

# A Multichip Module Design Process for Notebook Computers

Paul D. Franzon, North Carolina State University

Robert J. Evans, IBM Personal Computer Company

**Multichip modules can obtain significant performance improvements with a reduced footprint. This article examines the packaging-system design process, with application to choosing an MCM card for a 386SL notebook computer.**

**M**any of today's notebook computers — some weighing seven pounds and reaching three inches in thickness — hardly have notebook proportions. The ideal notebook computer would have the performance of the best desktop machine yet still be less than half an inch thick, with the area of a sheet of letter-sized paper. To achieve these goals, advanced packaging in the form of multichip modules must be considered. First, however, we examine some related trends in notebook design.

The thickness of a notebook computer is limited primarily by the height of its disk drives. As disk drives are replaced by (flash) memory cards,<sup>1</sup> the computer's thickness is limited by the bulk of the printed circuit board (PCB) that holds the electronic components, the keyboard size, and the screen dimensions. Minimizing thickness will then require that the PCB be populated on one side only. To maximize the number of slots reserved for memory cards and other cards (like a modem), one must minimize the PCB area. System size will continue to shrink as keyboards are replaced with pen-based notepads, placing additional constraints on the area available for the PCB.

Area can be reduced through greater integration, but this can be taken only so far. Advanced packaging options like multichip modules (MCMs) offer significant size reductions over conventional, single-chip packages. We analyze the effect of applying such advanced packaging options to a typical notebook computer design. For this analysis, we modified the publicly available Intel Mustang design.<sup>2,3</sup>

Here, we present a process for deciding on the appropriate MCM packaging and illustrate this process by applying it to a modified Mustang design. The main intent of the article is instructive; we do not attempt to define the best approach for packaging a notebook. We adapted the design process (but not the application) from Doane and Franzon.<sup>4</sup>

## System aims

In the design used for this analysis, two large chips contain the majority of the functionality. The 80386SL contains the 32-bit integer CPU, memory management, bus control, and buffering, while the 82360SL companion chip contains

## System-packaging decision process

1. Determine the system requirements and goals.
2. Express the goals and requirements in terms of performance and cost factors.
3. Determine the suitable alternatives for packaging, partitioning, and floorplanning.
4. Evaluate the performance and cost of each alternative.
5. Make the final decisions.

control functions for the system, I/O, peripherals, power management, and a majority of the glue logic required by the system. The memory system is configured with a 32-kilobyte cache, backed up by 2 megabytes of dynamic random-access main memory. The remainder of the design consists of the usual notebook peripherals: an LCD/CRT video controller, I/O ports, and a DC power supply. We consider how to minimize the area occupied by these electronic components by partitioning part of the design into an MCM.

The accompanying sidebar shows the five-step process we used in making the packaging decisions. The first step consists of determining the requirements and goals of the system.<sup>5</sup> Requirements represent the specifications that abso-

lutely must be met to ensure a properly operating and useful design. They become constraints that limit the design choices. The absence of any of these items in a design alternative immediately rules out any further consideration of that alternative. These items must be measurable; a design alternative either meets the specification or it does not. Goals are desired items that are not absolutely necessary for a successful design. These items are judged on their relative performance against other goals, through a trade-off analysis. They can consist of directional, rather than absolute, goals.

The next sidebar shows the requirements and goals for our design. CPU speed appears under requirements, as the 25-megahertz speed is both measur-

able and necessary. Minimum area appears under goals, as it does not have absolute limits but is used to judge the relative goodness between the available alternatives.

The second, third, and fourth decision steps occur somewhat concurrently. They involve evaluating a set of design alternatives against the requirements and goals. The packaging design alternatives involve the following elements:

- packaging technology or technology mix,
- chip partitioning between different packaging technologies, and
- floorplan (or relative position) of chips on the packages.

To minimize design time and time-to-market (a goal), we did not consider alternative chip designs. To evaluate the design alternatives against the requirements and goals, you must first express the goals in terms of factors that can be numerically evaluated for each design alternative. We consider these factors next.

## Packaging effects

Fundamentally, packaging limits system performance and increases its cost.

## Requirements and goals

### Requirements

1. 25-megahertz CPU speed
2. 2-megabyte main-memory size
3. 64-kilobyte cache-memory size
4. Passively air-cooled (no fan)
5. Upgradable to future CPU core designs
6. At least two PCMCIA slots for solid-state memory cards
7. Area equal to larger of keyboard or screen dimensions

### Goals

1. Minimum area and weight
2. Maximum number of PCMCIA slots
3. Maximum battery life
4. Low production cost
5. Minimum design-cycle time (quick to market)
6. Maximum field reliability

