Some Key Issues in Microelectronic Packaging

Guy V. Clatterbaugh, Paul Vichot, and Harry K. Charles, Jr.

Military and space electronics are tending toward increased system performance, i.e., higher speed, higher circuit density, and higher functionality. Recent reductions in government spending on space and military hardware have also made cost reduction a key consideration. As electronics approach physical size and performance limits, practical considerations such as wireability, thermal management, electromagnetic compatibility, and system reliability become dominant issues in system design. Resolving such issues requires the use of sophisticated analysis and computational methods.
(Keywords: Electronic packaging, Multiconductor transmission line analysis, Printed circuit board thermal analysis, Wireability.)

INTRODUCTION

In recent years, the electronics industry has discovered that the major economic advances made in high-performance electronic circuitry have come with increased integration. The industry is rapidly converging toward true wafer-scale integration, i.e., toward an entire system fabricated on one silicon substrate. Every 5 years or so, we see wafer foundries processing larger silicon wafers with smaller line geometries. Today, 12-in. wafers are being processed with 0.35-μm lines. By the year 2010, 16-in. wafers and 0.10-μm lines will likely be the standard.

Even with the increased ability to build larger chips with greater functionality, very few systems function as a single chip. Instead, the trend is to take advantage of recent technology advances in electronic packaging to pack chips side by side, sometimes even in three-dimensional stacks, to achieve still higher levels of functionality and performance while reducing volume and weight.

The quest to achieve better performance (higher speed and integration) has placed pressure on manufacturers and has forced integrated circuits (ICs) closer together. High-speed computer systems require that the central processing unit and the memory and controllers be proximal to minimize interconnection delays. The increased functionality of these chips has increased the number of inputs and outputs (I/Os). This, in turn, has imposed routing demands that have changed the way we think about fabricating printed wiring boards (PWBs) and building IC packages. The higher speeds and higher densities also create thermal and thermal stress problems that challenge the limits of today's microelectronic materials and manufacturing methods.

In this era of dwindling federal budgets, cost is a major concern. Previously, high performance and high reliability were demanded of military and space electronic systems; cost was often not an issue. However, the military now mandates as much use of commercial off-the-shelf hardware as possible, and the watchword from NASA is "smaller, lighter, and cheaper." So while the emphasis on cost has increased, the emphasis on performance and reliability has remained high. With this heightened demand for more complexity, cost-reduction opportunities are being strained.

The primary objective of microelectronic packaging is to design an electronic system that will satisfy the requirements for a particular application at a reasonable cost. Because of the many options available to interconnect and house an electronic system, the choice of a packaging technology for a given application is not always straightforward. Selection criteria may include one or more technology drivers:

- Wireability
- Yield
- Cost
- Heat transfer characteristics
- Electromagnetic performance
- Mechanical toughness
- Reliability

These fundamental design considerations affect speed, functionality, junction temperatures, volume, and weight. The primary goal is to select the most cost-effective yet reliable interconnection technology, which requires a quantification of the technology drivers. The resolution of these key design issues often depends on the use of sophisticated analysis methods.

This article discusses some of the quantitative tools developed and used at APL to assist the electronic packaging engineer in selecting reliable, cost-effective solutions for packaging high-performance electronic systems. After briefly examining various packaging options, we will detail five of the seven selection criteria already noted: wireability, yield, cost, heat transfer, and electromagnetic performance.

MICROELECTRONIC PACKAGING OPTIONS

Microelectronic packaging is a branch of engineering that deals with the design of methods for the fabrication and manufacture of interconnected miniature electronic systems (e.g., ICs and discrete and integrated passive devices) and the reliability of those systems. Specifically, microelectronic packaging involves routing signals while maintaining signal integrity, distributing ground and power to ICs, dispersing dissipated thermal energy while maintaining structural and material integrity, and protecting the circuit from environmental hazards.

Almost all of the methods for packaging ICs include the use of a PWB with connectors that provide the real-world I/Os to an electronic circuit. The traditional packaging approach involves the use of single packages. The chief advantage of single-chip packages is the ability to fully test the IC before interconnecting it to the underlying substrate. Such packaged devices are either through-hole mounted or surface mounted to the PWB.

Surface-mounted components do not require via holes to go through the entire board. Instead, packages and other surface-mounted components can be soldered to both sides of the PWB, thereby increasing circuit density. This approach is called surface-mount technology (SMT). The addition of area-array–style packages such as ball-grid arrays (BGAs) and chip-scale packages (CSPs) is making SMT competitive with the highest-density packaging technologies, although the reliability of some area-array packages is still questionable.

A newer packaging technology involves the attachment of more than one device onto a high-density interconnection substrate, which is then mounted in a large package. This provides both I/O pins and environmental protection. This so-called multichip module (MCM) technology is further characterized by the substrate technologies used to interconnect the attached ICs.

MCM-D represents thin film "deposited" metal and dielectric multilayers. MCM-D substrates have the highest wiring densities of all MCM technologies owing to the sophisticated semiconductor processing equipment used to manufacture them. MCM-C refers to multilayered "ceramic" substrates, fired from stacked alternating layers of screened metal inks and unfired ceramic sheets. This technology yields a moderately dense wiring capacity. MCM-L refers to multilayer substrates made from stacked, metallized PWB "laminates," which are individually patterned and then laminated. Once considered a low-density interconnect technology, MCM-L is rapidly approaching the density of MCM-C and MCM-D technologies.

Direct chip attach (DCA) or chip-on-board (COB) technology entails mounting the ICs directly to the PWB. A plastic encapsulant, which is "globbed" over the bare IC and then cured, provides environmental protection. ICs can be interconnected to the substrate using either flip-chip, tape-automated bonding, or wire bonding methods. DCA technology is particularly economical for systems that are limited to 10 or fewer ICs, since larger numbers of chips can affect system yield and DCA assemblies can be difficult to rework.

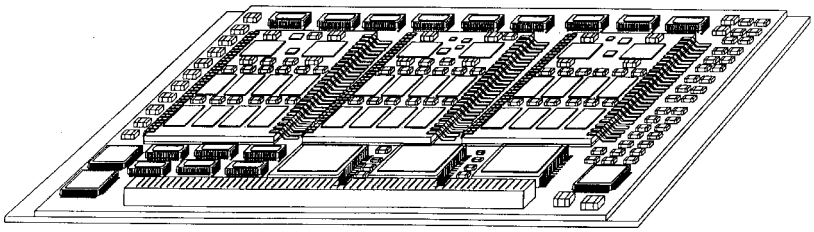

Figure 1. High-density surface-mount assembly with multichip carrier modules.

An advantage common to both the DCA and MCM packaging options is the elimination of the IC package interconnection level, which allows closer proximity (shorter propagation delays) and reduced lead inductance. The primary disadvantage with both methods is the difficulty in purchasing fully tested ICs (the "known good die" problem; see the section entitled Yield). Other disadvantages of DCA and MCM-L technologies include poor thermal management owing to the low thermal conductivity of PWB laminates and a poor coefficient of thermal expansion match between the die and the substrate. Solving the last problem usually requires an interposer substrate such as molybdenum for wire bonded die and an underfill epoxy for flip-chip die.

The multichip carrier module (MCCM, Fig. 1) attempts to marry all the positive aspects of DCA with MCM technology. The MCCM is simply a small MCM on a thin metal carrier that can be bonded or mechanically attached to a PWB. The metal bottom acts as both a heat spreader and a stress interposer for the MCM substrate. The MCCM has peripheral leads for wire bonding, soldering, or tab bonding to a PWB. The bare ICs are protected using a glob-top material.

We return now, for the remainder of this article, to a discussion of some key issues in the selection of packaging technologies and an examination of quantitative tools used to analyze them at APL.

SELECTION CRITERIA AND ANALYSIS

Wireability

The objective of a wireability analysis is to determine, before layout, the ability to route a PWB. Often the board size is fixed for a particular application, and the board cost may be constrained as well. Wiring analysis can predict the number of layers required to wire the components for a particular circuit design and ascertain the routability of high I/O, single-chip packages. Since the routing effort and the layer count are directly proportional to board cost, wireability analysis is a useful tool for projecting substrate cost.

Basic Wireability Concepts

The basic concepts behind wireability analysis include wiring demand, wiring capacity, average wire length, and connectivity. Wiring demand D is the amount of wiring required to interconnect a given circuit. Wiring capacity C is the amount of wiring available for interconnection. Wiring demand is related to wiring capacity through the equation

$$D = \varepsilon C \, , \tag{1}$$

where ε is wiring efficiency. Typically, wiring efficiency is near 50%, since a substrate cannot be wired using 100% of its available capacity.

