The Evolution of Microprocessor Packaging

Ravi Mahajan, Assembly Technology Development, Intel Corp.
Ken Brown, Assembly Technology Development, Intel Corp.
Vasu Atluri, Assembly Technology Development, Intel Corp.

Index words: packaging, technology, building blocks, interconnect, power delivery/removal

ABSTRACT

Microprocessor packaging is undergoing major changes driven by technical, business, and economic factors. From the traditional role of a protective mechanical enclosure, the modern microprocessor package has been transformed into a sophisticated thermal and electrical management platform. Furthermore, microprocessor architecture and design techniques can have significant impact on the complexity and cost of packaging. The need to optimize the total solution (chip, package, board, and assembly) has never been more important to maximize microprocessor performance and minimize cost. It is important to point out that the package represents a way of connecting the microprocessor to the motherboard. In this capacity, it enables the fine feature, silicon-level interconnects to be connected to the motherboard, i.e., the package assists in a space transformation in a controlled and economically viable manner. The key to packaging is to ensure that it enables and optimizes microprocessor performance. In its early evolution, the influence of the package on microprocessor performance was limited; however, as the microprocessor evolves to provide increasing performance, the package must evolve to keep up, and packaging design must ensure that it optimally enables the microprocessor. Package performance, in this context, implies a clear understanding and optimization of the package's electrical, thermal, and mechanical characteristics to enable overall electrical performance and power dissipation and to ensure mechanical robustness. Recent advances in microprocessor packaging indicate a migration from wirebond (where the chip or die is interconnected to the package only on the periphery of the die) to flip-chip (where the die is interconnected to the package using the entire die area); and from ceramic to organic packages, with cartridge and multichip technologies emerging as key form-factors. With the emergence of the segmented market (mobile, desktop, server and associated sub-segments), we see a significant proliferation of packaging types tailoring functionality and costs to the different markets. To address this proliferation, Intel focuses on packaging building blocks that can be configured for different applications. This paper traces the evolution of Intel's microprocessor packaging technologies, delineates the technical and business drivers, and highlights emerging trends. It then highlights the technical challenges faced by packaging developers now and in the future, and in a broad sense, it ties them into the challenges highlighted in the semiconductor industry technology roadmaps. Finally, it provides an introduction to the other papers in this issue of the *Intel Technology Journal*, which deal in greater detail with some of the technical challenges discussed in this paper.

INTRODUCTION

In unit volume, microprocessors account for a small percentage (~ ≤ 1%) of the semiconductor components sold worldwide. However, their technical and economic significance are far greater. Microprocessor packaging represents the technology envelope of this discipline. To better understand this statement, we present a historical perspective of microprocessors and follow with a review of the motivators and technology directions for this component of the semiconductor industry.

THE EVOLUTION OF PACKAGING

In the Beginning: The Mechanical Enclosure

For many years, wirebonding and ceramic packages were the base assembly technologies for microprocessors because of their versatility and reliability. This was also true for Intel. Intel's 4004 microprocessor and later, the 8080, 8086, and 8088 microprocessors were all housed in ceramic dual-in-line packages (DIP) that used wirebond connections. By today's standards, these microprocessors had few Input/Output (I/O) pins (less than 40) and delivered very modest performance

(<20MHz). The primary function of the package was to provide space transformation (i.e., fan out) of the I/Os in order to ease board routing and protect the chip from mechanical damage and from the environment. These were simple, single-layer packages. Figure 1 shows a typical example.

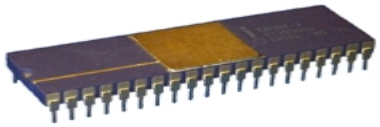

Figure 1: A 40 Lead DIP package used to package the 8088 and 8086 microprocessors

The Transition

In the i286™ and i386™ microprocessor generations, the number of I/O pins increased (~50 to 100) as greater functionality was incorporated into the microprocessor. This necessitated the use of Pin-Grid Array (PGA) packages in which a larger number of I/O connections could be accommodated in a small area. Also, in the i386 generation, it became evident that the increasing clock frequencies (a staggering 33MHz at the time) and simultaneous I/O switching could cause unwanted electrical coupling in the package manifesting itself as noise problems. Consequently, design and modeling competencies were substantially enhanced to account for these factors leading to the first use of multilayer ceramic packages at Intel. Figure 2 shows the i386 microprocessor package.

