



PHOTONIC INTEGRATED CIRCUITS

ISSUE III 2026

AN ANCEL BUSINESS COMMUNICATIONS PUBLICATION

PICMAGAZINE.NET

The background of the cover is a high-resolution, colorful image of a microchip wafer. The wafer is divided into a grid of individual dies, each with its own intricate circuitry. The colors are vibrant and varied, ranging from deep blues and purples to bright yellows and oranges, likely due to the different materials and processes used in the manufacturing. The perspective is slightly angled, showing the curvature of the wafer.

AUTOMATED VISIBLE-RANGE PIC TESTING AT WAFER AND DIE SCALE



www.EVGroup.com

INDUSTRY LEADING WAFER PROCESSING EQUIPMENT FOR PHOTONIC INTEGRATED CIRCUITS

■ Optical and Nanoimprint Lithography

■ Temporary and Permanent Wafer Bonding

■ NILPhotonics® Competence Center



EVG®7300
Multifunctional UV Nanoimprint Lithography System

GET IN TOUCH to discuss your process development and manufacturing needs

www.EVGroup.com

Industrialisation of Photonics

➤ The photonics industry is approaching an inflection point. After decades of research, development and early commercialisation, the focus is rapidly shifting toward scalability, manufacturability and system-level integration.

As demand accelerates across AI infrastructure, high-performance computing, sensing, telecommunications and emerging quantum applications, photonic integrated circuits are being asked to deliver not only performance, but also the consistency, reliability and production maturity expected of modern semiconductor technologies.

A central theme emerging across the industry is bandwidth. The continued growth of AI workloads, data centre traffic and machine-to-machine communication is driving the search for new device architectures capable of operating at frequencies once considered beyond practical reach.

Advances in electro-optic modulation, novel material platforms and ultra-compact device designs are pushing integrated photonics deeper into the sub-THz and THz domains, creating opportunities for faster, more energy-efficient communication systems.

At the same time, the challenge of scaling production has become equally important. The transition from laboratory innovation to high-volume manufacturing requires a fundamental shift toward data-driven engineering, process standardisation, automated workflows and closed-loop feedback systems.

As photonic devices become more sophisticated, success increasingly depends on the ability to establish robust

design environments, manufacturing control systems and test methodologies capable of supporting predictable yields and repeatable performance.

Integration continues to be another defining trend. Future systems will rely on increasingly complex combinations of photonics, electronics, advanced packaging technologies and heterogeneous architectures.

The boundaries between traditionally separate technology domains are becoming less distinct as designers seek higher levels of performance, density and energy efficiency. This convergence is creating new opportunities while simultaneously introducing new manufacturing and design challenges.

The importance of testing and validation is also growing. As devices become more compact and system architectures more interconnected, characterisation and process monitoring are evolving from verification activities into critical enablers of manufacturing scale. The ability to rapidly correlate design, process and performance data will play a decisive role in accelerating innovation while maintaining quality.

Underlying all these developments is a broader industry transformation.

The future of integrated photonics will not be determined solely by individual device breakthroughs, but by the maturity of the ecosystems that support them. Design methodologies, manufacturing infrastructure, packaging capabilities, supply chains and standards must evolve together if photonics is to fulfil its potential as a foundational technology for next-generation computing and communications.

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

Graphic Design & Multimedia Assistant

Harvey Watkins harvey.watkins@angelbc.com

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, Burnshall Road, Coventry CV5 6SP, UK. T: +44 (0)2476 718 970
E: info@angelbc.com W: picmagazine.net




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 6 times a year, February, April, June, August, October and December by Angel Business Communications Ltd, Unit 6, Bow Court, Fletchworth Gate, Burnshall 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.

Contents

COVER FEATURE

Automated visible-range PIC testing at wafer and die scale

Quantum Transistors has developed a highly automated electro-optic test platform based on MPI Corporation's probe system to characterize visible-spectrum photonic integrated circuits (PICs) at wafer and die level.



24 Photonics at the heart of AI data centers and beyond

AI-driven computing requires efficient optical interconnects to overcome data-center bottlenecks.

28 Enabling THz bandwidth on-chip with plasmonic modulators

Plasmonic electro-optic devices are emerging as a route to sub-THz and THz-class bandwidths, offering a compact, low-power alternative to traditional silicon photonic modulators for next-generation data center interconnects.

32 Ncodin's nanolasers eye AI infrastructure

Delivering a revolutionary energy-per-bit, the French start-up's nanolasers are a compelling source for on-chip communication in tomorrow's data centres.

34 Making terabit optics: What 400G per lane really demands from the photonic engine

As AI infrastructure pushes toward 400G per lane and beyond, the photonics industry faces a critical transition in materials, modulation formats and manufacturing capacity.

38 Building the fab of the future for high volume PIC manufacturing


Photonic integrated circuits are moving from technical promise to industrial expectation.

44 Precision single-wafer HBr wet etch optimizes compound semiconductor optoelectronics

A precision single-wafer HBr wet etch process developed by JST is delivering plasma-comparable uniformity for compound semiconductor fabrication, offering a cost-effective and scalable approach for next-generation optoelectronic and photonic devices.

NEWS

- 06 CORNERSTONE highlights UK SiPh growth
- 07 PhotonVentures on PIC investment
- 08 VIVA develops VCSEL smart glasses
- 09 Miniaturised lasers advance quantum PICs
- 10 EU proposes Chips Act 2.0



06

Luceda Photonics Design Platform

Luceda IPKISS

First-time-right
photonic IC design software

Automate and integrate all aspects of
your photonic design flow in one tool,
using one standard language.



Luceda Design Kits

Gain access to a wide range of PDKs, ADKs, and TDKs.
Design and tape out to your foundry of choice.

OFC 2026

Visit us at our booth
to learn more!


Booth
#965




New Product Luceda DRC

Identify and resolve design issues
instantly within your design environment
with a single click, and proceed to
successful tape-outs with confidence.


CONTACT US TO
TRY IT TODAY!



Component / Circuit
Layout



Design Rule Checking
in Luceda DRC



DRC-Clean
Layout

CORNERSTONE highlights UK SiPh growth

New research calls for domestic pilot line investment to strengthen sovereign silicon photonics capabilities and accelerate commercial scale-up.

NEW MARKET research from the CORNERSTONE Photonics Innovation Centre has highlighted significant growth potential for the UK silicon photonics sector, while warning that investment in domestic scale-up infrastructure will be critical to securing long-term competitiveness and sovereign capability.

The study found strong industry support for a UK silicon photonics pilot line, with 74% of respondents saying it would accelerate innovation and 79% agreeing it would significantly enhance the UK's sovereign technology capabilities.

According to the research, investment in domestic pilot line infrastructure could contribute £2.9bn to the UK economy by 2040 and create approximately 2,850 new jobs.

The findings also showed that 76% of respondents believe improved UK scale-up infrastructure would accelerate company growth.

The survey gathered responses from 100 UK-based decision-makers currently developing, deploying, or planning to develop silicon photonics technologies.


While 77% said they are already developing or planning silicon photonics activity within the UK, many highlighted manufacturing and prototyping challenges linked to overseas dependency.

Almost one-third of respondents identified tariff-related costs as a barrier to silicon photonics prototype development, reinforcing calls for

stronger domestic fabrication and commercialisation capability.

CORNERSTONE noted that global demand for silicon photonics is accelerating due to growth in AI infrastructure, data centres, quantum technologies and optical networking, with current market growth projections estimated at approximately 25–30% CAGR.

The organisation said the findings support recent recommendations from the UK Council for Science and Technology calling for a national photonics roadmap and investment in a dedicated silicon photonics pilot line to help transition UK innovation into scalable commercial production.



PhotonVentures on PIC investment

The integrated photonics-focused venture capital fund outlines the key factors it considers when evaluating photonic chip startups, from scalability and ecosystem alignment to market readiness and execution.

PHOTONVENTURES has shared its perspective on what investors look for in photonic integrated circuit (PIC) startups, highlighting the importance of commercial readiness, scalability, and ecosystem engagement alongside technical innovation.

The European venture capital fund, which focuses on integrated photonics companies from seed to Series A stages, argues that strong technology alone is not enough to attract investment.

Instead, startups must demonstrate a clear application, customer demand, and a realistic path from prototype to scalable production.

According to PhotonVentures, successful photonics companies are those that consider manufacturability, packaging, and system integration from an early stage.

The firm believes startups that engage with foundries, packaging providers, and system integrators early in their development are better positioned to overcome technical challenges and accelerate commercialization.

The fund also stresses the importance of founding teams combining deep technical expertise with business ambition, market awareness, and execution discipline.

Teams capable of balancing performance goals with customer requirements and commercial realities are viewed as more likely to succeed in the long term.

PhotonVentures notes that integrated photonics presents unique investment challenges due to long development cycles, high capital requirements, and reliance on specialized infrastructure.

As a result, investors place significant emphasis on timing, ecosystem maturity, and evidence that a technology can scale beyond the laboratory.

The company also highlighted the role of initiatives such as the Global Photonics Engineering Contest in helping identify emerging talent and technologies.

Such programs can provide early validation opportunities and expose startups to industry stakeholders, partners, and potential investors.

As integrated photonics continues to expand into areas such as AI, quantum technologies, communications, and sensing, PhotonVentures believes that execution, ecosystem integration, and commercial focus will increasingly distinguish investment-ready companies from purely research-driven ventures.

Organic device emits and harvests light

RESEARCHERS from the Institute of Science Tokyo and collaborating institutions have demonstrated an organic semiconductor device that can simultaneously harvest light and emit it, marking a step toward multifunctional optoelectronic systems and power-generating displays.

The team developed the device using multi-resonance thermally activated delayed fluorescence (MR-TADF) materials, engineered to minimise non-radiative energy losses that typically reduce efficiency in organic semiconductors.

Using OLED-associated materials v-DABNA and QAO in a layered

structure, the researchers achieved both 1.36% power-conversion efficiency and 2.0% light-emission efficiency within a single device.

The system also produced bright red emission at 1,000 cd/m² while operating at 3.2V, compatible with standard lithium-ion battery voltages.

According to the researchers, the work demonstrates simultaneous light emission, energy harvesting and photodetection within a single organic platform, establishing a new framework for multifunctional optoelectronics.

The team said the technology could support future applications including

self-powered displays, wearable electronics, transparent photovoltaics and integrated sensor systems, where lightweight and flexible organic materials offer advantages over conventional inorganic semiconductor platforms.

The system also produced bright red emission at 1,000 cd/m² while operating at 3.2V, compatible with standard lithium-ion battery voltages

VIVA develops VCSEL smart glasses

The VIVA project is developing photonics-enabled smart glasses that use VCSEL-based eye tracking without cameras.

AN EU-FUNDED research project is developing a new generation of smart glasses that use photonics-based eye tracking to create more intuitive and privacy-focused wearable devices.

The project, called VIVA, is using VCSEL technology to monitor eye movements without cameras, enabling lightweight smart glasses capable of responding to a user's gaze in real time.

The system is designed to support applications in automotive safety, industrial environments and future consumer wearables.

Unlike conventional eye-tracking systems that rely on cameras and

image processing, the VIVA platform uses laser-based sensing to detect microscopic eye movements through reflected light signals.

The approach aims to reduce hardware size, improve wearer comfort and enhance privacy by avoiding image capture altogether.

One of the project's first demonstrators includes auto-focal smart glasses that automatically adjust lens focus depending on where the user is looking.

The system is designed to recognise reading behaviour and other gaze patterns with high-

speed detection and low false-trigger rates.

The consortium is also integrating meta-optics and fast signal-processing technologies to support compact, low-power operation suitable for all-day wearable use.

Coordinated by Bosch Sensortec, the VIVA project brings together partners across Europe specialising in photonics, sensing, optics, AI and advanced manufacturing.

The initiative reflects growing interest in photonics-enabled human-machine interfaces and next-generation wearable sensing technologies.

PIC market to hit \$24.5B by 2033

THE GLOBAL PHOTONIC integrated circuit (PIC) market is poised for significant expansion, driven by surging demand from artificial intelligence infrastructure, hyperscale data centres, and next-generation telecom networks.

According to new research from DataM Intelligence 4 Market Research LLP, the market was valued at USD 10.79 billion in 2025 and is projected to reach USD 24.51 billion by 2033, growing at a compound annual growth rate (CAGR) of 10.8%.

This represents an incremental opportunity of nearly USD 13.7 billion over the forecast period. The growth reflects a broader shift in how data is transmitted and processed.

As AI workloads scale and data centre traffic intensifies, traditional electrical interconnects are increasingly

constrained by power consumption and bandwidth limitations.

PICs are emerging as a critical alternative. By integrating optical components such as lasers, modulators, detectors, and transceivers onto a single chip, these devices enable faster data transfer with significantly lower energy requirements.

This makes them particularly attractive for high-performance computing environments where efficiency and throughput are paramount.

The report highlights AI optical interconnects as a major growth driver, as companies seek to overcome bottlenecks in chip-to-chip and rack-to-rack communication.

At the same time, advances in silicon photonics manufacturing are improving scalability and reducing costs, accelerating commercial

deployment across cloud and telecom infrastructure.

High-bandwidth upgrades in global data centres are also contributing to increased adoption, as operators transition toward optical solutions to meet rising performance demands.

Industry observers note that while the technology continues to mature, challenges remain in areas such as packaging, thermal management, and large-scale integration.

However, ongoing innovation and investment are expected to address these barriers over the coming years.

With momentum building across both technical and commercial fronts, photonic integrated circuits are increasingly positioned as a foundational technology for the next era of data-driven computing.

Miniaturised lasers advance quantum PICs

Researchers have developed compact, PIC-based laser systems integrating novel Faraday isolators and laser-fabricated photonic packaging for quantum technology applications.

A CONSORTIUM of industry and research partners has demonstrated new laser-based approaches for building miniaturised, robust beam sources for quantum technology, combining photonic integrated circuits with advanced optical materials and precision glass packaging techniques.

The work, carried out under the HiPEQ project, focuses on reducing the size and complexity of laser beam sources used in quantum systems, which are typically large, sensitive and difficult to deploy outside laboratory environments.

The new prototypes measure just $22 \times 9 \times 6$ cm and are designed to support multiple wavelengths for quantum sensing and communication applications.

A key development in the project is the use of novel terbium-based crystals

with a significantly higher Faraday effect than conventional materials.

These crystals enable compact optical isolators that prevent back-reflected light from destabilising laser operation, a critical requirement for narrow-linewidth quantum light sources.

The crystals were grown using a laser-based optical floating zone process, allowing precise control of temperature gradients and material composition to achieve high-purity monocrystalline structures suitable for photonic integration.

In parallel, the team developed a glass-based packaging platform using selective laser-induced etching.


This approach enables μm -precision alignment of fibre coupling structures,

isolators and optical components within a monolithic housing, significantly reducing manual assembly and alignment requirements.

The laser systems integrate photonic integrated circuits, optical fibres, beam splitters and isolators into a unified module, with demonstrated operation at both blue and red wavelengths.

Researchers report that the approach improves robustness while paving the way for more automated assembly of complex photonic systems.

Further optimisation is still required, particularly in coupling efficiency and optical integration, but the results point toward more scalable and deployable laser sources for next-generation quantum photonic systems.



EU proposes Chips Act 2.0

The European Commission has unveiled Chips Act 2.0, a new legislative proposal aimed at strengthening semiconductor manufacturing, reducing strategic dependencies and boosting Europe's capabilities in advanced chip technologies.

THE EUROPEAN COMMISSION on Wednesday proposed Chips Act 2.0, a new package of measures designed to strengthen the European Union's semiconductor industry and reduce its reliance on foreign suppliers for critical technologies.

Building on the original European Chips Act, the proposal seeks to reinforce Europe's existing strengths in semiconductor production while expanding capabilities in advanced chip manufacturing and design.

The Commission said the initiative aims to improve supply chain resilience and address strategic vulnerabilities in key parts of the semiconductor value chain.

According to the Commission, the EU remains dependent on third countries in several critical areas, including

advanced chip manufacturing and semiconductor design.

Ensuring a stable supply of semiconductors is increasingly important for securing critical infrastructure, emerging technologies and Europe's broader digital economy.

Ensuring a stable supply of semiconductors is increasingly important for securing critical infrastructure, emerging technologies and Europe's broader digital economy.

The proposed legislation is intended to help the bloc maintain its position in the global semiconductor ecosystem while supporting the development of next-generation chip technologies.

It also seeks to strengthen Europe's ability to respond to disruptions in global supply chains.

Chips Act 2.0 forms part of a wider digital strategy that includes the Cloud and AI Development Act and the EU Open Source Strategy.

Together, the initiatives are intended to support a more competitive, secure and resilient European technology sector.

The proposal will now move through the EU legislative process before it can be adopted and implemented across member states.

Monash develops room-temperature valleytronic chip

RESEARCHERS at Monash University have developed a nanoscale valleytronic circuit that can generate, direct and detect light-based information on a single integrated chip, marking a step forward for photonic data processing technologies.

The device combines ultrathin two-dimensional materials with engineered metasurfaces to control light at the nanoscale, enabling information encoding through the "valley" degree of freedom, a quantum property that can be used to carry and process data.

For the first time, the team demonstrated a fully integrated system capable of creating optical signals,


routing them through nanostructures and converting them into electrical outputs within a compact chip-based platform.


Importantly, the system operates at room temperature, addressing a key limitation in many quantum and photonic technologies that typically require cryogenic conditions.

In a demonstration, the researchers successfully encoded and processed two separate images simultaneously, highlighting the device's ability to handle parallel streams of optical information.

The team said the approach could enable new classes of compact,

energy-efficient computing systems, with potential applications in quantum technologies, secure communications and AI-driven optical processing.






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



Automated visible-range PIC testing at wafer and die scale

Quantum Transistors has developed a highly automated electro-optic test platform based on MPI Corporation's probe system to characterize visible-spectrum photonic integrated circuits (PICs) at wafer and die level. Delivering exceptional repeatability, throughput, and measurement accuracy, the platform enables rapid PDK development, process optimization, and known-good-die qualification for scalable quantum photonic computing.