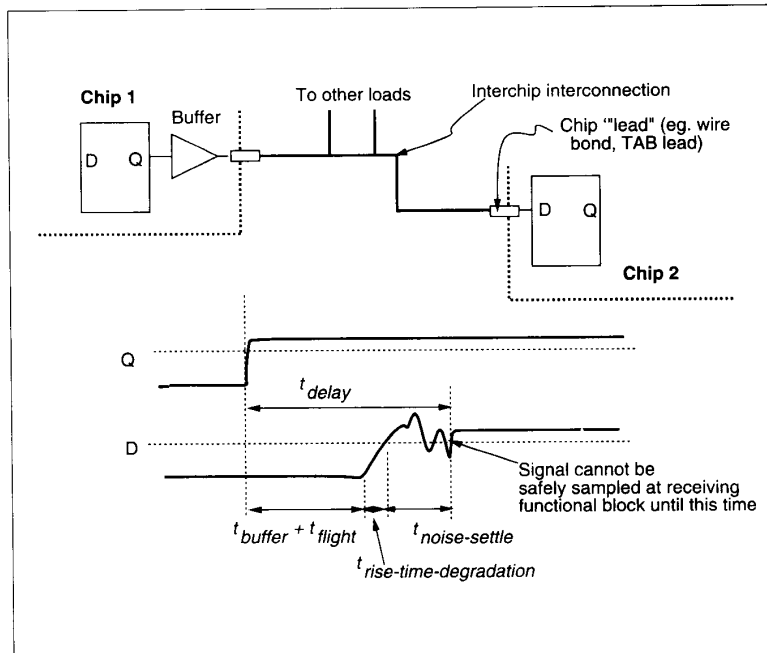


Figure 1. Elements of the time delay associated with packaging.<sup>4</sup>

This section lists these fundamental factors, defines them when appropriate, and shows how they interrelate.

**Performance factors.** The following performance factors relate the packaging-design alternatives to the system-performance parameters.

**Size and weight.** Size can mean any combination of area, height, volume, and form factor. The size and weight of a particular packaging option is especially important in portable systems. In larger systems, the size requirements are usually thought of as two-dimensional measurements. Here, both area and height are important. Although MCM packaging provides substantial area and weight reductions, you should consider the overall system in a size and weight determination. For example, moving the chips closer together in an MCM design increases the heat density. Consequently, the desire to avoid a large, heavy heat-removal subsystem could limit minimum chip spacing.

**Power consumption.** The energy consumption for a CMOS circuit can be expressed as  $E = CV^2f$ , with  $C$  representing the total capacitance being switched by the circuit,  $V$  the voltage swing, and  $f$  the frequency at which the circuit toggles. The switched capacitance depends on the length of the interconnect between successive circuits and the loads provided by those circuits. For package interconnect circuits, MCMs reduce this capacitive load by providing shorter, less capacitive lead attachments to the chips and by allowing shorter and narrower interconnections between chips. Minimizing power consumption maximizes battery life and reduces the size of heat sinks.

**Time delay.** The total delay associated with the packaging can be expressed as the sum of the driving buffer delay, the flight time of the signal along the interconnect, the rise-time degradation, and the settling time of the signal (Figure 1). Using MCM packaging substantially reduces all these delay contributors. The buffer delay is reduced because the capacitive load is reduced. The time of flight delay is given as

$$t_{\text{flight}} = \frac{l}{c/\sqrt{\epsilon_r}}$$

where  $l$  is the length of the interconnect,  $c$  is the speed of light in a vacuum, and  $\epsilon_r$  is the effective relative dielectric constant of the interconnect. When chips are placed close together in MCM technology, lengths are substantially shorter than in single-chip implementations. The various technologies differ substantially in terms of their value for  $\epsilon_r$ . The rise-time degradation is mainly determined by the inductance and capacitance of the chip leads, as well as line losses. Rise-time degradation decreases as the attachment leads become shorter, especially with solder-bump flip-chip attachments. The noise-settling delay is determined by how much time must elapse before electrical noise caused by a transition settles so that the signal can be clearly recognized as a 0 or a 1 level. (We discuss electrical noise later.) In computer designs, the interconnect delay partially determines the memory-access delay times. In the case of this 25-megahertz notebook design, the interconnect delay requirements are easily met and require no special considerations.

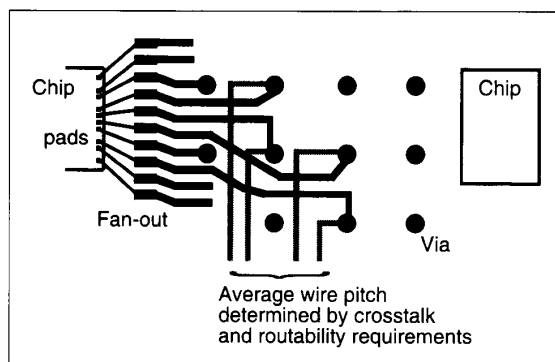
**Interconnection density.** This density in PCBs or MCMs is inversely proportional to the average pitch of the wires. Fundamentally, the main advantage of MCM technology is the high interconnection density achieved through fine line wiring. High interconnection densities allow closely placed chips to be directly attached to the substrate, making for small, fast systems. Figure 2 illustrates the reasons for this. First, you must match the pitch of the pads on the chip, which can be spaced as close as 75 microns. If this is impossible, a fan-out structure is needed. Second, the chips must be wired together. The average wire pitch determines how much inter-

chip area is needed for wiring and thus how closely chips can be spaced. The via size and pitch, routability, and crosstalk considerations determine average wire pitch.


