This was addressed by accelerated power and thermal cycling of fully assembled sealed modules with bubble-leak readouts at frequent intervals. The seal adhesion was also assessed by lid shearing and visual inspection to ensure the required joint strength.

Results and conclusions

Successful results were obtained in all stress cells. No degradation was observed in the C4 fatigue, C4 migration, LGA pad, seal integrity/reliability, or SCC cells after stressing. No thermal degradation was observed in any thermal test cell: PC, T/C, T/A, T/A-P/C, and T/H. No defects were detected, either at time zero or after severe stressing, in the glass-ceramic integrity cell, affirming sound design and well-controlled manufacturing processes. Fails observed in stress cells tested under severe stress conditions occurred as expected. The fail statistics were analyzed and used for failure-rate projections, which showed failure rates well below typical values for high-end server products. Samples from all construction analysis cells monitoring thermal gaps, thermal void fractions, and C4 wettability met specifications, with very tight distributions demonstrating a reliable and robust design which exceeded lifetime requirements.

11. Summary

An advanced multichip module was developed utilizing a high-performance glass-ceramic-multilayer ceramic substrate with copper wiring. Seven thousand eighteen C4 pads per chip were connected using a 200- μ m-pitch C4 solder pad array, and each die was connected with 190 m of high-speed copper wiring with bandwidth for operation at 41.6 GB/s. A state-of-the-art module cooling solution was developed utilizing SiC heat spreaders to support 156 W of heat removal per die. A total of 624 W of module cooling was achieved through a thermal conduction path comprising a conductive adhesive, SiC heat spreader, the IBM advanced thermal compound, and a copper hat. An advanced land grid array (LGA) connector was developed which provides 5184 I/O connections to the board with a minimum 1-mm I/O-pad-to-pad pitch. Combined with the high electrical interconnection density on chip, the advanced module and packaging technology makes this compact MCM a powerhouse for performance in an effective cost-performance system solution. The engineering integration of materials, processes, and specifications, as well as the teamwork of module and system designers, engineers, and production teams, resulted in an industry-leading module solution that has been received positively by customers and e-business users around the globe.

Acknowledgments

A project of this magnitude and complexity cannot be accomplished without the hard work of many individuals. The authors would particularly like to recognize the efforts of the following people for their contributions to the development of the advanced MCM for high-performance UNIX servers: Bradley McCredie, Joel Tendler, Glenn Daves, Robert Sherer, Ron Hering, Jasvir Jaspal, Giulio DiGiacomo, Guy Messina, Eric Kastberg, Joe Ross, Dave Long, Andy Brendler, Scott Bradley, Ray Jackson, Juan Jeri, Jac Burke, Steve Delaurentis, Scott Langenthal, Bob Schwartz, Bor Zen Hong, Harry Harrington, Arnold Terpening, Mark Laplante, and Renee Weisman. Many departments in the IBM Microelectronics Division contributed to the project. Here we acknowledge the key contributions of the LGA Development team (leaders: Mark Hoffmeyer, John Colbert, Marie Cole, Benson Chan, Bill Brodsky, Dan Massey, Charles Perry, Mark Plucinski, John Corbin) and the analytical departments GUED and LUED.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of The Open Group or Dow Corning Corporation.

