Chip Scale Review®

ChipScaleReview.com

The Future of Semiconductor Packaging

Volume 30, Number 2

Spring 2026

Enabling next-generation advanced integration and rapid innovation for photonics, quantum, and beyond with multi-column e-beam lithography

- Cracking the HVM code for glass core substrates
- Fine-grain copper for low-temperature hybrid bonding
- Pushing NAND packaging to power an AI-driven, data-hungry world
- Meeting AI requirements by scaling down technology and scaling up integration
- Production scaling of silicon photonics wafer testing using automated test equipment platforms
- Silicon photonics applications solutions and technology challenges in advanced co-packaged optics

LEENO

www.leeno.com

Fine Pitch Probe & Probe Head

Proven Mass Production Capability

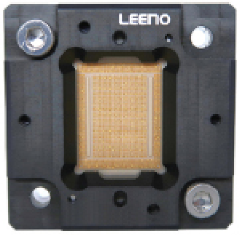

High Speed RF

Short Signal Path 0.80mm
with Spring Contact Probe!

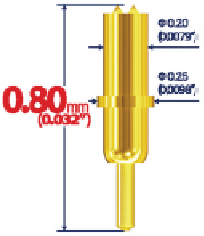

Through Continuous
Non-stop Innovation

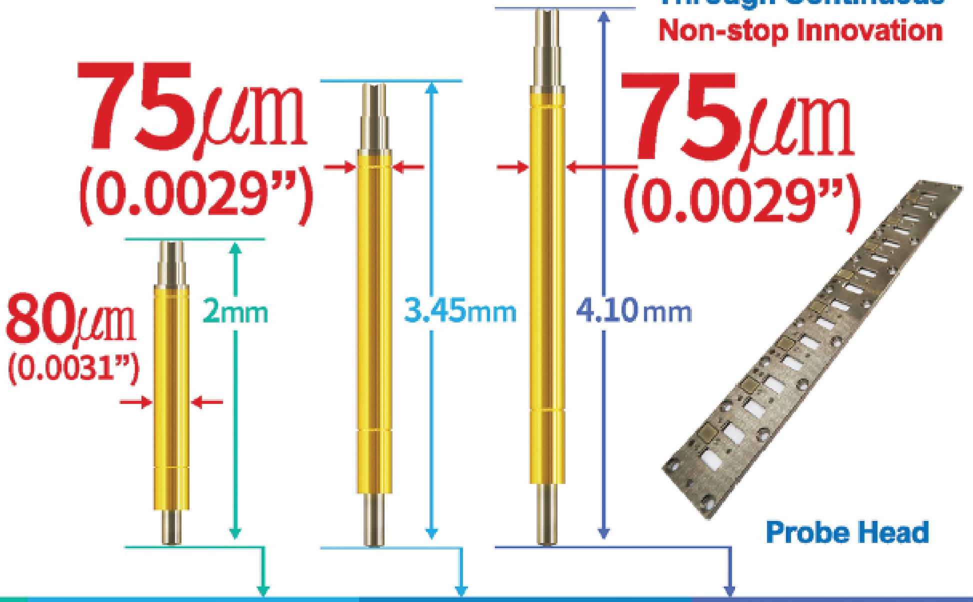

RF-300㎛ Pitch	97㎛ Pitch	92㎛ Pitch	92㎛ Pitch

5G MP 100GHz Coaxial Socket

Insertion Loss

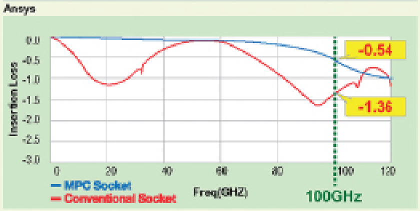

Return Loss

Crosstalk

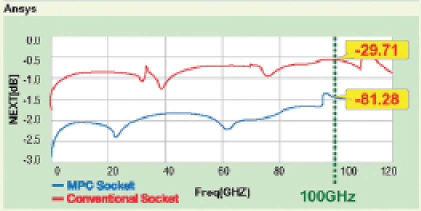

- Pitch: 0.35mm, Array

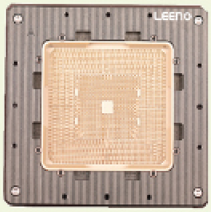

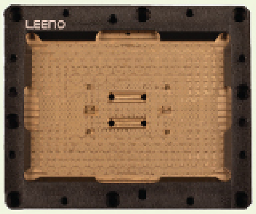

GLOBAL LEADER LEENO

HEAD OFFICE
10 105beon-gil MieumSandan-ro Gangseo-gu, Busan Metropolitan City, Korea

CONTACT
hskang@leeno.co.kr / +82-51-792-5646 / +1-408-784-5855
kimmj@leeno.co.kr / +82-51-792-5642

CONTENTS

Volume 30, Number 2
Spring 2026

DEPARTMENTS

TECHNOLOGY TRENDS

INDUSTRY EVENTS

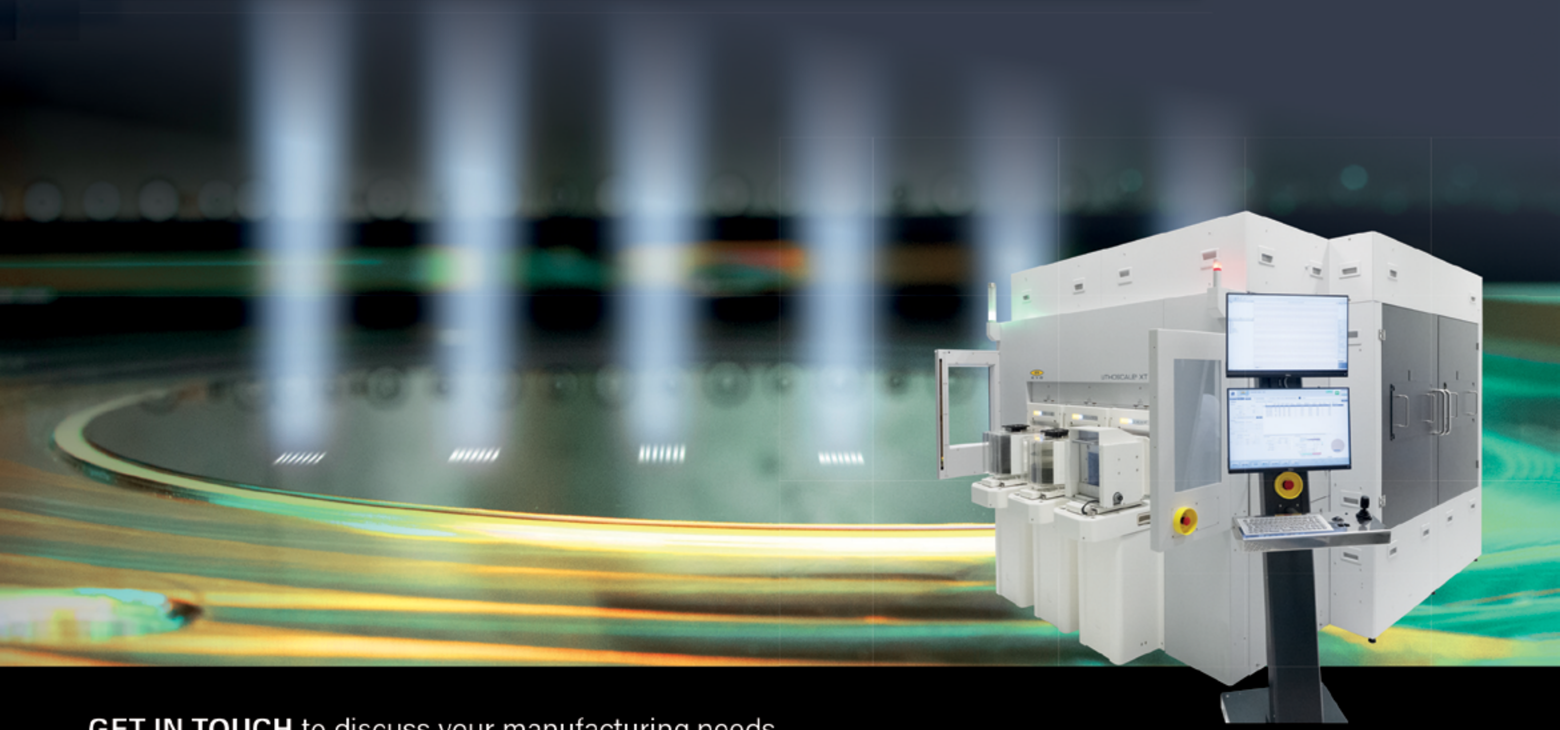

CONTENTS

Chip Scale Review
ChipScaleReview.com
The Future of Semiconductor Packaging

STAFF
Kim Newman
Publisher
knewman@chipscalereview.com

Lawrence Michaels
Managing Director
Editor-in-Chief
lmichaels@chipscalereview.com

Debra Vogler
Senior Technical Editor
debravogler@me.com

SUBSCRIPTION—INQUIRIES
Chip Scale Review
All subscription changes, additions, deletions to any and all subscriptions should be made by email only to
subs@chipscalereview.com

Advertising Production Inquiries:
Lawrence Michaels
lmichaels@chipscalereview.com

Subscriptions in the U.S. are available without charge to qualified individuals in the electronics industry.
Chip Scale Review is published four times a year with issues in Winter, Spring, Summer and Fall.

P.O. Box 2165
Morgan Hill, CA 95038
E-Mail: info@chipscalereview.com

Printed in the United States

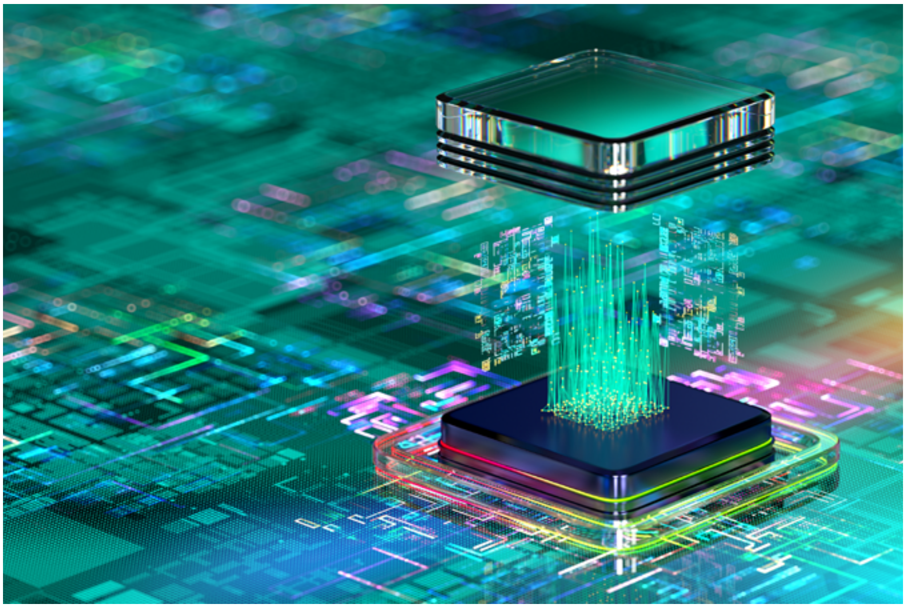

FEATURE ARTICLES

TECHNOLOGY TRENDS

Meeting AI requirements by scaling down technology and scaling integration

By Pratyush Kamal, Todd Burkholder [Siemens EDA]

3D integrated circuits (ICs) have been in commercial applications for optical image sensors for more than a decade. 3D IC configuration has also been a standout factor in AMD's resurgence. 3D ICs are now enabling the next generation of server class central processing units (CPUs) and artificial intelligence (AI) accelerators—both for training and inference—meeting the needs of both ever larger compute resources and higher compute efficiency. 3D ICs are seen as key "economic" enablers for very expensive 2nm and smaller process nodes.

At the same time, AI is redefining how we do 3D IC designs, relying primarily on traditional tribal knowledge about solutions and workflows that are human-in-the-loop automations. An understanding of six key trends that lie behind this shift will help leading-edge teams manage their design workflows so they can adopt and leverage 3D IC innovation.

The key to AI performance scaling: 3D IC

Between 2012 and 2018, AI compute-demand doubled every 3.4 months. More recently, this pace has slowed down to doubling every 7 months. This is pushing transistor density and bandwidth demand beyond what monolithic systems-on-chips (SoCs) can deliver. At the same time, the industry has run into a fundamental physical limit for feature sizes. Advanced lithography caps reticle size at roughly 858mm^2—the largest mask pattern that can be printed on a wafer; however, increased defectivity and poor yield in advanced technology nodes require dies to be much smaller to be economically viable. Side-by-side integration of split dies in 2.5D advanced packages is approaching the limits of die-to-die interface technology, as well as the form-factor limits of end systems.

As a result, scaling next-generation AI and high-performance computing systems requires a fundamental shift in architecture—where performance gains come from proximity in three dimensions (**Figure 1**). Logic and memory must be disaggregated then re-integrated across wafers, vertical stacks, and advanced packages, while ensuring thermal, electrical, and mechanical reliability.

Trend 1: Technology scale-down and integration scale-out

The advanced packaging roadmap aims for feature size reduction to achieve higher integration density, while boosting performance with lower power

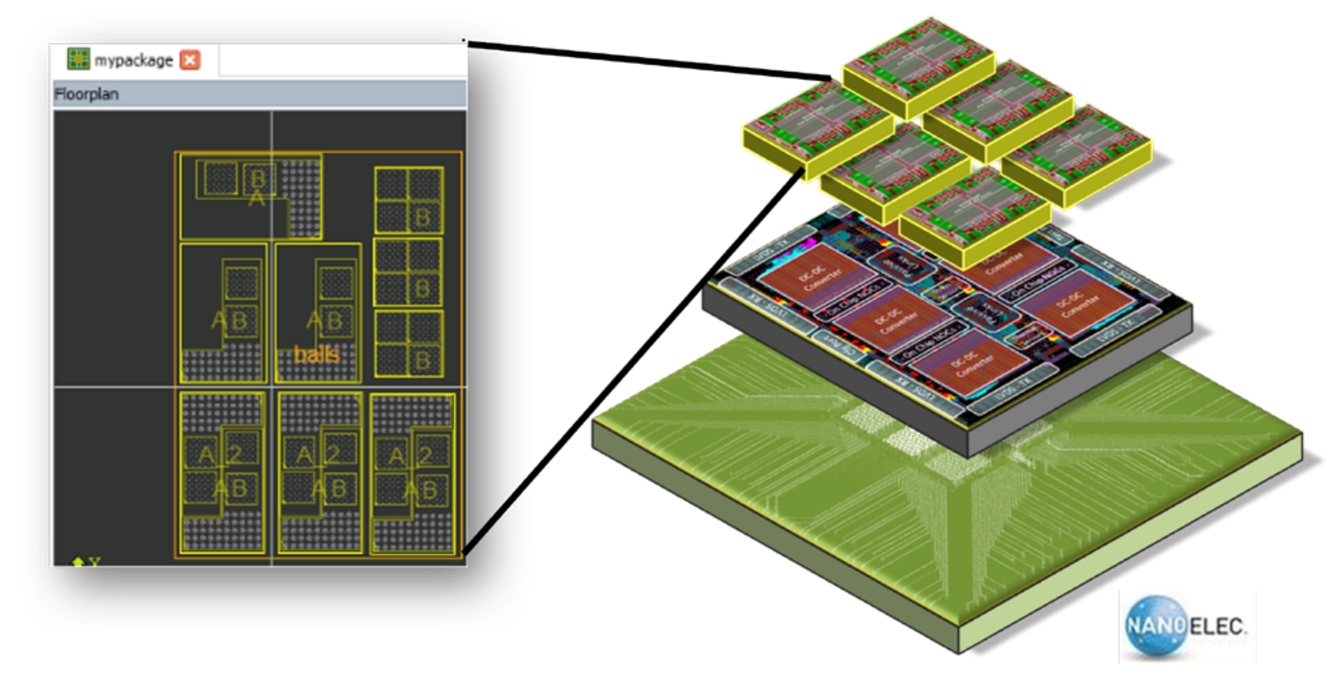

Figure 1: Arrayed blocks are used to construct a chiplet.

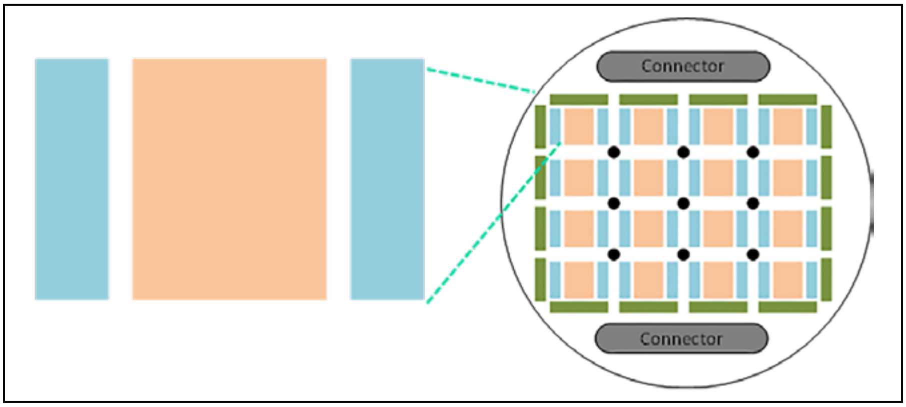

Figure 2: A wafer-scale package for server class application.

(**Figure 2**). As thermal-compressed bonds reach their limits, hybrid bonds will push 3D interconnects to 1µm and below. AI and high-performance computer (HPC) suppliers are also exploring wafer- and panel-level architectures for denser computing. Examples include Cerebras' WSE-3 accelerator, integrating trillions of transistors, and TSMC's modular SoW-X, which assembles pre-tested dies onto reconstructed wafers. This shift necessitates electronic design automation (EDA) tools capable of wafer-scale, multi-physics-system simulation for accurate power, thermal, signal, and mechanical modeling.

Trend 2: Let there be light: Unlocking 3D CPO for AI systems

Optical interconnects are crucial for large-scale compute engines due to their low attenuation. As AI systems grow, electrical links face rising power consumption and bandwidth density limits. Co-packaged optics (CPO) integrate electronic and photonic ICs within the same package, shortening link distances to millimeters, which improves signal integrity and bandwidth. Foundries like TSMC, with its Compact Universal Photonic Engine (COUPE™), demonstrate the practicality of stacking electronic dies directly on photonic dies. However, CPO demands rigorous device validation and testing because 3D architecture complexities introduce thermomechanical stresses and thermal-optical effects that impact alignment and reliability.

Trend 3: Cooling silicon "skyscrapers" with microfluidics inside the package

Stacking memory vertically as a multi-layer "silicon skyscraper" shortens interconnects, enabling higher bandwidth. The challenge is dynamic random access memory's (DRAM's) temperature sensitivity (often limited to 85°C, compared to logic at 125°C). In stacked systems, thermal headroom becomes a primary concern. Aggressive cooling strategies are needed to keep high-bandwidth memory (HBM) within specification while maintaining high logic power density. Emerging microfluidic approaches address this by directly bringing coolant micrometers from active transistors, using etched micro-channels to significantly improve heat transfer. This approach makes accurate 3D models of micro-channel networks and multi-physics solvers for coupled electrical, thermal, and mechanical effects essential.

Trend 4: Material innovations advancing AI, HPC, and 6G applications

Recent advancements in manufacturing and materials science are significantly improving semiconductor chips and electronics packaging. First, glass substrates are gaining traction for large-area and high-frequency designs. Mechanically, glass's thermal expansion coefficient (~3.2ppm/°C) is similar to silicon, reducing package warpage by nearly 50%. Electrically, its low dielectric constant supports reliable signaling for data rates up to 224Gbps and 6G radio frequencies (RF). Manufacturing techniques like laser-induced deep etching and fine-pitch through-glass vias (TGVs) make glass practical for heterogeneous integration. However, glass's brittleness makes it paramount to accurately model interface stress and the risk of fractures, especially as substrates are thinned. Second, polymer-based packaging materials are now critical for reliability, performance, and cost. As chips transition to 2.5D- and 3D-stacked chiplets, mechanical stress becomes a major challenge. Polymers absorb these stresses, preventing yield collapse and reliability failures.

Trend 5: Rising need for mmWave antenna-in-package design

As 6G research progresses, antenna-in-package (AiP) adoption will reach an inflection point. Operations in sub-THz and >100GHz regimes impose constraints on interconnect loss, antenna efficiency, and phase accuracy. As wavelengths shrink, antenna elements can be integrated directly inside a package. This eliminates board-level routing, enabling dense, phased-array architectures with advanced beamforming, tighter RF-to-antenna coupling, and lower losses. Recent AiP architectures increasingly leverage 3D integration—stacking multiple dielectric layers, embedded ground planes, and parasitic elements close to the RF IC, thereby minimizing electrical path length and suppressing parasitic loss. However, integrating antennas into the package tightens manufacturing and design tolerances. Micron-scale variations at sub-THz frequencies can degrade system efficiency. The proximity of RF, digital logic, and power delivery networks also increases susceptibility to coupling, electromagnetic interference (EMI)/electromagnetic compatibility (EMC) issues, and thermally-induced RF drift.

Trend 6: 3D AI for 3D IC design

As 3D IC complexity grows, traditional automation and point optimizers struggle with both scale and design-space explosion. Large-scale packages demand new approaches for pathfinding, prototyping, and signoff to achieve optimal system technology co-optimization (STCO) (**Figure 3**). This drives a shift toward AI-native workflows for 3D ICs, where large language models (LLMs), optimization engines, and domain-specific AI operate across the entire

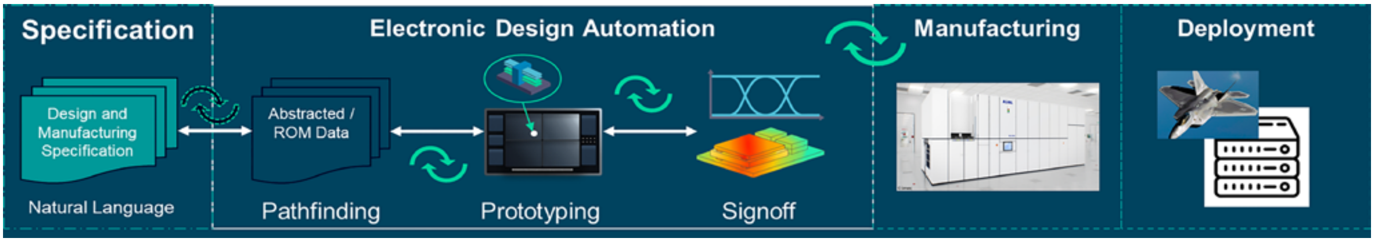

Figure 3: Specification-to-deployment: A new design paradigm.

design lifecycle. EDA vendors are already applying AI to design exploration, routing, and multi-physics optimization.

