AP ADVANCED PACKAGING

SPECIAL EDITION

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

ADVANCEDPACKAGING.NEWS

EMPOWERING ADVANCED FCBGA SUBSTRATE DESIGN RULES FOR AUTOMOTIVE APPLICATIONS

EVG®

invent
innovate
implement

www.EVGroup.com

FUSION AND HYBRID BONDING FOR HETEROGENEOUS INTEGRATION

- Enabling advanced 3D device stacking and chiplet integration for CMOS image sensors, memory and 3D system-on-chip (SoC)
- High-volume production equipment and process technology for Wafer-to-Wafer (W2W) and Die-to-Wafer (D2W) hybrid bonding
- GEMINI® FB Automated Production Wafer Bonding System delivers industry leading alignment accuracy and bonding performance
- Heterogeneous Integration Competence Center™ serving as leading-edge innovation incubator for EVG customers and partners

GET IN TOUCH to discuss your manufacturing needs
www.EVGroup.com

VIEWPOINT

BY SARAB CHOPRA
EDITOR

Advancing the Future of Advanced Packaging

Welcome to the first issue of Advanced Packaging, the newest addition to the Angel Business Communications portfolio. Joining our established titles, Compound Semiconductor, Silicon Semiconductor, PIC Magazine, and Power Electronics, this publication is dedicated to one of the most important and fast-moving areas in modern electronics. As traditional device scaling slows down, system-level performance is becoming the real measure of progress, and advanced packaging is now right at the centre of that shift.

Packaging is no longer something that happens at the end of the process. It is now a key part of how performance, efficiency, and integration are achieved. The move toward heterogeneous integration is changing the way systems are designed, making it possible to combine different process nodes, materials, and functions in a single package. Technologies such as 2.5D interposers, 3D stacking, and chiplet architectures are opening the door to higher bandwidth, better power efficiency, and more compact designs. These are exactly the capabilities needed for applications like AI, high-performance computing, automotive systems, and edge devices.

A big part of this shift comes from breaking systems into smaller building blocks. Instead of relying on one large, complex chip, designers are turning to chiplets that each handle a specific function. This approach improves yield, lowers costs, and speeds up development. It also gives much more flexibility, allowing companies to combine logic, memory, RF, and even photonics in ways that were not practical before. At the same time, it helps build a more adaptable and resilient supply chain.

Interconnects are becoming just as important as the chips themselves. As data rates continue to rise and energy efficiency becomes more critical, older packaging approaches are starting to fall short. New techniques such as through-silicon vias, hybrid bonding, and very fine pitch redistribution layers are helping to push interconnect density further while maintaining signal integrity. These advances reduce latency, increase bandwidth, and cut power loss, which is essential for today's data-heavy workloads.

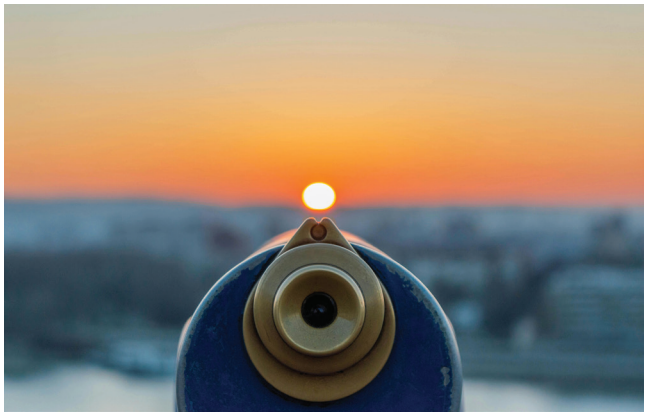

Materials themselves are playing a bigger role than ever. From low-loss dielectrics to high-conductivity metals and advanced substrates, every layer in the package has to be carefully engineered. Bringing all of these materials together in complex structures requires precise manufacturing and close collaboration across the supply chain.

Testing and inspection are also evolving to keep up. Ensuring that each chiplet is a known good die before assembly is critical, and new testing strategies are needed to check not only individual components but how they work together inside the package. As systems become more complex, so too must the methods used to validate them.

Advanced packaging is no longer just supporting innovation; it is driving it. By connecting design choices directly to system performance, it is reshaping the semiconductor landscape and enabling the next generation of technology. As we launch this publication, we look forward to exploring the ideas, challenges, and breakthroughs that will define the future of advanced packaging.

Contents

Cover Story

Empowering advanced FCBGA substrate design rules for automotive applications

Autonomous systems are forcing substrate design into new territory. As copper traces shrink, reliability becomes the real battleground for next-generation automotive packaging.

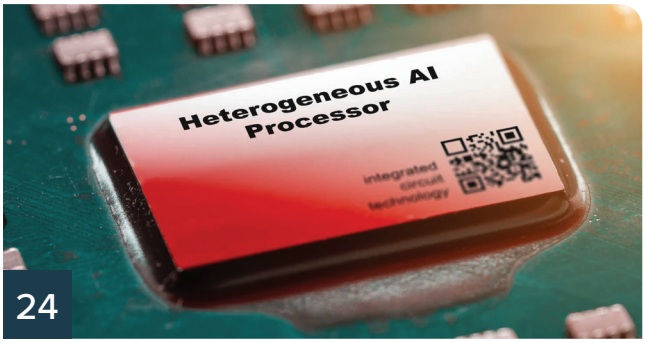

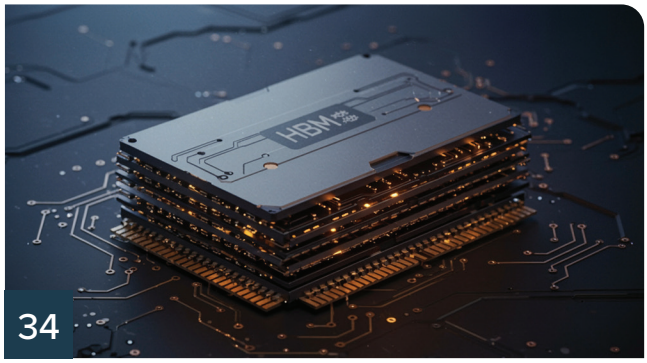

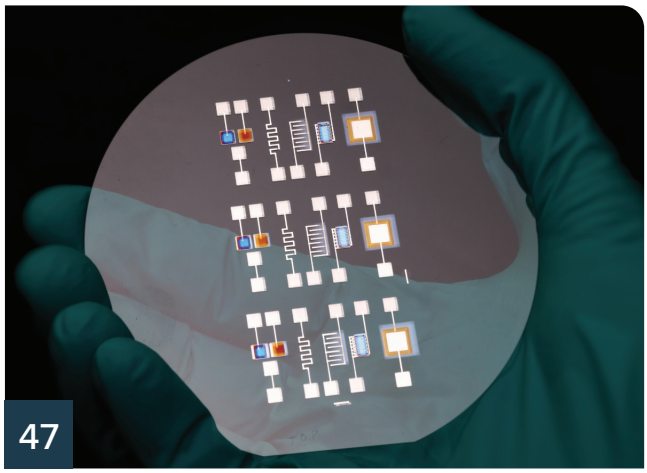

NEWS

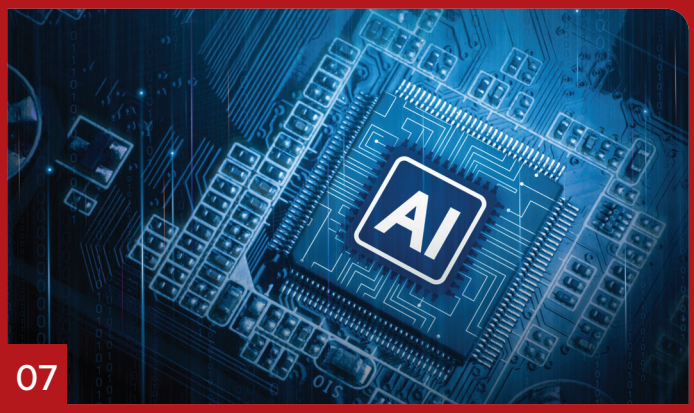

ADVANCED PACKAGING

Editor
Sarab Chopra sarab.chopra@angelbc.com
Contributing Technical Editor
Richard Stevenson richard.stevenson@angelbc.com +44 (0)1923 690215
Sales & Marketing Manager
Shehzad Munshi shehzad.munshi@angelbc.com +44 (0)1923 690215
Design & Production Manager
Mitch Gaynor mitch.gaynor@angelbc.com +44 (0)1923 690214
Publisher
Jackie Cannon jackie.cannon@angelbc.com +44 (0)1923 690205

Sales and Product Manager
James Cheriton james.cheriton@angelbc.com +44 (0)2476 718970

CEO Sukhi Bhadal sukhi.bhadal@angelbc.com +44 (0)2476 718970
CTO Scott Adams scott.adams@angelbc.com +44 (0)2476 718970

Published by
Angel Business Communications Ltd, 6 Bow Court, Fletchworth Gate, Burnsall Road, Coventry CV5 6SP, UK. T: +44 (0)2476 718 970
E: info@angelbc.com W: ADVANCEDPACKAGING.NEWS

Angel BUSINESS COMMUNICATIONS

Advanced Packaging is published four times a year on a controlled circulation basis. Non-qualifying individuals can subscribe at: £105.00/e158 pa (UK & Europe), £138.00 pa (air mail), $198 pa (USA). Cover price £4.50. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher. Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing.US mailing information: Advanced Packaging, ISSN 1096-598X, is published 4 times a year, March, May, August and December by Angel Business Communications Ltd, Unit 6, Bow Court, Fletchworth Gate, Burnsall Rd, Coventry CV5 6SP. UK. The 2026 US annual subscription price is $198. Airfreight and mailing in the USA by agent named Air Business Ltd, c/o Worldnet Shipping Inc., 156-15, 146th Avenue, 2nd Floor, Jamaica, NY 11434, USA. Periodicals postage paid at Jamaica NY 11431. US Postmaster: Send address changes to Advanced Packaging, Air Business Ltd, c/o Worldnet Shipping Inc., 156-15, 146th Avenue, 2nd Floor, Jamaica, NY 11434, USA. Printed by: The Manson Group. We strive for accuracy in all we publish; readers and contributors are encouraged to contact us if they recognise an error or omission. Once a magazine edition is published [online, in print or both], we do not update previously published articles to align old company names, branding, marketing efforts, taglines, mission statements or other promotional verbiage, images, or logos to newly created or updated names, images, typographic renderings, logos (or similar) when such references/images were accurately stated, rendered or displayed at the time of the original publication. When companies change their names or the images/text used to represent the company, we invite organizations to provide Angel Business Communications with a news release detailing their new business objectives and/or other changes that could impact how customers/prospects might recognise the company, contact the organisation, or engage with them for future commercial enterprise.

Tesla eyes packaging entry

Elon Musk's Terafab vision may sidestep advanced-node manufacturing challenges by targeting advanced packaging, potentially reshaping semiconductor supply chains without directly challenging TSMC's technology lead.

TESLA'S proposed Terafab mega-fab is raising questions across the semiconductor industry, particularly around its potential impact on foundry leader TSMC.

While the vision aims to integrate logic, memory, and packaging, industry observers say competing at cutting-edge nodes such as 2nm would be highly challenging due to the complexity of yield control, process integration, and equipment requirements.

At advanced nodes, the transition from FinFET to GAAFET transistor architectures demands significant upgrades across materials, tools, and manufacturing processes.

Combined with reliance on scarce EUV lithography systems and deep engineering expertise, these barriers reinforce TSMC's entrenched position in leading-edge fabrication.

Instead, attention is shifting to advanced packaging as Tesla's most viable entry point into the semiconductor value chain.

With AI chip production increasingly constrained by packaging capacity, Tesla could target back-end processes through in-house development, panel-level packaging, or partnerships with established players.

Such a move would allow the company to secure supply for its own AI and automotive chips while avoiding the steepest technical hurdles.

Terafab is expected to focus on chips for electric vehicles, humanoid robotics, and space-based AI systems, reflecting Tesla's broader ambitions across mobility and computing.

However, Musk has indicated that Tesla and its affiliated companies will continue sourcing chips from existing suppliers, including TSMC, Samsung, and Micron, while encouraging capacity expansion.

In the near term, Terafab is unlikely to challenge TSMC's leadership in advanced-node manufacturing. Over the longer term, however, even partial in-house capabilities could strengthen Tesla's bargaining power within the semiconductor supply chain, signalling a shift toward greater vertical integration among major technology players.

External light sources emerge as key to scaling CPO

EXTERNAL LIGHT SOURCES (ELS) are emerging as a critical enabler for scaling co-packaged optics (CPO) in next-generation AI datacentres, as the industry looks to overcome mounting thermal, reliability and integration challenges.

As AI workloads continue to expand, traditional electrical interconnects are increasingly unable to meet the bandwidth and energy efficiency demands of modern data centre architectures.

CPO has been widely identified as a solution, integrating optical components directly with GPUs and switch packages to deliver high-speed, low-power connectivity. However, the approach introduces new challenges, particularly around thermal management and laser stability.

ELS architectures address this issue by relocating temperature-sensitive lasers away from high-power processing units. This separation improves wavelength stability, enhances system reliability and simplifies serviceability, making ELS an increasingly attractive approach for both scale-out and scale-up optical systems.

Recent industry activity highlights growing momentum behind this shift. Sivers Semiconductors, O-Net Technologies and Enablence Technologies have announced a collaboration to develop an advanced ELS module designed to support CPO deployment in AI and high-performance computing environments.

The solution combines laser arrays with photonic distribution technologies to enable scalable, high-density optical interconnects.

The move reflects a broader industry trend, as data centre operators and technology providers seek practical pathways to scale optical connectivity without compromising performance or efficiency.

With market forecasts pointing to rapid growth in CPO adoption over the coming decade, ELS solutions are expected to play a central role in enabling reliable and manufacturable architectures.

As AI infrastructure continues to evolve, the shift towards ELS-based designs signals a growing recognition that optical innovation at the system level will be essential to sustaining performance gains in increasingly complex computing environments.

AI chip growth fueled by advanced packaging innovations

Advanced packaging is emerging as a key semiconductor growth driver, with AI demand boosting opportunities in high bandwidth memory and co packaged optics.

ADVANCED PACKAGING is attracting attention as a major growth driver in the semiconductor industry, shifting focus from traditional transistor shrinking to how components like memory and logic dies are integrated for higher performance.

Industry analysts highlight opportunities in co-packaged optics and high-bandwidth memory, technologies that enhance chip efficiency for AI and high-performance computing applications.

Companies well-positioned to benefit include BE Semiconductors, Teradyne (NASDAQ: TER), and Taiwan Semiconductor (NYSE: TSM).

According to Eric Bleeker and Austin Smith of the AI Investor Podcast, these technical areas are expected to generate substantial revenue in 2026 and beyond.

Co-packaged optics alone could contribute billions in sales as demand grows for heterogeneous integration and advanced interconnect packaging.

Investors are closely monitoring these developments, as advanced packaging offers a pathway to participate in high-growth segments of the semiconductor market, including AI accelerators, high-bandwidth memory, and next-generation photonics.

Co-packaged optics alone could contribute billions

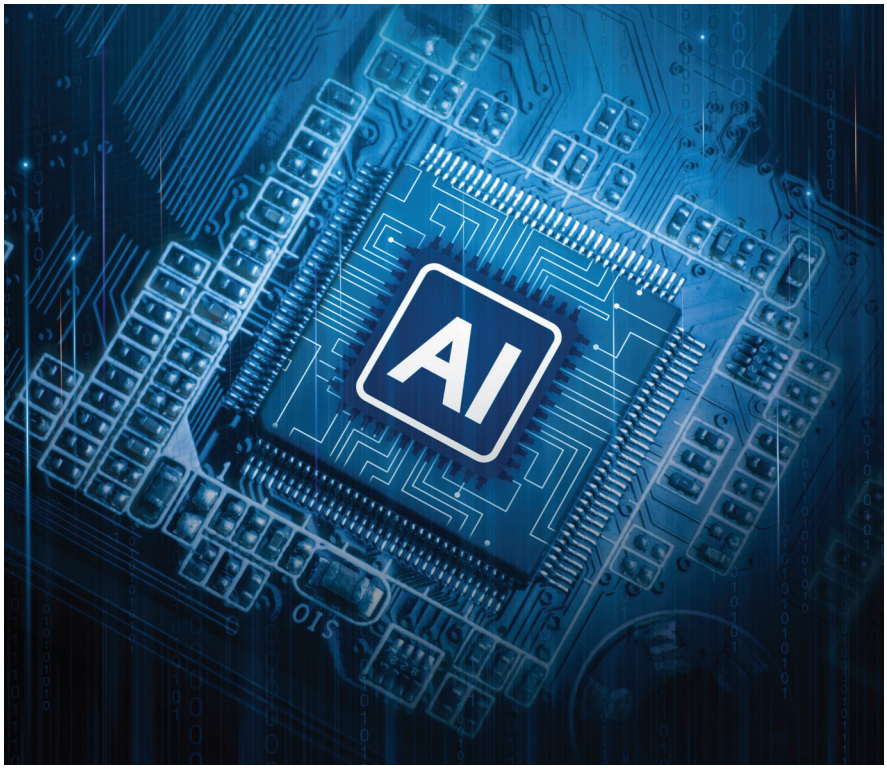

Rapid growth set for advanced packaging

THE ADVANCED PACKAGING technologies market is poised for significant expansion as the semiconductor industry continues to evolve, driven by rising demand across multiple sectors.

According to a report by The Business Research Company, the market is expected to reach $15.82 billion by 2030, growing at a robust CAGR of 14.6%.

Growth is fueled by increased adoption of heterogeneous integration, rising local semiconductor production, and expanding demand from automotive electronics.

Key trends shaping the market include 3D IC stacking, fan-out wafer-level packaging, system-in-package (SiP) solutions, miniaturisation, enhanced thermal management, and improved design reliability.

These innovations are helping manufacturers create higher-performance, energy-efficient chips suitable for advanced computing, AI applications, and compact electronics.

Analysts note that as semiconductor design complexity increases, advanced packaging technologies will remain a critical driver of performance, efficiency, and industry growth.

Malaysia advances IC design and packaging

Malaysia is stepping up its semiconductor strategy, focusing on IC design, advanced packaging, and high-value technologies. Initiatives by Mimos Bhd and MOSTI aim to develop talent, strengthen IP, and support commercialisation, shifting the country from assembly-focused manufacturing to technology ownership.

THE MINISTRY of Science, Technology and Innovation (MOSTI), through Mimos Bhd, is accelerating efforts in integrated circuit (IC) design, advanced packaging, and other high-value semiconductor technologies.

The initiative is intended to help local companies transition into original design manufacturers (ODMs) and technology owners, reducing reliance on assembly-based operations.

Mimos has implemented strategies spanning R&D infrastructure, intellectual property development, technology transfer, and workforce training.

Programs through Mimos Academy include IC design, wafer fabrication, reliability engineering, and artificial intelligence, targeting the manpower needs of Malaysia's growing technology sector.

Shared facilities at Mimos' Semiconductor Technology Centre and Industry Technology Innovation Centre (ITIC) in Kulim provide wafer fabrication, IC testing, materials analysis, reliability testing, and co-design capabilities.

ITIC conducted 19 technical training sessions and seminars in 2025, engaging roughly 400 participants from 60 industry and academic institutions.

Proof-of-concept projects with local electronics companies provide access to shared equipment and technical expertise, strengthening the domestic value chain and accelerating commercialisation.

With over 1,400 patents filed, including 330 in semiconductors, Mimos supports technology transfer through licensing, co-development, and spin-offs.

Beyond semiconductors, MOSTI also backs startups via Cradle Fund Sdn Bhd, offering CIP Spark and CIP Sprint grants for prototype development and commercialisation.

Between 2020 and 2025, 101 products were commercialised under the Malaysian Research Accelerator for Technology and Innovation (MRANTI).

These efforts collectively aim to position Malaysia as a hub for advanced packaging, IC design, and semiconductor innovation, creating a robust ecosystem for technology-driven growth.

Advanced packaging market to hit $31.8B by 2032

THE GLOBAL advanced semiconductor packaging market is entering a robust growth phase, driven by demand for higher-performance, power-efficient, and compact semiconductor applications.

According to the latest analysis by QYResearch, the market is projected to expand from US$ 19.31 billion in 2025 to US$ 31.81 billion by 2032, representing a compound annual growth rate (CAGR) of 7.5%.

Advanced packaging encompasses a variety of integration techniques beyond traditional methods, including 2.5D integration, 3D integrated circuits (3D-IC), fan-out and fan-in wafer-level packaging, flip chip, and system-in-package (SiP) technologies.

These approaches enable higher interconnect density, improved electrical performance, and enhanced thermal management, making them critical for next-generation semiconductors.

Analysts say that as applications like AI accelerators, high-performance computing chips, and compact consumer electronics continue to demand higher efficiency and smaller form factors, advanced packaging will remain a key driver of industry innovation and market expansion.

The market is projected to expand from US$ 19.31 billion in 2025 to US$ 31.81 billion by 2032, representing a compound annual growth rate (CAGR) of 7.5%.

CPO advances AI infrastructure scaling

Integration of photonics with computing is set to overcome bandwidth and energy limits in next-generation data centres.

THE RAPID GROWTH of artificial intelligence is accelerating the shift toward large-scale, distributed compute architectures, driving demand for new interconnect technologies such as co-packaged optics (CPO).

CPO integrates optical engines closer to switch and processor chips, replacing traditional front-panel transceivers to reduce signal loss, improve bandwidth density, and lower power consumption. This approach is emerging as a key enabler for scaling AI infrastructure, particularly as copper interconnects reach their physical limits.

Recent advances in electronic and photonic integration have already delivered significant gains, including a 64-fold increase in bandwidth density and a fivefold improvement in energy efficiency over the past five years. However, challenges remain in bringing CPO to high-volume manufacturing.

A key issue is the mismatch between the bandwidth density of AI chiplets and current optical interconnects, highlighting the need for more compact and scalable photonic integrated circuits. Fiber array design, wavelength stability, and thermal management also present ongoing technical hurdles.

Manufacturing complexity is another barrier, with CPO systems requiring the integration of lasers, photonic circuits, and fiber arrays across multiple process nodes.

Active alignment techniques, while necessary for performance, can limit throughput and increase costs. In addition, conventional permanent bonding methods reduce repairability, prompting industry efforts to develop detachable fiber array solutions.

Testing and yield optimisation are also critical concerns. The adoption of “known-good” optical engines, verified through wafer-level testing prior to final assembly, is seen as essential for improving manufacturing efficiency.

However, not all integration approaches are compatible with current testing methods, particularly in advanced 3D configurations.

Despite these challenges, co-packaged optics is widely viewed as a cornerstone technology for future AI and hyperscale data centre systems.

Continued progress will depend on close collaboration across the semiconductor and photonics ecosystem, spanning design, packaging, and system integration.

Japan’s advanced packaging market to rocket by 2033

JAPAN’S ADVANCED packaging substrate market is on track for strong growth as semiconductor makers adopt more complex chip designs.

Valued at USD 2.18 billion in 2024, the market is projected to reach USD 3.56 billion by 2033, growing at a CAGR of 6.2%.

Advanced packaging substrates are critical components in modern semiconductor devices, providing both structural support and reliable electrical pathways that allow chips to interface efficiently with electronic systems.

As semiconductor devices continue to shrink while increasing in computational power, the demand for high-density, high-performance packaging solutions is rising.

The market’s expansion is being driven by the proliferation of advanced computing technologies, including AI accelerators, high-performance processors, and next-generation memory modules, all of which require substrates capable of handling higher interconnect density and improved thermal performance.

In addition, Japan’s semiconductor ecosystem, with its strong manufacturing base and focus on quality, positions it as a key player in advanced packaging innovation.