BY LAWRENCE VAN DER VEGT – HEAD OF PHOTONICS AT MPI CORPORATION, AND THOMAS FERREIRA DE LIMA – PRINCIPAL PHOTONIC ENGINEER, BENJAMIN STREKHA – PHOTONIC ENGINEER AND IGAL BAYN – VP PHOTONICS AT QUANTUM TRANSISTORS INC.

Introduction to diamond-based quantum processors

QUANTUM TRANSISTORS is a quantum computing startup developing a scalable quantum processor architecture based on solid-state spin qubits in engineered diamond structures integrated with photonic interconnects. Their approach relies on creating and controlling color centers in synthetic diamond, which act as stable quantum bits whose electron spin states can be initialized, manipulated, and read out using optical and microwave techniques. These spin qubits exhibit long coherence times and can operate at relatively higher temperatures compared to superconducting qubits. The architecture heterogeneously integrates these diamond qubits into photonic integrated circuits (PICs), enabling photons to mediate entanglement between spatially separated qubits. This photonic coupling allows the system to scale beyond nearest-neighbor

interactions and supports modular quantum processor designs.

Unlike many other quantum architectures that face challenges in scaling, their technology offers the potential for large-scale, fault-tolerant quantum computing systems deployable in standard data-center environments, bridging the gap between laboratory prototypes and commercially viable quantum processors.

Scaling up development of visible PICs

Designing, testing, and packaging PICs at industrial-scale levels for operation at visible wavelengths requires developing a new process development kit (PDK) based on a compatible photonic foundry process. While the device geometry family can be generally inspired from telecom near-infrared designs, the devices engineered by QT have 2–3 times smaller features that approach the same order

of magnitude as the critical dimension that a state-of-the-art foundry can offer. As a result, fabrication variations within each wafer and across wafers have an amplified impact on device under test (DUT) performance variability. Furthermore, coupling to a visible PIC is, in the best-case scenario, 3x less tolerant to misalignment offsets compared to infrared ones.

Quantum Transistors' main objective of the automated probe system is to accurately measure performance variations across thousands of DUTs per wafer over several wafers and correlate the results with numerical modeling. This allows them to control, refine, and qualify their new PDK for visible photonics. Second, these measurements are used to compute a score for each product die according to custom metrics. QT uses the probe system to monitor the performance of each die across several packaging steps after singulation, resulting in fully validated known-good dies (KGD) available for heterogeneous packaging.

The primary challenges associated with testing these devices stem from the need to maintain high throughput, precision, and repeatability in the visible wavelength range. Device and system sensitivity, along with environmental conditions, directly influence the accuracy and repeatability of measurements, which are essential for reliably determining the performance and quality of the DUT. By maintaining tight control over these factors through advanced test automation, the testing process produces consistent, high-quality data that accurately reflects the true performance of the device across its lifecycle.

High accuracy, repeatability, programmability and compliance with eye-safety requirements together

with the best-in-class technical support were key criteria in selecting MPI Corporation, driven by the scalable architecture of its probe system and test solutions. The resulting measurement data was subsequently leveraged to refine and optimize QT's PDK and identify good dies prior to packaging.

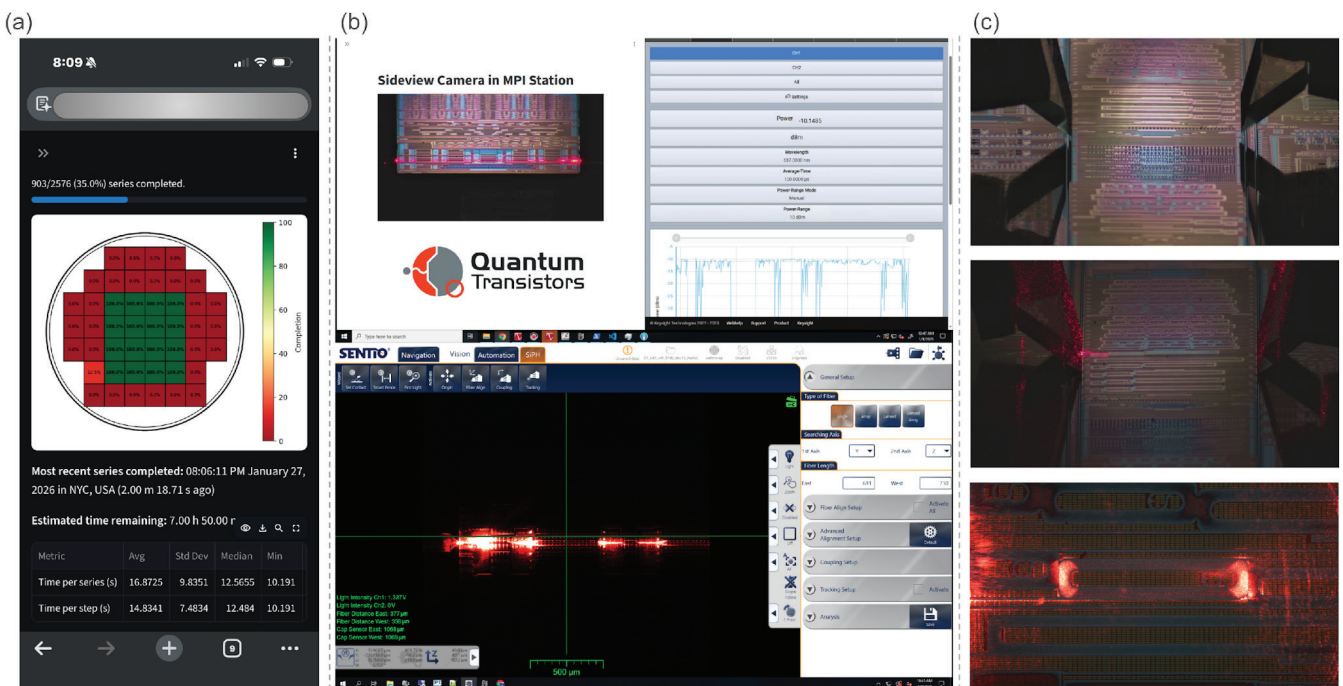
Automation requirements

A critical requirement was the system's capability to automatically compensate for temperature variations within a ± 2 °C window, while supporting continuous operation over multiple days. Equally important was the ability to pause and resume measurement sequences without impacting data integrity or system repeatability. While these requirements are generally manageable under standard conditions, the specific measurement environment introduced additional complexity. This included maintaining tight control over Polarization Extinction Ratio (PER) and Polarization Control (POL) from the laser source to the DUT. For example, one of the waveguide transition DUTs was measured via edge coupling using a single-mode fiber terminated by a lensed tip with a mode field diameter (MFD) of 1.66 μm , resulting in a measurable mean insertion loss (IL) of 0.015 dB ($\sigma = 0.003$ dB) across the wafer – note the extremely low variance in the order of milli-dBs required to measure such device. All these parameters were maintained with minimal alignment overhead and over several days per wafer while situated in a Midtown Manhattan office tower near a busy subway station.

Probe system description

To meet these stringent requirements, QT developed a dedicated test environment based on MPI's TS3500, configured for visible wavelength operation and integrated with SENTIO® control software via their Python SDK (Figure 1). This

► **Figure 1:** Trench and edge coupling onto electro-optic integrated circuits in the visible spectrum. (a) Wafer level dashboard monitoring measurement progress. (b) Real-time power monitoring with SENTIO® and External Power Meter. (c) Side and top view of live trench coupler measurement.



setup incorporated optimized optical fiber probes designed for trench, edge and grating coupling in the visible spectrum, best-in-class test instrumentation, and a highly optimized, automated software orchestrator that resulted in measurement throughput of 1 min per full wavelength scan per DUT.

➤ **Figure 2:** QT's test setup and automation pipeline confirmed the alignment stability and repeatability of the probe system in the visible range to be on par with the performance stability of the tunable laser. Monitors connected to the MPI station allow for continuous live monitoring of chuck and probe arm movement, wavelength sweeps, and power readings.

The configuration (Figure 2) shown below defines the measurement platform used to characterize the optical performance of QT's grating- and edge-coupler-based PICs. A tunable laser provides a wavelength-dependent optical source, with the signal split into two parallel paths: one for continuous optical power and wavelength monitoring and the other for device-under-test (DUT) characterization.

The probe system used by QT is configured with one West-positioned optical hexapod and one East-positioned optical hexapod, both equipped with undermount nanopositioners, each fitted with an arm and fiber holder. The hexapods perform coarse alignment, while the nanopositioners enable fine alignment with nanometer-level accuracy. No RF or DC positioners were used in this optical-to-optical (O–O) configuration. The dark box option provides an integrated light-isolation enclosure specifically designed for high-sensitivity optical measurements. It establishes a controlled, low-stray-light environment that minimizes ambient light interference and suppresses disruptive airflow and temperature gradients, increasing the environmental stability through the duration of the experiments. In addition to shielding the device under test from external illumination, it also serves as a light-absorbing environment that reduces internal reflections and suppresses parasitic optical resonances. This is particularly

important for measurements in the visible and near-infrared wavelength ranges, both at ultra-low power levels and during high optical power testing.

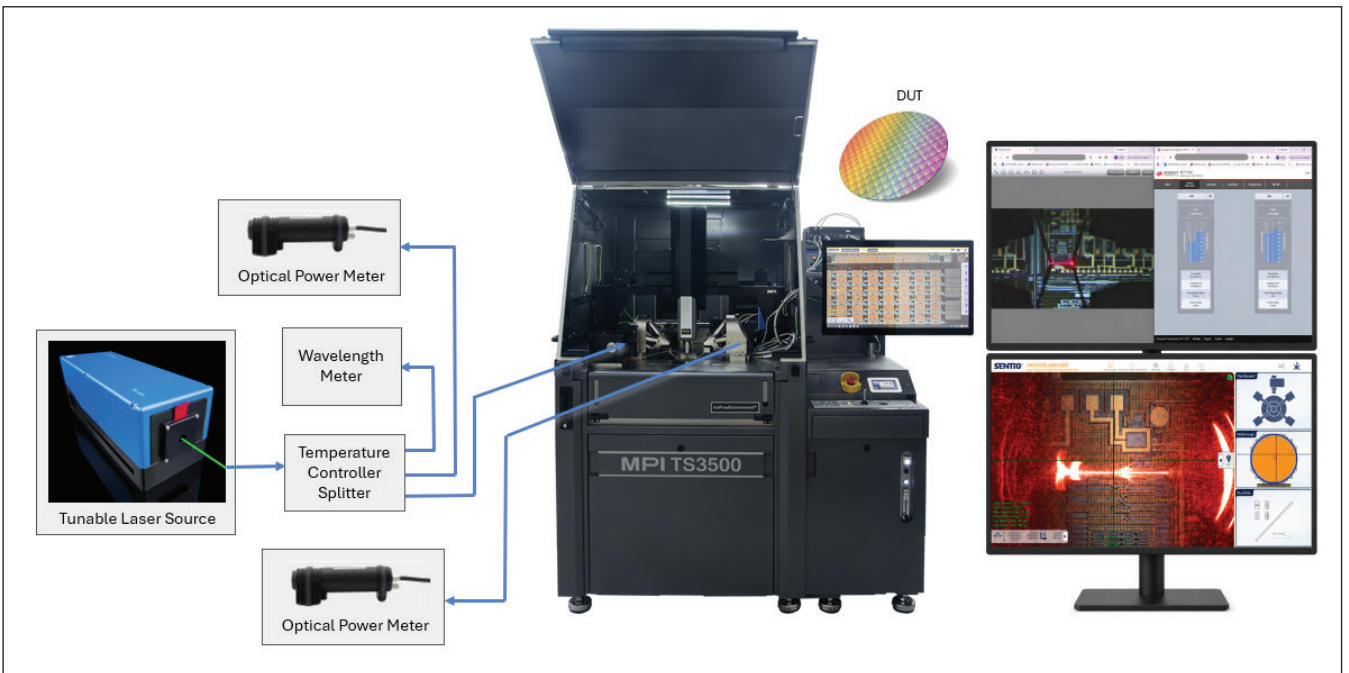
By improving optical isolation and environmental stability, the enclosure enhances measurement fidelity and repeatability while also supporting compliance with laboratory and manufacturing safety standards.

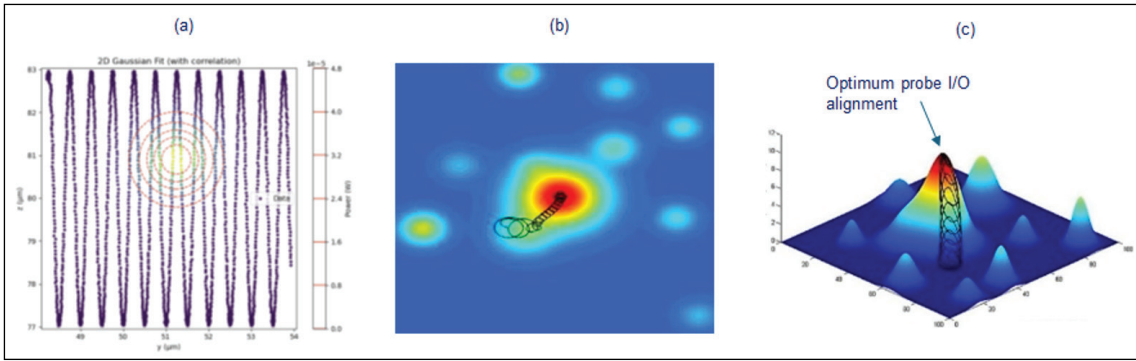
The calibration process for the physical fiber probe-to-wafer distance is fully automated and both temperature-controlled and temperature-compensated. It ensures precise and repeatable calibration of the probe-to-DUT distance and angular alignment across varying thermal conditions. The temperature sensitivity of the laser-to-fiber coupler, passive components, splitters, and fiber probe are also mitigated by this parallel path monitoring. This approach maintains alignment accuracy from ambient temperature up to +150°C, compensating for thermally induced mechanical drift and expansion effects to preserve consistent probe positioning.

Repeatability verification test

The alignment subsystem supports logic-driven, fully automated optical probe positioning and DUT alignment. This is achieved by a stationary, power-stabilized, and polarization-optimized laser source in combination with a single-channel optical power meter. The analog output of the power meter is directly interfaced with the SENTIO® environment of the MPI TS3500, enabling closed-loop feedback for high-precision, high-speed automated alignment.

(Figure 3) Image (a) shows the coarse scanning path of the optical I/O, confirming a Mode Field





➤ **Figure 3:** Hexapod multistage alignment using a hexapod for coarse scans and a nanopositioner for nanometer area scans and gradient search (b,c). This ensures optimum optical probe position prior to data collection. (b,c, Sample figures courtesy of MPI)

Diameter (MFD) of 1.66 μm , followed by fine alignment to achieve the optimum probe position. Note that the MFD can be fit after each alignment procedure and the probe to DUT distance can be adjusted to also approach the waist of the coupler's beam.

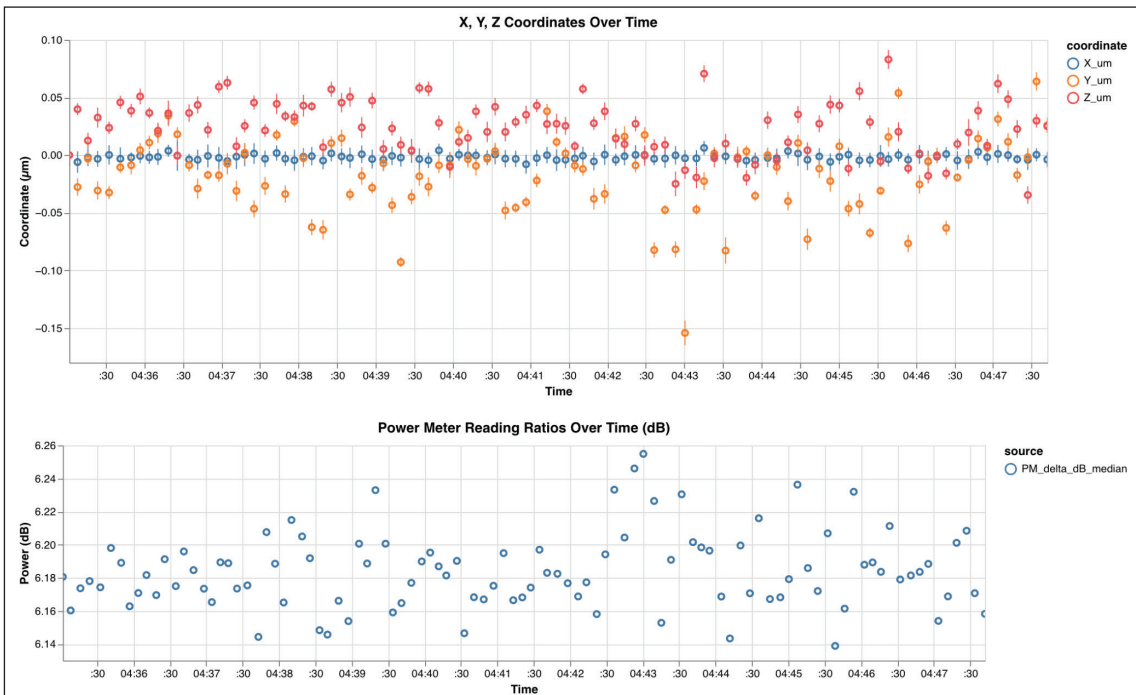
Running measurements with this optimized process demonstrated high accuracy and repeatability, as validated by one coarse-alignment measurement and 100 fine-alignment measurements (see graphs below). The consistency observed across successive runs confirms stable probe positioning and repeatable system performance. Each alignment took approximately 5 seconds, and consecutive alignments were spaced by 2 second intervals.