**MCM-PCB connection density.** The density of wires available to connect an MCM to a PCB is generally lower than the density of wires within either of these approaches. For example, pin spacing in a typical socket is limited by how close holes can be drilled on the PCB. Furthermore, dense, high-pin-count connectors are expensive. These issues must be carefully considered in partitioning a design among a collection of MCMs and PCBs. Heller and Mikhail<sup>6</sup> present an extensive discussion on the analysis of interconnection and connection functions of packaging.

**Electrical noise.** The electrical noise within a system originates from three main sources — reflection noise, crosstalk noise, and simultaneous switching noise (SSN). The system also produces radiation noise, commonly referred to as electromagnetic interference (EMI) noise, which can adversely affect neighboring systems if not controlled. Reciprocally, the system itself is susceptible to outside sources of EMI.

Reflection noise originates from changes of impedance along the signal interconnect path. This is analogous to the reflections seen from a light beam as it encounters a sheet of glass: at each change in medium, the “impedance” changes and a portion of the light is reflected back along the path. This noise can be minimized by making the impedance as constant as possible along the interconnect paths (which requires using ground and power planes, among other considerations), by keeping the



**Figure 2. Spacing between chips is determined by interconnection density properties. Black wires run on the surface layer, gray wires on the layer underneath.**



**Figure 3.** Heat from the chips can be removed through the substrate or through the backside of the chips.

interconnects as short as possible, and by using terminations with matching impedance where needed.

Crosstalk noise results from the electromagnetic coupling of two parallel signal lines. This noise generally increases with any increase in length of the parallel sections and any decrease in line spacing or the rise time of the signal. Crosstalk requirements often force the line spacing to be increased beyond the minimum manufacturable pitch.

SSN originates when several large driver circuits are switched at the same time, producing a current spike on the power and ground lines. This current flows through the inductance of the chip attachment leads and becomes a voltage spike, which appears on adjacent quiet signal lines attached to the same power and ground pins, producing noise and possibly false switching. The chip attachment methods offered in MCMs result in a much smaller lead inductance than found in single-chip packages, thus reducing the SSN for a given circuit. Decoupling capacitors are typically placed between power and ground connections to help control SSN.

The signal lines also act as antennas and produce EMI noise when signals change on the lines. Meeting government EMI requirements can be difficult and often involves increasing the metal shielding of the system enclosure (in a notebook; this is done by coating or impregnating the plastic case with metal), which is counterproductive to the low-weight system goal. Using MCMs makes this goal much easier to obtain by shortening the many "antenna" lengths in the interconnect circuitry. (Bakoglu<sup>7</sup> provides a basic discussion of electrical delay and noise, and Doane and Franzon<sup>4</sup> discuss the process of electrical design. Messner et al.<sup>8</sup> discuss many MCM-specific considerations, as do Johnson, Teng, and Balde.<sup>9</sup>)

*Chip temperature and thermal dissipation.* Heat removal is often a critical issue in MCM systems, even CMOS systems. If the heat is not removed at an adequate rate, chip temperatures rise, which slows operation and reduces reliability. Heat is a performance-related issue because achieving the desired chip temperature affects system size and interchip distances. The higher the heat density, the greater the effort required to remove the heat. Heat is removed through the MCM substrate and/or through the backside of the chip into a circulating fluid, which is air or a liquid coolant (some air-cooled alternatives are shown in Figure 3). Heat removal is much more efficient when the coolant flow is forced by a fan or pump.

Thermal issues often force trade-offs with other performance factors. It might be necessary to increase chip spacing to reduce power density. Materials with low dielectric constants are often poor conductors of heat. To avoid using a high dielectric constant material, it might be necessary to place copper slugs or thermal vias in the heat removal path. However, their presence beneath the chips reduces the amount of signal routing. (Thermal slugs are larger than thermal vias. Beneath a chip, you might find a single copper thermal slug or an array of copper thermal vias, depending on the packaging technology.) For a more detailed discussion of MCM thermal issues and thermal modeling and design, see Doane and Franzon,<sup>4</sup> Messner et al.,<sup>8</sup> and Johnson, Teng, and Balde.<sup>9</sup>

**Cost factors.** The following cost factors relate packaging-design alternatives to system-cost parameters.

*Production.* These costs include manufacturing, assembly, test, and repair and include such items as MCM substrates, bare chips, MCM-PCB connec-

tors, and heat removal mechanisms. Table 1 provides some typical MCM (and other) packaging costs obtained from vendor quotes. However, be careful in using this data because the cost structure of an evolving technology can change rapidly.

Figure 4 shows a simplified description of a manufacturing process for MCM-based products. The testability and yield aspects can have a significant impact on the production costs of MCM systems. If the substrates and chips are not fully tested and burned in before assembly, the final yield may be low or expensive repair steps may be needed. (Traditionally, full-speed testing and burn-in are traditionally not done before the chips are packaged. With single-chip packaging, the bare chips are tested at a slow speed because of tester limitations.) This is particularly true if the MCM contains several large, complex chips. Unfortunately, it is currently expensive to test bare chips at full speed. One approach is to mount the chips by using tape-automated bonded (TAB) attachment and testing them through the TAB leads. Generally, if the MCM contains more than one high-cost chip, it is preferable to use fully pretested chips rather than face the potential cost impact of low yield or high repair rates. On the other hand, if there are no high-cost chips or the percentage of chips that pass a full-speed test is high, there is no need to pretest the chips before assembly. Vardaman and Ng<sup>10</sup> further discuss these trade-offs.

