




ADVANCED PACKAGING

ISSUE II 2026

 AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION





ADVANCEDPACKAGING.NEWS



SHAPING THE NEXT DECADE OF ADVANCED PACKAGING INNOVATION

MYCRONIC

Think big. Print nano.



Quantum X align

High-precision 3D printing of aligned micro-optics

Scale photonic integration and co-packaged optics with freeform micro-optics printed directly on fibers, chips, and wafers.



Packaging moves to the centre of the chip

➤ For decades, semiconductor progress was largely defined by what happened on the silicon. Packaging was important, but it was often viewed as the final stage of the process: a way to protect the chip, connect it to the outside world and prepare it for integration into a system.

That relationship has changed.

As transistor scaling becomes more challenging and computing demands continue to grow, more of the performance equation is now being determined beyond the individual die. Chiplets, heterogeneous integration, 2.5D and 3D architectures, hybrid bonding and photonic integration are all changing what a semiconductor package needs to achieve. The package is no longer simply carrying the chip. Increasingly, the package is part of the chip's performance.

This issue explores what that shift means when advanced packaging moves from the laboratory to volume manufacturing.

One of the recurring themes across our feature content is that achieving an impressive demonstration is only the beginning. A process that works on a small number of devices can encounter an entirely different set of challenges when it has to operate continuously, repeatedly and economically across thousands or millions of units. Accuracy, repeatability, process control and yield become just as important as the underlying technology.

Hybrid bonding is a good example. It promises extremely fine-pitch interconnects and greater integration density, but it also demands exceptionally clean surfaces, tight control of flatness and precise vertical and lateral alignment. As pitches shrink, there is simply less room for process variation. Surface preparation, inspection and bonding therefore become interconnected parts of the same manufacturing challenge.

The same principle applies to photonic integration. Bringing optical and electronic components closer to the processor can address the bandwidth and energy challenges facing AI infrastructure, but it introduces another layer of precision into the package. Alignment, optical coupling, thermal

stability and long-term process repeatability all have to be controlled at production scale.

This is why I believe one of the most important conversations in advanced packaging today is not simply what can we integrate?, but how can we manufacture that integration reliably?

The answer will increasingly depend on how early packaging enters the design process. Packaging cannot be treated as something that happens after the architecture has been finalised. When a system contains multiple dies, different materials, high-density interconnects and potentially optical components, decisions made at the design stage can have a direct impact on assembly complexity, yield and cost.

There is another important shift taking place: the boundary between development and manufacturing is becoming less distinct. The ability to transfer a process from R&D into high-volume production without effectively starting again is becoming a competitive advantage. In-line inspection, process monitoring, automation, traceability and shared process platforms can all help reduce the gap between proving a technology and producing it reliably.

Our other feature on optical memory pooling reinforces this broader lesson from a different perspective. Disaggregating memory at the rack level does not remove the physical challenges inside the package and module. Electrical routing, thermal management, optical stability, power delivery, manufacturing and serviceability all still have to be solved at the physical level.

Ultimately, advanced packaging is becoming a systems-engineering discipline.

The winners will not necessarily be the technologies that achieve the smallest pitch, the highest bandwidth or the most impressive laboratory demonstration. They will be the technologies that can bring those gains together with yield, reliability, manufacturability and economics.


That is the real challenge for the next generation of advanced packaging and increasingly, it is where the future of computing will be decided.

Contents

Cover Story

Shaping the next decade of advanced packaging innovation

As advanced packaging evolves to meet the demands of AI and high-performance computing, manufacturers face growing challenges around yield, precision, throughput and process stability.




14

20 From lab to volume: What advanced packaging actually requires

Heterogeneous integration, hybrid bonding and photonic integration are carrying Moore's Law beyond 2D wafer-level manufacturing

24 Optical memory pooling still has a physical realisation problem

As AI systems evolve, balancing HBM cache, DDR5 capacity and optical I/O will make module granularity critical to building memory-pooling architectures that can be efficiently cooled, validated and scaled



24

28 Empowering violet lasers


More powerful violet lasers, realised by improving design and introducing superior thermal management, will target laser direct imaging for advanced semiconductor packaging

32 Agentic AI reshapes PCB and advanced packaging design

Phil Bishop, VP of Corporate Marketing, Global Customer Success Team at Cadence, talks to Phil Alsop, Contributing Editor of Advanced Packaging Magazine, about how agentic AI orchestration, combined with trusted EDA and system design and analysis tools, can help engineering teams manage the growing complexity of PCB and advanced package development.

36 Glass gains attention in the AI hardware race

As AI accelerates demand for more powerful semiconductor architectures, advanced packaging has become a critical battleground for improving performance, bandwidth and energy efficiency.



40 Streamlining the concept-to-production journey as advanced packaging gains momentum


An interview with Phil Alsop, contributing editor of Advanced Packaging, and Percy Gilbert (top image), Senior Vice President of Engineering at SkyWater Technology, explores how the company enables customers to move from early-stage concepts to manufacturable technologies through its Technology-as-a-Service (TaaS) model.

44 From nanometre alignment to high-volume manufacturing

Building on the key themes explored at the AngelTech Advanced Packaging International Conference 2026, this report examines how chiplets, heterogeneous integration and co-packaged optics are pushing advanced packaging towards a precision manufacturing discipline.

NEWS

- 06 Photonic chip market to hit \$9.3bn
- 06 Google AI push boosts packaging demand
- 07 NHanced advances photonic packaging
- 08 Thailand advances packaging plans
- 09 Industry backs open CPO standards
- 09 AI fuels advanced packaging growth
- 10 Advanced Packaging bottleneck hits Apple
- 10 BBCube platform advances AI chip packaging
- 11 FAMES to present 3D integration advances



Editor

Sarab Chopra sarab.chopra@angelbc.com

Contributing Technical Editor

Richard Stevenson richard.stevenson@angelbc.com +44 (0)1923 690215

Sales and Product Manager

James Cheriton james.cheriton@angelbc.com +44 (0)2476 718970

Senior Sales and Product Development Executive

Ranjodh Avern ranjodh.avern@angelbc.com +44 (0)2476 718975

Sales Executive

Nathan Diehl nathan.diehl@angelbc.com +44 (0)2476 7823123

Design & Production Manager

Mitch Gaynor mitch.gaynor@angelbc.com +44 (0)1923 690214

Graphic Design & Digital Media Assistant

Harvey Watkins mitch.gaynor@angelbc.com +44 (0)1923 690214

Publisher

Jackie Cannon jackie.cannon@angelbc.com +44 (0)1923 690205

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



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. Angel Business Communications Ltd © Copyright 2026. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. 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. © Copyright 2026. 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.

Photonic chip market to hit \$9.3bn

AI data centres, co-packaged optics and silicon photonics are expected to drive strong growth in the global photonic chip market through 2032.

The global photonic chip market is projected to grow from approximately \$4 billion in 2025 to \$9.3 billion by 2032, according to a new report from MarketsandMarkets.

The market is expected to expand at a compound annual growth rate of 12.9% between 2026 and 2032, driven by increasing bandwidth requirements in AI data centres, the adoption of co-packaged optics (CPO) and continued commercialisation of silicon photonics.

Optical interconnects are identified as the largest application segment, reflecting growing demand for higher-bandwidth and more energy-efficient connections across AI and cloud infrastructure.

North America currently represents the largest regional market, while thin-film lithium niobate (TFLN) is identified as the fastest-growing material platform.

The report highlights the physical limitations of electrical interconnects in terms of bandwidth, power consumption and signal integrity as an important factor accelerating photonic integration.

Generative AI and the rapid expansion of GPU clusters are further increasing the importance of optical connectivity within and between data centre systems.

Silicon photonics currently holds the largest share among material platforms, supported by its compatibility with CMOS manufacturing and an expanding PDK and foundry ecosystem.

TFLN, meanwhile, is gaining momentum for high-performance electro-optic modulation in areas including coherent communications and microwave photonics.

Hybrid integration remains the leading integration approach by revenue,

while heterogeneous integration is forecast to grow rapidly as technologies including flip-chip bonding, wafer bonding and micro-transfer printing mature.

These techniques are becoming increasingly important for integrating III-V active devices with silicon photonic platforms and supporting CPO architectures.

MarketsandMarkets also identifies the transition from 400G towards 800G and 1.6T connectivity, greater availability of photonic PDKs and multi-project wafer services, and the commercialisation of TFLN as significant trends shaping the sector.

Asia Pacific is forecast to be the fastest-growing regional market, while continued investment in photonics manufacturing, AI infrastructure and semiconductor R&D is expected to support expansion across North America and Europe.

Google AI push boosts packaging demand

GOOGLE IS REPORTEDLY planning to produce between 12 million and 15 million ninth-generation Tensor Processing Units (TPUs) by 2028, a scale that could rival Nvidia's projected AI accelerator shipments and significantly increase demand for advanced semiconductor packaging.

According to a research note from Fubon, the TPU v9 will feature a four-die architecture, reflecting the industry's shift toward multi-die chiplet designs for AI processors.

Scaling production to the reported volumes is expected to more than

double Google's packaging capacity requirements compared with 2027.

The report suggests that advanced packaging capacity could become a key constraint, with TSMC potentially unable to meet the anticipated demand alone.

Intel Foundry is reportedly positioned to play a role after securing orders to produce millions of Google TPUs, following successful evaluations of its advanced packaging technologies.

The reported partnership highlights the growing importance of advanced

packaging in next-generation AI hardware, where technologies such as Intel's EMIB and TSMC's CoWoS platforms are enabling the integration of multiple high-performance compute dies within a single package.

While Google has not officially confirmed its production plans, the reported expansion underscores how advanced packaging has become a critical enabler for scaling AI infrastructure and next-generation accelerator deployments.

NHanced advances photonic packaging

NHanced Semiconductors and the University of Florida will present advanced packaging approaches using hybrid bonding and silicon interposers for photonics, CPO and AI interconnects.

NHanced Semiconductors and the University of Florida will present research on the role of hybrid bonding and silicon interposers in advanced photonic packaging at the 2026 SMTA International Conference in Rosemont, Illinois.

The paper, titled Advanced Packaging for Photonics: Hybrid Bonding and Silicon Interposer Integration, will be presented by NHanced Semiconductors Vice President Dr. Charles Woychik and Dr. Navid Asadizanjani of the Florida Semiconductor Institute at the University of Florida.

The work focuses on fine-pitch Cu-Cu hybrid bonding and high-density silicon interposers for 2.5D heterogeneous integration of chiplets, including silicon photonics devices.

According to the researchers, hybrid bonding can provide dielectric-to-

dielectric bonding and direct copper interconnects at pitches below 10 μm , with potential scaling to below 5 μm .

Eliminating conventional micro-bumps can reduce interconnect height, parasitic capacitance and inductance, improving signal integrity and energy efficiency for high-speed photonic-electronic interfaces.

Silicon interposers provide another part of the integration strategy, using high-density redistribution layers and through-silicon vias to bring heterogeneous chiplets together within 2.5D architectures.


Combining the two technologies could provide the alignment accuracy required for low-loss optical coupling while helping address the electrical, optical and thermal requirements of co-packaged optics and AI interconnects.

However, significant manufacturing challenges remain. These include bonding defects, overlay accuracy, thermal expansion mismatch and heat dissipation within increasingly dense device stacks.

Addressing these issues will require closer co-optimisation of materials, process integration and thermal management.

The researchers see these advanced packaging approaches as an important foundation for scaling high-density, energy-efficient photonic-electronic integration in next-generation data-centric systems.

The 2026 SMTA International Conference will take place from 25–29 October at the Donald E. Stephens Convention Center in Rosemont, Illinois.



Thailand advances packaging plans

Thailand's new Siam Silica strategy aims to establish two advanced packaging facilities by 2030 as part of a broader effort to strengthen the country's semiconductor ecosystem.

Thailand has unveiled its Siam Silica semiconductor roadmap, placing advanced packaging at the centre of its strategy to become a leading semiconductor hub in Southeast Asia.

Announced by the Ministry of Higher Education, Science, Research and Innovation (MHESI), the initiative identifies advanced packaging as one of five strategic focus areas, alongside photonics fabrication, silicon design, quantum photonics and power devices.

The roadmap targets the establishment of two advanced semiconductor packaging facilities, one fabrication plant and eight integrated circuit design companies by 2030.

The strategy is designed to move Thailand beyond its traditional role in semiconductor assembly, packaging and testing by expanding into higher-value manufacturing and design activities.

The government is also launching workforce development programmes to build the engineering and technical skills needed to support the country's growing semiconductor sector.

Thailand is also seeking closer collaboration with ASEAN partners to strengthen regional semiconductor supply chains while attracting investment in advanced manufacturing capabilities.

Industry analysts have welcomed the roadmap but note that Thailand faces strong competition from established regional leaders such as Malaysia and Singapore.

They argue that sustained investment in technology, talent and manufacturing infrastructure will be essential if the country is to achieve its long-term ambitions.

By prioritising advanced packaging within its national semiconductor strategy, Thailand is aiming to strengthen its position in the global semiconductor value chain while supporting growing demand from AI, high-performance computing and next-generation electronic systems.



Industry backs open CPO standards

OCP initiative brings industry leaders together to develop interoperable silicon photonics architectures for next-generation AI infrastructure.

A 19-company industry coalition has formally launched the Open Silicon Photonics for AI Systems initiative within the Open Compute Project (OCP), aiming to establish an open architecture for deploying co-packaged optics (CPO) across large-scale AI infrastructure.

Led by Lightmatter, the initiative brings together companies including Celestica, Dell Technologies, Flex, Foxconn Interconnect Technology, Global Unichip Corporation, Hyve Solutions, Keysight, Qualcomm and Quanta Cloud Technology.

The group has published a foundational architecture white paper designed to guide development of a shared CPO blueprint for Modular Hardware System (MHS) and Open Rack v3 (ORv3) platforms.

The architecture is intended to support AI clusters scaling from 72 to more

than 1,024 nodes while enabling multi-vendor interoperability.

The initiative comes as increasing AI accelerator density and rising SerDes speeds place greater pressure on conventional electrical interconnects.

With SerDes rates approaching 448G, the coalition says copper reach is becoming increasingly constrained, accelerating the transition towards optical connectivity within AI systems.

A key objective is to create a technology-agnostic infrastructure that can accommodate different photonic technologies and successive generations of optical engines without requiring major changes to the surrounding rack infrastructure.

The proposed architecture is also intended to support both scale-up

and scale-out networking, including protocols such as NVLink and UALink, while allowing multiple fabric planes to coexist on common passive optical infrastructure.

“Interconnect has become as fundamental to AI infrastructure performance as compute,” said Nick Harris, founder and CEO of Lightmatter.

