Alternative Packaging Gains Ground

MCMs, BGAs Become Widespread for CPUs; Micro-BGA on Deck

by Linley Gwennap

Two years ago *(see **071304.PDF**)*, we first looked at emerging packaging technologies such as multichip modules (MCMs), tape automated bonding (TAB), and ball-grid arrays (BGAs). As we predicted, many microprocessors have since adopted these packages instead of the traditional PQFP and PGA packages. The highest-volume adopter has been Intel, which chose a two-chip MCM for its P6 processor and a TAB package for its mobile Pentiums.

Multichip modules have gained the most processor design wins and show promise for the future. TAB is shipping in the highest volume, thanks to Intel, but may decline in the future. New packages, such as Tessera's micro-BGA, may displace TAB. Larger BGAs are becoming prevalent in high-end microprocessors due to their ability to efficiently handle higher pin counts.

MCMs Stride Forward

Two years ago, the biggest obstacles to high-volume MCM production were the lack of known-good die (KGD) and high cost. Since then, vendors have made significant headway in both areas.

Intel, clearly on top of the latest packaging trends, has led the way in solving the KGD problem. The company's SmartDie program *(see **0809MSB.PDF**)* provides KGD at about the same price as packaged parts. Other chip vendors have followed suit; more than 70, including several microprocessor vendors, deliver tested bare die. These products reduce the possibility that a single bad die will render an entire MCM unusable.

One of the leaders in the MCM field is MicroModule Systems. MicroModule wisely realized that its MCM business required widespread KGD to grow, and it developed a test fixture that uses a pressure attachment to test die without bonding. This technique, along with others, has helped foster the KGD boom.

MicroModule builds the HyperSparc MCM for Ross *(see **0806MSB.PDF**)*. It also builds and markets Pentium modules that contain a complete CPU/cache subsystem. Hal's Sparc64 processor, built by Fujitsu, and IBM's Power2 are other processors manufactured today in multichip modules.

MicroModule is building MCMs at relatively low cost. The company uses silicon or aluminum substrates with large (10-micron) features and low density compared with CPUs; thus, they yield well despite their large size. MicroModule also has a competitive advantage in its substrate fab and assembly facility, acquired at a fire-sale price from Digital, which built VAX 9000 CPUs there. The fab was originally built by the ill-fated Trilogy.

Can MCMs Change Processor Design?

Designers of high-end processors are continually restricted by the maximum number of transistors that can fit onto a single die. Most designs result in a huge die that is expensive and barely manufacturable; for example, all five next-generation RISC processors, as well as Pentium Pro (the P6), require about 300 mm^2 or more of die area in advanced 0.5-micron processes. Before tape out, vendors frequently must trim their designs after discovering that they are too large to be built at all. Even when reduced to a buildable size, these large chips are very costly, because manufacturing cost increases with the square of the die size.

The same processor could be designed more quickly if spread across multiple chips, since none of the individual die would be close to the manufacturing limit. The cost of building multiple die is less than the cost of an equivalent area packed into a single die, as the yields of the smaller chips will be higher. In theory, breaking up a processor into smaller chips could lead to lower cost and faster time to market.

Alternatively, the use of multiple chips could allow designers to break Moore's Law, as Intel has done with Pentium Pro *(see **0906VP.PDF**)*. With its two-chip design, the Pro includes 21 million transistors, far more than can be packed into a single-chip processor today. Taking this concept further, the CPU itself could be divided among multiple chips. A multichip CPU could contain more features, more memory, and more processing power than is possible for a single-chip CPU.

Most vendors have not followed this path for several reasons. KGD and cost are big issues, but MicroModule and others have made progress in these areas, as noted above. Another crucial problem is that multichip designs traditionally operate at lower clock speeds than single-chip processors, because of the extra time required to drive signals from one die to another.

As chips use smaller and smaller metal traces, however, this advantage is starting to flip around. The thin metal traces used in 0.35-micron IC processes have a fairly high resistance, particularly when signals must be routed from one end of a large die to the other. An MCM allows a signal to hop off the die and run to another die across a wide, low-resistance metal trace. In some cases, the signal can switch faster than if it went through a

