CS COMPOUND SEMICONDUCTOR
CONNECTING THE COMPOUND SEMICONDUCTOR COMMUNITY

Data goes blue with GaN interconnects

avicena

VOLUME 27 ISSUE IV 2021 @compoundsemi www.compoundsemiconductor.net

INSIDE

SUPERIOR SWITCHING

New wavelength-blockers for datacom and telecom unite the strengths of InP and silicon photonics

SOARING SHARE PRICES

During the last twelve months many companies within the CS industry have outpaced the NASDAQ

BUILDING BETTER DEEP-UV LEDS

Tunnel junctions tackle the transparency conductivity conudrum that has held back progress in deep-UV LEDs

Global mega trends require best performance III-V materials

AIX 2800G4-TM

- 3D sensing
- Augmented reality
- Autonomous driving
- Cloud computing
- Internet of things
- Next generation displays

Discover the MOCVD tool of record for AsP-based applications

Made in Germany

AIXTRON

info@aixtron.com
www.aixtron.com

VIEWPOINT

By Dr Richard Stevenson, Editor

Delivering data with microLEDs

WITHIN INDUSTRY, developing a new device doesn't just take a great deal of toil. It also needs a lucrative application that promises a great return on all the cash that's been poured into the project.

What's rarely predicted, but actually quite common, is substantial deployment of a device beyond its intended application. Take the GaN-based laser: developed to provide the optical source for the successor to the DVD player and recorder, this device is now having a second lease of life as a key component in laser welding systems for 'yellow' metals, such as copper. And who would have thought that the VCSEL, an obvious candidate for data transmission, would have gone on to enjoy a ramp in sales as the source for facial recognition in mobile phones.

Right now, almost all developers of the microLED are viewing this tiny chip as a direct-emitting pixel for an upcoming generation of displays that will set new benchmarks for efficiency and contrast. But Avicena, a start-up in Mountain View, California, is marching to a different beat – it sees the microLED as a great solution to the data bottleneck in computer systems. According to this west-coast venture, microLED-based optical links are set to speed data transfer within chips, between chips and across data centres.

This application is sure to raise a few eyebrows within the optolectronics industry. After all, much effort has already been devoted to developing various classes of laser for the task of moving data. But Avicena is advocating that for transporting data at distances from a millimetre up to ten metres, all these devices fall short in many regards: they are too large, they demand a high drive power, their performance drops off at elevated temperatures, they are pricey and their yield is poor.

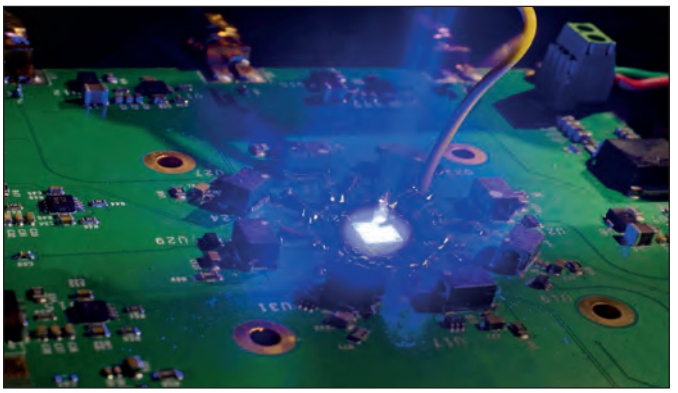

By optimising the microLED, including its doping and quantum-well structure, the team at Avicena have increased the speed of this device so that it is capable of delivering data at 4 Gbit/s using on-off modulation. This is plenty fast enough, given that data transfer would take place along thousands of parallel lanes, each with its own microLED.

You might question this approach, viewing the routing of blue light through an optical fibre as not ideal, due to high levels of absorption and dispersion. But these are not show stoppers, thanks to the short link lengths; and blue emission is the sweet spot for a silicon detector, allowing CMOS-compatible fabrication of efficient receivers with integrated photodiodes.

According to Avicena, 10 Tbit/s links based on microLEDs, fibre and photodetectors could be produced for one-to-two orders of magnitude less than the cost of today's optical link technologies employed in data centres.

While I'm hopeful that the microLED will revolutionise displays once the technology for mass transfer has been refined, it's great to see that there are also other opportunities for this miniature light source.

