PIC PHOTONIC INTEGRATED CIRCUITS
CONNECTING THE PHOTONIC INTEGRATED CIRCUITS COMMUNITY

ISSUE IV 2025

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

PICMAGAZINE.NET

TAMING REFLECTIONS:

GE SHAPE ENGINEERING UNLOCKS ULTRA-LOW BACK REFLECTION

EVG®

invent
innovate
implement

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

VIEWPOINT

BY SARAB CHOPRA
EDITOR

Scaling up PIC innovation

AS THE new Editor of *PIC Magazine*, I am pleased to introduce myself at a moment when Photonic Integrated Circuits are advancing with exceptional speed and maturity.

Stepping into this role has given me a fresh appreciation of how far the PIC community has come in refining circuit architectures, strengthening design processes, and pushing the limits of what integrated platforms can reliably deliver.

This issue comes at a time when the sector is moving forward with renewed clarity and purpose. Research results are transitioning into deployable PIC-based products at an accelerated pace, commercial timelines are tightening, and global demand for compact, high-performance integrated circuits continues to grow. Entering the field today means joining a community firmly focused on scale, efficiency, interoperability, and long-term PIC development strategies.

This edition places strong emphasis on the transition from demonstration-level PICs to full industrial readiness. Several features explore how the sector is preparing for high-volume PIC production and what this shift demands from foundries, circuit designers, process engineers, and packaging specialists. The challenges are substantial and increasingly interconnected. They include improving platform uniformity, strengthening design-for-manufacture principles, enhancing wafer-level testing, and ensuring closer alignment between PIC architectures and advanced packaging. Across the ecosystem, however, the message is one of confidence.

Contributors highlight growing investment in automation, tighter process control, improved simulation workflows, and deeper collaboration across the supply chain—efforts aimed at meeting commercial expectations for performance, yield, and scalability. We also examine how regional strategies are influencing the global PIC landscape. Europe's position is a recurring focus, especially as the region works to consolidate its PIC value chain, expand access to open-foundry services, and build a more resilient environment for PIC design, fabrication, and packaging. Policy coordination, shared infrastructure, and strategic funding programmes

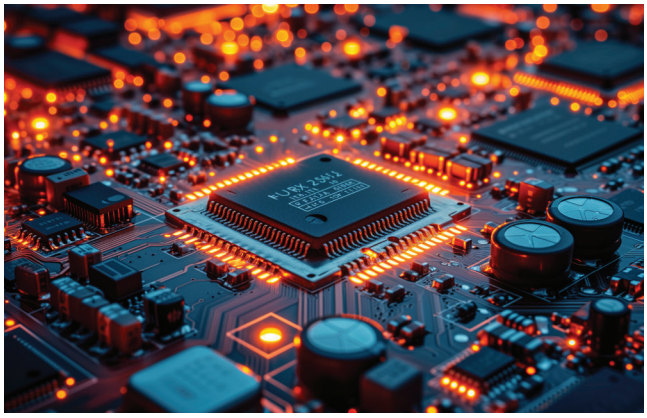

continue to play a crucial role in ensuring that PIC innovation is matched by the capacity to manufacture and scale within the region.

Another central theme in this issue is the increasing importance of ecosystem-level coordination. Progress in hybrid and heterogeneous integration, advances across established material platforms such as InP, SiN, and thin-film lithium niobate, improvements in PIC test strategies, and the introduction of new packaging approaches all depend on how effectively the PIC ecosystem adapts to rapid technological change. Features in this edition underline how essential co-design, interoperability, transparent communication, and evolving standards have become for the next stage of PIC maturity.

As I begin my role as Editor, I am inspired by the discipline, momentum, and ambition that define this field. Working on this issue has strengthened my enthusiasm for what PIC technologies can achieve and the new directions they are poised to unlock

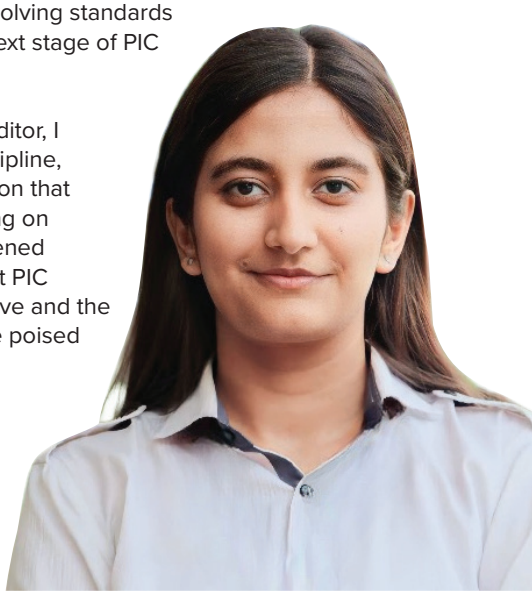

PIC PHOTONIC INTEGRATED CIRCUITS

COVER STORY

Ge Shape Engineering Unlocks Ultra-Low Back Reflection in High-Speed Ge-on-Si Photodetectors

Shape-engineered Ge-on-Si photodetectors deliver -36 dB optical return loss, ~60 GHz bandwidth, 0.95 A/W responsivity, and <30 nA dark current

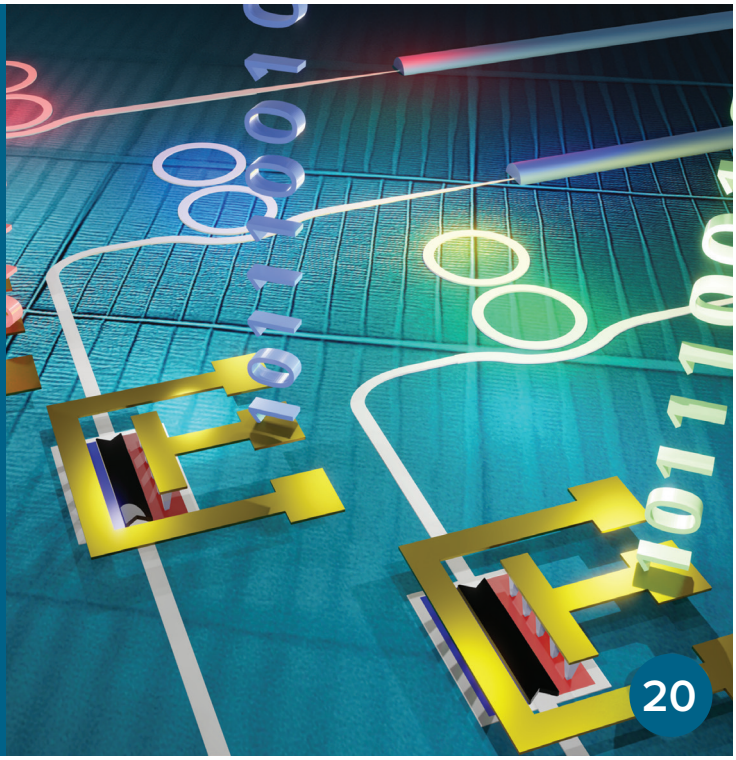

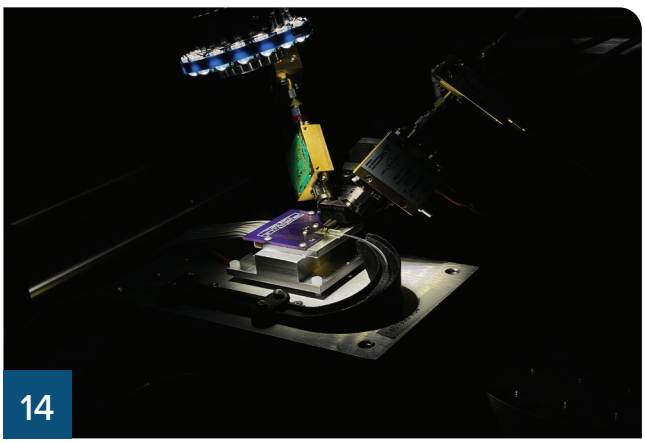

NEWS

PIC PHOTONIC INTEGRATED CIRCUITS

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

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

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

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

Angel BUSINESS COMMUNICATIONS

PIC Magazine is published four times a year on a controlled circulation basis. Non-qualifying individuals can subscribe at: £105.00/e158 pa (UK & Europe), £138.00 pa (air mail), $198 pa (USA). Cover price £4.50. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher. Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd © Copyright 2025. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing.US mailing information: Pic Magazine, ISSN 1096-598X, is published 4 times a year, March, May, August and December by Angel Business Communications Ltd, Unit 6, Bow Court, Fletchworth Gate, Burnsall Rd, Coventry CV5 6SP. UK. The 2025 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 post-age 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 2025. 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 objec-tives 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.

Photonics start-ups driving Europe's innovation surge

Europe's photonics start-up scene is booming, with 240 new companies identified since 2024, bringing the total to 556 active firms, according to new data from Photonics21.

THE European Technology Platform's updated database highlights start-ups developing technologies that tackle major societal challenges, from life-saving medical diagnostics and surgical tools to quantum communication chips, precision climate and energy sensors, and lasers powering advanced manufacturing.

Healthcare remains the leading sector, with one in five start-ups focused on medical applications such as cancer diagnostics, surgical imaging, and personalised medicine.

Other key sectors include components and advanced materials (16 percent) and Industry 4.0 (14 percent), where photonics supports automation, robotics, and digital factories. Berlin-based start-up DiaMonTech exemplifies the impact of these innovations. Founded in 2015, the company has developed a non-invasive glucose monitoring device using mid-infrared light, offering a pain-free alternative for millions of diabetes patients across Europe.

Catalina Plesmann, Head of Marketing, said, "Our technology uses light instead of needles to make glucose monitoring completely pain-free. Europe is where our science was born and where we want to scale."

Dr Lutz Aschke, President of Photonics21, said the growth reflects a dynamic and young ecosystem, with 35 percent of firms founded since 2020 and 84 percent under ten years old. He added that connecting founders with investors remains crucial to help these companies scale within Europe rather than overseas.

Photonics21 supports this through initiatives like Tech Tour Photonics, where EU-based entrepreneurs pitch directly to international investors. This year, 77 companies were evaluated, and 39 were selected to present, highlighting the pipeline of investment-ready photonics ventures emerging across Europe.

Europe's photonics industry, currently valued at €124.6 billion, employs over 430,000 people and is projected to reach €175 billion by 2027, reinforcing the continent's position as a global leader in photonics innovation across AI, quantum, health, green energy, and secure communications.

MicroAlign's ultra-precise fibre arrays boost PIC integration

Eindhoven-based startup MicroAlign is making waves in the photonic integrated circuit (PIC) sector with its ultra-precisely aligned fibre arrays, tackling one of the biggest bottlenecks in chip-to-fibre integration.

THE COMPANY has developed ready-to-integrate fibre arrays assemblies of optical fibres encapsulated in glass blocks capable of tenfold higher alignment accuracy and 75 times lower fibre-to-chip coupling loss than traditional methods. These advancements are crucial for PIC developers seeking high-performance, scalable optical networks. "Every photon counts in photonic quantum computing, and these fibre arrays set a new benchmark for connectivity efficiency," says MicroAlign CEO Simone Cardarelli.

Founded in 2021 by Cardarelli and CTO Marco Fattori, the startup emerged from PhD research at TU Eindhoven. Embedded in the PhotonDelta ecosystem, MicroAlign has leveraged shared expertise and local networks to refine its technology and target high-end photonic applications. While the company initially pivoted toward quantum computing, its fibre array solutions are broadly applicable to telecom, sensing, and other PIC-driven industries.

This year, MicroAlign delivered 100 fibre arrays to paying customers and is preparing to scale production with a 500-square-meter facility slated to open in 2026. The plant aims to produce up to 100,000 fibre arrays annually, supporting industry demand for high-precision, low-loss integration.

Cardarelli highlights the technical challenges behind their approach: "Fibre behaves chaotically when handled, so automating this delicate process requires precision engineering and deep expertise." Looking ahead, MicroAlign aims to become the gold standard for PIC fibre connectivity, potentially supporting the next generation of high-performance photonic systems.

CITC joins TNO to strengthen chip packaging in Nijmegen

THE Chip Integration Technology Centre (CITC) will become part of TNO, marking a major step in advancing chip packaging research and innovation. The integration, effective January 1, 2026, aims to secure CITC's long-term continuity and strengthen the Netherlands' semiconductor ecosystem, connecting Nijmegen with other hubs such as Eindhoven.

CITC, founded in 2019 by TNO and TU Delft, has developed into a broad R&D hub for advanced chip packaging technologies. By joining TNO's High Tech Industry unit, CITC will expand research efforts, accelerate innovation, and maintain its presence at the Noviotech Campus in Nijmegen.
Arnaud de Jong, TNO's director

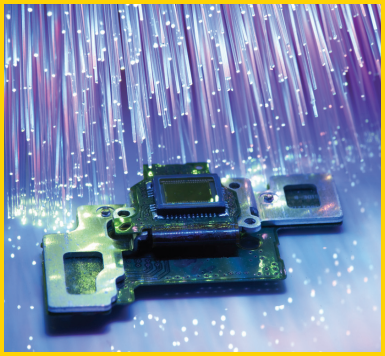

of High-Tech Industry, said, "Integrating CITC allows us to intensify our work in advanced chip packaging and contribute to a strong, future-proof semiconductor ecosystem in Nijmegen." Jeroen van den Brand, CITC general manager, added that the move strengthens CITC's international position as a centre of expertise.

TU Delft remains closely involved, ensuring that the collaboration continues to support the growth of chip integration and packaging technologies across the region.

GUC and Ayar Labs partner on co-packaged optics

Global Unichip Corp. (GUC), a leading provider of advanced ASIC design services, has announced a strategic partnership with Ayar Labs to integrate co-packaged optics (CPO) into next-generation compute architectures aimed at AI, HPC and hyperscale data centre applications.

THE collaboration combines GUC's advanced packaging and ASIC design expertise with Ayar Labs' TeraPHY™ optical engines, positioning the two companies to address the growing limitations of electrical interconnects as AI model sizes and bandwidth requirements continue to surge.

According to GUC CTO Igor Elkanovich, the shift toward optical I/O is becoming unavoidable as system bandwidth demands outpace electrical signalling capabilities. "The CPO revolution is at our doorstep," he said. "Integrating Ayar Labs' optical engines into our advanced packaging flows is a critical step. Our new joint design allows us to address the architectural, power, signal integrity, mechanical and thermal challenges of CPO integration, ensuring customers have access to robust, high-bandwidth and power-efficient solutions."

At the core of the partnership is a new XPU multi-chip package (MCP) architecture that replaces traditional electrical interfaces with optical engines mounted directly onto the organic substrate. The design delivers more than 100 Tbps full-duplex optical bandwidth, representing an order-of-magnitude leap over current-generation XPUs.

The companies are leveraging UCIe-S (64 Gbps) links between the optical engines and I/O chiplets, while UCIe-A (64 Gbps) provides connectivity between the I/O chiplet and the main AI die via local silicon interconnect bridges. The MCP design includes enhancements addressing power and signal integrity at scale, as well as thermal optimisations. A newly engineered stiffener enables detachable fibre connections while maintaining mechanical stability and warpage performance.

GUC plans to share further details on the technology during its presentation, "Advanced Packaging Technologies for Modular and Powerful Compute," at the 2025 TSMC Open Innovation Platform (OIP) Forum on November 18 in Hsinchu, Taiwan.

Access the full potential of hybrid multi-chip integration

Fiber to laser

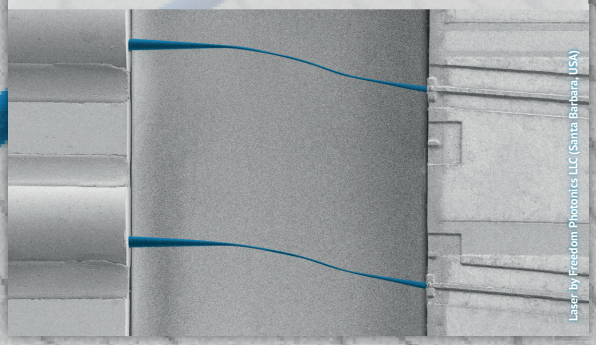

Laser by Freedom Photonics LLC (Santa Barbara, USA)

Laser to SOI chip

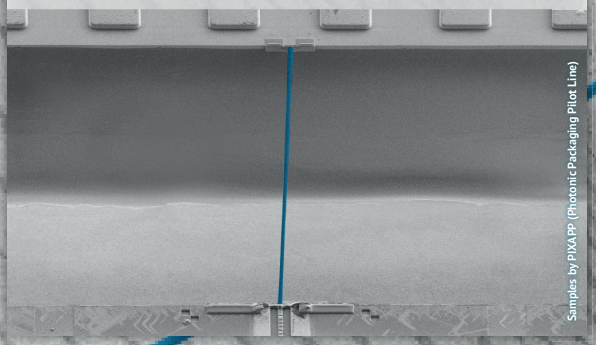

Samples by PIXAPP (Photonic Packaging Pilot Line)

Lens on laser

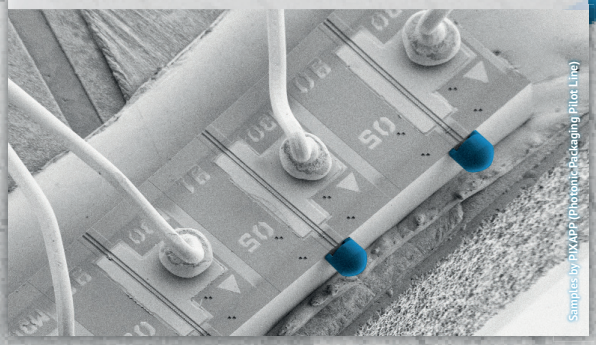

Samples by PIXAPP (Photonic Packaging Pilot Line)

Large scale micro-optics

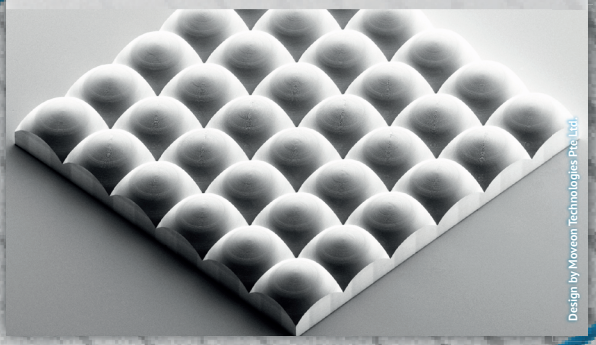

Design by Moveon Technologies Pte. Ltd.

vanguard AUTOMATION

MYCRONIC

High Tech Campus Eindhoven invests in facilities for TNO's 6-Inch photonic chip pilot line

High Tech Campus Eindhoven (HTCE) will host a new industrial pilot line for photonic chips, developed by TNO as part of the EU's PIXEurope project.

Construction of the building and cleanroom facilities will begin in February 2026 on the north side of the campus. Once operational, the facility will serve as both a testbed and fabrication site for advanced Indium Phosphide (InP) photonic chips, enabling full-scale manufacturing at a 6-inch wafer level.

The new pilot line represents a major step toward industrial-scale photonic chip production in the Netherlands and Europe. By combining research and manufacturing capabilities, the facility aims to accelerate the transition from innovative ideas to energy-efficient and reliable real-world technologies. With a capacity to produce up to 10,000 wafers per year, it will significantly increase throughput and efficiency for integrated photonic devices.
Ton van Mol, Managing Director at TNO, described the initiative as a game-changer for Dutch companies and a key part of the PhotonDelta ecosystem that positions the Netherlands as a leader in photonics. Otto van den Boogaard, CEO of HTCE, said the campus is proud to invest in and provide the facilities for the pilot line, reinforcing its role as a hub for breakthrough technology and innovation.

HTCE's investment reflects its commitment to projects that strengthen the region's technological ecosystem and national earning capacity.

The collaboration between TNO, PhotonDelta, and HTCE underlines the importance of local infrastructure in scaling up Europe's photonics manufacturing capabilities and supporting the growth of the integrated photonics industry.

Aluvia Photonics secures funding from PhotonDelta

ALUVIA PHOTONICS, a company developing aluminium oxide (AlOx) PICs, has secured new funding from PhotonDelta to expand its technology and deepen collaborations within the PhotonDelta ecosystem.

The company says this support strengthens its ability to accelerate ongoing projects and explore new opportunities with customers and suppliers in the region.

Aluvia's AlOx platform aims to bring compact, scalable waveguides and amplifiers to the photonic chip ecosystem, targeting telecom, quantum, and UV-visible applications where conventional solutions are limited. Complementary to existing photonic chip technologies, Aluvia says its platform can be combined seamlessly with other platforms, helping to give them new functionality and extend their performance. In this way, the technology aims to strengthen the ecosystem and open new opportunities for designers and manufacturers alike.

"We're delighted to have Aluvia in the PhotonDelta ecosystem and now also as one of our promising startup portfolio," said Erwin Holtland, investment manager at PhotonDelta. "We're excited about the opportunities now possible with their technology and see strong potential across our portfolio. ."