References

1. J. M. Tendler, J. S. Dodson, J. S. Fields, Jr., H. Le, and B. Sinharoy, "POWER4 System Microarchitecture," *IBM J. Res. & Dev.* **46**, No. 1, 5–26 (January 2002).
2. P. Walling, A. Tai, and H. Hamel, "High Bandwidth Low Latency Chip-to-Chip Interconnects Using High Performance MIC Glass Ceramic POWER4[®] MCM," *Proceedings of the Tenth IEEE Topical Meeting on Electrical Performance of Electronic Packaging*, 2001, pp. 299–302.
3. R. R. Tummala, J. U. Knickerbocker, S. H. Knickerbocker, L. W. Herron, R. W. Nufer, R. N. Master, M. O. Neisser, B. M. Kellner, C. H. Perry, J. N. Humenik, and T. F. Redmond, "High Performance Glass Ceramic/Copper Multilayer Substrate with Thin-Film Redistribution," *IBM J. Res. & Dev.* **36**, No. 5, 889–904 (September 1992).
4. P. J. Brofman, S. K. Ray, and K. F. Beckham, "Electrical Connections to the Thermal Conduction Modules of the IBM Enterprise System/9000 Water-Cooled Processors," *IBM J. Res. & Dev.* **36**, No. 5, 921–933 (September 1992).
5. G. Katopis, D. Becker, and H. Stoller, "First Level Package Design Considerations for the IBM S/390 G5 Server," *Proceedings of the IEEE 7th Topical Meeting on Electrical Performance of Electronic Packaging*, October 1998, pp. 15–16.
6. A. Plachy, B. Fasano, R. Indyk, D. O'Connor, and S. Reddy, "Glass Ceramic Substrates for Flip Chip Packages," *MicroNews* **4**, No. 4, 25–27 (July 1998).
7. *Microelectronics Packaging Handbook*, R. R. Tummala, E. J. Rymaszewski, and A. J. Klopfenstein, Eds., Van Nostrand Reinhold, New York, 1989, pp. 14, 30, 32, 369, 720.
8. L. W. Herron, R. N. Master, and R. R. Tummala, "Method of Making Multilayered Glass-Ceramic Structures Having an Internal Distribution of Copper-Based Conductors," U.S. Patent 4,234,367, 1980.
9. P. L. Flaitz, A. M. Flanagan, J. M. Harvilchuck, L. W. Herron, J. U. Knickerbocker, R. W. Nufer, C. H. Perry, S. N. Reddy, and S. P. Young, "Method and Means for Co-Sintering Ceramic/Metal MLC Substrates," U.S. Patent 5,130,067, 1992.
10. E. H. Laine and P. M. O'Leary, "IBM Chip Packaging Roadmap," *IBM MicroNews* **5**, No. 3, 22–25 (1999).
11. G. Katopis, D. Becker, H. Smith, and H. Stoller, "MCM-C/D Design for the CMOS Implementation of the S/390 System," *Proceedings of the 47th IEEE Conference on Electronic Components and Technology*, 1997, pp. 479–485.
12. M. J. Ellsworth, H. Hamel, and E. D. Perfecto, "A High Density, High Performance MCM-D/C Package: A Design, Electrical and Process Perspective," *Proceedings of the 46th IEEE Conference on Electronic Components and Technology*, 1996, pp. 821–828.
13. P. J. Brofman, K. J. Puttlitz, K. A. Stalter, and C. Woychik, "Flip-Chip Die Attach Technology," *Area Array Interconnection Handbook*, K. J. Puttlitz and P. A. Totta, Eds., Kluwer Academic Publishers, Boston, 2001, p. 343.
14. S. K. Ray, K. Beckham, and R. Master, "Flip-Chip Interconnection Technology for Advanced Thermal Conduction Modules," *Proceedings of the 41st IEEE Conference on Electronic Components and Technology*, 1991, pp. 772–778.
15. K. A. Stalter, R. A. Jackson, and D. C. Linnell, "Low-Cost, High Reliability Flip-Chip Removal for Multichip Modules," *Proceedings of the 49th IEEE Conference on Electronic Components and Technology*, 1999, pp. 446–450.
16. N. G. Koopman, T. C. Reily, and P. A. Totta, "Chip-to-Package Interconnections," *Microelectronics Packaging Handbook*, R. R. Tummala, E. J. Rymaszewski, and A. J. Klopfenstein, Eds., Van Nostrand Reinhold, New York, 1989, pp. 361–391.

17. D. W. Wang and K. I. Papathomas, "Encapsulant for Fatigue Life Enhancement of Controlled-Collapse Chip Connection (C4)," *IEEE Trans. Components, Hybrids, Manuf. Technol.* **16**, No. 8, 863–867 (December 1993).
18. S. L. Buchwalter, M. E. Edwards, D. Gamata, M. A. Gaynes, and S. K. Tran, "Underfill: The Enabling Technology for Flip-Chip Packaging," *Area Array Interconnection Handbook*, K. Puttlitz and P. A. Totta, Eds., Kluwer Academic Publishers, Boston, 2001, pp. 458, 464.
19. T. Lombardi, F. Pompeo, J. Coffin, D. Plouffe, and C. Reynolds, "Rapid Cure Encapsulant for Use in Ceramic Chip-Carrier Applications," *IBM MicroNews* **5**, No. 4, 26–28 (1999).
20. K. J. Puttlitz and L. S. Goldmann, "Product Connector Technology," Ch. 21 of *Area Array Interconnection Handbook*, K. J. Puttlitz and P. A. Totta, Eds., Kluwer Academic Publishers, Boston, 2001.
21. A. S. Tetelman and A. J. McEvily, *The Mechanics of Fracture*, Wiley, New York, 1967.
22. J. N. Humenik and J. E. Ritter, "Stress Corrosion of Alumina Substrates," *J. Mater. Sci.* **14**, No. 5, 626–632 (1979).

