PIC PHOTONIC INTEGRATED CIRCUITS

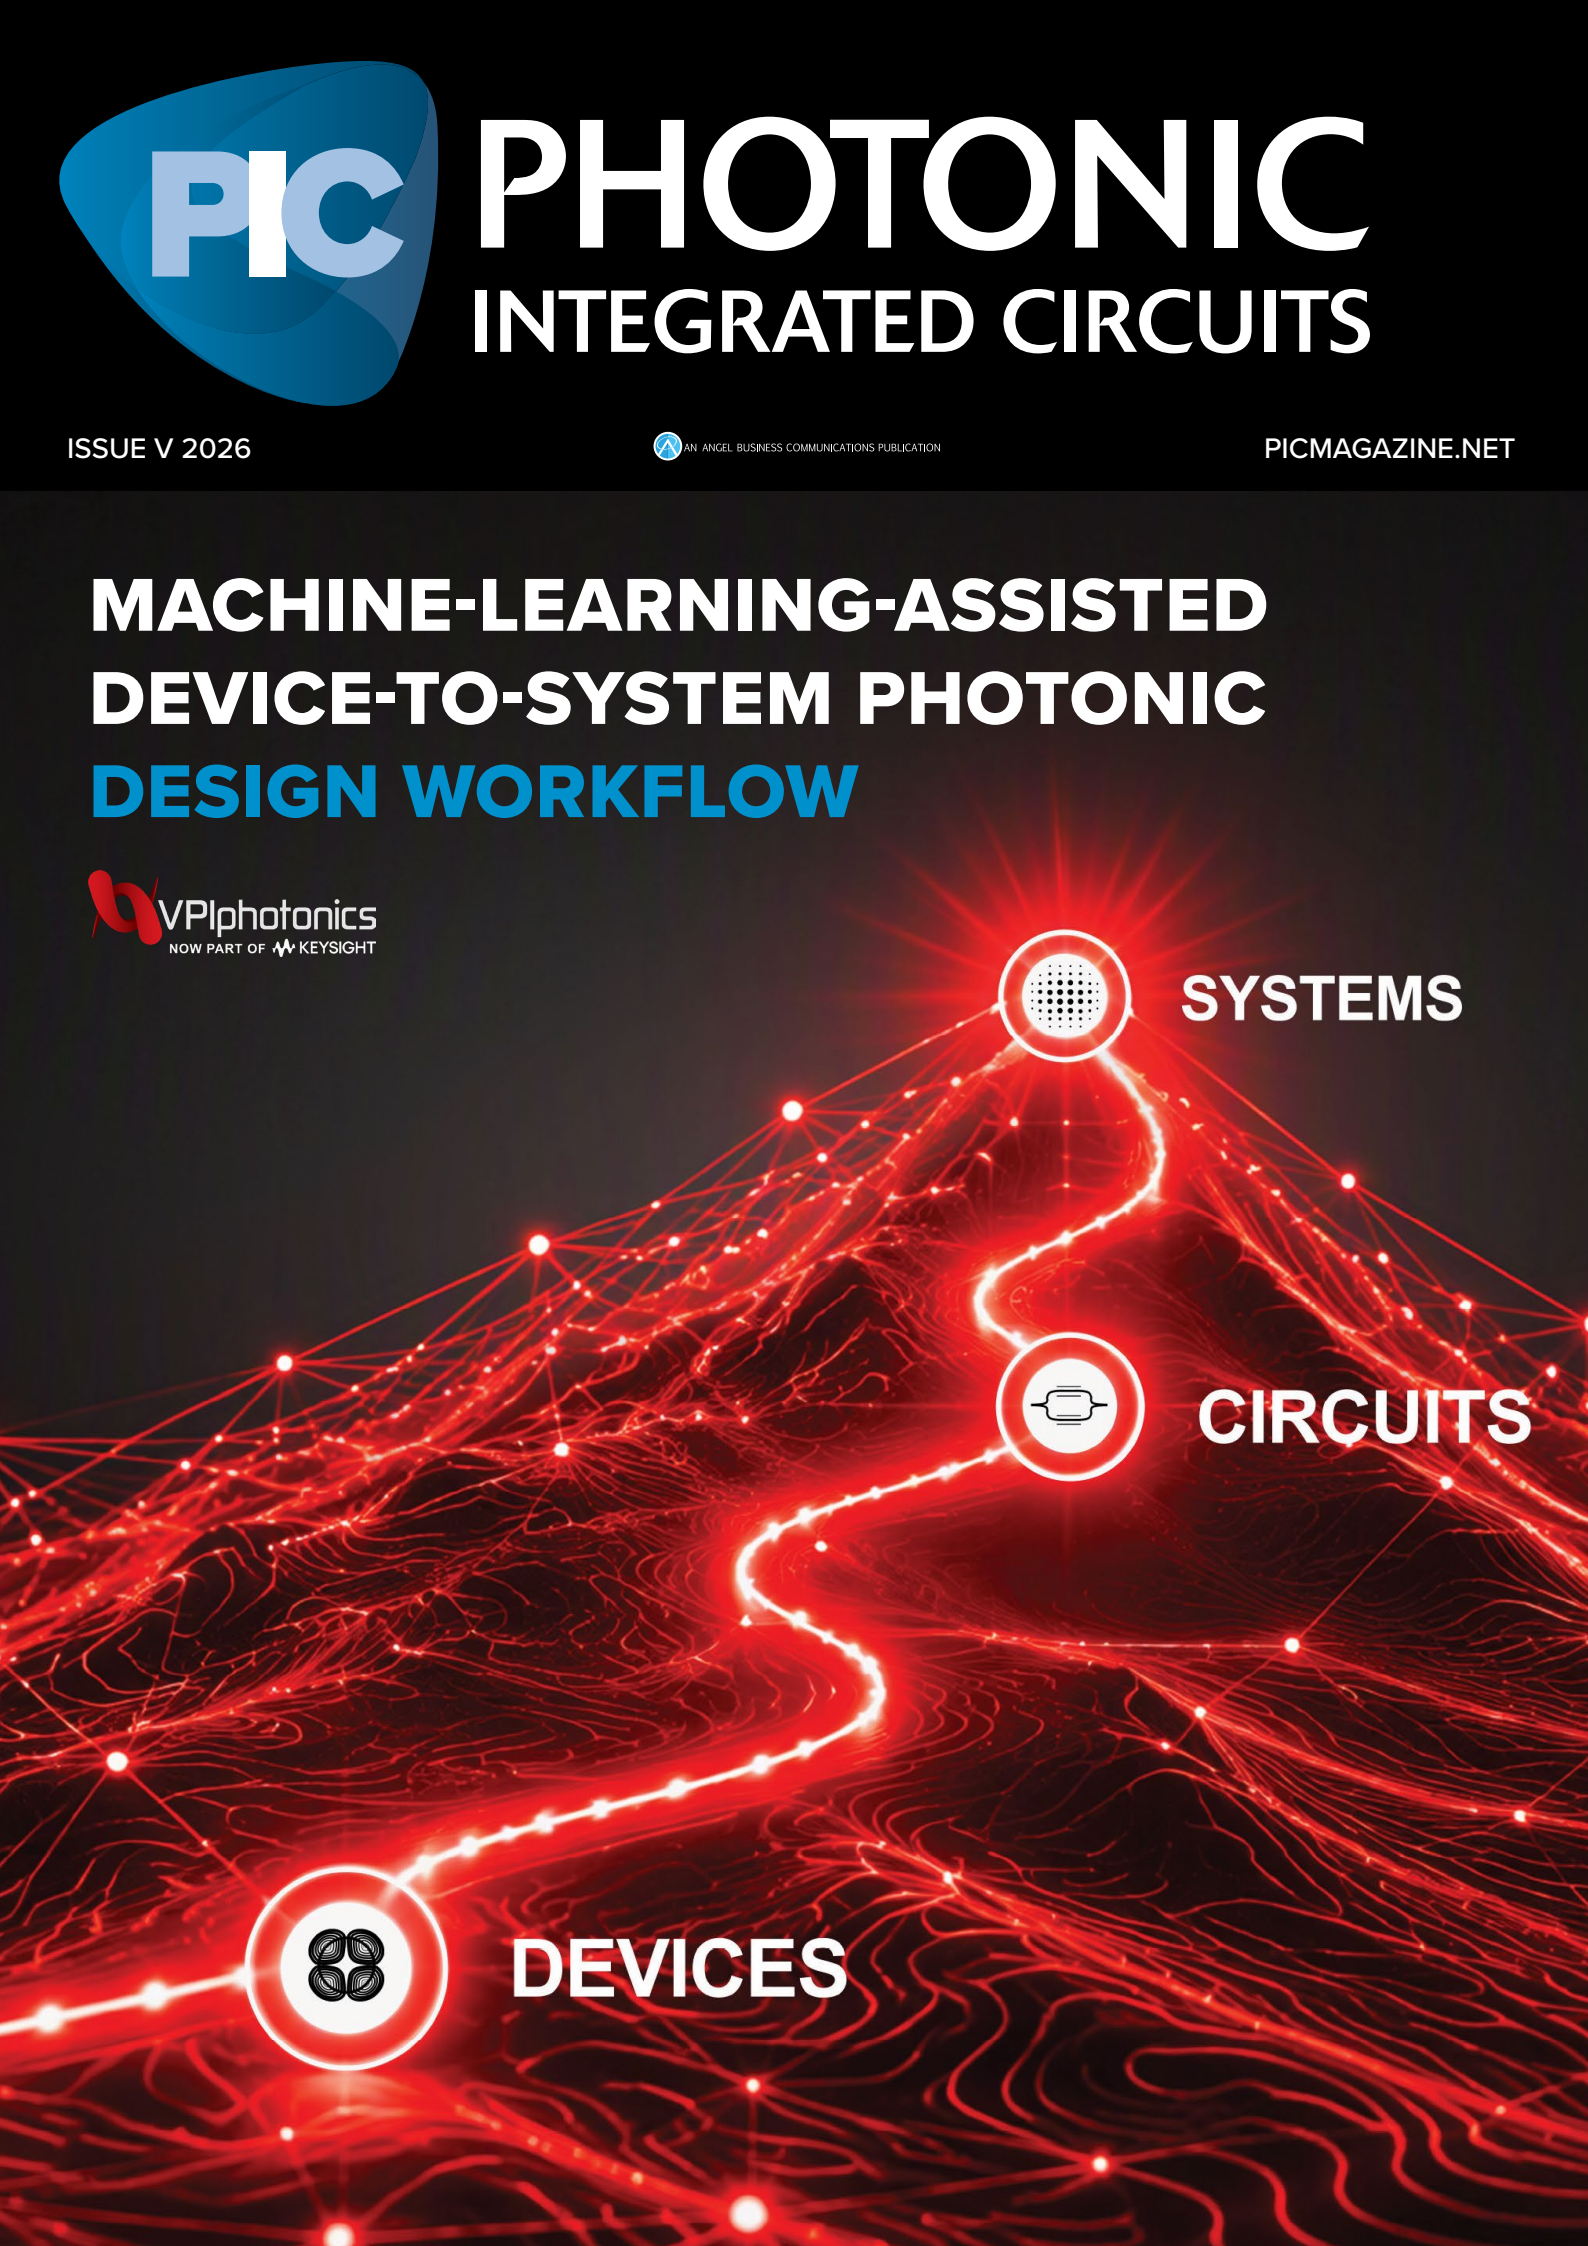
ISSUE V 2026

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

PICMAGAZINE.NET

MACHINE-LEARNING-ASSISTED DEVICE-TO-SYSTEM PHOTONIC DESIGN WORKFLOW

VPIphotonics
NOW PART OF KEYSIGHT



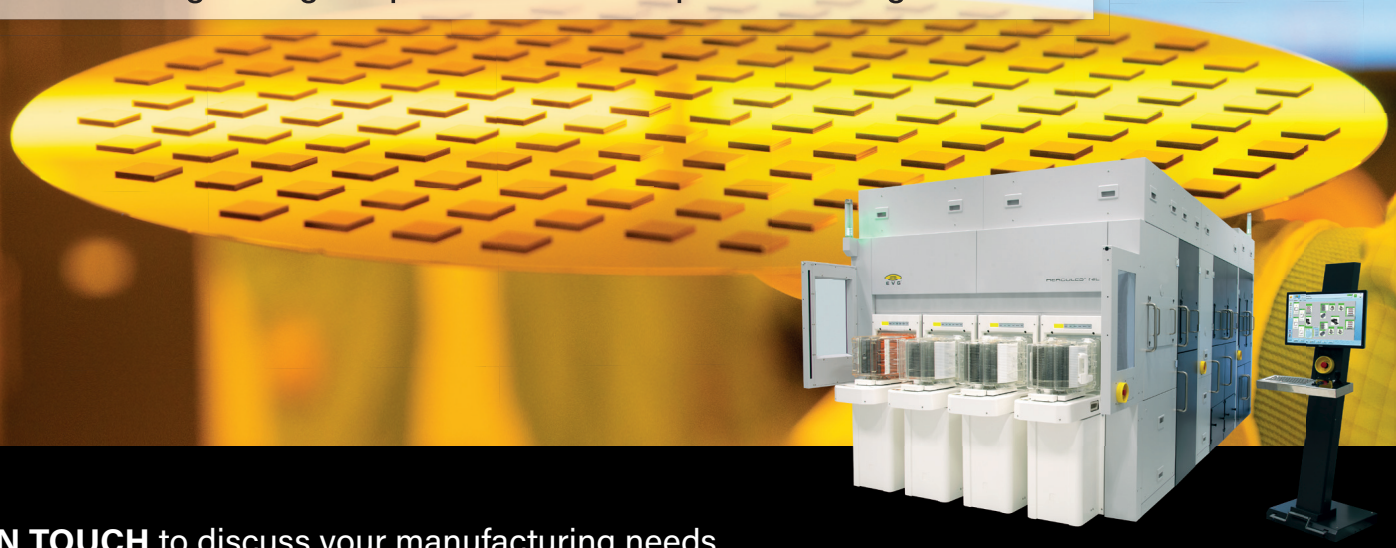
EVG®

invent
innovate
implement

www.EVGroup.com

INTEGRATION SOLUTIONS FOR NEXT-GEN CO-PACKAGED OPTICS

- Enabling customers to manufacture scalable CPO architectures from R&D to HVM
- Hybrid Bonding for low-parasitic PIC/EIC integration and HBM manufacturing
- Heterogeneous integration for laser integration, advanced photonics, and thermal-management materials
- Nanoimprint Lithography for efficient optical coupling and beam shaping
- Advancing next-gen optical I/O and 3D photonic integration



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

VIEWPOINT

BY SARAB CHOPRA, EDITOR



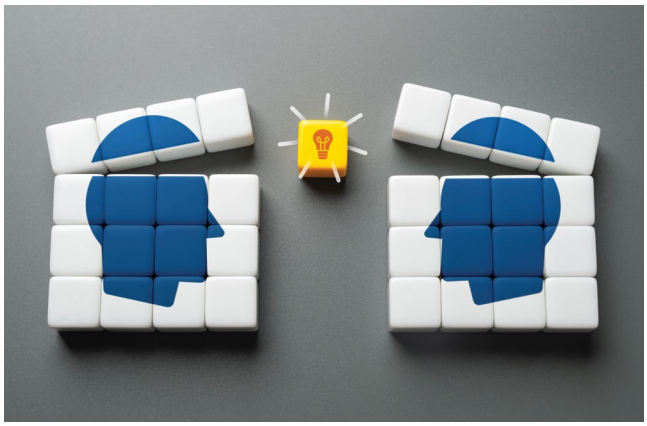
Building the AI photonics era

Photonic integrated circuits are moving from specialist technologies into the infrastructure underpinning a growing range of applications. As artificial intelligence drives demand for faster data movement, lower latency and improved energy efficiency, photonics is becoming central to the way information is generated, processed and transported. Yet the opportunity extends well beyond AI data centres. The same progress in integration, materials and manufacturing is opening possibilities across telecommunications, sensing, healthcare, quantum technologies and industrial systems.

This issue of PIC Magazine explores an ecosystem in transition. Advances in silicon photonics, silicon nitride and compound-semiconductor technologies are expanding the range of functions that can be integrated on chip. At the same time, heterogeneous integration is bringing together materials and devices with complementary strengths, creating new routes to performance that are difficult to achieve with any single platform. The challenge is no longer simply to demonstrate an impressive component in the laboratory. It is to turn promising devices into robust, repeatable products that can be manufactured and deployed at scale.

That shift places packaging and assembly firmly in the spotlight. Optical coupling, fibre alignment, thermal management and mechanical stability can all determine whether a photonic design performs as intended outside controlled laboratory conditions. As architectures become more complex, these considerations must be addressed earlier in the design process, rather than treated as downstream manufacturing problems. Closer collaboration between PIC designers, foundries, packaging specialists, equipment suppliers and end users will be essential.

The rise of co-packaged optics is one example of this convergence. By bringing optical engines closer to electronic processing, CPO has the potential to address bandwidth, power and interconnect constraints in high-performance computing. But its progress will depend on more than optical performance alone. Manufacturing flows, test strategies, serviceability, standards and supply-chain



readiness all need to mature together. Open, interoperable approaches can help the wider ecosystem move beyond isolated demonstrations towards solutions that customers can adopt with confidence.

Design automation and simulation are also becoming increasingly important as photonic systems grow in scale. Accurate models, efficient workflows and dependable testing can help reduce development cycles and bridge the gap between device-level innovation and system-level requirements. Alongside this, advances in metrology, process control and manufacturing automation will be critical to improving yield and consistency.

No single technology or organisation will define the next phase of photonics. Progress will come from connecting research with production, aligning component capabilities with system needs, and building shared approaches to reliability and scale. The key question is how quickly the industry can make complex photonic integration repeatable, accessible and commercially sustainable.

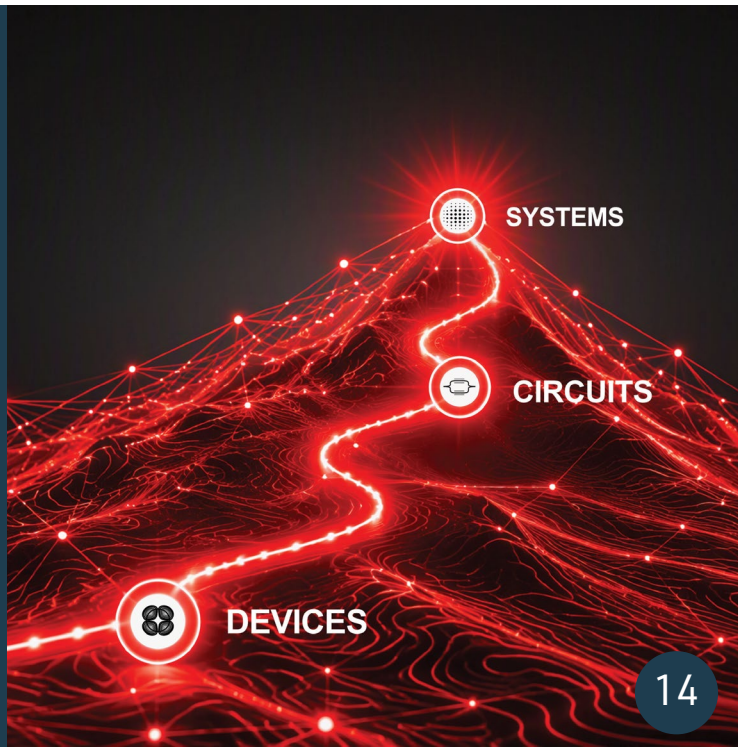
Across the pages of this issue, the common thread is clear: the future of photonics will be shaped not only by what can be integrated, but by how effectively that integration can be designed, packaged, tested and manufactured. That is the work now turning photonic potential into practical infrastructure.

Contents

Cover Feature

Machine-learning-assisted device-to-system photonic design workflow

An integrated device-to-system simulation workflow combines electromagnetic modelling and machine learning to accelerate the inverse design and optimization of microring resonator demultiplexers for high-capacity data centre interconnect receivers.



20 End-to-end process flow under one roof

Turning a photonic chip into a product means bringing together photonics, electronics, packaging and test. Handoffs between separate actors are where processes often stall. End-to-end process flow control under one roof may offer the best way through.

24 Barium titanate emerges as a route to faster AI optical links

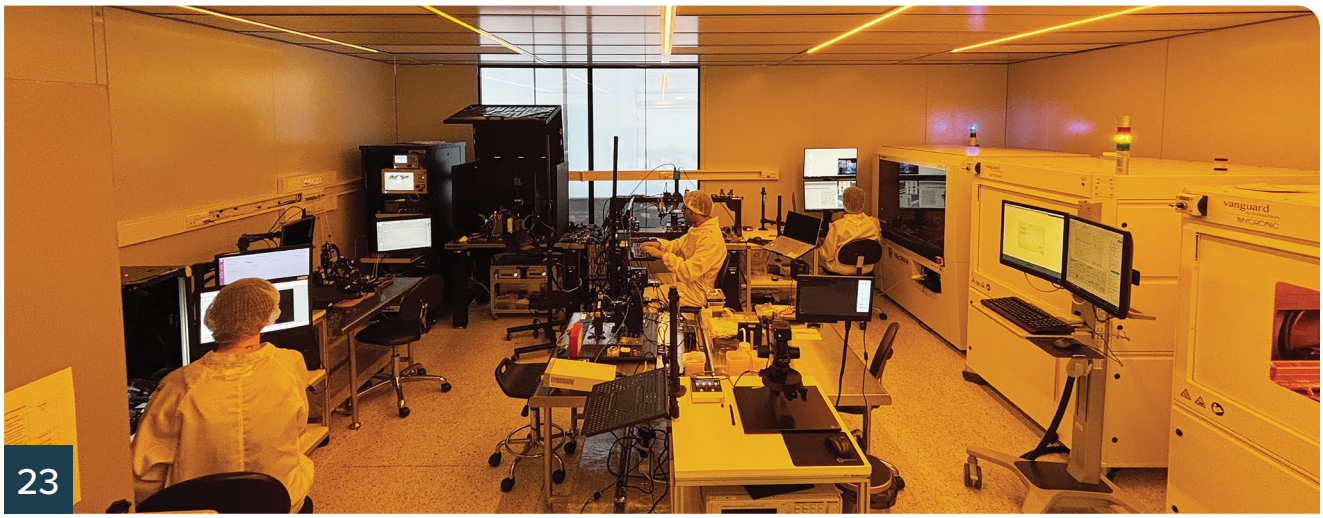
A barium titanate modulator built on a commercial silicon photonics line has carried a record net 1.6 T — and production-grade metal-oxide epitaxy is the quiet enabler that finally makes it manufacturable.

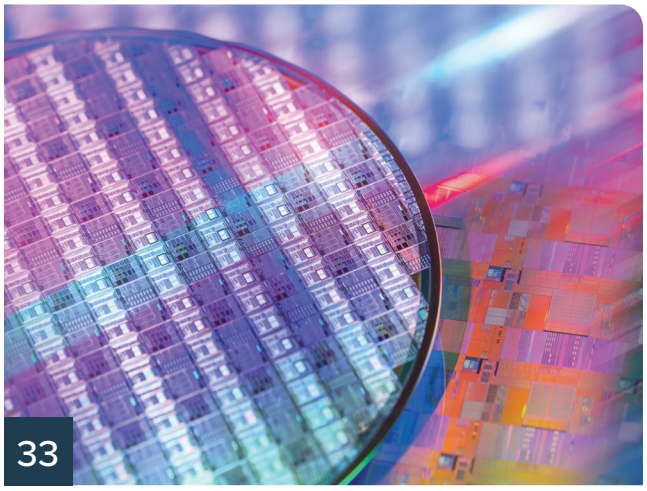
30 End-to-end nanoimprint lithography technology for photonic integration

Nanoimprint lithography enables high-precision, repeatable fabrication of micro-optics for photonic integration, with inkjet coating and accurate alignment supporting scalable manufacturing for CPO, microLEDs and VCSELs.

36 Why AI infrastructure must compute while data moves

AI's next infrastructure bottleneck is not arithmetic alone but data movement. Optalysys explores how photonic Compute-in-Transit could turn the data path itself into a useful computational resource for future AI systems.





33

40 Enabling high-volume silicon photonics through advanced test

Advanced photonic test is becoming critical to scaling silicon photonics for AI and HPC, with automated optical alignment, integrated electrical and optical measurements, and wafer-level testing.

44 Exorcising the DSP: A back-to-basics route to co-packaged optics

Every high-speed optical module hides a digital signal processor that quietly eats power and adds delay. Mark Holm argues the way forward is simpler modulation, many wavelengths and a single comb source.

48 Scaling VECSEL innovation for the next generation of photonics

Vexlum is expanding its presence in the UK with a new R&D laboratory focused on supporting customers and strengthening collaboration across the country's growing photonics and quantum ecosystem.

NEWS

06 JEDEC sets SiPh reliability standard

07 Ultrafast laser targets AI photonics

08 Fujitsu advances diamond photonics

09 MIT advances flexible silicon photonics

10 OIF highlights AI optical interconnect work

11 Hybrid PIC market set to grow by 2035

12 Photonics packaging market set to grow



12

PIC PHOTONIC INTEGRATED CIRCUITS

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

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

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

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

Angel BUSINESS COMMUNICATIONS

PIC Magazine is published six 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: Pic Magazine, 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 Pic Magazine, 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.

JEDEC sets SiPh reliability standard

JESD264 establishes industry-wide qualification guidelines for silicon photonics as PIC and CPO adoption accelerates.

JEDEC has published JESD264, its first industry-wide reliability standard for silicon photonics, establishing common qualification and reliability requirements as silicon photonics moves towards wider commercial deployment.

The standard addresses reliability across four key areas: silicon photonics dies and integrated chipsets, individual photonic components, assembly and laser integration, and manufacturing process control and traceability.

The development of a common framework addresses a growing need for consistent reliability testing as photonic integrated circuits (PICs) combine optical and electronic functions within increasingly complex architectures.

Unlike conventional semiconductor devices, PICs incorporate optical elements including waveguides, couplers, splitters, modulators and photodetectors.

Their reliability can therefore involve optical-specific effects such as changes in insertion loss, optical power degradation, wavelength drift and coupling misalignment, alongside conventional electrical and material-related failure mechanisms.

These challenges become more significant as silicon photonics moves into co-packaged optics (CPO), 2.5D and 3D integration and multi-chip packages.