He added that the white paper provides the foundation for developing a shared MHS-based CPO rack architecture compatible with hyperscaler deployment models.

The coalition expects the first specifications emerging from the initiative to be submitted in the fourth quarter of 2026, as the industry works towards moving CPO from early deployments into more widely adopted AI infrastructure.

AI fuels advanced packaging growth

The global advanced chip packaging technologies market is expected to reach \$87.6 billion by 2030, growing at a compound annual growth rate (CAGR) of 14.8% from \$38.6 billion in 2024, according to a new report from BCC Research.

The report identifies artificial intelligence and high-performance computing as the primary drivers behind the market's expansion, as semiconductor companies increasingly rely on advanced packaging technologies to deliver higher bandwidth, lower latency and greater interconnect density than conventional packaging approaches can provide.

Asia-Pacific remains the largest regional market, valued at \$19.9

billion in 2024, supported by the concentration of leading OSATs, foundries and integrated semiconductor manufacturing ecosystems across Taiwan, South Korea, Japan and China.

Among the fastest-growing technology areas are 2.5D and 3D integration, chiplet architectures, high-bandwidth memory (HBM) packaging, hybrid bonding, Fan-Out Panel-Level Packaging (FOPLP) and co-packaged optics (CPO), all of which are becoming increasingly important for AI accelerator and data centre deployments.

The report also highlights growing investment across the advanced packaging value chain. TSMC continues to expand CoWoS capacity while

integrating packaging and testing into its Foundry 2.0 strategy, and SK Hynix is investing \$12.9 billion in new HBM packaging capacity.

Meanwhile, government-backed initiatives including the U.S. and EU CHIPS Acts are accelerating domestic advanced packaging capabilities in response to supply chain resilience and geopolitical priorities.

BCC Research notes that advanced packaging is evolving from a manufacturing process into a strategic differentiator, enabling heterogeneous integration and next-generation AI systems while creating significant opportunities for foundries, OSATs and packaging technology providers.

Advanced Packaging bottleneck hits Apple

Strong production of Apple's 2 nm A20 Pro processor is reportedly being held back by limited DRAM supplies, highlighting advanced packaging as a key manufacturing bottleneck.

A shortage of mobile DRAM is reportedly slowing production of Apple's next-generation A20 Pro processor by delaying advanced packaging, despite strong manufacturing yields on TSMC's 2 nm process.

According to reports, Apple's A20 Pro is the first processor to use TSMC's Wafer-Level Multi-Chip Module (WMCM) packaging technology, which integrates the application processor and high-performance DRAM at the wafer level.

The approach replaces TSMC's long-standing Integrated Fan-Out (InFO) packaging for Apple's flagship mobile chips.

While production of the A20 Pro on TSMC's N2 node is said to be

progressing smoothly, limited DRAM availability is preventing completed processor wafers from moving into the WMCM packaging stage.

The report claims TSMC is currently holding around \$1 billion worth of finished processor wafers awaiting memory integration.

The packaging technology is being qualified on pilot production lines at TSMC's Advanced Packaging Fab 3 in Longtan and is expected to ramp at its AP7 facility in Chiayi.

The reported bottleneck also highlights the growing role of advanced packaging in semiconductor manufacturing.

As logic and memory are increasingly integrated within a single package, packaging capacity and component availability are becoming as critical as wafer fabrication in determining production schedules.

The delay could reduce the time available for final device assembly ahead of Apple's next iPhone launch, although suppliers are reportedly working to secure sufficient DRAM to meet expected demand.

The situation underscores how advanced packaging technologies such as WMCM are becoming central to next-generation mobile processor manufacturing.

BBCube platform advances AI chip packaging

Researchers at the Institute of Science Tokyo have developed BBCube, a next-generation semiconductor integration platform that combines advanced chip packaging, high-density interconnects and thermal management to address key challenges in AI hardware.

The platform integrates three core technologies: a proprietary high-density chip-on-wafer (COW) process, bumpless chip interconnection technology and a multiscale thermal analysis method, supporting advanced 2.5D and 3D heterogeneous integration for AI accelerators and high-performance computing systems.

BBCube uses via-last through-silicon vias (TSVs) to connect chips without conventional metal bumps, enabling finer-pitch interconnects and reducing chip-to-chip spacing to 10 μm .

According to the researchers, the architecture has the potential to deliver up to 16 times greater aggregate signal bandwidth within the same interconnection area while maintaining signal quality comparable to conventional micro-bump designs.


To address heat generation in densely integrated packages, the team also

developed a waffle wafer structure that reduced thermal resistance by approximately 52% in simulations.

A complementary multiscale thermal analysis technique enables heat distribution across an entire chip to be evaluated at 1 μm resolution using 100 million analysis points.

Together, the three technologies form a unified platform designed to support compact, high-performance AI systems with improved connectivity, efficient cooling and scalable heterogeneous integration.

The researchers believe BBCube could help accelerate the development of future AI accelerators and high-performance computing devices as semiconductor packaging complexity continues to increase.



FAMES to present 3D integration advances

The European pilot line will showcase progress in 3D integration, hybrid bonding and heterogeneous integration technologies at ESSERC 2026.

The FAMES Pilot Line and the SiNANO Institute will present the latest advances in advanced semiconductor technologies, including 3D integration and heterogeneous integration, during a joint workshop at the European Solid-State Electronics Research Conference (ESSERC 2026) in Palma de Mallorca, Spain, on 7 September.

The half-day workshop, FD-SOI Pilot Line Insights and Perspective, will highlight technical progress achieved by the €830 million FAMES initiative, one of the five pilot lines established under the European Chips Act.

Presentations will span advanced FD-SOI technologies, embedded non-volatile memories, RF components,

power management ICs and advanced integration technologies.

For the advanced packaging community, the programme will feature new developments in 3D sequential integration, including low-temperature process innovations enabling vertically integrated devices and the first millimetre-wave RF circuits fabricated using a full 3D sequential integration process.


The workshop will also highlight the growing role of heterogeneous integration in AI systems.

Researchers will present advances in small-pitch hybrid bonding and through-silicon via (TSV) technologies,

which are being developed to support next-generation computing architectures that reduce data movement, latency and power consumption.

Additional presentations will cover Tyndall National Institute's magnetics-on-silicon platform, including micro-transfer-printed inductors designed for integration into interposers and future buried power delivery networks.

According to the organisers, the workshop will provide attendees with an in-depth look at FAMES' latest technical achievements while offering opportunities to engage with researchers developing Europe's next generation of semiconductor technologies.



20
27

Sheraton Brussels
Airport Hotel

Brussels, Belgium

12-14 April 2027



REGISTER NOW

Advanced Packaging International Conference 2027

The 2nd Advanced Packaging International builds on the huge success of the 2026 launch event and the momentum driving the next era of semiconductor innovation. Over two days, industry leaders and academic experts will deliver around 40 presentations across key themes:

Conference attendees will gain insights into the materials, methods and technologies reshaping chip-level packaging and system integration. From cutting-edge thermal solutions to the reimagining of wirebond and leadframe formats, the event provides a comprehensive look at how packaging is enabling system-level innovation.

The two-day conference chaired by Peter O'Brien also includes expert panels, tabletop exhibition and two networking evenings creating multiple avenues for peer-to-peer engagement and technical exchange across the packaging ecosystem, from researchers and R&D teams to product engineers and system architects.

ap-international.net



SCAN ME

80+ EXHIBITORS

800+ ATTENDEES

4 CO-LOCATED CONFERENCES

120+ PRESENTATIONS





20
27

The European Event for The Semiconductor Supply Chain

REGISTER NOW

Conference Themes

- Chiplets and Heterogeneous Integration
- Scaling and Pilot Lines
- Precision Equipment for Scalable Advanced Packaging
- Materials for Packaging Performance
- Thermal Management and Reliability
- Design and Simulation


Why Attend?


- Unbeatable location: The Sheraton Brussels Airport Hotel is 20 yards from the airport terminal — fly in, walk to the conference.
- Networking: Two drinks receptions and a conference dinner provide exceptional networking opportunities.
- Full supply chain: Four co-located conferences bring the entire compound semiconductor, PICs, power electronics, and advanced packaging supply chain together under one roof.
- 800+ attendees: Join a growing community of decision-makers, engineers, and researchers.

THE WHOLE SEMICONDUCTOR SUPPLY CHAIN UNDER ONE ROOF



PREVIOUS SPONSORS INCLUDE





Shaping the next decade of advanced packaging innovation

As advanced packaging evolves to meet the demands of AI and high-performance computing, manufacturers face growing challenges around yield, precision, throughput and process stability. Mycronic outlines how its portfolio of testing, surface preparation, die bonding and photonic integration technologies is helping bridge the gap between emerging packaging architectures and high-volume production.

BY LAURA HORAN, GILLES LAMBINET, OVE SÖRENSEN, IRVING WANG, PER-ERIK GUSTAFSSON, MYCRONIC

The demand for advanced packaging, driven by AI and computing megatrends, is pushing manufacturing into a new phase of complexity and scale. This demand can no longer be satisfied through silicon design alone, and the industry is responding with novel manufacturing technologies such as chiplets, heterogeneous integration, 2.5D and 3D packaging, as well as the integration of silicon photonics and co-packaged optics (CPO).

Advanced packaging architectures are critical to enable the enhanced levels of performance required by the megatrends, but they also introduce significant manufacturing challenges. The pace of innovation has also accelerated significantly. Development cycles are shortening and new technology generations are transitioning into production faster than before. This leaves

manufacturers less time to develop, validate and industrialize new processes, materials and integration strategies, increasing the need for flexibility and rapid innovation validation.

Volume production, yield optimization, testing, and the integration of multiple dies rapidly drive-up manufacturing complexity. System performance is increasingly determined by how diverse technologies are integrated within a single package, and each technology may require vastly different manufacturing processes. Success now depends on achieving technical breakthroughs while also delivering high process stability, yield, throughput, and quality as production scales.

Mycronic has long recognised these challenges and has assembled a value chain of companies

dedicated to solving the manufacturing challenges of the next decade. Mycronic's strategic acquisitions in the photonics and advanced packaging space offer customers a broad value proposition and technologies across multiple critical stages of the manufacturing flow, from testing and surface preparation to advanced assembly and photonic integration.

Hprobe ensures that only known-good magnetic devices enter advanced packaging flows

One of the key requirements for advanced packaging is to ensure that only fully functional devices enter the packaging flow. As device complexity and packaging costs increase, identifying performance issues early in the manufacturing process becomes essential. For magnetic devices such as Magnetoresistive Random-Access Memory (MRAM) and magnetic sensors (e.g. Hall, AMR, GMR, TMR), this requires testing both at wafer level and after packaging to verify device functionality, process stability and final product performance.

Unlike conventional semiconductor devices, MRAM and magnetic sensors must be characterized under precisely controlled magnetic field conditions. Their performance depends not only on electrical parameters but also on their response to magnetic field amplitude and orientation, making dedicated magnetic test equipment essential to manufacturing.

Hprobe, a Mycronic company, develops automated magnetic test equipment dedicated to magnetic semiconductor devices, see Figure 1. Its patented three-dimensional magnetic field technology enables vectorial magnetic field generation above 200 mT in any direction, reaching over 500 mT in-plane and more than 250 mT out-of-plane across a large, highly uniform test area. Combined with magnetic field sweep rates up to 2 T/s, precise field calibration, and seamless integration with




industry-standard wafer probers and automated test equipment (ATE), the platform supports both development and high-volume manufacturing on 200 mm and 300 mm wafers

Wafer-level magnetic testing enables manufacturers to characterize devices early in production to identify process variations, monitor production quality, and screen known-good devices before costly packaging operations take place. Early magnetic testing reduces manufacturing risk while accelerating process development and yield optimization.


In advanced packaging, the packaged devices must also be validated under controlled magnetic field conditions before deployment. Advanced sensing devices benefit from adjustment of the magnetic sensitivity after packaging using a magnetic tester. Hprobe's packaged-device test solution combines precise three-dimensional magnetic field control with full 360° capability, automated calibration, and field mapping at the device-under-test (DUT) position to ensure accurate, repeatable measurements.

The system operates from -40°C to +150°C using CDA cooling and can be used with external ATE. This integration enables parallel testing of multiple devices under controlled magnetic fields. The

➤ Figure 1:
a) Inside the Hprobe Test Head: Magnetic Field Generation and Wafer-Level Probing
b) Hprobe ATE integrated with TEL Precio XL wafer prober



➤ Figure 2: a) More-than-Moore 3D Stacked Chip Design using Surfx plasma as a key enabling technology b) Surfx plasma head and the plasma glow in purple c) Surfx turn-key semiconductor equipment with Atomflo™ plasma technology



► Figure 3. Overnight stability test of a one-micron die bonder

complete test process is software driven, allowing users to define 3D magnetic field patterns, align the device and execute semi-automatic test sequences from a single programming environment.

Hprobe helps semiconductor manufacturers ensure consistent magnetic device performance throughout production, complementing Mycronic's broader portfolio of advanced packaging and manufacturing solutions.

Surfx tackles the yield challenge by maintaining surface quality

Every advance in AI computing requires one simple yet non-negotiable requirement, the chips must work. As interconnect pitches shrink and bond lines thin, the margin for error keeps closing. Often the condition of the chip surface determines the difference between a high-yield or struggling production line.

At Surfx, a Mycronic company, we focus on this challenge. Nearly every bonding and coating step in advanced packaging depends on starting from a clean, chemically ready surface. Traditional approaches based on wet chemistry or vacuum plasma add cycle time, queue time, and additional production line complexity. With AI demand compressing manufacturing timelines, the industry requires another solution.

Our answer is the Atomflo™ atmospheric plasma platform, shown in Figure 2. Unlike vacuum based systems, Atomflo™ operates without a vacuum chamber, without ion bombardment, and without disrupting the flow of production. It activates surfaces in seconds, right at the point of use, inside the tool rather than beside it. The technology generates a dramatically higher concentration of reactive radicals while eliminating ion bombardment and maintaining an essentially particle-free process. For 300 mm wafers and emerging larger panel formats, substrates can be treated in-line, at production speed, eliminating the cycle time bottleneck of batch vacuum systems.

As semiconductor manufacturers move toward hybrid bonding, flux-less thermocompression

bonding (TCB), and other advanced interconnect technologies, process sensitivity has increased significantly. These applications demand surface preparation methods that are both effective and gentle. Unlike many conventional atmospheric plasma systems, Atomflo™ produces no electrostatic discharge (ESD) and avoids energetic ion bombardment that can damage sensitive structures. This combined with high surface activation, makes it particularly well suited for advanced logic and memory devices powering the AI revolution.

