




ADVANCED PACKAGING

ISSUE I 2026

 AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

ADVANCEDPACKAGING.NEWS

CLOSING THE AI BANDWIDTH GAP WITH CO-PACKAGED OPTICS



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)





Packaging at the AI inflection point


➤ The semiconductor industry is entering a new era where system performance is increasingly defined by integration rather than transistor scaling alone. As artificial intelligence, high-performance computing and data-intensive applications accelerate demand for bandwidth, efficiency and compute density, advanced packaging is becoming the foundation for the next generation of electronic and photonic systems.

The industry is moving beyond traditional packaging approaches toward architectures that combine multiple dies, materials and technologies within a single system. Co-packaged optics, heterogeneous integration, wafer-level assembly and three-dimensional integration are no longer emerging concepts; they are becoming essential pathways to overcome the limitations of conventional scaling.

A key driver of this transformation is the growing challenge of data movement. AI workloads are pushing electrical interconnects toward their physical limits, creating demand for new architectures that deliver higher bandwidth with lower energy consumption. Co-packaged optics is emerging as a critical solution by bringing photonic and electronic components closer together, reducing signal loss and enabling higher-density optical communication.

However, moving these technologies into high-volume production requires more than device innovation. Advanced packaging must address a new generation of manufacturing challenges, including optical alignment, thermal management, interface standardisation, assembly precision and scalable test methodologies. The complexity of integrating electrical, optical and mechanical functions within a single package requires new approaches across the entire manufacturing ecosystem.

Heterogeneous integration is another area where packaging is becoming a strategic enabler. Future systems will combine specialised materials such as silicon, silicon nitride, thin-film lithium niobate and III-V semiconductors, each selected for its unique performance advantages. Technologies such as micro-transfer printing demonstrate how advanced packaging can enable selective integration of high-



performance devices without forcing incompatible materials into a single process flow.

As these architectures become more complex, manufacturing readiness is becoming as important as technical performance. Yield, throughput, reliability and process control will determine whether promising technologies can transition from laboratory demonstrations to commercial production. Testing is also evolving from a final verification step into a critical part of the design and manufacturing strategy, providing the data needed to improve processes and accelerate scale-up.

This shift represents a fundamental change in the role of packaging. The package is no longer simply a connection between finished chips; it is becoming the platform where system-level innovation takes place. The ability to integrate processors, memory, photonics and specialised accelerators will increasingly determine the performance and efficiency of future computing systems.


The next phase of semiconductor advancement will depend on collaboration across design, manufacturing, equipment and materials ecosystems. As AI continues to reshape technology requirements, advanced packaging will move from an enabling technology to a defining technology one that determines how the next generation of computing infrastructure is built.

Contents

Cover Feature

Closing the AI bandwidth gap with co-packaged optics

As frontier AI models push data center infrastructure to its limits, co-packaged optics is emerging as the packaging and photonics engineering response to a bandwidth challenge copper can't solve.



14

18 Optical engine fragmentation: The critical bottleneck to CPO volume production

The optical engine specification remains undefined, so testing cannot scale. Without scalable testing, CPO cannot enter volume production.

24 Micro-transfer printing for heterogeneous integration

Micro-transfer printing (MTP) is moving from single-device demonstrations on individual chips to wafer-scale validation.

29 Precision, reliability and integration in advanced packaging

Advanced packaging for AI, chiplets and photonics demands solutions for nanoscale alignment, thermal management and long-term reliability.

34 Advanced packaging moves to the center of system innovation

ECTC 2026 highlighted how heterogeneous integration, hybrid bonding, photonic packaging and thermal management are reshaping semiconductor innovation.

37 Building standards for the next era of advanced packaging


AP Magazine Editor Sarab Chopra speaks with Peter O'Brien, Head of Photonics Packaging at Tyndall National Institute.

40 Unlocking yield improvements in advanced packaging

Materials-driven failure analysis offers a systematic, multidisciplinary framework for understanding why defects occur, how they propagate, and what conditions exacerbate them.

46 How materials will unlock the future of advanced packaging

Advanced packaging is enabling the next era of semiconductor performance through 3D integration, chiplets and trillion-transistor systems.



54



37

50 Silicon photonics needs a new packaging strategy to reach volume

The schedule for silicon photonics has accelerated due to demands from AI and large data centers.

52 Unleashing the intelligence revolution

TSMC's 2026 Europe Technology Symposium revealed how AI is reshaping the semiconductor industry, driving advances in leading-edge logic, advanced packaging, photonics and system-level integration.

57 Motion design for the advanced packaging era

Advanced packaging has created new process bottlenecks related to the internal mechanics of the packaging equipment.

NEWS

06 Advanced packaging market to top \$120B
06 AI reshapes back-end equipment

07 SEMI updates assembly database

09 IEEE backs micro-transfer printing

10 TSMC expands Arizona packaging

10 NSW invests in advanced packaging

11 Report examines future packaging materials

11 Shanghai backs SiPh innovation

12 AI packaging boom drives investments



07

Editor

Sarab Chopra sarab.chopra@angelbc.com

Contributing Technical Editor

Richard Stevenson richard.stevenson@angelbc.com

Sales and Product Manager

James Cheriton james.cheriton@angelbc.com

Senior Sales and Product Development Executive

Ranjodh Avern ranjodh.avern@angelbc.com

Sales Executive

Nathan Diehl nathan.diehl@angelbc.com

Design & Production Manager

Mitch Gaynor mitch.gaynor@angelbc.com

Publisher

Jackie Cannon jackie.cannon@angelbc.com

+44 (0)1923 690215

+44 (0)2476 718970

+44 (0)2476 718975

+44 (0)2476 7823123

+44 (0)1923 690214

+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.

Advanced packaging market to top \$120 billion

Yole Group forecasts advanced packaging will become the dominant segment of the semiconductor packaging industry, driven by AI, HBM, and high-performance computing demand

THE GLOBAL advanced packaging market is expected to more than double in value over the next six years, growing from \$55 billion in 2025 to more than \$120 billion by 2031, according to Yole Group's latest Status of the Advanced Packaging Industry 2026 report.

The report highlights a structural shift in the semiconductor industry, with advanced packaging moving beyond its traditional backend role to become a key enabler of AI computing, high-bandwidth memory (HBM), power delivery, thermal management, and overall system performance.

By 2031, advanced packaging is expected to account for the majority of semiconductor packaging revenue.

Yole notes that pricing power has returned to parts of the market, although it remains concentrated in AI packaging, HBM packaging, memory testing, and high-end substrate technologies where capacity and

qualification barriers are highest.


In contrast, mature packaging segments continue to face pricing pressure from commodity inflation, currency fluctuations, and increasing competition. The report also identifies several critical supply chain bottlenecks.

TSMC's CoWoS capacity, Ajinomoto's ABF substrate materials, and Nittobo's T-glass products remain key enablers of leading-edge AI packaging, while companies including ASE, Amkor, PTI, SK hynix, Ibiden, and Unimicron are investing in capacity expansions.

However, many of these projects are not expected to come online until 2027–2028, suggesting supply constraints will persist in the near term. Yole further highlights the growing influence of geopolitics on the packaging ecosystem. While countries including the US, India, Japan, Malaysia, Vietnam, South Korea, and several European nations are investing in advanced packaging

and substrate manufacturing, leading-edge production remains concentrated in Taiwan, Japan, and South Korea, supported by established materials, equipment, and manufacturing ecosystems.

According to Yole Group, understanding pricing dynamics, capacity expansion, and supply chain bottlenecks will be increasingly important for semiconductor manufacturers, OSATs, and investors as advanced packaging continues to underpin the next generation of AI and high-performance computing systems.



AI reshapes back-end equipment

ADVANCED PACKAGING and AI infrastructure are transforming the semiconductor back-end equipment industry, with suppliers increasingly focusing on process integration, regional manufacturing, and strategic partnerships, according to Yole Group's latest Status of the Back-End Equipment Industry 2026 report.

Yole forecasts the back-end equipment market to reach \$9.5 billion in 2026, growing at a compound annual growth rate (CAGR) of 6.3% through 2031.

Demand is being driven by high-bandwidth memory (HBM), larger and more complex AI packages, and tighter

integration between logic and memory devices.

The report identifies thermocompression bonding (TCB), hybrid bonding, die attach, inspection, and metrology as key technology areas supporting next-generation AI packaging.

It also highlights increasing collaboration among equipment vendors, chipmakers, and OSATs to improve yield, scalability, and manufacturing readiness for advanced packaging.

Regionalization is emerging as another defining trend, with equipment suppliers expanding manufacturing and

service operations across Southeast Asia, the United States, Europe, and India in response to localization policies, export controls, and supply chain diversification.

At the same time, Chinese equipment vendors are expanding capabilities in selected back-end technologies, contributing to a more competitive global market.

According to Yole, future leadership in the back-end equipment sector will depend not only on tool performance but also on ecosystem partnerships, process integration, regional flexibility, and long-term customer support.

SEMI updates assembly database

The 2026 edition of SEMI's Worldwide Assembly & Test Facility Database expands coverage to more than 820 facilities and adds new insights into advanced packaging, AI, photonics and compound semiconductor capabilities

SEMI and TechSearch International have released the 2026 edition of the Worldwide Assembly & Test Facility Database, expanding coverage of the global semiconductor back-end manufacturing landscape and adding new data on advanced packaging, compound semiconductors, AI and photonics.

The updated database tracks more than 820 assembly and test facilities operated by integrated device manufacturers (IDMs) and outsourced semiconductor assembly and test (OSAT) providers, up from 750 facilities in the 2025 edition.

According to SEMI, the expansion provides greater visibility into facility capabilities, manufacturing technologies and end markets.

Among the new additions are facility-level classifications for advanced

packaging, substrate and interposer platforms, including silicon interposers, redistribution layers (RDL), laminate and ceramic substrates, as well as expanded information on assembly, test and module capabilities.

The database also identifies more than 360 facilities with compound semiconductor capabilities covering silicon carbide (SiC), gallium nitride (GaN), gallium arsenide (GaAs) and indium phosphide (InP).

More than 170 facilities support AI and high-performance computing (HPC) applications, while over 70 facilities focus on photonics and optical markets.

"AI infrastructure, photonics and compound semiconductors are expanding the range of capabilities

required across semiconductor assembly and test," said Clark Tseng, senior director of Market Intelligence at SEMI. "The 2026 database gives users a more detailed view of which facilities support these technologies and end markets."

In addition to expanded facility profiles, the 2026 edition includes updated OSAT revenue and capital expenditure rankings, helping users benchmark the industry's leading outsourced assembly and test providers.

The database is intended to help semiconductor manufacturers, equipment and materials suppliers, investors and analysts evaluate global manufacturing footprints, identify potential partners and monitor changes in the semiconductor back-end supply chain.



IEEE backs micro-transfer printing

Researchers say micro-transfer printing could enable scalable heterogeneous integration of diverse materials on silicon photonics platforms, paving the way for more advanced photonic integrated circuits

A NEW study published in the Journal of Lightwave Technology highlights micro-transfer printing (MTP) as a promising approach for enabling heterogeneous integration in silicon photonics, offering a pathway to more advanced photonic integrated circuits (PICs) for artificial intelligence, communications and sensing applications.

The research, led by scientists at Ghent University–imec, examines how MTP can overcome one of silicon photonics' longstanding challenges: integrating diverse material systems onto CMOS-compatible photonic platforms. While silicon photonics

benefits from mature semiconductor manufacturing processes, conventional silicon materials cannot deliver all the optical functions required for advanced photonic systems, particularly efficient on-chip light sources.

According to the study, MTP enables thin-film devices, or "coupons," fabricated from materials such as indium phosphide (InP), gallium arsenide (GaAs) and lithium niobate (LiNbO₃) to be transferred onto silicon photonic wafers. This allows each material system to be optimized independently before being integrated into a single photonic platform.

The paper highlights several recent demonstrations of the technology, including integrated silicon photonic engines combining optical and microwave signal processing, InP lasers integrated with silicon nitride waveguides, lithium niobate modulators on silicon photonic circuits, and heterogeneous electronic-photonic receiver platforms. The authors also discuss ongoing efforts to develop a pilot manufacturing line for large-volume industrial production while identifying remaining challenges related to yield, reliability, throughput and ecosystem development.



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



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



TSMC expands Arizona packaging

TSMC is adding advanced packaging capabilities to its Arizona semiconductor hub as rising AI chip demand drives expansion of US-based manufacturing capacity

TAIWAN SEMICONDUCTOR

Manufacturing Co. (TSMC) is expanding its Arizona semiconductor operations with plans that include the company's first advanced packaging facility in the United States, as rising demand for AI chips drives new investments in manufacturing capacity.

TSMC has committed significant additional investment to expand its Arizona site, where its first fab is already operational and additional fabs are under construction.

The expanded campus will include wafer fabrication facilities, advanced packaging capabilities and an R&D center, strengthening the company's US-based semiconductor supply chain.

The advanced packaging facility will support growing demand for AI and high-performance computing (HPC)

chips, where technologies such as heterogeneous integration and multi-die architectures are becoming increasingly important.

TSMC Chief Financial Officer Wendell Huang said the company is encouraged by progress at the Arizona site, noting that the first fab has achieved yields comparable to TSMC's leading facilities in Taiwan. The second fab is preparing for equipment installation, while construction of a third fab is underway.

Beyond Arizona, TSMC continues to expand advanced manufacturing and packaging capacity in Taiwan, where the company is developing additional leading-edge and advanced packaging facilities to meet long-term AI semiconductor demand.

The company acknowledged challenges including workforce availability and

infrastructure requirements as it scales its US operations. TSMC said it will work with government partners to address these constraints and support continued expansion.

The investment highlights the growing strategic importance of advanced packaging as AI accelerators become more complex.

As chiplet-based designs and high-bandwidth memory integration gain adoption, packaging capacity is becoming a critical part of the semiconductor supply chain alongside leading-edge wafer fabrication.

TSMC's Arizona expansion marks a step toward establishing a broader US semiconductor ecosystem that combines advanced chip manufacturing with next-generation packaging capabilities.

NSW invests in advanced packaging

THE NEW SOUTH WALES (NSW)

Government has announced the construction of a new 7,000 m² facility in Bradfield City that will house Australia's first commercial Semiconductor Advanced Packaging Facility, marking a significant investment in the country's semiconductor manufacturing capabilities.

The facility, known as the Second Building, forms part of the state's Advanced Manufacturing Readiness Facility (AMRF) programme and is intended to provide Australian companies with access to advanced semiconductor packaging, microelectronics and precision manufacturing services.

According to the NSW Government, the facility will feature ISO Class 5 cleanroom laboratories, advanced



microscopy, dicing, taping and electronics capabilities.