Figure 2: A 132L ceramic PGA package used for the i386™ microprocessor

Emergence of the Electrical/Thermal Platform

The i486 microprocessor was also housed in a PGA package with 168 leads. In addition to the basic function of connecting the I/Os, advanced electrical design concepts were incorporated. These included the use of power and ground planes as well as the inclusion of integrated capacitors in the package. These features transformed the package from a simple mechanical enclosure to a multilayer electrical distribution and signal-routing management platform.

The Intel® Pentium® processor helped in advancing this trend. In addition to the electrical features, the high-power dissipation, in the order of 15W, of the Pentium processor hastened the deployment of advanced thermal solutions such as an integrated heat spreader. These features, while effective, were expensive. An early version of the Pentium processor package is illustrated in Figure 3.

Figure 3: The Intel® Pentium® processor package in a ceramic PGA with a heat spreader

The Need for Increased Integration

The next-generation microarchitecture, commonly referred to as the P6 microarchitecture, introduced in the mid 1990s, represented a new era of performance and complexity. The microprocessor architecture called for a dedicated cache chip connected to the microprocessor via a Backside Bus (BSB). The first implementation of this architecture was on the Intel® Pentium® Pro processor where the microprocessor and cache chips were housed in a dual-cavity ceramic PGA package and connected by wire bonding. Because of the special requirement in the I/O configuration and because of the electrical performance of the cache memory of this microprocessor, custom Static Random Access Memories (SRAMs) were used, an expensive solution.

The second-generation implementation of the same microarchitecture utilized a cartridge form-factor. In this instance, the microprocessor and cache chips were housed in separate component packages and were connected using a standard printed circuit board. To start with, the microprocessor was assembled using Plastic Land Grid Array (PLGA) packages with wirebond technology, which later transitioned to Organic Land Grid Array (OLGA) packages that utilized Controlled Collapse Chip Connection (C4) technology. Aside from the transition from peripheral interconnect to area array interconnect, this packaging transition also marked the use of a high-performance package substrate technology,

i.e., OLGA technology. Plastic Quad Flat Package (PQFP) technology using wire bonding was used for the cache chips. This approach had two advantages over the dual-cavity ceramic PGA solution. First, it enabled the use of commodity Pipeline Burst SRAMs (PBSRAMs) thereby reducing cost. Second, the cartridge solution also allowed caches and other features to be customized for different market segments. The dual-cavity ceramic PGA and cartridge are shown in Figure 4. Figure 5 shows the portfolio of products packaged in a cartridge format.

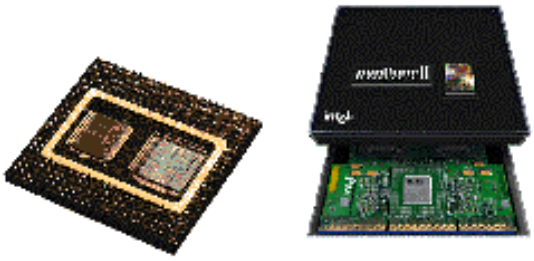

(a) (b)

Figure 4: (a) First-generation implementation of the P6 microarchitecture in a dual-cavity ceramic PGA (b) The Single Edge Cartridge Connector (SECC) cartridge is a second-generation form-factor

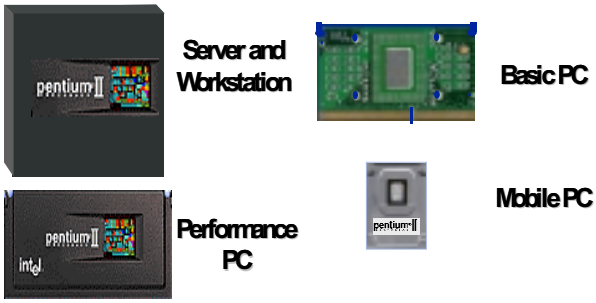

Figure 5: Portfolio of products utilizing cartridge packaging

Although the cartridge form-factor was an effective technical solution, the emergence of the cost-sensitive Personal Computer (PC) market demanded even more aggressive cost/performance packaging solutions.

Silicon Integration: Back to Single-Chip Packaging

Silicon feature scaling and the integration of the Level 2 (L2) cache directly into the microprocessor die were key enablers to lower the cost of packaging. Without the need for the multidie package or cartridge to service the high-speed BSB, it was possible to move back to single-chip packaging. Several single-chip packages were developed with form-factors based on market segmentation requirements.