CONTENTS

VOL:27 ISSUE IV 2021

COMPOUND SEMICONDUCTOR
CONNECTING THE COMPOUND SEMICONDUCTOR COMMUNITY

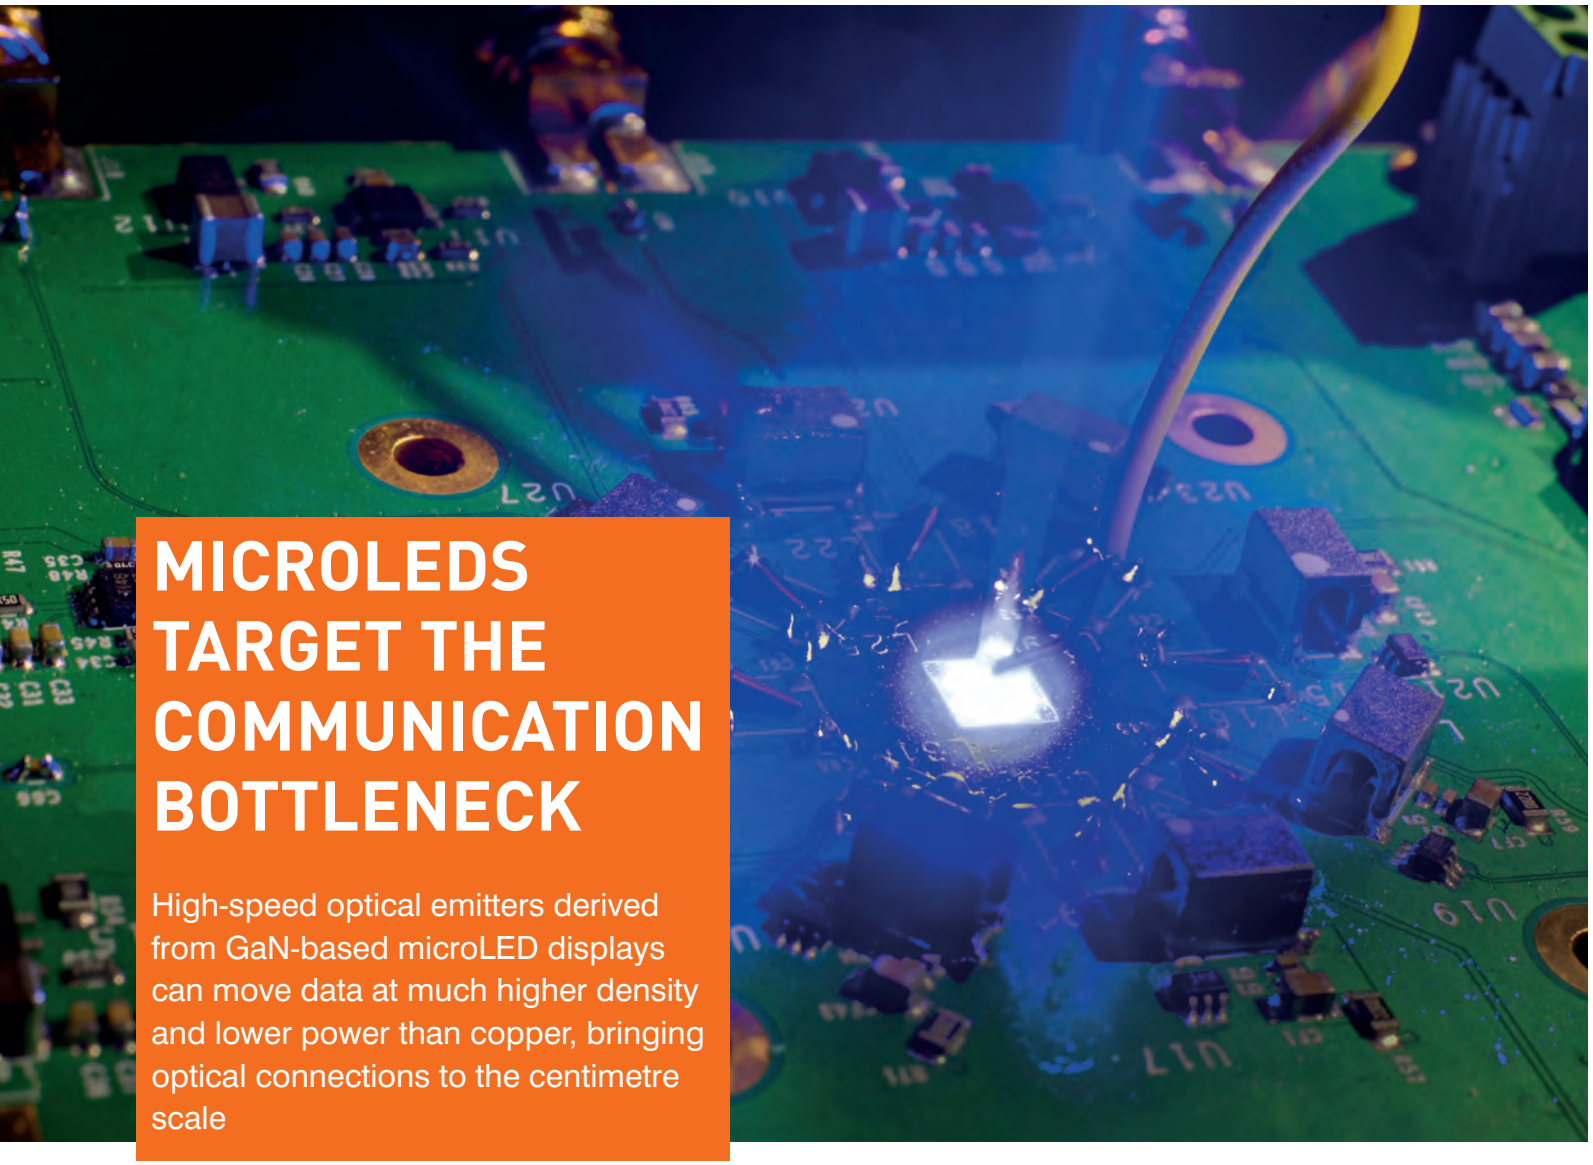

MICROLEDS TARGET THE COMMUNICATION BOTTLENECK

High-speed optical emitters derived from GaN-based microLED displays can move data at much higher density and lower power than copper, bringing optical connections to the centimetre scale

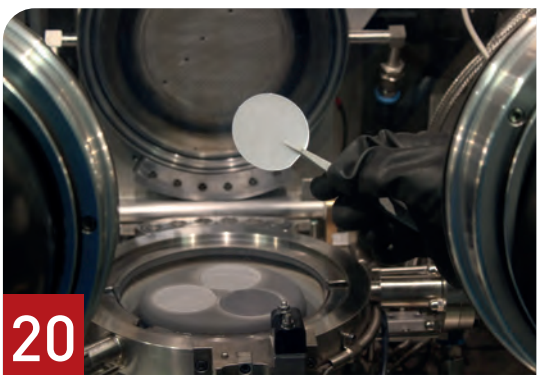

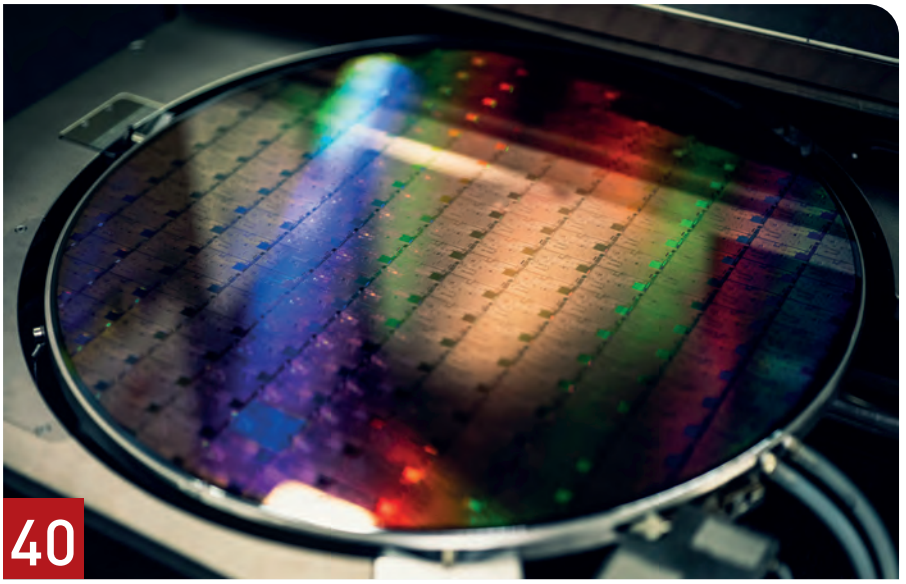

NEWS

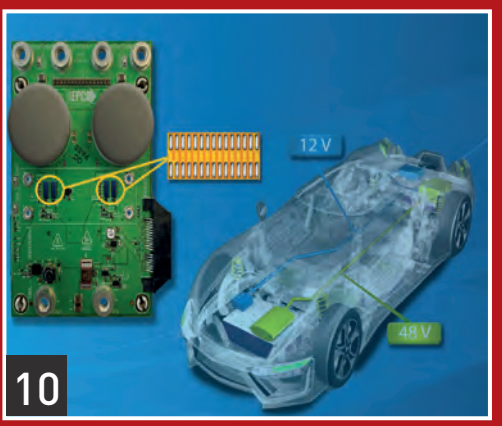

NEWS ANALYSIS

From the industry transition to 200 mm wafers to future market leaders, Omdia analyst, Richard Eden, tells Rebecca Pool, what he expects will happen next

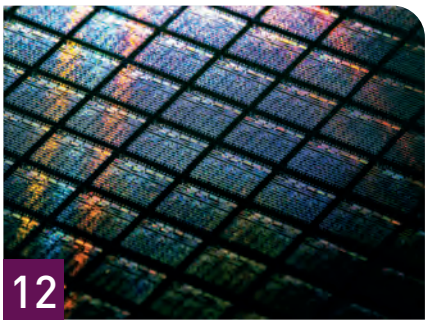

RESEARCH REVIEW

Editor Richard Stevenson richardstevenson@angelbc.com +44 (0)1291 629640
Contributing Editor Rebecca Pool editorial@rebeccapool.com
News Editor Christine Evans-Pughe chrise-p@dircon.co.uk
Sales Executive Jessica Harrison jessica.harrison@angelbc.com +44 (0)2476 718970
USA Representatives Brun Media
Tom Brun tbrun@brunmedia.com +001 724 539-2404
Janice Jenkins jjenkins@brunmedia.com +001 724-929-3550
Publisher Jackie Cannon jackie.cannon@angelbc.com +44 (0)1923 690205
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
Circulation Director Scott Adams scott.adams@angelbc.com +44 (0)2476 718970

Chairman Stephen Whitehurst stephen.whitehurst@angelbc.com +44 (0)2476 718970

Chief Executive Officer Sukhi Bhadal sukhi.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, Unit 6, Bow Court, Fletchworth Gate, Burnsall Road, Coventry CV5 6SP, UK. T: +44 (0)2476 718 970 E: info@angelbc.com

Angel BUSINESS COMMUNICATIONS

recycle When you have finished with this magazine please recycle it.