Wiring capacity, which is a function of the minimum signal line pitch P_s that can be fabricated on a given substrate technology, is usually normalized to a given square dimension such as 1 in. or 1 cm. Sometimes it is normalized to the size of pitch P_p of the chip carrier packages or bare die attached to the PWB. In this case (neglecting via holes and through-holes), for N_ℓ signal layers, the total wiring capacity C_ℓ is given by

$$C_\ell = P_p N_\ell / P_s \, . \tag{2}$$

To obtain the wiring demand, we must estimate the average length $\bar{L}$ per interconnection. This length between any two points is not "as the crow flies," since alternate layers of signal line are usually routed in x or y orthogonal streets. This length is sometimes referred to as the Manhattan length, which is given by $|x_1 - x_2| + |y_1 - y_2|$. The wiring demand is simply the number of pin-to-pin connections N_{pp} times the average wiring length $\bar{L}$. If we consider that for an average of N_{pins} number of pins per net there are $(N_{pins} - 1)/N_{pins}$ number of wires, then for $N_{I/Os}$ number of I/Os, the wiring demand is simply

$$D = \left(\frac{N_{pins} - 1}{N_{pins}} \right) N_{I/Os} \bar{L} \, . \tag{3}$$

To estimate average wire length, Rose et al.[1] assume that chips on a substrate interact with their nearest neighbors or second-nearest neighbors. This yields an average length of $\bar{L} = (P_p + 2P_p)/2 = 1.5P_p$. Rickerts[2] estimates this length to be approximately $\bar{L} = 0.77 P_p N_{chips}^{0.245}$, where N_{chips} is the number of chips to be interconnected.

Rent's Rule

Estimating the wiring demand in the absence of a preliminary layout requires the use of Rent's Rule, which is given by the relation

$$N_{I/Os} = ag^b \ , \tag{4}$$

where

$N_{I/Os}$ = number of chip, package, or substrate I/Os
g = number of circuits contained within the chip, package, or substrate,
a = Rent's coefficient, or the average number of connections per circuit or I/O, and
b = Rent's exponent.

Approximate Rent's Rule exponents and coefficients for various devices are given in Table 1.

Rent's Rule is especially useful in determining required circuit density and optimal device placement for a particular circuit design. In the following example, we extend the model described by Schmidt[3] for a PWB with an I/O connector on one edge to estimate the maximum wiring required by a four-sided MCM with I/Os on all sides. The total number of gates (or circuits) G on an MCM containing an arbitrary number of circuits can be found from the relation

$$G = \sum_{i=1}^{m} N_i g_i \ , \tag{5}$$

where N_i is the number of type i components, and m is the number of different devices. (This assumes that all of the circuits have the same Rent's relationship). The total package I/O count $N_{I/Os}$, assuming no interaction between circuits of different types (a worst-case wiring scenario), is given by

Table 1. Rent's Rule coefficients and exponents for various device technologies.

Chip/system	Rent's coefficient	Rent's exponent
Dynamic random-access memory (RAM)	6.20	0.085
Static RAM	6.00	0.120
Microprocessors	0.82	0.450
Random logic (gate arrays)	1.90	0.500
Computer systems	2.50	0.600
Chip/model level	1.40	0.630
Board/system level	82.00	0.250

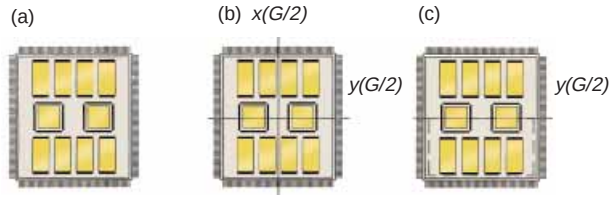

Figure 2. Circuit example to illustrate Rent's Rule technique for estimating the number of multichip module package leads. (a) Example of a multichip module. (b) Imaginary lines divide circuits such that half the logic gates fall on either side of the two lines. (c) Imaginary box incorporating half the total number of logic gates in the module.

$$N_{I/Os} = aG^b \ . \tag{6}$$

Consider the circuit shown in Fig. 2a, which contains two gate array devices with 5000 gates each and eight 64-MB static random-access memory (SRAM) chips. Using Table 1, we estimate the number of gate array chip I/Os to be 190 each and the SRAMs to be 52 leads each. The Rent's exponent for the gate arrays is assumed to be 0.5, the theoretical value for purely random logic. Since the SRAM circuits have a different Rent's exponent, we can develop an equivalent number of circuits for them that have the same Rent's exponent and coefficient as the gate arrays. We can do this by solving the equation $1.9(g_{SRAM})^{0.5} = 52$. Although the SRAMs have 64 million circuits, we have estimated the equivalent number of random logic circuits to be approximately 536.

Now applying Eq. 5, the total number of equivalent random logic gates G is approximately 15,992. The estimated total number of package leads can then be estimated from Eq. 6, i.e.,

$$N_{I/Os} = 1.9(15{,}992)^{0.5} \approx 240 \text{ leads.}$$

Thus, we can assume an MCM package with approximately 60 leads on a side.

In this example, we would like to estimate the maximum number of wires crossing a horizontal line. We assume that the maximum number of vertically traveling leads will occur at the coordinate point where there are $G/2$ gates above and $G/2$ gates below. The horizontal and vertical coordinates will be referred to as $y(G/2)$ and $x(G/2)$, respectively. For our example, these coordinates are illustrated in Fig. 2b.

To estimate the number of leads whose wires will cross this horizontal line, we place an imaginary box around the lower portion of the circuit as depicted in Fig. 2c. Using Rent's Rule, the number of leads whose wires will exit this imaginary box is simply

$$N = a(1/2G)^b = (1/2)^b N_{I/Os} \ . \tag{7}$$

Also, since we can assume that the number of leads leading to the package I/O pins located beneath this line will not affect the wire congestion in the vertical direction, we must reduce the number of package leads by a factor of $(1 - N'_{I/Os}/N_{I/Os})$, where $N'_{I/Os}$ is the number of package leads available to the $G/2$ circuits without crossing over the imaginary line $y(G/2)$. Thus, the maximum number of vertically going leads can be given by

$$N_h = (1/2)^b N_{I/Os}\left(\frac{1 - N'_{I/Os}}{N_{I/Os}}\right). \quad (8)$$

The number of wires crossing that horizontal line W_h is given by

$$W_h = aN_h . \quad (9)$$

In our example, the number of package leads N'_{Leads} available without crossing the line $G/2$ is 120. Therefore, assuming an average of 1.9 leads per the maximum number of vertically going wires crossing the horizontal line $y(G/2)$, the number of wires is estimated to be 161. By the symmetry of this particular example, the maximum number of horizontally traveling wires crossing the vertical line $x(G/2)$ would also be 161. If the module substrate is approximately 1.5 in. wide and the line-to-line pitch is 10 mils, then the wiring capacity is 150 wires per layer. However, with an assumed routing efficiency of 50%, the actual wiring capacity is 75 lines per layer. Thus, for a total of 322 wires, at least five signal layers would be required to route this design.

Ball-Grid Array Package

A BGA package has an advantage over quad flat packs in that its component footprint will grow more slowly with pin count. This increase in package pin density, however, can severely impact the routability of the PWB, as we will now show.

As a measure of wireability, we use the average wire or trace length $\bar{L}$. It has been shown that an average wiring length of $\bar{L} > 2.5P_P$ is easy to wire, whereas $\bar{L} < 2.0P_P$ is difficult to wire.[1] For our purposes here, we will consider $\bar{L} < 1.2P_P$ to be unwireable.[4]

Now, consider a square array of 256 I/O BGA packages (0.885 in^2) arranged on a 1-in. pitch. Using Eqs. 1, 2, and 3, we can express the average length as

$$\bar{L} = \varepsilon\frac{P_p N_{layers}}{P_s N_{I/Os}}\frac{(N_{pins})}{(N_{pins} - 1)}P_p . \quad (10)$$

For our example, we assume typical PWB wiring board design guidelines with 10-mil signal lines and 10-mil spaces (i.e., $P_s = 0.02$). For a 1-in. package pitch P_c, an average of 3 pins per net (N_{pins}), 6 signal layers (N_{layers}), a routing efficiency of 0.5, and 256 package I/Os, we get an average length of

$$\bar{L} = (0.5)\frac{(1)(6)}{(0.02)(256)}\frac{(3)}{(3-1)}P_p \approx 0.88P_p .$$

This is clearly an unacceptable number for routing BGAs on a 1-in. package pitch using conventional PWB design rules. Increasing the package pitch or the number of layers can increase the average length $\bar{L}$ to a reasonable value. However, only by varying the PWB design guidelines (reducing line widths and via diameters) can we improve routability without significantly increasing cost (more layers) or decreasing circuit densities.