*Reliability, repairability, and maintainability.* System reliability determines the number of field failures, repairability determines the repair time required to correct failures, and maintainability determines the time spent on regular maintenance. All three impact the overall life-cycle costs of the product. System reliability is ensured by eliminating possible failure mechanisms at the time of the design, and repairability and maintainability are enhanced by making the system as modular as possible, with fast and accurate field diagnostics available. By reducing the number of solder joints, MCMs provide for substantial reliability improvements. However, care must be taken so that a small MCM-based system, coupled with an advanced cooling system, is still easy to maintain and repair.

*Design and prototyping.* The cost of designing, prototyping, and debugging a system can be larger for MCM alternatives. The current industry infrastructure for obtaining qualified (that is, fully pretested) bare die is small, although some chip makers (including Intel<sup>11</sup>) are beginning to offer certain parts in unpackaged, bare-die formats. Correct-by-construction design methods and design for test become more important when dealing with MCM products because they are more difficult to probe and diagnose for faults. However, fault diagnosis is also difficult for PCB-based designs of comparable performance. Investment in good design methodology pays off no matter what the technology.

*Design-cycle time.* The design-cycle time is a cost factor because being late to market can substantially reduce profits. The design-cycle time varies in proportion to the complexity of the product; it is now less than a year for most low-end computers. The advantages in being the first to market a product include the early establishment of sales and market share, increased time for enhancing the production process to increase the yields and minimize costs, and being able to work on your next-generation design while your competitors are still working on their first generation.

Introducing new technology into the middle of a tight design cycle is obviously undesirable. If you are thinking about using MCM technology, it is important to go through the design cycle once with a noncritical product so that potential problems can be identified and resolved before full-scale design and production.

Both design cost and design-cycle time are reduced by reusing existing designs as much as possible, which we achieved by using the Intel-provided Mustang design.

## Requirements and goals

At the start of this article, we discussed how the system size and height goal — together with the system goal of maximizing the number of solid-state memory and other PCMCIA (Personal Computer Memory Card-International Association) slots — led us to the packaging goal of minimizing the area occupied by the system's electronics. We concluded that these components had

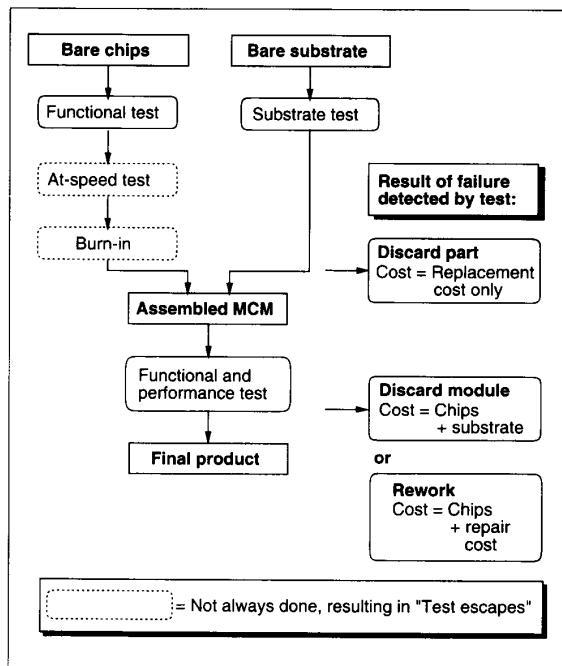
**Table 1. Typical high-volume manufacturing vendor costs for various packaging and interconnect elements. NRE stands for nonrecurrent engineering costs, which cover the design and tooling of special equipment.**

Package Type	Typical Cost
Plastic package	\$ .05-\$5.00
Ceramic PGA < 144 pins	\$5,000 NRE + \$ .10 per pin (e.g., 144 pins: \$14.00 or more)
Ceramic PGA > 144 pins (e.g., 500 pins: \$50.00 or more)	\$25,000 NRE + \$ .10 per pin
Cofired ceramic MCM	\$3.00 per sq. inch per layer (e.g., 10 layers: \$30.00 per sq. inch)
Thin-film MCM	\$60.00 per sq. inch (expected to decrease to around \$20.00)
Laminate MCM	\$3.00 per sq. inch
PCB	\$1.00 per sq. inch
CMOS chip wafers	\$25.00-\$150.00 per sq. inch

to fit on the smallest possible single-sided board.

This packaging-area-factor goal, together with the system requirement for an upgradable CPU, led us to consider partitioning the CPU core as a separa-

ble daughter card, whether MCM or PCB. This collection of components represents the set that would have to be replaced in a true performance-oriented CPU upgrade. (Though you can upgrade a computer by replacing just the



**Figure 4. A simplified manufacturing process for MCM-based designs.<sup>4</sup>**

CPU chip with a faster one, the best performance improvement is obtained by replacing CPU memories with proportionally faster ones.) This group also contains the two highest I/O count parts and has high levels of interconnectivity. Thus it could benefit more from the greater interconnection densities offered by the MCMs than most of the rest of the system. The video core (not considered here) also has high levels of interconnectivity and thus could benefit from advanced packaging.