Sundiode develops stacked 3-colour microLEDs on a single wafer

SUNDIODE, a Silicon Valley based company developing micro-LED technologies for display applications has announced the achievement of fully stacked 3-colour (RGB) microLED pixel devices on a single wafer. In the stacked-RGB pixel technology patented by Sundiode and developed in collaboration with KOPTI (Korea Photonics Technology Institute), a single pixel features three independently controlled microLED subpixels that are stacked vertically to allow full-colour emission from essentially the entire area of the pixel.

This pixel technology results in a very compact pixel structure and a substantial reduction in the pixel-transfer processing requirement for microLED display fabrication. In addition, the operation of a full-colour micro-display consisting of a pixel-array typically sized smaller than a penny is significantly enhanced due to increased utilisation of the extremely small pixel area.

The picture below shows a diced array of stacked-RGB pixels (left most). R, G, and B subpixels of three adjacent pixels of an array with each subpixel lit up separately (round-robin colours, middle three), and the subpixels of the center pixel lit up simultaneously to 5400K white (right most).

The stacked-RGB pixel technology overcomes a particularly difficult obstacle to mass commercialisation of microLED displays: a need for laborious pixel-transfer processing. Whereas fabricating a microLED display using conventional planar-RGB technologies typically requires transferring discrete R, G, and B subpixels using a pick-and- place process, the stacked-RGB pixel technology substantially or even entirely removes such a requirement.

To achieve the stacked-RGB pixel device, epitaxial and fabrication technologies were developed on sapphire substrates that allow stacking multiple LED junctions where each LED is independently controlled.

This breakthrough multi-junction LED technology is pivotal in enabling full-colour generation. The well-defined spectral peaks of the three colours lend themselves to excellent colour-saturation characteristics and thus a very large colour space.

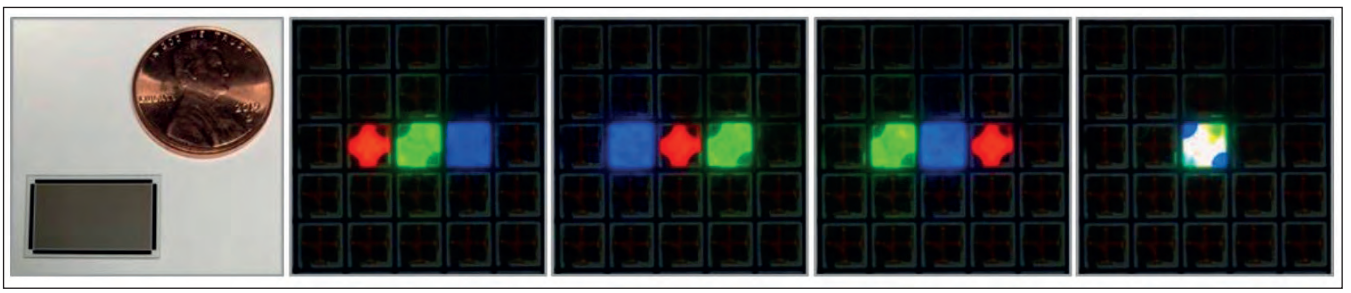

Nitride Semiconductors makes micro UV-LED breakthrough

NITRIDE SEMICONDUCTORS has succeeded in miniaturising the micro UV-LED chip for microLED displays. The company is now developing mass production technology.

Micro LED displays are being developed

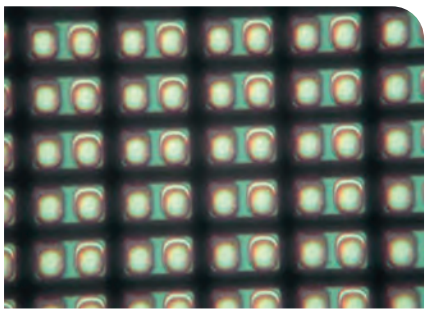

by Apple and many other companies across the world. The red, blue, and green microLED chips that these companies are working on are under development. But it is proving difficult to miniaturise the red LED chip to 50 μm or less.

Nitride Semiconductors' group is taking another approach by developing displays using micro UV LEDs to excite red, blue, and green phosphors. In this way, the cost of micro UV LED chips can be reduced, says Nitride.

Nitride previuously developed 385 nm micro UV LEDs that are 16 μm x 48 μm in size, and have a chip spacing as wide as 10 μm in the horizontal direction and 30 μm in the vertical direction. About 3.4 million chips can be harvested from a 4-inch wafer.

The newly developed micro UV-LED chip has a chip size of only 12 μm x 24 μm, the chip spacing is 5 μm in both the vertical and horizontal directions. About 14 million chips can be obtained from a 4-inch wafer, cutting costs by a factor of four.

When making a 25 mm square size display, 300,000 micro LED chips are required, but 11 micro LED displays can be made from one wafer. All chips are controlled by an IC driver to display a full-colour screen.

Shifting from blue to red for better microLEDs

A NEW microLED developed at KAUST (King Abdullah University of Science and Technology) can efficiently emit pure red light and may help in the quest to develop full-colour displays based on just a single semiconductor.

Full-colour micro-displays can be created by combining red, green and blue (RGB) microLEDs. Now, a KAUST team of Zhe Zhuang, Daisuke Iida and Kazuhiro Ohkawa have worked to develop a more efficient red LED.

The emission colour of an LED is determined by the material properties of the semiconductor. For example, nitride semiconductors can be used to make blue and green microLEDs, whereas phosphide semiconductors are generally used for red light. But combining different semiconductors in this way makes construction of RGB microLEDs more difficult and expensive. Besides, the efficiency of phosphide microLEDs reduces significantly with shrinking chip size.

Red-light emitting InGaN can be created by increasing the material's indium content. But this tends to lower the efficiency of the resulting LED because there is a mismatch between the separation of atoms in the GaN and InGaN, which causes atomic-level imperfections.

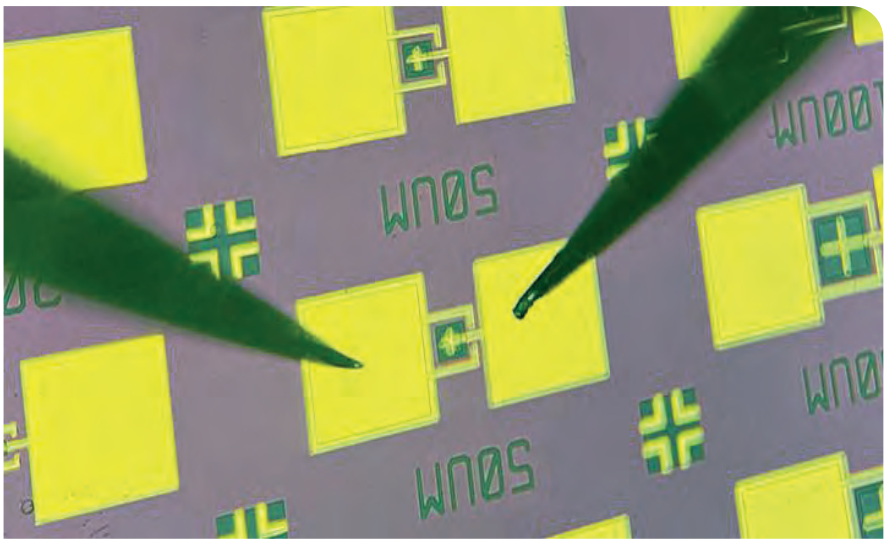

Moreover, damage to the sidewalls of an InGaN microLED induced during the fabrication process makes the new device less efficient. "But we have a chemical treatment to remove the damage and retain the high crystal quality of the InGaN and GaN sidewall interface," explains Zhuang.