Yield

One reason for the slow growth of MCM technology has been the unavailability of "known good dies" (KGDs). Because bare dies are hard to test at the wafer level, sources of KGDs have been slow to develop. Currently, the availability of KGDs from manufacturers has been limited to easy-to-test devices such as memory chips. Some chip suppliers, however, are procuring bare dies from manufacturers, testing them, and burning them in using a number of clever but cost-intensive schemes.

To deal with the demand and lack of availability of KGDs, manufacturers and chip suppliers have turned to several different solutions, including

- Exotic techniques: For example, first a dielectric is deposited, via holes are opened in the dielectric, the dielectric is metallized, and then the wafer is patterned with enlarged bond pads. Dies are wire bonded from the enlarged portion of the pad to a temporary substrate that is used solely for testing. After testing, the chip is removed from the temporary substrate and wire bonds are removed as well.
- Micro-sockets: Miniature sockets with tiny gold bump contacts are mounted on connectorized boards. The test die is then placed topside down into the socket so that the gold bumps are contacting the die.
- Lot sampling: This entails packaging and burning in a specified number of dies and determining a statistical probability of failure for a given lot.
- Wafer test structures: This approach calls for the attachment of test structures to wafers with features that are an order of magnitude more susceptible to common failure modes. These structures are packaged and burned in. If all the dies in a wafer lot pass

functional testing and the test structures survive burn in, then the dies are sold as KGDs.

Despite some progress by KGD suppliers, it is unlikely that all of the devices typically used in military or space applications will become available as KGDs. The question then becomes: Can reliable and cost-effective systems be manufactured using MCM and DCA technologies without access to KGDs? The best way to answer this question is by examining yield statistics.

The yield Y_m of an MCM or DCA assembly depends on the number of chips N_c and the chip yield Y_c after burn-in. This relation is given by

$$Y_m(\%) = 100 Y_c^{N_c} . \quad (11)$$

From Eq. 11 we can see that as the number of chips increases, the chip yield must improve to maintain board yield.

Depending upon the manufacturing methods, modules can often be repaired. The yield after n reworks is given by

$$Y_m(\%) = Y_c^{N_c(1-Y_c)^n} . \quad (12)$$

For example, a board with 20 chips and a chip yield of 90% will result in about 1 module in 10 working the first time, neglecting other components. With the same chip yield, the probability of having a good board after one rework rises to 81%. For a board with 30 ICs and the same chip yield, the probability of a good board would be about 4%. The probability after reworking the board would be 73%. In both cases, the probability of a board working after a second rework is better than 96%.

Sometimes, in cases where the number of modules required is low but the performance must be high (for example, in space applications), starting several modules to yield one is an acceptable practice. The probability that one module will be good after starting N_m modules with a module yield of Y_m is given by

$$P_1(\%) = 100\left[1 - \left(1 - Y_m\right)^{N_m}\right]. \quad (13)$$

For the case of a 30-chip module with a chip yield of 90% and an initial build of 3 modules, the probability of success after one rework is 98%.

Recently, Rozel[5] presented extensive chip yield data from Texas Instruments, Inc. (Table 2). The data indicate that most of their devices will experience a less than 1% fallout in subsequent module testing.

Table 2. MCM die yields for various devices.

Die type	Test performed (level tested)	Yield (%)
1-Mbit SRAM	DC, AC, multipattern (die)	99.4
4-Mbit dynamic RAM	DC, AC, multipattern (die)	99.4
Application-specific IC	Power, continuity (die)	99.0
Microprocessor	Scan, emulation (module)	65.0
Glue	Functional (module)	99.2

Note: Reprinted with permission from Ref. 5, a CMP Media, Inc., publication.

A graphical cost analysis illustrating the trade-offs associated with purchasing KGDs versus typical dies is shown in Fig. 3. The graph shows the relationship between number of chips used in a module versus module cost for various yield/rework and test scenarios. The cost model used to generate Fig. 3 is based on Eqs. 11 and 12 and assumes that KGDs are 3 times more expensive than ones that are only functionally tested. Depending on the volume of chips purchased, KGDs could cost anywhere from 2 to 10 times as much as untested dies. Rework cost is on a per-chip basis; module test cost increases with the number of chips per module.

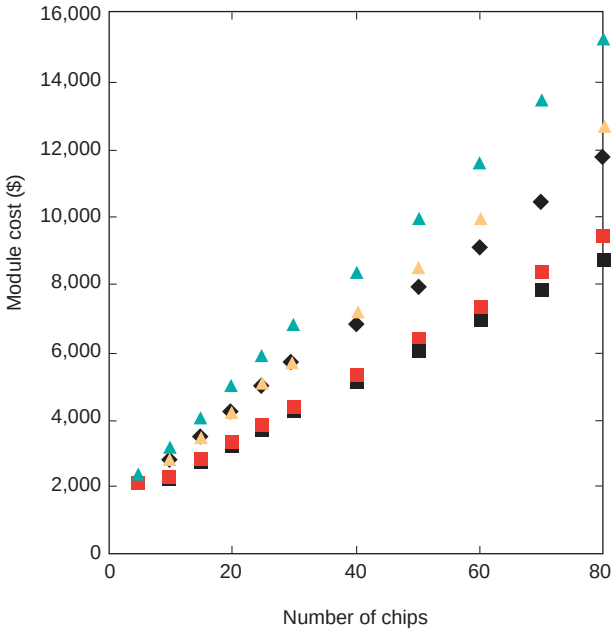

Figure 3. Multichip module cost as a function of chip yield, including effect of rework and test costs. (Diamonds = 99% chip yield/rework test cost high; black squares = 99% chip yield/rework test cost low; yellow triangles = known good die, no rework; green triangles = 90% chip yield/rework test cost high; red squares = 90% chip yield/rework test cost low.)

For this cost model, both high-yield dies (99%) and lower-yield dies (90%) are more cost-effective than KGDs when module rework and test costs are low. When repair and test costs are high, modules built with high-yield dies (99%) are still more cost-effective than KGDs, except for those built with less than 30 chips.

Cost

In an environment of limited capital for modern electronic packaging equipment, it becomes an economic necessity to predict which technologies will become dominant players in the near future. In addition to the capital expense, establishing a manufacturing capability also requires significant process development. For example, gearing up one's manufacturing capability to support a primarily tape-automated bonding technology when it first arrived on the scene 10 years ago would have turned out to be a major mistake. On the other hand, investing in SMT 10 years ago would have proved to be profitable.

Whereas choosing a technology based on performance criteria requires engineering analysis, the selection of cost-effective technologies for future spaceborne systems (e.g., electronic communications, missile guidance, underwater sensors) requires some cost modeling analysis. APL developed a cost analysis software package that employs wiring models discussed here and elsewhere[6,7] to compare performance advantages and fabrication costs among various electronic packaging approaches. The following example illustrates how this program was used to compare the manufacturing costs of three packaging technologies—MCM-L, standard SMT, and SMT using CSPs—to fabricate a conventional 25-chip digital signal processing unit.

The assumptions used for this analysis include aggressive line spacing, a wire bonded MCM-L, CMOS (complementary metal-oxide semiconductor) random logic, no repair requirements, and no hermeticity requirements. The program determines substrate layer count, estimates internal and external substrate I/Os using Rent's Rule, factors in chip dimensions and spacing, determines average and longest trace length, and includes a penalty factor for via type.

Cost parameters include the substrate fabrication in units of dollars/layer/cm^2, chip interconnection in units of dollars/interconnect, and package or encapsulation costs in units of dollars/mounted chip or package (Table 3). Repair costs were not considered here. Other costs associated with manufacturing microelectronic systems such as design, bare-board tests, system tests, and environmental tests were not considered because they are similar for all three of the packaging technologies in this example.