Some of these microprocessor packaging form-factors included

- low-profile and high-density pinned packages for mobile applications (Figure 6)
- pinned packages for sockets in desktops (Figure 7); the package substrate is referred to as the Flip-Chip PGA substrate, another version of the organic substrate technology

(a)

(b)

Figure 6: Views of low-profile micro-PGA for mobile socket applications; the micro-PGA uses an OLGA substrate surface mounted to an interposer

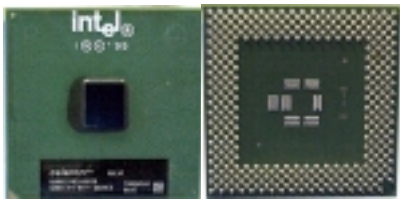

Figure 7: Pinned package for desktop socket; the package technology is referred to as Flip-Chip PGA (FC-PGA)

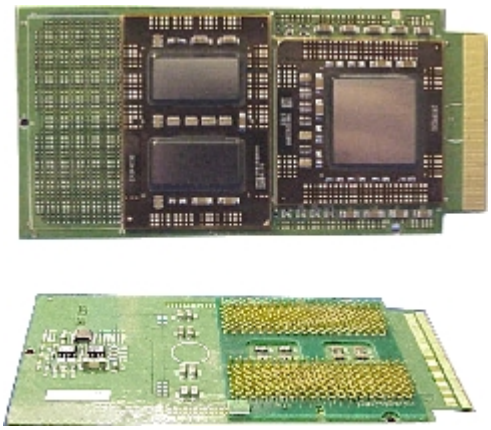

Figure 8: Intel® Itanium™ processor packaging shows how different elements of packaging can be combined

The focus at Intel has been to create packaging technology building blocks that can be combined to provide multiple features and form-factors, while minimizing piece part, process, and manufacturing costs. An example of this can be seen in Figure 8, which illustrates the packaging for the Intel® Itanium™ processor. However, before we discuss Intel's focus, we present an account of the technical and business drivers as well as the emerging directions for packaging technology.

TECHNICAL AND BUSINESS DRIVERS

Microprocessor packaging requirements are closely and intricately tied to the performance and architecture of the microprocessor. Figure 9 depicts the evolution of microprocessor/cache/bus architecture. Using this evolution as a framework, we examine five major drivers:

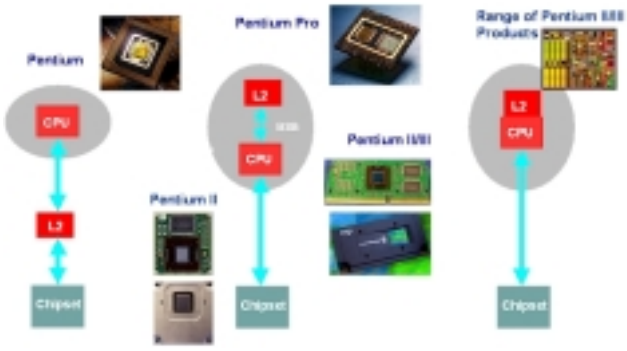

Figure 9: Evolution of the microprocessor and cache architecture from the i486™ microprocessor to the Intel® Pentium® Pro microprocessor

Driver #1 : Connecting the Cache

As the performance of microprocessors increased, the need to supply data and instructions to the microprocessor increased accordingly. This information normally resides in the main memories, such as the Dynamic Random Access Memory (DRAM) and disks, and it is channeled to the microprocessor via the bus, a parallel set of interconnects running between the microprocessor and the memory. The wider (i.e., more data lines) and faster the bus, the more data can be transferred at a given time. Since the days of the i486™ microprocessor, the speed requirements for data to be transferred to the microprocessor have exceeded the speed of the main memories (DRAMs). As a result, an L2 cache system utilizing fast Static RAMs (SRAMs) was added to the microprocessor architecture. This L2 cache stores frequently used data thereby reducing the need for frequent access to the external main memories. Consequently, this speeds up execution and leads to enhanced performance.

As described in the previous section, the Intel® Pentium® II/Pro microprocessor architecture had a dedicated BSB connecting the microprocessor to the L2 cache to further enhance the performance. Initially, this architecture was implemented by connecting the microprocessor and cache inside a ceramic package using wire bonding. This required custom SRAMs and expensive packaging. The implementation evolved to the use of a cartridge form-factor whereby commodity PBSRAMs were connected to the microprocessor by using a printed circuit board.

In the later silicon technology generations, improved Very Large Scale Integration (VLSI) density made it practical to integrate the cache into the same microprocessor chip. This accomplished two major objectives:

- It lowered the cost by eliminating the need for external PBSRAMs and the cartridge.
- It gave higher performance because of a full-speed BSB integrated into the silicon.