In the next section, we discuss a number of alternative technologies for packaging the CPU core. First, we relate the system goals of battery life, reliability, and size to the packaging performance factors of power consumption, chip temperature, and interconnection density.

For our purposes, it is useful to measure power consumption in terms of the energy required by the CPU core to execute one million instructions. When operating at full speed, the CPU executes instructions at a rate of about 6 million instructions per second (6MIPS). The contribution that packaging makes to the power consumption is due mainly to the power dissipated in charging and discharging the interconnect of the memory paths. Based on the properties of

the Intel processors and the memory configuration in this design, we estimated that the cache memory path is exercised 3.6 times and the main DRAM memory path is exercised 0.88 times for each instruction executed. These figures, along with the  $CV^2$  transition energy for each memory path, can be used to determine the approximate amount of energy dissipated per instruction, due to CPU core packaging effects.

The most difficult factor to analyze is thermal dissipation. At 6 MIPS, the CPU core of this system design typically dissipates 8.5 watts. The "low-power" chip design does more to reduce the average (rather than the peak) power; we must use peak power when calculating system temperatures. This section dissipates more heat than any other group of chips contained in the system. Meeting the stated reliability goal with no cooling fans is difficult for all packaging options, even with the low heat densities of the single-chip packages. This is also complicated by the fact that the height goals restrict the air flow further and prevent us from using large heat sinks. Moreover, analyzing the heat-flow patterns is complex. However, the net result is that every effort has to be made to provide multiple, good, heat

paths away from the CPU and controller chips. In particular, in two of the MCM technologies to be considered, copper thermal slugs or vias must be used beneath the CPU and controller chips, making it difficult to route lines underneath these chips.

Figure 5 shows the major interconnection paths within the CPU core and their wire counts. It also shows the floor-plan we considered. Here, the size is evaluated by dividing the wire count by the product of the average wire pitch and the number of wiring-plane pairs. (In a wiring-plane pair, one plane is typically reserved for *x*-axis running wires and another for *y*-axis running wires. This orthogonality prevents wires from running on top of each other to reduce crosstalk noise.) An allowance is made for any under-the-chip routing and for chip-to-MCM and MCM-to-PCB connection area requirements. These interconnect requirements, along with the interconnect densities suggested in Table 2 for the various MCM substrates, can be used in a routability analysis to determine the final size of the required daughter card.

The MCM substrate technology alternatives that we considered for packaging of the daughter card consist of PCBs, laminate MCMs, high-temperature cofired ceramic MCMs, and deposited thin-film MCMs, with wire and via sizes and densities as shown in Table 2. Briefly, a laminate MCM is essentially an advanced PCB with fine wires and small drilled via holes; a cofired ceramic MCM uses much the same technology employed to make ceramic pin grid arrays; and a thin-film MCM contains very fine lines made by using chip multilayer deposition metallization processes on deposited organic (usually polyimide) insulators.

## Daughter card packaging

We now evaluate a set of packaging alternatives for the daughter card, applying the performance and cost factors. The factors that differ most among the alternatives are the size, as determined by interconnect density, and the production cost. We also consider the power consumption and SSN properties for each design. The former has to be evaluated for one of our goals, and

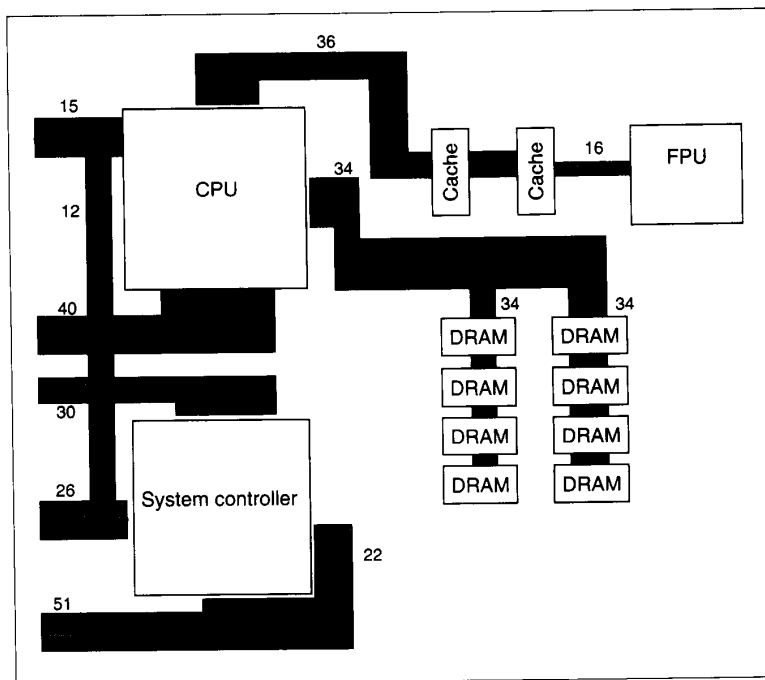


Figure 5. The interconnection pathways for the CPU core and the number of signal wires in each pathway.

