SILICON SEMICONDUCTOR

ISSUE V 2026

AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

SILICONSEMICONDUCTOR.NET

HOW MATERIALS WILL UNLOCK THE FUTURE OF ADVANCED PACKAGING

EVG®

invent
innovate
implement

www.EVGroup.com

FUSION AND HYBRID BONDING FOR HETEROGENEOUS INTEGRATION

- Enabling advanced 3D device stacking and chiplet integration for CMOS image sensors, memory and 3D system-on-chip (SoC)
- High-volume production equipment and process technology for Wafer-to-Wafer (W2W) and Die-to-Wafer (D2W) hybrid bonding
- GEMINI® FB Automated Production Wafer Bonding System delivers industry leading alignment accuracy and bonding performance
- Heterogeneous Integration Competence Center™ serving as leading-edge innovation incubator for EVG customers and partners

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

VIEWPOINT

BY PHIL ALSOP
EDITOR

Redefining scaling in the Age of Intelligence

The semiconductor narrative has entered a moment of profound recalibration. Two recent announcements — Huawei's τ scaling vision and TSMC's "Intelligence Revolution" roadmap — illuminate a striking convergence: the industry is no longer defined by transistor shrink alone, but by how effectively time, data movement and system integration are orchestrated across the stack.

For decades, Moore's Law provided both a roadmap and a comfort blanket. Performance, cost and efficiency improvements followed predictably from geometric scaling. But as physical and economic limits tighten, the industry is being forced to confront a more complex reality — one that Huawei has embraced out of necessity and TSMC is addressing from a position of technological leadership.

Huawei's τ scaling framework is notable not simply for its technical ambition, but for its philosophical clarity. By reframing progress around time constants — latency, propagation delay and system responsiveness — the company is effectively redefining what "scaling" means. The introduction of LogicFolding, with its vertical distribution of logic to compress signal paths, exemplifies this shift. It is an architecture-first response to a physics problem, where improvements in performance are extracted not by making transistors smaller, but by making systems faster in how they behave over time.

This is not merely a workaround for constrained access to leading-edge lithography. It is a recognition that even unrestricted scaling would no longer be sufficient on its own. Interconnect delay, energy per bit and system-level inefficiencies have already overtaken transistor switching speed as dominant constraints in many workloads. Huawei's focus on τ, therefore, aligns closely with the realities of modern chip design — particularly in AI, where data movement is the true bottleneck.

At first glance, TSMC's roadmap might appear to follow a more traditional trajectory, with 2nm-class nodes, nanosheet transistors and future CFET structures extending the geometric paradigm. Yet the underlying message from its 2026 Europe Technology Symposium tells a broader story.

The foundry's emphasis on advanced packaging, photonics and wafer-scale integration underscores the same conclusion Huawei has reached: system-level innovation is now the primary engine of progress.

Technologies such as CoWoS, SoIC and System-on-Wafer are not incremental enhancements; they are fundamental enablers of AI-era computing. By integrating logic, memory and interconnect more tightly, these approaches effectively reduce "time" in a different form — shortening the distance and latency between computation and data. TSMC's work in co-packaged optics goes further still, attacking the data movement problem at the physical layer with dramatic gains in bandwidth and energy efficiency.

In this sense, Huawei's τ scaling and TSMC's system-level roadmap are less divergent paths than parallel responses to the same underlying challenge. Both recognise that performance is now governed by how quickly and efficiently information can traverse increasingly complex architectures. Whether framed as time constants or interconnect optimisation, the objective is the same: minimise delay across every layer of the system.

There is, however, a critical distinction in emphasis. Huawei is positioning itself as a conceptual disruptor, proposing a new scaling doctrine born from constraint. TSMC, by contrast, is acting as an ecosystem enabler, extending existing paradigms while integrating new dimensions of scaling. One is redefining the rules; the other is expanding the playing field.

For the broader industry, the implication is clear. The future will not be dictated by a single scaling law, but by the integration of multiple approaches — geometric, vertical and temporal. Success will depend on an ability to co-optimise devices, architectures and systems in ways that were previously unnecessary.

Contents

COVER FEATURE

How materials will unlock the future of advanced packaging

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

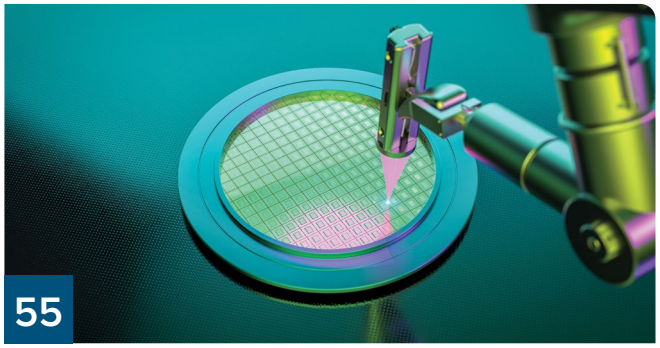

NEWS

Publisher Jackie Cannon jackie.cannon@angelbc.com +44 (0)1923 690205
Editor Phil Alsop phil.alsop@angelbc.com
Sales & Marketing Manager Shehzad Munshi shehzad.munshi@angelbc.com +44 (0)1923 690215
Director of Logistics Sharon Cowley sharon.cowley@angelbc.com +44 (0)1923 690200
Design & Production Manager Mitch Gaynor mitch.gaynor@angelbc.com +44 (0)1923 690214
Graphic Design & Multimedia Assistant Harvey Watkins harvey.watkins@angelbc.com

Chief Executive Officer Sukhi Bhadalsukhi.bhadal@angelbc.com +44 (0)2476 718970
Chief Technical Officer Scott Adams scott.adams@angelbc.com +44 (0)2476 718970
Directors Jackie Cannon, Sharon Cowley

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

Silicon Semiconductor is published ten times (6x print and 4x video issues) a year on a controlled circulation basis. Non-qualifying individuals can subscribe at: £50.00 per annum (UK) , €60.00 per annum (Europe), $90 per annum (air mail)(USA). Cover price £4.50. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher. Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd © Copyright 2026. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing. US mailing information: Silicon Semiconductor, ISSN 1096-598X, is published 5 times a year by Angel Business Communications Ltd, Unit 6, Bow Court, Fletchworth Gate, Burnsall Rd, Coventry CV5 6SP. UK. Airfreight and mailing in the USA by agent named World Container INC 150-15, 183rd St, Jamaica, NY 11413, USA. Periodicals Postage Paid at Brooklyn, NY 11256. POSTMASTER: Send address changes to Silicon Semiconductor, Air Business Ltd, c/o World Container INC 150-15, 183rd St, Jamaica, NY 11413, USA.

We strive for accuracy in all we publish; readers and contributors are encouraged to contact us if they recognise an error or omission. Once a magazine edition is published [online, in print or both], we do not update previously published articles to align old company names, branding, marketing efforts, taglines, mission statements or other promotional verbiage, images, or logos to newly created or updated names, images, typographic renderings, logos (or similar) when such references/images were accurately stated, rendered or displayed at the time of the original publication. When companies change their names or the images/text used to represent the company, we invite organizations to provide Angel Business Communications with a news release detailing their new business objectives and/or other changes that could impact how customers/prospects might recognise the company, contact the organisation, or engage with them for future commercial enterprise.
Printed by: The Manson Group. ISSN 2050-7798 (Print) ISSN 2050-7801 (Online) © Copyright 2026.

Europe launches RESOLVE to advance semiconductor innovation

Eighteen research and technology organisations have launched RESOLVE, a pan-European initiative to advance semiconductor technologies, improve energy efficiency, and strengthen industrial competitiveness.

EIGHTEEN research and technology organizations have committed to RESOLVE, a transnational initiative designed to develop the next generations of electronic components and systems and accelerate their industrial adoption across European territory. The initiative is a concrete response to the sovereignty challenges now posed by quantum technologies and semiconductors in global technological competition.

As major global powers accelerate their investments in disruptive technologies, Europe must transform its assets into industrial and sovereignty levers. RESOLVE addresses this dynamic by proposing an unprecedented framework.

It aims to position Europe as the global leader in future artificial intelligence (AI) solutions, to improve by a factor of 1,000 the energy efficiency of electronic systems by 2032, and to strengthen European semiconductor production capacity, thereby consolidating the continent's technological sovereignty. To achieve this, the project relies on transnational coordination between industry players and 18 research and technology organizations, enabling resource pooling and accelerating the transition from research to market.

Building on the foundations laid by the European Chips Act 1.0 and its pilot lines, this Pan-European initiative reflects the next step Europe now needs: turning strong R&D assets into industrial scale-up, tighter value-chain coordination, and lasting technological sovereignty. RESOLVE embodies a strong vision: making Europe a leader in energy-efficient, secure, and high-performance semiconductor technologies, committed to its industrial competitiveness and strategic independence.

A systemic ambition for Europe at the service of technological mastery and industrial impact

RESOLVE is therefore part of a political ambition for a technological Europe, structured around three priorities.

- **Sovereignty:** developing cutting-edge technologies in key areas, such as advanced memories, power electronics, radio frequencies, photonics, sub-2 nm technologies, and advanced packaging, to reduce dependence on non-European players
- **Competitiveness:** positioning Europe as a leader in manufacturing semiconductors and high-value-added electronic systems in strategic markets (automotive, aerospace, defense, industry, data centers, etc.)
- **Sustainability:** making energy efficiency a growth driver in support of Europe's climate ambitions to ensure the large-scale deployment of digital infrastructures.

Designed to strengthen European value chains

RESOLVE is structured around two inseparable pillars:

- **"R&D and technology maturation":** 15 key technology areas identified across the semiconductor and electronic systems value chain to demonstrate feasibility and foster industrial adoption.
- **"From laboratory to industry / From industry to market":** to ensure a rapid and effective transfer to European industry, including early involvement of industrial partners, alignment with product roadmaps, and co-development of prototypes.

The expected outcomes, both technical and economic, will primarily benefit companies in the strategic sectors of the European economy — data centers and AI infrastructure, automotive, defense, space, security and Industry 4.0 — by creating the conditions for new European champions to emerge.

Strengthening Europe's semiconductor industry to sustain its technological leadership

RESOLVE represents a complementary way of building and deploying critical technological capabilities and sovereign industrial foundations in Europe. By bringing together world-unique technological infrastructures into a critical mass, accessible to any European player, RESOLVE weaves a continuum from research to product commercialization, while generating concrete impact on employment, training, and the resilience of value chains, from materials to final systems.

The initiative is thus meant to consolidate and demonstrate Europe's ability to innovate at scale, combining scientific excellence, industrial coordination, and political will. Its ambition is to further strengthen Europe's semiconductor strategy building on proven assets in the areas of energy-efficient, secure, and high-performance semiconductor technologies to guarantee lasting technological leadership and sovereignty.

RESOLVE represents a complementary way of building and deploying critical technological capabilities

DTU Nanolab expands silicon DRIE infrastructure with cassette-loading etch system from Samco

DTU Nanolab is strengthening its silicon etching capabilities with the installation of Samco's RIE-800iPBC DRIE system, enhancing throughput, repeatability and process control for advanced microfabrication.

SAMCO INC. has announced that DTU Nanolab, the national nanofabrication facility at the Technical University of Denmark (DTU), will expand its silicon deep reactive ion etching (DRIE) capabilities with the installation of the RIE-800iPBC cassette-loading system. The system is scheduled for shipment in mid-2026 and will further strengthen DTU Nanolab's dry etch infrastructure for silicon-based micro- and nanofabrication.

Unlike load-lock-only configurations, the RIE-800iPBC features cassette-to-cassette wafer handling, enabling improved operational efficiency and stable process conditions across multiple wafers. The system is designed for advanced silicon DRIE using the Bosch process and supports high-aspect-ratio etching required in a wide range of research and development activities.

During the evaluation process, the system demonstrated strong performance in key process metrics including across-wafer uniformity, repeatability, precise sidewall profile control, and reliable endpoint detection. These parameters are essential for reproducible silicon micromachining in applications such as MEMS devices, microfluidic structures, silicon photonics platforms, and other high-precision microfabrication processes carried out at DTU Nanolab.

The new DRIE tool will serve a diverse user base of academic and industrial researchers who rely on stable and repeatable silicon etching performance. By expanding cassette-loading DRIE capacity, DTU Nanolab enhances both throughput and process robustness within Denmark's national nanofabrication ecosystem.

"The expansion of our silicon DRIE capability is an important development for DTU Nanolab," said Jonas Michael-Lindhard, Process Engineering at DTU Nanolab. "Uniformity, repeatability, and well-controlled sidewall profiles are critical for many of our users. The new cassette-loading system will improve workflow efficiency while maintaining the high process quality required across multiple research applications."

"We are pleased that DTU Nanolab selected the RIE-800iPBC to strengthen its silicon processing infrastructure," said Tsukasa Kawabe, President and COO of Samco Inc. "Reliable silicon etching with strong process control and repeatability is essential for advanced microfabrication. Samco remains committed to delivering plasma etching systems that combine precision, stability, and operational efficiency for research and production environments alike."

The RIE-800iPBC will be delivered alongside another Samco plasma etching system scheduled for installation at DTU Nanolab in 2026, further expanding the facility's dry etch capabilities across multiple material platforms.

Unlike load-lock-only configurations, the RIE-800iPBC features cassette-to-cassette wafer handling, enabling improved operational efficiency and stable process conditions across multiple wafers

Deep silicon etching capability to support advanced microfabrication at DTU Nanolab.

Siemens expands U.S. production as AI and data centre demand accelerates

Five-year investment programme strengthens domestic supply chains, supports AI-driven infrastructure demand, and signals long-term commitment to American industrial growth.

SIEMENS has announced it has reached $1 billion in U.S. manufacturing investments over the past five years, marking a significant milestone in its strategy to expand domestic production and support critical infrastructure sectors.

The investments, with projects coming online throughout 2026, span multiple states and industries, including industrial automation, transportation, and electrical infrastructure. The programme is designed to strengthen local manufacturing capacity, improve supply chain resilience, and create high-quality jobs across the United States.

According to the company, its U.S. manufacturing footprint now supports a network of more than 16,000 suppliers, many of them small and mid-sized businesses, underlining the broader economic impact of its localisation strategy.

A key driver behind the investment is accelerating demand for advanced infrastructure, particularly linked to AI, data centres, and electrification.

Siemens has been scaling production of electrical systems and automation technologies required to support these sectors, including recent expansions in regions such as the Carolinas aimed at meeting data centre growth.

The company also emphasised sustainability as a core component of its manufacturing expansion. Several facilities incorporate carbon-neutral operations, on-site renewable energy systems, and electrified processes, aligning with Siemens' broader target to achieve net-zero carbon emissions by 2030.

Beyond infrastructure and energy systems, the investment reflects a wider industrial shift toward regionalised production. As geopolitical pressures and supply chain disruptions reshape global manufacturing, companies are increasingly prioritising local capacity to ensure reliability and responsiveness to demand.

For Siemens, the U.S. remains a critical growth market. The $1 billion milestone signals not only a scaling of physical production, but also a deeper alignment with long-term trends in electrification, digitalisation, and intelligent infrastructure.

With projects continuing to ramp through 2026 and beyond, the focus now shifts to execution and workforce development, as manufacturers race to meet demand from AI-scale digital infrastructure and next-generation energy systems.

Semiconductor EUV photoresist market to hit $10.8 billion by 2033

Rising adoption of advanced process nodes, AI-driven chip demand and next-generation lithography is set to drive rapid growth in the EUV materials sector.

ACCORDING to latest research by Growth Market Reports, the global Semiconductor EUV Photoresist market size reached USD 1.32 billion in 2024, and the market is expected to grow at a robust CAGR of 23.7% over the forecast period.

By 2033, the market is projected to reach USD 10.8 billion, driven by the relentless demand for advanced semiconductor manufacturing and the transition to sub-7nm nodes. The primary factor fueling this growth is the increasing adoption of extreme ultraviolet (EUV) lithography in the production of next-generation integrated circuits and memory devices, enabling higher transistor density and improved performance.

The semiconductor EUV (Extreme Ultraviolet) photoresist market is emerging as a critical segment within the global semiconductor manufacturing ecosystem. As chipmakers transition to advanced process nodes below 7nm and even 3nm, EUV lithography has become essential for achieving higher precision and transistor density. EUV photoresists play a vital role in enabling this next-generation patterning technology.

Market drivers

One of the primary drivers of the EUV photoresist market is the accelerated adoption of advanced semiconductor nodes. Leading foundries and integrated device manufacturers are investing heavily in EUV lithography systems to improve chip efficiency and reduce power consumption.

Another key factor is the explosive growth of artificial intelligence and machine learning applications. These technologies require powerful processors, which in turn depend on highly advanced semiconductor fabrication processes. EUV photoresists are indispensable in achieving the precision required at such nanoscale levels.

Innovation in EUV photoresist chemistry is playing a crucial role in shaping the market. Traditional chemically amplified resists are being enhanced to improve sensitivity, line edge roughness, and resolution

Additionally, the increasing demand for smartphones, data centers, and electric vehicles is pushing semiconductor manufacturers to adopt more sophisticated lithography techniques, further supporting market growth.

Technological advancements in EUV materials

Innovation in EUV photoresist chemistry is playing a crucial role in shaping the market. Traditional chemically amplified resists are being enhanced to improve sensitivity, line edge roughness, and resolution.

Researchers are also focusing on metal oxide-based resists, which offer improved absorption of EUV light and better pattern fidelity. These advancements are helping overcome challenges such as stochastic defects and photon shot noise, which have historically limited EUV adoption.

Furthermore, collaboration between semiconductor equipment manufacturers and chemical companies is accelerating the development of next-generation resist materials tailored for sub-5nm nodes.

Challenges in the market

Despite strong growth potential, the EUV photoresist market faces several challenges. One of the major issues is the high cost of EUV lithography systems and associated materials, which limits adoption to only leading semiconductor manufacturers.

Another challenge is the technical complexity of developing resists that meet stringent performance requirements at extremely small geometries. Issues such as defect control, sensitivity balance, and resolution limits continue to pose hurdles for material developers.

Supply chain constraints for high-purity raw materials also impact production scalability, making consistent availability a concern for manufacturers.

The semiconductor EUV photoresist market is expected to experience sustained growth as the industry moves toward sub-3nm and even more advanced nodes. As demand for faster, smaller, and more efficient chips continues to rise, EUV lithography will become increasingly central to semiconductor manufacturing.

In the coming years, breakthroughs in material science and lithography precision are expected to unlock new opportunities, making EUV photoresists one of the most strategically important materials in the semiconductor value chain.

Global Electronics Association launches Global Electronics Policy Council

New body will unite electronics leaders to coordinate advocacy on trade, manufacturing, workforce and technology policy.

THE GLOBAL ELECTRONICS ASSOCIATION has formed the Global Electronics Policy Council (GEPC), a new body uniting leading electronics companies from around the world to advance a coordinated policy agenda across every major region of the electronics supply chain.

The GEPC builds on the Association's longstanding electronics industry advocacy work over its nearly 70-year history. For the first time, companies spanning the full electronics value chain - from PCB manufacturers and EMS providers to OEMs, semiconductor suppliers, wire harness, and advanced packaging firms - will have a single, structured forum to translate industry consensus into coordinated government engagement.

GEPC launches against a backdrop of escalating policy pressure on the electronics industry, including tariff volatility, export controls, and competing domestic investment mandates. The Global Electronics Association's own trade flows research underscores the urgency: global electronics trade totaled $4.5 trillion in 2023, with supply chains more globally interdependent than any other industry. That interdependence makes coordinated advocacy not just useful, but essential.

"No single company or country can navigate this environment alone," said Thomas Cetta, Senior Vice President, Jabil and chairperson of the GEPC. "The Global Electronics Policy Council gives the industry the structure and discipline to speak with one voice on the challenges that matter most and to engage governments with the credibility and accountability that comes from real organizational commitment."

A council built for action, not just alignment

Unlike informal industry coalitions, the Global Electronics Policy Council is governed by formal bylaws, a defined leadership structure, and regional execution arms spanning North America, Europe, East Asia, and India/ Southeast Asia.

The Council will produce an approved global policy agenda and annual advocacy plan, issue formal policy positions and testimony, and deliver quarterly reporting on government engagement activity.

Inaugural members include AT&S, Flex, Jabil, Plexus, TSMC, and TTM Technologies representing a deliberately balanced cross-section of the electronics value chain.

A clear and ambitious policy framework

The Council's advocacy will be organized around five priorities drawn from the Global Electronics Association's 2026 Policy Agenda: safeguarding predictable access to global markets; investing in domestic manufacturing capacity and capability; building robust workforce pipelines for the electronics industry; supporting industry-led technical and sustainability standards while rightsizing regulation; and accelerating technology leadership through collaborative R&D. Regional councils will execute against this global framework.

"The GEPC reinforces an essential aspect of our industry: a strong, connected global electronics manufacturing community," said Chris Mitchell, VP Global Government Relations, Global Electronics Association. "This Council will advance a policy agenda that strengthens supply-chain resilience, accelerates innovation, and secures trusted access to global markets for the 3,200+ member companies."

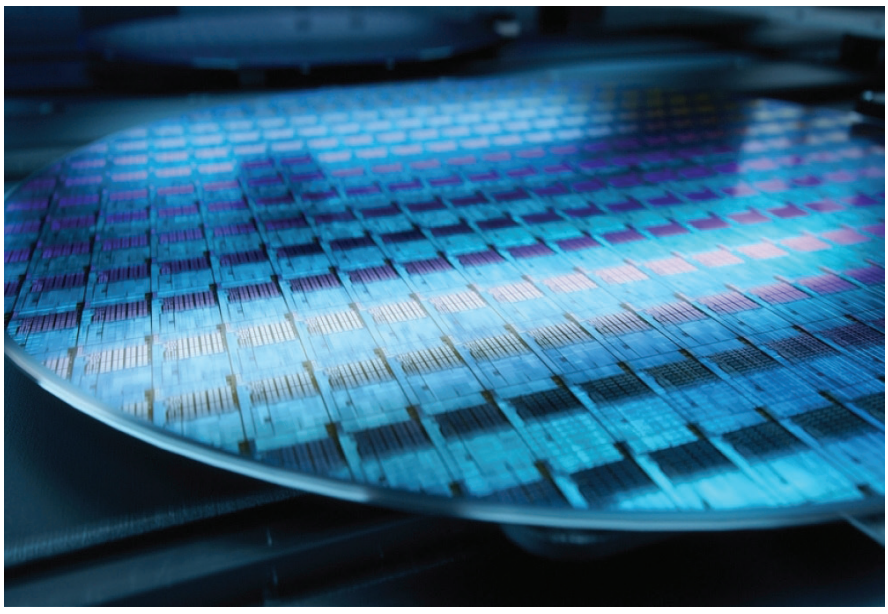

How materials will unlock the future of advanced packaging

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

BY JON KEMP, CEO, QNITY

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

This move from "shrink" to "shrink and stack" is forcing the entire ecosystem to think differently. Advanced packaging is the strategic enabler that makes this shift possible — the link that lets manufacturers push transistor architecture and density while holding the line on performance, efficiency, and reliability. When architectures go vertical, you can't bolt on the materials, interconnects, and packaging at the end. They have to be designed in from the very first step.

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

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

But those advances remain within the bounds of traditional scaling.