(Figure 4) Given an MFD of 1.66 μm , an excess loss of 0.02 dB corresponds to an approximate


positioning margin of ~ 60 nm from the peak coupling position in the YZ plane. It is worth noting that the excess loss in dB is inversely proportional to the square of the MFD value – assuming a typical SMF-28 MFD of ~ 10 μm , the excess loss due to a similar misalignment would be ~ 0.5 mdB!

The two-tiered optical power meter (OPM) setup effectively mitigates variability in the tunable laser output, which itself exhibits an intrinsic stability of < 0.03 dB. The two optical power meters jointly compensate for any variations in the test setup—both optical and mechanical—in real time.

The optical referencing of the two power meters was performed as a function of wavelength, ensuring that each wavelength step was normalized to a 0 dB reference point. This process was electrically synchronized with the tunable laser source and optically supported by the wavelength



➤ **Figure 4:** West positioned hexapod with nanopositioner measurement data on input of device under test (DUT). 100 YZ alignment followed by power measurements resulted in a 0.023 dB standard deviation optical insertion loss from OPM2-OPM1 (lower graph). The nanocube position readings (upper graph) after alignment remained within ± 60 nm of the average position in YZ and the X was not moved. The error bars around each position reading shows the accuracy of the readout. Similar measurements and plots for the east hexapod resulted in a 0.019 dB standard deviation optical insertion loss.



► **Figure 5:** (a) Shows the WAFT probes installed on the TS3500-IFE, while (b) shows a WAFT-to-WAFT calibration measurement operating around 630 nm and 5 μm separation. (c) Alignment routines over areas of 8x8 μm² with recorded position and analog power readings allow for the calculation of effective MFD when the WAFT is coupled through the trench to an inverse taper on the west (top figure) and east (middle figure) sides of a die. This is repeated at various separations for two waveguide thicknesses in different layers (bottom figure) to build a baseline for calibrating measurements to the same effective WAFT-to-edge distance.

meter, enabling accurate wavelength tracking and consistent power calibration across the full measurement range.

This referencing is typically performed as a one-time process step; however, periodic re-referencing may be required depending on environmental stability and test instrumentation specifications. In this QT setup, no re-referencing was necessary over the duration of the test runs. (For intervals see MPI application note AN62425).

This function is enabled by the 2x2 port thermally insensitive splitter, which provides temperature-independent power and polarization splitting ratios. QT used a special splitter constructed by a dielectric prism that maintains a reasonably flat splitting ratio across a large array of visible wavelengths which remains unchanged over the ±2 °C operating range.

Any wavelength variations from the tunable laser source during measurement runs were continuously monitored and recorded using an optical wavelength meter. This setup ensured accurate tracking of the operating wavelength and the input optical power alongside the DUT measurement. In QT’s approach, the DUT’s optical measurement was corrected ex post facto while preserving

measurement integrity throughout the test sequence resulting in an accuracy of ±1.5pm/±0.02dB (1σ). This approach minimizes the number of components in the optical path and obviates the need for real-time correction mechanisms for wavelength and power fluctuations. Active regulation of input power and wavelength is possible with feedback mechanisms at the expense of complexity and measurement speed.

Optical coupling mechanisms

Quantum Transistors used a range of custom grating-coupled and edge/trench-coupled fiber probes that accommodated wafer-level and die-level testing at scale with efficient optical coupling. Grating coupler devices emit light vertically with a spot size that matches that of a polished optical fiber, making alignment easier and fast. However, its coupling loss is high and very sensitive to wavelength, polarization, temperature, and fabrication variability. Edge couplers, on the other hand, have higher bandwidth and tolerance to temperature and polarization, but their spot size is lower than the fiber tip MFD, degrading coupling efficiency.


This edge coupler tradeoff was addressed using a wavelength-optimized Waveguide Array to Fiber Transposer (WAFT) fiber probe developed by Teem Photonics, specifically engineered to convert spot sizes between fiber and probe ends, and consequently increase optical port density on the die’s beachfront (Figure 5). They can also produce top-coupling WAFTs (TC-WAFT), which have an angled mirror finish at the probe facet designed for dual use: top coupling to grating couplers and trench-coupling into edge couplers on a wafer. Teem Photonics collaborated with MPI to produce custom-fit WAFTs matching the hexapod probe arm holders, resulting in stable coupling performance across varying wavelengths within the visible range. For the grating coupler coupling and die edge coupling QT used polished and lensed polarization maintaining fibers produced by Orbray Co., Ltd. of Japan (Figure 6).

End-to-End data collection and analysis pipeline

QT built a comprehensive data acquisition pipeline that scales the needs of production wafer testing and KGD qualification, providing automated control and minimizing manual intervention. Two key insights were instrumental. First, QT codified the measurement procedure into a deterministic “workflow” document that orchestrates measurement sequencing, alignment routines, and data validation steps, as well as decision trees for instrument fault self-recovery and automatic pausing and resuming of the data collection. The workflow can be inspected, simulated, and version-controlled before execution and attached to the metadata after execution. It also enabled test DUTs to be measured alongside calibration DUTs at the same environmental

conditions. Secondly, QT’s automation suite synchronizes instrument and telemetry data acquisition with timestamped logs that are streamed to an on-premises NoSQL database platform. This ensures low-latency data storage, logging, traceability, replication, backups, and controlled access across the company. With this data and telemetry acting as a system of record, post-processing routines can be run to normalize, filter, and correlate datasets into human-readable reports – often while the probe system is in the middle of scanning a wafer.

With this pipeline, QT demonstrated workflow throughput of approximately 6,000 DUTs per 24h, each with O-O high quality measurements, which is enough to verify wafer-scale device performance against design specifications, identifying fabrication-induced variations of key performance metrics. Consistent acquisition and automated data analysis at this level is critical for validating design parameters and quickly iterating over PDK elements and process adjustments with foundry partners, which is essential for meeting the stringent requirements of scalable quantum photonic and atomic systems, where even sub-dB-level deviations can significantly impact system performance and overall fidelity.



➤ **Figure 6:** MPI designed fiber holders to accommodate (a) lensed fibers and (b) polished fibers from Orbray Co., Ltd. of Japan that QT used for grating coupled measurements (c) or edge-coupled measurements (d).



➤ **Authors from left to right:** Lawrence van der Vegt, Thomas Ferreira de Lima, Benjamin Strekha, Igal Bayn




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





Integrated photonics confronts testing, packaging and scale

As photonic integrated circuits move toward larger-scale deployment in AI systems, communications and sensing, the industry is increasingly confronting the practical challenges of testing, integration and manufacturability. Discussions at PIC International 2026 highlighted how these pressures are reshaping optical metrology, packaging strategies and next-generation photonic architectures.

BY SARAB CHOPRA, EDITOR, PIC MAGAZINE

WHILE photonic integrated circuits continue to advance in capability and complexity, speakers at this year's PIC International made it clear that the industry's next challenge lies not simply in improving device performance, but in understanding how to test, integrate and manufacture these systems at scale.

Across the conference, attention increasingly focused on the growing disconnect between the rapid evolution of PIC architectures and the ability of existing testing and packaging strategies to keep pace.

As photonic devices incorporate larger numbers of optical functions, operate closer to physical limits and move into new application spaces, conventional approaches to characterization are becoming increasingly inadequate.

At the same time, the application landscape for integrated photonics is expanding rapidly. Alongside the continuing growth of silicon photonics for communications and AI infrastructure, researchers are now exploring nonlinear photonic platforms capable of extending integrated photonics deep into the ultraviolet.

Presentations during PIC International 2026 examined many aspects of this transition, including the challenges of high-volume PIC testing, the impact of co-packaged optics on optical metrology, and the development of aluminium nitride photonic circuits for far-ultraviolet light generation.

What became increasingly clear throughout the event is that the PIC industry is moving beyond the phase where success is measured primarily by demonstrating isolated

device functionality. Instead, the sector is now confronting the realities of scaling integrated photonics into manufacturable technologies that can satisfy the reliability, cost and performance requirements of communications, AI infrastructure, sensing and emerging quantum systems.

In many respects, integrated photonics is beginning to encounter challenges that mirror those faced by the semiconductor industry decades earlier. As circuits become denser and functionality becomes increasingly interconnected, the difficulty of understanding system behaviour rises sharply. Tiny process variations can alter performance, thermal effects become more difficult to manage, and packaging increasingly influences overall device characteristics.

At the same time, expectations placed on photonic systems continue to

rise. AI-driven data centres require dramatically higher bandwidths and lower power consumption, sensing applications demand greater precision and stability, and new nonlinear photonic platforms are pushing integrated optics into spectral regions that were previously inaccessible.

Against this backdrop, several presentations at PIC International suggested that the future competitiveness of the sector may depend less on isolated breakthroughs in device physics, and more on the industry’s ability to build robust ecosystems around testing, packaging, automation and integration.

Testing under pressure

Photonic integrated circuits promise many compelling advantages over discrete optical systems. They enable compact, multifunctional architectures with high-density integration, wafer-scale fabrication and increasingly sophisticated electro-optic functionality.

But according to Marc-André L’Aliberté, Marketing Director at APEX Technologies, these same advantages are also creating major challenges for optical characterization and production testing.

Speaking at PIC International, L’Aliberté explained that PIC functions are highly interdependent, making observability increasingly difficult.

“Functions cannot be isolated, effects are coupled, and one variable can

create multiple impacts,” he remarked, while noting that modern PICs operate increasingly close to physical performance limits.

This complexity is amplified by wafer-scale manufacturing, where small process variations can generate significant device-to-device performance differences. Thermal sensitivity, power fluctuations and packaging constraints further complicate characterization.

According to L’Aliberté, the challenge is not that the required measurements are new. Engineers still need to measure insertion loss, return loss, linewidth, chromatic dispersion, polarization effects and phase behaviour, all familiar quantities within photonics.

What has changed is the level of performance now required.

“Measurements are not new, but now devices operate closer to physical limits,” he explained.

One consequence is that no single measurement technique can provide a complete understanding of PIC behaviour.

For spectral analysis, grating-based optical spectrum analysers offer measurement speeds five-to-ten times faster than interferometric systems. However, interferometric approaches deliver dramatically higher resolution, reaching approximately 5 MHz and enabling the observation of side modes and comb lines.

Similarly, power and loss measurements involve trade-offs between spectral and spatial information. Tunable laser and power-meter systems provide fast, cost-effective insertion-loss measurements, while optical frequency-domain reflectometry offers high spatial resolution and distributed backscattering analysis.

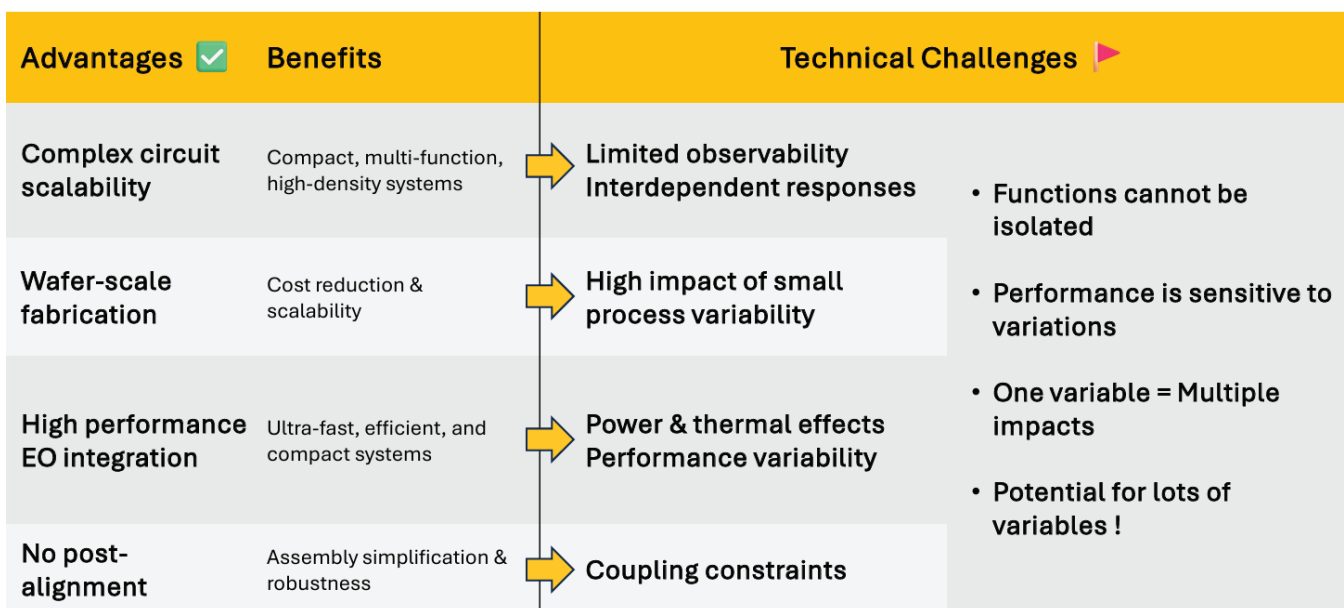
Polarization characterization introduces further complications. Depending on the application, engineers may prioritise speed, sensitivity, spectral coverage or measurement accuracy. As a result, testing increasingly relies on combinations of polarimeters, polarization state generators, polarization scramblers and optical vector analysers.

L’Aliberté argued that future PIC testing will depend heavily on integrating complementary measurement techniques into unified automated architectures.

“A complete testing solution requires application-specific tests combined and automated,” he explained, highlighting the growing importance of synchronized instrumentation, thermal control, probe stations and automated handling.

He also distinguished between the priorities of research laboratories and manufacturing lines.

In research environments, engineers require highly versatile systems capable of exhaustive characterization



and benchmarking. Production environments, by contrast, prioritise speed, robustness and scalable test strategies.

The emerging approach, according to L'Aliberté, is to first characterise the full parameter space during development, before reducing the test set through correlation and optimisation for production deployment.

That transition from exhaustive laboratory characterisation to streamlined production testing is becoming one of the defining engineering challenges for integrated photonics.

Historically, many photonic devices were produced in relatively modest volumes, allowing engineers to tolerate lengthy and highly specialised measurement procedures. But the emergence of silicon photonics for datacentres, co-packaged optics and sensing is rapidly changing those economics.

Wafer-scale manufacturing now demands test methodologies capable of operating at high throughput while maintaining increasingly stringent performance tolerances. This creates an inherent tension between measurement depth and manufacturing efficiency.

As one speaker noted during discussions at the conference, every additional measurement improves understanding of the device, but also increases testing time, equipment complexity and production cost.

For highly integrated PICs, the problem becomes especially acute because optical functions are no longer isolated building blocks. Losses, polarization effects, phase shifts and thermal interactions often influence one another simultaneously, making simplified testing strategies difficult to implement.

The growing importance of statistical process control also featured prominently in conversations around the conference. Rather than relying solely on pass-fail testing, manufacturers increasingly need large datasets capable of revealing subtle process drifts before yields begin to deteriorate.

This trend is likely to accelerate as PICs become larger and more heterogeneous, incorporating lasers, modulators, detectors and electronic control functions onto common platforms.

Co-packaged optics and the scaling problem

Another major theme at PIC International concerned the rapid emergence of co-packaged optics for AI and high-performance computing infrastructure.

Matthew Adams, Director of Product Management at VIAVI Solutions, explained how the integration of silicon-photonics engines directly beside switch ASICs is reshaping optical test requirements.

“Co-packaged optics is rapidly reshaping next-generation switch architectures,” said Adams.

The attraction of CPO is clear. By moving optics closer to the switching silicon, system designers can significantly improve bandwidth density and power efficiency, both critical metrics for AI data centres.

However, this tighter integration substantially complicates optical testing.

According to Adams, PIC test architectures must now adapt to three major trends: higher levels of functional integration; the use of external continuous-wave laser sources; and dramatic increases in lane counts.

Compared with conventional pluggable optics, co-packaged systems involve far more optical fibres per device under test. At the same time, testing increasingly shifts toward wafer- and die-level qualification.

“PIC yields drive parametric testing,” explained Adams.

Unlike pluggable modules, where much of the functionality is self-contained, co-packaged architectures require external digital and optical sources during testing.


As lane counts scale from 400G towards 1.6T and beyond, testing throughput is becoming a major bottleneck.

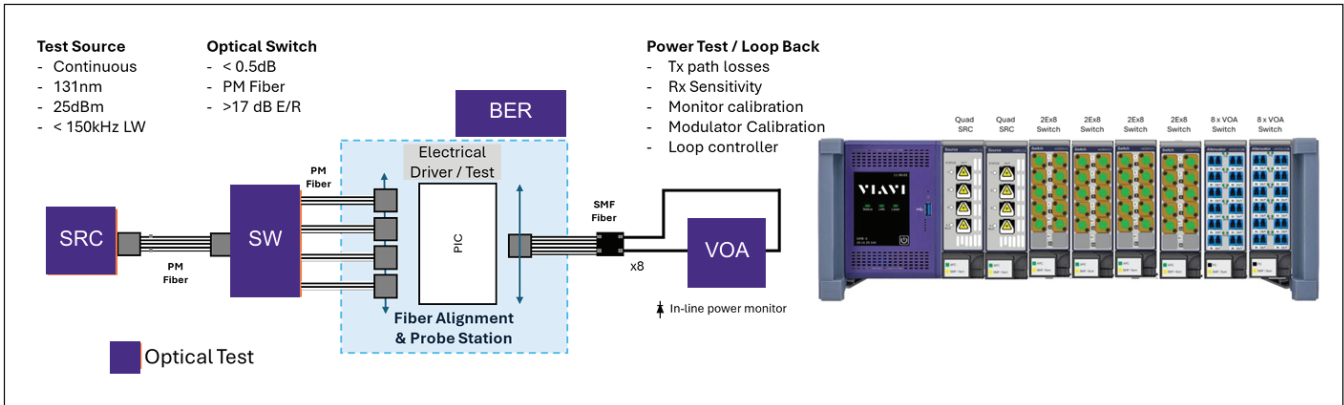
“Test times scale with the number of paths,” warned Adams, while noting that yield pressures are simultaneously driving tighter performance tolerances.

