Processors Put Pressure on Packages

A Review of the Packages Housing Modern PC Processors

by Dennis Herrell

With no end in sight to increases in frequency, microprocessor vendors are being forced to develop packages that handle higher power, have better signal integrity, and fit into a smaller space. Making things worse, the precipitous decline in the average selling prices of PCs requires that packages become less costly and have higher assembly yields.

Today, there are two basic types of packages used for PC processors, the pin-grid array (PGA) and the ball-grid array (BGA), as Figure 1 shows. Various material combinations and configurations are used to realize these packages. For manufacturing and marketing reasons, it is often desirable for packages to be removable (pluggable). This requirement has traditionally been met using a PGA, although pluggable card modules with BGA-mounted processors are also used.

Delivering power to the processor is usually the first consideration in package design. Once a solid power-delivery system is in place, signal fidelity presents only a small problem. Other elements, such as cooling and thermomechanical unreliability, present more difficult problems to overcome.

Garbage Power In, Garbage Signals Out

The highest-performance PC processors in desktop systems are approaching a power dissipation of 50 watts. Power dissipation (W) is related to the operating frequency (f), the power-supply voltage (V), and the chip capacitance (C) by the formula $W = C V^2 f$. For PC processors, operating frequency is rapidly approaching 1 GHz, and supply voltage is falling toward 1.5 V. The capacitance being charged or discharged in a clock cycle is typically on the order of 20 nanofarads. To lower power dissipation and improve device reliability, the power-supply voltage has been reduced with each succeeding semiconductor generation. (Lowering the voltage keeps electric fields in the ever-thinner gates from compromising the insulator, and causing damage such as threshold shifts from hot-carrier injection.) Even at lower voltages, however, power dissipation is increasing, partly as a consequence of larger die and thinner insulators (more capacitance), but also from dramatic increases in operating frequency.

High power at low voltage requires that high currents be delivered to the processor. A 50-W processor requires a supply current of 33 amps at 1.5 V. As Figure 2 shows, this current must flow from the power-regulator circuits, through the motherboard, through connections on the package to the processor die and return through a similar path.

Every link in this chain must have low resistance, else unacceptable voltage drops will occur along the way. For correct operation of the processor's circuits, no more than 5–10% voltage drop can be tolerated. To keep resistance low, high-conductivity metals and multiple parallel paths for current are used. This requires that most of the connections to the chip and the package be reserved for power and ground.

Over the past 10 years, the most commonly used package has been the ceramic pin-grid array. These CPGAs have used either tungsten or molybdenum as the conductors. These refractory metals are now giving way to copper, which has between two and five times lower sheet resistance (ohms per square). Copper becomes an option when the package substrate is changed from cofired alumina-ceramic to printed-circuit board (PCB) construction using organic laminates. Intel is leading this charge, with much of its production now switching to OLGAs (organic land-grid arrays). Others are being more conservative: AMD, for example, is still using ceramic BGAs for its newest processor, Athlon.

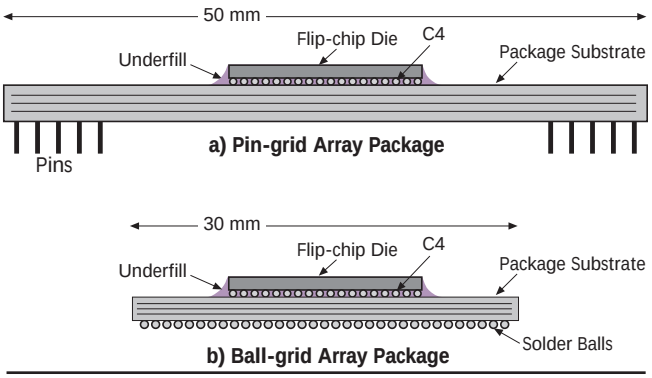

Figure 1. The two most common PC-processor package types are the pin-grid array (a) and the ball-grid array (b). Most processors are being mounted to the package substrate with C4 bonding.