It is expected to support companies developing technologies in quantum computing, photonics, communications and medical devices. The Second Building is being delivered by Hansen Yuncken adjacent to the existing AMRF in Bradfield City.

Together, the first AMRF facility and the

new building represent a A\$300 million investment by the NSW Government in advanced manufacturing infrastructure. The announcement coincides with planning approval for Burrah Park, a A\$1 billion industrial estate within the Western Sydney Aerotropolis.

While the estate is expected to create more than 6,300 jobs across logistics and advanced manufacturing, the new semiconductor packaging facility is intended to strengthen Australia's domestic capability in semiconductor assembly and advanced manufacturing technologies.

NSW Minister for Planning and Public Spaces Paul Scully said the investment would help attract advanced industries to Western Sydney and create opportunities for businesses to expand their manufacturing capabilities in the region.

AI packaging boom drives investments

LG Chem, Asahi Kasei and Sumitomo Chemical are expanding their advanced semiconductor packaging materials businesses

MAJOR CHEMICAL suppliers are accelerating investments in next-generation packaging materials, targeting technologies critical to future AI and data center growth.

LG Chem has entered the semiconductor stripper market after qualifying as a supplier for Amkor Technology, strengthening its electronic materials portfolio.

The company is expanding its focus on high-value semiconductor materials, including packaging adhesives, low-dielectric materials, thermal management solutions and glass substrate technologies, as part of its broader R&D strategy.

Asahi Kasei has completed a new SUNFORT dry film photoresist slitting line at its Tainan, Taiwan facility, increasing capacity to support

advanced semiconductor packaging demand.

The company is also developing new packaging materials, including dry film photoresists for AI servers and photosensitive polyimide films for redistribution layers and substrate insulation applications.

Sumitomo Chemical is moving deeper into next-generation packaging through a joint venture with Samsung Electro-Mechanics to develop glass core substrates.

The new company, Glassem, will focus on the research, production and sales of advanced glass substrates, aiming to establish a commercial supply system from 2027.

Glass core substrates are emerging as a key technology for future AI and high-performance computing packages

due to their dimensional stability, low warpage and ability to support high-density wiring.

Companies including Intel, TSMC and several Asian suppliers are advancing glass substrate development as demand grows for larger and more complex AI chip packages.

Meanwhile, Chinese companies are also expanding their glass substrate capabilities, with firms such as BOE, Lens Technology and Dinglong developing TGV processes, glass substrate production lines and supporting consumables.

The wave of investment highlights a broader shift in the semiconductor industry, where advanced packaging materials are becoming a critical bottleneck and a major growth opportunity as AI-driven computing demand accelerates.



Materials innovation to power AI's advanced packaging era

A new IDTechEx report highlights the critical role of materials, hybrid bonding, and thermal management in enabling the next generation of advanced semiconductor packaging

A NEW REPORT from IDTechEx explores the materials and manufacturing technologies expected to shape advanced semiconductor packaging over the next decade, as demand for AI and high-performance computing (HPC) drives increasingly complex package architectures.

The report, *Materials and Processing for Advanced Semiconductor Packaging 2027–2037: Technologies, Players, Forecasts*, examines the material choices and process innovations required to improve manufacturing yield across 2.5D and 3D packaging technologies.

Among the key areas covered are silicon, organic, and glass interposers used in 2.5D packaging. While silicon

remains the preferred option for high-performance applications due to its fine routing capabilities, the report notes that organic interposers and glass core substrates are gaining attention as potential alternatives for larger AI and HPC packages because of their scalability and cost advantages.

The study also examines the growing adoption of Cu-Cu hybrid bonding for 3D integration, comparing wafer-to-wafer (W2W) and die-to-wafer (D2W) approaches.

Although D2W bonding offers greater flexibility for heterogeneous integration and chiplet-based designs, it also introduces stricter requirements for alignment accuracy, contamination control, and manufacturing precision.

Thermal management is identified as another key challenge, particularly as higher package power densities and stacked architectures create localized hotspots.

The report highlights growing interest in next-generation thermal interface materials, including metal- and graphene-based solutions, to improve heat dissipation and long-term reliability.

According to IDTechEx, advances in materials, process control, and manufacturing techniques will be essential to supporting future generations of AI, HPC, and chiplet-based semiconductor devices as advanced packaging continues to evolve.

Shanghai backs SiPh innovation

SHANGHAI is accelerating the development of its silicon photonics ecosystem through the Zhangjiang High-Tech 895 Incubator, which is supporting startups in optical computing, optical interconnects, and next-generation AI infrastructure.

The incubator has nurtured more than 1,600 technology companies since its launch in 2015, helping startups move from laboratory research to commercial deployment through funding, industry partnerships, and entrepreneurial support.

Among the companies benefiting from the ecosystem is optical computing specialist Lightelligence, which relocated its operations from Boston to Shanghai in 2020.

The company cited Zhangjiang's integrated semiconductor supply chain, manufacturing capabilities, and support for emerging technologies as key factors behind the move.

Lightelligence CEO Shen Yichen said silicon photonic chips could account for more than 30% of chips used in computing centers within five years, compared with less than 3% today, driven by growing demand for AI infrastructure.

The ecosystem also supports startups such as optical computing company Lightning AGI, which leverages Zhangjiang's close network of design, tape-out, testing, packaging, and application development capabilities to accelerate commercialization.

Beyond silicon photonics, Zhangjiang has built a comprehensive

semiconductor ecosystem spanning chip design, advanced packaging, testing, optical interconnects, thermal management, and power delivery technologies, positioning Shanghai as a key innovation hub for next-generation AI hardware.





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





Closing the AI bandwidth gap with co-packaged optics



As frontier AI models push data center infrastructure to its limits, co-packaged optics is emerging as the packaging and photonics engineering response to a bandwidth challenge copper can't solve.

BY DR. HIMANI SUHAG KAMINENI, DIRECTOR, ADVANCED PACKAGING LAB (APL), GF LABS; DISTINGUISHED MEMBER OF TECHNICAL STAFF (DMTS)

AI INFRASTRUCTURE is hitting a bandwidth wall as workloads drive explosive growth in compute demand. Large language models now reach into trillions of parameters and the IEA reporting annual data center power consumption is projected to reach approximately 950 TWh by 2030, roughly doubling today's levels. The speed that data moves between processors, memory and the network determines whether that compute is useful as AI workloads scale. The copper interconnects that have long linked processors to the network can't keep up.

As data rates rise, copper faces physical limits in reach, bandwidth density and energy efficiency. Copper-based interconnects can operate reliably over distances under two meters, but the AI clusters that modern workloads demand, where thousands of processors function as a coordinated system, can no longer operate under those constraints.


Co-packaged optics (CPO) create a fundamental shift in how data moves to alleviate the demand from AI. By integrating photonic and electronic components at the package level, CPO enables terabit-scale optical I/O that raises GPU utilization

and reduces energy consumption without sacrificing efficiency.

When copper interconnects hit the bandwidth ceiling

Traditional data center architectures have switches inside each server rack, connected to processors by short copper cables. That works when servers operate mostly independently, but AI clusters require thousands of processors working in constant coordination, which pushes switches further from the servers. The distance between servers and switches increases dramatically, so when high-speed data transfer is required, copper becomes impractical.

Optical networking extends interconnect reach from board-level distances to rack- and datacenter-scale (>10 meters to 100+ meters), while delivering bandwidth density on the order of ~1 Tb/s per millimeter in advanced co-packaged optics solutions. The energy benefits are measurable, reducing interconnect power from ~15 picojoules (pJ) per bit (pluggables) to ~5 pJ per bit today, with roadmaps targeting under 1 pJ per bit with CPO.



CPO redefines the AI interconnect stack

CPO moves the engine from the front panel into the package itself. In the architecture, a photonic integrated circuit (PIC) and electronic integrated circuit (EIC) are co-integrated on a shared substrate.

Electrical signal paths that previously ran centimeters across a circuit board are reduced to hundreds of microns inside the package. The reduction in signal travel distance directly translates to lower power consumption, higher bandwidth density and lower latency. Vertical component stacking enables compact, high-density system architecture that could not be achieved with traditional pluggable module designs.

CPO delivers system-level gains across performance, power, size and integration. CPO packaging enables designers to mix-and-match the best component for each product and integrate them in a single package, even if they were designed in different manufacturing processes.


Scaling CPO to production requires rethinking the packaging design, manufacturing process and the test strategy simultaneously.


Silicon photonics as the future of data center connectivity

Silicon photonics provides the manufacturing foundation CPO needs to scale. By building photonic devices and CMOS logic together on 300mm wafers, the industry can produce optical components at the massive volumes AI data center demand requires.

The industry is advancing through generations of silicon photonics capability from 100G to 200G to 400G per wavelength. Each generation requires co-development of the packaging features and photonic devices to connect them electrically, including copper pillars and the fiber interface components that connect them optically.

Increased bandwidth doesn't have to result in sacrificing power efficiency. Each new generation of CPO improves efficiency by shortening electrical paths, improving the components that convert





electrical signals to optical ones and reducing signal losses at each connection point.

Beyond 400G per wavelength, silicon alone hits its own physical limits. The next generation of silicon photonics platforms is incorporating new optical materials that can carry more data per wavelength, with a goal to reach multi-terabit per wavelength operation.

Advanced packaging extends integration beyond photonics

The packaging techniques, including 3D heterogeneous integration (3DHI), that make CPO work are not limited to photonics. Wafer-to-wafer (W2W) and die-to-wafer (D2W) bonding are uniquely positioned to solve data center problems, but the same capabilities apply across a much wider range of products.

In the data center alone, the applications extend beyond the optical transceiver. The bonding techniques that connect photonic and electronic ICs can integrate amplifier and driver circuits directly with silicon photonics, cutting the signal losses that grow with every additional interface at high data rates.

Outside the data center, the applications look entirely different. Bonding micro-LED structures to CMOS logic opens new paths for display integration. Applying this to power devices and image sensors is enabling next-generation platforms in automotive, industrial and consumer applications.

The advanced packaging capabilities built to meet the most demanding requirements in AI data center connectivity are proving valuable as they are applied to adjacent markets. Investing in the packaging building blocks for CPO is an investment in a broader advanced packaging platform.

Integrated CPO manufacturing demands end-to-end process control

Scaling co-packaged optics from lab demonstration to high-volume production introduces significant manufacturing challenges. Whereas chips are single

devices made in a single process, CPO is a system built across manufacturing stages from photonic wafer fabrication, optical I/O integration, wafer-level optical testing, EIC and PIC assembly and final module test.

Optical testing introduces a new level of complexity. At the wafer level, testing requires coupling light in and out of devices at scale, with the alignment precision and throughput that high-volume manufacturing demands.

The dependency between stages has a direct implication for how CPO gets manufactured. Splitting the flow across multiple suppliers creates gaps in yield accountability that are difficult to close at data center volumes. Knowing where yield is lost and being able to fix it requires visibility across the entire chain.

The defining packaging challenge


AI scale-up networks are growing faster than copper-based architectures can accommodate, and the data center infrastructure industry is responding with a generational shift in how processors connect to the network. Co-packaged optics is at the center of that shift, converging silicon photonics process capability, advanced packaging techniques, heterogeneous integration and manufacturing infrastructure.

As these systems grow, the engineering decisions made at the package level increasingly determine system-level performance. Closing the bandwidth gap that AI infrastructure faces is the defining packaging challenge of this moment, and co-packaged optics is how the industry can meet it head on.


REFERENCES

- <https://www.technologyreview.com/2026/01/07/1130795/what-even-is-a-parameter/>
- <https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary>


Fiber to laser




Laser to SOI chip



Lens on laser



Lens on chip



Connect the World by lighting the way
- the essential link in next-generation optical connectivity!



vanguard
AUTOMATION
MYCRONIC

Optical engine fragmentation: The critical bottleneck to CPO volume production

The optical engine specification remains undefined, so testing cannot scale. Without scalable testing, CPO cannot enter volume production, regardless of how mature the photonics become.

BY COLLINS SUN, TECHNICAL DIRECTOR AND BARRY YEH, PROJECT SECTION MANAGER, GLOBAL SALES & OPERATING CENTER, WINWAY TECHNOLOGY


The bottleneck has moved: from physics to specification

MOVING THE OPTICAL ENGINE (OE) next to the switch ASIC to shorten electrical paths, cut power and buy higher bandwidth density the logic of CPO is no longer in dispute. Yet the clearer the technical roadmap becomes, the more the production timeline keeps slipping. The question is not whether a single device can be built, but whether it can be built and tested reliably, cheaply and at scale. Trace those difficulties backwards and they all point to the same source: the optical engine specification has yet to converge. When the spec is inconsistent, the connectors, fiber array units (FAUs) and test equipment downstream cannot be standardized. Test is the most painful symptom; the undefined OE specification is the underlying disease.

Billions invested, yet the spec remains open

Light replacing copper is a matter of physics: copper loss is measured per centimeter (around 1 dB/cm at 53 GHz), optical loss per kilometer (below 0.2 dB/km) a difference of more than four orders of magnitude, and light is immune to electromagnetic interference. Once bandwidth is pushed to the limit by AI, letting the electrical signal travel only the final few millimeters is the only way forward. That is why Cisco, Nokia, AMD, Marvell, Google and Intel have each poured billions into investing in or acquiring silicon photonics companies over the past two years ostensibly to buy technology, but in reality, to buy an entry ticket to help define the optical engine specification.

➤ **Figure 1:**
From OE
to module
Optical engine
diversity
varies across
vendors.



A settled direction, however, does not mean manufacturing and test are ready.

Every integration choice is a specification fork

At the heart of CPO is the optical engine, which integrates a photonic integrated circuit (PIC) and an electronic integrated circuit (EIC) to convert between electricity and light in both directions. But integration comes in degrees: a mature CPO places the OE and the compute die together on an interposer, not merely side by side on a substrate which layer the optics sit on, and how they connect, is itself a specification, and on that there is still no consensus.

Modulator selection is the clearest illustration: no single modulator can satisfy footprint, bandwidth, thermal stability, and linearity all at once. A Mach-Zehnder modulator offers the best linearity and maturity but a large footprint; a micro-ring modulator is tiny enough to pack many along the package edge yet demands tight thermal control. So each is a trade-off and every vendor takes its own path according to its system goals. From the light source through modulation and routing to I/O coupling, every stage carries a similar divergence, and the shape and interface of the optical engine never settle into a common form.

One further point is non-negotiable: the silicon photonics waveguide is itself single-mode, so CPO can only use single-mode fiber (SMF), with a core only about 9 μm across. That figure directly sets the hardest hurdle on the test line the nano-scale alignment between fiber and chip.