To address these challenges, the presentation examined modular metrology solutions spanning waveguide characterization, parametric testing and functional system validation.

Early-stage testing of silicon microrings employs swept-wavelength and polarization-sensitive measurements, while larger functional test systems support increasingly complex multi-lane architectures.

The growing importance of external optical injection sources for CPO testing was also highlighted. These sources require high output powers to overcome splitter losses, narrow linewidths to accurately characterise modulators, and high polarization extinction ratios.





One example discussed during the presentation delivered more than 24 dBm optical power with instantaneous linewidths as low as 33 kHz.

Packaging and connector integrity are becoming equally critical. The rise of co-packaged optics represents one of the clearest examples of how integration is fundamentally altering the photonics landscape.

For many years, pluggable optical modules provided a relatively clean separation between optics and electronics. Optical transceivers could be independently developed, tested and replaced, while electrical switching hardware evolved on a separate trajectory.

Co-packaged optics disrupts that model by moving optical engines directly beside high-performance switch ASICs. This shortens electrical interconnect distances and improves energy efficiency, but it also tightly couples optical and electronic design constraints.

As a result, optical testing increasingly becomes intertwined with system-level validation.

Instead of evaluating isolated optical modules, engineers must now characterise highly interconnected systems involving optical, electronic and thermal interactions simultaneously. This significantly complicates manufacturing workflows.

The increase in fibre density also presents substantial practical challenges. Large lane counts require highly parallel optical interfaces, tighter alignment tolerances and more sophisticated automated handling systems.

Even connector cleanliness becomes more critical as fibre counts increase. Small particles or scratches that might once have affected a single channel can now compromise much larger portions of the system.

Several presentations during the conference suggested that automation will become essential for managing this complexity. Automated inspection, alignment and synchronization are increasingly moving from desirable features to operational necessities.

At the same time, external laser sources are becoming more important within co-packaged optical architectures.

Unlike many traditional pluggable systems, some co-packaged approaches distribute optical power externally, placing greater demands on linewidth stability, polarization control and optical power management.

This shift is influencing the design of optical test equipment itself.

Test platforms increasingly require modularity, high port counts and the ability to support multiple insertion points throughout the manufacturing flow.

Beyond datacentres, many delegates also viewed co-packaged optics as part of a broader movement toward tighter photonic-electronic integration across the industry.


Whether in AI hardware, sensing systems or future quantum technologies, the ability to combine optics and electronics more closely is widely seen as essential for improving system efficiency and reducing interconnect bottlenecks.

Nonlinear photonics moves into the ultraviolet


As photonic systems move toward increasingly dense fibre-array interfaces, automated inspection and cleaning systems are becoming essential for identifying scratches, particulates and mobile contamination on connectors.

While many presentations focused on communications and AI infrastructure, another talk highlighted the expansion of integrated photonics into entirely new wavelength domains.

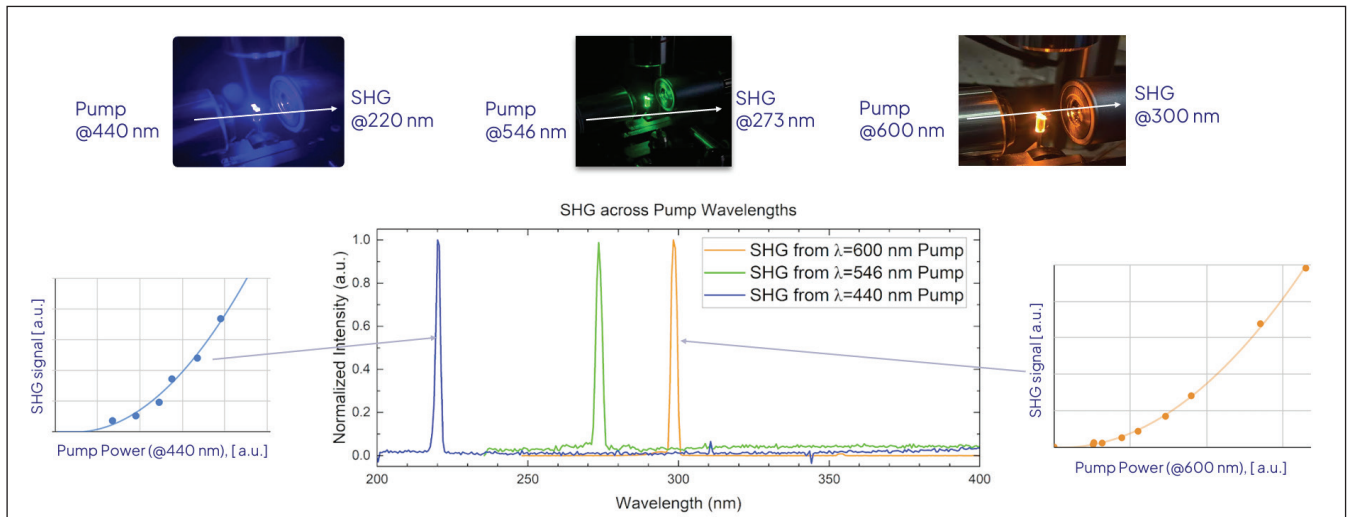
Scott Burroughs of Uviquity presented progress in aluminium nitride photonic integrated circuits designed for nonlinear frequency conversion into the far-ultraviolet spectrum.



AFM scan of MOCVD grown AlN-on-sapphire with RMS roughness of 0.5 nm (courtesy of Adroit Materials)



MOCVD grown AlN-on-Sapphire (1400 nm thick) wafer with Metricon measured intrinsic loss of 2.7 dB/cm at 447 nm



The work targets wavelengths around 222 nm, a region attracting significant interest because of its ability to inactivate pathogens while remaining safe for human skin and eyes.

Potential applications extend beyond disinfection to semiconductor inspection, environmental sensing, defence technologies and advanced photonics.

According to Burroughs, researchers have now demonstrated the generation of 222 nm far-UVC light on a scalable aluminium nitride photonic platform.

The approach relies on second-harmonic generation, in which photons from an integrated blue laser source are frequency-doubled within nonlinear AlN waveguides.

Efficient second-harmonic generation depends critically on phase matching, ensuring that generated harmonic waves remain constructively aligned during propagation.

Burroughs reviewed several phase-matching techniques, including quasi-phase matching, birefringent phase matching and modal phase matching.

Aluminium nitride is particularly attractive for this application. The material combines a wide 6.2 eV bandgap with transparency spanning approximately 0.2 μm to 5.5 μm , strong second-order nonlinearity and compatibility with scalable III-nitride semiconductor processing.

Researchers have already demonstrated low-loss AlN-on-

sapphire material grown by MOCVD, with intrinsic losses measured at approximately 2.7 dB/cm at 447 nm.

Fabrication employs standard e-beam lithography and chlorine-based dry etching to define waveguides.

Using this platform, second-harmonic generation spanning wavelengths from 214 nm to 300 nm has now been demonstrated.

According to Burroughs, multiple generations of samples representing thousands of devices and measurements have been produced.

The presentation also discussed a compact hybrid photonic package integrating a commercial blue laser diode with an AlN frequency-conversion chip.

Future work will focus on reducing waveguide losses, enhancing nonlinear coefficients using scandium-doped AlN structures, and pursuing heterogeneous integration of GaN laser sources directly onto the PIC platform.

Although much of the conference focused on communications and AI infrastructure, the work on aluminium nitride also illustrated how integrated photonics is steadily broadening beyond its traditional application areas.

Historically, many PIC platforms have concentrated on infrared wavelengths relevant to fibre-optic communications. Extending integrated photonics into the ultraviolet introduces an entirely different set of material and fabrication challenges.

Optical losses generally rise sharply at shorter wavelengths, fabrication tolerances become tighter, and suitable nonlinear materials are limited. Generating far-UVC light on-chip therefore represents a significant technical milestone.

The attraction of the 222 nm spectral region is considerable. Unlike conventional germicidal UV sources, far-UVC light appears capable of efficiently inactivating pathogens while remaining safe for occupied environments.

This has generated growing interest in applications ranging from hospitals and public transport systems to clean manufacturing environments.

At the same time, ultraviolet integrated photonics may ultimately find roles in spectroscopy, semiconductor metrology, environmental monitoring and advanced sensing.

One particularly notable aspect of the presentation was the emphasis on scalability.

Rather than demonstrating a laboratory-scale optical experiment, the work focused on creating a semiconductor-compatible photonic platform capable of supporting large numbers of devices and future packaging integration.

That distinction reflects a broader shift occurring throughout integrated photonics.

Increasingly, success is not determined solely by whether a device can function

in principle, but whether it can be manufactured reproducibly, integrated efficiently and operated reliably outside carefully controlled laboratory conditions.

A changing industry

Across PIC International 2026, a consistent message emerged: the photonics industry is entering a new stage of maturity.

For many years, the emphasis was on demonstrating the fundamental capabilities of integrated photonics. Today, attention is shifting toward scalability, reliability, manufacturability and application-specific optimisation.

Testing is becoming more complex as devices integrate larger numbers of optical functions onto single chips. Co-packaged optics is pushing measurement systems toward greater levels of automation and parallelisation. Meanwhile, advances in nonlinear platforms such as aluminium nitride are opening entirely new application spaces beyond traditional communications.

What united many of the presentations at PIC International was the recognition that future progress in integrated photonics will depend increasingly on engineering discipline rather than isolated proof-of-concept demonstrations.

The industry now faces challenges associated with repeatability, thermal stability, packaging yield, process variability and cost-effective manufacturing at scale. These are less visible than breakthroughs in headline device performance, but they are likely to determine how rapidly PIC technologies move into mainstream deployment.

At the same time, the breadth of applications discussed during the conference reflected the growing confidence surrounding integrated photonics.

No longer confined primarily to telecom infrastructure, PIC technologies are now influencing AI systems, advanced

sensing, quantum technologies, medical applications and ultraviolet photonics.

That expanding diversity may ultimately become one of the industry's greatest strengths. Different application spaces place different demands on materials, architectures and manufacturing methods, encouraging the development of multiple PIC platforms rather than a single dominant technological approach.

As PIC technology continues to evolve, the companies and research groups most likely to succeed may be those capable not only of advancing device physics, but also of solving the less glamorous yet increasingly critical challenges associated with testing, packaging and large-scale integration.

In that sense, PIC International 2026 suggested that the future of integrated photonics will be shaped as much by systems engineering and manufacturing infrastructure as by the photonic chips the photonic chips themselves.

DCS DATA CENTRE SOLUTIONS ROADSHOW

30 SEATS. 1 DAY. VIP STRATEGY DISCUSSION

THE DCS ROADSHOW 2026 is an exclusive executive forum, limited to 30 senior leaders responsible for data centre ownership, development, power strategy, and delivery across the UK.

The audience includes operators, hyperscale infrastructure teams, utilities, developers, EPCs, OEMs, infrastructure investors, and a limited number of independent advisors.

WHAT MAKES THE ROADSHOW UNIQUE?

- **Peer-to-peer Learning:** First-hand insights from industry leaders
- **No Vendor Sales Pitches:** Strategy-led discussions only
- **Curated Networking:** Build meaningful, high-value connections
- **Intimate Format:** Just 30 delegates for focused collaboration

LIMITED PLACES AVAILABLE

Apply for your complimentary pass to the **DCS Roadshow 2026** in Cardiff at: <https://datacentreroadshow.com/events/cardiff-2026/register>

Interested in speaking or attending? If you have any questions about attending as a delegate or speaker, please reach out to info@datacentreroadshow.com



03 SEPT 2026

Photonics at the heart of AI data centers and beyond

AI-driven computing requires efficient optical interconnects to overcome data-center bottlenecks. imec advances silicon photonics technology to address such a need. Furthermore, ultra-low loss SiN photonics is considered as a key technology to support the next generation quantum-based computing. The same platform can also be used for augmented and virtual reality and bio sensing.


BY LEILI SHIRAMIN, PORTFOLIO MANAGER-INTEGRATED PHOTONICS, JORIS VAN COMPENHOUT, FELLOW, PHILIPPE ABSIL, VICE PRESIDENT, IMEC

► **Figure 1:** S21 measurement of Ge Photodiode showing 3-dB bandwidth of 110GHz.

THE exponential growth of data generated across industries, combined with the rapidly increasing complexity and scale of AI models, has created significant bottlenecks in both training and inference workloads. As AI systems continue to evolve toward larger parameter counts and more demanding computational requirements, traditional computing infrastructures struggle to provide the necessary bandwidth, latency, and energy efficiency. This challenge has accelerated answering the need for large-scale GPU clusters and high-performance computing (HPC) systems to support next-generation AI factory interconnects. Consequently, both scale-out interconnects, which connect multiple servers and racks, and scale-up interconnects, which enable communication within a server or accelerator domain, is re-evaluated to meet these growing performance demands.

Current architectures for building optical communication links generally rely on two major approaches: the “Fast and Narrow” architecture and the “Wide and Slow” architecture. In the Fast and Narrow approach, data transmission is achieved through very high symbol rates over a relatively small number of lanes. This architecture requires digital signal processing (DSP), high-speed DAC/ADC components, and advanced high-speed electronic circuitry (Driver, TIA) to compensate for signal impairments and maintain reliable communication performance.

While this method can achieve high aggregate bandwidth density, it often comes at the cost of increased power consumption, higher latency, and larger system complexity due to the electronic processing involved. In contrast, the Wide and Slow architecture distributes data transmission across a larger number of parallel optical lanes operating at lower symbol rates. Because each lane operates more slowly, the architecture eliminates the need for DSP blocks, DAC/ADC modules, transimpedance amplifiers (TIAs), and high-speed driver circuits. As a result, the Wide and Slow approach offers advantages in terms of lower power consumption, reduced latency, and simplified electronic design. These characteristics make it an attractive solution for future AI and HPC interconnects where energy efficiency and communication latency are becoming design constraints.



However, although the Wide and Slow approach simplifies the electronic subsystem, the associated complexity shifts toward the optical and packaging domain, particularly at the fiber-to-chip interface and in high-density fiber coupling environments. Maintaining low insertion loss, minimizing crosstalk, and achieving scalable manufacturability in dense optical interconnect systems require advanced packaging and coupling techniques. In addition, other set of optical devices and components, such as 32-channel or 64-channel wavelength-division multiplexing (WDM) filters, multiplexers become necessary to manage the large number of optical channels.

Imec, as a technology provider, is developing solutions for both photonic integration architectures. To address the bandwidth requirements of next-generation 400G-per-lane pluggable optics and co-packaged optics (CPO), modifications of existing silicon photonics technology platforms are necessary. For Fast & Narrow, the focus is on achieving the target specification of 110 GHz bandwidth for both modulators and photodiodes and on for Wide & Slow architecture combining active devices operating at several tens of GHz with scalable multi-channel low loss passive photonic components is considered.

At the component level, germanium (Ge) photodiodes continue to remain a preferred technology choice, with the potential to support performance beyond 400G (>110GHz bandwidth). Imec's recent results, demonstrate an electro-optical 3 dB bandwidth of 110 GHz (Figure 1), while maintaining a responsivity exceeding 0.9 A/W and dark current below 15 nA. The device also exhibits low parasitic capacitance (<10 fF) [1]. When combined with Cu-oxide hybrid bonding technology featuring similarly low parasitic capacitance (<10 fF) and a co-designing with transimpedance amplifier (TIA), these photodiodes enable highly sensitive optical receiver architectures.

For even higher receiver sensitivity, avalanche photodiodes (APDs) are being considered as a future addition to the silicon photonics platform. Initial APD results are highly promising,

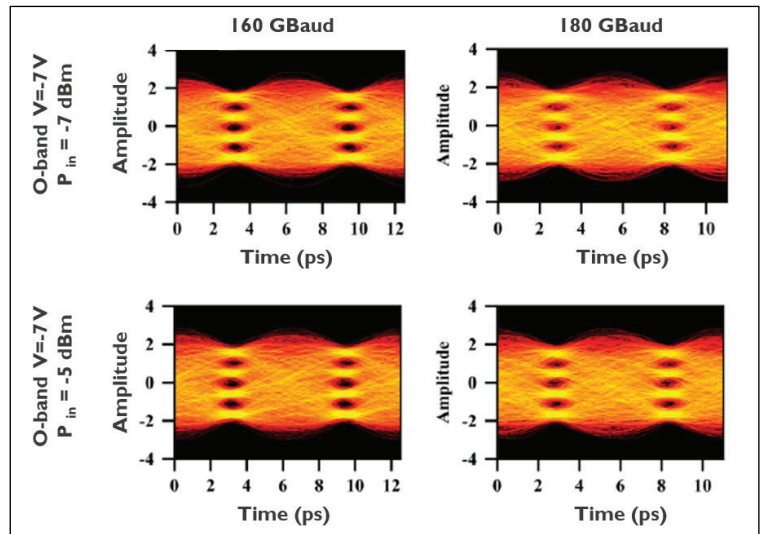


Figure 2: Eye diagram measurements corresponding to responsivity values of 2 A/W and 1.5 A/W for the O-band and C-band APDs, respectively at different data rates of 160 GBaud PAM4 and 180 GBaud PAM4.

demonstrating operation up to 180 Gbaud PAM4 in both the O-band and C-band while maintaining bit error rates (BER) below the hard-decision forward error correction (HD-FEC) threshold (Figure 2) [2].

On the transmitter side, several modulation technologies are being explored. Candidate technologies include lithium-niobate-on-silicon (LNO-on-Si) Mach-Zehnder modulators (MZMs), III-V-on-silicon electro-absorption modulators (EAMs), as well as silicon-organic hybrid (SoH) modulators. Each technology offers unique advantages and trade-offs. The primary challenges are related to heterogeneous integration complexity, contamination constraints associated with LNO and III-V materials, and passivation, poling complexity, and long-term reliability concerns for SoH devices.