With manufacturers increasingly integrating heterogeneous chip components and adopting 2.5D and 3D packaging techniques, the need for robust and scalable substrates has never been greater.

Analysts note that continued R&D investment and collaboration across material suppliers, OSATs, and semiconductor fabs will be essential to sustain growth and maintain Japan’s competitive edge in the global advanced packaging market.

The forecast underscores the growing importance of advanced packaging technologies in supporting the next generation of high-performance, energy-efficient, and compact semiconductor devices.

Valued at USD 2.18 billion in 2024, the market is projected to reach USD 3.56 billion by 2033

Grand Process Technology & Taiwan Tech team up on advanced packaging

Grand Process Technology and National Taiwan University of Science and Technology have launched a five-year partnership to develop AI-driven semiconductor advanced packaging technologies while strengthening equipment innovation and industry talent development.

GRAND PROCESS TECHNOLOGY has signed a five-year cooperation agreement with National Taiwan University of Science and Technology to establish an advanced semiconductor packaging platform focused on equipment innovation, materials development, and talent training.

The Hsinchu-based semiconductor equipment and chemical materials supplier said the collaboration will support next-generation packaging technologies driven by growing demand for artificial intelligence (AI) and high-performance computing. The partners plan to invest NT$50 million (US$1.6 million) in joint research and development.

Grand Process Technology provides equipment and materials used in advanced packaging processes such as chip-on-wafer-on-substrate (CoWoS), with customers including Taiwan Semiconductor Manufacturing Company and Advanced Semiconductor Engineering.

Under the agreement, the university and industry partner will collaborate on advanced electroplating and X-ray inspection technologies aimed at improving process stability, materials reliability, and manufacturing yields.

The initiative will also include scholarships and hands-on industry training opportunities to strengthen Taiwan's semiconductor talent pipeline.

Grand Process Technology Chair Chang Hung-tai said increasing AI computing requirements are accelerating the shift from single-step packaging processes toward integrated system-level solutions, driving demand for more precise equipment and optimised material design.

The university added that joint R&D outcomes are expected to generate new patents, combining Grand Process Technology's existing portfolio of 149 patents with academic research strengths in semiconductor manufacturing, silicon photonics, integrated circuit design, and advanced materials.

Grand Process Technology reported revenue of NT$6.51 billion last year, marking nearly 60% year-on-year growth, supported by rising advanced packaging demand and order visibility extending into the first half of 2026.

ASE expects advanced packaging sales to double on AI demand

ASE TECHNOLOGY Holding, the world's largest outsourced semiconductor assembly and test provider, expects its advanced packaging revenue to double in 2026 to approximately $3.2 billion, driven by strong demand from AI chipmakers.

The company also plans to increase capital expenditure this year beyond the $5.5 billion invested in 2025, as it expands cleanroom capacity and builds new facilities to support advanced packaging growth.

Much of the anticipated expansion is linked to outsourcing from TSMC, which has faced capacity constraints amid surging AI chip demand.

As leading-edge packaging requirements for AI processors intensify, more on-substrate and full-process packaging work is expected to shift to OSATs.

Analysts indicate that packaging for GPUs from Nvidia could be a primary growth driver in 2026, with a stronger ramp anticipated in the second half of the year as next-generation products enter volume production.

Advanced packaging opportunities are also emerging from Advanced Micro Devices, particularly as multiple CPU and non-GPU products transition to 2.5D architectures.

Industry observers suggest full-process packaging could begin contributing meaningful revenue in the latter half of 2026, with further upside in 2027.

Additional potential upside may come from custom AI accelerators, including programs linked to hyperscale customers.

With AI processors increasingly relying on complex 2.5D integration, high-density substrates and heterogeneous integration, ASE's expansion signals a broader structural shift in the advanced packaging landscape, as OSATs take on a greater share of high-value packaging traditionally concentrated within foundries.

Foundry 2.0 fuels 17% growth

The rise of Foundry 2.0 is reshaping the semiconductor industry, as AI demand and advanced packaging drive integrated manufacturing growth.

THE SEMICONDUCTOR INDUSTRY is experiencing a structural transformation as the Foundry 2.0 era gains commercial traction, with global foundry revenues rising 17% year-on-year in Q3 2025 to approximately $84.8 billion, according to a report by Counterpoint Research.

Unlike traditional wafer-only foundries, Foundry 2.0 encompasses pure-play fabs, non-memory IDMs, OSAT firms, and photomask suppliers, reflecting the increasing importance of integrating manufacturing, advanced packaging, and system-level design to meet the demands of AI and high-performance computing (HPC) workloads.

The growth is largely driven by sustained demand for AI accelerators, GPUs, and high-value compute silicon, supported by front-end wafer fabrication and back-end advanced packaging technologies such as chip-on-wafer-on-substrate (CoWoS).

China's domestic vendors also contributed to the expansion, benefiting from local subsidies and policy support.

Leading the charge is Taiwan Semiconductor Manufacturing Company (TSMC), which posted Q3 revenue of around $33.1 billion, exceeding prior guidance, fueled by strong adoption of its 3nm processes and high utilisation of 4/5nm capacity.

Analysts note that as silicon complexity increases, advanced packaging capabilities are becoming strategic differentiators, enabling integrated solutions that improve performance and energy efficiency.

However, tight capacity at leading-edge nodes, geopolitical risks, and supply chain dependencies could temper growth in upcoming quarters.

Overall, the surge underscores how AI-driven workloads and sophisticated packaging are reshaping investment and production priorities across the semiconductor ecosystem, marking a clear shift toward a more integrated Foundry 2.0 model.

KLA gains from packaging boom

KLA is seeing strong growth as AI demand drives advanced chip packaging, though competition and valuation concerns remain.

KLA is seeing strong growth as AI demand drives advanced chip packaging, though competition and valuation concerns remain. KLA Corporation is benefiting from rising demand for advanced semiconductor packaging as AI and high performance computing push chip complexity higher.

The company reported advanced packaging systems revenue of about $950 million in 2025, up more than 70 percent year on year. It expects further growth in 2026 in the mid to high teens, supported by demand for process control tools.

Across the industry, the wafer fabrication equipment market is forecast to reach the low $120 billion range in 2026. Advanced packaging is expected to account for around $12 billion, reflecting its growing role in chip production.

KLA faces competition from ASML and Applied Materials, both of which are also benefiting from AI driven demand.

Shares in KLA have risen about 26 percent this year, outperforming the broader technology sector. The company expects third quarter fiscal 2026 revenue of around $3.35 billion but flagged supply constraints and tariffs as near term pressures.

Analysts expect continued earnings growth, though some caution that the stock trades at a premium to the wider sector.

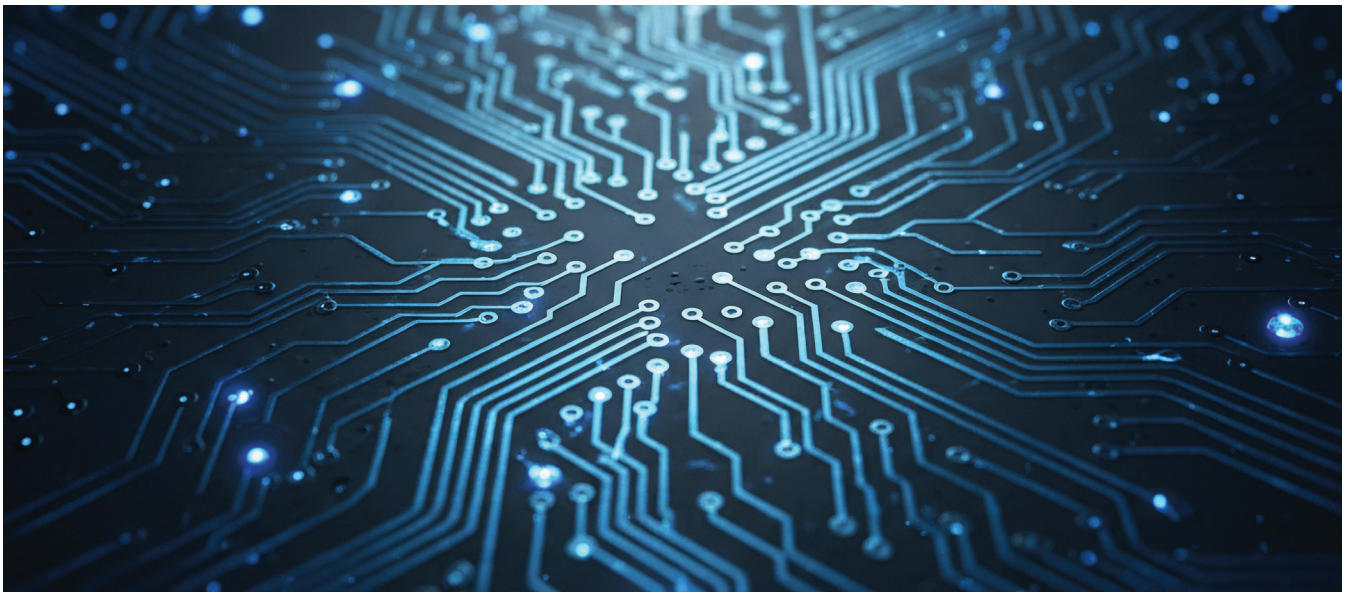

Empowering advanced FCBGA substrate design rules for automotive applications

Autonomous systems are forcing substrate design into new territory. As copper traces shrink, reliability becomes the real battleground for next-generation automotive packaging.

BY JAIMAL WILLIAMSON, TEXAS INSTRUMENTS

HISTORICALLY, semiconductor packaging had a moniker as the stepchild of the semiconductor ecosystem, with back-end assembly and packaging processes relegated as an afterthought or modest necessity for innovation, especially when compared to front-end wafer fabrication processes that embodied Moore's law and high technological investment. Now, back-end packaging has risen like a phoenix and taken its place within the semiconductor pantheon as a quintessential differentiator, along with the Internet of Things and artificial intelligence (AI). High-density interconnect routing enabled through flip-chip ball grid array (FCBGA) substrates (Figure 1) is now commonplace to meet trillions of operations per second (TOPS) in devices.

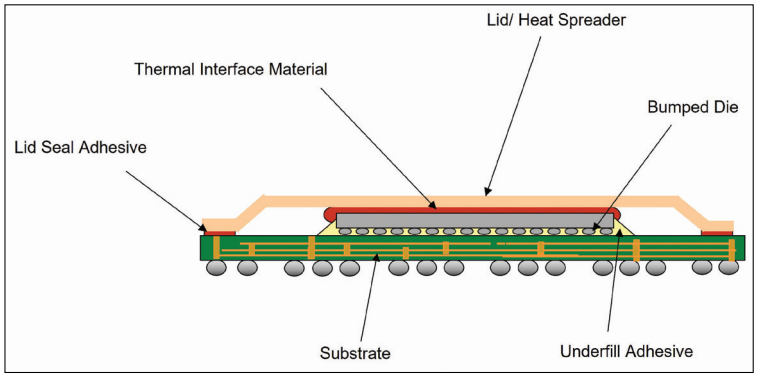

➤ Figure 1: Example of an FCBGA package

For example, devices with high electrical functionality combine serializer/deserializer and double-data-rate memory technologies, laying the groundwork for high-density interconnects routed in substrates for a myriad of applications. Applications supporting AI, autonomous driving, aerospace and space, data center networking, and high-performance computing (HPC) require TOPS to satisfy demanding operational specifications. Advanced packaging design rules (that is, routed in FCBGA substrates) maximize interconnect density through the distribution of fine copper (Cu) traces and spaces (the distance to adjacent flip-chip bond pads or traces), augmenting electrical performance. Ultimately, implementing advanced FCBGA substrate design rules can significantly impact silicon entitlement, package form factor and cost.

Advanced substrate design rules can push the boundaries of substrate supplier capabilities where latent process defects or shifts can impact package assembly manufacturing and reliability yields. For example, Cu traces yielded at the low end of the width specification can present challenges of low crack resistance during temperature cycling because of a reduced cross-sectional area [1] and local coefficient of thermal expansion mismatches. On the other hand, Cu traces that yield at the high end of the width specification can present

obstacles during temperature-humidity-bias testing, since a wider trace is closer to an adjacent trace or pad. This scenario increases the potential for ion migration.

Understanding the impact of advanced substrate design rules during component-level reliability (CLR) and board-level reliability (BLR) becomes even more paramount for automotive applications. With Tier 1 suppliers and original equipment manufacturers driving the enviable requirement of zero defects for automotive applications, a challenge for packaging engineers is to dig deep into the fundamentals of materials science. This entails understanding the chemical structure-to-transient material property relationships of organic material sets to thermomechanical changes in various metallic material properties within the package construction.

Examining the implementation of advanced FCBGA package design rules in literature and across substrate supplier and outsourced semiconductor assembly and test roadmaps, the proliferation of devices featuring finer Cu lines and space for automotive devices is lacking. For example, high-performance FCBGA devices such as 2.5D devices that use high-density routing through organic or silicon interposers for HPC applications are not subjected to the more stringent qualification conditions defined by the Automotive Electronics Council (AEC) Q-100 standard [2]. High-interconnect-density substrate routing is a prerequisite for FCBGA package design in order to power the demand for ultra-functionality on silicon and packages supporting different levels of driving automation.

As higher levels of driving automation increase from Society of Automotive Engineers J3016 standard Levels 2 and 3 (advanced driver assistance systems) to Level 4 and 5 (autonomous driving), the implementation of advanced substrate design rules becomes necessary to accommodate TOPS in devices. Yet there is a gap in exploring advanced substrate designs featuring finer Cu trace lines and space for automotive applications.

For this reason, Texas Instruments conducted a detailed study to extract package-level reliability data empirically. The study focused on FCBGA package reliability to understand the risks associated with finer Cu lines and space from an automotive chip-to-package interaction readiness standpoint. The reduced cross-sectional area of fine Cu lines make it susceptible to cracking under temperature-cycling environments (Figure 2). The Cu lines (or traces) are typically sandwiched between a silica-filled polymeric Ajinomoto build-up film-based dielectric material, where differences in coefficient of thermal expansion drive stress at the internal FCBGA substrate layers in temperature-cycling environments. Performing both board and component temperature cycling enables full quantification of the reliability margins of fine Cu trace routing.

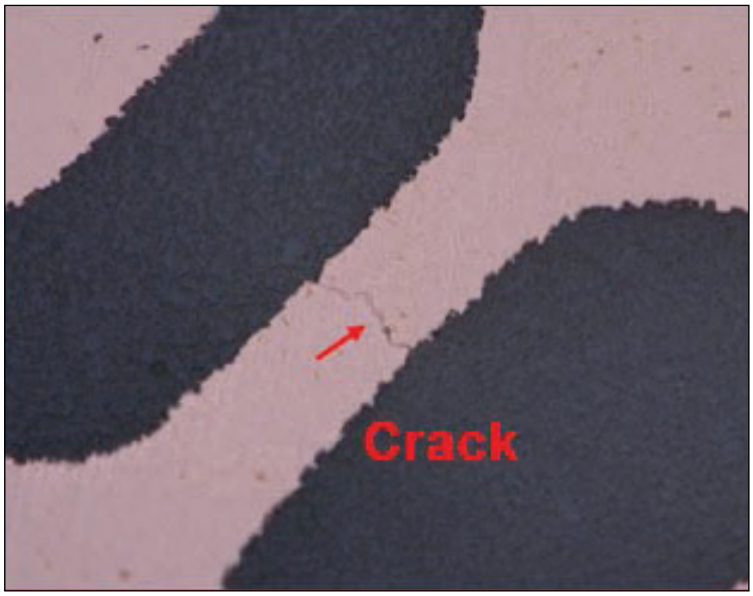

➤ Figure 2: Example of a copper trace crack during temperature cycling

Results and discussion

Performing both component- and board-level stress testing enabled a comprehensive perspective of the reliability of advanced design rules listed in Table 1. Employing an aggressive hammer test for CLR testing initiated an accelerated response based on an understanding of failure modes. This hammer test included a sequential series of aggressive moisture-soak-as-preconditioning substitute, multiple reflows at peak lead-free temperatures, and temperature cycling beyond condition B (–55°C to 125°C). BLR testing was based on Joint Electron Device Engineering Council (JEDEC) JESD22-A104 condition G, soak mode 4 conditions [3]. Recorded BLR failures occurred while in-situ monitoring after the first interruption for a period of a specified nanoseconds and an increase in daisy-chain resistance, followed by multiple additional interruptions within a percentage of the first cycles to failure.

Subjecting an equal sample size for legs 1 and 2 to the CLR hammer test conditions correlated to the failure mode produced from extended testing of typical AEC Q-100 temperature cycle condition B, and thus could accelerate the accumulation of test failures. After

Leg	Cu trace width	Cu trace space	Core thickness
1	8µm	8µm	<400µm
2	11µm	11µm	<400µm

➤ Table 1: Evaluation legs subjected to package reliability testing

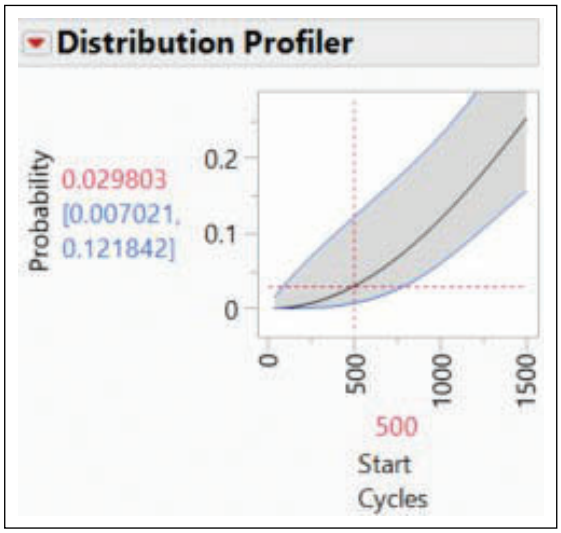

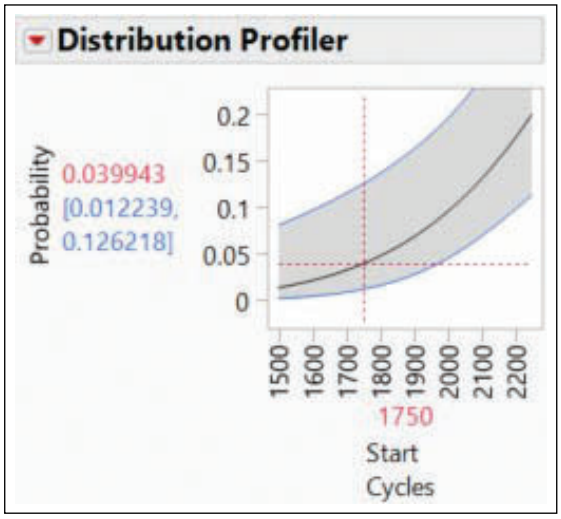

➤ Figure 3: Failure probability vs. first cycle to failure for legs 1 (above) and 2 (below)

test failures at three different temperature cycle read points, a Weibull analysis quantified the performance between legs 1 and 2. As indicated in Table 1, the only difference in the substrate design is the Cu trace width and space. With this being the case, the cross-sectional area of the Cu trace played a significant role in the CLR results.

For leg 1, the first cycle to failure was 3.5 times earlier than leg 2. Figure 3 shows the failure probability at a 95% confidence interval at first cycles to failure between legs 1 and 2. The figure also highlights the comparative results with similar failure probability plotted at a 95% confidence interval. The failure mode in both cases was Cu trace cracking at a similar inner layer of the substrate. Figure 4 shows the failure of the Cu trace crack.

With respect to characteristic life (alpha) and shape parameter (beta) values, the statistics of legs 1 and 2 provide a degree of clarity on the implications of the advanced substrate design rule of finer Cu traces as width reduces. Per Figure 5, leg 1 shows a slope (beta is greater than one, but less than four) that is not as steep as leg 2 (see Figure 6), indicating earlier wear out, higher variation and a lower degree of predictability. Leg 2 results represent a steeper slope (where beta is greater than four), indicating a more predictable and consistent life span for this failure mode given its wider Cu trace cross-sectional area. The beta value for leg 2 aligns with fatigue or an aging mechanism.

Figure 6 shows the characteristic life plot for leg 2. The tangible differences in slope between legs 1 and 2 are conspicuous, and therefore provide an explanation for the performance differences in characteristic life as observed in Figures 5 and 6.

Measuring warpage at both 30°C and 250°C revealed the contribution of package warpage and Cu density, if any, as a function of the different advanced substrate design rules between legs 1 and 2. These temperatures provide a contrast in warpage at room temperature and peak reflow temperature as associated with a lead-free surface-mount technology (SMT) process. Figure 7 shows that there was not any statistical difference between legs 1 and 2 based on overlap of circles from the JMP plot. When unmounted parts (parts not constrained to the printed circuit board) undergo temperature cycling, the package is free to conform to its natural state. As such, measuring warpage upon heating and cooling at 30°C captured any hysteresis effects and thermomechanical changes of the organic material sets as they surpassed their glass transition temperature. Similar to the results in Figure 7, the measurements in Figure 8 show no statistical difference between legs 1 and 2 upon heating and cooling.

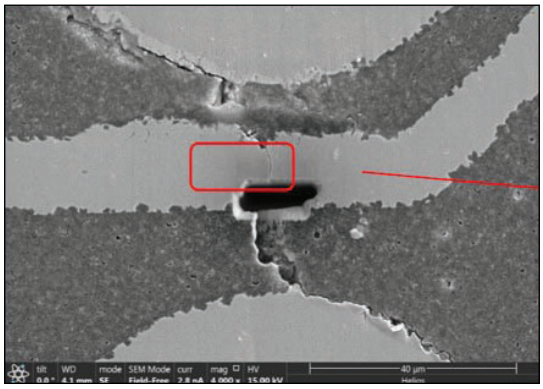

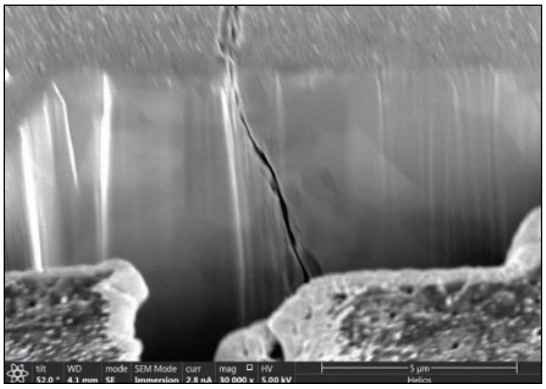

➤ Figure 4: Failure analysis of Cu trace cracking at the inner substrate layer

Evaluating legs 1 and 2 at JEDEC standard JESD22-A104 condition G, soak mode 4 conditions during BLR complements the CLR results previously reported. Similarly, Weibull analysis was the primary metric to quantify reliability performance between the two FCBGA advanced package rule designs as presented in Table 1. As expected, leg 1 failed first due to the reduced cross-sectional area of the Cu trace width. Failure analysis confirmed Cu trace cracking.

Although not reported in this study, multiple nets were monitored in-situ during BLR, with a separate net designed to isolate second-level interconnect reliability at BGA connections. The net reported in this study, as extracted from in-situ BLR monitoring, corresponds to a continuous loop path from the BGA pad through the package and first-level connection and routed back to the BGA pad. This net was also monitored for package integrity including Cu trace reliability.

Leg 1 failed approximately 1.2 times earlier than leg 2 at first cycles to failure. Figure 9 illustrates the characteristic life at 63.2% fails for legs 1 and 2 under BLR testing. Comparing Figure 9a and 9b, leg 2 has approximately 1.53 times higher cycles to failure than leg 1 at 63.2% fails.