PhotonDelta's funding provides Aluvia with a strategic boost, the company adds, as it prepares for a new investment round in 2026 aimed at scaling up manufacturing. The support aligns with PhotonDelta's mission to strengthen the Dutch integrated photonics ecosystem and attract private co-financing for industrial-scale development.

LuxQuanta raises €8 million for quantum cybersecurity

The company plans to use the investment to expand its teams, scale production, and advance research in integrated photonics – a key technology for developing scalable, quantum-safe solutions

LUXQUANTA, a spinoff of ICFO focusing on quantum cybersecurity, has announced the closing of its €8 million Series A funding round. Led by Big Sur Ventures, and with support of A&G as the main investor, the round also includes new investors GMV, Wayra, and the EIC Fund, alongside renewed commitments from existing investors Corning and GTD. The investment is further bolstered by the soft financing from the European Investment Committee (EIC) through the EIC Accelerator programme, following the €2.5 million grant awarded to LuxQuanta in March 2024.

The company plans to use the funding to scale production, advance research and development in quantum technologies and integrated photonics, expand its commercial, technical, and operational teams, and accelerate international market expansion.

The rise of quantum computing poses a significant risk to current encryption systems, but transitioning to quantum-safe infrastructures is a complex, long-term process. If public and private entities want to keep confidential data protected, immediate action is necessary.

LuxQuanta says its continuous-variable quantum key distribution (CV-QKD) technology is poised to democratise quantum-safe communications, delivering robust, scalable, and accessible solutions for telecommunications, governments, datacentres, financial institutions, critical infrastructures, and industrial and energy sectors.

The company's flagship product NOVA LQ, launched in its second generation in March 2025, is designed to integrate seamlessly into existing optical

networks while offering unmatched security and scalability. According to LuxQuanta, the product also offers high secret key rates, standards-based interoperability, and operational reliability.

"This Series A funding is a powerful validation of our vision to safeguard global communications in the quantum era," said Vanesa Díaz, CEO of LuxQuanta. "With this investment, we will scale our operations, enhance our technology, and expand our global footprint. The journey to quantum-safe networks is critical, and it starts now. We're committed to making quantum security accessible and reliable for organisations worldwide."

Ismael Almazán, partner at Big Sur Ventures, commented: "LuxQuanta's proven CV-QKD technology and successful deployments position it as a leader in quantum cybersecurity. Their focus on integrated photonics is a game-changer, and we're thrilled to support their mission to set a global standard for quantum-safe infrastructure."

LuxQuanta is headquartered in Barcelona with a business unit in Madrid. The company plans to relocate to larger facilities, ramp up production, and selectively expand into new markets in Europe and beyond. The funding will also fuel advancements in integrated photonics, LuxQuanta adds, a key technology for reducing costs and enhancing scalability to meet the growing global demand for quantum-safe solutions.

LuxQuanta's proven CV-QKD technology and successful deployments position it as a leader in quantum cybersecurity. Their focus on integrated photonics is a game-change

PIC INTERNATIONAL CONFERENCE

ANGELTECH

20-22 APRIL 2026
BRUSSELS
BELGIUM

CONNECTING, EDUCATING & INSPIRING THE PIC INDUSTRY

The 11th PIC International builds on the strengths of its predecessors, with around 40 leaders from industry and academia delivering presentations falling within four key themes:

Those attending these sessions will gain greater insight into device technologies while learning of the latest opportunities and trends within the PIC industry. Delegates will also discover significant advances in tools and processes that deliver enhanced yield and throughput.

Attendees at this two-day conference will also meet a wide variety of key players within the community, from investors and analysts to fab engineers and managers.

PIC International is part of AngelTech, which delivers a portfolio of insightful, informative, highly valued chip-level conferences. Bringing together an Innovate Summit and four conferences with more than 120 presentations, more than 700 delegates and over 80 exhibitors,

AngelTech is the premier global event covering compound semiconductors, photonic integrated circuits, power electronic technologies and advanced packaging.

With a significant overlap between the four conferences, attendees and exhibitors are exposed to the full relevant supply chains and customer and supplier bases.

Key Themes for 2026

Foundations of PIC design: materials, devices and processes
While PICs are well established, they still face technological limitations. How can new materials and devices improve their performance and expand their range of functionalities, and how can we accelerate these innovations to market?

Connectivity and scalability for secure, high-speed data networks
Data communications networks face numerous challenges, including soaring AI-driven demands, high energy consumption, and the possibility of future quantum technologies breaking current encryption methods. How can PICs help to solve these issues and underpin the high-speed, energy-efficient, quantum-secure networks of the future?

Emerging applications: photonics for sensing, imaging and beyond
Integrated photonics has a wide range of potential applications. From self-driving cars to miniaturised molecular sensors and non-invasive healthcare, how are PICs making these possibilities a reality?

Future computing: PICs for photonic processing, quantum computers, and neural networks
Computing power is integral to the modern world, but established technologies have limits, and novel systems could herald more powerful devices. From fully photonic processing to photonics for neuromorphic and quantum computing, how are PICs transforming the way we process data?

ANGELTECH

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

SPECIAL Event Discount:

Register before 31 JANUARY 2026 and get 10% OFF!

Registration to AngelTech allows you to access both days of **PIC International**, **CS International, Advanced Packaging** and **PE International.**

EVENT SPONSORS

PLATINUM SPONSORS

Ansys part of Synopsys | APEX Technologies | CRAFT Next Generation Photonic Chips | ficontec | imec embracing a better life | LIGENTEC | LIGHTIUM | LUCEDA PHOTONICS

nanoscribe | PhotonDelta | PI | PICadvanced | QCi Quantum Computing Inc. | vanguard automation MYCRONIC | VPIphotonics

GOLD SPONSORS

ASMPT enabling the digital world | Canada | EPIC European Photonics Industry Consortium | Fraunhofer Heinrich Hertz Institute | ISP System | KRELLTECH Energize Photonics | NAMICS | OPTICA

phix | set | SMART PHOTONICS

MEDIA PARTNERS

electronic specifier .com | Electro Optics | electropages | Microwave Journal | PHOTONICS spectra | PIC Photonic Integrated Circuits | PEI Power Electronics International | procurement pro | Semiconductor Review | SIJ

ORGANISERS

Angel Business Communications | ANGEL EVENTS

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

+44 (0) 2476 718970
info@picinternational.net
picinternational.net

ORGANISED BY

Angel Business Communications | ANGEL EVENTS

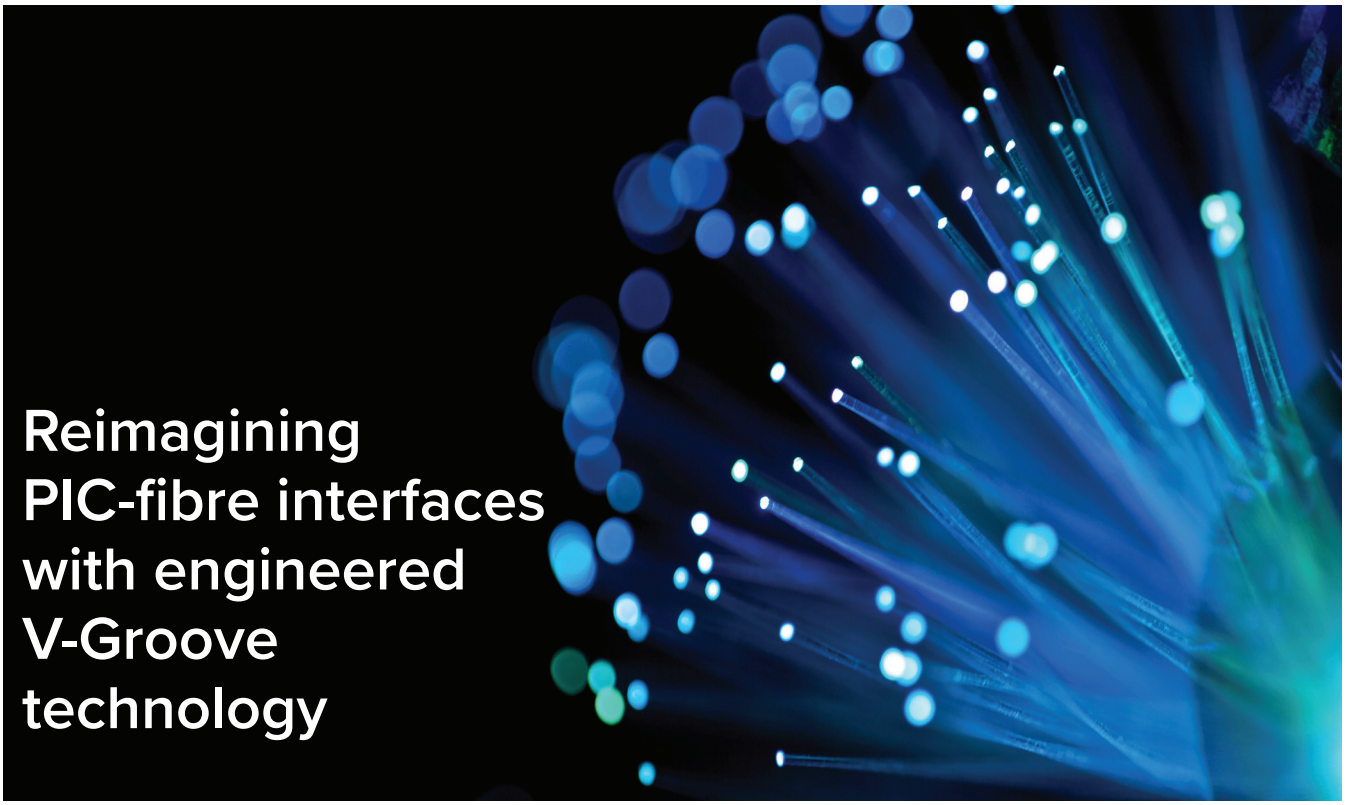

Reimagining PIC-fibre interfaces with engineered V-Groove technology

Engineered V-Groove technology is redefining fibre alignment in photonic integration, offering a cost-effective, scalable solution that bridges performance with manufacturability for next-generation optical interconnects.

BY: AURI RIPOLL, MARKET INTELLIGENCE & STRATEGIC MARKETING MANAGER, DIRK HAUSCHILD, SENIOR STRATEGIC MARKETING EXPERT, AND ZHICHAO HE, MARKETING MANAGER, FOCUSLIGHT

IN THE rapidly evolving world of photonics, PICs have emerged as a transformative force across industries – from ultra-fast datacentres to next-generation AI computing. But despite their potential, a fundamental challenge threatens to slow these chips' adoption: how they connect to the outside world, especially via optical fibres.

At the heart of every photonic system lies the need to align optical fibres to PICs with sub-micron precision. Even a small misalignment can drastically degrade signal quality, introduce insertion losses, or compromise system performance.

Fibre-to-chip interfaces are often overlooked but are mission-critical. They must meet the dual demands of sub-micron precision and high-volume manufacturability, a combination that traditional fabrication techniques struggle to deliver. As PICs become more complex and integrated with AI workloads, datacentres, and quantum systems, the cost of even minor inefficiencies in these interfaces increases exponentially.

This is precisely where engineered V-Groove technology steps in – a quiet innovation with far-reaching implications.

A traditional solution for connecting PICs to fibres involves blade dicing – a mature and well-understood process in which a rotating blade sequentially cuts grooves into a substrate to hold individual fibres in place. But despite its widespread use, the method presents serious limitations for modern PIC applications.

First, there's tool wear. As the blade cuts through glass or silicon, its geometry degrades, leading to inconsistent groove depth and pitch variation, especially across large channel counts (>16). The cumulative result is misalignment between the fibre and the PIC waveguide.

Second, sequential processing, whereby each groove is cut one at a time, introduces positional errors, with tiny inconsistencies accumulating—especially over large arrays. Add to this the impact of thermal drift, substrate warping, and material flaking, and the result is a high failure rate and increased need for active alignment and inspection.

Figures 1-3 show measurements that were made with commercially available components, illustrating some characteristic problems with current methods.

Cumulative errors

The singular random pitch and position errors limit the number of channels per product and require a 100 percent test procedure (both incoming quality control and outgoing quality control) to avoid defective products at functional inspection after assembly is finished. Due to the sequential positioning of the blade, individual cutting characteristics, and degradation of the blade, achieving a defect-free product means using a selection process that will affect the yield and the resulting cost structure of the product.

During assembly of connectors and PIC-fibre interfaces, the height and roll angles – key parameters defining how well the fibre is aligned to the groove – should be reproducible across the fibre array. However, current methods offer insufficient reproducibility. Not only does this result in lower overall performance and increased variation in performance within the array of channels, but it also requires some compensation for errors, increasing the assembly time and cost.

In addition, the larger variation of parameters has implications for using reproducibility and repeatability as a basis for a good process capability index (Cpk) – a statistical measure of how consistently a process can produce output conforming to specified limits.

The combination of errors and the distribution of functional parameters give a clear indication about the process capability and resulting yield in production. The errors and variation of parameters described above lead to a parameter distribution that is not comparable with a typical Gaussian normal distribution which is based on the variation of blade geometry and positioning during V-groove processing. Figure 3 shows a typical distribution of functional parameters of a sequential process that does not show a single “centre of gravity” value with an even distribution on either side.

This observed non-Gaussian distribution is due to the limitations inherent in sequential blade dicing; this process suffers not only from blade wear but also from cumulative 2D positioning errors and environmental temperature drift. These factors lead to systematic degradation of groove geometry, affecting reproducibility and the overall process capability index. In addition, uncontrolled sharp edges and glass flakes from the cutting process can damage delicate optical fibres, further reducing yield.

These issues also compromise compatibility with emerging co-packaged optics (CPO) and multi-channel transceivers, where even a 1 µm deviation can degrade link performance.

Additionally, traditional V-groove arrays are often limited by their reliance on specific materials, usually constrained to glass types that are compatible with mechanical cutting. This restricts compatibility with a broader range of PIC materials, such as silicon, indium phosphide, silicon nitride, and lithium niobate, which are increasingly used in advanced photonic systems.

Finally, traditional blade-diced grooves often struggle to support high-density integration. As industry trends shift toward greater channel counts and denser packaging formats, the limitations of sequential processing become even more pronounced. Blade cutting is not only slower but also more likely to introduce thermal stress and microfractures, especially when pushing toward edge-to-edge array configurations.

Innovation unveiled

Focuslight’s Engineered V-Groove technology solves these problems by rethinking the process from the ground up.

Instead of cutting grooves one by one, Focuslight uses a proprietary wafer-level simultaneous structuring approach. All grooves and alignment features are structured simultaneously across the entire wafer, processing all grooves and assembly features at the same time and thus avoiding cumulative errors inherent to sequential dicing. The process is conducted in a temperature-controlled environment, compatible with large substrates, and allows integration of concave/convex geometries, smooth transition zones, and customised v-angles. 100 percent in-line wafer-level testing ensures statistical reproducibility and Gaussian distribution of positional accuracy.

The benefits are immediate and profound, eliminating tool wear and

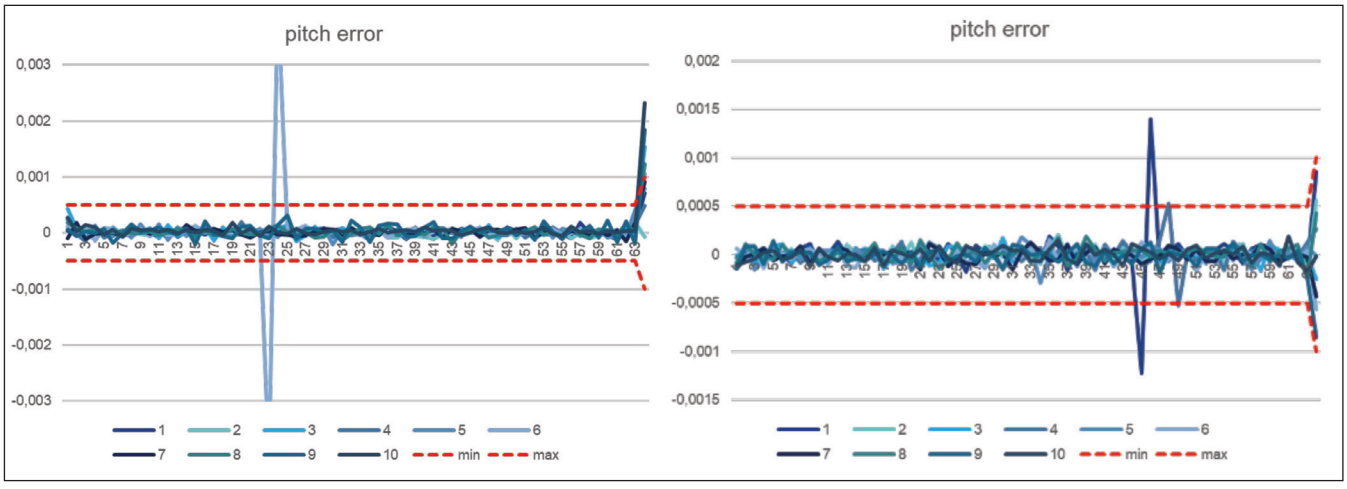

➤ Figure 1: Single channel errors observed in typical >16 channels blade diced v-grooves.

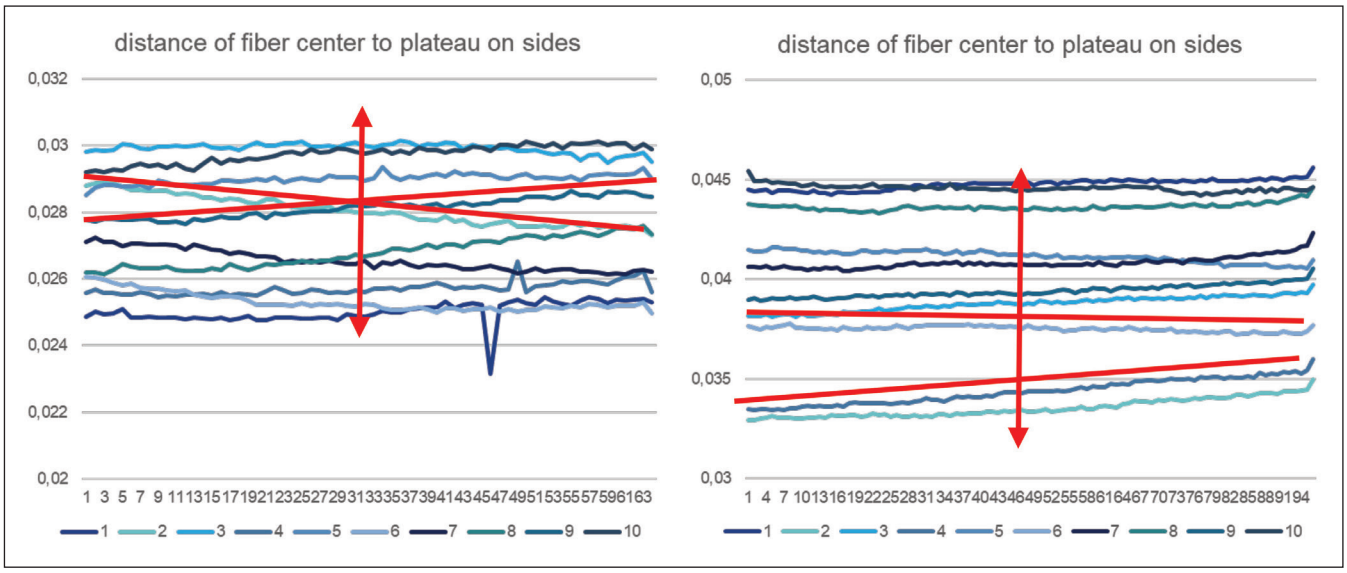

Figure 2: Variation of the position and orientation of fibre arrays within blade diced V-grooves.

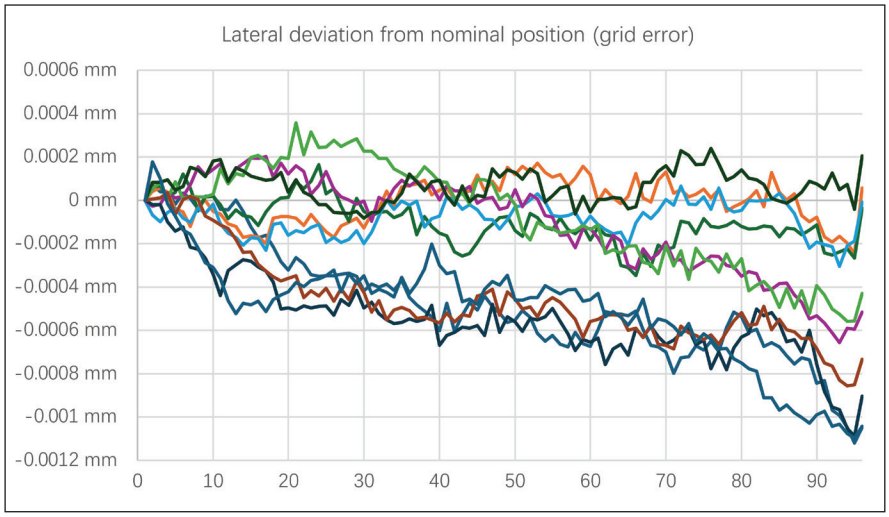

Figure 3: Random fibre pitch variation without a single centre of gravity value.