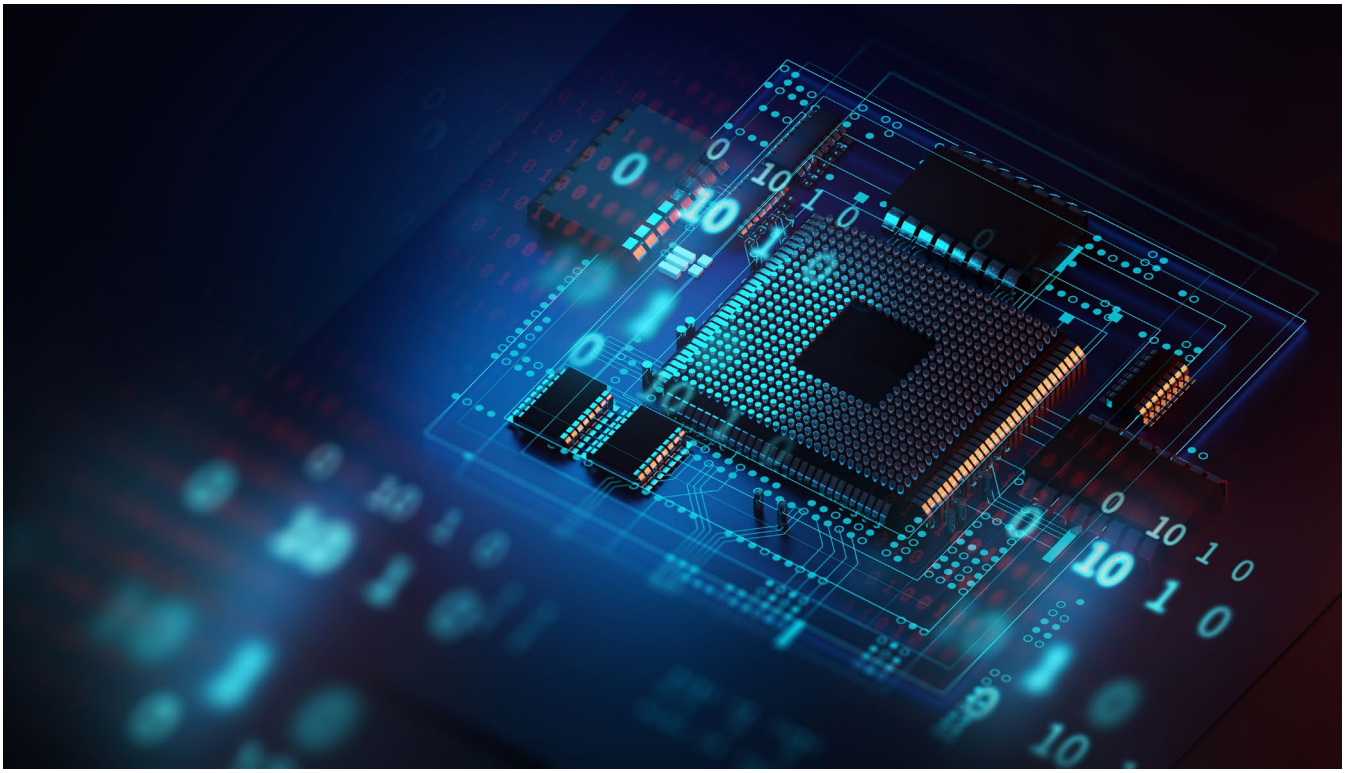
Thermal expansion, package warpage, bonding and optical alignment can affect performance when photonic and electronic components, lasers, fibres and substrates are integrated into increasingly compact assemblies.

JESD264 is intended to provide a common basis for qualification across the supply chain, including chip designers, foundries, OSATs, module manufacturers and system developers.

Standardised reliability practices could help improve consistency in evaluating silicon photonics devices and accelerate their adoption in high-volume applications.

The publication comes as major semiconductor manufacturers increase their investment in silicon photonics and optical interconnect technologies for AI and high-performance computing.

As photonic integration moves closer to the processor and package, reliability qualification is becoming an increasingly important part of the technology's development.



Ultrafast laser targets AI photonics

Researchers from Russia and China have developed an ultrafast micro-laser for photonic chips targeting artificial intelligence applications, achieving a reported response speed of 4.15GHz.

SCIENTISTS from HSE University in St Petersburg and the Shandong Academy of Sciences developed the laser with the aim of supporting faster photonic chips for AI systems, including applications in data centres and neural networks.

The researchers directly measured the maximum response speed of the micro-laser and used a standardised measurement method to assess its performance when operated using either optical or electrical excitation.

Tests showed a maximum response speed of 4.15GHz, which the researchers described as the highest recorded for this type of micro-laser.

The development addresses growing demand for higher computing performance and faster data transfer as AI workloads become increasingly complex.

Photonic integrated circuits, which use light to transmit and process information, are being investigated as an alternative or complement to conventional electronic approaches.

According to the researchers, faster photonic devices could support the development of AI chips capable of integrating multiple computing functions on a single chip while improving data transfer between components.

The work could therefore contribute to the development of higher-speed photonic architectures for future AI and data-centre systems.

OIF defines 1.6T coherent interface

THE Optical Internetworking Forum (OIF) has published its Implementation Agreement (IA) for 1600ZR Coherent Interfaces, establishing a common technical framework for transmitting a 1.6T Ethernet client over a single coherent wavelength.

The specification is designed for single-span, amplified, point-to-point Dense Wavelength Division Multiplexing (DWDM) links extending up to 120km. OIF said the agreement doubles the capacity of its 800ZR specification while maintaining the reach and interoperability required for data-centre interconnect (DCI) applications.

The move comes as AI training, inference and cloud workloads drive increasing traffic within and between data centres. Higher optical capacity is needed without proportionally increasing power consumption, cost or equipment footprint.

According to OIF, 1600ZR combines 1.6T transmission with DCI-class reach and high port density, providing a basis for pluggable coherent solutions from multiple vendors.

“1600ZR is a defining milestone for coherent networking,” said Karl Gass, OIF Physical & Link Layer Working Group Optical Vice Chair. “It doubles the capacity of 800ZR while preserving the reach, interoperability and client-like density that have made the ZR ecosystem so fundamental.”

The 1600ZR specification defines the client mapping, frame format, forward-error correction (FEC), modulation and optical characteristics required for interoperable implementations. It follows OIF’s earlier 400ZR and 800ZR work, which helped move coherent optics into routers and switches.

“1600ZR continues the trajectory OIF set with 400ZR, which fundamentally changed how coherent optics could be deployed,” said Ian Betty, OIF Board Member at Ciena. “Delivering 1.6T over a single wavelength marks another major step forward—and provides a critical building block for the AI infrastructure now being deployed.”

The specification also highlights the industry’s continued effort to increase capacity per wavelength, rather than relying solely on additional wavelengths or increased system footprint.

For photonic system designers, this places greater emphasis on high-speed coherent transmitters, receivers, modulators and associated optical components.

OIF members were scheduled to demonstrate multi-vendor ZR interoperability at ECOC 2026, held from 21–23 September in Málaga, Spain. The demonstrations were intended to showcase coherent implementations from across the ecosystem and provide an opportunity to assess the practical deployment of the new 1.6T specification.

This places greater emphasis on high-speed coherent transmitters, receivers, modulators and associated optical components.

Fujitsu advances diamond photonics

Fujitsu has developed a diamond-spin quantum computing prototype integrating tin-vacancy centres with photonic circuits, using optical connectivity to support scalable quantum architectures.

FUJITSU has developed a working diamond-spin quantum computer prototype that integrates tin-vacancy (SnV) colour centres with photonic integrated circuits (PICs), creating an architecture designed to support scalable and modular quantum computing.

The prototype operates at -271.6°C and uses SnV centres in diamond as qubits. The centres offer stable quantum states and can emit single photons, making them suitable for both quantum information processing and optical communication between quantum modules.

A key element of the architecture is the integration of the diamond-based qubits with photonic circuitry. Fujitsu developed a heterogeneous bonding process to attach diamond substrates containing implanted tin to alumina and silicon-dioxide substrates.

The diamond is then thinned from several hundred micrometres to several hundred nanometres, enabling integration into a compact photonic platform.

The resulting structure incorporates nanometre-scale diamond crystals containing SnV centres alongside alumina optical waveguides. The waveguides can collect photons emitted by the SnV centres during qubit readout and provide an optical interface for transferring quantum information.

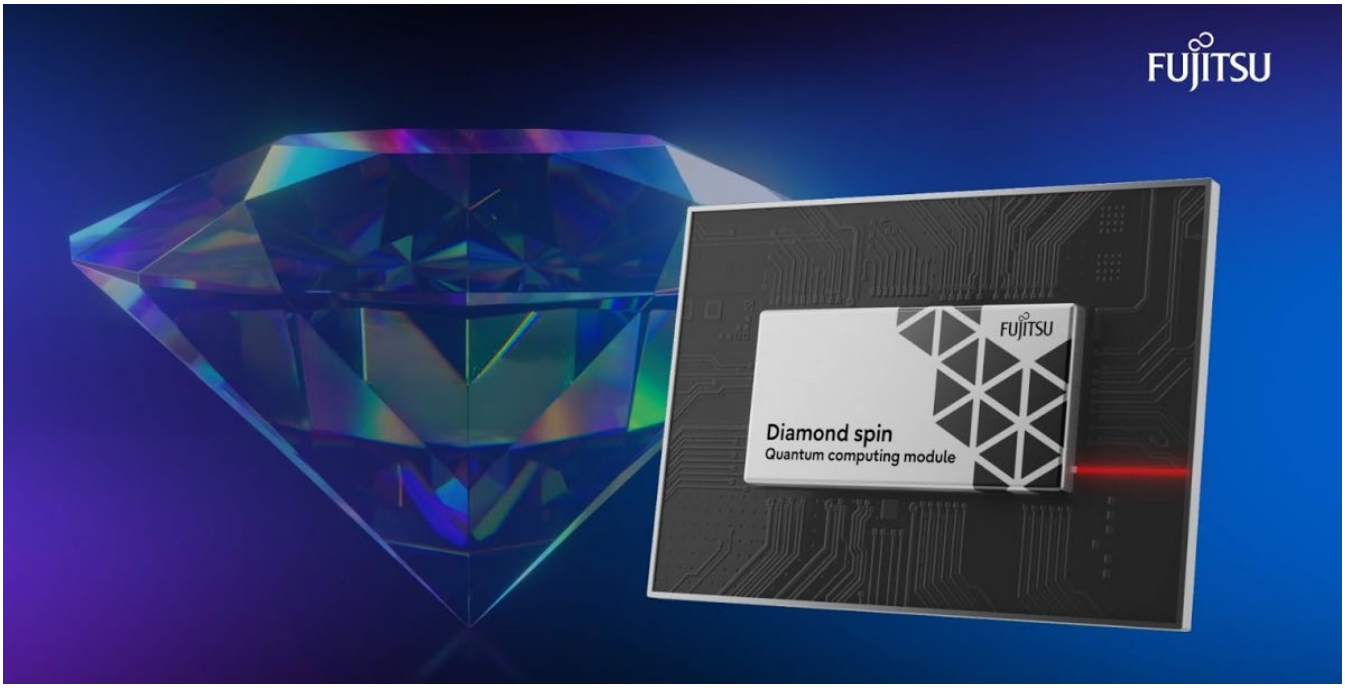
This optical connectivity could allow individual quantum modules to be linked without requiring all qubits to be placed on a single large processor. Such a modular approach is intended to address one of the key challenges in scaling quantum computing systems.

The prototype also incorporates optical, microwave and radio-frequency control. Fujitsu has developed a circuit-conversion mechanism that translates quantum gates into the physical control sequences required to operate the diamond-spin qubits.

The system can be operated through Fujitsu's Hybrid Quantum Computing Platform, providing an interface between the quantum processor and higher-level quantum applications.

The work follows a research collaboration between Fujitsu, Delft University of Technology and QuTech that began in 2020. Fujitsu plans to develop a multi-module diamond-spin quantum computer prototype by 2027, while also investigating the combination of diamond-spin and superconducting quantum technologies.

The integration of diamond quantum devices with PICs highlights the potential of photonic technologies to provide the optical readout and interconnect infrastructure required for future modular quantum processors.



MIT advances flexible silicon photonics

MIT researchers have developed a wafer-scale fabrication platform that combines silicon photonics with mechanical flexibility and optical transparency, potentially extending photonic integrated circuits (PICs) beyond conventional rigid, planar chip architectures.

THE APPROACH enables flexible and transparent silicon-photonic devices to be fabricated on 300-mm wafers using established semiconductor manufacturing processes.

This could provide a route towards scalable production of PICs designed for applications where optical circuits need to conform to curved or moving surfaces.

Conventional silicon photonics relies on rigid semiconductor substrates, limiting where photonic circuits can be deployed. While flexible and transparent photonic devices have previously been demonstrated, many have relied on specialised laboratory-scale fabrication techniques.

The MIT platform instead seeks to retain the manufacturing advantages of large-diameter semiconductor wafers while introducing mechanical flexibility into the finished photonic devices.

Silicon photonics integrates optical functions including waveguides and other nanoscale structures onto semiconductor substrates, enabling light to be generated, routed and processed within compact systems.

Applying these structures to a flexible and transparent platform could expand the range of physical architectures available for integrated photonics.

The researchers demonstrated the mechanical performance of the devices by repeatedly bending them around cylinders with different diameters.

The chips remained operational after thousands of bending cycles, including tests involving a cylinder approximately as narrow as a small screw. This indicates that the photonic structures can withstand repeated mechanical deformation without significant degradation in their optical performance.

Optical measurements also examined the transparency of the devices. The researchers reported limited haze and distortion when viewing through the material, an important characteristic for applications where the photonic circuit must operate without substantially obstructing the underlying optical path.

The combination of wafer-scale fabrication, optical transparency and mechanical flexibility could enable PICs to be integrated into applications that are difficult to address with conventional rigid silicon-photonic platforms.

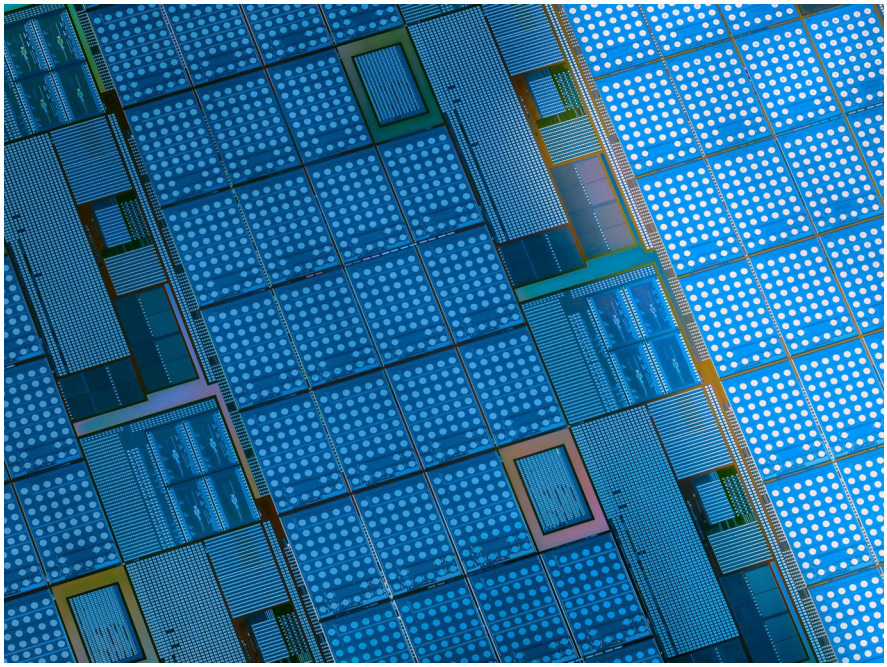
Potential areas include conformable optical sensors, wearable systems and transparent augmented-reality displays.

For PIC manufacturing, the use of 300-mm wafers is particularly significant. Rather than treating flexibility and transparency solely as specialised device characteristics, the approach attempts to integrate them into a fabrication flow compatible with established semiconductor infrastructure.

This could provide a foundation for further development and scaling of flexible photonic technologies.

The technology could also introduce new challenges for future photonic integration. Flexible PICs may require new approaches to optical coupling, electrical interconnects, packaging and assembly, particularly where devices must maintain performance while being bent or attached to curved surfaces.

By combining the manufacturing scalability of silicon photonics with mechanical flexibility and transparency, the MIT platform could help broaden the physical form factors of PICs and support the development of integrated optical systems that move beyond conventional flat-chip architectures.



OIF highlights AI optical interconnect work

OIF highlighted its work on high-speed electrical and optical interfaces for AI infrastructure at two industry events in October.

AT THE Ethernet Alliance Technology Exploration Forum (TEF) 2026, OIF experts were scheduled to participate in discussions covering 448G electrical signalling, AI optical interconnects, coherent optics, and test and measurement.

The sessions were intended to address the challenges of developing interoperable technologies as AI infrastructure places increasing demands on bandwidth, power efficiency and system scalability.

OIF also scheduled a half-day workshop at the OCP Global Summit on October 13, bringing together industry experts to examine technologies needed for open, scalable and sustainable AI infrastructure.

Planned topics included 448G-per-lane electrical interfaces, energy-efficient interconnects, coherent optics and optical network management.

The workshop was also set to explore different approaches to optical integration, including near-package optics (NPO), co-packaged optics (CPO) and retimed transmitter linear receiver (RTLR) architectures.

Discussions were expected to consider how these technologies could address the increasing connectivity requirements of AI systems, from package-level interfaces through to data-centre interconnects.

Interoperability was another key focus, with OIF highlighting the importance of multi-vendor solutions as AI infrastructure becomes increasingly complex. The organisation said its work is aimed at supporting common electrical, optical and management interfaces that can scale with evolving AI architectures.

The events were expected to bring together representatives from network operators, system and component suppliers, and test equipment companies to discuss the development and deployment of next-generation interconnect technologies.

3D integration reshapes CPO

THE GROWING bandwidth and power demands of AI infrastructure are accelerating the shift from pluggable optical transceivers towards co-packaged optics (CPO), according to IDTechEx.

The research firm expects CPO to reach a cumulative market value of $20 billion by 2036, as optical engines move closer to network switch ASICs to reduce electrical interconnect distances, latency and power consumption.

Early CPO architectures have placed photonic integrated circuits (PICs) and electronic integrated circuits (EICs) side-by-side in 2D or on an interposer in 2.5D configurations. However, increasing bandwidth requirements are creating greater pressure on package real estate, driving interest in vertical 3D integration.

Monolithic integration of PICs and EICs into a single die could minimise interconnect distances, but differences between photonic and electronic manufacturing requirements create performance trade-offs. IDTechEx therefore expects heterogeneous 3D integration, in which separately optimised photonic and electronic dies are vertically stacked, to gain traction.

Cu-Cu hybrid bonding is emerging as a key technology for this approach, enabling interconnect pitches below 10µm while reducing parasitics compared with conventional microbumps. However, tighter pitches increase alignment, manufacturing yield and cost challenges.

Thermal management also remains a major obstacle, particularly because heat-generating electronics are positioned close to temperature-sensitive photonic components. IDTechEx expects 3D heterogeneous integration using hybrid bonding to account for a significant share of the CPO market as optical and electronic integration continues to evolve.

Cu-Cu hybrid bonding is emerging as a key technology for this approach, enabling interconnect pitches below 10µm while reducing parasitics compared with conventional microbumps. However, tighter pitches increase alignment, manufacturing yield and cost challenges.

Hybrid PIC market set to grow by 2035

AI, data centers, 5G and sensing applications are expected to drive growth in hybrid photonic integrated circuits over the next decade.

THE GLOBAL hybrid photonic integrated circuit (PIC) market was valued at $2.45 billion in 2025 and is forecast to reach $6.82 billion by 2035, according to a new report from ResearchAndMarkets.com.

The market is projected to grow at a compound annual growth rate of approximately 13.6% between 2026 and 2035, driven by demand for high-bandwidth data center interconnects, telecommunications infrastructure, LiDAR and sensing technologies, and photonic systems for quantum computing.

Data communication accounted for 45.2% of market revenue in 2025, making it the largest application segment. Growing AI workloads and increasing data volumes are driving demand for higher-capacity and more energy-efficient optical connectivity in data centers and high-performance computing environments.

The report identifies silicon photonics as a key material platform, alongside hybrid approaches including silicon-III-V and silicon nitride-III-V integration and thin-film lithium niobate.

North America was the largest regional market in 2025, while Asia Pacific is forecast to be the fastest-growing region, supported by 5G deployment, semiconductor manufacturing investment and expanding cloud infrastructure.

The report identifies companies including Intel, Broadcom, Infinera, Lumentum, Nokia, Marvell, Coherent and others as participants in the hybrid PIC market.

Photonics targets 400Gbps links

INTEGRATED PHOTONICS is advancing as demand for higher-bandwidth, lower-power optical interconnects grows across communications and data centre applications.

Imec is developing integrated photonics platforms covering device design, wafer fabrication, packaging and testing. Its process design kits (PDKs) and device libraries allow customers to develop photonic integrated circuits (PICs) using established manufacturing processes.

The push towards higher data rates is also driving new device and material approaches. Imec has developed silicon photonics technologies targeting electro-optical bandwidths above 100GHz and data rates of up to 400Gbps per lane.

Heterogeneous integration is playing an increasing role, with germanium-on-silicon used for photodetection and thin-film lithium niobate being integrated with silicon photonics for high-speed modulation. Imec and Ghent University have demonstrated lithium-niobate modulators integrated onto silicon photonic platforms using micro-transfer printing.

These developments highlight the growing importance of combining different materials and functions within manufacturable PIC platforms. As bandwidth increases, fabrication must be coordinated with optical coupling, packaging, thermal management and testing.

The move towards platform-based fabrication could help bridge the gap between photonic research and volume deployment, providing a pathway towards higher-speed optical interconnects for future data centres and communications systems.



Photonics packaging market set to grow

AI infrastructure, co-packaged optics and wafer-level integration are expected to drive growth in the global photonics packaging market through 2036.

THE GLOBAL photonics packaging market is set for significant development over the next decade as AI infrastructure, co-packaged optics (CPO), advanced semiconductor packaging and other emerging applications increase demand for higher levels of photonic integration, according to a new report from ResearchAndMarkets.com.

The Global Photonics Packaging Market 2026-2036 report identifies AI infrastructure as a key catalyst, with increasing bandwidth requirements placing greater pressure on conventional pluggable optical architectures. CPO, which integrates optical engines alongside switch or compute chips, is highlighted as an important approach for reducing electrical connection distances and improving power efficiency.

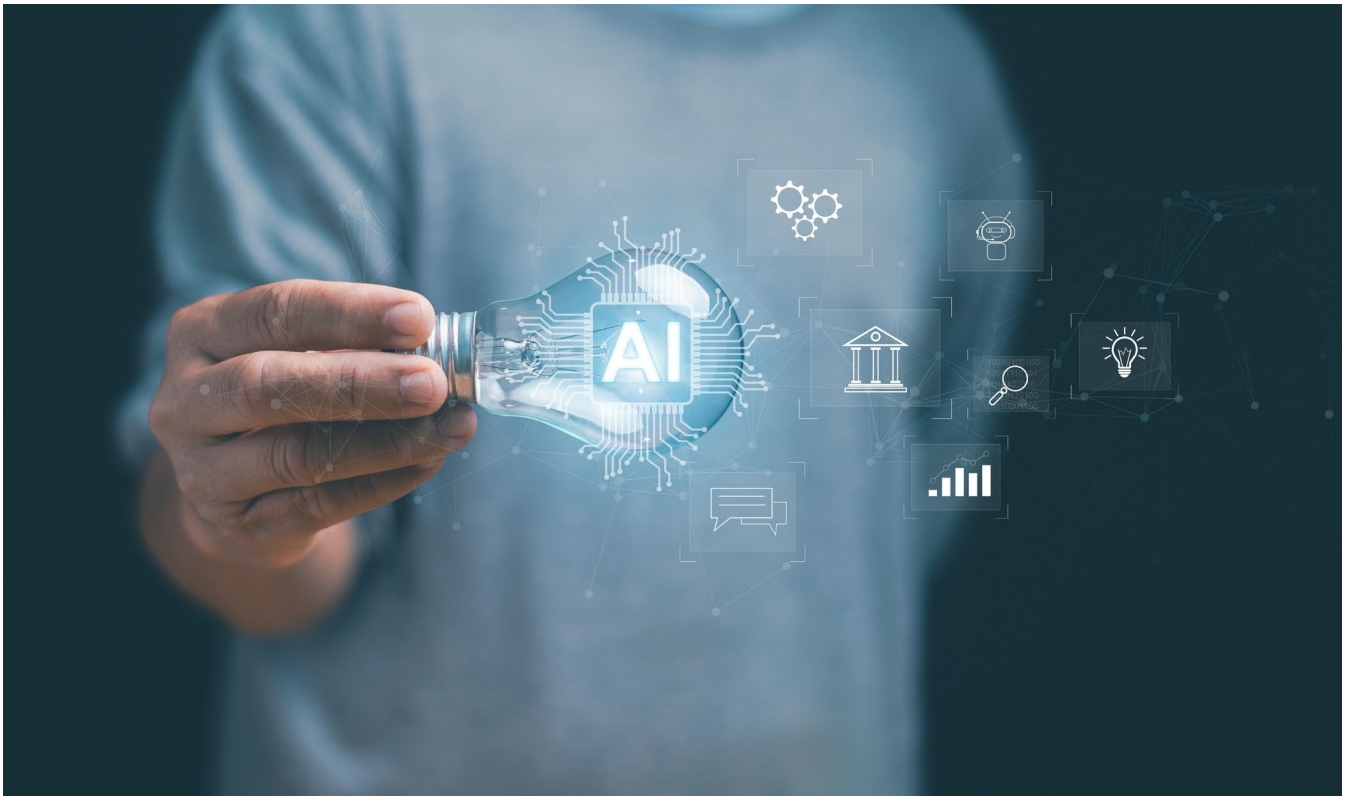
The report also points to a broader shift from module-level assembly towards wafer-level and heterogeneous integration. Technologies covered include 2.5D and 3D packaging, fan-out wafer-level packaging, micro-bump and Cu-Cu hybrid bonding, alongside photonic wire bonding and advanced fibre-to-chip coupling.