**Table 2. The packaging alternatives analyzed in this study and their interconnection density properties.**

Alternative	MCM Type	Min. Wire Width and Space	Via Pitch	Average Wire Pitch
A	PCB with single-chip plastic surface-mount packages	150 $\mu\text{m}$ (6 mils)	2,500 $\mu\text{m}$ (100 mils)	1,000 $\mu\text{m}$ (40 mils)
B	PCB with TAB-mounted die	150 $\mu\text{m}$ (6 mils)	2,500 $\mu\text{m}$ (100 mils)	1,000 $\mu\text{m}$ (40 mils)
C	Six-layer laminate MCM with wire-bonded die	75 $\mu\text{m}$ (3 mils)	1,250 $\mu\text{m}$ (50 mils)	640 $\mu\text{m}$ (25 mils)
D	Seven-layer cofired ceramic MCM with TAB-mounted die	125 $\mu\text{m}$ (5 mils)	125 $\mu\text{m}$ (5 mils)	375 $\mu\text{m}$ (15 mils)
E	Thin-film MCM with wire-bonded die	25 $\mu\text{m}$ (1 mil)	25 $\mu\text{m}$ (1 mil)	75 $\mu\text{m}$ (3 mils)

the latter is the most important noise source in this design, due to the large number of switching bus drivers.

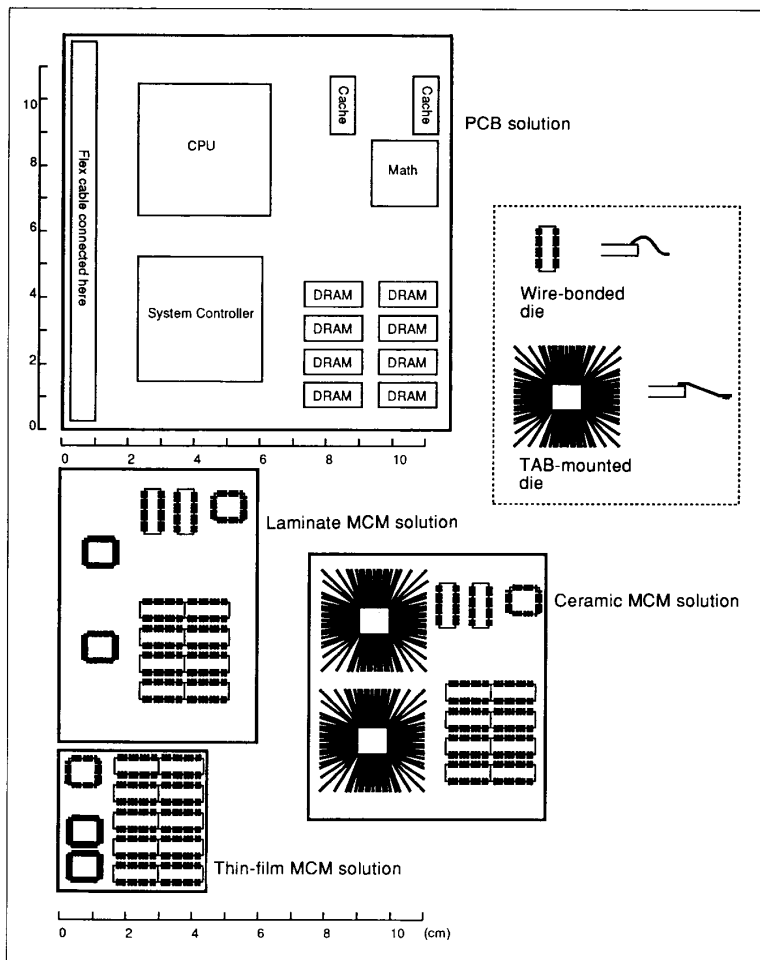
was calculated. The results appear in Table 3 on the next page. The power dissipated due to the interconnect is

about 2 percent of total dissipation. Table 3 also shows the estimated production cost (determined by combining vendor

**Alternative A: Single-chip PCB.** Figure 6 shows the layout for this alternative (and the alternatives that follow). The design requires a six-layer board: four signal layers, a power layer, and a ground layer. The signal lines are 150 micrometers wide (6 mils) and the vias are 300  $\mu\text{m}$  (12 mils) wide, placed on a 2.5-millimeter (100-mil or 0.1-inch) grid. These drilled through-hole vias pass through all layers of the board. A via must be used when the signal path changes layers or turns a corner (which requires a layer change). The large size of these vias and the wide spacing between them force the average wire pitch to be much wider than the minimum pitch. A routability analysis arrives at an average wire pitch of 1,000  $\mu\text{m}$  (40 mils).

The low interconnect capacity of this option results in a final daughter card size of 120  $\times$  120 mm. This is due mainly to two technology-related reasons. The first is the size of CPU and controller single-chip packages. Their leads are soldered to pads on the top of the board. If the lead pitch is less than 0.5 mm, solder bridges become likely. Considerable fan-out is required, making the single-chip packages much larger than the die size. For example, the 227-pin CPU package is approximately 40  $\times$  40 mm, while the actual CPU chip is about 13  $\times$  13 mm (520  $\times$  520 mils). The second reason for the large average wire pitch arises from the large size and spacing of the vias.