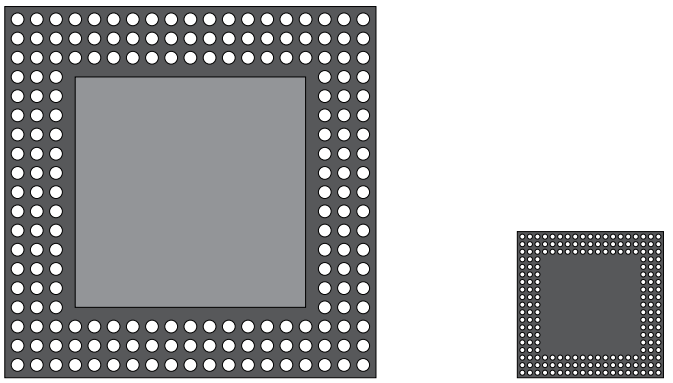

Figure 1. At left is a typical perimeter BGA, in this case a 192-pin device that, with a 1.27-mm lead pitch, measures 25 mm on a side. At right is a 192-pin micro-BGA with a lead pitch of 0.5 mm. It measures less than 10 mm on a side.

thin metal trace across a large monolithic IC. As process features continue to shrink, the advantage could swing solidly to the MCM method.

Many processor vendors are considering a multichip CPU for next-generation projects, but none has publicly committed to such a design. The advantages of potentially better performance, lower cost, and reduced design time are compelling. The risk is in committing to an unproven technique; if a problem arises, engineers might have to put Humpty Dumpty back together again, probably in a more advanced IC process. But the risks appear manageable, and a vendor willing to take the plunge could gain a significant competitive advantage.

TAB Packaging a Hit in Notebooks

TAB's future is less clear. Two microprocessors, HyperSparc and Sun's MicroSparc, have shipped in TAB packages, but their vendors both rejected TAB for subsequent versions of these products.

Although the SPARC vendors have been unimpressed by TAB, Intel has taken to it wholeheartedly—and so have some of its customers. Mobile versions of Intel's Pentium processor *(see **091301.PDF**)* are available in both PGA and TAB packages; Intel calls the latter a tape carrier package (TCP). The company says it is shipping about half of its Mobile Pentiums in the TAB package; this volume is concentrated among several large customers. TAB parts require an expensive new machine to be added to a PCB assembly line, so smaller notebook vendors still prefer the PGA package.

This assembly expense negates TAB's cost advantage. TAB proponents claim that the package reduces cost compared with a ceramic PGA, which is much more expensive than most plastic packages. But Intel says that it does not see a cost advantage in TAB. Sun claims that the TAB package used for MicroSparc provided a small cost benefit at the chip level, but extra cost at the board level erased any net cost advantage.

For portable systems, TAB delivers the weight advantage of a very thin plastic package while providing better cooling and greater lead density than a PQFP or PBGA. For desktops, however, there is little benefit, and no desktop microprocessor is currently using TAB.

BGAs Move to Perimeter Leads

In the microprocessor world, Motorola and IBM have been the prime movers behind BGA packaging. Motorola invented the BGA *(see **071203.PDF**)* and is now using it for the PowerPC 603, 604, and 620. IBM is also sourcing PowerPC chips in BGA packages and uses BGAs in some of its RS/6000 workstations for its Power2 processor. All of these products use ceramic BGAs because their power dissipation is too high for plastic packages.

Typical BGAs use a lead pitch of either 1.5 mm or 1.27 mm (50 mils). A 525-lead package, which is roughly the lead count used by most next-generation CPUs, requires a 25×25 array of leads, resulting in a package size of about 32 mm on a side. Unless the package is placed on a PC board with many routing layers, however, it is nearly impossible to route signals to the inner leads.

To solve this problem, vendors have moved to a perimeter array design, as Figure 1 shows. In this design, leads are restricted to the outer three or four rows of the array, depending on the number of routing layers available. This design makes it much easier to route signals to the BGA. Because of this, perimeter BGAs are rapidly replacing standard BGAs for all but the smallest lead counts.

A variation on this design is the so-called Super-BGA from Amkor. This package places the die face down, so a copper lid can be mounted directly to the back side, avoiding the high cost of ceramics. An internal heat coupling shifts out heat through the perimeter solder-ball array as well as through the top of the package. Thermal performance can be further extended by using an add-on heat sink. Internal power and ground planes improve electrical performance while providing EMI shielding. Several x86 processor vendors are considering this package for future products.

Tessera Shrinks Ball-Grid Array

One downside to perimeter BGAs is the increased package size. With only four rows of leads, the 525-lead package from the above example now requires a 37×37 array of leads on a 47-mm square package. Even a simple 200-lead device needs a 17×17 array in a 22-mm package. Tessera has developed its micro-BGA (µBGA) package to reduce these sizes. The µBGA uses fine pitches of 0.5–0.7 mm. With a 0.5-mm lead pitch, for example, the above packages lose nearly 85% of their area! In many cases, the limiting factor for package size becomes the size of the die.

For More Information

Contact John Novitsky of MicroModule Systems (Cupertino, Calif.) at 408.864.6182; fax 408.864.5950 or send e-mail to *novitsky@mms.com*.