Similar to the CLR results, the beta value (greater than four) post-BLR testing for leg 2 aligns with fatigue or an aging mechanism. Leg 1 has a beta value less than four, which is consistent with it having earlier first cycles to failure. Again, the reduced cross-sectional area of the finer Cu trace produced with leg 1 parts is likely the reason for the earlier first cycle to failure.

Substrate suppliers typically have a Cu trace width tolerance of approximately ±2µm, where inherent manufacturing variation plays a role in the characteristic life of the distribution of parts. Nonetheless, both BLR and CLR results show high margin in meeting AEC Q-100 conditions.

Conclusion

This study demonstrated the efficacy of advanced FCBGA design rules compatible with specific Cu trace width and space values and material sets through BLR and CLR conditions. Extracting package-level reliability margins befitting higher levels of driving automation was the focal point of this empirically based study.

The study subjected an equal distribution of FCBGA parts to accelerated CLR based on an understanding of common failure modes generated during standard AEC Q-100 conditions and BLR at JEDEC standard JESD22-A104 condition G, soak mode 4 conditions, tested to device failure. Weibull analysis provided suitable statistics to understand the wear-out mechanism of the Cu trace crack failure mode.

As anticipated, the finer Cu trace parts (leg 1) exhibited the first cycles to failure. The Weibull

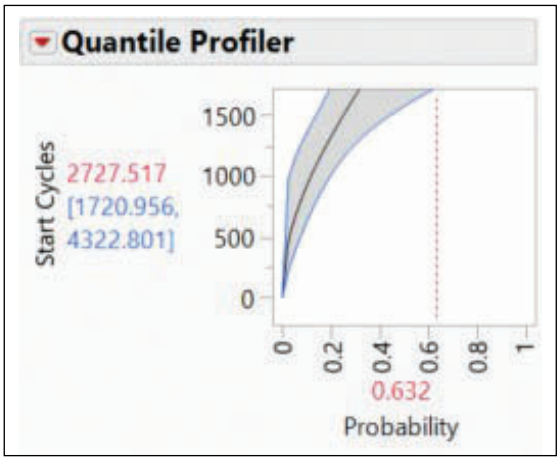

➤ Figure 5: Characteristic life at 63.2% fails for leg 1

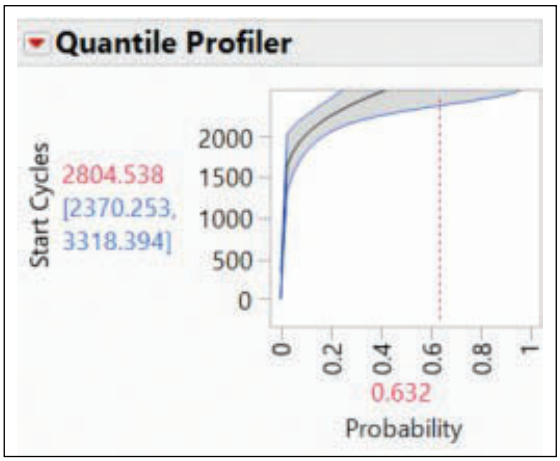

➤ Figure 6: Characteristic life at 63.2% fails for leg 2

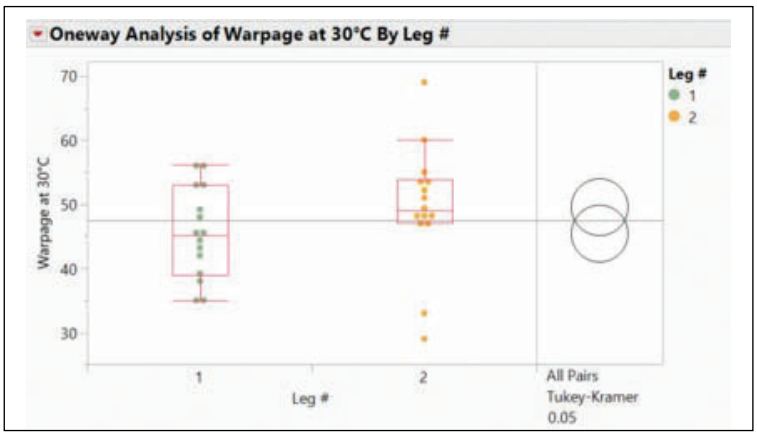

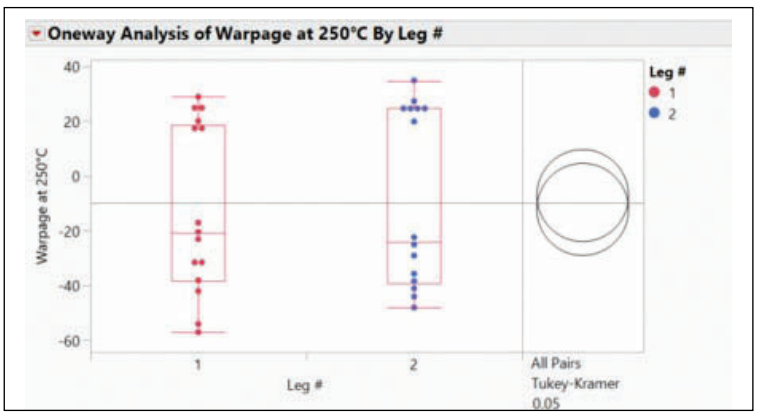

➤ Figure 7: Comparison of package warpage between legs 1 (above) and 2 (below) at 30°C and 250°C

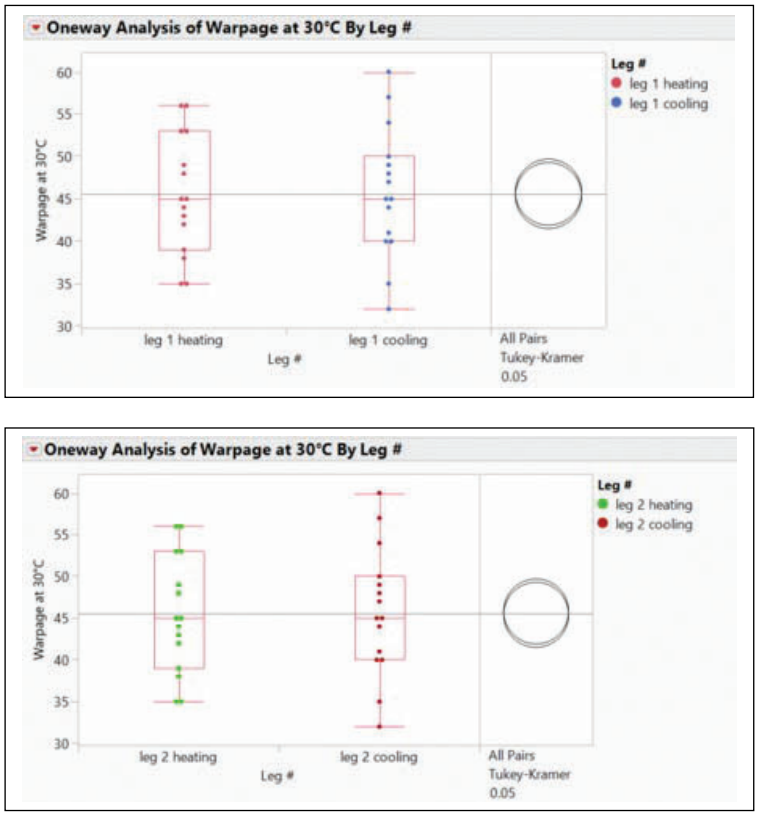

➤ Figure 8 (above): Comparison of package warpage of leg 1 (above) and 2 (below) heated and cooled to 30°C

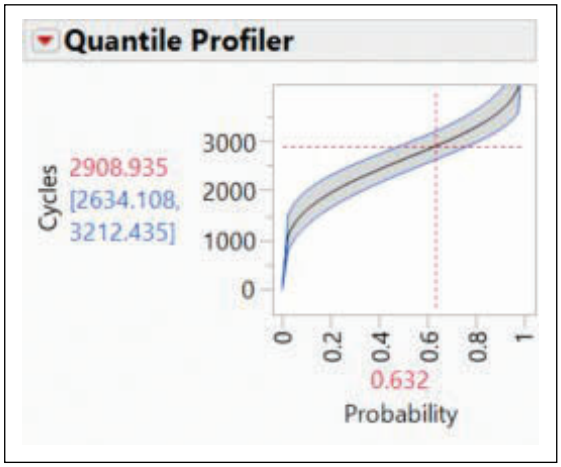

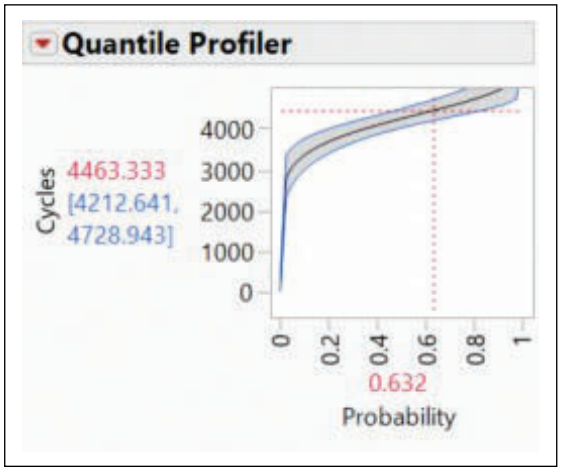

➤ Figure 9 (right): Characteristic life at 63.2% fails for legs 1 (above) and 2 (below) under BLR conditions

shape parameter (beta) value of less than four was consistent with the early cycles to failure that occurred during CLR and BLR testing, coupled with a wider variation of parts failing.

Given the reduced cross-sectional area of the Cu trace and its impact on reliability margin, packaging engineers must go the extra mile in understanding substrate supplier process control because of its contribution to chip-to-package (CPI) interaction. Studying the effects of CPI by measuring package warpage across a lead-free SMT profile found no statistical differences.

Developing an automotive mindset to achieve zero defects requires constructive interaction across substrate suppliers, assembly sites and material suppliers, coupled with a strong discipline canvassing literature. For the latter task, there is usually a precedent to set a path for exploration and deeper learning. The crux is empowering advanced substrate design rules as an inevitable feature of flip-chip packaging to facilitate higher levels of driving automation.

Acknowledgment

The author would like to acknowledge Lee McNally for his input on Weibull analysis.

Jaimal Williamson is a Packaging Engineer and Senior Member Technical Staff within Texas Instruments' Packaging Technology Solutions group, Dallas, TX, USA. His focal areas as a lead technologist include flip-chip platform development of advanced CMOS Si nodes, qualification, and productization to support multiple automotive, industrial, aerospace, and defense applications. He has authored and co-authored 35+ journal articles, conference papers, and technical magazine articles, as well as 35+ U.S. patents in the field of semiconductor packaging. He received a BS in Chemistry from Grambling State University and a MS in Polymers (via School of Materials Science and Engineering) from the Georgia Institute of Technology.

REFERENCES

- Yu, Wei, Faxing Che, Vance Liu, Raymond Chen, Sam Ireland, Yeow Chon Ong, Hong Wang Ng, and Gokul Kumar. "An Overview of Substrate Copper Trace Crack Through Experiments, Characterization, and Numerical Simulations." Micromachines 2025 16, no. 428.
- "Failure Mechanism Based Stress Test Qualification for Integrated Circuits in Automotive Applications." Rev J. Automotive Electronics Council: Aug. 11, 2023.
- "Temperature Cycling." JESD22-A104F.01. JEDEC: Arlington, Virginia, April 2023.

Advanced Wafer and Panel-Level Packaging Tools for Leading Edge Fabs

Panel-Level Packaging Solutions

- Metal Deposition
- Flux Cleaning
- Bevel Etching/Cleaning

Wafer-Level Packaging Solutions

- Metal Deposition
- Photoresist & Solvent Coating
- Photoresist Development
- Wet Chemical Etching
- Backside Cleaning/Etching
- Scrubber Cleaning
- Stress-Free Polishing
- Photoresist Stripping & Metal-Lift Off (MLO)

ACM RESEARCH

LEARN MORE

Atomic Layer Deposition:
Interfacial engineering to enable high performing package architectures

Atomic layer deposition quietly enabled the high-k era. Now AI-compute and extreme aspect-ratio interconnects demand atomic control again - at scale. Advances in deposition technology reveal that semiconductor winners will master interfacial control, not just dimensions.

TYLER J. MYERS, PHD, SENIOR TECHNICAL SALES ENGINEER, FORGE NANO, INC.

ATOMIC LAYER DEPOSITION (ALD) rarely gets top billing in semiconductor headlines. It is the kind of technology that does its most important work out of sight. Yet if you trace the industry's most pivotal inflection points, ALD keeps showing up as a critical enabler to realize next-generation devices.

Since Gordon Moore uttered his eponymous law, scaling was a story of shrinking – how can we fit more precious chips on a wafer? The obsession with smaller transistors, tighter pitches and higher density drove the semiconductor industry's technology development, including the first adoption of ALD for logic manufacturing. That era is not over, but it is no longer the full picture. The performance and energy demands of AI-era compute, coupled with heterogeneous integration carrying a growing share of system-level innovation, have pushed the industry toward a new reality: **The next winners will not be defined solely by who shrinks the fastest at a reasonable yield. They will be defined by who controls matter most precisely - at the atomic level - and who can do it at scale.**

ALD: the "quiet enabler" of miniaturization

ALD became essential long before it gained any notoriety. When Intel introduced high-k dielectrics to their manufacturing line in 2007 they demanded uniformity and defect control that conventional deposition struggled to deliver.

The sequential, self-limiting ALD method provided the process control needed to fabricate thin, conformal and pinhole-free films. ALD provided a route to much thinner gate oxides, and ultimately devices with lower power consumption, enabling Intel's transition from the 65 to 45nm node.

Today, a similar dynamic is repeating, but the challenge is many-fold. As device architectures are becoming increasingly complex, the traditional scaling conversation has moved to the package, and the performance demands are outpacing what systems can currently deliver.

This complexity requires the optimisation of more than size; it requires the strict control of every

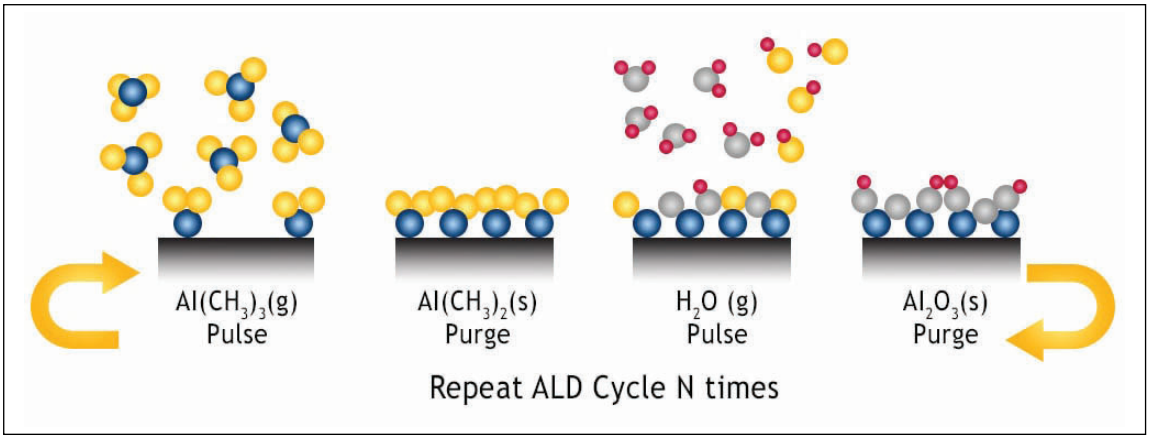

➤ Figure 1. Illustrative schematic of the ALD process. Material is deposited using sequential reactions that react only until all surfaces sites are saturated, resulting in precision thickness control, 100% step coverage, and high-quality materials.

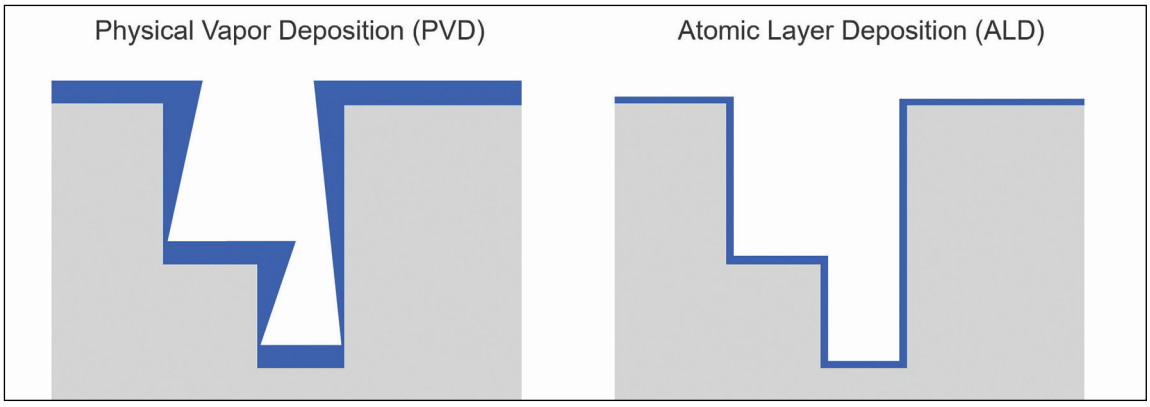

➤ Figure 2. Illustrative surface coverage examples of physical vapor deposition (PVD) and atomic layer deposition (ALD). PVD tends to have higher deposition rates with a trade-off in uniformity. ALD-grown materials are 100% conformal and uniform due to its self-limiting nature.

material, and in consequence, every interface, down to the nanometre. At those bounds, small deviations are more than noise. A few nanometres of non-conformal barrier at a critical junction, a thermal interface with too large non-uniformity or a void that becomes an electromigration hotspot, can spell disaster.

Packaging is becoming the architecture

Innovation in the advanced packaging space primarily involved ways to protect the die and route signals outward. Back-end-of-the-line (BEOL) work was generally seen as a separate set of fabrication steps. Now the package architecture increasingly defines performance. The systems driving AI acceleration and high-bandwidth memory no longer rely on a single monolithic chip. They rely on many chips with varying functions connected in close proximity, with massive interconnect density and tight power/thermal constraints.

As primary integration strategies like chip-on-wafer-on-substrate (CoWoS) and other 2.5D/3D approaches become mainstream, the interconnect stack begins to look more like front-end scaling than back-end assembly. Through-silicon vias (TSVs), through-glass vias (TGVs), and fine-pitch redistribution layers must support higher I/O counts, lower latency, and better signal integrity without sacrificing yield.

This is where deposition physics becomes strategic. To raise interconnect density, vias and trenches are pushed toward higher aspect ratios (AR), i.e. a feature's height divided by its width. But as AR increases, conventional, directional deposition methods struggle.

Line-of-sight processes, particularly physical vapour deposition (PVD), can lose effectiveness at surprisingly modest aspect ratios. In practice, many PVD approaches fail above ~10:1, and in some cases even lower.

The result can be a chain reaction of manufacturability issues:

- Non-conformal barriers inviting copper diffusion
- Poorly uniform seed layers preventing robust electroplating
- Voids and pinch-off driving yield loss and reliability risk

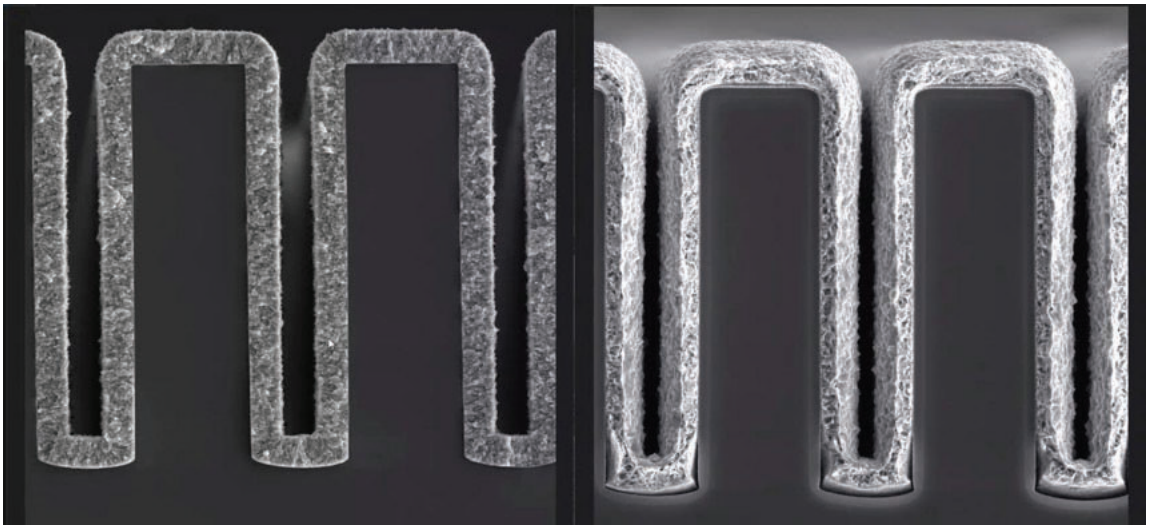

➤ Figure 3. Copper electrodeposition in 4:1 AR blind silicon vias utilizing (left) ALD-coated SiO_2/TiN/Ru barrier/seed and (right) PVD-coated Ti/W barrier/seed. Numerous voids, trench narrowing and bread-loafing is visible when plating on a PVD seed layer.

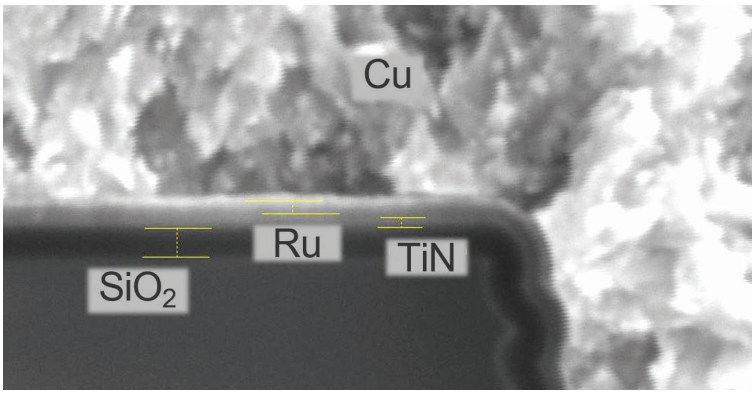

➤ Figure 4. ~100nm barrier/seed deposited by ALD on a blind Si via. The film is highly conformal with 100% step coverage and excellent uniformity at the trench corner and on the Bosch scallops. Cu adheres well and conformally on the Ru seed layer.

With packaging becoming synonymous with new device architectures, reliable interconnect fabrication is one of the new choke points.

The atomic-control requirement

It is tempting to treat advanced packaging challenges as a sole matter of making features narrower. But at high ARs and high-performance demands, it becomes harder to deliver the right material at the right thickness, throughout the entire structure.

The problems become more complicated than just, "Do I have a solution to enable via filling?" They now include:

- Can I deposit seeds conformally from top to bottom?
- Can I create the conditions for void-free electrodeposition in extreme structures?
- Can I keep interfaces chemically stable and electrically reliable?
- Can I do all of this without consuming too much via volume?

Those questions are fundamentally answered by atomic-level control. And ALD excels when it is engineered to meet manufacturability constraints.

The barrier/seed bottleneck: where packaging scaling gets stuck

In today's high-density interconnects, copper remains a workhorse of metal. But copper brings an old problem into a new geometry: it needs a diffusion barrier to prevent electromigration into underlying substrate, and a seed layer that supports smooth, continuous electrodeposition. In deep, narrow features, successful barrier/seed deposition is not guaranteed and decides whether the copper plates cleanly or fails.