Surfx plasma technology has become a key enabler for next-generation packaging architectures aligned with the More-than-Moore roadmap, supporting advanced 3D integration strategies that push performance beyond traditional semiconductor scaling. Today, Surfx proudly serves leading global logic and memory manufacturers, including companies driving AI, with solutions tailored to the fastest-growing advanced packaging applications.

It's a small, precise intervention that lets manufacturers hold yield steady while chip architecture evolves toward smaller pitches. Surfx's contribution to advanced packaging is the proof that sometimes the biggest unlock is controlling the surface nobody sees.

Integrating optoelectronics using die bonding for packaging at scale with MRSI

The rapid evolution of AI architectures, driven by growing bandwidth demands, is accelerating the adoption of co-packaged optics (CPO) and near-packaged optics (NPO). The industry is increasingly assembling systems where multiple optical and electronic components are combined within a common package. This trend is driving new requirements for assembly precision, throughput, and manufacturing stability.

For more than 40 years, MRSI, a Mycronic company, has provided precision die bonding equipment for optoelectronic and microelectronic packaging. Its systems are built to support well-established processes, including eutectic and epoxy die attach on singulated substrates and wafers, delivering

placement accuracy as tight as one micron. This combination of accuracy, process flexibility, and production capability positions MRSI to support both conventional pluggable transceivers and emerging, highly integrated photonic assemblies.

Artificial intelligence is rapidly changing the scale and architecture of data center networking. AI training and inference require large numbers of accelerators to exchange data continuously across scale-up and scale-out networks, making high-bandwidth, high-reliability optical interconnects critical to overall computing performance.

NVIDIA describes optical interconnects as the critical backbone for scaling AI clusters and offers products for infrastructure operating at speeds of up to 1.6T.

In novel CPO and NPO architectures, photonic ICs, electrical ICs, interposers, and micro-optical components are packaged into highly integrated optical engines using advanced processes, including wafer- and panel-level packaging. MRSI addresses these requirements through a portfolio of high-accuracy, flexible die bonding solutions. The die bonding systems have placement accuracy down to 0.5 micron, with proven long term process stability required for 24/7 high-volume manufacturing (see Figure 3). MRSI enables manufacturers to scale production to new generations of products. As optical connectivity moves closer to the compute engine, precision die bonding is becoming an essential manufacturing capability for the next generation of AI infrastructure.

Vanguard Automation unlocks scalable photonic integration using 3D nano-printing

One of the critical challenges to advanced packaging is photonic integration, a key process for manufacturing electro-optical engines that underpin transceivers, CPO, light engines, and external light sources (ELSFP). Integrating active components such as lasers with passive photonic integrated circuits (PICs) typically requires highly precise, time

One of the critical challenges to advanced packaging is photonic integration, a key process for manufacturing electro-optical engines that underpin transceivers, CPO, light engines, and external light sources (ELSFP).


consuming, active alignment processes that can be difficult to scale across multiple optical channels without compromising loss budgets or product reliability.

Vanguard Automation, a Mycronic company, has developed 3D nano-printing technology that enables scalable photonic integration while utilizing passive alignment processes. The solution uses highly precise direct-write 3D laser lithography to fabricate 3D freeform optical interconnects directly onto photonic devices at chip and wafer level. The optical interconnects compensate for process misalignments, while maintaining high optical performance, allowing manufacturers to integrate active components with PICs at low cost.

“Photonic integration remains one of the most significant barriers to high-volume production. By replacing complex active alignment processes with passive alignment solutions enabled by 3D nano-printing, we are helping customers achieve yield and manufacturing throughput required for next-generation optical systems.” Dr. Laura Horan, Head of Product Management, Vanguard Automation.

Vanguard offers a complete portfolio of optical interconnects, including Photonic Wire Bonds (PWBs) and facet-attached micro-lenses (FaMLs), to support scalable photonic integration across a wide range of packaging scenarios.

➤ Figure 4:
a) vanguard SYMPHONY fully automated photonic integration and packaging platform. b) SEM image of a single-mode fibre (SMF) array integrated to an AMF chip using Photonic Wire Bonds (PWBs).



PWBs are freeform optical waveguides fabricated between photonic components to compensate placement misalignments of up to $\pm 20 \mu\text{m}$ while maintaining efficient coupling by modefield matching. PWBs are increasingly recognised as the preferred solution for hybrid integration of III-V lasers with PIC platforms. In recent demonstrations, PWBs achieved insertion losses lower than 1 dB per interface and supported stable laser operation.

FaMLs perform spot-size conversion and beam shaping, relaxing mechanical alignment tolerances to $\pm 15 \mu\text{m}$, and reducing sensitivity to lateral and angular misalignment. This eliminates the need for bulk optics, reducing system footprint and packaging complexity.

The fully automated vanguard SYMPHONY platform (SONATA1000 and REPRISE1000), as seen in Figure 4, combines machine-vision systems, post-processing, industrial-grade materials (VanCore series) and software (BrightWire3D) to support reliable photonic integration. The result is a software-defined manufacturing process that transitions seamlessly from prototyping to volume production without product-specific tooling. The resulting optical interconnects are compatible with standard manufacturing processes and withstand solder reflow temperatures, mechanical shock, vibration, and extended environmental stress testing, demonstrating the reliability for manufacturing deployment.

Mycronic group shaping the advanced packaging value chain

The advanced packaging industry is currently discussed as individual technologies, yet the real challenge is integrating all of them into a stable, repeatable, high-yield manufacturing chain. The manufacturing sequence for a typical advanced package clearly illustrates this challenge, from testing to surface preparation, electronic integration and photonic packaging, and failure at any stage can significantly impact manufacturing yield.

Mycronic's solutions portfolio addresses multiple stages of the advanced packaging flow through complementary capabilities. These capabilities complement Mycronic's broader portfolio in semiconductor and electronics manufacturing, including photomask equipment, PCB and substrate assembly and test, and high volume assembly and dispensing technologies. Customers increasingly seek partners capable of supporting them throughout the complete product and process lifecycle.

"We have the ability to work closely with customers in early

The advanced packaging industry is currently discussed as individual technologies, yet the real challenge is integrating all of them into a stable, repeatable, high-yield manufacturing chain.

development, while also having the capability to deliver advanced production systems at scale, with high delivery precision and stable yield in the field," says Per-Erik Gustafsson, Head of Group Business Control & Corporate Planning, Mycronic. "That combination is becoming increasingly critical, as customers need speed, flexibility, yield and reliability at the same time."

This requires tight collaboration, strong process understanding, and the ability to transition from development to stable production with short lead times.

"At the same time, flexibility remains critical. Customers need to adapt quickly to new architectures, materials and processes without compromising production. The ability to bridge emerging technologies into stable, high-volume manufacturing will become increasingly important, and Mycronic is well positioned to support that transition."

Ultimately, long-term success in advanced packaging will not be determined by technology trends alone. Successful execution on the factory floor, where precision, stability, yield, and throughput define what can be manufactured at scale, will be equally important. Mycronic is looking to the future to remain ahead of the latest packaging needs and has positioned itself as a company at the convergence of the most critical manufacturing technologies enabling advanced packaging for the future megatrends.



➤ Authors from left to right: Laura Horan, Gilles Lambinet, Irving Wang and Per-Erik Gustafsson

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)



From lab to volume: What advanced packaging actually requires



Heterogeneous integration, hybrid bonding and photonic integration are carrying Moore's Law beyond 2D wafer-level manufacturing. Everything now depends on the package, and on whether a working device can be built reliably at volume.

**BY TOBIAS GLEICHMANN, SENIOR COMMUNICATIONS MANAGER,
TRAVIS SCOTT, SENIOR PRODUCT MANAGER AND SENIOR BUSINESS
DEVELOPMENT MANAGER, FINETECH GMBH & CO. KG**

For most of the semiconductor era, packaging played a supporting role. Performance came from the transistor; the package protected the die and presented it in a form the next assembly step could handle. That order has been reversed. Moore's Law is slowing, cost per node is rising, and performance now comes from integration rather than from shrinking a single chip. Chiplets, heterogeneous integration, 2.5D and 3D stacking and co-packaged optics are already in volume production. Within a few years the parts that matter most, AI accelerators, high-bandwidth memory, optical engines, will not be buildable without advanced packaging, and companies still treating it as a back-end afterthought will be locked out of the products that define the market.

The bottleneck has moved from design to assembly: from making a single device work to making many devices and technologies work together at the same yield and cost as the wafer scaling that precedes it. Three factors decide whether a process survives the move to volume: accuracy, process flexibility and repeatability. Miss any one and the process does not scale, however well it works in the lab.


Why the package is now the bottleneck

Once a system is built from many dies joined by ever denser interconnects, from wire bonds and bumps to 2.5D, 3D stacking, hybrid bonding and optical interfaces, the connections carry as much of the performance as the dies themselves. The bandwidth roadmaps climbing from 800G to 1.6T and 3.2T are really roadmaps for the package, and

factory-floor tolerances have tightened to match. The same shift is dissolving the old split between front end, back end and test.

Silicon photonics shows why. A laser diode can be aligned to a photonic integrated circuit with sub-micron accuracy in the lab, but volume production behaves differently: small deviations, thermal effects and drift accumulate across many units, making yield less predictable. What matters shifts from whether the device can be built to whether the process holds optical coupling, alignment stability and repeatability at production rate. Static placement is not enough. Active alignment, measured live in six or twelve axes while the component is positioned, holds the coupling that a fixed placement would lose to thermal and process variation, on every unit rather than only in the lab.

Scaling is not only about going smaller. Wafer- and panel-level packaging push toward larger formats, 300 mm wafers and glass-core panels up to 600 mm square, while components range from over 100 mm down to single VCSELs or photodiodes of a few hundred microns. Larger substrates lower cost per unit only if the process holds across the whole area. Maintaining sub-micron bonding and uniform control across a 300 mm wafer or a 600 mm panel, without stress building up, is a different proposition from a small coupon in the lab, and most processes that look ready in the lab fail their first serious test in production.



➤ Figure 1: The packaging gap: as transistor scaling slows, advanced packaging carries the growth in interconnects, with wire bonding giving way to flip chip and on to hybrid bonding and photonic integration. Credit: Finetech GmbH & Co. KG.

Highly versatile architectures will require multiple bonding technologies

Scaling forces many purpose-built technologies together within a single device, a wide range of materials, bonding methods and processes that must all function together at scale.

Established methods such as eutectic and thermocompression bonding are joined by others moving into volume on 300 mm wafers and panels: plasma-assisted hybrid bonding, ultrasonic and laser-assisted bonding, UV-curing adhesives, each chosen for its thermal budget, throughput and interconnect performance.


Combining several in one device is where it gets difficult, each with its own tooling, thermal control


and accuracy, and each switch a chance for variation.

A dedicated machine is often faster for one fixed task, so consolidation is not always the answer. It pays off when the process chain matters more than any single step: when each transfer risks contamination or damage, when frequent tool changes across a high-mix line cost more than they save, or when a fragile, high-value assembly cannot afford repeated load and unload. Then one automated flow with interchangeable heads keeps the line stable.

The same system also has to span extremes of scale, from large multi-die assemblies down to single micron-scale emitters, without a change of

➤ Figure 2: A ladder of rising interconnect density: from solder ball and wire bond, through bump and pillar, to interposer-based 2.5D, 3D stacking and hybrid bonding, and on to optical interfaces, each step raising bandwidth and tightening tolerances. Credit: Finetech GmbH & Co. KG.





➤ Figure 3: No single bond fits all. Advanced integration draws on a portfolio of bonding technologies, here: ultrasonic friction welding, plasma-prepared hybrid bonding, laser-assisted bonding and UV-curing adhesives, each suited to different materials and tolerances. Credit: Finetech GmbH & Co. KG.

platform, since splitting the process across separate tools reintroduces the transfers it is meant to avoid.

Hybrid bonding pushes density further, and is less forgiving

Hybrid bonding changes the interface itself. It forms direct die-to-die or die-to-wafer connections at extremely fine pitch, with no bump in between, and makes the densest 2.5D and 3D architectures possible: better signal integrity, lower power, higher integration in a smaller footprint. It is also the least forgiving, demanding high flatness, co-planarity and very clean, activated surfaces. At these pitches the vertical axis matters as much as the lateral one, and reproducible Z-control and bond-line thickness monitoring decide whether a stacked structure holds across the whole wafer, not just at a test site.

Surface preparation here is part of the bond. And the effort going into scaling hybrid bonding to mass

manufacturing will not stay confined to today's devices: photonic integration depends on the same fine-pitch placement, surface cleanliness and process control, and will inherit much of what hybrid bonding is working out at volume.


Bond quality depends on the surface, not just placement

Placement accuracy gets the attention, but in many processes the limiting factor is the bonding interface. Contamination or native oxides can weaken a bond long before any test reveals it, which is why surface preparation belongs inside the bonding sequence: an in-line plasma activation stage immediately before the bond, with no window for the surface to degrade before the parts meet. As assembly scales, material testing that keeps failed parts out of the flow will matter as much to yield as placement accuracy.

MicroLED displays show how quickly packaging becomes a yield problem. Tiny, fragile emitters leave little margin for contamination, misalignment or excessive force, so a stable process depends on placement accuracy, surface preparation, low-force bonding and repeatable handling alike. What looks like a placement task is really a combined challenge of interface quality, force control and handling at scale.

Materials add constraints. Thin silicon and compound-semiconductor wafers warp easily, optical surfaces are unforgiving of contamination and misalignment and combining them in one architecture brings coefficient-of-thermal-expansion (CTE) mismatch that device size, forces and temperatures only narrow further. Cryogenic and high-reliability work adds another layer, with interconnects such as indium or gold bumps needing exact force and temperature control.

➤ Figure 4: A bond without solder: cross-section of a thermo-compression joint, where heat and force form a direct substance-to-substance connection between chip and substrate. Courtesy of AEMtec GmbH.



Processes fail in the transfer to production, not in the lab

What catches teams out emerges in production, not development. A process can perform perfectly in the lab and misbehave the moment it runs at volume. Often the cause was there from the start: a bonding strategy not validated with production in mind carries cost and risk forward from the first prototype, resurfacing as yield problems, qualification delays and missed commitments.

The silicon photonics case is the same story at scale: a process stable on a research system drifts once it moves to a different tool chain, software or production rhythm. This process drift costs time and money. Failure analysis carried out earlier, on R&D equipment that can reproduce production conditions, catches weaknesses that would otherwise appear only after ramp, pulling the cost of failure forward to the cheapest point in the product's life.

The cause usually lies in the environment, not the process. When development and manufacturing use different systems, software and process logic, every transfer introduces variation. Production seldom fails inside a single step; it fails in the handoffs between steps, tools and teams, and those handoffs multiply as integration deepens. The complaint heard on many lines captures it: we can make it work, but it takes too many steps.