Zhuang's team created and characterised a series of square devices with a side-length of 98 μm or 47 μm. Their 47 μm-long devices emitting at a peak wavelength of 626 nm have an external quantum efficiency – the number of photons emitted from the LED per electron injected into the device – of up to around 0.87 percent. According to the team, the colour purity of the red micro-LED is optimum because it is very close to the primary red colour defined by the industrial standard known as Rec. 2020.

"The next step is to increase the efficiency of the red microLED with even smaller chip sizes, maybe below 20 micrometers," says Zhuang.

"Then we hope to integrate RGB nitride-based LEDs for full-colour displays."

'Investigation of InGaN-based red/green micro-light-emitting diodes' by Z. Zhuang *et al*. Optics Letters **46** (2021)

Showa Denko and Infineon sign SiC wafer agreement

INFINEON TECHNOLOGIES has concluded a supply contract with the Japanese wafer manufacturer Showa Denko K.K. for an extensive range of SiC material including epitaxy.

The German semiconductor manufacturer has thus secured more base material for the growing demand for SiC-based products.

"Our broad and fast growing portfolio demonstrates Infineon's leading role in supporting and shaping the market for SiC-based semiconductors, which is expected to grow 30 to 40 percent annually over the next five years," says Peter Wawer, president of the Industrial Power Control Division at Infineon.

"The expansion of our supplier base with Showa Denko for wafers in this growth market marks an important step in our multisourcing strategy. It will support us to reliably meet the growing demand mid to long term. Furthermore, we plan to collaborate with Showa Denko on the strategic development of the material to improve the quality while cutting costs at the same time."

"We are proud to be able to provide Infineon with best-in-class SiC material and our cutting-edge epitaxy technology," says Jiro Ishikawa, Senior Managing Corporate Officer from Showa Denko K.K.

The contract between Infineon and Showa Denko K.K. has a two-year term with an extension option.

SDK expects the contract will enable Infineon to apply SDK SiC materials to various power semiconductor products and the two companies to accelerate improvement in quality of products by bringing together the two companies' knowledge.

Due to the homogeneity in properties and low density of surface defects, SDK SiC epi-wafers materials, which were launched into the market in 2009, have been adopted by electronic device manufacturers as parts of various devices including power supply for servers of cloud computing systems, inverters for railcars and solar power generation systems, and converters installed in quick charging stands for EVs.

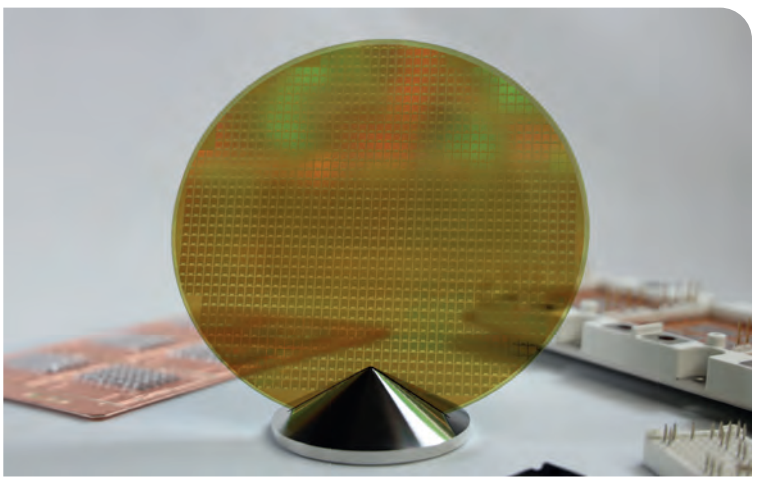

Infineon introduces CoolGaN IPS family

INFINEON has added the new CoolGaN IPS family of integrated power stage (IPS) products to its broad portfolio of wide bandgap power devices. The initial IPS portfolio consists of half-bridge and single-channel products, targeting low-to-medium power applications, including chargers and adapters and switched-mode power supplies.

The 600 V CoolGaN half-bridge IPS IGI60F1414A1L is ideally suited for compact and lightweight designs in the low-to-medium power range. Coming in a thermally enhanced 8x8 QFN-28 package, it enables systems with high power density.

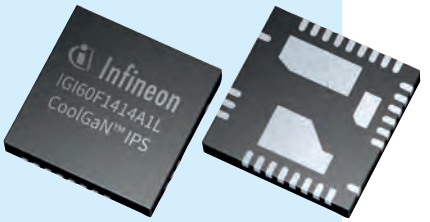

The product combines two 140 mΩ / 600 V CoolGaN e-mode HEMT switches with dedicated galvanically isolated high- and low-side gate drivers out of Infineon's EiceIVER family.

The IGI60F1414A1L is easy to control due to the isolated gate driver with two digital PWM inputs. The integrated isolation function, the clean separation of digital and power ground, and the reduced complexity of the PCB layout are crucial in achieving shorter development time, lower system bill-of-material and lower total cost. The gate driver's input-to-output isolation is based on Infineon's proven on-chip coreless transformer technology. This is said to guarantee high speed and excellent robustness even for extremely fast switching transients with voltage slopes exceeding 150 V/ns.

The switching behavior of IGI60F1414A1L can be easily adapted to the needs of different applications by means of a few passive gate path components. This allows slew rate optimisation to reduce electromagnetic interference efforts, steady-state gate current setting, and negative gate drive for robust operation in hard-switched applications.

Furthermore, due to the system-in-package integration and the highly accurate and stable propagation delay of the gate drivers, the IGI60F1414A1L enables the lowest possible system dead-times.

Navitas to go public

GAN POWER chip company Navitas Semiconductor has entered into a definitive agreement to combine with Live Oak Acquisition Corp. II, a publicly-traded special-purpose acquisition company.

The transaction, which values the combined entity at a pro forma equity value of $1.4 billion, will result in Navitas becoming a publicly-traded company on a national exchange.

Driven by increasing demand for connectivity, electrification away from fossil fuels, and efficient sustainable energy sources, Navitas predicts GaN ICs can address markets estimated to grow to over $13 billion in 2026. Markets include mobile, consumer, enterprise (data centre, 5G), renewables (solar, energy storage) and EV / eMobility.

Navitas GaNFast power ICs integrate GaN power and drive plus protection and control to deliver simple, small, fast and efficient performance. Navitas is in mass production and ramping shipments to many major OEMs and aftermarket suppliers, including Dell, Lenovo, LG, Xiaomi, OPPO, Amazon, Belkin and dozens of others. Over 18 million GaNFast power ICs have shipped, with zero reported field failures.

Navitas has a proprietary process design kit and over 120 patents granted or pending. Navitas estimates that GaN can impact up to 2.6 Gtons of CO_2 reduction annually by 2050.

Gene Sheridan, co-founder and CEO of Navitas, commented: "Navitas was formed with the vision to revolutionise the world of power electronics while addressing significant sustainability challenges for our planet. Not only has Navitas' world-class team invented and patented revolutionary new technology, but we have also overcome all the key hurdles associated with successfully bringing it to market. We are proud to enter the public capital markets with strong operating momentum and investor partners who share our enthusiasm for our long-term mission."

"We are excited to partner with Navitas," said Rick Hendrix, CEO of Live Oak, "This is the most compelling opportunity we have seen in the semiconductor industry, and we are delighted that Navitas' solutions contribute meaningfully to reduced carbon emissions through more efficient power delivery."