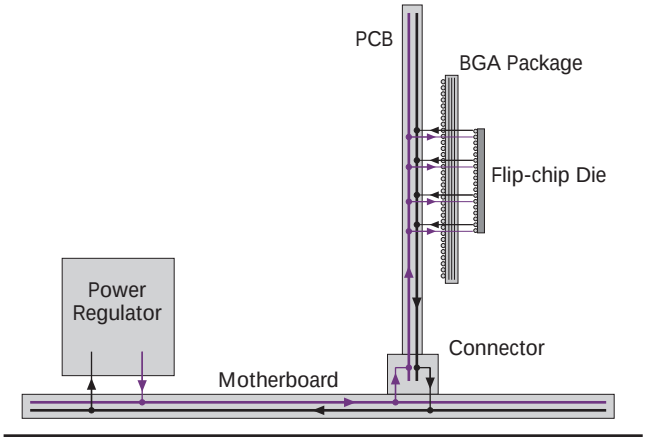

Figure 2. To present a stable power and ground supply to the processor circuits, low resistance and inductance must be maintained all the way from the power regulator to the processor die.

Electromigration is another problem being closely modeled by chip, package, and system designers. Electromigration is the movement of metal atoms in the conductors as a consequence of high currents; over time, such movement can lead to circuit defects. One of the main advantages of copper interconnect over aluminum is the approximately 10× higher current that can be carried by small-cross-section conductors without electromigration failures.

Supplying current is only a small part of the power delivery problem. An even bigger problem arises from transient currents. As the PC goes through its gyrations—heavy demand, light demand, sleep, etc.—current can vary rapidly, from a few milliamps to tens of amps within a few nanoseconds. The current surge between the microprocessor and power-regulator circuits creates spikes in the supply voltage. If this occurs, a back-EMF (dv) is developed across the inductance (L) of the power-delivery connections in direct proportion to the rate of change of the current ($dv = L\, di/dt$).

With just 1 nanohenry of inductance—corresponding to approximately 1 mm of wire and a current slew of 1 A per nanosecond—dv can exceed 1 V. Clearly power-distribution systems must be designed to have inductances measured in picohenries, not nanohenries.

Unfortunately, it is not possible to design regulator circuits to squelch these resistive and inductive voltage spikes. The regulator circuits are too slow and too far away from the processor to respond to nanosecond variations at the microprocessor. Regulators typically operate by chopping the incoming supply voltage at approximately 1 MHz, and the round-trip delay from sensing the on-die voltage to delivering that information to the regulator is several nanoseconds.

As a consequence, the primary way to manage surge currents and reduce voltage fluctuations is to use decoupling capacitors distributed throughout the power-delivery system from motherboard to die. Various capacitors are normally implemented on the processor die and built onto the package and PCB. But even the best capacitor layout cannot completely solve this problem.

For a many years, processor packages came with discrete ceramic capacitors mounted on the surface of the package. These were typically small surface-mount devices with capacitance values ranging from a few to a few hundred nanofarads. These capacitors provided low impedance for frequencies up to 50 or 100 MHz, but they were not effective at frequencies of hundreds of megahertz. It took a long time for this fact to be realized, and many of today's several-hundred megahertz microprocessor packages still retain the vestigial pads for mounting on-package capacitors.

For microprocessors operating at more than 200 MHz, the only serviceable capacitor is an on-die capacitor, or one that is very close. Digital was perhaps the first to recognize this need when it tried placing a discrete capacitor directly over the Alpha chip. This approach worked, but Digital soon switched to capacitors integrated directly in the silicon. All PC-processor manufacturers now use on-die capacitors.

To be effective, on-die capacitors must be at least 10 times as large as the chip's total switching capacitance, or typically about 200 nF. On-die capacitors can be built from wells and from arrays of unswitched devices, but they are often augmented by gate-oxide capacitors placed in unused silicon. Even that is not sufficient in some cases, and some designers add silicon area just for these capacitors.

Good Signals Require Short Connections

With a solid power-delivery system in place, ensuring high signal fidelity into and out of the microprocessor is usually not a huge challenge. Normally, only a few hundred connections carry signals, and package traces are usually short enough that small differences in characteristic impedance (Z) between the package and the PCB can be tolerated. As a consequence, the main design criterion for signals is to minimize the number of conductor layers to keep costs down. A good rule of thumb is that each layer (one conductor plus one insulator) adds about $1 to the cost of the package.