This cost is usually undercounted. A second development loop consumes engineering time, qualification material and expensive wafers, and every re-qualification means material reworked or written off, cost and time spent twice to reach a result the line should have reached once.

Continuity from prototype to volume is the fix

Companies that scale well treat research and manufacturing as one continuous effort, aligning the two early, choosing development tools and process logic that already reflect production, and working with partners who scale with them from R&D into production rather than handing the process over at the factory door. The requirements are consistent: repeatability, sub-micron placement, the right alignment strategy, automation and an unbroken route from prototype to production.

Process continuity matters most: the same hardware principles, bonding technologies, software and recipes throughout. Then a validated process transfers without a second round of development, learning cycles shorten and yield settles earlier. This favours platforms that bring several bonding processes together rather than spreading them across unrelated tools, for lower cost, less scrap and fewer chances for drift, not only speed.

Continuity also depends on what the line can see. In-line inspection, bond-line measurement and closed-loop correction catch drift as it happens,

while MES integration, traceability and data logging make a qualified process auditable and repeatable at volume. Systems that stay easy to operate and expose their process data keep development transparent, which is what lets a new process be understood, trusted and qualified quickly.


Production tooling is moving this way, and the systems are starting to appear. It matters less that one machine does everything than that several machines built on the same platform, sharing the same software and process logic, form one ecosystem. Within such a family a process transfers from one system to the next without being re-learned; across unrelated machine generations that continuity has to be engineered, not assumed. Where the platform is shared, a manufacturer can cover even complex packaging end to end, for reliability rather than reach: fewer transfers, fewer changeovers, fewer chances for drift.


Plan packaging from the first design, not after it

Heterogeneous integration is changing how systems are designed, and with it how companies compete. Devices are becoming multi-die, multi-material and increasingly photonic, integrated at 300 mm wafer and soon 600 mm panel scale. Precision and scalability are no longer back-end concerns; they mark the difference between a company that reaches volume and one that stalls at the prototype.

The lesson is to plan the link between design and packaging from the start. Even with careful design, performant architectures and modern materials create challenges that only surface in assembly, so wafer layout, interconnect design and assembly architecture are best defined with bonding requirements in view. Early contact between design and packaging engineers prevents costly re-spins and yield losses later. Teams that leave packaging until the device is fixed give themselves less room, and often see the cost two years later in yield and missed ramps. For advanced packaging, scalability is a design constraint from the first prototype. Packaging is becoming the first decision in the flow.

➤ Figure 5: Wafer-level integration: VCSELs, photodiode arrays and NTC thermistors assembled directly on wafer. Courtesy of PHIX B.V.





Optical memory pooling still has a physical realisation problem



As AI systems evolve, balancing HBM cache, DDR5 capacity and optical I/O will make module granularity critical to building memory-pooling architectures that can be efficiently cooled, validated and scaled.

BY DR. MOH KOLBEHDARI, SR. DIRECTOR OF IC/PACKAGING, SOCIONEXT US

Accelerators need high-bandwidth memory close to compute, but long-context and multi-user workloads can generate key-value caches that exceed practical local HBM capacity. Host DRAM offers greater capacity at lower cost, while storage provides even more capacity but introduces too much latency for frequently accessed inference data [1-3].

Optical memory pooling offers an attractive response: separate memory capacity from the accelerator, connect it through a low-latency optical fabric, and allow multiple servers or accelerators to share a larger memory resource.

Recent public architectures combine three complementary technologies [1-2]:

- DDR5 for economical, terabyte-scale capacity
- HBM3E for a high-bandwidth cache or fast memory tier
- Optical I/O for rack-scale reach and memory sharing

This is a compelling system concept. However, solving the rack-level connectivity problem does not eliminate the physical-realization challenges inside the memory module.

The central engineering question is: **At what point does a disaggregated-memory architecture become too highly integrated within each package-plus-board module?**

This question extends beyond module partitioning. The physical realization must also close DDR5 channel topology and achievable data rate, HBM cache behavior, package and PCB routing, power delivery, thermal distribution, optical stability, manufacturing variation, test, reliability, and serviceability. A more distributed implementation may improve some of these boundaries, for example by shortening DDR channels, reducing the number of PHYs and routes concentrated around one ASIC, distributing heat, and creating smaller units for validation and manufacturing

learning. However, these benefits must be weighed against duplicated optical interfaces, control logic, packaging, power management, and system-level coordination. The objective is therefore not maximum distribution, but identifying the physical granularity that provides the best combined bandwidth, thermal, electrical, manufacturing, reliability, and cost result.

HBM for speed, DDR5 for capacity

The purpose of combining HBM and DDR5 is not to create two equivalent memory pools.

DDR5 supplies most of the capacity. HBM acts as a high-bandwidth cache for active regions of the DDR address space, although the published architecture also allows HBM to be configured as additional pooled memory [1].

A recent publicly available Marvell technical paper describes a Photonic Fabric Memory Module containing

two 36 GB HBM3E stacks, eight DDR5 interfaces supporting as much as 2 TB of DDR5 capacity, 7.2 Tbps of optical connectivity, a photonic integrated-circuit interposer, a fiber-array unit, and a memory-fabric ASIC. Sixteen such modules form a 32 TB shared-memory appliance [1].

The hierarchy is straightforward:

- HBM provides bandwidth
- DDR5 provides capacity
- Optics provides distance and sharing

The HBM cache is especially important because DDR5 alone may not sustain the peak optical-link bandwidth under every access pattern. The published evaluation provisions the HBM tier with bandwidth above the optical-link requirement so that the optical fabric can remain utilized when the DDR subsystem would otherwise become the limiting factor [1-2].

Caching, however, does not remove the DDR5 implementation burden. System performance still depends on cache locality, address interleaving, replacement policy, write behavior, and the proportion of transactions that ultimately reach the DIMMs.

The DDR5 topology creates a difficult choice

A public block diagram may show four DIMMs on either side of the central package, but physical placement does not reveal electrical channel topology [1].

There are two broad possibilities.

One DIMM per channel

If eight DIMMs use eight independent DDR5 interfaces, the architecture avoids the severe loading associated with placing several DIMMs on one channel.

However, the ASIC must support eight high-speed DDR interfaces. Even when one common controller complex manages them, the implementation still requires multiple channel data paths, PHYs, training engines, clocks, command/address resources, and package connections.

That creates substantial costs in:

- ASIC die-edge area
- Controller and PHY power
- Package pin count
- Package escape routing
- Board layer count
- Timing calibration

- Test coverage
- Power-delivery design

DDR5 modules also contain local power-management circuitry and additional active components. That reduces some motherboard rail-management burden, but it adds module-level power, monitoring, and thermal considerations [6].

The public technical description identifies eight DDR5 interfaces, which strongly suggests extensive parallel memory connectivity. It does not, however, disclose the supported DDR transfer rate, exact channel architecture, DIMM rank configuration, controller power, or full-population electrical margins [1].

Multiple DIMMs per channel

Using fewer controllers and PHYs would reduce ASIC area and I/O concentration, but connecting two or especially four DIMMs to one channel introduces a different problem.

Additional DIMMs increase capacitive loading, command/address loading, discontinuities, stubs, reflections, training complexity, and trace-length variation. These effects generally reduce the maximum practical operating rate.

The trade-off is visible in commercial server platforms. Intel documentation, for example, identifies lower maximum DDR5 rates for two-DIMM-per-channel configurations than for equivalent one-DIMM-per-channel configurations on several Xeon generations [4-5].

Four DIMMs per channel would be considerably more challenging at high DDR5 rates. It should not be assumed merely because four physical DIMMs appear on one side of a drawing.

The implementation therefore faces a fundamental trade-off:

- Fewer channels reduce ASIC complexity but increase channel loading.
- More channels preserve bandwidth but increase ASIC, package, power, and routing complexity.

The PCB remains part of the memory path


Optics may carry the transaction across racks, but the final movement from the memory ASIC to DDR5 still occurs electrically across the module PCB.

That path includes:

- ASIC package escape
- Vias and layer transitions
- Long board traces
- DIMM connectors
- Command/address and clock distribution
- Data and strobe matching
- Reference-plane transitions
- Return-path control

The propagation delay of the PCB trace is only one portion of total access latency. A cache-miss transaction may pass through:

Host CXL interface → optical NIC
→ optical fabric → photonic receiver
→ memory/fabric ASIC → HBM cache
lookup → DDR controller and PHY → DIMM → return path.



The published work reports low-latency behavior based on hardware emulation and system simulation. The authors also state that end-to-end validation using physical appliance hardware and inference workloads remains pending [1].

That distinction matters because simulation can validate architecture and protocol behavior, while physical hardware must also close:

- DDR signal integrity
- Simultaneous-switching noise
- Package and PCB power integrity
- DIMM temperature
- Airflow
- Optical alignment
- Photonic temperature stability
- Manufacturing variation
- Long-duration traffic behavior

Thermal management is distributed, not eliminated

A module containing a central ASIC, two HBM stacks, integrated photonics, eight DDR5 DIMMs, and multiple local power-management circuits creates several thermal regions [1-6].

The central 2.5D package requires concentrated heat removal. The surrounding DIMMs need airflow or conduction paths that are not obstructed by the central heat sink. DIMM power-management devices and DRAM components can generate localized heating, while elevated DRAM temperature may increase refresh activity or force performance throttling.

The optical region adds another constraint. Modulators, detectors, fiber coupling, and laser illumination must remain stable while adjacent digital and memory devices dissipate heat.

HBM reduces repeated DDR traffic when cache locality is favorable, but it does not remove DIMM thermal requirements. Under cache misses, streaming workloads, cache replacement, or write-through behavior, the backing DDR remains active.

The complete package-plus-board thermal field therefore matters more than any individual component temperature.

Has the module become a board-scale monolith?

The industry moved toward chiplets because putting every function

into one die eventually created unacceptable pressure on area, yield, power, development cost, and design complexity.

The same principle can apply one level higher.

A system may be disaggregated across racks yet remain highly concentrated inside each memory module. Integrating eight DDR interfaces, two HBM stacks, high-bandwidth optical I/O, switching logic, cache management, and a large board-level memory subsystem into one realization object can reproduce many of the same challenges that chiplets were intended to relieve.

The issue is not whether the architecture is innovative. It clearly is.

The issue is whether the chosen module granularity is the most practical one.

Concentrated versus distributed implementation

Consider a conceptual alternative in which the capacity and bandwidth of one concentrated module are partitioned across multiple smaller modules connected in parallel through the optical fabric. This is not a prescription for a fixed module count. The number of modules, DDR5 interfaces, DIMMs, HBM or cache capacity, and optical bandwidth per module are design variables.

Each smaller module could contain fewer DDR5 interfaces and DIMMs, a smaller memory/fabric ASIC, proportionally allocated high-bandwidth cache capacity, fewer package and board routes, and a lower optical-bandwidth allocation.

The distributed modules could collectively provide the same target capacity and aggregate bandwidth.

This finer-grained architecture could offer several advantages:

- Shorter DDR routes
- Fewer DDR PHYs per ASIC
- Smaller package and die size
- Distributed power and heat
- Simpler board stack-up
- Easier airflow
- Better fault isolation
- Incremental capacity expansion
- More manageable design verification

Shorter and less heavily loaded DDR channels may also improve electrical margin and create an opportunity to sustain higher practical DDR5 data rates. This is not automatic; the achievable rate must be demonstrated through channel analysis, training behavior, full-population testing, and hardware measurement.

A failure or thermal limit in one unit would affect only part of the memory pool. Manufacturing learning could also proceed on a smaller realization object rather than on the largest possible module.

But modularity has costs.

- A distributed implementation may duplicate
 - ASIC control logic
 - Optical interfaces
 - Fiber-array attachments
 - Package substrates
 - Power-management circuitry
 - Clocking
 - Test infrastructure
 - Cache-management resources

The software and fabric must also stripe memory efficiently across modules, preserve ordering and consistency, balance traffic, and prevent one slower module from creating tail latency.

The right answer may therefore be neither maximum integration nor maximum partitioning.

It may be the smallest module that provides acceptable economics while keeping electrical, thermal, optical, manufacturing, and validation complexity within a manageable boundary.

What physical validation must show

Before any optical pooled-memory architecture can be judged at production scale, several implementation factors must be addressed, including the sustained DDR5 rate with every DIMM populated and whether each DIMM has a dedicated interface.

The associated controller and PHY area and power costs must also be considered, alongside how much traffic is served from HBM, how workload-dependent the cache hit rate is, and what bandwidth can be sustained during HBM misses.

Thermal management is another key challenge, particularly how the ASIC, HBM, photonics, DIMMs and DIMM PMICs can be cooled simultaneously, and how temperature variations affect DDR margins, refresh behaviour and optical stability. Finally, the architecture must account for module or optical-path failures and assess whether a distributed implementation could deliver better total cost, manufacturability, reliability and serviceability than a single concentrated module.

Optical memory pooling may become an important new AI-infrastructure tier. But optical reach alone does not prove physical scalability.


The deeper lesson is familiar:

- Disaggregation should not stop at the rack boundary.
- The package and module must also be partitioned at a scale that can be routed, powered, cooled, manufactured, tested, and repaired.

A system architecture becomes real only when its chosen physical granularity is practical.


REFERENCES

- [1] J. Ding, Y. Nishant, C. Ambati, J. Kamati, and T. Diep, "A Photonic-CXL Memory Appliance for Scalable KV Cache Management in LLM Inference," arXiv:2607.27187, July 2026. <https://arxiv.org/abs/2607.27187>
- [2] J. Ding and T. Diep, "Photonic Fabric Platform for AI Accelerators," arXiv:2507.14000, July 2025. <https://arxiv.org/abs/2507.14000>
- [3] Marvell, "HCDS Best Paper Award Winner: Bridging the Memory Hierarchy Gap with Photonic Fabric™ Technology," April 16, 2026. <https://www.marvell.com/blogs/hcbs-best-paper-award-winner.html>
- [4] Intel, "DDR1DPC: Maximum Intel DDR5 1 DPC DIMM Speed," Eagle Stream Platform Electrical Data Sheet, Document ID 814095. <https://edc.intel.com/content/www/us/en/design/products-and-solutions/processors-and-chipset/s/eagle-stream/platform-electrical-data-sheet/ddr1dpc-maximum-intel-ddr5-1-dpc-dimm-speed/>
- [5] Intel, "DDR2DPC: Maximum Intel DDR5 2 DPC DIMM Speed," Eagle Stream Platform Electrical Data Sheet, Document ID 814095. <https://edc.intel.com/content/www/us/en/design/products-and-solutions/processors-and-chipset/s/eagle-stream/platform-electrical-data-sheet/ddr2dpc-maximum-intel-ddr5-2-dpc-dimm-speed/>
- [6] Micron Technology, "Micron's DDR5 Technology Enablement Program Empowers an Ecosystem," describing on-module PMIC/local voltage regulation and other DDR5 module features. <https://www.micron.com/about/blog/memory/dram/microns-ddr5-technology-enablement-program-empowers-ecosystem>



Note: This figure illustrates architectural and physical tradeoffs. The optimal point in the design space must be determined through electrical, thermal, optical, manufacturing, cost, and system-level analysis.