This is the current trend for future microprocessors. Consider, for example, the die shown in Figure 10. These are similar die except some have integrated caches and some don't. For a small increase in die size, it is possible to accomplish the two objectives mentioned above.

As silicon features shrink, this mitigates the initial chip size penalty of adding the L2 cache.

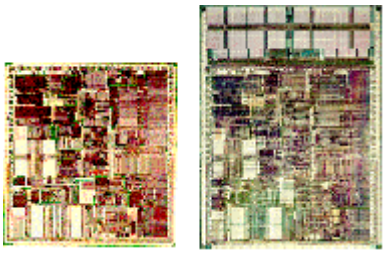

(a) (b)

Figure 10: P6 architecture die with and without integrated L2 cache

Future microprocessors may also integrate part of the chipset into the same silicon thereby further reducing the interconnect complexity and costs at the system level.

Driver #2: Connecting the Bus

Although the incorporation of the L2 cache alleviates data traffic on the microprocessor bus, an increase in the bandwidth of the microprocessor bus is still necessary in many applications where I/O bandwidth is important, such as graphics and servers.

To enable high performance, the microprocessor bus speed evolved from an 8 MHz Industry-Standard Architecture (ISA) bus on the original PC-Advanced Technology (AT) to a 100 MHz bus on today's microprocessors. Moreover, there are clear indications that we need to further increase this speed and bandwidth. Aside from the raw speed, additional challenges in high-end systems include the use of multiple processors on a single microprocessor bus. This requires the support of several electrical loads on the same bus thereby necessitating very precise electrical designs to achieve the desired performance. Figure 11 illustrates this configuration.

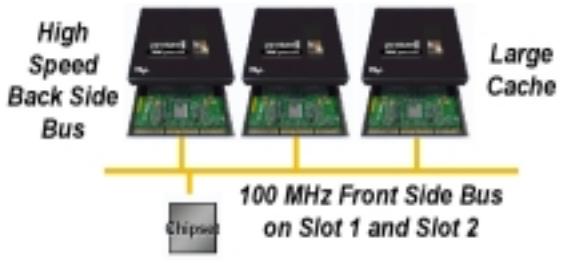

Figure 11: Multiprocessor on a microprocessor bus

To manage these high-performance electrical environments, the interconnections from the silicon through the package to the system board must be designed as an integral unit in order to ensure the desired electrical characteristics. From a technology viewpoint, this requires high-conductivity interconnect traces, low capacitance, and matched impedance for high I/O speed to minimize noise generation. The design of the I/O drivers/receivers on the silicon must also account for package as well as system interconnects. Careful matching of impedance, voltage signal levels, and timing is essential to guarantee performance.

Driver #3: Power Delivery

The third driver is delivering power to the chips. A key element that enabled the advances in silicon technology and the resultant density and performance improvements from generation to generation is the scaling of the supply voltage. While this approach is beneficial to silicon scaling and power dissipation, the challenge of delivering power to the silicon chip via the package is increased. There are two elements to power delivery:

(a) DC (average) supply current

As the supply voltage was scaled, the integration of additional functions and operations at higher clock speeds kept the power dissipation high. As a result, the average supply current increased significantly. This high current was delivered from the power supply on the system motherboard to the chip through the package. As illustrated in Figure 12, the typical supply currents were in the 10 – 20 A range, a range that is expected to increase for future processors. To handle this high current, the package must provide a very low resistance path, in the order of < 1mili-ohm.

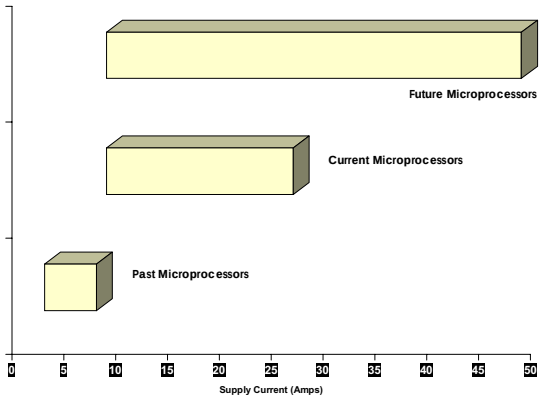

Figure 12: Supply current trends for microprocessors illustrating the effects of voltage scaling and ever-increasing performance

(b) AC (di/dt switching) current

Even more challenging is the management of the switching current. The high clock speed circuits and power conservation design techniques such as clock gating and SLEEP mode result in fast, unpredictable, and large magnitude changes in the supply current. The rate

of change of many Amps per nanosecond of this switching current far exceeds the ability of the power supply and the voltage regulator to respond. If not managed, these current transients manifest themselves as power supply noise that ultimately limits how fast the circuits can operate. This is further compounded by the reduced noise margin in the Complementary Metal-Oxide-Silicon (CMOS) logic circuits that result from power supply voltage scaling.