Maximum signal bandwidth deteriorates with inductive (L) discontinuities, which cause reflections and crosstalk; to a first-order term, bandwidth scales as the ratio of L/Z. Pin or solder-ball inductances are relatively small, typically 5 nH and 1 nH, respectively, so these are not usually a limiting factor (e.g., 5 nH in a 50-ohm system limits rise times to no faster than 100 ps). Inductive and capacitive crosstalk between adjacent signals is managed by placing power and ground connections in close proximity to the signal paths. Very high bandwidth signals, such as Rambus and AMD's Lightning Data Transfer (LDT), demand stricter impedance control and more careful signal-line placement.

Getting the Heat Out

Semiconductor manufacturers typically recommend restricting junction temperatures to below 100° C; higher temperatures can impair reliability. To keep junction temperatures below this level, heat from processor circuits must be dissipated into the surrounding air. This is usually accomplished by moving air across a heat sink attached to the processor.

There are very few design methods available for cooling high-power microprocessor junctions below 100° C, especially in PCs, where ambient air temperature can be as high as 40–50° C. State-of-the-art heat sinks (which would more correctly be called heat exchangers) have large, massive structures with narrow air channels, often directly attached to, or integrated with, a fan. Future design advances that improve laminar air flow should allow heat to be removed at a rate of 5 kW per liter at reasonably low noise levels.

Perhaps the biggest cooling challenge is the connection of the heat sink to the processor with sufficiently low thermal resistance. With most, if not all, flip-chip packages, the only effective way to remove heat is through the back of the die, as Figure 3 shows. Because silicon is a good conductor of heat, only a small thermal gradient exists from the circuit side to the back of the die. Effectively coupling a heat sink to the

back of a die, however, is not easy; the difficulty arises in creating a reliable yet compliant interface. A compliant layer is typically created by a thermally conductive layer of grease on the back of the die. The grease is formulated with conductive particles to transfer heat effectively.

Early flip-chip packages contained the grease beneath a lid over the microprocessor; the heat sink was pressed against the lid. Today, vendors are turning to lidless packages to reduce the number of thermal barriers between the die and heat sink. Lidless packages have a downside, however: the critical task of connecting the die to the heat sink now lies with the OEM, or even the consumer, who has less control for making good thermal contact without contamination.

Processors Can Die of Fatigue Too

Thermomechanical reliability adds a further constraint to package design. Large thermal gradients can develop throughout the package system, and there is a wide range of thermal-expansion coefficients associated with the materials involved. Silicon, for example, has a coefficient of about 3 parts per million per degree C, while PCB materials have over 20 ppm/°C. Operating temperatures can vary from room temperature to nearly 100° C. The reliability of flip-chip solder joints exposed to repeated heating and cooling cycles was among the early problems faced by the industry.

To avoid solder-joint failure from stress and fatigue, IBM introduced the idea of filling the space between the die and the package substrate with an organic underfill to take the brunt of the mechanical strain between the die and package, as Figure 1 shows. This has been a highly successful solution for flip-chip onto ceramic substrates, which have a coefficient of expansion of less than 10 ppm/°C. Only recently have underfills been developed that are capable of handling the larger disparity of flip-chip on PCB substrates.

CBGA packages soldered to a PCB must also withstand the thermal cycling strain across their solder balls. This requirement limits the dimensions of CBGAs to approximately 30 × 30 mm to keep the relative movement between the center and edge solder-ball connections to within tolerable limits. To relieve stress on these solder joints, various techniques have been developed to extend the height of the solder connections, making them more compliant.

Last, package flex in response to thermomechanical stresses can affect the cooling performance of the package system. Because the thermal interface between the back of the die and the heat sink is constantly flexing, as Figure 4 shows, thermal grease can be partially pumped away, increasing the thermal resistance between the die and the heat sink.