The packaging-related power consumption of the standard PCB option



**Figure 6. Layout for alternative design options.**

**Table 3. Summary of sizes, package-related energy consumption, and production costs for the packaging alternatives.**

Alternative	Size (mm)	Power per MIPS (W/MIPS)	Cost (\$/part)
A	120 × 120	0.027	\$536
B	120 × 110	0.026	\$736
C	82 × 64	0.015	\$785
D	62 × 60	0.015	\$950
E	44 × 45	0.014	\$1,060

**Table 4. Summary of the comparison ratings of each packaging alternative based on the performance and cost factors.**

Goal	Weight	A	B	C	D	E
Minimum size	10	3	3	8	8	10
Low production cost	10	10	9	8	5	3
Minimum weight	8	5	5	6	6	6
Low power	7	5	5	6	6	6
Minimum design cost and cycle time	6	10	8	8	5	5
Thermal dissipation	5	3	3	8	10	8
Total weighted ranking		280	258	338	300	290

quotes). On-board decoupling capacitors are required to control the SSN.<sup>3</sup>

**Alternative B: PCB with bare-mounted CPU core chips.** The first bare-chip option to be considered uses the same PCB layout, with the chips TAB-mounted on the board. This does not substantially reduce the size of the board because of the chip lead fan-out issue. TAB lead frames must be used to provide the fan-out. The outer lead pitch of the frames is the same as the wire pitch on the board. The resulting size reduction is minimal, producing a 120 × 110-mm PCB daughter card. Thermal requirements dictate the use of thermal slugs, preventing routing beneath the chips. Table 3 reports the package-related power dissipation and the production cost calculations for all alternatives. The power dissipation for alternative B is only slightly lower than in A, which is attributable to the reduced capacitance of the TAB frame leads. The requirements for decoupling capacitors remain unchanged due to the large inductance of the TAB leads. The production cost is higher due to the premium that chip manufacturers currently charge for fully tested bare die.<sup>11</sup>

**Alternative C: Laminate MCM with wire-bonded chips.** To achieve substantial size reductions, it is necessary to reduce both the minimum and average line and via pitches. The laminate MCM considered here uses 75- $\mu$ m (3 mil)-wide lines and spaces, giving a minimum pitch of 150  $\mu$ m (6 mils). As the average pad pitch on the chips is also 150  $\mu$ m, no fan-out is needed for pitch-matching purposes. If a via size of 150  $\mu$ m is used, the average wire pitch will be 225  $\mu$ m (9 mils), which is four times better than the standard PCB. (Note that 150- $\mu$ m vias are currently available only in prototype form.) Through-hole vias are used again, as in the PCB options, because they are cheaper to make than vias that connect only two layers. This option results in an 82 × 64-mm, six-layer daughter card, which is four times smaller than the standard PCB option.

Once again, thermal requirements dictate the use of thermal slugs beneath the CPU and controller, preventing routing beneath these chips. The lower power dissipation is attributable to the reduced capacitance of the wire-bonded chip leads and to the reduced distance between the CPU and the memories.

Less decoupling capacitance is required due to the reduced inductance of the wire-bond leads.

**Alternative D: Cofired ceramic MCM.**

The cofired ceramic options have one advantage over the PCB and laminate MCM options. The ceramic substrate is a sufficiently good thermal conductor so that thermal vias are not required. The signal via size and pitch is also the same as the wire width and pitch respectively, and vias can be made to connect only two layers. Thus the average wire pitch is much smaller than in the laminate. However, the cofired ceramic substrates here have a minimum line width and space of 125  $\mu$ m (5 mils) and a minimum pad pitch of 250  $\mu$ m (10 mils). Thus the chips have to be mounted in a TAB format to provide the fan-out needed for pitch-matching. The resulting seven layer (four signal layers) MCM is 62 × 60 mm.

The additional cost over the laminate option is due to the extra costs of the TAB lead frames and the ceramic substrate itself. It is necessary to use discrete miniature capacitors on this MCM to control SSN.

**Alternative E: Thin-film MCM.** In this alternative, the line width, spacing, via size, and via pitch all have a minimum specification of 25  $\mu$ m (1 mil). Thus, even with two signal layers and with thermal via arrays underneath the CPU and controller, interconnection density is not an issue and the chips can be spaced as close as assembly considerations permit. The final size of the thin-film MCM is limited only by the size of the bare die and the space required between the chips for the wire-bond lead attachments. If wire bonding is used, a size of 44 × 45 mm is achieved; if solder-bump flip-chip lead attachments are used, the size is 41 × 44 mm. However, we reject the last option due to the lack of reasonably priced solder-bumped die at the time of writing. (This is an infrastructure issue rather than a production cost issue. If chip manufacturers provided the solder-bumped die, the cost might actually be less than for the TAB or wire-bonded alternatives.) This thin-film alternative requires five layers, consisting of two signal layers, one surface pad layer, one power layer, and one ground layer.

Although the thin-film MCMs could be mounted in a standard ceramic PGA



package, this is expensive. An alternative would be to mount the MCM in a quad plastic flat-pack-style package, with a special heat spreader and heat sink attached. Although we did not evaluate this solution thoroughly, it might be difficult to dissipate the heat in a small volume with this solution. The lower power dissipation value can be attributed mainly to the reduced distance between the CPU and the memories. If the power and ground layers are closely spaced, they effectively form a distributed decoupling capacitor, negating the need for an on-MCM mounted discrete decoupling capacitor for this design.