Figure 4: Wafer-level structured substrate with engineered V-grooves.

geometric limitations of blade dicing, and reducing production time by 40 percent compared to sequential methods. Additionally, this approach offers flexibility for standard and application-adapted designs and mechanical positioning support, and allows customisation with a variable number of channels, and the potential to include assembly features to facilitate the packaging process.

Moreover, this process is fully compatible with wafer-level metrology. Focuslight inspects every unit on the wafer before dicing, ensuring quality control without needing downstream inspection.

Key performance specifications include a lateral grid error of ±0.25 µm in the high-accuracy version of the technology, and ±0.5 µm as standard, with a vertical alignment error of ±0.35 µm (high accuracy), or ±0.7 µm (standard). The single-mode fibre coupling efficiency is greater than 99.5 percent per channel. Additionally, the process is compatible with materials including fused silica, Borofloat 33, S-TIH53, N-BK7, and silicon, and supports pitches of 127 µm, 250 µm, and custom sizes.

In addition to standard features, Focuslight's platform enables customers to design grooves with application-specific mechanical or optical features. These can include thermal compensation features to reduce drift, optical blocking or guiding elements,

recessed or elevated plateau areas for dual-layer packaging, and built-in fiducials for vision alignment systems.

Focuslight's platform also allows for scalable production with reduced waste, minimal edge chipping, and significantly reduced risk of microcracks, making it ideal for use in sensitive optical or high-voltage environments.

With wafer-level control, Focuslight can produce V-groove arrays scalable to 96 channels and higher, tailored to application needs, while ensuring consistency across every groove.

Designed for the real world

What truly differentiates engineered V-grooves from other high-precision technologies is their readiness for production. Focuslight's process is built on an infrastructure originally developed for high-volume micro-optics. With substrate sizes ranging from 0.4 to 300 mm and standard thickness tolerances of ±0.05 mm, it supports both custom and standardised product configurations.

Further capabilities include a fully customisable groove depth and step dimensions of ≥0.1 mm ± 0.02 mm. The solution also offers edge chipping of ≤0.1 mm on outer dimensions (minimised internally) and either a grinded or polished surface finish, depending on optical grade. Focuslight has also developed a wafer-level testing strategy that includes automatic quality mapping prior to back-end separation.

This ensures that 100 percent of units are tested, that the grid deviation and plateau height are within defined tolerances, and that there is no need for post-assembly active alignment.

> Looking ahead, Focuslight continues to invest in next-generation manufacturing technology, including AI-driven process control and data-driven yield optimisation. These enhancements will further improve the consistency, cost-effectiveness, and sustainability of photonics manufacturing

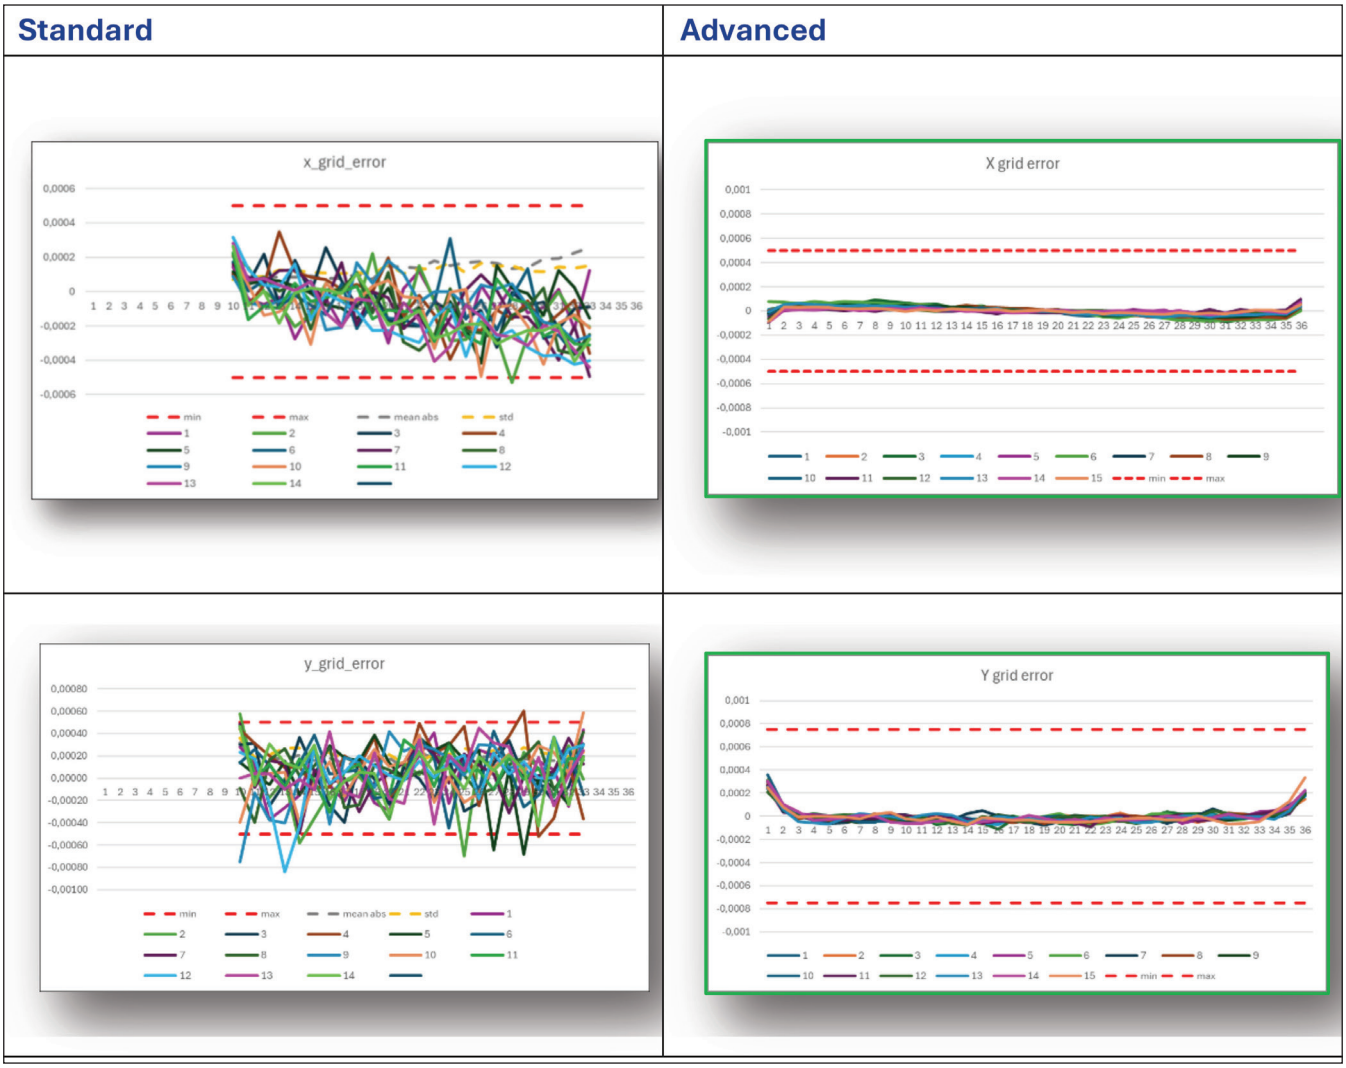

➤ Figure 5: Wafer test result showing Gaussian distribution of grid error.

The resulting process achieves high process capability (Cpk), low variation, and a reliable statistical distribution of performance characteristics – ideal for automation.

With its low-defect, high-throughput process and compatibility with a wide variety of substrates, Focuslight's engineered V-groove platform enables cost-effective production of high-performance optical systems across the full spectrum of commercial photonics applications.

As photonic integration continues to push the boundaries of how fibres and chips are brought together, achieving reliable coupling is no longer only about precision – it is equally about scalability, manufacturability, and long-term system stability. We are entering an era of platformisation in photonics; much like CMOS enabled the rise of ubiquitous electronics, scalable packaging and interfacing solutions will be key to the next wave of optical systems. By removing complex alignment steps and minimising tolerance errors, Focuslight is enabling a future in which PIC modules are manufactured like electronic integrated circuits: reliably, rapidly, and affordably. By shifting from sequential blade dicing to wafer-level simultaneous structuring, the industry gains a platform that unites submicron alignment accuracy with volume production capability.

This convergence opens the door to new design freedoms, from customised geometries to integration with advanced PIC materials. These advantages will support the expansion of photonics into consumer devices, healthcare, edge computing, and satellites, where reliable and scalable interfaces are essential. Looking ahead, Focuslight continues to invest in next-generation manufacturing technology, including AI-driven process control and data-driven yield optimisation. These enhancements will further improve the consistency, cost-effectiveness, and sustainability of photonics manufacturing.

As photonic systems evolve toward higher density and broader application spaces – from AI accelerators to quantum communication – the role of engineered V-grooves will be central in ensuring that performance and scalability advance hand in hand.

- *Usage Statement: All images included in this article are the intellectual property of Focuslight Technologies and are provided with full publication rights for use in PIC Magazine Issue 2025.*

FURTHER READING

➤ 1. Focuslight Technologies, V-Groove Capability Guidelines 2025

OFC

The future of optical networking
and communications

Register Today

It's all here at OFC.
The brightest minds. The boldest innovations.

Technical Conference: 15 – 19 March 2026
Exhibition: 17 – 19 March 2026
Los Angeles Convention Center
Los Angeles, California, USA

OFCConference.org

IEEE ComSoc IEEE Communications Society
IEEE Photonics Society
OPTICA

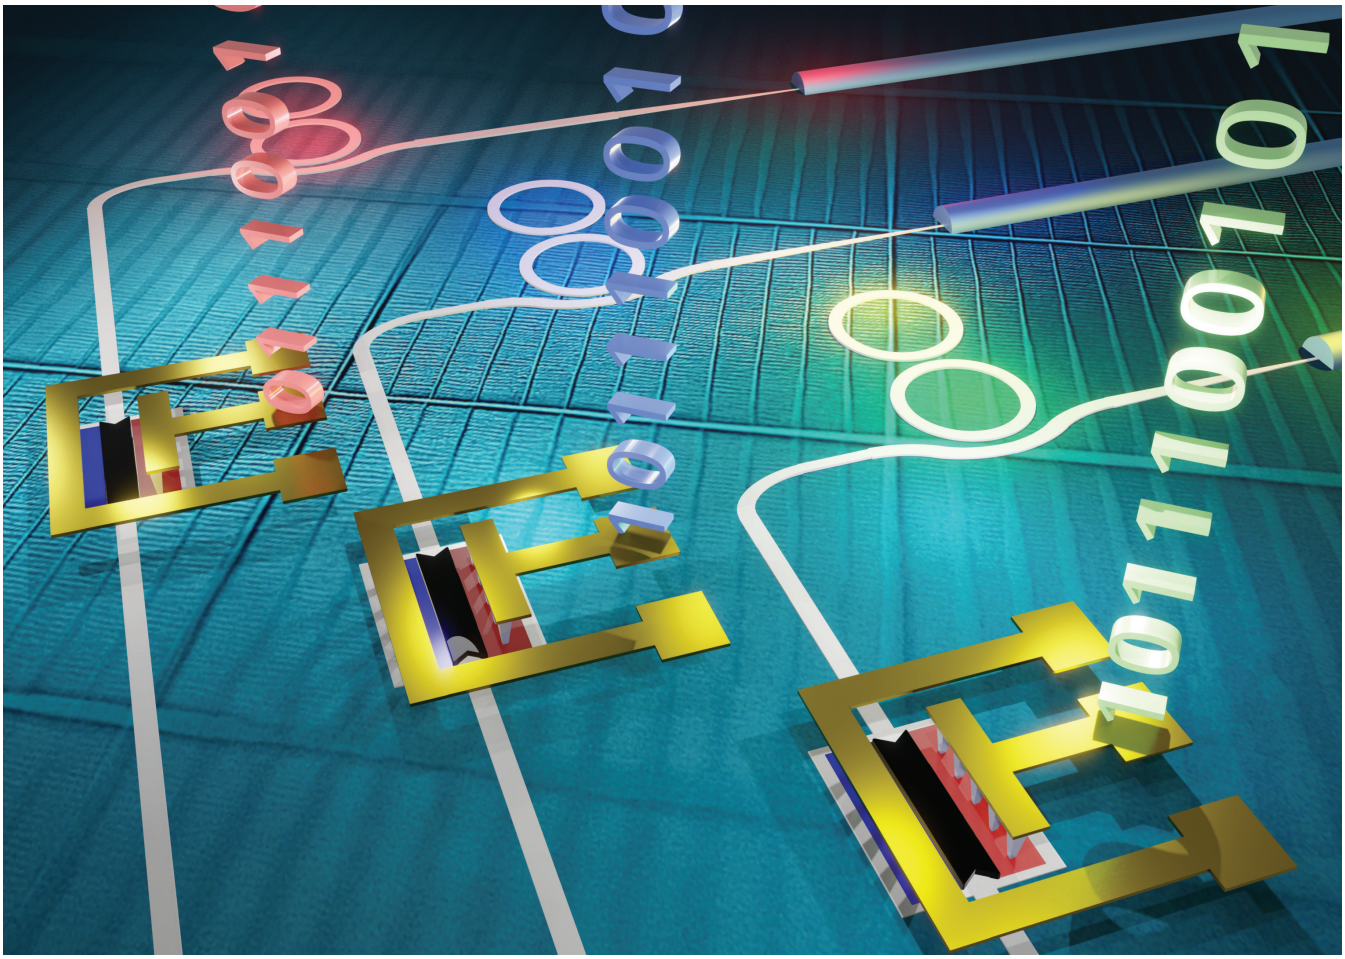

Taming reflections: Ge shape engineering unlocks ultra-low back reflection in high-speed Ge-on-Si photodetectors

Shape-engineered Ge-on-Si photodetectors deliver -36 dB optical return loss, ~60 GHz bandwidth, 0.95 A/W responsivity, and <30 nA dark current - achieved through innovative Ge shaping in a 300 mm CMOS silicon photonics flow, enabling robust, IEEE-compliant solutions for AI data centers and co-packaged optics.

BY YUSHENG BIAN, SUJITH CHANDRAN, ABDELSALAM ABOKETAF, WON SUK LEE, QIDI LIU, MASSIMO SORBARA, BOB MULFINGER AND RYAN SPORER, GLOBALFOUNDRIES; EDGAR HUANTE-CERON, RANOVUS

AS LINK SPEEDS surge for artificial intelligence (AI) clusters and cloud data center interconnects (DCI), back reflection from optical components has emerged as a subtle but escalating threat to link stability and signal integrity. Reflections drive standing waves, jitter, and even laser instabilities; in dense links, they also inflate bit error rate (BER) and degrade signal to noise ratio (SNR). IEEE 802.3 standards enforce tight specifications on both transmitter (TX) and receiver (RX) reflectance to ensure link stability and compliance [1] - shifting some margin from TX to RX can relieve costly TX-side constraints at 100G+/lane. From the field's perspective, ORL is the system-level sum of Fresnel reflections and Rayleigh backscatter along the light path; controlling it at the component level (including photodetectors (PDs) is essential to tame end-to-end reflection budgets in modern links..

Our contribution: shape engineered Ge on Si PIN PDs

In a CMOS monolithic silicon photonics flow [2-9], we redesigned the Ge absorption geometry - moving beyond

the conventional rectangular mesa (Fig.1(a)) to angled, convex/concave, quadrilateral, and pentagon shapes that steer reflected light away from the Si waveguide taper, as shown in Figs.1(b)-(e)). The result is a family of PDs that achieve ORL down to -36 dB while preserving high responsivity (~0.95 A/W), low dark current (<30 nA median at -1 V), and ~60 GHz 3 dB opto-electrical (OE) bandwidth (BW) in the O band [10].

Why it matters: Reducing PD originated reflections relaxes upstream laser/isolator demands and makes it easier to meet stringent ORL budgets in pluggables and future co packaged optics (CPO) - without exotic process steps.

Design: how "shaping" suppresses reflection

Problem: Butt coupling a Si waveguide (WG) taper into a Ge absorption region introduces effective index discontinuities and mode profile mismatch - prime triggers for back reflection at the Si - Ge interface.

Approach: We use geometrical shaping of the Ge region (including angled tips and asymmetric facets) so any reflected field misses the incoming waveguide axis. Finite difference time domain (FDTD) optimization tuned the tip angle, facet lengths, and taper overlap to minimize coherent back coupling while sustaining absorption and contact layouts compatible with foundry design rules.

Outcome: Several candidates deliver ORL < −30 dB, with our lead design (Fig. 2(b) - AngPD_c) reaching -41 dB at 1310 nm. The statistical spread across 10 dies/design remains tight, indicating process tolerance.

Measurement: how we quantified ORL - and everything else

ORL characterization was performed at wafer level using grating coupler test structures and an Optical Vector Analyzer (OVA) with Optical Frequency Domain Reflectometry (OFDR) to map reflection events in the time/space domain. OFDR's micrometre scale spatial resolution is particularly useful in photonic integrated circuits (PICs) where multiple weak reflectors coexist within millimetres.

For readers less familiar with OFDR:

In a CMOS monolithic silicon photonics flow, we redesigned the Ge absorption geometry - moving beyond the conventional rectangular mesa to angled, convex/concave, quadrilateral, and pentagon shapes that steer reflected light away from the Si waveguide taper

it's an interferometric technique using a swept laser and Fourier processing to recover a reflection vs. distance profile with high dynamic range and spatial resolution, ideal for short devices and PICs; it complements Optical Time-Domain Reflectometry (OTDR) and Optical Low-Coherence Reflectometry (OLCR).

High speed response was measured on a 70 GHz lightwave component analyzer (LCA). DC metrics (I V, dark current, responsivity vs. bias/power/temperature) were gathered on temperature controlled probe stations.

Results at a glance

Optical return loss (ORL) (Fig.2(c)):

- Reference (rectangular Ge): higher reflection at PD-WG interface (time domain peaks)
- Shape engineered PDs: multi dB ORL reductions across all variants; best -36 dB => -41 dB at 1310 nm; tight lot to lot dispersion.

Dark current & responsivity (-1 V) (Fig.3):

- Dark current: median <30 nA; most devices <100 nA, consistent with high quality Ge epitaxy and junction control (Fig.3 (c)).
- Responsivity: >0.85 A/W median; up to ~0.95 A/W, matching or exceeding the reference PD (Fig.3 (d)).

Linearity & high power handling:

- Photocurrent remains linear up to several milliwatts; as expected, response begins to roll off at higher powers due to carrier recombination/screening, but increasing reverse bias expands the linear region (e.g., 10% roll off point shifts from ~3.35 mW to ~4 mW at room temperature) (Figs.3 (e) – (g)). The power handling vs. temperature trade off follows known high carrier density dynamics in Ge PDs.

OE bandwidth:

- Normalized OE frequency response shows ~60 GHz 3 dB bandwidth at ~0.2 mA photocurrent; still >50 GHz at ~1 mA (Fig.4). Bandwidth rolls off moderately at higher photocurrents, consistent with carrier screening; we're exploring implant/profile refinements to further stabilize high power bandwidth.