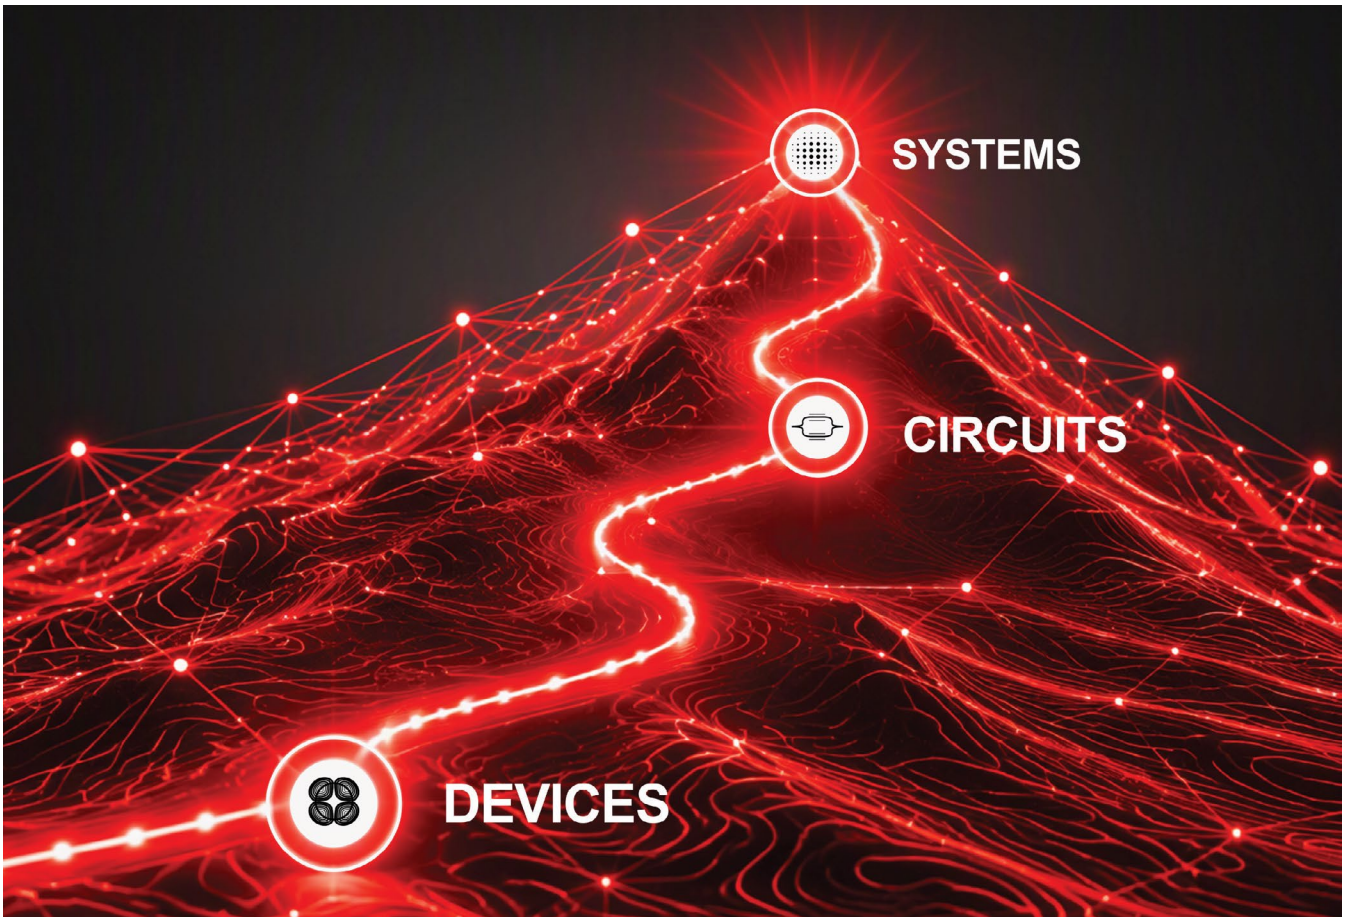
For photonic integrated circuits (PICs), the report examines the integration of electronic and photonic components across 2D, 2.5D and 3D architectures. It also covers optical engine configurations, external laser integration and packaging approaches for silicon photonics.

Beyond datacom and AI infrastructure, the analysis covers photonics packaging requirements for augmented reality displays, automotive FMCW LiDAR, quantum technologies, medical imaging, defence and industrial sensing.

Standardisation is identified as another factor influencing the adoption of photonics packaging at scale. The report highlights the ongoing development of process and assembly design kits, CPO fibre interfaces and electrical specifications linking switch ASICs with optical engines.

The study includes market forecasts through 2036, alongside analysis of the global supply chain, packaging technologies, regional ecosystems and competitive landscape. It profiles organisations across the photonics packaging value chain, including foundries, OSATs, equipment suppliers, PIC designers and module manufacturers.





Machine-learning-assisted device-to-system photonic design workflow



An integrated device-to-system simulation workflow combines electromagnetic modelling and machine learning to accelerate the inverse design and optimization of microring resonator demultiplexers for high-capacity data centre interconnect receivers.

BY EUGENE SOKOLOV, ANDRZEJ POLATYNSKI, HANNES LÜDER, SOHRAB SANGINI, ELIAS GIACOUMIDIS, ANDRÉ RICHTER, VPIPHOTONICS INC., USA; VPIPHOTONICS GMBH, GERMANY

THE RAPID development of hyperscale data centers over the past few years, driven by the massive growth of artificial intelligence (AI) and machine learning (ML) applications [1], has led to a search for new architectural solutions for optical interconnects, as traditional pluggable optical transceivers struggle to address the growing bandwidth and power consumption demands [2]. While co-packaged and near-packaged optics (CPO and NPO) solutions help meet industry requirements and improve scalability, challenges and opportunities for optimization remain [3], [4].

Shorter-reach intra-data-center interconnects (DCIs) commonly rely on direct-detection links based on wavelength-division multiplexing (WDM) and 4-level pulse amplitude modulation (PAM-4) [5]. High-capacity transmission is primarily supported through high-speed silicon microring modulators (MRMs) on the transmitter side [6]. On the receiver side, high-order microring resonators (MRRs) enable efficient, robust, and compact add-drop filters, making them particularly attractive as demultiplexers for WDM optical transmission systems [7]. In both cases, photonic integrated circuit (PIC) building blocks must

be carefully designed and optimized for their target application.

However, MRM- and MRR-based PICs are highly sensitive to fabrication tolerances, design choices, and environmental variations, including thermal fluctuations [8], [9]. This sensitivity is further amplified by CPO architectures, where photonic components are placed close to high-power electronics, introducing significant thermal crosstalk. As a result, device behavior and system-level performance become tightly interdependent, making their co-optimization a crucial task that requires understanding how overall system performance depends on the individual PICs and underlying photonic devices [10]. For high-order MRR filters, this task is especially challenging: the relationship between physical design parameters, such as coupling coefficients, and system-level metrics, such as passband bandwidth and shape factor, is highly nonlinear, and conventional optimization relies on iterative electromagnetic-based device-level or circuit-level simulations, resulting in high computational cost and long design cycles. Moreover, system designers typically specify filters in terms of performance metrics rather than the physical parameters that device designers can control.

Having appropriate design automation tools at each level does not by itself resolve this problem, since the overall workflow can become fragmented, requiring manual file exchanges, custom scripting, and repeated tuning iterations between circuit- and system-level teams. We present a unified simulation framework that seamlessly integrates component characterization, custom process design kit development, and photonic circuit design with optical system simulation [11]. To illustrate the proposed workflow, we use the MRR-based demultiplexers [12] for 800G WDM links as an example and show how deep neural networks (DNNs) can be effectively employed for rapid MRR optimization and inverse design.

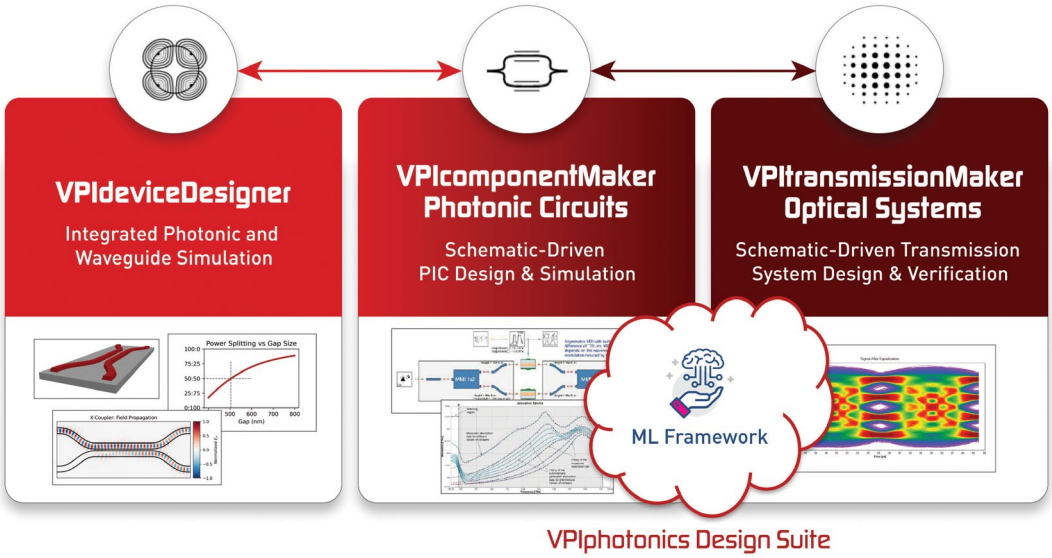
Device-to-system workflow

The core device-to-system workflow (Figure 1) consists of electromagnetic (EM) simulation of a photonic device, extraction of the EM solver response and import into a circuit-level compact simulation model, and validation of the resulting model in a target application [11].

At the device level, the 3D beam propagation method (BPM) and the eigenmode expansion (EME) method provide sufficient accuracy and reasonable simulation times for optimizing integrated photonic devices. The finite-difference time-domain (FDTD) method delivers the most accurate results and can serve as a rigorous final verification step, but its high computational and memory costs make it impractical for larger components and extensive optimization.

In contrast to device-level simulations, which may take several minutes for a single device configuration at a single wavelength even with fast BPM or EME methods, the corresponding circuit-level compact models are fast and typically require only a few seconds of wall-clock time. This advantage enables optimization routines and tolerance studies that require thousands of simulation runs and therefore accumulate significant total analysis time.

For devices well described by analytic or behavioral models, such as the waveguides and directional couplers that form an MRR, the most direct way to build a compact model is to extract device characteristics (e.g., effective index, coupling coefficient, loss) from EM solver results and pass them as parameters to a circuit-level model. These parameters are often wavelength-dependent arrays or functional expressions rather than single values, so both tools must support them natively. For passive linear devices that are harder to describe analytically, a full scattering-matrix (S-matrix) approach can be used instead. Both approaches require seamless, standardized interfacing between device- and circuit-level simulators, along with



➤ Figure 1: Block diagram of the cross-layer photonic design framework based on the corresponding simulation tools [13] – [17].

advanced post-processing of the exported data, including interpolation, extrapolation, regularization, and enforcement of passivity and/or causality.

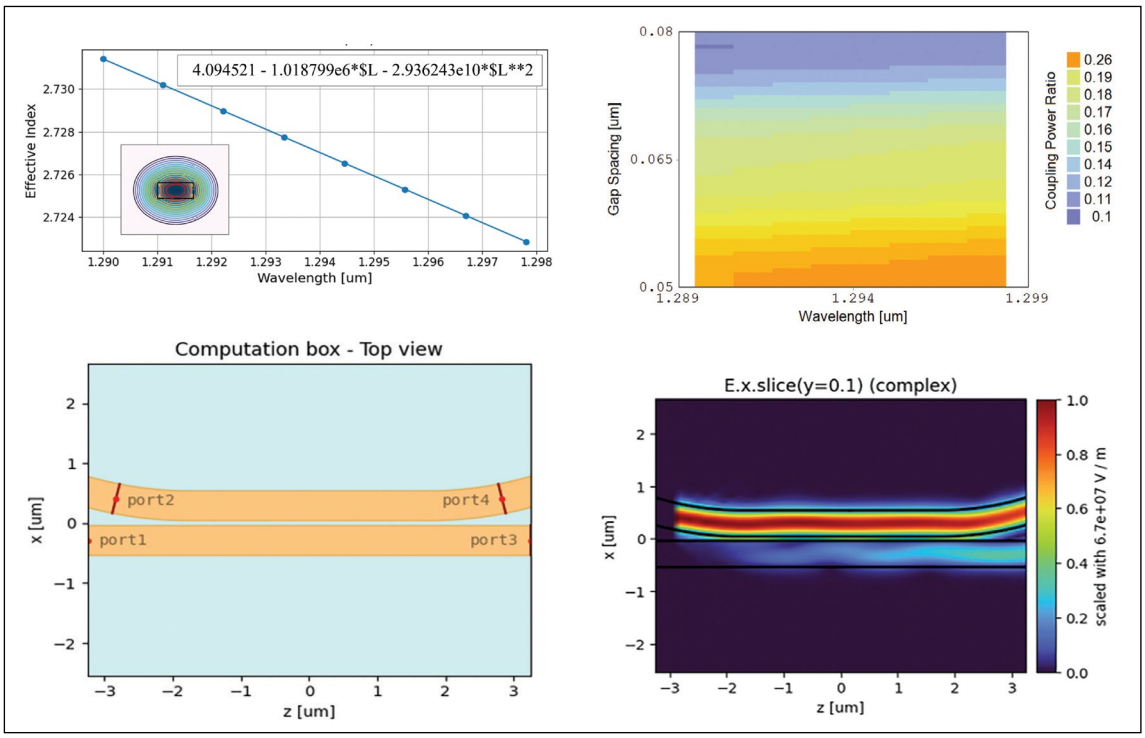
To complete the design process of the PIC components, it is crucial to evaluate their performance in a transmission system. Since the device/circuit design is often addressed by a different engineering team than the system design, there is a gap between the physical device properties (such as geometrical and material properties) and the system-level performance metrics (3-dB bandwidth, shape factor, extinction ratio, insertion loss, channel spacing, symbol error rate (SER), etc.). To bridge this gap, we enhanced the workflow with ML-based inverse design, translating between device- and system-level specifications. The speed of the compact models makes it practical to generate training data from tens of thousands of simulations, enough to train accurate DNN models. In inverse design mode, the DNN takes system-level specifications as input and directly returns the corresponding compact-model or layout parameters, replacing brute-force or heuristic parameter sweeps and reducing both computation time and manual effort. Because the DNN operates on the compact models rather than underlying device physics, the framework remains accessible to users ranging from component designers to system engineers.

ML-based inverse design of MRR demultiplexers for 800G WDM links

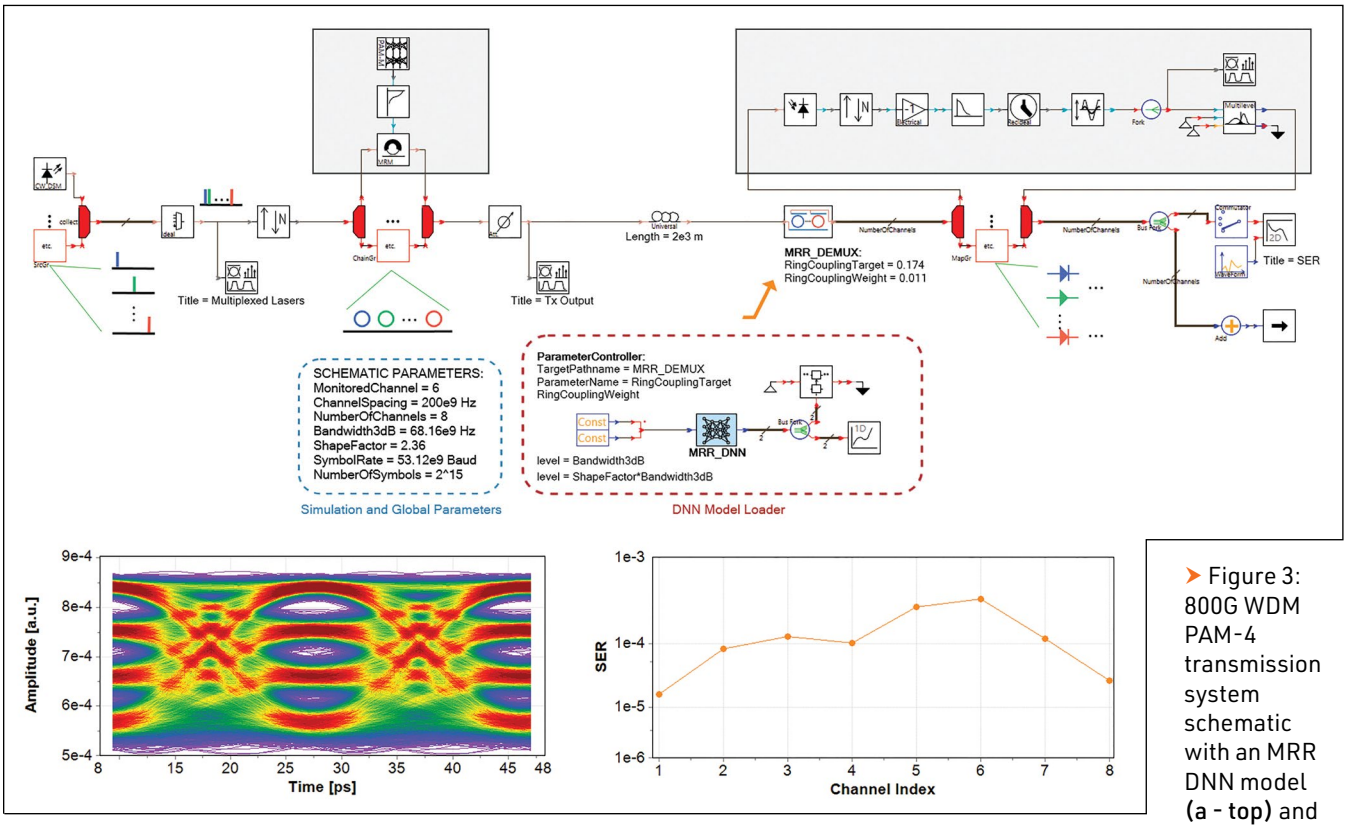
As an example, we implement and optimize an MRR-based add-drop filter used for signal demultiplexing in a WDM PAM-4 optical DCI [5], [6]. Following the workflow described above, we first characterize the devices comprising the MRR in VPIdeviceDesigner [13]. To simulate the ring-based structures, we split them into coupling regions, where ring waveguides are coupled either to bus waveguides or to adjacent ring waveguides, and uncoupled straight or bent waveguide regions.

The waveguide is a 500 nm × 220 nm silicon wire embedded in silica cladding. A finite-difference (FD) mode solver computes the effective index of the fundamental TE mode as a function of wavelength (Figure 2a). This dispersion curve is automatically fitted with a Python expression in wavelength, which can be pasted directly into the effective index parameter of the waveguide model in VPIcomponentMaker Photonic Circuits [15].

Further, the coupling regions are characterized with BPM (Figure 2b–d). The cross-coupling power ratio is automatically extracted from the T41 (or T32) wavelength-dependent port-to-port transmission elements of the device's S-matrix. Sweeping the gap from 50 to 80 nm across the full wavelength range gives the coupling power ratio as a function of both gap spacing and wavelength (Figure 2c). These



➤ Figure 2: Effective index vs. wavelength for an isolated waveguide and corresponding fitting as a Python expression with a placeholder $L for wavelength and mode field profile in the insets **(a - top left)**; ring-to-bus waveguide coupling region top view **(b - bottom left)**; calculated coupling power ratio vs. gap spacing for 8 channel wavelengths **(c - top right)** and an example of the corresponding BPM simulation results when light propagates through a ring waveguide **(d - bottom right)**.



➤ Figure 3: 800G WDM PAM-4 transmission system schematic with an MRR DNN model **(a - top)** and corresponding simulation results: eye diagram at the receiver for channel #6 **(b - bottom left)** and the SER performance for all 8 channels **(c - bottom right)** for the optimum MRR parameter settings.

values are exported as a text file and imported directly into the circuit-level model.

With the waveguide and coupling region characterized, the circuit-level model of the third-order MRR can be assembled [12]. Compared with single-ring designs, third-order microring filters provide improved passband flatness and greater out-of-band signal rejection. Adapting the filter to the demultiplexer requires two steps. First, the resonances are aligned with the target channel frequencies using a standardized swept-tuning routine. Second, the coupling power ratios are optimized to achieve the desired passband shape (bandwidth, shape factor, and ripple) and therefore the system performance. For a symmetric three-ring filter, only two coupling ratios are independent: bus-to-ring and ring-to-ring.

These two coupling ratios, or the corresponding gaps, could be optimized directly to minimize the SER at the receiver. However, system designers do not specify filters in terms of coupling ratios or gaps, but in terms of passband metrics such as the 3-dB bandwidth, the shape factor (the ratio of the 20-dB to the 3-dB bandwidth), and passband ripple, which together determine the signal distortion and resulting SER. Even when only the 3-dB bandwidth and shape factor are used as target metrics, the mapping from coupling ratios to passband metrics is highly nonlinear and generally nontrivial.

To obtain this inverse mapping, we trained a DNN model implemented in VPItoolkit ML Framework [17] that takes 3-dB bandwidth and shape factor as inputs and returns coupling power ratios. We ran the training simulation approximately 50,000 times to build up a dataset relating filter performance to physical coupling parameters, and then used it to train a DNN model and optimize its hyperparameters. For validation, the coupling parameters predicted by the DNN were applied to the physical model, and the resulting 3-dB bandwidth and shape factor were compared with the target values originally given as DNN inputs. The prediction error was below 1%, confirming that the DNN reliably maps performance targets to device parameters.

The trained DNN model is then embedded directly into the full WDM system simulation using VPItransmissionMaker Optical Systems [16]: an 800G WDM PAM-4 optical link with 8 channels at 200 GHz spacing, modulated by cascaded MRMs, transmitted over 2 km of single-mode fiber (SMF), and demultiplexed via MRRs before signal detection and SER estimation (Figure 3a). Instead of setting the demultiplexer's ring coupling power ratios manually, the MRR DNN model loader block determines them using the desired 3-dB bandwidth and shape factor as inputs, and a parameter controller module passes the DNN output values to the demultiplexer. The optimization can therefore operate directly on system-level filter specifications.

Using the built-in golden-section search algorithm, we then determined the 3-dB bandwidth and shape factor resulting in the lowest aggregate SER. The search converges to a 3-dB bandwidth of around 68 GHz and a shape factor of 2.4, yielding a predicted

ring-to-bus coupling target of 0.174 and a ring-to-ring coupling of 0.002. The resulting DNN-optimized filter design is validated at the full system level using standard PAM-4 WDM performance metrics. The simulation typically completes a single full-system run in ~30 seconds. Results include an eye diagram and SER across channels for the final optimized configuration (Figure 3b,c).

Conclusion

To meet growing bandwidth and power-efficiency demands, new optical interconnect architectures and their optimization are needed. The described workflow combining device-, circuit-, and system-level simulation capabilities provides a streamlined data flow between different levels of abstraction – EM simulation, compact circuit models, and full transmission systems – allowing individual devices to be optimized from the perspective of system-level performance metrics. The workflow's flexibility allows extension, as illustrated by using a DNN model for component inverse design.

The framework was applied to MRR-based demultiplexers in an 800G WDM PAM-4 link. Waveguide and coupling-region characteristics obtained from device-level EM simulations were transferred into circuit-level models of a third-order MRR filter, which were then evaluated and optimized within the full link simulation. A DNN, trained on these compact models, enabled inverse design of the ring coupling parameters from system-level filter specifications, with a prediction error below 1%. This example shows how the workflow carries device-level results all the way to system-level performance, and how separating model training from system-level optimization keeps the overall design process simple.

REFERENCES

- [1] Gelvanovska-Garcia, N. and Rossotto, C.M., [Advancing Cloud and Data Infrastructure Markets], World Bank Publications-Books (2024).
- [2] Hitchcock, I., and Cahoon, M., "Hyperscaler Data Center Buildout: A Sustainability Bane, Boon, or Both?" NI R 25-04. Durham, NC: Nicholas Institute for Energy, Environment & Sustainability, Duke University. https://nicholasinstitute.duke.edu/publications/hyperscaler-data-center-buildout-sustainability-bane-boon-or-both, accessed October 9, 2026.
- [3] Maniotis, P., Kuchta, D.M., "Exploring the benefits of using co-packaged optics in data center and AI supercomputer networks: a simulation-based analysis," J. Opt. Commun. Netw. 16(2), A143-A156 (2024); https://doi.org/10.1364/JOCN.501427
- [4] Tan, M., Xu, J., Liu, S. et al., "Co-packaged optics (CPO): status, challenges, and solutions," Front. Optoelectron. 16, 1 (2023); https://doi.org/10.1007/s12200-022-00055-y
- [5] Wang, Y., Novick, A., Parsons, R., Wang, S., Jang, K., James, A. Hattink, M. et al., "Scalable architecture for sub-pJ/b multi-Tbps comb-driven DWDM silicon photonic transceiver," Next-Generation Optical Communication: Components, Sub-Systems and Systems XII, vol. 12429, pp. 271-288. SPIE (2023); https://doi.org/10.1117/12.2649506
- [6] Yuan, Y., Peng, Y., Sorin, W.V. et al., "A 5×200 Gbps microring modulator silicon chip empowered by two-segment Z-shape junctions," Nat Commun 15, 918 (2024); https://doi.org/10.1038/s41467-024-45301-3
- [7] Schwelb, O., "Transmission, group delay, and dispersion in single-ring optical resonators and add-drop filters – a tutorial overview," J. Lightwave Technol. 22(5), 1380-1394 (2004); https://doi.org/10.1109/JLT.2004.827666
- [8] Padmaraju, K., Bergman, K., "Resolving the thermal challenges for silicon microring resonator devices," Nanophotonics. 3(4), 269-281 (2014); https://doi.org/10.1515/nanoph-2013-0013
- [9] Jayatilleka, H., Shoman, H., Chrostowski, L. and Shekhar, S., "Wavelength tuning and stabilization of microring-based filters using silicon in-resonator photoconductive heaters," Opt. Express. 23(19), 25084-25097 (2015); https://doi.org/10.1364/OE.23.025084
- [10] Amiralizadeh, S., et al. "Co-packaged optics technology evaluation for hyperscale data center fabric switches," 2025 European Conference on Optical Communications (ECOC), IEEE, (2025); https://doi.org/10.1109/ECOC66593.2025.11263202
- [11] Sokolov, E., Polatynski, A., Lüder, H., Sangini, S., Giacoumidis, E., Richter, A., "Integrated device-to-circuit simulation workflow for Si/SiN-based PICs," Proc. SPIE 14213, Optical Design Automation II, 1421306 (2026); https://doi.org/10.1117/12.3114849
- [12] Xia, F., Rooks, M., Sekaric, L., and Vlasov, Yu., "Ultra-compact high order ring resonator filters using submicron silicon photonic wires for on-chip optical interconnects," Optics Express, Vol. 15, No. 19, pp. 11934-11941 (2007); https://doi.org/10.1364/OE.15.011934
- [13] VPIdeviceDesigner, https://www.vpiphotonics.com/Tools/DeviceDesigner/, accessed October 9, 2026.
- [14] VPIphotonics Design Suite, https://www.vpiphotonics.com/Tools/DesignSuite/, accessed October 9, 2026.
- [15] VPIcomponentMaker Photonic Circuits, https://www.vpiphotonics.com/Tools/PhotonicCircuits/, accessed October 9, 2026.
- [16] VPItransmissionMaker Optical Systems, https://www.vpiphotonics.com/Tools/OpticalSystems/, accessed October 9, 2026.
- [17] VPItoolkit ML Framework, https://www.vpiphotonics.com/Tools/MLFramework/, accessed October 9, 2026.