"The capital raised through this transaction will allow Navitas to accelerate that vision as they expand from mobile and consumer markets into even more power-intensive applications like data centers, solar energy and electric vehicles – all while delivering a significant CO_2 reduction as part of their net zero initiative."

Navitas was originally funded by the company's management team, along with venture capitalists.

EPC Launches DC-DC demo board for 48V mild hybrid cars

EPC has announced the availability of the EPC9137, a 1.5 kW, two-phase 48 V – 12 V bidirectional converter demonstration board based on its GaN FETs. The board operates with 97 percent efficiency in a very small footprint, according to the company.

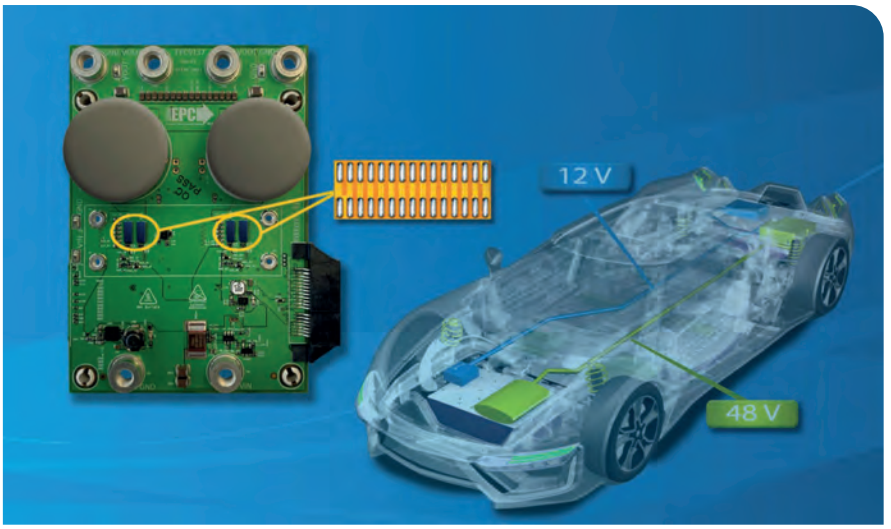

The design of board is scalable; that is, two converters can be paralleled to achieve 3 kW or three converters can be paralleled to achieve 4.5 kW. The board features four EPC2206 100 V eGaN FETs and is controlled by a module that includes the Microchip dsPIC33CK256MP503 16-bit digital controller.

By 2025, one in every ten vehicles sold worldwide is projected to be a 48 V mild hybrid. 48 V systems boost fuel efficiency, deliver four times the power without increasing engine size, and reduce carbon-dioxide emissions without increasing system costs. These systems will require a 48 V – 12 V bidirectional converter, with power ranging from 1.5 kW to 6 kW. The design priorities for these systems are size, cost, and high reliability.

EPC eGaN FETs can operate with 97 percent efficiency at 250 kHz switching frequency, enabling 800 W/phase compared to silicon-based solutions, which are limited to 600 W/phase due to the limitation on the inductor current at 100 kHz maximum switching frequency. By using GaN FETs, it is possible to reduce the number of phases from five to four for a 3.5 KW converter while increasing efficiency. The efficiency of a four-phase GaN converter operating at 250 kHz is 1.5 percent higher than a five-phase silicon MOSFET-based converter operating at 100 kHz.

Overall, the DC-DC converter is three times faster, greater than 35 percent smaller and lighter, and offers greater than 1.5 percent higher efficiency compared with silicon MOSFET solutions. And the overall system cost is less. Additionally, the excellent efficiency and thermal performance of GaN FETs enables air cooling instead of water cooling and the small size of the GaN FETs strongly reduce heat-dissipating aluminium housing for additional system cost saving.

"eGaN FETs provide the fast switching, small size, and high efficiency needed to further reduce the size and weight of 48 V – 12 V automotive power system converters. The demonstrated superior reliability of GaN FETs make them ideal for this very demanding application," said Alex Lidow, CEO of EPC. "The EPC9137 is an ideal example of the capabilities of GaN FETs to increase frequency and efficiency to allow smaller inductance for less phases and higher power density."

The EPC9137 demonstration board is priced at $510.72/each and is available for immediate delivery from Digi-Key.

GaN power market to surpass $1 billion in 2026

THE GaN power market doubled in 2020 compared to 2019 and is poised to surpass the $1 billion mark in 2026, according to Yole Développement. It forecasts that the markets for telecom and datacom, and automotive and mobility, will contribute in the mid- to long-term to overall growth, benefiting from GaN's ascension in fast chargers.

According to Yole, the GaN consumer power supply market will be the main driver, growing from almost $29 million in 2020 to around $672 million in 2026 with a CAGR of 69 percent.

Following the first small-volume adoption of GaN-based power supplies by Eltek, Delta, and BelPower in recent years, Yole expects a larger penetration of GaN. This market is expected to grow from $9.1 million in 2020 to more than $223 million in 2026. "The automotive and mobility market is also paying lots of attention to GaN, following big incentives for the electrification of cars and the interest in increasing driving range through system efficiency optimization," asserts Yole analyst Poshun Chiu.

Players such as EPC, Transphorm, GaN Systems, Texas Instruments and Nexperia are AEC qualified. The major IDM STMicroelectronics, through partnership and acquisition, is also targeting GaN for EVs. Starting from 2022, GaN is expected to penetrate in small volumes in applications such as on-board chargers and DC/DC converters, mainly related to sampling by OEMs and Tier-1s. Yole expects the automotive and mobility market to reach more than $155 million in 2026.

In 2020, the power GaN market doubled thanks to an impressive penetration of GaN devices in fast-charger applications. The adoption of GaN in the smartphone market is fueled by system compactness, high efficiency, and adapter multifunctionalities.

Fast charging is likely to be the killer application for the GaN power device market. So far, at least 10 smartphone OEMs have launched more than 18 phones with an inbox GaN charger. This growth will continue in the aftermarket as well, with companies like Apple, Xiaomi, and Samsung opting for an out-of-the-box charger solution.

In the long term, in cases where GaN has proven its reliability and high-current capabilities at a lower price, it can penetrate the more challenging EV/HEV inverter market and the conservative industrial market.

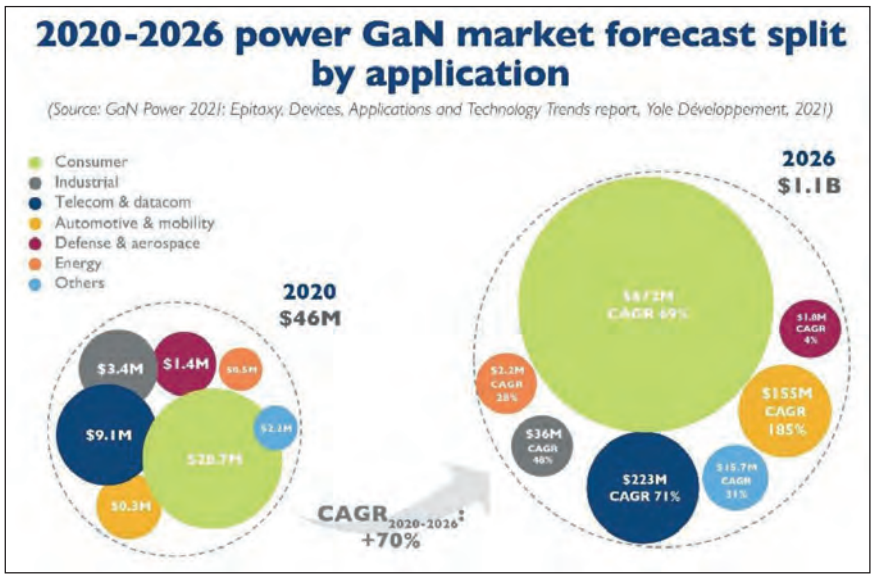

SiC: a glimpse of what's to come

From the industry transition to 200 mm wafers to future market leaders, Omdia analyst, **RICHARD EDEN**, tells **REBECCA POOL** what he expects will happen next

ASK RICHARD EDEN, senior market analyst at Omdia, what single development has had the biggest impact on the SiC market in recent months, and his answer is the launch of Infineon Technologies' CoolSiC CIPOS Maxi: the world's first SiC IPM. This 1200 V integrated power module comprises a six-channel, silicon gate driver with six 1200 V CoolSiC MOSFETs, and signalled the end of the power semiconductor company's roll-out of CoolSiC products in 2020.