➤ Figure 1: Concentrated module versus distributed implementation The same aggregate memory capacity can be realized with different physical granularities. The optimum depends on ASIC cost, DDR signaling, cache behavior, thermal density, optical duplication, yield, and serviceability.



Empowering violet lasers

More powerful violet lasers, realised by improving design and introducing superior thermal management, will target laser direct imaging for advanced semiconductor packaging

**BY SHINJI YOSHIDA AND COLLEAGUES FROM
NUVOTON TECHNOLOGY CORPORATION JAPAN**

THROUGHOUT the history of the semiconductor industry, there's been a relentless drive to improve semiconductor performance. As well as reshaping device architectures, this eternal goal refines the way that chips are packaged and manufactured. Within the silicon industry, conventional transistor scaling is approaching fundamental limits. Due to this, innovation is increasingly shifting towards advanced packaging, where processors, memory devices and specialised chips are assembled three-dimensionally and interconnected by extremely fine patterning. This transition is placing an unprecedented demand on manufacturing equipment, and in particular on the lithographic technologies employed to form redistribution layers and high-density interconnects on large-area substrates.

Packaging at the forefront

There are a number of requirements associated with advanced semiconductor packages. They include finer features and resolutions, a higher overlay accuracy between multiple layers, large substrates and high productivity.

Efforts to excel on these fronts are stretching the capabilities of traditional steppers using photomasks. While the photolithography equipment that relies on ultraviolet radiation from mercury lamps and solid photomasks has served the silicon industry well, there are intrinsic challenges when substrate sizes grow and product variants multiply. Mask preparation is now costly and time consuming, as it is increasingly difficult to ensure a high overlay accuracy across large areas, because patterns on solid photomasks are fixed and fail to compensate for distortion and deformation of the substrate.

Due to these constraints, there's renewed attention in laser direct imaging (LDI). This is a maskless exposure technique, with circuit patterns written directly by scanning a focused laser beam. Thanks to elimination of photomasks, LDI enables rapid design changes, excellent resolution and overlay accuracy, and efficient support for high-mix manufacturing. In addition, this solution complies with a global trend for regulating and phasing out mercury-based light sources.

Despite these advantages, today's adoption of LDI for next-generation semiconductor packaging is limited to pilot-line equipment. Holding back full-scale industrial deployment of LDI are concerns over throughput, which is governed by the power of the laser light source. Existing semiconductor lasers emitting in the UV or violet have enabled early deployment of LDI systems, but their output powers are insufficient for high-volume production of advanced semiconductor packages.

The light-source bottleneck

In practical LDI systems, laser diodes must deliver a high continuous-wave output power while maintaining a stable emission wavelength and delivering reliable long-term operation. Simply increasing drive current is not a viable solution, as this increases self-heating, in turn increases the junction temperature of the laser diode, and leads to degradation of laser performance and an acceleration of output power degradation during long-term operation. To overcome this bottleneck, there must be a holistic approach, addressing both the laser diode and the thermal path that removes heat from the active region.

At Nuvoton Technology Corporation Japan we are addressing these issues during our development of a high-power InGaN-based edge-emitting laser diode that emits at 405 nm – that's a wavelength well matched to photoresist sensitivity, and widely used in LDI systems. By jointly optimising the device geometry, epitaxial waveguide structure and package thermal design, we have realised a high-output power and an improved wall-plug efficiency. These strengths have enabled improved long-term reliability through suppression of the junction temperature in violet laser diodes.

Managing heat at its source

A major limitation of the high-power operation of InGaN-based laser diodes is the rise in junction temperature at high current densities. Elevated junction temperatures increase carrier leakage from the quantum well, leading to a higher threshold current and a reduced slope efficiency. In addition, high junction temperatures accelerate device degradation during aging, due to the thermally activated formation of defects, which act as non-radiative recombination centres.

Compounding these matters, the unwanted loss mechanisms are amplified in violet and UV lasers, due to reduced band offsets in the III-N quantum well and increased internal optical loss, making high-power operation particularly sensitive to junction temperature. Consequently, it is imperative to suppress the rise in junction temperature under high-power operation, as this holds the key to high performance and long-term reliability.

One of our key design choices when working towards this objective is our laser chip geometry. We employ a wide stripe and a long cavity. By

expanding stripe width and extending the cavity length, we reduce the optical power density at the facet and the current density in the cavity – and ultimately distribute heat generation over a larger volume. There's a reduction in the thermal resistance of the laser chip, helping to suppress the rise in junction temperature.

For the devices described in this article, stripe widths are up to 45 μm , and cavity lengths up to 1400 μm . In addition to enabling a higher output power, this geometry mitigates degradation mechanisms exacerbated at high optical intensities. As emission heads towards the violet and UV, thermal management is increasingly critical in InGaN-based lasers. Shorter wavelengths are accompanied by higher photon energies, and typically higher operating voltages – both contribute to increased heat generation for a given optical output. At the same time, it is more challenging to ensure carrier confinement in the quantum well, due to reduced band-offsets, making device performance particularly sensitive to junction temperature.

There's a strong interplay between key laser characteristics. As the temperature rises, more carriers spill over from the active region, leading to an increase in threshold current and a rapid deterioration in slope efficiency. In high-power operation, this detrimental feedback loop between temperature and laser performance threatens to dominate device behaviour, and incremental design changes are insufficient for addressing these concerns. This issue is one of the reasons why high-power scaling of violet laser diodes has historically lagged behind that of devices emitting at longer wavelengths.

Our approach to improving the performance of violet lasers involves trimming its current density and thermal resistance through a widening of the stripe and an extension of the cavity length. However, just making these geometrical changes does not guarantee high-power performance. In short-wavelength lasers, internal optical loss plays an outsized role in determining wall-plug efficiency. There's the possibility that free-carrier absorption in heavily doped cladding and guiding layers, as well as scattering losses associated with imperfect optical confinement, offset gains resulting from a reduced current density and a lower thermal resistance of the laser chip.

To address this, we treat waveguide optimisation and thermal design as interdependent challenges. We know that reducing internal loss not only improves efficiency directly – it also cuts dissipated power for a given output, thereby alleviating the thermal load on the package. Meanwhile, if we lower the thermal resistance of the entire device, including the package, we will preserve the optical and electrical benefits of the optimised waveguide under practical operating conditions. Due to this strong coupling

between optical design and thermal management, it is critical for progress at the device level to be matched by advances in packaging.

Optimising the waveguide

Thermal management alone is insufficient to ensure success; a high wall-plug efficiency requires low internal optical loss and strong confinement of carriers and photons. We meet these requirements with a laser structure employing a single quantum-well active region to minimise threshold current, surrounded by carefully engineered guiding layers. The heterostructure features high-refractive-index InGaN waveguide layers to confine the optical mode, and AlGaIn cladding layers to provide carrier confinement and minimise optical leakage.

Our design also addresses free-carrier absorption. Arising from dopants and impurities in the waveguide, this impediment to efficiency is a major source of internal loss in short-wavelength lasers. We mitigate this with a design that minimises the spatial overlap between the optical mode and highly doped regions. It's a strategy that reduces free-carrier absorption without compromising carrier injection. Through optimisation of dopant concentrations and impurity profiles during growth of the heterostructure, produced by MOCVD, internal losses are reduced without compromising electrical performance. Our emitters also incorporate an electron-blocking layer – this helps suppress carrier overflow from the active region, and supports high efficiency, even at elevated drive currents.

Rethinking the package

While device-level optimisation improves intrinsic performance, it is the package that governs how effectively heat is extracted. In industrial laser applications, TO-CAN packages, such as TO9, are widely used. They provide a compact form factor and ease of integration. However, these packages have traditionally suffered from relatively high thermal resistance. Moving to much larger packages could address this issue, but it will complicate system design, so we have focused on trimming thermal resistance within the TO9 footprint.

To accomplish this, we redesigned the package as a monolithic copper structure, replacing conventional iron-based base plates and minimising internal thermal interfaces (see Figure 1). With this redesign, package thermal resistance is roughly two-thirds of that of a conventional iron-based TO-9 package. Behind this reduction is a substantially higher thermal conductivity, realised by replacing iron-based materials with copper.

We have made further gains by mounting the laser chip *p*-side down, as this shortens the heat path from the active region to the submount. Note that the choice of submount material is equally critical. We have evaluated various high-thermal-conductivity materials, and found that a diamond submount provides a substantial reduction in thermal resistance compared with conventional ceramic materials. Our use of a diamond submount, offering a thermal conductivity several times higher than that of typical ceramic materials, delivered an additional reduction in junction-to-case thermal resistance.

Taken together, the redesigned copper package and *p*-side mounting on diamond have doubled the overall heat extraction capability of the laser, and enabled stable high-power operation within a compact TO-9 package.

Power performance and efficiency


As one would expect, the combined effect of a tailored device geometry, waveguide optimisation, and a low-thermal-resistance package has produced marked improvement in high-power performance. Operating at a case temperature of 25°C, our 405 nm lasers deliver a continuous-wave output exceeding 12 W (see Figure 2). When the output is 5.2 W, wall-plug efficiency is 43.9 percent.


For industrial applications, high efficiency must be maintained at elevated temperatures. We find that even at a case temperature of 50°C, our devices exhibit wall-plug efficiencies above 40 percent, when delivering an output power of more than 5 W. Behind this success is the device's thermal robustness, essential for integration into compact LDI modules, where active cooling capacity is limited and there may be fluctuations in ambient temperature.

In addition to enabling high output powers and thermal robustness, our work highlights the flexibility of our device platform, which can be optimised for various performance priorities.


If we design devices with a reduced stripe width, we can enhance wall-plug efficiency by decreasing the threshold current – but this goes hand-in-hand with a fall in the maximum output power, due to increased current density and a higher thermal resistance. When making this trade-off, the laser geometry must be tailored to specific application requirements. While a high output power is essential for an LDI system seeking maximum throughput, applications

► Figure 1. Schematic of the low-thermal-resistance TO-9 package, featuring a monolithic copper base, *p*-side-down mounting, and a diamond submount, designed to enhance heat extraction from the laser chip.





► Figure 2. Optical output power and wall-plug efficiency measured under continuous-wave operation at a case temperature of 25°C for a 405 nm InGaN-based laser diode with a cavity length of 1400 µm and a stripe width of 45 µm.



► Figure 3. Long-term output stability of the laser diode during aging at 6 W and a case temperature of 25°C, demonstrating stable operation without catastrophic optical damage.

requiring a moderate output power are better served with geometries optimised for peak efficiency.

The key message is that our underlying design principles can support both high-power, thermally robust operation and high-efficient operation at lower output powers. What's required is to optimise the device geometry for the required characteristics.

Reliability at high power

When laser diodes operate at high output powers, there's severe stress on their facets and internal structures. This raises concerns related to long-term reliability. To look into this matter, we conducted aging tests under automatic current control, involving an initial output power of 6 W and a case temperature of 25°C. These conditions are representative of demanding industrial operation beyond typical commercial levels.

During extended operation, exceeding 1900 hours, we did not observe any catastrophic optical damage, and recorded minimal output degradation (see Figure 3). Based on the low degradation rate during our investigation, we can conclude that our device design provides a sufficient margin for long-term operation at this power level. These results underscore the importance of reducing optical and thermal stress through appropriate design of the device and its package.

Implications for LDI and beyond

The performance of our lasers overcomes a key barrier to the wider adoption of LDI in advanced semiconductor packaging. By delivering a substantially higher optical power from a compact semiconductor laser, we are enabling engineers in silicon fabs to shorten exposure times and boost throughput, without sacrificing system stability or lifetime. This leads to a step change in productivity, allowing LDI to provide the throughput required for industrial-scale manufacturing on large substrates.

Beyond LDI, there are opportunities for high-power, high-efficiency violet laser diodes in a number of other industrial applications, including additive manufacturing, precision processing and illumination systems – all benefit from short-wavelength, compact light sources.

As semiconductor manufacturing continues to evolve, the demands placed on light sources are only going to increase. Meeting them is our next step. We are targeting further improvements in efficiency, extension to other ultraviolet and visible wavelengths, and integration into application-optimised modules. Such efforts will build on successes to date, demonstrating that a coordinated approach – spanning epitaxial design, device geometry and thermal packaging – will enable semiconductor laser diodes to meet the stringent requirements of next-generation manufacturing technologies.

FURTHER READING

- S. Yoshida *et al.* "High-efficiency and high-power InGaN-based laser diode at a wavelength of 405 nm" *Proc. SPIE* **13896** 138960H (2026)



Agentic AI reshapes PCB and advanced packaging design



Phil Bishop, VP of Corporate Marketing, Global Customer Success Team at Cadence, talks to Phil Alsop, Contributing Editor of Advanced Packaging Magazine, about how agentic AI orchestration, combined with trusted EDA and system design and analysis tools, can help engineering teams manage the growing complexity of PCB and advanced package development. As hyperscale data centres deploy massive AI clusters and industries such as automotive and power electronics push towards increasingly intelligent, high-performance systems, Bishop discusses why connected design workflows, multiphysics analysis and intelligent automation are becoming critical to faster, higher-quality design realisation.

Phil Alsop: AI infrastructure, automotive electronics and high-performance computing systems are becoming increasingly complex. What are the biggest design challenges facing PCB and advanced packaging teams today?

Phil Bishop: One of the biggest challenges is that the design process is complex and the teams are often disparate. PCB layout, system architecture, packaging and analysis may all sit within different groups. Advanced packaging is now critical to high-performance electronics, but the

handoffs between those different teams can slow the overall workflow. A major opportunity is therefore to bring those teams together and allow data and design intent to move more effectively across the entire system design process.

For PCB and advanced packaging teams, improving that end-to-end workflow is becoming fundamental to increasing throughput.

Phil Alsop: Cadence describes AuraStack as an agentic AI platform for PCB and advanced packaging design.

What distinguishes agentic AI from the AI-assisted tools engineers are already familiar with?

Phil Bishop: AI-assisted tools can help optimise an individual step of the design process. Cadence Agentic AI goes further by looking across the entire workflow.

That means working from the schematic and design intent through architecture, PCB development and packaging, while coordinating the different tools and datasets involved.

Multiphysics simulation is another important part of that approach. The AuraStack AI Super Agent acts as an assistant to the lead engineer throughout the PCB and package design process, allowing electrical, thermal, mechanical and other analyses to be performed as the design evolves.