Franz-Keldysh (FK)-based GeSi EAMs and silicon-based micro-ring modulators (MRMs) and MZMs are also being investigated as alternative solutions. However, FK-based GeSi EAMs are limited to C-band operation due to their material bandgap properties, making them more suitable for scale-up connectivity applications only. Silicon-based modulators continue to attract industrial interest because they avoid additional heterogeneous integration complexity and eliminate the need for dedicated fabrication toolsets within CMOS fabs. Nevertheless, achieving

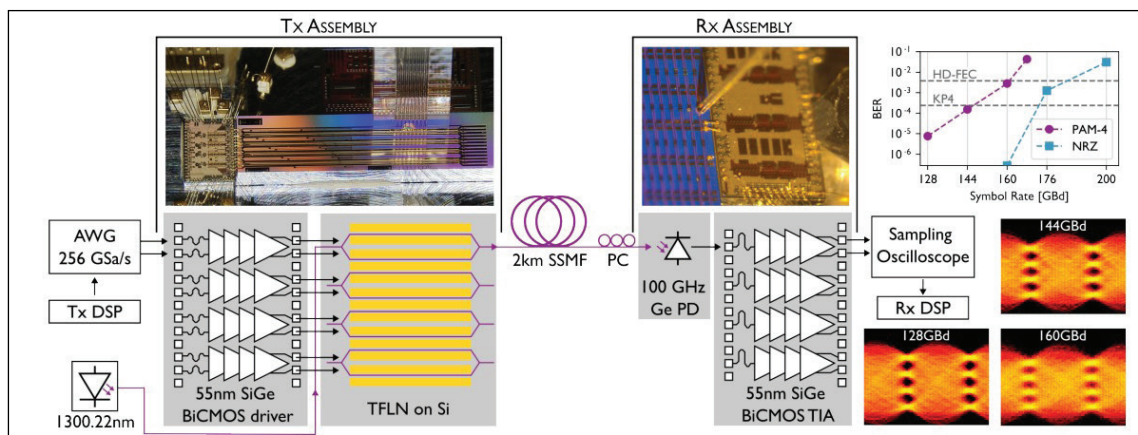




Figure 3: Transmission measurements showing transmit and receive assemblies along with eye diagrams and Bit Error Rate curves.

► **Figure 4:**
Above:
S21
measurement
showing 3-dB
bandwidth
exceeding 110
GHz.

Below:
TSV Test
structure



the required performance while managing trade-offs among critical device parameters remains challenging for 110 GHz-class devices.

The current bandwidth performance of several modulators in imec technology platforms is summarized as follows:

- Silicon MZM (O-band): 75 GHz
- Silicon micro-ring modulator (O-band): >67 GHz [3]
- GeSi EAM (C-band): 110 GHz [4]

Imec continues to further optimize these modulators, with primary focus currently on Si MZMs and LNO-based modulators.

Integration of emerging materials such as lithium niobate onto silicon photonics platforms is being pursued through several approaches, including die-to-wafer (D2W) bonding, wafer-to-wafer (W2W) bonding, and micro-transfer printing (mTP). The results presented in this work for both LNO and III-V integration utilize micro-transfer printing as one of the key heterogeneous integration techniques.

The current status of the LNO modulator (an under-development device in imec) demonstrates a 3 dB bandwidth of 70 GHz, a $V_{\pi} \cdot L$ of 2.5 V.cm, and insertion loss (IL) below 3 dB. The target specifications are 110 GHz bandwidth, $V_{\pi} \cdot L$ below 2 V.cm, and insertion loss below 1 dB. The next phase


of development will focus on reducing $V_{\pi} \cdot L$ to below 1 V.cm while maintaining bandwidths exceeding 110 GHz.

Using micro-transfer-printed LNO device integrated on silicon photonics together with high-speed germanium photodiodes, imec demonstrated a 2 km PAM4-modulated O-band optical link operating at data rates up to 320 Gb/s while achieving BER values below the HD-FEC threshold. In addition, the demonstrated circuit incorporates imec's custom-designed traveling-wave drivers and TIAs, enabling complete end-to-end link characterization (Figure 3) [5].

In the area of passive photonic device development, imec has demonstrated ultra-low-loss polarization beam splitters with insertion loss as low as 0.019 dB and extinction ratios reaching 29 dB. Additional developments include broadband directional couplers supporting 100 nm bandwidth and low-insertion-loss 32-channel wavelength-division multiplexing (WDM) filters with 100 GHz channel spacing. Results for 64-channel devices are expected to be published soon.

The introduction of 3D integration features into the technology platform is also critical for next-generation compact products. Through-silicon vias (TSVs) represent one of the key enabling technologies for efficient power delivery and high-speed signal transmission. Depending on RF loss, bandwidth, and interconnect density requirements, slightly different integration flows are adopted. Figure 4 presents the extracted 3 dB bandwidth of a TSV test structure consisting of Metal + TSV + backside redistribution layer (BSRDL) + TSV + Metal. The measured bandwidth exceeds 110 GHz, making the technology highly suitable for next-generation 3D-interconnected chips used in pluggable optics and co-packaged optics systems.

From the perspective of future optical interconnect evolution, III-V integration on silicon photonics is becoming essential to provide access to gain sections required in certain data center applications or architectures. To address this requirement, imec is actively developing the integration of both laser sources and semiconductor optical amplifiers (SOAs) as part of its roadmap. Initial experimental results demonstrate the strong potential of this heterogeneous integration platform.



► **Figure 5:** Left: Bright-field microscope image of a modelocked laser and Dark-field microscope image of the same laser under 50 mA gain current with forward-biased saturable absorber. Right: RF comb at optimal mode-locking point and Fundamental RF line, measured with resolution bandwidth of 100 Hz.

Beyond optical interconnects, imec is also developing a low loss silicon nitride (SiN) platform to enable emerging application domains. One important development is a fully integrated laser operating at 800 nm, targeting applications such as augmented reality/virtual reality (AR/VR) systems and optical atomic clocks. Figure 5 shows the fabricated device together with butt-coupled GaAs-based amplifiers operating at 800 nm integrated with saturable absorbers and silicon nitride [6].

These devices successfully generate RF frequency combs after transfer printing, as illustrated in Figure 5, including a zoomed-in view of the fundamental RF tone exhibiting an extinction ratio of approximately 50 dB, currently limited by the measurement noise floor. These results demonstrate the strong potential of the proposed integration methodology for high-power laser applications utilizing the substrate as an efficient heat dissipation pathway.

As the next step toward further technology enablement and increased maturity level across multiple application domains, imec will continue advancing chip-to-chip interconnection technologies to help address emerging industry challenges.

REFERENCES

- [1] C. Coughlan et al, OFC 2026.
- [2] A. Shahin et al, OFC 2026.
- [3] A. Ostrovskis et al, Opt. Express, 2025.
- [4] C. Bruynsteen et al, ECOC 2025.
- [5] J. Declercq et al, ECOC 2025
- [6] M. Kiewiet et al, LPR, 2025.



- Authors from left to right: Leili Shiramin, Joris Van Compenhout, Philippe Absil



ANGELTECH

CS INTERNATIONAL CONFERENCE

PIC INTERNATIONAL CONFERENCE

PE INTERNATIONAL CONFERENCE

AP INTERNATIONAL CONFERENCE

AIS ANGELTECH INNOVATE SUMMIT

SAVE THE DATE

12-14 APRIL 2027

Sheraton Brussels Airport Hotel

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

or email:

info@angeltech.net

Enabling THz bandwidth on-chip with plasmonic modulators



Plasmonic electro-optic devices are emerging as a route to sub-THz and THz-class bandwidths, offering a compact, low-power alternative to traditional silicon photonic modulators for next-generation data center interconnects.

BY STEPHAN KOCH, VP MARKETING & SALES AT POLARITON TECHNOLOGIES IN ZURICH, SWITZERLAND. POLARITON IS NOW PART OF MARVELL SEMICONDUCTORS WITH A PLAN OF BRINGING PLASMONIC MODULATORS TO DATA CENTER APPLICATIONS

GROWING DEMAND for higher data rates in intensive data center operations is accelerating the development of advanced system components, including electro-optic modulators designed to support THz-regime operation beyond conventional communications applications.

Next-generation modulation is increasingly a competition among material platforms, several of which have credible paths to adoption.

Traditional III-V components, together with emerging thin-film lithium niobate (TFLN), barium titanate (BTO), and polymer-powered devices, are all expected to play a role. In an environment constrained by production capacity, multiple platforms are likely to coexist, with a clear advantage for solutions that combine cost efficiency with a scalable supply chain.


Plasmonic modulators use polymers as the non-linear material that

produces the desired phase shift of coherent light. Their micrometer-scale dimensions are central to device performance, particularly in relation to bandwidth and power dissipation.

Plasmonics is silicon photonics

The modulator's non-linear material is arranged coplanar to the silicon photonics waveguide within a slot bounded by two plates fabricated from silicon-compatible metals. Gold is commonly used, although other wafer-fab-compatible alternatives are also feasible. The plasmonic modulator is added during post-processing of standard silicon photonics wafers, which can be sourced from multiple silicon foundries [1], and requires only a limited number of lithographic steps. This approach enables electro-optic devices that provide:

- bandwidth suitable for sub-THz operation;
- reduced signal processing for equalization, or alternatively greater tolerance to noise in the communication channel;
- micrometer-scale active devices that support efficient use of chip area;
- lower switching power.




Within the active device, the electrical electrodes also function as the optical waveguide. This compact design is effective because the optical and electrical driving fields strongly overlap, resulting in high modulation efficiency and enabling short modulator lengths. Electrically, the modulator can be represented as an RC circuit with a corner frequency in the THz range, determined solely by the RC time constant.

From a physical perspective, the device width is smaller than the diffraction limit of the light used for communication, typically 1.3 or 1.55 μm . As a result, the relevant operating principle is plasmonic rather than purely photonic, allowing operation at smaller dimensions. The photonic mode is tapered into a plasmonic mode that propagates along the metal-dielectric interface of the slot before being converted back into a photonic mode.

These structures are sometimes referred to in the industry as plasmonic organic hybrid (POH) devices. In this article, however, we use the term plasmonic electro-optic modulator because the underlying physics relies on surface plasmon polaritons (SPPs). In 2025, researchers at ETH Zurich reported a record for an electro-optic modulator, achieving a 3 dB bandwidth of 997 GHz [2].

The prime application in the AI era

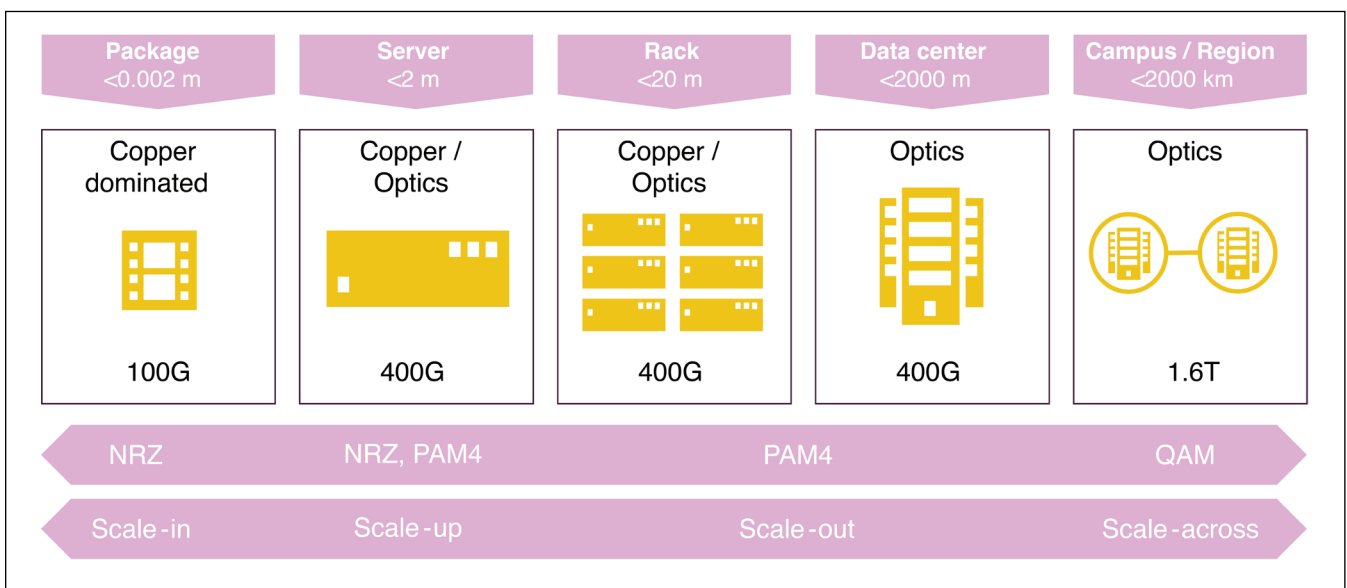
Plasmonic modulators have long been developed for fiber-optic



communication, among other applications, because of their substantial bandwidth. The current increase in demand from data centers has placed silicon photonics at the center of efforts to address segments in which optical links can replace copper interconnects. Copper will nevertheless continue to coexist with optical technologies within servers and racks because of its cost and

power-dissipation advantages, while optical links are expected to gain share wherever greater reach and bandwidth are required.

Within data centers, system-component dimensions, integration density, power dissipation, and communication protocols are increasingly important design considerations. As a result, equipment is becoming more



Criteria	At	Silicon MZM	TFLN MZM	Silicon Organic Hybrid (SOH) MZM	Plasmonic MZM
Enough bandwidth	400G and 800G per lane	-	-	-	Strong
Less DSP resources needed	400G	-	Strong	-	Strong
Micrometer scale modulators		-	-	-	Strong
Switching power benefit		-	-	-	Strong
Scale-up fitness	400G	-	-	-	Strong
Scale-out fitness	400G	Strong	Strong	Strong	Strong
Scale-across fitness	1.6T	Strong	Strong	Strong	Strong

specialized by application, and standards and implementation approaches are becoming more segmented. The industry commonly describes these segments as scale-in, scale-up, scale-out, and scale-across. At short reach, most communication remains copper-based; optical links become more prevalent as distances increase across racks and clusters. Hyperscalers have a clear interest in extending copper connectivity at the rack level for as long as practical, although the transition toward optical interconnects is becoming increasingly difficult to avoid.

Scale-across is a relatively recent segment definition, introduced approximately one year ago, and its precise boundaries remain under discussion across the industry. Historically, coherent communication was used primarily to connect hubs and data centers across regions. More recently, it has begun to move inside data-center premises, particularly in implementations that require lower signal-processing overhead and can benefit from the shorter link distances.

Silicon photonics implementations

Optical communication can be implemented using a range of modulators that are more or less native to silicon photonics. This article compares Mach-Zehnder modulator (MZM) implementations in terms of their fitness for current and future communication speeds. The next generation of optical communication is expected to support 400G per lane, with data transmitted using four-level

pulse-amplitude modulation (PAM4) and SerDes operating at 224 GBd, corresponding to 448 Gbit/s transmitted per lane. The principal advantages of plasmonic modulators arise from their micrometer-scale implementation. Smaller active devices can support very high operating speeds while reducing the energy required for charging and discharging. Their compact footprint also provides an integration advantage when many lanes must be placed in close proximity, as in co-packaged optics for scale-up applications. In addition, the flat transmission spectrum contributes to improved signal quality by reducing distortion along the communication path.

At the Max-Planck-Institute for Plasma Physics in Greifswald, in collaboration with the Chalmers University of Technology in Gothenburg, they are pursuing higher magnetic fields in confinement fusion experiments by driving the electron cyclotron emission experiments well beyond the 100 GHz range

Exploring THz domain applications

The pull from AI-driven innovation is benefiting more applications. Several entities are pushing the operation into the THz regime. At Leapwave in Madrid, they are developing dielectric wires for communication and test & measurement, hence enabling boundaries to 500 GHz and more. Interfaces between the components are essential, and they have developed ultra-wideband approaches for that.

Another company, AttoTude in Menlo Park, is combining THz radio and the innovation of THz interconnect. The playground is the Terahertz gap, which is the part of the spectrum between what has traditionally been addressed by RF engineering and what is served with photonics.

Further, the EU project ECO-eNET explores ultra-high-speed wireless links to bridge gaps between fiber networks using THz frequencies. Instead of converting signals to electrical, the approach keeps them in the optical domain to minimize losses. The goal is to achieve 200–300 GHz carrier frequencies.

Finally, at the Max-Planck-Institute for Plasma Physics in Greifswald, in collaboration with the Chalmers University of Technology in Gothenburg, they are pursuing higher magnetic fields in confinement fusion experiments by driving the electron cyclotron emission experiments well beyond the 100 GHz range. This is the result of the availability of cost-effective telecommunication technology for sub-THz-to-optical up-conversion.

Polymer development progresses fast

Electro-optic material manufacturers are also gaining increasing relevance, as their materials provide the essential non-linear response required for modulation. These materials are primarily responsible for converting electrical signals into optical signals, and vice versa, with new classes of polymers emerging as particularly well suited to this function.

Polymers are organic materials with strong optical and electrical properties, developed over more than two decades and now progressing toward commercial adoption. Earlier

development focused primarily on dielectric performance, followed by substantial advances in easy of manufacturing and thermal stability at elevated temperatures.

Moreover the pace of polymer development exceeds that of new crystal engineering, creating significant potential for continued performance improvements over the coming years.

Enough bandwidth for generations

Silicon photonics has historically faced scrutiny in high-speed optical communication because it lacks native high-speed modulators. However, continued advances in both silicon photonics and integrated material platforms now enable 400G-per-lane implementations and position silicon as a strong foundation for the next several generations of transceivers. While multiple material platforms are competing for market adoption, silicon-based approaches retain

FURTHER READING

- [1] [Silicon Photonics Dominates Optical Transceiver Sales | Stephan Koch post | 2026](#)
- [2] [Tiny component for record-breaking bandwidth](#)

important advantages in cost structure and supply-chain scalability. The shift toward smaller photonic integrated circuits (PICs), including ring resonator modulators (RRMs) and micrometer-scale plasmonic modulators, can also reduce wafer demand, supporting more efficient resource use in an industry constrained by manufacturing capacity.