Fragmented specs make standardized test impossible

Traditional IC test scales because packages and contacts are highly standardized: the same equipment and the same interface test a huge variety of chips. CPO removes that premise when every vendor's optical engine looks different, test reverts to something close to bespoke, with each new design demanding a reworked interface and a retuned machine. The "one method to test millions of units" that volume production depends on cannot exist

Category	Items requiring standardization
Mechanical	OE size, keep-out zone, optical port location, FAU interface
Optical	Fiber pitch, coupling method, SMF interface, alignment tolerance
Electrical	EIC interface, power delivery, high-speed signal map
Thermal	OE thermal control, micro-ring temperature control range
Test	Probe access, golden FAU, test sequence, known-good OE criteria
Reliability	Alignment drift, thermal cycling, optical power stability

until the spec converges. This disease shows up in three places precisely the three pain points WinWay focuses on: optical alignment, testing efficiency and signal integration. The three are interlocked, and all three are amplified by the undefined upstream spec.

Concretely, the convergence that volume production needs spans six domains. Until the industry agrees on each, no test platform can be standardized across vendors (see chart above).

Alignment and test flow: where throughput collapses

Consider alignment first. The core is only 9 μm, and tolerances stack: fiber roundness, core concentricity, and the depth and angle of the V-groove layer up to a total of about 3.8 μm roughly 40 percent of the core diameter, leaving very little margin. An optical engine has to align not one fiber but a whole array; a mainstream OE today carries around 32 fibers, and active alignment takes roughly 5 seconds per fiber. As future OEs scale to higher fiber counts, that alignment time only accumulates further.

What makes this a production-cost problem rather than a minor inconvenience is that alignment time is added on top of the actual test not absorbed within it. Consider a CPO device integrating 32 OEs, each with 32 fibers. Alignment cannot run in parallel: at 5 seconds per fiber, each OE takes about 160 seconds, and 32 OEs in series consume roughly 85 minutes. If the genuine electrical and optical test of that device takes, say, 60 minutes, the total occupancy of the test cell becomes about 145 minutes and nearly 60 percent of that is alignment overhead that exists only because the fiber-to-chip coupling is so unforgiving. Every device therefore carries a test cost that is not one test time but the true test time plus a large, irreducible alignment penalty. As fiber counts rise, that penalty grows, and units per hour (UPH) fall accordingly.

Traditional package test is pick and place; a CPO device, once in the socket, still has to be aligned fiber by fiber, with no plug and play at deployments running into millions of GPUs, a throughput-collapsing bottleneck. Worse still, the optical-electrical arrangement differs at


every station from wafer to module: at wafer level the electrical probes come down from the top while the optical probes enter from the bottom, the arrangement reverses at die level, and high-end packages with contacts on the top surface require both faces to be probed at once double-sided probing.


Two directions are being pursued to attack the alignment penalty. The first is a golden FAU: a master fiber-array reference, characterized once to a known-good position, against which all production FAUs are aligned to a fixed reference rather than re-optimized from scratch each time collapsing a per-fiber search into a single transfer step. The second is self-alignment at the die level, where the coupling structures are designed so that the fiber array settles into its optimal position mechanically or through guided features, reducing the active search the machine must perform. Both attack the same throughput wall from different angles, and both can meaningfully shorten alignment time. But a reference part is only meaningful if the engines being tested look sufficiently alike, and a self-alignment scheme is only reusable if coupling geometries are shared so the ceiling on either approach is set, once again, by whether the optical engine specification converges.

Final test determines product value, not just yield

CPO also pushes a problem that once belonged to “chip yield” down to package test. To stack up bandwidth, some customers integrate multiple optical engines on a single IC a single ASIC ringed by six, eight or more OEs and these can only be verified one by one once the package is complete. The final value of the product is decided only at this last test step.

➤ Figure 2: Light vs electricity Copper loss measured per centimeter, optical loss per kilometer.





➤ **Figure 3:**
CPO test flow

Test therefore shifts from pass-or-fail to a grading exercise: confirm how many OEs in the package work, then sort the device into bins by that number, each corresponding to a different spec and price. A device targeting eight channels but found to have only six working OEs is not necessarily scrapped; it is down-binned and shipped as a six-channel product. This “integrate first, test later, then grade” model turns test from a quality gate into something that directly determines product mix and margin. And because the device is only verified after packaging with the ASIC, the multiple OEs, the interposer and the FAU all already invested a single OE failing at final test costs not just that OE but the whole assembly. Test is no longer merely confirming yield; it is protecting an investment already committed.

➤ **Figure 4:**
FAU alignment
Stacked tolerances reach $3.8\ \mu\text{m}$, about 40% of the $9\ \mu\text{m}$ core.

Four extremes converging on a single interface


The third pain point manifests in the test interface itself. CPO and next-generation advanced packaging are pushing the package to extremes: sizes toward

$200 \times 200\ \text{mm}^2$, package warpage up to $0.6\ \text{mm}$, pin counts approaching 30,000, data rates toward 448Gbps, and per-device power past 4,000 W with some customers’ requirements reaching into the tens of thousands of watts. These pressures bear down on the same interface at once.

Current is the sharpest. In some customers’ requirements the current at a single contact already reaches 6A, and such contacts are distributed across a package carrying tens of thousands of contacts. A conventional spring probe conducts through a single point of contact, too small an area for this current density, generating large amounts of Joule heat at best drifting contact resistance and losing measurement accuracy, at worst burning out the contact or melting the solder ball; meanwhile the preload force required by so many contacts warps the socket housing, degrading contact consistency.

CPO also rewrites the role of clearance. A conventional socket sets aside avoidance space only as the case demands, never as a necessity; with CPO, where the probes only carry current and the optical signal travels through the fiber array, the mechanism must actively make room for the alignment apparatus to reach the package from the side or top. Clearance shifts from optional to a structural necessity, and socket design shifts from “how to arrange the pins” to “how to leave a clean channel for the light while arranging them”.

The next-generation test interface must satisfy four requirements simultaneously: stable and low contact resistance across the full current range;




sufficient current capacity per contact without localized thermal runaway; a housing that maintains geometric integrity under the aggregate preload of tens of thousands of contacts; and a contact geometry that distributes current and heat over area rather than concentrating them at a point.

WinWay's HyperSocket addresses all four extremes. The architecture pairs a spring probe which provides the controlled compression travel needed to accommodate package warpage up to 0.6 mm with a conductive elastomer layer that expands the electrical contact from four crown points into a distributed surface. This surface-contact geometry reduces contact resistance, raises sustainable current capacity per contact to levels compatible with the 6 A per-pin requirements already appearing in leading customers' specifications, and eliminates the localized Joule heating that causes resistance drift and, in extreme cases, solder ball damage.

Beyond current handling, HyperSocket's housing is engineered to resist the cumulative warpage force generated by high pin-count package application: the compliance is distributed across the elastomer rather than concentrated at the probe tip, so housing deformation does not degrade contact consistency across the array. For CPO-specific configurations, the socket geometry preserves a defined clearance channel a structural necessity, not an option.

The path to volume: converge the optical engine spec, then scale the test


Before the CPO ecosystem matures, the industry need not simply wait and in practice, it is not waiting. Co-packaged copper (CPC) routes high-speed signals out over copper cable mounted on top of the package, preserving the electrical interface while deferring the optical alignment challenge. For system architects, CPC sustains today's bandwidth specifications without requiring the nano-scale fiber alignment infrastructure that CPO demands;



➤ Figure 5: Package roadmap based on TSMC

Advanced packaging is going to have larger packages, higher pin count and higher heat density.

for test engineers, it means the socket interface remains predominantly electrical, and existing high-speed test methodologies require only incremental adaptation rather than a fundamental redesign. WinWay's test socket platforms already support CPC configurations. The same HyperSocket architecture that handles CPO's surface-contact and clearance requirements is equally applicable to CPC's high-current, large-package demands there is no product gap between today's CPC ramp and tomorrow's CPO transition. CPC and CPO will coexist in the data center stack for the next several years, deployed in different layers according to reach and bandwidth requirements. The transition is not a cliff edge but a gradient, and the test infrastructure needs to serve both simultaneously. But coexistence is




➤ Figure 6: Four converging extremes on one test interface

Ultra-large package, high pin count, high speed, high heat density will happen all at once in advanced package testing.


not production readiness: nano-scale alignment, the throughput dragged down by fiber-by-fiber alignment, the multi-OE yield that can only be graded after packaging, the test interface pushed to its limits not one is an isolated engineering problem; all stem from an optical engine specification that remains undefined. Once the spec fragments, the downstream cannot be standardized, test cannot scale, and volume production cannot follow. That CPO can be built has never been the question; whether it can be tested affordably, and to a standard is. And the answer begins with defining the optical engine specification.

REFERENCES

- [Fig. 1] Official Website Of Company.
- [Fig. 2] WinWay CPO Forum, "The Evolution of Co-Packaged Optics (CPO): Advanced Testing Methodologies for Silicon Photonics." 2026
- [Fig. 4] WinWay CPO Forum, "The Evolution of Co-Packaged Optics (CPO): Advanced Testing Methodologies for Silicon Photonics." 2026
- [Fig. 5] TSMC Technology Symposium. 2026
- [Fig. 6] WinWay CPO Forum, "The Evolution of Co-Packaged Optics (CPO): Advanced Testing Methodologies for Silicon Photonics." 2026
- [Fig. 7] SEMICON Taiwan Forum, "Built for Complexity: Advanced Packaged Test Solution for High-Speed, Large & High-Pin-Count Applications." 2025
- [Fig. 8] WinWay CPO Forum, "The Evolution of Co-Packaged Optics (CPO): Advanced Testing Methodologies for Silicon Photonics." 2026



➤ **Figure 7: HyperSocket™**
Beyond CPO applications, HyperSocket provides the soft, stable contact required by the rigorous testing of next-generation advanced packaging.



➤ **Figure 8: Dual-architecture configuration on a single chip** CPC as a transitional topology toward CPO.



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

From lab demonstration to pilot-line validation: Micro-transfer printing for heterogeneous integration

Micro-transfer printing (MTP) is moving from single-device demonstrations on individual chips to wafer-scale validation. At imec and Ghent University, the TRANSVERSE pilot line is evaluating whether this technique can meet the manufacturing metrics that matter for heterogeneous photonic integration: overlay, transfer yield, throughput, backend compatibility and post-transfer device performance and reliability.

BY DR. EMADREZA SOLTANIAN, R&D ENGINEER, DR. EMIEL DIEUSSAERT, BUSINESS DEVELOPER, DR. MAXIMILIEN BILLET, R&D PROJECT LEADER, DR. SARAH UVIN, R&D PROJECT LEADER, PROF. DR. BART KUYKEN, CO-LEAD, PROF. DR. GUNTHER ROELKENS, CO-LEAD, TRANSVERSE




➤ Figure 1: Schematic of coupon fabrication and micro-transfer printing process: a) Source wafer of thin film devices, b) Thin film devices coupon definition, c) Tether definition and coupons encapsulation, d) Releasing coupons by undercut etching the release layer, e) Coupon picking using a PDMS stamp, f) Coupon printing onto the target circuit, g) Parallel printing of coupons from multiple source wafers onto a target wafer.

The packaging problem in advanced photonic systems

ADVANCED PHOTONIC packages increasingly require functions that cannot be realized efficiently in a single process. Lasers, modulators, detectors and low-loss routing each favor different materials. No single platform can provide all of these functions with best-in-class performance. Silicon (SOI) enables dense routing and detector integration, but lacks efficient light sources and ultra-high-speed modulation. Silicon nitride (SiN) offers ultra-low-loss routing, but no active devices. Thin-film lithium niobate (TFLN) enables ultra-high-speed electro-optic modulation, while III-V materials are required for lasers, amplifiers and efficient photodetectors. Because these materials cannot be co-processed in a common front-end flow, they must be combined at the packaging level. Current approaches include flip-chip assembly on wafer level, wafer bonding and butt coupling of chips on module level. These are used in production today, but each comes with trade-offs in alignment complexity, integration density, throughput or process compatibility.

Limits of existing integration approaches

A recurring limitation of existing approaches is that they often handle much more material than the final function requires. Large dies or full wafers may be bonded or assembled even when the required non-native devices occupy only a small fraction of the target chip area. This adds cost, increases process complexity and can couple the yield of otherwise independent platforms. Flip-chip assembly is mature and selective, but scaling to large device counts increases placement time and alignment overhead. Wafer



➤ Figure 2: Wafer-scale printing of TFLN modulators onto a SiPh wafer. a) A 200mm SiPh wafer. b) microscope images of printed modulators on a 200mm wafer. c) cross-sectional schematic of the TFLN modulators on SiPh circuit. d) Alignment accuracy of printed coupons over the wafer. e) demonstrated performance of wafer-scale printing of TFLN modulators on a 200mm SiPh wafer.

bonding enables dense integration, but requires co-processing of dissimilar materials. Multi-chip integration on module level offers flexibility, but introduce optical loss and packaging complexity. As photonic systems become more complex, these limitations become more relevant. There is a need for an integration approach that combines:

- Selective placement of devices
- Highly efficient source and target material use
- Sub-micron placement accuracy
- Wafer-scale compatibility
- Compatibility with a wide variety of materials
- The integration on complex wafers without disrupting front-end processes

What micro-transfer printing adds

Micro-transfer printing offers a different integration strategy. Devices are fabricated and tested on their native source wafers, typically at high density. Only known-good, ultra-compact coupons are then transferred and printed onto the target wafer.

This allows III–V devices, TFLN modulators and silicon-based photonic circuits to be processed independently and combined afterwards, without forcing process compromises. The key question is no longer whether this concept works at device level, but whether it can meet manufacturing metrics such as:

- Overlay accuracy within wafer and from wafer-to-wafer
- Throughput per print cycle and per wafer
- Transfer yield
- Post-transfer device performance & reliability
- Compatibility with standard backend processing

This is the focus of the TRANSVERSE pilot line.

Micro-transfer printing in a nutshell

MTP is a parallel pick-and-place technique for integrating microscale devices from one wafer onto another. Devices are first fabricated on a source wafer and released using a sacrificial layer. A compliant stamp picks up these released device coupons and transfers them onto a target wafer. Multiple devices can be transferred in a single step, with sub-micron alignment.