As Eden puts it: "Infineon is a juggernaut in power semiconductors, especially in discrete and modules... the IPM was anticipated a while ago and we always thought it was the logical thing to do."

"We will see more companies introducing SiC IPMs – someone like On Semiconductor could be next – but Infineon was the first company to get this finished and released," he adds.

According to Eden, the IPM offers the benefits of SiC MOSFETs, over silicon MOSFETs and IGBTs, but with the ease of having a co-packaged gate drive IC. Critically, the move could see IPMs being used in both domestic and more demanding industrial applications.

But 2020 wasn't all about Infineon. Other key SiC developments centre on new generations of existing products, such as ST Microelectronics' fourth generation 750 V SiC FET platform, as well as increasing production capacities from Rohm Semiconductor, X-Fab, II-VI, and of course Wolfspeed, the Cree company. Indeed, Eden reckons that Wolfspeed's Mohawk Valley Fab in New York state – set to be the world's first 8-inch SiC fab with production scheduled for 2022 – will be a game-changer.

"Many other semiconductor companies buy Wolfspeed wafers, so this facility could trigger the transition from six- to eight-inch wafers in the industry," says Eden. "Also, many companies currently using six-inch wafers are working with old, eight-inch silicon wafer equipment, so the move to the larger wafer size might not be too difficult."

"Yields will be poor with these eight-inch wafers initially, and it will take a year or two to get them back up with this larger wafer size," he adds.

Following the concerted efforts to ramp SiC wafer production, twinned with the many industry supply agreements, Eden believes the SiC wafer supply shortages that plagued the industry from 2017 to 2019 are no longer an issue, for now. Analyst forecasts predict strong market growth, which he reckons could strain future supply levels. "I know companies such as Rohm, Cree and II-VI are trying to add capacity beyond when they think they will need, but we will have to wait and see," he says.

Eden also reckons that the current pandemic hasn't affected SiC market demand significantly. As he says, factories temporarily closed during the first major lockdown of 2020, but have since recovered with the market having also grown slightly. "The silicon semiconductor market was affected more severely than the wide bandgap segment," he says.

Market leaders

Earlier this year, and for the first time, Eden, with Omdia, ranked SiC semiconductor supplier market shares, with intriguing results. For the overall power discrete and module markets in 2018

and 2019, Wolfspeed lead the way, followed by STMicroelectronics, Infineon, On Semiconductor and Rohm. But individual markets tell a different story.

For the SiC rectifier device market alone, Wolfspeed had a huge lead in 2018 and 2019, capturing more than 40 percent of the market in each year. "Wolfspeed has this lead as it is the oldest incumbent supplier of SiC discretes and has been making these for around 20 years," highlights Eden.

Infineon, Rohm, STMicroelectronics and On Semiconductor are ranked from second to fifth, respectively, again largely thanks to longevity. "Rohm is the newest supplier here, but is now established as it was the first company to mass produce SiC MOSFETs in 2010, and developed SiC rectifiers around the same time," says Eden.

For the SiC power MOSFET markets, STMicroelectronics takes pole position, which according to Eden, is thanks to the company's significant contract to supply devices to battery electric vehicle maker Tesla. Yet despite its Tesla-dependence, he expects ST to maintain its lead.

"Currently at fourth place [following ST, Wolfspeed and Rohm] Infineon would surely be higher in the rankings if it hadn't focused on developing SiC JFETs ahead of SiC MOSFETs during the mid-2010s," he adds.

The full SiC and hybrid SiC module markets are quite different, with Mitsubishi Electric and On Semiconductor leading each market in 2019, respectively. Germany's Vincotech, Infineon and Wolfspeed also feature highly in these markets with China players, StarPower Semiconductor and CRRC Times Semiconductor also creeping into the module supplier top five. According to Eden, many module suppliers tend to buy in power transistors and assemble these in the modules, rather than manufacture the chips themselves. Also, going forward, he expects to see more China-based players enter module markets.

"StarPower is currently the largest Chinese power module manufacturer and is still quite small in the overall market," he says. "But there are other Chinese companies making power modules, such as Nanjing SilverMicro Electronics, and they will become more prevalent following the Chinese government drive to use silicon carbide for electric vehicles and industrial applications."

Eden doesn't believe that the US-China trade tensions have particularly affected SiC markets. But as he adds: "There is a trend for Chinese companies to want to buy Chinese products, as strongly influenced by the Chinese government."

So what does Eden reckon will come next for the SiC market? For the time being, he expects that SiC technology will happily co-exist with silicon and GaN, capturing market share from silicon throughout this decade. However, beyond this timeframe may bring change. "Right now, the GaN market is behind the SiC market but ultimately, by the 2030s and 2040s say, GaN might end up having the larger market share than SiC."

Indeed, thanks to wafer costs, Eden reckons GaN will always offer a cheaper alternative to SiC. He highlights how a silicon wafer currently costs around $10, a GaN-on-silicon wafer costs around $25, while SiC wafers currently come in at up to $1000 each.

"The substrate is the determining factor for device cost, and I think that once yield, design and other challenges are resolved, GaN devices are going to be cheaper than SiC devices," says Eden. "Silicon carbide is potentially a stop-gap during the 2020s until gallium nitride is perfected – and then this technology will have a bigger market for discrete power semiconductors and power modules." However, at the end of the day, silicon, SiC, GaN and, potentially, Ga_2O_3 semiconductors will co-exist, each occupying market niches, depending on the cost and performance benefits that each offer.

World supplier market share estimates for SiC discretes and modules in all applications

(percentage dollar revenues)

		2018	2019	DIFF '19 vs. '18
1	Wolfspeed	26.0%	22.3%	-3.8%
2	STMicroelectronics	14.2%	21.7%	7.4%
3	Infineon Technologies	10.7%	10.5%	-0.2%
4	ON Semiconductor	4.6%	6.7%	2.1%
5	ROHM Semiconductor	8.4%	5.8%	-2.5%
6	Vincotech	6.2%	5.3%	-0.9%
7	Mitsubishi (inc. Powerex)	4.0%	3.8%	-0.3%
8	CRRC Times Semiconductor (inc. Dynex)	2.0%	1.4%	-0.5%
9	Microchip	1.6%	1.3%	-0.3%
10	Fuji Electric	0.8%	0.9%	0.1%
	Others	21.4%	20.4%	-1.0%
	Total Market Size =	**$703.1 million**	**$888.1 million**	
	Growth Rate =		**26.3%**	