The main difference is therefore not simply applying AI to an individual task, but optimising across the complete workflow.

Phil Alsop: Cadence also offers AI super agents spanning chip, package and PCB development. Why is it important to connect those traditionally separate design domains?

Phil Bishop: The AuraStack AI Super Agent focuses on PCB and packaging, but Cadence also has super agents addressing other parts of the design flow.

The ChipStack AI Super Agent addresses areas such as verification and RTL creation, InnoStack focuses on digital and physical design, while the ViraStack AI Super Agent addresses analogue design, including layout migration and analogue performance optimisation.

Together they provide a broader suite of agentic capabilities.

The AuraStack AI Super Agent can, in some cases, sit above these different domains because it operates at the system level. For example, it can call other technologies where 3D-IC integration or analysis is required.

The important point is that system design increasingly cannot be treated as a series of isolated domains. Chip, package and PCB behaviour are interconnected, so the design workflow needs to reflect that.

Phil Alsop: The AuraStack AI Super Agent places a strong emphasis on integrated electrical, thermal, mechanical, signal integrity and power integrity analysis. Why is multiphysics co-optimisation becoming so important?

Phil Bishop: Power integrity, signal integrity, thermal analysis, EM-IR and other forms of multiphysics analysis have always been necessary. What changes with an agentic workflow is the ability to bring those analyses together and apply them continuously throughout development.

The system also builds what we call a mental model around the design intent and the environmental constraints in which the PCB and package will operate.

A 5G system, for example, will have very different requirements from an automotive electronic system. Thermal

behaviour, vibration, power integrity and other factors will therefore be assessed against the intended operating environment.


The value comes from linking agentic workflow optimisation with principled simulation technologies and the underlying Cadence tools. That allows the system to consider multiple physical effects together rather than treating them as separate checks at the end of the process.

Phil Alsop: In power-dense applications such as data centres, EVs and energy infrastructure, how does earlier visibility into thermal and power integrity issues improve system performance and reliability?

Phil Bishop: The earlier you can analyse thermal performance, power integrity and signal integrity, the better you can optimise the design.

Historically, these analyses were sometimes performed relatively late. A design team might complete much of the work before handing the system to a separate analysis group.

With an integrated multiphysics workflow, analysis can happen incrementally throughout development. At each stage, engineers can ask whether the design is delivering the performance required by the end application.



That is particularly important in environments such as data centres, where power consumption and thermal behaviour are critical.

Instead of discovering an issue near the end of the design cycle, engineers can evaluate performance continuously and adjust the architecture much earlier.

Phil Alsop: Cadence has cited significant productivity improvements from agentic AI. Beyond reducing manual work, how do you see AI changing the role of engineers?

Phil Bishop: AI is not a replacement for good engineers. What it does is allow engineers to move up a level and focus more attention on architectural and system-level decisions.

Agentic solutions are only as good as the design intent, constraints and hierarchy provided to them, and that knowledge ultimately comes from experienced engineers.

The opportunity is to augment engineering expertise and allow teams to explore far more alternatives than would have been practical with traditional workflows.

A strong system engineer can evaluate different architectures earlier instead of discovering late in the process that another approach may have been better.

There is a useful analogy with software development. Engineers gradually

moved away from working directly with machine code and assembly languages as higher-level abstractions became available. AI represents another step in that direction.

Engineers still need deep technical knowledge, but natural-language interaction and agentic automation can allow them to spend less time on lower-level tasks and more time defining what they want the system to achieve.

Phil Alsop: Cadence has worked with companies including NVIDIA, TSMC, Socionext, FORVIA HELLA and Schneider Electric. What early lessons are emerging from these collaborations?

Phil Bishop: One of the clearest benefits is workflow optimisation.

With the AuraStack AI Super Agent, we are seeing productivity improvements by connecting different elements of PCB and packaging design, sharing data more efficiently and applying multiphysics analysis across the workflow.

The emphasis varies between collaborations. With Socionext and NVIDIA, for example, much of the value has been around process and workflow optimisation. With TSMC, the focus has been more on packaging, including substrate design, routing and placement.

Across these examples, the common theme is improving end-to-end throughput.

Previously, different people or teams might have had to invoke separate tools for individual parts of the design process. AI-assisted tools could improve those local tasks, but agentic AI enables optimisation across system design, PCB design, chip design and packaging together.

That broader orchestration is where we see an important difference.

Phil Alsop: Does identifying problems earlier also help shorten the design cycle and improve time to market?

Phil Bishop: Absolutely. There are two major benefits.

The first is getting to market sooner. If the workflow becomes significantly more productive, customers can shorten development cycles and ultimately reach revenue earlier.

The second benefit is design quality.

When tools are connected and optimisation happens continuously, engineers can identify errors earlier and improve the quality and reliability of the final product.

So, this is not just about time to market. It is also about improving what is ultimately produced.

A workflow that can look at multiple issues simultaneously, run multiphysics simulation throughout the process and optimise continuously can improve both turnaround time and the quality of the end system.

Phil Alsop: How do you expect agentic AI to change PCB and advanced packaging development over the next few years?

Phil Bishop: We are still at a relatively early stage.

At the Design Automation Conference, there were many companies exploring different forms of agentic EDA. A lot of the early activity has focused on verification, while other developments are beginning to address IC physical design.

The important point is that agentic systems still need strong underlying tools and principled simulation. If you



move too far away from those core technologies and rely only on natural-language techniques or large language models, you are unlikely to produce a reliable engineering result.

The agents need to orchestrate trusted design and simulation technologies.

Over time, I expect agentic AI to become increasingly common across EDA. These systems should improve as they learn from more design activity, while experienced lead engineers and system architects will become even more important because they will define the constraints and prompt the systems at a deeper level.

We will see improvements not only in turnaround time but also in design quality across more areas of the EDA flow.

Phil Alsop: Is there one area of electronic system design where you

think agentic AI will create the greatest value?

Phil Bishop: The overall system workflow is probably one of the biggest opportunities.

Bringing PCB and package design together and then being able to simulate the complete system is potentially transformative because, for many years, those processes have largely depended on handoffs between different teams.

Now we have the opportunity to optimise across the entire flow.

I also think analogue design will be an important area. Analogue layout migration and optimisation are inherently complex processes, and complexity is exactly where agentic technologies can have significant impact.

For decades, there has been an ambition to build a single tool that does everything. A better model is to allow

Bringing PCB and package design together and then being able to simulate the complete system is potentially transformative because, for many years, those processes have largely depended on handoffs between different teams.

multiple specialised tools to exchange data intelligently.

That intelligence and orchestration layer is where technologies such as the AuraStack and ViraStack AI Super Agents can make a major difference.



PHOTONICS ASSEMBLY



Advanced PIC packaging in scalable volumes



- For many PIC platforms, using visible to mid-IR wavelengths
- The world's most accurate die bonding for light source co-integration
- Proven track record in CPO, AR/VR, quantum, and LiDAR areas
- From prototype to volume manufacturing
- Standard building block offering, including:
 - Lensed and PM fiber arrays
 - Spot size converters
 - Characterization packages
 - Free space optics

www.phix.com



Glass gains attention in the AI hardware race



As AI accelerates demand for more powerful semiconductor architectures, advanced packaging has become a critical battleground for improving performance, bandwidth and energy efficiency. Among the materials attracting growing industry attention is glass, offering a unique combination of dimensional stability, electrical performance and scalability for next-generation packaging. In this conversation, Phil Alsop, Contributing Editor of Advanced Packaging Magazine, speaks with Colin Schmucker, Global Sales and Product Manager for SCHOTT's Semicon Glass Solutions, about why glass is emerging as an enabling material for advanced packaging, the role of SCHOTT's "Semicon next" knowledge hub, and the opportunities and challenges that lie ahead.

Phil Alsop: SCHOTT recently launched "Semicon next" as a knowledge hub. What prompted this initiative, and what conversations are you hoping to drive within the semiconductor industry?

Colin Schmucker: Interest in glass for semiconductor applications has grown rapidly over the past few years, but many engineers and decision-makers are still unfamiliar with its capabilities as an advanced materials platform. When SCHOTT established its Semicon


Glass Solutions business in 2024, one of our priorities was to help educate the industry about where glass can add value.

The semiconductor industry has always benefited from open knowledge sharing, so Semicon next was created to provide an accessible educational resource. The first modules introduce the fundamentals of glass and explain why it is attracting so much attention. Future content will dive deeper into

specific applications, materials science and manufacturing technologies as the industry continues to evolve.

Phil Alsop: Advanced packaging has become one of the biggest areas of innovation in AI hardware. Where do you see glass fitting into this evolution?

Colin Schmucker: There are two major areas where glass is already making an impact.



The first is as a carrier for advanced fan-out packaging, particularly fan-out panel-level packaging. As AI devices become larger and more complex, manufacturers need extremely stable carrier materials capable of supporting high-volume processing while maintaining excellent dimensional stability. Glass provides that foundation.

The second area attracting significant attention is glass-core substrates (GCS). Here, glass forms the core of an organic substrate stack rather than acting as a temporary carrier. As AI processors continue to grow in size and complexity, glass offers the dimensional stability needed to manufacture larger substrates while maintaining flatness and reliability throughout processing.

Phil Alsop: What makes glass such an attractive material for these next-generation semiconductor packages?

Colin Schmucker: Several material properties come together.


One of the most important is stiffness. Modern packages combine many different materials, each expanding differently as temperatures change. A stiff glass core helps minimise warpage, which directly improves reliability.

Glass also offers highly tunable coefficients of thermal expansion (CTE), allowing manufacturers to better match the thermal behaviour of surrounding materials within the package.

Another major advantage is manufacturability. Glass supports extremely precise through-glass vias (TGVs), enabling high interconnect densities that are becoming essential for advanced AI devices.

Finally, its electrical properties are particularly attractive. Glass has a relatively low dielectric constant, very low electrical loss and an exceptionally smooth surface. Together these reduce signal attenuation and improve high-speed signal transmission, helping designers extract maximum performance from increasingly demanding AI systems.

Phil Alsop: Chiplet-based architectures are becoming central to heterogeneous integration. How does glass support this trend?



Colin Schmucker: Chiplets are really part of the broader move toward 2.5D and 3D heterogeneous integration. Although individual dies become smaller, the overall package becomes larger and more complex because multiple functional blocks are integrated together.

That requires finer redistribution layers, higher interconnect densities and exceptional dimensional control.

Glass provides the flatness, structural stability and processing precision needed to manufacture these increasingly dense interconnect structures. Whether used in glass-core substrates or advanced fan-out packaging, it enables manufacturers to integrate many active and passive components into a single high-performance package while maintaining tight manufacturing tolerances.

Phil Alsop: AI workloads demand enormous amounts of high-speed data movement. How does glass contribute to improved signal performance?

Colin Schmucker: Glass performs extremely well electrically. Its low dielectric loss minimises signal degradation as signals travel through the substrate, while its exceptionally smooth surface reduces losses caused by surface roughness, particularly at higher frequencies.

Looking further ahead, one of the most exciting opportunities is co-packaged optics. Glass isn't simply an excellent electrical substrate it also

has tremendous potential as an optical platform.

As optical communication moves closer to processors, glass could support optical interposers and optical bridges capable of integrating waveguides alongside electronic components. That opens the door to dramatically increasing chip-to-chip and package-to-package bandwidth while reducing overall power consumption.

Phil Alsop: Another characteristic frequently associated with glass is dimensional stability. Why is that becoming increasingly important?

Colin Schmucker: As interconnect dimensions continue to shrink, manufacturing tolerances become much tighter. Producing ultra-fine redistribution layers requires an extremely flat and stable substrate throughout processing.

Glass naturally provides that stability because it is manufactured as a highly homogeneous material with very consistent properties across the entire panel. It can also be polished to extremely low total thickness variation, allowing precise layer alignment throughout multiple processing steps.

This stability is equally valuable for advanced die stacking and hybrid integration, where maintaining accurate overlay between successive layers is essential for both manufacturing yield and long-term reliability.

Phil Alsop: Beyond signal integrity, glass is also recognised for its

dimensional stability. How does that benefit increasingly dense semiconductor packages?

Colin Schmucker: As package complexity increases, manufacturers are pushing towards ever finer redistribution layers and interconnects. Achieving these ultra-fine geometries depends on having an extremely flat and dimensionally stable substrate.

Glass naturally provides that stability because it is manufactured as a highly homogeneous material with virtually no point-to-point variation. It also maintains its flatness throughout thermal processing and can be polished to extremely tight thickness tolerances. These characteristics help manufacturers achieve the precise overlay required for advanced redistribution layers, die stacking and high-density interconnects while maintaining long-term reliability.


Phil Alsop: Looking specifically at interconnect density, what role do you see glass playing in enabling the next generation of semiconductor packages?

Colin Schmucker: The trend is straightforward smaller line widths and spaces enable higher interconnect density, which ultimately improves package performance.

Glass contributes in several ways. Through-glass vias support very fine interconnect structures, while the material's flatness allows increasingly precise redistribution layers to be built on top. Glass is also proving valuable as a carrier for hybrid copper bonding, where precise alignment is critical for stacking HBM memory and other advanced die-to-die architectures. As bonding pitches continue to shrink, the dimensional control offered by glass becomes increasingly important.

Phil Alsop: Glass has generated considerable excitement, but it is still not widely used as a substrate material. What challenges remain before broader adoption?

Colin Schmucker: It's worth remembering that glass is already well established in several areas of semiconductor manufacturing. It's been used for many years in MEMS, optoelectronics, sensor packaging and



wafer processing, and more recently as a carrier for advanced fan-out packaging.

Glass-core substrates (GCS) represent a newer opportunity, but they introduce new manufacturing challenges. There isn't yet a standard production process that everyone follows. Each substrate manufacturer is developing its own process flows, equipment and material combinations.

The industry has largely solved many of the fundamental engineering challenges, such as processing and handling glass, but attention is now shifting towards yield optimisation, cost reduction, process maturity and quality control. Manufacturers need confidence that glass-core substrates can be produced consistently, economically and with the long-term reliability required for AI systems.

Phil Alsop: Does accelerating adoption depend on greater collaboration across the semiconductor ecosystem?

Colin Schmucker: Absolutely. No single company can deliver this alone.

At SCHOTT, we understand glass extremely well, but advanced substrates involve many different materials, manufacturing processes and equipment suppliers. Glass manufacturers, substrate producers, equipment vendors and packaging specialists all need to work together to solve processing challenges and optimise manufacturing.

Collaboration is equally important for developing industry standards. Panel-level packaging is a good

example. Standardising panel sizes allows equipment suppliers to develop manufacturing tools that can serve the wider industry rather than supporting numerous incompatible formats. That kind of cooperation ultimately accelerates adoption for everyone.