Wafer scale validation: The TRANSVERSE pilot line

The TRANSVERSE pilot line addresses the step between laboratory demonstration and manufacturable process flow. The aim is not only to print individual devices, but to validate MTP on 200 mm wafers with controlled overlay, yield, repeatability and backend compatibility. Because each pick-and-place cycle takes roughly one minute, throughput depends on transferring many coupons in parallel. The pilot line is developing fully automated reticle-level printing as a route towards printing a full 200 mm wafer in under an hour, while respecting key manufacturing metrics such as:

- Overlay and placement accuracy across the wafer (< 500 nm 3 sigma)
- Transfer yield
- Throughput per print cycle and per 200mm wafer
- Repeatability across print cycles and wafers

In addition, the workflow includes backend steps such as passivation and metallization to connect with the silicon photonics platform, allowing evaluation of device performance after full processing, not just after placement. Initial results indicate that MTP can

➤ Figure 3: (a) Microscope image of an evanescently coupled InP UTC photodiode at C-band. (b) Frequency response of the device shown in (a) with a $2 \times 12 \mu\text{m}^2$ active region, showing bandwidth suitable for very high-speed photonic links.



be operated as part of a wafer-scale process flow, rather than only as a sequence of isolated device demonstrations. Reliability evaluation is still ongoing, but no fundamental process-related showstoppers have been identified so far.


Device-level proof points

Pilot-line validation only matters if the printed devices still perform as intended after transfer and backend processing. TRANSVERSE therefore evaluates both placement-level metrics and device-level results, including high-speed modulation and photodiode uniformity.

Demonstrated: High-speed TFLN modulators

TFLN modulators fabricated on their native wafers are transfer-printed onto silicon and silicon nitride circuits. TFLN is attractive because it combines low optical loss with very fast electro-optic modulation, but it is difficult to integrate directly into standard silicon photonics process flows.

In this work, more than 600 pre-patterned coupons, incorporating TFLN waveguides, were printed with a transfer yield above 95% across multiple wafers. Here the yield is defined as the number of



➤ Figure 4: Examples of transfer-printed III-V amplifiers integrated on silicon photonic wafers, with representative optical characteristics.



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



➤ Figure 5: Estimated MTP throughput as a function of the number of coupons transferred per print cycle, assuming a print cycle time of approximately one minute.

successfully printed coupons without any breakage. Measurements after integration show:

- V_{π-L} of 2.8 V/cm
- Electro-optic bandwidth exceeding 70 GHz
- Low insertion loss of <2 dB

These values are consistent with the performance of the native devices. Results are obtained across multiple printed devices, indicating that alignment, transfer and backend processing do not introduce significant degradation.

Demonstrated: Photodiode uniformity

MTP of uni-traveling carrier (UTC) photodiodes (PDs) provides a second, complementary validation, implying that very small coupons less than 100 μm in size can be printed. The printed devices resulted in a waveguide-coupled photodiode with a responsivity of 0.3 A/W at 1550 nm with bandwidths going to 155 GHz at -1 V bias. The dark and saturation currents of this device are 10 nA and 4.5 mA, respectively.

Following this proof-of-concept demonstration,

wafer-scale photodiode integration show uniform dark currents within and across wafers. These dark currents are also on par with the dark currents on the source wafer indicating minimal process-induced degradation. The current results are encouraging print yields of 99.9% and device yields of 97% across 500 photodiodes on a single target wafer show that the process holds up at scale. The photodiode results show that high yield is achievable; the next challenge is to combine this with higher throughput through fully automated reticle-level printing.

Ongoing: III–V lasers, semiconductor optical amplifiers and electro-absorption modulators

III–V lasers and semiconductor optical amplifiers have already been transfer-printed and shown to remain functional after integration. However, their full pilot-line evaluation, including yield statistics, array-level uniformity and reliability, is still ongoing. In parallel, the STARLIGHT project will evaluate micro-transfer printing of InP electro-absorption modulators.

What still needs to mature

While the current results demonstrate feasibility at pilot-line level, several aspects still need further development before production deployment. These include:

- Increasing throughput through reticle-level printing
- Co-design source and target for array printing and placement optimization
- Tighter integration with backend packaging processes
- Extended reliability testing and failure analysis under relevant operating conditions


In addition, broader validation across device types, particularly lasers and amplifiers, is still in progress.

Conclusion

Heterogeneous photonic integration is becoming a packaging challenge as systems combine low-loss routing, high-speed modulation, detection and light generation from different material platforms. Micro-transfer printing addresses this challenge by allowing devices to be fabricated on their optimal substrates, tested before integration and printed only where they are needed on a target wafer.

The TRANSVERSE pilot line shows how MTP can move beyond individual demonstrations toward controlled 200 mm wafer-scale validation. High-speed TFLN modulators and wafer-scale photodiode integration provide concrete proof points, including preserved device performance and encouraging yield data.

The next step is to combine these results with higher throughput, standardized layouts and extended reliability qualification, so that MTP can mature from a validated integration method into a production-ready process.





AP International 2026 reveals the future of semiconductor packaging

The inaugural AP International conference in Brussels, chaired by Peter O'Brien from Tyndall and Tyler Myers from Forge Nano, brought together packaging engineers, decision makers equipment suppliers, materials specialists, foundries, OSATs and researchers to examine how chiplets, heterogeneous integration, thermal management and manufacturing scale are reshaping semiconductor systems.

BY SARAB CHOPRA, EDITOR, ADVANCED PACKAGING MAGAZINE

As AI infrastructure scales and semiconductor architectures become increasingly heterogeneous, advanced packaging is rapidly emerging as one of the industry's most strategically important enabling technologies. Discussions throughout AP International 2026 demonstrated how packaging is no longer viewed as a secondary assembly process, but as a central determinant of system performance, manufacturability and long-term reliability.

The growing adoption of chiplet-based architectures, co-packaged optics, 2.5D and 3D integration, fan-out wafer-level packaging and hybrid bonding is fundamentally reshaping manufacturing requirements across the semiconductor ecosystem. These architectures

promise substantial gains in bandwidth density, power efficiency and system-level integration, but they also introduce unprecedented assembly and reliability challenges.

A recurring message from speakers was that advanced packaging is increasingly inheriting the precision requirements traditionally associated with front-end semiconductor fabrication. Overlay tolerances are shrinking from microns to nanometres, while package architectures are becoming far more thermally and mechanically complex. The result is that packaging engineers are now confronting problems that span optics, materials science, motion control, thermal engineering and reliability physics simultaneously.

Heterogeneous integration drives new alignment demands

One of the strongest themes at the conference was the growing importance of heterogeneous integration. Combining logic dies, memory, photonic integrated circuits, ASICs and advanced substrates into tightly integrated packages is allowing semiconductor companies to overcome some of the scaling limitations associated with traditional monolithic device architectures.

However, this transition is dramatically increasing assembly complexity.

Presentations from PI (Physik Instrumente) highlighted how modern heterogeneous systems require sub-

micron alignment between optical, electrical and mechanical interfaces. In many advanced packaging configurations, even extremely small positioning errors can significantly degrade performance.

This becomes especially critical in silicon photonics, where optical waveguides may only measure 150–200nm in width. At these dimensions, sub-micron positional deviation can produce insertion losses that exceed device operating margins. The challenge becomes even greater when coupling multi-channel fibre arrays, where all optical paths must be optimised simultaneously while compensating for vibration, thermal drift and process variation.

Speakers repeatedly stressed that conventional passive alignment approaches are becoming insufficient for next-generation packaging flows. Instead, the industry is moving toward active alignment systems capable of continuously adjusting positioning in real time using optical or electrical feedback signals.

These systems increasingly rely on six-degree-of-freedom positioning platforms capable of controlling X, Y and Z translation as well as rotational correction. The ability to dynamically compensate for mechanical movement, thermal expansion and curing-induced stress is becoming essential for maintaining coupling efficiency and production yield.

Packaging evolves into a precision engineering discipline

The conference repeatedly reinforced the idea that packaging is evolving into a front-end-like precision manufacturing discipline.


In traditional semiconductor assembly, packaging tolerances were comparatively forgiving relative to lithography and wafer fabrication processes. That distinction is rapidly disappearing.

Hybrid bonding and wafer-level integration are pushing overlay requirements into the nanometre regime. Wafer-to-wafer bonding processes now target overlay accuracy below 200nm across full 300mm wafers, while die-to-wafer assembly systems are being expected to achieve throughput exceeding 1,000 units per hour without compromising placement precision.

This transition is forcing manufacturers to rethink motion control, metrology and assembly architectures.

Presenters described how advanced alignment platforms increasingly combine long-range servo stages with piezoelectric nanopositioners capable of sub-nanometre resolution. Closed-loop capacitive sensors are being integrated directly into assembly systems to continuously monitor and correct drift during operation.

Flexure-guided nanopositioners, parallel-kinematic hexapods and hybrid coarse/fine positioning architectures




are becoming increasingly important because they eliminate backlash, minimise friction and improve long-term repeatability. Some systems now integrate firmware-level optimisation algorithms capable of autonomously identifying optimal coupling positions in milliseconds.

The increasing sophistication of these positioning systems reflects the growing difficulty of maintaining process stability at advanced packaging scales.

Hybrid bonding introduces new manufacturing challenges

Hybrid bonding emerged as one of the most technically demanding areas discussed during the event.

As semiconductor manufacturers pursue higher interconnect density and lower power consumption, hybrid bonding



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



Decoding TSMC's COUPE Technology



is becoming increasingly attractive for next-generation chiplet integration. However, these processes introduce substantial assembly complexity.

Presenters explained that thermal compression bonding generates thermo-mechanical drift during assembly, requiring active compensation throughout the bonding cycle. High-density bump arrays with pitches below 40µm leave effectively no tolerance for placement error, while wafer bow and local substrate warpage further complicate alignment.

The challenge becomes even greater in multi-die and multi-material systems. Differences in coefficient of thermal expansion between silicon, glass, organic substrates and photonic materials can generate substantial stress during thermal cycling. These stresses can distort package geometry, reduce co-planarity and compromise optical or electrical interconnect integrity.

Maintaining flatness across the package substrate therefore becomes critically important. Even relatively small warpage can prevent precise optical alignment or generate stress concentrations capable of damaging fine-pitch interconnects.

To address these issues, manufacturers are increasingly exploring low-CTE materials, advanced underfill chemistries, stiffener-ring architectures and highly controlled adhesive deposition techniques. However, presenters acknowledged that achieving acceptable yield and repeatability remains difficult as package complexity continues increasing.

Co-packaged optics increases integration density

Co-packaged optics (CPO) was another dominant topic throughout the conference, particularly as AI systems continue demanding higher bandwidth density and lower interconnect power consumption.

By placing photonic components directly alongside switching ASICs and accelerators, co-packaged optics architectures aim to overcome the bandwidth and power limitations associated with traditional pluggable optical transceivers. However, this approach substantially increases manufacturing and reliability complexity.

Presenters described how co-packaged optics systems integrate ASICs, photonic integrated circuits, fibre arrays, silicon bridges, TSV interposers and advanced substrates into highly compact assemblies requiring extreme positional accuracy.

In many cases, fibre-array placement accuracy below 0.2µm is required to maintain acceptable coupling performance. Multiple flip-chip bonding stages, high-density interconnect structures and multi-material integration further increase the risk of assembly-induced defects and reliability degradation.

TSMC's COUPE architecture was discussed as an example of the industry's broader effort to bring optical connectivity physically closer to compute resources. However, speakers emphasised that these architectures can only become commercially scalable if alignment automation, thermal management and manufacturing

repeatability improve significantly. The conference made clear that co-packaged optics is no longer simply a research topic. Instead, it is rapidly becoming a practical engineering challenge tied directly to future AI infrastructure deployment.

Thermal management becomes a limiting factor

Thermal management emerged as one of the most pressing constraints facing advanced packaging technologies. The convergence of photonic devices with increasingly power-dense AI processors is creating extremely challenging thermal environments.

While modern ASICs generate substantial heat during operation, optical components remain highly temperature-sensitive and can experience performance degradation even under relatively modest thermal variation.

This creates conflicting design requirements within tightly integrated packages. Presenters highlighted how 2.5D and 3D integration architectures further complicate thermal behaviour by increasing power density and restricting heat dissipation pathways. Vertically stacked dies, TSV structures and dense interconnect layers can all contribute to localised thermal hotspots.

As a result, conventional air-cooling strategies are increasingly viewed as insufficient for future high-performance systems. Several advanced cooling approaches were discussed during the conference, including liquid cold-plate cooling, integrated heat spreaders and high-performance thermal interface materials incorporating graphite and boron nitride. These solutions aim to improve heat extraction while minimising thermal gradients across sensitive photonic-electronic assemblies.

However, speakers noted that thermal management is no longer simply a cooling problem. Thermal behaviour now directly affects alignment stability, package warpage, material stress and long-term reliability, making it deeply interconnected with nearly every aspect of advanced package design.

Reliability concerns grow for harsh-environment applications

Beyond manufacturability and

throughput, reliability remained a major area of concern throughout the conference. Presentations from Tyndall National Institute focused heavily on the environmental stresses faced by advanced packages operating in harsh conditions such as aerospace and defence applications.

Space-qualified systems may experience severe thermal cycling, radiation exposure, vacuum operation, mechanical shock and high-frequency vibration simultaneously. These environments create major challenges for both electronic and photonic assemblies.

Researchers outlined extensive qualification methodologies involving thermal shock, temperature cycling, humidity exposure, vibration analysis, vacuum testing and radiation evaluation.

Accelerated stress testing is increasingly being used to identify potential failure mechanisms before deployment.

Particular attention was given to the evaluation of automotive-grade commercial off-the-shelf (COTS) components for possible space applications. Automotive-qualified devices offer attractive cost and availability advantages compared with traditional space-grade components, and early testing results showed encouraging reliability performance across several evaluated devices.

However, photonic systems remain significantly more difficult to qualify for harsh environments.

Photonic packaging introduces unique reliability risks

Unlike conventional electronic assemblies, photonic devices are highly sensitive to positional stability and environmental contamination.

Conference presentations highlighted several reliability risks associated with photonic packaging, including cracked optical elements, solder-joint degradation and vibration-induced fibre damage following environmental stress testing.

Hermetic sealing remains one of the industry's most persistent challenges. Many photonic modules contain organic materials that can outgas under vacuum conditions, potentially contaminating optical surfaces and degrading long-term performance.

At the same time, delicate optical structures are vulnerable to thermo-mechanical stress caused by coefficient-of-thermal-expansion mismatch between package materials. Even relatively minor dimensional changes can disrupt optical alignment.

Presenters also noted that photonic systems frequently contain brittle glass structures and precision fibre assemblies that are especially susceptible to vibration and shock damage. Maintaining alignment stability over long operational lifetimes therefore remains a major qualification challenge for photonic packaging technologies.

As optical integration density increases, these reliability concerns are expected to become even more difficult to manage.

Economics and scalability remain major barriers

While much of the conference focused on technical capability, production economics emerged as an equally important discussion point.