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

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

The angstrom era

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

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

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

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

That makes thermal management critical. A trillion transistors on a single chip system demands enormous energy and throws off unprecedented heat. Power delivery is already a major challenge. But thermal management may be the bigger test — because if you can't remove heat efficiently, you don't have a high-performing, reliable system. You have a limit.

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

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

Advanced packaging as the invisible infrastructure

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

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

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

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

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

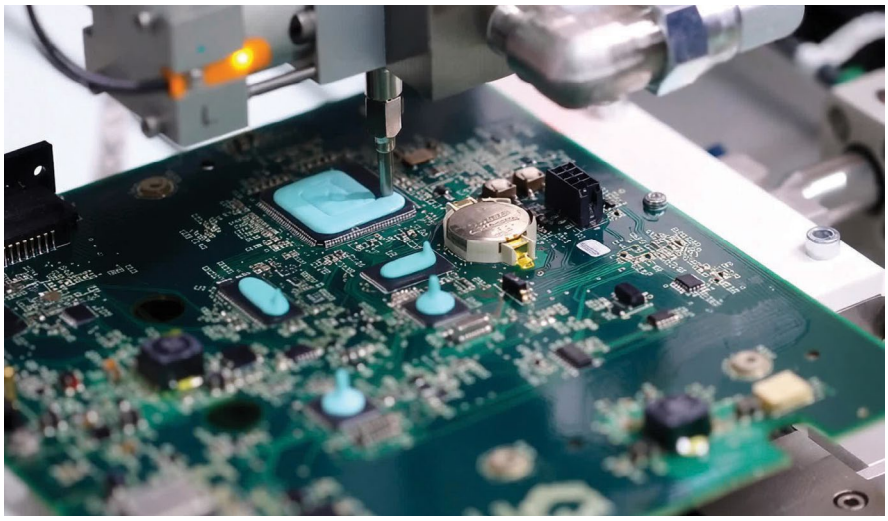

these solutions keep data moving fast even across the most densely packed architectures.

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

Qnity's approach to advanced packaging

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

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

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

By combining materials expertise, process knowledge, and system-level integration, Qnity's innovations let technologies like 3D stacking, chiplet integration, and heterogeneous integration operate at scale. That's how manufacturers get past the real pain points — chip-to-chip communication, interconnect efficiency, thermal management, and high-density scaling.

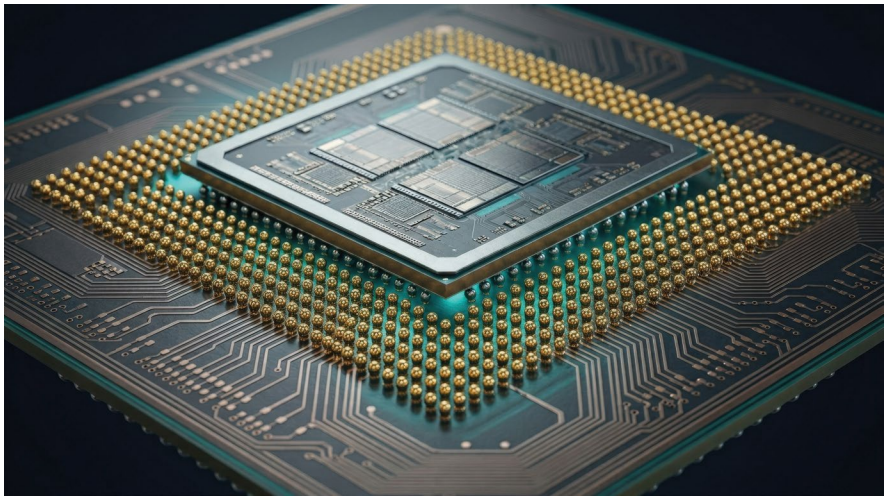

The future of advanced packaging

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

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

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

Our technology solutions enable effective communication, manage thermal loads, and support heterogeneous integration — helping manufacturers realize the full performance potential of transistor systems in the trillions, while preparing for what lies beyond 2030.

The stack era begins

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

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

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

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

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

SEMI Foundation and the U.S. National Science Foundation launch first four regional nodes

THE SEMI Foundation, serving as the Hub Operator for the National Network for Microelectronics Education (NNME), has launched the first four Regional Nodes of the NNME, marking a major national milestone in the effort to build America's microelectronics and semiconductor workforce at unprecedented scale.

Funded by the U.S. National Science Foundation Directorate for Technology, Innovation and Partnerships (NSF TIP) in partnership with the U.S. Department of Commerce and aligned with the CHIPS and Science Act, the NNME is designed to serve as national infrastructure for microelectronics talent development across the United States.

The activation of the first Regional Nodes marks the launch of a connected national ecosystem designed to align industry demand with education, training, and career pathways that prepare new learners for high-demand careers powering the future economy. The four Regional Nodes are:

- **NNME Southwest**, led by the Arizona Commerce Authority, serving Arizona, Southern California, Colorado, New Mexico, and Utah;
- **NNME Pacific Intermountain**, led by Boise State University, serving Idaho, Washington, Oregon, Montana, Utah, Colorado, Nevada, Northern California, and Hawaii;
- **NNME Northeast**, led by NY Creates, serving Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Pennsylvania, Rhode Island, Vermont, and Virginia; and
- **NNME South**, led by the University of Texas at Austin, and serving Arkansas, Mississippi, Georgia, Florida, Texas, New Mexico, Louisiana, Oklahoma, Utah, and Alabama.

Collectively, the four NNME Regional Nodes activate a national network of more than 325 performer organizations, including K-12 school districts, colleges and universities, workforce development organizations, economic development agencies, community-based organizations, and semiconductor employers working together to build stronger pathways into microelectronics careers for all Americans.

"America's ability to lead in semiconductors depends on whether we can build and sustain the workforce needed to power innovation here at home," said Senator Todd Young. "The National Network for Microelectronics Education represents the kind of national, industry-connected workforce strategy needed to strengthen America's competitiveness, expand opportunity, and prepare the next generation of talent for high-demand careers in microelectronics and advanced manufacturing."

Under the NNME structure, the SEMI Foundation and NSF expect to support the four Regional Nodes with potential funding opportunities of up to $20 million per node over five years to accelerate scalable talent solutions aligned to industry demand and establish the gold standard for microelectronics education and workforce training across America.

Building the workforce required for america's semiconductor expansion

"The launch of these Regional Nodes constitutes the activation of national

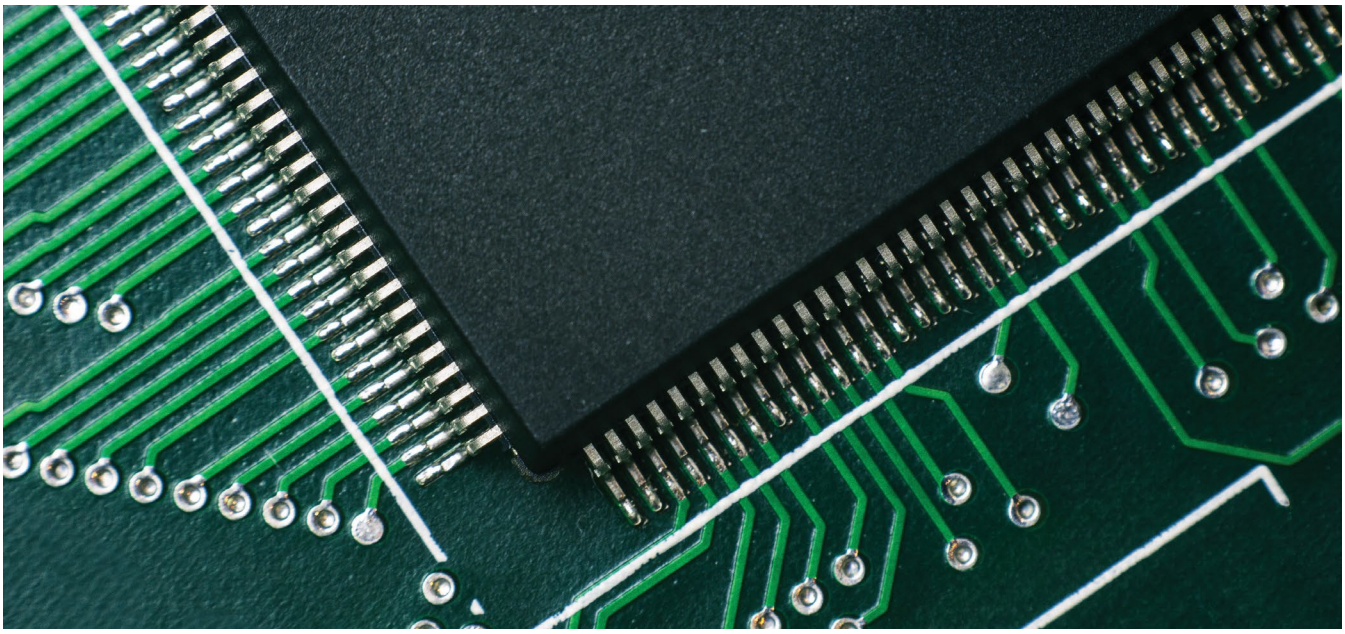

infrastructure built to meet the most consequential economic and technological challenge of our time," said Shari Liss, Vice President of Workforce Development and Initiatives at SEMI and the SEMI Foundation.

"NNME is bringing together industry, educators, workforce organizations, and regional leaders to build a talent engine capable of preparing the next generation of semiconductor workers," said Jennifer Ellis, Director of the NNME.

"This is about creating opportunity, accelerating innovation, and ensuring all Americans can participate in the future of the microelectronics economy," Michelle Williams, Executive Director of the SEMI Foundation, added.

The announcement comes as the semiconductor industry experiences historic expansion across the country, driven by advanced manufacturing investments, supply chain growth, and increasing demand for skilled talent. A recent national landscape analysis conducted by the SEMI Foundation in collaboration with McKinsey & Company reinforces the scale and urgency of that challenge. By 2030, the United States is projected to face a shortfall of approximately 127,000 to 157,000 semiconductor and microelectronics workers.

The NNME was established to help close that gap through a nationally aligned, regionally activated strategy focused on expanding awareness, accelerating workforce readiness, modernizing education and training systems, and connecting learners directly to employment opportunities across the semiconductor ecosystem.

"America's leadership in semiconductors depends on our ability to develop our talent," said Erwin Gianchandani, NSF Assistant Director for Technology, Innovation and Partnerships. "The NNME represents an important investment in America's innovation capacity, talent readiness, and long-term competitiveness. These Regional Nodes will strengthen regional ecosystems to prepare the next generation of America's semiconductor design and manufacturing workforce."

Connecting industry, education, and workforce systems at national scale

The NNME Hub will work closely with employers and the NNME Regional Nodes to translate national workforce standards, technical competencies, and industry needs into scalable education and training programs tailored to regional demand.

Industry partners across the network will play a central role in shaping curriculum, validating skills, supporting work-based learning, and connecting learners directly to careers throughout the semiconductor and microelectronics ecosystem.

"Micron is proud to support the NNME and excited about what this network means for the future of American workers," said April Arnzen, EVP and Chief People Officer at Micron Technology. "Building the semiconductor workforce our country needs isn't something any one company or institution can do alone. It takes intentional collaboration across industry, education, and workforce systems. NNME is exactly that kind of partnership: a connected, national effort that creates real pathways for more Americans to access meaningful, well-paying jobs in one of the world's most important industries."

The NNME will support workforce development across the full semiconductor ecosystem, including semiconductor manufacturing, advanced packaging, facilities operations, equipment maintenance, integrated circuit design, materials, and emerging technologies.

As the NNME continues to expand nationwide, the Hub and its partners will work to strengthen shared standards, scale innovative workforce models, and build a resilient talent pipeline capable of supporting America's long-term economic growth and technological leadership.

Global semiconductor equipment billings increase 14%

SEMI has announced in its Worldwide Semiconductor Equipment Market Statistics (WWSEMS) report that global semiconductor equipment billings increased 14% year-over-year to US$36.55 billion in the first quarter of 2026. First quarter 2026 billings registered a 1% quarter-over-quarter growth.

Record quarterly billings were driven by continued AI-related investment, including capacity expansion and technology upgrades supporting leading-edge logic, DRAM, and advanced packaging.

"The strong start to 2026 reflects continued industry investment in the capacity and infrastructure needed to support AI-driven semiconductor growth," said Ajit Manocha, SEMI President and CEO. "Record first-quarter billings highlight ongoing momentum in leading-edge manufacturing and advanced packaging."

Compiled from data submitted by members of SEMI and the Semiconductor Equipment Association of Japan (SEAJ), the WWSEMS report is a summary of the monthly billings figures for the global semiconductor equipment industry.

The following are quarterly billings data in billions of U.S. dollars, with quarter-over-quarter and year-over-year changes by region:

European Commission presents Chips Act 2.0 at SEMI Europe Policy Forum

The European Commission has officially presented the new European Chips Act 2.0, marking a major step forward in Europe's ambition to strengthen its semiconductor competitiveness, resilience and technological sovereignty. Following the announcement, officials presented the new policy at the SEMI Europe Policy Forum in Brussels, a high-level gathering hosted by SEMI Europe bringing together leading semiconductor industry representatives, policymakers and key stakeholders from national governments.

SEMI Europe welcomes the revised framework as an important signal of the European Union's continued commitment to supporting the semiconductor industry and securing Europe's position in the global technology landscape. Building on the foundation established by the first European Chips Act, the updated framework aims to reinforce the entire semiconductor value chain and address the structural challenges

that could impact Europe's long-term competitiveness.

"We welcome the Chips Act 2.0's ambition to expand the 'first-of-a-kind' concept, accelerate the industrial uptake of R&D, streamline permitting procedures, and stimulate demand across the semiconductor value chain," said Laith Altimime, President of SEMI Europe. "These are critical measures to strengthen Europe's semiconductor ecosystem, enhance technological sovereignty, and build a more resilient and globally competitive European supply chain."

Coinciding with the Chips Act 2.0 publication, the Policy Forum once again demonstrates SEMI Europe's role as a leading platform for dialogue between industry and policymakers at a critical moment for Europe's semiconductor ambitions.

"Semiconductors are a global industry, and SEMI remains committed to facilitating international collaboration essential to strengthening resilient and competitive supply chains," said Ajit Manocha, President and CEO of SEMI. "We applaud government collaborations with industry such as the European Chips Act 2.0 and investments that build on regional strengths to accelerate innovation."

In this regard, the SEMI Europe Policy Forum highlighted the importance of coordinated action between governments, industry and global partners to support Europe's semiconductor ambitions. The updated Chips Act 2.0 framework reflects several priorities consistently advocated by SEMI and the broader industry, reinforcing Europe's commitment to innovation, manufacturing and long-term competitiveness across the semiconductor ecosystem.

SEMI and Global Net Corp. release Glass Core Substrate market report

SEMI has announced the availability of a new industry research report focused on the glass core substrate market and trends, produced in collaboration with Global Net Corp. (GNC). The Glass Core Substrate Market and Development Trends Report examines the emerging market for glass core substrates, a potential next-generation packaging technology attracting increased attention as AI and high-performance computing (HPC) drive demand for larger and more advanced semiconductor packages.

"As chipmakers look for new ways to improve system performance beyond traditional device scaling, advanced packaging has become a critical area of innovation," said Clark Tseng, Senior Director of Market Intelligence, SEMI. "Glass core substrates are being evaluated as one possible solution for future high-end packages because they may help support larger package sizes, finer interconnects, and improved dimensional stability compared to conventional package-substrate materials. This new report helps industry stakeholders understand where glass core substrates may fit in the next phase of packaging technology and what challenges must be addressed to achieve broader adoption."

The report highlights increasing industry activity and investment around glass core substrates as companies prepare for the next phase of advanced packaging—providing an independent view of the technology's market potential, development status, key players, and remaining barriers to commercialization. Under the report's market-development scenarios, initial production could begin around 2028 in selected high-performance applications, with adoption expanding over time across larger and more complex package architectures. Based on the average of its Positive, Base,
and Negative scenarios, the report projects a compound annual growth rate (CAGR) of 67.2% from 2028 to 2040. It also maps an increasingly active global development ecosystem, with companies and research organizations across Asia, North America and Europe advancing glass core substrate technologies.

Key topics covered in the report include:

- Market outlook and adoption scenarios for glass core substrates
- Technology drivers and barriers to commercialization
- Expected applications in AI, HPC, advanced processors, co-packaged optics, and image sensors
- Company activities across substrates, glass materials, equipment, processing, inspection, and related supply-chain segments
- Development hubs, consortia, and supply-chain structure
- Forecast scenarios through 2040

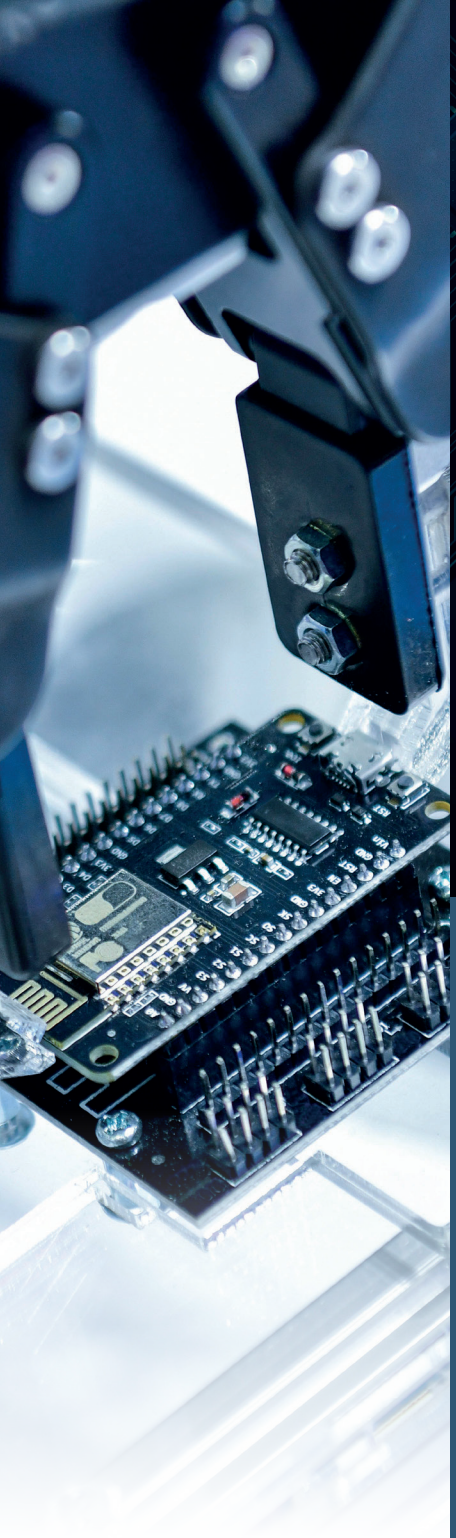

Upgrade to measure trace oxygen and moisture before contamination spreads

GEN-7 SERIES

The Gen-7 SERVOPRO DF-700 series

delivers ultra-trace oxygen and moisture analysis:

- Easy, reliable digital operation with larger screen and buttons
- Measure moisture down to 55ppt with TDLAS technology
- Future-ready with long-life solid-state drive and improved filter

Earlier insight.
Faster action.
Higher purity.

Ready to upgrade?
Find out more at:
servomex.com/gen7

Analysis that **empowers**

SERVOMEX
a spectris company

Money builds fabs. People run them.

The fabs are coming. The tools are being installed. The question the industry — and policymakers — need to answer with the same urgency they brought to capital investment is this: who, exactly, is going to run them?

BY TINA RUPP, NATIONAL LEADER, SEMICONDUCTOR – ACTALENT

THE CHIPS and Science Act has done something remarkable: by unlocking billions of dollars in federal support, it has made it economically viable for companies to build semiconductor fabs on American soil. With up to 35% total federal support available when combining tax credits and discretionary CHIPS Act funding for qualifying projects, the financial math that once made domestic manufacturing a non-starter has fundamentally changed. There are nearly one hundred manufacturing projects now underway across the United States.

And yet, standing in the middle of this historic industrial revival, I keep hearing the same sentence from clients: “I need more people.”

Every time I hear it, I ask the same question back: “Why? What will they be doing? When do you actually need them to be productive in their role?” The pause that follows tells me everything. The CHIPS Act solved the financial capital problem. The human capital problem is another story.

Facility development outpaces talent

Federal investment can fund construction, tools and equipment. It cannot manufacture the engineers and technicians needed to operate those tools, and that distinction is now coming into sharp relief as facilities approach operational readiness.

There has long been a talent gap problem in semiconductor manufacturing. What the CHIPS Act has done is magnified that existing problem by increasing the scale and urgency of manufacturing needs. Companies that offshore their manufacturing didn’t just move production — they moved the knowledge base, the training pipelines, the career pathways and the institutional memory that takes decades to build. The CHIPS Act provides funding to reassemble the physical infrastructure in a matter of years. It provides no equivalent mechanism for rebuilding the human infrastructure.

The result is a collision that I see playing out in real time with clients across the country: facilities coming online ahead of the workforce capable of running them.

The rude awakening already happening

There is a geographic reality that rarely enters the policy conversation. Semiconductor fabs require enormous amounts of land. Large tracts of affordable land tend to exist in areas with limited existing infrastructure — fewer housing options, fewer services, fewer residents who might fill those estimated 2,000-person headcount growth targets. Companies are not just competing for scarce technical talent nationally; they are asking that talent relocate to communities still being built around the facilities themselves.

Layer onto that a deeper problem: many companies reshoring to the United States are expecting to find the same depth of semiconductor-specific expertise they’ve come to rely on in Taiwan, South Korea, Japan and Europe. They are discovering that expectation was naïve. Those ecosystems were cultivated over generations. The skilled technician workforce, the process

engineering bench strength, the institutional knowledge embedded in experienced teams — none of that exists at scale in the United States right now, and it cannot be conjured by a tax credit.

Workforce strategy is not simply an HR function

The companies I work with that are navigating this well share a common trait: they don't treat talent as a number. They treat it as a sequence of institutional bottlenecks to be solved.

The most effective approach I've seen breaks overwhelming headcount targets — those 2,000-person ramps that paralyze talent acquisition teams — into phased, time-bound objectives tied directly to facility milestones. Rather than planning against abstract future numbers, the focus shifts to near-term operational triggers: what needs to be accomplished, by whom, and by when.

That last point is one of the most consistently overlooked in semiconductor workforce planning. Bringing someone on 30 to 60 days before the triggering event — the tool install, the cleanroom qualification, the production ramp — makes the difference between a contributor and a liability on day one. Companies that wait until the last moment to hire are competing in a crowded market for the same candidates, paying premium rates and still getting people who need months of ramp time.

The harder conversation, and the one more clients need to have, is about the root cause behind the hiring request. When a manager says, "I need an engineer," the real question is: what business bottleneck does that engineer resolve? What work is not getting done today and what is the consequence of it staying undone? Companies that approach talent with that level of operational specificity make better decisions, move faster and get more from their investment in workforce development.

What a real strategy actually requires

Upskilling programs matter. Community college partnerships matter. Veterans' pathways and apprenticeship models matter. But none of these are sufficient on their own, and none of them operate on a timeline compatible with facilities that are opening now.

A national talent strategy for semiconductor manufacturing has to acknowledge several things simultaneously: pipeline gap is structural and will take years to close; near-term workforce needs require creative solutions including selective high-skilled immigration pathways; offshore and nearshore staffing models can play a legitimate role for non-core functions, freeing domestic talent for the roles that require it most; and companies which treat workforce development as a strategic priority (versus a simple matter of procurement) will be the ones that hit their ramp targets.