Contact Amkor (Chandler, Ariz.) at 602.851.5000; fax 602.821.2389.

Contact Tessera (San Jose, Calif.) at 408.894.0700; fax 408.894.0768.

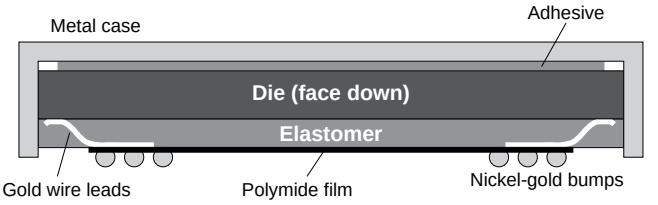

Figure 2. Tessera's micro-BGA is barely larger than the die itself. The elastomer provides a cushion between the die and the PC board, while the gold bond wires flex enough to handle different thermal expansion rates between the die and the board.

The µBGA is only slightly larger than the die itself, as Figure 2 shows, yet the package has several advantages over simply mounting a bare die onto a PC board, a technique called chip-on-board (COB). In the µBGA, an elastomer layer, 6–7 mils thick, sits between the die and the board, providing a mechanical interface that cushions the die during insertions. Signals are routed from the die using thin gold bond wires that flex along with the elastomer. As the figure shows, this design works best for a die with a perimeter pad ring.

This mechanism also protects the die from shear forces caused by thermal expansion. Because silicon and FR4 (the material used in most PC boards) expand at different rates when heated, a bare die soldered to a board can crack and literally tear apart when the temperature changes too much. The bond wires in the µBGA flex to avoid this problem.

The optional metal case further protects the die during handling and draws heat from the die during operation. A heat sink can be mounted on the µBGA when greater power dissipation is required.

The underside of the µBGA is a thin polymide film that routes signals from the bond wires to the solder balls that connect to the board. Because the signals travel only a few millimeters from the die to the board, lead inductance is minimized, allowing faster signals. The small size of the package also reduces cost: Tessera believes that a 208-pin µBGA package will cost about $3, or less than 1.5 cents per pin, in 1997. PQFP and other low-cost packages run about 2 cents per pin for similar pin counts.

Hitachi, Others License µBGA

To date, Tessera's package has not reached production. It has been licensed by Hitachi and Shinko Electronics (Nagano, Japan), both of which plan to sample products in early 1996. Hitachi will offer the package as an option for its gate-array line; Shinko is a packaging house for other vendors' chips. In addition, packaging leader Amkor has made an equity investment in Tessera and will provide packaging services to that company and possibly others. Other companies have developed packages similar to Tessera's, including Motorola's Slicc and Sandia's mini-BGA, but none is close to production.

One drawback to the µBGA is that the fine lead pitches are difficult for many board manufacturers to handle, particularly the smaller board makers. To reduce cost, some vendors may put the µBGA parts on a small daughterboard, limiting the amount of board space produced using the higher-cost fine-pitch process. The µBGA is also unproved in volume production; it remains to be seen if the part can meet its goals, particularly on cost, and what other problems might crop up.

MCMs Have Many Possibilities

Over the coming years, the packaging landscape will continue to change as processor vendors scramble to deal with increasing pin counts, higher frequencies, and rising costs by investigating new packaging alternatives. We expect greater use of MCM packaging for high-performance microprocessors. For low-cost desktop systems, today's PGA and PQFP packaging seems adequate, with some PGAs being replaced by ceramic BGAs or packages like Amkor's SuperBGA.

The competition will be more intense in the portable area, where size and cost are very important considerations. If it can prove itself in volume production, the µBGA may challenge TAB, as it has a smaller footprint and is just as light and thin. By using a metal enclosure, the µBGA can dissipate more heat than a typical plastic package, matching the thermal characteristics of TAB. Finally, if Tessera's projections are accurate, the µBGA could deliver a significant cost advantage over TAB.

Like TAB, the µBGA requires a more expensive board-level manufacturing process, which could delay its adoption. Instead of placing several µBGA parts on a small fine-pitch daughter card, the same components could be combined in an MCM, further reducing footprint and improving cost. The MCM could use a PGA or QFP enclosure, making it easy to mount on the board and rework, if necessary, using standard tools.

At least one notebook maker—Matsushita's Panasonic division—is using an MCM that combines a 486 CPU with an L2 cache. We expect that more notebook vendors will adopt MCMs in the next year, while TAB adoption levels off. Other packaging alternatives, such as the µBGA and chip-on-board, will also be tried as notebook vendors seek to reduce system size and cost. ♦