To mitigate this undesirable noise effect, the package must provide a very low inductance path for the switching current. In addition, charge storage devices, in the form of capacitors native to the silicon chip and augmented by capacitors on the package, are also necessary in some designs.

Driver #4: Dissipating Power

Dissipating high power, and managing high-power density, is another challenge. With the high density of integration and high clock rates, advanced microprocessors dissipate a significant amount of power in a very small physical area. Figure 13 illustrates the problem .

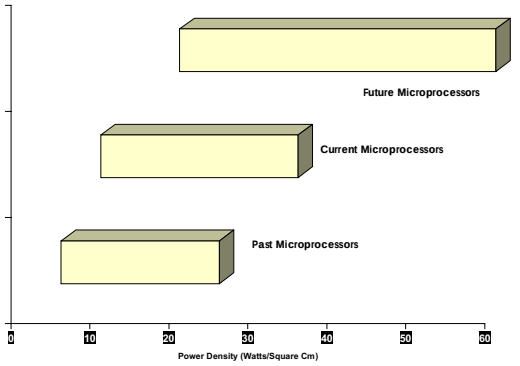

Figure 13: Power density trends of microprocessors

Another factor that exacerbates the thermal management problem is that local areas of the die, depending on where different functions are executed, have higher power densities than the average power density. The challenge to packaging is to ensure that the thermal path from die to the environment is optimized to allow for effective spreading and ducting of heat to the environment. In a broad sense, thermal management involves

- accurately estimating the power dissipation requirements, including power, on-die power distribution, and die temperature expectations,
- managing the thermal performance of all interfaces in the thermal path from die to the environment,
- providing and optimizing the spreading characteristics of all the thermal elements from die to the environment including the package heat-spreader and the heat sink, and finally,
- managing the thermal environment in the chassis by ensuring that the local air temperature is as low as possible to provide a better environment for the microprocessor to dissipate heat.

It is clear that the increasing thermal challenge requires advanced thermal management to ensure chip functionality and reliability.

Driver #5: Silicon Density and Die Shrinks

As silicon technologies advance, the size of the physical feature that can be fabricated gets smaller. Correspondingly, a given amount of circuitry can be built in a smaller die area. Both Intel and the rest of the semiconductor industry employ an aggressive die shrink or die compaction strategy to reduce the silicon area. This approach has two major benefits. First, by reducing the die area, more die can be fabricated on the wafer, hence the cost is lower. Second, a reduction in die area results in higher speed and lower power dissipation for the same speed. This trend is illustrated in Figure 14.

Figure 14: Die shrinks driven by advances in silicon technologies necessitate corresponding improvements in chip-to-package interconnectivity

As the die size shrinks, the number of I/O connections does not change. Furthermore, the number of power supply connections is often increased to support the performance increases brought on by the die shrink. These factors result in a decrease in the bonding pitch for wire bonding or bump pitch for flip-chip. In order to keep pace with this trend, the package geometries and the assembly technologies must also evolve. Today, very fine pitch wire bonding has brought wire bonding down to a pitch of 65 microns. The pitch used on flip-chip arrays is considerably larger, currently in the range of 200 microns, as it utilizes the entire surface of the die to lay out the bumps. Nevertheless, this bump pitch still has to be scaled to keep pace with silicon scaling and die size reduction.

Driver #6: Socketability

Socketability is a business requirement. The reasons for socketability include Original Equipment Manufacturer (OEM) inventory control, the impact of tax and duty, and manufacturing flexibility.

From a technical standpoint, socketability is not desirable. In most cases, it makes the package larger, more expensive, and the performance is lower. Nevertheless, the quest for a low-cost, high-performance socketable package is a strong business-driven requirement.

THE TECHNOLOGY AGENDA

To meet these challenges, the Assembly Technology Development group within Intel has been engaged in defining and creating many new technologies to serve as building blocks as well as integrating the design and analysis environments. The key building blocks are as follows**:**

1. A packaging technology that has high electrically conductive metallurgy that minimizes the IR drop and acts as a high-current conduit to deliver power from the power supply to the chip, such as copper conductors in organic packages.
2. Low-inductance connections from chip to the package and from the package to the socket, such as flip-chip interconnects.
3. Low-capacitance insulator materials, such as organic packages.
4. Advanced thermal-interface materials and a focus on thermal design to manage the high-power density.
5. An integrated analysis, design, and validation environment that enables dynamic trade-offs between chip and package design and layout in the interconnect continuum that includes Computer Aided Design (CAD) tool suites, test vehicles, etc.
6. Predictive models especially in power delivery, power dissipation, and Electromagnetic Interference (EMI).