Companies currently trying to manage 2,000-person ramps with small internal recruitment teams are going to feel the cost of that choice. The semiconductor industry is margin-obsessed for good reason: missing targets in this environment is existential. The companies that recognize they need strategic talent partners — not just staffing vendors — and make those partnerships early, are building a structural advantage over those still operating as if traditional hiring approaches will scale.

The limiting factor was never the fab

The CHIPS Act proved that domestic semiconductor manufacturing is economically viable again.

The fabs are coming. The tools are being installed. The question the industry — and policymakers — need to answer with the same urgency they brought to capital investment is this: who, exactly, is going to run them?

It is the semiconductor industry's most important strategic challenge right now. And the clock is already running.

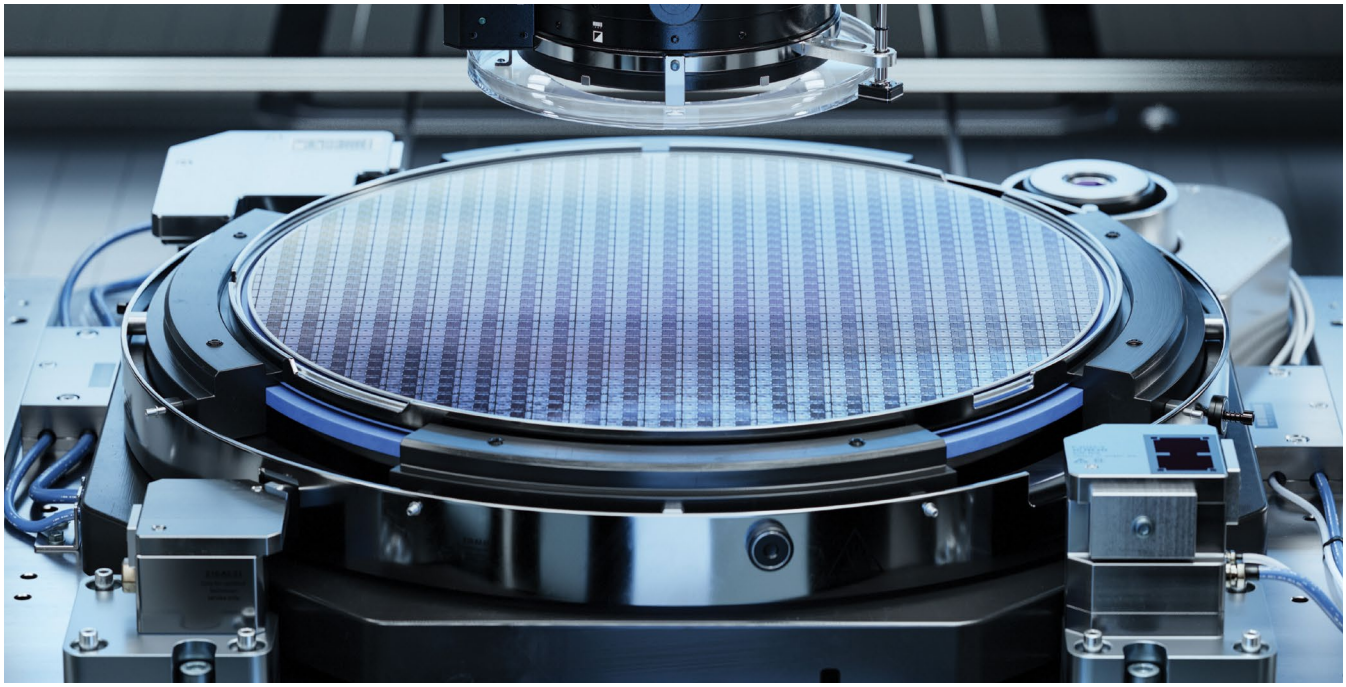

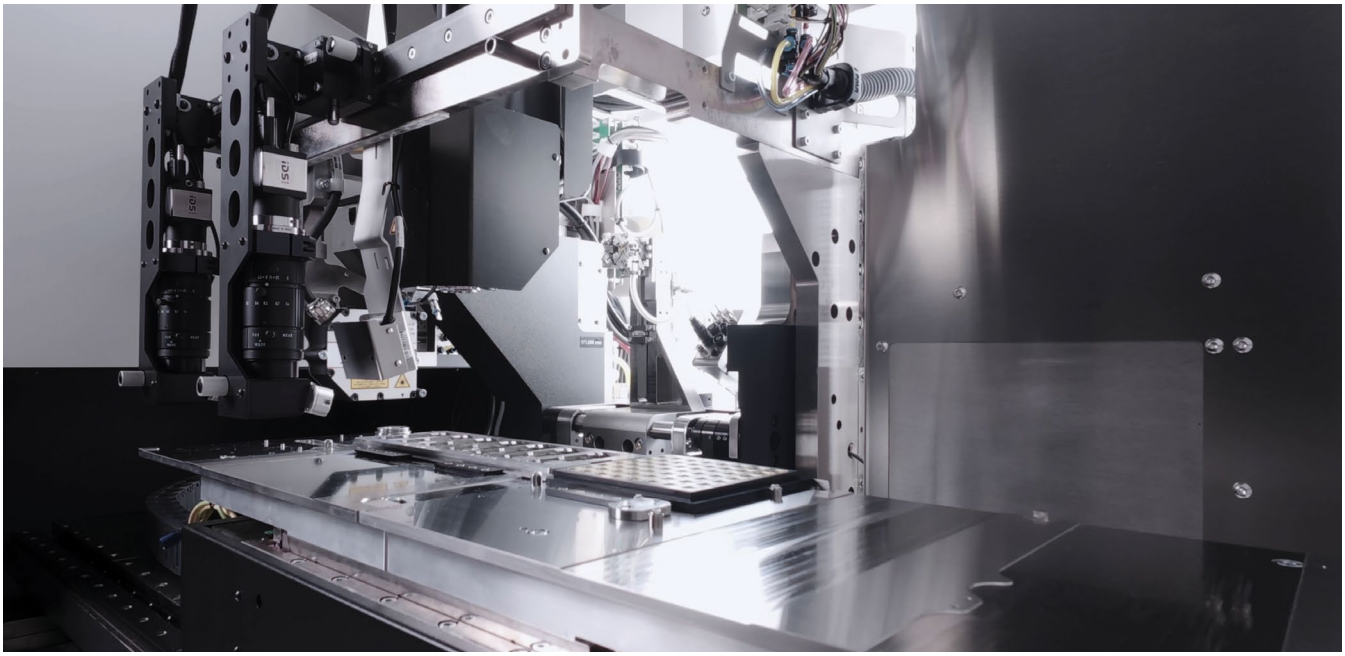

Silicon photonics needs a new packaging strategy to reach volume

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

BY SYLVAIN DULPHY, FINETECH

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

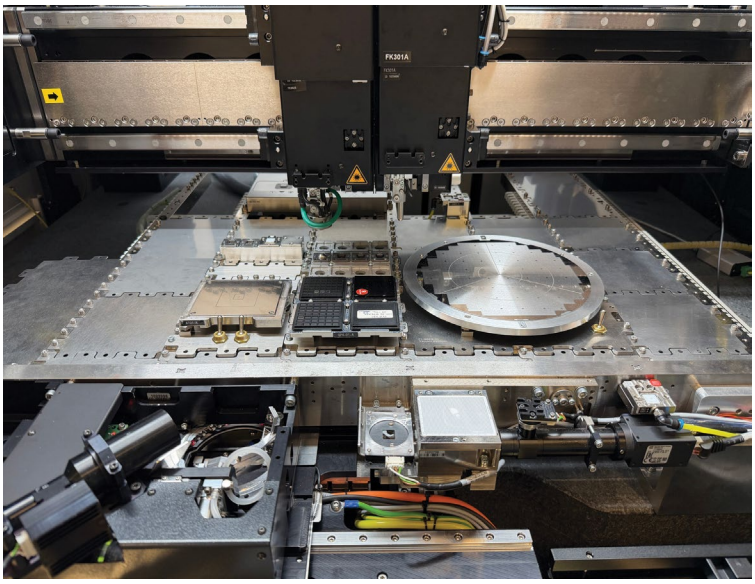

The primary constraint has shifted: it is no longer front-end lithography only, but the packaging process itself. Reaching the 2030 targets depends on a fundamental shift in how the industry views the “backend” process.

The value shift to fifty percent

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

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

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

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

Moving beyond pluggables

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

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

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

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

Solving optical alignment

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

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

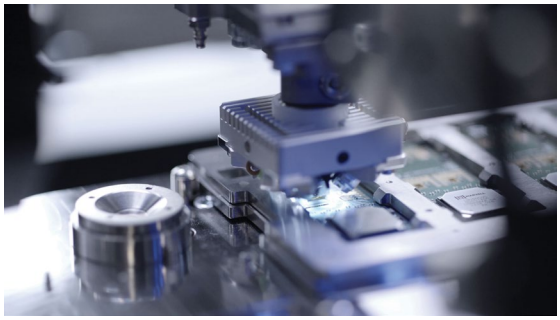

designed to guide light beams without requiring powered checks for every unit. Mastering passive methods is the only way to reach the cost structures needed for mass adoption.

The supply chain gap

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

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

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

We built a global ecosystem. Now we need a way to keep it moving

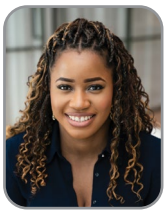

How collective intelligence will define the next era of the semiconductor industry.

BY THELMA ONYEKA, CO-CEO & CO-FOUNDER OF BEEBOLT

GLOBALIZATION has been this steadily expanding pool of opportunity for decades. For the high-tech industry, trade growth meant that companies gained access to capabilities, materials, and markets that no single region could provide alone. It's an interconnectedness that's accelerated innovation and helped build the modern semiconductor ecosystem.

But the same networks that enabled this growth have also made the sector uniquely vulnerable. Supply chains, trade systems, and regional alliances are now so tightly linked that disruptions move through value chains, and if it impacts one, it impacts many.

A quick Google will prove the point, but there's no shortage of disruptions and challenges semiconductor firms face.

The chip shortage of 2021/22 is the most obvious example of a recent - fairly cataclysmic - disruption specific to the high-tech industry. Its fallout resulted in the automotive industry halting operations altogether, not accounting for the long-term ripple effects.

But the reality is, this is only a single, well-known example within a pool of increasingly prevalent disruptions. Export controls, tariffs swings, rare earth bottlenecks, memory shortages, lead time challenges - I could go on, but this is an all-too familiar list of potential disruptions for firms and fabs and everyone in between.

Headline or no headline, industry-stopping or not, the level of impact doesn't change the fact that in 2026, they're near-daily occurrences. And it's having a significant effect on our global structures.

To keep pace with the scale and complexity of this ecosystem, the industry needs a new model of intelligence - one built for the ecosystems that connect them.

Why isn't there a solution?

To address the most obvious question first.

It's not that we don't have one, per se, it's that we don't have a viable one.

The systems currently in place are designed to aid optimization within an organization - not within an entire ecosystem. They're designed to close gaps and enhance decision-making across departments - not across firms and regions.

It's less about a lack of intelligent technology, and more about a lack of collective intelligent technology. We face shared problems that don't have shared solutions.

This is where AI applications alone fall down. AI models are designed to optimize for a single objective within a single organization. They can't collate and analyze relevant data points from across a system which is inherently reliant on the actions and decision-makings of one another. AI can't reconcile conflicting incentives across organizations, create shared situational awareness, or coordinate collective action.

A far more connected intelligence model is required to meet the needs and demands of a rapidly developing industry.

The solution is us

A mode of intelligence designed for groups, Co-Intelligence™ enables people, organizations, and systems to develop shared understanding and act in coordination across entire ecosystems.

Think of it as connective tissue. The model is a layer of intelligence which sits between industry players to close the gaps around strategy, market insight, collaboration, visibility and so on. Where most intelligence systems today are built for internal use, Co-Intelligence is a structural shift in how information is created, interpreted, and acted upon across complex systems.

It's dynamic, not static; shared, not siloed; and critically, it's ecosystem-wide, not organization-centric. And for an industry that's historically struggled with these core principles, the potential is significant.

Here's a scenario.

A specialty gas supplier in Japan detects an early production constraint that could reduce output by 10% in the coming weeks. On its own, it looks minor - but that gas is essential for lithography steps across fabs in Taiwan, South Korea, and the US. Traditionally, each company would only see its own slice of the issue, and by the time the shortage became visible downstream, fabs would already be adjusting schedules and OEMs searching for alternatives.

With Co-Intelligence, that early signal becomes a coordinated response. The supplier's alert is shared securely across the ecosystem, AI models map the impact on fab capacity, logistics, demand and so on, meaning partners can synchronize maintenance, inventory, and production plans. What would have become a multi-quarter disruption could be resolved before it even reaches the market.

This is the practical potential - and huge real-life value - of Co-Intelligence: a small upstream signal becomes collective action across the entire network. Minimal to no disruption. No headline news. No stock performance decline. BAU.

And this isn't a vague concept. The model, in some form, is already taking shape in the industry.

Early adoption by early adopters

In an early form, yes, but Co-Intelligence is already making its way in the market.

TSMC has begun building shared intelligence layers with its supplier network, integrating readiness data, logistics signals, and customer forecasts into a common environment, meaning risks can be identified earlier and production schedules to be synchronized more tightly across partners.

Intel is taking a similar approach through its IDM 2.0 strategy, while Samsung is using AI-enabled intelligence to harmonize demand signals and production planning across its semiconductor, mobile, and consumer electronics divisions.

Regions are moving in the same direction - the US CHIPS Act programmes and EU Chips Act being two examples. These emerging forms of Co-Intelligence demonstrate there's a strong appetite for a shared intelligence infrastructure.

The long-term, far-reaching impact is revolutionary, much of which we can't measure yet. Innovation cycles will shorten, recovery from disruptions will be exponentially faster, risk and associated costs will be reduced, and overall value will accelerate.

The high-tech industry will reap the rewards, but it's just as valuable to see the potential in our wider systems. On a fundamental level, Co-Intelligence can speed up the end-to-end lifecycle of semiconductor manufacturing. But the resulting impact can shift how we view global ecosystems; it has the potential to boost economies, create more equitable societies, develop critical infrastructure, establish better healthcare and so much more.

> The semiconductor industry has reached a point where its greatest strength is its greatest risk

Conclusion

The semiconductor industry has reached a point where its greatest strength is its greatest risk.

The systems that enable growth can no longer absorb the pace, scale, and complexity of modern disruption. We've built an ecosystem where every breakthrough depends on hundreds of interdependent actors, yet the intelligence guiding those actors remains fragmented, siloed, and fundamentally misaligned with how the industry actually operates.

We benefit together, we break together.

Co-Intelligence turns that on its head. It doesn't eliminate complexity, it makes complexity manageable, predictable, and strategically advantageous.

With collective intelligence, we can transform isolated signals into shared understanding, isolated decisions into coordinated action, and isolated optimizations into ecosystem-wide resilience.

Faster recovery, shorter innovation cycles, reduced risk, and greater value creation are only the beginning. The future of the high-tech industry - and the global systems it powers - depends on our ability to think, act, and adapt collectively.

It really is as simple as that. The 2026 Semiconductor Supply Chain Survey is currently live on Conductor™ to gather insights for collective action. Add your views anonymously for a stronger pulse on the semiconductor ecosystem.

Unleashing the intelligence revolution

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

Key takeaways from TSMC's 2026 europe technology symposium

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

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

A trillion-dollar industry arrives

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

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

What is particularly important in this context is the evolving nature of AI workloads. The industry is transitioning from a phase dominated by model training to one increasingly defined by inference, where trained models are deployed at scale to generate outputs in real time. This shift is critical because inference workloads are far more widely distributed, spanning cloud data centres, edge devices and embedded systems.

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

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

The evolution of transistor technology

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

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

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

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

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

The advanced logic roadmap: From N2 to A13

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

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

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

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

A14, scheduled for production in 2028, offers substantial gains over N2, including up to 15% speed improvement at the same power or up to 30% power reduction at the same speed. Logic density is increased by approximately 1.23 times, while overall chip density sees similar gains. Early product-like demonstrations indicate strong performance improvements and high SRAM yields, underscoring the maturity of the technology.

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

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

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

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

Scaling AI compute through advanced packaging

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

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

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

This scaling trajectory is remarkable, effectively enabling exponential increases in compute capability within a single package. The number of transistors integrated within these systems is expected to grow dramatically, further reinforcing the role of packaging as a key driver of performance.

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

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

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

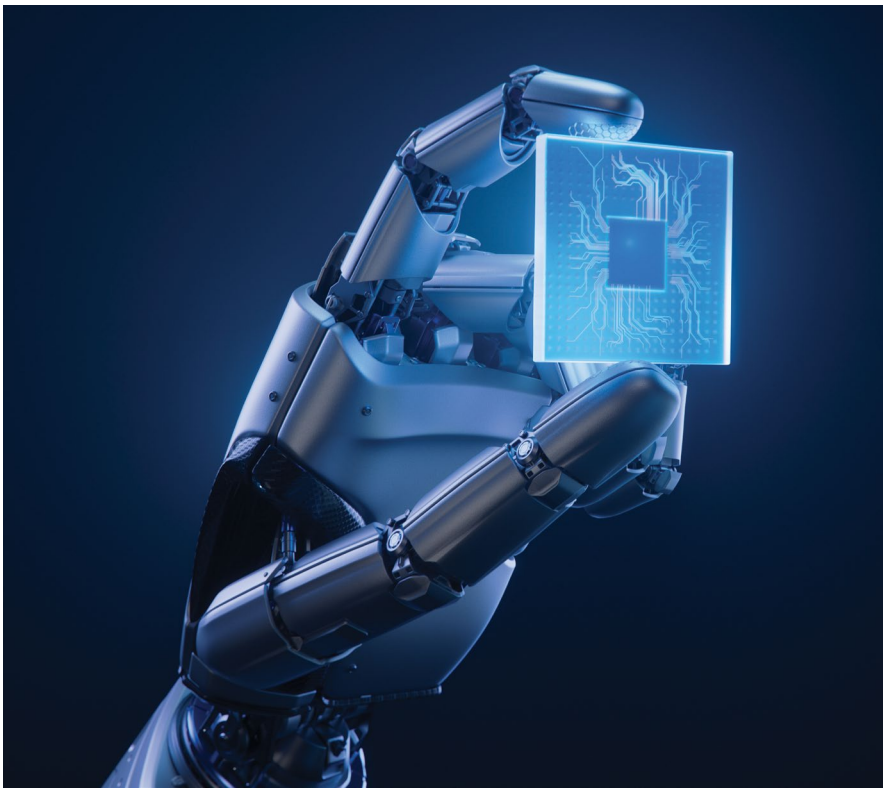

Photonics and the future of data movement

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

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

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

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

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

Automotive, robotics and physical AI

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

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

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

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

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

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

Specialty technologies and broader innovation

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

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

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

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

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

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

Manufacturing scale and sustainability

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

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

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

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

Startups and ecosystem innovation

A notable feature of the symposium was the "Innovate with TSMC" demo zone, which showcased emerging technologies from startup partners. These included innovations in sensing, edge AI, connectivity and robotics.

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

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

Conclusion: A new era of system-level innovation

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

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

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

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

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

Seeing beneath the surface

Improving power device reliability with large-area EBIC.

BY GREG JOHNSON, CARL ZEISS MICROSCOPY AND ANDREAS RUMMEL, KLEINDIEK NANOTECHNIK

THE DEEP junction structure of IGBTs plays a vital role in their reliability.

Power semiconductor devices, such as insulated-gate bipolar transistors (IGBTs), are key to the electrification of the planet. They are used in electric and hybrid vehicles, industrial motor drives, renewable energy systems, uninterruptible power supplies in data centers, and more. IGBTs have complex vertical structures with multiple p/n junctions buried deep beneath the surface (**Figure 1**). These p/n junctions control the flow of electric current in the circuit. Understanding this subsurface electrical behavior is essential for improving device reliability and identifying failure mechanisms.

Electron beam induced current (EBIC) is a powerful technique for probing these deeply buried junctions.

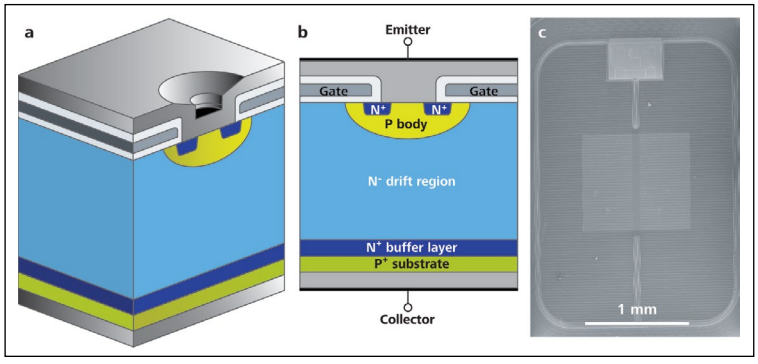

➤ **Figure 1:** Depictions of an IGBT device: a) 3D cartoon showing structure, b) cross-sectional view labeled with implants, c) top-down SEM view.

When combined with advanced scanning electron microscopy (SEM), EBIC enables spatially resolved electrical characterization across large device areas. However, extracting meaningful data requires careful attention to signal interpretation and sample preparation.

This article discusses what EBIC is and why it is important for depth-profile analysis of p/n junctions. It will cover the commonly ignored signals that affect data interpretation. Additionally, it will explain why mechanical polishing provides a superior preparation method compared to traditional etching, enabling more failure analysis (FA) labs to have access to the technique.

EBIC-SEM integrated failure analysis workflow

EBIC is a technique performed inside an SEM that measures current generated when an electron beam interacts with a semiconductor. As the beam penetrates the material, it creates electron-hole pairs. In regions where electric fields exist—such as depletion zones at p/n junctions—these carriers are separated and collected, producing a measurable current.

Unlike purely imaging-based methods, EBIC provides functional electrical insight, making it particularly valuable for FA of power devices. EBIC is uniquely suited for mapping depletion regions,

identifying leakage paths and defects, locating electrically active junctions and characterizing subsurface device behavior. For vertical devices like IGBTs, this is critical because the most important electrical activity occurs deep within the structure of the device, far below surface metallization and passivation layers.

While EBIC can be implemented on a variety of SEM platforms, ZEISS field emission (FE) SEMs such as the GeminiSEM series offer specific advantages that significantly enhance EBIC performance. These advantages include stable low-kV operation, signal stability and noise reduction, large-area imaging and an integrated electrical failure analysis workflow.

Stable low-kV operation: ZEISS FE-SEMs deliver stable imaging and beam control at very low landing energies (sub-1 kV). This is critical because lower beam energies reduce penetration depth, enabling surface-sensitive junction analysis. Higher energies increase penetration but can blur spatial resolution and complicate interpretation. The ability to precisely tune beam energy allows users to probe different depths within the device, effectively turning EBIC into a depth-profiling technique.

Signal stability and noise reduction: EBIC currents are often extremely small (nanoamp or lower). ZEISS FE-SEMs provide high beam stability, low noise imaging conditions and consistent signal generation across large scan areas. This stability is essential for distinguishing real EBIC signals from background noise.