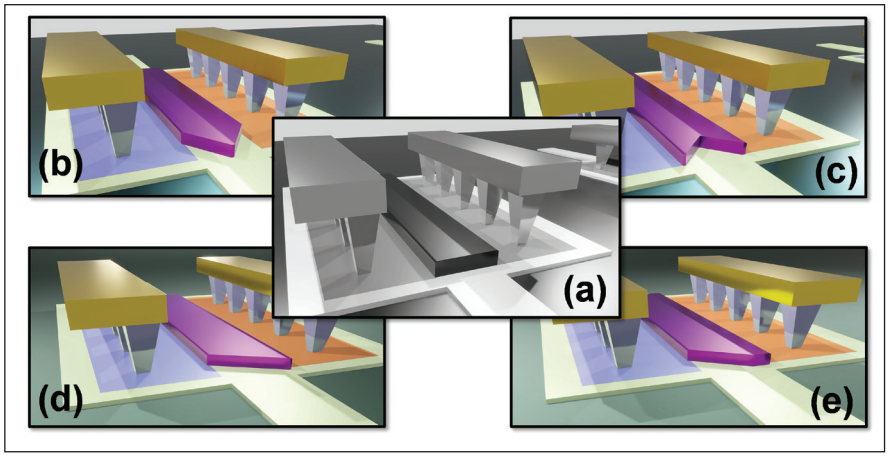

Figure 1. Ge-on-Si PD shape variations 3D perspective views of Ge-on-Si photodetector designs illustrating geometry-driven reflection control: (a) Reference PD with rectangular Ge; (b) convex-shaped Ge; (c) concave-shaped Ge; (d) quadrilateral-shaped Ge; (e) pentagon-shaped Ge. These alternative geometries redirect reflected light away from the Si waveguide taper to minimize ORL

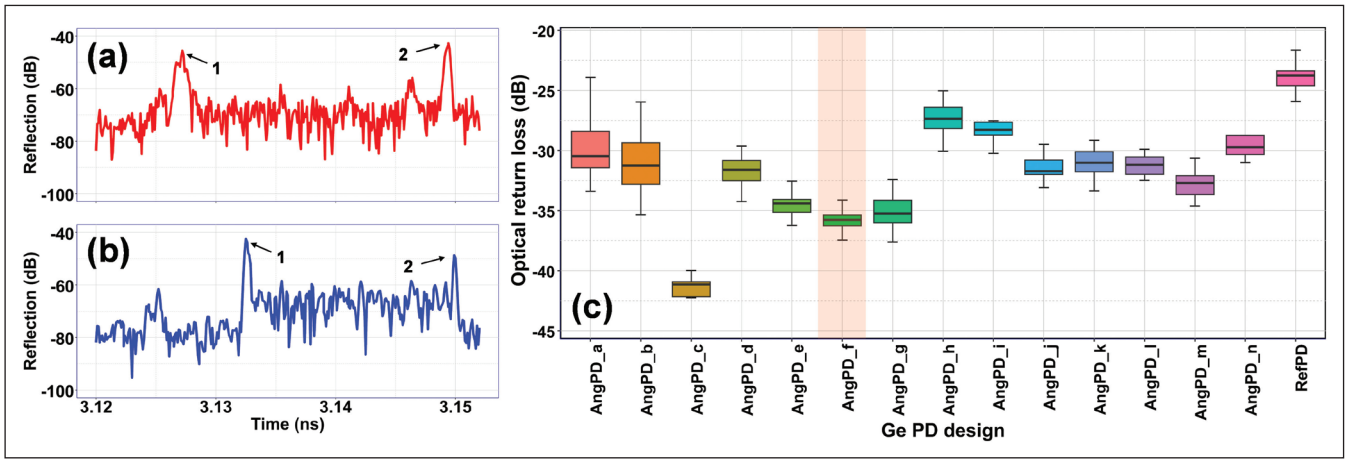

➤ Figure 2. ORL characterization and statistical performance (a) OFDR reflection spectra comparing a reference PD with rectangular Ge and a representative angled PD. The PD-WG interface peak is strongly suppressed in the shaped design. (b) Statistical distribution of measured ORL values for all PD geometries at 1310 nm, showing consistent reductions below -30 dB and best-case performance near −41 dB.

Key point: Across ORL, DC, and RF metrics, no speed/responsivity penalty accompanies reflection suppression - the trade off is effectively neutralized through geometric shaping.

System impact: cleaner receivers, simpler lasers

In Ethernet and AI interconnect roadmaps, ORL and reflectance allocations are hotly debated because every dB of reflection control buys laser stability and may reduce dependency on isolators or complex TX side strategies. By attenuating the PD's own reflection, receiver chains become easier to integrate and more robust, especially for co packaged optics and high radix fabrics where many short optical paths amplify sensitivity to internal reflections.

This work aligns with ecosystem moves to monolithically integrate more of the optical stack and to package at scale as well as system level efforts in CPO and application specific optical engines. Minimizing PD reflection is a small change with big system level dividends.

Where this sits in the literature - and what's next

Ge on Si PDs are a backbone device for SiPh, prized for CMOS compatibility, high responsivity, and strong bandwidth -attributes charted in review literature and sustained by ongoing advances in epitaxy and junction engineering. Our contribution specifically targets the under discussed reflection dimension, demonstrating that geometry alone can reclaim >10 dB of ORL while preserving hallmark performance.

Next steps include:

- Co-optimization with couplers, junctions, and electrical parasitics to further improve responsivity and bandwidth across a wider operating range.
- Integration with TIA and eye-diagram analysis to assess link-level performance and confirm compliance with high-speed standards.
- Module-level validation (e.g., with isolator-free or reduced-isolation TX) to quantify system-level penalties

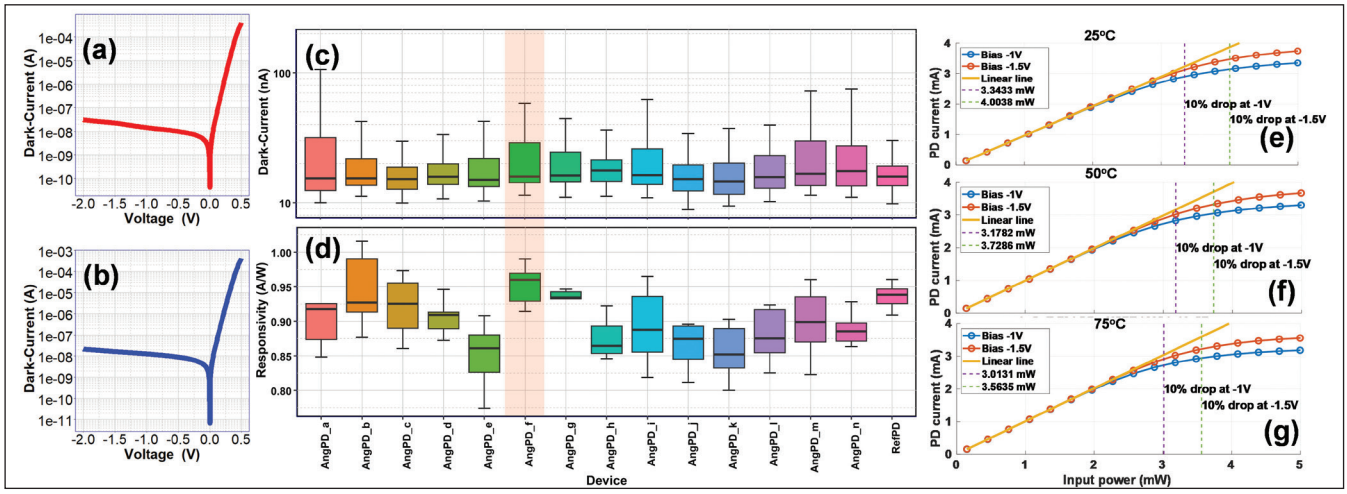

➤ Figure 3. DC characteristics and high-power behavior (a)–(b) I–V curves for reference and angled PD designs, confirming preserved diode behavior. (c) Dark current distribution at -1 V reverse bias, with median values <30 nA. (d) Responsivity at -1 V bias, exceeding 0.85 A/W for most designs and reaching ~0.95 A/W. (e)–(g) Responsivity versus input optical power across temperatures (25 °C, 50 °C, 75 °C) for reverse biases of -1 V and -1.5 V, highlighting improved linearity and power handling with increased bias.

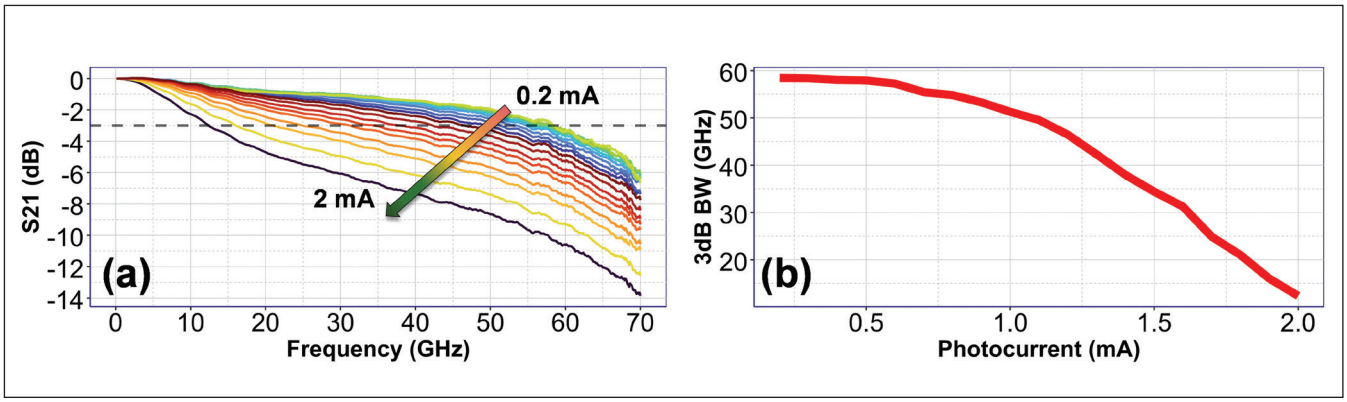

➤ Figure 4. High-speed performance (a) Normalized opto-electric frequency response of the angled PD at varying input optical powers (photocurrent range: 0.2 mA to 2 mA). (b) Extracted 3 dB EO bandwidth as a function of photocurrent, demonstrating ~60 GHz at low photocurrent and >50 GHz at 1 mA, with moderate roll-off at higher powers due to carrier screening.

recovered by low-reflection PDs - particularly in O-band pluggables and CPO engines.

Practical design takeaways (for readers building receivers)

- **Start with geometry**: If you see stubborn reflection peaks at the PD-WG interface in OFDR, try angled or asymmetric Ge tips to dump eflected power off axis. This is a layout-level adjustment using the existing mask set and remains fully compatible with foundry design rules.
- **Validate statistically:** Measure dozens of dies; reflection is phase sensitive and can hide behind path length variations. OFDR's spatial mapping reveals true contribution.
- **Co design bias & power:** Expect bandwidth roll off at higher photocurrent due to screening; allocate headroom in reverse bias to extend linearity and BW in deployment.
- **Think system budgets:** A few dB better ORL at the PD can shift margin in IEEE style allocations and de risk both RX and TX. Coordinate with optics + packaging teams early.

Acknowledgement

The authors would like to thank Michelle Zhang, Judson Holt, Kevin K. Dezfulian, Mankyu Yang, Brian Popielarski, Ming Gong, Shantanu Pal, Javier Ayala, Patrick Snow, Helen Wong, Mrunal Shah, Felix Beaudoin, Frank Pavlik, Crystal Hedges, Frieder Baumann, Takako Hirokawa, Andy Stricker, Oscar Restrepo, Madhuchhanda Brahma, Brett Yatzor, Yu Zhang, Michal Rakowski, Hanyi Ding, Kate McLean, Teodor Stanev, Paul Webster-Pact, Ah Fatt Tong, Jonathan Rullan, Glyn Braithwaite, Vikas Gupta, Rick Carter, Ken Giewont, Ted Letavic, Kevin Soukup, and the rest of the GlobalFoundries team for their development engineering support for Fotonix™ technologies. We also extend our sincere thanks to Hatef Shiran, Ana Villafranca, Kyle Murray, Mohammad Karimi, Akram Hajebifard, Prova Christina Gomes, John Martinho, and the rest of the Ranovus team for their invaluable support with measurements and data analysis.

FURTHER READING

- [1] IEEE 802.3, IEEE 802.3 ETHERNET
- [2] Y. Bian et al., "3D Integrated Laser Attach Technology on a 300-mm Monolithic CMOS Silicon Photonics Platform," IEEE J. Sel. Topics Quantum Electron., Invited Paper, vol. 29, no. 3, pp. 8200519, 2023. DOI: 10.1109/JSTQE.2023.3238290
- [3] V. Stojanović et al., "Monolithic silicon-photonic platforms in state-of- the-art CMOS SOI processes," Opt. Express, vol. 26, no. 10, pp. 13106–13121, 2018.
- [4] K. Giewont et al., "300-mm monolithic silicon photonics foundry technology," IEEE J. Sel. Topics Quantum Electron., Invited Paper, vol. 25, no. 5, Sep. 2019, Art. no. 8200611.
- [5] M. Rakowski et al., "45nm CMOS silicon photonics monolithic technology (45CLO) for next-generation, low power and high speed optical interconnects," in Proc. Opt. Fiber Commun. (OFC) Conf. Exhibi., T3H.3, 2020.
- [6] Y. Bian et al., "Towards low-loss monolithic silicon and nitride photonic building blocks in state-of-the-art 300mm CMOS foundry," in Proc. Frontiers Opt. (FiO), pp. paper FW5D. 2, Sep. 2020.
- [7] T. Letavic et. al., "Monolithic Silicon Photonics," Proceedings of BCICTS, 2a.1, Invited Talk, 2023.
- [8] T. Hirokawa, et al., "Latest Progress and Challenges in 300mm Monolithic Silicon Photonics Manufacturing," in Proc. Opt. Fiber Commun. (OFC) Conf. Exhibi., Th3H.2, Invited Talk, 2024.
- [9] Y. Bian et al., "300-mm monolithic CMOS silicon photonics foundry technology," Conference on Lasers and Electro-Optics (CLEO), ATu3H.1, Invited Talk, 2024.
- [10] Y. Bian et al., "Monolithically integrated high-performance Ge-on-Si PIN photodetectors with nearly 60 GHz EO bandwidth, 0.95 A/W responsivity, <30 nA dark current and -36 dB optical return loss enabled by Ge shape optimization," in Proc. Opt. Fiber Commun. (OFC) Conf. Exhibi., Th3E.6, 2025.

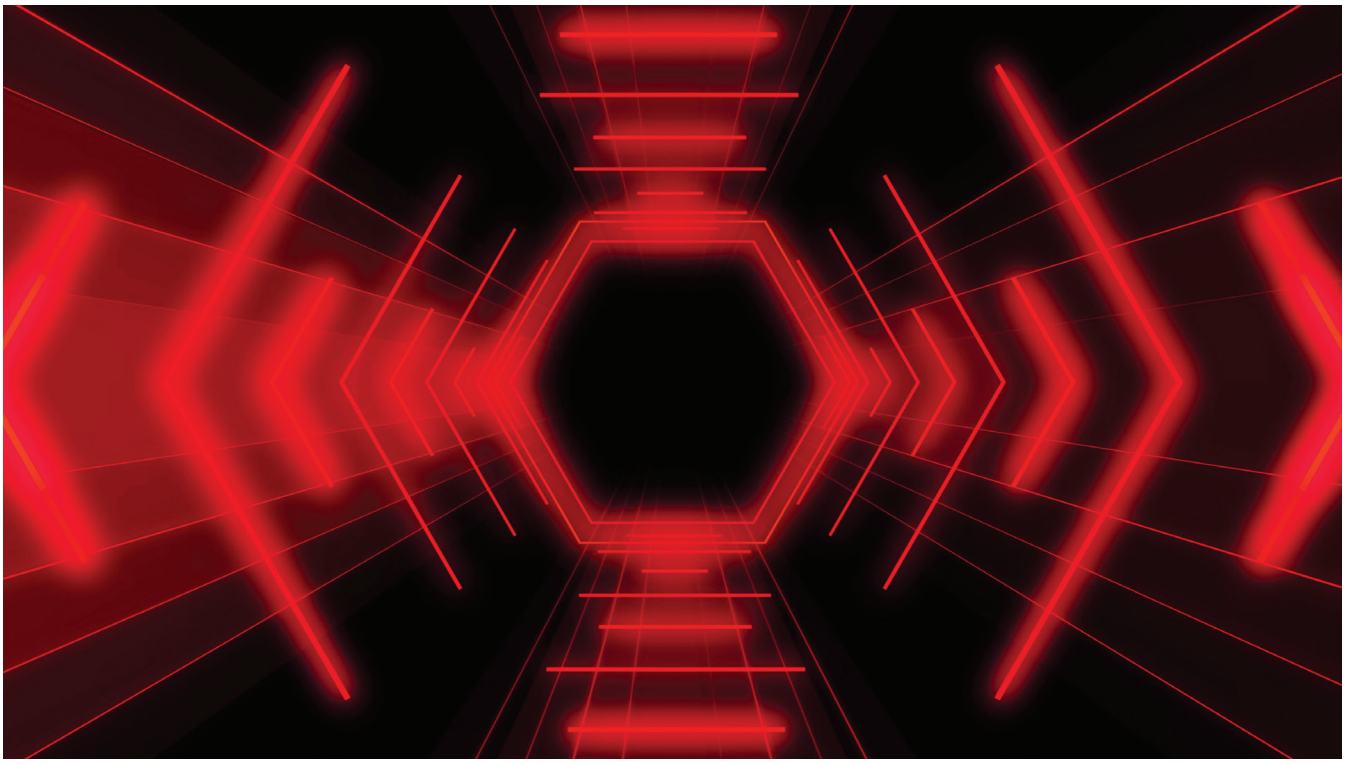

The race to revolutionise silicon photonics with seamless III-V integration

Growing III-V nano-ridge lasers directly on silicon sidesteps the complexities of existing integration methods and paves the way for scalable, cost-efficient photonic integrated circuit production

BY BERNARDETTE KUNERT AND JORIS VAN CAMPENHOUT FROM IMEC

IN TODAY'S data-driven world, the demand for faster and more efficient data computation and transmission is growing at an unprecedented pace. On top of the everyday expectation of instant access to high-resolution images, seamless video streaming, and immersive augmented reality experiences, the rapid expansion of AI and machine learning introduces even greater challenges. These technologies rely on massive volumes of machine-to-machine data exchange, further amplifying the need for high-speed, low-latency communication.

To keep up with these demands, optical data transmission must extend beyond traditional long-haul networks, moving closer to the chip level. This has implications for fibre-to-the-x deployments and chip-to-chip optical interconnects, in which ultra-fast, low-power data transfer is becoming increasingly critical. As data volumes continue to surge, innovative solutions to optimise transmission efficiency are paramount. A particular focus is reducing power consumption per transmitted bit to ensure computing infrastructure is both scalable and sustainable.

One key technology in this transformation is silicon photonics, which leverages the well-established CMOS fabrication process to enable large-scale production of optical systems.

However, a longstanding barrier to the full-scale deployment of silicon photonic integrated circuits (PICs) is that they lack one critical element: the light source. Due to its inherent material limitations, silicon cannot efficiently emit light. The industry has so far filled this gap by using III-V compound semiconductors – materials prized for their exceptional optoelectronic properties. But integrating these crystalline materials into silicon photonics has proven to be a formidable engineering challenge.

For this reason, most datacom products currently rely on III-V light sources that are separately processed on native III-V substrates and later attached to silicon photonic chips using micro-mounted laser packages or high-precision flip-chip assembly techniques. While these hybrid integration approaches are effective, they are also costly and tricky to scale, making PIC devices an expensive, low-volume solution.

To address this, the industry is exploring alternative integration methods. For example, micro-transfer printing, a back-end-of-line technique, involves the parallel transfer of prefabricated III-V components onto silicon photonic wafers, significantly improving throughput. Another heterogeneous integration approach employs die-to-wafer bonding of unprocessed III-V device chips onto silicon photonic wafers, followed by III-V device patterning and CMOS-based back-end-of-line metal interconnect formation.

The latter method has gained commercial traction and is now available in at least two commercial manufacturing lines. Yet its dependence on complex bonding techniques and expensive III-V substrates remains a bottleneck. Additionally, the substrates are discarded as waste during the manufacturing process, raising concerns about health, safety, and environmental sustainability.

The ultimate goal is the direct monolithic growth of III-V materials on silicon wafers, which would eliminate the need for external III-V substrates and intricate assembly processes. However, silicon and III-V materials have significant differences in their lattice structures, leading to strain in the III-V layers during this monolithic growth. When this strain releases, it leads to defects such as misfit and threading dislocations forming in the III-V film during deposition. These crystal defects do not just appear at the interface with silicon, but also penetrate the entire device stack, degrading its performance and rendering it commercially unusable.

Nano-ridge engineering

For decades, researchers have worked to control the inevitable strain release and block the propagation of dislocation defects through the device stack to the active layers. Some of the techniques explored include growing very thick transition buffers between the silicon and III-V, employing annealing treatments to fix some of the defects, and adding strained superlattice layers to control and isolate defects, accommodating them while limiting their wider propagation. However, these approaches have met with limited success.

A game-changing development emerged with the successful demonstration of InAs quantum dot lasers monolithically grown on silicon. These zero-dimensional nanostructures confine charge carriers, the electrons and holes, more tightly than the more common quantum wells, in which carriers can move in two dimensions. Electrons and holes in quantum dots are therefore less likely to encounter defects, meaning that quantum-dot structures exhibit greater defect tolerance than traditional lasers based on multi-quantum well gain regions. Thanks to these properties, quantum-dot lasers have achieved reliability levels that make them a promising technology for future PIC applications.

One fabrication method that has driven recent advancements in monolithic growth and in integrated photonics more broadly, is selective-area growth (SAG). This technique allows the deposition of III-V material exclusively within predefined silicon oxide patterns, minimising the need for extensive material removal during fabrication, and making the process more efficient than conventional two-dimensional growth. SAG can also leverage aspect ratio trapping (ART), which involves creating deep features that can trap relaxation defects and suppress their propagation through the device. Since only selected regions of the silicon substrate are overgrown, rather than the whole surface, this approach also reduces strain and mitigates common challenges, including crack formation and silicon wafer warpage.