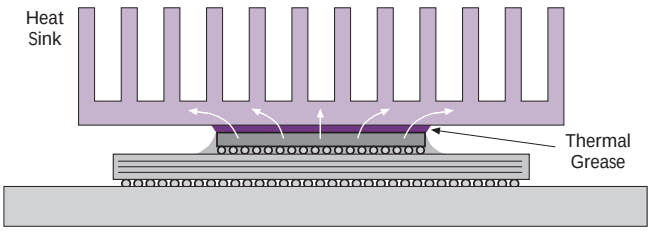

Figure 3. A heat-conductive thermal grease layer is often used to improve heat transfer from the backside of the chip to the heat sink, which dissipates the heat created by the processor into the air.

C4 Revolutionizes Chip Connections

For connecting processor chips to package substrates, the major advance in recent years has been area-array solder joints. Lower supply voltages (e.g., 1.5 V) and higher currents (exceeding 20 A) make it increasingly difficult when using peripheral connections to avoid voltage sag at the center of the die. At the same time, it is difficult to avoid electromigration with power supplied from the edge of the die only.

To avoid voltage sag, area connections are required. Most processor manufacturers are now using variants of the C4 (controlled collapse chip connect) technology developed by IBM for mainframes. C4 is a method of connecting across the entire face of the die with small solder bumps, which are on the order of 75 microns in diameter and placed on a pitch as small as 225 microns. The microprocessor is then flip-chip attached to the package by reflowing the solder connections.

With flip-chip technology and two thick layers of metal on the die, a very robust power-delivery system can be realized. Of the many thousands of flip-chip connections that can be made on a processor die measuring 100–200 mm^2, only a few hundred are required for signals. This leaves most of the C4 connections (1,000 to 2,000) for power and grounds, allowing very low inductance and resistance connections.

A Similar Solution for Mounting the Package

The body of a package acts as a space transformer by bridging from the tiny chip connections to the much larger PCB connections, which are typically on a pitch of 50 to 100 mils (1.27 to 2.54 mm). Although smaller pitches are possible, cost establishes the lower limit. Typical PCB motherboards have just four layers of metal, usually configured as signal-power-ground-signal. The requirements for routing signals from under the package to other components establishes the PCB-connection pitch. Usually, this pitch is no finer than a 5-mil trace (125 microns) plus a 5-mil space.

Plated through-hole vias, which connect the various metal levels in the PCB, are normally on a minimum pitch of

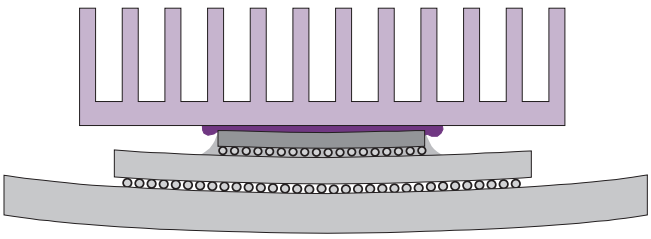

Figure 4. Variations in thermal coefficients cause flexing of the components in response to temperature. Over time, this flexing can pump thermal grease out of the processor-die-to-heat-sink interface, degrading heat transfer. This can cause junction temperatures to rise, possibly leading to catastrophic failure of the part.

50 mils (1.27 mm). Note that the package is also the interface between the metric system of measurement used on chips and the English system still used in the PCB world.

PGAs are the most popular package type used in PCs today. PGAs typically contain just the microprocessor and are up to 50 mm square. PGAs usually have a peripheral array of pins (in 4 or 5 rows), allowing the device to be plugged into a socket. Pluggability is an important feature, because it allows the PC manufacturers to insert the microprocessor just prior to final shipment, which has positive logistic and financial ramifications. It also facilitates customer upgrades.

The second-most-popular form factor is the SECC2 (single-edge contact cartridge) processor card, popularized by Intel's Pentium II. The card has the advantage of allowing other high-speed components, such as L2-cache SRAMs, to be mounted close to the processor. The microprocessor is most often soldered to the card in a BGA package. The SRAMs are usually packaged in either BGAs or quad flat pacs (QFPs).