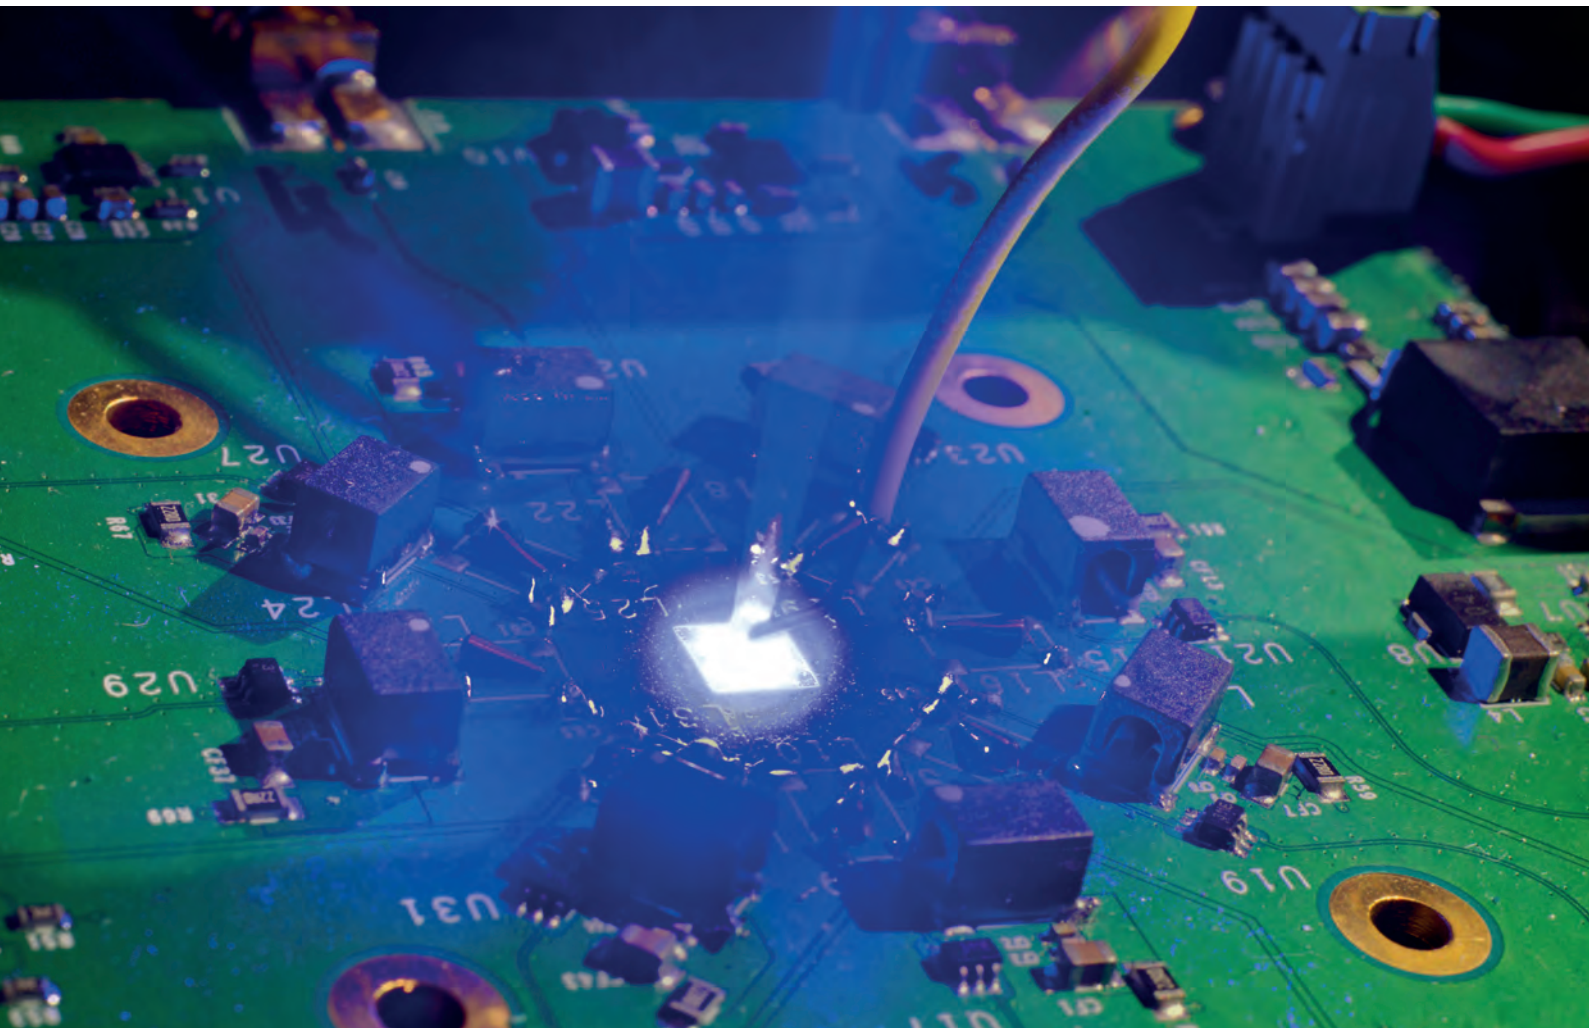

Easing the chip-to-chip communication bottleneck by leveraging microLED display technology

High-speed optical emitters derived from GaN-based microLED displays can move data at much higher density and lower power than copper, bringing optical connections to the centimetre scale

BY BARDIA PEZESHKI, ROB KALMAN, ALEX TSELIKOV AND CAMERON DANESH FROM AVICENA

MOST OF THE ENERGY consumed in computing systems is not in the computation, but in moving data, and the longer the distance, the greater the challenge in terms of energy and density. At longer length scales, fibre optic links have replaced copper, but at short distances the significant amount of energy required to convert data back and forth between photons and electrons makes optical interfaces prohibitive.

Although it may raise a few eyebrows, at these shorter length scales, optimized optical emitters derived from GaN microLEDs could be a promising candidate for optical communications by leveraging their success in the display industry. Such a move could transform the $400 billion computer hardware industry and enable

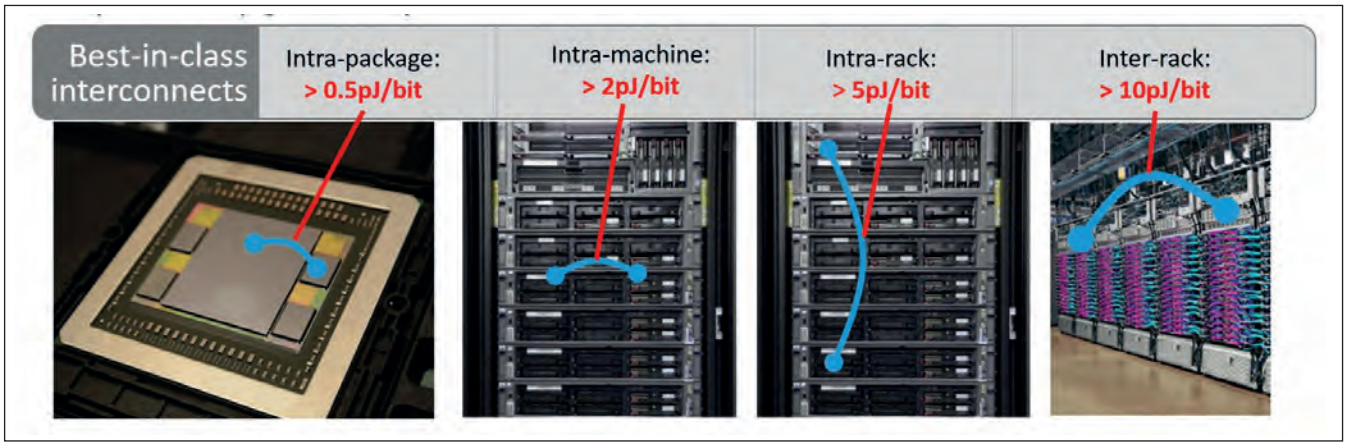

➤ Figure 1. Moving data takes energy. The further you go, the more pico-joules you'll need to move each bit.

entirely new architectures for parallel computing, machine learning, and processors.