EMPOWER THE
INTELLIGENCE REVOLUTION

ASMPT enabling the digital world



FROM PROCESS DEVELOPMENT TO HIGH-VOLUME MANUFACTURING **FOR SILICON PHOTONICS & CO-PACKAGED OPTICS**

asmpt.com
amicra.semi.asmpt.com

End-to-end process flow under one roof



Turning a photonic chip into a product means bringing together photonics, electronics, packaging and test. Handoffs between separate actors are where processes often stall. End-to-end process flow control under one roof may offer the best way through.

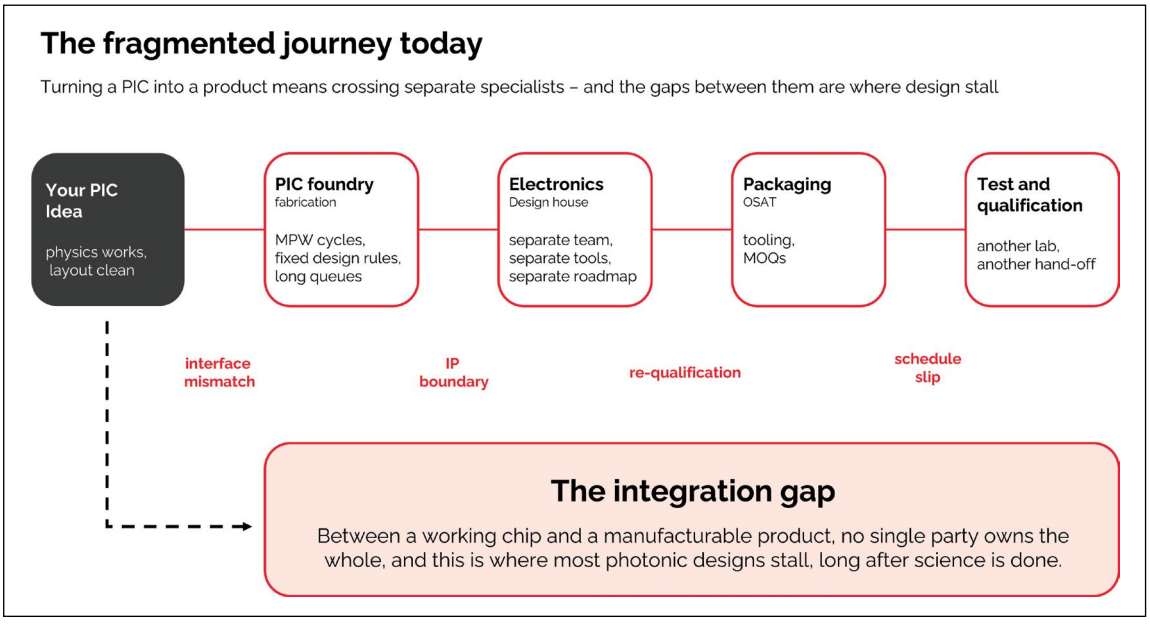
BY ANTONIO TEIXEIRA, CO-FOUNDER AND CSTO; PICADVANCED

FOR SIMPLICITY, let's consider a company that sees an opportunity and wants to see it developed, and this opportunity has in its core a PIC. The product can be, for example, a photonic integrated circuit that will move data faster, bring higher accuracy to sensing applications, or shrink into a single chip an optical system that today fills ten times the space. The process, as always, starts with a block diagram and some simulations, followed by a design that makes sense and is feasible in the chosen platform. On paper, the hard part looks finished: the physics holds up, the simulations agree, the layout is clean. However, when we start going forward, more and more difficulties start to pop up and quickly, it is understood that the chip design and concept were the easy part.

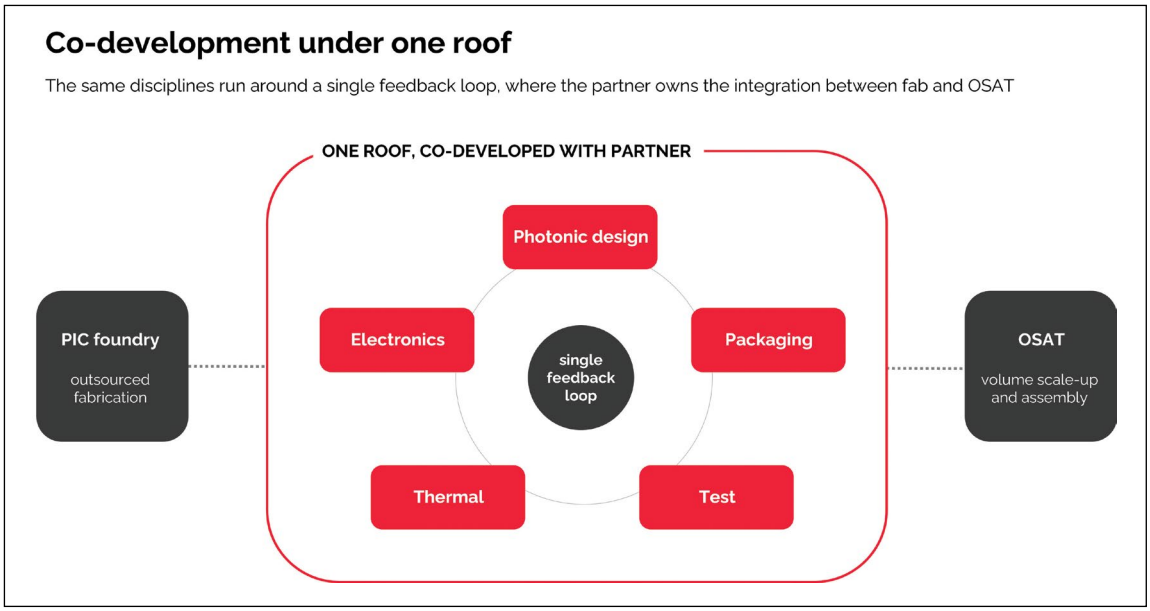
Building a photonic product is not the same as building a photonic chip, because a PIC will usually require electronic control and, eventually, optical I/Os. It needs electronics to drive its components and read its monitors, a packaging structure that holds optical and electrical connections in micron-accurate alignment, a thermal design that stops components from drifting, and a system or subsystem test regime that proves that the whole assembly behaves– not once at room temperature, but repeatably, across the product's working life, standards and client requirements. Each of those steps is a demanding discipline in its own right, and each, in today's industry, due to the required investments and specialization tends to sit inside a different company.

The handoff chain

The journey from idea to product runs as a relay race between specialists who rarely share a building or a schedule (see Figure 1). The PIC itself is fabricated by a foundry, on dedicated runs or multi-project wafer cycles, with fixed design rules, with variable waiting times and queues. The driving electronics come from a separate design house, which must understand the requirements while working with its own tools, roadmap and timeline. Typically, the assembly and packaging stages are also outsourced – sometimes to an outsourced semiconductor assembly and test provider, or OSAT – with specific tooling and minimum order quantities written for high-volume parts, and the more tailored and medium and low-volume packaging houses, besides being expensive, are not fully available and also have their own timelines. Testing and



➤ Figure 1: The fragmented journey – crossing several separate specialization domains, where interface mismatches, IP boundaries, minimum order quantities, re-qualification and schedule slip are part of the stalling mechanisms.



➤ Figure 2: Seamless multidisciplinary co-development under one roof – the same disciplines run concurrently around a single feedback loop, with the partner acting as accelerator between the outsourced foundry and the OSAT.

qualification may involve yet another party and may require some feedback to the packaging entity in order to improve and validate the final requirements. Only then does a system integrator try to make the result fit the application it was meant for with high-level system testing and comparison with the required results.

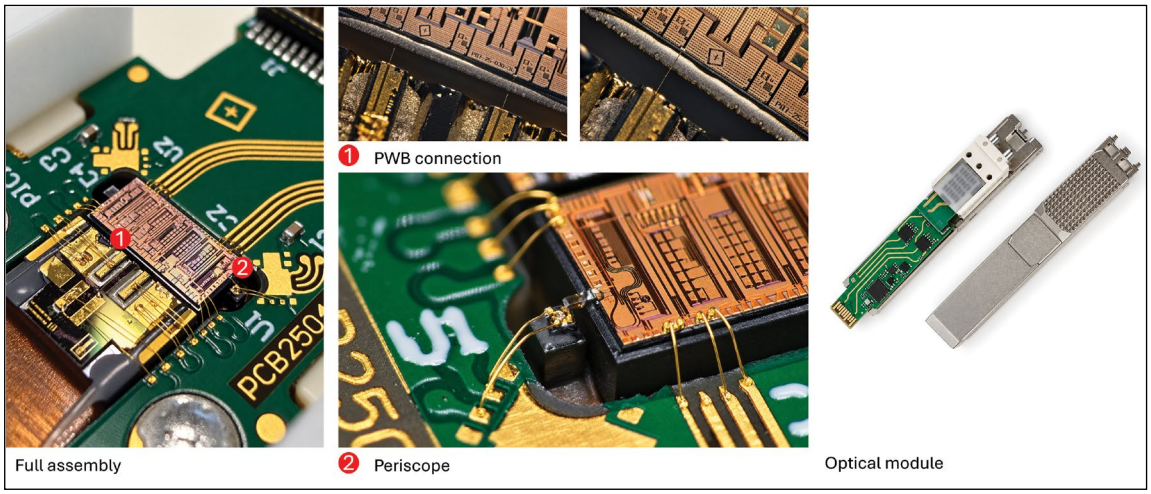
Although this approach is feasible and potentially valuable because each party specialises in its own domain, efficiency can be lost in the transitions, simultaneously in time alignment and requirement completion. Also, intellectual property, process control and process feedback for continuous improvement can be a huge barrier to a smooth and seamless flow. What has been seen as a simple project becomes very complex and demanding for the company specifying the product, as well as for the companies that are developing each of the processes, since they have to fully understand the requirements that are passed from different levels of complexity and also different languages and environments. Individually, none of these is fatal; collectively, they open the gap between a chip that works and a product that ships.

The backdrop of these interdependencies is that this obstacle course is particularly demanding for those trying to push boundaries in photonic integration. The organisations with the sharpest application insight (start-ups chasing a new architecture, system vendors with a specific performance target, network operators who know exactly what their traffic demands, research groups spinning out a breakthrough) are usually the ones least equipped to run this complex full cycle. They have the idea, but not a multidisciplinary team spanning photonics, electronics, packaging and test, nor the volumes that persuade a large-scale supplier to adhere to the journey of changing or adding new processes, which is a task required for most of these more specific developments. The established transceiver vendors, understandably, optimise for their own high-volume catalogues, and a customised design at low or medium volume is an awkward fit for a business built on scale. That leaves the innovator to become a project manager across multiple suppliers who do not talk to one another, and none of whom fully grasps the idea behind the product.

End-to-end process flow control and internal co-development, not coordination

The topic that we are addressing is already being adopted by several actors of each of the steps, especially by adding functions in the neighbourhoods. Packaging already incorporates some electronics development, as well as testing. However, these processes may be taken even further. It may start by refusing to treat design, electronics, packaging and test as a chain of separate contracts. Instead, they operate within a single internal co-development process: what we call working under one roof.

The distinction matters, because the real value lies in co-design throughout the different steps of the process. When the packaging engineer sits beside the photonic designer, package constraints shape the chip layout from the first day rather than the last, commonly defined as design for packaging (DFP). When the test engineer is in the room while the architecture is still being developed, the device is born testable instead of being forced onto a probe station months later, design for testing (DFT). When the electronics designer knows all implications stemming from the PICs and the surrounding electronics and packaging, e.g. the package's parasitic budget from the outset, the electronic control circuits are matched to the system or subsystem rather than compensating for it later. When the thermal modelling and the electronics design inform each other continuously, parameter drift is engineered out early rather than discovered late. Problems that would otherwise surface at a



➤ Figure 3: From chip to product: a photonic-based optical transceiver for access networks.

handoff, expensive and slow to fix, or even requiring a recall and redo of the full or partial flow, are caught while they are still cheap and, therefore, more suitable to be in time for the market window.

In practice, this means treating manufacturability and testability as architectural questions rather than downstream constraints. How a device will be packaged, aligned and measured shapes the design from the first day, and settling those questions alongside the PIC itself is what keeps them from becoming late-stage problems.

This is the model we have built our business around. Rather than acting as one more link in the relay, we position ourselves as an accelerator between the foundries and the OSATs (Figure 2): fabless, so the wafers are still made by the expensive and specialised fabs, and deliberately not chasing commodity volume, which the OSATs handle at scale. What we own is everything in between – architecture, electronics, advanced packaging, thermal management and test carried from concept through to a qualified pre-series. In other words, we take ownership of exactly the integration gap that used to swallow projects whole.

“Under one roof” does not mean that every manufacturing step must happen in one building. It means that one integrated engineering team owns the interfaces, the trade-offs and the outcome. Foundries and volume manufacturing partners remain essential parts of the ecosystem; the difference is that the customer no longer has to manage the engineering gaps between them.

Consider what this means in practice for an access-network transceiver, the kind of module that carries broadband into homes (Figure 3). Getting one to market the traditional way means aligning a photonic chip, custom driver and receiver electronics, a fibre-coupled package and a full qualification programme, each from a different supplier, each on a different timeline. Handled as a single co-development, the same job becomes one iterative conversation: the package is designed around the chip, the electronics around the package, the test plan around all three, with one team adjusting each as the others evolve. The result is a faster and smoother path to a pre-series that a real network can actually deploy, and the same logic scales up to the dense, high-speed interconnects that AI datacentres are now demanding.

The feedback loop

For us, co-development is not the same as coordination. A prime contractor who simply subcontracts the same fragmented chain still inherits every seam in it; the customer gains a single point of contact but none of the underlying integration. True co-development collapses the loop because design, foundry liaison, packaging and test share one team and one feedback loop. For example, a problem found in test on a Monday can reshape the layout by Friday, instead of triggering a three-month round trip across three companies, each waiting on the others and synchronizing their own roadmaps. Tighter loops mean fewer and faster iterations; fewer iterations mean lower cost and earlier time to market. And, above all, someone owns the outcome from end to end. The economics matter as much as the engineering. A late interface problem can trigger another fabrication cycle, a package redesign, new tooling or a fresh qualification round. Preventing one such iteration can be worth far more than optimising the cost of any individual component. The relevant metric is therefore not simply engineering cost, but the cost and time required to reach a manufacturable product.

There is a human dimension to this as well, because a chain of arm’s-length transactions loses a little information at every boundary, and a little trust with it. A co-development relationship works the other way: the customer’s deep knowledge of the application and the partner’s deep knowledge of integration reinforce each other, pulling in the same direction rather than negotiating across a contract. The customer keeps their intellectual property and their differentiation (the reason their product deserves to exist), while the partner supplies the plumbing that turns that differentiation into

something manufacturable. Neither has to become the other to get a product out of the door.

None of this works without real capability behind it. Our cleanroom and automated optical and electrical test stations (Figure 4), and the advanced packaging platforms that sit alongside them, are at the service of this model, not the model itself. The model is what turns a decade of assembled expertise into a way for someone else's idea to reach the market.

The fair objections

Two objections deserve a straight answer. The first is dependency. A chain of separate suppliers gives a customer leverage and a second source, while a single integration partner concentrates that risk in one relationship. The trade-off is real, and it is why the deliverable of a co-development should be a transferable package - a documented architecture, a qualified process and a test strategy that another party can pick up - rather than a black box only one company can reproduce. Nor does co-development imply keeping production under one roof indefinitely: the objective is to remove integration risk before scale, and once the architecture, package, test strategy and manufacturing process are stable, production can move to the manufacturing partner best suited to the volumes required. Integration risk is worth removing; lock-in is not worth acquiring.

The second is best-of-breed: a specialist at each step may well beat a generalist at that step. Often true, and eventually irrelevant if the parts do not meet. The argument for co-development is not that any single discipline is performed better under one roof, but that the interfaces between them are owned by someone. Where a design is conventional, the interfaces are standard and the volumes are high, best-of-breed wins on merit, and a commodity module built to a settled specification is better served by the scale economics of a large transceiver vendor. Where the interfaces are the hard part - a genuinely new design, low to medium volumes, risk concentrated in the seams – ownership of those seams beats excellence at the ends. That is precisely where most photonic products live today.

The bottleneck is shifting

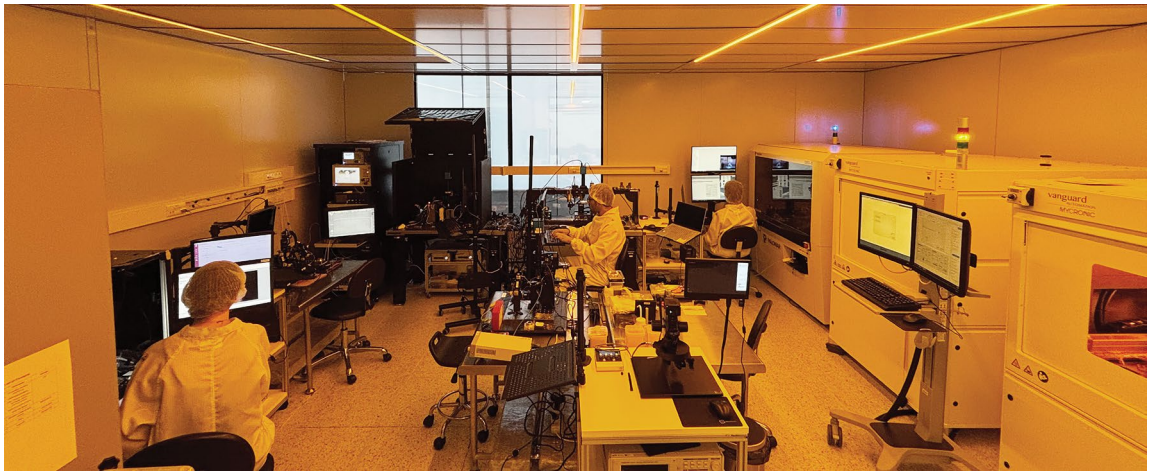
As photonic integration spreads, pushed by AI-driven bandwidth growth, co-packaged optics, optical sensing and the first stirrings of quantum photonics, the bottleneck in the field is moving. As foundry access, mature process platforms and process design kits make PIC development increasingly accessible, a growing share of product risk is migrating downstream, from proving the photonic function to integrating it into a reliable, testable, thermally stable and manufacturable system.

And manufacturable ultimately means manufacturable at acceptable yield. A design that performs brilliantly once but cannot be assembled and tested repeatably is not yet a product. Yield is therefore not simply a manufacturing metric; it is an architectural constraint that needs to enter the co-design loop early.

In the next decade, the organisations that thrive will not necessarily be those that can fabricate the cleverest chip, but those that can co-develop across every discipline that stands between that chip and a customer. That points towards a genuine democratisation of photonics.

The role of EDA tools is also a relevant matter, since they are the ones able to incorporate most of the processes in one single tool and therefore able to pass requirements in a standardized manner to the relevant actors.

A founder with a good idea and no cleanroom, or an operator with a precise need and no photonics team, should be able to bring a PIC-based product to market by plugging into a partner who carries the integration burden for them. Building a photonic product is still hard – the physics, the electronics, the packaging, the heat and the testing all remain hard-won. The difference is that the effort no longer has to be scattered across a dozen companies and several regions. It can live under one roof. And for the engineer with the genuinely good idea, that changes everything: the chip can go back to being the interesting part.



➤ Figure 4: Inside the pipeline: advanced packaging and testing in the PICadvanced cleanroom.

Barium titanate emerges as a route to faster AI optical links

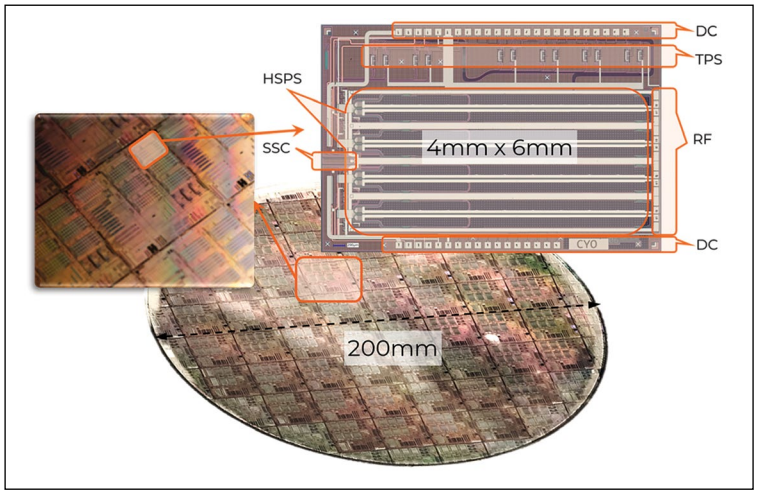


A barium titanate modulator built on a commercial silicon photonics line has carried a record net 1.6 T — and production-grade metal-oxide epitaxy is the quiet enabler that finally makes it manufacturable.

BY DR SABINA HATCH, SALES DIRECTOR, DCA INSTRUMENTS OY

AT THIS YEAR'S Optical Fiber Communication Conference (OFC) in Los Angeles, a post deadline paper delivered a result the silicon photonics community has chased for more than a decade. A four-channel transmitter chip, built on a commercial 200 mm silicon photonics line, pushed a net 1.6 terabits per second — four lanes of 448 Gb/s — through 2 km of standard fibre. Researchers from McGill University in collaboration with Ciena and the Swiss company Lumiphase achieved it not with an exotic laboratory assembly, but with a 5.1 × 3.5 mm^2 die whose active ingredient is a ferroelectric oxide: barium titanate, or BTO [1].

The same week, the same McGill–Lumiphase partnership reported a companion result: a BTO dual-polarisation IQ modulator carrying a net 1 Tb/s on a single wavelength, both over an 80 km link modelled on the ZR standard and over a 2 km "coherent-lite" link of the kind hyperscalers are eyeing for the inside of the data centre [2].

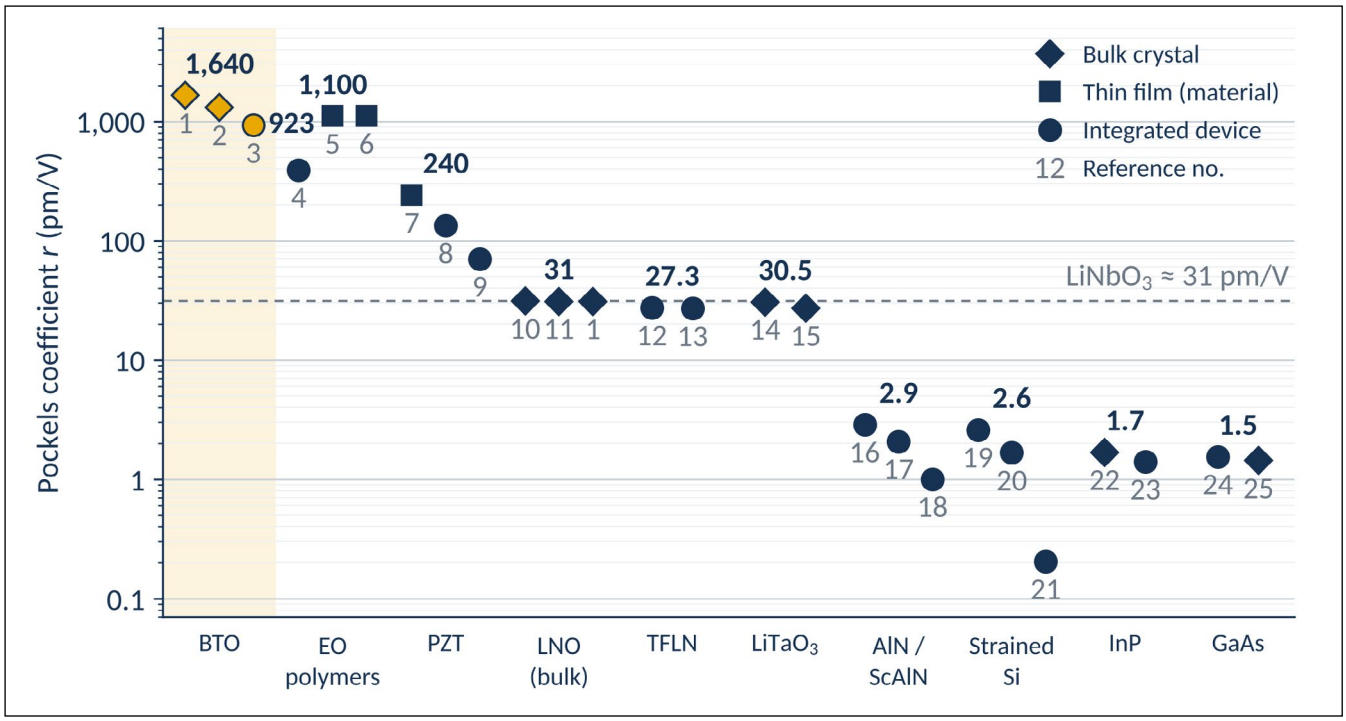


➤ Figure 1: A fully processed 200 mm BTO-on-silicon-photonics wafer, monolithically integrated at wafer level, with a zoom-in of a single reticle and a die shot of a C-band dual-polarisation IQ modulator supporting up to 130 GBd DP-16QAM in a 4 × 6 mm^2 die.

Why does this matter? Because AI has moved the bottleneck. Every new generation of accelerators multiplies the data that must flow between chips, boards and racks, and the optical links that carry that traffic must deliver more bits at lower energy and in less space with every product cycle. The laser, the fibre and the detector have largely kept pace. The component that has not is the modulator — the device that imprints data onto light. And the reason it has not is a materials problem that silicon photonics has carried since its beginning, roughly fifteen years ago.

Silicon remains an excellent platform for guiding light, and it comes with the full manufacturing weight of the CMOS industry behind it. The catch is its crystal symmetry, which rules out the Pockels effect — the clean, field-driven switching mechanism the best modulators rely on. Without it, silicon modulators have to shuttle charge carriers in and out of the light path instead, and that approach is inherently slower, lossier, and more prone to distortion.

Thin-film lithium niobate (TFLN) brought the Pockels effect back and pushed bandwidth to new records. But its Pockels coefficient is only moderate, so low-voltage devices still end up long. Lithium niobate also can't be grown directly on silicon — it has to be sliced from bulk crystals and bonded onto a wafer, a process most mainstream foundries simply aren't built for. Lithium itself brings contamination risks that make it a poor fit for standard CMOS lines. Indium phosphide modulators are fast, but they're stuck on small, costly 3- to 6-inch wafers.