More compact PICs also support higher levels of integration, including 32 or 64 lane architectures, which are becoming important for co-packaged optics (CPO) and near-packaged optics (NPO) implementations.

Polymers are also important enablers for high-speed modulators,

offering strong dielectric and electro-optic properties that extend into the THz region. The rapid development of new polymer generations, together with the emergence of multiple material providers, is contributing to a broader reshaping of the industry.

Micrometer-scale plasmonic PICs provide the bandwidth required for 400G, 800G, and 1.6T transmission per lane, supporting multiple future generations of photonic transmitters. When combined with application-specific drivers and digital signal processors, they can also deliver advantages in switching power.



WEBINAR

Contact: Jackie Cannon

E: jackie.cannon@angelbc.com

YOUR EXPERTISE. OUR AUDIENCE. REAL IMPACT.

A PIC webinar helps you share knowledge, generate quality leads and build global brand awareness within the industry.



BRAND AWARENESS

Increase global visibility and credibility



LEAD GENERATION

Generate and nurture qualified leads



THOUGHT LEADERSHIP

Position your experts as industry leaders

OUR 8-WEEK WEBINAR PROMOTION PROGRAMME INCLUDES



Promotion in
PIC magazine
promotion



Targeted exposure
through newsletters to
30,000+ professionals



Promotion in our
active LinkedIn
Group



Featured on
the *PIC*
website



Moderated by
the editor of
PIC

Ncodin's nanolasers eye AI infrastructure

Delivering a revolutionary energy-per-bit, the French start-up's nanolasers are a compelling source for on-chip communication in tomorrow's data centres

BY RICHARD STEVENSON, EDITOR, CS MAGAZINE

SUPPORTING increasing adoption of AI is the building of more and more data centres. This roll-out of essential power-hungry infrastructure already accounts for hundreds of billions of dollars per annum, and total costs are forecast to exceed a trillion dollars a year by the middle of the next decade.

These eye-watering sums offer a fantastic opportunity for companies with technologies that can enhance critical hardware, especially if these products deliver breakthroughs in performance and efficiency.

One of the biggest issues in today's datacentres surrounds the transfer of data between chips. Although widely deployed for decades, copper is a bottleneck to higher speeds. And that's not its only downside, with other significant weaknesses including a limited reach and losses that lead to a significant contribution to overall energy consumption.

Offering attractive alternatives that address all these issues are various forms of optical links. There are a number of options for the light source – including edge-emitting lasers, VCSELs and microLEDs – and photons can be routed through optical fibres or waveguides in wafers.

Amongst these competing solutions, French start-up Ncodin, which has just raised €16 million, believes it has the winning formulation: miniature lasers that launch their emission into waveguides in silicon-based wafers.

Optical interposers formed by this approach, featuring incredibly small lasers and photodetectors – both are created by bonding InP epiwafers to silicon substrates and subsequent photolithography – are claimed to deliver unrivalled efficiencies. Data transfer is incredibly fugal, at less than 0.1 pJ/bit, and integration can exceed 10,000 mm², enabling breakthroughs

in performance on two fronts – what is possible per Watt, and per dollar.

A pioneering PhD

The origins of Ncodin can be traced back to the labs at the Centre for Nanosciences and Nanotechnologies, located in the southern suburbs of Paris, where co-founding CEO, Francesco Manegatti, developed these nanolasers during his PhD days.


Enthused by the emergence of a start-up ecosystem – in 2017, colleagues from this lab founded the quantum computing spin-off Quandela – Manegatti shifted his post-graduation plans from joining a big company to trying to commercialise the nanolaser technology he'd developed, a vision shared by his supervisor, now chief scientific officer and co-founder Fabrice Raineri.

Efforts in this direction began in late 2019 and continued during the pandemic, when these entrepreneurial partners started to try and secure initial funding and develop a first roadmap towards commercialisation.

The next steps involved: working with the support of CNRS on an R&D project; establishing a business plan; and bringing in third co-founder, Bruno Garbin, who switched roles from a postdoc in the laboratory to Ncodin's CTO.

Founding of the start-up followed in 2023, with Ncodin raising €3.5 million in a pre-seed funding round that closed in March 2024. Since then, headcount has mushroomed from four to 30 employees, and the company has established its own headquarters.

"We still exploit the facilities of CNRS, because this lab has the largest



➤ Ncodin's silicon photonic chip integrating nanolasers, the core technology showing record energy efficiency below 0.1 pJ/bit. The wires provide a connection between the controller (an FPGA) and the nanolasers.

academic clean room of France. It's a 3,000-squared-metre clean room, where they have a lot of collaboration with private entities," explained Manegatti.

Ncodin uses this facility to develop its chip technology, and make proof-of-concept devices that are shared with partners.

"In parallel, we are working on the industrialisation of our technology, to reproduce it in a CMOS pilot line on 300 millimetre [wafers]."

The start-up has always focused on optical interconnects. Neuromorphic architectures initially garnered consideration, but the primary goal is to complement copper and aid the scaling of infrastructure and process architecture.

"For extremely short reach, copper is best," argues Manegatti. "It's super resilient, it's super robust, you can deliver a high bandwidth."

But for connections beyond a few centimetres, links should switch to the optical domain, where Ncodin's solution addresses the challenge of transferring tens of terabits per second of bandwidth while consuming just a fraction of a picojoule per bit.

This is an attractive option for AI workloads, which are fulfilled through continuous communication between the compute element, which could be a GPU or an ASIC, and the high-bandwidth memory. For these tasks, nano-lasers promise to play a role in fulfilling demands for extremely fast retrieval of data, as well as the writing of data in the memory.

Numerous nano-lasers

To produce its devices, the French start-up draws on external expertise for the growth of its epiwafers. Multiple coupons are crafted in these epiwafers, prior to hybrid bonding to silicon-based wafers. Subsequent lithography and patterning define the dimensions of lasers and photodetectors, both operating at telecom wavelengths. The emission from the lasers, which have a footprint that is 500 times smaller than their conventional cousins, evanescently couples into the waveguides in the underlying silicon wafers.



➤ Since its founding in 2023, headcount at Ncodin has increased to 30 employees.


Nanolasers produced by this process have an output of up to a few hundred microwatts – that's more than enough for the intended task – and an efficiency of around 15 percent. These low emitting powers prevent self-heating issues for these miniature sources, which are capable of delivering robust operation at temperatures approaching 100°C.

Thanks to the incredibly high density that can be realised with Ncodin's technology, the intention is to connect every input and output pin of every ASIC and every high-bandwidth memory to a nanolaser.

"Our goal is, of course, to add redundancy, to extend even further the lifetime and essentially the reliability of our chip," adds Manegatti.

The nanolasers are directly modulated, as it's not critical to have a high data rate per channel, thanks to the use of so many channels.

According to modelling by Ncodin, it's possible to realise data rates of up to 64 Gbit s⁻¹. However, the nanolasers currently operate at 16 Gbit s⁻¹.



➤ Ncodin's core technology is the production of telecom lasers with a footprint that is 500 times smaller than their conventional cousins.

"This is more than enough," argues Manegatti. "High-bandwidth memory, for example, is driven at 8 gigabits-per-second today."

Before Ncodin's optical interposers are deployed in data centres, they need to have proven reliability. Efforts in this direction are on-going, according to Manegatti, who remarks: "We are industrialising the solution, so the most important KPIs will be extracted during this project." Note, though, that his team have already demonstrated lifetimes that are well beyond a few thousand hours.

The business plan

To produce its technology in volume, Ncodin will partner with a dedicated foundry to manufacture wafers packed with optical interposers.

"We'll sell the wafers to our customers, so chipmakers like Nvidia, Qualcomm, AMD, Intel, *et cetera*," says Manegatti. These household names will employ partners to test, assemble and package entire systems that combine Ncodin's wafers with a variety of chips, such as those for memory and processing.

Manegatti says that the main goal for now is to finish the industrialisation project, undertaken with a partner, and start qualification of its product.

"We're going to open an office in Silicon Valley, to be closer to our customers and partners, and to be able to engage in an even deeper relationship with them."

These plans, which make a lot of sense, promise to enable III-Vs to play yet another role in tomorrow's technology infrastructure.

Making terabit optics: What 400G per lane really demands from the photonic engine



As AI infrastructure pushes toward 400G per lane and beyond, the photonics industry faces a critical transition in materials, modulation formats and manufacturing capacity. Thin-film lithium niobate is emerging as a strong contender for next-generation optical engines, but its success will depend as much on foundry readiness and supply-chain maturity as on device performance.

BY FRÉDÉRIC LOIZEAU CO-FOUNDER & CRO, LIGHTIUM

➤ **Figure 1:** AI data center interconnect schematics. Scale-up: intra-rack links below 500 m. Scale-out: rack-to-rack and rack-to-building links up to 2 km. Scale-across: inter-data center connectivity from 2 to 20 km.

The bottleneck driving everything

AS GPU clusters have grown from thousands to hundreds of thousands of accelerators, the optical fabric connecting them has become the limiting variable for continued scaling. More bandwidth per fibre, more reach per lane, more density per rack: interconnects have become the bottleneck of AI infrastructure, and pressure to deliver more comes from every direction.

Three distinct connectivity layers define how AI infrastructure actually moves data (Figure 1). Scale-Up is the tight interconnect between accelerators within a single compute domain, where latency and bandwidth density are essential. Scale-Out is the intra-datacentre fabric, the Spine/Leaf networks connecting racks and buildings, increasingly spanning distances that approach 2 km as datacentre footprints grow. Scale-Across links connect physical facilities on a shared campus, operating over distances of 2 to 20 km. All three layers are expanding in reach and throughput simultaneously, and all three are converging on the same bottleneck: per-lane optical speed.

Optical modules are inherently parallel: Total throughput equals lane speed multiplied by lane count. Doubling the per-lane speed halves the required number of lasers, modulators, and fibres at a given aggregate bandwidth, reducing cost, improving packaging density, and lowering system power. The industry is now transitioning from 200G to 400G-per-lane. That shift will simultaneously reshape modulation schemes, material platforms, and manufacturing capacity.


What won at 200G, and why

Understanding the transition requires treating two dimensions separately: the modulation scheme and the material platform.

On the modulation side, the 200G generation in datacentres is built almost entirely on IMDD (intensity modulation with direct detection), using pulse amplitude modulation (PAM4) at 116 GBaud. It encodes data in optical intensity; the receiver measures optical power directly, yielding a structurally simple, low-cost, power-efficient link. Coherent modulation, on the other hand, simultaneously encodes data in amplitude, phase, and polarization, thereby requiring a local-oscillator laser at the receiver. Spectral efficiency is approximately four times that of IMDD, but cost and power are substantially higher. In the 200G era, coherent is confined to long-haul DCI and telecom. It has no presence within the data centre's short reach.

On the material side, four platforms are relevant (Table 1):

InP electro-absorption modulated lasers (EML) integrate a DFB laser and an electro-absorption modulator on a single die. They are the 200G workhorse: at 116 GBaud, the modulator operates comfortably within its bandwidth envelope, and



► Table 1: Photonic platform capabilities and application fit at 400G per lane

	SiPh MZM	InP EML	InP MZM	TFLN MZM
Modulation format	PAM4 or IQ	PAM4 only	PAM4 or IQ	PAM4 or IQ
EO bandwidth	50–70 GHz	60-100 GHz	80–100 GHz	>110 GHz
Drive voltage	V_{π} 2–5 V	Drive swing < 1 V	V_{π} 1.5–4 V	V_{π} < 2 V
Modulator linearity	Moderate (carrier nonlinearity)	-	Moderate (active bias needed)	High (Pockels)
Athermal operation	No	No	No	Yes
Waveguide propagation loss	~2 dB/cm	-	~1–3 dB/cm	< 0.5 dB/cm
Fits 200G IMDD?	✓	✓	Capable, but overkill	Capable, but ecosystem not ready
Fits 400G IMDD?	BW limited; needs GeSi/heterogeneous	✓	Capable, but supply constrained	✓
Fits 400G CL?	✓ RX side; TX V_{π} too high	No IQ modulation	Capable, but supply constrained	✓
Fits 1600G Coherent?	BW insufficient	No IQ modulation	✓ Incumbent platform	✓

the monolithic laser integration is mastered at high volume, making the device compact and mature.

InP Mach-Zehnder modulators (MZMs), either monolithically integrated with the laser, SOA, and detectors or combined with silicon photonics for the receiver part, enable full coherent modulation, but the cost and integration complexity exceed the 200G IMDD requirements.

Silicon photonics (SiPh) uses CMOS-process waveguides on silicon wafers with an external InP laser source. The availability of germanium detectors makes silicon photonics very attractive for the receiver function, and recent progress on SiGe-based electro-absorption modulators (EAMs) opens a path towards 400G-per-lane IMDD. The 300 mm wafer is the other key manufacturing advantage, while the MZI modulator bandwidth is adequate for 200G. While transceivers were commonly introduced by the more costly InP-based platforms, silicon photonics was used to significantly reduce costs at 200G.

Finally, **TFLN** leverages the exceptionally strong electro-optic Pockels effect; modulation bandwidth exceeds 100 GHz, and drive voltage is sub-1V, but industrialization was not mature enough to catch the 200G deployment window. It also lacks detectors natively, so TFLN is currently confined to the transmitter side of the photonic engine.

In the 200G landscape, InP EML and SiPh capture virtually the entire datacentre short-reach IMDD market. InP modulators serve long-haul coherent while TFLN was not mature enough to enter the market. The dominant platforms won because 200G IMDD made relatively modest material demands: 50

GHz of electro-optic bandwidth, manageable 2-3 V drive voltages, and forgiving link budgets for sub-2 km reach. At 400G, that room disappears entirely.


Why 400G is a fundamentally different problem

The 200G-to-400G transition significantly raises the bar for device bandwidth, link budget, and modulation efficiency. The core physical problem is dispersion. IMDD at 400G PAM4 requires roughly 226 GBaud, approximately double the 200G rate, and chromatic dispersion tolerance scales inversely with the square of the baud rate. Hence, IMDD reach at 400G per lane collapses to under one kilometre at moderate link loss.

The Scale-Across segment (2 to 20 km) that IMDD covered confidently at 200G is now stranded, and coherent modulation must be considered. However, deploying a tunable laser, a dual-polarisation IQ modulator, and a full four-dimensional DSP for a 2-20 km datacentre campus link solves the wrong problem. The cost, power, and complexity overhead of conventional coherent links are unjustified. The correct architecture is Coherent Lite: Single-polarisation IQ modulation at 113 GBaud, the same baud rate as today's 200G IMDD PAM4, with a fixed DFB replacing the tunable laser and a simplified receiver architecture. No breakthrough in device bandwidth required. But IQ modulation is now the requirement, and that single change rewrites the material selection logic entirely.

What 400G will be built on

InP EML is the natural home for 400G IMDD. Demonstrations have progressed rapidly, from the



➤ **Figure 2:** Wafer-scale electro-optical characterization of TFLN dies on a 200 mm wafer. Measuring every die in-situ before dicing enables known-good-die selection and provides the statistical yield data required by a PDK-based design flow.

first 400G EML at OFC 2025 to a 99 GHz electro-optic bandwidth result presented at OFC 2026. The material boundary, however, is fixed by physics. The electro-absorption mechanism modulates only intensity. There is no IQ modulation path from the InP EML; therefore, there is no Coherent or Coherent Lite path.

InP MZM modulators can perform IQ modulation. Monolithic integration of the laser, IQ modulator, and detector on a single InP die provides the functionality Coherent Lite requires. The constraint is supply. InP PICs are significantly larger than EMLs, roughly 100-150x larger. InP wafers are predominantly produced in 3- to 4-inch diameters, with 6-inch being on the horizon. Small wafer and large die: the volume ceiling is hard, and is amplified by geopolitical influences and export restrictions.

The supply constraint deepens when you consider that InP EML and InP PIC draw from the same substrate supply. As 400G IMDD deployment accelerates, EML demand will grow substantially.

This is not a device performance question. It is a wafer economics question, and it creates a structural opening for alternative modulator platforms in the Coherent Lite segment that InP PICs will find difficult to fill at volume.


Silicon photonics faces its own material transition at 400G. The plasma dispersion effect that underpins silicon MZI modulators reaches fundamental bandwidth and drive-voltage limits well before 226 GBaud. For Coherent Lite, silicon photonics is best suited to the receiver side thanks to the germanium photodetectors. On the transmitter side, it requires an external laser and a high-performance IQ modulator made of another material. The silicon photonics community has responded with heterogeneous integration, using GeSi or III-V absorbers on the silicon platform to recover bandwidth that plasma-dispersion silicon cannot reach. IMEC's demonstration at OFC 2026 of a >110 GHz GeSi electro-absorption modulator on a 300 mm silicon photonics platform is a credible result on this path to date. The architectural boundary, however, is the same as for InP EML: a GeSi EAM modulates only intensity. It addresses 400G IMDD on silicon but has no path to IQ modulation and therefore no path to Coherent Lite. SiPh at 400G is evolving into an integration substrate, with the high-performance modulation function delegated either to GeSi absorbers for IMDD or to alternative material platforms for coherent applications.

TFLN is in a different category entirely. The electro-optic Pockels effect in lithium niobate creates a direct linear relationship between applied electric field and refractive index shift, with no charge carriers involved. The response is instantaneous and temperature independent. The result is a set of modulator performance parameters that no other platform simultaneously achieves: greater than 100 GHz electro-optic bandwidth with no roll-off from carrier dynamics; drive voltage below 1V, reducing thermal load and relaxing driver IC specifications; waveguide propagation loss below 0.5 dB/cm, preserving link budget in multi-channel configurations where loss accumulates; and intrinsic linearity from DC to millimetre-wave frequencies, critical for the signal fidelity that 16QAM demands (Figure 2).