Large area imaging: power devices often require analysis across millimeter-scale regions. ZEISS systems support wide field-of-view imaging and uniform signal response across large scan areas. This enables large-area EBIC mapping, which is essential for identifying spatially distributed defects or variations in junction behavior.

Stable low-kV operation: ZEISS FE-SEMs deliver stable imaging and beam control at very low landing energies (sub-1 kV). This is critical because lower beam energies reduce penetration depth, enabling surface-sensitive junction analysis. Higher energies increase penetration but can blur spatial resolution and complicate interpretation. The ability to precisely tune beam energy allows users to probe different depths within the device, effectively turning EBIC into a depth-profiling technique.

An integrated electrical FA workflow: the ZEISS FE-SEMs integrate well with nanoprobing systems, such as the Kleindiek Nanotechnik PS8e prober shuttle. The integration enables one- or multi-probe EBIC measurements, electron beam absorbed current (EBAC) and resistance contrast imaging. This creates a complete electrical FA environment where structural and electrical information can be correlated directly within the same tool.

The challenge of EBIC measurement interpretation

Despite its strengths, EBIC is not a straightforward measurement. In the paper, “Low Impact Analysis of Junctions in Power Devices,” presented at the 51st International Symposium for Testing and Failure Analysis (ISTFA) conference in November 2025, it was noted that raw EBIC images can be misleading

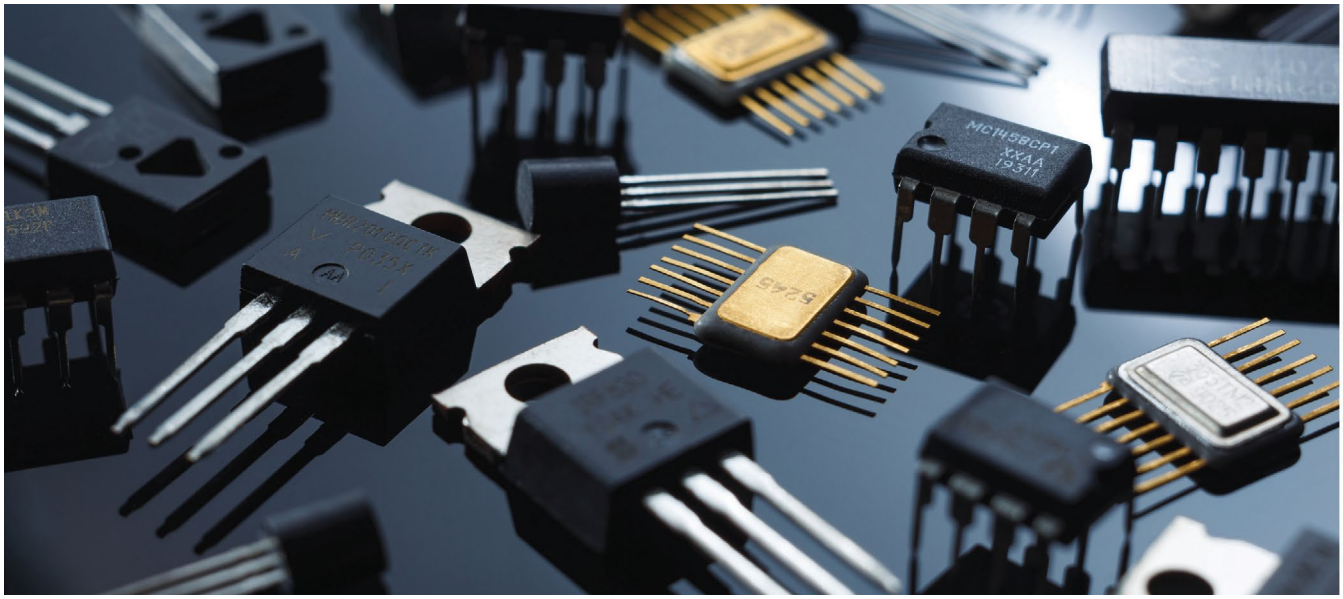

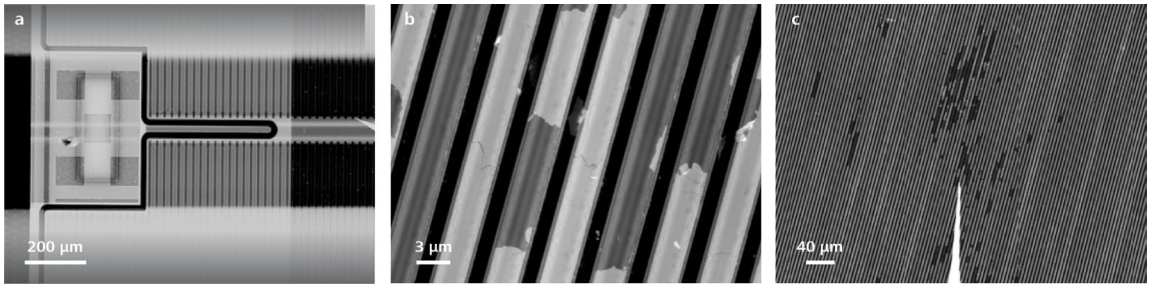

➤ Figure 2: Depictions of power device EBIC measurements from previous work: a) 30 kV EBIC image, top-down, on as-received IGBT; b) top-down EBIC on lightly etched SiC MOSFET; c) low magnification image from same work.

if interpreted naively[1]. A common assumption is that bright regions in EBIC images correspond directly to meaningful electrical activity. However, this is often incorrect. Several additional signal sources contribute to the measured current—and failing to account for them can lead to data misinterpretation.

One of the most overlooked contributors to EBIC measurements is secondary electron emission (SEE). In a typical EBIC process, a probe is placed on the device, the sample is grounded, and the SEM beam generates both electron-hole pairs and secondary electrons. These emitted secondary electrons can be picked up by the probe but are not related to junction behavior. This is especially the case for metallic areas within the sample. For example, strong signals observed at the edge of the chip or metal-covered regions often originate from SEE rather than true EBIC effects. These regions may appear bright but contain no useful electrical information.

Another frequently ignored factor is that electric fields at p/n junctions are vector quantities. Depending on the probe placement, carriers may be driven toward or away from the probe. This results in positive or negative current contributions. Thus, EBIC signals are not simply “strong vs. weak”—they can invert depending on geometry. Without being aware of this, interpretation becomes ambiguous.

EBIC signals strongly depend on electron beam landing energy. Low kV results in shallow penetration and limited interaction with deeper junctions. High kV results in deeper penetration and broader interaction volume. As the beam’s energy increases, the apparent width of junction-related signals increases; signal peaks may shift or merge, and contrast can decrease if the beam penetrates beyond the depletion region. Ignoring this dependence can lead to incorrect conclusions about junction size or location.

A key insight from the ISTFA 2025 paper is that maximum signal intensity does not directly correspond to junction position. Instead, useful information emerges when signals are analyzed relative to a baseline. Subtracting background contributions reveals true junction-related features. This approach helps to transform the data into meaningful electrical maps.

By systematically varying beam energy and analyzing EBIC profiles, a consistent pattern emerges. At low energies, only shallow portions of depletion regions are detected, and at higher energies, deeper portions become visible. This signal broadening reflects increased electron penetration. Combining electron transport simulations enables estimation on junction depth, mapping of depletion region geometry and quantitative correlation between beam energy and electrical response. This transforms EBIC from a qualitative imaging method into a quantitative characterization tool.

Sample preparation is critical for EBIC

Sample preparation is one of the most critical factors in EBIC analysis. Acid etching has traditionally been used to expose subsurface features, but it

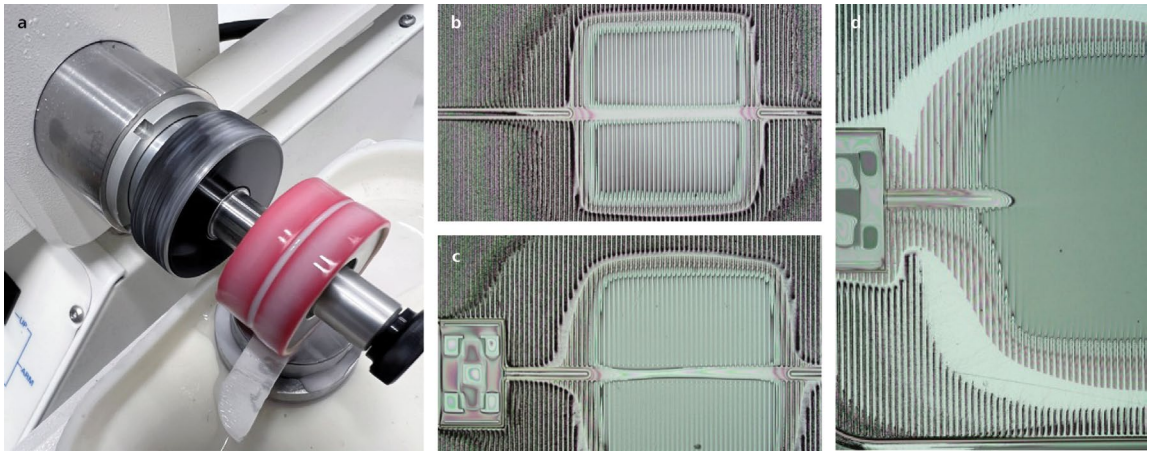

➤ Figure 3: View of the chip under dimple polishing: a) view of the dimple polisher, b) samples after three minutes, c) after an additional one minute, d) completed work.

introduces several problems. Non-uniform material removal results in inconsistent surfaces, and buried regions remain blocked by metal or dielectric layers. It also results in surface roughness variability, which complicates interpretation. Etching produces surfaces where EBIC signals are incomplete, spatially inconsistent and difficult to interpret at low magnification **(Figure 2)**.

Carefully controlled mechanical polishing, especially dimpling techniques, offers significant advantages **(Figure 3)**. It creates a gradual depth profile across the sample exposing surface structures at the edges and deep junctions near the center, enabling continuous observation across multiple device layers in a single scan.

Unlike etching, polishing produces smooth, controlled surfaces and reproducible depth gradients—essential to correlating EBIC signals with physical structures. Large portions of the device remain intact, which is critical for large-area mapping. With polishing, there is a more gradual exposure of junctions, allowing for clear correlation between structure and EBIC response and systematic analysis of depth-dependent behavior—resulting in better interpretation of the signals. Finally, mechanical polishing offers environmental and safety benefits since there is no need for hazardous chemicals or complex waste management. This makes EBIC more accessible to a broader range of FA labs.

Enabling large-area, depth-resolved EBIC

In summary, with controlled mechanical polishing, stable low-kV SEM operation and careful EBIC signal interpretation, it becomes possible to perform large-area, depth-resolved electrical characterization of power devices **(Figure 4)**.

This approach provides:

- Wide-field mapping of junction behavior
- Identification of electrically active regions
- Quantitative estimation of junction depth
- Improved understanding of failure mechanisms

Conclusion

EBIC is a uniquely powerful technique for probing the internal electrical behavior of semiconductor devices, particularly for complex vertical structures like IGBTs. However, its full potential is only realized when both instrumentation and methodology are carefully optimized.

> Unlike etching, polishing produces smooth, controlled surfaces and reproducible depth gradients—essential to correlating EBIC signals with physical structures. Large portions of the device remain intact, which is critical for large-area mapping.

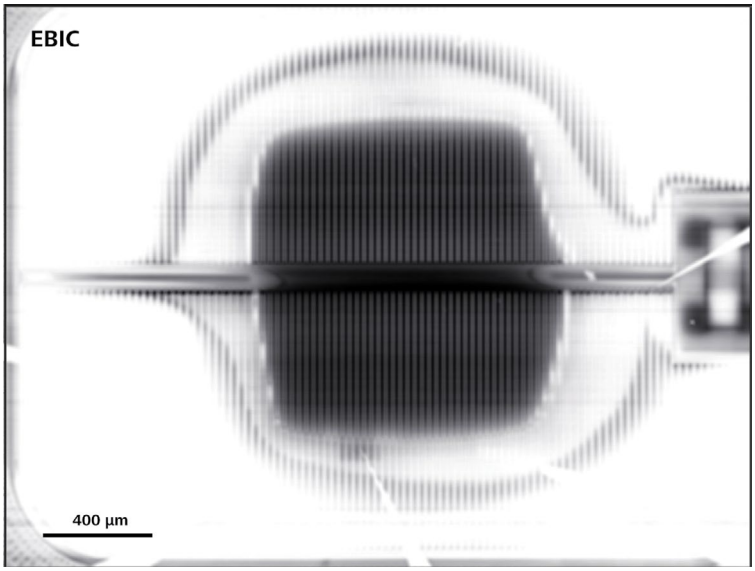

➤ **Figure 4:** Low magnification EBIC analysis showing signals from the entire width of the chip. The analysis is taken from the location on one of the metallized emitter lines, as shown in the lower center.

ZEISS FE-SEM platforms provide a strong foundation for EBIC through stable low-kV operation, high signal fidelity and seamless integration with nanoprobing systems. These capabilities enable precise control over beam interaction and support large-area analysis.

Equally important is a correct understanding of EBIC signal formation. Contributions from secondary electron emission, field directionality and beam energy effects are often overlooked but critically shape the measured signal. Proper interpretation requires moving beyond simple image contrast and adopting a more rigorous, physics-based analysis.

Sample preparation plays a critical role. Mechanical polishing offers significant advantages over acid etching by providing uniform, large-area access to subsurface structures while improving reproducibility and interpretability.

These advances transform EBIC into a quantitative technique for evaluating p/n junctions across large areas, supporting improved reliability and deeper insight into advanced power semiconductor devices.

REFERENCES

➤ [1] G. Johnson, H. Stegmann, H. Kim <Carl Zeiss Microscopy>; A. Rummel <Kleindiek Nanotechnik>; L. McCluskey, J. Jokisaari <Ted Pella, Inc.>, “Low Impact Analysis of Junctions in Power Devices,” ISTFA 2025: Conference Proceedings from the 51st International Symposium for Testing and Failure Analysis, Pasadena, California, USA, 2025.

Silicon Labs builds a future-ready SAP strategy

An interview with Radhika Chennakeshavula, CIO and vice president of Silicon Labs, explaining the company's decision to select Rimini Street as its strategic partner to maximize the value of its SAP ECC 6.0 investment. This collaboration provides the US-based semiconductor manufacturer with long-term SAP maintenance and professional services to accelerate modernisation without costly upgrades or business disruption. The project went live ahead of schedule and has delivered significant savings which have enabled 'frictionless' investment in AI innovations across the organisation.

IN AN ERA when enterprise IT strategies are increasingly shaped by vendor roadmaps, cloud-first mandates and aggressive modernization timelines, Silicon Labs has taken a deliberately different path. Rather than rushing into a wholesale platform migration, the semiconductor company has chosen to extend the life of its SAP ECC 6.0 environment through a strategic partnership with Rimini Street - freeing up significant cost savings to invest in artificial intelligence (AI), data platforms and innovation at the edge.

This approach reflects a broader shift underway across parts of the semiconductor industry: a move away from transformation for its own sake, towards more measured, value-led decision-making that balances operational stability with forward-looking capability.

Beyond vendor timelines: A strategic decision

At first glance, changing SAP support providers could be perceived as a tactical move. However, Silicon Labs' leadership makes clear that the decision was rooted in long-term business value rather than short-term operational convenience.

The core motivation was to decouple IT strategy from vendor-driven timelines - particularly as SAP customers face pressure to migrate away from ECC systems in line with evolving support policies. For Silicon Labs, the question was not whether modernization would happen, but when it would deliver the greatest impact.

"We didn't want to change our business processes driven by vendor-imposed deadlines," explains **Radhika Chennakeshavula**. "Instead, we wanted to ensure our critical systems remained fully supported while giving ourselves the time to move into a modern environment when it genuinely makes sense."

This perspective proved particularly prescient in a period of corporate evolution. The flexibility gained from

extending ECC support has allowed the organisation to reassess its roadmap in light of external developments, including significant strategic activity. The underlying philosophy is clear: modernization should follow business readiness, not compliance deadlines.

Why not RISE with SAP - yet?

Many SAP customers are pursuing RISE with SAP as a pathway to cloud transformation, but Silicon Labs has taken a more selective approach. The company has not rejected RISE outright; rather, it has chosen to delay adoption until it aligns more closely with its operational needs and business priorities.

Experience across the semiconductor sector has shown that RISE can deliver strong value in manufacturing-intensive environments, particularly where sustainability initiatives, production optimisation and supply chain integration are central concerns. However, Silicon Labs operates under a fabless model, meaning many of these benefits are less immediately relevant.

At the same time, the company's existing environment remains stable and fit for purpose. A migration undertaken purely for support reasons would have delivered limited near-term return on investment. By postponing the transition, Silicon Labs has preserved both flexibility and focus, ensuring that any future move is driven by opportunity rather than necessity.

Preserving value while enabling innovation

A central challenge in delaying large-scale transformation is maintaining the value of existing systems while continuing to innovate. Silicon Labs has addressed this through a dual strategy that stabilises its ERP core while unlocking funding for new initiatives.

By transitioning SAP support to Rimini Street, the company achieved approximately 50% savings on annual support costs. Rather than absorbing these savings into general budgets, Silicon Labs has redirected them into innovation.

This financial shift has enabled the organisation to launch a broad set of AI pilots and proofs of concept, while also beginning the development of a unified enterprise data platform. These

initiatives are laying the groundwork for more advanced analytics, automation and machine learning capabilities.

Radhika emphasises that while AI is now essential, it must be implemented in a disciplined way. New technologies cannot simply be layered onto existing operations without clear purpose and planning. By funding experimentation through operational savings, Silicon Labs has created a controlled and sustainable pathway to innovation.

Stability, support and operational confidence

The move to Rimini Street has also delivered significant improvements in operational stability and support experience.

SAP ECC environments are both complex and mission-critical, particularly in finance functions where compliance and regulatory requirements are stringent. Maintaining reliability in these systems is therefore non-negotiable.

Through its new support model, Silicon Labs has achieved a more predictable cost structure with a fixed multi-year contract, eliminating the annual cost increases typically associated with vendor support. The partnership also provides comprehensive coverage across the entire SAP landscape, rather than limiting support to specific components.

Equally impactful has been the responsiveness of the support service. Issues are addressed rapidly, often within minutes, reducing downtime and easing the burden on internal IT teams. This has translated into a more confident and less reactive operational posture.

There is also a cultural dimension to this shift. The leadership team challenges the assumption that older systems are inherently obsolete, suggesting instead that well-functioning platforms should be viewed as valuable assets. In this context, SAP ECC is not a legacy system to be replaced at all costs, but a "heritage" platform that continues to deliver business value.

Accelerated implementation and risk mitigation

The transition to third-party support was executed ahead of schedule, facilitated by a carefully structured overlap period that minimised risk.

Rimini Street aligned its billing to begin only after existing SAP support ended, allowing Silicon Labs to operate both support arrangements in parallel for several months without incurring additional cost. This dual-support period provided an opportunity to validate service performance, test responsiveness and ensure operational continuity.

As confidence grew, the organisation was able to complete the transition smoothly and ahead of plan. This phased approach eliminated many of the traditional concerns associated with switching support providers, including potential downtime or service gaps.

Self-funding innovation: A new financial model

One of the most notable aspects of Silicon Labs' strategy is its adoption of a self-funding model for innovation. Rather than seeking additional budget

for new initiatives, the IT organisation has committed to generating the necessary funding through cost optimisation.

The savings achieved through the SAP support transition have played a central role in this model, but they are part of a broader effort to rationalise the company's technology landscape. Reducing tool sprawl and eliminating redundant software has further freed up resources.

This approach has made it significantly easier to invest in new capabilities. By keeping initiatives within existing budget envelopes, the IT team can move more quickly and with greater autonomy. At the same time, the model aligns closely with financial leadership priorities, as it combines innovation with fiscal discipline.

"Save at the core, innovate at the edge"

Silicon Labs encapsulates its strategy in a simple but powerful principle: "save at the core, innovate at the edge."

The ERP system represents the operational backbone of the organisation, ensuring that business processes run smoothly and reliably. However, it is not a source of competitive differentiation. Customers do not choose semiconductor suppliers based on ERP platforms; they choose them based on product innovation and performance.

By reducing expenditure on core systems and reallocating those resources to edge innovation, Silicon Labs is aligning its IT strategy with its business objectives. Investments are focused on areas that directly enhance product development, engineering productivity and customer engagement.

This includes the deployment of advanced development tools, generative AI platforms and sales enablement technologies. It also extends to enabling experimentation, allowing teams across the organisation to explore new ideas and capabilities without being constrained by budget limitations.

Enhanced monitoring and operational agility

The partnership has also driven improvements in system monitoring and operational visibility.

Previously, the organisation lacked proactive monitoring capabilities, relying instead on users to report issues after they occurred. This reactive model created inefficiencies and often left global teams - particularly in different time zones - dealing with problems before IT was even aware of them.

With the introduction of RiminiWatch, Silicon Labs now benefits from real-time monitoring, automated alerts and proactive incident management. IT teams are able to identify and address issues before they escalate, improving both system reliability and user satisfaction.

This shift has been particularly valuable for international operations, where time zone differences previously created delays in response and resolution.

Exploring agentic AI in ERP environments

Looking forward, Silicon Labs is exploring the potential of agentic AI to fundamentally reshape how users interact with ERP systems.

Traditional SAP interfaces can be complex and difficult to navigate, often requiring specialised training and reliance on detailed reports. These reports themselves can become unwieldy, presenting users with large volumes of data that are difficult to interpret quickly.

The company's vision is to introduce AI-driven agents that enable natural language interaction with ERP data.

Instead of navigating menus or analysing reports, users would be able to ask questions and receive immediate, context-aware responses.

This concept has broad applicability, from querying demand forecasts to accessing supply chain information or analysing customer activity. The goal is to make ERP systems more intuitive, accessible and responsive to user needs.

While SAP's native AI capabilities are primarily focused on cloud-based platforms, Silicon Labs is exploring ways to bring similar functionality to its ECC environment through partnerships and innovation. This includes working with third-party providers and emerging technology firms to develop agent-based solutions.

A vision for the future

The long-term vision for Silicon Labs is an ERP environment where complexity is abstracted away and users interact with systems through intelligent, conversational interfaces.

Rather than forcing employees to adapt to rigid software structures, the organisation aims to create a more natural and efficient user experience. This involves reducing reliance on traditional interfaces and reports, while enabling real-time access to data through AI-driven tools.

Importantly, this transformation does not require immediate system replacement. By layering new capabilities onto existing platforms, Silicon Labs is extending the value of its current investments while preparing for future evolution.

Strategic flexibility in a changing industry

The semiconductor industry is characterised by rapid technological change, supply chain complexity and evolving market dynamics. In this context, maintaining flexibility is essential.

Silicon Labs' decision to extend its SAP ECC environment has preserved that flexibility, allowing the organisation to respond to new opportunities and challenges without being constrained by rigid transformation timelines.

Recent developments, including significant corporate interest and valuation uplift, further highlight the importance of maintaining strategic agility. By avoiding premature commitments, the company remains well positioned to adapt its technology strategy as circumstances evolve.

Conclusion: A pragmatic blueprint for modernisation

Silicon Labs' approach to ERP strategy offers a compelling alternative to conventional transformation models.

By focusing on timing, context and business value, the company has created a balanced roadmap that combines operational stability with innovation.

The decision to extend SAP ECC support has not been about delaying progress, but about enabling it under the right conditions.