➤ Figure 2: How the leading electro-optic materials compare. The Pockels coefficient – the strength of the linear electro-optic effect – of thin-film BTO integrated on silicon (≈923 pm/V) is roughly 30 times that of lithium niobate (≈30 pm/V), while strained silicon manages only a marginal effect; bulk BTO crystals exceed 1,000 pm/V. A larger coefficient permits shorter phase shifters, lower drive voltages and lower energy per bit.

What the industry has wanted all along is a material with a huge Pockels effect that can be manufactured on silicon, in a silicon fab. Barium titanate is that material — and the OFC results suggest its moment has arrived, in two markets at once: the terabit interconnects of AI data centres, and the fast, low-loss optical switches of photonic quantum computers.

Why barium titanate wins

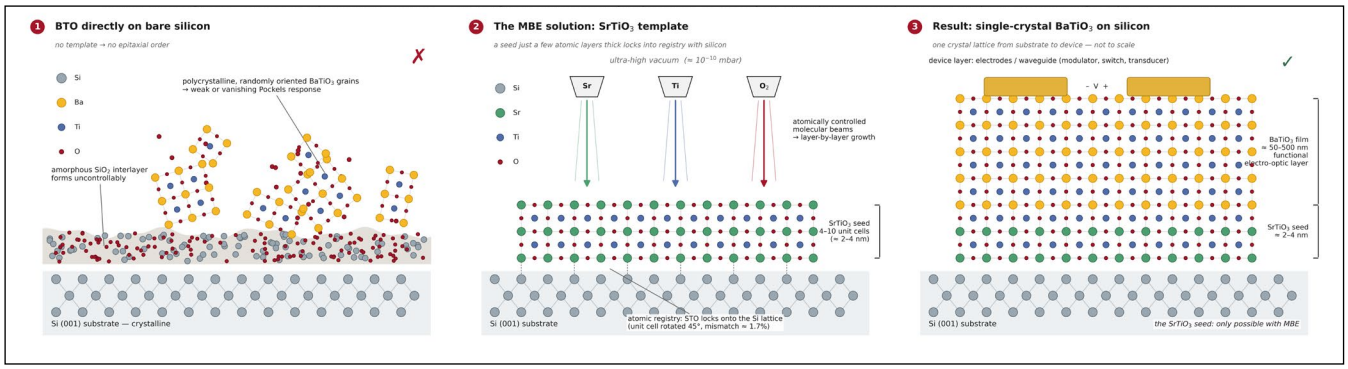
The Pockels effect is easiest to picture as a material whose refractive index — the "stiffness" it presents to light — shifts instantly in response to an electric field. Apply a voltage and the light travelling through speeds up or slows down by a whisker; remove it and the material relaxes, essentially instantaneously. Nothing physical has to move and no charge has to be injected or swept out, so the effect is fast, linear and adds almost no optical loss.

How strongly a material's refractive index shifts is captured by its Pockels coefficient, and this is where BTO is extraordinary. In 2019, a team led by IBM Research in Zurich measured a coefficient of roughly 923 pm/V in thin BTO films integrated on silicon [3], about thirty times that of lithium niobate — and demonstrated data modulation at 50 Gb/s. In bulk crystals the coefficient exceeds 1,000 pm/V, among the largest of any known oxide. A larger coefficient is not an academic nicety: it means the same optical phase shift can be achieved with a shorter electrode, a lower voltage, or both. Shorter means more channels per die; lower voltage means fewer joules per bit. For interconnects whose power budget is now the limiting line item in AI infrastructure, that trade is the whole game.

A fair-minded reader will ask about the fine print, because there is some. In 2025, researchers at ETH Zurich measured BTO and lithium niobate continuously from 100 MHz to 330 GHz and confirmed that BTO's Pockels coefficient, unlike lithium niobate's, falls with frequency [4]. Crucially, the same study showed that BTO remains a remarkable electro-optic material across the entire range, and that devices can be engineered for a flat frequency response despite the dispersion. The OFC hardware is the proof: the Lumiphase-built modulator in the 1.6 Tb/s demonstration shows a 6 dB electro-optic bandwidth of 105 GHz — well clear of any previous non-plasmonic BTO device — and maintained its performance at 85 °C, the kind of case temperature a pluggable module actually experiences, with only a marginal penalty attributed to the electronics. Thin-film BTO devices have even demonstrated penalty-free operation up to 250 °C [5], far above the 120 °C Curie point of the bulk crystal.

The hard part: growing a crystal that does not want to grow

If BTO is so good, why has it taken this long? Because the material is brutally difficult to make well. The Pockels effect lives in the perfection of the crystal: BTO must be deposited as an ordered, single-crystal-like film with the right atomic



➤ Figure 3: Why molecular beam epitaxy is indispensable. BTO will not grow as an ordered crystal on bare silicon; a strontium titanate (STO) seed layer, just a few atomic layers thick and grown by MBE in ultra-high vacuum, locks into registry with the silicon and presents the perovskite template on which crystalline BTO can then form. The schematic shows the layer sequence from silicon substrate to functional BTO device film.

arrangement, the right oxygen content, and clean interfaces. Deposited onto bare silicon, it simply will not comply — the atomic spacings and surface chemistry are wrong, and the film grows disordered, with its electro-optic strength largely lost.

The solution, established in the IBM-lineage research that underpins today's devices, is a seed layer. A film of strontium titanate (STO), only a few atomic layers thick, is first grown on the silicon surface. STO bridges the structural gap: it locks into registry with the silicon beneath while presenting a perovskite template above, onto which BTO then grows as a well-ordered crystal.

That seed layer is the crux of the whole supply chain because there is essentially one way to make it: molecular beam epitaxy (MBE). MBE assembles films in ultra-high vacuum, atom by atom, with independent control over each element's flux and over the oxygen environment — gentle and precise enough to coax a crystalline oxide onto silicon without oxidising the interface into useless amorphous glass. Faster deposition methods cannot match this feat. Sputtering, for instance, is a workhorse of the semiconductor industry and attractive for throughput, but it cannot nucleate the ordered STO template on bare silicon; where sputtering is used for the BTO layer itself, it relies on an epitaxially seeded template underneath. No MBE-grown seed, no crystalline BTO on silicon — it is as close to a hard dependency as materials science offers.



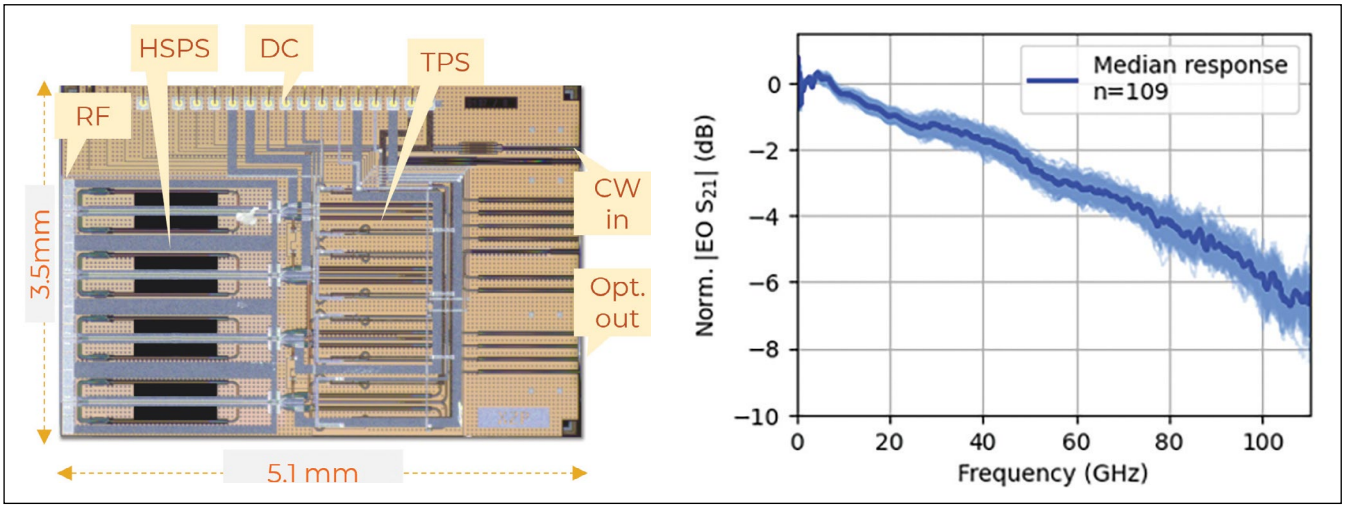
➤ Figure 4: Production metal-oxide MBE: DCA's M-series platform supports wafers up to 300 mm with an integrated equipment front-end module (EFEM), fully automated wafer transfer, 24/7 batch processing and closed-loop process feedback – the fab-compatible architecture that carries oxide epitaxy from laboratory to volume manufacturing.

For years that dependency looked like a dead end, because MBE carried a reputation as a delicate research technique: slow, manual, a tool for papers rather than products. That reputation is out of date. DCA Instruments, the Finnish ultra-high-vacuum specialist founded in 1989, has delivered more than 300 systems worldwide and in 2019 built the world's first production metal-oxide MBE for 300 mm wafers — the wafer size on which modern silicon photonics economics depend. Its production platforms pair the physics of MBE with the furniture of a modern fab: an integrated equipment front-end module for standard wafer handling, fully automated transfer and 24/7 batch processing, ultra-low-particle operation, and closed-loop process feedback. Production statistics presented by DCA at PIC International 2026 — three-sigma uniformity figures across 100-wafer runs of solid-source MBE-grown BTO — show what that translates to on the factory floor. And for the throughput question that always follows, a hybrid-MBE route that delivers titanium from a metal-organic precursor is raising growth rates, with reported device performance so far comparable to the all-solid-source process [6].

The pattern is familiar from other corners of the industry: a once-exotic technique matures into production infrastructure, and a material follows it out of the laboratory. EUV lithography needed its light source; BTO-on-silicon needed production-grade oxide MBE.

What terabit-class modulators buy

Return to the OFC numbers with that context, and their significance sharpens. The 1.6 Tb/s chip is not a hero experiment; it reads like a product preview.



Four Mach-Zehnder modulators with 1.75 mm phase shifters sit beside integrated germanium monitor photodiodes, on-chip terminations, and thermal tuners, in a die small enough for the tightly packed optics of a 1.6T pluggable transceiver. The electrical drive came from a 225 GBaud serialiser/deserialiser fabricated in 3 nm CMOS — exactly the class of silicon that will sit beside these modulators in real modules. All four channels ran 448 Gb/s PAM4 below the hard-decision FEC threshold over 2 km, the first time any silicon photonics platform has carried a net 400 Gb/s lane with PAM4, and pushing the same chip to PAM6 and PAM8 signalling yielded net rates of 2.02 and 2.16 Tb/s.

The coherent result matters for a different reason. As lane rates climb past 200 Gb/s, simple intensity-modulated links strain, and the industry is examining "coherent-lite" architectures — stripped-down coherent links that promise better sensitivity inside the data centre without the full power cost of long-haul signal processing. The BTO dual-polarisation IQ modulator demonstrated a net 1 Tb/s per wavelength in both an 80 km ZR-style system and a 2 km coherent-lite system, with a packaged half-wave voltage of just 0.4 V at 1 GHz. A detail worth savouring: on the coherent prototype, the operating points are held not by miniature heaters, as in conventional silicon photonics, but by BTO's own Pockels effect — under 5 mW, against 13 mW per phase shift for the equivalent thermal tuners. It is an option rather than an obligation: later iterations, including the IMDD transmitter, use silicon thermal tuners where that is the simpler choice. But the arithmetic changes with scale. A transmitter has a handful of phase shifters; an optical switch matrix has thousands, every one of them holding a bias indefinitely. There, static power is the design constraint, and a material that holds its state on a voltage rather than a current stops being a nicety and becomes the reason the switch fabric is buildable at all.

From material to market

Lumiphase, a Swiss company founded in 2020, has pioneered and industrialised the technology: epitaxial BTO grown on silicon by MBE is integrated at wafer scale into a 200 mm commercial silicon photonics flow alongside the silicon, silicon nitride and germanium layers of a standard platform. The OFC devices — the four-channel PAM4 transmitter and the dual-polarisation IQ coherent modulator — are based on epitaxial starting material grown on metal-oxide MBE systems supplied by DCA Instruments.

Lumiphase is not alone, and the breadth of the field is itself evidence that BTO has crossed from curiosity to platform. PsiQuantum, the Palo Alto quantum computing company, published its manufacturable photonic quantum computing platform in Nature in 2025 [7] — built in GlobalFoundries' 300 mm fab — and named BTO electro-optic switches as the key next-generation component for the fast, low-loss photon routing that fault-tolerant photonic quantum computing demands. Its MBE-grown BTO phase shifters achieved a loss–voltage product of 0.33 dB·V, a figure of merit that opens a path to the large switching networks a useful quantum machine will need. That integration work also reported a 95% functional yield for BTO devices in the 300 mm flow — about as strong a manufacturability statement as a new material can earn.

A third player closes the loop on supply. La Luce Cristallina, based in Austin, Texas, sells BTO-on-silicon as a product: in 2025 it opened a wafer fabrication facility — equipped with deposition systems from DCA Instruments — and now offers 200 mm BTO wafers and epitaxial STO-templated substrates to device makers who want the material without building an epitaxy operation of their own. Its route applies RF sputtering to the thicker BTO layer on an epitaxially templated surface, trading some growth finesse for throughput — a sign of a supply chain beginning to differentiate, the way mature materials markets do.

Three companies, three business models — a vertically integrated device maker, a quantum computing developer, and a merchant wafer supplier — on two continents, addressing two very

➤ Figure 5: The record-setting chip. Micrograph of the four-channel BTO/silicon-photonics transmitter die (5.1 × 3.5 mm²) presented at OFC 2026, alongside its measured electro-optic response – a 6 dB bandwidth of 105 GHz. All four channels achieved 448 Gb/s PAM4 below the hard-decision FEC threshold over 2 km, for a net 1.6 Tb/s.

different end markets. What they share is the same enabling step: precision metal-oxide epitaxy, on production-grade tools.

What comes next

The next chapter is about volume. The 200 mm demonstrations are done; the equipment for 300 mm exists; and the customers are explicit about their roadmaps — hyperscalers planning 448 Gb/s lanes and co-packaged optics, and quantum developers whose utility-scale machines call for millions of optical switches apiece. Industry forecasts cited by the wafer suppliers put the optical interconnect market above $34 billion by 2030, and every credible path to those volumes runs through foundry-compatible materials.

Expect three developments. First, BTO modulators will move from post deadline papers into transceiver and co-packaged-optics product programmes, where their short phase shifters and sub-volt drive directly attack the energy-per-bit and bandwidth density targets AI infrastructure now sets. Second, epitaxy throughput will keep climbing — hybrid MBE and sputter-on-template routes are early signs of a maturing toolbox — pulling wafer costs down the same curve every successful semiconductor material has ridden. Third, the quantum and datacom threads will reinforce each other: every yield point learned in a 300 mm photonics flow serves both.

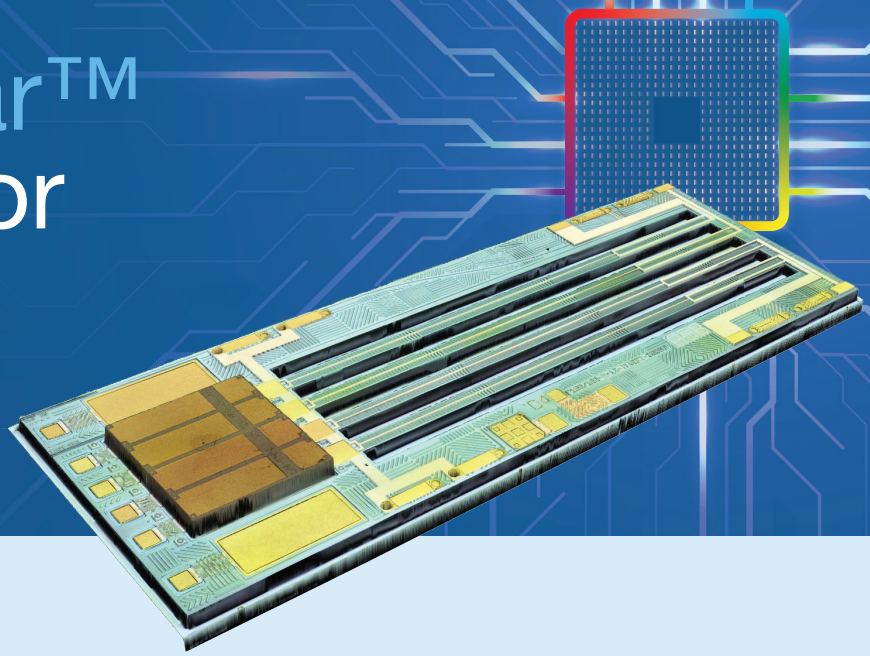
Fifteen years ago, silicon photonics struck a bargain: accept a mediocre modulator in exchange for the world's best manufacturing base. Barium titanate, grown by production metal-oxide MBE, dissolves that bargain. The light, at last, has a material that can keep up with the silicon — and the machines that grow it are already on the fab floor.

REFERENCES

- [1] B. Qiu, C. St-Arnault, J. Cauchon, C. Minkenberg, T. Kornher, C. Laperle, W. Diels, S. Abel, N. Ben-Hamida, B. Beggs, F. Pelletier, L. Czornomaz, X. Chen, F. Eltes and D. V. Plant, "Barium Titanate Enabling Net 1.6T (4x448 Gbps PAM4) On a Silicon Photonics Platform," in Optical Fiber Communication Conference (OFC) 2026, 2026.
- [2] B. Qiu, C. St-Arnault, S. Bernal, A. Nikic, K. Zhang, Y. Hu, P. Stark, W. Diels, F. Eltes and D. V. Plant, "Barium Titanate DP-IQM Enabling Net 1 Tbps/λ ZR and Coherent-Lite Data Center Networks," in Optical Fiber Communication Conference (OFC) 2026, 2026.
- [3] S. Abel, F. Eltes, J. E. Ortmann, A. Messner, P. Castera, T. Wagner, D. Urbonas, A. Rosa, A. M. Gutierrez, D. Tulli, P. Ma, B. Baeuerle, A. Josten, W. Heni, D. Caimi, L. Czornomaz, A. A. Demkov, J. Leuthold, P. Sanchis and J. Fompeyrine, "Large Pockels effect in micro- and nanostructured barium titanate integrated on silicon," Nature Materials, vol. 18, p. 42–47, January 2019.
- [FIGURE 2] [1] Boyd, Nonlinear Optics (2008); [2] Zgonik et al., Phys. Rev. B (1994); [3] Abel et al., Nat. Mater. (2019); [4] Kieninger et al., Optica (2018); [5] H. Xu et al., Adv. Mater. (2021); [6] H. Xu et al., Mater. Horiz. (2022); [7] George et al., ACS Appl. Mater. Interfaces (2015); [8] Ban et al., Opt. Mater. Express (2021); [9] Alexander et al., Nat. Commun. (2018); [10] Jazbinšek & Zgonik, Appl. Phys. B (2002); [11] Glazer & Cox, Int. Tables Crystallogr. D (2006); [12] Q. Li et al., Crystals (2022); [13] Chelladurai et al., Nat. Mater. (2025); [14] C. Wang et al., Nature (2024); [15] Casson et al., J. Opt. Soc. Am. B (2004); [16] T. Xu et al., Photonics Res. (2025); [17] G. Yang et al., Nat. Commun. (2024); [18] Xiong et al., Nano Lett. (2012); [19] Chmielak et al., Opt. Express (2013); [20] Chmielak et al., Opt. Express (2011); [21] Jacobsen et al., Nature (2006); [22] Suzuki & Tada, Jpn. J. Appl. Phys. (1984); [23] Bach et al., Appl. Phys. Lett. (1983); [24] Berseth et al., J. Appl. Phys. (1992); [25] Yariv & Yeh, Optical Waves in Crystals (1984).
- [4] D. Chelladurai, M. Kohli, J. Winiger, D. Moor, A. Messner, Y. Fedoryshyn, M. Eleraky, Y. Liu, H. Wang and J. Leuthold, "Barium titanate and lithium niobate permittivity and Pockels coefficients from megahertz to sub-terahertz frequencies," Nature Materials, vol. 24, p. 868–875, March 2025.
- [5] A. Messner, F. Eltes, P. Ma, S. Abel, B. Baeuerle, A. Josten, W. Heni, D. Caimi, J. Fompeyrine and J. Leuthold, "Plasmonic Ferroelectric Modulators," J. Lightwave Technol., vol. 37, p. 281–290, January 2019.
- [6] S. Hatch, "Enabling Foundry Ready BaTiO3 Photonics: Advances and 300mm Wafer Manufacturing," 2026. PIC International 2026, Brussels. Barium Titanate on 300 mm silicon | DCA Instruments Oy
- [7] K. Alexander, A. Benyamini, D. Black, D. Bonneau, S. Burgos, B. Burridge, H. Cable, G. Campbell, G. Catalano, A. Ceballos, C.-M. Chang, S. S. Choudhury, C. J. Chung, F. Danesh, T. Dauer, M. Davis, E. Dudley, P. Er-Xuan, J. Fargas, A. Farsi, C. Fenrich, J. Frazer, M. Fukami, Y. Ganesan, G. Gibson, M. Gimeno-Segovia, S. Goeldi, P. Goley, R. Haislmaier, S. Halimi, P. Hansen, S. Hardy, J. Horng, M. House, H. Hu, M. Jadidi, V. Jain, H. Johansson, T. Jones, V. Kamineni, N. Kelez, R. Koustuban, G. Kovall, P. Krogen, N. Kumar, Y. Liang, N. LiCausi, D. Llewellyn, K. Lokovic, M. Lovelady, V. R. Manfrinato, A. Melnichuk, G. Mendoza, B. Moores, S. Mukherjee, J. Munns, F.-X. Musalem, F. Najafi, J. L. O'Brien, J. E. Ortmann, S. Pai, B. Park, H.-T. Peng, N. Penthorn, B. Peterson, G. Peterson, M. Poush, G. J. Pryde, T. Ramprasad, G. Ray, A. V. Rodriguez, B. Roxworthy, T. Rudolph, D. J. Saunders, P. Shadbolt, D. Shah, A. Bahgat Shehata, H. Shin, J. Sinsky, J. Smith, B. Sohn, Y.-I. Sohn, G. Son, M. C. M. M. Souza, C. Sparrow, M. Staffaroni, C. Stavrakas, V. Sukumaran, D. Tamborini, M. G. Thompson, K. Tran, M. Triplett, M. Tung, A. Veitia, A. Vert, M. D. Vidrighin, I. Vorobeichik, P. Weigel, M. Wingert, J. Wooding and X. Zhou, "A manufacturable platform for photonic quantum computing," Nature, vol. 641, p. 876–883, February 2025.

POET

POET Blazar™
The Laser for the AI Era



- **Multi-channel external cavity laser** built on the patented POET Optical Interposer™
- **Designed to power both CPO and high-bandwidth**, chip-to-chip data communications links
- **OCI-compliant**, manufacturing-ready
- **300 milliwatts+ per channel via MOPA architecture** with single-mode operation and temperature stabilization
- **Separates the gain chip from the wavelength selection** — eliminating the power-law yield degradation that cripples DFB arrays

Customer sampling has begun. Contact sales@poet.tech to learn more.

OUR PARTNERS

MITSUBISHI NTT Adtran Lumilens

Visit Us at ECOC 2026 in Málaga, Spain!

poet-technologies.com
sales@poet.tech
#POETpowersAI

@POETtech
@POETTechnologies
@POETtechnologies

End-to-end nanoimprint lithography technology for photonic integration



Nanoimprint lithography enables high-precision, repeatable fabrication of micro-optics for photonic integration, with inkjet coating and accurate alignment supporting scalable manufacturing for CPO, microLEDs and VCSELs.

BY ANDREA KRONAWITTER, BUSINESS DEVELOPMENT MANAGER, EV GROUP

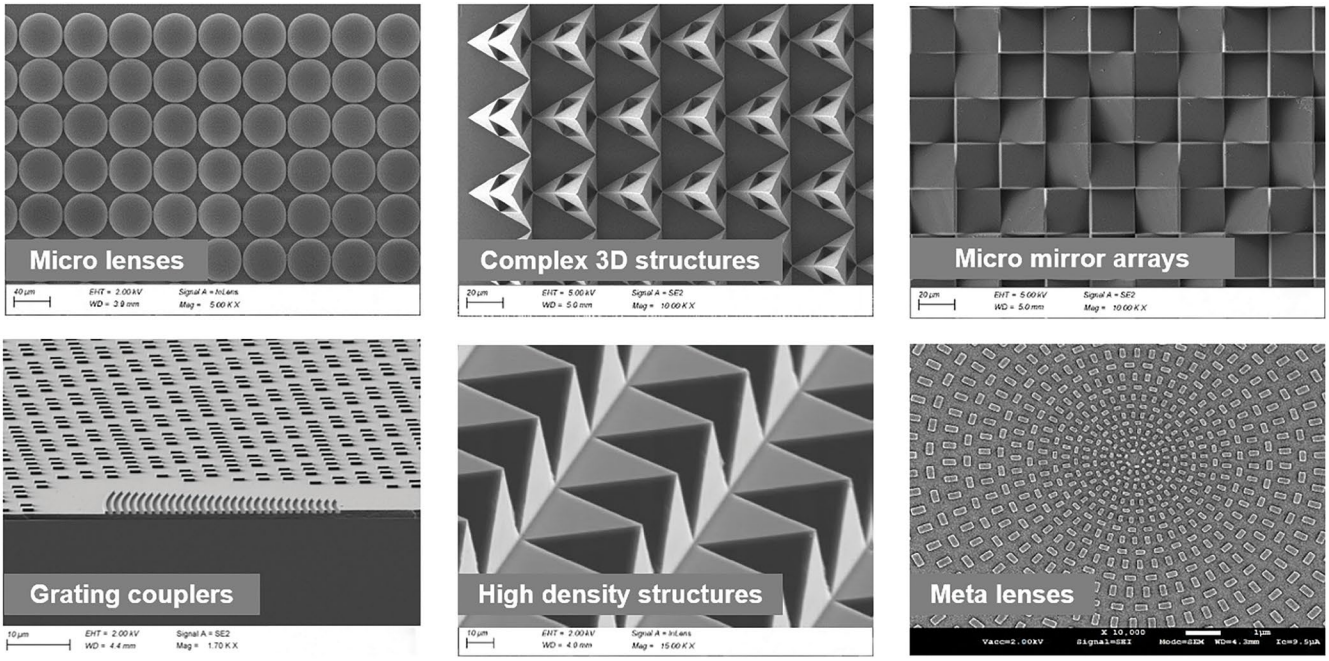
THE RAPID growth of optical interconnects, co-packaged optics (CPO), and AI-driven data infrastructure is accelerating the demand for highly precise and scalable micro-optical components. As photonic integration continues to increase, optical elements must not only deliver superior performance, but also be manufactured with the accuracy, throughput, and cost efficiency required for high-volume production.