Phil Alsop: Finally, as AI processors continue to evolve, what milestones should the industry watch for that will signal glass becoming a mainstream packaging technology?

Colin Schmucker: The first major milestone will be the commercial deployment of glass-core substrates in production devices. Whether that begins with high-performance AI processors or smaller applications remains to be seen, but it will demonstrate that the technology is ready for commercial manufacturing.

The next milestone will be the introduction of larger glass-core packages that fully exploit the material's advantages for high-performance computing. Those larger form factors represent one of the strongest value propositions for glass.

Ultimately, I think the industry is moving in that direction. Every major player is investigating glass in some form, and the technical benefits are becoming increasingly clear. There is still work to do around manufacturing and process optimisation, but the momentum is undeniable. As the industry continues to push for greater performance, bandwidth and efficiency, glass is well positioned to become one of the enabling materials for the next generation of advanced semiconductor packaging.



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

Streamlining the concept-to-production journey as advanced packaging gains momentum



An interview with Phil Alsop, contributing editor of Advanced Packaging, and Percy Gilbert (top image), Senior Vice President of Engineering at SkyWater Technology, explores how the company enables customers to move from early-stage concepts to manufacturable technologies through its Technology-as-a-Service (TaaS) model. The discussion spans advanced sensing, heterogeneous integration and quantum-adjacent systems, while Ross Miller (bottom image), Chief Strategy Officer, provides additional insight into advanced packaging capabilities including chiplets, 2.5D and 3D integration, packaged optics and quantum computing, a market projected to reach \$80 billion by 2030 as applications extend beyond the data centre.



Phil Alsop: We are going to cover technology as a service and advanced packaging. Starting with TaaS, SkyWater positions this as distinct from standard manufacturing services. Structurally, what makes it different from joint development agreements or custom process engagements?

Percy Gilbert: The biggest advantages come down to access, speed and end-to-end support. From an access standpoint, we work with materials that larger fabs would not typically allow into production environments. We have the

flexibility to bring in these more exotic materials and process them on real production tools.

That leads directly to speed. Customers are developing and experimenting on the same toolsets that will ultimately be used for manufacturing, so learning cycles are faster and more relevant. From an end-to-end perspective, we are not just developing processes. We are enabling platforms with full design infrastructure, including PDKs, so customers can move from innovation to product realisation.

In this space, we are dealing with new materials, new integration schemes and new processes. Our goal is to take these from proof of concept to manufacturable solutions quickly, which differs significantly from traditional joint development alliances that focus on scaling established technologies.

In conventional environments, especially at scale, anything that introduces risk to production is avoided. Exotic materials or unconventional approaches are not allowed because they could disrupt yield or tool



stability. Our model is different. We can bring those materials in and apply production-level thinking to early-stage development, helping customers bridge the gap between concept and manufacturability.

Phil Alsop: Turning to advanced packaging, it is evolving rapidly as a key enabler for next-generation semiconductors. What is driving the shift away from monolithic scaling?

Ross Miller: Advanced packaging is a natural evolution as transistor scaling approaches physical limits. Demand for performance has not slowed, so the industry is finding new ways to deliver it.

Packaging enables higher density integration, combining different types of chips into new architectures that improve performance, reduce power and increase flexibility. AI is the primary driver, but we are also seeing strong interest from areas like quantum computing. That will require complex integration and higher qubit counts, which advanced packaging can enable.

Phil Alsop: Percy, you describe innovation occurring deep in the technology stack before standards or volume production exist. What does that mean in practice?

Percy Gilbert: Traditional foundry models rely on well-characterised processes that can be combined to produce an end product. In our case, those baselines often do not exist.

We may start with entirely new materials or device concepts, working from first principles to understand how they behave in process tools. That includes developing recipes, defining specifications and establishing integration pathways.

A key challenge is translating laboratory concepts into CMOS-compatible flows. For example, developing an etch process for a material that has never been used in a CMOS environment. The question becomes how to integrate that material into a manufacturable process.

Working deep in the stack means solving those integration challenges, ensuring that novel materials and device concepts can ultimately be scaled into real products.

Phil Alsop: Ross, chiplet architectures are a major focus. What are their key advantages in terms of performance, yield and flexibility?

Ross Miller: Chiplets offer clear benefits in density, size, weight and power. They enable tighter integration and shorter interconnects, which improves performance and reduces power consumption.

They also provide flexibility. You can combine different process technologies in a single system, which is not possible with monolithic designs. For example, integrating compound semiconductor devices, advanced logic and sensors into a single heterogeneous architecture.

Phil Alsop: Percy, when innovation happens at the process and integration level, IP sensitivity increases. What are the main risks?

Percy Gilbert: The two main categories are process knowledge and integration decisions. At this stage, IP includes tool recipes, process flows and interactions between materials and design rules. These are highly sensitive.

Integration decisions are equally critical. Choices around thermal cycles, planarisation or stress management can determine whether a design is viable. These decisions form a key part of the IP.

This differs from traditional design-level IP such as layouts or system architectures. Here, the IP is embedded in how the process itself is created and controlled.

Phil Alsop: Ross, what challenges remain for widespread adoption of advanced packaging?

Ross Miller: Yield is a major challenge. These processes are still relatively early in their lifecycle, and the business model itself introduces complexity by combining components from different sources.

Ensuring known good die, validating integration and testing full heterogeneous systems are all ongoing challenges. Infrastructure is another issue. There is a need for significant investment to scale advanced

We may start with entirely new materials or device concepts, working from first principles to understand how they behave in process tools. That includes developing recipes, defining specifications and establishing integration pathways.

packaging capacity, particularly in Western markets.

The third challenge is knowledge. Many customers are still developing the expertise needed to adopt these technologies effectively and cost-efficiently. That is where collaboration becomes critical.

Phil Alsop: Are you confident yield challenges will be resolved?

Ross Miller: Yes. It is part of the natural evolution of any emerging technology. Advances in metrology, equipment and architecture will improve yield over time. These are engineering challenges, not fundamental barriers.

Phil Alsop: Percy, how do you manage IP protection without slowing collaboration?

Percy Gilbert: It comes down to structured team environments and need-to-know access. Teams are organised around specific projects with controlled access to data and tools.

Within each team, communication is open and efficient. Across teams, there are clear boundaries to protect IP. This allows us to maintain speed while ensuring proper separation between projects.

Phil Alsop: Ross, what role do co-packaged optics play in addressing bandwidth and energy constraints?

Ross Miller: They are critical. Data centre power consumption trends make it clear that efficiency improvements are essential.

Optical communication is a key lever, and integrating optics at the package level improves both bandwidth and energy efficiency. Co-packaged optics will be fundamental to scaling high-performance and edge computing systems.

Phil Alsop: Percy, when working on technologies that may not reach volume production for years, how do you define success?

Percy Gilbert: It depends on the project. Process stability, control and repeatability are always essential. Beyond that, success metrics can include early reliability, specific performance targets or even a single critical capability.

In some cases, achieving a defined metric, such as a specific etch quality,

is the entire objective. For projects moving towards production, we also look at readiness for scaling, even if that scale is relatively small in absolute terms.

Phil Alsop: Ross, how important is advanced packaging for quantum computing?

Ross Miller: It is foundational. Scaling quantum systems requires sophisticated packaging solutions.

It is not just about adapting existing architectures. Quantum systems operate under entirely different conditions, including temperature and material requirements.

Advanced packaging is essential to move from laboratory demonstrations to scalable, utility-level quantum computing.

Phil Alsop: Finally, how important is collaboration as these technologies evolve?

Percy Gilbert: It is critical. Our model depends on close collaboration with customers and across technical disciplines.

Many of the challenges we face are not fully defined at the outset, so solving them requires collective expertise.

Ross Miller: Collaboration is at the core of what we do. Unlike conventional foundry models, innovation here spans the entire technology stack. Customers bring device-level expertise, and we bring process and manufacturing expertise.

Together, we form integrated teams that turn new ideas into real products.

Percy Gilbert: That collaborative environment is one of the main reasons customers choose to work with us. It enables open innovation while maintaining the structure needed to deliver results.



PHOTONICS ASSEMBLY






Advanced PIC packaging in scalable volumes



- For many PIC platforms, using visible to mid-IR wavelengths
- The world's most accurate die bonding for light source co-integration
- Proven track record in CPO, AR/VR, quantum, and LiDAR areas
- From prototype to volume manufacturing
- Standard building block offering, including:
 - Lensed and PM fiber arrays
 - Spot size converters
 - Characterization packages
 - Free space optics

US + EU BASED
MANUFACTURING



www.phix.com

MYCRONIC

Bringing tomorrow's electronics to life



Surfx[®]
MYCRONIC

MRSI
MYCRONIC

Hprobe
MYCRONIC

vanguard
AUTOMATION
MYCRONIC

➤ Zeina Abdallah Senior Research Associate, University of Bristol.



From nanometre alignment to high-volume manufacturing

Advanced packaging has become a new frontline of semiconductor performance and manufacturing innovation. At the inaugural Advanced Packaging International Conference, held on 21–22 April at the Sheraton Brussels Airport Hotel, speakers revealed how chiplets, heterogeneous integration and co-packaged optics are intensifying demands for precision, thermal control, metrology and reliability as the industry moves towards high-volume production.

BY SARAB CHOPRA, EDITOR, ADVANCED PACKAGING MAGAZINE

ADVANCED PACKAGING is becoming a critical determinant of system performance.

That was one of the central messages to emerge from the inaugural Advanced Packaging International Conference 2026, where semiconductor manufacturers, equipment suppliers, research organisations and technology developers examined how packaging

is evolving beyond conventional assembly.

Across two days of presentations, speakers addressed chiplets, heterogeneous integration, hybrid bonding, photonic integration, co-packaged optics (CPO), advanced substrates, thermal management, metrology, reliability and manufacturing automation.

Although they approached these challenges from different positions in the supply chain, a common theme emerged. Advanced packaging is moving from an assembly function towards a highly controlled, multi-physics manufacturing discipline.

The transition is being driven by several pressures simultaneously. Moore's Law scaling is becoming increasingly expensive, while AI and

high-performance computing demand greater bandwidth and energy efficiency. At the same time, more functionality is being distributed across multiple dies and material platforms.

The result is a package containing logic, memory, analogue functions, photonics and potentially several material systems, all of which must operate together. Different materials expand at different rates. High-performance logic generates significant heat. Optical interfaces can be sensitive to displacement, while fine-pitch interconnects leave little room for placement error.

GlobalFoundries examined heterogeneous integration and CPO in the context of AI bandwidth. Physik Instrumente (PI) addressed precision and process control in photonic alignment, while PhotonBridge demonstrated a self-aligning alternative. Finetech explored advanced bonding and the convergence of front-end and back-end manufacturing. AIM Photonics focused on the relationship between photonic design and assembly. Fraunhofer IZM highlighted the role of collaborative pilot lines, IBM Research Zurich demonstrated a scalable approach to optical assembly, and Tyndall National Institute examined reliability in increasingly complex packages.

Together, these perspectives exposed the industry's central challenge: turning increasingly complex package architectures into repeatable and economical manufacturing processes.

Heterogeneous integration becomes a scaling strategy

Heterogeneous integration allows designers to select the most appropriate process technology for each functional block instead of forcing every function onto one monolithic die.

A compute die can use an advanced logic process, memory can be optimised for density, and photonic functions can be produced using silicon photonics or other material platforms. These components can then be combined through 2.5D or 3D architectures.

This can improve performance, functionality and manufacturing flexibility, but it also transfers complexity

into the package. Components produced using different processes, materials and thermal budgets must operate as one system. Their electrical connections, mechanical interfaces, thermal paths and optical interfaces all have to work together.

GlobalFoundries' presentation, Heterogeneous Integration and Co-packaged Optics: Meeting AI Bandwidth Demand, placed this challenge in the context of rapidly rising AI requirements.

For Himani Kamineni, Director of the Advanced Packaging Lab within GlobalFoundries' CTO organisation, the shift represents a change in where semiconductor scaling takes place.

"System-level innovation has really become the new scaling vector."

Conventional transistor scaling alone cannot address the combination of data movement, bandwidth, energy consumption and compute requirements emerging from AI infrastructure. Kamineni identified reach, bandwidth density, energy efficiency and compute efficiency as key requirements driving architectural change.

As data movement becomes a greater constraint, the physical distance between processing, memory and optical interfaces becomes increasingly significant. Packaging architecture can

therefore influence system performance rather than merely determine how a completed semiconductor connects to the outside world.

CPO is one consequence of this trend. Bringing photonic engines closer to switching or compute devices can reduce electrical losses and support greater bandwidth density. GlobalFoundries treats devices, integration, packaging and test as interconnected elements of its silicon photonics roadmap.

"We're co-optimizing our devices, integration, packaging, and test to deliver solutions at scale with the real AI bandwidth and workloads," said Kamineni.

TSMC is developing a similarly broad integration model. Its 3DFabric platform encompasses 2.5D and 3D integration, while its COUPE approach targets optical I/O integration through advanced semiconductor packaging.

A CPO assembly can include an ASIC, photonic integrated circuits, fibre arrays, bridges or interposers, through-silicon vias, redistribution structures and advanced substrates. Each element has its own thermal, mechanical and processing characteristics.

Electrical power produces heat; heat changes thermal distribution; thermal gradients cause mechanical



Testing is also moving earlier in the manufacturing process. As packages incorporate more valuable dies and additional process steps, manufacturers cannot afford to complete an entire module before discovering that one constituent component has failed

deformation; and deformation can displace optical components and alter coupling. The package may remain mechanically intact while its optical performance changes.

CPO is therefore a multi-physics manufacturing problem in which electrical, optical, thermal and mechanical behaviour must be considered together.

Two routes to photonic alignment

Photonic integration provides one of the clearest examples of conventional assembly tolerances being replaced by functional alignment requirements.

Fibre-array assembly can be particularly demanding because several optical channels must be aligned simultaneously. Movement may need to be controlled across multiple axes while the system measures the resulting optical response.

David Forer, Director Semiconductor Strategy, Global, at PI, examined sub-

micron active alignment and the role of precision motion and feedback in photonic assembly.

Rather than relying exclusively on mechanical reference features, active alignment measures the functional response while positioning takes place. The component is moved, optical power or coupling efficiency is measured, and the position is adjusted.

“Submicron alignment is the decisive technology for yield and throughput and cost control,” said Forer.

Precision cannot be defined simply by the resolution of a positioning stage. It also depends on repeatability, stability, measurement accuracy and response speed. The manufacturing system must establish not only whether a component reached its intended coordinates, but whether its actual position delivers the required performance.

PhotonBridge presented an alternative based on designing alignment into the package.