By reinvesting savings into AI and data initiatives, Silicon Labs is building a future-ready foundation without compromising current performance.

In doing so, it demonstrates that digital transformation does not have to be a binary choice between legacy and modernity. Instead, it can be a continuum - one in which organisations maximise existing assets while selectively embracing new capabilities.

For many in the semiconductor industry and beyond, this pragmatic, value-driven approach may provide a more sustainable path forward in an increasingly complex technological landscape.

Closing the gap: How intelligent sensing is reshaping wafer process control

With semiconductor fabs under increasing pressure to improve performance and maximize productivity, verifying critical tool conditions while minimizing downtime is becoming increasingly important. Here, **Vidya Vijay, Director of Business Development at Nordson Test and Inspection**, examines how intelligent in-situ sensing is helping manufacturers to achieve these objectives while supporting the wider shift toward more predictive semiconductor manufacturing.

AS SEMICONDUCTOR manufacturing advances into ever smaller nodes and more complex architectures, the margin for mechanical variation is rapidly diminishing. In modern fabrication environments, even the slightest deviation in wafer positioning or process geometry can translate into measurable differences in yield, reliability, and overall equipment effectiveness.

At the same time, fabs are under increasing pressure to maximize tool uptime, reduce operator dependency, and move toward more connected, autonomous production models, particularly as workforce constraints add growing pressure to reduce manual intervention. These competing demands expose the limitations of traditional setup and maintenance approaches, particularly where manual checks and operator judgement remain central.

In response, the industry is turning its attention to intelligent sensing technologies that align more closely with the precision demands of modern advanced semiconductor manufacturing.

Why modern process control demands a new approach

For decades, key setup and maintenance procedures in wafer processing have relied on manual intervention. Verifying wafer centering, checking transfer accuracy, or measuring critical gaps between components has historically required tools to be taken offline, opened, and physically inspected. While long accepted, this approach introduces inherent inefficiencies. Maintaining vacuum integrity in production is not simply a matter of convenience, it is fundamental to process stability, as any interruption can introduce contamination risk, trigger cleaning cycles, and require time-intensive pump-down before production can resume. What begins as a routine check can

rapidly evolve into extended and costly downtime.

Beyond this, manual procedures are often iterative and dependent on operator experience. Subtle differences in interpretation, sometimes as minute as tens of microns, can lead to meaningful variation in process outcomes. As fabs scale and process tolerances tighten, this variability becomes increasingly difficult to manage.

The growing importance of micron-level control

In advanced semiconductor processes, precise wafer positioning across X, Y, and Z axes is critical. Even minor deviations, particularly in vertical gap, can influence uniformity, contribute to edge defects, or introduce backside contamination. In more severe cases, misalignment can lead to wafer handling issues or unexpected tool alarms. In multi-chamber tools, the challenge extends further. It is not enough for each chamber to be accurate in isolation; they must also be matched to one another to ensure consistent results across production lots. Without this level of consistency, subtle variations can accumulate, impacting yield and complicating process optimization.

The implications can also be significant during recipe development, where engineers need confidence that any observed variation is driven by process chemistry rather than mechanical inconsistency. Without reliable measurement, isolating these factors becomes increasingly complex.

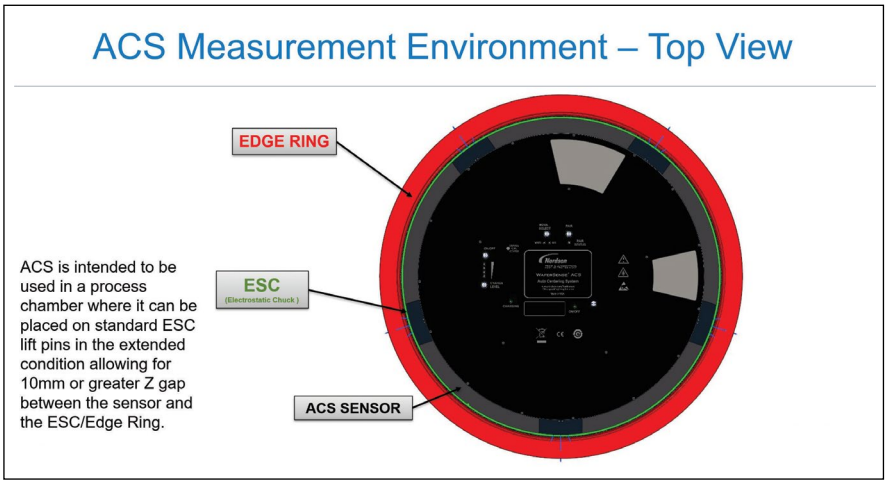

Measurement versus interpretation

One of the more understated challenges in traditional approaches is the reliance on visual feedback. In many cases, inspection methods provide only an image or qualitative indication of alignment, leaving interpretation to the operator. While useful, this still introduces subjectivity. One engineer may judge a deviation differently from another, and even small differences can be critical in high-precision processes. As a result, visualization alone is often insufficient to ensure repeatable outcomes.

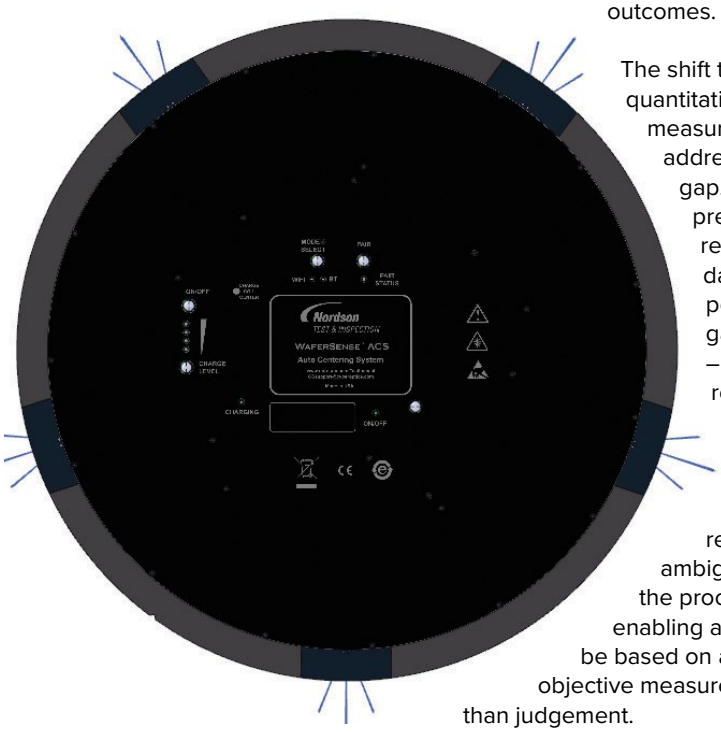

The shift toward quantitative, in-situ measurement addresses this gap. By providing precise, repeatable data on wafer position and gap conditions – alongside real-time visual feedback – modern sensing technologies remove ambiguity from the process, enabling adjustments to be based on accurate and objective measurement rather than judgement.

Bringing measurement inside the process

This need for greater precision and repeatability is driving adoption of advanced in-situ solutions. Nordson's WaferSense® Auto Centering Sensor (ACS), for example, is designed to replicate the form of a production wafer and move through the tool path while capturing real-time X, and Y offset data for centre transfer positions without disassembly or requiring the equipment to be opened to ambient conditions.

By measuring under true process conditions, including vacuum, such systems provide a more accurate view of tool behaviour while preserving process integrity. They also deliver micron-level X, and Y data that can be applied to robot alignment, wafer handoff calibration, chamber setup and process optimisation, with far less downtime.

From preventative routines to predictive control

With access to real-time, in-situ measurement data, fabs are increasingly able to move beyond fixed preventative maintenance schedules. Traditionally, tools are serviced after a predefined number of wafer lots, regardless of their actual condition. While effective to a degree, this approach can result in both unnecessary intervention and missed early warning signs.

The more meaningful shift is not simply faster measurement, but a more informed basis for decision-making. When engineers have visibility on subtle changes in wafer positioning, gap conditions, or component wear as they emerge, maintenance becomes

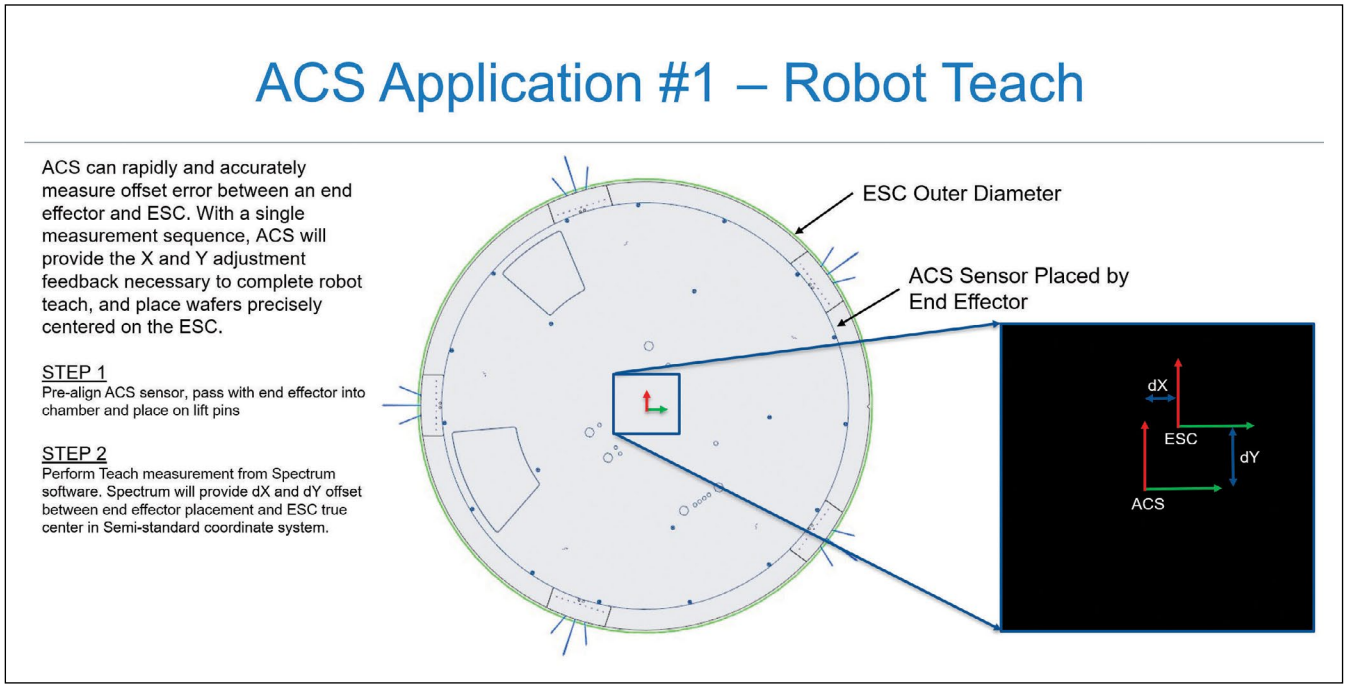

less about adhering to a fixed routine and more about responding to actual tool behaviour. This allows for earlier and more precise intervention, before variation begins to affect yield or develop into more serious and expensive failures.

Over time, this creates scope for a more refined maintenance strategy overall. Rather than relying on fixed schedules alone, fabs can make better-informed decisions about when maintenance is genuinely required, helping to reduce unnecessary disruption while maintaining tighter control over process consistency. This capability is particularly valuable when qualifying new process recipes, where both mechanical and chemical variables must be tightly controlled to ensure consistent results.

Keeping uptime up – and cost of disruption firmly down

In high-value semiconductor manufacturing, time is one of the most critical resources. Any reduction in maintenance duration has a direct impact on productivity, but the real gains come from minimizing disruption altogether. By eliminating the need for manual intervention and vacuum break procedures, in-situ sensing technologies can dramatically reduce setup and verification times. In some instances, processes that previously required several hours of intervention can be reduced to just minutes, allowing tools to return to production far more quickly. In the case of our own aforementioned Wafersense ACS, for example, we're typically seeing reductions of 30 to 70 percent in setup time.

Equally important is the avoidance of more severe downtime events, particularly as with every hour that hardware is down, manufacturers can be exposed to significant financial cost ranging from tens of thousands to millions of dollars. A misaligned wafer or incorrect gap condition can lead to handling errors or, in extreme cases, wafer breakage within the tool. Recovery from such events can take days, involving cleaning, recalibration, and full process requalification. By improving measurement accuracy and reducing reliance on manual estimation, advanced sensing solutions help mitigate these risks, supporting more stable and predictable operation.

Software as the bridge to smarter manufacturing

As with many developments in semiconductor manufacturing, the hardware is only part of the story. The value of measurement lies in how effectively it can be interpreted and applied.

Real-time software platforms are playing an increasingly important role, providing live visibility into tool conditions and enabling immediate response. Engineers can monitor measurements as they are captured, diagnose issues more quickly, and make adjustments with greater confidence. What's more, measurement data can be fed into automated workflows, supporting closed-loop control and enabling more connected, responsive manufacturing environments. This is a valuable evolution that reflects the industry's wider move towards smarter production, predictive metrology and advanced, data-driven automation.

The quieter but critical shift shaping future success

While industry attention often focuses on major advances in electronics, some of the most impactful developments in manufacturing are, in fact, those that improve the precision and repeatability of the underlying processes. Intelligent wafer-based sensing represents precisely this type of shift, helping to close a long-standing gap in semiconductor manufacturing.

As fabs continue their progression toward greater automation and tighter tolerances, in step with the wider evolution of electronics, this ability to measure and act on real conditions in real time will become increasingly important. In an industry where success is measured in microns, and often fractions of them, access to this level of insight is no longer an optional refinement, it is a requirement for success.

Powerful inspection solutions for wafer and chip level packaging.

From root cause analysis in the lab, to fast process analysis in the fab, Comet's leading 2D and 3D X-ray technology provides powerful inspection solutions for wafer and chip level packaging, enhanced with incredible AI-based software.

	Height	Diameter
TSV	105 µm	10 µm
Void	50 µm	6 µm

Visualization of 3D X-ray inspection results of TSV fills in a wafer. Wafer designed for Comet by Fraunhofer IZM_ASSID, it contains voids for illustrative purposes.

Contact us to learn more, or explore our entire portfolio

comet

yxlon.comet.tech

Chip manufacturing at a sustainability crossroads

Stephen Russell, Senior Technical Fellow, Sustainability, TechInsights, explains how the semiconductor industry is facing a critical sustainability inflection point with the growing demand and complexity of integrated circuits (ICs). Success will require making sustainability a core business and manufacturing priority rather than a constraint.

ARTIFICIAL INTELLIGENCE (AI) is driving one of the biggest technology booms in decades. But behind the rapid growth of AI systems lies a difficult reality: the semiconductor devices powering this revolution are becoming increasingly resource-intensive to manufacture, creating a sustainability challenge that the industry can no longer treat as secondary.

From advanced logic chips to high-bandwidth memory (HBM), the newest generation of semiconductors requires more process steps, more materials, and significantly more energy than previous generations. As the industry pushes toward smaller process nodes and more complex packaging architectures, manufacturing emissions are rising alongside performance gains.

This challenge is explored in depth in a recent sustainability e-book from TechInsights that examines how semiconductor manufacturing complexity is reshaping the industry's environmental footprint, and what companies can do about it.

Chip complexity drives up carbon consequences

One of the clearest examples is the transition from older manufacturing nodes, such as 90nm technology, to leading-edge 2nm production. Under comparable assumptions for wafer size, fab location, capacity, utilization, and abatement, TechInsights analysis indicates that as nodes approach 2nm, total manufacturing emissions can be roughly three times those at 90nm, while holding wafer size, location, capacity, utilization, and abatement constant. That gap is driven by process complexity and the energy demands of the leading-edge tool chain. While modern chips deliver dramatically improved computing power and energy efficiency during operation, the manufacturing process itself becomes substantially more carbon intensive. Advanced nodes demand far more lithography, deposition, etch, and inspection steps. Each additional process consumes energy, chemicals, ultrapure water, and specialized materials **(Figure 1)**.

AI and the HBM boom

AI is not just creating demand for more chips; it is changing the definition of what a chip is. Modern AI accelerators are multi-die systems that combine reticle-sized logic dies, HBM stacks, and advanced packaging. Each additional die, stack, and interconnect layer multiplies embodied emissions. In fact, leading-edge accelerators are expected to integrate roughly 250 HBM dies by 2030, making memory the dominant source of embodied carbon in many AI hardware systems. The industry is facing a paradox: chips are becoming more efficient in use, yet more environmentally demanding to produce.

At the same time, demand is growing faster than efficiency gains can offset. AI infrastructure

expansion, data center growth, and increasing global semiconductor consumption continue to drive overall emissions upward. According to TechInsights modeling cited in the e-book, by 2030, manufacturing emissions from AI GPU production are projected to rise more than twelvefold, from about 1.8 million MT CO2e in 2025 to 21.6 million MT CO2e (Figure 2).

Policy and reporting expectations are tightening

The growing complexity is changing how companies approach sustainability reporting. Historically, environmental reporting in semiconductors focused on broad facility-level metrics. Today, regulators, customers, and supply-chain partners increasingly require far more granular data. Companies are now being asked to understand emissions at the product, process-node, and supplier level. This is especially the case in Europe, where reporting requirements are raising the bar for climate disclosure and supply-chain transparency. In the United States, California climate disclosure rules are increasing pressure on large companies to account more rigorously for emissions and climate-related risk.

This is particularly important because semiconductor supply chains are global. A chip designed in the United States may be fabricated in Taiwan, packaged in Malaysia, and integrated into products sold in Europe. As environmental regulations tighten, especially within the European Union, companies throughout the entire supply chain are being pushed to provide more detailed emissions data, even if they are not directly located within those regulatory regions.

Scope 3 emissions, which are indirect greenhouse gas emissions produced throughout a company's supply chain, spanning from raw material extraction to the end-of-life treatment, have become especially challenging. Tracking the carbon footprint of materials, manufacturing steps, and outsourced suppliers across multiple countries requires far greater visibility than many companies previously maintained. Teams now need credible Scope 3 answers that can support design choices, procurement conversations, customer requests and emerging regulatory expectations.

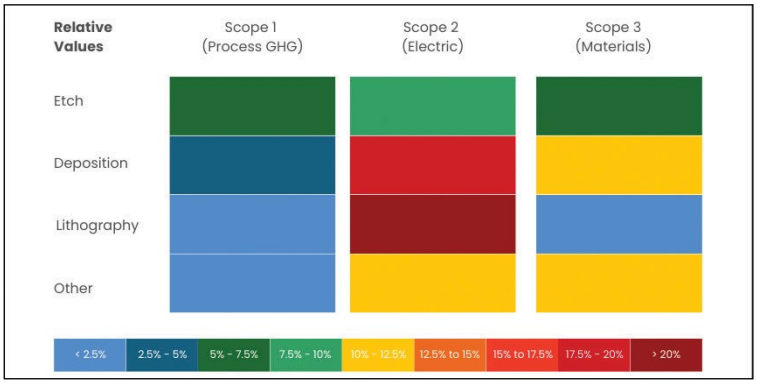

➤ Figure 1: Sustainability matrix for 3nm logic illustrating processing carbon hotspots. Scope 2 is concentrated in lithography.

Better data is the foundation for lower-carbon manufacturing

To make those decisions, companies need more than facility-level averages. TechInsights' EcoInsights platform was developed to help semiconductor companies better understand the environmental impact of semiconductor manufacturing. The platform provides detailed analysis of manufacturing carbon, process complexity, and supply chain considerations, helping organizations make more informed sustainability decisions.

A key focus is identifying "high leverage" opportunities where companies can reduce emissions most effectively. That may include process optimization, material choices, manufacturing location decisions, or supplier selection. By improving visibility into the environmental impact of specific manufacturing flows, companies can move beyond generalized sustainability goals toward more actionable strategies.

The EcoInsights modeling suite (Figure 3), for example, quantifies embodied emissions and water use across wafer fabrication, grounded in process

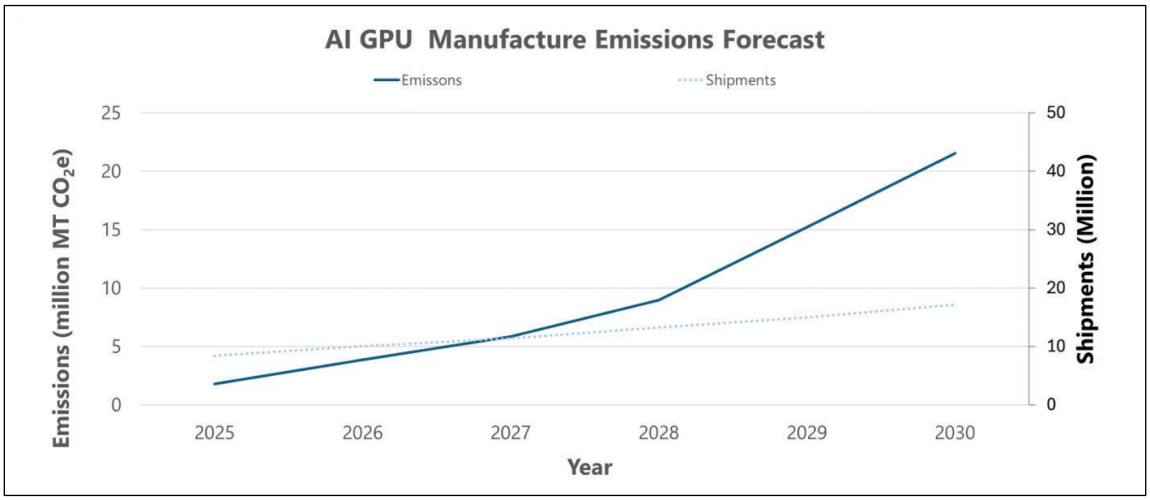

➤ Figure 2: AI GPU Manufacture Emissions Forecast, emissions plotted to left y-axis, shipments to right y-axis. Manufacturing emissions rise significantly faster than shipments.

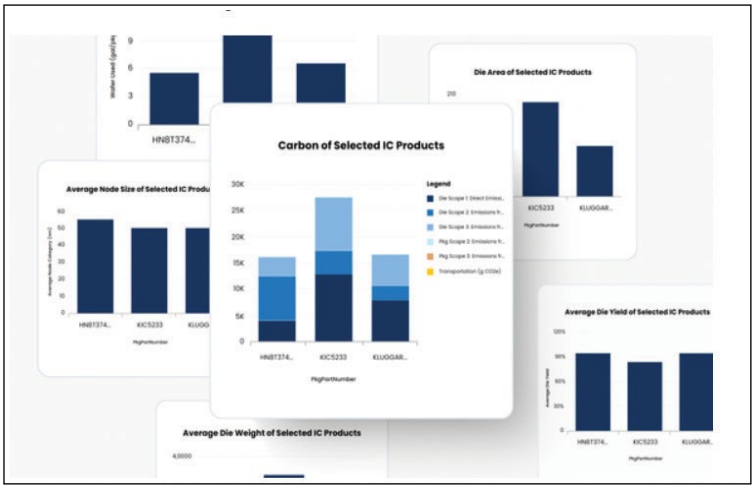

➤ Figure 3: Comparative Analysis: SK hynix HN8T374Z-JKX141, Samsung KLUGGAR-HHD-B0G1, and KIOXIA K1C5233. Source: Screenshots from TechInsights' Carbon Analyzer Module, 2024.

flows, equipment sets, and region-specific electricity factors. It helps teams answer questions like:

- Which tool families dominate emissions for a given node and device type?
- How much of the footprint is electricity (Scope 2) versus process gases (Scope 1) versus materials and chemicals (Scope 3)?
- How sensitive is the outcome to fab location, utilization, and abatement efficiency?
- Where will the hotspots move as the industry advances from 3nm to Angstrom-class nodes, or from 200-layer to 800-layer 3D NAND?