Received March 28, 2002; accepted for publication June 24, 2002

John U. Knickerbocker *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (knickerj@us.ibm.com)*. Dr. Knickerbocker is an IBM Distinguished Engineer and is Director of World Wide Development in the InterConnect Products (ICP) organization within the IBM Microelectronics Division at the IBM East Fishkill facility. He received B.S. and M.S. degrees in ceramic engineering in 1977 and 1978, respectively, from the State University of New York at Alfred and his Ph.D. in ceramic engineering in 1982 from the University of Illinois at Champaign–Urbana. He subsequently joined IBM at the East Fishkill facility, where he has held a series of engineering and management assignments and currently focuses on flip-chip packaging technology development, including thermal solutions, chip-to-carrier interconnections, ceramic and organic chip carriers, chip carriers to boards and cards, PWBs and microcards, lead reduction and elimination, and intellectual property development. Dr. Knickerbocker has received an IBM Corporate Award and three Division Awards, including IBM Outstanding Achievement Awards and an IBM Patent Portfolio Award. He has authored or co-authored 96 U.S. patents or patent applications and 16 papers and publications; he has served as the IBM representative on the SEMATECH Focus Technical Advisory Board for Packaging and has served with multiple universities on technical advisory boards. He is a member of IMAPS and a Fellow of the American Ceramic Society.

Frank L. Pompeo *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (pompeo@us.ibm.com)*. Mr. Pompeo joined the IBM East Fishkill facility in 1982 after receiving his B.Eng. degree in metallurgical engineering in 1982 and his M.S. degree in materials science in 1987, both from Stevens Institute of Technology. Currently a Senior Engineering Development manager for module interconnection technology, he has held a series of management and executive technical staff assignments focusing on ceramic and organic chip carrier and PWB interconnection processes. Mr. Pompeo holds 25 patents and has co-authored 11 external publications in the areas of module interconnections and assembly. He has received one IBM Outstanding Achievement Award and two IBM Division Excellence Awards for his module packaging contributions.

Alice F. Tai *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (alicetai@us.ibm.com)*. Dr. Tai joined IBM in 1987 after receiving her Ph.D. in chemistry from the University of California at Los Angeles. She has held technical and managerial positions in the plating and thin-film processing areas. Dr. Tai is currently a Senior Engineer responsible for the Regatta packaging application and program management. She holds one patent and has published several internal technical papers.

Donald L. Thomas *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (thomasdo@us.ibm.com)*. Mr. Thomas joined IBM in 1988. He holds a B.S. degree in electrical and computer engineering from Clarkson University and a master's degree from Syracuse University in computer engineering. He is currently a Senior

Engineer responsible for the technical marketing strategy for first- and second-level packaging that the IBM Microelectronics Division delivers to the IBM Server Group. Mr. Thomas was the lead packaging designer and subsequently the packaging application and design manager for the previous five generations of IBM mainframe servers.

Roger D. Weekly *IBM Server Group, 11400 Burnet Road, Austin, Texas 78758 (weekly@us.ibm.com)*. Mr. Weekly joined IBM in 1976 after receiving a B.S. degree from Worcester Polytechnic Institute in 1975 and an M.S. degree from Stanford in 1976, both in electrical engineering. He initially worked with power-supply development for IBM low- to mid-range systems in printers, personal computers, and UNIX-class workstations. As part of this activity, he was lead designer of a set of CMOS ICs for switched-mode power supply control. In 1986 Mr. Weekly taught at Tuskegee University, Alabama, under the IBM Faculty Loan Program. He subsequently spent several years during the emergence of the Fibre Channel Standard developing fiber optics adapters and switches. In 1996 he moved to the POWER3 processor development group, with responsibility for the electrical application support of this processor family in IBM products. Since 1998 Mr. Weekly has been involved in the first- and second-level packaging of IBM POWER processors and their support chips. He has received several IBM Outstanding Technical Achievement Awards for his work in these areas.