These building blocks have been in development at Intel for the past several years. Accordingly, a number of significant technology transitions is already underway.

Transitions

Away from Wire Bond and Ceramics

Versatile, ceramic package technology can be expensive. Furthermore, as performance increases, the physical characteristics of ceramic packages may become limiting. Specifically, a ceramic material based on Al_2O_3 has a relatively high dielectric constant ($\varepsilon_r \sim 7$-8). Additionally, because of the high-temperature processing, metallization is limited to refractory metals, such as Molybdenum and Tungsten, which are quite resistive. Wire bond connections have relatively poor electrical characteristics because of their high inductance. As described above, high-resistive, inductive, and capacitive structures are not conducive to high performance. Ceramic substrate suppliers are addressing some of these limitations with advances in their materials.

Use of Organic Packages

A key thrust pursued at Intel, was the transition from ceramics to organic laminate packages. It started in 1996 with the introduction of the Intel® Pentium® processor in the Plastic Pin Grid Array (PPGA) package. Today, all of Intel's microprocessors are in organic laminated packages.

In contrast to ceramic packages, organic laminated packages utilize epoxy resin dielectric materials ($\varepsilon_r \sim 4.2$) and copper conductors. These low dielectric characteristics and copper result in substantial improvements in power distribution and signal transmission characteristics.

The organic package is also indigenous to our latest flip-chip packaging. The attributes of this package provide significant boosts to power distribution and signal routing on the chip. The table below contrasts the physical and electrical characteristics of a typical silicon circuit versus that of a flip-chip OLGA. In short, the routing density is much higher on the silicon, but the electrical characteristics are much better on the organic package. Hence, a judicious use of these capabilities in an interconnect continuum can result in optimal product performance and cost.

Conductor	Pitch (µ)	Material	Thk (µ)	Sheet Rho (mΩ)	Insulator
Si	0.5	Al-Cu	0.5	~85	Oxide
C4-OLGA	70	Cu	17	~0.01	Epoxy

In order to meet the tight pitch demands for chip-area array interconnect (C4 discussed in next section), it was necessary to construct a new organic package. This package uses a laminated core set of layers with high-density "build-up" layer(s). The high-density layers are used to match the pitch on the die. This package is illustrated in Figure 15.

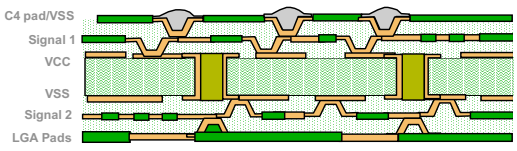

Figure 15: Cross-sectional view of organic flip-chip package illustrating core and high-density build-up layers

Refinements to this technology have been developed to allow alternative package form-factors, as described earlier, based on market segment needs.

Another advantage of organic packaging is that the Coefficient of Thermal Expansion (CTE) of the package is more closely matched to that of the motherboard as compared to a ceramic package. This ensures that the stresses induced in the package-to-motherboard interconnects are significantly lower, resulting in more reliable interconnections even when the interconnect count is large. This is especially true of Ball Grid Array (BGA) connections where interconnect reliability is of significant importance.

C4 Flip-Chip

In order to optimally access the superior electrical characteristics of the organic package, we must also establish a high-density, high-performance method to connect the chip to the package. To that end, a solder-bump-based C4 area array flip-chip capability was developed to replace wire bonding.

In contrast to traditional wire bonding, C4 utilizes an area array method for interconnection. The C4 bumps can be placed anywhere on the die, even over active circuitry. This enables the placing of many more bumps as virtual vias (through the thick electrical connections) connecting the metallization on the chip to the metallization on the package. In fact, the metallization on the package can be visualized as metal layers augmenting the metal layers on the chip. The primary benefit of this approach is in power/ground distribution and clock routing. The C4 connections, in combination with the electrical characteristics of the copper-based organic packages, result in a superior electrical environment where maximum performance can be realized. As an example of this implementation, consider the same die, P6 architecture, in both the wire-bonded and C4 versions as shown in Figure 16. The substantial increase in power/ground connections ensures maximum performance. Additionally, a native C4 die design eliminates the need for bond pads, which results in a small die, ~0.012 inches smaller per side.