AI itself can be part of the solution

While AI is currently a major driver of semiconductor demand and energy consumption, many industry experts believe AI technologies could eventually help improve sustainability practices across manufacturing and infrastructure management. AI is already being explored for applications such as weather prediction, infrastructure optimization, and industrial efficiency improvements, areas that could have meaningful long-term sustainability benefits.

AI-assisted tools can also help companies manage the complexity of the data problem. Helping to identify proxy data, flag emissions hotspots, compare scenarios, and support faster analysis across complex manufacturing and supply-chain datasets. TechInsights is developing AI agents designed to help customers navigate the enormous volume of semiconductor manufacturing and sustainability data available today. These tools aim to help users identify proxy data, uncover patterns, and better analyze environmental impacts across complex supply chains.

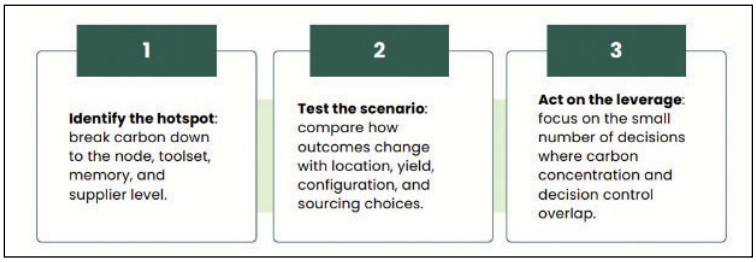

➤ Figure 4: Turning pain points into measurable actions.

Conclusion

The semiconductor industry's sustainability challenge is unlikely to disappear anytime soon. Manufacturing complexity will continue increasing and AI demand shows little sign of slowing. But improved data visibility, more precise reporting, and emerging AI-assisted optimization tools may help the industry begin addressing a problem that is becoming impossible to ignore.

The first step is turning pain points into measurable action **(Figure 4)**. Identify the hotspot, compare how outcomes change with location, yield, configuration, and sourcing choices, and then take action. Prioritize the decisions where carbon concentration and decision control overlap. The companies that make the most progress will be those that connect carbon data to real manufacturing and sourcing decisions. That means identifying emissions hotspots, testing how outcomes change by location, yield, configuration, and supplier, and prioritizing areas where teams have both high carbon impact and real decision control.

For a deeper look at the data, trends, and solutions shaping semiconductor sustainability, readers can explore the full e-book from TechInsights here: https://www.techinsights.com/carbon-age-ai-chips-earth-day-ebook.

AI-assisted tools can also help companies manage the complexity of the data problem. Helping to identify proxy data, flag emissions hotspots, compare scenarios, and support faster analysis across complex manufacturing and supply-chain datasets. TechInsights is developing AI agents designed to help customers navigate the enormous volume of semiconductor manufacturing and sustainability data available today. These tools aim to help users identify proxy data, uncover patterns, and better analyze environmental impacts across complex supply chains

SOS platform helps to address the semiconductor industry's data 'emergency'

An interview with Pedro Pires, AI Product Manager at Keysight Technologies and Phil Alsop Editor of Silicon Semiconductor, focusing on SOS Enterprise, the company's design data management platform.

PA: *What is SOS Enterprise, and how does it extend the capabilities of Keysight's existing SOS Core platform?*

PP: SOS Enterprise is Keysight's design data management platform built for global, multi-site engineering organizations that require secure collaboration, governance, and AI-ready infrastructure at scale. It extends SOS Core – which delivers high-performance version control, seamless EDA integration, and metadata-driven IP reuse for small to mid-sized teams – by unifying distributed engineering environments under a single, traceable system of record.

Where SOS Core focuses on streamlined, reliable data management for departmental or single-site teams, SOS Enterprise adds enterprise-grade collaboration, and end to end IP lifecycle traceability. It connects data across sites and domains, catalogues and governs IP for reuse, manages bills-of-materials and workspaces, and prepares engineering data for analytics, machine learning, and agentic AI workflows.

In short, SOS Core empowers teams to work efficiently and avoid data conflicts; SOS Enterprise transforms that operational data into a strategic asset that powers innovation at organizational scale.

PA: *Why is engineering data governance emerging as a critical bottleneck for AI adoption in semiconductor design?*

PP: AI / ML pipelines depend on structured, accessible, and trustworthy data. Yet the reality across most semiconductor organizations is that design data remains largely unstructured and siloed, making it unusable without substantial engineering effort, that often does not scale well.

Without standardized metadata, lineage, and well-defined APIs, teams cannot leverage historical design data for predictive analytics, yield optimization, or design automation. Manual, inconsistent governance further exposes data to security risks and compliance gaps. There is no clear path from raw engineering archives to curated, labeled datasets that ML models require.

Moreover, semiconductor companies are often subject to very stringent regulations, which require assets to be traceable to the root inputs and experiments to be fully reproducible, years after the set up.

This is why governance is often seen as the bottleneck: until organizations can catalog, version, and contextualize engineering data – and ensure that data is discoverable and traceable – AI remains aspirational rather than operational. Platforms that solve the governance layer solve the AI-readiness layer simultaneously.

PA: *What specific challenges do fragmented design files, verification data, and IP libraries create for engineering teams today?*

PP: The challenges are multi-dimensional. First, fragmented toolchains and isolated data prevent automation and insight across design,

verification, and test. Without a unified data model, cross-domain traceability is impossible, and engineers cannot reliably determine which design version generated a given simulation result.

Second, inconsistent versioning, missing lineage, and weak dependency tracking cause rework and re-spins. Design knowledge gets trapped in silos, metadata is inconsistent, and IP blocks become hard to find. Engineers repeatedly recreate existing blocks, and organizations lose institutional knowledge.

Third, the productivity cost is substantial: Legacy storage systems slow down data operations, version conflicts arise when multiple engineers work on the same block, and teams must rely on manual coordination across global locations. For engineering leaders, this translates into limited visibility into project progress, data utilization, and compliance adherence – forcing reactive rather than proactive decision-making.

Taken together, these challenges create a compounding drag on engineering velocity that grows worse as design complexity and team distribution increase.

PA: *How does SOS Enterprise establish a "single system of record," and why is this important for enterprise-scale semiconductor development?*

PP: SOS Enterprise establishes a single system of record by unifying distributed engineering environments – spanning analog, digital, RF, and system domains – under one traceable data platform. It does this through real-time multi-site synchronization, and deep integration into EDA tools and workflows. Every artifact, dependency, and file is tracked for reproducibility and confidence, from schematic through to system-level integration, and all of them live in this single source of truth substrate that connects all steps of the development lifecycle.

For enterprise-scale development, this is essential for several reasons. Global design teams suffer from slow sync, version drift, and lack of visibility across projects and sites, all of which force manual coordination. There is also the challenge of knowledge accessibility and discoverability. It is not uncommon that redundant work is performed due to the lack of awareness that a given IP was already developed by someone else in the company. A single system of record eliminates these inefficiencies by providing a knowledge platform, consistent workspaces and global visibility. It also establishes the metadata-rich foundation required for governance, auditability, and compliance – requirements that are now board-level concerns in regulated markets such as automotive, aerospace, and defense.

Beyond operational efficiency, a unified system of record converts historical project archives into reusable, machine-readable datasets, laying the groundwork for AI-driven design workflows.

PA: *Can you explain how automated data lineage and traceability directly improve AI model reliability in chip design workflows?*

PP: AI model reliability hinges on knowing exactly what data a model was trained on, where that data came from, and how it has changed over time. In chip design, this means tracing the full chain from a specific design version, through the simulation parameters that produced a dataset, to the model that consumed it.

SOS Enterprise captures this complete design-to-simulation lineage automatically. Structured metadata and versioning ensure that datasets fed into ML pipelines are accurately labeled, contextualized, and reproducible. If a model produces unexpected results, teams can trace back to the exact design revision and simulation conditions that generated the training data.

This traceability also supports continuous learning loops. As design data evolves, model retraining can be automated with confidence that the underlying data reflects known, auditable changes rather than uncontrolled drift. The result is AI models that are not only more reliable, but also explainable and auditable – a critical requirement in industries where regulatory compliance governs design decisions.

PA: *What role does SOS Enterprise play in helping organisations meet increasingly complex compliance and regulatory requirements?*

PP: SOS Enterprise delivers compliance-ready design data management with built-in governance capabilities. This includes enterprise-grade role-based access control, encryption, end-to-end audit trails, and automated approval workflows.

For organizations in regulated industries – automotive (ISO 26262), aerospace and defense (ITAR, DO-

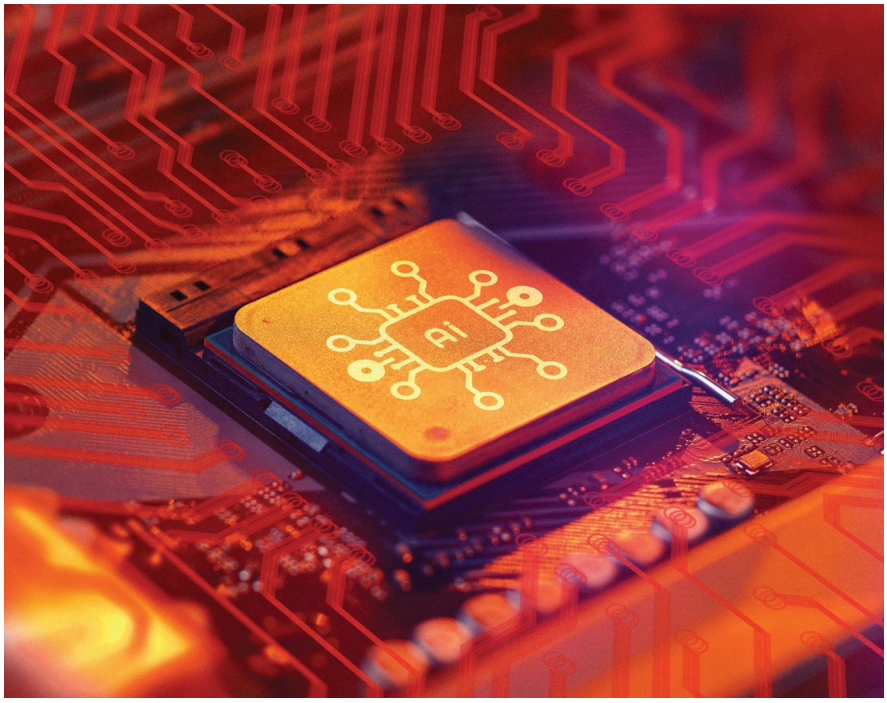

254), healthcare (medical device regulations), and government programs – SOS Enterprise provides the lineage, auditability, and secure collaboration needed to satisfy compliance mandates without manual overhead. Design data access can be controlled by geography, project, and role, and every action on the data is logged and traceable.

This positions governance not as a bottleneck, but as a strategic accelerator: it is the reason global enterprises can standardize on SOS for controlled collaboration and audit confidence, reducing the risk of data loss, IP breaches, and certification delays.

PA: *How does the platform address the needs of distributed engineering teams operating across multiple sites and geographies?*

PP: SOS Enterprise is architected from the ground up for multi-site, global collaboration. It delivers real-time synchronization across sites, ensuring that every workspace reflects the latest design state. Its Smart-Cache technology creates distributed networks of data and user workspaces that remain lightweight regardless of project data size, minimizing data duplication while maximizing performance.

For distributed teams, this means no more version drift, no more slow syncs, and no more manual coordination across time zones. Engineers at any site work against the same, consistent system of record. Access controls can be configured per site, per role, and per geography to satisfy regional compliance and data sovereignty requirements.

The platform also integrates with corporate IT systems – SSO, LDAP, and PAM – ensuring that multi-site deployment aligns with existing enterprise infrastructure rather than creating a parallel administration burden.

PA: *In what ways does SOS Enterprise reduce manual engineering workflows, and what measurable productivity gains can users expect?*

PP: SOS Enterprise reduces manual workflows at multiple levels. Version control and check-in/check-out operations are automated and integrated directly into EDA design environments – so engineers can manage data without leaving their design tools. Metadata capture, lineage tracking, and dependency management happen automatically rather than through manual documentation.

SOS Enterprise is architected from the ground up for multi-site, global collaboration. It delivers real-time synchronization across sites, ensuring that every workspace reflects the latest design state. Its Smart-Cache technology creates distributed networks of data and user workspaces that remain lightweight regardless of project data size, minimizing data duplication while maximizing performance

Approval workflows, release tagging, and IP publishing follow standardized, automated processes that replace ad-hoc email and spreadsheet-based tracking. For CAD managers, pre-built configuration templates and centralized administration eliminate per-project manual setup.

In terms of measurable outcomes, customers can expect fewer re-spins caused by data misalignment, faster tapeout cycles through reliable IP reuse and shorter verification cycles through traceable design-to-simulation linkage. Collectively, these gains translate into shorter time-to-market and lower engineering cost per project.

PA: *How does improved IP reuse across teams and locations translate into faster development cycles and reduced costs?*

PP: Enterprises lose significant engineering effort annually recreating existing design IP. The root cause is that data is stored but not organized or discoverable – metadata is inconsistent, naming conventions vary by team, and there is no centralized catalog of validated blocks.

SOS Enterprise transforms past work into future value through a metadata-driven architecture that captures design lineage and enables searchable, governed IP repositories. Engineers can discover proven IP blocks, understand the validation status and dependencies, and consume them through standardized workflows. This turns the IP library from a passive archive into an active productivity multiplier.

The impact on development cycles is direct: reusing a validated analog block or RF front-end module avoids weeks of re-design and re-verification. At scale, across dozens of projects and hundreds of engineers, this compounds into measurably shorter tapeout timelines and reduced engineering costs. The platform also enables variant creation and configuration management, so teams can adapt existing IP to new process nodes or design constraints without starting from scratch.

PA: *What security mechanisms are built into SOS Enterprise to protect sensitive semiconductor IP, particularly in regulated industries like aerospace, defence, and automotive?*

PP: SOS Enterprise provides a comprehensive security framework designed for the most demanding environments. At its foundation is granular role-based access control – which governs who can access, modify, and distribute design data, configurable by project, site, geography, and organizational role.

In addition, data in transit is protected through encryption, while end-to-end audit trails record every access, modification, and distribution event to provide the forensic visibility needed for compliance reviews and security investigations. Automated approval workflows also enforce policy gates before IP can be released or shared, helping prevent unauthorized distribution.

For aerospace and defense customers, this supports compliance with frameworks such as ITAR and internal

classified-data handling protocols. For automotive, it aligns with ISO 26262 traceability requirements. The platform's on-premises and sovereign deployment options ensure data sovereignty for organizations that cannot place engineering data in shared cloud environments.

These mechanisms collectively protect IP not just from external threats, but from the internal risks of process non-compliance, accidental data exposure, and ungoverned collaboration.

PA: *Have early customer deployments revealed any quantifiable benefits in terms of project orchestration, operational efficiency, or AI integration?*

PP: SOS's architecture is trusted by some of the world's leading semiconductor companies for design data management at enterprise scale. Customers consistently report reduced design rework through reliable version control and dependency tracking, faster tape-out timelines enabled by IP reuse and cross-site collaboration, and improved audit readiness that streamlines compliance certification processes. Even at the level of discoverability alone, companies can save a significant amount of time.

On the AI readiness front, the industry is slowly converting data infrastructure to build up to fully agentic workflows and we are yet to know of specific performance KPIs measured on this topic.

PA: *Looking ahead, how do you see the role of data management platforms evolving as AI becomes more deeply embedded in semiconductor design and verification?*

PP: Data management platforms will evolve from operational infrastructure into the strategic intelligence layer of semiconductor engineering. The platform will become the foundation on which AI copilots, predictive analytics, and autonomous design agents operate.

This evolution is already underway. As AI models move from research to production in EDA workflows – optimizing placement, expanding verification coverage, predicting yield, and eventually driving agentic automation of entire design sub-flows – the dependency on structured, governed, and traceable data will only intensify. The organizations that invest now in organizing engineering data, building lineage, and exposing it through standardized APIs will be the ones able to capitalize on each successive wave of AI capability.

For Keysight, this is the core of our roadmap: connecting the data management platform to AI/ML pipelines, enabling the transition from design data to design intelligence, and ensuring that every dataset, model, and experiment is catalogued and reusable. The semiconductor industry is entering an era defined by data, collaboration, and intelligence, and the data platform is the foundation upon which all of that rests.

Huawei outlines “τ Scaling” vision for the post-Moore semiconductor era

At ISCAS 2026, Huawei’s He Tingbo delivered one of the semiconductor industry’s most closely watched keynote speeches, presenting what she described as a new path for semiconductor evolution in the post-Moore era.

FRAMED around the concept of “τ Scaling,” the presentation outlined Huawei’s efforts over the past six years to continue improving chip performance despite the slowing of traditional geometric scaling and restrictions surrounding access to leading-edge lithography technologies.

The speech offered a detailed look at Huawei’s emerging semiconductor strategy, covering everything from smartphone processors and advanced packaging to AI systems, optical interconnects and future chip architectures. More significantly, it represented a broader philosophical shift away from conventional transistor scaling as the sole engine of progress.

Instead of relying primarily on shrinking geometries, Huawei argues that future semiconductor advances will increasingly depend on optimizing time-related characteristics across devices, circuits, chips and systems.

Beyond moore’s law

He began by acknowledging the historical importance of Moore’s Law in driving decades of computing progress.

For years, geometric scaling enabled the semiconductor industry to deliver faster processors, more capable smartphones and increasingly powerful computing systems at economically viable costs. However, the industry now faces mounting physical and economic challenges as transistor dimensions approach practical limits.

According to He, these challenges emerged even earlier for Huawei because of restrictions affecting access to advanced manufacturing technologies and lithography equipment.

“Geometric scaling plateaued for us six years ago,” she explained. “We started to rethink the nature of Moore’s Law and electronic systems.”

This led Huawei to focus on what it calls “time scaling,” or τ scaling, named after the time constant τ = RC that governs many electrical characteristics within semiconductor systems.

The company’s central argument is that semiconductor progress has never been purely about shrinking physical dimensions. Geometric scaling historically delivered benefits because it also improved timing characteristics

— faster transistors, lower latency, shorter signal paths and higher operating frequencies.

Huawei's proposal is therefore to shift optimization priorities away from purely geometric shrinkage and toward reducing delays throughout the electronic system stack.

Time as the new scaling metric

In practical terms, Huawei believes improvements can still be achieved across devices, circuits, chips and systems even without aggressive node shrinks.

At the device level, He explained that performance gains can still come from RC optimization techniques, including improvements to front-end and back-end processes. Technologies such as high-k metal gate structures and strained silicon were cited as examples already used successfully in previous Huawei products.

The broader idea is to optimize systems according to time-domain behavior rather than relying solely on transistor dimensions.

Huawei argues that this "τ-centric" methodology can apply across twelve orders of magnitude, ranging from picosecond-scale transistor switching to second-scale system-level operations.

At the circuit level, the company is targeting reductions in propagation delay through improved interconnect design, lower parasitic resistance and capacitance, shorter signal paths and optimized clock distribution.

At the system level, Huawei sees architecture and system optimization becoming increasingly important drivers of future performance gains.

The emergence of "LogicFolding"

The centerpiece of the presentation was Huawei's introduction of a new methodology called "LogicFolding."

According to He, LogicFolding represents a fundamentally new design methodology for digital systems based on vertically stacked active logic layers.

The concept involves distributing critical logic paths across multiple vertically integrated planes to shorten physical signal paths, reduce parasitic effects and improve timing performance.

Huawei claims this approach enables simultaneous optimization of power, performance, density and cost.

Traditional digital circuits rely heavily on managing timing relationships between sequential logic stages. Critical-path timing remains one of the dominant constraints limiting clock frequency and overall chip performance.

Huawei's LogicFolding approach attempts to address this by physically compressing propagation distances between flip-flop stages.

By moving logic vertically rather than spreading it horizontally across increasingly large monolithic dies, Huawei says it can significantly reduce signal wiring lengths, improve clock integrity and lower timing margins.

The company argues that this allows chips to run faster while also improving efficiency.

Hybrid bonding and advanced packaging

Achieving LogicFolding required substantial advances in hybrid bonding and packaging technologies.

Huawei stated that successful implementation depended on extremely aggressive bonding pitch scaling, specifically reducing hybrid bonding pitch below 2µm while maintaining alignment overlay errors under 0.5µm.

According to He, the ratio between hybrid bonding pitch and top metal pitch needed to remain below three for LogicFolding to work effectively.

Huawei also described innovations involving through-silicon vias (TSVs), including sub-1.5µm critical dimensions and sub-6µm pitch structures.

The company claims these technologies have achieved extremely low failure rates and high repair yields through redundancy techniques.

Many of these technologies, Huawei said, are expected to enter wider production deployment from 2027 onwards.

Kirin 2026: The first LogicFolding processor

Huawei used the keynote to formally unveil its first commercial implementation of LogicFolding technology: Kirin 2026.

Scheduled for release later in 2026, the processor reportedly transitions from a conventional single-layer architecture to a double-layer folded design.

Huawei claims the shift produced dramatic gains in transistor density. According to the presentation, transistor density increased from 155 million transistors per square millimeter to 238 million in a single architectural generation.

The company also reported a 41% improvement in SoC performance-core power efficiency alongside a

near-13% increase in maximum clock frequency.

He emphasized that these gains were achieved despite industry expectations that smartphone processors had already entered a performance saturation phase after previous Kirin generations.

The implication was clear: Huawei believes LogicFolding provides a viable route for continued performance scaling even as traditional transistor shrinkage slows.

SRAM and circuit-level folding

Huawei also provided technical examples illustrating how LogicFolding can improve specific subsystems.

One example focused on SRAM performance. As SRAM arrays increase in size, interconnect delays and communication overhead increasingly dominate total latency, often exceeding intrinsic transistor switching delays.

Huawei stated that in modern 1Mb SRAM arrays, more than 70% of total latency now originates from interconnect and communication delays.

Applying LogicFolding to SRAM structures reportedly reduced physical distances between memory arrays, peripheral circuits and processor cores while simultaneously lowering RC delays.

According to Huawei, this resulted in significant reductions in memory latency and energy per bit, while boosting SRAM operating frequency by more than 40%.

The company argued that such improvements would be difficult to achieve through process scaling alone.

Huawei also highlighted gains in processor clock distribution. In one example, folding architectures reportedly reduced clock buffer counts by over 50%, cut clock skew by 25% and shortened clock tree wire lengths by approximately 30%.

AI systems and "τ systems"

While much of the keynote focused on mobile processors, Huawei also devoted considerable attention to AI infrastructure.

He argued that AI systems represent a fundamentally different scaling problem from smartphones.

Modern AI systems consist of hundreds or thousands of heterogeneous chips operating in massively parallel configurations. In these environments, data movement increasingly dominates both energy consumption and system cost.

Huawei cited estimates suggesting more than 80% of energy in AI systems is consumed by moving data, while more than 70% of system cost is associated with data storage.