Several presentations emphasised that alignment processes are becoming one of the largest contributors to photonic manufacturing cost. In some photonic assembly flows, alignment alone may account for as much as 80% of total device cost.

This creates major scalability concerns as AI and optical interconnect deployments continue accelerating. Traditional serial alignment methods, which may require several minutes for individual coupling operations, are incompatible with future high-volume manufacturing requirements. Reducing "cost per placement" is therefore becoming a major strategic priority across advanced packaging

supply chains. Automation is increasingly viewed as the only viable solution. Real-time feedback systems, autonomous optimisation algorithms and high-speed multi-axis alignment platforms are being deployed to reduce manual intervention while improving throughput consistency.

The industry's broader challenge is no longer simply achieving high-performance packaging in laboratory environments. It is achieving that performance repeatedly, reliably and economically at manufacturing scale.


Packaging becomes central to future semiconductor roadmaps

A clear conclusion emerging from AP International 2026 was that advanced packaging is no longer supporting semiconductor innovation from the sidelines, it is actively driving it.

The rise of AI accelerators, optical interconnects, chiplet architectures and heterogeneous integration is fundamentally changing the relationship between packaging and system performance. Precision alignment, thermal engineering, reliability qualification and manufacturing automation are now central to the semiconductor roadmap itself.

As package architectures continue increasing in density and complexity, the ability to simultaneously manage nanoscale precision, thermal stability, manufacturability and long-term reliability will become one of the industry's defining competitive differentiators.

The discussions at AP International 2026 demonstrated that the future of semiconductor scaling may depend just as much on packaging innovation as on advances in transistor technology itself.



Advanced packaging moves to the center of system innovation

At ECTC 2026, advances in heterogeneous integration, hybrid bonding, advanced substrates, photonic packaging, and thermal management underscored a fundamental shift in semiconductor innovation, one where system performance is increasingly defined by packaging and integration technologies rather than transistor scaling alone.

BY SARAB CHOPRA, EDITOR, ADVANCED PACKAGING MAGAZINE

Researchers, engineers, and industry stakeholders across the semiconductor ecosystem, with technical sessions spanning advanced interconnects, substrate technologies, thermal management, power delivery, photonics integration, and reliability engineering. This year's program reflected the continued acceleration of system complexity driven largely by artificial intelligence (AI), high-performance computing (HPC), and data-centric applications.

A consistent theme across the technical program was the increasing role of

advanced packaging as a system enabler rather than a downstream integration step. As transistor scaling faces increasing physical and economic limitations, innovation is shifting toward package-level architectures that integrate multiple dies, functions, and materials into unified systems.

Heterogeneous integration and AI-driven requirements

Heterogeneous integration remained central to many presentations at ECTC 2026, reflecting the industry's response to growing AI workload demands. Increasing memory bandwidth

requirements, rising data movement energy costs, and thermal constraints are driving the adoption of multi-die architectures combining logic, memory, analog, and emerging photonic components within single packages.

A representative contribution in this area explored fan-out panel-level packaging (FOPLP) approaches for HPC and AI applications. The work demonstrated a 600 mm panel interposer flow subsequently partitioned into 300 mm sub-panels, enabling compatibility with existing assembly infrastructure while improving



manufacturing throughput. The study highlighted key process challenges including warpage control, layer-to-layer alignment, and uniformity in redistribution layer (RDL) fabrication at large panel scales. The results indicated that panel-level integration can support higher interconnect density and improved cost efficiency, while maintaining reliability through optimized process partitioning and laser-based singulation techniques.

Across multiple sessions, it was evident that system design methodologies are also evolving, with package-level considerations increasingly incorporated earlier in the design cycle through closer collaboration between chip designers, package architects, and system engineers.

Hybrid bonding advances toward manufacturing readiness

Hybrid bonding continued to be a major focus area, with discussions shifting further toward manufacturing implementation rather than feasibility.

Recent developments highlighted progress in alignment accuracy, defect control, surface preparation, and metrology techniques. This reflects the technology's transition from research demonstration toward early-stage high-volume manufacturing readiness.

A related study on chip-to-wafer hybrid bonding with inter-die gap fill integration (IDGF) demonstrated a process flow enabling 3D stacking while managing critical reliability risks. The work focused on void formation, wafer thinning challenges, and delamination mechanisms during multi-stack integration. Simulation and experimental validation showed that optimized gap-fill materials and process sequencing can significantly improve structural stability while enabling higher stacking densities.

These developments are particularly relevant for AI-driven architectures where bandwidth density and vertical integration are increasingly critical.

Power delivery and thermal constraints

Power delivery was widely discussed as an increasingly critical constraint in advanced system design. AI and HPC devices are now operating at



power densities that place significant demands on package-level electrical infrastructure.

Sessions addressing power distribution networks, voltage regulation strategies, and low-impedance package designs emphasized the importance of minimizing resistive losses and improving current delivery efficiency. Power delivery is increasingly being treated as a first-order design parameter rather than an afterthought.

Thermal management is similarly becoming a key limiting factor. With rising heat flux in advanced devices, conference presentations highlighted continued progress in liquid cooling approaches, advanced thermal interface materials, vapor chambers, and embedded cooling techniques. Importantly, thermal considerations are now being integrated into early-stage architectural planning alongside electrical and mechanical design.

Substrate innovation and glass-based approaches

Substrate technology remains a critical area of innovation as package sizes grow and interconnect density increases. Traditional organic substrates are facing challenges related to warpage control, dimensional stability, and routing density.

Glass substrates were a notable area of interest at ECTC 2026, with multiple presentations focusing on manufacturing techniques, through-glass vias, metallization processes,

and reliability characterization. While still emerging, glass-based substrates are increasingly viewed as a potential solution for next-generation high-density packaging platforms due to their mechanical stability and electrical performance advantages.

Photonic integration and co-packaged optics

Another key theme at ECTC 2026 was the growing role of photonic integration in addressing data movement bottlenecks. As system bandwidth requirements increase, electrical interconnects face limitations in both energy efficiency and scaling potential.

Co-packaged optics and photonic integration technologies were discussed with increasing emphasis on manufacturability, reliability, and system integration challenges. The focus is shifting from conceptual demonstrations toward practical deployment pathways, particularly in data center and AI infrastructure applications.

Manufacturing, reliability, and modeling

Across the conference, there was strong emphasis on manufacturability and reliability as enabling factors for emerging technologies. Topics included warpage modeling, defect detection, thermomechanical reliability, and advanced material characterization.

One fault isolation study in advanced fan-out RDL interconnect systems highlighted the growing difficulty of identifying nanometer-scale defects in dense redistribution networks. The

work demonstrated that conventional inspection methods, such as 3D X-ray and focused ion beam analysis, face significant limitations in large-scale interconnect architectures. The proposed fault localization approach improves failure identification efficiency, supporting faster diagnosis in high-density AI packages.

Predictive modeling techniques, including multiphysics simulation and digital twin approaches, are becoming increasingly important in bridging the gap between design and production. These tools are enabling earlier identification of potential failure mechanisms and improving development efficiency for complex package architectures.

Conclusion

ECTC 2026 highlighted a semiconductor industry undergoing a clear structural shift toward system-level innovation. As traditional scaling approaches face increasing constraints, advanced packaging is becoming



a central platform for integrating heterogeneous technologies into high-performance systems.

Technologies such as hybrid bonding, advanced substrates, co-packaged optics, panel-level fan-out integration, and integrated power and thermal solutions are no longer peripheral areas of research. They are now central to

enabling the next generation of AI, HPC, and data-intensive applications.

The overall direction emerging from the conference is clear: future semiconductor performance gains will depend not only on transistor innovation, but increasingly on how effectively diverse technologies are integrated at the package and system level.



**The 2027 IEEE 77th
Electronic Components
and Technology Conference**

As the premier event in the semiconductor packaging industry, ECTC addresses new developments, trends and applications for fan-out & fan-in packages, 3D & 2.5D integration, TSV, WLP, flip-chip, photonics, LEDs, materials and other integrated systems packaging topics.

June 1 – 4, 2027
Gaylord Rockies Resort & Convention Center
Denver, Colorado, USA




CALL FOR PAPERS OPENS AUGUST 17!

Abstract submissions for the 77th ECTC are due by
October 5, 2026. To submit, visit: www.ectc.net

We welcome previously unpublished, non-commercial abstracts in areas including, but not limited to:

Applied Reliability	Materials & Processing
Assembly and Manufacturing Technology	Thermal/Mechanical Simulation & Characterization
Emerging Technologies	Packaging Technologies
Electrical Design & Analysis	Photonics
Interconnections	Interactive Presentations





Building standards for the next era of advanced packaging

AP Magazine Editor Sarab Chopra speaks with **Peter O'Brien, Head of Photonics Packaging at Tyndall National Institute**, about how standardisation, chiplets and photonic-electronic integration are shaping the future of advanced packaging.

AS SEMICONDUCTOR architectures become increasingly complex, advanced packaging is emerging as a critical technology enabling the next generation of high-performance computing, AI infrastructure and photonic-electronic integration. From chiplets and advanced interposers to thermal management and digital modelling, the industry is moving towards highly integrated solutions that require new approaches to design, manufacturing and collaboration.

Advanced Packaging Magazine Editor Sarab Chopra, spoke with Peter O'Brien, Head of Photonics Packaging at Tyndall National Institute, to discuss Tyndall's role in advancing semiconductor packaging, the importance of standardisation, how pilot lines are helping bridge the gap between research and manufacturing, and the technologies shaping the future of advanced packaging.

SC: For readers who may be unfamiliar with Tyndall, could you introduce the organisation and explain your role in advanced packaging?

POB: Tyndall is based in Cork, Ireland, and is primarily a semiconductor research centre. We work across wafer fabrication, semiconductor design, packaging, reliability testing and application development.

My area is advanced packaging, where we take semiconductor chips and develop packaging solutions for different applications. We also have significant reliability testing capabilities, allowing us to test packaged devices under different conditions and perform failure analysis.

Beyond packaging, we have application teams working across areas such as telecommunications and biomedical technologies.

A key part of Tyndall's model is our close relationship with industry. We have several industry partners physically located on-site, including packaging equipment companies and technology providers. We also collaborate closely with major technology companies such as Meta and Intel.

The objective is to create an environment where research and industry can work together closely to accelerate innovation in semiconductor technologies.

SC: What are the current priorities for Tyndall in advanced packaging?

POB: One of our biggest priorities is standardisation.

When we established the packaging group around 15 years ago, we quickly recognised a significant opportunity

because many areas, particularly photonics, were highly customised.

Integrated photonics, whether based on silicon photonics or indium phosphide technologies, has traditionally been very bespoke. The customer often defines the entire ecosystem around their specific application requirements.

That flexibility is useful during early-stage technology development, but it becomes a challenge when you want to scale. Scaling requires standards, similar to what has enabled the semiconductor industry to grow.

Our focus has been moving towards more standardised approaches in design, fabrication and packaging.

One example is the development of packaging building blocks, which we refer to as Assembly Design Kits, or ADKs. They are similar in concept to Process Design Kits, or PDKs, used in semiconductor foundries. The objective is to create validated and repeatable processes that address the majority of users while still allowing customisation when required.

By following standards, companies can reduce costs, shorten development timelines and enter the ecosystem more easily.

SC: How important is the convergence of photonics and electronics for the future of advanced packaging?

POB: This is one of the most important developments happening today.

Historically, photonics has often been treated as a separate technology domain with its own ecosystem. However, integrated photonics increasingly needs to become part of the wider microelectronics landscape.

By aligning photonics with electronics, we can benefit from the scale of the semiconductor industry, including manufacturing infrastructure, equipment, workforce expertise and established supply chains.

We are already seeing major semiconductor foundries and technology companies incorporating photonics into their roadmaps. The industry is moving towards a future

where photonics technologies are manufactured and integrated within semiconductor environments.

Photonics will not replace electronics, but it provides significant advantages in areas such as bandwidth, performance and energy efficiency.

The challenge now is identifying the most important technologies, standardising them and integrating them into the semiconductor ecosystem so they can scale effectively.

SC: Scaling technology from research into manufacturing remains a major challenge. How does Tyndall help bridge that gap?

POB: Pilot lines are designed specifically to bridge the gap between research and industrialisation.

A technology may be successfully demonstrated in a research environment, but commercialising it requires a much broader ecosystem. Companies, particularly SMEs, often struggle to access the infrastructure, expertise and supply chains needed to scale.

Historically, pilot lines focused on demonstrating that hundreds or thousands of components could be manufactured. However, what we have learned is that a successful pilot line needs to do more than manufacture parts, it needs to build an ecosystem.

That means bringing together material suppliers, equipment providers and industrial partners.

For example, if we are developing new substrates, the objective is not necessarily for Tyndall to manufacture those substrates ourselves. Instead, we want industry partners to become part of the supply chain and help establish a scalable manufacturing route.

The same applies to equipment. Packaging equipment is often highly customised, but more standardised manufacturing tools will be essential as the industry grows.

Ultimately, scaling is not only about producing larger volumes. It is about creating the environment that allows industry to scale successfully.


SC: How does Tyndall collaborate with industry partners, and what support do companies typically seek?

POB: The first step is engagement. Semiconductor technology is a global industry, so we spend significant time travelling, meeting companies and understanding their challenges.

A major part of our role is communicating the importance of ecosystem development, standardisation and supply chain creation.

When companies approach us, they often have a product concept but may not have deep expertise in areas such as packaging. They may ask: How should we approach this design? Which manufacturing route should we follow? Which suppliers should we work with?

Our role is not simply to package devices for them. Instead, we guide



them towards established processes and connect them with the right industrial partners.

Training is also a very important part of our work. Many companies need their workforce to understand advanced packaging technologies, even if they ultimately outsource manufacturing to an OSAT.

They need engineers who understand the technology well enough to evaluate suppliers, manage manufacturing relationships and ensure requirements are being met.

SC: What trends do you see as most transformative in advanced packaging?

POB: There are several major trends shaping the future of advanced packaging.

The first is the continued integration of photonics and electronics. Technologies such as co-packaged optics require photonics manufacturing processes to become compatible with semiconductor manufacturing environments.

The goal is not to create separate photonics factories, but to enable photonics technologies to be manufactured within existing semiconductor ecosystems.

Another major trend is the rise of chiplet architectures.

Instead of creating one very large chip containing every function, systems are increasingly being divided into smaller chiplets manufactured using different processes and then integrated together.

This requires advanced interconnect technologies and increasingly sophisticated interposers. Materials such as silicon, glass and hybrid substrates are becoming increasingly important.

Thermal management is another major challenge. As package densities increase, managing heat becomes essential for maintaining performance and reliability.