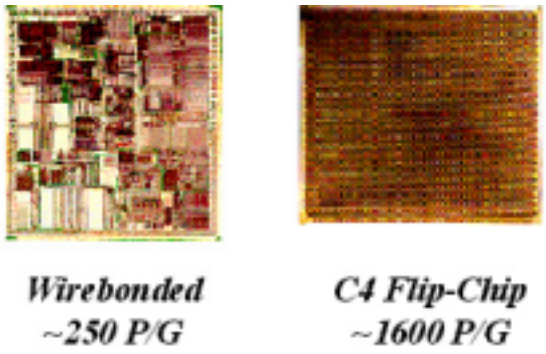

Figure 16: P6 architecture microprocessor implemented as wire-bonded die and in C4 Flip-Chip

In this paper, we discuss packaging and technology building blocks as a key concept to cost effectively meet a wide range of both form-factor (surface-mount, high-density pinned, low-density pinned, etc.) and performance needs for packaging by a judicious combination of these building blocks.

Figure 17 further illustrates this concept. A single bumped die is mounted on either a surface mount (OLGA) or a pinned substrate. The OLGA substrate can subsequently be surface mounted directly to a board, mounted to a pinned substrate for socketing, or mounted to a cartridge (to be combined with other chips).

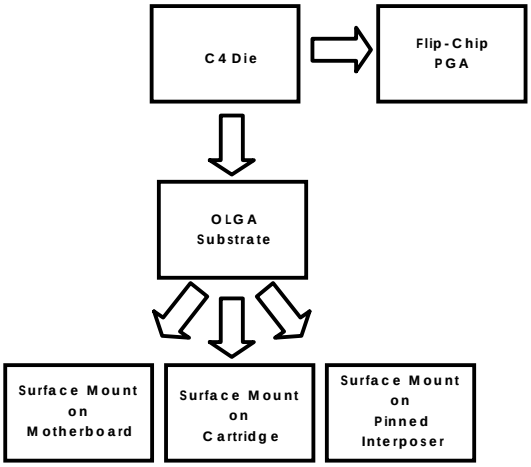

Figure 17 : Building-block technologies

FUTURE CHALLENGES

As microprocessors continue to improve in performance, technical challenges in packaging will also increase. A comprehensive view of these challenges can be found in the 1999 International Technology Roadmap for Semiconductors (ITRS) [1]. This roadmap discusses the need for improved materials and assembly processes as well as a need to have integrated simulation tools and

methods to assess the reliability of the integrated die-package-motherboard system. Since the design environment and the assembly processes and reliability aspects of packaging fall outside the scope of this paper, we limit this discussion to the technical challenges in packaging as they impact microprocessor performance. Technically, the challenges fall into three broad categories: power delivery, power removal, and also the provision of viable, i.e., appropriately scaled with optimal performance characteristics, interconnection strategies between silicon and board.

Power delivery challenges are highlighted in Figure 12. To move forward, the focus will have to be on continuing to understand and optimize the electrical path from power supply to the die. With increasing demand for performance, the general separation of market segment requirements and constraints, and the shortening of the time-to-market, it is expected that the power delivery solutions will continue to be challenging.

Power removal, i.e., thermal management of the processor, is another increasingly challenging aspect of packaging. As shown in Figure 13, the average power density of processors is expected to increase. The problem will be exacerbated by the need to manage local power densities on die. The development of cost-effective and technically viable thermal management solutions that maintain die temperature at acceptable levels will be key to ensuring future success. This can be accomplished through development and deployment of effective spreader solutions and thermal interface materials. Controlled assembly processes to manage the thermal interfaces are also a key to successful design. Finally understanding and managing the die power, power distribution, and the thermal environment in the chassis are important.

Silicon technology in the future is expected to scale aggressively, which will require intelligent space transformation methods from packaging. Ensuring that the interconnects have refined electrical characteristics so that packaging provides the appropriate space transformation while enabling the required electrical performance will be essential to future development.

SUMMARY

In this paper, we discuss the evolution of microprocessor packaging from a simple protective enclosure to a more technically complex and challenging platform that enables optimal microprocessor performance. The general strategy adopted within Intel to address continuing challenges is to develop building blocks that can be effectively combined to meet the needs of current and future microprocessor packaging. The remaining papers in this issue of the *Intel Technology Journal* discuss different aspects of these challenges in greater detail.

The 2nd paper discusses the FC-PGA package, i.e., flip-chip technology on organic pin grid array substrates. The paper looks at the motivations that led to the development of this package technology, its characteristic features and capabilities, and some of the issues that were successfully addressed during the development and deployment of this technology into high-volume manufacturing.