The next step is orchestration by generative and agentic AI. In early architecture, LLMs can convert human-readable specifications into machine-ready formats, thereby accelerating analysis. During implementation, AI can explore vast design spaces, learning from prior designs to converge faster. At signoff, LLMs synthesize cross-domain results into concise summaries, closing the loop between analysis and decision making. This "3D AI" paradigm scales human expertise, enabling faster exploration, deeper insights, and confident decision-making.

Becoming a 3D AI expert

Designers now have more factors to worry about: material choices, architectural decisions, and tightly-coupled interactions across electrical, thermal, and mechanical domains—with far less margin for late-stage correction. Getting key trade-offs correct earlier in the flow has become essential.

Our focus is on helping 3D IC design teams manage 3D IC complexity through multi-physics-aware, die-to-system design flows powered by AI (**Figure 4**). These workflows are built to provide earlier visibility into cross-domain interactions, so engineers can understand the implications of their decisions sooner and evaluate alternatives with confidence.

Biographies

Pratyush Kamal is the Director of 3D IC R&D Solutions at Siemens EDA, San Diego, CA. Kamal is an experienced SoC and systems architect and silicon technologist providing technical leadership for advanced packaging and new foundry technology programs. Kamal previously held various positions at Google and Qualcomm as SoC designer, SoC architect, and systems architect. He also led 3D IC research at Qualcomm, focusing on both wafer-on-wafer hybrid bonding and monolithic 3D design integrations. Email pratyush.kamal@siemens.com

Todd Burkholder is a Senior Editor at Siemens DISW, Portland, OR. For over 25 years, he has worked as editor, author, and ghost writer with internal and external customers to create print and digital content across a broad range of EDA technologies. Todd began his career in marketing for high-technology and other industries in 1992 after earning a Bachelor of Science degree at Portland State University, and a Master of Science degree from the University of Arizona.

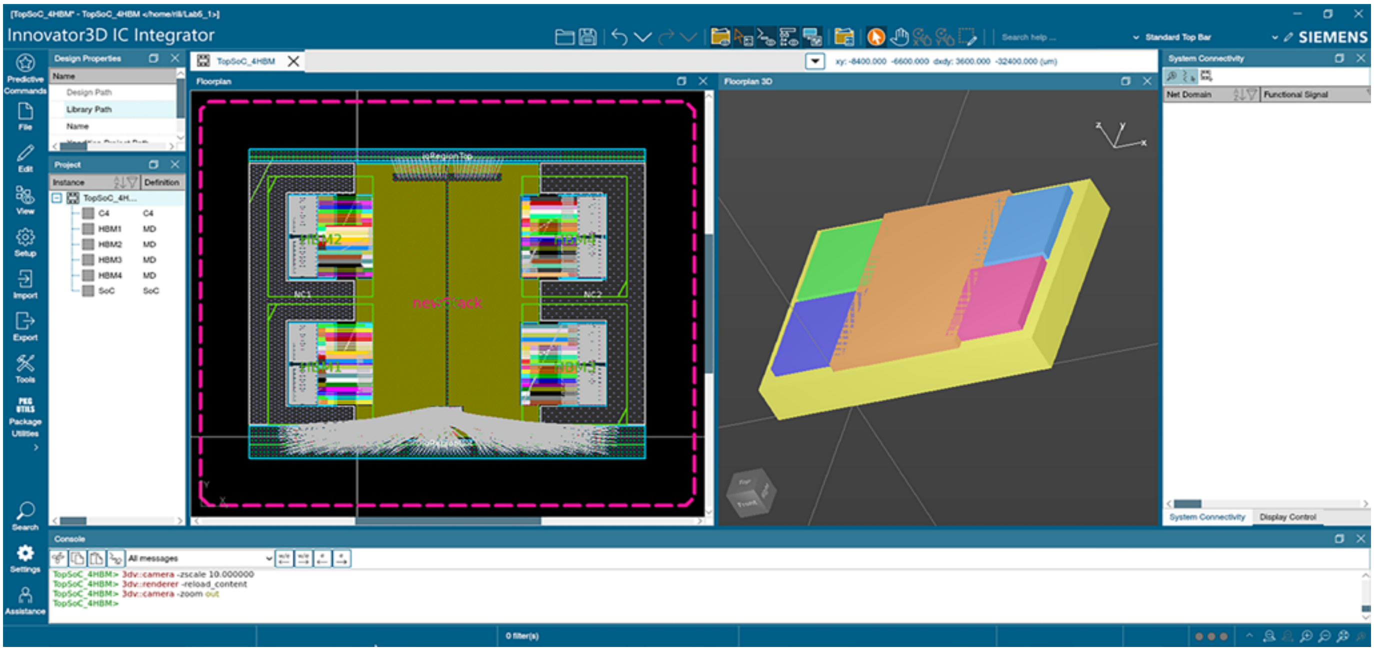

```
TopSoC_4HBM> 3dv::camera -zscale 10.000000
TopSoC_4HBM> 3dv::renderer -reload_content
TopSoC_4HBM> 3dv::camera -zoom out
TopSoC_4HBM>
```

Figure 4: Innovator3D IC solution suite cockpit.


Amkor
Amkor

Cracking the HVM code for glass core substrates

By Bilal Hachemi [Yole Group]

The semiconductor industry is growing rapidly to the one-trillion-dollar market value cap due to the aggressive demand from artificial intelligence (AI). This pull has passed through the die level and even reached the board level, having moved on from the interposer and integrated circuit (IC) substrate levels. These two are key to most of the advanced packages of commercial AI accelerators today and are vital enablers in the AI accelerator roadmap, making advanced packaging at all its levels a central and versatile platform to unlock new levels of performance and speed at an acceptable performance/cost ratio.

It has been left to the advanced packaging players to bring many innovations to real trends, such as heterogeneous integration, chiplet adoption, thermal management, power delivery for the next-generation AI accelerators, and finer pitch, among others. At Yole Group, we think, these trends are squeezing the advanced packaging players' capabilities, capacity, cost-effectiveness, flexibility, and versatility. A part of this enablement comes from the IC substrate level, which is often hidden or ignored when we talk about the value chain. However, the dynamics have changed, and the IC substrate level has become increasingly critical for the semiconductor industry year by year, especially after the 2021 shortage of Ajinomoto build-up film (ABF) for IC substrates. Since then, the perception of IC substrates as a mere mechanical support has changed to that of a platform that can optimize power efficiency, increasing computing power by adding more Si dies and even embedding Si bridges or integrated voltage regulators (IVRs) for power delivery.

Today, the IC substrate industry is being pushed beyond organic Bismaleimide-Triazine (BT) and build-up substrates if it wants to be part of the AI ecosystem for the coming decades. Organic build-up substrates, present in 97% of AI accelerators today, are confronting the warpage wall, which has become the main challenge to retaining dominance as the foundation for the next-generation AI accelerators (**Figure 1**). Their inability to maintain dimensional stability at the large form factors required for xPUs (i.e., a specialized processing unit—where 'x' represents a variable application domain), including graphics processing units (GPUs), neural processing units (NPUs), data processing units (DPUs), and central processing units (CPUs), has catalyzed the industry to analyze alternatives, such as increasing the organic core thickness, adopting a multicore approach, or even changing the core material itself. Currently, this mainly comprises a resin and fiberglass sandwich between two copper-clad laminate (CCL) layers. Glass has emerged as the candidate material with the highest potential, though it also carries significant risks if not handled correctly. Its use was largely initiated by Intel in September 2023 and Absolics' fab investment in November 2022.

Compared to organic cores, glass offers a unique combination of extreme flatness, a wide range of coefficient of thermal expansion (CTE) values (depending on the glass composition), and dielectric properties closer to those of silicon (which reduces the CTE mismatch). However, its adoption for high-volume manufacturing (HVM) is filled with challenges related to fragility, manufacturing complexity, and strong market demand. Multiple pilot lines have been set up between 2022 and 2026 from existing IC substrates players and newcomers looking to book their spot in the emerging glass core substrates (GCS) ecosystem, as deeply analyzed in Yole Group's dedicated report, "Status of the Advanced IC Substrates Industry 2025," (2026 edition coming soon) [1].

GCS represents an opportunity for existing manufacturers to increase their market share and ranking, while it is an entry opportunity for other business models to be part of the value chain, including raw glass suppliers, laser drillers for through-glass vias (TGVs), TGV inspection tools, display markets,

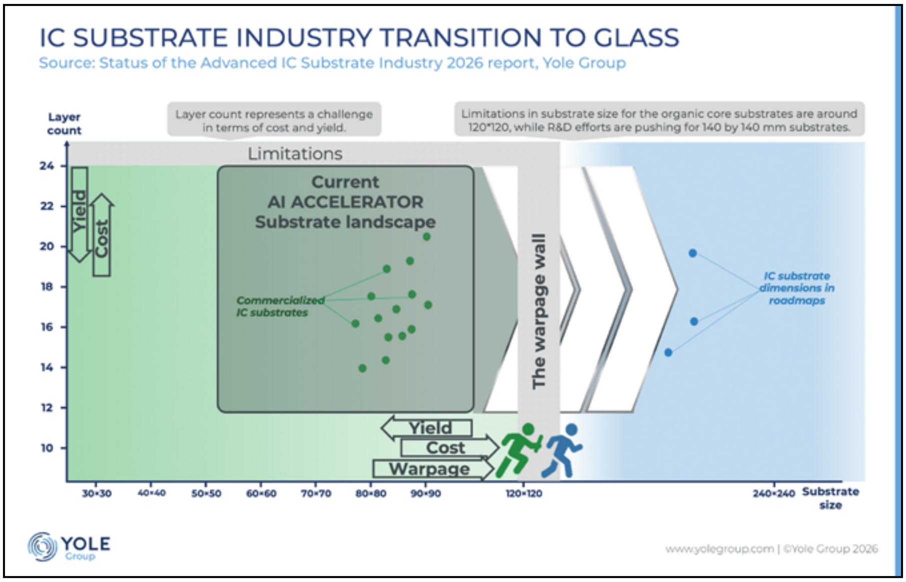

IC SUBSTRATE INDUSTRY TRANSITION TO GLASS

Source: Status of the Advanced IC Substrate Industry 2026 report, Yole Group

Figure 1: IC substrate industry transition to glass. SOURCE: [1]

etc. Achieving the required scalability for the HVM phase requires a versatile process that integrates substrate edge warp and reliability error (SeWaRe)-free singulation, sub-micron inspection, and inline metrology, at high yield, to navigate the landscape of TGV formation, metallization, and inspection, in addition to large-format panel processing (510x515mm^2).

The industry's shift to glass is expected to result in at least two co-existing solutions to scale alongside the chiplet revolution. This is based on the expected failure or slowness of organic polymers to meet AI demand efficiently and quickly at an affordable cost. As AI GPUs grow from standard sizes of 50–70mm range to 100mm x 100mm and beyond, the physical limitations of organic cores become challenging, mainly due to the warpage and reduced yield induced by the CTE mismatch between organic materials and the silicon dies they carry. When a package cycles through temperatures exceeding 250°C during reflow, the IC substrate and the die expand at different rates, exerting mechanical stress that can snap microscopic solder bumps or cause the circuit to delaminate.

Glass provides a sophisticated solution to this thermo-mechanical crisis. Its CTE can be precisely engineered within the range of 3.0ppm/°C to 8.0ppm/°C, aligning it much closer to silicon's 2.6ppm/°C than organic materials. This alignment reduces pattern distortion by 50% to 70% and minimizes warpage by up to 80% compared to silicon-based interposers. Furthermore, glass features atomic-level flatness, often measured at 1.0µm across large areas, which is 5-10 times better than the 5–10µm deviation typical of organic cores. This flatness is not merely a mechanical advantage; it is vital for the creation of very advanced packages that integrate not only logic dies but also high-bandwidth memory (HBM) and optical engines into a unified system. It is a critical requirement for the ultra-fine lithography needed to pattern sub-2µm lines and spaces (L/S), enabling the high-density interconnects for next AI and HPC architectures.

The mechanical rigidity of glass also permits both a reduction and an increase in package thickness depending on the application. Standard substrates are being challenged by the need for low-profile packages, while the emerging co-packaged optics (CPO) packages will push for thicker cores in order to include embedded waveguides. Glass core substrates can be roughly half the thickness of traditional alternatives while providing the necessary mechanical strength during testing without fatigue. However, this rigidity comes with the trade-off of high brittleness, making the glass susceptible to thermal shock during the reflow process. In addition, thinner glass opens the door to glass emerging at the interposer level, with glass interposers a cost-effective performant

solution. Manufacturers must optimize their heating and cooling ramps to alleviate CTE-mismatch-induced thermal stress at the copper-glass interface.

The economic drivers are equally compelling. While silicon interposers offer superb interconnect density, they are limited by the reticle size of lithography tools and the high cost of wafer-scale processing. Glass, by contrast, leverages the panel format of the display industry, enabling processing on large panels, yet another application of glass in interposers. According to Yole Group's report, this shift to panels increases area utilization from approximately 55% to over 80%, allowing manufacturers to achieve lower per-unit costs once HVM is reached.

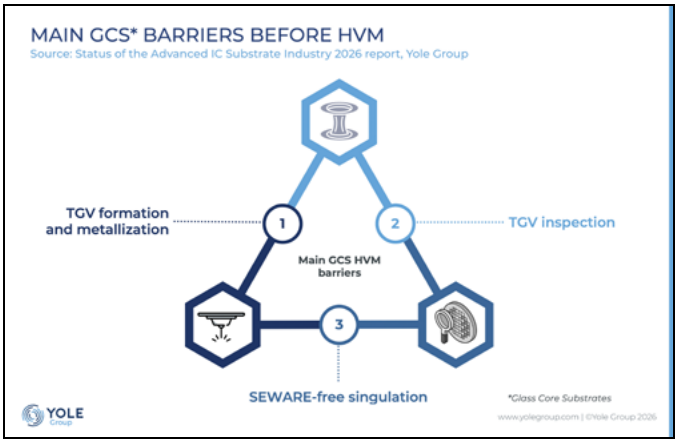

Figure 2: Main GCS* barriers before HVM. SOURCE: [1]

The main triad to resolve before HVM

GCS has successfully attracted many players. More than 12 pilot lines have been set up, but three main process steps need to be at the HVM level of yield before moving to the HVM phase: TGV formation and metallization, singulation, and inspection (**Figure 2**).

Through-glass via process: The anchor process

The formation of TGVs is the defining unit process for GCS, analogous to the role of through-silicon vias (TSVs) in Si interposer technology. Two principal approaches compete for HVM adoption, each with distinct process control implications. Direct laser ablation, typically using femtosecond or picosecond pulses, can achieve high aspect ratios (ARs) but tends to introduce microcracks that require extensive post-process cleaning.

The second approach—laser exposure and selective wet chemical etching or laser-induced deep etching (LIDE) by LPKF—has emerged as the most promising route for advanced TGV fabrication. It operates as a two-step process. A laser first modifies the internal structure of the glass volumetrically, and the modified zones are then selectively removed by wet chemical etching, producing smooth, crack-free via structures. This decoupling of energy deposition from material removal is critical to achieve the high selectivity ratios.

The practical design space for TGV aspect ratio (AR) ranges from approximately 1:4 for less aggressive applications, to 1:20 for the most demanding high-density interconnect architectures. The resulting via geometry is often hourglass-shaped, with the taper angle representing one of the most critical process outputs. Taper angles in the range of 1.8° to 5° are targeted depending on downstream physical vapor deposition (PVD) Cu seed layer deposition requirements. Too narrow a taper and the seed layer conformality degrades catastrophically, too wide and the via pitch density benefit is lost. Critically, the taper angle is coupled to the etch rate; higher etch rates generally produce larger tapers, while lower rates yield steeper sidewalls, creating a direct tension between throughput and geometric precision that the process control system must navigate continuously. If TGV drilling defines the latent via geometry, it is the wet chemical etching step that realizes it, and it is here that HVM scalability faces its most immediate process control challenges. In a production environment, the key considerations are the choice of chemistry to achieve the highest possible selectivity and etch rate on glass compositions, optimize fluid dynamics for uniform etching across the full panel area, and control temperature and flow precisely to maintain process capability over extended production campaigns.

Once TGVs are formed and etched, metallization must establish electrical connectivity through structures whose geometry, governed by the upstream laser and etch processes, imposes hard constraints on the selection of deposition techniques. Physical vapor deposition (PVD) for barrier and seed layers provides adequate step coverage only up to aspect ratios of approximately 1:5; beyond this threshold, the conformality of sputtered films degrades to the point where the seed layer becomes discontinuous at the via waist, creating voids during subsequent electroplating that become latent reliability defects.

For higher ARs, particularly in glass thicknesses approaching 1.1–2mm, where TGV metallization becomes exceptionally challenging, alternative deposition strategies are required. Electroless copper plating offers conformal coverage, but demands a high-quality activation layer. Atomic layer deposition (ALD) achieves excellent step coverage, but carries a cost penalty that may be prohibitive for all but the highest-value applications. Optimized electroplating, potentially employing pulse-plating waveforms and advanced electrolyte chemistries, offers a cost-effective middle ground, but bath conditions must be tightly controlled to ensure uniform metal fills throughout the full via depth.

The interdependence between via geometry and metallization highlights a critical strategic fact: Design-process co-optimization must begin at the earliest stage of package development, not after the TGV process is nominally defined. Via dimensions, taper angles, aspect ratios, and any additional features such as blind vias or embedded cavities must be specified with the explicit knowledge of their impact on metallization yield, cycle time, and cost.

The inspection challenge is unique to glass

TGV inspection is the critical technical bottleneck in GCS manufacturing because it requires the simultaneous verification of thousands of micron-scale structures across large transparent panels at HVM speeds. Unlike traditional 2D surface inspection used for organic substrates, TGV quality is primarily defined by Z-axis parameters and internal integrity, which are complicated by the material's transparency and high reflectivity.

Conventional automated optical inspections (AOI) systems often face a compromise between resolution and throughput when dealing with the complex 3D structure of vias. To address this, the industry is adopting specialized hardware. Process control requires 100% inspection to ensure all vias are open and within tolerance because a single missing or malformed via can cause catastrophic package failure. Detecting internal voids and copper fill integrity is critical for yield optimization, particularly as TGV aspect ratios increase to 20:1 and beyond. Many startups have emerged aiming to solve this problem.

Possible SeWaRe-free GCS after singulation?

SeWaRe error, defined by [2], is a reliability failure mode that originates at the free edge of GCS after singulation, which leaves micro-defects or cracks. The build-up process drives tensile opening stresses at the exposed glass edge. Under thermal cycling or with moisture, small edge flaws can propagate as cracks, causing latent yield and reliability issues. To reach HVM, there must be SeWaRe-free substrates after singulation, which makes the singulation step even more critical than TGV formation. Mechanical and laser dicing are being tested, but they need to be faster to ensure a high-yield, high-throughput process.

Summary

Glass core substrates are no longer a speculative technology. The convergence of TGV formation, advanced wet etch process control, and metallization strategies, from electroless deposition and ALD to pulse electroplating, has created a credible manufacturing pathway. The remaining gap is manufacturing maturity, the ability to sustain material advantages across millions of units with predictable yield, cost, and reliability [1].

Biography

Bilal Hachemi is a Senior Technology & Market Analyst, Semiconductor Packaging, at Yole Group, France. He contributes on a day-to-day basis to the analysis of packaging technologies, their related materials and manufacturing processes. Previously, Bilal carried out experimental research in the field of nanoelectronics and nanotechnologies, focusing on emerging dielectric materials and their ferroelectric applications. He (co-) authored several papers in high-impact scientific journals and participated in several international conferences. Bilal obtained a PhD in nanoelectronics in 2022 from the Grenoble Alpes University (France) and he studied at IAE Grenoble for a Management Master degree. Email Bilal.hachemi@yolegroup.com

Reference

1. Status of the Advanced IC Substrate Industry 2026. https://www.yolegroup.com/product/report/status-of-the-advanced-ic-substrate-industry-2026/?utm_source=Article&utm_medium=email&utm_campaign=%20Status-of-the-Advanced-IC-Substrate-Industry_2026
2. F. Weil, "Empirical investigations on die edge defects reductions in die singulation processes for glass-panel-based interposers for advanced packaging," Proc. IEEE Electronic Components and Technology Conference (ECTC 65th), p. 1991, 2015.


X-RAY ON
X-RAY ON
X-RAY ON
X-RAY ON
XM8000-7 Pro
Nordson
Test & Inspection


Nordson
TEST & INSPECTION

Enabling next-generation advanced integration and rapid innovation for photonics, quantum, and beyond with multi-column e-beam lithography

By Kenneth P. MacWilliams, Regina Freed, Ted Prescop, Andre Linden, Shani Williams [Multibeam Corporation]

The semiconductor industry is in the midst of a profound packaging inflection, driven by artificial intelligence (AI) accelerators, advanced graphics processing units (GPUs), high-performance central processing units (CPUs), and emerging quantum and automotive compute platforms. Technologies such as 2.5D interposers, 3D stacking, hybrid bonding, chiplet-based architectures, flip-chip integration, and advanced redistribution layers (RDL) are enabling unprecedented levels of heterogeneous integration. Companies like NVIDIA, AMD, and Google are pushing the limits of chip-to-chip bandwidth and power efficiency to meet the demands of AI training clusters, edge inference, data centers, and next-generation consumer devices. At the same time, growth in electric vehicles, high-performance computing, and quantum systems is amplifying the need for packages that deliver higher interconnect density, lower latency, and dramatically improved power-per-bit metrics.

Today's advanced packaging approaches, however, face significant limitations. Silicon interposers and embedded bridge die solutions add cost, complexity, and yield risk. Bridge-die approaches—such as embedded silicon bridges—also introduce alignment challenges, yield sensitivities, and additional process steps. Hybrid bonding and fine-pitch micro-bumping require extremely tight overlay accuracy and defect control. As a result, novel integration schemes are at the forefront of the advanced packaging world.