Yet, despite these advantages, SAG based on ART has a drawback; the narrow deposition patterns that are critical for effective defect reduction also limit the amount of III-V material that can be integrated. This constraint restricts the range of devices that can be realised.

Enter nano-ridge engineering (NRE), a cutting-edge integration approach

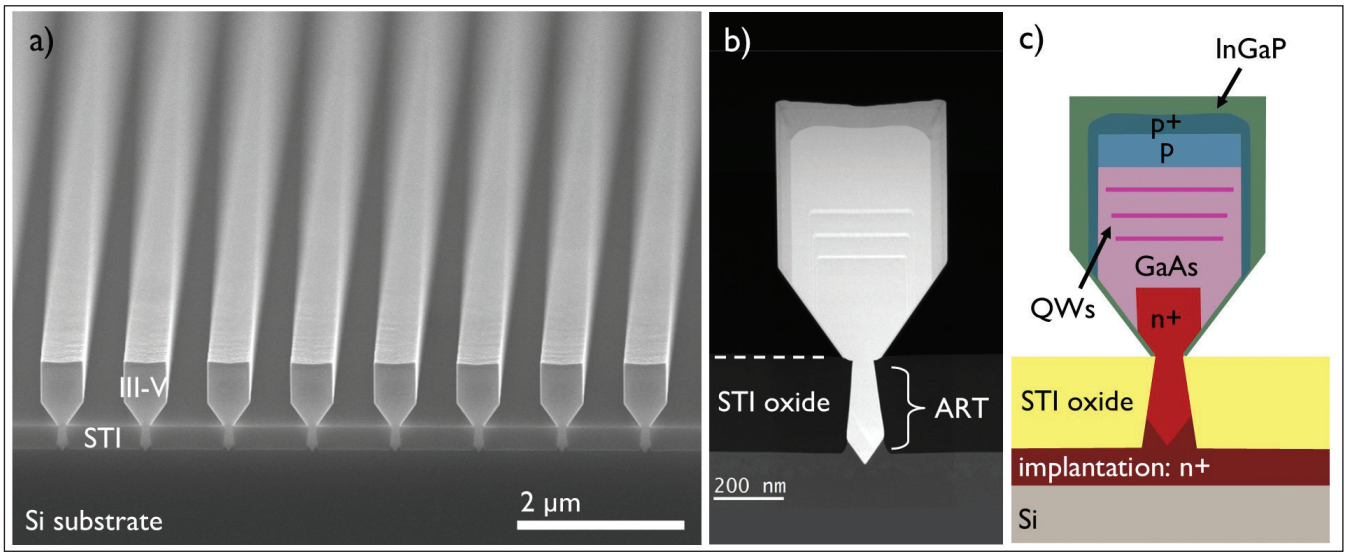

➤ Figure 1. (a) Scanning electron microscope image of a III-V nano-ridge array grown by MOCVD on a 300 mm patterned silicon dioxide/silicon wafer. The oxide pattern was fabricated by applying a shallow trench isolation (STI) process. (b) High-angle annular dark-field scanning transmission electron microscope image of a III-V nano-ridge laser in cross-section highlighting the different III-V material systems. (c) Schematic of the nano-ridge device stack (QW: quantum well; n^+: highly *n*-doped region; p: *p*-doped region; p^+: highly *p*-doped region).

pioneered by imec. By refining SAG with MOCVD – the preferred deposition technique to achieve selectivity – and enhancing ART through the use of very narrow, elongated trenches, NRE enables efficient defect suppression while significantly increasing the volume of usable III-V material. This is achieved through continuous growth beyond the trench pattern and carefully controlled nano-ridge shape engineering.

This breakthrough not only expands the design possibilities for III-V-based devices, but also redefines their functionality. The nano-ridge itself can serve as a waveguide in optoelectronic components, such as light-emitting diodes, lasers, modulators, and photodetectors. Additionally, when multiple nano-ridges are contacted in parallel, they can provide sufficient current to power a heterojunction bipolar transistor – a type of transistor that is often used as a power amplifier in RF communications.

Imec's recent milestone – the demonstration of nano-ridge lasers fully processed on 300 mm silicon in a CMOS prototyping line – marks a major step toward scalable, cost-efficient, and environmentally sustainable integration of high-quality III-V devices on silicon.

Fabrication challenges

To make this vision a reality, imec's researchers had to overcome three major hurdles. The first challenge was ensuring defect-free nano-ridge growth on 300 mm silicon substrates. ART has been proven effective in reducing misfit defects across various material systems, particularly for GaAs inside narrow trenches. However, for this approach to work, complete strain relaxation must occur within these trenches before uniform nano-ridge growth can extend beyond them. To achieve this, the imec team optimised the MOCVD conditions to facilitate efficient strain release at the III-V/silicon interface and gliding of the threading dislocation inside the trench.

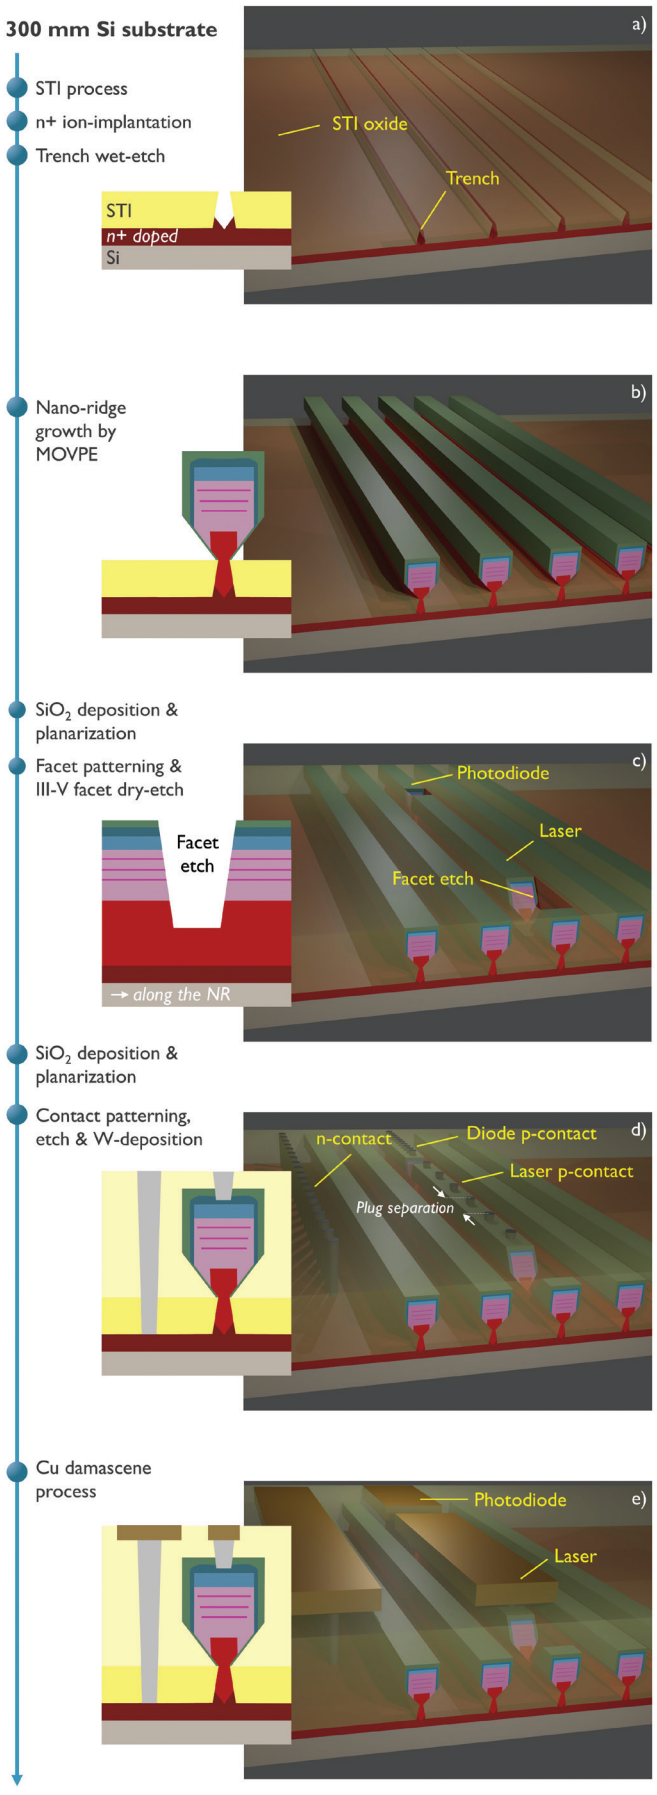

III-V nano-ridge laser diode + read-out photodiode

➤ Figure 2. Simplified schematic of the 300 mm process flow executed in imec's CMOS prototyping line (W: tungsten; Cu: copper). For more details see Y. D. Konick et al. Nature **637** 63 (2025)

Once a uniform *n*-doped GaAs base-ridge was established, we grew an InGaAs/GaAs multi-quantum well stack to serve as the laser's optical gain medium, followed by a *p*-doped GaAs contact layer. We then capped the final nano-ridge waveguide with an InGaP layer to enhance carrier confinement and to reduce carrier losses at the III-V surface (see Figure 1).

Measurements of the unprocessed nano-ridges (using cathodoluminescence techniques) confirmed an impressively low density of misfit defects – fewer than 6 x 10^4 per cm^2 – in the nano-ridge waveguide. This is a remarkable achievement for such a thin III-V stack.

The second challenge was designing a low-loss contact approach. It was relatively straightforward to perform electron injection via a silicon layer implanted with n^+ ions and the *n*-doped GaAs inside the trench. However, establishing a robust *p*-contact on top of the nano-ridge proved more difficult. Using a continuous metal stripe on top of the nano-ridge would have led to excessive optical losses. Instead, we distributed isolated *p*-plugs along the nano-ridge.

To enable efficient hole injection, these metal plugs were selectively punched through the InGaP layer to reach the *p*-doped GaAs. This periodic *p*-contact design introduced a unique beating-mode formation, supporting discrete waveguide modes with minimised overlap with the metal plug array – a phenomenon that also supported single-mode laser operation. To create the Fabry-Pérot laser cavity, we etched facets into the nano-ridge using a dry-etch process.

Finally, integrating and processing III-V materials within a CMOS prototyping line presented an additional set of challenges. Most semiconductor manufacturing tools are not designed to handle III-V materials, or even tolerate the risk of contamination from III-V traces on 300 mm silicon wafers. Furthermore, we had to develop many of the necessary fabrication processes from scratch. To do this, we leveraged imec's extensive expertise in III-V material processing, establishing a controlled manufacturing environment tailored to accommodate these non-standard CMOS materials.

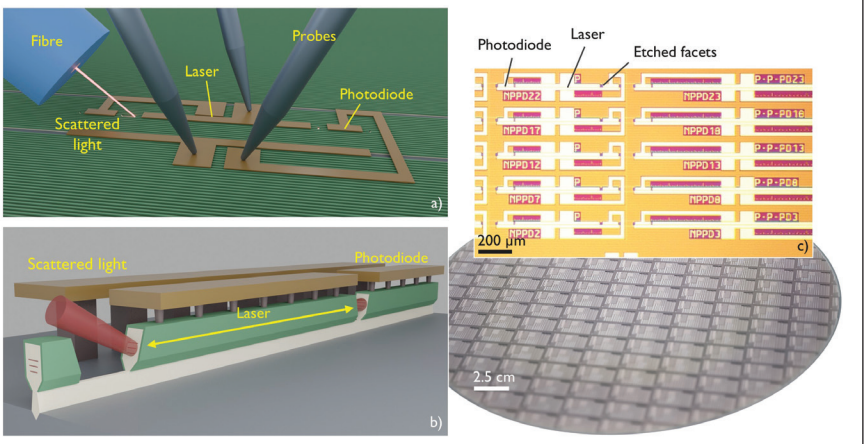

Figure 3. (a) Schematic of the wafer-scale readout test setup: light scattered upwards from the left facet can be collected by a multimode fibre. The three electrical probes drive the nano-ridge laser and photodiode. (b) Schematic (ignoring the oxide) of the nano-ridge cavity formed by two etched facets with an inline photodiode collecting light emitted from the right laser facet, while the light from the left facet is scattered. (c) An optical top-view image of processed devices. A photo of a completely processed 300 mm wafer is shown in the background.

In light of the inherent uncertainties associated with novel sequential process steps and device operability, we devised a mask layout that incorporated a diverse set of physical device parameters. This included variations in laser cavity lengths, pattern trench widths, and *p*-plug separations to systematically assess their impact on device functionality and increase the probability of a successful demonstration.

To enable rapid, wafer-scale characterisation of these nano-ridge laser devices, we introduced an innovative approach: integrating a nano-ridge photodetector in line with the laser. This coupled device layout provides an elegant and efficient method for obtaining wafer-scale performance statistics, ensuring a comprehensive evaluation of device operation across the wafer. Further details on the 300 mm process flow and measurement configuration are provided in Figures 2 and 3.

Nano-ridge lasers

Leveraging the wafer-scale fabrication, we put thousands of GaAs nano-ridge devices – including lasers, photodetectors, and test structures – to readout test, as shown in Figure 3. The results were highly promising: over 300 functional nano-ridge lasers exhibited emission at approximately 1025 nm at room temperature, as evidenced by a clear threshold behaviour in the photodiode current (see Figure 4).

This milestone is particularly noteworthy considering the challenges of pioneering a completely new 300 mm process flow. Risks such as short circuits, imperfect facet quality, and high contact resistivity could easily have hindered performance. Instead, these findings mark a major step forward, proving that nano-ridge lasers can be reliably fabricated at scale, and bringing III-V integration on silicon closer to real-world applications.

To validate these findings, we conducted additional measurements, including wafer-scale measurements detecting radiated laser light from the top of the wafer, as well as die-level analysis, for example of the emission spectra from cleaved facets of diced wafers. These tests confirmed the laser operation, revealing a threshold current as low as 5 mA, a slope efficiency of up to 0.5 W A^{-1}, and a maximum total output power of 1.75 mW. These are encouraging values, given the nano-ridge's sub-micrometre cross-section. Furthermore, the periodic grating structure formed by the *p*-contact plugs played a crucial role in stabilising single-mode operation in the Fabry-Pérot resonator, achieving a side-mode suppression ratio exceeding 30 dB.

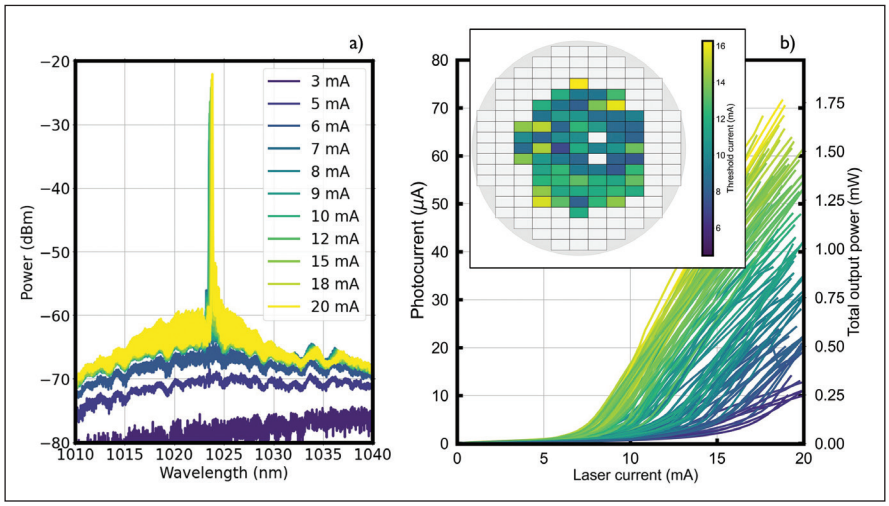

➤ Figure 4. (a) Optical output spectra collected at a cleaved facet of a 1.4 mm-long laser cavity with an optical fibre for different drive currents. (b) Photodiode current (left) and related total output power (right) versus the laser drive current of a 2 mm-long device. The inset shows a wafer map of the die distribution containing the operational 2 mm-long laser devices. The colour code indicates the corresponding threshold current.

A particularly notable observation is that lasing occurred only in devices where the *p*-plug separation exceeded 3 μm, highlighting that, if the plug density is too high, the absorption losses in the top metal contacts have a significant impact on the laser threshold.

However, increasing the *p*-plug separation introduced a device performance trade-off as it also led to higher current densities at each contact point. While initial reliability tests showed continuous-wave lasing for over 500 hours, further investigation revealed that the high current density at the *p*-plugs induced localised defects, ultimately leading to device failure.

A more promising observation was that, aside from the region surrounding the *p*-plugs, the rest of the nano-ridge waveguide retained high crystalline quality, with no additional misfit defects forming. This finding suggests that further optimisation of the contact design could significantly enhance the long-term reliability of these nano-ridge lasers, paving the way for silicon photonics applications.

Future focus

To the best of our knowledge, this demonstration marks the first successful realisation of a fully processed III-V nano-ridge laser on a 300 mm silicon wafer. By enabling monolithic III-V deposition directly onto silicon, this approach eliminates the dependence on III-V substrates and complex bonding techniques, leveraging the advanced fabrication capabilities of a CMOS pilot line. This breakthrough is pivotal in facilitating high-volume scalability, improved yield, and reduced production costs – key factors for the widespread adoption of silicon photonics. However, while this achievement represents a major milestone, the nano-ridge engineering technology remains in its early development phase.

This work is part of a broader pathfinding mission at imec to advance III-V integration processes towards higher technological maturity. In the near term, the focus remains on hybrid approaches, such as flip-chip assembly and transfer printing. In the mid-term, we expect heterogeneous methods based on die-bonding techniques to further enhance integration efficiency. Ultimately, the long-term goal is to achieve direct epitaxial growth of III-V materials on silicon – a concept that this work demonstrates is possible.

Current R&D efforts in our NRE technology focus on refining the contact approach to mitigate localised high-current injection spots, thereby extending device lifetime. Additionally, we are investigating alternative GaAs-based optical gain media to red-shift the laser emission wavelength, expanding the potential range of applications.

Another promising avenue involves leveraging the waveguide nature of III-V nano-ridges to enable efficient coupling into silicon photonic waveguides. This would open up the possibility of developing external cavity lasers, further enhancing the versatility of this platform.

Beyond device-level innovations, significant work remains to optimise fabrication processes and improve throughput. Achieving full-scale commercialisation will require close collaboration with semiconductor tool suppliers to improve processing techniques and ensure compatibility with existing CMOS manufacturing infrastructure. Encouragingly, ongoing global efforts in heterogeneous integration, combined with the growing importance of III-V process technologies in CMOS environments, are set to drive continuous improvements in tool capabilities and manufacturing efficiency.

By addressing these challenges, imec aims to accelerate the development of III-V-on-silicon lasers, bringing them closer to large-scale deployment in future-generation silicon photonics applications.

FURTHER READING

- Y. D. Konick *et al*. “GaAs nano-ridge laser diodes fully fabricated in a 300-mm CMOS pilot line” Nature **637** 63 (2025)
- C. I. Ozdemir *et al*. “Low Dark Current and High Responsivity 1020nm InGaAs/GaAs Nano-Ridge Waveguide Photodetector Monolithically Integrated on a 300-mm Si Wafer” J. Light. Technol. **39** 5263 (2021)
- Y. Shi *et al*. “Novel adiabatic coupler for III-V nano-ridge laser grown on a Si photonics platform” Opt. Express **27** 37781 (2019)
- P.-Y. Hsieh *et al*. “Degradation Mechanisms of Monolithic GaAs-on-Si Nano-Ridge Quantum Well Lasers” J. Light Tech. **43** 5811 (2025)

LUCEDA PHOTONICS

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.

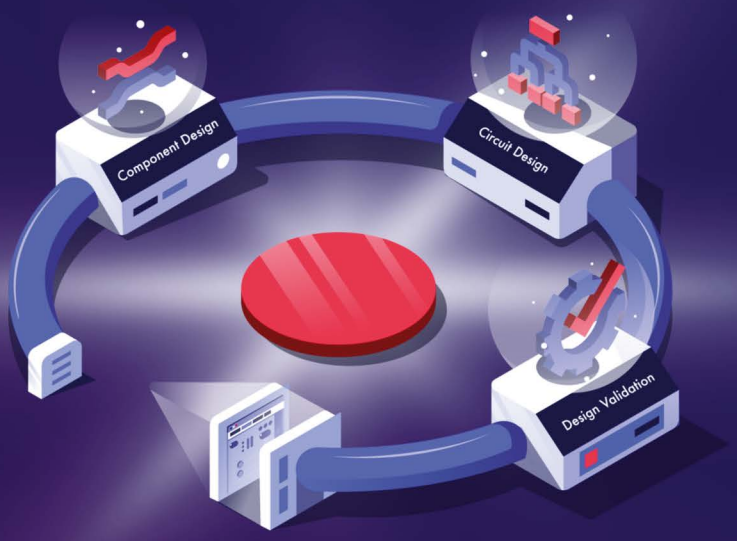

Luceda Design Kits

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

SMART PHOTONICS · amf ADVANCED MICRO FOUNDRY · leti CEA tech · CompoundTek · imec · OpenLight · AUSTRALIAN SILICON PHOTONICS · AIM PHOTONICS · Fraunhofer HHI · imec PIX4life · CORNERSTONE · ALUVIA · IMECAS · Tower Semiconductor · SiEPIC FAB · VLC PHOTONICS A Hitachi Group Company · ANT NANOSOI · csem · LIGENTEC · phix PHOTONICS ASSEMBLY · ihp · LIONIX INTERNATIONAL · 南智光电 IOPTEE · ALTER · SITRI · Tyndall National Institute · LXT · SiPhotonIC TECHNOLOGIES · UMEC

Luceda Modules

Design advanced photonic components and complement Luceda IPKISS with specialized modules.