Heat affects electrical performance, optical alignment, mechanical stability and long-term reliability. Developing new thermal materials and cooling approaches will therefore be critical.

SC: How is Tyndall addressing energy efficiency and sustainability in advanced packaging?

POB: One important area is the development of digital twins for semiconductor packages.

A digital twin combines the physical characteristics of a real-world package into a comprehensive model. For advanced packaging, this means modelling multiple physical domains simultaneously, including thermal, optical, electrical and mechanical behaviour.

These domains are highly interconnected. A change in temperature can affect mechanical expansion, optical alignment and electrical performance.

Using multi-physics modelling, we can predict how changes in materials or designs will influence the complete system before manufacturing begins.

This is particularly valuable when companies consider alternative materials or more sustainable substrate options. Instead of relying on trial and error, they can simulate the impact of different choices.

With advances in AI and machine learning, these models are becoming increasingly powerful, allowing companies to predict both technical performance and manufacturing outcomes.

Cost modelling is another important development. Companies need to understand not only whether a technology works, but whether it can be manufactured economically at scale.

By modelling different manufacturing routes, companies can make better decisions before committing significant investment.

SC: What future breakthroughs or milestones in advanced packaging are you most excited about?

POB: One exciting area is the move towards surface-level optical packaging.

Traditionally, optical fibres connect at the edge of a chip in a one-dimensional arrangement. However, future systems

will require optical connections on the top surface of chips, similar to how electrical connections are already arranged.

Moving to two-dimensional optical arrays will significantly increase optical I/O density, which is essential for future AI systems requiring enormous data movement.

Another important development is advanced interposer technology.

Future chiplet-based architectures will require thousands of high-density interconnects between different semiconductor devices.

Glass interposers are particularly interesting because they can potentially use large panel processing techniques, similar to those developed for flat-panel display manufacturing.

This could enable very large, high-volume packaging solutions for applications such as AI data centres, where package sizes continue to increase.

SC: What is the long-term vision for Tyndall's role in advanced packaging?

POB: The long-term goal is to help create a scalable and globally competitive ecosystem for advanced packaging.

That means developing standardised technologies, connecting research with industry, supporting supply chain development and ensuring companies have access to the skills and infrastructure required.

Advanced packaging is becoming central to the future of semiconductor technology. By bringing together photonics, electronics, materials, manufacturing and modelling, we can enable the next generation of high-performance computing systems.

Advanced packaging is becoming central to the future of semiconductor technology.

Unlocking yield improvements in advanced packaging through materials-driven failure analysis



Materials-driven failure analysis offers a systematic, multidisciplinary framework for understanding why defects occur, how they propagate, and what conditions exacerbate them.

BY DR. PRADYUMNA (PRADY) GUPTA, FOUNDER & CHIEF SCIENTIST, INFINITA LAB | FOUNDER & CEO, INFINITA MATERIALS

THE SEMICONDUCTOR INDUSTRY is at a pivotal moment. As digital transformation accelerates, driven by artificial intelligence, edge computing, 5G/6G connectivity, and electrification, advanced packaging has moved from being an incremental enhancement to a strategic imperative. Today's heterogeneous systems, ranging from 2.5D and 3D ICs to fan-out wafer-level packaging (FOWLP), deliver unprecedented performance density, improved power efficiency, and reduced form factors. Yet these benefits come with an unavoidable challenge: Yield pressure.

Yield is no longer merely about defect counts; it's about understanding how materials behave across scales, interfaces, and processes. Traditional defect inspection methods, while still valuable, are no longer sufficient on their own. The key to unlocking sustainable yield improvements lies in materials-driven failure analysis: a multidisciplinary approach that links material characteristics, structural integrity, process interactions, and comprehensive failure mechanisms into a unified diagnostic framework.

This article examines how materials science, advanced characterization, cross-functional collaboration, and

predictive analytics converge to identify hidden failure drivers, enhance process accuracy, and deliver measurable gains in advanced packaging yield.

Why yield matters more than ever in advanced packaging

Advanced packaging technologies, including FOWLP, 3D TSV stacks, chiplets, and integrated passive devices, are increasingly used to meet the demand for higher performance and smaller footprints. These approaches promise reduced interconnect length, improved signal integrity, and enhanced thermal performance. But they also introduce complex materials interfaces, thermal stresses, and multi-physics interactions that are highly sensitive to process variability.

In traditional packaging, yield problems could often be traced to specific process steps. In advanced packaging, yield excursions can originate from subtle variations in material properties, interfacial degradation, or microscale stress accumulation that may not be readily observable with optical or electrical inspection alone.

The economic implications are significant. Yield improvements directly influence usable wafer output, which in turn affects revenue per wafer, cost

of goods sold (COGS), and profitability. Even small incremental yield gains in high-value advanced packages can generate substantial financial returns. This is where materials-driven failure analysis becomes indispensable: it transforms yield investigation from reactive troubleshooting into a proactive engineering discipline.

What materials-driven failure analysis actually means

At its core, materials-driven failure analysis is the systematic investigation of failure modes through the lens of materials behaviour and interactions. Rather than simply identifying that a failure occurred, this approach seeks to understand why it occurred at a fundamental level:

- What specific material property variations contributed to the defect?
- How did interfacial stress accumulate under processing conditions?
- Did thermal expansion mismatch cause microcracking or delamination?
- Were contamination or unintended residue responsible for electrical anomalies?

In advanced packaging, these questions often revolve around a handful of recurring themes:

- **Interfacial integrity:** Particularly

between dielectrics, metal interconnects, solder joints, and underfill materials.

- **Thermo-mechanical stress:** From thermal cycling during reflow and operation, especially in stacked components.
- **Material anisotropy:** Variations in material properties depending on orientation or deposition method.
- **Environmental susceptibility:** Moisture ingress and electromigration in dense interconnect networks.

Materials-driven analysis involves correlating microscale material behavior with macroscale system performance. It assumes that to improve yields at scale, engineers must first understand the underlying multiscale physics.

Bringing the right tools to the investigation

To implement materials-driven failure analysis, organisations must expand beyond conventional electrical tests and optical inspection. Modern failure analysis blends multiple characterization techniques:

- **Electron Microscopy and X-ray Imaging**
Scanning electron microscopy (SEM), focused ion beam (FIB) cross-sectioning, and X-ray computed tomography (XCT) enable visualization of sub-micron defects, voids in solder joints, delamination, and dielectric cracking. These tools reveal structural anomalies that are invisible in traditional 2D microscopes.
- **Nanoindentation and Mechanical Property Mapping**
These methods provide insights into local stiffness, hardness, and modulus variations, particularly at interfaces where materials with differing elastic properties meet. Mapping these mechanical gradients helps identify stress concentration zones prone to failure.
- **Thermal Imaging and Infrared Thermography**
Because advanced packages often suffer from localized heating, thermal imaging helps pinpoint hotspots and thermal patterns that correlate with material degradation or poor thermal conduction paths.
- **Elemental and Phase Analysis**
Techniques such as energy-dispersive X-ray spectroscopy (EDS) and electron backscatter diffraction (EBSD) reveal

compositional variations, grain orientation, phase transformations, and unintended inclusions all of which influence reliability.

- **In-situ Stress/Strain Measurement**
Advanced optical and X-ray-based methods enable stress mapping during thermal cycling or mechanical loading, providing a direct correlation between process conditions and material responses.

By combining these tools into an integrated workflow, engineers can transform raw data into actionable insight.

Case study: Interfacial delamination in a 3D package

Consider a common failure mode of advanced packages: interfacial delamination between an organic substrate and a redistributed layer (RDL). On the surface, delamination appears as an area of poor adhesion detectable through acoustic microscopy. But mixed materials analysis often reveals the root cause lies in:

- A mismatch in the coefficient of thermal expansion (CTE) between layers
- Micro-scale voids generated during lamination

- Localised moisture trap points exacerbate thermal stress

Materials-driven analysis might proceed as follows:

- Thermo-mechanical modelling identifies zones of highest mechanical stress under reflow and service conditions.
- Nanoindentation and modulus mapping detect unexpected stiffness gradients near the interface.
- XCT imaging reveals micro-void clusters that act as stress concentrators.
- EDS analysis uncovers contamination or residue at the interface that reduced adhesion strength.

With this information, engineers can adjust the substrate prep process, revise laminate materials, or modify reflow profiles to minimise stress buildup — leading to measurable yield improvements.

This kind of diagnostic precision is only possible when material behavior is the starting point, not an afterthought.

Materials databases and predictive analytics

Modern yield improvement practices increasingly incorporate materials



databases and predictive models. By cataloguing materials properties, including modulus, thermal conductivity, moisture uptake, and ageing behaviour, organizations can develop multi-factor predictive models that estimate failure probability under given process conditions.

Machine learning tools can also be employed to detect patterns in historical failure data, linking material signatures to defect types. For example, correlating shifts in specific viscoelastic properties with increased solder joint fatigue can pre-emptively flag future yield risks.

These data-driven methods reduce reliance on iterative trial-and-error and accelerate actionable insight across design, engineering, and production teams.

Collaborative failure analysis workflows

High-volume manufacturing environments benefit from integrating failure analysis into cross-functional workflows. Materials insights must be communicated upstream to design and process engineers, and downstream to production teams.

A robust feedback loop ensures that lessons from failure analysis directly influence design rules and process controls.

Increasingly, organisations are standardising failure reporting and remediation frameworks, enabling accountability and trend identification at enterprise scale.

Collaborative platforms unify data from lab analysis, field returns, reliability



testing, and statistical process control, creating a closed-loop system of continuous improvement.

Execution challenges and cultural shifts

Materials-driven failure analysis is not simply a set of tools; it is a mindset. Organisations that excel in yield improvement cultivate a culture where:

- Engineers embrace multidimensional investigation, not one-dimensional metrics
- Materials science is treated as a core competency, not a peripheral function
- Cross-disciplinary dialogue is encouraged between materials, reliability, process, and design teams

Investing in analytical infrastructure and talent is part of this shift, but so is enabling domain expertise to evolve through applied learning rather than rote troubleshooting.

Benefits beyond yield

The value of materials-driven yield improvement extends beyond defect reduction. Organisations

that understand their materials at a fundamental level benefit from:

- More predictable process windows
- Reduced time-to-market for new packages
- Higher confidence in reliability performance
- Fewer field failures and warranty costs
- Better readiness for next-gen materials like low-k dielectrics, novel underfills, and heterogeneous substrates

The approach builds resilience into the entire product lifecycle, from research to mass production to service.

Conclusion: A new imperative for advanced packaging

Advanced packaging is indispensable for meeting future compute, connectivity, and power requirements. Yet without a deliberate focus on materials behavior and failure mechanisms, yield challenges will remain a primary barrier to scalability and profitability.

Materials-driven failure analysis offers a systematic, multidisciplinary framework for understanding why defects occur, how they propagate, and what conditions exacerbate them. By leveraging advanced characterization, predictive analytics, and collaborative workflows, organisations unlock yield not as a temporary improvement but as an ongoing engineering discipline.

In an era when performance expectations are high and process margins are tight, the differentiator will not be inspection alone, but insight - insight into material behaviour, interface integrity, and how the tiniest anomalies can create the biggest impacts.

WEBINARS

Specialists with 30 year+ pedigree and in-depth knowledge in overlapping sectors



For more information contact:

Jackie Cannon **T:** 01923 690205 **E:** jackie@angelwebinar.co.uk **W:** www.angelwebinar.co.uk
T: +44 (0)2476 718 970 **E:** info@angelbc.com **W:** www.angelbc.com

Expertise: Moderators, Markets, 30 Years + Pedigree

Reach: Specialist vertical databases

Branding: Message delivery to high level influencers via various in house established magazines, websites, events and social media

Angel 
 BUSINESS COMMUNICATIONS



**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



ANGELTECH

INTERNATIONAL CONFERENCE

12-14 APRIL 2027

Sheraton Brussels Airport Hotel Belgium

THEMES ANNOUNCED

CHIPLETS AND HETEROGENEOUS INTEGRATION

How are chiplets and heterogeneous integration changing the way advanced systems are designed, manufactured and packaged for performance, yield and scalability?

SCALING AND PILOT LINES

What role will pilot lines, process scale-up and manufacturing readiness play in moving advanced packaging technologies from development into volume production?

PRECISION EQUIPMENT FOR SCALABLE ADVANCED PACKAGING

How can European automation and packaging equipment providers help deliver the precision, throughput and reliability needed for next-generation advanced packaging?

MATERIALS FOR PACKAGING PERFORMANCE

How are European and Japanese advances in substrates, underfills and connectors enabling higher-density, more reliable and more manufacturable advanced packages?

THERMAL MANAGEMENT AND RELIABILITY

How can the industry overcome rising thermal and reliability challenges as advanced packages become more powerful, compact and complex?

DESIGN AND SIMULATION





How can design and simulation tools help optimise package architectures, reduce development risk and accelerate the path from concept to production?

If you are interested in speaking at AP International 2027, contact: **James Cheriton**
james.cheriton@angelbc.com
+44 (0)1923 690210

<https://ap-international.net/>



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.



How materials will unlock the future of advanced packaging



As transistor scaling reaches its physical limits, advanced packaging has become the critical enabler of next-generation semiconductor performance, powering the shift from 2D scaling to 3D integration, chiplets, and trillion-transistor systems.

BY JON KEMP, CEO, QNITY

FOR DECADES, the story of our industry was a single word: shrink. Moore's Law made transistors smaller, and their size didn't compromise speed and power. This isn't the case anymore. We've reached the physical limits of what shrinking alone can deliver, and a new era is taking its place — one I'd sum up in a different word: stack. The industry is moving from flat, 2D chip designs to vertical, 3D architectures, and that shift puts pressure on materials, integration, and reliability that the old playbook was never built to handle.

This move from "shrink" to "shrink and stack" is forcing the entire ecosystem to think differently. Advanced packaging is the strategic enabler that makes

this shift possible — the link that lets manufacturers push transistor architecture and density while holding the line on performance, efficiency, and reliability. When architectures go vertical, you can't bolt on the materials, interconnects, and packaging at the end. They have to be designed in from the very first step.

The customer base is widening, too. More companies across the ecosystem now have a direct stake in advanced packaging and 3D integration, and OEMs are increasingly in the room for investment, material selection, and design decisions. More complexity means more demand — for innovation, for specialized talent, and for advanced

manufacturing capacity across the board.

We've seen significant transitions before. Most recently, it was the shift to gate-all-around (GAA) transistor architecture — a shift powered by a multi-year effort to develop and optimize the materials needed to meet the stringent demands of 2nm logic at high yields. When manufacturers debuted 2nm in 2025, it quickly became the most competitive process node among the world's biggest tech companies, prized for the transistor density and energy efficiency that have become critical gains for AI.