A cost comparison among the various packaging approaches investigated here is presented in Table 4. The analysis assumes a minimum quantity of 1,000 units. In Table 5, we compare cost and performance for these three packaging technologies. The results indicate that from a purely performance point of view, MCM-L is preferable to CSP and SMT. From a cost standpoint, in 1000-unit quantities, MCM-L is also preferred; however, for larger, 10,000-unit quantities, SMT with CSP become more cost-effective.

Heat Transfer

Characteristics in DCA and MCM-L Substrates

Of all the MCM substrate technologies, MCM-L has always been the least expensive to fabricate. In the past, the lower density of MCM-L has inhibited its use for high-performance MCMs. However recent technology advances have resulted in both smaller lines and smaller spacing (2 mils typically) and the development of laser-drilled microvias (8-mil diameter and less). In addition, some DCA applications are emerging that require components to be mounted on both sides of the high-density PWB. With the increased density of MCM-L and the desire to use both sides of

Table 3. Comparison of cost parameters for a 25-chip signal processing unit.

Technology	Low volume	High volume
Board/substrate		
PWB laminate	8 (dollars/layer/cm^2)	3–5 (dollars/layer/cm^2)
Interconnect		
Wire bond	0.02–0.05 (dollars/interconnect)	0.002–0.01 (dollars/interconnect)
SMT solder attach	0.02–0.03 (dollars/interconnect)	0.002–0.01 (dollars/interconnect)
Package		
Die attach and glob topping	1.25–1.35 (dollars/mounted chip or pkg.)	0.40 (dollars/mounted chip or pkg.)
SMT	3.00–6.00 (dollars/mounted chip or pkg.)	2.75–5.50 (dollars/mounted chip or pkg.)
CSP	6.00–12.00 (dollars/mounted chip or pkg.)	5.75–11.50 (dollars/mounted chip or pkg.)

Table 4. Cost comparison (in $) for MCM-L, CSP, and SMT: 25-chip signal processing unit.

Module type	Substrate	Assembly	Chip	Test	Package	Total
MCM-L	2812	95	2037	975	10	5929
CSP	3093	172	2287	975	20	6547
SMT	3271	172	2662	975	14	7094

Note: Assumptions = agressive spacing, wire bonded MCM-L, CMOS random logic, no repair, nonhermeticity.

Table 5. Cost/performance comparison for MCM-L, CSP, and SMT: 25-chip signal processing unit.

Module type	Volume	Average delay (cm^3)	Longest delay (ps)	Weight (ps)	Cost per module ($)	
					1,000 units	10,000 units
MCM-L	34.4	297	518	64.1	5929	4880
CSP	44.0	328	570	80.1	6547	3785
SMT	821.0	770	1344	690.0	7094	3547

Note: Assumptions = wire bonded MCM-L, quad flat packs for SMT, mature CSP technology.

the substrate for component attachment, the need to more accurately predict board performance has increased.

For thermal analysts, one of the most challenging problems is estimating board temperatures when conduction through the PWB is the primary heat transfer path. When air is plentiful, increasing airflow over components is a means to cool them. In space applications, however, heat must be transported through the board and through attached heat sinks to an external radiator. Heat pipes are an option, but the associated hardware adds weight and volume that MCM-L and DCA technologies seek to reduce.

The key to optimizing system weight and volume is to provide as much heat sinking as necessary to meet component temperature requirements. The analyst must therefore be able to accurately predict device temperatures, which is difficult owing to the extreme inhomogeneity of the PWB substrate. A typical PWB substrate has laminates consisting of a matrix of glass fibers and a polymer such as polyimide or epoxy. The laminates are usually coated on one side or two with a copper foil that range from 0.25 oz (0.35 mil thick) to 2 oz (2.8 mils thick). The copper foil on the laminates is etched into lines, planes, pads, etc. The laminates are then bonded to each other using relatively thick adhesive layers known as prepreg. The finished board is typically sprinkled with 1-mil-thick copper-plated via holes that connect conductors on different layers. Because of this complexity, performing an analysis using modern thermal analysis methods (e.g., finite element, finite difference), even on the smallest multilayer PWB, would require solving matrix equations with millions of matrix elements. Solving such equations would take an inordinate amount of time, even for today's high-speed computers.

One solution to this problem of complexity is to (1) categorize certain areas of the PWB with respect to line density, via hole density, and plane density, (2) perform a detailed analysis on each area yielding lumped equivalent thermal conductivities, and (3) perform a simplified thermal analysis using those lumped element properties. Since copper ground planes, copper signal lines, and plated via holes can impact both lateral and vertical thermal conductivity, it is useful to break down the PWB into representative sections such as

- High via density/high line density; typical of areas surrounding chips where chip I/Os are interconnected to the substrate
- High line density/low via density; typical of spaces between devices usually used for line routing
- High via density/low line density; typical of areas under chips (high via density is used to conduct heat away from heat-dissipating chips)
- Low via density/low line density; typical of stagnant areas along board edges and corners where available space is rarely used for routing interconnections
- Thermal vias; typically placed under heat-dissipating components to enhance heat conduction

A conventional cross section of a PWB section for a 10-layer board is illustrated in Fig. 4. We chose this

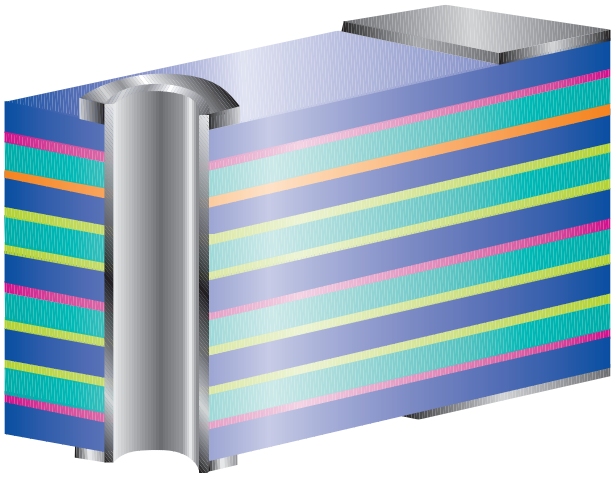

Figure 4. Cross section of a typical 10-metal-layer PWB.

board since it afforded a reasonable number of pad and signal layers (six) as well as four planes for controlling signal line impedance and assisting in lateral heat transfer. Thermal contours for one of the typical PWB sections (high line density/low via density) analyzed in this study using the finite element method are illustrated in Fig. 5. Results from the analysis of several representative PWB sections, also analyzed using the finite element method, are presented in Table 6.

Preliminary Thermal Analysis of Command and Data Handling DCA Assembly

We used the estimated thermal conductivities listed in Table 6 for the eight prototype PWB sections to perform a thermal analysis of a COB assembly intended for future space missions. The proposed component layout and housing of the board, which is a command and data handling (C&DH) system used in all spacecraft, is illustrated in Fig. 6. This COB design is a prototype for future APL Space Department electronic systems and is partially supported by a grant from NASA. The new design offers a 4:1 reduction in size over a previous SMT design. Components are mounted to both sides, making traditional heat sinking methods from beneath the device impractical. Heat transfer to the housing will occur along a narrow strip at the four edges of the board and out through the mounting gussets that are located on the bottom portion of the housing.

Table 6. Finite element method calculated in-plane and *z*-axis thermal conductivities for typical PWB sections for 10-layer command and data handling COB assembly.

Situation analyzed	Thermal conductivity (W/m-°C) *x–y* plane	*z*-axis
High line density/high via density	14.6	
High line density/low via density	25.6	
High line density/no vias	31.9	
High via density/low line density	12.2	
No lines/low via density	21.7	
No lines/no vias	23.6	
10-mil thermal vias/50-mil pitch		35.8
10-mil filled thermal vias/50-mil pitch		42.1

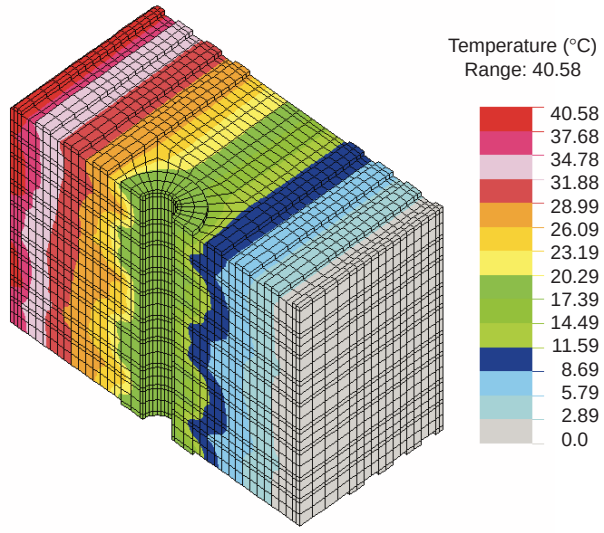

Figure 5. Thermal analysis results for a high line density/low via density section of a 10-metal-layer PWB.