With an inherent mission to enable continuous innovation across a range of evolving semiconductor technologies, our platform continues to evolve. The second-generation platform was designed to address the demands of a rapidly advancing industry—a system thoughtfully designed and enhanced with the productivity and flexibility to meet the demands of a rapidly advancing industry. This platform's versatility is evident in its ability to process a wide range of substrate sizes and materials, making it suitable for emerging applications such as quantum technologies and integrated photonics, as well as for co-packaged optics in advanced data centers. This adaptability allows for rapid innovation on the chip and supports two critical transitions—scaling small batch manufacturing from concept to volume production, and the transition from small to larger substrates for special applications.

The technology continues to prove its value in advanced integration. By facilitating higher chip-to-chip bandwidth and improved power efficiency, it addresses the growing demand for high-bandwidth, low-latency, and low-power solutions in next-generation systems. Further, the technology also reduces production cycle times, enhances yield, and lowers the overall cost of ownership for semiconductor fabs. As the industry moves toward sub-chiplet architectures and backside power delivery—and as interconnect pitches shrink and routing density increases—lithographic constraints in the backend become more pronounced, and challenges to performance, cost, and cycle time associated with traditional lithography methods are heightened.

High-productivity direct-write electron beam lithography (EBL) is emerging as a critical enabler for the next generation of heterogeneous integration. By combining high-productivity multicolumn architectures with adaptable beam sizing, EBL delivers high resolution and nanometer-scale precision across a wide field of view and large depth of focus—capabilities that are essential for advanced RDL, fine-pitch interposers, and 3D interconnect layers. Unlike laser-based systems, e-beam patterning is not diffraction-limited, enabling significantly finer features and tighter overlay control. The result is improved yield, reduced defectivity, and greater process flexibility, particularly for 2.5D integration schemes where cost efficiency and throughput are paramount.

Multi-column direct write enables approximately a 10× reduction in chip-to-chip latency, translating to ~40× improvements in bandwidth density and ~100× improvements in power per bit. This allows chip-to-chip interconnects to approach on-chip interconnect performance levels—ushering in a new class of packaging often referred to as advanced integration (**Figure 1**).

The gains noted above are particularly relevant for AI accelerators and data center GPUs, where memory bandwidth and energy efficiency directly impact system-level performance and operating cost. By enabling finer-pitch interconnects and more efficient routing, EBL helps accelerate roadmaps for flip chip, hybrid bonding, and 3D stacking, while also supporting emerging approaches such as backside power delivery.

Multi-column EBL can offer a higher-yielding alternative to embedded bridge die architectures. By enabling direct, fine-pitch interconnects across larger substrates without the need for intermediate silicon bridges, manufacturers can reduce assembly complexity, eliminate additional die handling steps, and improve overall yield. The technology's adaptability also supports curvilinear routing and customized interconnect topologies, which are increasingly required for chiplet-based designs.

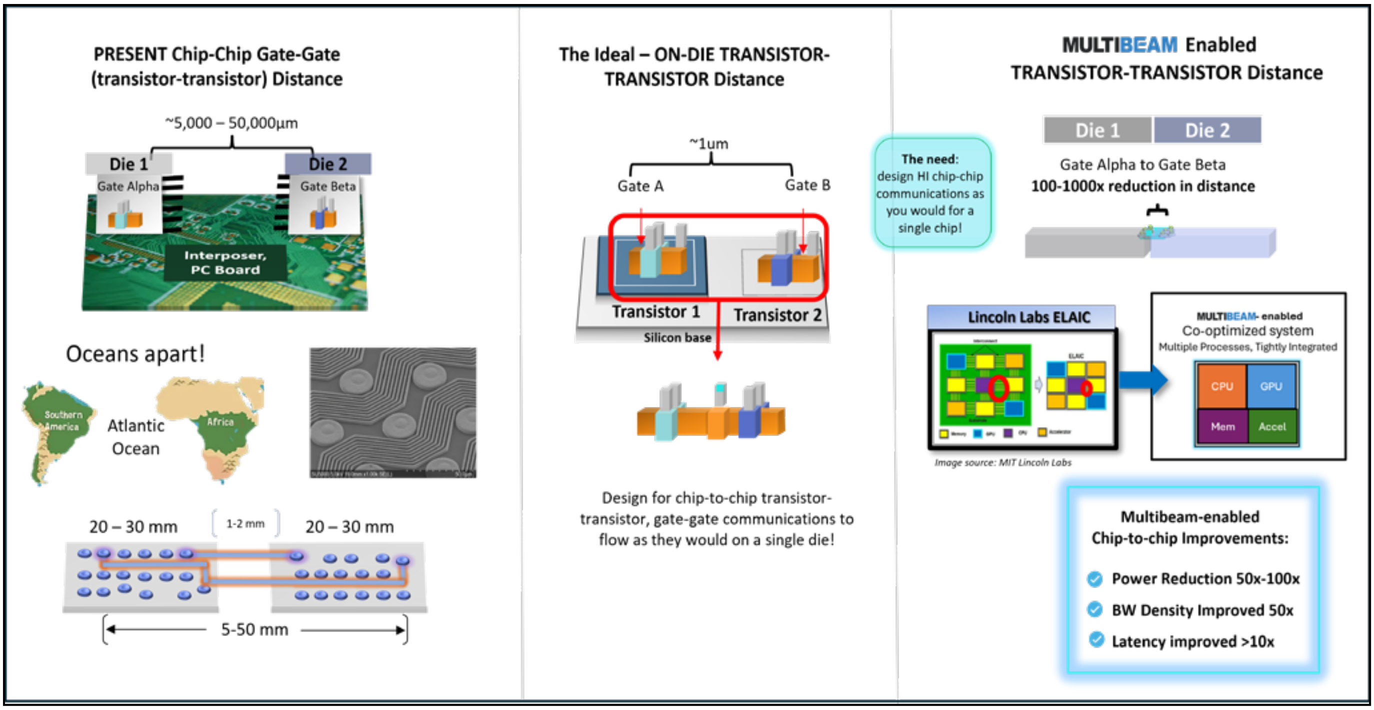

Figure 1: Chip-chip bandwidth, latency and power improvements with multi-beam-enabled advanced integration.

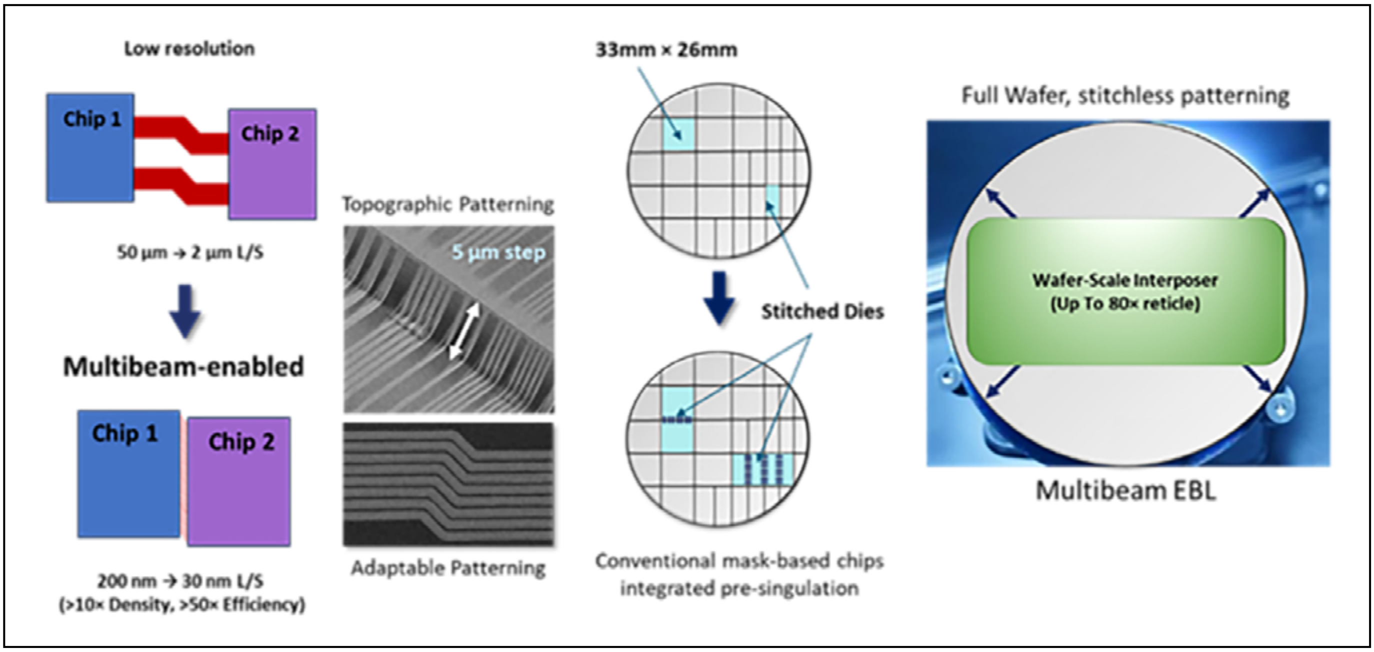

Figure 2: Enabling adaptable patterning, high-density interconnects, patterning over topography, and stitch-free wafer-scale interposers for photonics packaging.

As AI, quantum computing, data center expansion, and electric vehicle (EV) electrification continue to accelerate, the demand for higher performance, lower power, and more tightly-integrated semiconductor systems will only intensify. Multibeam direct-write e-beam lithography is proving to be a foundational technology in this transition—enabling not just incremental improvements in 2.5D packaging, but paving the way toward scalable wafer-scale 3D integration and beyond. By driving cost efficiency, yield enhancement, and unprecedented interconnect performance, it is redefining what is possible in advanced packaging and helping the industry meet the compute demands of tomorrow.

As advanced packaging extends beyond traditional logic and memory integration, quantum and photonics applications are introducing entirely new patterning demands in the backend. These emerging markets are not only pushing for higher

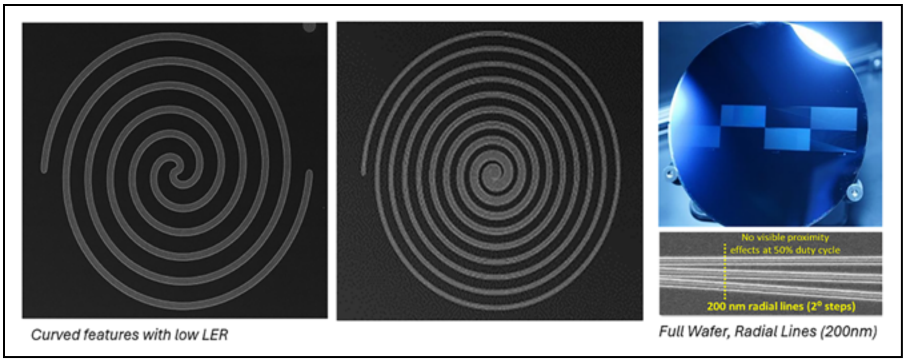

Figure 3: Multibeam-demonstrated patterns: standard and tapered spiral with low line edge roughness (LER), and full-wafer, 200nm radial lines.

interconnect density and lower power—they also require fundamentally different geometries, materials, and integration schemes that strain the limits of conventional lithography.

Photonics advanced packaging: Enabling precision, flexibility, and scale

Silicon photonics and co-packaged optics are becoming critical for AI clusters, data centers, and high-speed networking as traditional electrical I/O approaches its physical limits. The transition to optical interfaces requires tight control of waveguides, gratings, couplers, and micro-optical structures—many of which are inherently curvilinear and very sensitive to line edge roughness (LER) and dimensional variation (**Figure 2**).

Photonics packaging presents several challenges that are difficult for conventional lithography methods to overcome. Some of which include fabricating precision gratings and curvilinear patterns at nanometer scale, maintaining uniform critical dimensions across large areas, handling fluctuating topography and diverse materials (Si, SiN, InP, polymers, glass, etc.), and scaling from R&D prototyping to high-volume manufacturing without redesigning the process.

Our second-generation flexible e-beam lithography platform directly addresses the constraints listed above. Because it is not diffraction-limited like laser direct imaging (LDI), it can generate ultra-fine, smoothly-curved geometries required for low-loss waveguides and grating couplers. The ability to natively pattern curvilinear features without fracturing into Manhattan approximations helps to improve optical performance and reduce scattering (**Figure 3**). Furthermore, our second-generation system enables: 1) Pattern flexibility for custom photonic layouts and rapid design updates; 2) Support for multiple substrate sizes and materials, enabling heterogeneous photonic integration; 3) Large depth of focus, accommodating warped or non-uniform substrates common in advanced packaging flows, and eliminating the need for expensive chemical mechanical polishing (CMP) processing; and 4) Wide field of view enabling package sizes that exceed traditional reticle limits.

Flexibility is particularly important because co-packaged optics architectures continue to evolve at a rapid pace. Instead of waiting for mask cycles or being constrained by optical lithography limits, photonics developers can iterate quickly while maintaining production-grade precision.

Quantum advanced packaging: Precision meets iteration

Quantum computing platforms—whether based on superconducting qubits, trapped ions, spin qubits, or photonic qubits—impose some of the most stringent patterning requirements in the ecosystem. Feature geometries are often complex, non-Manhattan style, and highly sensitive to nanoscale variation. Yield is especially critical because even minor defects can degrade coherence times or qubit fidelity.

Many of the patterning demands that come with integrating quantum systems present challenges to conventional lithography. Some of these requirements include nanoscale precision and uniformity for Josephson junctions and resonator structures, intricate feature generation for qubit arrays, and perhaps most critically—rapid iteration cycles because architectures evolve and the ability to transition from lab-scale fabrication to scalable manufacturing.

High-productivity multi-beam e-beam lithography is uniquely positioned to bridge R&D and manufacturing gap for quantum systems. Unlike single-beam research tools that sacrifice throughput, multi-column architectures maintain nanoscale resolution while enabling higher wafer-per-hour productivity—critical for scaling quantum device fabrication beyond experimental volumes (**Figure 4**).

When it comes to integrating quantum systems, our platform enables: 1) Rapid prototyping and iteration without mask dependencies; 2) Ability to pattern highly-intricate, curved, and non-Manhattan geometries; and 3) Scalable manufacturing capability to move from small arrays to large qubit counts.

As quantum companies push toward fault-tolerant systems requiring thousands—or millions—of qubits, reproducibility and backend integration will become central bottlenecks. Multi-beam lithography enables two critical manufacturing and innovation advantages: 1) Tighter control of critical dimensions improving device uniformity; and 2) A credible path from innovative quantum architecture concepts to optimized quantum systems that can be manufactured, replicated, and deployed at scale, thereby accelerating roadmap timelines.

Beyond optical and laser direct imaging limitations

Laser direct imaging and other optical direct-write techniques face inherent diffraction limits that restrict resolution and pattern fidelity, particularly for sub-micron redistribution layers and fine optical structures. Depth-of-focus limitations can degrade yield on non-planar substrates common in heterogeneous integration. Additionally, complex curved geometries often require data fracturing that increases write time and introduces variability.

In contrast, multi-beam e-beam lithography offers: 1) Significantly higher resolution without optical diffraction constraints; 2) Greater depth-of-focus tolerance; 3) More precise control of critical dimensions and line edge roughness (LER); and 4) Adaptable beam sizing to optimize

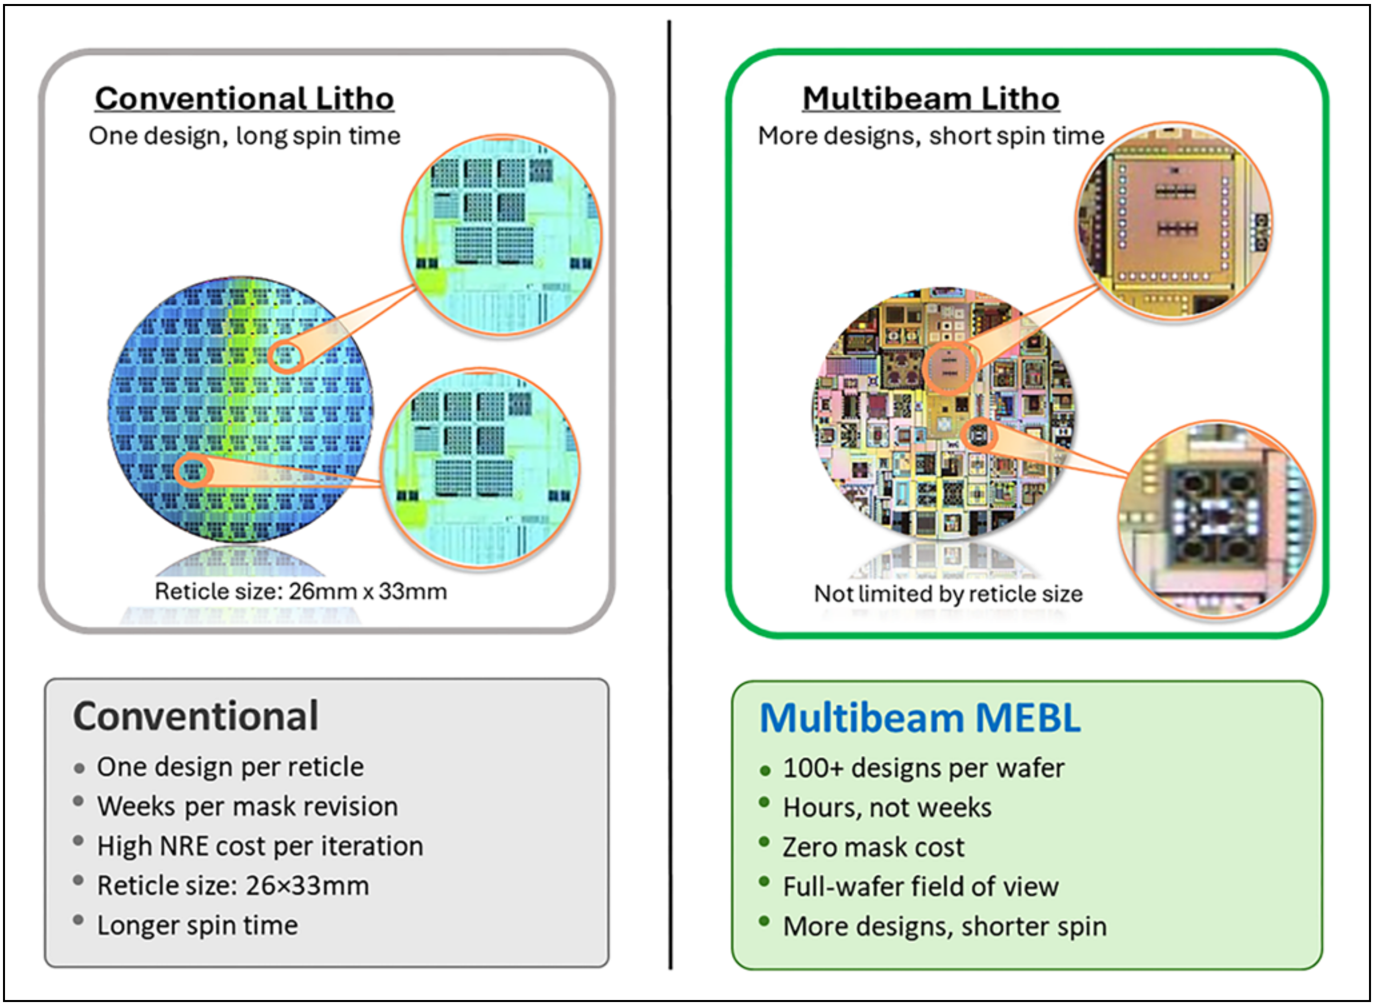

Conventional

- One design per reticle
- Weeks per mask revision
- High NRE cost per iteration
- Reticle size: 26×33mm
- Longer spin time

Multibeam MEBL

- 100+ designs per wafer
- Hours, not weeks
- Zero mask cost
- Full-wafer field of view
- More designs, shorter spin

Figure 4: Multibeam maskless direct-write lithography compared with conventional optical-based lithography. The technology eliminates the cost and time burden and enables rapid design iteration.

throughput for varying feature scales. This combination allows direct-write e-beam to deliver both performance and manufacturability—an increasingly rare pairing in advanced packaging.

Cost modeling

When it comes to fabricating and integrating advanced systems, it is not solely a matter of technical optimization. Controlling cost is critical—and the lithography process is one of the biggest contributors to overall manufacturing cost. However, lithography brings value by enabling scaling, and scaling drives reductions in cost per chip-chip interconnection. In other words, "Moore's Law" has evolved from a metric of cost-per-transistor on chips in the past, to a new metric of cost-per-interconnection between chips. Advances in lithography equipment drive improvements in both performance

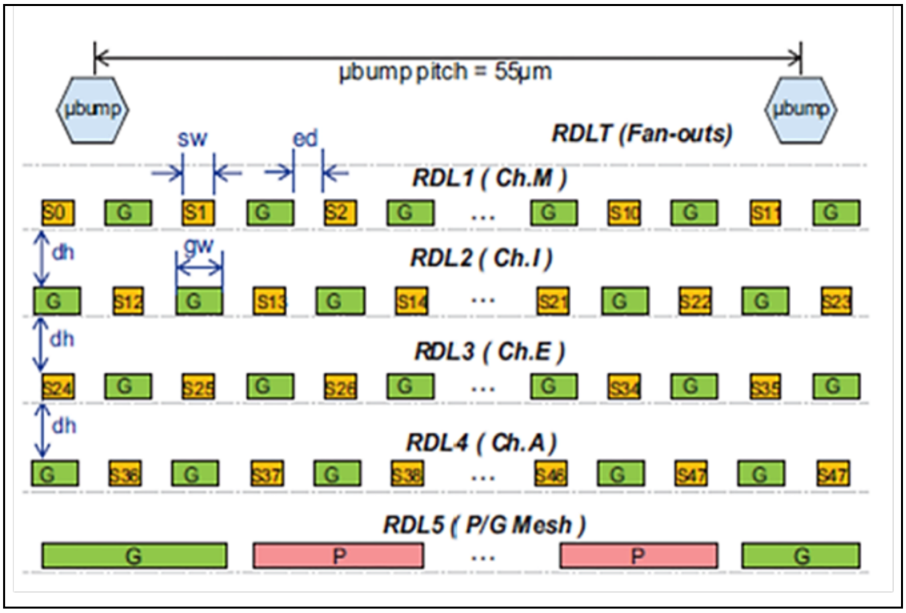

Figure 5: CoWos-R® ASIC to HBM3E with 4x RDL layers: Cross-section view of HBM3 high-speed interconnection of a single slice. SOURCE: [1]

	Cost per Wafer				
	Labor Cost	**Capital**	**Material**	**Total (Direct Cost)**	**Total Cost per Pkg**
4 Dual Damascene RDLs	$7	$189	$101	$298	$22.90
1 e-beam Defined RDL	$3	$201	$29	$233	$17.92
				Cost reduction of 1 e-beam RDL over 4 standard dual damascene RDLs	*-22%*

Table 1: Reducing from 4x RDL layers to 1x e-beam layer results in approximate cost savings of 22%.