Michael G. Nealon *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (nealon@us.ibm.com)*.

Harvey C. Hamel *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (hhamel@us.ibm.com)*. Mr. Hamel received a B.S.E.E. degree from Wayne State University and an M.S.E. degree from Syracuse University, joining IBM in Poughkeepsie, New York, in 1964 as a designer of special circuits. He began work on wide-band amplifiers, microwave microstrip filters, and directional couplers, and his work has included designing Josephson circuits and servo systems and developing techniques to analyze the stability of switching regulators. After a time spent modeling bipolar structures, he has worked in the general area of electrical packaging analysis and design. Mr. Hamel's interest is in frequency-dependent lossy coupled structures and how they relate to signal bandwidth, crosstalk, and power integrity. His present work is concerned with the modeling of the electrical limitations and tradeoffs associated with MCMs and interconnections. He holds numerous patents in the packaging and interconnection field.

Anand Haridass *IBM Server Group, 11400 Burnet Road, Austin, Texas 78758 (anandh@us.ibm.com)*. Mr. Haridass received the B.E. degree from KREC, Surathkal, India, in 1995 and the M.S.E.E. degree from the Georgia Institute of Technology, Atlanta, in 1997. That same year, he joined the IBM Somerset Design Center in Austin, Texas, and was involved in first-level packaging for PowerPC processors designed for Apple. Since 1998, Mr. Haridass has been an

integral part of the IBM Server Packaging Team, involved in modeling, simulation, and measurement of packages and issues related to high-speed signal integrity. He is pursuing a Ph.D. degree in the field of electromagnetic modeling at the University of Texas through the IBM Work-Study Program.

James N. Humenik *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (humenik@us.ibm.com)*. Dr. Humenik is a Senior Technical Staff Member at the IBM East Fishkill facility. He received a B.S. degree in 1970, an M.S. degree in 1972, and a Ph.D. degree in 1974, all in ceramic engineering from the University of Illinois at Urbana-Champaign. He joined IBM at the East Fishkill facility, where he has held a series of management and technical staff assignments focusing primarily on the materials and processes used to fabricate microelectronic packages. Dr. Humenik has authored or co-authored ten papers and holds 19 U.S. patents. He is a Fellow of the American Ceramic Society, where he also serves as the IBM representative.

Richard A. Shelleman *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (shellema@us.ibm.com)*. Dr. Shelleman received his Ph.D. degree in ceramic science from The Pennsylvania State University in 1988. He joined the IBM Microelectronics Division in East Fishkill in the packaging development group in 1988 as a Staff Engineer. He is now an Advisory Engineer in the InterConnect Products group, responsible for glass-ceramic sintering development. Dr. Shelleman is a member of the American Ceramic Society and the Materials Research Community in IBM.

Srinivasa N. Reddy *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (snreddy@us.ibm.com)*. Dr. Reddy is a Senior Engineer in Ceramic Packaging Development. He received B.E. and M.E. degrees in metallurgy from the Indian Institute of Science and a Ph.D. in materials science from Marquette University. Before joining IBM in 1982, he worked as a postdoctoral researcher at Ohio State University and as a research engineer at the Amoco Research Center. His specialty is the application of thermodynamic and kinetic principles to the development of high-temperature processes. Dr. Reddy holds 28 patents and has been a leading technical contributor in the development of high-performance glass-ceramic technology.

Kevin M. Prettyman *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (kmpretty@us.ibm.com)*. Dr. Prettyman joined IBM in 1981. He received a B.S. degree in chemistry from Brigham Young University, M.E. and Ph.D. degrees in materials science and engineering from The University of Utah, and an M.B.A. from the New York University Stern School of Business. He has worked as a process engineer in Bond Assembly and Test, as a development engineer in a Materials Development Group, as a development engineering manager in Sintering, Measurements, and Test, and as a product line manager and product development team leader in the Interconnect Products Business Line Management Area. He has been a leader in new materials understanding and qualification and has helped bring to market the IBM high-performance glass-ceramic (HPGC) product line. Dr. Prettyman holds ten U.S. patents and has eight technical publications.