Key components enabling advanced photonic integration include micro lens arrays (MLAs), diffractive optical elements, beam-shaping structures, and optical coupling features that improve light collection, transmission efficiency, and optical alignment. These optical structures are used in applications ranging from CPO to microLED displays, vertical-cavity surface-emitting laser (VCSEL) arrays, distributed-feedback (DFB) lasers, and photodetectors.

Standard optical lithography excels at manufacturing planar electronic devices through layer-by-layer patterning, but it is less suited for fabricating the three-dimensional structures that are essential for controlling light. Nanoimprint lithography (NIL) addresses this challenge by directly replicating micro- and nanostructures from a patterned mold into a functional imprint material. Unlike traditional lithographic approaches, NIL enables the formation of true three-dimensional optical structures in a single patterning step.

Advantages of nanoimprint lithography for photonics

The ability to replicate 3D structures is particularly valuable for photonic applications, where optical performance depends strongly on surface geometry, curvature, aspect ratio, and shape accuracy. Structures such as micro lenses, coupling elements,



➤ Figure 1: Various types of 3D structures that can be replicated with exceptional fidelity using nanoimprint lithography (Source: EVG).

mirrors, beam-shaping optics, and freeform optical surfaces can be replicated with exceptional fidelity while maintaining high throughput and excellent wafer-to-wafer repeatability (Figure 1).

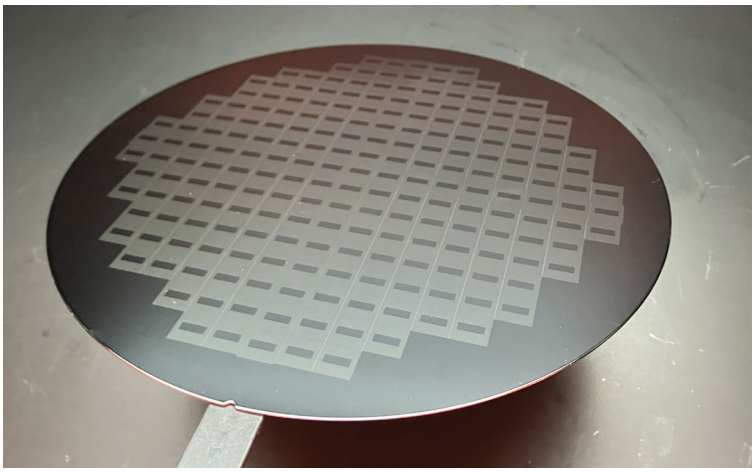
In addition, NIL is not constrained by the sophisticated optical projection systems required in advanced photolithography. As a result, it represents one of the most cost-effective approaches for manufacturing nanoscale and microscale structures across large substrate areas.

Compared with alternative fabrication techniques such as diamond turning, laser direct writing, or electron-beam lithography, NIL offers significant advantages in scalability, throughput, and cost of ownership while maintaining sub-micron resolution and excellent pattern fidelity [1]. The technology has already been proven as a high-volume-manufacturing capability in consumer optical sensors and AR waveguide manufacturing, and it is increasingly being adopted for advanced photonic packaging applications [2].

Inkjet coating enhances NIL performance

The performance of a nanoimprint process is strongly influenced by resist thickness uniformity and residual layer control. Variations in the residual layer can adversely affect both optical functionality and subsequent pattern transfer processes, particularly when NIL structures are used as etch masks (also referred to as sacrificial NIL).

To address these challenges, NIL can be combined with advanced inkjet coating technology. Rather than applying a continuous resist layer across the entire wafer, precisely metered droplets of imprint material are deposited only where material is required. This selective deposition approach enables highly uniform material distribution while significantly improving control over the residual layer thickness. The result is greater flexibility for different structure designs and substantially reduced material consumption.



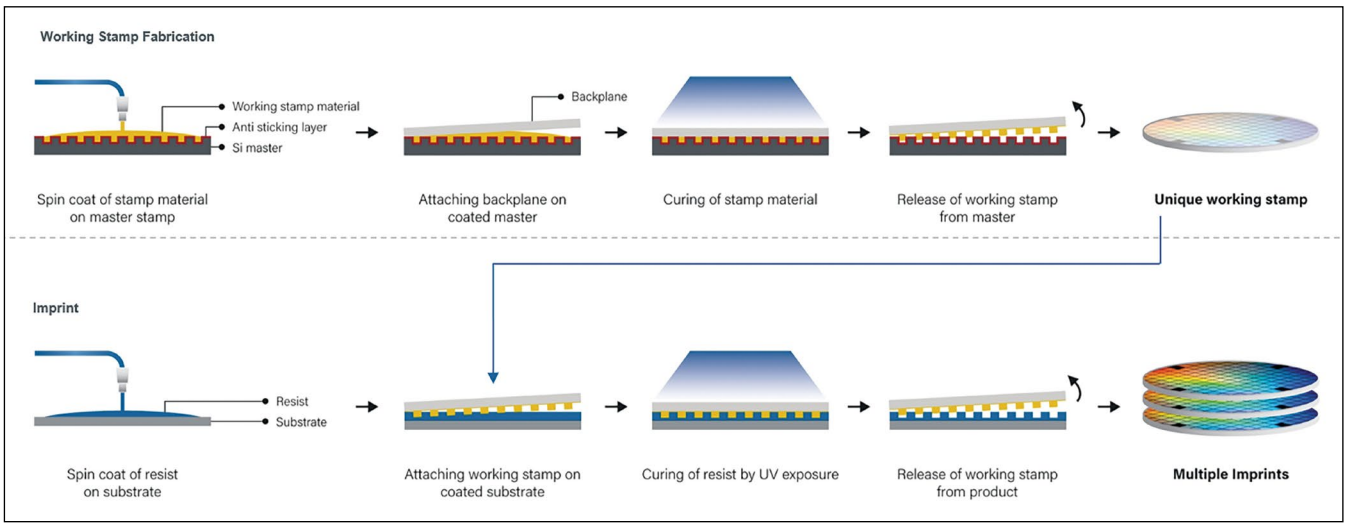
➤ Figure 2: MLA Master produced via reflow (Source: EVG).

High-precision optical packaging for microLEDs and VCSELs

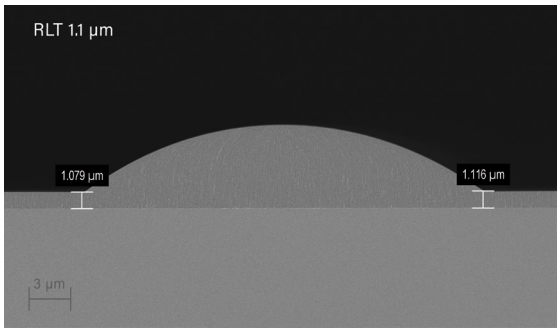
A key benefit of NIL is the ability to integrate optical structures with precise alignment relative to existing device features. Accurate imprinting of optical coupling elements onto silicon photonic devices, VCSELs, and microLED emitters is essential for maximizing optical efficiency and overall device performance.

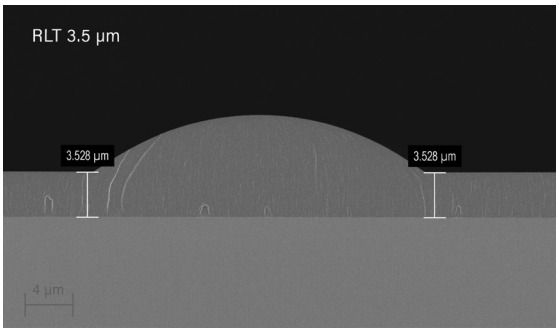
To demonstrate this capability, EV Group (EVG) began a feasibility study at its NILPhotonics® Competence Center, an open-access innovation incubator and process development facility located at its headquarters in St. Florian, Austria. First, EVG fabricated an EVG MLA reflow master, a highly precise patterned template containing lenses surrounded by bonding frames (Figure 2). Because NIL replicates the structures defined by the master, the quality of the master directly influences the quality of the resulting imprints. Any aberrations or unwanted surface roughness in the master can also be replicated during the imprint process, making high-quality master fabrication crucial to achieving the desired optical structures.

To produce the master, an EVG6200 mask alignment system was used to form pillars, followed



➤ Figure 3: EVG UV-NIL process flow.





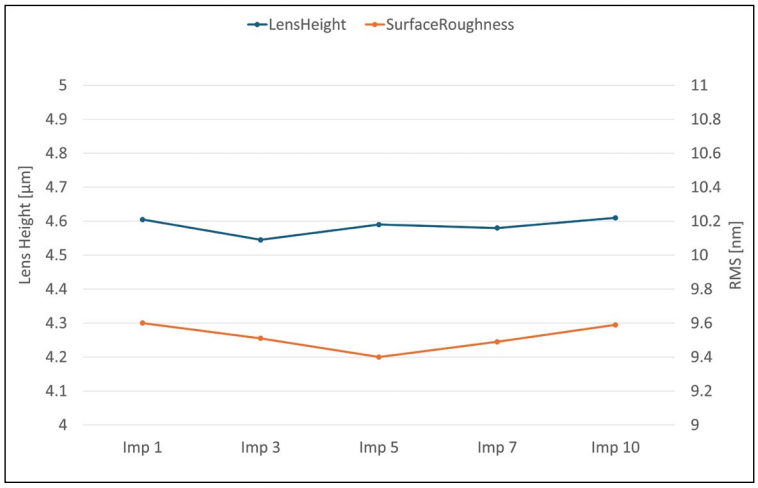
➤ Figure 4: Lenses with residual layer thicknesses of approximately 1.1 µm and 3.5 µm (Source: EVG).

by a reflow process to form extremely small (<5µm) lenses. Additionally, inkjet coating was applied for selective area coating to prevent the NIL polymer from covering electrical connecting points.

This master was then subsequently used for creating multiple imprints. As shown in Figure 3, the UV-NIL process involves pressing a patterned mold (working stamp) into the UV-curable resin. Once the resist has conformed to the molded topography, exposure to light solidifies the material. Subsequently, the mold is released, leaving behind a precise replica of the original structure.

The results of the study demonstrated that NIL technology enables imprinting of MLAs with a height of approximately 5µm with sub-micrometer alignment accuracy across the entire wafer, allowing precise MLA placement directly above individual microLED pixels. Such alignment performance is critical for these technologies, where even small placement errors can significantly impact optical efficiency and image quality.

The resist was coated via inkjet coating. Using inkjet coating, the residual layer thickness can be tailored to specific customer layouts and optical requirements, with the study demonstrating thicknesses of approximately 1.1 µm and 3.5 µm (Figure 4). Such precise control over residual layer thickness is also important for other photonic applications requiring sacrificial NIL processes, where a thin and uniform residual layer is essential for reliable downstream etching.



➤ Figure 5: Lens surface roughness and height over 10 imprints, demonstrating high reproducibility (Source: EVG).

In addition, the feasibility study focused on achieving a low surface roughness, which is important for minimizing unwanted light scattering and maintaining optical performance. A constant lens surface roughness of less than 10-nm RMS was achieved, with results across 10 imprints demonstrating the high reproducibility of the NIL process (Figure 5).

Taken together, these process capabilities contribute to more consistent device performance and enable the fabrication of highly uniform optical structures across the entire substrate, supporting both high-yield manufacturing and scalable production.

A unique capability of NIL is the simultaneous fabrication of optical and other support structures of different heights within a single imprint step. In addition to the MLA structures themselves, bonding frames have been replicated concurrently during the same process cycle (Figure 6). This eliminates additional lithography and alignment steps, simplifies the overall process flow, and improves manufacturing efficiency while maintaining high placement accuracy between optical and mechanical features.

Beyond microLED packaging, the same manufacturing approach can be applied to VCSEL arrays, where accurately aligned micro-optics improve beam shaping, optical coupling efficiency, and overall system performance. NIL therefore provides a versatile platform for a broad range of photonic packaging applications.

End-to-end process development from master to device

A key differentiator of EVG's approach is the ability to provide a complete NIL ecosystem. Rather than focusing solely on the replication step, EVG supports the entire process chain, including in-house master fabrication, working stamp replication, imprint process development, material optimization, dedicated equipment, and process transfer to manufacturing. Having these capabilities available within a coordinated process development environment can be particularly important for advanced photonics applications, where, for example, specialized requirements may make it difficult to source suitable masters for high-resolution structures and extremely small lenses.

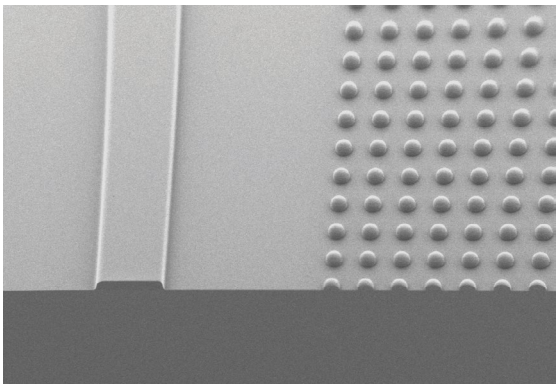
Through its NILPhotonics Competence Center, EVG provides in-house master fabrication as part of the overall process development flow. This enables customers to significantly reduce development cycles while ensuring seamless process integration from R&D to high-volume manufacturing.

Materials development provides another example of the benefits of this integrated approach. By bringing materials development and process optimization together within the same development environment, EVG can more quickly optimize materials and process parameters to meet the requirements of specific NIL applications. At the same time, EVG maintains an open materials platform that enables it to work with third-party materials suppliers as well as leverage its in-house materials development capabilities.

This end-to-end approach also enables multiple process requirements to be addressed together as photonic devices move toward greater levels of integration. For example, the study described here demonstrates the simultaneous fabrication of optical and bonding frame structures of different heights, with MLAs and bonding frames replicated within the same process step. EVG can also leverage expertise developed through its wafer bonding business in areas such as high-accuracy alignment and process control and apply this know-how to NIL to support deeper integration of NIL into photonic device and packaging processes. This cross-functional expertise becomes increasingly important as advanced NIL applications require more demanding alignment accuracy performance.

Conclusion

As photonic integration expands to support applications such as CPO and optical interconnects, manufacturers will increasingly need processes that can translate complex optical designs into high-volume production while meeting demanding performance requirements. This includes the ability to fabricate high-resolution 3D structures with precise alignment, while maintaining consistent surface roughness and residual layer thickness with high reproducibility from wafer to wafer.

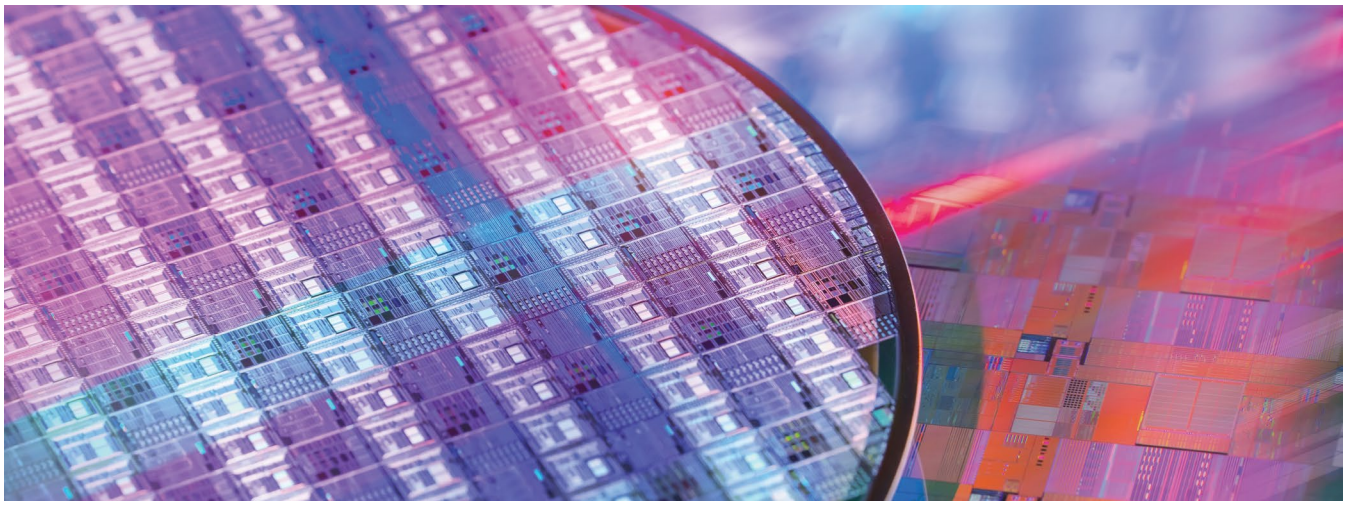


➤ Figure 6: MLAs and bonding frames imprinted in the same process step (Source: EVG).

Nanoimprint lithography provides a scalable manufacturing approach capable of meeting these requirements, but achieving the necessary precision, repeatability and process control requires more than imprint equipment alone. It requires expertise across the complete process flow, from master and stamp development through materials, coating, imprinting and process transfer to manufacturing. Building on decades of NIL experience and the process-development capabilities of its NILPhotonics Competence Center, EVG provides this end-to-end approach. The micro lens array study described here demonstrates how these capabilities can be combined to achieve the precision, repeatability and process control needed to move advanced photonic devices from development toward high-volume manufacturing.

REFERENCES

- [1] C. Thanner, R. Breyer and J. Litterscheidt, “Optimized UV Nanoimprinting Processes for Fabrication of High Fidelity Patterns,” 2020 International Semiconductor Conference (CAS), Sinaia, Romania, 2020, pp. 129-132, doi: 10.1109/CAS50358.2020.9268031.
- [2] M. Eibelhuber, C. Thanner, D. Treiblmayr, et al. “Towards AR waveguides with refractive index 2.0 utilizing nanoimprint lithography,” Proc. SPIE 11765, Optical Architectures for Displays and Sensing in Augmented, Virtual, and Mixed Reality (AR, VR, MR) II, 117650H (29 Mar 2021); https://doi.org/10.1117/12.2579505



2027

Sheraton Brussels Airport Hotel

Brussels, Belgium

12-14 April 2027

ANGELTECH

PIC INTERNATIONAL CONFERENCE

REGISTER NOW

Photonic Integrated Circuits International Conference 2027

The 12th PIC International builds on the huge success of the 2026 event and the strengths of its previous editions, with around 40 leaders from industry and academia delivering presentations across five key themes:

Conference attendees will gain greater insight into device technologies while learning about the latest opportunities and trends within the PIC industry. Delegates will also discover significant advances in tools and processes that deliver enhanced performance and higher yields.

As the longest-running PIC-dedicated event, PIC International hosts key decision makers from across the PIC supply chain, from materials and equipment suppliers, foundries and advanced packaging houses, module and subsystem suppliers, through to OEMs and equipment makers.

picinternational.net

80+ EXHIBITORS

800+ ATTENDEES

120+ PRESENTATIONS

4 CO-LOCATED CONFERENCES

SCAN ME



PIC 2027

The European Event for The Semiconductor Supply Chain

REGISTER NOW

Conference Themes

- Industrialisation and global scaling of PIC manufacturing
- Bridging the gap between innovation and adoption
- PICs as an enabling platform across industries
- Evolution of the PIC value chain and manufacturing ecosystem
- Defining standards, reliability, and commercial readiness

Why Attend?

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

THE WHOLE SEMICONDUCTOR SUPPLY CHAIN UNDER ONE ROOF

PIC INTERNATIONAL CONFERENCE | AP INTERNATIONAL CONFERENCE | PE INTERNATIONAL CONFERENCE | CS INTERNATIONAL CONFERENCE

SPEAKERS AND SPONSORS FOR 2027

CONTACT INFO | E: INFO@PICINTERNATIONAL.NET W: PICINTERNATIONAL.NET

Angel BUSINESS COMMUNICATIONS

Why AI infrastructure must compute while data moves

AI's next infrastructure bottleneck is not arithmetic alone but data movement. Optalysys explores how photonic Compute-in-Transit could turn the data path itself into a useful computational resource for future AI systems.

BY FLAVIO BERGAMASCHI, VP ALGORITHMS & IMON KUNDU, DIRECTOR OF PHOTONICS, OPTALYSYS





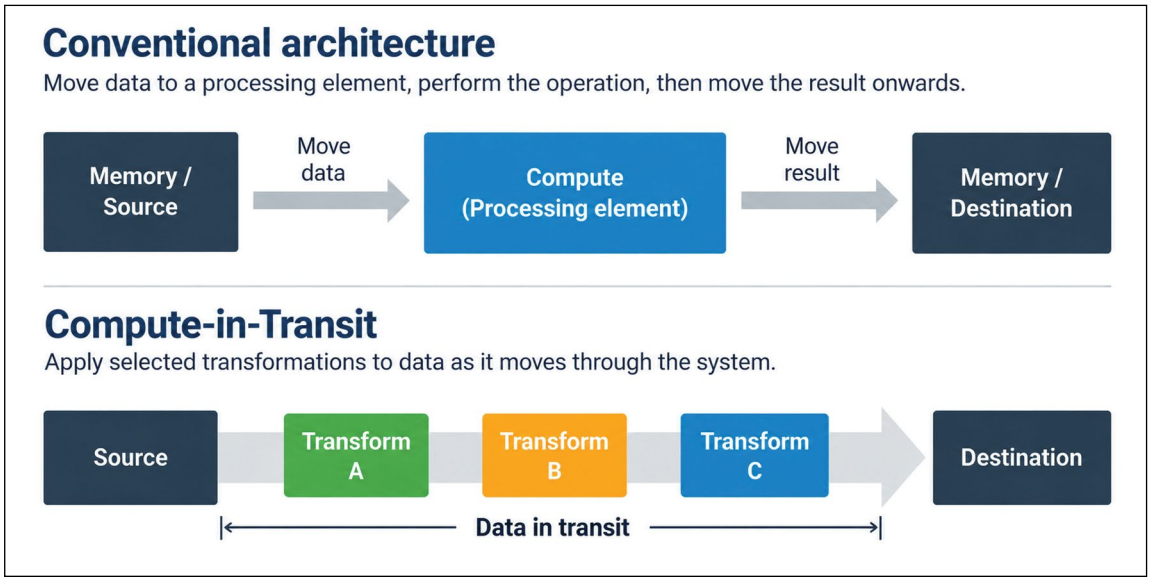
ARTIFICIAL INTELLIGENCE has become one of the strongest drivers of computing infrastructure. GPUs have grown more capable, memory systems faster and interconnects wider as increasingly large models demand more from every part of the system.

But AI scaling is increasingly constrained by data movement and memory limits. Improving the rate at which we perform arithmetic only addresses part of the problem.

AI workloads create and transform large volumes of intermediate data — data that must move between memory and processing elements and, increasingly, across distributed computing resources. The cost of moving data is becoming as much of an architectural concern as the cost of computing on it.

What we need are smarter systems.

Most computing infrastructure follows a familiar pattern: move data to where computation takes place, perform the operation, then move the result onwards. Faster memory, specialised accelerators and higher-bandwidth interconnects have made that process considerably more



➤ Figure 1: Conventional architectures move data to discrete processing elements; Compute-in-Transit applies selected transformations along the data path itself.

efficient, but the underlying relationship remains largely unchanged.

Computation and data movement are still treated as distinct, separate activities.

This raises a different architectural question: what if some computation could happen while the data was already moving?

Making the data path do useful work

Compute-in-Transit is an architectural model in which selected transformations are applied to data as it moves through a system, rather than only once it reaches a separate processing element.

It's not just about moving computation closer to the data, as in near-memory or in-network approaches, but developing new compute capability as the data path itself becomes part of the computation.

Rather than always asking how quickly we can transport data to another processor, we ask: which transformations could take place on the way? For which parts of the workload can movement and useful computation become part of the same process to ease the burden on processors?

Finding the right fit in an AI workload

So where would the Compute-in-Transit approach make sense?

An AI model contains many different types of computation, dependencies, and intermediate data, and they do not all suit the same execution model. Looking at individual operators tells us what computation is taking place, but it does not necessarily tell us enough about how data moves between operations.

Instead, we can examine the workload as a dataflow.

Mapping a workload in this way reveals where data enters a region of computation, which intermediate results are created and consumed inside it, what data leaves, and which operations take place along the way.

That makes it possible to identify areas where compatible transformations could be organised into a continuous stream.

For suitable parts of a workload, serial and vectorised photonic computing could then apply selected transformations as data propagates through the transport path, rather than transporting it to another processor, carrying out the operation and moving it again.

Compute-in-Transit is not intended to force an entire AI workload into a new architecture — it offers another execution model for the portions of a workload whose dataflow makes it useful.

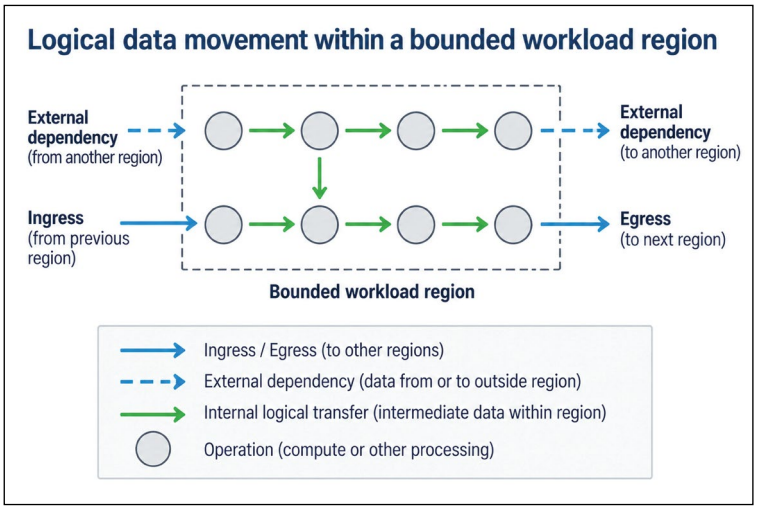


Figure 2: A bounded workload region can be analysed in terms of ingress, external dependencies, internal logical transfers and egress. This exposes where data is created, consumed and moved before any architectural mapping is attempted.

How photonics makes this possible

Photonics offers unusual properties that are well suited to this type of architecture.

An optical signal carries information in both amplitude and phase, and passive photonic structures can transform those properties as the signal propagates.

Interference and diffraction can therefore implement linear mathematical operations within the optical path itself. Rather than requiring a signal to be received, processed electronically and then transmitted again, selected transformations can occur in the same physical path that is already carrying the data.

Electronics remain essential; control, programmability, error management and nonlinear functions still need electronic logic. The opportunity comes from combining the strengths of both domains.