The 3rd paper discusses the technical complexity of interconnect design to achieve optimal electrical performance. This paper also discusses the design analysis and synthesis techniques used to ensure optimal electrical design.

The 4th paper discusses challenges in power removal. Thermal solutions that are optimized for cost and performance and tailored to meet different market segment needs are a key enabler to successful microprocessor deployment. This paper discusses some of the considerations that must be taken into account to ensure successful thermal design.

Underfill processes and underfill material development are a major component of flip-chip packaging processes. The 5th paper discusses a novel method of accomplishing this objective.

Ensuring that packaging continues to meet high standards of reliability is a key to success. The 6th paper discusses how our assessment of reliability has evolved during the past few years. Intel has moved from a stress-based certification strategy to a more fundamental mechanism-based methodology that allows for a better linkage between stress testing and the end-user environment.

Finally, the 7th paper talks to the practical problem of managing the thermal environment during microprocessor testing. The goal of testing is to effectively assess performance and reliability without introducing artifacts due to the testing process. This paper examines how this goal can be accomplished and looks at some of the unique issues that should be addressed.

CONCLUSIONS

High-performance and cost-effective microprocessor technologies require a holistic approach that comprehends the interconnect continuum including the silicon, the package, and the system. By properly exploiting the attributes of these regimes, optimal performance and cost can be realized. This review of the evolution of packaging reinforces the view that it will be a technically challenging and rewarding area of focus.

ACKNOWLEDGEMENTS

The authors thank Bill Siu for providing the essential framework of this paper. We also acknowledge the support of Scott Shirley, Susan Coartney, and Keith Carnes in procuring some of the graphics for this paper.

REFERENCES

[1]"International Technology Roadmap for Semiconductors," 1999 Edition, *Semiconductor Industry Association*.

AUTHORS' BIOGRAPHIES

Ravi Mahajan is currently the manager of a Thermal and Mechanical Analysis group in Assembly Technology Development. His group provides design and analysis support in the development of advanced packaging for Intel processors and chipsets. He also leads a group of thermal experts from different organizations within Intel. This group is chartered with developing and maintaining the thermal roadmap for Intel microprocessors. Dr Mahajan received his B.S. degree in Mechanical Engineering from the University of Bombay (1985), his M.S. degree in Mechanical Engineering from the University of Houston (1987), and his Ph.D. degree in Mechanical Engineering specializing in fracture mechanics from Lehigh University (1992). Dr. Mahajan holds several patents in packaging, has edited two conference proceedings for the Society of Experimental Mechanics, and is currently one of the editors of an Intel Internal Technology Journal. His e-mail is ravi.v.mahajan@intel.com.

Ken Brown is manager for non-CPU Packaging Development in the Assembly Technology Development group at Intel. He has worked in the semiconductor industry for 19 years, and currently serves as Intel's representative to the SIA International Technology Roadmap for Semiconductors (ITRS) packaging committee. His first 17 years were with Digital Equipment Corporation, where his last position was as Manager of the Package Design & Development organization. He moved to Chandler and ATD about two years ago following Intel's acquisition of Digital Semiconductor. He received a B.S.M.E. degree (1981) from the University of Massachusetts, an MBA from Northeastern University (1985), holds two patents in packaging, has presented at several packaging conferences, and authored a chapter on packaging for a graduate text book. His e-mail is ken.m.atd.brown@intel.com.

Vasu Atluri is currently a Silicon Integration Manager in Assembly Technology Development. His group is responsible for ensuring that all technical issues are properly managed to enable a successful interface between silicon and assembly technology development groups. He is also responsible for the design of test chips that help validate performance and reliability of silicon after it has been packaged. He also helps in defining design rules for product designs.

Dr Atluri received a B.S. in Chemical Engineering from Osmania University (1982), anM.S. in Materials Science and Engineering from University of Arizona (1987), an M.S. in Chemical/Metallurgical Engineering (1988), and a Ph.D in Materials Science and Engineering from University of Arizona (1998) with a specialization in silicon fabrication. Dr. Atluri has more than twenty technical publications in conference proceedings and journals. He has presented at several technical conferences. He has filed for couple of patents in silicon processing. Dr. Atluri is currently one of the editors of Intel's Assembly and Test Technology Journal, which is Intel Corporation's Internal Technology Journal. Dr. Atluri is an active member of IEEE and is currently serving as Vice-Chair for IEEE Phoenix Section. Dr. Atluri maintains active participation in a wide array of academic forums associated with electronic packaging. His e-mail is vasudeva.atluri@intel.com.