Traditional approaches relied on PVD stacks. However, as aspect ratios rise, PVD struggles to maintain continuity, suffering from thickness non-uniformity. The failure modes are familiar to those in high-AR metallisation: poor adhesion, large void formation, and top-heavy deposition profiles.

Already, it is easy to see that ALD can solve the baseline functional problem with higher AR interconnects. Even at aspect ratios as low as 4:1, attempted copper electrodeposition on PVD barrier/seed layers in blind silicon trenches show familiar failure mechanisms. The copper grown on an ALD stack enables clean, uniform, and void-free electrodeposition due to ALD's ability to grow materials conformally in any geometry.

ALD, however, doesn't just enable functional electrodeposition. It can improve the entire stack through precision thickness control and careful material selection.

Let's take Forge Nano's all-thermal metal barrier/seed solution as an example of how ALD redefines what is manufacturable for interconnect fabrication.

What an ALD barrier/seed stack enables

Forge Nano has devised a metal barrier/seed stack comprised of all-thermal grown SiO_2, TiN and Ru for promotion of conformal, void-free copper electrodeposition in TSVs and TGVs (only TiN and Ru for TGVs).

A conformal ALD stack addresses the scaling and performance bottleneck in four practical ways:

- **All-thermal ALD supports higher AR**
  Thermally-grown ALD films mean the deposition conformality is not limited by geometry. This ensures the roadmap to high AR and high-density interconnect structures is not constrained by the barrier/seed layer.
- **Ultrathin barrier/seed films preserve via volume**
  In high-AR structures, every nanometre counts. A thick, non-uniform barrier/seed stack consumes valuable space, raising overall resistance. ALD's precision thickness control offers thinner films that free up volume for the conductive fill. The narrower the via, the more important fine

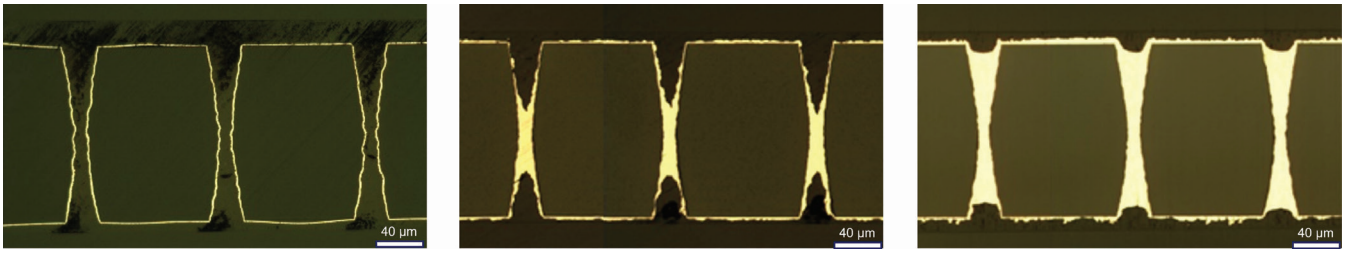

➤ Figure 5. Electrodeposition of Cu on Forge Nano's ALD barrier/seed in 10:1 AR TGVs. The copper starts plating conformally on all surfaces, then begins filling from the middle of the via and proceeds symmetrically to fill the entire via, free from voids, narrowing or bread-loafing.

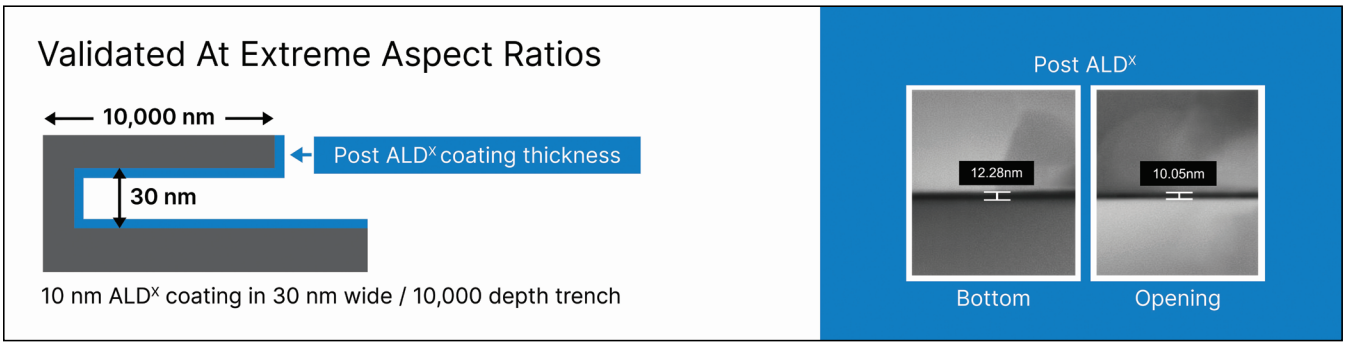

➤ Figure 6. Horizontal trenches with starting AR of 330:1 were coated conformally with ALD and production-competitive speeds. As the gap narrows, chemical was being transported to the bottom of the trench with effective aspect ratios of 1000:1. Images courtesy of C2Mi.

thickness control becomes, especially in high density vias where copper also helps with thermal management.

- **Low-resistivity metals improve overall electrical outcome**
  When the barrier and seed become a meaningful fraction of the cross-sectional area, their resistivity matters. Low resistivity materials grown by ALD, like TiN and Ru, can create an interconnect pathway with higher overall conductivity than stacks that rely heavily on higher-resistivity layers. These materials will become crucial for enabling the superconducting needs of interconnects in quantum computing applications.
- **100% step coverage enables, conformal and void-free copper electrodeposition**
  Copper electroplating reveals where the seed is continuous and where it is not. With conformal ALD barrier/seed layers, copper nucleates and grows uniformly, reducing void formation and pinch-off risk in vias and trenches.

This barrier/seed stack has been demonstrated on silicon and glass vias exceeding 35:1 aspect ratio, with the goal of enabling next-generation packaging interconnects where directional deposition fails.

For packaging audiences, the headline is straightforward:
If you want to move to the 30:1–35:1 AR regime and beyond, atomic-level control of your interfaces becomes the deciding factor for yield and reliability.

And that brings us to the most important part of the ALD conversation: can it scale?

Why ALD faced resistance

It is tempting to talk about ALD as though it is inevitable. What's not to love about atom-by-atom control? In practice, ALD spent much of its history on the margins because everything else was faster and cheaper.

For many layers, directional deposition and other high-throughput approaches delivered "good enough" films at a fraction of the cost and cycle time. ALD, by contrast, came with very real adoption barriers: slower deposition rates, purges and exposures that stretch dramatically in high aspect-ratio features, and process trade-offs between precursor utilisation and throughput.

The result was an industry mindset that treated ALD as a last resort. In other words, a tool you reached for only when the fast, economical options had unequivocally run out of road.

Advanced packaging in the AI-era is changing that mindset. As interconnect densities increase, "good enough" films are no longer viable. The problem remains, though, that the industry can absorb higher process cost when the alternative is yield, reliability and roadmap risk, but the throughput tax can be more difficult to accept.

Which raises the real question for the next era: if the winners are truly those who control matter the most precisely, how is it possible to stay ahead if the costs are ballooning?

Breaking the deposition bottleneck

This is where process innovation matters as much as deposition principle.

ALD became essential long before it gained any notoriety. When Intel introduced high-k dielectrics to their manufacturing line in 2007 they demanded uniformity and defect control that conventional deposition struggled to deliver. The sequential, self-limiting ALD method provided the process control needed to fabricate thin, conformal and pinhole-free films

In today's high-density interconnects, copper remains a workhorse of metal. But copper brings an old problem into a new geometry: it needs a diffusion barrier to prevent electromigration into underlying substrate, and a seed layer that supports smooth, continuous electrodeposition. In deep, narrow features, successful barrier/seed deposition is not guaranteed and decides whether the copper plates cleanly or fails

Forge Nano's patented ALDx approach targets the practical bottlenecks that historically limited ALD adoption in production contexts, particularly at high aspect ratio. Rather than treating ALD as a slow but perfect process, the goal is to deliver ALD-quality films with manufacturing-relevant efficiency.

In Forge Nano's metal barrier/seed work, the toolset focuses on fast dosing and rapid cycling, using clever valve and chamber design optimised to deliver throughputs rivalling traditional processes with orders of magnitude less precursor consumption.

With single-wafer deposition rates as high as 12 nm/min and precursor usages 100x lower, it becomes realistic to think of ALD as an architecture-enabling platform technology for advanced packaging, rather than a niche process reserved for only the most unforgiving geometries.

That shift matters because it separates two futures:

- One where packaging scaling stalls because deposition cannot keep up with manufacturing needs
- Another where packaging becomes a design space where engineers can push architectures without worry that execution is impossible due to OPEX constraints

A wider lens: from 35:1 interconnects to 1000:1 structures

Thus far, we have anchored the packaging narrative at a bottleneck being addressed today: barrier/seed metallisation and the 35:1 interconnect regime that challenges directional deposition.

But advanced packaging is not limited to TSVs and TGVs. The same interplay can be extended to other packaging aspects that will show up across the package in new ways.

One compelling example is deep trench capacitors, where surface area translates to charge storage.

The ability to conformally coat deeper trenches without defects translates to higher capacitance in the same footprint. Forge Nano has shown turbulent-flow capability that can coat aspect ratios up to 1000:1, with production-level robustness, and that result has been independently validated by Canadian research institute, C2Mi.

This is not a claim that packaging will become 1000:1 overnight. It is, however, a signal that the industry's assumed limits are moving *quickly*, and that deposition constraints once treated as immovable can become engineering problems with practical solutions.

The AI race is becoming an atomic engineering race

So, the truth is laid bare: the most meaningful performance gains now do not come from another geometry shrink. They come from system architecture and the reliability of the materials that make that architecture manufacturable. That means the competitive advantage shifts.

In the next generation:

- Some teams will optimise routing and assembly around deposition limits.
- Others will *change the limits* by innovating the materials and interfaces that define those processes.

The second group wins more than yield; they win freedom. They can:

- Push interconnect density higher
- Reduce resistance and improve signal/power integrity
- Enable architectures competitors cannot easily copy

ALD will not be the only tool in that toolbox, but it will be what you reach for when you need atomic control, conformality, and repeatability in critical structures.

ALD's second act will be louder than the first

ALD helped enable the high-k era without demanding the spotlight. Advanced packaging in the AI-era may change that.
As interconnect aspect ratios climb and heterogeneous integration becomes a core lever for system performance, the industry needs deposition approaches that do not compromise film quality or interface control. At the same time, those approaches must be compatible with scale because "it works in a lab" does not win markets.

The next generation of semiconductors will reward those who can innovate at the atomic level and deliver that innovation through manufacturable processes. That is where ALD moves from a quiet enabler to a competitive separator.

Atomic precision is not a slogan. It is becoming the price of admission.

FORGE NANO

Single Wafer ALD at Batch Speeds

Production-Speed ALD at 1000:1 Aspect Ratio - A Semiconductor First

12 nm/min Deposition Rate	>25% Device Enhancement
1000:1 Aspect ratio | Defect-free	90% Chemical Efficiency

Validated At Extreme Aspect Ratios

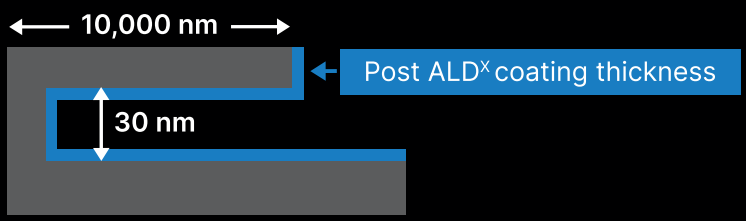

10 nm ALD^X coating in 30 nm wide / 10,000 depth trench

Post ALD^X

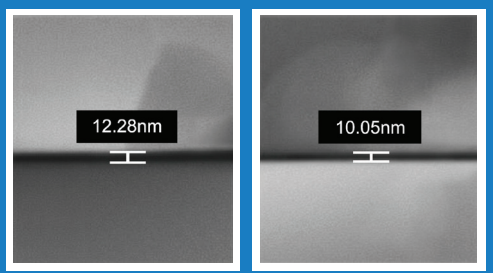

Bottom Opening

Independently validated on production representative wafers by C2MI

Discover Forge Nano's Semiconductor ALD tools and applications

forgenano.com

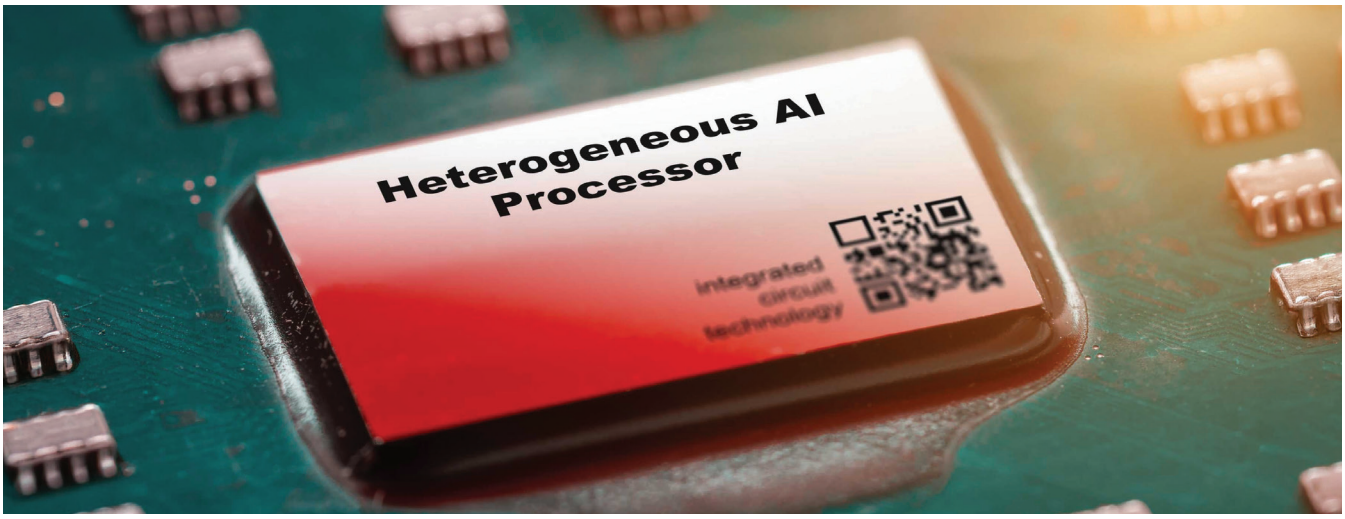

How PI Enables Advanced Packaging

How precision motion, alignment intelligence, and integrated subsystems turn heterogeneous integration into manufacturable reality.

DAVID FORER – DIRECTOR MARKET STRATEGIES – SEMICONDUCTOR, PHYSIK INSTRUMENTE

Introduction: packaging is now the performance lever

ADVANCED PACKAGING has moved from a back-end afterthought to a front-line driver of system performance. As computing becomes increasingly modular chiplets, stacked memory, co-packaged optics, and heterogeneous integration, the package is no longer just protection; it is the interconnect fabric and, in many designs, part of the signal path. Industry investment reflects this shift. One recent example is the creation of new advanced packaging and photonics centres aimed at scaling 3D integration and silicon photonics for high bandwidth, energy-efficient systems.[1] That scaling effort exposes a hard truth: the best architectures are only as good as the assembly technology that can build them at yield.

At leading-edge nodes, assembly tolerance budgets collapse. Optical waveguide cores in silicon photonics can be on the order of 0.2 micrometres, so coupling optics to photonic integrated circuits quickly becomes a nanometer class alignment problem rather than a conventional pick-and-place task.[2] Similar precision pressures are showing up across advanced packaging steps: aligning die to interposer interfaces, placing micro bump arrays, controlling coplanarity for bonding, and holding position through adhesive cure or thermal excursions. These demands are fundamentally motion control problems multi-axis, high resolution, repeatable, and fast.

PI (Physik Instrumente) enables this transition by supplying the motion, sensing, and control subsystems that allow equipment makers and device manufacturers to execute advanced packaging workflows with the necessary combination of accuracy and throughput. PI's contribution is not limited to a catalog of components; it includes integrated multi axis platforms and firmware level alignment intelligence that can run optimization loops directly on the controller. The result is a practical path from 'it aligns on the bench' to 'it aligns in high volume manufacturing'.

Where advanced packaging stresses the factory

Advanced packaging processes tend to concentrate risk into a small set of operations: alignment, bonding, and verification. Each is a high-sensitivity step where misalignment can permanently lock in performance loss or cause outright failure. In photonics packaging, for example, active optical alignment is frequently the gating operation. Because many assemblies include multiple optical channels and multiple optical elements, manufacturers may need to find and then re- optimize alignment several times along the build flow.[4] That repetition drives both cycle time and cost, and it magnifies any shortfall in automation robustness.

From a motion standpoint, the manufacturing challenge has three defining characteristics. First, alignment is inherently multi degree of freedom. Translation and rotation are coupled; optimising X Y often changes the best Z and tilt angles. Many workflows, therefore require sub-micron, and in some cases nanometer, control across as many as six degrees of freedom. [3] Second, the system must be stable while the process acts on the assembly. Adhesive cure, thermal cycling, or mechanical clamping can introduce drift, so the motion system must either hold position with high stiffness or actively compensate. Third, the factory cares about seconds, not minutes. Serial, axis-by-axis tuning scales poorly as channel count increases; what is tolerable for a single fibre becomes impossible for an array.

Packaging equipment also faces tight envelope constraints. Tools must fit inside compact production stations, and in some applications, they must operate in controlled environments (clean, vacuum, or temperature-managed).

Assembly systems must maintain alignment under vibration and shock while minimising footprint.[5] Meeting all these requirements simultaneously is difficult with conventional stacked stages and PC driven control loops.

PI's motion hardware foundation: stiffness, resolution, and dynamics

PI's portfolio for semiconductor and photonics manufacturing covers precision motion and positioning technologies from nanopositioners to multi-axis robots.[6] For advanced packaging, three hardware themes are particularly relevant: parallel kinematic multi-axis platforms, hybrid coarse/fine positioning, and ultra-smooth long travel stages for wafer and panel level operations.

Parallel kinematics is exemplified by PI's hexapods: six struts arranged as a Stewart platform provide motion in all six degrees of freedom. Compared with stacking individual linear and rotary stages, a hexapod can deliver a smaller footprint, higher stiffness, and better dynamic behavior.[7][8] High stiffness matters directly in packaging, because the system must hold an aligned position while bonding forces, tooling contact, or curing processes act on the assembly. PI provides hexapods across a range of sizes, including compact and vacuum-compatible variants for controlled environments.[9] Many platforms use direct drive actuation and high-resolution encoders, supporting smooth motion and nanometer-class step size.[10]

At the subsystem level, PI has recently highlighted compact photonics alignment platforms that combine multiple degrees of freedom with substantial travel in a small package. For example, an alignment system with 40 mm travel in a roughly 5 x 7 x 4 inch envelope, using crossed roller bearings and direct drive linear motors for fast, backlash-free motion.[11][10] These design choices map well to advanced packaging realities: a tool must move far enough to acquire alignment, then converge quickly without overshoot.

Hybrid motion architectures bridge the gap between coarse travel and ultra-fine control. PI's piezo-driven flexure nanopositioners provide sub-nanometer resolution and rapid response and are commonly paired with a longer range motorised platform.[12][8] In practice, the motor stage performs acquisition and gross alignment, while the piezo stage executes high-speed scans, dithers, and drift compensation near the optimum. This approach is especially valuable for active optical alignment, where the peak coupling region can be extremely narrow.[2]

Wafer and panel scale workflows often require long travel motion with exceptional straightness and repeatability. PI's semiconductor motion solutions include air bearing stages that provide frictionless motion and nanometer straightness for inspection or lithography like scanning operations.[17] Such platforms can move a 300 mm wafer rapidly and then stop with sub-micron precision at points of interest.[18] Where bonding demands planarity or levelling, tip/tilt and levelling stages can be integrated to maintain the required surface orientation.[9]

Alignment intelligence: moving faster by moving smarter

Mechanical precision alone does not deliver production throughput. PI explicitly treats software and control algorithms as core enablers of photonics and semiconductor assembly automation.[19] Accordingly, PI implements alignment routines and optimisation logic directly in high-performance motion controllers, so the system can search and converge without relying on slow, external loops.

A centrepiece is PI's Fast Multichannel Photonics Alignment (FMPA) capability, implemented as firmware-level commands that execute complex alignment sequences across multiple axes and channels.[20] In traditional serial alignment, a tool optimises one channel or one degree of freedom at a time, then repeats because axis interactions shift the optimum.[21] PI's approach is to treat alignment as a coupled, multi-variable problem and to optimise multiple channels simultaneously to find a global maximum.[22] In PI's reporting, multi-channel alignment systems can reduce alignment time by about 99% compared to legacy approaches often described as a 100x speed up that turns minutes into seconds.[20][23] For factories, this is the difference between a viable per-unit cost and an expensive bottleneck.[24][25]

➤ H-811.F2 Miniature Hexapod | Ideal for Fiber Alignment

Controller resident scan execution is another practical accelerator. To acquire initial coupling, often called 'find first light' the controller can run continuous spiral or sinusoidal area scans without stop start overhead, avoiding delays caused by host communication and mechanical settling.[26][27] Once a signal is detected, the controller transitions into gradient-based searches that dither position and use signal change to drive convergence to the optimum.[28][29] Because these routines can operate concurrently on different degrees of freedom, the system can converge faster than sequential tuning and can better handle coupled sensitivities.[30] Importantly for packaging, such closed-loop strategies can continue during adhesive cure or thermal effects, maintaining alignment as the assembly evolves.[31]

PI further shortens the loop by integrating metrology inputs directly into the motion platform. Many alignment systems accept high resolution analog feedback from photodetectors or power meters, enabling the controller to treat optical signal as the figure of merit.[32] In PI's F 712 alignment system, for example, the optical power can be read directly and used in the optimization routine, removing the need for slower external measurement workflows.[33] Firmware level functions can map coupling

F-713.MAx Compact, High-Speed XYZ Photonics Alignment System

landscapes and compute the optimum using embedded evaluation logic.[34] PI has also described EtherCAT based controller architectures that support rapid signal analysis and even onboard machine learning to accelerate detection of the best alignment point.[35]

Taken together, these capabilities turn alignment from a manual craft into a deterministic, automated process: the system searches, optimises, locks, and if necessary, continuously corrects. This is why PI describes its approach as combining speed, nanoscale performance, and industrial robustness to improve manufacturing economics.[36]

From components to production subsystems: what 'enablement' looks like

PI's value in advanced packaging is most visible when its motion hardware and alignment intelligence are delivered as integrated subsystems. Equipment builders often prefer a validated motion and control module they can integrate, rather than assembling a solution from individual parts. PI addresses this through modular photonics alignment platforms, controllers, and application-level software that together constitute a turnkey motion subsystem.[6][8]

A widely cited example is the F 712 and F 713 family of photonic alignment engines. These systems combine a 6 axis motorised platform with nested piezo nano scanners and ship with the necessary control electronics and alignment firmware.[8][37] In the dual-sided configuration (F 712.HA2), two opposing alignment units can align input and output optical interfaces simultaneously, effectively handling multi-channel coupling in six degrees of freedom.[37] PI positions this approach as the benchmark for high-throughput silicon photonics packaging, enabling device-level alignment in about one second by leveraging embedded scan and parallel optimisation routines.[23][37] PI has also communicated manufacturing readiness by noting increased production capacity and reduced lead times for these systems, reflecting deployment scale and maturity.[41][6]

Not every packaging task needs a full 6 DOF capability. For applications where angular alignment is less critical, PI offers compact high-speed XYZ alignment systems such as the F 713.MA series.[40] These modules target use cases like fibre to laser or fibre to photodiode alignment, where rapid translational optimisation dominates the cycle time. The key is that the same controller-level alignment primitives can be reused across different mechanical platforms, simplifying tool development and reducing validation effort.