- Luceda IP Manager
- Luceda AWG Designer
- Link for Ansys Lumerical
- Link for Tidy3D
- Link for Dassault Systèmes Simulia

DEVICE SIMULATION

- Link for Siemens EDA
- Link for Check Mate DRC

www.lucedaphotonics.com

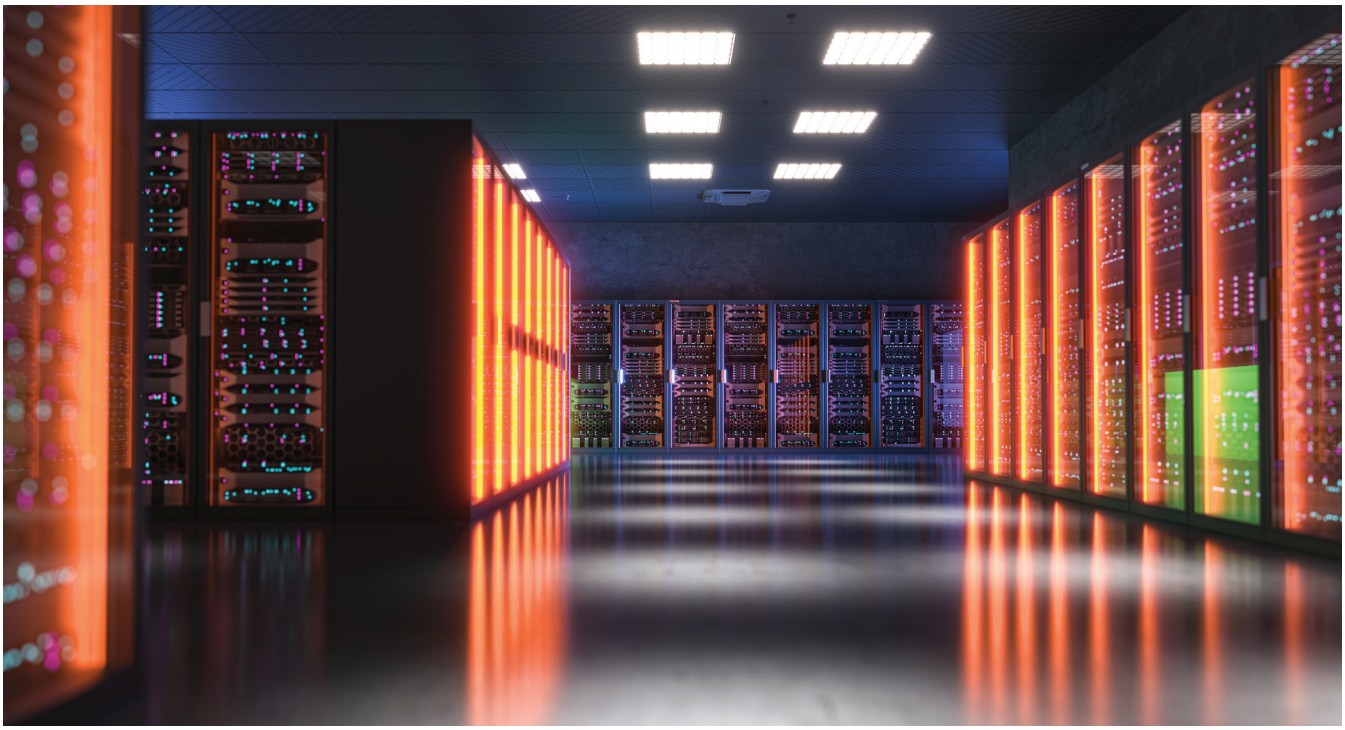

Giving AI room to grow with ultra-compact silicon modulators

Scaling AI infrastructure requires ultra-fast communications across distant processors, but conventional silicon modulators are bulky and power-hungry. A new, ultra-compact silicon microring modulator has demonstrated record-breaking transmission speeds at low power, offering a solution that could give AI space to grow.

BY ALIREZA GERAVAND, LESLIE A. RUSCH AND WEI SHI, CENTRE FOR OPTICS, PHOTONICS AND LASERS (COPL), UNIVERSITÉ LAVAL, CANADA

AI IS fundamentally transforming every sector, from healthcare and finance to automotive and consumer technology. This revolution is being powered by enormous computing clusters – thousands of computing nodes spread across multi-acre campuses—and as AI models grow in size, they require exponentially more processors working in parallel. OpenAI's GPT-4, for instance, used 25,000 interconnected computing nodes over 90 days for training, but next-generation clusters are projected to involve 300,000 or even a million nodes in the near future.

To operate at this scale, massive volumes of data must be exchanged rapidly and efficiently, both between compute nodes within a cluster and between clusters. Optical interconnects form the backbone of these systems, providing high-speed, low-loss data transmission through optical fibres.

This relies on the efficient conversion of electrical data from processors into light—a process that is enabled by silicon photonics technology, which also benefits from mature manufacturing capabilities.

However, the unprecedented communication demands between AI processors are creating a performance bottleneck – one that risks slowing the pace of AI advancement. Today's optical links in AI clusters often rely on simple intensity modulation and detection – essentially turning a light on and off to send data. While cost-effective, this approach faces increasing challenges in meeting AI's soaring requirements for both speed and reach. Next-generation silicon photonics-based technologies must deliver far higher data capacity and longer reach, while scaling seamlessly as AI clusters expand. The solution is coherent optics, a technology that already dominates long-distance applications. Leveraging the coherent properties of light, we can go beyond the simple on/off flashlight approach and instead modulate both the amplitude and phase of light.

This advanced modulation dramatically increases transmission capacity, which explains its widespread adoption in long-haul networks. It also offers greater robustness to link impairments and can reach longer distances at the same signal power.

However, conventional coherent optics has so far been bulky, power-hungry, and challenging to integrate with co-packaged optics (CPO). The primary culprit is the Mach-Zehnder modulator, the most commonly used modulator in coherent systems.

To achieve truly massive data rates, all available wavelengths must be exploited through wavelength-division multiplexing (WDM). While Mach-Zehnder modulators can handle multiple wavelengths, they require many additional components, adding complexity, size, and cost. To overcome these limitations, we have developed a transmitter for coherent optical interconnects, featuring ultra-compact, energy-efficient coherent modulators integrated directly on silicon chips.

Rethinking coherent silicon optics

Central to our innovation is the microring modulator (MRM), a tiny structure etched onto silicon chips. Although MRMs are known for their compact size and low power operation, they have historically been considered unsuitable for coherent data transmission, because they can introduce an unwanted frequency modulation (or "chirp") that gets in the way of advanced modulation. This effect also complicates modulator biasing and operation.

Previous studies, including research from Bell Labs, have proposed mitigating the chirp through a mirrored structure in which a pair of MRMs operate in counterbalance. While promising, this approach had, until now, never been demonstrated in a complete coherent transmitter employing advanced modulation formats at high speeds.

However, the rapid growth of AI has intensified the race to deploy MRMs in the demanding datacentre environment.

To develop a better understanding of MRMs, we conducted an extensive study of their dynamics [1] to determine how to fully exploit their potential. Building on this, we have now proposed an MRM-based transmitter capable of performing complex modulation. As illustrated in Figure 1, our design nests two counter-balanced MRM structures within a higher-level Mach-Zehnder interferometer configuration. Complex modulation encodes data in two dimensions – in phase and quadrature – with a dedicated counter-balanced MRM structure handling each data stream.

Our design effectively eliminates the drawbacks of MRMs to allow advanced modulation at unprecedented speeds. Harnessing the energy-efficiency and compactness of this technology, our proposed device miniaturises coherent silicon modulators and eases the challenges of co-integration with advanced electronics.

Compared to traditional travelling-wave Mach-Zehnder modulators, our approach offers significantly higher modulation energy efficiency and a much smaller footprint.

Furthermore, thanks to the natural wavelength selectivity of MRMs, the proposed architecture is inherently compatible with WDM. This feature is particularly valuable for implementing multi-wavelength transmitters driven by optical frequency combs, eliminating the need for additional multiplexing or demultiplexing components.
In rigorous laboratory tests conducted at the Centre for Optics, Photonics and Lasers (COPL), our device achieved staggering performance metrics, with data transmission rates exceeding 1T (terabits per second) at 180 Gbaud over distances of up to 80 km. This marks the fastest MRM-based transmission reported to date, achieved with exceptionally low modulation energy consumption – just 10.4 femtojoules per bit.

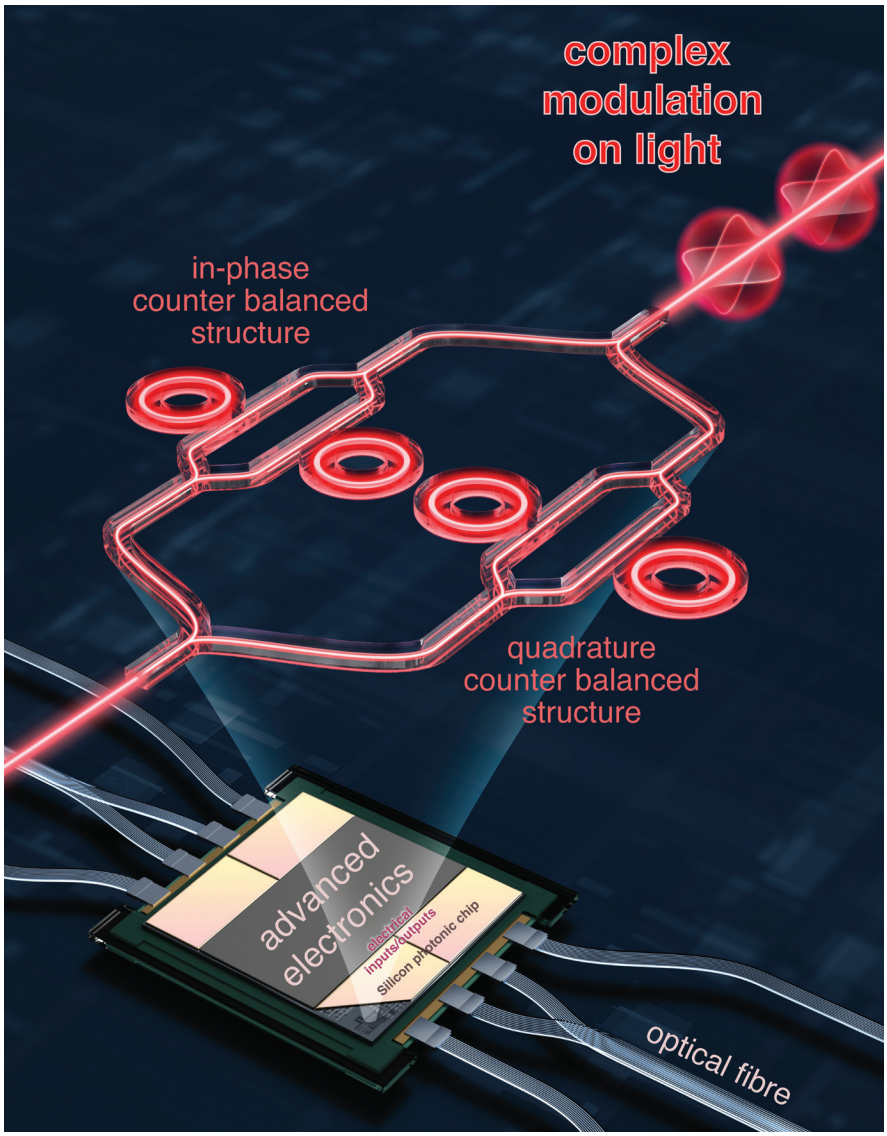

➤ Figure 1: Illustration of the ultra-compact silicon photonic modulator integrating microring modulators, modulating both the amplitude and phase of light. It enables fast connectivity for electronics through co-packaged optics technology.

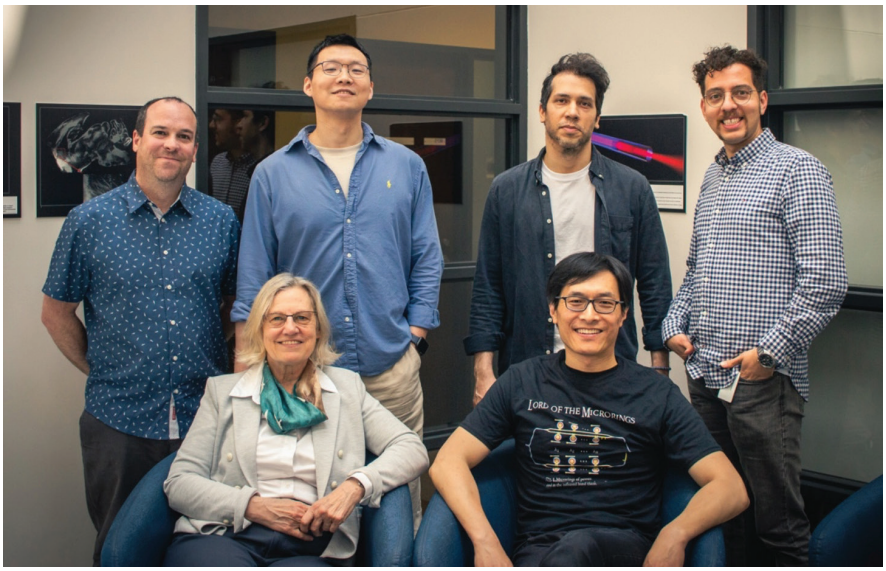

➤ Figure 2: The authors are: (standing, from left to right) Simon Levasseur, Zibo Zheng, Farshid Shateri, Alireza Geravand. (seated, from left to right) Leslie A. Rusch, Wei Shi.

Owing to its compact on-chip footprint, we also demonstrated a record shoreline bandwidth density exceeding 5T per millimetre. This metric is particularly significant, as it underscores the device's suitability for CPO systems, enabling tight integration with advanced electronic chips.

Achieving such record-breaking performance was not without challenges. Operating MRMs at their resonance wavelength while using high laser powers requires specific stabilisation methods, which have been underexplored in previous research.

Additionally, evaluating the novel device at these speeds demanded state-of-the-art laboratory setups. However, our interdisciplinary team and cutting-edge lab facilities enabled us to overcome these obstacles.

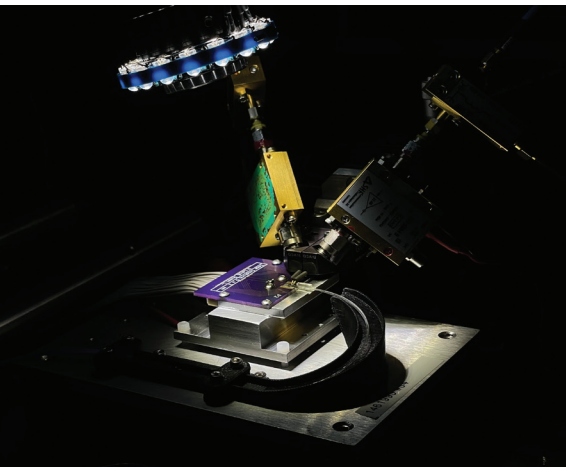

➤ Figure 3: A packaged MRM-based coherent modulator undergoes high-speed data transmission experiment at Université Laval's COPL laboratories.

Vision for the future

As well as being fascinating in its own right, our new coherent silicon photonic modulator is a significant technological advancement for AI infrastructure, since it enables processors to communicate as if they were only centimetres apart, even across large datacentres spanning tens of kilometres. Dramatically reducing the size and power consumption of coherent optics enables datacentres to expand their cluster sizes by employing coherent optical interconnects.

Furthermore, coherent optics enables dynamic network optimisation within datacentres through optical switching technology, thanks to the additional link budget provided by utilising coherent links. This flexibility allows physical networks to adapt rapidly, optimising resources based on workload demands – crucial for next-generation AI hardware. On the manufacturing side, our silicon photonic modulator technology is fully compatible with standard semiconductor fabrication methods, enabling economical mass production. However, several hurdles remain before widespread adoption will become practical. Chief among them is reducing the complexity of the digital signal processing required for coherent detection. Promising pathways include operating in the less dispersive O-band [2] and integrating analogue coherent technologies to minimise reliance on energy-intensive digital processing.

Our lab has a strong track record of advancing silicon photonics from concept to widely adopted practice. Nearly a decade ago, we published the first demonstration of high-speed multi-level modulation in silicon MRMs – overcoming early scepticism and inspiring subsequent research efforts. This technique has since been adopted in commercial AI products, including recent implementations by companies such as NVIDIA.

Looking ahead, our research opens pathways for new architectures of compact coherent interconnects. Future devices based on this breakthrough have the potential to transform how data transmission systems are constructed by facilitating the integration of coherent optical technology and electronics. Scalable, energy-efficient coherent optical interconnects could usher in an era in which vast clusters of processors, memory, and storage systems are seamlessly integrated, enhancing AI performance while reducing environmental impact.

As silicon photonics continues to mature rapidly, innovations like ours will play a pivotal role in shaping the future landscape of high-performance computing and AI infrastructure.

FURTHER READING

➤ [1] A. Geravand, Z. Zheng, F. Shateri, S. Levasseur, L. A. Rusch, and W. Shi, "Ultrafast coherent dynamics of microring modulators," Nature Photonics, vol. 19, no. 7, pp. 740–750, Jul. 2025.

➤ [2] A. Geravand, E. Weckenmann, Z. Zheng, J. Vallée, S. Levasseur, L. A. Rusch, and W. Shi, "Micrometer-Scale Silicon Modulator for O-Band Coherent Interconnects Beyond 100 GBaud," IEEE Journal of Selected Topics in Quantum Electronics, vol. 31, no. 3, pp. 1-9, May 2025.

POET

POWERING AI & THE FUTURE OF DATA CENTERS

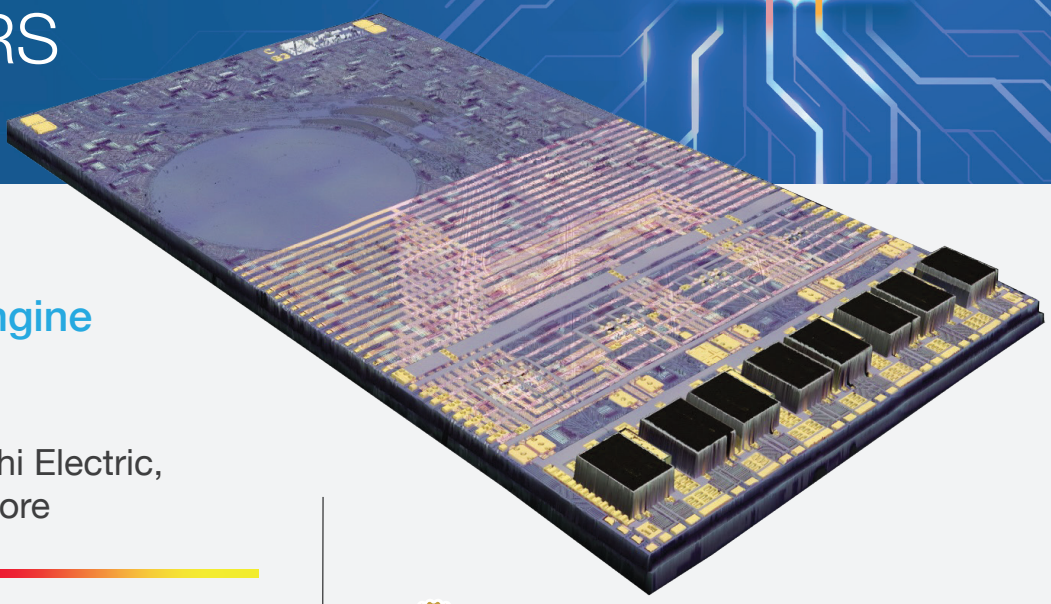

TERALIGHT
1.6T 2xFR4 Optical Engine

Partnerships with Mitsubishi Electric, Foxconn, Luxshare, and more

High-volume manufacturing with wafer-scale assembly and testing

Optical engines for 800G,1.6T and beyond for pluggable transceivers and CPO

Light sources and light engines for AI hardware systems

Low cost, energy-efficient, silicon-based hybrid integration platform

2025 Lightwave Award
Elite Score for Optical Transceivers

2024 Best in Artificial Intelligence
(Global Tech Awards)

2024 AI Innovator of the Year
Gold Prize (Merit Awards)

2024 Best Optical AI Solution
(AI Breakthrough Awards)

LEARN MORE

poet-technologies.com

sales@poet.tech

#POETpowersAI

@POETtech

@POETTechnologies

@POETtechnologies

Putting the UK at the forefront of silicon photonics

The CORNERSTONE foundry at the University of Southampton offers both academic and industrial users a flexible solution for fabricating silicon photonics devices, as well as broader start-up support to accelerate the commercialisation of the technology. An interview with Graham Reed, Professor of Silicon Photonics and Director of the Optoelectronics Research Centre at the University of Southampton, and Laura Hiscott, contributor at PIC Magazine.