As a result, Huawei believes reducing communication time — not merely increasing raw compute capability — is becoming the defining challenge of AI infrastructure.

The company's AI strategy is centered around what it calls "SuperNode" systems, including the Ascend 910C and newly announced Ascend 950 platforms.

Huawei explained that these systems are optimized differently for training and inference workloads, with each architecture tuned according to its own τ characteristics.

Inference systems, for example, place greater emphasis on low latency and rapid token generation, while training systems prioritize throughput and overlapping communication with computation.

Unified bus and "system as one chip"

One of the most notable AI-related announcements was Huawei's new Unified Bus (UB) architecture.

Traditional multi-node AI systems often rely on multiple protocol layers and

complex interconnect conversions, introducing latency, reliability challenges and cost overheads.

Huawei's UB architecture attempts to unify communication protocols across both intra-box and inter-box interconnects using a peer-to-peer design.

The company claims this approach eliminates many protocol translation stages while simplifying deployment of ultra-large AI clusters.

Huawei also introduced memory-semantic communication capabilities within UB, enabling direct peer-to-peer data transfers without higher-layer protocol encapsulation.

According to He, this architecture allows AI systems to behave more like "System as One Chip," reducing latency across large-scale deployments.

Optical interconnect and HiONE

Huawei also outlined its optical interconnect strategy for future AI infrastructure.

As AI bandwidth requirements rise into the multi-terabit-per-second range, the company argues that conventional electrical interconnects become increasingly impractical due to cable bulk, thermal constraints and signal integrity limitations.

To address this, Huawei developed HiONE, a high-density optical interconnect engine delivering 8Tbps bandwidth per AI chip.

The technology reduces electrical SerDes reach from approximately one meter to just a few centimeters while extending optical communication distances to around 100 meters.

Huawei believes this distributed optical approach enables high-density, gigawatt-scale AI data centers without concentrating excessive thermal and power density into individual racks.

Rethinking 2.5D packaging

Another major theme involved overcoming what Huawei described as the "2.5D scaling dilemma."

In traditional 2.5D AI chip packaging, logic dies sit centrally within the package while memory stacks, interconnects and power delivery structures surround the edges.

As die sizes increase, compute capability scales quadratically with chip dimensions while edge-based memory bandwidth, power delivery and IO scale only linearly.

Huawei argues this mismatch creates a fundamental bottleneck for future AI scaling.

Its proposed solution is "System Folding," where memory, power delivery and optical IO move vertically onto chip surfaces rather than remaining confined to package edges.

The company believes this architecture will allow all system resources to scale proportionally with compute density.

Huawei predicts that this approach could increase hardware integration density by more than 100 times by 2035.

Challenges ahead

Despite the ambitious vision, He acknowledged that major technical challenges remain.

One significant obstacle involves developing entirely new EDA tools and methodologies capable of supporting free-form folded logic design.

Huawei stated that it has already begun developing preliminary toolchains and plans to disclose further details in future technical publications.

Thermal management was identified as another critical challenge.

As chip power densities continue rising, Huawei believes thermal optimization must span every level of the system hierarchy, from individual devices to gigawatt-scale AI infrastructure.

The company discussed work involving integrated capacitors, thermal interface optimization and advanced cooling techniques.

A new semiconductor narrative

The keynote concluded with Huawei framing τ scaling as a universal and sustainable post-Moore scaling law.

According to He, Huawei has already designed and mass-produced 381 chips between 2020 and 2026 using this broader design philosophy across multiple industries and applications.

The company believes τ scaling provides a practical roadmap for continued semiconductor evolution over the coming decade, even as traditional geometric scaling becomes increasingly difficult and expensive.

Huawei's projections include transistor density exceeding 400 million transistors per square millimeter, CPU frequencies surpassing 5GHz by 2031 and continued gains in smartphone and AI system performance.

Whether the broader semiconductor industry embraces Huawei's τ scaling framework remains to be seen. However, the presentation clearly signaled Huawei's intention to position itself not merely as a semiconductor company adapting to constraints, but as one attempting to redefine the future direction of semiconductor architecture itself.

In an era where the limits of Moore's Law are becoming increasingly apparent, Huawei is betting that time — not just geometry — will define the next generation of computing.

Whether the broader semiconductor industry embraces Huawei's τ scaling framework remains to be seen. However, the presentation clearly signaled Huawei's intention to position itself not merely as a semiconductor company adapting to constraints, but as one attempting to redefine the future direction of semiconductor architecture itself

Keeping it cool at the nano level

Semiconductor manufacturing operates at extraordinary precision. In advanced fabrication processes, temperature stability within fractions of a degree can influence wafer alignment and process consistency. Here, **Ben Kitson, head of business development at chemical etching specialist Precision Micro**, explains why cooling hardware plays a critical role in semiconductor fabrication and how chemically etched flow plates help maintain stable operating conditions inside these complex systems.

THE SEMICONDUCTOR industry itself is vast and still expanding. Recent analysis suggests the global market could reach $1.6 trillion by 2030. Behind those numbers sit fabrication plants that cost many billions of pounds. TSMC, for example, is investing more than $165 billion in semiconductor fabs in Arizona.

Inside these facilities, silicon wafers pass through specialised process tools that use light, plasma and chemical reactions to build microscopic circuit patterns layer by layer. The structures formed are measured in nanometres. At this scale, heat becomes a serious engineering challenge.

Thermal stability inside fabrication tools

Fabrication equipment generates heat from multiple sources. High power lasers, plasma process chambers and control electronics all contribute to rising temperatures inside the tool. If that heat is not carefully managed, even minor temperature changes can influence the behaviour of materials and the positioning of wafer layers.

The linear motors accelerate extremely fast, at rates of up to 150 metres per second squared. That is comparable to a car accelerating from zero to 100 km/h in just 0.1 seconds. Cooling these motors is extremely important for output and position accuracy.

Cooling systems are therefore an essential part of semiconductor manufacturing. Liquid coolant circulates through the machine, absorbing heat and carrying it away from critical components. Within these systems sit cooling plates and compact heat exchangers that help regulate temperature and keep the equipment operating within tight limits.

At the centre of these assemblies are flow plates containing networks of fluid channels. These channels guide coolant through the system and control how efficiently heat is removed. Their geometry must be extremely consistent as microchannel cooling performance has been shown to depend strongly on precise channel design. Variations in channel size or surface quality can disrupt fluid flow and reduce cooling performance.

Precision manufacturing of flow plates

Producing these plates requires manufacturing methods capable of creating intricate features in thin metal sheets. One approach is photochemical etching, a process that uses light-sensitive masks and chemical solutions to remove material from metal surfaces.

Because the metal is dissolved rather than cut, the process avoids introducing mechanical stress and produces clean channel edges. All features are created simultaneously across the plate, which helps maintain consistent geometry across large production runs.

Microchannels can also be produced using machining or laser techniques. These methods typically create features sequentially and may introduce burrs that require additional finishing. Photochemical etching is often well suited to thin flow plates where repeatability and surface quality are important.

In many semiconductor cooling assemblies, etched flow plates are stacked and brazed together to form compact heat exchangers. These plates are extremely thin and tightly integrated into the system design.

By comparison, the plates used in larger industrial heat exchangers such as printed circuit heat exchangers are much thicker and are diffusion bonded into solid blocks. Flow plates used in semiconductor cooling applications can be only a fraction of that size and thickness.

Producing these components reliably requires careful control of materials and process conditions. Semiconductor equipment manufacturers expect consistent performance from every component within the system. Even small variations in channel geometry can affect cooling efficiency and thermal stability.

The semiconductor industry is often associated with advances in chip design or breakthroughs in lithography, and Intel's $32 billion investment is a good example of the scale of investment needed to support that progress. Yet the machines that produce those chips depend on a network of supporting technologies working quietly in the background.

Cooling systems are among the most important of these. Within the thermal management systems that support semiconductor fabrication equipment, precision engineered flow plates help maintain stable temperatures while wafers are processed into finished devices.

Without that stability, the precision required for modern semiconductor manufacturing would be difficult to achieve.

ANGELTECH

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

SAVE THE DATE

12-14 APRIL 2027

Sheraton Brussels Airport Hotel

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

or email:

info@csinternational.net

info@picinternational.net

info@peinternational.net

info@angel-tech.net

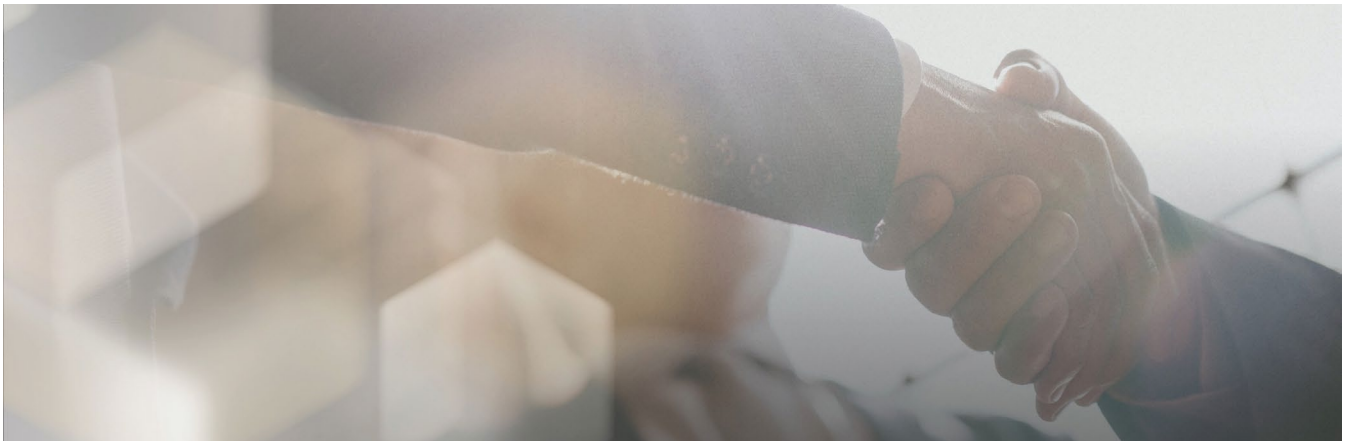

A new era for EUV lithography: How Multi-Trigger Resist technology could redefine semiconductor scaling

Dinesh Bettadapur, Irresistible Materials' CEO, discusses the company's strategic partnership with TOK, the evolution of Multi-Trigger Resist technology, and the future of EUV lithography.

THE SEMICONDUCTOR industry is entering one of the most demanding phases in its history. Extreme ultraviolet (EUV) lithography has become the backbone of advanced chip manufacturing, yet the challenges associated with pushing resolution, reducing roughness, and improving sensitivity continue to intensify. As the industry transitions from low-NA to high-NA EUV - and eventually toward hyper-NA systems - the pressure on resist materials has never been greater.

Against this backdrop, Irresistible Materials (IM), a UK-based specialist in advanced resist chemistry, has taken a bold step. The company recently announced a major strategic partnership with Tokyo Ohka Kogyo (TOK), one of the world's most established photoresist manufacturers. The collaboration combines IM's innovative Multi-Trigger Resist (MTR™) platform with TOK's manufacturing scale and deep industry relationships.

In a wide-ranging conversation, Dinesh Bettadapur - who joined the company 18 months ago - outlined the motivations behind the partnership, the science behind MTR, and the broader implications for the future of EUV lithography.

A strategic partnership built for scale

When Dinesh joined Irresistible Materials, one of the first mandates was clear: secure additional funding and strategic support needed to accelerate the company's trajectory. That search led to TOK, a company with decades of experience in photoresist development and high-volume manufacturing.

The partnership is two-pronged. First, TOK has made a significant investment in IM, providing the financial foundation required to scale the company's technology. Second, the two organisations have entered into a joint development agreement focused on productising and commercialising the MTR platform.

From IM's perspective, the collaboration brings together complementary strengths. IM contributes a multidisciplinary technical team that has spent years developing the MTR platform from first principles, along with a management team experienced in scaling technology companies and navigating successful exits. TOK, meanwhile, brings world-class R&D capabilities, deep expertise in resist chemistry across multiple lithography generations, and manufacturing facilities capable of supporting high-volume production.

The result, Dinesh says, is a "very synergistic relationship" that positions both companies to accelerate the introduction of next-generation resist materials into the market.

Understanding Multi-Trigger Resist technology

At the heart of Irresistible Materials' value proposition is its Multi-Trigger Resist platform - a technology designed to combine the best attributes of chemically amplified resists (CAR) and metal-oxide resists (MOR), while avoiding their limitations.

The best of CAR and MOR - without the drawbacks

Chemically amplified resists have long been the workhorse of lithography, relying on catalytic reactions to

achieve high sensitivity. However, the uncontrolled regeneration of acid in CAR systems can lead to pattern blur, line-edge roughness, and resolution limits.

MOR resists, on the other hand, offer high opacity and small molecular structures that support excellent resolution. But their metal content introduces contamination risks and complicates integration into semiconductor fabs.

MTR aims to bridge these worlds. Like CAR, it uses a catalytic mechanism - but with a crucial difference. IM has developed proprietary molecules that prevent uncontrolled acid regeneration, enabling a more controlled reaction and improved pattern fidelity. Like MOR, MTR is based on small molecules with high opacity at EUV wavelengths, supporting high resolution and strong image contrast. Yet unlike MOR, it is fully organic and metal free.

This combination, Dinesh explains, gives MTR "the best of both worlds," along with unique differentiators that make it a compelling candidate for next-generation EUV processes.

Meeting the challenges of EUV lithography

EUV lithography presents a unique set of challenges for resist materials. As Dinesh notes, the EUV light source produces a sparse photon distribution, making photon absorption efficiency critical. Resist materials must capture as many photons as possible to achieve the required sensitivity.

Beyond photon scarcity, EUV lithography introduces issues such as pattern collapse, thin-film etch transfer limitations, and shallow depth of focus. These challenges become even more pronounced as the industry moves toward high-NA systems, where the depth of focus shrinks further and patterning tolerances tighten.

IM believes that MTR is well-positioned to address these issues. Its high opacity improves photon absorption, while its controlled catalytic behaviour and small molecular dimensions support tighter resolution and reduced roughness. The company has demonstrated the ability to pattern extremely fine pitches - such as 26 and 28 nm - at doses below 30 mJ/cm^2, a level of sensitivity that could significantly reduce cost of ownership for EUV fabs.

Why small molecules matter

One of the defining characteristics of MTR is its use of small molecules rather than polymer chains. This design choice is central to the platform's performance.

Dinesh offers a simple analogy: small molecules act like smaller pixels. Just as a higher pixel density enables sharper images, smaller molecular "pixels" allow the resist to form finer, more precise patterns. Combined with the material's high opacity at EUV wavelengths, this leads to improved signal-to-noise ratio and better pattern fidelity.

This molecular architecture also contributes to reduced linewidth roughness, a critical parameter for advanced nodes where even sub-nanometre variations can impact device performance.

Speed of innovation: A competitive advantage

In the world of EUV resist development, speed matters. The qualification process at major semiconductor manufacturers is long and iterative, often spanning multiple phases of 9–12 months each. Customers provide a set of target metrics - resolution, linewidth roughness, sensitivity, defectivity, depth of focus, and CD uniformity - and suppliers must iteratively refine their formulations to converge on the optimal balance.

This process requires rapid formulation development. Each iteration must respond to customer feedback, address shortcomings, and move closer to the required performance envelope.

For a company like Irresistible Materials, agility is a strategic advantage. Larger competitors may have more resources, but smaller companies often excel at rapid experimentation and fast decision-making. IM's ability to quickly generate new formulations allows it to progress through qualification cycles more efficiently, improving its chances of adoption.

The role of AI in resist development

While resist formulation has traditionally relied on human expertise and intuition, AI and machine learning are beginning to play a role. IM is already exploring these techniques, using data from its extensive lab/fab testing to identify trends and guide formulation decisions.

Dinesh emphasises that the effectiveness of AI depends heavily on the quality of the input data. Training models requires not only large datasets but also the right kind of data - accurate, consistent, and representative of the relevant chemical and process interactions.

As the company continues to expand its dataset and refine its models, AI is expected to become an increasingly important tool in accelerating

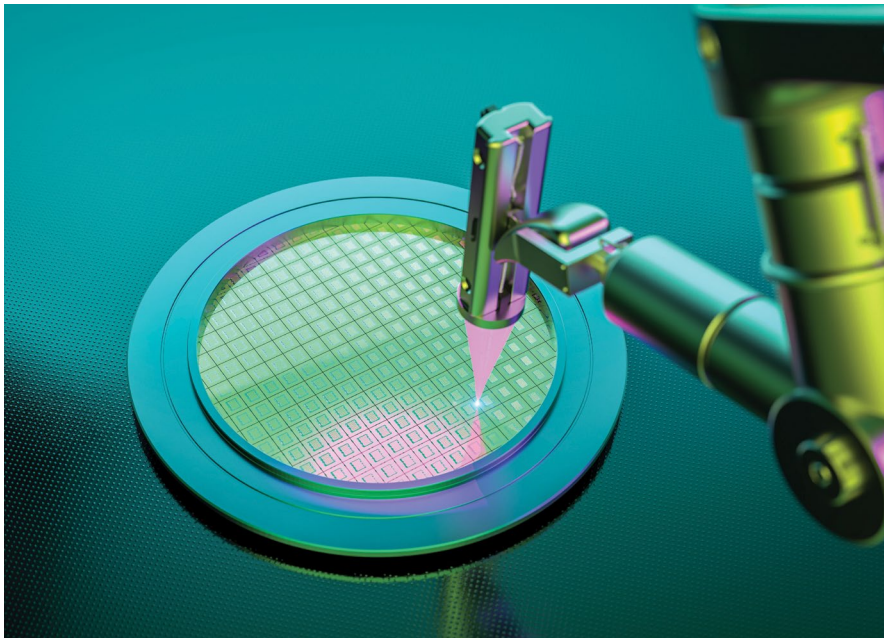

formulation development and predicting performance outcomes.

Accelerating the path to High-Volume manufacturing

One of the most significant benefits of the partnership with TOK is the ability to bridge the gap between research and high-volume manufacturing. TOK's long history in the resist business - spanning i-line, KrF, ArF, and now EUV - gives it deep expertise in scaling materials from lab-scale innovation to production-ready products.

TOK's state-of-the-art R&D facilities and world-class manufacturing infrastructure provide the foundation needed to bring MTR to market at scale. Just as importantly, TOK's long-standing relationships with major semiconductor manufacturers offer a direct pathway to customer adoption.

The response from industry partners has been overwhelmingly positive. Organisations such as ASML and imec, key players in the lithography ecosystem, have expressed strong support for the IM-TOK collaboration, recognising the potential impact of combining IM's innovation with TOK's manufacturing capabilities.

Strengthening the lithography ecosystem

Irresistible Materials has built a broad partnership ecosystem that includes tool makers, materials suppliers, software vendors, and research institutes. The addition of TOK strengthens this network, enhancing IM's credibility and expanding its reach.

TOK itself has independent relationships with many of the same partners, creating a multiplier effect. The combined ecosystem supports faster feedback loops, more robust testing environments, and a clearer path to commercialisation.

In an industry where collaboration is essential, given the complexity of EUV lithography and the interdependence of tools, materials, and processes, this expanded network is a significant strategic asset.

Sustainability: PFAS-Free and Metal-Free by design

Sustainability is becoming a central concern in semiconductor manufacturing. PFAS chemicals - often referred to as "forever chemicals" - are under increasing regulatory scrutiny due to their persistence in the environment and potential health risks. Many countries have already enacted bans or restrictions, prompting the semiconductor industry to seek PFAS-free alternatives.

MTR is inherently PFAS-free, positioning it well for current and future regulatory landscapes.

Metal contamination is another concern. Metal-oxide resists, while effective, introduce contamination risks that can impact yield and profitability. MTR's fully organic, metal-free formulation eliminates this risk, offering a cleaner and more fab-friendly alternative.

These sustainability attributes are not just regulatory advantages - they also align with the industry's broader push toward greener manufacturing practices.

Cost of ownership: A critical metric

EUV lithography is expensive. Tool costs, process complexity, and throughput limitations all contribute to high cost of ownership (CoO). Resist materials play a significant role in this equation, particularly through their impact on sensitivity and resolution requirements.

A high-speed resist with high sensitivity can reduce exposure time, improve throughput, and lower overall CoO. IM's ability to pattern tight pitches at sub-30 mJ/cm^2 doses positions MTR as a strong candidate for reducing EUV manufacturing costs.

Given that cost of ownership remains one of the biggest concerns among EUV adopters, this advantage could be a decisive factor in customer evaluations.

A roadmap for the future of EUV

Irresistible Materials designed the MTR platform specifically for EUV lithography, with scalability in mind. The company believes the platform is extensible across low-NA, high-NA, and eventually hyper-NA EUV systems.

As the industry pushes toward ever-smaller features and more demanding patterning requirements, resist materials must evolve in parallel. IM's roadmap focuses on continuous innovation, ensuring that MTR keeps pace with customer needs and lithography advancements.

Dinesh acknowledges that challenges remain, but also sees tremendous opportunity. EUV adoption is strong, and the industry is committed to making it successful

Dinesh acknowledges that challenges remain, but also sees tremendous opportunity. EUV adoption is strong, and the industry is committed to making it successful. With only a handful of leading-edge manufacturers driving demand, meeting the needs of even a subset of device applications could secure significant market traction.

The partnership with TOK provides the resources, scale, and industry access needed to advance that vision.

Conclusion: A platform poised to shape the next era of scaling

As EUV lithography becomes the foundation of advanced manufacturing, the materials that enable it must evolve rapidly. Irresistible Materials' Multi-Trigger Resist platform represents a promising step forward, combining the strengths of existing resist technologies while addressing their limitations.

With TOK as a strategic partner, IM is positioned to accelerate development, scale manufacturing, and engage more deeply with the global lithography ecosystem. The road ahead will undoubtedly present challenges, but the company's innovative approach, strong partnerships, and commitment to rapid innovation give it a compelling opportunity to influence the future of semiconductor scaling.

As Dinesh puts it, the goal is simple: continue to rapidly innovate, align with customer roadmaps, and deliver materials that meet the evolving needs of EUV lithography. If the company succeeds, MTR could become a foundational technology for the next generation of chips and a key enabler of the industry's ongoing pursuit of smaller, faster, and more efficient devices.

WEBINAR

YOUR EXPERTISE. OUR AUDIENCE. REAL IMPACT.

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

BRAND AWARENESS
Increase global visibility and credibility

LEAD GENERATION
Generate and nurture qualified leads

THOUGHT LEADERSHIP
Position your experts as industry leaders

OUR 6-WEEK WEBINAR PROMOTION PROGRAMME INCLUDES

- Promotion in *Silicon Semiconductor* magazine promotion
- Targeted exposure through newsletters to 55,000+ professionals
- Promotion in our active LinkedIn Group
- Featured on the *Silicon Semiconductor* website
- Moderated by the editor of *Silicon Semiconductor*

LET'S CREATE IMPACT TOGETHER

Our editors work with you to shape the content, engage your audience and deliver measurable results.

Contact: Jackie Cannon

E: jackie.cannon@angelbc.com

T: +44 (0)1923 690 205

SILICON SEMICONDUCTOR

The “Silicon Crossroads”: A visit to BOI Thailand as the nation bets big on chips

Backed by a national strategy targeting $79 billion in investment, Thailand is rapidly repositioning itself within the global semiconductor value chain. From advanced packaging and photonics to talent development and industrial infrastructure, a recent tour of the country’s emerging ecosystem reveals a nation determined to become a strategic force in the semiconductor industry.