	Cost per Wafer				
	Labor Cost	**Capital**	**Material**	**Total (Direct Cost)**	**Total Cost per Pkg**
8 Dual Damascene RDLs	$15	$379	$202	$595	$45.79
1 e-beam Defined RDL	$3	$201	$29	$233	$17.92
				Cost reduction of 1 e-beam RDL over 8 standard dual damascene RDLs	*-61%*

Table 2: Reducing from 8x RDL layers to 1x e-beam layer results in approximate cost savings of 61%.

and cost per interconnection by enabling scaling to smaller feature sizes.

To highlight the cost benefit of scaling, we looked at CoWos-R® application-specific integrated circuit (ASIC) to HBM3E as a test case. Referring to **Figure 5**, we see that with a minimum linewidth of 2µm, a stack of four layers of RDL are needed.

By scaling the linewidth to 400nm, the same bandwidth can be achieved with only one layer of RDL. The single lithography step capital cost is higher, but the overall cost is much lower, resulting in an overall cost savings of over 20% (**Table 1**).

Interconnection density can be increased significantly beyond the above example, further improving cost per interconnection while improving performance. Reducing 8x RDL layers to 1x e-beam layer results in a cost savings of over 60% (**Table 2**).

Additionally, minimizing the number of layers simplifies the overall process flow. With fewer process steps, there are fewer opportunities for defects. This

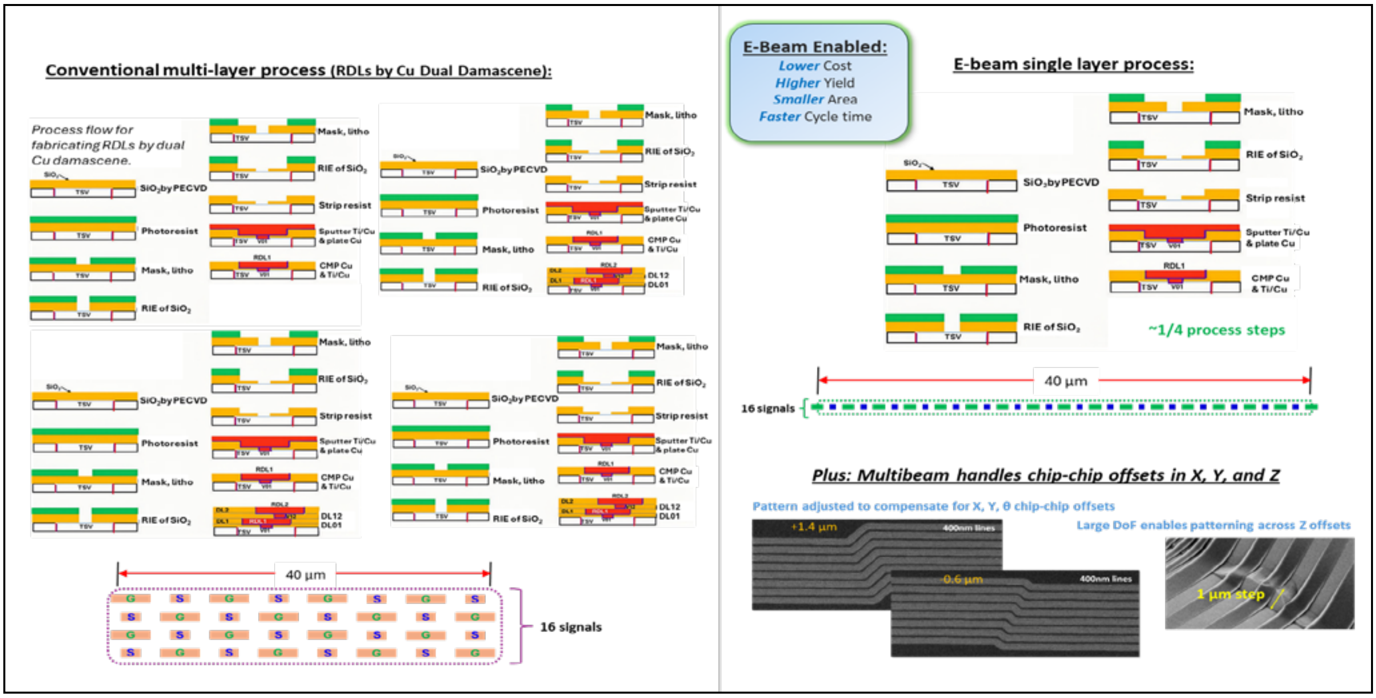

Figure 6: Conventional multi-layer process of producing RDLs by Cu dual damascene compared to E-beam single-layer process: A reduction from 4x RDL layers to 1x e-beam layer.

improves manufacturing yield, and even small yield improvements have a big impact on cost. A summary of the cost, yield, size, and cycle time benefits of scaling are highlighted in Figure 6.

Summary

AI, data center expansion, EV electrification, and quantum innovation are collectively driving an industry-wide shift toward higher performance, lower power, and tighter integration. As chip-to-chip interconnects approach on-chip performance levels—enabled by dramatic reductions in latency, bandwidth density improvements, and power-per-bit gains—the backend becomes as strategically important as the frontend.

For photonics and quantum advanced packaging in particular, our second-generation system is not merely enabling an incremental improvement. It was built to serve as a foundational enabler—supporting precision gratings and curvilinear optical structures, rapid quantum device iteration, nanoscale uniformity, and scalable manufacturing that encourages rapid innovation on the chip.

As heterogeneous integration extends into true 3D architectures and advanced integration paradigms, this technology is poised to accelerate innovation well beyond traditional 2.5D packaging—helping define the next generation of compute infrastructure.

Reference

1. S-F. Yang, W-C. Wang, Y-T. Lin, C-C. Hung, H-Y. Tung, et al., "Signal integrity designs at organic interposer CoWoS-R® for HBM3-9.2Gbps high speed interconnection of 2.5D-IC chiplets integration," 1098-1103. 10.1109/ECTC51529.2024.00176.

Biography

Kenneth P. MacWilliams is President and board member of Multibeam Corporation, headquartered in Santa Clara, CA. He has 10+ years of experience in semiconductor device research and 20 years in process and equipment development and management, helping launch multiple novel equipment platforms and processes at companies like Applied Materials, Novellus Systems (acquired by Lam Research), Veeco, and Yield Engineering Systems (YES). He holds PhD and MS degrees in Electrical Engineering from Stanford University. He is an NSF Fellow with over 100 publications, and more than a dozen patents. Email kmacwilliams@multibeamcorp.com

Co-authors

Regina Freed - Chief Technology Officer | Ted Prescop - Vice President, Technology
Andre Linden - Vice President, Sales | Shani Williams - Sr. Marketing Manager


AI
Hot spot
Multi T_j Function

Silicon photonics applications solutions and technology challenges in advanced co-packaged optics

By Mike Tsai, Steven Lin, Yih-Jenn Jiang, Don-Son Jiang [Siliconware Precision Industries Co., Ltd]

In this article, a new co-packaged optical (CPO) packaging that integrates electronic integrated chips (EICs) and photonic integrated chips (PICs) into an optical engine (OE) and bundles multiple OE modules with a switch chip is introduced. This work compares OE architectures with horizontal and vertical integration, using a fan-out redistribution layer (FO-RDL), 3D through-silicon via (TSV) stacking, and hybrid bonding to combine EIC/PIC. A fiber array unit (FAU) aligns PIC waveguides to couplers via grating or edge coupling. As electrical and thermal performance remain major challenges, the insertion loss and heat sink warpage are analyzed. This article proposes warpage mitigation and assembly strategies for wafer-level and package-level bonding. The results offer advanced CPO and OE solutions for future networking and high-performance computing (HPC) markets.

Introduction

Moore's Law continues to increase transistor density, enabling faster data transmission speeds—for example, doubling bandwidth in switch routers from 51.2–102.4Tbps. This drives demand for silicon photonics, which offers higher speeds, lower power use, and smaller sizes compared to traditional optical systems. Key applications include transceivers, networking, HPC for cloud computing, switch routers and artificial intelligence (AI). Silicon photonics uses light to represent data (0s and 1s), and optical fiber offers longer reach and lower power consumption than copper cables as shown in **Figure 1**.

Optical connectivity evaluation trend

Traditional pluggable modules are easy to use and maintain, but lack bandwidth. It is easy for the signal to lose strength as it goes through the circuit board to the switch IC; pluggable modules also consume a lot of power, and are expensive. To meet the demand for faster data transmission, systems are moving from copper cables (0.04Tbps) to optical pluggable modules (1.6Tbps). Now, there is a trend towards even more advanced CPO solutions that integrate optics directly into the package to achieve over 6.4Tbps per OE, and up to 51.2, or even 102.4Tbps per package.

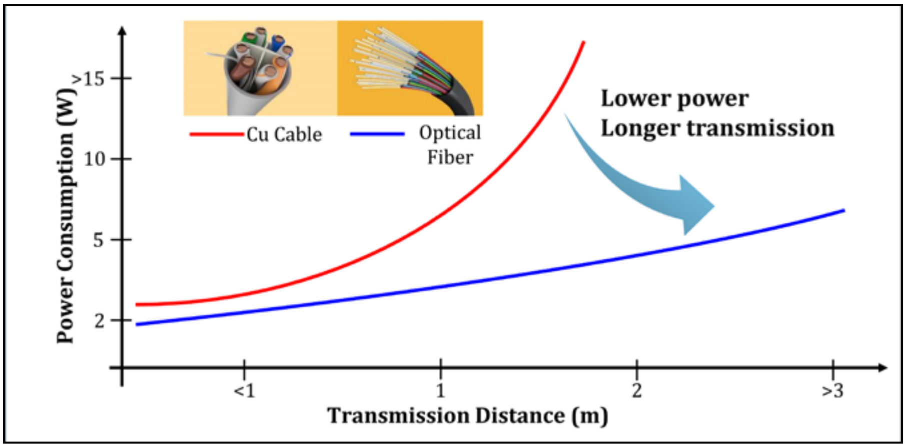

Figure 1: Power consumption between Cu and optical.

The conventional electronic signals lose strength as they travel through the printed circuit board (PCB) to reach the main processor. CPO solves this loss of signal strength by integrating the OE module directly alongside the processor. Because photonic circuits are built to convert electrical signals into optical signals, these circuits—along with EICs and PICs—are combined into a single OE module. CPO uses light to transmit signals, minimize loss and improve speed. Also, CPO places multiple optical modules around the processor, connecting them on the same board, which definitely shortens the signal path, reduces power loss, and improves heat dissipation. Driven by the growing need for AI, CPO is becoming the preferred solution over traditional pluggable optical modules for networking switches. The architecture of a pluggable module and CPO for a networking switch are shown in **Figure 2** and **Figure 3**.

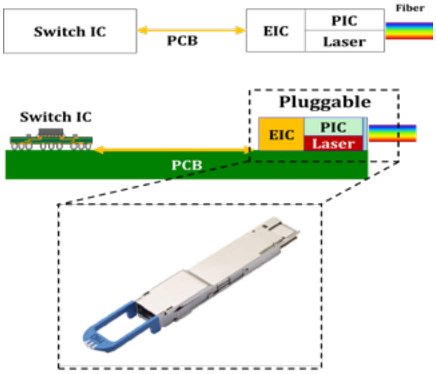

Figure 2: Pluggable module architecture for a networking switch.

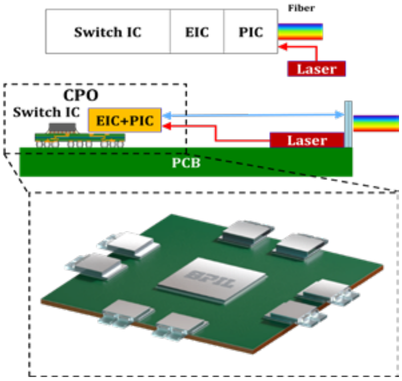

Figure 3: CPO architecture for a networking switch.

Item	Side-by-Side (WB)	Fan-Out	3D IC	Hybrid Bond	Monolithic EPIC
Package	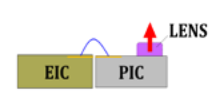 EIC PIC LENS	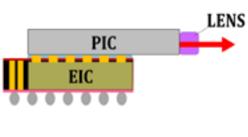 PIC EIC LENS	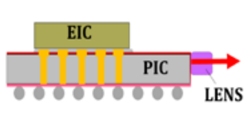 EIC PIC LENS	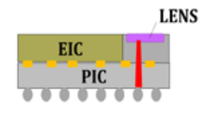 EIC PIC LENS	 EPIC LENS
Coupling Method	Grating Coupling	Edge Coupling	Edge Coupling	Edge Coupling Grating Coupling	Edge Coupling

Figure 4: Various OE structure solutions.

OE module

CPO is far more than photonic-electronic chip co-design; it encompasses advanced packaging, laser integration, optical packaging, structural engineering, and thermal solutions. The OE—as the heart of a CPO system—incorporates electronic EIC and PIC, converts an electrical signal to an optical signal, and vice versa. Currently, several OE structures have been demonstrated in the market. As shown in **Figure 4**, there are various types of OE structures from conventional, to those with high performance in terms of bandwidth speed. The structures shown in **Figure 4** range from: 1) A side-by-side wire bonding package; 2) A fan-out structure using a redistribution layer (RDL) for interconnection; 3) 3D stacking with µbumps and TSV; 4) A hybrid bond solution to approach less cross talk; and 5) In the future, monolithic methods to integrate EIC and PIC into the same silicon chip could be achievable.

The different performance indices are compared for a combination of the EIC and PIC components within an OE in **Table 1**. A key factor is minimizing the distance between these components for better electrical signal transmission. Traditional side-by-side wire bonding causes signal loss. More advanced methods—like using micro-bumps, 3D stacking with TSVs, and hybrid bonding with direct copper to copper bonding—reduce this loss and improve performance.

Optical alignment is crucial for connecting fibers to the optical engine (OE) module and can be done actively or passively. Active alignment precisely positions the optical components (waveguides and fibers) to maximize signal strength, then secures them with glue and tests performance for both edge coupling (EC) and grating coupling (GC). This procedure is challenging because of the tiny size of the optical waveguides. Passive alignment uses a V-shaped

Item	Side-by-Side (WB)	Fan-Out	3D IC	Hybrid Bond	Monolithic EPIC
Width/ Space (um)	22/44	2/2	0.56/0.56	0.4/0.4	0.4/0.4
Routing Density	Poor	Better	Good	Best	Best+
Band width	Poor	Better	Good	Good	Best
Crosstalk	Long	Middle	Middle	Low	TBD

Table 1: Various OE performance comparison results.

Coupling Method	Edge Coupling (EC)		Grating Coupling (GC)
	Passive Alignment	Active Alignment	Active Alignment
Structure	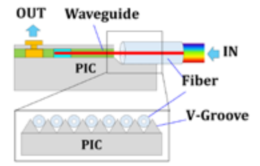	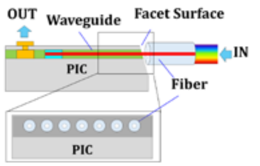	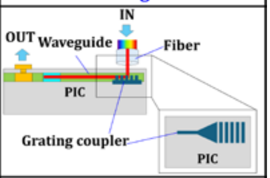
Pros	Easy for fiber alignment process	High coupling performance	Small form factor
Cons	V-Groove Special design & particle free, high accuracy	Need active alignment with high accuracy	Optical conversion efficiency during reflect

Table 2: Optical coupling method comparison.

groove to guide the fiber into place, thereby simplifying the process. Table 2 compares the two alignment methods.

OE module integration approach

This article details the development of an OE module using a fan-out co-packaged optics (FO-CPO) structure using PIC and EIC chips. The PIC is placed on top and vertically connected to the FO module with a special overhang design to improve the fiber connection. The EIC is embedded within the FO module to handle primary signal transmission. A key process challenge is attaching the PIC to the top redistribution layer (RDL) using wafer-form die attachment with an overhang design. We evaluated both mass reflow (MR) and thermal compression bonding (TCB) methods, finding that TCB provides superior solder height control, mitigating non-wetting risks associated with high wafer form FO warpage and resulting in improved solder joint quality, as illustrated in Figure 5.

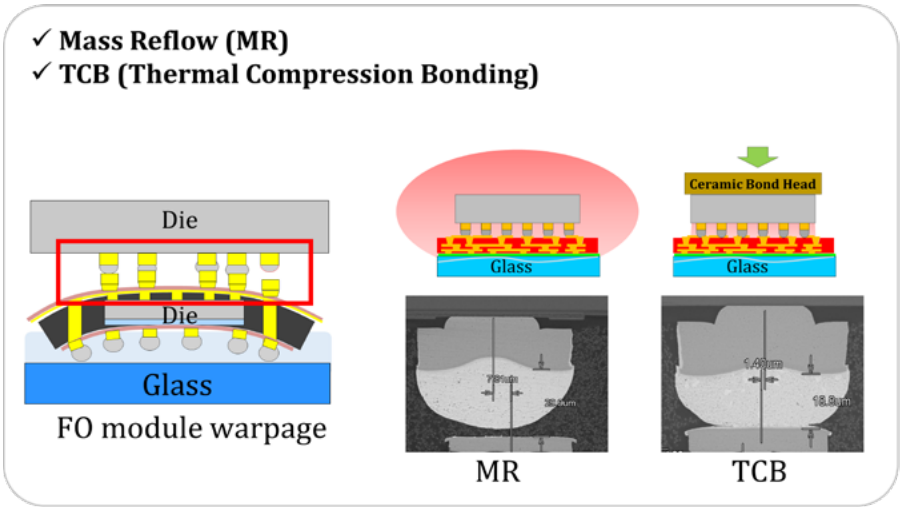

Figure 5: Wafer-level OE module warpage challenge comparison (MR vs. TCB).

The CPO packaging solution directly integrates OE modules and main chip ICs onto the substrate. This can cause warpage issues as the different components are mounted. We utilized simulation tools to understand these warpage behaviors and identified ways to improve the process, considering the layers of the chip, optical engine, and substrate. For larger chips, controlling warpage during the reflow process is critical. We developed solutions to predict and adjust warpage behavior, including “smile, cry and flat” shapes. By proper selection of materials, designing support tools, and optimizing the bonding equipment, one can ensure that the chip module and substrate warpage remain properly aligned, as shown in Figure 6.

A significant challenge with large CPO packages is controlling warpage, which is addressed through heat sink design optimization. We employed simulation tools to evaluate various heat sink shapes under both room temperature (RT) and high temperature (HT) conditions, aiming for stable package warpage while ensuring sufficient space for the OE modules and acceptable package stress. For instance, increasing the adhesive area of the heat sink improves its bonding strength and reduces substrate warpage. Throughout the manufacturing process, we fine-tune the heat sink’s design, shape, and how it dissipates heat via simulation to meet both thermal and warpage specifications. The heat sink simulation design of experiments (DOE) is detailed in Table 3.

The result of DOE led to warpage improvement at RT and HT, allowing OE modules to be placed securely within specifically-shaped cavities (Table 3). Leg4 and Leg5 results indicate that the special shapes can reduce warpage by around 30%. However, the package stress and heat sink shape also influence the optimal number of OE modules that can be integrated into a single CPO package.

Summary

This article explored integrating EIC and PIC components into OE modules using various packaging techniques: wire bonding, fan-out, 3D IC, and hybrid

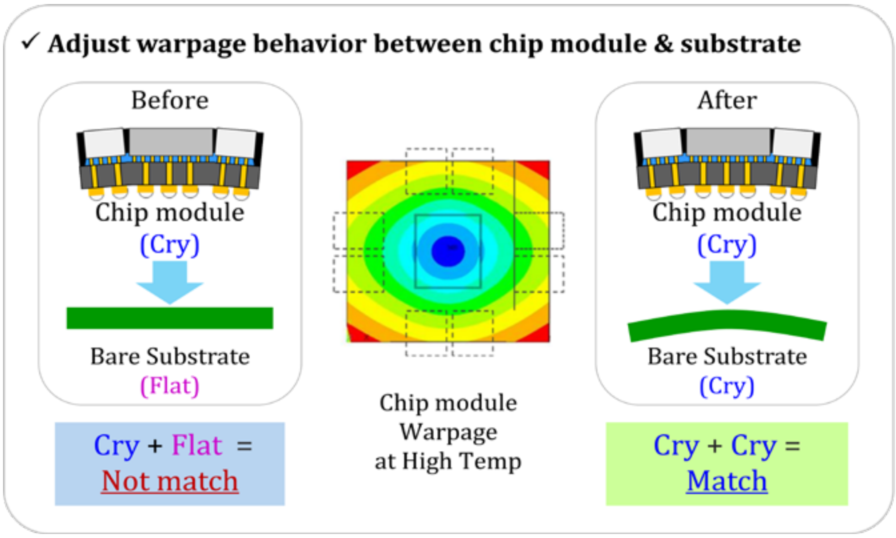

Figure 6: Chip-module-to-substrate warpage control comparison (before vs. after).

Heat Sink DOE				Simulation			OE Module Floorplan (Quantity trend)	
DOE	Purpose/ Factor / Floorplan			Die Stress Ratio	RT Warpage Ratio	HT Warpage Ratio		
Leg1	Without HS			0.60X	-1.00A	+1.00B	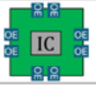	≧8pcs
Leg2	Standard HS			1.00X	-0.63A	+0.70B		No space
Leg3	HS shape study	Shape#1		0.95X	-0.90A	+0.94B	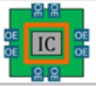	≧8pcs
Leg4		Shape#2		1.08X	-0.68A	+0.72B	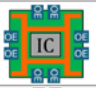	≧8pcs
Leg5		Shape#3		1.10X	-0.67A	+0.71B	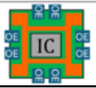	≧8pcs

Table 3: Heat sink (HS) simulation DOE comparison.

bonding. We examined both passive and active alignment methods for FAU attachment to meet assembly and system integration needs. Warpage challenges during wafer handling and large-chip assembly were addressed by evaluating different bonding technologies (mass reflow vs. thermal compression bonding) to optimize bond joint quality. Large-chip warpage was further managed through material selection, tooling design, and substrate choices to ensure good stand-off height (SOH) control during reflow. Heat sink optimization—considering design and material factors—enabled better warpage control and allowed for maximizing the number of OE modules on the CPO floor plan to meet bandwidth requirements. By integrating optical connections into CPO packages, silicon photonics overcomes the limitations of traditional electronic transmission, promising faster data speeds to support the growing demands of AI data centers and HPC.