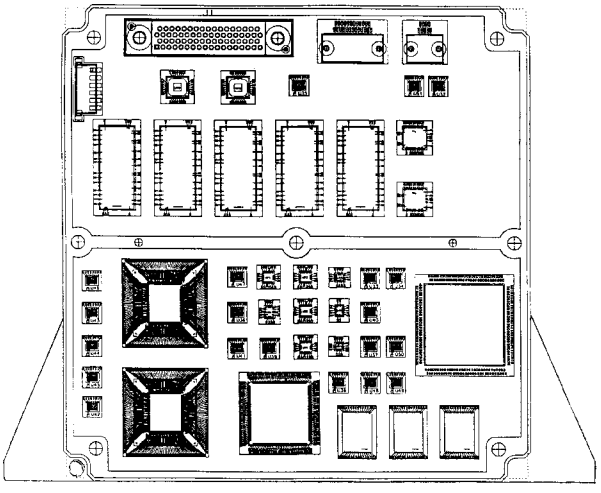

Figure 6. Component layout for the command and data handling (C&DH) COB circuit.

The finite element model of the C&DH PWB is depicted in Fig. 7; the various colors represent the different thermal conductivities used to model particular sections of the PWB. The board was modeled using thick, three-dimensional, second-order plate elements. In all, 831 elements and 2,568 nodes were used. The immense reduction in computational effort can be appreciated by comparing the representative section of Fig. 5, which was composed of 19,236 elements and 22,069 nodes. Isothermal contours for the C&DH COB assembly are presented in Fig. 8.

Electromagnetic Performance

New device technologies are rapidly pushing beyond our ability to package and interconnect these device while still preserving signal integrity. High-speed silicon-based computers are reaching clock speeds of 500 MHz with subnanosecond rise times. Gallium arsenide digital logic used in the fastest supercomputers is clocked at frequencies approaching 1 GHz, with rise times in the 150- to 250-ps range. The next large leap in computer processing speed may likely come from the superconducting, Josephson junction–based ICs such as superconducting quantum interference devices[8] (SQUIDs) and quantum flux parametrons[9] (QFPs). The rise times for superconducting devices are extremely fast and range from 10 to 100 ps.

Fast-rising digital signals typically have harmonic content well above the fundamental clock frequency of the circuit. The maximum significant frequency of a digital signal can be approximated by the equation

$$F_{\max} = \frac{0.5}{t_r}, \qquad (14)$$

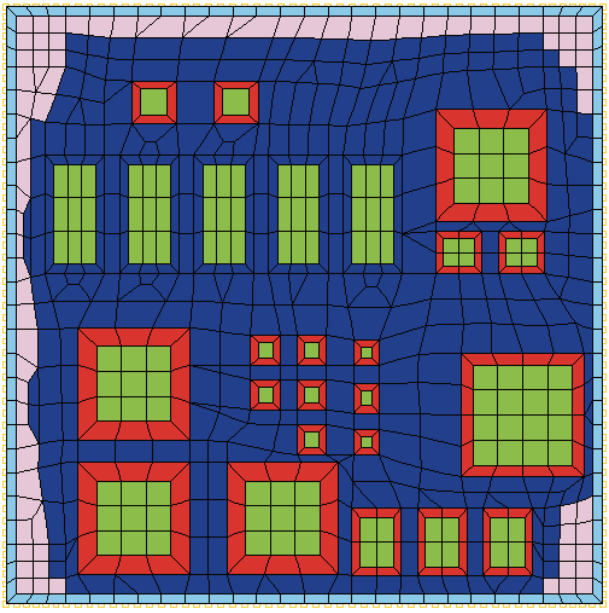

Figure 7. Finite element thermal model for the C&DH COB circuit.

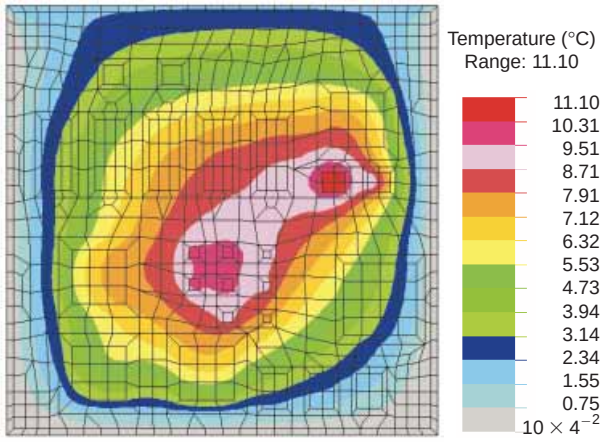

Figure 8. Isothermal contours for the C&DH COB circuit.

where t_r is rise time. Thus, for superconducting circuits with rise times of 10 ps, the maximum harmonic content will exceed 50 GHz.

At typical high-speed clock rates between 100 and 500 MHz, characterizing the electrical performance of various substrate technologies is a two-step process. The first step involves calculating transmission line parameters for the capacitance and inductance. For a lossy medium (poor conductors and lossy dielectrics), signal line resistance and conductance are also calculated. For multiconductor transmission lines, the capacitance, inductance, resistance, and conductance take the form of matrices.

Calculating the matrix elements for these matrices requires the use of an electromagnetic simulation program that can (1) simulate the electromagnetic fields in a multiconductor environment, and (2) calculate charge and current distributions from which the capacitance, inductance, resistance, and conductance matrices may be extracted. Field simulation programs employ either boundary element or moment methods when the field equations are cast in integral equation form, or they use finite element methods when the field equations are more conveniently expressed in differential equation form.

The second step involves casting the transmission line equations into matrix form as a set of n-coupled equations given by

$$\begin{bmatrix} 0 & L^{-1} \\ C^{-1} & 0 \end{bmatrix} \begin{bmatrix} \frac{\partial v}{\partial x} \\ \frac{\partial i}{\partial x} \end{bmatrix} + \begin{bmatrix} L^{-1}R & 0 \\ 0 & C^{-1}G \end{bmatrix} \begin{bmatrix} i \\ v \end{bmatrix} = - \begin{bmatrix} \frac{\partial i}{\partial t} \\ \frac{\partial v}{\partial t} \end{bmatrix}, \qquad (15)$$

where C, L, R, and G are the square n x n capacitance, inductance, resistance, and conductance matrices, respectively. The terms v and i represent the $n \times 1$

voltage and current vectors, respectively. The variable x represents the length along the transmission line, and t represents time.

For lossless transmission lines ($R<<L$ and $G<<C$), Eq. 15 can be written in the coupled form, i.e.,

$$\begin{bmatrix} \frac{\partial v}{\partial x} \\ \frac{\partial i}{\partial x} \end{bmatrix} = -\begin{bmatrix} L & 0 \\ 0 & C \end{bmatrix} \begin{bmatrix} \frac{\partial i}{\partial t} \\ \frac{\partial v}{\partial t} \end{bmatrix}. \tag{16}$$

These voltages and currents can be decoupled to yield the following eigenvalue equations:

$$\left\{ [L][C] - \lambda [I] \right\} v = 0 , \tag{17}$$

and

$$\left\{ [C][L] - \lambda [I] \right\} i = 0 , \tag{18}$$

respectively, where $[I]$ is the identity matrix and λ is the set of n eigenvalues.

Solving these characteristic equations for the associated eigenvalues and eigenvectors by following the procedure similar to that described by Chang[10] allows one to diagonalize the L and C matrices, yielding the set of n-uncoupled equations

$$\begin{bmatrix} \frac{\partial v}{\partial x} \\ \frac{\partial i}{\partial x} \end{bmatrix} = -\begin{bmatrix} L' & 0 \\ 0 & C' \end{bmatrix} \begin{bmatrix} \frac{\partial i}{\partial t} \\ \frac{\partial v}{\partial t} \end{bmatrix}, \tag{19}$$

where L' and C' are the diagonalized inductance and capacitance matrices, respectively.