PI's motion control ecosystem also benefits from its multi-axis control heritage and the addition of ACS Motion Control, which PI acquired to strengthen controller capabilities for demanding automation tasks.[41] For tool makers, this matters because advanced packaging equipment is increasingly mechatronic: multi-axis motion, synchronisation, metrology integration, and factory communications must work together as a system. Subsystem solutions reduce integration risk and shorten time to market.

Application impact across advanced packaging and photonics

PI technologies appear across advanced packaging workflows wherever precision placement and alignment determine yield. In silicon photonics, high-throughput wafer testing and die-level assembly depend on the ability to align multiple optical channels quickly and repeatably.[43][44] PI's automation content emphasises that active optics alignment can compress test and packaging cycles substantially, enabling practical scaling of silicon photonics manufacturing.[2][3] The same toolchain supports assembly steps like coupling fibre arrays to PICs, aligning lenses or micro optics, and maintaining alignment through fixation.[3]

In semiconductor packaging, PI's motion platforms contribute to a broad set of equipment classes, including mask aligners, wafer dicing systems, lithography and inspection stages, and bonding-related tools.[47] During die attach or flip chip operations, for example, the requirement is not only placement accuracy but also repeatability and stable holding during bonding. Where needed, a coarse platform can bring parts into proximity while a fine nanopositioner performs final alignment and correction. At the wafer scale, air bearing stages support scanning and metrology steps that feed forward into bonding and assembly decisions.[17][18]

Beyond semiconductors, PI's active alignment expertise is used in packaging high-density optical and sensor modules. PI notes that active alignment methods are used in applications such as camera module assembly and LiDAR related optics, where quality metrics can be optimised in real time during positioning.[4] These adjacent markets reinforce a broader point: when performance depends on precise alignment of multiple elements,

Not every packaging task needs a full 6 DOF capability. For applications where angular alignment is less critical, PI offers compact high-speed XYZ alignment systems such as the F 713.MA series

the manufacturing solution tends to converge on multi-axis motion plus a robust optimisation loop.

Conclusion: enabling manufacturability at the nanometer scale

Advanced packaging is, at its core, a manufacturability problem under extreme tolerance budgets. Chiplet architectures, co-packaged optics, and dense interconnect schemes create value only when assembly can be done quickly, repeatably, and at yield. PI enables this by providing a toolkit that spans high stiffness multi-axis mechanics, nanopositioning for last micron adjustments, long travel scanning stages for wafer level operations, and controller resident alignment intelligence that converts complex multi-variable alignment into an automated process.

The practical outcome is measurable in cycle time and scalability. PI reports alignment time reductions on the order of 99% for multi-channel photonics alignment compared to traditional methods, bringing per device alignment into a production friendly time budget. [20][23] Equally important, integrated subsystems such as the F 712 and F 713 families help equipment builders reduce integration risk by delivering validated hardware and firmware as a unit.[8][37] As advanced packaging continues to evolve, motion and alignment subsystems will remain essential infrastructure, and PI's integrated approach positions it as a key enabler for the next generation of heterogeneous systems.

REFERENCES

- [1] GlobalFoundries to open Advanced Packaging and Photonics Center in New York - PIC Magazine News https://picmagazine.net/article/120944/GlobalFoundries_to_open_Advanced_Packaging_and_Photonics_Center_in_New_York
- [2] [19] [36] [43] [44] Photonics Automation for Packaging, Active Optics Alignment https://www.pi-usa.us/en/expertise/photonics-packaging-automation-active-optics-alignment
- [3] [5] [45] [46] SiPh Testing, Assembly, and Packaging https://www.pi-usa.us/en/expertise/markets/photonics/silicon-photonics/photonic-packaging
- [4] [20] [21] [22] [24] [25] [26] [27] [28] [29] [30] [31] [42] Active Alignment https://www.physikinstrumente.com/en/expertise/technology/controllers-software/active-alignment
- [6] [8] [13] [23] [37] [38] [41] Higher Throughput in Photonic Chip Manufacturing: PI Halves Lead Time for F-712.HAx Fiber Alignment Systems https://www.physikinstrumente.com/en/about/news-press/news-detailpage/lead-time-f-712
- [7] [12] A Wide Range of Precision Motion and Positioning Products from PI https://www.azonano.com/nanotechnology-video-details.aspx?VidID=1190
- [9] [17] [18] [47] Motion Control Nanopositioning | Semiconductor Applications | PI https://www.pi-usa.us/en/products/motion-control-solutions-for-the-semiconductor-industry
- [10] [11] [32] [35] PI to demonstrate new PIC alignment system at Photonics West - PIC Magazine News https://picmagazine.net/article/120934/PI_to_demonstrate_new_PIC_alignment_system_at_Photonics_West
- [14] [15] [16] [33] [34] [39] Item # F-712.HA2, F-712.HA2 High-Precision 12-Axis Fiber Alignment Stages On PI (Physik Instrumente) L.P. https://motionsystems.pi-usa.us/item/fiber-alignment-stages---photonics-alignment/f-712-ha2-high-precision-alignment-systems/f-712-ha2
- [40] F-713.MAx Compact, High-Speed XYZ Photonics Alignment Systems https://www.pi-usa.us/en/products/photonics-alignment-solutions/f-713max-compact-high-speed-xyz-photonics-alignment-systems

Hybrid bonding for chiplets: Unlocking BEOL-Level system integration

As AI pushes power and bandwidth limits, chiplets promise a new scaling path. Hybrid bonding goes beyond fine-pitch I/O, enabling BEOL-direct integration that reshapes system design across wafers and dies.

FUMIHIRO INOUE, PROFESSOR, YOKOHAMA NATIONAL UNIVERSITY

Why packaging now drives system-level scaling

AI DEMAND accelerated by large language models is turning data-centre electricity consumption into a societal constraint. For decades, the semiconductor industry improved performance and energy efficiency mainly through scaling—shrinking devices and interconnects. Today, the rapid rise in AI compute demand is colliding with the practical limits of traditional scaling, forcing the industry to find new levers. Advanced packaging has therefore moved from “supporting technology” to a central driver of system progress.

Front-end-of-line innovation is not over. Nanosheet (gate-all-around, GAA) architectures, and longer-term concepts such as CFET that stack devices vertically, suggest that integration density per unit area can still increase. Yet the capital required is immense, and only a few companies can execute at scale. This reinforces industry consolidation and makes supply capacity and lead time more dependent on a small number of players. At the same time, higher process complexity translates into yield pressure. As scaling slows, integration often pushes die sizes larger, reducing dies per wafer and amplifying the yield penalty. Moreover, next-generation integration approaches—such as CFET and backside power delivery network (BSPDN)—carry strong design and thermal constraints, making it harder to assume that energy efficiency will improve “automatically” with each node as it often did in the past.

In this environment, chiplet architectures enabled by advanced packaging are gaining momentum. The core idea is to partition functions once integrated into a monolithic SoC and then re-integrate them

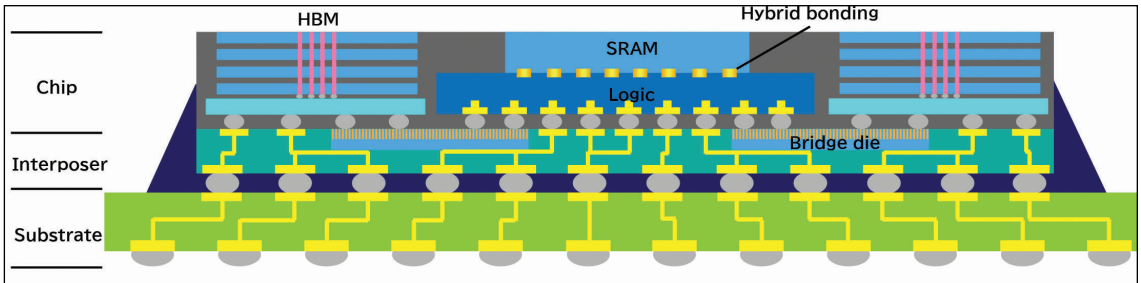

Hierarchy		Horizontal routing	Vertical interconnects	
			Through-vias	Bond joints
Chip (die level)		Top BEOL global routing	TSV (Through Si Via)	Microbumps / hybrid bonding
Interposer	Si interposer (wafer-level)	Top BEOL global routing	TSV	Microbumps with Cu pillars
	Organic RDL (on glass carrier)	Redistribution Layer, RDL (semi-additive process, formed on carrier)	TMV (Through-mold via) =Mega-pillar	
	Bridge (on glass carrier)			
Package substrate		Build-up layers (laser direct imaging) or RDL (semi-additive)	Micro via(Through Substrate via) or TGV (Through-Glass via)	C4 bumps

➤ Figure.1 Chiplet integration is inherently layered across the stack.

Comparison axis	Microbumps (solder / Cu-pillar)	Hybrid bonding (dielectric + Cu–Cu)
	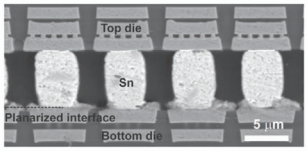	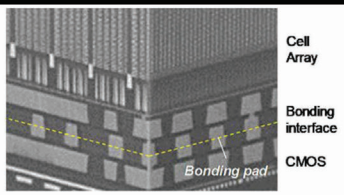
Primary role	I/O connection between dies	**BEOL-direct connectivity** across die boundary (BEOL-level integration)
Typical pitch regime	**Tens of µm** class (practical)	**Sub-µm to few µm** class (technology-dependent)
Interconnect resistance / latency	Higher R and parasitics; longer vertical path	Lower parasitism; shorter path → better latency/energy potential
Thermal path	Underfill + bumps introduce thermal impedance	Thinner interface; potential for improved heat spreading (depends on stack & TIM)
Interface thickness	Thicker joint (bump + underfill)	**Ultrathin** bonded interface (no solder joint)
I/O density scalability	Limited by bump collapse/short risk and assembly constraints	Scales with BEOL-like patterning + planar bonding (limited by overlay/defects)
Key process bottlenecks	Assembly yield, underfill voids, electromigration / fatigue	CMP planarity, particles, surface chemistry, bond-front control
Design freedom	Mostly connect I/Os → architecture largely fixed	Enables **exporting wiring** across chips; co-optimize architecture + interconnect

➤ Figure.2 Why hybrid bonding goes beyond microbumps. Compared with solder microbumps, hybrid bonding enables finer pitch and BEOL-like connectivity across die boundaries, turning chiplet integration into an architectural design tool.

in an optimized way. By manufacturing functional blocks—xPU, SRAM, SerDes, modems—separately, each die can be smaller, and yield can improve. Different nodes can be selected per function, enabling IP reuse and cost optimization.

For AI systems, the most powerful benefit is often interconnect redesign. In GPU–HBM configurations, latency, bandwidth, and interconnect power strongly shape system performance. Chiplets allow the interconnect structure itself to be rebuilt—for example, by embedding fine-pitch bridge dies only where ultra-dense interconnect is required. This is why chiplets have become a leading path for “scaling beyond lithography” in the AI era.

Chiplets are a stack of layers, not a single technology

Chiplets are often introduced as a concept and a design philosophy. In practice, the technical landscape is broader and more layered than it appears. The term “chiplet” now spans device vendors, foundries, OSATs, and substrate manufacturers. Many groups develop under the same banner, even though they operate at different layers and historically produced different product categories. As a result, technologies that used to be distributed across distinct layers can appear to blur under a single label, which can hinder clear discussion.

Interposers and substrates, for example, are frequently treated as separate categories, yet they share many development items: RDL formation, microbump fabrication, warpage reduction, and more. When discussing chiplets, it is therefore essential to clarify which layer is being addressed and how a given development activity fits into the overall integration stack.

Chiplet integration also faces several manufacturing bottlenecks. A prominent one is chip-to-chip I/O density. Solder-based microbumps remain widely used, but their practical pitch often sits around several tens of micrometres. Thermo-compression bonding relies on simultaneous heating and pressing, which increases short-circuit risk at finer pitches and raises concerns around resistance and thermal performance. Economics add another layer of difficulty: chiplets are often motivated by cost and yield optimisation, yet advanced packaging can become expensive, demanding careful optimisation across design, process, equipment, and yield. Thermal management and test—especially in relation to Known Good Die (KGD) and system-level inspection—remain tightly coupled to architecture choices.

Among these many challenges, this feature focuses on hybrid bonding as a key technology to move beyond the limitations of microbumps in chip-to-chip interconnect.

What hybrid bonding really changes

Hybrid bonding is frequently described as the next step after microbumps. In reality, its value extends beyond fine-pitch I/O replacement. Its deeper potential lies in design freedom: hybrid bonding can make chiplet partitioning and re-integration practical at the BEOL level.

Technically, hybrid bonding is rooted in the same foundation as damascene copper interconnect formation. It bonds wafers (or dies) without an adhesive layer and without solder bumps. Planarized dielectric surfaces come into contact and mechanically bond, while facing copper pads are simultaneously aligned. Subsequent thermal treatment promotes copper self-diffusion and forms

robust Cu–Cu bonds. In short, dielectric bonding and metal interconnect bonding are achieved in a single bonding operation—hence “hybrid”.

Architecturally, hybrid bonding is compelling because it can be approached with a BEOL-like mindset. If the interface were simply an I/O connection, it might be treated as a separate design domain, as microbumps often are. Hybrid bonding, however, can span BEOL hierarchies across the die boundary. It effectively allows interconnect optimization and architectural thinking once confined within a monolithic SoC to be “exported” beyond the die boundary. Signals can be moved vertically across chips over the shortest physical distance, turning chiplet integration into an architectural tool. This aligns with the “CMOS 2.0” framing often used to describe a new era of system scaling.

Wafer-level hybrid bonding has already progressed from development into high-volume applications. CMOS image sensors [1] were an early driver, and adoption has expanded beyond that domain. In NAND, for example, the CMOS circuit wafer and the memory array wafer can be manufactured separately and then integrated via hybrid bonding in a CMOS Bonded to Array (CBA) scheme [2]. In such applications, bonding has reached regimes below 1 µm pitch, well beyond what is practical with microbumps.

Scaling barriers: CMP, metrology, overlay, reliability and test

Scaling hybrid bonding into robust and economical manufacturing is a “total engineering” challenge spanning tightly coupled layers. Key issues can be grouped into:
(1) surface preparation, (2) surface evaluation and characterization, (3) physical/chemical mechanism control, (4) overlay (alignment) error and root-cause analysis, (5) interconnect reliability, and (6) testability and inspection flow.

The most critical step is often CMP and associated upstream/downstream processes that bring copper and dielectric into a highly planar state. In hybrid bonding, subtle topography variation or copper recess non-uniformity can drive voids, non-bonded regions, and electrical failure. Control must extend across wafer- and die-scale flatness, edge roll-off, dielectric roughness, and copper pad erosion and recess—far tighter than typical BEOL CMP requirements.

Metrology can become a bottleneck. AFM provides high-precision local information but is limited in the throughput and sampling area. This motivates broader-area, higher-throughput optical metrology and inline inspection, as well as process schemes that manage interface state from cleaning/activation to bonding, including sensitivity to queue time.

Mechanism understanding is also evolving. Increasing attention is being paid to the coupling between bonding mechanics (bond-front or bond-wave propagation) and surface chemistry. Bond-front stability can influence defect formation, including particle-driven voids and local non-bonded regions. As a result, “how bonding progresses” is increasingly treated as an engineering variable rather than an uncontrollable outcome [3].

Overlay requirements tighten sharply as pitch moves into the sub-micron range. Following the initial contact (often triggered near the wafer centre), a bond wave propagates rapidly—on the order of 10 seconds across the wafer—typically as concentric fronts. Meanwhile, maintaining <25 nm misalignment across the full wafer becomes necessary, a far more stringent requirement than conventional scanner lithography [4]. Because the interface effectively “locks” upon contact, the overall placement–metrology–correction loop ultimately determines success.

Reliability evaluation is advancing both mechanical integrity (interfacial strength and fracture modes) and electrical stability (contact resistance variation and degradation under thermal/stress conditions). Meanwhile, hybrid bonding introduces test constraints: because planarity is critical, direct probing copper pads using conventional probe cards can be difficult, and pad choices may be restricted. Test strategy, therefore, needs to be integrated into development rather than treated as an afterthought.

Reliability evaluation is advancing both mechanical integrity (interfacial strength and fracture modes) and electrical stability (contact resistance variation and degradation under thermal/stress conditions).

Die-Level hybrid bonding: The hard part of chiplets

Wafer-level hybrid bonding will likely expand and continue scaling in pitch and yield. In parallel, chiplet growth depends on extending hybrid bonding to die-level integration. Many products require selective assembly and KGD usage, which pushes integration towards die-to-wafer (D2W) hybrid bonding.

A widely recognized early example is AMD’s 3D V-Cache, where a cache die is hybrid-bonded onto a logic die using foundry-level 3D integration capability. Yet adoption remains limited, and broader application still faces significant hurdles.

	Wafer-to-wafer (W2W)	Die-to-wafer (D2W) / die-level
	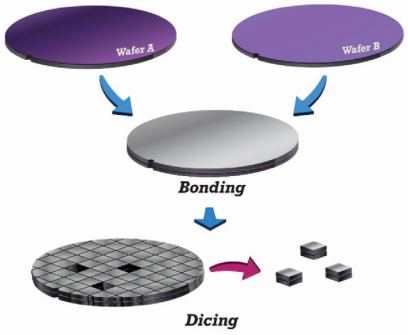	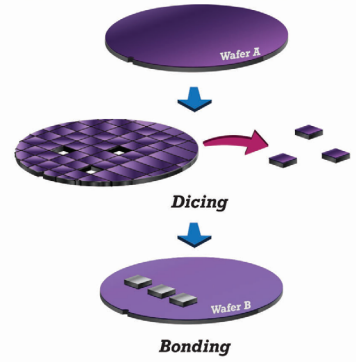
Best suited for	High-volume, matched die formats	Heterogeneous integration with KGD flexibility
Process flow	Bond first, then thin/dice	Thin and singulate before bonding
Residue / particle sensitivity	Important	**More critical**, especially after singulation and handling
Activation-to-bond time lag	Typically shorter and better controlled	Longer and more variable in sequential assembly
Overlay challenge	Wafer-level alignment	Die placement + alignment
Yield loss drivers	Wafer non-uniformity, particles, overlay	**Added die-level variability**, residues, queue time, handling damage
Key manufacturing challenge	Wafer-scale uniformity	**Managing die-level variability**

➤ Figure.3 Wafer-level vs die-level hybrid bonding: the process reality. Die-level flows add thinning/singulation, surface handling, residue/particle risks, activation-to-bond time lag, and stricter overlay control–often with limited opportunity for self-alignment once contact occurs.

Die-level hybrid bonding must overcome almost all wafer-level challenges—surface preparation, characterization, mechanism control, overlay, reliability—while also addressing die-specific problems. A representative example is thinning and singulation. Wafer-level flows can often thin and dice after bonding; die-level flows typically thin and singulate first. If the hybrid-bonding surface is formed upfront, it must survive interactions with temporary bonding materials or backgrinding tape, and singulation. Residues, contamination, and micro-particles that are less critical at wafer level can become direct yield killers at die level.

Singulation therefore requires careful engineering. Plasma dicing can reduce residues and particles, but cost and equipment constraints matter. Blade dicing continues to improve as well, with approaches aimed at minimizing particle generation. Beyond singulation, die bonding tools are under active development, including handling and transport schemes that minimize or avoid contact with sensitive surfaces—approaching non-contact pick-and-place [5].

Another die-level bottleneck is the activation-to-bond time lag. In wafer-to-wafer flows, this interval can often be kept short because both wafers are processed and bonded in close succession, for example within a wafer-to-wafer cluster bonder. By contrast, sequential die-level assembly involves repeated transport, alignment, and bonding steps, resulting in longer and more variable queue times for each die. Surface deactivation during these delays can reduce bond strength and ultimately impact yield. Moreover, standardized methods to quantify these effects—particularly robust bond-strength metrology—are not yet fully established, which further slows development.

Hybrid bonding is also discussed as a potential future enabler for high-bandwidth memory stacking. Yet achieving very high yields across extremely tall stacks (for example, ~18–20-high) and managing demanding configurations such as face-to-back hybrid bonding with Through-Si vias (TSV), remain a major challenge.

Outlook: From cloud AI to edge/physical AI

Hybrid bonding changes more than metrics such as resistance, interface thickness, thermal pathways, and pitch. Once BEOL-direct bonding becomes practical, it reshapes the boundary between interconnects and architecture and has the potential to update design philosophy itself. When aligned with the chiplet concept, hybrid bonding becomes not merely a bonding technique but a starting point for new integration structures—a bridge between front-end and back-end technologies that is likely to become indispensable for future 3D integration.

Significant barriers remain, especially for die-level hybrid bonding. Process-driven variability—introduced by thinning, singulation, handling, residue control, and activation-to-bond time lag—can severely impact yield. A promising strategy is to migrate the hardest aspects of die-level assembly into wafer-like workflows. One such direction is reconstructed die-to-wafer (D2W) hybrid bonding,

which reconstitutes multiple dies on a carrier into a unified “pseudo-wafer”. This allows subsequent planarization, interconnect, and bonding to leverage wafer-level process maturity while reducing reliance on individual die handling. Intel has described a closely related concept under the term quasi-monolithic chips (QMC)[6]. Gap fill, planarization control, and placement-error propagation remain key challenges, but the direction is clear: wafer-like reconstitution may open new paths for chiplet-scale integration.

This direction is particularly compelling for edge and physical-AI systems, where constraints on form factor, power, heat dissipation, and long-term reliability are often tighter than in cloud deployments. By enabling finer-pitch, shorter interconnects and tighter co-integration without resorting to large interposers or oversized substrates, wafer-like reconstitution can help reduce interconnect energy, improve thermal paths, and support compact, robust modules suited to harsh operating environments and extended lifetimes.

Hybrid bonding will not be a sprint; it is a medium- to long-term competition spanning process, metrology, design, and test. The prize, however, is substantial: chiplets not merely as a cost optimization tool, but as a platform that liberates system-level design freedom.

FURTHER READING

- Y. Kagawa et al., “Novel stacked CMOS image sensor with advanced Cu2Cu hybrid bonding,” 2016 IEEE International Electron Devices Meeting (IEDM), San Francisco, CA, USA, 2016, pp. 8.4.1-8.4.4, doi: 10.1109/IEDM.2016.7838375.
- M. Tagami, “CMOS Directly Bonded to Array (CBA) Technology for Future 3D Flash Memory,” 2023 International Electron Devices Meeting (IEDM), San Francisco, CA, USA, 2023, pp. 1-4, doi: 10.1109/IEDM45741.2023.10413718.
- R. Sato .et. al, 2025 Jpn. J. Appl. Phys. 64 03SP55 DOI 10.35848/1347-4065/adbafc
- K. Ryan et al., “Integration, Materials and Equipment Innovations to Enable 100 nm Pitch W2W Bonding for Memory-to-Logic and Logic-to-Logic 3D Stacking,” 2025 IEEE 75th Electronic Components and Technology Conference (ECTC), Dallas, TX, USA, 2025, pp. 542-546, doi: 10.1109/ECTC51687.2025.00096.
- Y. Yoshihara et al., “Degradation Mechanisms in Die-to-Wafer Hybrid Bonding Governed by Surface Activation Lifetime and Moisture Evaporation,” in IEEE Transactions on Components, Packaging and Manufacturing Technology, doi: 10.1109/TCPMT.2025.3636933.
- A. Elsherbini et al., “Enabling Next Generation 3D Heterogeneous Integration Architectures on Intel Process,” 2022 International Electron Devices Meeting (IEDM), San Francisco, CA, USA, 2022, pp. 27.3.1-27.3.4, doi: 10.1109/IEDM45625.2022.10019499.