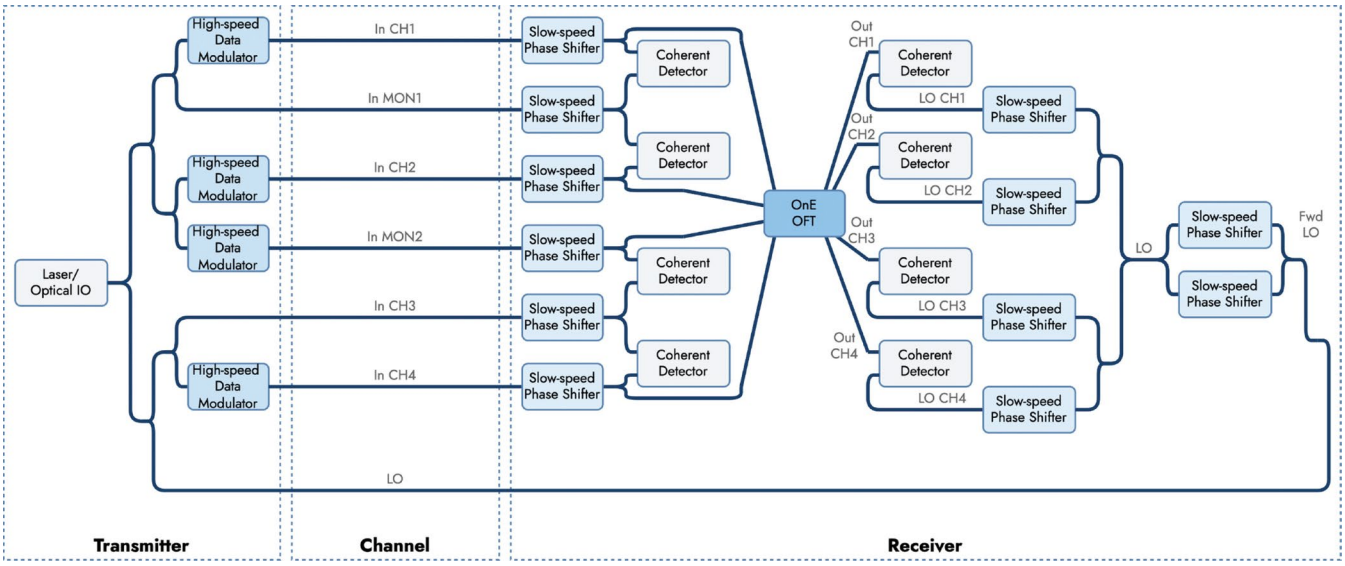
Where a workload exposes compatible vectorised dataflows, blocks of data can be presented as streams to compatible mathematical operations while remaining in transit.

For those parts of the workload, the transport path stops behaving purely as a link between processors. It begins doing useful computational work of its own.

Putting Compute-in-Transit into hardware

At Optalysys, we have been developing the technologies required to put this architectural idea into hardware.

An optical signal carries information in both amplitude and phase, and passive photonic structures can transform those properties as the signal propagates.



➤ Figure 3: Serial photonic computing combines streamed data movement with programmable transformations in the optical path, while electronic logic retains control and supporting functions.

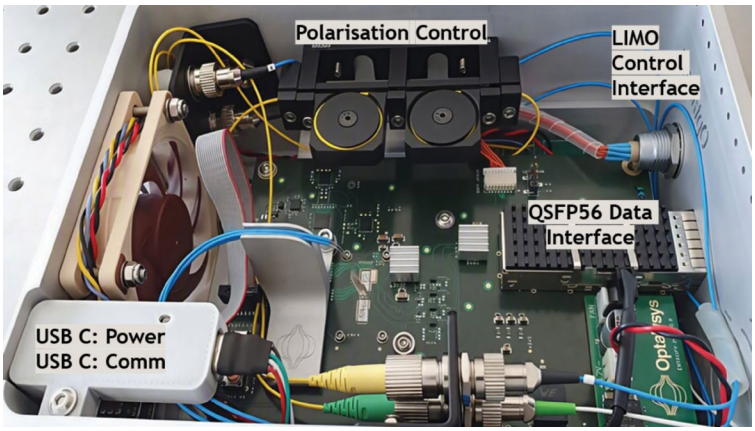
Our current prototype packages a photonic integrated circuit into a pluggable module connected through a standard QSFP56 port to an AMD Alveo V80 FPGA.

At its heart is our proprietary passive diffractive device, which performs an optical Fourier transform.

High-speed modulators, phase control and balanced photodetectors form the photonic path, while FPGA-based electronic logic manages serialisation, synchronisation, backpressure and error correction. A microcontroller handles calibration and signal-integrity control.

Together, these components create a programmable serial pipeline capable of supporting operations including Fourier and number theoretic transforms, multiplication, multiply-accumulate and modular arithmetic.

The prototype has processed hundreds of gigabytes of data streamed through the interface, with transform-stage results consistent with the native CPU implementation. The current module draws 9.7 W, with its photonic devices operating at 14 GBaud and the surrounding driver and receiver electronics supporting up to 56 GBaud.



➤ Figure 4: Optalysys programmable photonic prototype, showing the pluggable module and its photonic and electronic subsystems.

These figures are important not because the prototype is intended to replace a GPU, but because they demonstrate that programmable mathematical transformations can be integrated directly into a high-speed data path and exercised on sustained data streams.

From prototype to integrated platform

Programmability is an important part of the architecture. The current module uses foundry process-design-kit components, helping portability across silicon-photonics platforms, while the photonic function is exposed through software and our FPGA intellectual property rather than being fixed around a single algorithm. We are developing an implementation that combines the photonic circuit, analogue front end and high-speed digital logic.

Our current designs are based on a 32 GBaud link with a transceiver energy efficiency of 2.68 pJ/bit, excluding the laser. Custom front-end circuits add 0.3–1 pJ/bit to perform operations including multiply-accumulate, Fourier transforms, and floating-point operations. At that link rate, the design performs the equivalent of 384 INT4 GOPs natively. Our next design target is 112 GBaud, increasing that figure to 1.344 INT4 TOPS. These are engineering targets for the integrated platform, not measurements from the current prototype.

Photonics doesn't need to replace the GPU to change the architecture of AI infrastructure. GPUs remain well suited to dense computation, while CPUs provide the control and flexibility needed elsewhere in the workload.

But as moving data is becoming such a major part of the problem, simply moving it faster cannot be the whole answer.

Faster memory and interconnects improve the movement itself, but the familiar sequence remains: move the data, perform the computation, then move the result again.

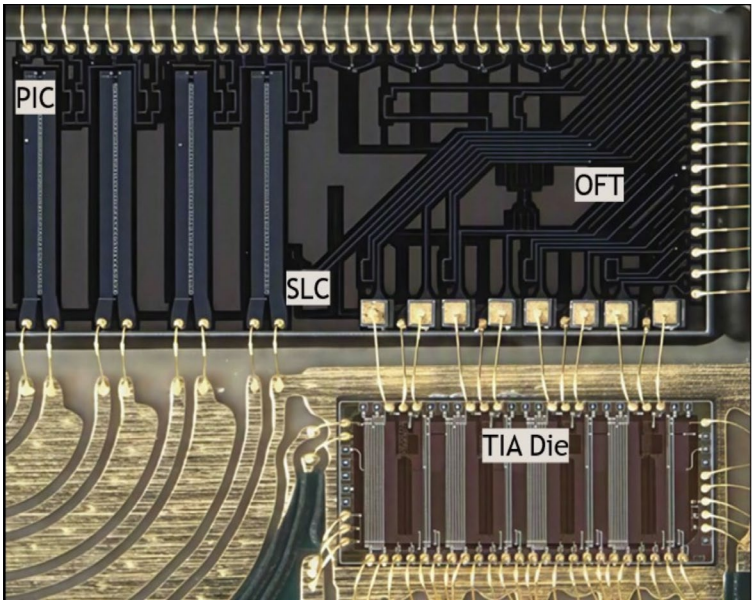
Compute-in-Transit offers another option for the parts of a workload that fit it: transformations can take place while data is already travelling between components, putting the photonic hardware in the data path rather than at the end of it.

Where an operation can be performed on the way, the data does not need to be sent somewhere else simply to perform that operation.

We will continue to explore which dataflows are worth mapping, how they should be represented and scheduled, and what interfaces are needed to connect them efficiently to the rest of the machine.

These are questions for computer architecture and software as much as for photonics. AI infrastructure will continue to need faster processors, memory and interconnects, but novel architectures and capabilities will be foundational to sustain its growth trajectory.

It's time to recognise photonics as more than connective tissue and explore the transformative potential of photonic Compute-in-Transit.



➤ Figure 5: Detail of the photonic circuit and detection path used in the current Optalysys prototype.

PIC PHOTONIC INTEGRATED CIRCUITS

BOOK YOUR REPRINT TODAY!

A reprint from **PIC Magazine** amplifies your editorial exposure and provides authoritative third-party validation. It's a powerful tool for **sales and marketing**, helping you showcase innovation with independent credibility. Use it across **events, pitches, newsletters, and social channels** to reinforce trust, extend visibility, and strengthen your position in the PIC industry.





Contact: **Shehzad Munshi**
shehzad.munshi@angelbc.com



Enabling high-volume silicon photonics through advanced test



Advanced photonic test is becoming critical to scaling silicon photonics for AI and HPC, with automated optical alignment, integrated electrical and optical measurements, and wafer-level testing.

BY AETHER LEE, GLOBAL PHOTONICS CENTER OF EXCELLENCE MANAGER, ADVANTEST

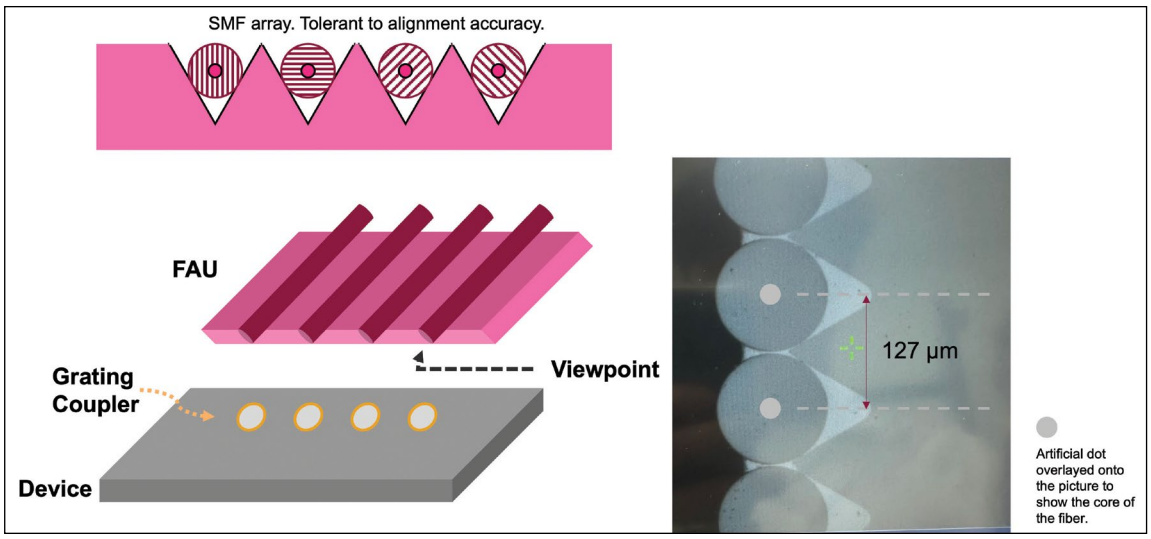
WITH THE explosive growth of artificial intelligence (AI) and high-performance computing (HPC), the requirements for data center connectivity are changing rapidly. AI training and inference systems must move enormous quantities of data among processors, memory and networking infrastructure, creating demands for bandwidth density and energy efficiency that are becoming increasingly difficult to meet with conventional electrical interconnects and traditional pluggable optical architectures. The industry is consequently moving toward silicon photonics and, increasingly, near- and co-packaged optics (NPO/CPO), where photonic components are placed close to or within the switch or ASIC package.

But scaling silicon photonics is not simply a matter of designing better photonic integrated circuits (PICs). As PIC complexity and manufacturing volumes increase, test is becoming a critical part of the manufacturing strategy. Accurate, repeatable wafer-level testing can identify known-good die (KGD) earlier, provide data for yield improvement and prevent costly downstream processing of defective devices. The challenge is doing all of this at production-worthy throughput.

Unlike conventional semiconductor test, PIC testing requires simultaneous management of two domains. Electrical signals must be applied and measured and light must be accurately delivered to and collected from structures that can be highly sensitive to position, angle, wavelength and polarization. Automated optical alignment, precise positioning and tightly integrated electrical and optical instrumentation are thus fundamental requirements for high-volume manufacturing (HVM).

Alignment becomes a manufacturing challenge

One of the most significant differences between electrical IC testing and photonic testing is the need



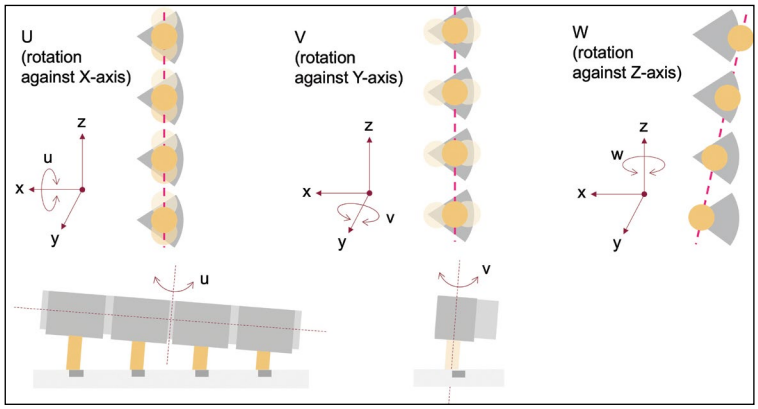
➤ Figure 1: The FAU shown here uses a visual/imaging process to bring the FAU to the grating couplers on the PIC, allowing it to look for light in a larger area.

to establish a repeatable optical connection to every device under test. Electrical wafer probing benefits from mature probe-card technology and well-defined contact structures. In photonics, however, a fiber array unit (FAU), as shown in Figure 1, must couple light into and out of grating couplers or other optical interfaces on the PIC. Small variations in position can translate directly into optical loss and measurement variation. The problem therefore extends beyond simply finding the device: the test system must find the point of optimum optical coupling quickly and consistently.

An HVM-capable approach begins by using machine vision to bring the FAU close to the grating couplers and establish "first light." From there, the system can perform an area scan followed by a gradient search to determine the optimum X-Y position. Fine positioning can then locate the local optical-power peak. In one implementation, a piezo-based nanopositioning stage supports a 100 × 100 µm scan range, with an area scan requiring a few seconds and a local gradient search approximately one second.

But X and Y are only part of the problem. Optimum coupling can require control across six degrees of freedom: X, Y and Z positioning plus U, V and W rotation. A high-performance digital imaging system, six-degree-of-freedom hexapod positioning platform and nanopositioning stage can therefore work together to provide coarse and fine alignment (see Figure 2).

The rotational axes are particularly important. Adjusting the W angle can help maximize beam coverage on the grating coupler and minimize channel-to-channel power variation. U and V alignment is more complex because both can influence insertion loss and center wavelength. Consequently, a production test methodology must optimize the complete optical interface rather than treating alignment as a simple two-dimensional positioning problem.

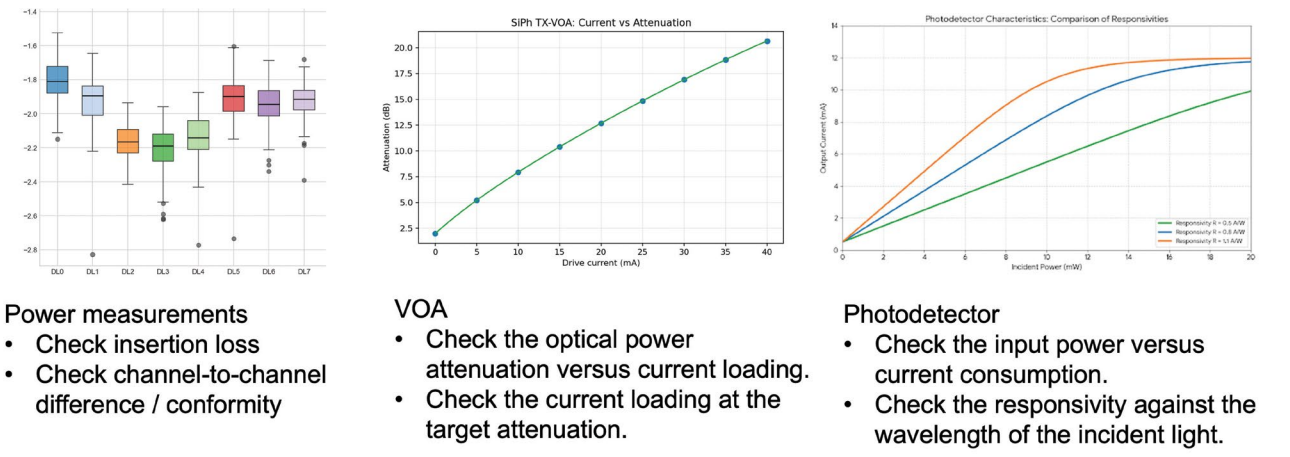


➤ Figure 2: The FAU actuator system shown here implements X, Y and Z positioning together with U, V and W rotation, providing optimal six-degree-of-freedom positioning and alignment.

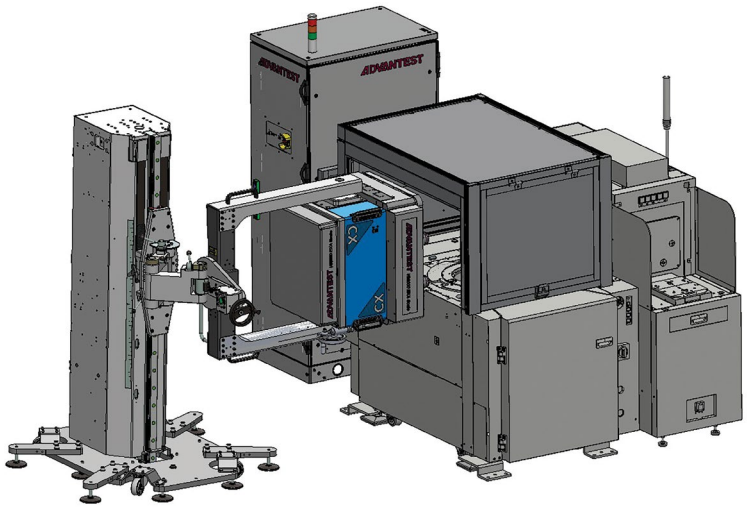
Making optical measurements repeatable

Finding maximum optical power is only useful if the resulting measurement accurately represents the device rather than variations in the test setup. That makes calibration, loss compensation and polarization management essential. An optical loopback, for example, can provide a known path through which transmitted light returns to an optical power meter, allowing alignment to be verified by measuring returned optical power. Multiple grating couplers can also be designated as reference channels to improve consistency across channels.

Polarization adds another variable. The optical response of a PIC can change with the polarization state of the incoming light, potentially making test results dependent on the test setup rather than solely on device performance. These capabilities matter because PIC production involves much more than a simple pass/fail optical continuity check. As shown in Figure 3, typical measurements can include insertion loss and channel-to-channel conformity as well as device-specific characterization. A variable optical attenuator (VOA), for example, can be tested for optical attenuation versus applied current and for the current required to achieve a target attenuation. Photodetector testing can correlate input optical power with electrical current and measure responsivity as a function of wavelength.



➤ Figure 3: The typical optical tests shown here are essential to ensuring successful PIC production.



➤ Figure 4: The V93000-TRITON Photonic Test Solution is a novel, integrated test solution optimized for high-volume, single-sided wafer-level probing to support scalable production of photonic devices.

Ring modulators present another example. Their intrinsic optical characteristics, together with those of the grating coupler, can be measured optically without electrically probing the wafer. Wavelength sweep data can be curve-fitted to derive the grating coupler's center wavelength, while removing the fitted baseline reveals ripple in the measurement.

Bringing electrical and optical test together

At the same time, photonic manufacturing cannot afford to treat electrical and optical testing as completely separate operations. The PIC is increasingly part of an electro-optical system, and HVM economics favor a test architecture that can handle both domains efficiently.

One opportunity is to use optical alignment time productively. Electrical alignment is performed before optical alignment, allowing continuity and potentially other purely electrical tests to run while automated optical probe alignment is taking place. This parallelism can turn what otherwise would be idle test time into productive measurement time.

Dark-current measurement is another important electrical test, particularly for photodetectors. Because typical monitor-photodiode dark currents can be in the nanoampere to sub-nanoampere range, instrumentation must provide the sensitivity and guarding needed to measure very small currents accurately.

The larger manufacturing opportunity comes from integrating optical instruments into the same automated test environment used for electrical testing. Rather than controlling lasers, power meters and other optical equipment through separate test programs, an optical control library can provide test-method interfaces within the semiconductor test software environment. The approach supports optical instruments from multiple suppliers while providing built-in calibration and compensation functions.

This convergence is important for throughput, but it also addresses another HVM requirement: consistency. As the number of instruments, measurements and test steps grows, manual setup and independent software environments create opportunities for configuration errors and test-to-test variation. Coordinated software control makes calibration, compensation and measurement sequencing part of an automated production flow.

Test earlier to protect yield

Wafer-level photonic test also has implications beyond test time itself. NPO, CPO and other advanced photonic systems combine multiple valuable components. A defective PIC discovered after assembly may mean losing not only that die but also the value added through packaging and integration. Establishing device performance before those later manufacturing stages therefore becomes increasingly important as integration complexity rises.

Industry work on NPO & CPO test methodology reflects this need for testing throughout the manufacturing hierarchy. Test can occur at PIC and electronic IC wafer level, after die singulation for known-good-die screening, at the assembled optical-engine level and ultimately at the completed co-packaged unit and system levels.

Wafer-level test provides the earliest opportunity to generate meaningful production data. Measurements of optical loss, wavelength behavior, detector responsivity, modulator performance and electrical characteristics can identify defective devices before expensive packaging steps while also feeding information back to wafer fabrication.

That data becomes particularly valuable during the transition from development to volume manufacturing. Production test is no longer simply an inspection step at the end of fabrication; it becomes part of the yield-learning loop. Stable, automated measurement enables manufacturers to distinguish actual process variation from variation introduced by alignment or instrumentation.

An ecosystem solution

No single company controls all of the technologies required for photonic production test. PIC designers define optical interfaces and device characteristics. Foundries determine process capabilities and manufacturing variation. Probe and positioning specialists provide precision optical access. Instrumentation suppliers provide lasers, power meters and other measurement technologies. Packaging companies must ultimately integrate known-good devices into increasingly complex assemblies, while automated-test specialists must bring these elements together in a repeatable production environment. An example

of such an environment is shown in Figure 4—the V93000-TRITON Photonic Test Solution, developed collaboratively by Advantest and FormFactor.

The importance of interoperability will only grow as silicon photonics moves toward higher volume. Industry efforts already envision test and measurement spanning wafer/die, component, optical-engine, co-packaged-unit and system levels, with an emphasis on ensuring that measurements performed by different vendors, on different equipment and in different facilities have consistent meaning.

This requirement changes the role of test. The objective is not simply to measure more photonic parameters. It is to create a manufacturing infrastructure in which optical and electrical measurements are automated, repeatable and scalable.

Silicon photonics promises to remove one of the major bottlenecks confronting AI infrastructure: moving ever-growing quantities of data without allowing interconnect power and complexity to become unsustainable. Realizing that promise at scale, however, depends on being able to manufacture sophisticated PICs economically.

This requirement changes the role of test. The objective is not simply to measure more photonic parameters. It is to create a manufacturing infrastructure in which optical and electrical measurements are automated, repeatable and scalable.

This makes test a key technology enabling the transition. Automated optical alignment, submicron positioning, multi-axis coupling optimization, polarization management, integrated optical instrumentation and parallel electrical/optical test can transform photonic characterization from a laboratory-intensive process into a repeatable production operation. Just as automated test became indispensable to the economics of semiconductor manufacturing, advanced photonic test will increasingly determine how quickly silicon photonics can move from promising technology to high-volume infrastructure.

PIC PHOTONIC INTEGRATED CIRCUITS

BOOK YOUR REPRINT TODAY!

A reprint from **PIC Magazine** amplifies your editorial exposure and provides authoritative third-party validation. It's a powerful tool for **sales and marketing**, helping you showcase innovation with independent credibility. Use it across **events, pitches, newsletters, and social channels** to reinforce trust, extend visibility, and strengthen your position in the PIC industry.





Contact: **Shehzad Munshi**
shehzad.munshi@angelbc.com



Exorcising the DSP: A back-to-basics route to co-packaged optics



Every high-speed optical module hides a digital signal processor that quietly eats power and adds delay. Mark Holm argues the way forward is simpler modulation, many wavelengths and a single comb source.

BY DR MARK HOLM, VP OF TECHNOLOGY, SOLINIDE PHOTONICS AB

ASK AN engineer what made the last decade of optical interconnect possible, and you will hear about lasers, modulators and silicon photonics. Ask what actually carried the signal across a lossy backplane, through two connectors and down a dispersive fibre, and the honest answer is the digital signal processor. The DSP equalises the channel, corrects the errors and lets us run lane rates that the raw physics would never tolerate on its own. It has been the quiet hero of every pluggable module since 100G.

It has also become the hidden cost. In a modern pluggable, the DSP is the hottest thing in the box, and in an AI datacentre, where every watt spent on the network is a watt not spent on compute, that matters. It adds delay as well as heat: the signal has to be sampled, equalised, decoded, and corrected before it goes anywhere. We have come to accept this as the price of speed. I think that is a mistake, and I think the industry's own roadmaps are beginning to say so.

Datarate: The treadmill that feeds the devil

The orthodox way to add bandwidth is to run each lane faster. Fifty gigabits became a hundred, a hundred became two hundred, and four hundred is on the slides. Each step costs more than the one before. The electrical channel gets harder, the equaliser gets deeper, the error correction gets stronger, and the copper reach gets shorter, which is why the optics have been forced off the front panel and onto the package in the first place. The DSP grows at every step, because it is the DSP that makes the step possible.

There is a second way to add bandwidth that costs none of this. Instead of running one wavelength faster, run more wavelengths. Dense wavelength division multiplexing has done this in long-haul networks for thirty years, and it is now arriving at the chip. Hold the lane rate at something the physics is comfortable with and multiply the lanes in colour rather than in clock. The pleasant side effect is that the simple formats come back into play: non-return-to-zero, a one and a zero, no equaliser worth the name, no forward error correction, and a clock forwarded on its own wavelength, so the receiver never has to recover one. A link like that is about as close to a wire as optics gets.



➤ Figure 1: The devil in the box: DSP as the hot spot of a pluggable module. Inside a DSP-based pluggable module the signal processor is the hottest component in the box. It equalises the electrical channel, recovers the clock and corrects errors, and every one of those jobs costs power and adds delay before a single bit reaches the fibre.