An uncoupled characteristic impedance and propagation velocity can be given by the following relationships:

$$z_k = \sqrt{\frac{L'_k}{C'_k}} \tag{20}$$

and

$$v_{Pk} = \frac{1}{\sqrt{\lambda_k}} , \tag{21}$$

where L'_k and C'_k are the mode equivalent inductance and capacitance, respectively, and λ_k is the kth eigenvalue of Eq. 17. The uncoupled transmission lines from Eq. 19 lead to the equivalent circuit model for the two line system, as shown in Fig. 9. The subcircuits in this model consist of voltage controlled voltage sources, voltage controlled current sources, and delay lines. The matrix $[M_v]$ in Fig. 9a represents the matrix of eigenvectors resulting from the solution of Eq. 17. The matrix $[M_I^{-1}]^T$ in Fig. 9a represents the transpose of the inverse of the matrix of eigenvectors obtained from the solution of Eq. 18. This model is easily implemented in a circuit simulation program such as SPICE.[11,12]

One main disadvantage of this type of time-domain simulation method is the inability to easily characterize frequency-dependent variations in the propagation medium of a transmission line as well as frequency variations in the loads. At very high frequencies, the line attenuation and propagation velocity will vary significantly with the frequency.

A more suitable approach for analyzing very high-speed digital circuits involves the application of a high-frequency simulator, a computer program that usually uses finite element methods to estimate frequency-domain scattering parameters (S-parameters) instead of time-domain circuit parameters. The simulator is also used to cascade multiports together and estimate their net effective frequency-domain

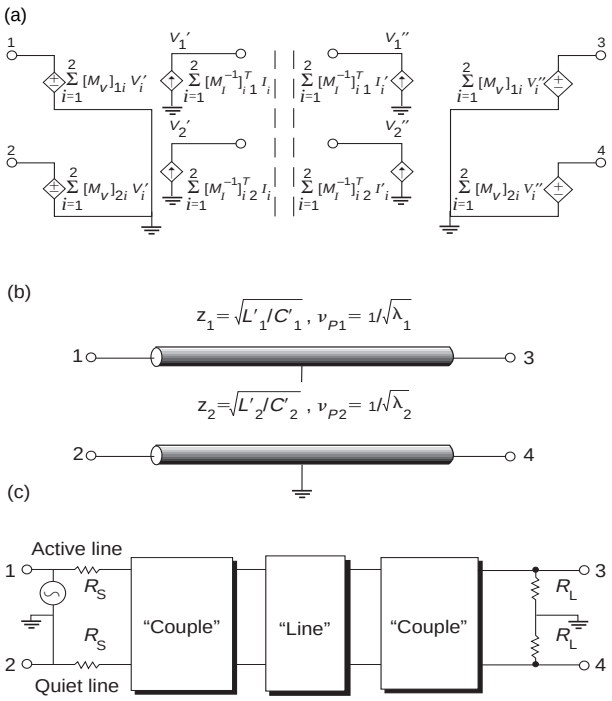

Figure 9. SPICE model used for time-domain analysis of voltage cross talk between adjacent transmission lines: four-port subcircuit (a) "couple," (b) "line," and (c) "net," where R_S and R_L are the source and load resistances, respectively.

S-parameters. The finite element formulation can account for skin effects, dielectric losses, coupling effects, and parasitic reactances caused by discontinuity structures such as transmission line bends and vias.

To obtain time-domain results, the frequency response for the complete network is computed. The impulse response is then calculated by performing an inverse Fourier transform on the frequency response. The time-domain waveforms transmitted through the circuit are then computed by performing a convolution on the circuit's impulse response with the desired input signal waveform.

Analysis of Vias in Low-Temperature Co-fired Ceramic MCM-Cs

This section illustrates the alternative frequency-domain approach for analyzing signal transmission in the time domain. Our example is the investigation of the degradation effect that a via has on a high-speed digital signal as it traverses through the via. The results from the via simulation presented here will subsequently be incorporated into a larger electrical model for a clock distribution circuit detailed in the next section. This clock distribution circuit is to be used in a superconductive network switching circuit. We can define a via as a vertical, electrically conductive path between one conductor layer and another within a multilayer circuit board. By using the complex boundary conditions of the via, we can estimate the S-parameters of this multilevel connection and then use the results in a frequency-domain simulator where the via is represented as a simple two-port network. When vias, represented as two-port networks, are cascaded with sections of transmission line networks for use in a frequency-domain simulator, an accurate and computationally efficient method for simulating high-speed signal transmission through the circuit board is achieved.

Vias are required in multilayer circuits to vertically translate signals between different conductor layers. As circuit elements, vias are normally modeled as small inductors. The magnitude of this inductance is usually approximated by

$$L = 5.08h\left[\ln\left(\frac{4h}{d}\right)+1\right], \qquad (22)$$

where L is the inductance in nanohenries, h is the via length in inches, and d is the via diameter in inches.

At very low frequencies this parasitic inductance has no effect on the signal as it propagates between layers. As the frequency content of the signal increases, the inductive reactance presented by this via may become significant.

We considered a low-temperature, co-fired ceramic tape from DuPont for a multilayer MCM substrate for high-speed superconducting ICs. The application is a high-speed (10-GHz) serial-to-parallel converter. The rise times of such devices are nominally 40 ps, with a maximum frequency content of approximately 12.5 GHz. The loss tangent for the DuPont 901 material is less than 0.002 for frequencies as high as 20 GHz. The gold conductor thickness for this fabrication process is typically 8 to 12 μm, and the surface roughness of these conductors may range from 0.33 to 0.86 μm. The skin depth for gold at 2.5 GHz is 1.6 μm, while at 10 GHz it is 0.79 μm. Since the skin depth is smaller than the metallization thickness, the series resistance will vary with frequency. For frequencies above 10 GHz, the skin depth is on the order of the surface roughness of the conductor, and the effective series resistance will increase even further. For high-speed digital signals, variations in resistance, propagation velocity, etc., with respect to frequency lead to dispersion and distortion of the pulse (e.g., rounding of pulse edges and spreading of pulse width).

The simulator used in this work was the Microwave Design System (MDS) High Frequency Simulator from Hewlett Packard. The via analyzed has a 203-μm (8-mil) diameter and a length of 610 μm (24 mils). Figure 10 compares the magnitude and phase of the S21 S-parameter for the inductance approximation (Eq. 22) and the results obtained from a three-dimensional finite element simulation. The approximation given in Eq. 22 for the parasitic inductance of a via is only accurate for determining the reflection or insertion loss caused by the via. Being an isolated lumped element, it does not accurately model the propagation delay (i.e., phase change) through the physical length of a via.

Although there is good agreement between the approximation of Eq. 22 and the finite element solution for the magnitude of S21 in Fig. 10a, as expected, the finite element solution gives a more accurate estimate of the phase of S21 as shown in Fig. 10b. The phase difference between the two approximation methods can be as much as 45° at 20 GHz. To obtain a more accurate lumped element model, the parasitic capacitance between the signal via and ground would have to be added.

Analysis of a Clock Distribution Network for a Superconducting Crossbar Switch Circuit

APL has been involved with the design and fabrication of a high-speed superconducting crossbar switch circuit for high-speed computer applications. One part of the development effort has been the design of a clock distribution network, which required a 50-Ω transmission line to feed four low-impedance superconducting devices. A fixed skew in the clock signals

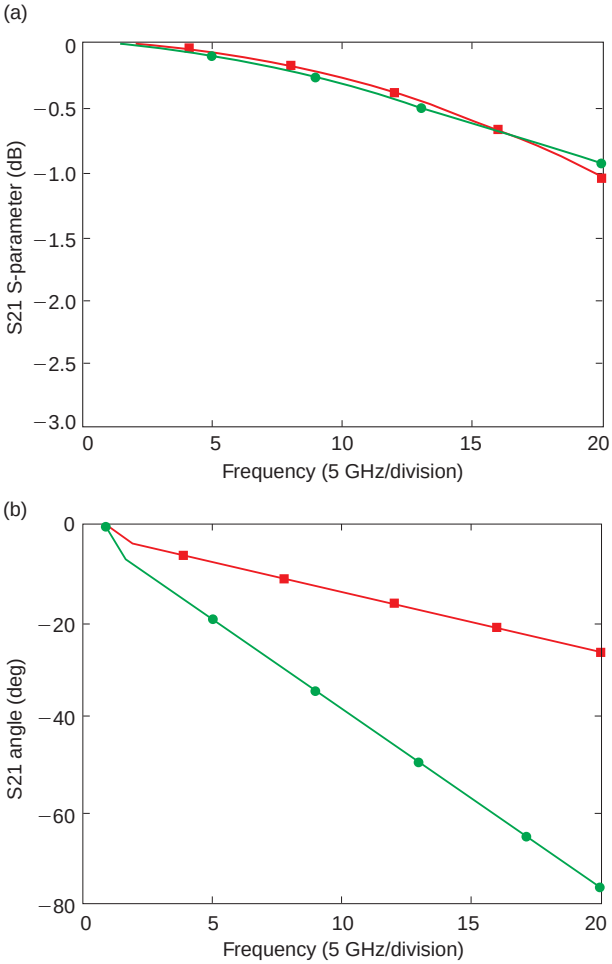

Figure 10. Estimated S21 S-parameter magnitude (a) and phase (b) for a multilayer ceramic via using an analytical approximation (red curves) and finite element method (green curves).