The SECC2 connects to the motherboard through an edge connector containing four rows of contacts. The card itself is an enhanced PCB, typically containing up to 10 metal layers for good routability, impedance control, and power delivery. The card is also a suitable format for multiprocessor configurations, allowing cards to be placed side by side on the motherboard. Unfortunately, cards cannot be placed very close together because of the space needed for heat sinks.

Although the card provides a fast connection to the L2, it is an expensive approach, and it is on its way out as manufacturers move to 0.18-micron processes and integrate the L2 onto the processor die. The card is expensive for a number of reasons. First, the SRAMs become a part of the microprocessor and must be acquired, stocked, and assembled onto the card. Second, the enhanced PCB substrate is expensive, and a large number of discrete resistors and capacitors must be assembled onto the card for signal termination and power regulation. Last, the card requires a connector with demanding geometrical and electrical specifications, an expensive option compared with a simple PGA socket.

A pluggable version of the BGA package is starting to appear in the form of a land-grid array. An LGA is essentially a BGA package with the solder balls replaced by a compressible connection that is part of a socket.

The BGA has much better arrangement of connections than the PGA, as Figure 5 shows. With a full array of connections, the centrally located pads can be devoted to power and ground. But the power and ground delivery capability of the PGA is being constantly improved, with more metal layers to reduce lateral resistance and with conductive copper metal instead of tungsten or molybdenum.

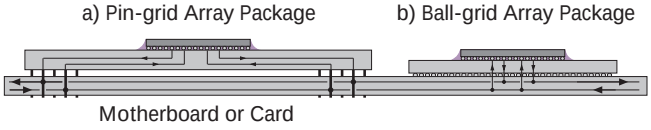

Figure 5. BGA packages (b) provide a more-direct and lower-impedance connection to the power and ground planes on the motherboard than do PGA packages (a). The result is more stable power delivery and better signal integrity for high-frequency signals.

Adding layers or changing from ceramic to organic packages costs more at present. A good intermediate solution is to add centrally located pins to the PGA, so power and ground connections can be made more directly to the core region of the processor die. In this way, the power and ground system of the PGA will approach that of the BGA.

Primarily for cost, but also for electrical and mechanical reasons, packages are made as small as possible. The minimum size is set by the connection constraints of the PCB and by the afore-mentioned thermal constraints. Since BGAs are smaller than PGAs with the same number of connections, BGAs have a cost advantage. In addition to size, package cost is proportional to the number of metal layers used. Consequently, designers attempt to minimize the number of layers.

Only Modest Package Evolution Seen

A well-honed packaging approach has evolved over the past 10 years in direct response to the demands of microprocessors and PC systems. The package types in highest-volume use today are the PGA and BGA, constructed either as a co-fired ceramic or, more recently, as an advanced organic PCB.

One motivation for organic packages is the anticipation of lower cost. Until fine-line PCB technology is fully developed, however, this goal will remain illusory. So far, all PCB packages have done is drive the price of ceramic packages down in response to the threat. But the advantages are compelling, and within the next few years, organic packages are likely to succeed in their quest to supplant ceramics.

Power-delivery challenges have resulted in more and more PC-processor manufacturers adopting flip-chip packages. Full-array connections between the die and the package, as well as between the package and the board, are a clearly superior method for meeting the low-impedance requirements for the power delivery and ground return.

Packages are likely to continue to hover around the 30 to 50 mm size, being driven primarily by the need to accommodate the required signal and power/ground connections. Pitches smaller than 1.27 mm for solder balls or 2.54 mm for pins will be exploited only when the cost of motherboards with more than four layers comes down.

The processor-card format might hang on for some time in multiprocessor systems, but, since it is more costly than PGA by at least a factor of two, it is likely that most PC processors will revert to PGAs once the L2 is integrated on the processor die. The PGA will probably be extended with the modification of additional pins (or pads) for central power and ground connections. M

Dennis Herrell recently retired from his position as an AMD Fellow, responsible for packaging R&D. Prior to AMD, he was VP of the packaging program at MCC and worked on silicon and superconducting technology development at IBM in Yorktown Heights. He holds a Ph.D. in physics from Cambridge University.