LH: *How was CORNERSTONE established and what role does it play in the UK semiconductor industry?*

GR: We started CORNERSTONE in 2014 based on a grant from the Engineering and Physical Sciences Research Council (EPSRC). The reason was twofold; it was an additional income stream to the university to underpin the cleanroom, but also, because we have such a good cleanroom, we thought it would make sense to open it up to other academics in the UK. So the target audience was UK academic researchers who could use the CORNERSTONE facilities to fabricate their silicon photonics devices. That's where we began. Now we started and always have been open source, which means we don't encumber the intellectual property of the users. We don't lock them in. We offer as much flexibility as we reasonably can. And along the way we've offered training in an ad hoc, unofficial way, because we often had users who were not specialised in photonics. For example, we had somebody who specialised in micro-electro-mechanical systems (MEMS), who wanted to marry MEMS and photonics together, so we trained them in photonics.

Over the years, the word spread without us advertising it much, so we gradually started to get more industry users. Last year, we had more industry users than academic and part of the reason for that is this open-source approach and flexibility. You can't necessarily get the same platforms everywhere. Start-ups also want to be as far enough along their prototyping journey as possible in order to make an informed decision about how to go to a scale-up foundry. So small companies in particular started to use CORNERSTONE more.

As we became aware of that, we started to try and give these customers more of what they wanted. When we were going to offer this just to academics, we weren't even thinking we would

have a PDK, but industry-oriented users wanted a PDK, so we now offer that in all our platforms.

We still have academic users. They often have ideas about wanting to start a company, but don't have funding yet. But equally, there are academics who just want to do their research, so they might use a Multi Project wafer (MPW), or even use an MPW once and then do a bespoke run where they've put the funding on their grants because they want to demonstrate certain device ideas.

As silicon photonics has grown, the users have asked for more diverse platforms, so silicon nitride and then platforms for longer wavelength applications. Finally, we now offer a lithography service. That's partly because we are in this unique position of having the only deep UV scanner in UK academia, as well as exceptional new e-beam facilities.

So a variety of groups are using us, even some of the bigger companies and sometimes defence-related people. We have quite a few international users as well; at the last count, we had customers in 24 different countries. And we were also invited into EUROPRACTICE, which is a big European brokerage where you can find people like TSMC and imec, so people could use us via that route.

LH: *How does CORNERSTONE meet the differing needs of its various users?*

GR: The needs of the users change with time. Something we've committed to in recent grants is that we will try to give the community what it wants. So we've added platforms because the community wanted them, and in the future they will want other things. When we were leading up to the last grant application they told us they wanted what we're calling an online portal, a one-stop shop for silicon photonics in the UK, if you like.

So people could put their contact details on it and use it to find collaborators, partners, training etc. We haven't started that yet, but it's something we plan to do within the programme, and it's the philosophy of trying to give the community what it's asking for.

LH: *Which application areas do you think will gain the most traction with the platforms you offer?*

GR: That's a very interesting question, because if I was asked that 30 years ago, I probably would have got it completely wrong in terms of what happened. When we started doing this, we were thinking about sensor interrogation systems, not about high-speed communications at all. But it quite quickly changed. The big success of silicon photonics so far commercially has been a few companies leading the way in datacentre-based communications, and gradually silicon photonics has taken more of the market share.

Thinking of where the UK could benefit, you might think that communications might not be the biggest opportunity, because it already has big multinational leading players – the Nvidias of this world, Intel, Cisco, for example. However, the advent of co-packaged optics means there's a whole semi-revolution happening again to the communications infrastructure, the optical input/output demands and so on.

Things will probably evolve quite a lot, especially with the demands of AI and the linkage of AI chips with photonics chips. So that whole area probably has a new lease of life itself.

But then you have the emerging applications, like imaging, LiDAR, healthcare, defence, environmental sensing, and quantum. My view is that's where the biggest opportunities for the small companies in the UK are, because they don't have the corresponding dominant multinationals that they have to displace. But communications has a big lead, because it's already essentially a commercial industry versus these emerging applications. It's a really exciting time. There's plenty of stuff going on. I wouldn't want to bet my house on any individual application at the moment.

LH: *CORNERSTONE focuses on silicon photonics, but this platform has traditionally had challenges with active components. Is the foundry exploring hybrid and heterogeneous integration?*

GR: Yes, we're looking at multiple different materials to integrate with silicon. Part of this is driven by people saying that silicon modulators are running out of steam. My view is there will be an all-silicon modulator-based solution to 400Gb/s per wavelength, which is the next industry milestone, maybe a year to 18 months away.

Maybe after that the other materials will play a bigger role. But a lot of people would disagree with that and say that you need those other materials. The reason I don't say that is because we've already done 200GBaud in an all-silicon solution, which isn't yet 400Gb/s, but it implies that it might be possible.

We've already done more than 300Gb/s in a PAM-4 format, and it's not as much of a leap as it used to be to reach 400Gb/s per wavelength. In fact there are some solutions at 400Gb/s per wavelength out there, but they also include significant Digital Signal Processing (DSP), which is very power hungry, and we have been taking the approach of minimising DSP to ensure low power.

At one point speed was the only thing that mattered. Then, once we got to 50Gb/s in about 2011, people started to think that getting the energy consumption down is more important. Then it became speed again, because of increasing data demands. Now there seems to be less emphasis on energy

“We've already done more than 300Gb/s in a PAM-4 format, and it's not as much of a leap as it used to be to reach 400Gb/s per wavelength. In fact there are some solutions at 400Gb/s per wavelength out there, but they also include significant Digital Signal Processing (DSP), which is very power hungry, and we have been taking the approach of minimising DSP to ensure low power”

than there was even two years ago. However, a lot of people are working on ring resonator-based modulators, partly because they are inherently lower energy than the Mach Zehnder modulators. But they come with other problems; at an ambient temperature, they would drift, so you either put some sort of feedback loop on or you control the temperature, and both of those things require additional energy. That's not to say you can't do that, but it's a question of the best trade-offs of speed and energy. There is a type of ring resonator where we also hold a world record for data rate. But equally, there are many types of ring resonator as well.

The other thing you could do to control ring resonators is something called optical trimming. When you use silicon to make a resonant structure, because it's a high-refractive-index material, it will drift more than a low-refractive-index material for a given temperature change. Secondly, it's also more susceptible to fabrication errors. So if your waveguide is a little bit too narrow or wide (say), it won't perform at quite the targeted resonant wavelength as it would if it was spot on accurate. Those sorts of variations are even observable in 12-inch wafer fabrication at the state of the art facilities.

Therefore, some people are looking at trimming, where you correct the manufacturing defects. It doesn't change the drift with temperature. But imagine you're trying to hit a resonant wavelength here, and you manufacture it, and your resonant wavelength is actually over there, if you only rely on thermal control, you've got to bring it all the way back over here first before you stabilise it. And that means that you use a lot more energy. If I can find a way of trimming it back to here, I've only then got to use the energy to keep it there. I haven't got to use the energy to actually tune it there in the first place. So we also have a unique, non-volatile way of trimming that's come out of our

➤ Graham Reed, Professor of Silicon Photonics and Director of the Optoelectronics Research Centre at the University of Southampton

research group. That's actually quite interesting for ring-based modulators too. So we'll see which way these things end up going.

Finally, some of the folks who are using ring resonators to deal with the low energy aspect have discovered that they have to use digital signal processing to actually retrieve the signals, which is very power hungry. With the Mach Zehnder-based modulators, you are sometimes less likely to have to do DSP. Hence there is yet another trade-off. There are so many trade-offs that you have to consider what will end up being the winning technology. I think the jury's probably still out in many ways.

LH: *Last year, CORNERSTONE won funding from the UK government. How is the foundry planning to use it?*

GR: Up to now, CORNERSTONE has literally just been a foundry, and we've been making things for other people based on their designs. But as I mentioned, we've tried to be flexible along the way and do a little bit of ad hoc training to help the start-ups with whatever they're trying to do. So when we heard this call was coming, we asked: could we put more of a UK infrastructure spin on this, and bring in other partners from around the UK to deliver some of these things in a more structured way?
For example, thinking about start-up support, we are working with the Future Worlds accelerator, and there's another accelerator from Silicon Catalyst involved who run the ChipStart programme. So that means that we have people who do accelerator support for a living, rather than us trying to do it in an ad-hoc way.

We've got STFC who do electronics training. We can do some of the photonics training, and Glasgow are a part of this too, but it's better to have a specialist doing the electronics training, because pretty much for every PIC you have, you've usually got an electronics chip to interface to the real world. Therefore, STFC have become a partner to do some of that training for the photonics people.

I mentioned earlier the portal and that extends to networking. We're doing a lot more engagement with government and policy work. There are people employed on the CORNERSTONE grant who are embedded in our policy unit at Southampton who are engaging with government as part of the programme. All these activities are things we wouldn't have done before as a foundry, and the overarching aim is to speed up the journey of a start-up.

We've already seen some of our start-ups getting funding. Obviously, we can't take credit for that, but I would hope that we've accelerated their journey. The core of this is still the foundry doing what it does, but now we've got people who can also support start-ups in the other ways, who will have better contacts and better pathways already established that we can effectively take advantage of.

LH: *Last year, CORNERSTONE announced it would be taking part in*

As silicon photonics has grown, the users have asked for more diverse platforms, so silicon nitride and then platforms for longer wavelength applications. Finally, we now offer a lithography service. That's partly because we are in this unique position of having the only deep UV scanner in UK academia, as well as exceptional new e-beam facilities

PIXEurope. Could you tell me about the foundry's role in that project?

GR: It's an interesting story. The PIXEurope consortium was already formed and had been in place for the best part of two years, before the UK was once again eligible for European funding. They knew how much money was available in the call and they had allocated it to participants.

We knew a lot of the partners from previous collaborations and they were very supportive of us being involved on a technical level, but there wasn't money to give us because they had allocated it already. If they'd given us money, they would have had to take it away from somebody else who'd been contributing for two years.

But it was important for us to get in, because it's a five-year funded programme, followed by five years where it's self-sustaining. Therefore, if the UK was not involved, it would have been outside of this programme for 10 years, which would have been really detrimental to the UK photonics landscape.

Now, whatever you get from that programme, there has to be matching funding from the home country. We asked DSIT if they could, on this one occasion, give us the equivalent of the matching funding, even though we wouldn't get any from PIXEurope. The consortium said, if they did that, they would let us in. But DSIT said we can't do that, because it would set a precedent. They suggested that we try and persuade Europe to increase the amount of funding, because the UK is paying into Europe now but we're not always able to get our equivalent share back out. So we were doing the technical bit with the consortium, and DSIT was doing the politics. They did a fantastic job, and persuaded Europe to change the budget.

At that point, we had to offer the consortium something they didn't already have. We combined with Cambridge, who we're already working with, because they are one of the leaders in graphene.

There's a lot of background work on photonic applications of graphene, and Cambridge were particularly keen to put graphene-based photonic devices onto silicon photonics. So we jointly went to the consortium and said: you don't have this. Would that be something you would want? And they said yes.

Now that the funding has been approved, I'm hopeful that we will do more than just the graphene, because I think we have a lot more to offer now. We don't want to displace other partners, because we're all working as a team, but I think there will be things that we can do that other partners cannot, perhaps in some of the emerging, longer-wavelength platforms, for example. We're saying to the consortium: everything that CORNERSTONE can do is available to you if you want it. It's very early days for that programme, but importantly for the UK we're in.

LH: *As the industry continues to evolve, how do you see CORNERSTONE's role evolving with it?*

GR: I would like to see silicon photonics evolve in such a way that the UK benefits from it. We started the silicon photonics group in 1989. There was a little bit of work in the US, theoretical work from about 1986, but there wasn't very much practical work at all. So we were more or less in at the very beginning of silicon photonics. So I personally would be very pleased to see the UK benefitting from it.

We are getting feedback from commercial users, and small companies in particular, that they need a flexible prototyping silicon photonics foundry for several reasons. Firstly, they may not be able to get into a commercial foundry, because at the early stages they won't have the volume. Secondly, they might not want to get into a commercial foundry too soon, because it might encumber their IP. And thirdly, they might not know which commercial foundry to get into, because they're not far enough along the prototyping journey to know which foundry suits their technology best.

Therefore, I think what the UK needs is that type of flexible prototyping service. Now it doesn't have to be CORNERSTONE. We can separate the questions of whether we need one, and

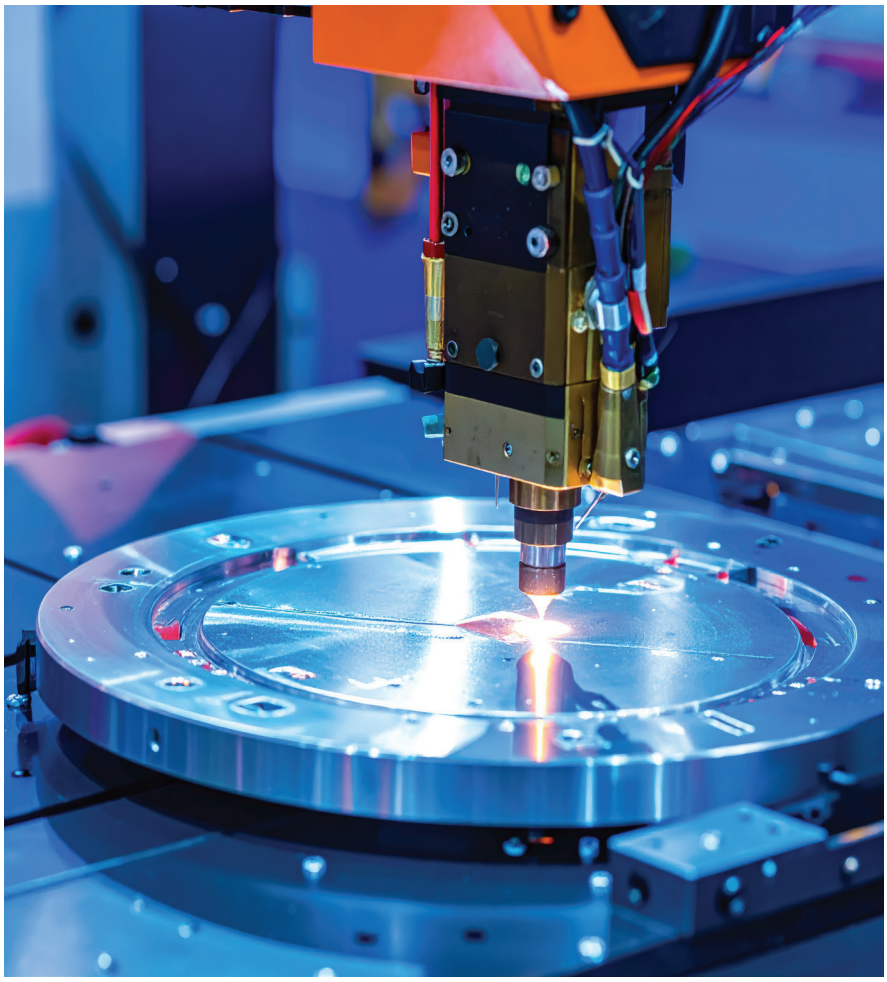

where or how it might be implemented. Even if it wasn't CORNERSTONE, it's better for the UK if there is such a service for start-ups and companies. But if that decision was taken, then CORNERSTONE might be in a good position to bid for that and offer a solution.

The motivation is that the UK does better out of silicon photonics than it will do otherwise, because, like many technologies, I think there's a real danger that this gets exploited all around the world and everywhere but the UK, even though we pioneered the technology

LH: *In theory, silicon photonics could be scaled up and become very cheap, like traditional electronic semiconductors, but we haven't reached that point yet. What are the barriers for silicon photonics companies in going to higher volumes?*

GR: For years people have talked about whether datacentre communications is a high-volume application or not. There is a growing number of datacentres, and they need increasing numbers of transceivers so it's obviously growing. But in terms of consumer products, it's not high volume, because it's very specialised and specific.

If you think about autonomous vehicles, there are already driverless taxis in some cities, and they're based in part on LiDAR. They've got big rotating things on the roof because it's a mechanical system.

Now you can implement LiDAR imaging via PICs, so instead of the system on the roof, you could deploy chips around your vehicle to give you the same, or in fact probably better, sensing as part of the driverless vehicle control system.

In terms of round figures, let's say you have 10 chips per vehicle. Most people have a vehicle, lots of people have two. That's a proper mass-market application.

In healthcare, if you had a chip that does cancer detection, you might be able to do it at home or at the doctor's surgery. You'd definitely have it on the wards of hospitals. That's potentially another mass market. And there are so many healthcare applications that you could probably measure multiple measurands with the same chip. That's going to be much higher volumes potentially.

It's a bit of a chicken and egg. For many years people were asking: what's going to be the really high-volume silicon photonics application? But datacentre communication has progressed so far now that I think everyone believes this is a commercially viable technology. Otherwise you wouldn't be seeing the proposals and implementation of co-packaged optics.

But by the same token, it gives confidence to emerging applications that the technology works and is viable. There are differences obviously, but it's no longer such a massive leap of faith that mass-market applications can be implemented in silicon photonics.

LH: *What advice would you give to researchers and start-ups working on integrated photonics?*

GR: Well obviously I would say do silicon photonics! But assuming they're already doing that, I would say aim high and be ambitious. If I look back to when I started doing silicon photonics, did I ever think we would get anywhere near 400G?

Not a chance, and I don't think anybody did. And even if I did think we could do it, I would have expected there would be competing technologies that would probably do it better. But as it turns out, silicon photonics performs much better than most of us expected. It depends on what your ambition is.

If you want to be an entrepreneur, the funding is much more likely to be available now, because the technology is now much more commercial. If you want to be a researcher, there are way more jobs available in research, in large and small companies, as well as universities.

You can work in areas of silicon photonics that interest you. In some ways I'm quite jealous of the new silicon photonics graduates these days, because they've got so much choice and opportunity. So aim high, change the world – because you can – but do something you're motivated by.

VPIphotonics
DESIGN AUTOMATION

PROFESSIONAL SIMULATION TOOLS & EXPERT DESIGN SERVICES

Integrated Devices, Waveguides & Fibers
VPIdeviceDesigner™

Photonic Integrated Circuits
VPIcomponentMaker™ Photonic Circuits

Fiber Amplifiers & Lasers
VPIcomponentMaker™ Fiber Optics

Transmission Systems
VPItransmissionMaker™ Optical Systems

Lab-proven DSP Algorithms
VPItoolkit™ DSP Library

Quantum Key Distribution
VPItoolkit™ QKD

We empower you to
define the cutting edge.

Request an evaluation

www.VPIphotonics.com

Beyond silicon: building a photonic foundation for next-gen computing

AI, quantum, and high-performance computing are outpacing the limits of traditional silicon photonics. Hybrid organic electro-optic technology offers a new foundation – one that combines manufacturability with breakthrough performance and efficiency.

BY BRAD BOOTH, CEO, NLM PHOTONICS

EVERY LEAP forward in computing eventually runs headlong into a wall. For today's AI workloads, that wall isn't just compute, but also the energy required to move data. Inside hyperscale datacentres, photonic interconnects can account for 20–40 percent of total facility power consumption. As per-lane speeds climb past 200G, the question is not just whether we can move the bits, but also whether we can afford the watts to do so.