References

1. S. Razdan, "Advanced 2.5D and 3D packaging technologies for next generation silicon photonics in high performance networking applications," 72nd Electronic Components and Tech. Conf. (ECTC), 2022.
2. J. Li, T. Hsu, M-H. Zhuan, S. Lin, T. Shih, N. Kao, Y-P. Wang, "Innovative fan-out embedded bridge structure for co-packaged optics," Proc. 73rd ECTC, 2023.
3. C-H. Lu, "High bandwidth and energy efficient electrical-optical system integration using COUPE technology," 74th ECTC, 2024.
4. H. Hsia, "Heterogeneous integration of a compact universal photonic engine for silicon photonics applications in HPC," 71st ECTC, 2021.
5. K. Muth, "High density integration technologies for SiPh based optical I/Os," 73rd ECTC, 2023.
6. A. Prasad, "Advanced 3D packaging of 3.2Tbs optical engine for co-packaged optics (CPO) in hyperscale data center networks," 74th ECTC, 2024, pp. 101-106.

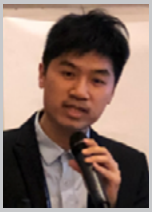

Biographies

Wen-Jung (Mike) Tsai is a Department Manager of Corporate R&D at Siliconware Precision Industries Co., Ltd., Taichung, Taiwan, R.O.C. He joined SPIL in 2011 and is in charge of new assembly package and new technology development, with a special focus on Silicon Photonics with CPO platform solutions for advanced assembly technology. He received a Masters Degree in Mechanical Engineering, National Chung Hsing University, Taiwan. Email: miketsai@SPIL.com.tw

Jui-Cheng (Steven) Lin is a Technical Manager in the Corporate R&D Center at Siliconware Precision Industries Co., Ltd., Taichung, Taiwan, R.O.C. He has over twenty-five years of job experience in the semiconductor industry. He received his MS in Electrical Engineering from the University of Southern California.

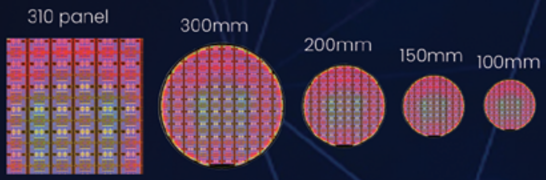
310 panel
300mm
200mm
150mm
100mm

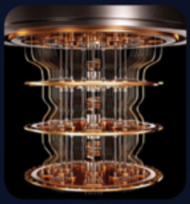


MULTIBEAM
MULTIBEAM
MULTIBEAM
MULTIBEAM

Production scaling of silicon photonics wafer testing using automated test equipment platforms

By HsuHao (Andy) Chang [Marvell]

The explosive growth of artificial intelligence (AI) infrastructure is driving an unprecedented demand for high-bandwidth optical interconnect technologies within hyperscale data centers [1]. As AI clusters scale to encompass tens of thousands of graphics processing units (GPUs) and specialized accelerators, the network bandwidth required to prevent data bottlenecks is increasing exponentially. Industry roadmaps indicate a rapid transition from 400G optical interconnects toward 800G, 1.6T, and beyond in the immediate future [2].

To enable the required aggressive scaling, optical engines are evolving toward significantly higher levels of integration. Silicon photonics (SiPho) has emerged as a critical enabling technology by integrating active and passive optical components—including modulators, waveguides, and photodetectors—directly onto silicon substrates [3]. Compared with traditional discrete optical assemblies, silicon photonics provides vastly higher channel density, superior power efficiency, and improved manufacturability at scale.

However, the increased integration of photonic integrated circuits (PICs) introduces formidable new manufacturing challenges. Modern PIC devices combine complex optical and electrical subsystems that must operate concurrently within tightly-controlled parametric limits. Consequently, ensuring device quality early in the manufacturing flow is no longer optional—it is an absolute economic necessity for maintaining yield and controlling high-volume production costs.

The economics of "shift-left" yield in optical manufacturing

A fundamental strategy for managing yield in complex semiconductor systems is shifting defect detection as early in the manufacturing flow as possible—a concept commonly referred to as "shift-left" yield (**Figure 1**). In optical module manufacturing, value and cost are added across three distinct stages of integration:

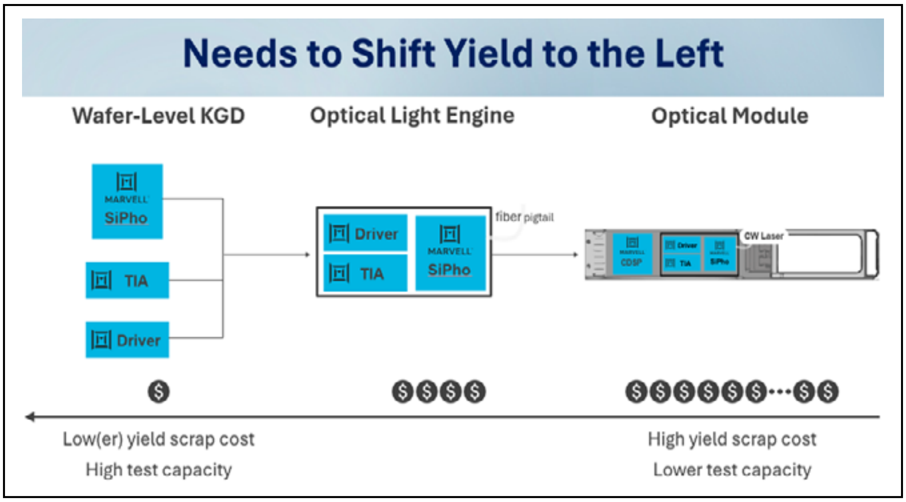

Figure 1: Optical module manufacturing flow highlighting the compounding cost of defects and the strategic placement of shift-left wafer-level screening. SOURCE: Marvell

- **Level 1 (wafer-level testing).** Individual photonic dies and electronic components are tested in wafer form.
- **Level 2 (optical engine assembly).** PICs are heterogeneously integrated or assembled with drivers, transimpedance amplifiers (TIAs), and optical fiber interfaces to form the optical engine.
- **Level 3 (optical module integration).** The optical engine is integrated with digital signal processors (DSPs), control electronics, power delivery networks, and printed circuit boards to create the final pluggable or co-packaged optical module.

As devices move through the stages listed above, the financial penalty associated with yield loss compounds drastically. A defect that escapes Level 1 testing might be a loss of a few dollars. However, if that same failure is only detected at Level 3, the manufacturer is forced to scrap a highly complex assembly containing expensive DSPs, advanced packaging substrates, and integrated micro-optics.

Wafer-level testing acts as the critical gatekeeper, enabling the early screening of defective devices to ensure that only known-good die (KGD) proceed to packaging. Furthermore, rigorous wafer-level parametric measurements provide rapid, high-fidelity feedback to the wafer fabrication facility. Process shifts or excursions can be detected immediately, accelerating yield learning cycles and optimizing the total cost of test.

The complexities of SiPho wafer probing

While wafer-level testing is a highly-mature discipline in conventional complementary metal-oxide-semiconductor (CMOS) manufacturing, SiPho introduces unique opto-mechanical hurdles. Unlike purely electrical ICs, SiPho testing requires simultaneous electrical biasing and optical signal measurement. Optical signals must be precisely coupled into and out of the device under test (DUT) through surface grating couplers or edge couplers, requiring sub-micron alignment tolerances between the wafer and the probe interface. Critical wafer-level parametric measurements include: 1) Mach-Zehnder modulator half-wave voltage (Vpi); 2) Optical insertion loss (chip loss); 3) Extinction ratio (ER); 4) Heater resistance and phase tuning characteristics; and 5) Photodetector dark current and responsivity.

Historically, SiPho engineering labs have used a range of optical test configurations, including manual alignment stages, as well as automated motorized or hexapod-based alignment systems, together with discrete optical instrumentation such as PXI-based chassis and custom characterization benches. These environments are highly effective for new product introduction (NPI), device bring-up, and detailed characterization, but they are not inherently optimized for the scale, test parallelism, and infrastructure efficiency required for high-volume manufacturing. In our previous work, we demonstrated that silicon photonics wafer testing can be implemented without active optical alignment, providing a more production-oriented path [4]. Looking ahead, as SiPho devices integrate more channels and require increasingly dense direct-current (DC) source measure unit (SMU) resources for combined optical-electrical testing, bringing the SiPho probe card architecture onto automatic test equipment (ATE) platforms will be a key enabler for future wafer-level production testing.

An ATE-based opto-electrical architecture

To resolve the bottleneck between laboratory accuracy and production throughput, an integrated SiPho wafer test architecture was developed using the Advantest V93000 system-on-chip (SoC) automated test equipment (ATE) platform. This production-ready infrastructure consists of three highly synchronized elements:

The ATE tester platform. The Advantest V93000 provides scalable electrical test resources including precision DC parametric measurement units (PMUs), high-speed digital channels, precise timing/control resources, and analog instrumentation. These resources enable accurate electrical biasing and parametric measurement of silicon photonics devices as the channel counts increase and even supporting multisite testing to increase the throughput.

The optical load board interface. The bridge between the ATE and the wafer is a specialized optical load board (**Figure 2**). This custom interface integrates blind-mate optical connectors alongside traditional radio-frequency (RF) and DC electrical signal routing. By utilizing highly-stable mechanical supports for the probe card and precise fiber connector interfaces, the load board ensures that external laser sources and optical

power meters can interface seamlessly with the DUT without sacrificing the mechanical integrity required for automated wafer probing.

The integrated SiPho probe card. The final and most critical mechanical interface is the SiPho probe card, which provides both simultaneous electrical contact through conventional probe needles or microelectromechanical systems (MEMS) contacts, and optical coupling. Optical signals are delivered through fiber arrays designed to passively align with the wafer grating couplers after electrical touchdown. By removing the need for active optical alignment on a die-by-die basis, this architecture enables a more fully-automated and manufacturing-compatible SiPho wafer test flow.

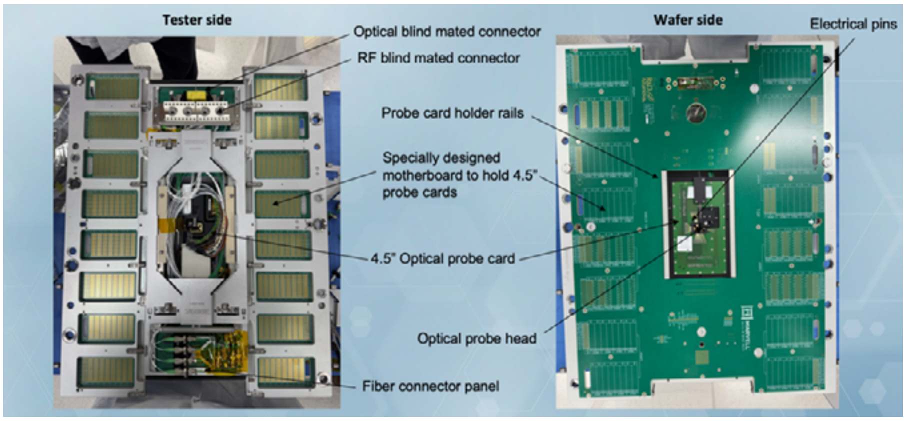

Figure 2: Overview of the optical load board architecture integrated with the Advantest V93000 ATE tester. SOURCE: Marvell/Advantest

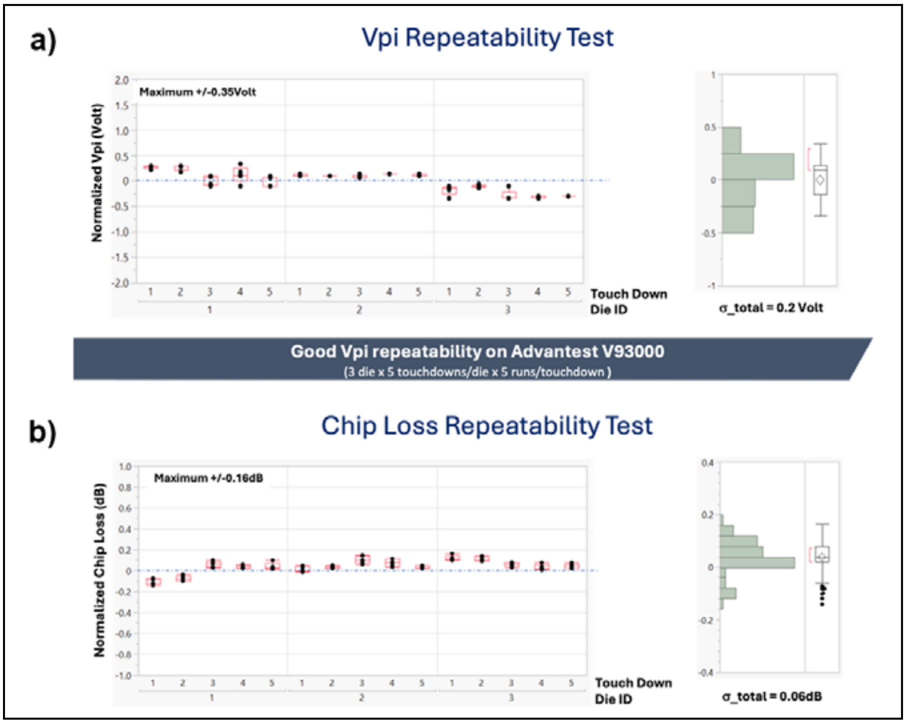

Figure 3: ATE platform repeatability experiments: a) Vpi repeatability test; and b) Chip loss repeatability test. SOURCE: Marvell/Advantest

	Production line: Advantest V93000
Parameter	σ_{total}
Vπ	2.7%
Chip loss	0.06 dB
Heater resistance	0.6%
Thermal diode current	2.8%
PD dark current	1.5%

Table 1: ATE platform measurement repeatability summary. SOURCE: Marvell/Advantest

Validating production readiness: Repeatability and correlation

A test platform is only as viable as its measurement stability. To validate the ATE platform for HVM SiPho testing, exhaustive repeatability experiments were conducted across multiple dies with repeated probe touchdowns (**Figure 3**). The system demonstrated exceptional mechanical and parametric stability. As detailed in **Table 1**, the ATE platform achieved a modulator Vpi repeatability of 2.7% and a chip loss repeatability of just 0.06dB.

Beyond repeatability, establishing “lab-to-fab” correlation is essential for production readiness. Correlation experiments were performed comparing the V93000 ATE platform against a reference PXI-based engineering optical test setup. The results confirmed strong statistical alignment (**Figure 4**):

Vpi correlation. The average difference between the two platforms was approximately 2.8%, with a normalized root mean square (RMS) error of just 0.22%.

Chip loss correlation. The average difference was a mere 0.46dB, with a normalized RMS error of 0.24dB.

The metrics listed above conclusively demonstrate that standard semiconductor ATE platforms can replicate laboratory-grade measurement accuracy while delivering the scalability, automation, and throughput required for HVM.

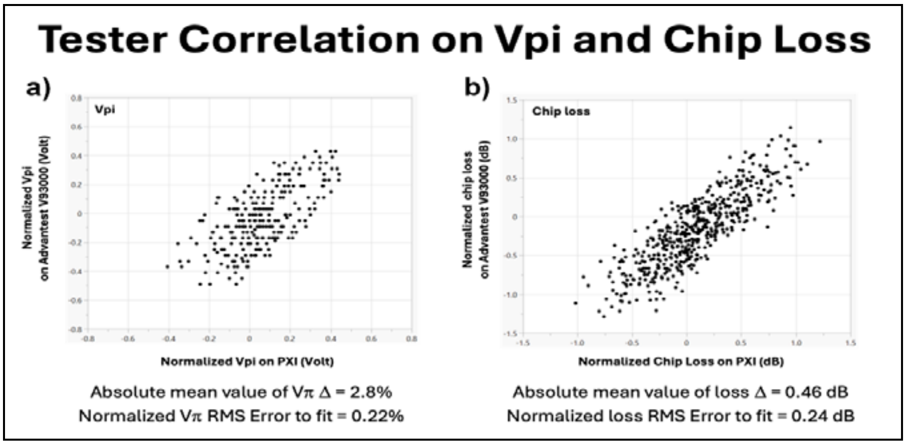

Figure 4: Correlation plots comparing a) Vpi and b) Chip loss between the engineering PXI system and the V93000 ATE tester. SOURCE: Marvell/Advantest

Summary

As the industry advances toward 1.6T and beyond, driven by co-packaged optics and higher-density electrical interfaces, silicon photonics will require much higher channel density, deeper integration with host application-specific integrated circuits (ASICs), and a corresponding step-function improvement in manufacturing test scalability. Scaling wafer-level test infrastructure in parallel with device complexity will therefore become a strategic requirement, not merely a technical consideration.

Integrating silicon photonics wafer test onto a proven ATE platform provides a scalable path for manufacturing traceability, automated wafer stepping, resource consolidation, and parallel test throughput. By leveraging the Advantest V93000 platform, Marvell has demonstrated an effective shift-left yield strategy that enables earlier screening of critical failure modes in production. The result is higher manufacturing efficiency, lower packaging loss, and a stronger production foundation for the next generation of high-volume, multi-terabit optical interconnect platforms supporting large-scale AI and cloud infrastructure.

References

1. M. Tan, J. Xu, S. Liu, J. Feng, H. Zhang, C. Yao, et al.,“Co-packaged optics (CPO): status, challenges, and solutions,” *Frontiers of Optoelectronics*, vol. 16, no. 1, 2023.
2. S. Shekhar, W. Bogaerts, L. Chrostowski, J. E. Bowers, M. Hochberg, R. Soref, et al., “Roadmapping the next generation of silicon photonics,” *Nature Communications*, vol. 15, 2024.
3. OIF, “Common Electrical I/O – CEI-224G” and “1600G Coherent,” Optical Internetworking Forum technical work pages, accessed 2026.
4. A. Yick, et al., “Efficient PIC testing enables data center capacity expansion for optical connectivity,” *Chip Scale Review*, 2025.

Biography

HsuHao (Andy) Chang is Director of Engineering at Marvell, Westlake Village, CA. He leads wafer-level test engineering for silicon photonics and high-speed interconnect products from design through new product introduction (NPI) and high-volume manufacturing (HVM). He also drives test and manufacturing strategy across silicon photonics, 400G-class drivers, and transimpedance amplifiers, helping translate next-generation optical platforms into scalable production. His work focuses on automated wafer-level test infrastructure, platform integration, and manufacturing readiness for complex photonic devices and high-speed interconnect technologies. Email hsuhaoc@marvell.com

SIGMA SERIES

WLCSP SOLUTION

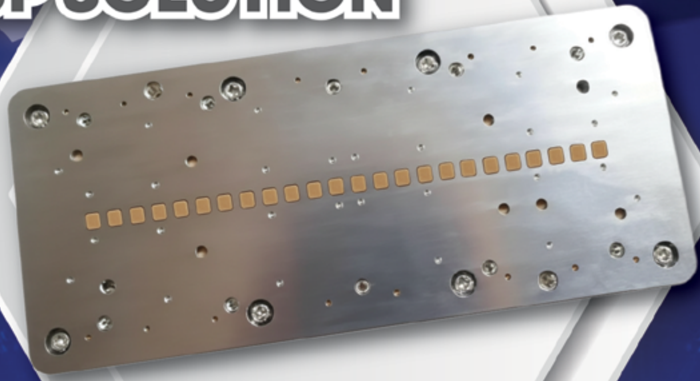

Min. Pitch ≧0.125mm

Pin Count ≦6000

Longevity >1000K

Key Features :

- *Pin coplanarity less than 100μm.*
- *Pin overdrive greater than 200μm.*
- *Short signal length, high compliance pins.*
- *Optimized for systems up to 6000 total I/O's.*
- *Suitable for WLCSP, WLBGA, Cu pillar and bare pad probing.*

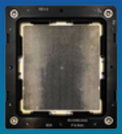

COAXIAL SOLUTION
High Speed 224Gbps
Pin Count >10,000

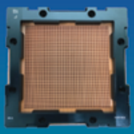

HIGH PIN COUNT SOLUTION
Min. Pitch 0.80mm
Pin Count Up to 20,000

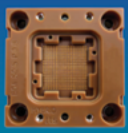

HYBRID ELASTOMER SOLUTION
Transmission Speed >200Gbps
Inductance <0.2nH

PoP SOLUTION
Min. Pitch <0.40mm
Stacked Device

CMOS SOLUTION
Min. Pitch <0.40mm
Imaging Device

RF & mmWAVE SOLUTION
Min. Pitch 0.35mm
Insertion Loss >100GHz@-1dB

Test Tooling Solutions Group

www.tts-grp.com sales@tts-grp.com

50 YEARS
1976 - 2026

SPEA

DOT800

Device-Oriented Mixed Signal Tester

THE POWER OF A MAINFRAME

The Footprint of a Test Head

DOT800C integrates the capabilities of two traditional testers in a single, test-head-only system, offering up to 8,320 channels.
A true zero-footprint solution, delivering 100% test efficiency.

- ✓ **Distributed-intelligence, multi-core architecture** ideal for asynchronous test on SoC and complex chip designs.
- ✓ **Universal instrument slots** to host digital, analog, RF, signal processing and power resources.
- ✓ Device-oriented modular architecture delivers **maximum performance with minimum investment.**
- ✓ Air-cooled design, plug-and-play instrument modules, easy access to all components, ensure the **easiest operation.**
- ✓ **Reliable, gentle contacting** with continuous connection between DUT and test resources during physical stimulus.

Power & Discretes
PMIC
ADC/DAC
BMS
MCU
MEMS & Sensors
Digital & Mixed Signal
Automotive
SerDes
Medical
Linear
5G, IoT & Wireless

DISCOVER HOW DOT800C CAN ENHANCE YOUR TEST PERFORMANCE AND EFFICIENCY.