## Making the final decisions

The fifth and final step of the decision process is to select the best packaging from among the analyzed alternatives. The analyses of the performance and cost factors discussed are (qualitatively) converted into numeric ratings for each alternative and listed in Table 4. The largest differentiation between the alternatives comes from the size and area factors, where the advantages of the MCMs over the standard PCB can be clearly seen. Using a thin-film MCM results in the smallest packaging, but even the laminate MCM produces a 75 percent savings over the PCB in the daughter card area. Note that bare-mounting the die on a PCB provides only a small area savings. This technology leads to a reduced size only when low I/O chips are used. The weight differences between the alternatives is a very small fraction of the total system weight.

Production costs differ significantly between the alternatives. The PCB is by far the cheapest to manufacture; however, the laminate MCM costs only slightly more than the PCB, mostly due to the premium price paid for the qualified bare die. As the infrastructure moves toward cheaper availability of pretested bare die, this price difference is expected to decrease. The thin-film option is the most expensive.

The overall ratings in Table 4 show that the six-layer laminate MCM with wire-bonded chips is the best choice for our stated goals and requirements. The strong differentiation among the alternatives in the size and production costs

factors most influenced the final decision. The thin-film option provides a slightly smaller daughter card, but the production costs are significantly higher.

**N**otebook computers are becoming faster and smaller, increasing the impact of electronic packaging on the performance and cost of the overall system. MCMs offer improvements over conventional, single-chip packages in several factors that affect the performance and cost of the system, especially in size reduction. ■

## Acknowledgments

We thank the reviewers for their helpful comments and acknowledge the National Science Foundation's support for part of this work under Grants MIP-901704 and DDM-9225755.

## References

1. "Flash Memory Card Redefines Portable Computing." Order No. 296839-001. Intel Literature Sales. Mt. Prospect, Ill., 1990.
2. "386SL Mustang Design Kit." Order No. Mustangdoc. Intel Corp., Santa Clara, Calif., 1990.
3. "386SL Microprocessor SuperSet. System Design Guide." Order No. 240816-001. ISBN 1-55512-130-6. Intel Literature Sales. Mt. Prospect, Ill., 1990.
4. D.A. Doane and P.D. Franzon, *Multichip Module Technologies and Alternatives: The Basics*. Van Nostrand Reinhold, New York, 1992.
5. C.H. Kepner and B.B. Tregoe, "The New Rational Manager." Princeton Research Press, Princeton, N.J., 1981.
6. W.R. Heller and W.F. Mikhail, "Package Wiring and Terminals." *Microelectronics Packaging Handbook*. R.R. Tummala and E.J. Rymaszewski, eds., Van Nostrand Reinhold, New York, 1989.
7. H.B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*. Addison-Wesley, Reading, Mass., 1990.
8. G. Messner et al., *Thin-Film Multichip Modules*. Int'l Soc. Hybrid Microelectronics, Reston, Va., 1992.
9. R.W. Johnson, R.K.F. Teng, and J.W. Balde, *Multichip Modules, System Advantages, Major Constructions, and Materials Technologies*. IEEE Press, Piscataway, N.J., 1991.
10. E.J. Vardaman and L.H. Ng, "A Cost-Performance Analysis of Multichip Module Interconnects." *Proc. 1991 Int'l Symp.*

*Microelectronics*, Int'l Soc. Hybrid Microelectronics, Reston, Va., 1991.

11. T. Costlow, "Intel Rolls Dice for MCM Makers." *Electronic Eng. Times*, Issue 691, May 4, 1992.



**Paul D. Franzon** is an assistant professor in the Department of Electrical and Computer Engineering at North Carolina State University. He has over eight years experience in electronic systems design and design methodology research and development. His research interests include design sciences/methodology for high-speed packaging and interconnect, high-speed and low-power chip design, and the application of Micro Electro Mechanical Machines to electronic systems.

Franzon received a BSc degree in physics and mathematics, and BE and PhD degrees in electrical engineering, all from the University of Adelaide, Adelaide, Australia. He is a member of the IEEE, ACM, and the International Society for Hybrid Microelectronics. He chairs the Education Committee of the IEEE Components, Hybrids, and Manufacturing Technology Society.



**Robert J. Evans** is an engineer with the portable products group at the IBM Personal Computer Company, Research Triangle Park, North Carolina. He is pursuing a PhD degree at North Carolina State University in the area of low-energy computing.

He received the BE degree in electrical and biomedical engineering from Vanderbilt University, Nashville, Tennessee, in 1980, and the ME degree in biomedical engineering from the University of Virginia at Charlottesville in 1982. He is a member of the IEEE and the National Society of Professional Engineers.

Readers can contact Paul Franzon at Electrical and Computer Engineering, North Carolina State University, Box 7911, Raleigh, NC 27695-7911; e-mail paulf@ncsu.edu; and Robert Evans at the IBM Personal Computer Company, Dept. 304, Bldg. 203, PO Box 12195, Research Triangle Park, NC 27709; e-mail rjevans@ralvm8.vnet.ibm.com.