of 12 ps between inputs was also required. Since the voltage level required to switch the QFP logic was an order of magnitude lower than the input signal, the design would allow for one reflection back to the source, which then must absorb all of the reflected energy. Multiple reflections could not be tolerated, as they would compromise signal integrity at the QFP inputs and could affect the skew.

Since the signal was a fast-rising digital pulse, a narrowband impedance transformer was not acceptable. Resistor-based transformers are normally wideband, but the heat dissipation of the resistors would raise the temperature of the superconducting ICs above the superconducting temperature.

After some preliminary design evaluations, a spider clock circuit was designed that featured a 50-Ω line terminating directly into four 26-Ω legs. In the spider design, a single impedance mismatch exists at the point where the 50-Ω input line transitions into the four 26-Ω lines. Provided that the circuit is impedance matched at both the driver and receiver ends, an undistorted digital signal (although diminished in amplitude) may be transmitted. One preliminary spider design is depicted in Fig. 11.

A spider clock distribution test circuit was fabricated at APL using low-temperature co-fired ceramic multilayer wiring board technology. This test circuit comprised an input pad on the top of an MCM substrate that was connected to a 50-Ω stripline layer through a via. The other end of the 50-Ω stripline was connected to another via, which connected the signal to four 26-Ω signal lines in parallel. The line lengths of the 26-Ω lines were varied to effect a time delay of 40 ps between the shortest and longest lines. The 26-Ω lines transitioned to vias that terminated at probe pads on the top MCM surface. Each 26-Ω line was terminated with a high-frequency 26-Ω chip resistor.

Figure 12 shows the simulation schematic used for modeling the spider circuit. The more complex via geometries are modeled first using a high-frequency simulator, as discussed in the previous section. The S-matrices obtained from these simulations are then incorporated into the frequency domain circuit simulator, as indicated by the labeled boxes in the figure. The first via circuit model, via (M6-M2), connects the signal line on the probe layer (M6) at the top of the MCM to the 50-Ω signal translation layer (M2). Via (M2-M4) is then the circuit model representing the vertical translation from the 50-Ω layer (M2) to the 26-Ω layer (M4), where the signal splits in four directions. From the 26-Ω distribution layer (M4), via

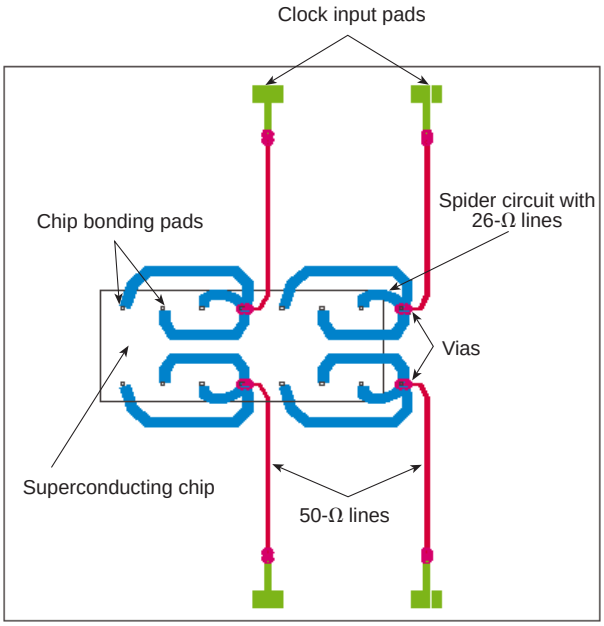

Figure 11. Candidate spider network for superconducting crossbar switch clock distribution circuit.

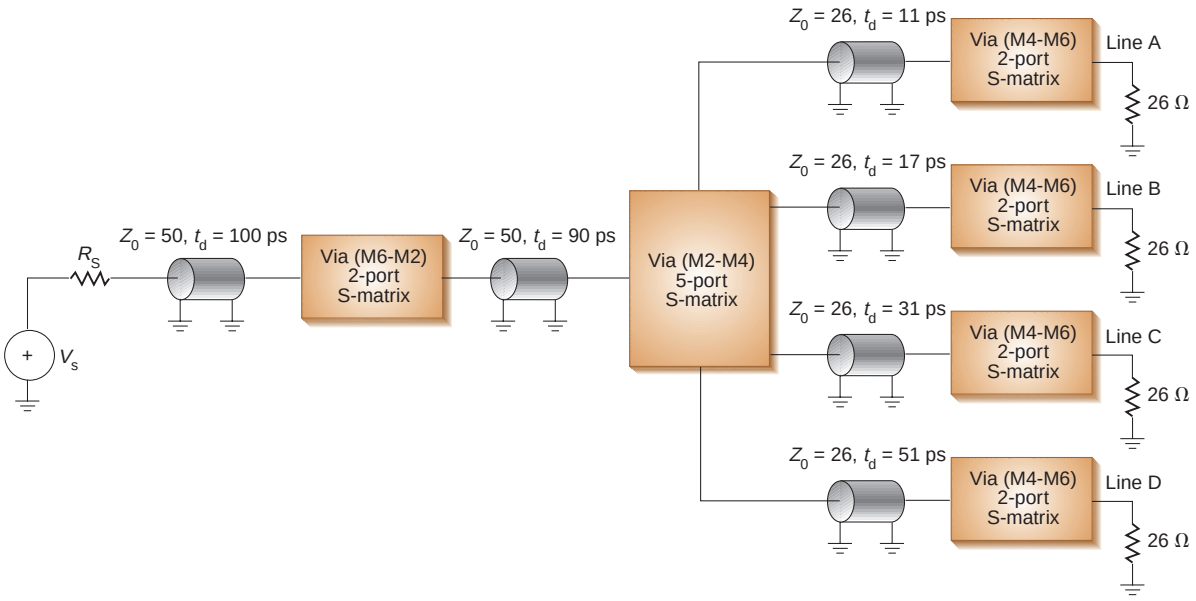

Figure 12. Circuit representation of a spider clock distribution network. (Z_0 = impedence, R_S = source impedence, t_d = propagation delay, and t_r = rise time; t_r = 20 ps and unit-step source = 0 to 2 V.)

(M4-M6) represents the vertical translation back to the uppermost layer (M6) of the MCM, where the output signals are probed across the 26-Ω chip resistance. The transmission line models in Fig. 12 characterize the horizontal propagation of the signals on the individual layers of the MCM. The excitation is a unit-step voltage source with a rise time of 20 ps.

Figure 13a shows the computed waveforms from the circuit simulation for a unit-step voltage source with a 10 to 90% rise time of 20 ps. These waveforms represent the signals transmitted to each of the four outputs of the spider circuit. Line A represents the shortest leg of the circuit with an equivalent propagation delay of 11 ps; line D represents the longest leg with a delay of 51 ps (see Fig. 12). The waveform for line A has a rise time of 38 ps; the waveform for line D has a rise time of 47 ps. The slower edge rate of line D may be attributed to transmission line attenuation since its path length is nearly 5 times greater than line A's. Figure 13b shows the effects of disparate signal edge rates on the overall skew of the circuit. The circuit was originally designed to have a time delay of 40 ps between the transmitted signals on line A and line D. Accounting for the unequal edge rates between lines A and D, the actual time delay measured is 44 ps.

SUMMARY

The relentless progression of IC manufacturers toward the fabrication of circuits with higher speeds and higher densities is constantly challenging the ability of packaging engineers to integrate ICs into highly functional, high-performance electronic systems. This article has discussed some of the methods developed to assist the electronic packaging engineer in finding solutions to the difficult problems associated with such high-performance electronics.