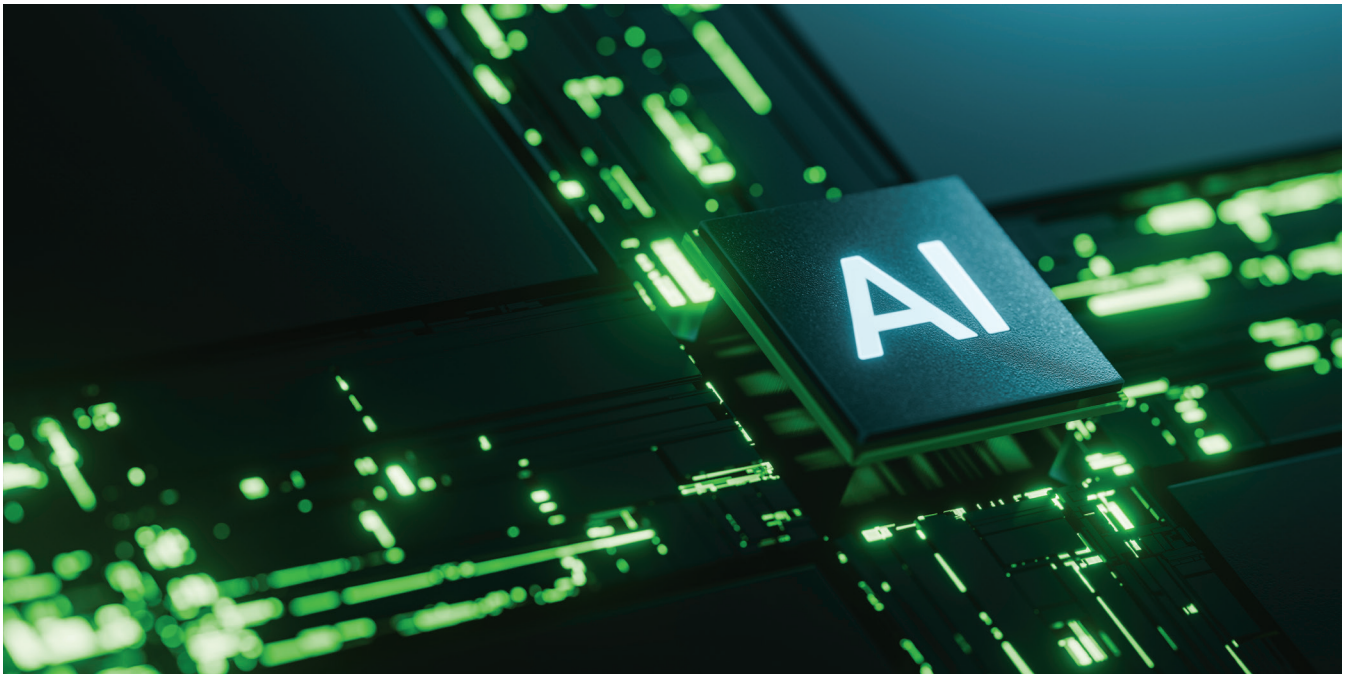

Fiber to laser

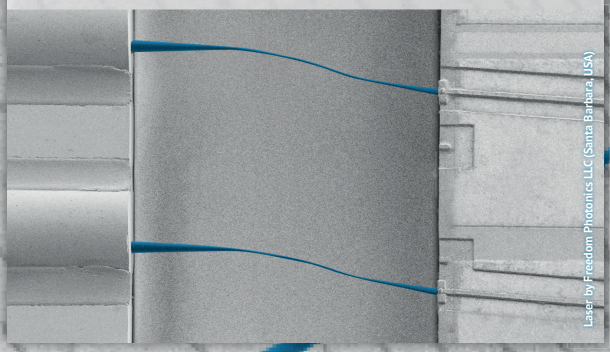

Laser by Freedom Photonics LLC (Santa Barbara, USA)

Laser to SOI chip

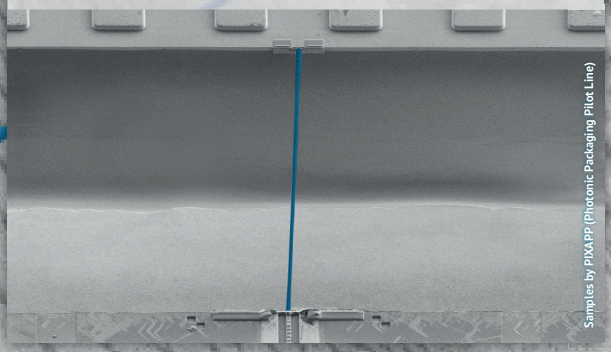

Samples by PIXAPP (Photonic Packaging Pilot Line)

Lens on laser

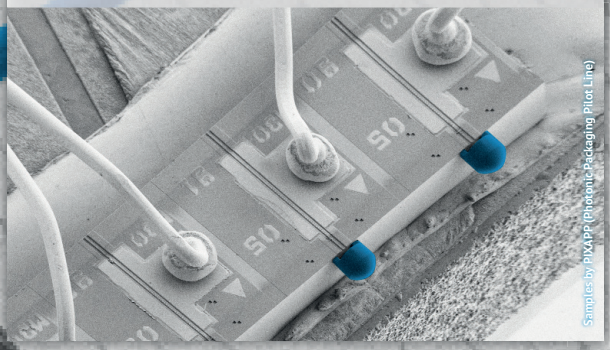

Samples by PIXAPP (Photonic Packaging Pilot Line)

Large scale micro-optics

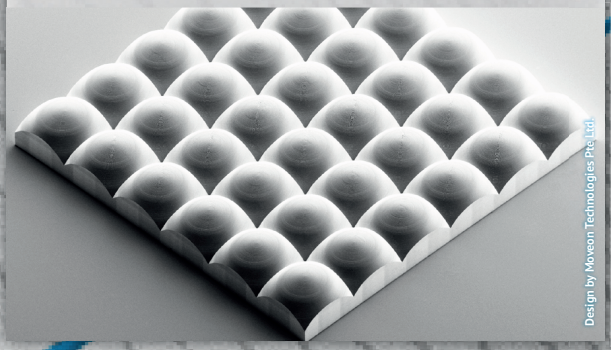

Design by Moveon Technologies Pte. Ltd.

Connect the World by lighting the way

- the essential link in next-generation optical connectivity!

vanguard AUTOMATION

MYCRONIC

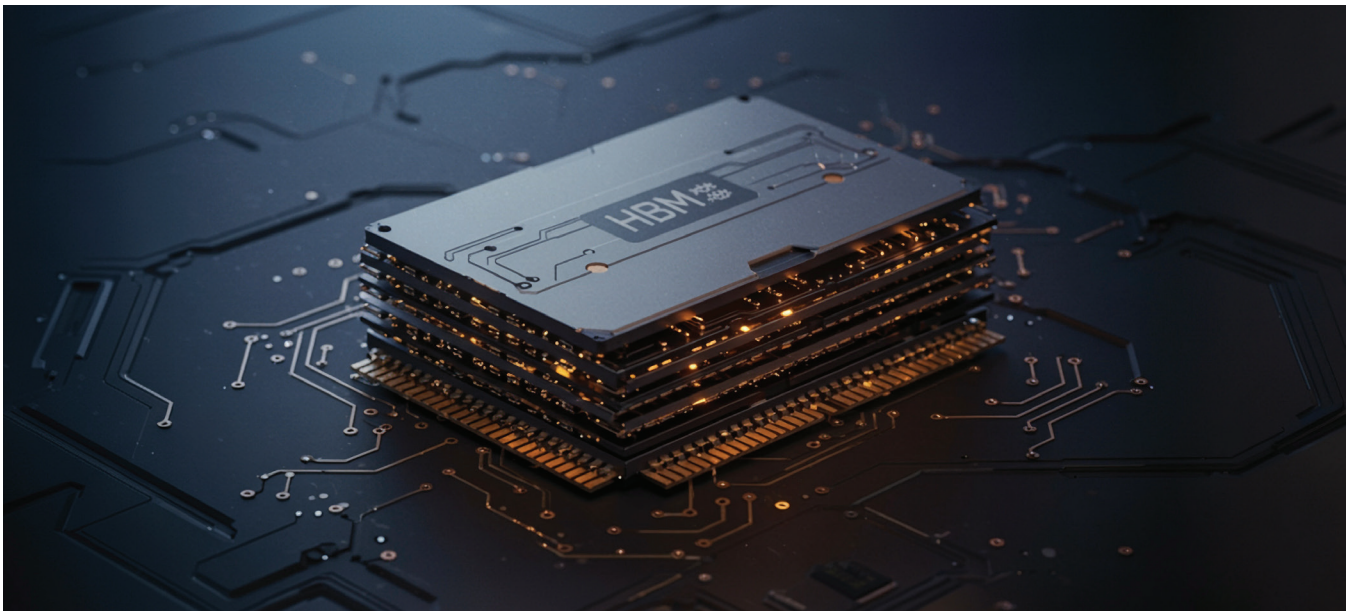

AOR TCB paves the way to higher HBM4 stacks

Current advanced packaging technology is reaching its limits in terms of the roadmap for high-bandwidth memory. Among ASMPT's solutions to enable the next generation of HBM is thermocompression bonding with active oxide removal (AOR).

DR. AMI EITAN, SVP, CHIEF SCIENTIFIC OFFICER, ASMPT

IN THE PAST, advances in hardware found their application; today, in the age of AI, it is the applications that place specific demands on the further development of chips and their processing. This is particularly true for high-bandwidth memory (HBM). HBM is increasingly becoming a potential bottleneck in the advancement of AI devices, in terms of both capability and capacity to meet the global demand for AI products. Therefore, memory suppliers will continue to push HBM performance and look for high-quality processes to enable the manufacture of HBM stacks. The number of memory dies to be stacked continues to increase, as do the potential die sizes. In April 2025, JEDEC officially released the HBM4 specification. It supports data rates of up to 8 Gb/s per pin across a 2048-bit interface, delivering aggregate bandwidth of up to 2 TB/s and supporting 4-Hi to 16-Hi DRAM stack configurations with per-die densities of 24 Gb or 32 Gb, enabling total stack capacities of up to 64 GB. The roadmap is in place, and various bonding technologies are being used in high-volume manufacturing. Currently, Mass Reflow-Moulded Underfill (MR-MUF) and Thermocompression with Non-Conductive Film (TC-NCF) are used in the high-volume manufacturing of HBM. However, when it comes to stack height, current advanced packaging technologies are reaching their limits.

Where the challenges lie

The key challenges in stacking more and more layers of dies are:

- Die warpage during the bonding process
- Chip gap height control
- High thermal and electrical conductivity
- A residue-free underfill process

All these aspects must be considered to ensure quality and reliability. In addition, advancing HBM performance requires continued pitch scaling of the interconnects between dies as the industry moves from HBM3E to HBM4 and beyond to HBM5. The two leading methods can be evaluated against these challenges as follows. Mass Reflow-Moulded Underfill (MR-MUF): Mass reflow melts all bumps at once, followed by moulded underfill encapsulating and filling gaps in a single, integrated step. The primary advantage is the superior thermal dissipation due to the high thermal conductivity of moulded underfill, which is ideal for high-power and tall HBM stacks. MR-MUF is a highly efficient process combining bonding and underfill in one step. It prowvides strong structural support for tall 12–16-Hi stacks. However, it is difficult to control the warpage, and the chip gap height and uniformity when heating the full stack.

In contrast, Thermocompression with Non-Conductive Film (TC-NCF) simultaneously forms

joints and fills gaps die by die. The key advantage of this technology with a pre-laminated non-conductive film is its ability to handle ultra-fine pitches of less than 20 µm. Its high alignment accuracy makes it ideal for increasingly dense I/O designs, such as fine-pitch HBM stacking and advanced 2.5D/3D packaging, where tight keep-out zones and precise alignment are critical. It enables good warpage control with thin dies, and a separate capillary underfill step is not required because the NCF fills the gap during bonding. Another advantage of TC-NCF is that it enables better control of the chip gap height. The disadvantages are lower throughput and higher cost, as well as lower thermal conductivity due to the polymer NCF.

With the transition to next-generation HBM, tighter chip gap height control and improved uniformity have become essential for precise stack height management. The HBM roadmap should therefore move towards hybrid bonding to improve thermal and electrical conductance and achieve very tight chip gap heights. This is in addition to pitch scaling, which requires high bonding accuracy while limiting solder volume. This is important because, as we move from HBM3E to HBM4 and then HBM5, pitch scaling for the joints between the dies is being driven in order to advance HBM performance. ASMPT proposes a new approach to enable the next generation of HBM: thermocompression bonding with active oxide removal. The thermocompression bonding solution, FIREBIRD TCB, is a highly capable platform offering high force capability, high accuracy and high-throughput processes. It provides all the necessary functions for the HBM device bonding process, for the short term and long-term roadmap of memory suppliers.

Residue-free fluxless bonding

AOR TCB™ is a new fluxless first-level interconnect (FLI) process with active oxide removal (AOR). Using a plasma-based approach, the AOR technology empowers 3D chiplet integration and the HBM devices with fine bump pitch roadmaps as well as new package architectures. By eliminating flux residue and the costly cleaning solutions typically associated with traditional methods, AOR will mark a new era in FLI bonding processes, enabling high-volume manufacturing and driving advancements in advanced packaging technology. Deploying the AOR approach aims to improve package interconnect yield at both finer pitch levels and overall larger package sizes. As stack heights increase, even tiny variations in chip gap, residue, or warpage can cascade into significant yield loss. For memory manufacturers operating on a large scale, yield is not just a technical metric; it directly affects profitability and time to market. HBM4 requires consistent bonding quality across every layer of the stack. AOR TCB™ reduces rework and scrap, thereby reducing cost per bit. Fluxless, residue-free bonding significantly reduces failures caused by voids or contamination, providing a measurable cost advantage, while better uniformity enables faster time-to-yield for new HBM nodes. The controlled chip gap and oxide-free bonding of AOR TCB™ shorten the optimisation phase when fabs transition from HBM3E to HBM4, enabling high-volume manufacturing (HVM) to stabilise more quickly.

In-situ cleaning

The fluxless HBM die stack TCB faces a significant challenge with formic-acid-based oxide removal, where salt crystal residues may persist and negatively impact the moulded underfill (MUF) process, by leaving voids and foreign material (FM). These residues and FM can impede underfill adhesion and cause reliability failures. An additional post-bond cleaning step is therefore essential to remove these residues before the MUF process. In contrast, AOR's plasma cleaning process produces only water vapour, creating a highly oxide-free surface on the bump and pad for joint formation. As the plasma-based approach generates no residues, no downstream cleaning operations are required to remove salts and other FM generated by the

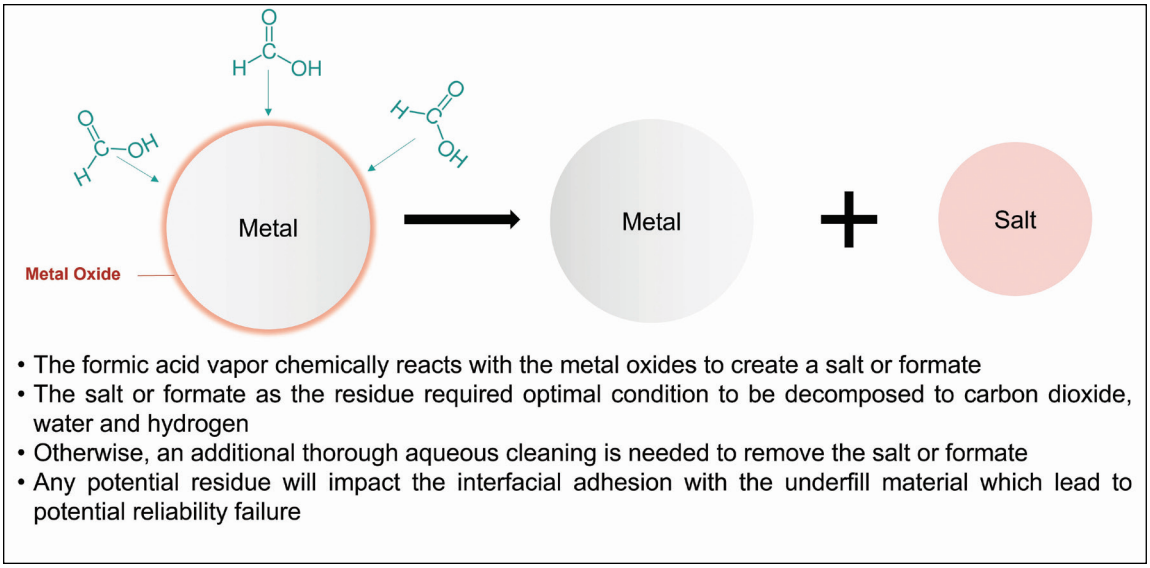

➤ Figure 1: Formic acid oxide removal process results in the formation of microcrystals that can be deposited on the die, wafer, or substrates being bonded, potentially leading to reliability issues. SOURCE: ASMPT

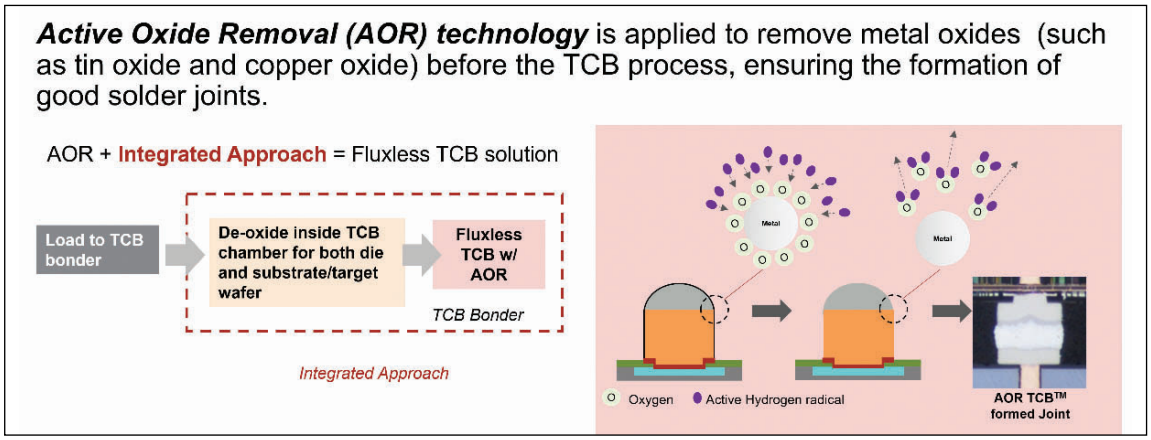

➤ Figure 2: Depicting the integration of the AOR process into the TCB bonder, this plasma-based approach generates no residues, ensuring all surfaces are clean for bonding and downstream processing. SOURCE: ASMPT

acid-based oxide removal process. A residue-free process is essential for meeting quality and reliability standards without increasing production costs. This exemplifies the distinct advantage that AOR plasma cleaning offers to the industry.

It is expected that the appropriate TCB bonder design equipped with a robust oxide removal process (AOR) will enable a clean and reliable HVM interconnection process for chiplet integration at bump pitches below 10 µm, with roadmap scaling toward 5 µm. Through a partnership with a major HBM supplier, a successful demonstration of a 16-Hi AOR TCB™ bonded sample with a maximum stack-up height of 775 µm has been achieved. This demonstrates how HBM4 technology can extend the current roadmap further with advanced AOR TCB™ bonding equipment. For this application, placement accuracy will be improved from traditional values of <2 µm to <0.8 µm. Furthermore, MUF process chip gap height control will be enhanced to address lower solder volumes on micro-bump architectures using next-generation equipment. The introduction of new bond head heaters will also showcase faster heating and cooling capabilities, further improving the quality and speed of fluxless bonding.

Enabling the progress of HBM

The future of HBM stacking technologies may depend on larger and thinner dies, lower electrical and thermal resistance, and tighter pitch. ASMPT offers solutions for a wide range of advanced packaging tasks and knows that its solutions empower the intelligence revolution. Two machines represent the pinnacle of HBM assembly technology: The FIREBIRD Series thermocompression bonding solution, which is the first to feature the innovative AOR TCB™ technology that ensures clean, residue-free surfaces for reliable interconnect formation, thereby improving package integrity and overall performance. The second is the LITHOBOLT™, a next-generation die-to-wafer hybrid bonding solution for 3D integration that delivers ultra-high precision control to ensure superior interconnect quality and high productivity. Hybrid bonding represents the most advanced integration approach currently available.

It is expected that the appropriate TCB bonder design equipped with a robust oxide removal process (AOR) will enable a clean and reliable HVM interconnection process for chiplet integration at bump pitches below 10 µm, with roadmap scaling toward 5 µm.

ASMPT conducts research and development in all technologies with the potential to enable the efficient and high-precision mass production of HBM and is also in constant dialogue with manufacturers. The AOR TCB™ approach is one of the most promising technologies in this field. Strong reliability performance has been demonstrated across various metallisation schemes, including solder-on-solder, solder-on-Ni/Au, and solder-on-Cu. ASMPT will further test the method for fine-pitch Cu-to-Cu bonding down to <5 µm. This initiative aims to bridge the gap between TCB and hybrid bonding, ensuring a smooth transition, with AOR serving as the key enabler over alternative fluxless technologies.

As challenges to HBM performance continue, further innovation and progression in HBM architecture and stacking will be required. In this context, advanced fluxless interconnect solutions such as AOR-enabled thermocompression bonding provide a scalable pathway to sustain yield, reliability and cost efficiency as stack heights increase and bump pitches move further into the sub-10 µm regime. Continued process optimisation and close collaboration across equipment suppliers, material providers and device manufacturers will be essential to translate these capabilities into stable high-volume manufacturing for next-generation AI and high-performance computing systems.

ANGELTECH

CS INTERNATIONAL CONFERENCE | PIC INTERNATIONAL CONFERENCE | PE INTERNATIONAL CONFERENCE | AP INTERNATIONAL CONFERENCE | AIS ANGELTECH INNOVATE SUMMIT

SAVE THE DATE

12-14 APRIL 2027

Sheraton Brussels Airport Hotel

To find out more about our sponsor and speaker opportunities, contact us today on:
+44 (0) 2476 718970

or email:

info@csinternational.net

info@picinternational.net

info@peinternational.net

info@angel-tech.net

Advanced packaging at the limit: Where wet chemical precision meets cost-efficiency

As advanced packaging pushes redistribution layer (RDL) geometries to sub-micron limits, manufacturers must balance extreme process precision with the economic realities of high-volume production. Innovations in wet chemical processing and batch spray technologies are emerging as critical enablers for achieving both yield reliability and cost-efficient advanced packaging at scale.

BERNHARD HAMMERL, BUSINESS DEVELOPMENT, SICONNEX

THE SEMICONDUCTOR industry is going through a massive shift. For decades, Moore's Law relied almost entirely on front-end scaling – the relentless push to shrink transistor nodes. But as we get closer to the physical limits of silicon atoms and face the massive costs of High-NA EUV (Extreme Ultraviolet) lithography, shrinking things at the transistor level is simply getting too expensive.

The real action, and the biggest driver of performance gains, has decisively shifted to Advanced Packaging. Concepts that used to be considered experimental niches – like Fan-Out Wafer Level Packaging (FOWLP), 2.5D silicon interposers, and true 3D heterogeneous stacking (chiplets) – are now the go-to standard for high-performance computing (HPC), artificial intelligence (AI) accelerators, and top-tier mobile processors.