Latency: The cost nobody puts on the datasheet

The watts are visible. The nanoseconds are not, and in an AI fabric they may matter more. Training a large model means thousands of accelerators exchanging partial results and waiting for the slowest one to answer. Every hop through a DSP and an error-correction block adds delay, and the delays compound across the fabric. The simplest NRZ link with a forwarded clock has the shortest electrical path from one chip to another that we know how to build. It is no coincidence that the people building the largest fabrics are the ones talking about dropping the DSP.

Shoreline: Bits per millimetre of edge

Once the optics moves onto the package a new limit appears. The electrical reach from the accelerator to the optical engine is now measured in millimetres, so the engine has to sit right at the edge of the die, and there is only so much edge. Engineers call it shoreline, or beachfront, and it is becoming the scarcest real estate in the datacentre. The question is no longer how fast one fibre can go. It is how many bits you can push through each millimetre of edge.

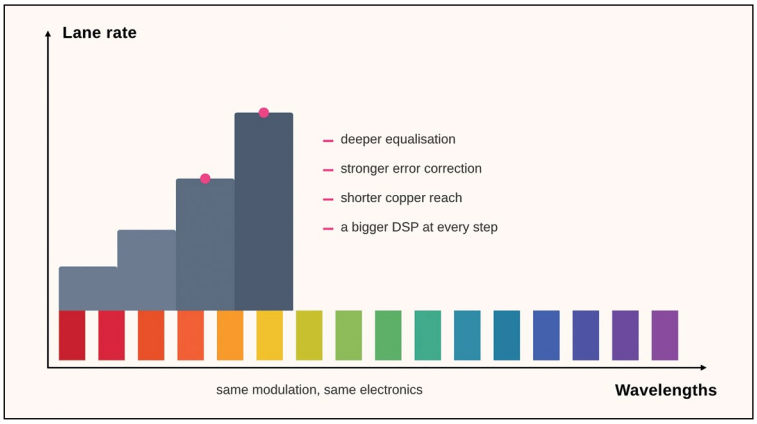
Many wavelengths per fibre is the obvious answer, and here the devil comes back in through the side door. Many wavelengths has always meant many lasers, and each laser brings its own driver, its own cooler, its own monitor and its own control loop. Sixteen wavelengths means sixteen of everything. Forty-eight means forty-eight. The electronics you threw out with the DSP creep back in, one laser at a time, and the shoreline you saved on the fibre side you spend again on the source side.

The source: The saving is in what you do not build

This is where optical frequency combs change the arithmetic. A comb takes one pump laser and one microresonator, a ring of silicon nitride less than a millimetre across, and produces a whole grid of wavelengths at once, evenly spaced and locked to one another. One pump, one temperature loop, one control circuit, and as many wavelengths as the ring will give you.

It is worth being precise about where the saving comes from, because it is easy to get wrong. A comb is not a more efficient laser. Line for line, a good semiconductor laser turns electrical power into light at least as well as a comb does. The saving is in everything that sits around the laser. A laser array carries a stack of control electronics for every wavelength; a comb carries one stack for all of them. That is why even a modest comb can hold its own against an array of lasers at the system level, despite being the less efficient light source. The electronics it does not need pay for the light it wastes.

A high-efficiency comb goes further. When the comb converts pump light into usable lines efficiently, with enough power to keep the links error-free,



➤ Figure 2: Two scaling axes: lane rate versus wavelength count. Two ways to add bandwidth. Running each lane faster (vertical axis) demands deeper equalisation, stronger error correction and shorter copper reach at every step. Adding wavelengths (horizontal axis) multiplies capacity while leaving the modulation format, and the electronics, unchanged.

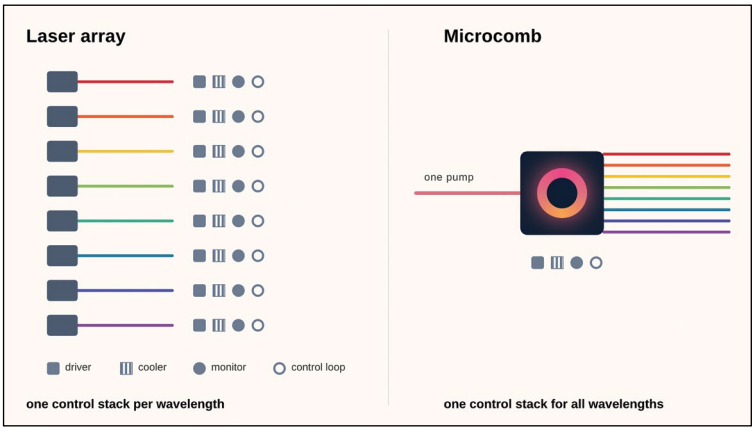
the removal of the per-line electronics stops being a trade and becomes a straight gain. In our own system modelling, a co-packaged engine fed by an efficient comb comes out at roughly half the energy per bit of the same engine fed by a laser array, and the whole of that difference is electronics that were never built. Against a DSP-based pluggable, the gap is an order of magnitude, which is the gap the industry is actually trying to close.

Why the industry's own roadmap points this way

NVIDIA has published a roadmap towards sub-picojoule-per-bit optical interconnect that is worth reading twice. Its first conclusion is that once a link goes DWDM, with simple modulation and a forwarded clock, the analogue circuits stop being the problem and the optical power becomes the largest single slice of the budget. Its second conclusion is a list of what to fix: better detectors, better rings, better processes, direct drive from the die. Every item on that list attacks the receiver, the modulators or the circuits. The source is not on it.

I think the source is the list. If optical power is the biggest slice and per-wavelength electronics is the biggest hidden multiplier, then a single efficient source that feeds every wavelength is the lever the roadmap is missing.

One objection deserves a direct answer. Running NRZ without error correction is not free: it demands real optical margin on every wavelength, because there is no DSP to rescue a weak channel. That might sound like an argument against the simple link. It is the opposite. It is precisely why the quality of the source decides whether the simple link is possible at all. A comb that delivers its lines at the right power, at the right spacing, in the right band, and holds them there as the package heats and cools, gives every wavelength the margin it needs. A poor comb, or a laser array that drifts, forces you to put the electronics



➤ Figure 3: Laser array with per-line control stacks versus a single microcomb. A laser array carries a driver, cooler, monitor and control loop for every wavelength; a microcomb carries one set for all of them. A silicon nitride microresonator, pumped by a single laser, generates an evenly spaced grid of wavelengths that can be modulated and received by simple ring resonators with no DSP in the path.

back. The source the roadmap leaves off its list is the thing that makes the rest of the roadmap work.

Back to basics

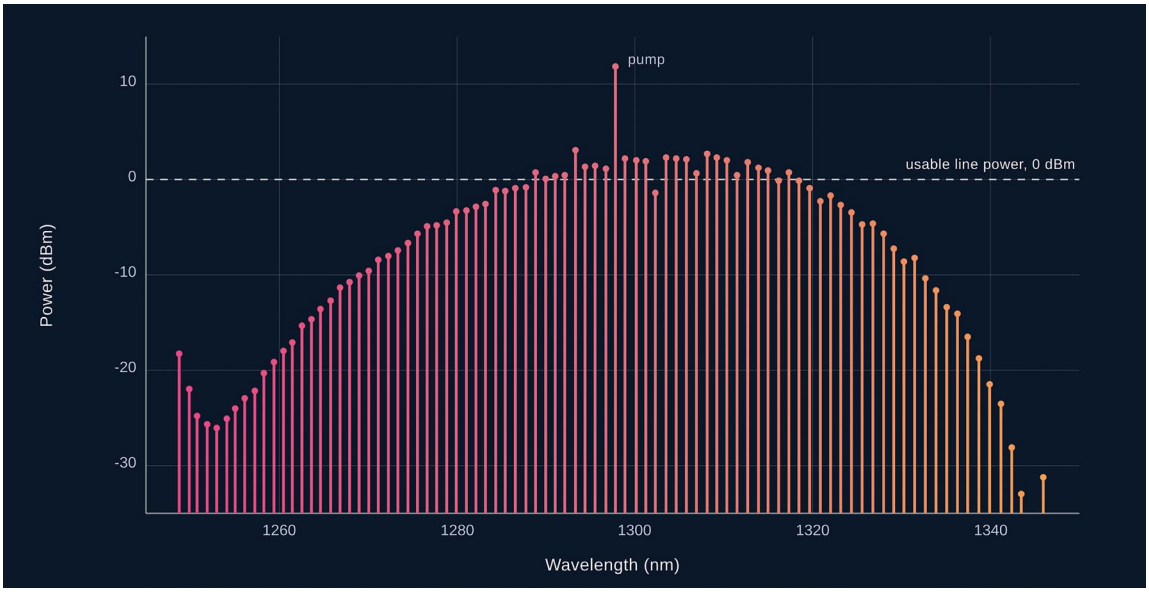
So here is the vision. NRZ, no FEC, lots of wavelengths, one comb. The modulator is a ring, the receiver is a ring and a photodiode, the clock rides on its own colour, and the DSP leaves the module. The link becomes as dumb and as fast as a wire, which is what the accelerator wanted all along.

The roadmap axis changes with it. Instead of counting lane rate, we count wavelengths: a few tens from a single source today, approaching a hundred within the decade, with the electronics staying the same size while the bandwidth multiplies. The external laser source, the form factor the industry is now standardising, becomes the natural home for the comb, feeding several engines from one pump.

Thirty years ago, we built optical links with a laser, a photodiode and not much else, and we were rather proud of how simple they were. We added the DSP because the channel got hard. The channel is about to get easy again, because it is a few millimetres of interposer and a short run of fibre. It is time to let the devil go.

FURTHER READING

- S. Cohen, B. G. Lee, S. Song, A. Rekhi, P. Sun, N. Mehta, N. Nedovic, G. Kalogerakis, L. Xu, T. H. Greer and C. T. Gray, A Roadmap Toward Sub 1pJ/b Optical Interconnect, IEEE Photonics Society Summer Topicals Meeting Series, Berlin, 2025.
- B. G. Lee, N. Nedovic, T. H. Greer and C. T. Gray, Beyond CPO: A Motivation and Approach for Bringing Optics Onto the Silicon Interposer, Journal of Lightwave Technology 41, 1152-1162, 2023.
- Ó. B. Helgason, M. Girardi, S. Sideris, L. Gantz and M. Holm, High-Efficiency Silicon Nitride Microcombs for Co-Packaged Optics, Optical Fiber Communications Conference (OFC) 2026, paper Th2A.13.
- N. Mehta, A. S. Rekhi, N. Nedovic, S. G. Tell, S. Song and C. T. Gray, A microring resonator-based clock-forwarded DWDM optical interconnect in a monolithic silicon photonics process, ECOC 2023.
- Optical Internetworking Forum, External Laser Small Form Factor Pluggable (ELSFP) Implementation Agreement, OIF-ELSFP-02.0, January 2025.



➤ Figure 4: Solinide O-band microcomb output spectrum. Measured output spectrum of a Solinide O-band silicon nitride microcomb, showing the evenly spaced lines and the flatness that a no-FEC NRZ link relies on. The dashed line marks the usable line power; the tall line is the pump. (Credit: Solinide Photonics AB.)

VIAVI Solutions



Scale with Confidence
Infrastructure Test for the AI Revolution

VIAVI MAP-300 is optimized for compact, high-throughput silicon photonics testing - from development through manufacturing - from wafer through packaging devices.

- Superior testing speed
- Lowest cost-of-ownership
- High-density hardware with multithreaded test execution
- Global footprint with expert local support
- Triggering with deep buffer capture

Learn More:
https://www.viavisolutions.com/en-us/products/map-300-multiple-application-platform



Scaling VECSEL innovation for the next generation of photonics



Vexlum is expanding its presence in the UK with a new R&D laboratory focused on supporting customers and strengthening collaboration across the country's growing photonics and quantum ecosystem.

In this conversation, **Sarab Chopra, Editor at PIC Magazine**, speaks with **Jussi-Pekka Penttinen, Co-founder, CEO and CTO of Vexlum**, about the company's VECSEL semiconductor laser technology, its approach to precision laser manufacturing, and the opportunities for VECSELs in photonic integrated circuits, quantum technologies and other emerging applications.

SC: Could you briefly introduce yourself and Vexlum?

J-P: My name is Jussi-Pekka Penttinen, and I am the CEO and co-founder of Vexlum. Vexlum is a semiconductor laser company focused on developing high-performance laser sources for demanding applications.

Our core technology is the VECSEL, or vertical external-cavity surface-emitting laser. These are optically pumped semiconductor lasers that can provide single-frequency, narrow-linewidth operation. This combination is particularly valuable for applications such as quantum technologies, where precise control of the optical source is extremely important.

The technology also gives us flexibility in terms of wavelength and allows us to address applications where conventional semiconductor lasers may not provide the required combination of power, beam quality, spectral performance and wavelength.

We are therefore developing VECSELs not simply as individual laser products, but as a platform that can support a range of applications and wavelengths.

SC: Why did Vexlum choose the UK for its first overseas R&D laboratory?

J-P: There are two main reasons. First, the UK has very strong activity in quantum technologies, and we already have a number of customers here. There is a significant ecosystem covering research organisations, technology developers and companies working towards commercial quantum systems.

There is also a strong emphasis on developing the supply chain in the UK. As quantum technologies move from research towards industrial

products, having the right suppliers and technologies available locally becomes increasingly important. For us, it therefore makes sense to have a presence close to customers and partners rather than supporting everything remotely from Finland.

The second reason is our collaboration with Stefan Trupp, the new managing director of Vexlum UK. We have the opportunity to bring Stefan into the company and make use of his valuable knowledge and experience of the UK photonics and quantum landscape.

The new R&D presence gives us a platform from which we can build those relationships further. It is not simply about establishing an office in another country. It is about being closer to the ecosystem and being able to participate more directly in its development.

SC: How will the new UK lab support the UK's growing photonics and quantum ecosystem?

J-P: The new UK location will allow us to provide better support to our existing customers, while also introducing new customers to our technology and showcasing our products.

Being physically closer to customers is particularly valuable when you are working with advanced laser technology. Customers often want to understand how the technology performs in their own systems, and they may have very specific requirements around wavelength, beam characteristics or integration.

The UK facility will also allow us to work more closely with customers as they develop new products and applications. When it comes to new capabilities and new quantum applications, collaboration becomes very important.

We see the laboratory as a way of creating a closer connection between Vexlum and the UK ecosystem. It allows us to learn more about the challenges customers are facing, while also giving them easier access to our technology and expertise.

Ultimately, we want to contribute to the development of the ecosystem rather than simply supplying a component into it.

SC: What makes VECSEL technology different from other laser technologies already available?

J-P: VECSELs are semiconductor lasers, so it would be easy to assume that they are similar to conventional diode lasers. That is actually one of the things we often need to explain to customers and collaborators. Although VECSELs are semiconductor lasers, they are quite different from diode lasers.

VECSELs are surface-emitting, which in this respect makes them more similar to traditional solid-state disk lasers, while retaining the gain versatility of semiconductor technology. That combination gives us some interesting characteristics in terms of beam quality and wavelength flexibility.

Another important difference becomes apparent when looking at visible wavelengths. With conventional approaches based on diode or fibre lasers, several components can be required to reach visible emission. Typically, you may need a seed laser, an amplifier and then a doubling cavity or another frequency-doubling stage.

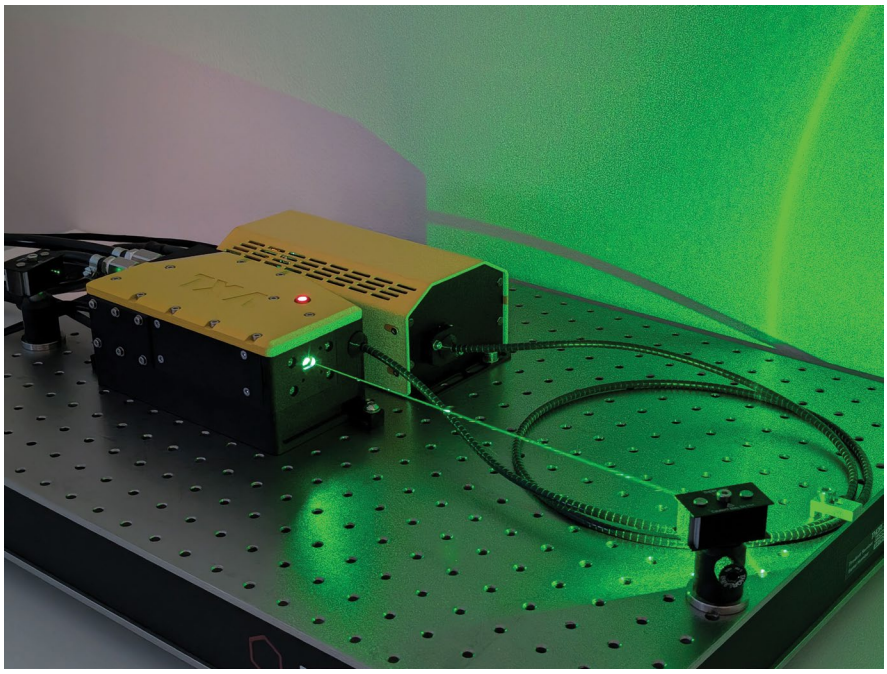
With our VECSEL architecture, the unique surface-emitting geometry and external cavity allow us to perform highly efficient intracavity frequency doubling. This means that we can reach visible wavelengths, including yellow, within a relatively compact laser architecture.

This is particularly interesting because of what is often referred to as the "yellow gap" in laser technology. Certain wavelengths in this region have historically been challenging to produce efficiently with conventional semiconductor laser approaches.

This is particularly interesting because of what is often referred to as the "yellow gap" in laser technology. Certain wavelengths in this region have historically been challenging to produce efficiently with conventional semiconductor laser approaches.

For VECSELs, yellow wavelengths are not an issue. In fact, the ability to address these wavelengths is one of the characteristics that differentiates the technology.

The combination of semiconductor gain, surface emission and the external cavity



therefore gives us a platform that is quite different from a conventional diode laser.

SC: Where do you see the biggest opportunities for VECSELs in integrated photonics?

J-P: For integrated photonics, VECSELs can serve as light sources that are coupled into waveguides. One of the major benefits is the excellent circular beam quality of our lasers, which comes from the combination of the surface-emitting geometry and the external cavity.

For photonic integrated circuits, coupling light efficiently from a source into the PIC is an important part of the overall system. The quality and characteristics of the optical beam can have a direct impact on how effectively that coupling can be achieved.

The circular beam profile of VECSELs can therefore be advantageous when coupling into waveguides and other photonic structures.

There is also an interesting relationship between the scaling of PIC technology and the scaling of laser sources. As integrated photonics allows more optical functions to be placed onto a chip, there is an increasing need to provide appropriate optical sources that can feed those systems.

This is particularly relevant in quantum technologies, where scaling the number of optical channels or beams is an important challenge.

SC: How important is in-house manufacturing to the performance and reliability of the technology?

J-P: It is crucial, particularly for quantum applications. Quantum technologies require a large number of different wavelengths, and our ability to supply those wavelengths comes from controlling the whole process.

We grow the wafers, process the gain chips and build the laser ourselves. The laser is essentially a platform, with the wavelength determined by the luminescence of the gain chip, which in turn is determined by the epitaxy recipe used for the wafer.

This level of control is very important because it gives us the ability to adapt the technology to different wavelength requirements while maintaining control over the underlying manufacturing process.

Having all these capabilities in-house also becomes increasingly important as the market develops. Quantum technology is becoming less of a purely research-driven field and more of an industrial product environment.

When technologies are primarily being used in research laboratories, the requirements can be different. As products move towards commercial deployment, customers start to care much more about reliability, consistency, repeatability and long-term performance.

Having direct access to the gain material and chip-processing capabilities is therefore crucial. It allows us to understand and control the technology at a much deeper level, rather than treating the laser simply as a finished component.

SC: Beyond quantum, which PIC applications are driving demand for your technology?

J-P: PICs are particularly valuable in quantum technologies because they provide a route to scaling. As quantum systems scale, they require more optical beams. Traditionally, this has often been achieved using bulk optics, which can result in very complicated and large systems.

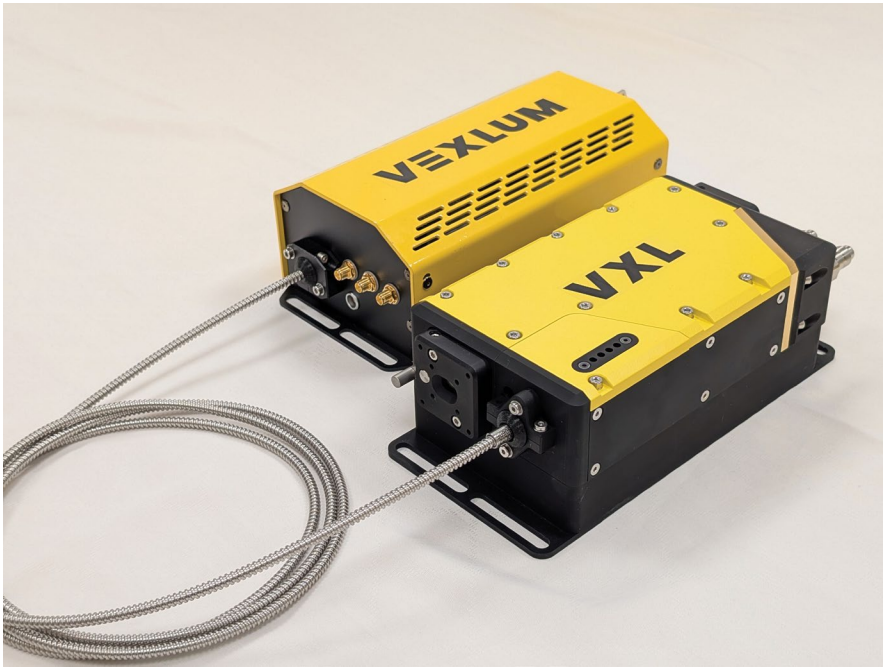
As you increase the number of beams, the complexity of those optical systems also increases. This is one of the reasons integrated photonics is so attractive. PICs provide an opportunity to put more of the optical functionality onto a compact platform.

For PIC-based quantum systems, one of the benefits of our technology is that we can generate a high level of power from a single laser and then distribute that light into tens or even hundreds of beams.

That ability to start with a suitable high-power source and then distribute the optical power is potentially very valuable as systems scale.

Beyond quantum, there are other applications where similar requirements exist. Semiconductor metrology is one example, where precise and stable optical sources can be important for measurement and inspection.

Biomedical applications are another area where our lasers, combined with PIC technology, could potentially provide benefits. The common theme is the need for precise optical sources that can be integrated into increasingly sophisticated systems.



SC: What are the major applications and technology trends driving demand for this approach?

J-P: We are a laser company focused on quantum technologies, so it is difficult for me to discuss potentially proprietary applications in detail. However, within quantum, the ability to divide beams and modulate signals is becoming increasingly important as systems scale.

The scaling challenge is a significant factor. As quantum systems become larger, you need to control more optical channels and maintain the required optical performance across them.

This is where the combination of high-performance laser sources and photonic integrated circuits becomes particularly interesting. The PIC can provide the functionality needed to manipulate and distribute the optical signals, while the laser provides the source.

We therefore see a growing opportunity for the two technologies to work together. Rather than thinking about the laser and the PIC as completely separate technologies, it is increasingly useful to consider them as parts of a larger photonic system.

SC: Looking at the broader picture, what is next for Vexlum and its role in the European photonics industry?

J-P: Vexlum is rapidly becoming a category leader in Europe for single-frequency lasers. The focus now is on scaling up volumes and bringing down costs, both for quantum applications and for other applications requiring precision lasers.

As these markets develop, the challenge is not only to demonstrate that the technology works. We also need to make it available at the volumes, cost and reliability levels required for commercial applications.

That is why scaling our production is an important part of our strategy. We want to move from technologies that are highly capable but produced at relatively low volumes towards products that can support the growth of emerging industries.

With these capabilities, we are enabling the wider quantum and semiconductor industries to move forward with new capabilities and develop new products.

We launched our new VXL product this year, and we are also developing new products for launch in the coming years. At the same time, we are scaling up our production.

The European photonics industry has a strong position in many areas, and we believe there is a significant opportunity to develop that further. For Vexlum, our aim is to contribute by providing precision laser technologies that can enable new systems and applications.

It is a very exciting time to be a European photonics company.



To promote your products and services contact:
Shehzad Munshi
T: +44 (0)1923 69015
E: shehzad.munshi@angelbc.com

EPIC
EUROPEAN PHOTONICS INDUSTRY CONSORTIUM
https://epic-photonics.com

IQE
http://www.iqep.com

LUCEDA PHOTONICS
Tel: +32 (0) 52 339 838
erwin@lucedaphotonics.com
www.lucedaphotonics.com/en

POET Technologies
Raju Kankipati
Tel: +1 510-648-0382
raju@poet-technologies.com
www.poet-technologies.com

VPIphotonics DESIGN AUTOMATION
Tel: +1-781-762-3901 / +49-30-398058-0
info_pic@VPIphotonics.com
www.vpiphotonics.com

PIC PHOTONIC INTEGRATED CIRCUITS
To advertise here call
Shehzad Munshi
+44 (0)1923 690215
shehzad.munshi@angelbc.com

Showcase your presence to the global PIC industry for 12 months

PIC MAGAZINE is the primary information resource for professionals working with PIC materials and devices.

Each issue of the magazine is distributed digitally to over 28,000+ professionals worldwide.

With the largest and most diverse global audience of qualified buyers served by any photonic focused magazine, secure a corporate partnership and buyers guide entry and receive 12 months of promotion and exposure to the PIC industry across the magazine portfolio including:

- 12 months of promotion in quarterly issues of the digital and PDF magazine
  featuring your company logo and company details
- 12 months of promotion and exposure on the PIC Magazine website **www.picmagazine.net** featuring a Buyers Guide entry and hyperlinked company logo
- Sponsorship of one e-news alerts sent out to over 15,000 industry professionals with text and a standard web banner.

To find out more, contact:
Shehzad Munshi
E: shehzad.munshi@angelbc.com
T: +44 (0)1923 690215





VPIphotonics
DESIGN AUTOMATION

PROFESSIONAL SIMULATION TOOLS & EXPERT DESIGN SERVICES

Integrated Devices, Waveguides & Fibers
VPIdeviceDesigner™

Photonic Integrated Circuits
VPIcomponentMaker™ Photonic Circuits

Fiber Amplifiers & Lasers
VPIcomponentMaker™ Fiber Optics

Transmission Systems
VPItransmissionMaker™ Optical Systems

Lab-proven DSP Algorithms
VPItoolkit™ DSP Library

Quantum Key Distribution
VPItoolkit™ QKD

We empower you to
define the cutting edge.

Request an evaluation

www.VPIphotonics.com