Benjamin V. Fasano *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (bfasano@us.ibm.com)*. Dr. Fasano joined IBM in 1992 after receiving B.S., M.S., and Ph.D. degrees in ceramic engineering from Rutgers University. He developed high secondary emission overcoat and glass materials for plasma display panels at the IBM Kingston, New York, facility. Dr. Fasano transferred to the Packaging Development group for the IBM high-performance glass-ceramic modules using thin films in 1985, and was responsible for raw materials and supplier qualifications for ceramic and metallization powders and organics used in the manufacturing process. He also helped develop the tape casting process for greensheets. Dr. Fasano has been the technical team leader for the high-performance glass-ceramic product development team since its inception. His current interests include high-frequency packaging and structures. He holds 36 U.S. patents in the packaging field.

Sudipta K. Ray *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (skray@us.ibm.com)*. Dr. Ray is a Senior Engineer at the IBM East Fishkill facility. He joined IBM at the Thomas J. Watson Research Center in 1979 after receiving a B.S. degree from Calcutta University, an M.S. degree in physics from Delhi University, and a Ph.D. degree in solid-state physics from the University of Virginia. Dr. Ray has held a series of management and technical staff assignments focusing primarily on the materials and processes used for first- and second-level microelectronic package interconnections, and thin-film wiring used on ceramic packages. He has authored or co-authored more than 25 papers and holds 30 U.S. patents. Dr. Ray is a Senior Member of the IEEE Computer Packaging and Manufacturing Technology Society, where he also serves as a member of the Advanced Packaging Committee.

Thomas E. Lombardi *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (lombard@us.ibm.com)*. Dr. Lombardi received B.S. and M.S. degrees in ceramic engineering from Rutgers University in 1984 and 1986, respectively. He subsequently joined IBM and worked as a development engineer in the areas of lamination and sintering of alumina and glass-ceramic packages. Since receiving his Ph.D. degree from Rutgers University in 1996, he has worked on a variety of projects related to bond and assembly development. He is currently a Senior Engineer responsible for underfill activities, C4 chip joining, and fine-pitch development on advanced ceramic packages. Dr. Lombardi is a member of the American Ceramic Society and an author or coauthor of nine patents.

Kenneth C. Marston *IBM Microelectronics Division, 2455 South Road, Poughkeepsie, New York 12601 (kmarston@us.ibm.com)*. Mr. Marston received a B.E.M.E. degree from the State University of New York at Stony Brook and an M.S.M.E. degree from the University of Minnesota, with a concentration in thermal sciences. He has 11 years' experience in Advanced Module Manufacturing, including development of temperature control systems for module test and burn-in equipment; process responsibility for temperature stress, solder reflow operations, and encapsulation; and project management for several products manufactured in Poughkeepsie. He is currently an Advisory Engineer.

Patrick A. Coico *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (coico@us.ibm.com)*. Mr. Coico is an electronic module development engineer at the IBM Microelectronics facility in East Fishkill, New York. He joined IBM in 1985. His areas of expertise include multichip module design, hermetic and nonhermetic sealing, land grid array interconnects, and package mechanical considerations. Much of his effort over the past 14 years has been focused on the large modules used in the IBM Enterprise Servers. Mr. Coico has made contributions that have enabled many innovations to these module designs. He received his B.S.M.E. degree from the State University of New York at Buffalo and his M.S.M.E. degree from the University of California at Berkeley. He is a registered Professional Engineer in the state of California.

Peter J. Brofman *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (brofman@us.ibm.com)*. Dr. Brofman is an IBM Distinguished Engineer in the InterConnect Products group at the IBM East Fishkill facility. He received a B.S. degree in materials engineering in 1975, an M.S. degree in physical metallurgy, an M.B.A. degree in 1978, and a Ph.D. degree in physical metallurgy in 1980, all from Rensselaer Polytechnic Institute. He subsequently joined IBM at the East Fishkill facility, where he has held a series of management and technical staff assignments all focusing on module interconnection processing. He has received two IBM Outstanding Achievement Awards, an IBM Outstanding Innovation Award, and an IBM Division Excellence Award; he has authored or co-authored 22 papers and 18 U.S. patents. Dr. Brofman serves as the IBM representative on the Semiconductor Corporation Technical Advisory Board for packaging and interconnects. He is a member of the American Society for Metals, International and the International Microelectronics and Packaging Society.