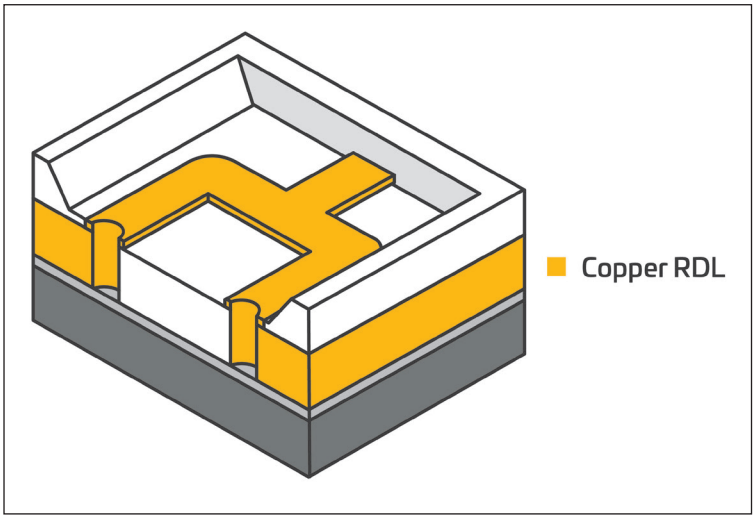

➤ A cross-section of an RDL stack showing the dielectric, barrier layer, seed layer and electroplated copper to visualize the structural complexity.Photo Credit: Siconnex.

But this new era of system-level integration comes with a catch: it demands the kind of strict precision, cleanliness, and defect control we used to only see in front-end-of-line (FEOL) manufacturing.

Redistribution Layers (RDLs): The critical highway

At the heart of these advanced setups are Redistribution Layers (RDLs). As die sizes shrink and input/output (I/O) pin counts go through the roof, finding the physical space to route signals off the chip becomes a massive bottleneck. RDLs act as microscopic highways that connect the tiny, high-density bumps on the silicon die to the larger pads on the package substrate.

And this is exactly where things get tricky, especially when you are working with standard 12-inch (300mm) wafers and the strict demands of copper.

The Copper Dilemma: Why RDLs are so demanding

Copper is still the undisputed gold standard for RDL metallization. Compared to older aluminum tech, copper gives you much better electrical conductivity and holds up a lot better against electromigration. This is key to keeping the Resistance-Capacitance

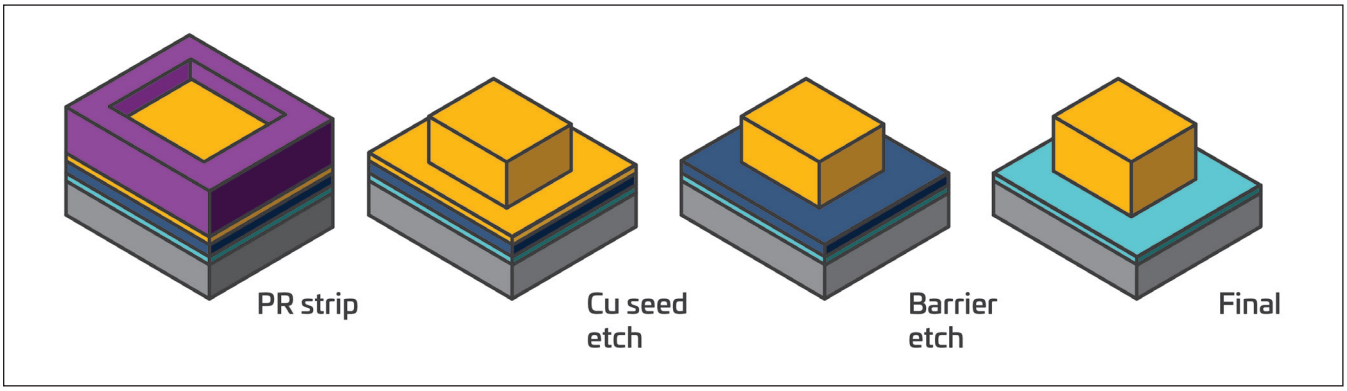

Photo Credit: Siconnex

delay (RC) low, which dictates how fast electrical signals can travel through the interconnects.

But working with copper at the packaging level gives integration engineers a really tough checklist:

- **Aggressive Scaling:** We are talking extremely tight line/space (L/S) resolution (often pushing sub-2µm in advanced nodes) while keeping the geometric shapes absolutely perfect.
- **Electrical & Mechanical Integrity:** You need low electrical resistance combined with a strong mechanical grip to various organic and inorganic dielectrics (like polyimide or silicon nitride).
- **Yield Reliability:** The process needs to be rock-solid and repeatable, ensuring the thousandth wafer turns out exactly like the first.

The main issue is that these high-end demands often clash hard with manufacturing realities. Moving from the additive process of electroplating to the subtractive wet chemical removal of excess layers is a major bottleneck. This is the make-or-break moment that decides if a highly valuable, nearly finished wafer turns into functional chips or just becomes expensive scrap.

The Standard Semi-Additive Process (SAP) flow and its pitfalls

To understand the holdup, let's look at the classic Semi-Additive Process (SAP) workflow used for copper RDLs:

- **PVD (Physical Vapor Deposition):** A blanket barrier layer (usually Titanium or Titanium Tungsten, Ti/TiW) is sputtered onto the wafer to stop copper from bleeding into the dielectric. A thin copper seed layer goes on right after.
- **Lithography:** A thick photoresist is applied, exposed, and developed to map out the exact negative patterns of the interconnects.
- **Electroplating (ECD):** Copper is electrochemically grown inside the resist of trenches to build the actual structural lines and vias.
- **Wet Chemical Etch:** This critical final stage requires extreme precision to remove the redundant metal. The resist is stripped, followed by a wet-etch of the exposed copper seed layer and the underlying Ti/TiW barrier layer.

These final wet chemical steps are the main reason for yield loss. The chemicals need to be aggressive enough to fully clear the seed and barrier layers, but gentle enough not to eat away the plated copper lines.

- **Under-etching:** Leaves behind metallic leftovers, which causes electrical shorts between adjacent lines.
- **Over-etching:** Leads to galvanic corrosion or structural undercut beneath the copper lines. This causes a loss of critical dimension (CD), mechanical weakness, and eventually, a failed device.

The agony of choice: Wet bench, single wafer, or batch spray?

When setting up a fab for these critical wet etch steps, process engineers hit a technological crossroad. There are three main ways to get the chemistry onto the wafer, and they all have different physical and economic trade-offs.

The Classic: Automated wet bench

Wet benches submerge entire cassettes of wafers into static or circulating chemical baths. They are relatively cheap and offer the highest throughput of all available technologies by processing large batches simultaneously. However, for fine-pitch RDLs, the fluid dynamics are often inadequate. The chemical mix in a bath is too static, leading to serious non-uniformity across the wafer, particularly centre-to-edge variations.

> Modern batch spray systems get around this by using in-situ Endpoint Detection (EPD). Using advanced optical sensors, the system can literally "see" the exact moment the copper seed is cleared, and the underlying barrier layer is exposed. By actively stopping the process at exactly the right moment, the etch step repeatability becomes incredibly robust. It absorbs upstream variances from lithography or deposition without sacrificing yield

Furthermore, wafers frequently drag contaminants from one bath to the next, increasing the risk of defects. While the large footprint of these systems is manageable if floor space is not a constraint, the high chemical consumption remains a significant drawback. Ultimately, for sub-2µm RDLs, the etch precision of a wet bench is insufficient to meet modern integration requirements.

The Sprinter: Single wafer spin process
In this technology, each wafer is processed one by one inside a closed chamber. The wafer spins on a chuck while nozzles spray fresh chemistry right onto the surface. This gives you excellent within-wafer (WIW) uniformity. However, the economics are prohibitive. It suffers from low throughput, takes up a massive footprint on the fab floor (due to the low productivity-per-square-meter), and burns through chemistry because it is a "single-pass" (straight to the drain) system.

The Hybrid Solution: Batch spray technology
This is where innovative designs are shaking up the market. The engineering concept is elegant: it combines the high-precision uniformity of a rotational spray process with the high-throughput economics of a batch tool, processing up to 50 wafers simultaneously. Beyond its impressive performance, this approach offers a compact footprint and significantly lower chemical consumption, making it a highly sustainable and cost-effective solution.

How batch spray offers the "best of both worlds"

Siconnex targets exactly this technical and economic sweet spot. But what happens physically inside the process chamber to allow batch processing to rival single-wafer precision?

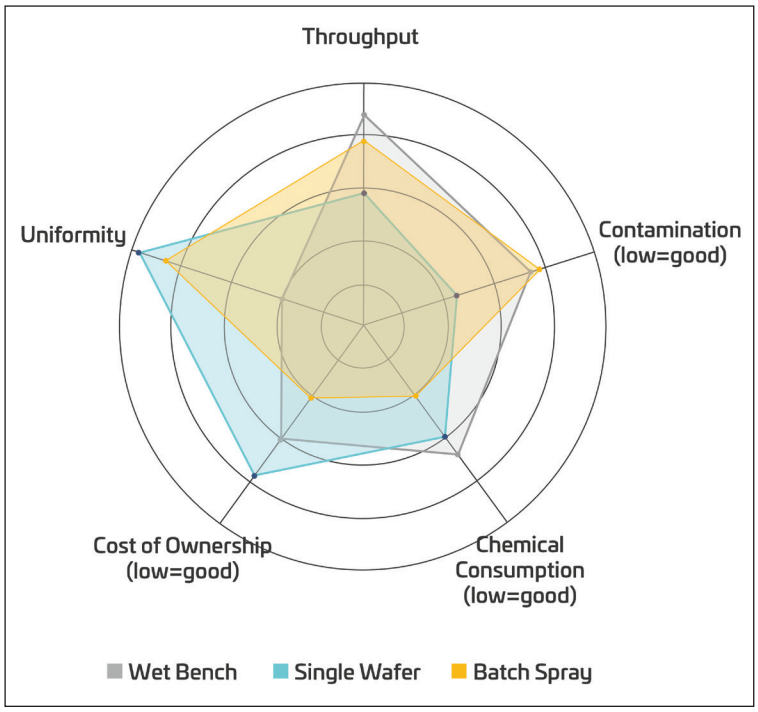

➤ Comparing Single Wafer, Wet Bench and Batch Spray technologies.

Photo Credit: Siconnex

Fluid Mechanics: Rotational dynamics meets spray power
Unlike the static immersion of wet benches, wafers in the Siconnex BATCHSPRAY® system are placed in a rotor inside a sealed chamber. As the rotor spins, an array of precisely calibrated nozzles actively sprays the chemistry onto the wafers.

This relies on advanced fluid mechanics. The rotation speed (RPM) and spray pressure are actively thin and control the hydrodynamic boundary layer on the wafer surface. Perfectly synchronised rotation and spray speeds create a homogeneous, constantly renewing liquid film. This physical agitation ensures that reaction byproducts are instantly swept away, preventing localised chemical depletion at the surface. The result is an etch rate non-uniformity of less than 3% across a 300mm wafer – rivalling the best single-wafer tools but achieving much higher throughput.

Ecology and Economy: Closed-loop recirculation
While burning through chemicals and managing waste are significant challenges for single-wafer systems, batch spray architectures fix this with closed-loop recirculation. Instead of flushing expensive, toxic chemistry to the fab's waste treatment plant, the etchant is continuously collected, filtered for particles, dynamically temperature-controlled, and reused. This closed-loop process locks in the absolute thermal and chemical stability you need for repeatable etch rates on fine-pitch RDLs. Beyond technical stability, this approach slashes consumable costs and massively shrinks the fab's environmental footprint – a crucial metric for modern Environmental, Social, and Governance (ESG) reporting.

Process Integration: Merging the toughest steps into one platform
In the high-stakes world of advanced semiconductor manufacturing, having great fluid dynamics and chemical efficiency is only half the battle. Process integration is fast becoming the ultimate differentiator. By combining photoresist stripping, copper seed etching, and barrier material etching in a single equipment platform, fabs don't just save floor space – they drastically cut down handling complexity, boost process stability, and reduce costs.

Siconnex tackles this head-on with its BATCHSPRAY® platform, a highly customizable wet processing system designed to streamline these advanced backend applications. The core of this system's efficiency lies in its four independently recirculating chemical tanks.

Each loop operates as a fully functional process module, enabling the management of four distinct concentrated chemistries within a single tool. This integrated architecture facilitates the seamless execution of critical Semi-Additive Process steps:

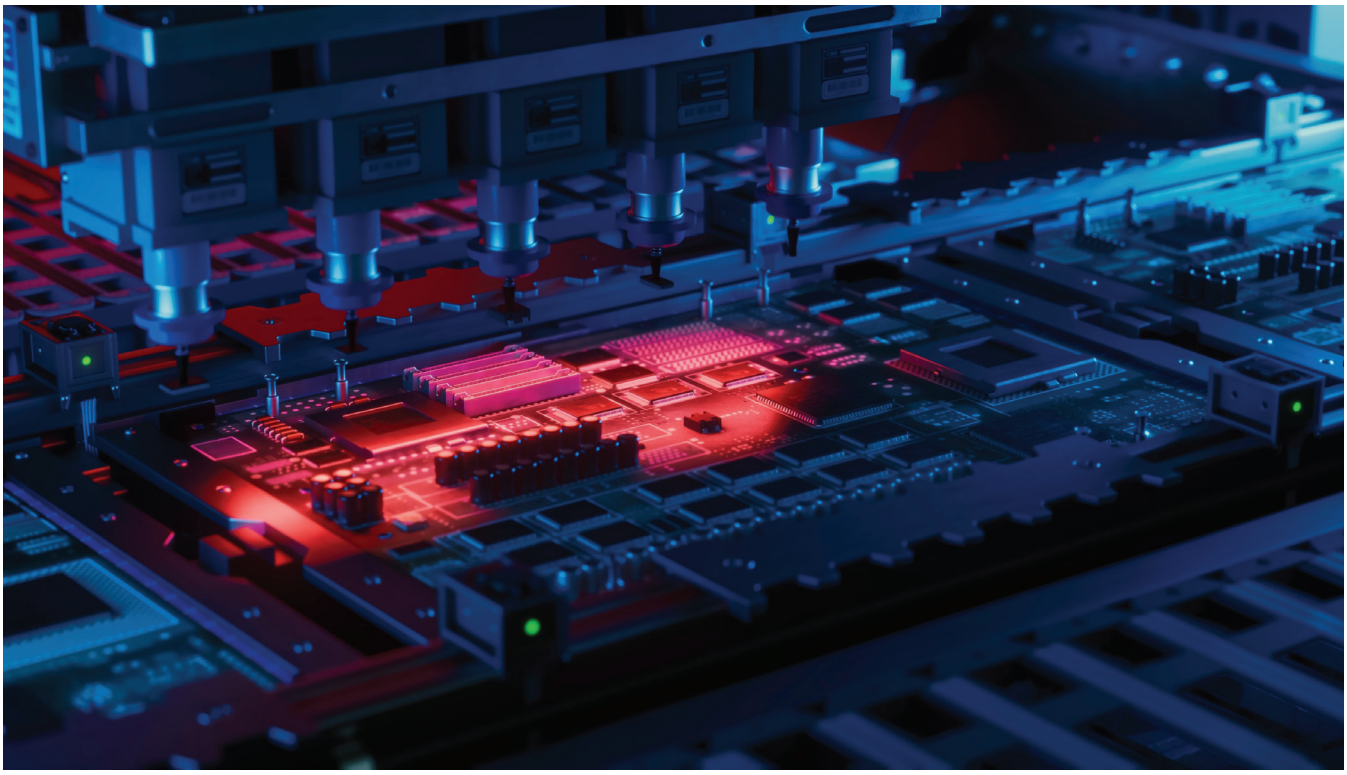

- Photoresist stripping
- Copper seed layer etching
- Barrier material etching

Each tank loop packs its own heater to dial in precise temperatures for optimized reaction kinetics. Integrated filtration systems keep the chemistry pure and particle-free – an absolute must for high-yield processing. Plus, optional inline concentration monitoring gives you real-time control over the chemical makeup, ensuring stable, repeatable results in High-Volume Manufacturing (HVM).

This modular, independent tank design takes the operational drawback out of chemical management while maxing out flexibility. Engineers can tweak parameters for each chemistry individually without sacrificing uniformity or throughput. With the demand for heterogeneous integration exploding, packing multiple wet processes into one configurable platform is a massive advantage.

The Yield Safety Net: In-Situ endpoint detection (EPD)

Process engineering operates on a simple truth: no previous process step is ever truly perfect. Sputtered seed layers and barriers will naturally vary in thickness across a wafer and from batch to batch. Blindly etching based on a set timer creates a big risk of over-etching or under-etching. Modern batch spray systems get around this by using in-situ Endpoint Detection (EPD). Using advanced optical sensors, the system can literally “see” the exact moment the copper seed is cleared, and the underlying barrier layer is exposed. By actively stopping the process at exactly the right moment, the etch step repeatability becomes incredibly robust. It absorbs upstream variances from lithography or deposition without sacrificing yield.

Scalability is the ultimate metric

As fabs scale their Advanced Packaging lines from pilot R&D to High-Volume Manufacturing, the physical footprint of equipment often becomes a major strategic roadblock. Cleanroom floor space is among the most expensive real estate on the planet.

Therefore, density is as critical as throughput. A standalone batch spray system occupies only a fraction of the floor space compared to an equivalent lineup of single-wafer tools. This compact footprint enables seamless capacity expansion within existing cleanroom boundaries, effectively bypassing the need for massive capital expenditures on costly facility extensions.

As geometric and electrical demands push copper RDLs to their limits, staying competitive requires fabs to master the exact intersection of wet chemical precision and cost-efficiency. When looking for a solution that secures both high-yield performance and overall profitability, integrated batch spray solutions offer a proven, physics-based approach that delivers the ideal balance.

By bridging the gap between the high cost of single-wafer processing and the inherent footprint challenges of wet benches, this architecture provides a versatile workhorse capable of meeting modern RDL requirements without compromising on economic efficiency.

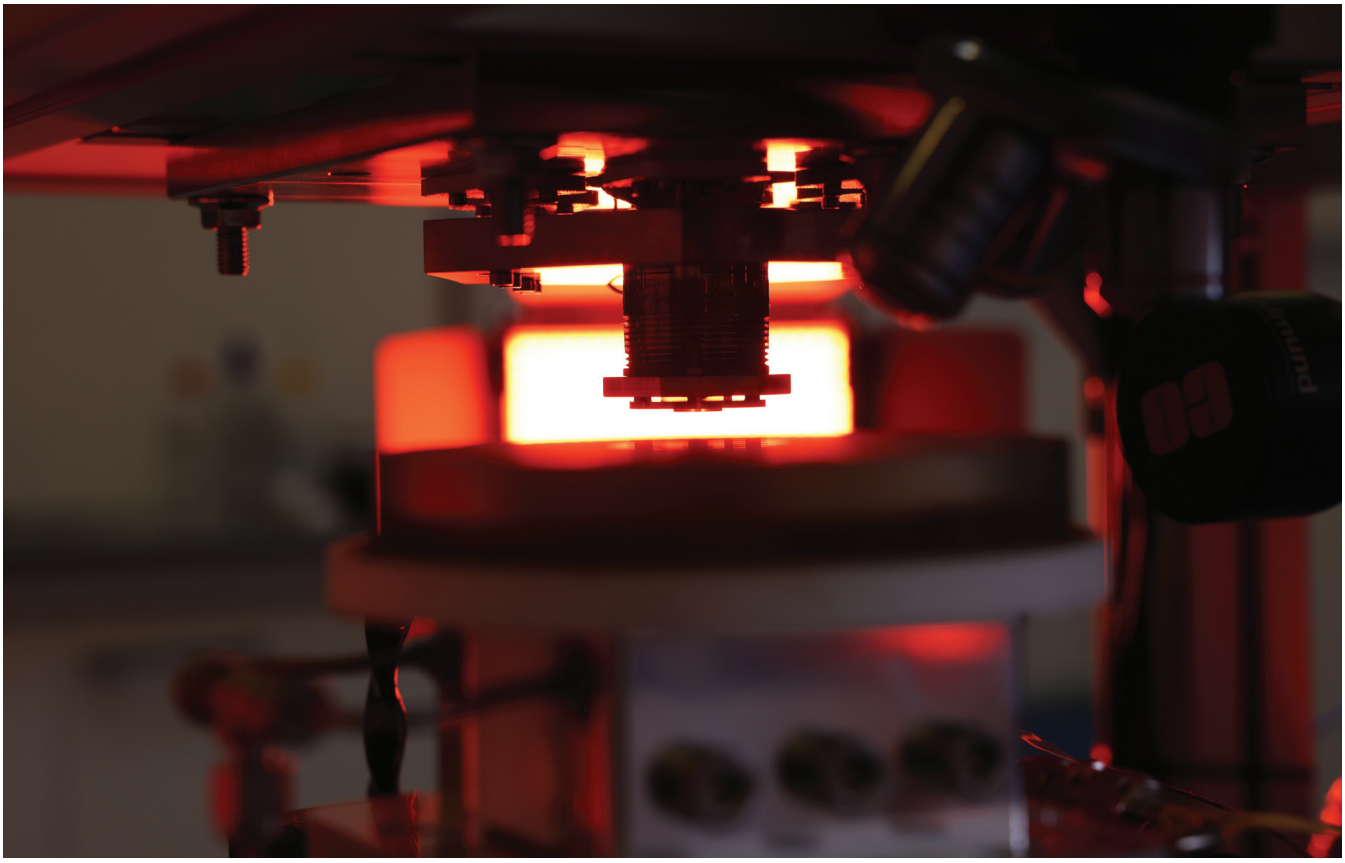

Maximising advanced packaging hinges on manufacturing process optimisation

Unlocking packaging's promise will require solving upstream manufacturing process bottlenecks and rethinking global semiconductor supply chain strategies.

BY DR. MAKSYM PLAKHOTNYUK, CEO AND FOUNDER, ATLANT 3D

AI CHIP revenue is soaring, driven by breakthrough models and unprecedented demand across the GPU and ASIC segments.

This rapid growth benefits semiconductor manufacturers, as each new AI generation requires ever-larger quantities of leading-edge silicon to deliver improved cost, performance, and efficiency per operation.

To keep pace with this demand – and transcend the physical limits of classic transistor scaling—advanced packaging technologies, especially 3D stacking and chiplet architectures, have become a vital frontier. These techniques integrate multiple chips and components (such as CPUs, GPUs, memory, and high-speed interconnects) in a single heterogeneous package.

By minimizing the physical distance between elements, advanced packaging improves data transmission rates and energy efficiency, a key factor in the escalating technological rivalry between the U.S. and China.

Leading foundries and outsourced semiconductor assembly and test (OSAT) players are responding. Notably, TSMC has announced a $100 billion U.S. investment plan – including a next-generation advanced packaging facility in Arizona – while Intel recently expanded its advanced packaging operations in New Mexico, aiming for domestic supply chain resilience.

Such moves are further catalyzed by policy incentives and tariff structures designed to onshore semiconductor manufacturing capacity.

Despite these advances, challenges remain. Advanced packaging is highly complex, involving tightly integrated stacks of diverse materials and structures – such as interposers, redistribution layers (RDLs), and multiple active and passive chips.

Conventional thin-film approaches like atomic layer deposition (ALD)

offer atomic-scale precision but generally lack the throughput or process simplicity required for large-scale packaging. Direct-write additive manufacturing techniques like DALP are emerging to address some of these bottlenecks, enabling more flexible and rapid construction of advanced packages with lower material waste, fewer processing steps and the same precision as ALD.