These parameters are achieved together, in a single material, without the trade-offs that constrain every alternative. The architecture consequence is direct: a TFLN Mach-Zehnder modulator can function as a single-arm intensity modulator for PAM4 IMDD or be configured as an IQ modulator for SP-16QAM Coherent Lite, using the same waveguide platform and fabrication process. No other modulator material covers both 400G pathways without a fundamental compromise.

Maturing TFLN supply chain enables volume production

TFLN missed the 200G window. Not because of what the material could deliver, but because of what



➤ **Figure 3:** A 200 mm LNOI wafer. The transition from 150 mm to 200 mm provides roughly 2.3x more dies per wafer and opens access to the global installed base of high-volume silicon foundries.


the foundry ecosystem could not. Substrate supply came from a small number of vendors, wafer sizes were limited to 150 mm, and the fabrication toolsets available at that diameter were research-grade: capable of producing excellent individual devices, but unable to deliver the process uniformity, throughput, and statistical characterization that production qualification demands. A module manufacturer qualifying a new component needs wafer-scale yield data, lot-to-lot consistency, and a PDK grounded in manufacturing statistics rather than nominal values from a handful of devices. None of that was available from 100- or 150-mm TFLN research lines at universities or RTOs.

Two structural shifts are changing this: First, 200 mm TFLN wafers are now available and compatible with production-grade semiconductor toolsets (Figure 3). The area advantage matters for economics, but the more important change is process control. Production equipment qualified on 200 mm substrates delivers within-wafer and wafer-to-wafer uniformity that is quantitatively superior to that of 150 mm research toolsets. Electro-optic bandwidth histograms, Vpi distributions, and propagation loss maps measured across a full production wafer provide the statistical foundation a reliable PDK requires and the qualification evidence a procurement team needs before committing volume to a new platform. Second, the smart cut TFLN wafer supply base is broadening. Multiple vendors are now producing wafers at commercially relevant quality levels and at quantities to satisfy the present and future demand of the industry. In photonics supply chains, single-vendor substrate dependency is a qualification risk that no serious customer accepts. A multi-vendor wafer market resolves that exposure and signals a level of ecosystem maturity that 150 mm TFLN never reached.

Co-design is not optional at 100+ GHz

The third shift is structural to the industry: The emergence of pure-play open-access TFLN foundries. At 400G-per-lane and 100 GHz modulation bandwidth, photonic circuit design cannot be decoupled from system architecture. Drive voltage, differential versus single-ended topology, RF electrode geometry, packaging approach, and co-packaging strategy all interact directly with modulator circuit performance. The coupling between the photonic circuit and its electrical and thermal environment is tight enough that a single reference modulator design, however well engineered, will not be optimal for every system architecture that deploys it.

Module manufacturers need to design their own photonic engine, tailored to their specific system context, and own the resulting IP. A foundry model that makes this possible fully separates fabrication from product design: The foundry provides a stable, well-characterized process, a statistically validated PDK, and a reliable manufacturing service. What the



► **Figure 4:** A 100+ GHz assembly-compatible test package, co-developed by PHIX BV and Lightium AG.

customer builds on that foundation belongs entirely to them. That combination of process access, design freedom, and IP ownership is not a commercial differentiator at 400G. It is a prerequisite for extracting the performance that TFLN physics makes available.

Conclusion

TFLN missed the 200G window because the infrastructure was not yet in place. The 400G window is opening under different conditions: System-level performance is validated, 200 mm production capability is coming online, the wafer supply base is maturing across multiple vendors, and pure-play open-access foundry services are available for the first time. The InP supply chain will be pulled strongly toward EML as 400G IMDD demand scales, creating a structural opening in the Coherent Lite segment that InP PIC will find difficult to fill at volume. TFLN is the only modulator platform that serves both 400G pathways from the same material system, without compromise. The question is no longer whether the material works. It is whether the foundry ecosystem scales in time to meet the opportunity window.

FURTHER READING

- X. Zhou, E. Mao, C. F. Lam, and H. Liu, "Next-Gen Datacenter and Campus Optics: Technology Choices beyond 200 Gbps per Lane IM-DD," *J. Lightwave Technol.* 44(3), 845–860 (2026). DOI: 10.1109/JLT.2026.3654031
- C. St-Arnault et al., "Net 1.6 Tbps (4×400 Gbps/λ) O-band IM/DD transmission over 2 km using uncooled DFB lasers and sub-1 V drive TFLN modulators," *OFC 2024*, paper Tu2D.2.
- C. Bruynsteen et al., "110 GHz GeSi Electroabsorption Modulator on a 300 mm SiPh Platform Enabling High-Density 400 G/lane IM/DD Links," *ECOC 2025*.
- K. Luke et al., "Wafer-scale low-loss lithium niobate photonic integrated circuits," *Opt. Express* 28, 24452–24458 (2020).
- C. Wang et al., "Integrated lithium niobate electro-optic modulators operating at CMOS-compatible voltages," *Nature* 562, 101–104 (2018).

Building the fab of the future for high volume PIC manufacturing

Photonic integrated circuits are moving from technical promise to industrial expectation. Customers are no longer asking whether PICs can work; they expect reliable delivery, traceable quality, and repeatable performance at scale. For the photonics industry, this marks a decisive shift: from research excellence to disciplined manufacturing.

BY BART THIESEN, MANAGING CONSULTANT HIGH-TECH, ITILITY AND ANDRE VAN DE GEIJN, MANUFACTURING IT MANAGER, SMART PHOTONICS

FOR YEARS, the photonic integrated circuit (PIC) industry has been driven by research, prototyping, and technical demonstration. Across Europe, companies, universities, and research institutes have proven that PICs enable powerful solutions in telecom, sensing, LiDAR, medical applications, quantum technologies, and datacentre connectivity. This research phase has created the foundation on which the industry stands today.

However, the demand has changed. It is no longer limited to whether a device can be designed and demonstrated. It is now also about whether PICs can be produced reliably, at acceptable cost and in high volume.

This marks the shift toward industrial maturity. Customers no longer accept demonstrations

alone. They expect predictable delivery, controlled processes, full traceability, repeatable quality, and the ability to scale rapidly without losing control of yield or performance. For Tier 1 customers, these are not preferences, they simply demand it.

SMART Photonics is one of the companies navigating this transition. As an independent pure-play InP foundry, it sits at the centre of the lab-to-fab challenge: enabling advanced photonic designs while building the manufacturing maturity required for scalable production. Itility supports this challenge by implementing proven semiconductor principles into a digital and operational backbone.

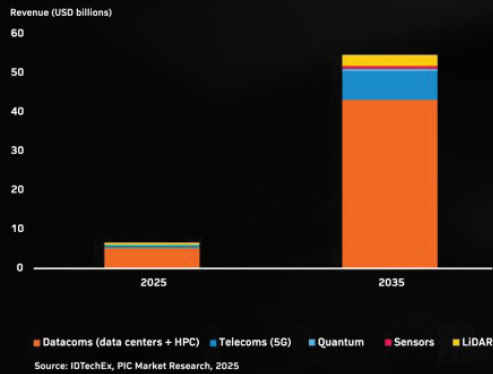
Market forces driving urgency

The urgency to scale is driven by strong market forces. AI workloads are increasing rapidly, while

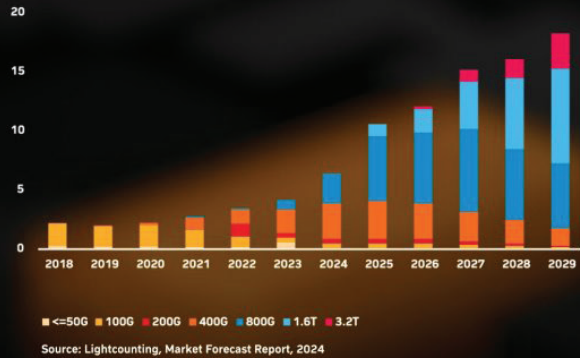
➤ **Figure 1:** A cleanroom operator working alongside advanced manufacturing equipment in a highly controlled production environment.



Photonic Integrated Circuit Market grows massively over the next decade



Need for faster data transmission drives this development



hyperscale datacentres require faster, denser and more energy-efficient data transfer. As copper-based systems approach their limits, PICs offer a compelling alternative toward higher-bandwidth, lower-power interconnects.

Industry expectations point to significant growth, with the PIC market projected to exceed 50 billion USD in the coming decade. Capturing this opportunity requires more than technology alone: it requires a manufacturing model that can deliver PICs with the reliability, repeatability, and discipline seen in high-volume semiconductor environments.

From research mindset to industrial reality

Scaling PIC manufacturing requires a shift in mindset. Many photonics organisations originate from research environments, where flexibility, speed and experimentation are essential. While this heritage is a strength, it becomes a constraint when moving to volume production. For example, in early-stage R&D manual workflows are common, data can be fragmented across tools and teams, tools operate in isolation, and processes depend on individual expertise. Critical knowledge often resides on laptops or in the minds of experienced researchers. This works for rapid experimentation but does not provide the right backbone required for high-volume manufacturing.

Transitioning to industrial manufacturing is therefore not a simple step. It requires time, investment, change management and a structured approach to standardisation, data, and process control. With demand increasing and customer expectations rising, such transformation must start early to ensure timely readiness to scale.


The focus shifts from identifying challenges to building the capabilities required to overcome them. At the core is a strong data foundation: data must be complete, consistent, and accessible across the organisation, enabling a single, trusted

view of operations and faster decision-making. Standardised workflows and well-defined processes are equally critical. Alignment on internal practices and use of industry standards ensures repeatability, comparability, and scalability, forming the basis for consistent manufacturing performance.

Design, manufacturing and testing systems should function as a connected toolchain that provides end-to-end visibility and enables insight into process and equipment behaviour. A further key step is the transition from individual-driven expertise to structured, organisation-wide knowledge. Design methodologies, process know-how and best practices must be captured and shared, reducing dependency on individuals. Together, these elements mark the shift from flexible R&D to disciplined manufacturing.

SMART Photonics is an independent pure-play Indium Phosphide (InP) foundry based in Eindhoven, the Netherlands. Founded in 2012 as a spin-off from TU Eindhoven and Philips, the company was created to make advanced Indium Phosphide photonic integration technology available through a foundry model. Smart's InP platform is especially relevant because it enables active optical functions, such as lasers and amplifiers, to be integrated with other photonic components on the same chip. Over

➤ **Figure 2:** Projected growth of the Photonic Integrated Circuit market, driven by rising demand for faster data transmission.



➤ **Figure 3:** Wafer-level image of an InP Photonic Integrated Circuit, combining active optical components, waveguide routing, and metal interconnects.

the years, SMART has grown from an R&D-driven photonics company into an important player in Europe's integrated photonics ecosystem, serving applications in data and telecom, sensing, LiDAR, and medical technologies. Its current journey reflects the wider challenge facing the industry: moving from innovation and early production toward mature, reliable, and scalable high-volume PIC manufacturing.

For a foundry as SMART Photonics, scaling is not only about producing more wafers. It means strengthening the full operating model across design enablement, manufacturing processes, equipment data, product qualification, testing, quality control, and customer delivery. SMART's transformation is therefore a practical example of the shift needed in the photonics industry.

From ambition to execution: Building maturity in practice

SMART Photonics has the ambition to operate as a Tier 1 photonics foundry. In practice, this means meeting customer expectations on predictable delivery, quality, traceability, and scalability. The semiconductor industry offers a valuable reference model. Decades of high-volume chip manufacturing have created mature practices in design, automation, process control, equipment integration, quality systems, traceability, and yield learning. However, these practices cannot simply be copied into photonics. PIC manufacturing has its own materials, process sensitivities, and test requirements.



The challenge therefore lies in translating semiconductor best practices into photonics context. This translation must be practical and incremental. It starts with controlled workflows,

Design methodologies, process know-how and best practices must be captured and shared, reducing dependency on individuals. Together, these elements mark the shift from flexible R&D to disciplined manufacturing

reliable data capture, and standardised processes, and then evolves toward integrated systems, automated execution and closed feedback loops across design, manufacturing, and test.

This is where Itility supports SMART Photonics. Itility brings 20+ years of experience in high-tech and semiconductor manufacturing environments, combining software engineering, data platforms, automation, cloud infrastructure, and secure engineering workflows. Together, SMART Photonics and Itility translate these foundations into an executable maturity roadmap for high-volume PIC manufacturing. Rather than a one-time transformation, this roadmap describes a step-by-step evolution moving fully away from manual, person-dependent ways of working toward standardised processes, automated equipment, and real-time data availability.

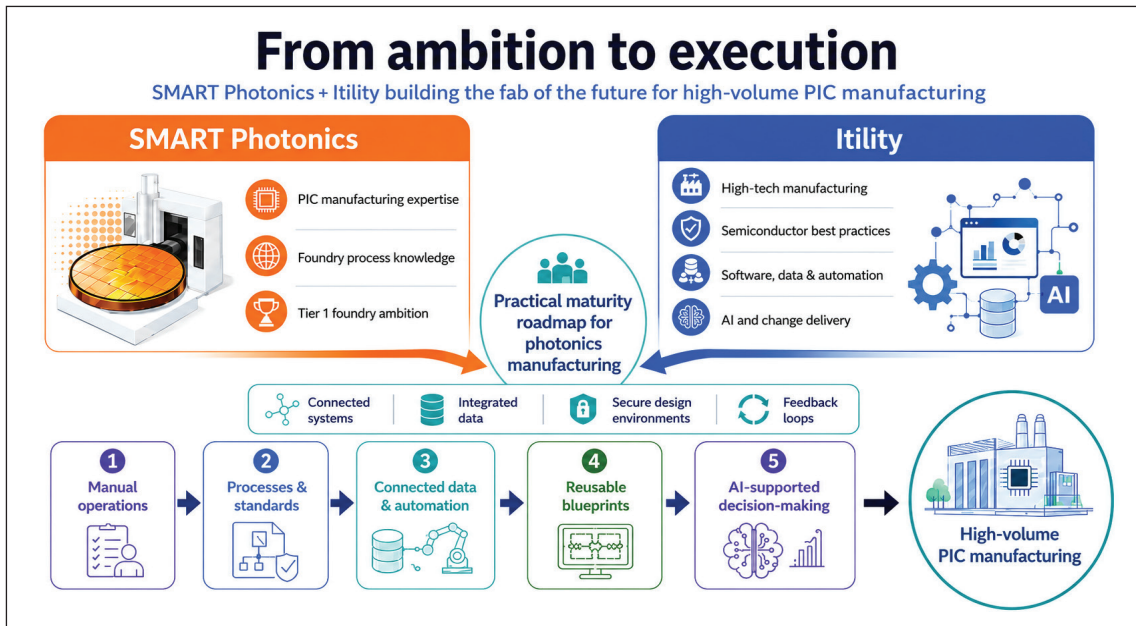
Itility is a digital engineering consulting company headquartered in Eindhoven in the Netherlands, with a strong presence in high-tech manufacturing and the semiconductor industry. The company supports customers translate technology into practical business and operational value by building customer-specific solutions using software, automation, data, cloud, and AI. Itility is renowned for working closely at customers to truly understand their processes and develop solutions that improve how teams operate, make decisions, and achieve higher maturity levels following industry best practices.



As maturity increases, data becomes structured and integrated across the organisation. This results in reusable blueprints, consistent execution and advanced capabilities as automated decision-making and AI-assisted engineering. The goal is not automation for its own sake, but controlled scaling: increasing output while improving predictability, quality, and speed of learning. In a mature PIC foundry, design, manufacturing, and testing are tightly interconnected. Design defines what is built, manufacturing determines how it is produced, and testing provides the feedback needed to continuously improve both. In the following sections, we will explore how each of these three core processes enable scalable, data-driven photonics manufacturing.

Core process 1: Design as a foundry discipline (EDA as backbone)

In high-volume PIC manufacturing, design is not an isolated creative activity. It is a foundry discipline that extends from PDK governance and customer



► **Figure 4:** From ambition to execution: combining SMART Photonics' PIC manufacturing expertise with Itility's digital engineering capabilities to build a maturity roadmap for high-volume PIC manufacturing.

design verification, to controlled tape-out and feedback from manufacturing and test.

A PIC foundry does more than manufacture customer layouts. It develops, maintains, and qualifies its own building blocks that customers use to create functional circuits. These building blocks are brought together in a Process Design Kit (PDK), which includes libraries, design rules, technology data, layout constraints, and simulation models. The PDK enables customers to develop manufacturable designs that align with the foundry process. Its maturity directly affects first-time-right design, manufacturability, yield learning, and customer confidence.

For this reason, the PDK must be treated as a controlled asset. Changes to building blocks, design rules or models need to be governed, versioned and traceable. Customers and internal engineering teams must work with consistent libraries and qualified workflows. Without this discipline, design freedom can easily translate into process risk, especially as volumes increase and more customers depend on the same foundry platform.

Tape-out is the critical handover point between design intent and manufacturing execution. In a mature foundry environment, tape-out cannot be treated as an administrative step toward mask creation but is a controlled release process.


To ensure quality and scalability, this entire process must be supported by a semicon industry best practice: an Electronic Design Automation (EDA) environment, which acts as the backbone of all design activities. First, it provides a secure collaboration environment for working with sensitive data, including customer IP and the foundry's own building blocks. Second, it creates a consistent and

managed workspace for engineers, where tools, libraries and workflows are made available in a controlled way. This allows engineers to focus on design quality rather than design tool configuration or version management. Third, it acts as a platform for structured verification, simulation, and automation.

The real value emerges when this design environment is connected to manufacturing and test data. Insights from the fab and test floor can be fed back into design rules, building blocks and simulation models. Over time, this creates the foundation for high-yield manufacturing by using advanced capabilities in design optimisation, automated verification, and AI-supported design workflows. This is what data-driven steering means in practice: not adding AI as a standalone feature, but first building the secure, connected, and automated design foundation that allows AI to create real value.


Core process 2: Manufacturing powered by data

Reliable, high-volume PIC production requires manufacturing to be supported by trusted operational data. In a mature fab environment,



► **Figure 5:** A close-up view of an integrated circuit, where compact optical components are brought together on a single chip.