But those advances remain within the bounds of traditional scaling.

What we're facing today is different: performance is no longer defined only at the transistor. It is increasingly defined at the package.

The future of chipmaking depends on how entire systems are integrated, and meeting the sophisticated advanced packaging requirements that every new "stack" demands. That's where our role at Qnity becomes even more pivotal. Leveraging decades of expertise, we're delivering the materials solutions to the transistor and advanced packaging challenges of the angstrom era.

The angstrom era

The next evolution in transistor architecture is the complementary field-effect transistor (CFET), a more advanced 3D design that will free up space, increase density, and deliver faster performance at lower power. Imec predicts its commercial introduction around 2033.

In parallel, the world's largest logic chipmakers are already moving towards angstrom-class nodes. But the industry's ultimate target of trillion-transistor systems on a chip package by 2030 cannot be achieved through scaling alone. Packaging is pushing boundaries today, but the dual demands of trillion-transistor density and more advanced architectures call for a fundamental leap — denser connections, more sophisticated signal routing, and packaging schemes capable of coordinating an exponentially complex web of chip-to-chip communication.

First, the industry has to reckon with physics. As transistors shrink and the power density of a chip increases, quantum effects like electron leakage and heat buildup grow more significant, and they put a hard ceiling on how far miniaturization can go.

Getting to the required density will lean on techniques like 3D stacking and chiplet integration — and those bring challenges of their own: managing power, heat and EMI, while wiring together a lot more chips in a lot less space.

That makes thermal management critical. A trillion transistors on a single chip system demands enormous energy and throws off unprecedented heat. Power delivery is already a major

challenge. But thermal management may be the bigger test — because if you can't remove heat efficiently, you don't have a high-performing, reliable system. You have a limit.

Forming all of those interconnects presents the third constraint. As transistor counts climb, so does the complexity of linking logic and memory. Traditional methods of making copper-based interconnects are running into limits on speed, efficiency, and latency that require new techniques.

As our partners and customers have reckoned with these difficulties, each challenge points to the same answer: advanced packaging. Done right, it addresses design complexity, power management, communication bottlenecks, and system integration all at once.

Advanced packaging as the invisible infrastructure

As transistor counts rise, the complexity of integrating logic, memory, and specialized processors into one cohesive system grows exponentially. But there's a deeper problem underneath the density problem, and it comes down to communication.


Here's something that doesn't get enough attention. Even as individual chips get faster, comparatively little industry-wide investment has gone into the technologies that let chips exchange data efficiently. In AI and high-performance computing, a significant share of data center energy goes to powering the connections between chips. Those connections generate heat through

electrical resistance, require cooling, and create bottlenecks.

Consider this: in some advanced AI systems, a processor can sit idle roughly 30% of the time, simply waiting for data to arrive from memory. That idle time is fewer computations, higher costs, and performance left on the table — and at scale, the waste is staggering. The idle power burned during that 30% across a modern AI compute cluster runs on the order of 200 MWh a year — enough to fully charge 2,500 electric vehicles.

Advanced packaging attacks those inefficiencies head-on. Techniques like 3D stacking and chiplet integration raise density by combining multiple chips into a single package, and they improve communication, cut latency, and optimize power distribution. Heterogeneous integration goes further still — combining chips with specialized roles, like AI accelerators, memory modules, and high-speed controllers, into one highly tuned system. Each component gets optimized for its job, and that's what powers the workloads defining this moment: real-time language model inference, autonomous vehicle perception, complex scientific simulation, and next-generation gaming engines.

Thermal management is becoming an even more critical part of the equation, too. Advanced packaging brings in thermal interface materials (TIMs) and interposers that dissipate heat efficiently and keep chips at their peak. Paired with high-density interconnects, chip-to-chip bonding, and emerging optical connections,



these solutions keep data moving fast even across the most densely packed architectures.

Put it together, and advanced packaging resolves the core bottlenecks — communication, power, and thermal management — that stand between today's chips and tomorrow's potential. It's the invisible infrastructure that modern semiconductors are built on.


Qnity's approach to advanced packaging

Building on a legacy of more than 50 years, Qnity offers a full spectrum of technology solutions across the advanced packaging ecosystem — from wafer processing and interposer integration to IC substrates and panel-level packaging. By pairing deep expertise in semiconductor front-end and back-end processes with expertise in PCB manufacturing, we help manufacturers design fit-for-purpose systems that maximize compute efficiency, power management, and performance-per-cost.

Advanced packaging demands precision at every stage, and our materials and process technologies are built for exactly that. Planarization pads, slurries, and cleans reduce defects across complex architectures. Fine-feature photoresists enable the high-resolution patterning that next-generation IC substrates require. Metallization solutions support high-density interconnects — including through-silicon vias (TSVs) for stacked memory (HBM) and microbump interconnects — so chips communicate efficiently and reliably.

Materials are only part of it. Our process development infrastructure supports end-to-end evaluation and optimization. Test vehicle builds help identify failure modes and refine design rules, so stacked chips, interposers, and IC substrates perform reliably under real conditions. Panel-level packaging and interposer demo capabilities extend that control across wafer- and system-level integration, supporting both IC substrates and next-generation PCB fabrication.

By combining materials expertise, process knowledge, and system-level integration, Qnity's innovations let technologies like 3D stacking, chiplet



integration, and heterogeneous integration operate at scale. That's how manufacturers get past the real pain points — chip-to-chip communication, interconnect efficiency, thermal management, and high-density scaling.

The future of advanced packaging

More advanced 3D architectures, larger substrates, multi-reticle fabrication, higher transistor densities — the angstrom era creates new pressure at every level.

Advanced packaging doesn't just support scaling compute clusters and AI systems over the next five years. It lays the foundation for the even more forward-looking high-density, high-performance chips that come after.

To get ready for what's next, we strategically differentiate our R&D investment to our top 10 innovation programs, spanning advanced-node fabrication, advanced packaging, IC substrate and PCB manufacturing, along with thermal management and EMI technologies. Our longer-term bets explore the emerging materials and technologies that could define the next round of breakthroughs. By enabling customer and industry roadmaps and anticipating long-term trends, we're focused on solving the immediate technical challenges and helping shape the next generation of semiconductor systems at the same time.

Our technology solutions enable effective communication, manage thermal loads, and support heterogeneous integration — helping manufacturers realize the full

performance potential of transistor systems in the trillions, while preparing for what lies beyond 2030.

The stack era begins

There's no longer any real doubt: advanced packaging has gone from a supporting technology to a strategic enabler. It bridges the gap between what's physically possible on a chip and what's practically achievable in a complete system.

This is a turning point in how our industry approaches design, integration, and performance. We must challenge our assumptions across transistor miniaturization, interconnect efficiency, thermal management, and system-level integration.

At Qnity, our success is defined by an ecosystem-wide perspective. We bring materials expertise, process development, application engineering, and system-level integration together to take on the full spectrum of technical challenges — from wafer processing and interposers to IC substrates and panel-level packaging.

The next generation of high-performance semiconductor systems won't succeed on transistor scaling alone. It will depend on sophisticated integration, advanced materials, and forward-looking process design.

If "shrink" built the last era, "stack" will help shape the next. And we'll be there — helping the industry unlock the full potential of the chip designs that power the AI, high performance computing, and advanced connectivity breakthroughs of the coming decade and beyond.




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



Silicon photonics needs a new packaging strategy to reach volume

The schedule for silicon photonics has accelerated due to demands from AI and large data centers.

BY SYLVAIN DULPHY, FINETECH

THE JOINT DISPLAY by Nvidia and Mercedes at CES 2026 of a commercial autonomous vehicle signaled that high-volume applications of photonic components exist today. Yet, a major obstacle remains before the industry achieves the mass manufacturing scale required by the 2030 roadmap.


The primary constraint has shifted: it is no longer front-end lithography only, but the packaging process itself. Reaching the 2030 targets depends on a fundamental shift in how the industry views the “backend” process.

The value shift to fifty percent

Semiconductor packaging was traditionally viewed as a commodity, a final step handled in lower-cost regions. In standard products, this stage represents roughly 15 percent of the total value. In advanced photonics for data centers and AI, that ratio has flipped. Packaging now holds up to **50 percent of the final product value**.

Advanced processes such as copper-to-copper hybrid bonding, championed by leaders like TSMC, are erasing the traditional division between standard packaging side by side assembly and wafer fabrication. These steps are migrating from standard assembly lines into high-end cleanrooms to address new steps like die to wafer bonding, chiplets, usage of interposer. At these scales, the precision required is absolute.

“If the packaging is compromised, the entire product is lost,” notes Sylvain Dulphy, Sales



Manager at Finetech. “Regardless of the quality of the 2-nanometer chip inside, if the packaging is not executed perfectly, the unit is discarded.” Packaging has turned from a protective shell into a functional performance requirement.

Moving beyond pluggables

To support the signal speeds AI workloads demand, the sector is transitioning from external pluggable optics to co-packaged solutions. The objective is to drastically reduce the distance the signal travels within the chip, “external” signals being optical.

“If the electrical signal path is too long, it effectively becomes a filter,” explains Dulphy. “It is similar to the concept of power in physics, which requires both force and velocity. You might have speed, but if the connection length diminishes your signal strength, your global performance is capped.”


To maintain both signal integrity and speed, connections must be immediate. This necessitates the use of bumps or metal-to-metal contacts via interposers and chiplets. While this integration allows for the miniaturization needed for high-volume products—from LiDAR in cars to medical tools—it introduces a new manufacturing hurdle linked to alignment in 3D of all those elements.

With today’s technologies in semiconductor foundry, this is achievable, but now we mix optical and electrical in the same chip with mixed materials. We therefore need to make assemblies of chip that are not thermally homogeneous so not compatible with same process profiles and we need to optically align them

Solving optical alignment

For silicon photonics to reach volumes similar to mobile phone parts, the industry must fix the optical alignment issue. Today, aligning a fiber to a light source often uses **active alignment of single elements**. This involves a slow loop of moving, turning on the laser, reading the signal, and adjusting again.

This cycle lacks the speed required for semiconductor volumes. It costs too much and takes too long. Research now points to **passive alignment** methods. The goal is to use intermediate parts



designed to guide light beams without requiring powered checks for every unit. Mastering passive methods is the only way to reach the cost structures needed for mass adoption.

The supply chain gap

As we look to 2030, the supply chain readiness poses a greater obstacle than the technology itself. The sector currently lacks standards, so each player invents their own way, which limits broad growth.

Furthermore, the equipment supply is under pressure. Developing a new class of machine from the ground up takes years. We cannot wish for the tools to exist; they must develop alongside the technology. This is causing a convergence where traditional front-end semiconductor tool makers are now looking at packaging steps. Partnering with an existing packaging specialist is probably the most effective combination.

The message is plain. Hardware limits software. The limits we solve today define what future AI code can do. But to solve them by 2030, tool makers, foundries, and institutes must set their schedules now. The volume is coming; the question is if the supply chain can handle it.





Unleashing the intelligence revolution

TSMC's 2026 Europe Technology Symposium revealed how AI is reshaping the semiconductor industry, driving advances in leading-edge logic, advanced packaging, photonics and system-level integration while accelerating the path to a trillion-dollar market.

Key takeaways from TSMC's 2026 europe technology symposium

AT ITS 2026 Europe Technology Symposium, TSMC delivered a clear and emphatic message: the semiconductor industry has entered a new era defined by artificial intelligence, and the pace of change is accelerating faster than previously anticipated. What emerged from the event was not just an update on process nodes or packaging technologies, but a comprehensive vision of how silicon will underpin an "Intelligence Revolution" that extends from data centres to autonomous systems and into the physical world.

From advanced logic and 2nm-class innovations to 3D integration, photonics, and system-level design, TSMC outlined a roadmap that reflects both unprecedented demand and increasing architectural complexity. At the same time, the company provided a broader market context, revealing how AI is reshaping the economics, priorities and trajectory of the semiconductor sector.

A trillion-dollar industry arrives

One of the most striking announcements from the symposium was the revised outlook for the global semiconductor market. While earlier projections suggested that the industry would reach the US\$1 trillion milestone by 2030, TSMC now expects that threshold to be exceeded as early as 2026. Looking further ahead, the market is forecast to reach approximately US\$1.5 trillion by the end of the decade.

This accelerated growth is overwhelmingly driven by high-performance computing and AI, which together are expected to account for roughly 55% of total semiconductor demand. Smartphones remain a significant segment at around 20%, while automotive and IoT applications each contribute approximately 10%. The remaining share is distributed across other emerging and legacy applications.

What is particularly important in this context is the evolving nature of AI workloads. The industry is transitioning from a phase dominated by model

training to one increasingly defined by inference, where trained models are deployed at scale to generate outputs in real time. This shift is critical because inference workloads are far more widely distributed, spanning cloud data centres, edge devices and embedded systems.

The concept of "tokens", whether words, images or data points generated by AI systems, is central to this transformation. As token generation increases, so too does the demand for compute resources. This creates a powerful feedback loop in which higher productivity drives further investment in AI infrastructure, which in turn increases demand for advanced silicon.

The implication is clear: semiconductor growth is no longer tied solely to traditional computing cycles. Instead, it is increasingly driven by the exponential scaling characteristics of AI itself.

The evolution of transistor technology

At the core of TSMC's roadmap is the continued evolution of transistor architectures. The industry has already

moved from planar transistors to FinFETs and is now firmly transitioning to nanosheet (or gate-all-around) structures. These developments have enabled continued scaling in performance, power efficiency and density, even as physical limits become more challenging.

Looking beyond nanosheet architectures, TSMC is exploring complementary field-effect transistor (CFET) designs, which involve vertically stacking n-type and p-type transistors. This approach represents a potential pathway for extending Moore's Law by increasing transistor density without requiring further horizontal scaling.

The company has already demonstrated early progress in this area, including the development of extremely compact SRAM bit cells with footprints roughly 30% smaller than comparable nanosheet designs. Demonstrations of CFET-based ring oscillators further highlight the viability of this approach for future logic designs.

These innovations point to a future in which vertical scaling becomes increasingly important, not only at the transistor level but across entire system architectures.

The advanced logic roadmap: From N2 to A13

TSMC's advanced logic roadmap continues to push the boundaries of process technology, with the 2nm

generation and beyond forming the foundation for next-generation AI, HPC and mobile applications.

The N2 process entered volume production in the fourth quarter of 2025, marking a major milestone in the industry's transition to nanosheet transistors. Follow-on variants, including N2P and N2U, provide incremental improvements through design-technology co-optimization. N2P is expected to reach volume production in the second half of 2026, while N2U, scheduled for 2028, offers further enhancements in power efficiency, speed and density.