Meanwhile, the infrastructure that houses AI chips must evolve. Next-generation AI accelerators and advanced packages require new datacenter designs – with upgraded power delivery, cutting-edge cooling, and higher rack densities – to avoid bottlenecks that would offset performance and sustainability gains from packaging improvements.

While advanced packaging reduces the energy needed for interconnects, the aggregate power draw for hyperscale AI datacenters continues to rise in line with massive computational demand.

Crucially, high-volume advanced packaging capacity in the U.S. is not expected to be available before mid-2027, and U.S.-designed chips like NVIDIA's Blackwell family continue to be routed to Taiwan for packaging.

This gap creates opportunities for other countries to accelerate local industry development, and exposes U.S. manufacturers to additional costs and potential tariffs if products must be re-imported after overseas packaging – potentially eroding domestic cost advantages.

Industry consensus is forming around the idea that advanced packaging may be as critical as transistor innovation itself for the future of AI and semiconductors. As NVIDIA CEO Jensen Huang noted,

"To meet [AI's] demand, advanced packaging has become as critical as transistor design in delivering the efficiency and power our customers require." But fully unlocking packaging's promise will require solving upstream

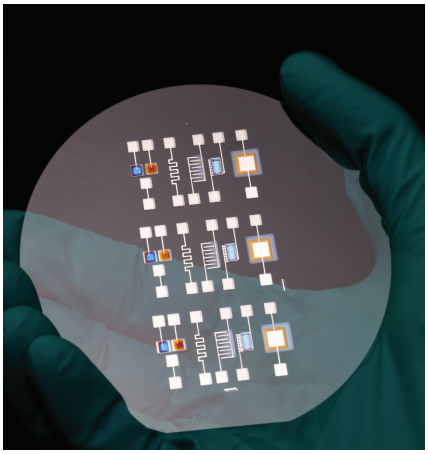

manufacturing process bottlenecks and rethinking global semiconductor supply chain strategies.

Advanced packaging holds the potential to shift power balances not just within AI, but also in fields such as defense, biomedical engineering, and energy, provided nations and firms can adapt their supply chains, infrastructure, and manufacturing processes for a new era of integrated, high-performance computing.

Targeting semiconductor packaging pain points with multiphysics simulation

Multiphysics simulation can be used in semiconductor packaging to predict performance and ensure packaging reliability. This article goes over the areas of packaging where simulation can improve R&D.

BY ANDY CAI, COMSOL, INC.

THE PERPETUAL DRIVE to shrink semiconductor components means ever-increasing complexity — the need for denser integration, tighter thermal margins, and stricter reliability requirements has become more severe. In advanced packaging, especially for emerging 3D chip architectures, this complexity is amplified by vertical stacking, high-density interconnects, and diverse material interfaces. Challenges like warpage, interconnect fatigue, and thermal stress aren't isolated issues; they're tightly coupled and difficult to manage through testing alone. Modeling and simulation has become more common in packaging for addressing such challenges. However, even with the growing adoption of simulation, there remains a general hesitancy to fully embrace it.

It's true that parameters like material behavior carry uncertainty, so numerical models must rely on additional physics assumptions and simplifications. It's also true that no model can capture every detail of the physical world, but the goal of simulation isn't to mirror reality perfectly. Simulation is meant to help teams address design challenges and gather critical insight that can guide development decisions. When built thoughtfully, models can reveal dominant effects and clarify cause-and-effect relationships.

Multiphysics simulation is particularly powerful in semiconductor packaging, where coupled effects (e.g., thermal, mechanical, and sometimes electrical effects) drive system performance and reliability. Moreover, simulation complements testing (rather than replacing it), helping teams reduce iteration cycles and make better decisions with fewer surprises.

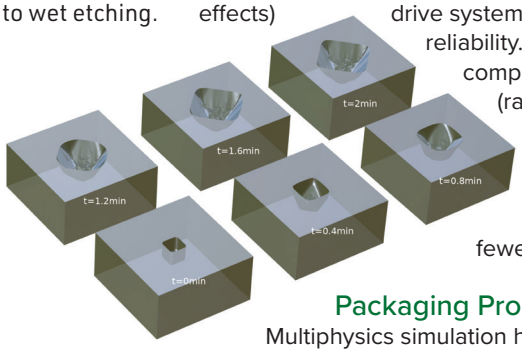

➤ Figure 1. The development of an etched groove shape over time due to wet etching.

Packaging Process Simulation

Multiphysics simulation has made a difference in various areas of semiconductor packaging. Below, we highlight some of these application areas.

Wet and Dry Etching

Both wet and dry etching are essential for the creation of features in advanced packaging, and both processes can be simulated accurately despite their complexities.

When simulating wet etching, chemical reactions as well as mass transport need to be considered. The interaction between fluid transport, diffusion, and reaction kinetics can be simulated and optimized to ensure etch uniformity or minimal undercutting. Conversely, dry etching can require significantly more complex physics couplings, as it involves plasma chemistry, ion transport, and directionally dependent material removal. Regardless, the right tools will enable you to model this process accurately.

Figure 1 shows an example of anisotropic wet etching of silicon, based on the different etching rates of silicon's crystal planes in a KOH solution. In the model, a small initial groove on the wafer surface comes into contact with the KOH solution, and pyramid-shaped grooves gradually form over time due to the crystal-plane–dependent chemical etching reaction.

Soldering Process

Solder joints and the reflow process used to create them are crucial in semiconductor packaging but can also be sources of significant reliability challenges, including thermal fatigue and warpage. To give one example, warping of a die from internal stresses can build up during the reflow process (Figure 2) and adversely affect a component's performance or cause early failure.

The melting and solidification behavior typical in the soldering process is often hard to predict, but it can be accurately simulated to provide insight into how thermal loads will evolve and

how stress will develop as materials transition from liquid to solid. Multiphysics modeling can replicate the metallurgical phase transitions and the residual stress state with multiple metal phase transformations. These visualizations can offer insight into the mechanical reliability of the solder joints and the operating conditions in which they will succeed. Solder joints are typically modeled with temperature- and rate-dependent viscoplastic (creep) constitutive laws, since stresses evolve during reflow and thermal cycling. Elastic–plastic and fatigue/damage models are added as needed to capture irreversible deformation and predict life.

Underfill

Underfill adhesives often exhibit non-Newtonian flow behavior, requiring rheological models and temperature-dependent properties. When these adhesives cure, they create residual stress and warpage. To get a complete understanding of underfill phenomena, it's best to use a multiphysics approach to assess both the chemical kinetics of the curing process and the resulting mechanical deformation within the same simulation space.

Grinding, Dicing, and Molding

Mechanical operations in semiconductor packaging usually involve grinding, dicing, and molding (Figure 3). While these are routine steps, they introduce stresses and carry a risk of cracking in the packaged product. In particular, the risk of cracking is highest near edges and interfaces. Multiphysics simulation can be used to assess the impact of cutting forces and thermal effects on structural integrity.

Package Reliability Simulation

When testing package reliability, it's critical to consider environmental factors such as ambient humidity and moisture ingress through diffusion into materials. Let's talk about the areas where multiphysics simulation can play a role in ensuring packaging reliability.

Humidity and Moisture Ingress

The presence of moisture has damaging consequences: hygroscopic swelling, corrosion, and delamination risks. A multiphysics approach supports coupled heat- and moisture-transfer (hygro-thermal) modeling, allowing users to simulate moisture, ingress over time, and AC interactions with temperature and mechanical stress. This modeling is particularly useful for predicting failure during preconditioning or accelerated life testing, as it can reduce the amount of prototyping and testing.

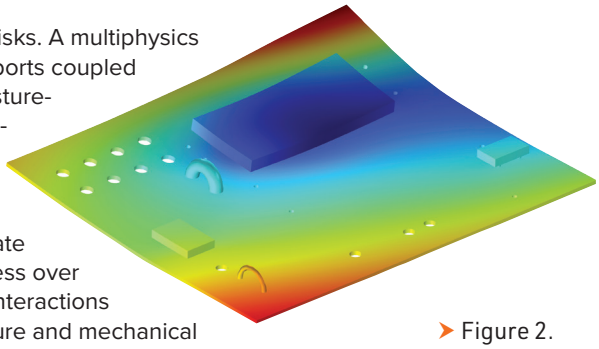

➤ Figure 2. Warpage in a semiconductor after reflow soldering.

Structural Damage and Failure

When things go wrong, multiphysics simulation can be used for detailed failure analysis. Whether the failure is due to interfacial delamination, crack initiation and growth, or even thermal fatigue of solder joints, there are methods for virtually assessing what went wrong and how it can be prevented. Users can create models that include fracture mechanics for crack risk estimation as well as simulate progressive crack growth. In Figure 4, you can see a stress singularity at the crack tip in a sample plate geometry.

Shock, Vibration, and Static Loads

When modeling semiconductor packaging, it is important to account for mechanical shock and vibration and evaluate how these loading conditions affect package reliability. Models can be used to analyze the vertical displacement of a motherboard from a response spectrum evaluation, as shown in Figure 5. This need is especially heightened for mobile or automotive applications where semiconductor packages must withstand various dynamic stresses.

Users can perform a variety of dynamic mechanical analyses, including frequency-domain studies, impact simulations, response spectrum evaluations, and random vibration modeling. Similar approaches can be applied to assess package response to static

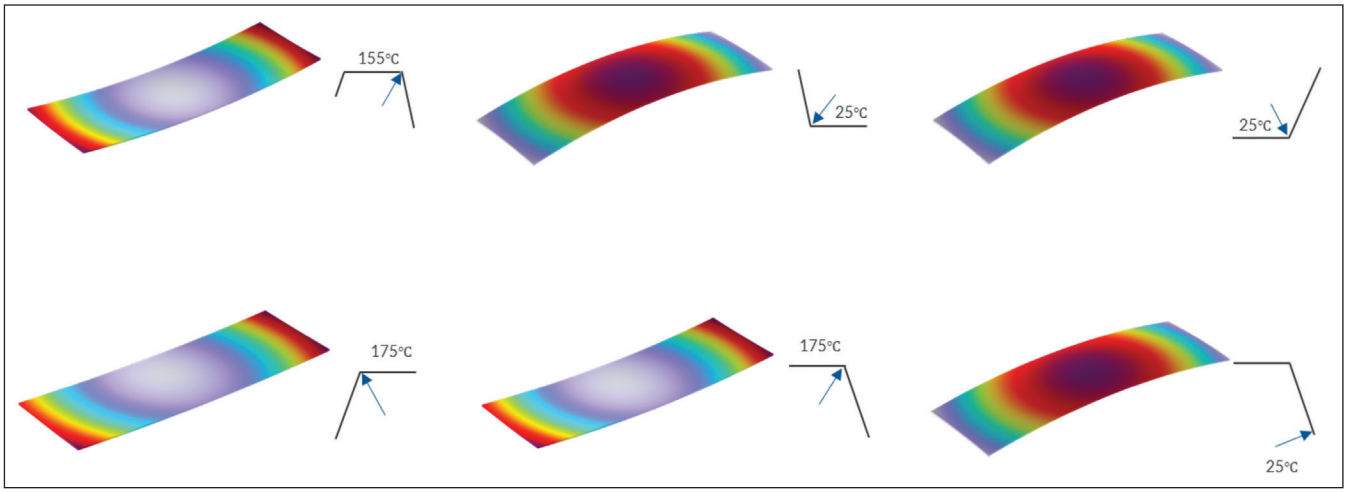

➤ Figure 3. Warpage evolution during the epoxy molding compound (EMC) molding process.

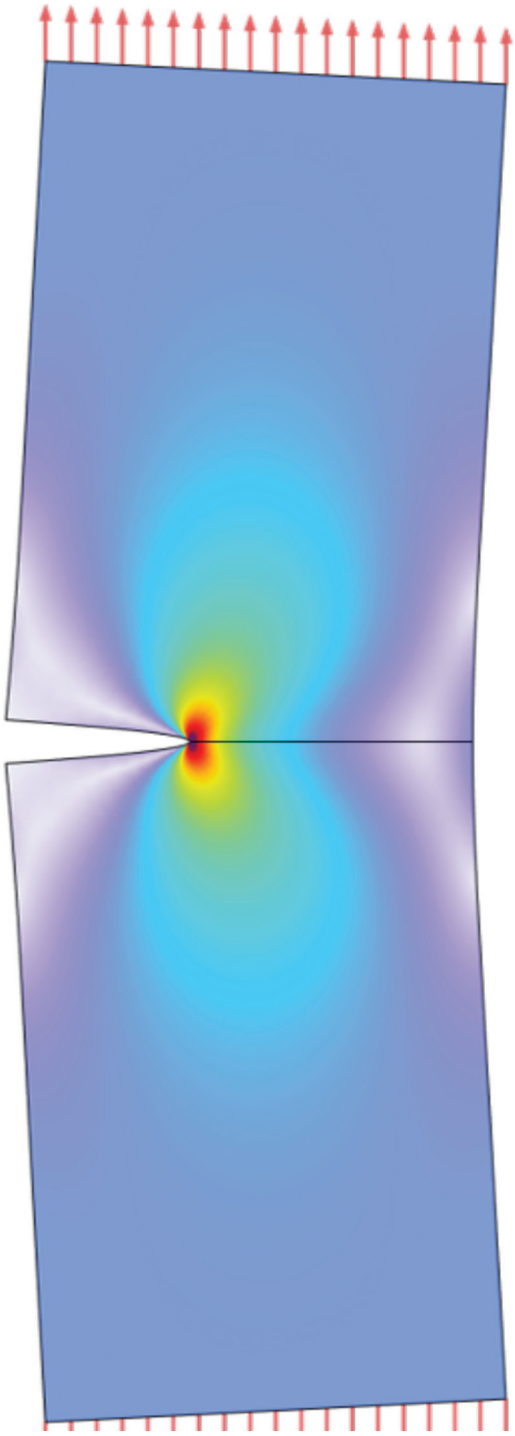

➤ Figure 4. A geometry showing the von Mises stress and the deformed shape of a cracked plate. Note that the displacement is exaggerated to illustrate the deformation under the applied load.

mechanical stresses, such as bending, torsion, or compression. These evaluations support design validation for a range of scenarios, such as a four-point bending test.

Thermal Management in Packaging

Thermal management is a multiphysics process that requires thermoelectromechanical simulation and CFD simulation. In advanced packaging specifically, these models require significant computational resources for computation. This is particularly true when working with 3D stacked layers or, in general, when including designs with a large number of components.

Simulation users benefit by coupling the three primary heat transfer modes in the same model: conduction, convection, and radiation. This coupling gives users the option to decipher the individual impact on different pieces of the overall system. For example, for a device in a package with air gaps, heat sinks, and a surrounding enclosure, users may want to include all three heat transfer modes but find that only one or two of the mechanisms will be the dominant dissipation or heating mechanism. Simulation enables users to easily turn on and off each mechanism and test it until they have a better understanding of how each one impacts the results.

In electronic packages, heat sources often arise from electromagnetic loss mechanisms such as resistive (ohmic) heating, eddy current losses, or dielectric heating in RF components. Simulation can couple these electrical and magnetic effects with thermal and mechanical responses, incorporating temperature-dependent material properties for metals and substrates.

Electromagnetic Performance in Packaging

Electromagnetics simulation can be used to analyze how packaging affects signal propagation, including high-frequency effects such as skin and proximity effects. It can also evaluate how different geometries or materials influence transmission-line performance. This capability is valuable for predicting and mitigating signal integrity issues early in the design process, before hardware fabrication, to help rule out potential problems. A multiphysics approach enables coupling of detailed finite element (FEM) electromagnetic models with circuit-level simulations, supporting mixed-domain modeling in which part of the signal path is analyzed in full-wave detail while the remainder is represented by circuit elements. Users can automatically generate SPICE-compatible circuit models from extracted data and simulate signal reflections, losses, or noise to trace their origins to specific physical structures, which can be useful for both design and verification.

Bringing Simulation to the Factory Floor

Let's switch gears and discuss a specific modeling tool and how it can spread the use of simulation throughout organizations.

Through simulation apps, models can be brought into the field or onto the factory floor. These custom-made simulation tools can present team members with the information most relevant to their work, as opposed to presenting a full software user interface or unrelated data. App users can easily modify input parameters and study the computational results, even if they don't have foreknowledge of the underlying model or simulation software.

In semiconductor packaging, complex models can be compiled into apps and shared with team members who are not simulation experts. Figure 6 shows a thermoelectric cooler simulation app

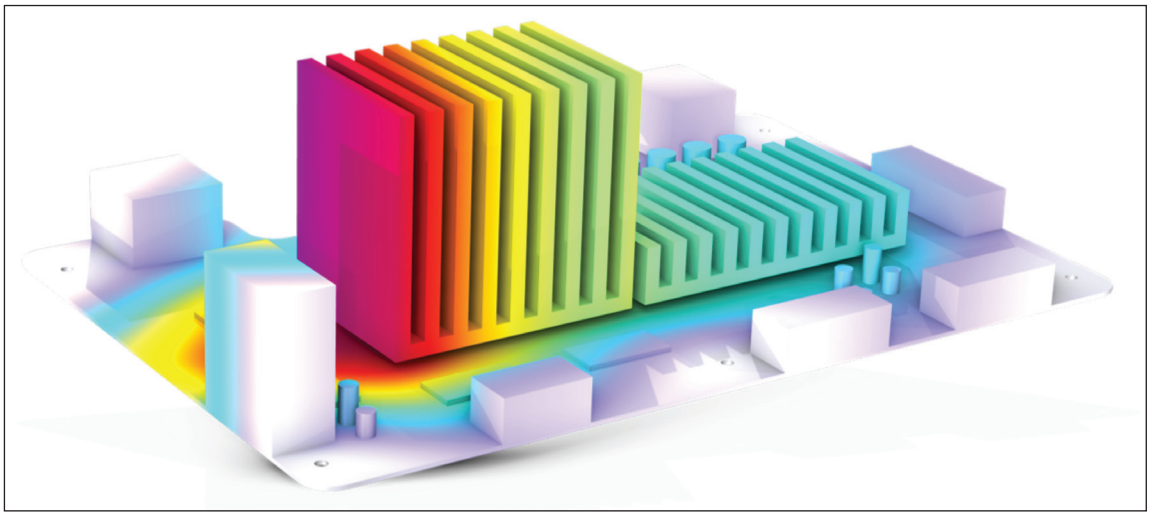

➤ Figure 5. A motherboard's vertical displacement as a result of a response spectrum evaluation.

with a specialized interface that contains only the parameters needed for specific analysis cases. This type of app can be used to examine current and temperature distributions for different inputs, enabling critical analyses to be performed easily and without the help of a simulation expert.

In this article, we have touched on common application areas within semiconductor packaging where multiphysics simulation can be used.

Multiphysics is how the world works, and semiconductor packages are proof of this point. With devices only getting more complex, it's only right that the tools we use to assess their design and operation have powerful, growing functionality.

Multiphysics simulation can help give R&D departments an understanding of where they need to move their design in tandem with their real-world experimentation.

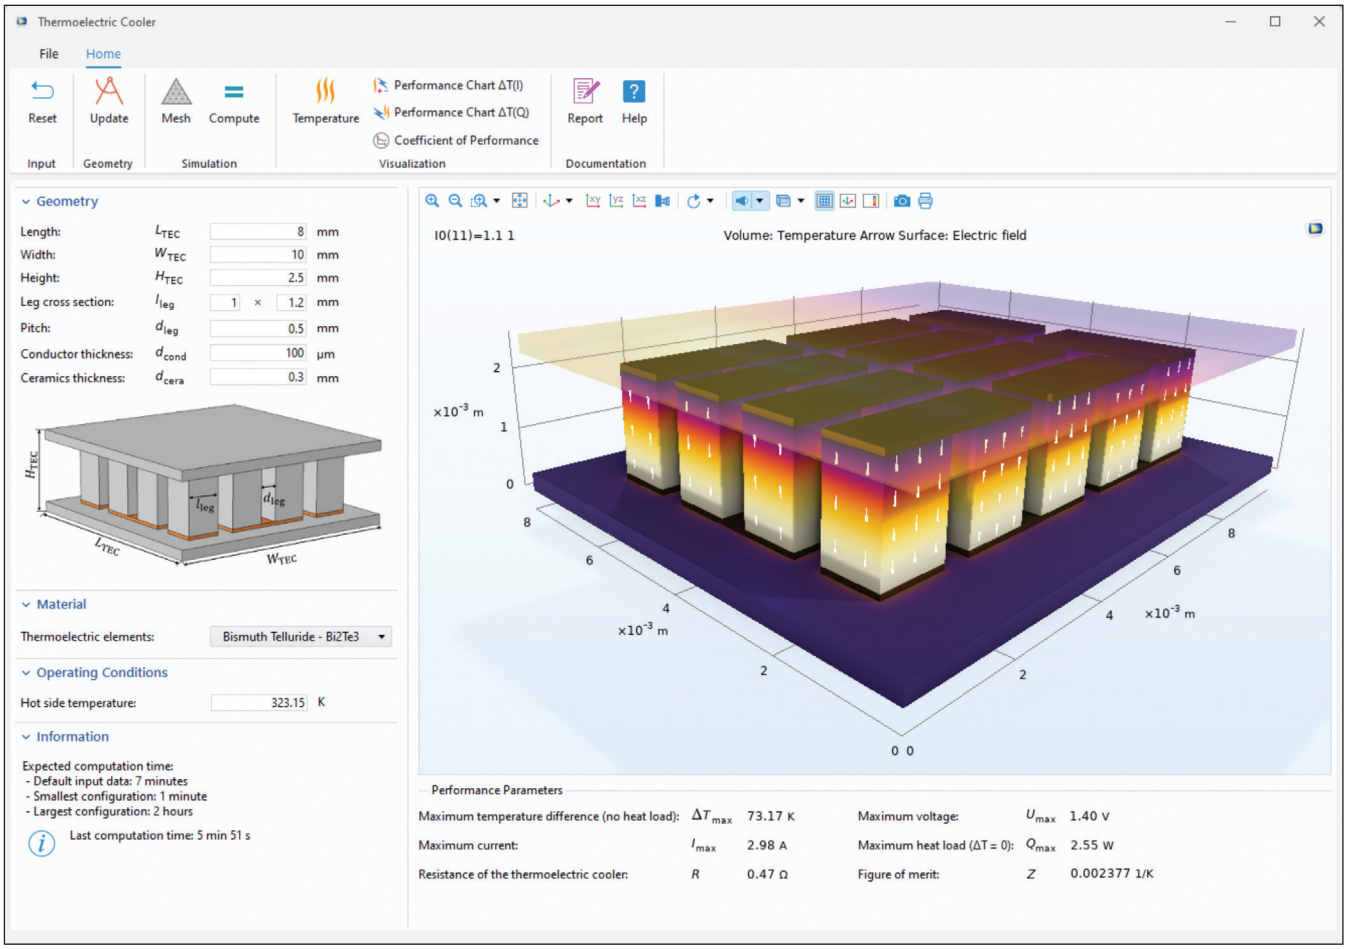

➤ Figure 6. A simulation app that enables users to analyze single-stage thermoelectric cooler designs by testing various geometries, thermocouple configurations, and materials.

EVG®

invent
innovate
implement

www.EVGroup.com

FUSION AND HYBRID BONDING FOR HETEROGENEOUS INTEGRATION

- Enabling advanced 3D device stacking and chiplet integration for CMOS image sensors, memory and 3D system-on-chip (SoC)
- High-volume production equipment and process technology for Wafer-to-Wafer (W2W) and Die-to-Wafer (D2W) hybrid bonding
- GEMINI® FB Automated Production Wafer Bonding System delivers industry leading alignment accuracy and bonding performance
- Heterogeneous Integration Competence Center™ serving as leading-edge innovation incubator for EVG customers and partners

GET IN TOUCH to discuss your manufacturing needs
www.EVGroup.com