Lewis S. (Lew) Goldmann *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (goldmann@us.ibm.com)*. Mr. Goldmann received a B.S. degree in pre-engineering and liberal arts from Queens College, City University of New York; a B.S. degree in mechanical engineering from Columbia University; and an M.S.M.E. degree from MIT. For the last twenty years he has been involved in the mechanical evaluation of electronic packages, interconnections, devices, and materials, and is currently a Senior Technical Staff Member in Ceramic Module Development. He previously contributed to the development of C4 interconnections and early versions of the IBM thermal conduction module.

David L. Edwards *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (dledward@us.ibm.com)*. Mr. Edwards joined IBM in 1983 after receiving his B.S. degree in mechanical engineering from Clarkson College and his M.S. degree in mechanical engineering from Cornell University. He has held several technical and managerial positions at the IBM Poughkeepsie and East Fishkill facilities, working in the areas of system cooling, ceramic module thermal/mechanical design, and module bond and assembly. Mr. Edwards is currently a Senior Engineer, specializing in the cooling of high-power ceramic modules. He has achieved the IBM 11th level Invention Plateau, and holds 26 U.S. patents and several foreign patents.

Jeffrey A. Zitz *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (jeffzitz@us.ibm.com).* Mr. Zitz is the lead microprocessor/ASIC packaging engineer at IBM East Fishkill Microelectronics facility. During his 14 years with IBM, he has applied his diverse engineering expertise in the areas of materials, mechanical, thermal, and statistical modeling, reliability, and manufacturing to IBM ceramic and organic chip carriers and computer systems. He led the development of bare-die flip-chip ceramic carriers, and, more recently, development of the Direct Lid Attach (DLA) package utilized by IBM and other flip-chip package manufacturers. His innovations appear across the IBM ASIC and PowerPC die product offerings, and have been key to higher package performance at lower cost. Mr. Zitz received his B.S.M.E. and M.S.M.E. degrees from Rensselaer Polytechnic Institute; he is a registered Professional Engineer in the State of New York.

Sushumna Iruvanti *IBM Technology Group, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (iruvanti@us.ibm.com).* Dr. Iruvanti is a Senior Engineer in the InterConnect Products Development group. He received his B.Tech. degree from Madras University, India, his M.S. degree from Clarkson University, and his Ph.D. degree from the State University of New York at Buffalo, all in chemical engineering. He has received an IBM Corporate Award and an IBM Outstanding Innovation Award for his work in the development of advanced thermal compounds and flat-plate cooling technologies. He holds 21 patents and has received seven IBM Invention Achievement Plateau Awards. Dr. Iruvanti is a member of AIChE and IMAPS.

Subhash L. Shinde *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (shinde@us.ibm.com).* Dr. Shinde joined IBM at the Thomas J. Watson Research Center in 1986 after receiving B.Tech. and M.Tech. degrees from the Indian Institute of Technology, Bombay, and a Ph.D. degree in materials science from Stanford University. At the Research Center, his work was concentrated primarily on electrical and magnetic properties of granular superconductors, and interfaces in such systems. Dr. Shinde edited a book in this area published by Springer-Verlag. He also worked on interfacial interactions in metal/ceramic systems related to the IBM electronic packages. In 1992 he moved to the IBM Microelectronics Division and initiated a project on low-temperature sintering aluminum nitride for thermally efficient electronic packages. He holds more than 15 patents in this area. Dr. Shinde's current work focuses on issues related to optoelectronic packaging and thermal management materials development. He is a member of the American Physical Society, the American Ceramic Society, and the Materials Research Society. He holds 33 patents, has published 26 papers, and is currently editing a book on high-thermal-conductivity materials.

Hai P. Longworth *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (longworth@us.ibm.com).* Ms. Longworth is a Senior Engineer in the Ceramic Packaging group. She is responsible for reliability assurance of glass ceramics, C4s, thin films, and packaging product qualification. She joined IBM in 1992 after receiving an Sc.D. degree in electronic materials engineering from the Massachusetts Institute of Technology. She received

an M.S. degree in metallurgical engineering from Michigan Technological University and B.S. degrees in physics and chemistry from Saigon University. She had worked as an analytical chemist and supervisor at Howmet Corporation in Michigan. Ms. Longworth holds one patent; she has published and presented 18 technical papers on microstructures vs. properties, electromigration, thin film, packaging reliability, and qualification.