► **Figure 6:** A cleanroom environment with operators and equipment used for high-precision manufacturing and process control at Smart Photonics



reliance on manual tool settings, individual expertise or fragmented reporting is no longer appropriate. Engineers and operators need access to the right information at the right time to understand process behaviour, equipment performance and execute actions in a standardized order on the shop floor.

In practice, this means actively supporting operators through intuitive interfaces, uniform instructions, and automated guidance, while ensuring that equipment recipes and configurations are automatically selected and consistently enforced for each wafer and process step. This reduces variability and human error, enabling repeatable execution across the fab. To reach this level of control, semiconductor best practices play a key role, requiring clear visibility into equipment behaviour, process performance and production flow, combined with the ability to translate this insight into effective actions, including taking the right corrective measures when deviations occur. Here an MES layer is essential to manage and optimise execution on the shop floor, delivering end-to-end traceability and data-driven process control. By adhering to the

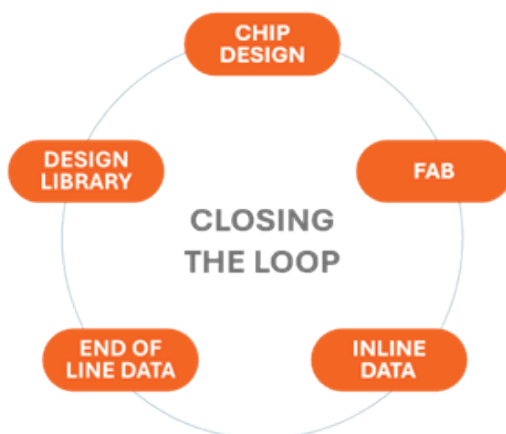
ISA-95 framework, it establishes a clear interface between business systems and manufacturing operations, enabling standardisation and future scaling.

With this foundation in place, the focus shifts from data collection to insight and control. Statistical Process Control (SPC) enables early detection of process variation and prevents drift before it impacts product quality. Overall Equipment Effectiveness (OEE) complements this by providing insight into availability, performance, and utilisation, allowing teams to track productivity and steer improvements through clearly defined KPIs.

To make this effective, manufacturing must be built on an end-to-end data integration strategy. The objective is not to build isolated dashboards or point solutions, but to create a structured data foundation that connects equipment, process, wafer, product, and quality information. This enables engineers to analyse process behaviour, identify bottlenecks and act on deviations with greater speed and confidence.

A practical way to implement this is through reusable blueprints. Instead of building isolated solutions for each tool or process, standardised templates can be applied across similar equipment and setups. These blueprints define interfaces, integrations, and data flows, ensuring consistent execution and enabling structured monitoring across the fab.

As manufacturing maturity increases, the fab evolves from retrospective reporting to proactive control. Insights are no longer used only after the fact; they are fed back into operations to support faster intervention, more stable processes, and more predictable output. This is the route toward scalable PIC manufacturing: not more data for its own sake, but data that directly improves control, quality, and throughput.



► **Figure 7:** Data driven feedback loop from chip design to fabrication.

Core process 3: Testing as the scaling engine

In high-volume PIC manufacturing, testing cannot be treated as the final checkpoint at the end of the process. It is a scaling engine. Testing is where design intent, process behaviour and product performance come together, and where the foundry learns whether its platform can deliver repeatable results at volume. In this way, it actively steers the full process flow. It supports yield improvement and enhances equipment productivity and enables a shift from experience-based decisions to consistent, data-driven control.

The first step toward higher maturity in testing is automation of test execution and data processing. Manual validation, inconsistent test setups, and fragmented result files create delays and uncertainty. Automated test workflows, calibrated measurement routines, and built-in data verification allow engineers to trust the data and act immediately. This reduces the need for retesting, shortens learning cycles, and improves the efficiency of engineering and production teams.

As test data quality and trust improve, testing becomes a feedback mechanism for both manufacturing and design. Measurement results can reveal process drift, layout sensitivities, building block behaviour, and product-specific yield patterns. These insights can then be fed back into process control, PDK updates, design rules, and simulation models. This changes the role of testing fundamentally. Instead of confirming quality after production, testing actively steers yield improvement and platform maturity.

Over time, better test automation and stronger data correlation reduce the overall test burden. Test times are shortened, retesting becomes less frequent, and discussions on data validity also diminish. This is not only an efficiency gain, but essential for meeting customer expectations in speed, reliability, and traceability.

In a mature PIC foundry, testing closes the loop across design, manufacturing, and product qualification, enabling faster optimisation cycles and creating the confidence needed for scalable high-volume production.

The hard truth: you cannot scale photonics in isolation

Even if a foundry executes perfectly, scaling PIC manufacturing will inevitably expose constraints elsewhere in the value chain. High-volume photonics is not only a fab challenge. It is a system-level industrialization challenge.

Upstream, material availability and supplier maturity can become limiting factors, particularly for specialised substrates (such as InP wafers). As

volumes continue to grow, bottlenecks typically shift downstream into testing, packaging, and assembly. These are often device-by-device processes, making them similarly difficult to automate and scale.

The challenge becomes even more critical with heterogeneous integration between photonics and electronics, such as in co-packaged optics. In these cases, performance and reliability must be verified within the fab flow before permanent integration since failures after integration are significantly more expensive and difficult to resolve.

Scaling photonics therefore requires the full ecosystem to mature together. Foundries, design houses, EDA providers, equipment suppliers, material suppliers, packaging partners, OSAT providers, system integrators and investors all play a role. If one part of the chain matures while others lag, bottlenecks will not disappear; they will simply move.

One potential consequence is that the industry may move towards more vertically integrated models, in contrast to the mature horizontal semiconductor ecosystem optimised for massive volumes, speed and innovation.

For Europe, this is a critical opportunity and a critical test. The region has strong photonics research, leading equipment expertise, specialised foundries, and an active ecosystem around integrated photonics. But industrial leadership will depend on more than technical excellence. It will require alignment on standards, interfaces, qualification methods, data exchange, and investment in scalable manufacturing infrastructure.

The concept of the “fab of the future” is therefore not a single breakthrough, but a disciplined and structured journey. Those who start this journey early and do so together with the right partners and support across the value chain, will be best positioned to capture the full potential of photonic integrated circuits at scale.

Credit: Itility / SMART Photonics



➤ Authors from left to right: Bart Thiesen and Andre van de Geijn

Precision single-wafer HBr wet etch optimizes compound semiconductor optoelectronics



A precision single-wafer HBr wet etch process developed by JST is delivering plasma-comparable uniformity for compound semiconductor fabrication, offering a cost-effective and scalable approach for next-generation optoelectronic and photonic devices.

BY ISMAIL KASHKOUSH, PHD, CHIEF TECHNOLOGY OFFICER, JST

THE RAPID expansion of optoelectronic technologies is reshaping semiconductor manufacturing requirements. Devices such as lasers, LEDs, and photodiodes are increasingly central to high-speed data transmission, sensing, and communications systems. As performance expectations rise, so does the demand for tighter process control in etching, even minor deviations can significantly impact optical efficiency, wavelength precision, and overall device yield.

For compound semiconductor materials, achieving high etch-rate uniformity, minimal surface damage, and strong wafer-to-wafer repeatability is especially challenging. Traditional approaches, including plasma etching, have long been favored for their precision. However, growing cost pressures, sustainability concerns, and the need for scalable high-volume manufacturing are driving renewed interest in advanced wet processing techniques.

Among these, hydrogen bromide (HBr)-based wet etching implemented on precision single-wafer platforms is emerging as a compelling alternative delivering plasma-comparable uniformity with improved cost efficiency and reduced process complexity (see Table 1).

Rising optoelectronics manufacturing demands

Optoelectronic devices rely heavily on compound semiconductors such as indium phosphide (InP), which offer superior electrical and optical properties

compared with silicon. These materials enable high-speed operation, low noise, and excellent thermal stability, making them ideal for advanced photonic applications.

However, these same advantages introduce fabrication challenges. Device performance is highly sensitive to structural variations, meaning that etch uniformity across the wafer must be tightly controlled. Non-uniform etching can lead to variations in layer thickness, feature geometry, and ultimately, device characteristics reducing yield and increasing manufacturing costs.

As shown in Figure 1, device architectures evolve and tolerances tighten, manufacturers require processes capable of delivering:

- Sub-2% within-wafer etch uniformity
- Consistent wafer-to-wafer repeatability
- Minimal surface damage and defects
- Flexibility across a range of etch depths and structures

Meeting these requirements consistently in production environments has proven difficult with conventional batch or plasma-based approaches alone.

A competitive wet-etch alternative

Plasma etching remains a cornerstone of semiconductor fabrication due to its directional control and ability to produce anisotropic profiles. However, it comes with inherent complexity.

Aspect	Dry Etch	Wet Process Etch
Cost of Ownership	Higher, More complex maintenance	Lower, Simpler maintenance
Flexibility	Better flexibility with certain substrates	Chemically limited
Speed	Slower process, slower set up	Faster process, faster setup

► Table 1: Plasma Etch vs. Wet Process Etch

Plasma systems require vacuum environments, precise power control, and sophisticated hardware to maintain process stability. These factors contribute to higher capital costs, increased energy consumption, and operational overhead.

In contrast, wet chemical etching offers a fundamentally different approach. By relying on purely chemical reactions rather than ion bombardment, wet etching enables:

- Damage-free material removal
- Lower energy consumption
- Simplified system architecture
- Reduced consumables and cost of ownership

Historically, wet etching has been associated with isotropic profiles and limited control, which has restricted its use in advanced applications. However, recent advances in single-wafer wet processing have significantly improved precision and uniformity bringing wet techniques into direct competition with plasma processes for certain applications.

HBr chemistry for high-uniformity etching

Hydrogen bromide has emerged as a particularly effective etchant for compound semiconductors, especially InP. Its chemical properties enable high selectivity, allowing precise material removal while preserving underlying layers and critical device features.

Key advantages of HBr-based wet etching include:

- High etch uniformity across the wafer
- Excellent selectivity for compound semiconductor materials
- Reduced risk of plasma-induced damage
- Faster etch rates suitable for high-volume manufacturing


When properly controlled, HBr chemistry can deliver uniformity levels comparable with plasma etching while maintaining the inherent benefits of wet processing.

A critical factor in achieving this performance is the method of chemical delivery. Uniform distribution of the etchant across the wafer surface ensures consistent reaction rates from center to edge, minimizing variations in etch depth and feature geometry. The combination of InP's material advantages and HBr's selective, uniform, and low-damage etching makes this chemistry highly suitable for fabricating advanced device structures.

JST Manufacturing has developed a proprietary HBr etch process engineered to maximize this uniformity by combining chemistry optimization with precise fluid delivery control, enabling highly repeatable results in production environments.

Precision single-wafer processing

The transition from batch to single-wafer processing has been instrumental in advancing wet etching capabilities. The main advantages of single-wafer



➤ **Figure 1:** Etch rate uniformity is critical for device performance and reliability. The wafer map provides a visual representation of etch rate distribution across the wafer, targeting < 2% etch rate uniformity.

processing are its superior process control, lower risk of particle redeposition, and highly consistent cross-wafer and wafer-to-wafer outcomes resulting in more reliable semiconductor devices with higher production yields. In this approach, each wafer is processed individually using automated chemical spray techniques. Fan nozzle designs ensure uniform distribution of the HBr solution, while real-time process control maintains consistent conditions throughout the etch cycle.

This level of control is essential as manufacturing tolerances continue to shrink. Single-wafer processing also provides greater flexibility, allowing systems to accommodate varying device requirements and etch depths without compromising performance.

JST's Ospray single-wafer wet processing system was designed to implement this HBr chemistry at production scale, using an automated chemical spray approach that ensures uniform distribution across the wafer surface. The system's integrated hardware and process control capabilities enable tight within-wafer and wafer-to-wafer uniformity, which is critical for advanced optoelectronic device fabrication. Table 2 illustrates the success of JST's HBr wet process using Ospray system, with each test achieving the desired <2% uniformity.

Demonstrated performance and uniformity

Production-scale implementation of HBr-based single-wafer etching has demonstrated impressive results. Uniformity levels as low as 1.24% have been achieved, with consistent wafer-to-wafer repeatability in the 1%-2% range across a broad range of etch depths—from submicron to greater than 4 μm.

These results are derived from testing JST's proprietary HBr process on the Ospray platform, where aggregated data from multiple runs confirmed stable performance under varying process conditions and etch depth requirements. The results represent a significant improvement over conventional wet etching systems, which typically achieve uniformity closer to 5%.

The reduction in variability directly translates to improved device consistency and higher manufacturing yields. Detailed process evaluations show:

- Tight control of etch depth across the wafer, with minimal deviation from center to edge
- Stable repeatability across multiple process runs
- Consistent performance across varying etch depths and device structures

For example, measured uniformity values across multiple trials remained within approximately 1.18% to 2.13%, confirming both the accuracy and robustness of the process. In addition, etch depth control was maintained across a range of approximately 0.55 μm to over 4 μm, demonstrating the Ospray system’s ability to adapt to different device requirements without sacrificing uniformity or process stability—a key requirement for evolving optoelectronic architectures. Even nanometer-level deviations can impact device performance in optoelectronics, underscoring the importance of maintaining such tight control.

Improved etch uniformity has a direct and measurable impact on manufacturing economics. By reducing variability, manufacturers can increase device yield, minimize scrap and rework, improve process predictability, and reduce overall cost per wafer.

In high-volume production environments, even small gains in uniformity can translate into substantial cost savings. Additionally, the lower energy requirements and reduced system complexity of wet processing contribute to a lower total cost of ownership compared with plasma-based alternatives.

Surface integrity and device performance

One of the most important advantages of wet etching—particularly for optoelectronics—is its ability to preserve surface integrity. Because the process does not rely on ion bombardment, it avoids the physical damage and defect generation often associated with plasma etching.

This is especially critical for optical devices, where surface defects can degrade performance by introducing scattering, absorption, or recombination losses. By minimizing damage, HBr wet etching helps maintain the intrinsic material properties required for high-efficiency optical operation.

As optoelectronic device designs continue to evolve, manufacturing processes must adapt accordingly. Precision single-wafer wet processing offers the flexibility needed to support a wide range of applications and requirements.

Key capabilities include:

- Adjustable etch depths without loss of uniformity
- Compatibility with different material systems
- Configurable process parameters for diverse device structures
- Scalability to meet increasing throughput demands

This adaptability ensures that the technology remains relevant as new device architectures emerge, including more complex multilayer structures and integrated photonic systems.

Enabling the next generation of optoelectronics

The combination of HBr chemistry and precision single-wafer processing represents a significant advancement in compound semiconductor manufacturing. By delivering plasma-level uniformity with the inherent advantages of wet processing, this approach addresses many of the key challenges facing the industry today.

For manufacturers, the benefits are clear:

- High uniformity and repeatability for consistent device performance
- Reduced defects and improved surface quality
- Lower cost of ownership and improved sustainability
- Flexibility to support future device innovations

As demand for optoelectronic devices continues to grow—driven by applications in data communications, sensing, and emerging photonic technologies—advanced etching solutions will play a critical role in enabling scalable, high-yield production.

Precision HBr wet etching on a single-wafer platform such as JST’s Ospray system demonstrates that precision single-wafer wet etching can meet these demands in a production environment, providing a robust, high-uniformity solution for advanced compound semiconductor manufacturing.

► **Table 2:**
HBr Wet Process Etch Rate Uniformity Using JST Ospray

	Test 1.1	Test 2.1	Test 2.2	Test 3.1	Test 3.2	Test 3.3	Test 4.1	Test 5.1	Test 5.2	Test 5.3	Test 6.1	Test 7.1
Mean delta (μm)	0.8826	1.971	1.860	3.695	3.359	4.107	0.5568	0.5917	0.5724	0.5473	2.365	0.9726
Max - Min (μm)	0.0218	0.0649	0.0479	0.1576	0.1248	0.1549	0.0191	0.0213	0.0178	0.0169	0.028	0.0385
Uniformity (%)	1.24%	1.65%	1.29%	2.13%	1.86%	1.94%	1.72%	1.80%	1.55%	1.54%	1.18%	1.98%



ANGELTECH

INTERNATIONAL CONFERENCE

12-14 APRIL 2027

Sheraton Brussels Airport Hotel Belgium

THEMES ANNOUNCED

INDUSTRIALISATION AND GLOBAL SCALING OF PHOTONICS

How can photonics transition from innovation-led development to high-volume, globally competitive manufacturing?

BRIDGING THE GAP BETWEEN INNOVATION AND ADOPTION

What is preventing photonic technologies from achieving widespread commercial deployment across industries?

PHOTONICS AS AN ENABLING PLATFORM ACROSS INDUSTRIES

How can photonics deliver measurable value across diverse sectors and become a foundational technology?

EVOLUTION OF THE PHOTONICS VALUE CHAIN AND ECOSYSTEM

How will partnerships, regional strategies, and supply chain shifts reshape the photonics industry?

DEFINING STANDARDS, RELIABILITY, AND COMMERCIAL READINESS

What is needed to ensure consistent performance, interoperability, and trust in photonic technologies?

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

<https://picinternational.net>







POWERING AI CONNECTIVITY

Key features of our Optical Interposer-based technology:

- Lower cost, silicon-based hybrid-integration platform
- Enabling next-generation interconnect solutions for 1.6T and above data rates
- Wafer-level chip-scale packaging (no wire bonds and no active alignments)




Our revolutionary products include:



POET Blazar

A high-power, multi-channel External Light Source. Designed to power both AI Scale-up and Scale-out Networks, including co-packaged optics.



POET Teralight

Tx and Rx optical engines for 1.6T and beyond pluggable transceivers.

Our Partners



FOXCONN

LITEON[®]

Adtran



poet-technologies.com

sales@poet.tech

[#POETpowersAI](https://twitter.com/POETpowersAI)



@POETtech



@POETTechnologies



@POETtechnologies