Visit www.spea.com and book now your free demo.

Pushing NAND packaging to power an AI-driven, data-hungry world

By Chiung Lee, Adrian Arcedera [Amkor Technology, Inc.]

The modern digital landscape is shaped not only by advances in computing power, but also by an explosive rise in data generation. Although discussions about artificial intelligence (AI) often focus on compute performance, memory plays an equally critical role. AI workflows—from training large datasets to running real-time inference—depend heavily on fast, scalable, and reliable memory.

AI models, cloud services, smart devices, and connected systems continuously create, process, and store information at unprecedented rates. In this environment, storage has become a foundational element of digital infrastructure. Without capable storage systems, AI platforms cannot operate efficiently. This trend is visible across multiple sectors. Consumer devices—from sensors to entertainment platforms—continuously generate and store data. In automotive systems, particularly autonomous and connected vehicles, massive volumes of stored information are essential for safe and reliable operation (**Figure 1**). The growing need to access and store large datasets is driving structural, long-term demand for NAND flash memory.

Unlike previous market cycles centered primarily on consumer electronics, today's growth is supported by a broader set of applications. As enterprise solid-state drives (SSDs), cloud infrastructure, communications networks, and AI-optimized platforms expand, NAND is shifting from supporting components to a core pillar of digital infrastructure.

Everyday data growth illustrates the trend of a data hungry world

The expansion of data generation is evident in everyday digital experiences. When streaming video, the transition from HD (1080p) to 4K resolution increases the data required for a typical two-hour movie by approximately 4.7X. This growth is driven by higher pixel density, increased bit depth required for high dynamic range (HDR) content, and higher frame-rate targets.

Photography provides another example. Image file sizes have increased by roughly 3X as smartphone camera resolutions have grown from 12MP to 48MP. Higher spatial resolution combined with enhanced color processing and expanded dynamic range significantly increases the data footprint. Audio formats have also evolved. The shift from compressed MP3 files to immersive audio formats, such as Dolby Atmos, can increase the file size of a three-minute track by up to 6X.

Continued scaling in memory technologies, combined with advances in codec development, will be essential to support the rising data intensity of modern digital media (**Figure 2**).

AI-driven shifts reshape the memory market

According to Yole Group's 2024 data, the stand-alone memory market rebounded to $170 billion, representing a 78% year-

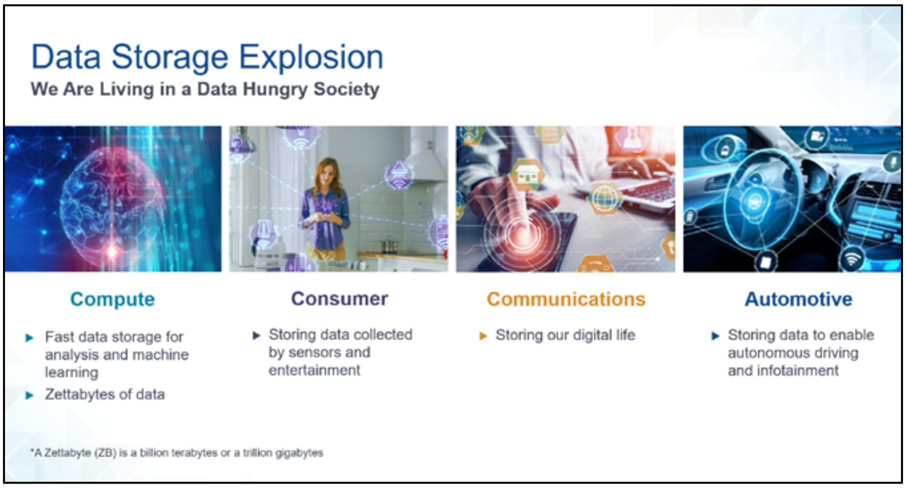

Data Storage Explosion

We Are Living in a Data Hungry Society

Compute

- Fast data storage for analysis and machine learning
- Zettabytes of data

Consumer

- Storing data collected by sensors and entertainment

Communications

- Storing our digital life

Automotive

- Storing data to enable autonomous driving and infotainment

*A Zettabyte (ZB) is a billion terabytes or a trillion gigabytes

Figure 1: Storage has become a foundational element of digital infrastructure. SOURCE: Amkor

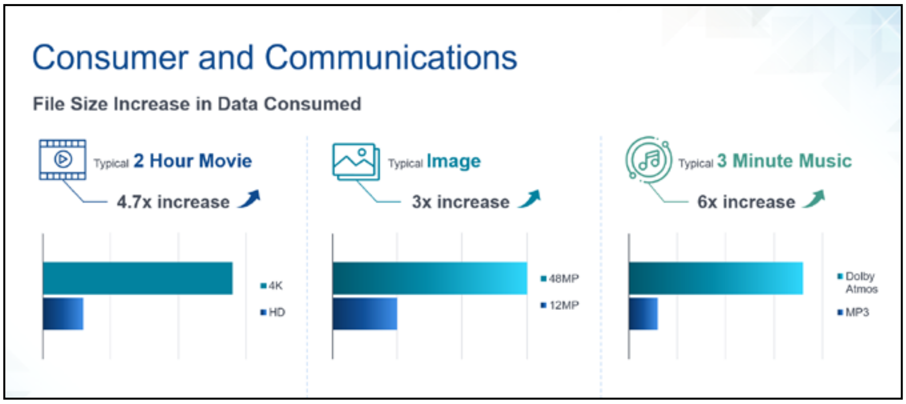

Consumer and Communications

File Size Increase in Data Consumed

Figure 2: The rising data intensity of modern digital media. SOURCE: Amkor

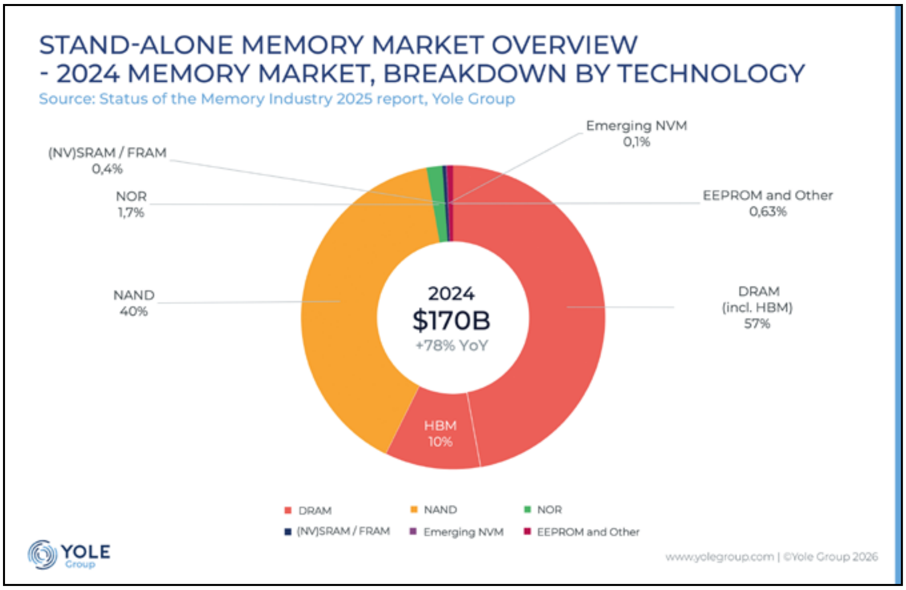

Figure 3: Memory market overview (2024), "Status of the Memory Industry 2025," Yole Group [1].

over-year increase after two years of decline (**Figure 3**) [1]. Dynamic random access memory (DRAM) (including high-bandwidth memory [HBM]) accounted for approximately 57% of revenue, followed by NAND at 40%.

Looking ahead to 2026, AI workloads are reshaping how memory suppliers allocate capacity and investment across DRAM, HBM, and NAND. Manufacturers are increasingly: 1) Expanding production of HBM and advanced DRAM nodes; 2) Prioritizing products aligned with AI performance requirements; and 3) Accelerating HBM ramp-ups, which now represent more than 10% of the total memory market.

While the steps listed above strengthen the HBM ecosystem, they also limit investment in traditional DRAM and NAND capacity. Despite these constraints, NAND demand continues to grow across multiple applications: 1) Data center SSDs: Critical storage infrastructure for AI servers; 2) PC and laptop SSDs: Increasing attach rates across consumer and enterprise devices; 3) Mobile devices: Higher storage tiers in premium smartphone models; and 4) Automotive systems: Electrification and intelligent vehicle architectures requiring robust solid-state storage. These diversified demand drivers support a stable outlook for NAND through 2026, even amid tighter supply conditions.

The evolution of NAND for a data-hungry world

As storage requirements continue to expand, the NAND industry is advancing both wafer-level scaling and high-density die-stack packaging technologies. Consider the physical footprint required to store 2TB of data in 2010, when only planar (2D) NAND was available (**Figure 4**). At that time, achieving 2TB would have required more than 50 Embedded MultiMediaCard (eMMC) chips, each in a package roughly 17 X 22mm in size.

The introduction of 3D NAND in 2013 dramatically improved storage density by enabling vertical stacking of memory cells. By 2020, devices featuring 100+ layer architectures significantly reduced the number of packages required to achieve the same storage capacity.

Looking ahead, suppliers are preparing for wafer-to-wafer hybrid bonding, which increases interconnect density while improving electrical performance at the wafer level. This technology is expected to extend vertical scaling beyond today's 100-layer devices toward architectures that may eventually reach 800 to 1,000 layers through multi-wafer stacking [2-3]. In parallel, higher-layer 3D NAND architectures combined with increasing quad-level cell (QLC) adoption are enabling a single eMMC device to reach 2TB of storage capacity.

Key packaging technology evolution: Thin, tall, and reliable

At the package level, the rapid growth of mobile and enterprise SSDs is driving demand for denser ball grid array (BGA) packages and higher-speed interfaces. These requirements demand tightly integrated die stacks and controller integration enabled by advanced packaging technologies (**Figure 5**).

Achieving thinner NAND dies to enable higher stacking counts—while maintaining high yield and reliability—requires careful coordination across multiple manufacturing processes. Through continuous innovation and disciplined process optimization, outsourced semiconductor assembly and test (OSAT) providers must establish best-in-class quality control across the entire packaging flow, from wafer preparation through final assembly.

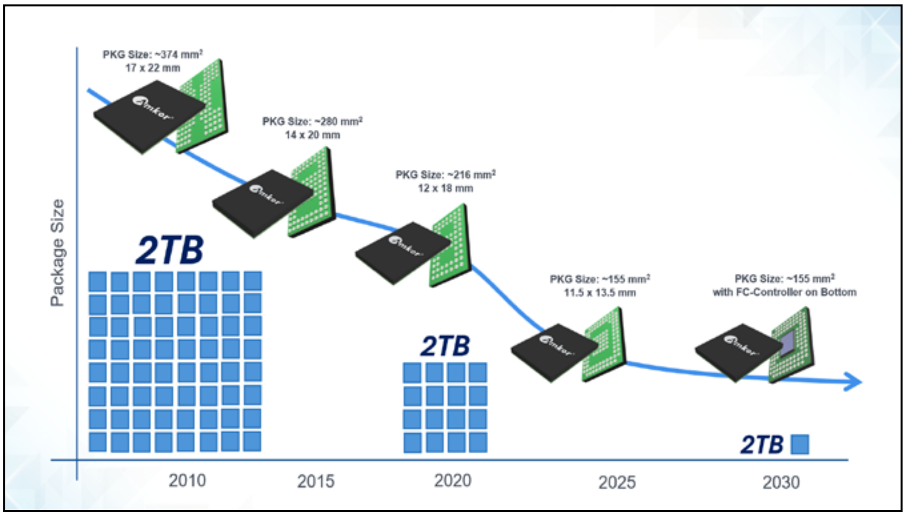

Figure 4: The physical footprint for storage required over time. SOURCE: Amkor

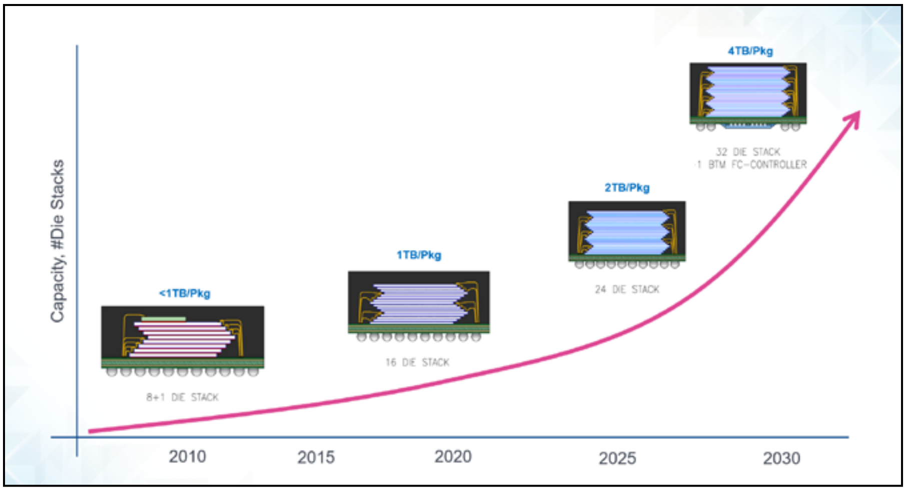

Figure 5: Growth in NAND capacity/number of layers over time. SOURCE: Amkor

Wafer preparation for high-layer 3D NAND

Whether assembling single-die devices or advanced packages containing 4, 8, 16, or even 32 stacked dies, the process begins with precise wafer preparation. Wafer preparation refers to the set of processes performed after wafer fabrication but before assembly, preparing wafers for singulation, inspection, and downstream packaging operations.

Modern wafer-preparation flows are engineered for 300+ layer NAND devices, addressing die stress and warpage to prevent cracking or chipping. Key capabilities include: 1) Laser grooving; 2) Ultra-fine wafer grinding; 3) Blade or stealth dicing; and 4) Automated optical inspection (AOI) for crack detection. **Figure 6** shows the typical wafer dicing technologies.

As wafers transition to ultra-thin profiles, thickness is reduced from approximately 700–800µm to as little as 25–50µm. At these thicknesses, careful thinning and backside conditioning become essential to maintain mechanical integrity. Fine grinding, polishing, and backside surface treatments remove subsurface micro-cracks and reduce stress concentrations, significantly improving die strength and handling margins.

To further protect fragile high-layer NAND wafers during singulation, saw streets are conditioned using techniques such as laser grooving or plasma pre-cut that reduce mechanical stress during dicing. Advanced methods including dice-before-grind (DBG) and stealth dice-before-grind (SDBG) are then used to minimize edge damage and chipping. Throughout this process, AOI plays a critical role by detecting die cracks in real time, enabling early removal of compromised dies before downstream assembly.

High-stack memory and controller integration

Once the NAND dies are prepared, the next critical step is die pickup and handling. As die thickness decreases, traditional handling approaches must be replaced with specialized thin-die tooling. This includes needle-less pickup systems with precisely-controlled vacuum force to minimize die bending and prevent fracture. Together with the wafer-level preparation steps ensure that each die is sufficiently thin, clean, and mechanically robust to withstand high-stack assembly, wire bonding, molding, and the rigorous reliability testing required for advanced NAND packaging.

As NAND architecture scales up to 32+ layer counts as the structure shown in **Figure 7** illustrates, integration challenges shift from optimizing isolated process steps to managing the die stack as a unified system. Achieving robust high-stack NAND packages and mixed flip-chip/wire bond controller attach configurations requires holistic control over interconnect design, mechanical reinforcement, and material selection. The entire packaging flow must be tightly coordinated across low-profile wire bonding, advance molding, inline inspection, and carefully balanced material systems to effectively control warpage and enable high-volume manufacturing.

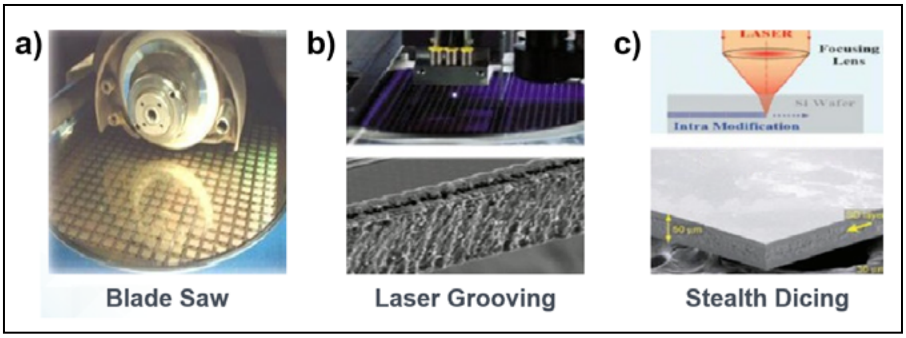

Figure 6: Typical wafer dicing technologies: a) (left) Blade saw; b) (center) Laser grooving; and c) (right) Stealth dicing. SOURCE: Amkor

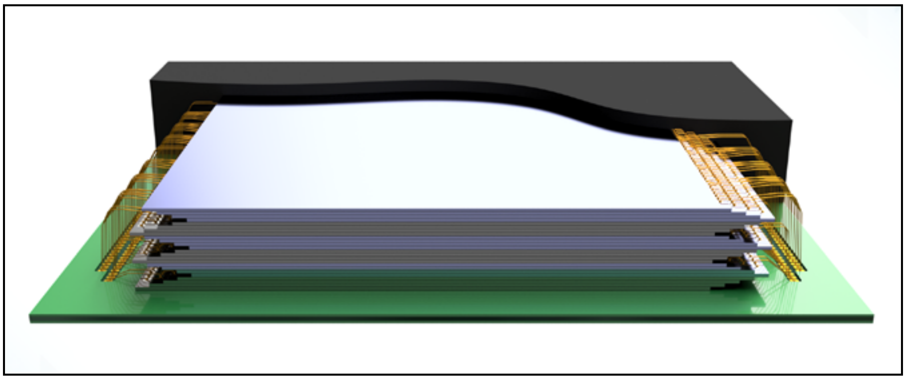

Figure 7: Example of NAND flash package architecture scaled up to 32+ layer counts. SOURCE: Amkor

Wire bond architecture for high-density integration

Wire-bond architecture plays a central role in enabling high-density integration within increasingly constrained package Z budgets, the total package thickness, as required by end-customer applications such as smartphones. Both ultra-low loop and ultra-long loop profiles are required to maintain clearance in tall stacks, while extended die overhang capability supports complex stair-step and zig-zag configurations without compromising reliability. Careful tuning of bond parameters and loop geometry ensures compatibility with different pad designs and prevents wire sweep or shorting during encapsulation.

Stack topology directly impacts stability and routing, with stair-step layouts ensuring wire clearance in tall stacks and zig-zag designs distribute stress. Customized die attach methods such as film-on-die (FOD), film-on-wire (FOW), or spacers prevents die sag and preserves wire bond integrity, especially in controller-under-NAND configurations, without adding Z-height.

Molding, warpage control, and reliability

Encapsulation is critical for warpage control at the package level. Advanced compression molding is favored to minimize flow-induced wire sweep and ensure uniform pressure distribution across tall stacks. Mold cap thickness and overall material stack balance must be carefully optimized to meet profile targets while controlling coefficient of thermal expansion (CTE) mismatch.

Throughout the flow, inspection and data analytics provide the feedback loop needed to sustain yield and reliability. Inline inspection at key stages and unit-level traceability enable early detection of anomalies and rapid correlation to materials, tools, or process conditions. This data-driven approach accelerates learning cycles and improves excursion containment in high-volume production.

Finally, reliability is reinforced through deliberate materials engineering. Low-stress die-attach films, appropriately-selected substrates, and optional electromagnetic interference (EMI) shielding are used to manage mechanical stress, thermal cycling behavior, and system-level performance requirements. Together, these integration strategies transform advanced NAND packaging from a collection of individual process steps into a cohesive, manufacturable platform capable of supporting next-generation memory density and performance.

OSAT providers: Strategic partners in a changing global landscape

Not all NAND flash manufacturers are integrated device manufacturers (IDMs) capable of handling the full value chain from intellectual property (IP) development and wafer fabrication, to final assembly and test. This is where OSAT providers play a critical role in the semiconductor ecosystem. In today's increasingly complex geopolitical environment, the value of OSAT partners with globally-distributed manufacturing capabilities has become even more significant.

Leading OSAT providers support semiconductor companies across consumer, industrial, and automotive markets while offering advanced capabilities such as: 1) High-stack memory packaging; 2) Unit-level traceability; 3) Wafer-level fan-out technologies; and 4) Turnkey test services. Only a limited number of providers operate multi-region manufacturing networks that provide both supply flexibility and resilience. Facilities distributed across Asia, Europe, and the United States are becoming increasingly important as semiconductor companies prioritize diversified and geopolitically-balanced supply chains.

Summary

The rapid growth of data across consumer electronics, automotive systems, enterprise infrastructure, and cloud computing continues to elevate the importance of NAND storage.

Higher-resolution media, connected devices, and data-intensive computing have made NAND a core element of modern digital systems. To support this growth, the industry is advancing both NAND device density and packaging technology. Technologies such as wafer-to-wafer hybrid bonding, taller die stacks, and improved reliability controls are enabling higher storage capacity within increasingly compact footprints.

Progress in NAND scaling, advanced packaging technologies, and global manufacturing ecosystems will be essential to meeting the storage demands of an increasingly data-driven world.

References

1. Yole Group, "Memory Market Outlook 2024–2030," https://files.futurememorystorage.com/proceedings/2025/20250807_BMKT-301-1_Lau.pdf
2. Semiconductor Engineering, "NAND Flash Targets 1,000 Layers," Dec. 2024; https://semiengineering.com/nand-flash-targets-1000-layers/
3. AllElectroHub / AllPCB, "3D NAND Hits 400+ Layers with Hybrid Bonding," Aug. 2025; https://www.allpcb.com/allelectrohub/3d-nand-hits-400-layers-with-hybrid-bonding

Biographies

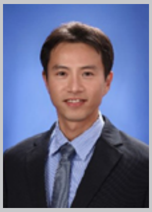

Chiung Lee is Director in the Memory, MEMS, and Sensor BU at Amkor Technology, Inc., Tempe, AZ, USA. He joined Amkor in 2015 and is currently in charge of NAND memory and fingerprint sensor products. Prior to Amkor, he was a Marketing Manager at STMicroelectronics, responsible for MEMS sensor products. He previously spent eight years at ITRI Taiwan developing a wide range of MEMS technologies. He holds a Master's degree in Mechanical Engineering from National Taiwan University. Email Chiung.lee@amkor.com

Adrian Arcedera is SVP of the Memory, MEMS, and Sensor BU, leading business and platform development at Amkor Technology, Inc., Tempe, AZ, USA. Adrian joined Amkor Technology in 1997 and has held multiple leadership roles in chip-scale products, led ChipArray® platform development, authored technical papers, holds multiple U.S. patents, and earned a Chemical Engineering degree from the University of the Philippines.