N2U is particularly notable as a balanced node designed to serve both AI/HPC and mobile applications. It delivers modest but meaningful gains over N2P, including improvements in speed, reductions in power consumption and incremental density scaling. Crucially, it maintains compatibility with existing design ecosystems, enabling efficient IP reuse and smoother migration paths.

Beyond the N2 family, TSMC introduced its A-series nodes, beginning with A14 and extending to A13. These nodes represent the next phase of nanosheet evolution and are designed to deliver full-node improvements in performance, power and area.

A14, scheduled for production in 2028, offers substantial gains over N2, including up to 15% speed improvement

TSMC's advanced logic roadmap continues to push the boundaries of process technology, with the 2nm generation and beyond forming the foundation for next-generation AI, HPC and mobile applications

at the same power or up to 30% power reduction at the same speed. Logic density is increased by approximately 1.23 times, while overall chip density sees similar gains. Early product-like demonstrations indicate strong performance improvements and high SRAM yields, underscoring the maturity of the technology.

A13, which follows in 2029, builds on A14 as a direct shrink. While its area reduction of around 6% may seem modest, its real strength lies in backward compatibility. By maintaining design rule compatibility with A14, A13 enables customers to migrate designs quickly and efficiently, reducing both time-to-market and development costs.

This emphasis on design continuity reflects a broader trend in semiconductor development,



where ecosystem considerations are becoming as important as raw performance metrics.

Scaling AI compute through advanced packaging

While transistor scaling remains essential, TSMC made it clear that future performance gains will increasingly depend on advanced packaging and system integration technologies.

CoWoS (Chip-on-Wafer-on-Substrate) continues to play a central role in enabling AI workloads. The technology allows multiple chips, including logic and high-bandwidth memory (HBM), to be integrated into a single package, significantly increasing compute density and bandwidth.

TSMC announced that it has achieved production of 5.5-reticle CoWoS packages with yields exceeding 98%, marking a significant milestone in manufacturability. Looking ahead, the company plans to scale CoWoS to 14-reticle configurations by 2028 and beyond 14 reticles by 2029, supporting up to 24 HBM stacks.

This scaling trajectory is remarkable, effectively enabling exponential increases in compute capability within a single package. The number

of transistors integrated within these systems is expected to grow dramatically, further reinforcing the role of packaging as a key driver of performance.

Complementing CoWoS is TSMC's System-on-Wafer (SoW) technology, which takes integration to an even higher level. SoW enables wafer-scale integration of logic and memory, with potential configurations supporting dozens of compute dies and up to 64 HBM stacks. The upcoming SoW-X platform, expected in 2029, is designed to meet the extreme demands of AI training workloads.

At the same time, TSMC continues to advance its SoIC (System-on-Integrated-Chips) technology, which provides high-density 3D stacking capabilities. Compared to traditional 2.5D approaches, SoIC offers significantly greater interconnect density and improved power efficiency. Future iterations are expected to achieve even tighter bonding pitches and higher performance.

Together, these technologies represent a shift towards system-level scaling, where performance gains are achieved not just through smaller transistors, but through more sophisticated integration of multiple components.

Photonics and the future of data movement

One of the most compelling areas of innovation highlighted at the symposium was the integration of photonics into semiconductor systems.

TSMC's Compact Universal Photonic Engine (COUPE) is designed to enable co-packaged optics (CPO), allowing optical communication components to be integrated directly into semiconductor packages. This approach addresses one of the key bottlenecks in AI systems: data movement.

Compared to traditional copper interconnects, CPO with COUPE offers dramatic improvements in both power efficiency and latency. Early implementations provide up to four times greater energy efficiency and a tenfold reduction in latency, with further gains expected as the technology matures.

The first production implementation, featuring a 200Gbps micro-ring modulator, is expected in 2026. Looking ahead, TSMC aims to scale the technology to support 400Gbps data rates and beyond, with bandwidth densities reaching approximately 4Tbps per millimetre by 2030.


These developments are critical for enabling large-scale AI systems, where the efficient movement of data between chips and across racks is becoming as important as compute performance itself.

Automotive, robotics and physical AI


The symposium also highlighted the growing importance of what TSMC describes as "physical AI" — applications where AI systems interact directly with the physical environment.

This includes autonomous vehicles, robotics and emerging applications such as humanoid robots. These systems require a combination of high-performance computing, real-time sensing, connectivity and power management, all integrated into highly reliable platforms.


TSMC is addressing these requirements with specialised process technologies such as N3A and N2A, which are




Fiber to laser




Laser to SOI chip



Lens on laser



Lens on chip



Connect the World by lighting the way
- the essential link in next-generation optical connectivity!



vanguard
AUTOMATION
MYCRONIC

designed to meet stringent automotive standards. The introduction of nanosheet-based automotive processes marks a significant step forward, enabling higher performance and efficiency in safety-critical applications.

At the system level, physical AI platforms combine multiple semiconductor components, including application processors, microcontrollers, sensors and communication modules. The integration of these elements reflects the increasing complexity of modern electronics, where multiple domains converge within a single system.

The company also highlighted the role of silicon in enabling humanoid robotics, illustrating how sensing, computation and control functions come together to create intelligent machines capable of interacting with the real world.

Specialty technologies and broader innovation

Beyond leading-edge nodes, TSMC continues to invest in a wide range of specialty technologies that support diverse applications.

In the RF domain, the company's advanced CMOS technologies are enabling significant improvements in power efficiency and performance for wireless devices. These innovations are particularly relevant for emerging applications such as AI-enabled wearables and augmented reality (AR) systems.

Memory technologies are also evolving, with MRAM and RRAM playing an increasingly important role in next-generation designs. These non-volatile memory solutions offer advantages in scalability, performance and energy efficiency, making them well suited to both AI and embedded applications.

In display technologies, TSMC's introduction of high-voltage FinFET processes represents a notable advancement, enabling more efficient and compact solutions for OLED displays and near-eye devices.

Together, these developments highlight the breadth of TSMC's innovation ecosystem, extending far beyond advanced logic nodes.

AI is not just another application driving incremental demand; it is reshaping the entire landscape, from device architecture to system integration and manufacturing strategy. As AI evolves from generative models to physical systems, the demands placed on silicon will only continue to increase

Manufacturing scale and sustainability

To support the surge in AI-driven demand, TSMC is significantly expanding its manufacturing capacity. The company plans to build nine new fab phases in 2026 alone, more than doubling its historical average.

Capacity for advanced nodes such as N2 and A16 is expected to grow rapidly, with a compound annual growth rate of approximately 70% between 2026 and 2028. At the same time, demand for AI-specific manufacturing, including large die wafers and advanced packaging, is increasing dramatically.

This expansion is underpinned by strong collaboration between research and development teams and manufacturing operations, enabling rapid scaling of new technologies.

Alongside capacity growth, TSMC is also focused on sustainability. The company has committed to achieving net-zero emissions by 2050 and is implementing initiatives in resource recycling and water stewardship. These efforts reflect the increasing importance of environmental considerations in semiconductor manufacturing.

Startups and ecosystem innovation

A notable feature of the symposium was the "Innovate with TSMC" demo zone, which showcased emerging

technologies from startup partners. These included innovations in sensing, edge AI, connectivity and robotics.

Companies such as ARIA Sensing, Hailo, Rigpa.AI and VSORA highlighted how advanced semiconductor technologies are enabling new applications in areas ranging from real-time perception to energy-efficient AI inference.

These collaborations emphasise the role of TSMC not just as a manufacturer, but as an enabler of innovation across the broader semiconductor ecosystem. By providing access to advanced process technologies and design tools, the company is helping startups bring cutting-edge ideas to market.

Conclusion: A new era of system-level innovation

The 2026 Europe Technology Symposium made one thing abundantly clear: the semiconductor industry is entering a period of unprecedented growth and transformation.

AI is not just another application driving incremental demand; it is reshaping the entire landscape, from device architecture to system integration and manufacturing strategy. As AI evolves from generative models to physical systems, the demands placed on silicon will only continue to increase.

TSMC's response is a comprehensive roadmap that spans advanced logic, packaging, photonics and system integration. At the same time, the company is scaling its manufacturing capabilities and strengthening its ecosystem to meet the needs of a rapidly changing industry.

The message from the symposium is ultimately one of convergence. Future performance gains will not come from any single innovation, but from the integration of multiple technologies across different layers of the stack.

In this new era, success will depend not just on transistor scaling, but on the ability to orchestrate complex systems that bring together compute, memory, connectivity and intelligence. TSMC's roadmap suggests that it is well positioned to lead this transformation, unleashing innovation at every level of the semiconductor ecosystem.

When mechanics become yield-critical: motion design for the advanced packaging era



Advanced packaging has created new process bottlenecks related to the internal mechanics of the packaging equipment. As alignment tolerances compress from microns to tens of nanometers, motion platforms stop being background automation and instead become yield-critical infrastructure.

BY JUSTIN BRESSI, BUSINESS DEVELOPMENT MANAGER, AEROTECH INC.

VARIATIONS that previously sat comfortably below the sensitivity of the packaging processes — vibration, thermal drift, mechanical compliance, encoder latency — now surface directly as bonding defects, layer misalignment, or via inconsistency.

This shift is evident in production environments for hybrid bonding, fan-out packaging, and glass interposer manufacturing. The discussion below reflects the engineering challenges emerging across these applications, and how system-level motion design is reshaping what precision means in semiconductor assembly.

SIS: *What's driving the spike in motion-related failures in next-generation packaging equipment?*

JB: The primary driver is the tightening of alignment tolerances from the micron domain into the nanometer domain. At this scale, position stability is not determined by encoder resolution alone — it depends on how mechanical stiffness, thermal behavior, controller latency, cabling dynamics, and metrology timing interact while executing precise motion profiles.

Processes such as wafer-to-wafer hybrid bonding, high-bandwidth memory stacking, and through-glass via drilling expose this clearly. Any

deviation in planarity, any vibration mode left undamped, or any drift over time becomes visible as yield loss. Effects once considered negligible now matter.

System-level precision is not a property driven by any one individual component; it is the emergent behavior of the entire motion system. Improvements must therefore be systemic, not incremental.

SIS: *Why can't legacy motion architectures simply be upgraded to meet these new tolerances?*

JB: Traditional upgrade paths — swapping to a higher-resolution encoder, adding compensation routines, or selecting a stiffer stage — often fail to fully address the problem once tolerances fall into the nanometer range. At this level, every source of error shares the same budget, and these errors interact and compound.


Legacy architectures often rely on stacked mechanical assemblies, PID-only control, and correction

applied after error is detected. These approaches work when tolerances allow room for adjustment. They break down when correction latency exceeds the mechanical response time of the system.

The shift required is conceptual: motion must be co-engineered with the process, rather than added on as an automation layer. Integration, not component improvement, becomes the mechanism by which stability is achieved.

SIS: *How do you achieve sub-micron alignment and planarity in high-throughput hybrid bonding tools?*

JB: Sub-micron alignment and planarity are achieved through the combined effect of mechanical architecture, metrology feedback, control strategy, and thermal stability. Direct-drive structures reduce compliance, air bearings provide near-frictionless



movement, high-resolution digital encoders or laser interferometers supply nanometer-scale feedback, and synchronization latencies are tuned to microseconds.

Thermal behavior matters as much as structural accuracy. Material selection and heat path design influence how the system behaves during extended operation, especially under the highly dynamic and demanding motion profiles needed to meet throughput targets.

Alignment is not produced by a single stage or sensor. It is the result of the entire motion stack being engineered as a unified system.

SIS: *Why is through-glass via (TGV) drilling considered a “worst-case” motion problem — and how is it being addressed?*

JB: TGV drilling simultaneously requires pinpoint accuracy, timing coordination between several subsystems, and extreme throughput. The substrate must be positioned precisely, the laser path must be synchronized with motion, and pulse placement must remain consistent even as the system accelerates through complex trajectories.


To achieve this in production environments, the servo stages and galvo scan heads must operate as a single coordinated motion system. Infinite Field of View (IFOV) blends the motion of both elements, eliminating stitching errors across the substrate. Position Synchronized Output (PSO) ensures each laser pulse is triggered at the correct spatial coordinate.

Setup and optimization are accelerated through machine-learning-based step-and-settle tuning (via Aerotech’s DrillOptimizer tool), which identifies efficient motion and pulse parameters automatically.

This allows TGV performance to be repeatable and scalable, not just achievable in isolated demonstrations.

SIS: *Can precision and speed be co-optimized, or will they always compete?*

JB: Precision and speed can conflict if they are treated as independent objectives. In practice, they are both



outcomes of mechanical stiffness, inertial sensing, and control design.

Increasing structural stiffness raises the natural resonance frequency of the system, allowing higher acceleration without inducing vibration. Real-time inertial feedback — such as integrated accelerometers — enables more aggressive servo tuning without overshoot. Control strategies that reflect the mechanical resonance profile maintain throughput without sacrificing accuracy.

The goal is not to balance speed against precision, but to engineer the system so that speed and precision are both reinforced by the system architecture.

SIS: *What challenges do equipment builders face when scaling from R&D tools to high-volume manufacturing systems?*

JB: A common issue is control platform discontinuity. Many tools begin development on one control architecture and shift to another for volume deployment, which forces re-tuning, re-validation, and software re-integration — often delaying high-volume introduction.

Consistent control architecture from prototype through pilot and into production avoids this disruption.

Software, motion behavior, and servo tuning should travel intact across the development cycle, and subsystems should scale rather than be redesigned.

The motion architecture should scale with the tool, not reset when production begins.

Where motion control is headed next

Motion control in advanced packaging is moving toward closer integration of motion, sensing, and real-time analysis. Rather than correcting drift or vibration after the fact, next-generation control layers will predict and compensate continuously, informed by models of mechanical and thermal variation.

Motion system data and in-process metrology are beginning to be fused natively, allowing deviations to be detected earlier and compensated automatically.

Machine-learning-driven auto-tuning will gradually reduce manual optimization, particularly as production tools accumulate operational datasets that can generalize across recipes and substrates. Meanwhile, the motion platform itself is increasingly recognized as part of the process architecture, rather than a mechanical subsystem — especially in chiplet assembly, hybrid bonding, and glass interposer manufacturing.

The next phase of advanced packaging will depend not just on accurate positioning, but on how effectively motion, sensing, and control evolve together as a unified system.



**SAVE THE
DATE
12-14
APRIL
2027**

**Sheraton Brussels
Airport Hotel**

To find out more about our sponsor and speaker
opportunities, contact us today on:

+44 (0) 2476 718970

or email:

info@csinternational.net

info@picinternational.net

info@peinternational.net

info@angel-tech.net

Advanced 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)