Some of the issues pertaining to the wireability of densely populated wiring boards and high I/O area array–style packages have been addressed and examples have been given. The message here is that higher density packages and devices will require design guidelines with smaller features such as lines, spaces, and vias.

Issues pertaining to manufacturing cost modeling were also presented. The KGD problem was examined in some detail. Several yield models were discussed and industry data provided that lend some insight into the problem. A cost model was described that showed how assembly cost may depend as much on repair and test costs as it does on the probability that a chip is a KGD. Cost comparisons were made among two prominent technologies, MCM-L and SMT, and a third more recent player, CSP. Cost analyses indicated that in smaller quantities, MCM-L would remain the least expensive laminate-based packaging technology, whereas in higher quantities, CSP was more cost-effective.

The difficult issue of thermally analyzing complex PWB laminates was presented, and methods for more accurately predicting device temperatures were provided. Finally, methods for analyzing signal integrity in a multiconductor environment were discussed. Several examples illustrating the use of these methods to solve current high-speed electronic packaging design problems were also given.

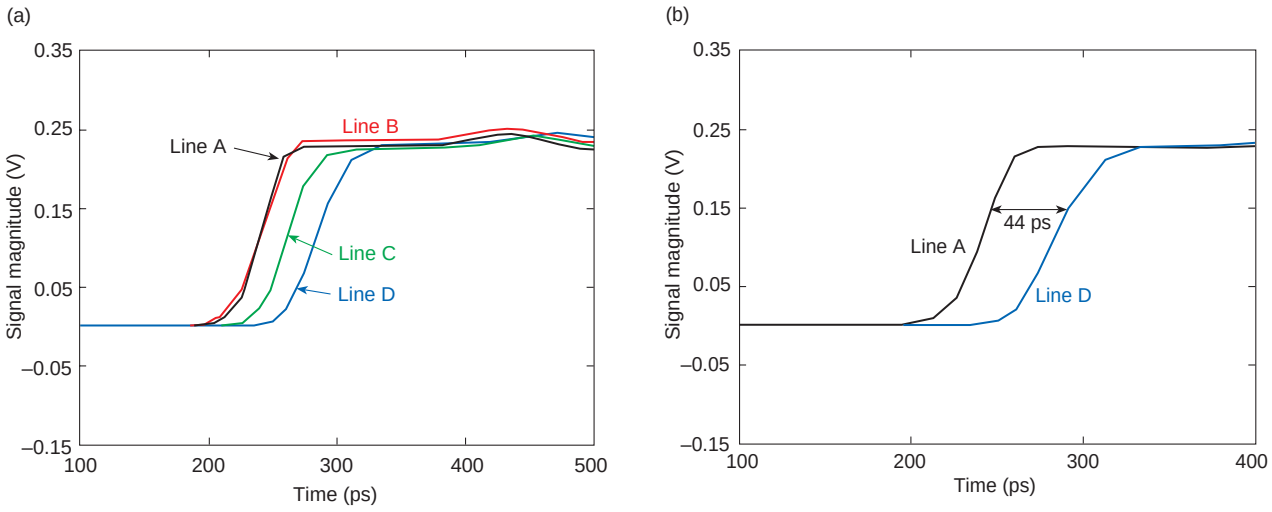

Figure 13. Simulated circuit performance for a spider clock distribution circuit. (a) Waveforms at 26-Ω loads for the four legs of the spider clock distribution circuit in response to a unit-step voltage with a 20-ps rise time. (b) Estimated signal skew between the shortest and longest legs.

REFERENCES

[1] Rose, K., Chiba, T., Heller, W. R., and Mikhail, W. F., "Package Wiring and Terminals," in *Microelectronics Packaging Handbook, Part I*, R. R. Tummala, E. J. Rymaszewski, and A. G. Klopfstein (eds.), Chapman & Hall, New York, pp. 130–195 (1997).

[2] Rickerts, R. G., "Design Tradeoffs," in *Electronic Materials Handbook, Vol. 1, Packaging*, ASM International, Materials Park, OH, pp. 18–44 (1989).

[3] Schmidt, D. C., "Circuit Pack Parameter Estimation Using Rent's Rule," *IEEE Trans. Comput.* CAD-1(4), 186–192 (Oct 1982).

[4] Amerson, R., and Kuekes, P., "A Twenty-Seven Chip MCM-C," in *SPIE Proc.* 2256, pp. 578–582 (1994).

[5] Rozel, L., "MCM Testing," *Electronic Engineering Times* (9 Oct 1995).

[6] Charles, Jr., H. K., "Cost Versus Technology Trade-Offs for Multichip Modules," *ISHM J. Microelectron. Electron. Packaging* 19(3), 295–300 (1996).

[7] Messner, G., "Cost-Density Analysis of Interconnections," *IEEE Trans. Compon. Hybrids Manuf. Technol.* CHMT-10, 143–151 (1987).

[8] Orlando, T. P., and Delin, K., *Foundations of Applied Superconductivity*, Addison-Wesley Publishing Co., Reading, MA (1991).

[9] Harada, Y., Nakane, H., Miyamoto, N., Kawabe, U., Goto, E., and Soma, T., "Basic Operations of the Quantum Flux Parametron," *IEEE Trans. Magn.* 23(5), 3801–3807 (Sep 1987).

[10] Chang, F. Y., "Transient Analysis of Low-Loss Coupled Transmission Lines in Non-Homogeneous Dielectric Medium," *IEEE Trans. Microwave Theory Tech.* MTT-18, 616–626 (Sep 1970).

[11] Tripathi, V. K., and Retting, J. B., "A SPICE Model for Multiple Coupled Microstrips and Other Transmission Lines," *IEEE Trans. Microwave Theory Tech.* MTT-33, 1513–1518 (Dec 1985).

[12] Clatterbaugh, G. V., and Charles, Jr., H. K., *Advanced Packaging for VLSI/VHSIC Applications: Electrical, Thermal, and Mechanical Considerations - An IR&D Report*, TG-1362, JHU/APL, Laurel, MD (Nov 1987).

THE AUTHORS

GUY V. CLATTERBAUGH received B.S. and M.S. degrees in physics from Drexel University. He has been a staff scientist with the Microelectronics Group and then the Electronic Services Group at APL since 1982. Mr. Clatterbaugh was appointed to the Principal Professional Staff in 1996. He has authored or co-authored over 60 papers concerning wire bonding, soldering, high-speed digital packaging, and multichip modules. Mr. Clatterbaugh's current interests include electro-optical interconnections, flip-chip bonding, and surface-mount soldering. He is currently a member of the International Microelectronics and Packaging Society. His e-mail address is guy.clatterbaugh@jhuapl.edu.

PAUL VICHOT received his B.S. degree in physics from Duke University in 1993 and an M.S. degree in electrical engineering from the University of Colorado at Boulder in 1997. He is currently working towards a Ph.D. in electrical engineering at the University of Colorado. Mr. Vichot recently joined APL's Strike Warfare Systems Group, where he has been working on the development of high-performance digital systems based on superconducting device technology. His current research interests include electronic packaging, microwaves, superconductivity, and digital systems. His e-mail address is paul.vichot@jhuapl.edu.

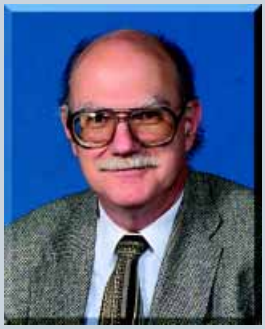

HARRY K. CHARLES, JR., received his B.S. in 1967 from Drexel University and his Ph.D. in 1972 from The Johns Hopkins University, both in electrical engineering. He joined APL in 1972 as a Postdoctoral Research Associate. He is currently Assistant Department Head for Engineering in the Technical Services Department and a member of the Principal Professional Staff. Dr. Charles has worked for 25 years in the microelectronics arena and is a specialist in electronic devices, systems, packaging, and reliability. His latest interests include multichip module design, fabrication, and testing; advanced interconnect; biomedical instrumentation; and novel sensors based on MEMS technology. He is a Fellow and former President of IMAPS (The International Microelectronics and Packaging Society), a Fellow of the IEEE, and a member of the Board of Governors of IEEE's Components Packaging and Manufacturing Technology (CPMT) Society. His e-mail address is harry.charles@jhuapl.edu.