LAPLACE

Fine-grain copper for low-temperature hybrid bonding

By Pingping Ye [MacDermid Alpha Electronics Solutions] Jing Xu [Applied Materials]

The demand for higher performance, increased bandwidth, and reduced form factors in electronic systems continues to accelerate, driven by the rapid evolution of high-performance computing (HPC), artificial intelligence (AI) accelerators, and high-bandwidth memory (HBM) architectures. Modern data-centric workloads—including large language models, real-time analytics, autonomous systems, and advanced graphics—require unprecedented data movement between logic and memory. As transistor scaling alone can no longer deliver proportional system-level gains, performance improvements increasingly depend on heterogeneous integration and advanced packaging technologies. In this landscape, interconnect density, bandwidth per unit area, and the energy efficiency of chip-to-chip communication have become primary design constraints.

These system-level pressures are pushing semiconductor packaging technologies toward increasingly aggressive interconnect density requirements. Shorter interconnect lengths reduce parasitic resistance and capacitance, thereby lowering latency and power consumption while enabling higher input/output (I/O) counts. At the same time, the need for compact three-dimensional (3D) integration demands vertical stacking solutions that minimize z-height without compromising mechanical robustness or electrical integrity. Consequently, advanced packaging is transitioning from a supporting role to a central enabler of overall system performance.

Traditional interconnect methods, including solder-based microbumps and copper-pillar bumps, have supported die-to-die, die-to-wafer, and wafer-to-wafer integration for many years. These approaches have enabled 2.5D integration using silicon interposers and early generations of 3D stacking with through-silicon vias (TSVs). However, as interconnect pitches approach and fall below approximately 10µm, these technologies encounter fundamental scaling limitations. Solder volume control becomes increasingly difficult at fine pitches, leading to height non-uniformity and alignment sensitivity. Reflow-induced defects, solder bridging, and electrical shorts become more prevalent as pitch size shrinks.

Moreover, reliability concerns intensify at advanced nodes. Electromigration risks increase due to higher current densities, while thermomechanical stress can cause fatigue and crack propagation over time. Underfill materials, which play an important role in enhancing mechanical stability, also add process considerations that must be carefully managed alongside stress-related reliability requirements. From a design perspective, the physical height of solder-based interconnects imposes constraints on achievable vertical scaling, limiting how tightly dies can be stacked in high-density 3D architectures.

As scaling trends continue and system bandwidth requirements grow, these limitations highlight the need for a fundamentally different interconnect paradigm. Hybrid bonding has therefore emerged as a promising solution to address these challenges and enable continued interconnect scaling.

Hybrid bonding as a scaling enabler

Hybrid bonding simultaneously forms a dielectric-to-dielectric bond, such as silicon dioxide – silicon dioxide (SiO_2–SiO_2), and a direct metal-to-metal bond, typically copper – copper (Cu–Cu), between mating surfaces. The bump-less nature of this interconnect architecture enables ultra-fine-pitch interconnects, reduced vertical height, improved electrical performance, and enhanced long-term reliability. As a result, it is increasingly viewed as a key enabling technology for next-generation heterogeneous integration and advanced 3D system architectures.

Commercial adoption of Cu–Cu hybrid bonding began with Sony's stacked complementary metal-oxide semiconductor (CMOS) image sensors in 2016, marking a major milestone in the high-volume manufacturing of direct copper-to-copper interconnect technology [1]. In these devices, pixel arrays and logic circuits were fabricated on separate wafers and vertically integrated using fine-pitch hybrid bonding, enabling significant improvements in pixel density, signal integrity, and functional partitioning. By eliminating conventional micro-bumps, Sony achieved tighter interconnect pitch, reduced parasitic resistance and capacitance, and improved form factor—demonstrating that hybrid bonding could meet both performance and yield requirements in mass production.

Following its success in image sensor applications, Cu–Cu hybrid bonding expanded into high-performance computing and advanced logic–memory integration. For example, AMD introduced 3D V-Cache™ technology, which vertically stacks static random access memory (SRAM) cache directly on top of logic dies using hybrid bonding to dramatically increase on-chip cache capacity without enlarging the processor footprint [2]. Similarly, Intel developed Foveros Direct, an advanced 3D stacking platform that leverages direct copper hybrid bonding to achieve ultra-fine interconnect pitch and high interconnect density for heterogeneous integration [3]. These implementations demonstrate hybrid bonding's scalability beyond imaging devices into bandwidth-intensive processor architectures, where interconnect density and energy efficiency are critical.

Fine-grain copper for low-temperature hybrid bonding

Despite the advantages noted above, most existing Cu–Cu hybrid bonding processes rely on post-bond annealing temperatures exceeding 300°C to achieve sufficient copper diffusion across the bonding interface and ensure low-resistance, mechanically-robust joints. Elevated temperatures promote grain growth and interfacial atomic interdiffusion, reducing void formation and enhancing bond strength. However, such high thermal budgets introduce several important challenges discussed below.

First, high annealing temperatures restrict compatibility with temperature-sensitive materials, including low-k dielectrics, polymer-based passivation layers, certain metal stacks, and prefabricated device layers. Degradation of these materials can result in altered electrical characteristics, dielectric breakdown risks, or mechanical instability. Second, thermal excursions above 300°C increase wafer-level and die-level stress due to coefficient-of-thermal-expansion (CTE) mismatch between bonded materials. This stress can lead to wafer warpage, delamination, or defect generation, particularly in large-diameter wafers and thin-die configurations used for advanced 3D stacking.

These constraints have motivated significant research into lower-temperature Cu–Cu hybrid bonding solutions. Approaches including grain engineering [4], surface activation, plasma treatment, chemical–mechanical planarization (CMP) optimization, passivation-free copper surfaces, and transient liquid-phase or diffusion-assisted bonding aim to reduce the required annealing temperature while maintaining electrical conductivity, mechanical strength, and long-term reliability. Achieving robust hybrid bonding below 300°C, or ideally near, or below 200°C, would substantially broaden process compatibility, reduce thermal stress, and accelerate the adoption of hybrid bonding across diverse advanced packaging platforms. Therefore, reducing the thermal budget of Cu–Cu hybrid bonding has become increasingly critical as advanced packaging integrates thinner wafers, novel dielectric materials, and complex heterogeneous stacks.

One extensively studied grain-engineering strategy involves nanotwinned copper, which is dominated by the (111) crystallographic orientation and exhibits enhanced atomic diffusivity along twin boundaries [5-6]. While nanotwinned copper enables stronger Cu–Cu bonding at lower temperatures than conventional electroplated copper, achieving consistent nanotwin microstructures in vias with openings smaller than 4µm and high aspect ratios is highly challenging.

Fine-grain copper offers an alternative and highly-manufacturable solution. Fine-grain copper contains a high density of grain boundaries that store excess grain-boundary energy that can be released during controlled thermal annealing. In this article, we focus on metastable fine-grain copper as a practical and scalable material platform for low-temperature Cu–Cu hybrid bonding.

A typical Cu–Cu hybrid bonding process (**Figure 1**) begins with damascene fabrication of copper pads embedded in a dielectric layer (such as SiO_2 or SiCN), followed by CMP to achieve a highly planar and coplanar copper/dielectric surface with controlled copper recessed depth. Strict surface topography control—typically within a few nanometers—is essential to ensure uniform contact during bonding.

The two wafers (or die and wafer) are then precisely aligned using high-accuracy alignment systems, with submicron overlay capability. Initial bonding occurs at room temperature through dielectric-to-dielectric contact driven by van der Waals forces, establishing mechanical pre-bonding. This is followed by a post-bond annealing step at elevated temperature. Due to the difference in CTE between copper and the dielectric materials, copper expands more than the dielectric during heating, allowing the copper pads to come into intimate contact. The elevated temperature further promotes copper atomic diffusion across the interface, enabling direct Cu–Cu metallurgical bonding.

Working mechanism of fine-grain copper in hybrid bonding

In fine-grain copper, an additional diffusion-enhancing mechanism arises from grain-growth-induced energy release. As metastable grains coarsen at elevated temperatures, the total grain-boundary area is reduced, releasing stored energy that locally enhances atomic diffusion at the bonding interface and promotes robust metallurgical bonding.

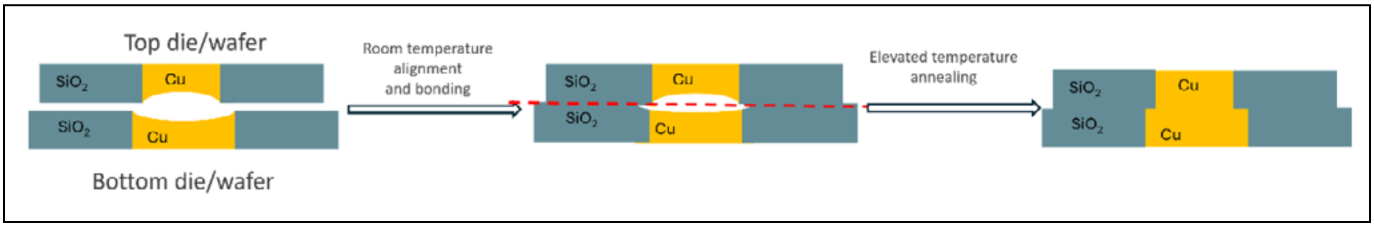

Figure 1: Hybrid bonding process flow. SOURCE: MacDermid Alpha Electronics Solutions

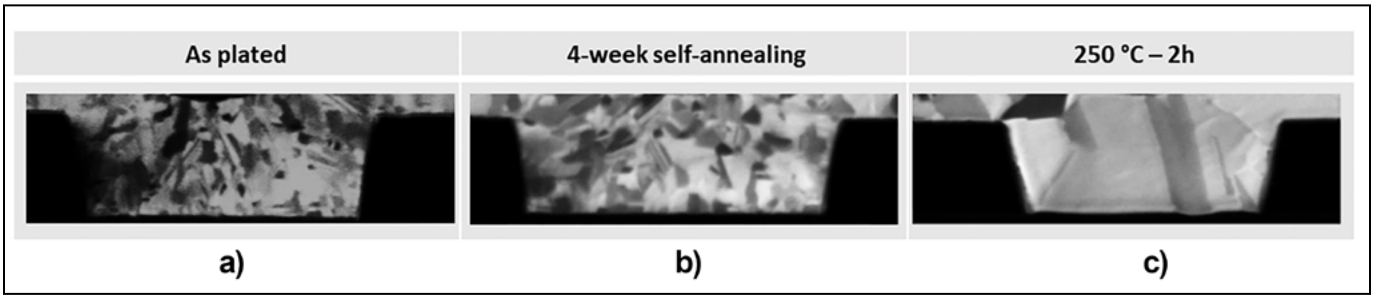

Figure 2: Ion images of Cu grain: a) As plated; b) Four weeks self-annealing; and c) 250°C annealing for 2h. SOURCE: Applied Materials

To effectively utilize this mechanism, the copper microstructure must satisfy two key conditions. First, it must remain metastable at room temperature, maintaining a fine-grain structure during typical wafer queue times of several weeks prior to bonding. This stability is achieved by suppressing grain-boundary mobility and minimizing the thermodynamic driving force for spontaneous grain growth. Microstructural evolution involves the movement of grain boundaries, and the velocity of a grain boundary is determined by its mobility and the driving force (ΔG) [7]. The driving force for growth originates from factors such as grain-boundary energy, stacking faults, dislocations, surface energy, and elastic strain, while Zener pinning provides an opposing force. It is therefore critical to reduce both grain-boundary mobility and the driving force to ensure room-temperature stability of the metastable fine-grain structure.

Second, the material must exhibit pronounced grain growth at bonding temperatures, typically in the range of 200–280°C, enabling rapid energy release and enhanced copper interdiffusion during bonding. At these elevated temperatures, the primary driving force is the reduction of total grain-boundary energy, combined with increased atomic mobility. This facilitates atomic rearrangement and migration, leading to grain coalescence and the formation of larger grains.

Experimental verification of low-temperature hybrid bonding

Experimental characterization verifies that the developed fine-grain copper films satisfy both structural stability at room temperature and the thermally-activated grain-growth requirements essential for hybrid bonding applications. The electroplated copper films were deposited using MacDermid Alpha GE 100 electroplating chemistry on the Applied Materials Nokota plating platform. All electroplating processes were conducted at room temperature to ensure simplified temperature control in mass production. The test vehicle fabrication and integration work were carried out in collaboration with the Applied Materials team.

Microstructural analysis indicates that copper films exhibit excellent stability under ambient conditions. Even after four weeks of self-annealing at room temperature, the average grain size remained approximately 100nm (**Figures 2a** and **2b**), demonstrating minimal spontaneous grain growth and confirming low driving force and low grain-boundary mobility at low temperatures. This stability is critical for maintaining high grain-boundary density prior to bonding.

In contrast, thermal annealing at 250°C for two hours resulted in pronounced grain coarsening (**Figure 2c**), clearly evidenced by a significant increase in grain-boundary mobility at elevated temperatures. The marked transition from microstructural stability at room temperature to rapid grain growth at bonding temperatures highlights the strong temperature dependence of the film's grain-boundary dynamics.

Overall, this pronounced temperature-dependent behavior aligns well with the requirements of low-temperature hybrid bonding processes: grain stability during storage and handling, coupled with enhanced atomic mobility and grain growth during bonding to promote interfacial diffusion and robust metallurgical joining.

The effectiveness of the metastable fine-grain copper films was further validated through comprehensive die-to-wafer hybrid bonding demonstrations assessing both electrical performance and interfacial integrity. Bonding was carried out at 250°C, substantially lower than the temperature typically required for conventional Cu–Cu bonding processes. Under these conditions, the bonded samples achieved greater than 90% electrical yield across a 2.5k daisy-chain test structure (**Figure 3**). C-mode scanning acoustic microscopy (C-SAM) revealed no detectable voids or delamination across the bonded area (**Figure 4**), indicating uniform contact and consistent bond quality.

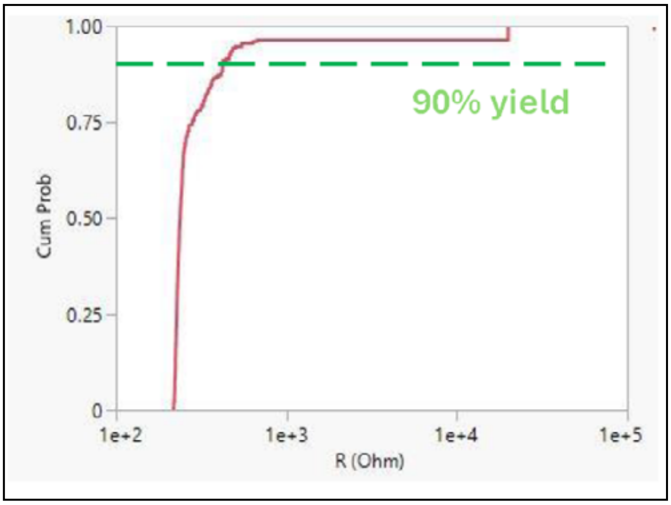

Figure 3: Bonding electrical performance of the D2W 250˚C - 2.5k-via DC. SOURCE: Applied Materials

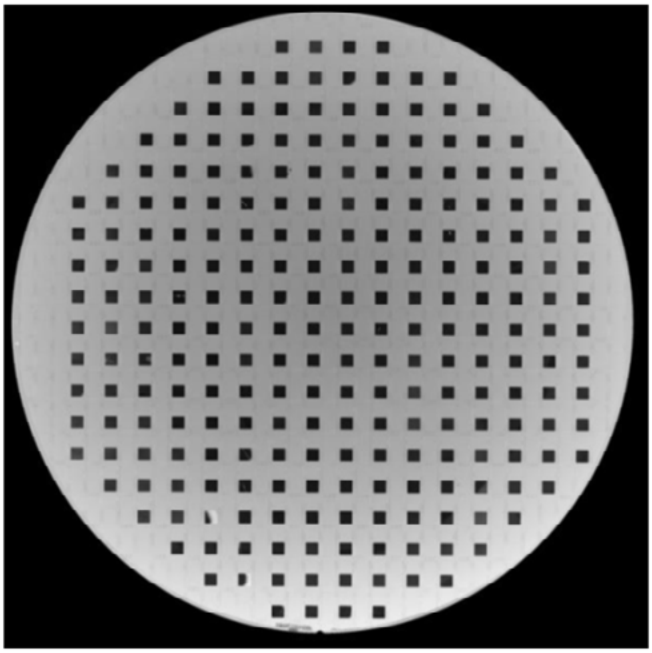

Figure 4: cSAM image for the bonded D2W HB TV after 250˚C. SOURCE: Applied Materials

Electron backscatter diffraction (EBSD) and focused ion-beam (FIB) cross-sectional imaging confirmed copper interdiffusion across the original bonding interface (**Figure 5**). The disappearance of a distinct interface and the presence of continuous grain structures across the bond line verify active

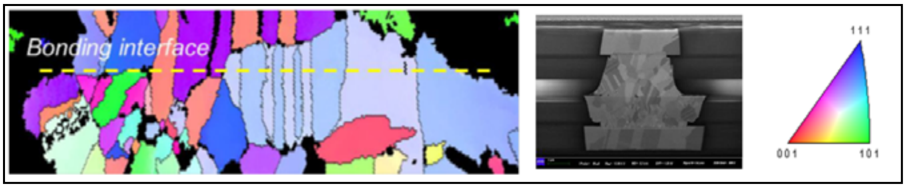

Figure 5: EBSD and SEMI of bonding structure. SOURCE: Applied Materials

grain growth and atomic diffusion during bonding, confirming the formation of a robust metallurgical Cu–Cu bond.

Collectively, the results discussed above demonstrate that metastable fine-grain copper provides a reliable and manufacturable pathway to low-temperature Cu–Cu hybrid bonding. By combining room-temperature microstructural stability with grain-growth-driven diffusion at bonding temperatures, fine-grain copper enables strong electrical and mechanical interconnects while substantially reducing the overall thermal budget.

Summary

The findings presented in this article establish metastable fine-grain copper as a validated, reliable, and practical material solution for enabling low-temperature Cu–Cu hybrid bonding. The unique combination of room-temperature stability and grain-growth-driven diffusion during bonding provides an effective mechanism for reducing the thermal budget of hybrid bonding processes. This capability preserves the integrity of sensitive device structures while enabling robust metallurgical bonding at significantly lower temperatures than conventional approaches.

As hybrid bonding becomes a cornerstone of advanced 3D integrated circuits (3D ICs) and heterogeneous system-in-package (SiP) architectures, reliable low-temperature bonding will support greater design flexibility, improved yield, and compatibility with temperature-sensitive materials such as low-k dielectrics, advanced transistors, and organic interposers.

Continued progress in materials engineering, surface preparation, alignment precision, and hybrid-via architecture development will further unlock the potential of fine-grain copper. Together, these advances will support continued scaling of interconnect density, improved signal integrity, and enhanced overall system performance in advanced 3D and heterogeneous integration platforms. In this context, metastable fine-grain copper represents a scalable, manufacturable, and forward-looking pathway toward enabling future high-density, low-temperature three-dimensional integrated systems, positioning it as a key enabler for next-generation high-performance electronic devices.

References

1. K. Kanda, et al., "Three-layer stacked CMOS image sensor technology using Cu–Cu hybrid bonding," IEDM Technical Digest, 2016.
2. J. A. Kahle, et al., "3D V-Cache™ technology: SRAM die stacking using hybrid bonding," IEEE Inter. Solid-State Circuits Conf. (ISSCC), 2022.
3. C. Auth, et al., "Foveros Direct: Direct copper-to-copper bonding for high-density 3D integration," IEEE Electronic Components and Tech. Conf. (ECTC), 2021.
4. L. Mirkarimi, C. Uzoh, D. Suwito, B. Lee, G. Fountain, T. Workman, et al., "The influence of Cu microstructure on thermal budget in hybrid bonding," ECTC 2022.
5. J. Juang, C. Lu, K. Chen, C. Chen, P. Hus, C. Chen, et al., "Copper to copper direct bonding on highly (111)-oriented nanotwinned copper in no-vacuum ambient," Scientific Reports 2018 8:13910.
6. P. Ye, J. Han, S. Braye, K. Whitten, V. Hayes, H. Khanna, et al., "Transforming WLP applications: Introducing cutting-edge next-generation nanotwinned copper," Jour. of Microelectronics and Electronic Packaging (2024) 21, 20–24.
7. C. Detavernier, S. Rossnagel, C. Noyan, S. Guha, C. Cabral, Jr., et al.,"Thermodynamics, and kinetics of room-temperature microstructural evolution in copper film," J. Appl. Phys. 94, 2874–2881 (2003).

Biographies

Pingping Ye is a Principal Research Scientist in the Wafer-Level Packaging Group at MacDermid Alpha Electronics Solutions, Shelton, CT, USA. Her work focuses on the development of specialty electroplating chemistries for advanced semiconductor interconnect metallization, including redistribution layer (RDL) and copper pillar technologies. In recent years, her research has centered on chemical process and plating solution development for hybrid bonding applications, supporting next-generation heterogeneous integration and advanced packaging architectures. She received her PhD in Chemistry from Carnegie Mellon University, and her MS in Materials Science and Engineering from the Chinese Academy of Sciences. Email Pingping.Ye@MacDermidAlpha.com

Jing Xu is a Senior Member of Technical Staff at Applied Materials, Kalispell, MT, USA. She has expertise in copper damascene interconnect and advanced packaging electroplating process development and integration. Her work encompasses electroplating chemistry optimization, feature scale fill behavior, and process window expansion, with a strong focus on interactions across unit processes and integration schemes for advanced logic and packaging applications. She supports technology development and manufacturing readiness through data-driven integration, cross-module optimization, and cross-disciplinary collaboration. She received her PhD from Case Western Reserve University, and an MS from Nanjing University.