BANGKOK – Walking through the halls of Thailand Board of Investment (BOI) 3 weeks back, one can feel the shift in the air. It is no longer just about tourism or automotive parts. There is a quiet, determined hum of ambition focused squarely on the trillion-dollar semiconductor industry.

I was respectfully invited by BOI to Bangkok to understand the noise. For years, Southeast Asia’s semiconductor narrative has been dominated by Malaysia’s packaging stronghold and Vietnam’s supply chain shift. But during my visit with BOI officials and a review of the newly unveiled National Semiconductor and Advanced Electronics Strategy, it became clear: Thailand is no longer content being just an assembly line. It wants to become a “Tech Owner.”

From “Made in Thailand” to “Created in Thailand”

Over strong Thai tea, senior officials of BOI *Ms. Tanita Sirisup, Deputy Secretary General and Ms Anin Meksukai – Director Foreign Investment Marketing Division* walked me through their “Semiconductor Roadmap 2030–2050.” The goal is to secure over THB 2.5 trillion (approx. $79 billion USD) in investment over the next 25 years and train more than 230,000 skilled workers. The BOI has introduced an “A1+” incentive category for front-end wafer fabrication and advanced packaging (chiplet, SiC, GaN), offering a 13 to 15 years corporate tax exemption. More importantly, they’ve partnered with Customs to prevent illicit transshipment – a mature move that signals Thailand takes supply chain integrity seriously. The Thai government has set a target to attract over 2.5 trillion Thai Baht.

But the more immediate headline is the pivot in value chain positioning.

Historically, Thailand has been known as a manufacturing hub for hard disk drives and basic automotive electronics. However, the new strategy, reviewed by the National Semiconductor and Advanced Electronics Policy Committee, explicitly aims to pull the country “upstream.”

“We are focusing on segments where we already have supply-chain depth,” *as Ms. Tanita Sirisup, Deputy Secretary General* from BOI explained during our briefing, pointing to five key target groups: Power, Sensor, Photonics, Analog, and Discrete chips. These are not bleeding-edge logic chips (like

CPUs) but the essential components driving the Electric Vehicle (EV) and AI data centre revolution—sectors where Thailand already hosts massive manufacturing bases.

The "Pillars" of Incentives

During this visit, I reviewed the BOI's playbook for luring global giants. It is arguably the most aggressive incentive structure in the region.

To move beyond Outsourced Semiconductor Assembly and Test (OSAT) into wafer fabrication, the BOI is offering the "A1+" super cluster. This includes:

- **Long Tax Holidays:** Up to 15 years of corporate income tax exemption for front-end wafer fabrication or advanced packaging (Chiplet, SiC).
- **The Talent War:** Recognizing that engineers are the bottleneck, the BOI has capped the personal income tax for foreign experts at 15%—a dramatic cut from the standard top rate—to lure talent from Taiwan, Korea, and the US.
- **Infrastructure:** Assurance of "green" energy stability and water security for fab-adjacent parks.

Ms Anin Meksukai – Director Foreign Investment Marketing Division proudly noted that this strategy is already bearing fruit. During a recent investment mission to the United States in April 2026, the BOI secured a major commitment from Phononic (a cooling chip manufacturer for Nvidia systems) to relocate upstream material production to Thailand by 2027.

"By 2030, the market will be worth $1 trillion," the BOI's *secretary general Narit Therdsteerasukdi* has stated publicly. "This roadmap ensures Thailand isn't just a consumer of technology, but a designer and producer".

Between 2018 and November 2025, Thailand attracted 1,748 investment promotion applications in the electrical and electronics sector, worth 1.17 trillion baht, representing 19 percent of total promoted investment — the largest single sector by value. The BOI joined SEMI as a member in March 2026 and is already in conversations about bringing SEMICON to Thai soil.

From talent to testing: Mahanakorn, TMEC and microchip

No strategy works without people. At Mahanakorn University of Technology (MUT), I saw boot camps that pack a full semiconductor module into 12 days. *Dr. Panavy Pookaiyaudom – President of MUT* (a legend in Thai engineering) showed me their cleanroom-in-a-classroom concept. MUT also hosts a National Semiconductor Training Centre, churning out technicians ready for fabs.

MUT has established semiconductor and electronics engineering programs specifically calibrated to meet the needs of the companies setting up operations nearby. Faculty members spoke candidly about the challenge: industry moves faster than curriculum can follow, which has forced closer collaboration with manufacturers. BOI-facilitated linkages between companies like those I would visit later in the week and universities like MUT mean that students are increasingly working with real process equipment, real design tools, and real project briefs from companies that intend to hire them.

A short drive away, the Thai Microelectronics Centre (TMEC) operates as Thailand's first IC foundry prototype line. They're developing MEMS pressure sensors for automotive and medical use.

No visit to Thailand's semiconductor ecosystem is complete without a trip to the Thai Microelectronics Center in Chachoengsao province, about 90 minutes east of Bangkok in the Eastern Economic Corridor. TMEC is, simply put, Thailand's most strategically irreplaceable asset in this industry. Founded in 1997 and operational since April 2004, it is the country's first and only research-and-development fab capable of small-volume production wafer work.

TMEC is the R&D bridge between academia and industry – and they recently signed an MOU with Lumentum International (Thailand) to co-develop photonics packaging mentions *Adisorn Tuantranont – Director TMEC.*

Speaking of industry, Microchip Thailand (their massive backend facility in Chonburi) is anything but sleepy. I toured their high-voltage test floors and saw how they're expanding capacity for SiC devices. *Arun Wellchan – Finance Director Microchip* told me they're now handling wafer-level chip-scale packaging (WLCSP) – a step up from traditional lead frame assembly.

Few names carry more weight in the microcontroller and mixed-signal space than Microchip Technology, and the Arizona-headquartered giant has made Thailand a meaningful node in its global production network. Microchip is a handful of US companies with 4,600 employees in Thailand, driving a combined 60 BOI-approved semiconductor projects worth more than 32 billion baht, a figure that reflects the depth of American confidence in Thailand's manufacturing credentials.

EIC semiconductor, UniEQ and Amata smart city

My first facility visit took me to the Lat Krabang Industrial Estate on Bangkok's eastern fringe, where EIC Semiconductor operates from within the I-EA-T Free Zone. EIC is one of those companies that rarely makes headlines precisely because it does exactly what it promises, at volume, with consistency. The firm manufactures and supplies high-quality discrete semiconductors, and its portfolio is staggering in breadth: over 10,000 parts across more than 50 package types, covering the full spectrum of diode technology from rectifiers and high-speed diodes to small signal, Zener, and transient voltage suppressors. A homegrown OSAT player that has quietly grown to 500 employees. EIC specializes in high-reliability packaging for automotive and industrial, and they're now qualifying copper clip bonding.

One surprise was UniEQ Integrated Technology, a Thai Chinese joint venture, which is a subsidiary of Foxsemicon (Fiti) Group in Thailand. Their *Project Management – Thailand Operation, Alex Sun* proudly showed me a tape-out done entirely in Thailand. UniEQ Integrated Technology showcased advanced precision engineering capabilities supporting semiconductor and electronics production environments. Conversation there with focused on automation, cleanroom systems, and supply-chain localization — areas increasingly important as companies seek greater resilience following years of global disruption.

No semiconductor ecosystem exists without industrial infrastructure, and in Thailand, that infrastructure story is largely the story of Amata Corporation. My visit to Amata Smart City the company's flagship integrated industrial development in Chonburi was one of the most visually arresting of the trip (being escorted by AMATA's own traffic police). Amata isn't just an industrial estate; it is an attempt to design the physical environment of advanced manufacturing from first principles, integrating energy systems, logistics, digital infrastructure, and workforce amenities into a coherent urban proposition for industrial tenants.

All this is happening inside Amata Smart City in Chonburi. Think of it as a living lab: 5G, smart grid, and an industrial IoT backbone that connects factories to a central data platform. Amata's executives *Kazushige Kobayashi – Section Manager Sales & Marketing and Claudia Nonthitipong- Bieri- Department Manager – Corporate Communications at AMATA Corporation* told me they're reserving land for a dedicated semiconductor park, complete with shared chemical and waste treatment – a gamechanger for small and mid-sized OSATs.

Together, these visits illustrated that Thailand's semiconductor ecosystem is not emerging from scratch. It is being built upon decades of manufacturing expertise in electronics, automotive systems, hard drives, and industrial automation.

SUBCON 2026 opening ceremony

Wednesday morning, May 13, I joined several thousand others at the Bangkok

International Trade and Exhibition Centre for the opening ceremony of SUBCON Thailand 2026 — the 20th edition of ASEAN's largest industrial subcontracting and business matching event, co-organized by the BOI, Informa Markets Thailand, and the Thai Subcontracting Promotion Association.

In his opening remarks, Thai Deputy Prime Minister and Finance Minister Ekniti Nitithanprapas struck a note that captured the event's evolved ambition. This was no longer simply a parts sourcing show. Thailand's manufacturers, he said, must move from being built-to-order parts suppliers to becoming industry solution co-creators and global corporate partners — a higher-value role that demands design capability, quality systems, and the confidence to compete internationally.

"SUBCON Thailand is a mechanism the BOI uses to put Thai entrepreneurs at the center of global supply chains — and to keep them there. The world is going through the biggest industrial shift in a generation." — Narit Therdsteerasukdi, Secretary-General, Thailand BOI

Standing next to them was the US-ASEAN Business Council, which has formed a semiconductor working group to help Thai SMEs qualify for global supply chains. Their message: "Thailand is no longer just an alternative to China – it's a primary destination for IC backend."

Lumentum, benchmark and the US-ASEAN business council

Of all the company visits on this trip, the one at Lumentum International (Thailand) in Nava Nakorn Industrial Estate, Pathum Thani, felt the most like a glimpse of where Thailand wants to go — not just where it is. Lumentum has been operating here since 2017, initially as an assembly and testing base for chips used in communications and telecommunications. Nine years later, the facility employs more than 6,000 workers, including over 700 engineers and scientists trained in photonics technology — a depth of local expertise that is genuinely rare in Southeast Asia.

I saw photonics chip assembly for 800G optical transceivers – cutting-edge work that requires submicron alignment. Their GM told me they chose Thailand because of BOI's supply chain incentives and the talent pipeline from MUT.

In 2024, exports from the Thai plant exceeded 14 billion baht, with output forecast to more than double in 2025. The most recent BOI-approved expansion — a new investment of more than 2.3 billion baht — is producing ultra-high-power chip-on-carrier light-emitting chips destined for AI processors (GPUs), laser-based medical devices, and intelligent vehicle systems. These are not commodity components. They are precision photonics devices at the frontier of AI infrastructure.

Benchmark Electronics represents a different dimension of Thailand's electronics ecosystem — the electronics manufacturing services (EMS) layer that sits between component suppliers and OEM brands. With operations anchored in Thailand, as highlighted by *Ken Hendrickson Vice President Quality / Engineering-Asia*, Benchmark brings sophisticated contract manufacturing capabilities to complex, high-mix, low-to-medium-volume production: the kind of work that demands engineering flexibility, quality systems, and deep customer partnership rather than brute-force volume.

At its Thai facility, Benchmark handles production for customers in industrial, medical, aerospace and defence, and advanced computing sectors. These are the most demanding EMS programs in existence traceability requirements, regulatory qualifications, and customer specifications that leave no margin for error. For Thailand's semiconductor supply chain narrative, Benchmark matters because it represents the integration capability that transforms components into systems the value-add layer that moves a country up the industrial sophistication curve.

My final formal engagement of the week was a roundtable convened with the US-ASEAN Business Council *Praab Pianskool, Chief Representative at US-ASEAN Business Council* a forum that put the industrial visits of the preceding days into sharper geopolitical relief. The Council brings together American companies with stakes across ASEAN, and in 2026, with global supply chains still rebalancing in the wake of tariff uncertainty and US-China technology tensions, Thailand's position as a trusted, US-aligned manufacturing destination has never been more commercially relevant.

They're facilitating something rare: a trilateral working group involving US chip designers, Thai OSATs, and Malaysian wafer fabs.

A nation positioning for the long game

What emerged from this tour was not simply a collection of isolated projects or investment announcements.

Thailand is attempting something much larger: the coordinated construction of a national semiconductor ecosystem.

The ingredients are increasingly visible — government coordination, industrial infrastructure, multinational participation, academic engagement, engineering development, and long-term policy planning.

Challenges certainly remain. Competition across Southeast Asia is fierce, particularly from Malaysia, Vietnam, and Singapore. Advanced wafer fabrication remains extraordinarily capital-intensive. Talent development must scale rapidly.

Yet Thailand's approach appears notably grounded in realism. Rather than chasing headlines, the country is leveraging existing manufacturing strengths while steadily climbing the semiconductor value chain.

For the global semiconductor industry, Thailand is no longer merely an electronics manufacturing destination.

It is becoming a strategic player worth watching.

Challenges certainly remain. Competition across Southeast Asia is fierce, particularly from Malaysia, Vietnam, and Singapore. Advanced wafer fabrication remains extraordinarily capital-intensive

Honey-like heat flow: A new heat transport regime discovered in ultrathin semiconductors

A study led by researchers from the Catalan Institute of Nanoscience and Nanotechnology (ICN2), the Universitat Autònoma de Barcelona (UAB), Eindhoven University of Technology (TU/e) and McGill University, describes a new regime of heat transport in two-dimensional materials. These findings, published in Nature Physics, open the door to new ways of controlling heat flow without altering the structure of materials, with potential applications in thermal management and thermoelectric energy conversion.

CONTROLLING heat flow is a major challenge for many technologies. In electronic and photonic devices, for example, **heat dissipation can limit the performance and efficiency**, as well as their potential for further miniaturisation. At the same time, two-dimensional (2D) materials, which are made of layers just a few atoms thick, have emerged as a promising platform in these fields. For example, 2D semiconductors are expected to be used in conduction channels of future transistors. However, their thermal behaviour remains difficult to predict and control.

Now, an international team of researchers led by ICN2, UAB, TU/e, and McGill has discovered a new regime of heat transport in ultrathin materials. The study shows that in **2D semiconductors**, in particular molybdenum disulfide (MoS_2) and molybdenum diselenide ($MoSe_2$), heat can behave in a completely new way, known as **hydro-thermoelastic transport**, where **thermal diffusion is highly impeded**. These findings, published in Nature Physics, could have a significant impact on the development of new strategies for thermal management in devices.

A combination of unexpected phenomena

Under normal conditions, heat spreads gradually from hot regions to cold ones. However, in these ultrathin materials, more complex effects occur. As **Dr Sebin Varghese**, first author of the paper, remarks: *"Our results challenge the conventional picture of diffusive heat transport and reveal a richer, more complex transport mechanism in ultrathin semiconductors."* One of the effects that occur is **phonon hydrodynamics**, whereby heat is carried collectively and **behaves like a viscous fluid**. At the same time, heating induces mechanical deformations in the material, which also affect how heat moves. Although these types of effects were already known, they had never been observed in this type of materials.

The interplay of these phenomena results in unexpected behaviour: **heat propagates much more slowly than predicted**, with the thermal diffusivity reduced by up to an order of magnitude. To reach these conclusions, the researchers used an advanced optothermal technique that enabled them to track heat flow in real time with nanometre resolution. **Prof. F. Xavier Alvarez** from the Department of Physics at the UAB, who led the theoretical part of the work, further notes: *"For the first time, we observe how mechanical stress can redirect – and even obstruct – the flow of heat in a material."*

Can heat flow "the opposite way"?

The experiments show that, in these ultrathin materials, **heat tends to remain concentrated** around the heated region for longer than expected. This happens because heating causes the material deformations that alter how heat moves through the material, even pushing the **heat flow in unexpected directions**.

As **Prof. Klaas-Jan Tielrooij** (ICN2 and TU/e), who led the study, explains: *"What surprised us most is that heat can, under certain conditions, resist leaving the hot region, which is due to contributions to the heat flux that point from cold to hot regions, rather than the conventional flux that points from hot to cold regions. This opens up a completely **new way to control heat flow intrinsically**, without the need to modify the material's structure."*

This discovery provides new fundamental insight into how heat is transported at the nanoscale and could pave the way for designing **electronic, photonic, and thermal devices** with new functionalities. The ability to control rather than simply dissipate heat could be pivotal for future technologies, from improving the thermal management of chips to making thermoelectric systems more efficient.

REFERENCES

➤ Varghese, S.; Tur-Prats, J.; Mehew, J.D.; Saleta Reig, D.; Farris, R.; Camacho, J.; Haibeh, J.A.; Sokolov, A.; Ordejón, P.; Huberman, S.; Beardo, A.; Álvarez, F.X.; Tielrooij, K.J. Controllable hydro-thermoelastic heat transport in ultrathin semiconductors at room temperature. Nature Physics. (2026). DOI: 10.1038/s41567-026-03297-1. https://www.nature.com/articles/s41567-026-03297-1

CLEAN AIR

INSPECTION SYSTEMS

COMPONENTS

ETCH

FLUID SYSTEMS

GAS DELIVERY SYSTEMS

GAS PURIFIERS

HYDROGEN GENERATORS

GAS ANALYSERS

PLASMA TECHNOLOGY

TEST

VACUUM EQUIPMENT

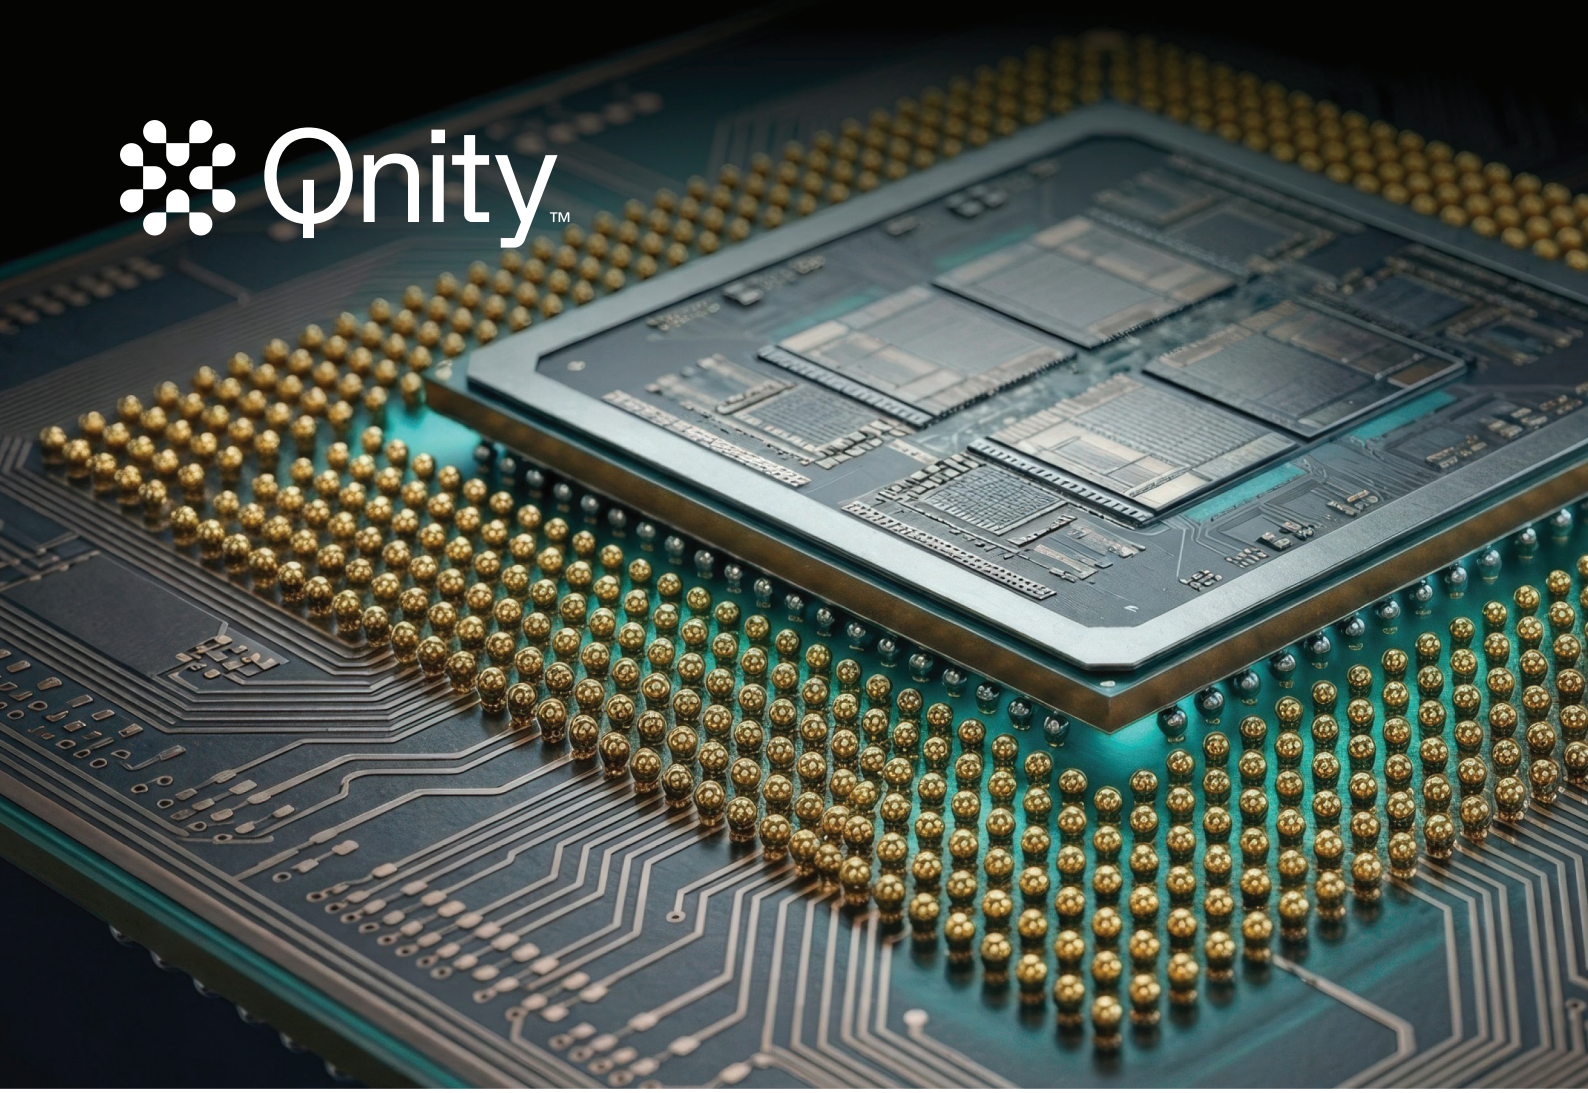

Shrink built the last era of semiconductor innovation. Stack will define the next.

Qnity is ready.

As chip architectures move to stacking, the materials and process technologies that enable 3D integration have never mattered more. Qnity's advanced packaging solutions span end-to-end—from wafer processing to interconnects to thermal management and beyond—giving manufacturers the yield, precision, and reliability that AI and high-performance computing demand.

Explore the Qnity Advanced Packaging Innovation Hub to see how we're enabling the next generation of heterogeneous design.

Scan to learn more.