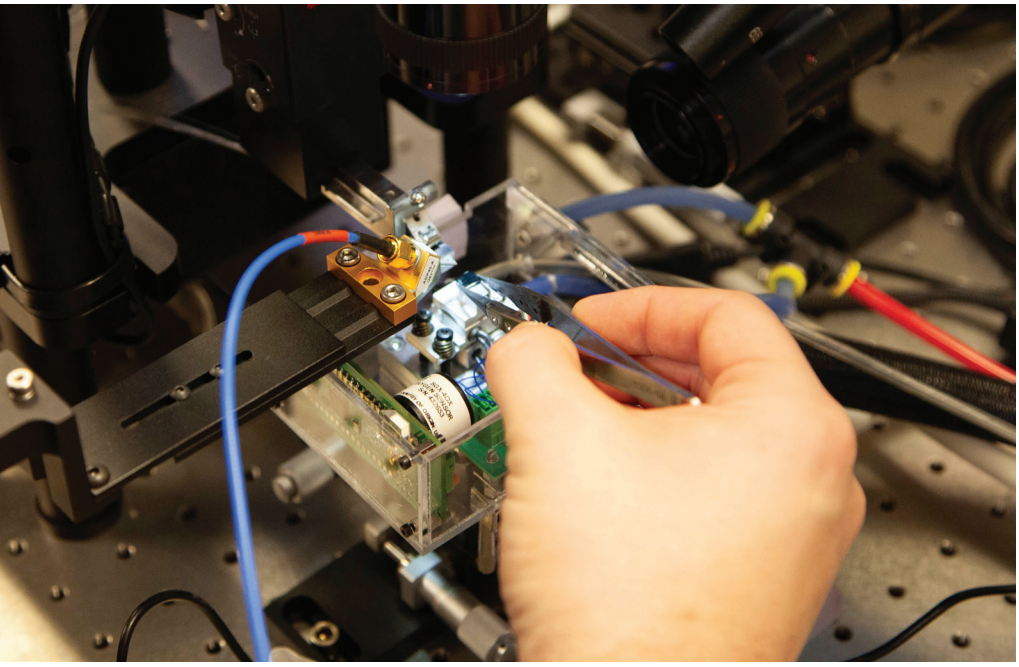

Traditional silicon photonics modulators work by shuttling charge carriers in and out of a waveguide to shift its refractive index. It's simple and CMOS-friendly, but it comes at a price: higher capacitance, higher drive voltages, and degraded signal integrity as speeds rise.

The consequence is a hard limit; whereas electronics developed for a long time according to Moore's Law, photonics does not have a similar rule. Devices get bigger instead of smaller as you push performance, consuming more power and wafer real estate and generating more heat that must be removed from the system. It's the wrong direction for an industry that needs to pack more bandwidth into tighter envelopes while cutting energy per bit.

Traditional silicon photonics modulators have carried the industry impressively far, serving brilliantly for 100G and 200G, while offering manufacturability that has allowed the photonics industry to build on a foundation of familiar fabs and processes. But physics, unfortunately, is unforgiving. At higher baud rates, the power penalties pile up, the signal quality declines, and silicon photonics' viability faces a cliff edge. Beyond 200G per wavelength, these issues can't be sidestepped by clever design alone.

Facing relentless bandwidth demand, but finite power budgets, many companies are seeking to avoid the current unsustainable trajectory by investigating alternative photonic platforms. The new high-bandwidth platforms being explored rely on the Pockels effect, which permits a lower electric field to modulate light very efficiently. It's a subtle difference with enormous impact for achieving higher bandwidths and reducing power consumption.

	Historical paradigm			Current paradigm	
Generation	**Gen 1** *1980s-1990s*	**Gen 2** *1999*	**Gen 2.5** *2007*	**Gen 3** *2017*	**Gen 4** *Future*
EO coefficient (r_{33})	< 30 pm/V	60-200 pm/V	100-500 pm/V	300-1100+ pm/V	1000+ pm/V
Composition	Guest-host	Guest-host	Neat/Binary	Neat/Binary	Neat/Binary
Device scale	Mesoscopic	Mesoscopic	Nanophotonic	Nanophotonic	Nanophotonic
Alignment method	Poling	Poling	Poling	Poling	Inherent
Orientational stability	From polymer	From polymer	From sidechains	Crosslinking (option)	Fully covalent network

➤ Generations of OEO materials, adapted from Dalton et. Al. 2023

Currently, researchers are pursuing two material approaches for new modulators: inorganic materials like lithium niobate or barium titanate, and organic materials like polymers, organic small molecules (monomers), or liquid crystals. Organic materials are introduced into the silicon photonics platform to provide high-performance modulation, replacing traditional *p-n* junction-based modulators. But not all Pockels-effect materials are created equal, especially as industry demands for performance, loss, manufacturability, reliability, size, and efficiency become more important.

Organic electro-optic (OEO) materials have emerged as a strong contender thanks to their ability to integrate with traditional silicon photonic manufacturing. This combination creates silicon organic hybrid (SOH) modulators, in which the p-n junction is replaced with a slot waveguide containing an OEO material, providing a > 10x improvement in modulation efficiency. This permits SOH modulators to become more energy efficient and dramatically smaller. Drive voltages plummet and bandwidth climbs. Thermal stability – once a stumbling block for organic materials – has also made advancements to levels compatible with commercial photonics manufacturing.

Crucially, the SOH approach is fab-friendly. The organic material is deposited in a post-fabrication process module without disrupting established front-end-of-line silicon photonics manufacturing flows.

For foundries, that means no modification to their traditional manufacturing or to their tooling. The post-processing step can either be part of a back-end-of-line step for a foundry or a pre-processing step for an outsourced semiconductor assembly and test (OSAT) firm.

Productisation

Designing OEO materials is not just about chasing the highest electro-optic (EO) coefficient. It's about balancing competing factors including wavelength of operation, responsivity, and loss. Over the past two decades, chemists have steadily improved these trade-offs through improved theoretical understanding and computationally guided design. Early-generation OEO materials were repurposed dyes blended in common polymers. Now, with the current third generation of materials, methods including chromophore engineering, tuning of sidechains for processibility and stability, and advanced poling techniques have unlocked stability and activity levels that rival or exceed inorganic materials.

At NLM, we've focused on building tuneable material families. Selerion-HTX, our first thermoset material, delivers EO coefficients of 150–450 picometres per Volt – between 5-15x that of lithium niobate. Additionally, Selerion-HTX showed for the first time that high-performance OEO materials could be crosslinked into a thermoset material with long-term stability at temperatures in excess of 120 degrees C. This combination of strong activity with exceptional thermal resilience makes HTX well suited to pluggable optics and long-lived telecom infrastructure, where devices must survive high operating temperatures, field deployment, and extended shelf life without performance drift.

Our newest material, Selerion-BHX, crosses the 1000 picometres-per-Volt threshold while maintaining long-term stability at 85 degrees C. It now holds the record for modulation efficiency in a slot-waveguide device, with a

➤ Testing performed in collaboration with Keysight demonstrating a 224G PAM-4 eye diagram on NLM's 1.6T DR8 SOH PIC

modulation efficiency of 0.038 Volt-millimetres at 1550 nm in a plasmonic-organic hybrid (POH) modulator. Selerion-BHX's record-breaking electro-optic activity enables ultra-compact modulators with very low drive voltages – ideal for co-packaged optics, optical I/O, and space-constrained systems like satellites, where footprint and energy savings are paramount.

Each material finds a careful balance of chemistry, stability, and manufacturability, allowing the different variants to be tailored to the needs of specific architectures and environments. One of the benefits of organic materials is that they can be readily modified and tuned to the performance of the application. With an intrinsic bandwidth capability in the tens of terahertz, there is a lot of performance headroom in the technology.

In a recent commercial milestone, data tests performed in collaboration with Keysight Technologies confirmed 224G PAM-4 (4-level pulse amplitude modulation) data transmission for an 8-channel silicon-organic hybrid PIC using Selerion-HTX. The tests demonstrated exceptionally low driver requirements: 0.31 Volt-millimetres on the best single channel, representing a 10-15x modulation efficiency improvement over traditional silicon photonics modulators.

These results represent real-world improvements in 200G performance and a path to 400G in a commercially available silicon photonics platform. While the rote understanding of organics in photonics was once "great on paper, unstable in practice", the story has now changed, and productization is well underway.

Impact

What does all of this mean for operators and system architects? When photonic links represent 20-40 percent of an AI datacentre's power budget and SOH devices reduce link power by 20-50 percent, total site energy use can be reduced by up to 15 percent. This substantial reduction translates into the ability to add racks and expand AI capacity without building new facilities. It's a path to sustainable growth in an era when grid availability is a significant gating factor. This means not only millions of dollars per year saved in energy costs, but also the option to defer construction of entirely new datacentres.

Any ability to rein in power consumption with each new generation of bandwidth requires photonic technology to be able to scale and eventually to be integrated tightly with the electronics.

This tighter integration will introduce greater thermal challenges in operating environments and in manufacturing. But this is a challenge that NLM Photonics is focused on solving with new generations of OEO materials, to continue supporting datacentres far into the future.

While AI is the most visible driver of alternative electro-optic material development, the benefits of OEO materials ripple across numerous other domains too.

Quantum computing systems, for instance, demand low-loss, high-responsivity links at cryogenic temperatures to maintain coherence of entangled states. Addressing this need, NLM Photonics has developed a material for a quantum customer's specific requirements.

Meanwhile, in the space and satellite sector, size, weight, and power (SWaP) considerations dominate payload economics. A smaller, lighter, more energy-efficient modulator directly translates into more capacity or reduced launch costs. NLM's Selerion family can meet these stringent SWaP constraints via either SOH or POH devices. Another potential application of NLM's

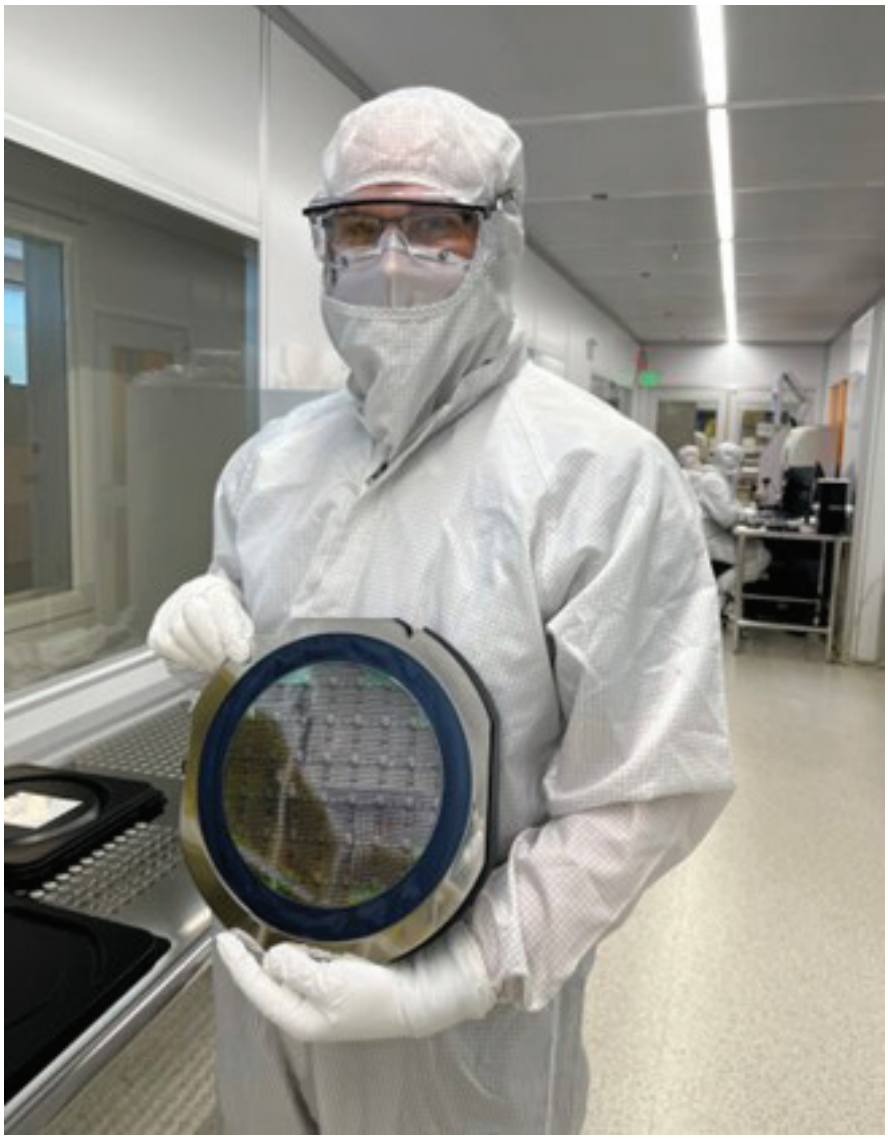

➤ Dr. Scott Hammond, Senior Director of Process Technology of NLM Photonics, in the NLM Photonics clean room with the company's wafer.

products is in telecom infrastructure, where coherent links have become the standard. Selerion-HTX has lower-loss performance in the C- and L-band, while delivering the modulation efficiency required for low driver swings with clean signals.

Finally, NLM's OEO materials permit the reduction of footprint and drive voltages, which are critical in chip-to-chip and co-packaged architectures, making them useful in optical I/O. For example, NLM's 1.6T SOH PIC is 40 percent smaller than a traditional 1.6T silicon photonics PIC, and NLM's 3.2T PIC is even smaller than the 1.6T PIC. SOH micro-rings can also be designed to be athermal, reducing thermal control needs.

Each of these domains underscores the same point: far from being just a *"datacentre play,"* stable, efficient OEO materials could become a foundational technology.

The road to adoption

When NLM Photonics launched in 2018, the idea of stable, scalable hybrid OEO modulators was still considered speculative. Seven years later, we've gone beyond supplying research materials to developing ecosystem partners, operational PICs, and our Selerion family – all driven by our R&D focus on next-generation materials, process technologies, and modulator designs.

We continue to collaborate with partners including universities, start-ups, hyperscalers, and foundries, and publish peer-reviewed papers confirming record performance and reliability. The progress has been steady, deliberate, and measurable. The next frontier isn't just higher bandwidth or lower voltage; it's reliability and manufacturability at scale. That means wafer-level processing – automating deposition, encapsulation, crosslinking, and poling so OEO can be added at the same scale as any other photonic step. It also means scaling from today's 1.6T DR8 PICs to next-generation devices capable of 3.2T and beyond, supporting 400G per lane and beyond.

For the technology to reach its full potential, we will need to work on application-specific tailoring, designing OEO materials optimised for highly technical and future-focused markets and applications. We will also need to continue our ecosystem collaboration, working with foundries, hyperscalers, and customers to ensure NLM's materials, processes, and devices align with demand for energy-efficient communication.

Ultimately, the vision is a photonic ecosystem where organics coexist with any photonic platform to reap the benefits of efficiency and scalability as bandwidth pushes forward and new markets gain traction. Together, they can unlock the headroom needed for AI, quantum, and whatever comes next.

As a small start-up, NLM Photonics' approach is to enable partners in the photonics ecosystem to adopt OEO technology and provide the benefits to their customers. Currently, we are working with our partners on the development of a Selerion-HTX ink that can be easily used by foundries. We're also working with our foundry partners to make the process technology and device designs available in their PDKs. Our goal is not to disrupt the existing supply chains, but rather to enhance them with higher-performance and lower-cost PICs.

As such, NLM Photonics' technology is designed to fit seamlessly with existing photonic platforms. Our 1.6T DR8 PIC, fabricated at Advanced Micro Foundry, used a standard commercial process flow and the post-processing step was performed at our facility. Over time, the post-processing will be transferred to partners to enable high-volume manufacturing.

We also engage with standards bodies, because interoperability and conformance are required for new technologies to succeed. Players in the photonics industry must feel that OEO fits into their existing ecosystems, from protocol to manufacturing to operational compliance. Our mission is to lower barriers to entry for our partners and their customers.

The rewards of OEO adoption can be profound: reduced power use in AI datacentres to support sustainable growth, enablement of quantum systems to scale using existing photonic platforms, and a communications future that can scale with the existing supply chain. Seven years of advancements in hybrid OEO development show that SOH modulators are the most viable path to 400G and beyond for datacentres, as well as an extremely promising foundation for a host of next-generation networking applications. The challenge ahead is scaling production, building trust, and integrating with the broader photonics ecosystem.

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

Avo Photonics
PIC Packaging Services
Tel: +1 215-441-0107
eillg@avophotonics.com
www.avophotonics.com

EPIC
EUROPEAN PHOTONICS INDUSTRY CONSORTIUM
www.epic-assoc.com

EXFO
Tel: +1 418-683-0211
francois.couny@exfo.com
www.exfo.com/en

ficontec
micro assembly machines
Tel: +49 4202 51160-435
info@ficontec.com
www.ficontec.com

IQE
http://www.iqep.com

LIGENTEC
info@ligentec.com
www.ligentec.com

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

lumerical
Tel +1(604) 733-9006
info@lumerical.com
www.lumerical.com

Optiwave
Tel:+1-613-224-4700*249
sawyer.ge@optiwave.com
www.optiwave.com

PI
Tel: +49 721 4846 1802
j.denzer@pi.de
www.physikinstrumente.com/en

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

PICMAGAZINE
To advertise here call
Shehzad Munshi
+44 (0)1923 690215
shehzad.munshi@angelbc.com

Showcase your presence to the global PIC industry for 12 months

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

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

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

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

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

KRELLTECH
ENERGIZE PHOTONICS

Advanced Photonic and Fiber Optic Polishing & Inspection Systems

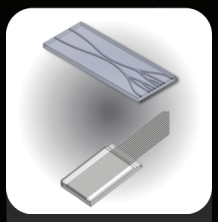
Waveguide & PIC


Bare Fibers & Arrays


Connector & Termini


Custom Probes & Devices

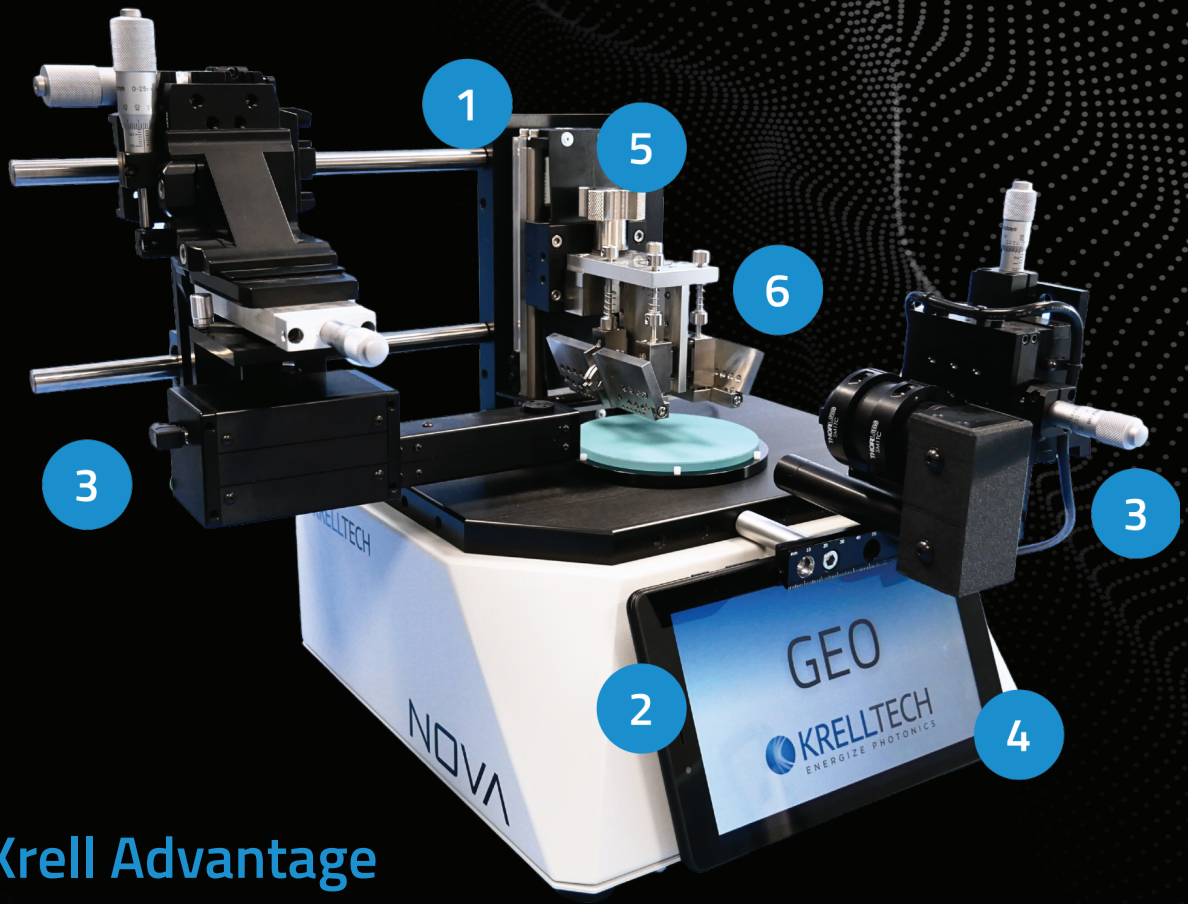

The Krell Advantage

1. Micron-level control
2. Flexible programming
3. In-line video monitoring
4. Field, lab & production configurations
5. Interchangeable fixtures
6. Exceeds industry specifications

INNOVATING PHOTONICS FOR 30 LIGHT YEARS

Have a demanding application?

LIGHTWORKS

Custom Solutions Team

www.krelltech.com | sales@krelltech.com

MADE IN USA KRELLTECH