SEMIEXPO HEARTLAND
WORKFORCE DEVELOPMENT THEATER


Kelly
GIVE YOUR TEAM

Highlights of the 27th Electronics Packaging Technology Conference (EPTC 2025)

*By Tang Gong Yue [Institute of Microelectronics (IME), Agency for Science, Technology and Research (A*STAR)]*

The 27th Electronics Packaging Technology Conference (EPTC 2025) was held December 2–5, 2025, at the Resorts World Convention Centre, Sentosa, Singapore. The conference featured a comprehensive technical program covering advanced packaging, heterogeneous integration, artificial intelligence (AI) infrastructure, and co-packaged optics, etc. Building on the strong momentum of previous years, EPTC 2025 continued with a four-day program format to accommodate the growing breadth and depth of technical content and community engagement. The conference attracted a record 800 attendees from 21 countries, reflecting the continued importance of EPTC as a key platform for the global electronics packaging community.

The conference opened with a welcome address by the General Chair, followed by a strong plenary program comprising four keynote presentations, two panel sessions, and a Heterogeneous Integration (HI) Workshop, setting the tone for four days of in-depth technical exchange and discussion.

Keynote addresses

EPTC 2025 featured four distinguished keynote speakers from industry and research institutes, who shared their perspectives on the future of advanced packaging, AI infrastructure, and heterogeneous integration. These keynote sessions (listed below) highlighted the increasing demand for heterogeneous integration, AI-driven packaging advancements, and sustainability considerations in semiconductor manufacturing.

	**Scaling AI infrastructure with advanced optical interconnects.** Speaker: Radha Nagarajan, Senior Vice President & Chief Technology Officer, Optical Engineering Group, Marvell.
	**Interconnect horizons: Wafer and panel innovation and industry partnerships to unlock AI's next leap.** Speaker: Audrey Charles, Senior Vice President of Corporate Strategy & Advanced Packaging; President, Lam Capital.
	**Rewiring edge AI system efficiency with advanced packaging.** Speaker: Pax Wang, Technology Development Director, UMC.
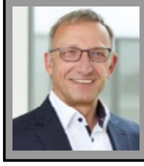	**Driving innovation through hetero-integration: Technologies, challenges and future directions.** Speaker: Prof. Dr. Harald Kuhn, Director, Fraunhofer Institute for Electronic Nano Systems (ENAS).

Panel sessions

Two panel sessions were organized to address timely challenges in the advanced packaging ecosystem. Both sessions generated lively exchanges between panelists and the audience, reflecting strong interest from the community in manufacturability, metrology, test, and system-level integration challenges. The sessions are described below.

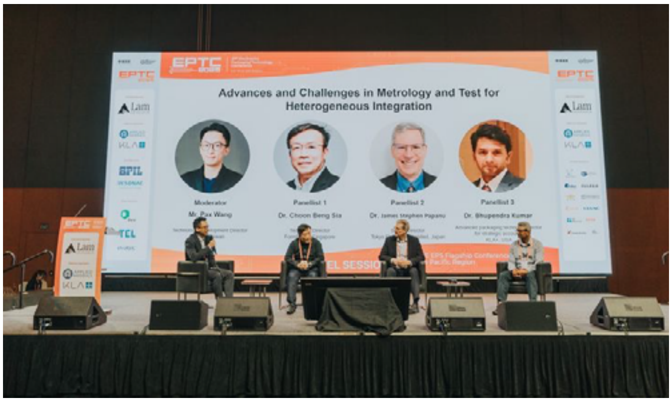

Panel session 1. "Advances and challenges in metrology and test for heterogeneous integration." The moderator was Pax Wang (UMC). Panelists were: 1) Dr. James S. Papanu (TEL); 2) Dr. Sia Choon Beng (FormFactor); and 3) Bhupi Kumar (KLA-Tencor).

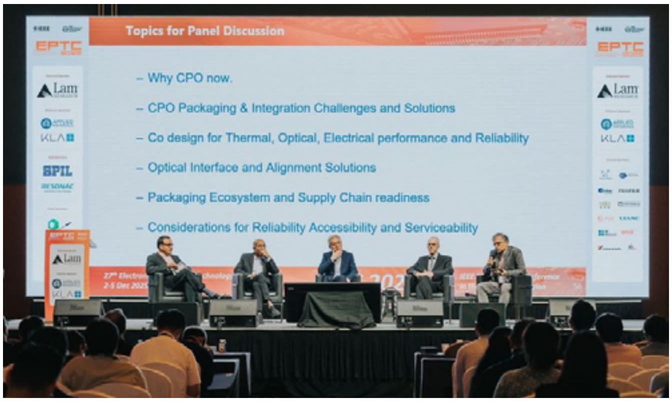
Topics for Panel Discussion

- Why CPO now.
- CPO Packaging & Integration Challenges and Solutions
- Co design for Thermal, Optical, Electrical performance and Reliability
- Optical Interface and Alignment Solutions
- Packaging Ecosystem and Supply Chain readiness
- Considerations for Reliability Accessibility and Serviceability

Panel session 2. "Co-packaged optics: The next inflection point for advanced packaging." The moderator was Dr. Surya Bhattacharya (IME, A*STAR). Panelists included: 1) Dr. Torsten Wipiejewski (Huawei); 2) Dr. Radha Nagarajan (Marvell); 3) Cindy Palar (Celestial AI); and 4) Dr. Jagadish CV (Advanced Micro Foundry).

Professional development courses (PDCs)

Day 2 opened with six professional development courses (PDCs) in the morning (see list below, including instructors' names). The courses were followed by the Electronic Packaging Society (EPS) luncheon and other technical programs in the afternoon.

1. Advanced substrates for chiplets, heterogeneous integration, and co-packaged optics: John H. Lau, Unimicron Technology Corporation, Taiwan.
2. Photonic components and packaging technologies for data center, communication, sensing, and displays: Torsten Wipiejewski, Huawei.
3. Advanced packaging for MEMS and sensors: Horst Theuss, Infineon Technologies.
4. Overview of characterization techniques for 3D heterogeneously-integrated circuit packaging: Ali Shakouri, Purdue University.
5. Current and future challenges and solutions in AI & HPC system and thermal management: Gamal Refai-Ahmed, AMD.
6. Design-on-simulation technology for reliability prediction of advanced packaging: K. N. Chiang, National Tsing Hua University, Taiwan.

Heterogeneous integration (HI) workshop

A dedicated HI Workshop was once again convened under the theme: "Interconnects – Design and Manufacturing of Complex HI Structures."

The workshop opened with an address by Kitty Pearsall, followed by technical sessions covering: 1) Design for manufacturability of >1kW Systems: Gamal Refai Ahmed (AMD); 2) Photonics design: Amr S. Helmy (University of Toronto); 3) HBI manufacturability: Loke Yuan Wong (Applied Materials); 4) TCB manufacturability: Li Ming (ASMPT); and 5) CPO manufacturability: Surya Bhattacharya (IME, A*STAR). The workshop concluded with a panel discussion entitled: "Why is HBI/CPO high-volume manufacturing adoption still slow despite all the buzz?" with moderator Wong Shaw Fong (Intel).

Technical program

The technical program at EPTC 2025 reached new milestones in both scale and diversity. There were 15 invited talks, 162 oral presentations, and 90 interactive presentations. Contributions came from 21 countries, underscoring the international nature of the conference.

Student engagement

In line with the IEEE EPS mission to nurture young professionals, 9 EPS Student Travel Grants were awarded to support student participation in EPTC 2025.

Industry participation

EPTC 2025 received strong industry support with a record 21 sponsors and 47 exhibitors. The exhibition area remained a vibrant focal point for technical exchange and networking.

Summary

With 800 participants, EPTC 2025 set a new attendance record, reflecting the growing relevance of advanced packaging, heterogeneous integration, and AI-driven system integration. The Organizing Committee expresses its sincere appreciation to all speakers, authors, sponsors, exhibitors, partners, and participants.

The 28th Electronics Packaging Technology Conference (EPTC 2026) will be held from December 1–4, 2026, at the Resorts World Sentosa in Singapore and the organizing committee looks forward to seeing you at EPTC 2026. For further updates, visit the EPTC website: www.eptc-ieee.net.

22ND INTERNATIONAL CONFERENCE & EXHIBITION ON
DEVICE PACKAGING
SHERATON GRAND AT WILD HORSE PASS PHOENIX, AZ
MARCH 2-5, 2026 • WWW.DEVICEPACKAGING.ORG

IMAPS DPC 2026: Conference wrap-up

By Jason Rouse, IMAPS Device Packaging Conference 2026 General Chair
[Taiyo America Inc.]

It is a pleasure to announce that the 22nd IMAPS Device Packaging Conference has cemented itself as a premier global electronics packaging event with engaging content, growing attendance, and valuable networking. This year, DPC saw another record number of attendees (1,090), 96 exhibitors (with a wait list), and more than 100 attendees from over 20 countries outside the US. Held for a second time at the unique Sheraton Wild Horse Pass, the conference kicked off with a full day of Professional Development Courses followed by an outdoor welcome reception at the nearby Koli Equestrian, where attendees mingled and enjoyed an incredible Arizona sunset.

Our strong technical program included keynote speakers from NVIDIA, Infineon, AMD and AGC, an evening panel on Addressing Barriers to Co-Packaged Optics and our Global Business Council session, “The Light at the End of the Fiber: Has the Promise of Silicon Photonics Finally Arrived.”

The profound effect of the artificial intelligence (AI) build-out on the electronics packaging community was evident in the keynotes from NVIDIA, Infineon, AMD, and AGC. AI scaling is causing the entire supply chain—from materials to equipment to packaging capacity—to accelerate production and quickly launch innovative market solutions. The need for higher density and efficient power conversion was highlighted in the keynote by Infineon. AGC built on the broader needs of improved optical packaging, highlighted by NVIDIA and AMD in their materials focused keynotes.

Co-packaged optics were further expanded during Tuesday evening’s panel and the Wednesday morning Global Business Council session. Speakers from Cisco, NVIDIA, Marvell, ASE, GlobalFoundries, FormFactor, and IDC, contributed valuable insights, technological steps, and challenges, as optics works to replace Cu interconnections.

Our Wednesday evening included the well-attended Poster Session where students, startups, and established members of our growing community highlighted developments of interest.

The robust social atmosphere coupled with the pleasant Phoenix weather fostered an opportunity for attendees to discuss the deep technical sessions of the prior two days. The evening concluded with the Backyard Olympics fundraising event.

The busy international exhibit hall drew in the entire industry supply chain (from equipment and materials, to manufacturing and services) and is a key complement to the global technical program.

Student participation was stronger than ever, and fun: ~131 university students joined in the technical program as attendees and speakers learned from experts in Professional Development Courses. A special student-industry workforce session saw students engaged with IMAPS leadership and corporate sponsors to learn more about our industry and the many career opportunities.

The University of Arizona Center for Semiconductor Manufacturing team enjoyed its time interacting with the conference attendees and fellow exhibitors, and telling them all about the research strengths, state-of-the-art facilities, and talent at the University of Arizona.

DPC also welcomed 26 students from Hamilton High School (Chandler, Arizona) who are enrolled in a semiconductor manufacturing CTE program—a first of its kind in the country. Partnered with the University of Arizona, the program offers diverse, hands-on pathways aimed at college credit and career readiness in high-demand technical fields. These students listened to presentations on career preparation and advanced packaging 101, followed by a tour through the exhibit hall to experience the supply chain firsthand.

"The continued momentum of the packaging community was evident in this year's program and record attendance. The explosive needs of artificial intelligence (AI) have highlighted the value of our community, and we look forward to a record 23rd edition," noted Jason Rouse, General Chair.

The Device Packaging Conference 2027 will return to the Sheraton next year, March 1-4.

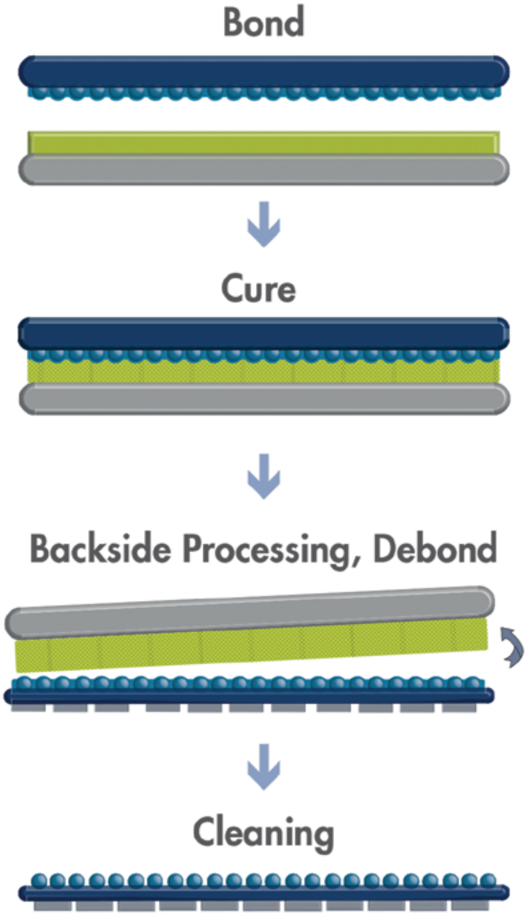

Highlighting the next-generation of advanced packaging technologies for high-performance computing and AI-enabled systems

By Bora Baloglu [ECTC 2026 Technical Program Chair]

On behalf of the Program and Executive Committees, I invite you to attend the 76th IEEE Electronic Components and Technology Conference (ECTC), which will be held May 26–29, 2026, at the JW Marriott & The Ritz-Carlton Grande Lakes Resort, Orlando, Florida. Sponsored by the IEEE Electronics Packaging Society, ECTC brings together more than 2,000 professionals from across the global microelectronics packaging industry, including manufacturers, design houses, foundries, material suppliers, universities, and investors.

This year marks the 76th year for the semiconductor industry's premier electronics packaging conference with expanded special sessions on critical technologies, and a series of professional development courses.

ECTC will deliver a technical program of more than 450 papers in 41 technical sessions, including five interactive presentations—one of which is a student session; 12 special sessions on critical technologies; a series of 16 continuing education unit (CEU)-approved professional development opportunities; and more than 135 exhibits representing industry-leading product and service companies from around the world. In addition, the conference provides comprehensive student engagement activities and multiple social events for networking opportunities.

ECTC history

ECTC began in 1950 as the Symposium on Improved Quality Electronic Components, held at the U.S. Department of the Interior and sponsored by the former American Institute of Electrical Engineers (AIEE), Institute of Radio Engineers (IRE), and Electronic Industries Association (EIA). In succeeding years, the conference has evolved, with various name changes, technical programs aligned with advanced electronics technologies, different locations, and co-sponsors. Today, the ECTC conference serves as a global platform for exploring leading-edge advancements in microelectronic packaging and component technologies, fostering innovation, and addressing industry challenges.

Topics in the ECTC 2026 program include advanced packaging technologies such as, wafer-level and fan-out packaging, 2.5D, 3D and heterogeneous integration, interposers, advanced substrates, materials modeling, reliability, interconnections, packaging for high-speed and high-bandwidth, photonics, quantum electronics, as well as flexible and printed electronics.

Keynote presentation to highlight advanced packaging technologies

On Wednesday, May 27, Dr. Tien Wu, CEO of Advanced Semiconductor Engineering (ASE), will present the keynote address, "Advanced packaging & the future of system optimization." This keynote will address how global artificial intelligence (AI) infusion is redefining performance, power, and integration requirements, with a key role played by advanced packaging at the leading edge of semiconductor innovation. As chip architecture complexity increases, the industry is progressing beyond device-level scaling toward system-level optimization. Advances in heterogeneous integration and packaging-enabled co-design are shaping more efficient, scalable, and resilient systems. Moving forward, system-centric strategies that align architecture, packaging, and collaboration across the ecosystem will be paramount to shaping the golden era of AI and semiconductors.

ECTC technical session

ECTC Luncheon presentation

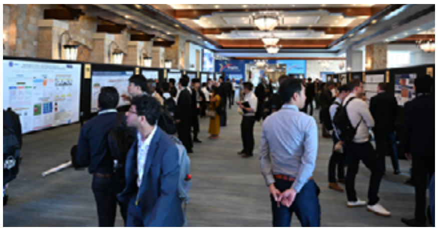

ECTC Interactive poster session

Special sessions to showcase critical technologies

ECTC 2026 includes a series of 12 special sessions that feature industry experts discussing technology status and roadmaps in key areas of interest, including: 1) Quantum infrastructure for AI applications; 2) New packaging technologies enabled by panel-level integration; 3) AI-enabled electronic design automation; and 4) Photonics-based systems for exascale computing. Other special sessions will cover: 1) System integration challenges for high-power components in AI applications; 2) Electrical-thermal-mechanical co-design; 3) Wafer-to-panel: Next-generation advanced packaging technology; and 4) Innovative materials for advanced packaging.

An IEEE EPS Seminar will be presented on, “Redefining system integration – The rise of organic substrates in the chiplet era.” Also, a plenary session will cover the challenges in data center power, “Efficiency is not enough – Are we solving the wrong problem in the data center energy use?” Finally, the IEEE EPS President’s Panel will also cover data centers with a discussion entitled, “Data centers in the age of AI – Challenges and solutions.”

Student competition and engagement programs

On Wednesday, six winning student teams will have the opportunity to attend ECTC 2026 with financial assistance. Students in Bachelors and Masters programs will compete on the topic: “Low-cost robust thermal solution for high power AI/datacenter processors;” while PhD students will compete in one of two topics: “Materials, interfaces, and processes for ultra-scalable interconnects,” or “Electromigration solutions for BGA interconnects in AI-focused packages.”

Startup competition

On Thursday, start-up companies will pitch their innovative ideas to a panel on the theme, “The light age: Strategic investment in photonics to power the next computing era.” Their presentations will be followed by an audience Q&A, jury deliberation, award presentation, and a networking event.

Professional development courses

Tuesday’s program includes a series of 16 professional development courses (PDCs). These four-hour courses on relevant electronic packaging topics are taught by world-class experts, enabling participants to broaden their technical knowledge base. Attendees will be awarded either CEU or professional development hours (PDH) credits. These courses are offered in conjunction with the co-located IEEE ITherm Conference, which focuses on thermal/thermomechanical issues in electronic systems.

Comprehensive technical program

ECTC 2026 will feature some 400 technical papers, presented in 41 technical sessions, including five interactive presentations, one of which is a student session. Authors from more than 20 countries will share their latest research on topics including: wafer- and panel-level packaging, co-packaged optics, hybrid bonding technology, pitch scaling with advanced bonding, thermal management, reliability, heterogeneous integration, wafer-to-wafer and chip-to-wafer bonding, novel substrate materials, AI and high-performance computing, signal integrity, advanced characterization and modeling, mmWave technology, redistribution layers (RDLs) and fan-out interconnects, 3D integration, digital twin and AI for advanced packaging, glass/ceramic/silicon substrates, advanced interconnects, power delivery, flexible electronics, and emerging technologies.

Exhibits

ECTC 2026 will feature exhibits on Wednesday and Thursday that showcase cutting-edge technologies and products from more than 135 leading companies in electronic components, materials, packaging, equipment and services. The exhibition hall provides excellent opportunities for networking during coffee breaks, luncheons, and evening receptions.

The conference also offers unique experiences for engineers, managers, students, and executives in the microelectronics packaging and components industry. Make plans now to attend the 76th annual ECTC and take part in these exciting technical and professional opportunities. On behalf of our sponsors, exhibitors, authors, speakers, PDC instructors, session chairs, and program committee members, as well as all the volunteers who help make the ECTC conference a success, we look forward to meeting you at the JW Marriott & The Ritz-Carlton Grande Lakes Resort, Orlando, Florida.

SPRING 2026 ADVERTISER INDEX

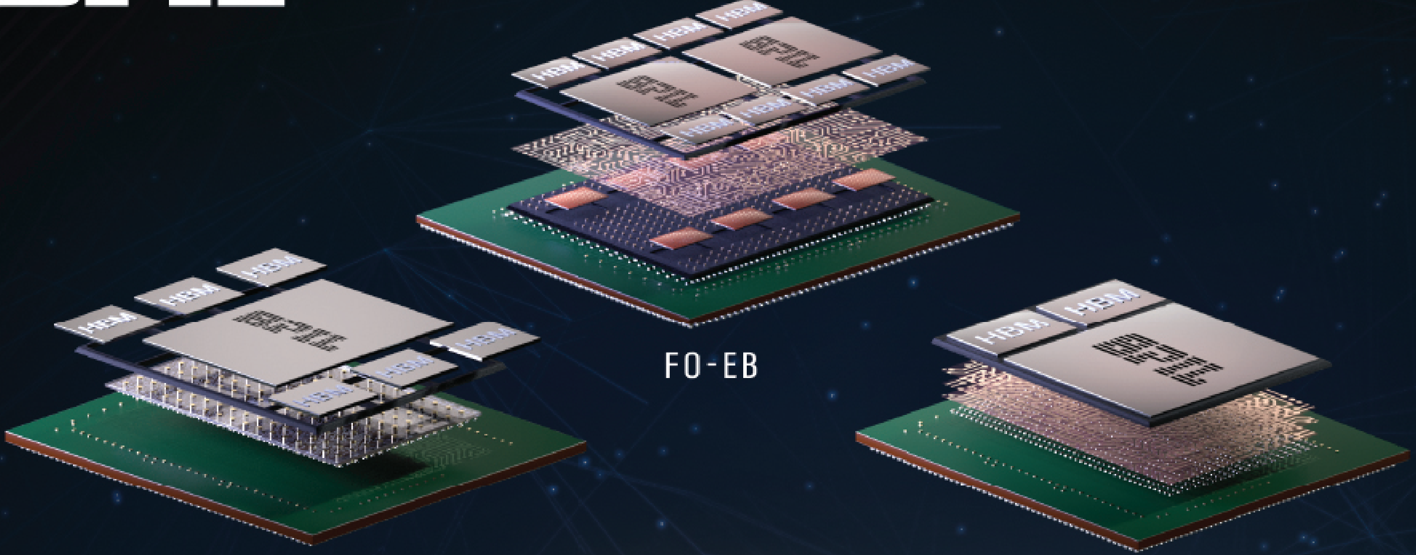
FO-EB
2.5D IC
FO-MCM