Paul Marshall, CEO of PhotonBridge, described a self-aligned cantilever architecture for combining III-V lasers with silicon photonics. MEMS-like silicon photonic waveguides extend from the silicon platform, while corresponding slots in the III-V component allow the laser to locate itself on the silicon photonics structure.

Marshall compared this to placing the III-V element onto the silicon photonics “like a Lego” structure that self-aligns.

“We can actually achieve coupling losses less than 0.4 dB across our process window.”

The approach is intended to replace lengthy active-alignment operations with passive mechanical alignment. This could be especially valuable when many lasers must be integrated and could potentially enable high-throughput assembly using conventional outsourced semiconductor assembly and test methods.





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

The company's silicon photonics platform uses six mask steps, according to Marshall, with the laser components manufactured separately and subsequently integrated.

PI and PhotonBridge therefore attack the same problem from different directions. One uses precision motion and functional feedback; the other changes the package architecture to reduce the precision demanded from the equipment.

Neither accuracy nor speed is sufficient in isolation. The manufacturing objective is to achieve the required functional performance through the simplest, fastest and most repeatable process.

Bonding moves precision into the interface

The same pressure is reshaping die bonding. Hybrid bonding combines dielectric-to-dielectric bonding with direct metal-to-metal connections, enabling finer interconnect pitches than conventional solder-bump approaches.

The electrical benefits can include shorter interconnects, reduced parasitics and greater connection density. However, as dimensions shrink, the bonding interface becomes a critical manufacturing surface. Particles, contamination, roughness and local topography can prevent uniform bonding, making planarisation, cleaning and surface preparation important contributors to yield.

Travis Scott, Senior Product Manager, Advanced Packaging Systems, at Finetech, placed hybrid bonding within a wider transformation in assembly.

"The scaling of Moore's Law is slowing down, the price per node is going up, and this is driving the next generation of scaling into advanced packaging."

Finetech's presentation illustrated how bonding processes are diversifying. Alongside hybrid bonding and thermal compression bonding, manufacturers are investigating laser-assisted, ultrasonic and ultraviolet-curable processes. Each offers a different balance of speed, temperature, accuracy and complexity.

The relevant process window depends on bonding architecture,



➤ David Volk, Product Manager Coating Systems, SUSS MicroTec Solutions.

pitch, overlay tolerance, die size and process sequence. Temperature adds another variable. During bonding and subsequent processing, materials expand and contract at different rates, so alignment established at one stage may change during later thermal cycles.

This helps explain why the traditional boundary between front-end semiconductor manufacturing and back-end packaging is becoming increasingly blurred. Advanced packaging now demands equipment, cleanliness, alignment and metrology that resemble aspects of wafer fabrication.

"It's definitely a very clear impulse from the industry that this is coming," said Scott.

Hybrid bonding therefore demonstrates a wider principle in advanced packaging: the interface cannot be separated from the materials, thermal history and mechanical behaviour of the package surrounding it.

Testing is also moving earlier in the manufacturing process. As packages incorporate more valuable dies and additional process steps, manufacturers cannot afford to complete an entire module before discovering that one constituent component has failed.

Finetech expects this convergence to extend towards wafer-scale and panel-level processing as substrates become larger and architectures more complex.

Warpage becomes a system-level constraint

Warpage has long been a concern in semiconductor packaging, but its importance increases as packages grow larger and interconnects become finer.

A heterogeneous package may combine silicon, glass, organic substrates, metals, adhesives and underfills. These materials have different coefficients of thermal expansion, stiffnesses and thermal properties. During bonding, curing and cooling, they do not expand or contract uniformly. Internal stresses develop and the package can deform.

Larger interconnects can absorb some deformation within the process window. Fine-pitch interfaces have less tolerance, while optical connections can be disrupted by very small relative movements. A package may therefore remain structurally intact but move far enough to reduce optical coupling or



YOUR ADVANCED PACKAGING SOLUTIONS

To cover as many advanced packaging processes as possible, we offer our BATCHSPRAY® Acid, which enables the combination of multiple process steps within a single system.

SUPPORTING Cu RDL FORMATION

Combine your processes



1. PR strip
2. Cu seed etch
3. Barrier layer etch

ADVANTAGES

Combining multiple processes boosts throughput while saving valuable floor space.

POST PLASMA DICING CLEAN

Tape frame processing



1. Wafers are diced
2. PR strip and polymer clean
3. sicOzone™ strip

ADVANTAGES

High throughput with 25 tape frames per sequence, gentle processing, no die peeling, and improved yield.

MORE
INFORMATION



BATCHSPRAY® Acid
Semi-automated or
fully automated.



➤ Peter O'Brien, Head of Research Group, Director of Photonics Packaging Pilot Line, Head of European Photonics Academy.



compromise an electrical interface. The solution is not necessarily to make every package as rigid as possible. Increasing stiffness can improve dimensional stability but transfer greater stress into brittle components and interfaces. More compliant materials may absorb displacement, but can introduce creep or long-term dimensional change.

The objective is predictable behaviour. If deformation can be modelled, measured and kept within the functional process window, it becomes a manufacturing variable that can be controlled. This places greater importance on package simulation, material characterisation and in-process metrology before final assembly.

Packaging moves upstream into design

Increasing manufacturing complexity is also forcing packaging considerations into the earliest stages of product development.

AIM Photonics' presentation on assembly design kits addressed the relationship between photonic integrated circuit design and packaging. The organisation spans PICs, interposers, electronic-photonic design automation, assembly and packaging, allowing it to approach the problem across several stages of development.

The principle behind an assembly design kit is straightforward: packaging constraints should be understood while the photonic device is being designed, not after fabrication.

David Haramé, Chief Operating Officer at AIM Photonics, explained:

"So you modify the PIC because you're doing the PDK and the ADK simultaneously."

The approach incorporates design rules, package-ready PIC layouts, optical characterisation and mechanical information into the development flow. Fibre interfaces, optical coupling structures and attachment points can affect the PIC layout, while thermal requirements can influence the placement of optical and electronic functions.

Development consequently becomes an iterative process connecting design, simulation, fabrication, assembly,

measurement and optimisation. Test structures can characterise fibre attachment and coupling behaviour, with the resulting information incorporated into subsequent design rules. Packaging knowledge is captured before production instead of remaining dependent on the expertise of individual assembly engineers.

Materials must be considered in the same way. Silicon, glass, organic laminates, metals, III-V compounds, silicon nitride and lithium niobate can have very different thermal, mechanical and processing characteristics. A process suitable for one component may damage another, while adhesive curing, moisture absorption and differences in thermal expansion can change dimensions over time.

Assembly sequence therefore becomes an engineering parameter. The package must be designed around the complete material stack and the order in which those materials can safely be processed.

Thermal management becomes spatial

Higher power density is making thermal management more complicated than keeping the overall package below a specified temperature. The distribution of temperature can be as important as the average.

A high-power compute die can create localised hotspots while neighbouring photonic components may have



➤ Joana Mendes, Researcher, Instituto de Telecomunicações.



➤ Christian Keil, Director Business Development & Sales, Lidrotec GmbH.

different temperature sensitivities. Thermal gradients can produce differential expansion, mechanical stress and optical displacement. Three-dimensional integration compounds the problem because stacked dies can place active layers farther from the primary cooling surface, while bonding layers and interposers add thermal resistance.

Cooling must therefore be considered as part of the complete package architecture. Heat spreaders, liquid cooling, thermal vias and advanced thermal interface materials can improve heat extraction, but they can also affect mechanical stiffness, stress and manufacturability.

The best thermal solution is not simply the one with the lowest thermal resistance. It must remove heat while preserving dimensional stability, reliability and compatibility with the rest of the package.

This is especially important for CPO, where a thermal gradient can change optical alignment and performance even when component temperatures remain within their individual operating limits.

Building a packaging ecosystem

The need to connect design, manufacturing and packaging was also central to the presentation by Ulrike Gnesch, Director at Fraunhofer IZM and head of the APEX pilot line.

Gnesch argued that increasingly complex packaging challenges cannot be solved effectively by organisations working in isolation.

“We cannot afford anymore to be in our silos as researchers.”

Fraunhofer IZM is developing the APEX pilot line as part of a European ecosystem intended to support advanced packaging development, technology transfer, manufacturing capacity and skills.

One example involves a consortium of 17 partners covering the value chain for fan-out wafer-level and panel-level packaging.

A common platform can establish capabilities that several companies can access, avoiding the need to repeat development work separately for each customer.

Fraunhofer IZM is also pursuing glass-core substrates and CPO through a consortium approach.

This includes multilayer glass substrates with through-glass vias, embedded waveguides and the co-assembly of electronic and photonic ICs.

Gnesch argued that European capabilities could focus on areas including high-power, high-frequency, photonic and biomedical applications rather than attempting to replicate the hyperscale semiconductor manufacturing models of Asia and the US. In this model, the packaging ecosystem itself becomes part of the competitive strategy.

Optical assembly seeks electronic scale

IBM Research Zurich offered another perspective on manufacturability.

Bert Offrein, Manager Co-packaged Optics and Principal Research Scientist at IBM Research Zurich, examined how optical connections could be incorporated into package assembly without repeating the highly manual processes that have historically made photonics expensive.

Earlier systems had demonstrated an important lesson: a technically viable optical package is not necessarily commercially viable.

“We must make sure that now optics comes in in a way that is scalable and manufacturable and cost-effective in the end,” said Offrein.

IBM is investigating polymer waveguides as one route to this goal. The approach uses a polymer waveguide interface between silicon photonics and optical fibre, with adiabatic coupling providing the transition between waveguide structures.

In one demonstration, a silicon chip was flip-chip attached to a glass substrate containing polymer waveguides, producing 100 simultaneous optical connections in a single attach operation.

That result is significant because conventional optical assembly can require individual alignment operations. The objective is to scale the number of optical connections in a manner more comparable with electrical assembly.

Offrein also identified a possible route towards offering the technology as

a packaging service, reinforcing the need for chip designers, foundries and packaging providers to build their processes around one another.

Reliability must measure performance, not just survival

Greater integration density also changes the reliability challenge.

A package may contain fine-pitch electrical connections, optical structures, fibre assemblies, organic materials and multiple interfaces. Each responds differently to thermal cycling, humidity, vibration and mechanical shock.

Finbarr Waldron, Principal Engineer at Tyndall National Institute, linked manufacturing complexity directly to reliability.

“With every advance in manufacturing complexity, it poses many new challenges for reliability.”

The problem becomes particularly severe in aerospace and space applications. Waldron described low-Earth-orbit temperature cycling between approximately -65°C and 125°C, potentially reaching thousands of cycles per year depending on orbital conditions. Vacuum, radiation,

mechanical shock and vibration add further stresses.

Such environments expose a fundamental issue: structural survival does not necessarily guarantee functional survival.

A fibre can remain mechanically attached while optical coupling degrades. An optical element can develop a defect without immediately producing an electrical failure, while vibration can alter alignment without physically separating components.

Reliability testing must therefore consider functional performance alongside structural integrity. The package has to survive, but it also has to continue operating within specification.

Waldron’s contribution also demonstrated why accelerated testing must reflect the interactions within the package rather than evaluate individual components in isolation.

Temperature cycling, vibration, humidity, vacuum and radiation can act on several interfaces simultaneously. Failure analysis therefore needs to distinguish between a component defect, an attachment failure and a gradual shift in system performance.

➤ Himani Suhag Kamineni ,
Head of Advanced Packaging
at GlobalFoundries.





SEP 17-19, 2026

Yashobhoomi (IICC), New Delhi

TRANSFORM TOMORROW



Register at www.semiconindia.org



OCT 13-15, 2026

Moscone Center, San Francisco, CA

TRANSFORM TOMORROW



Register at www.semiconwest.org

As commercial components are considered for more demanding applications, qualification also has to establish whether their original test regimes represent the conditions they will face in service. Advanced packaging expands the number of possible failure mechanisms, making construction analysis and physical analysis increasingly important alongside conventional electrical testing.

Closing the manufacturing loop

As tolerances shrink, inspection can no longer remain an activity performed only at the end of production. Alignment, overlay, surface roughness, coplanarity, warpage and bond quality increasingly need to be measured while manufacturing is taking place.

Active optical assembly can use coupling data to adjust component position. Bonding systems can use surface and overlay measurements to identify variation, while thermal measurements can expose deformation before it becomes a final-assembly problem.

The emerging manufacturing philosophy can be expressed as a continuous sequence: measure, calculate, correct and verify. Instead of following fixed coordinates and checking the finished result, equipment can increasingly measure the state of an assembly and respond during the process.

Process data can reveal drift and establish relationships between equipment conditions and final package performance. AIM Photonics'



➤ Paul Marchal, Chief Executive Officer Photon Bridge.

methodology shows how measurement can also feed back into design, with assembly results refining subsequent design rules. Metrology is consequently moving from final inspection into the production loop and, ultimately, into the design process itself.

This is both a technical and an economic requirement. Discovering a defect after several valuable components have been integrated represents a much greater loss than identifying process drift earlier.

The approaches presented at the conference connect different parts of this manufacturing loop. PI uses functional feedback to control active alignment. PhotonBridge seeks to remove an alignment step through package design. Finetech

combines placement, bonding and test. AIM Photonics feeds assembly measurements back into design rules.

Fraunhofer IZM connects capabilities across an ecosystem, while IBM aims to complete many optical connections in one operation. Tyndall tests whether the resulting package continues to perform throughout its operating life.

Together, they show that advanced packaging precision extends far beyond the motion stage. It includes control of thermal gradients, material interactions, warpage, interfaces, throughput and long-term reliability.

The next phase of semiconductor scaling will not depend solely on smaller transistors. It will also depend on the industry's ability to integrate different dies, materials and technologies with sufficient precision and stability to create repeatable products.

The central challenge is no longer simply how to connect more components. It is how to make the entire package behave as one system, and how to manufacture that system consistently, economically and at scale.



➤ David Hame, COO AIM Photonics and Associate Vice President for Process Development, EPDA, and Test Assembly and Packaging AIM Photonics.

ap-international.net

○ In 2027, AP International will be held on 13-14 April at the Sheraton Brussels Airport Hotel.





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

WE AUTOMATE YOUR TECHNOLOGY


Turning your manual process into a production-ready automated solution. From customizable **LIEB EFEMs** to fully integrated semiconductor manufacturing systems.

Our expertise includes handling of:

- Thin Wafers (100 µm)
- Highly warped Wafers
- Tape Frames
- Panels
- Masks



Taking you process from LAB to FAB.



Werner Lieb GmbH
Brückenstr. 32
96472 Rödental | Germany

Phone: +49 9563 7230-0
Email: info@werner-lieb.de

www.werner-lieb.de