Within the semiconductor industry, the days of enjoying rapid rates of progress on multiple fronts are long gone. There has been little increase in chip clock rates during the last two decades. There are also limits on the number of IC package input and output pins. Consequently, almost all high-performance chips utilize high-speed Serializers/De-serializers – known as SerDes – for input/output on the periphery of the die. Their role is to dramatically increase the bit rate compared with on-chip clock speeds so that all the information can be to squeezed through a limited number of pins. This takes energy, real estate, and chokes data flow – and the situation is only going to get worse with future system advances being realised primarily through new architectures that interconnect more chips rather than improvements in raw transistor performance.

Basic physics accounts for the decline in performance of longer electrical interconnects, which are limited by resistance and capacitance of electrical lines. When lines are longer, more capacitance needs to be charged and discharged through their resistance. Increasing density is also problematic – making wires thinner and closer together increases both resistance and capacitance, decreasing maximum interconnect lengths. And that's by no means the only issue to consider. There's crosstalk, non-linearities in the dielectric constant of the material causing distortions in the waveform, and the skin effect, which causes an increase in electrical resistance at higher frequencies.

By packaging chips on silicon interposers with a far higher density of lines, electrical data-pipes can be connected to very wide busses and operate close to the clock speed of the chips. However, for reasons already outlined, a high density limits the reach, with chips typically having to be placed edge-to-edge. It is common to co-package high-bandwidth memory chips adjacent to the processor (see the leftmost illustration in Figure 1), and communicate over a bus that is typically about 1000 lanes wide, running at only 1 or 2 Gbit/s. Note that Intel's AIB bus, TSMC's LIPINCON bus and the Open Compute Project BoW busses are all wide and slow, with hundreds or thousands of lanes, each operating at 1 Gbit/s to 16 Gbit/s.

Different optics

It has been known for decades that even very short optical interconnects promise significant power and density advantages over electrical interconnects. This advantage, based on 'quantum impedance transformation', hinges on the use of low-capacitance,

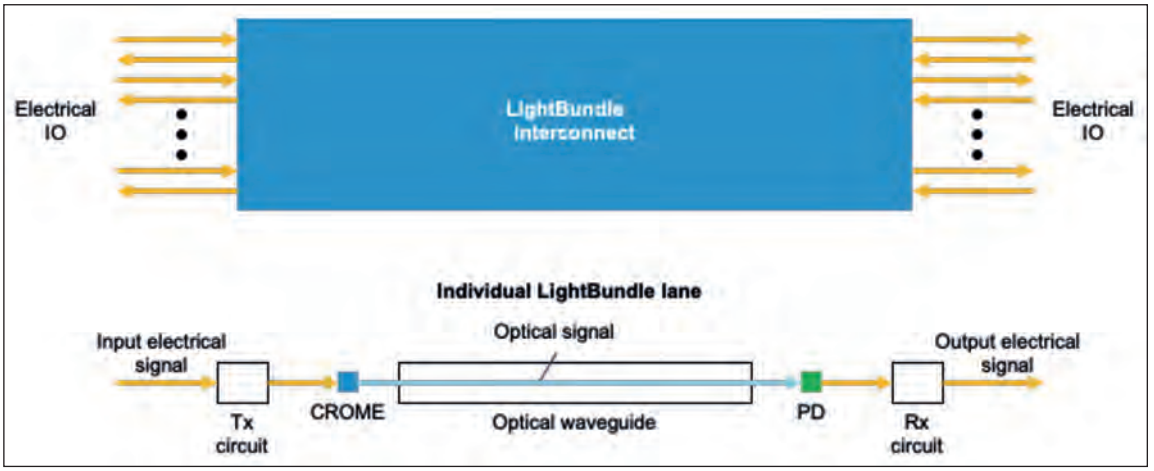

➤ Figure 2. Avicena's LightBundle interconnect with details of a single lane.

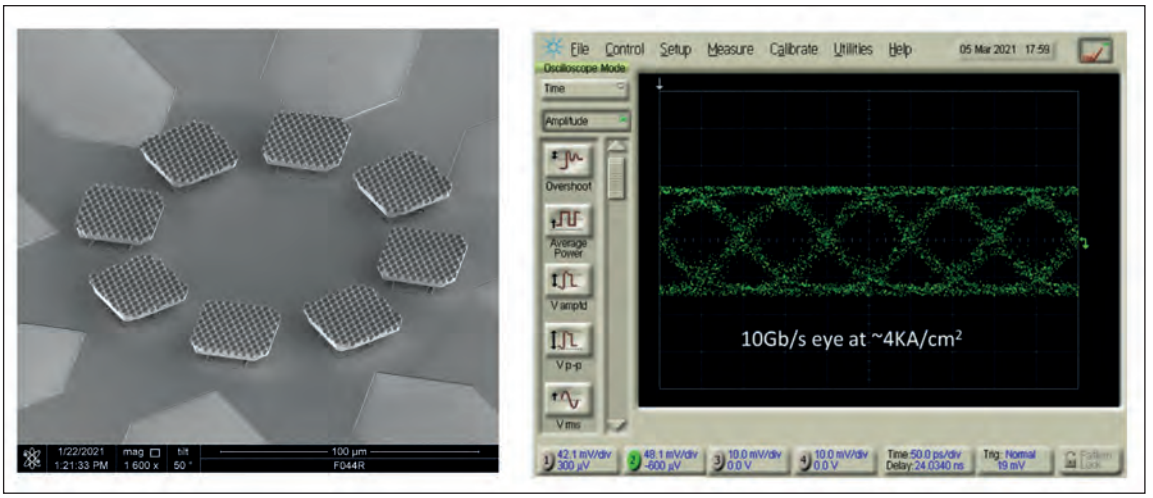

➤ Figure 3. Lifted-off high-speed light emitters on silicon, and open eyes at a 10 Gbit/s modulation rate. Large arrays can reach very high densities of data transfer.

high-quantum-efficiency optoelectronic devices. Unfortunately, despite decades of effort, such sources are still to emerge for short links. Edge-emitting lasers, VCSELs, quantum-well modulators and silicon photonics modulators all fall far short of what is needed for practical short optical interconnects (from 1 mm to 10 m). A litany of issues has prevented success, including a high drive power, large size, poor high-temperature performance, expensive and bulky packaging, and a poor yield.

These challenges are not insurmountable. What's needed is a very different approach; our team in Mountain View, California is pioneering a new type of optical emitter derived from high-speed microLEDs. These devices that we call CROMEs (cavity-reinforced optical micro-emitters) utilize very innovative epitaxial and device structures to achieve far shorter carrier lifetimes, and a much higher modulation bandwidth than their more common cousins deployed in ubiquitous lighting and display applications. We have demonstrated CROMEs that are fast enough for current and future high-performance IC interconnects.

A significant advantage of using visible light is that it allows the fabrication of low-capacitance, large-area CMOS detectors, which can be integrated with simple amplifier circuits to form fast, extremely low power receivers. Thanks to the large area of the detectors, alignment is simple while the packaging cost is low.

There are several options for moving the light between chips. This can be carried out with monolithically fabricated waveguides on an interposer or with various kinds of multicore and imaging fibres. The use of high-density multicore fibres is well established, having been employed for decades in imaging applications, and used in borescopes, medical endoscopy and other 'display' applications.

Operating in the visible spectrum with optimized CROMEs and CMOS-compatible photodetectors that

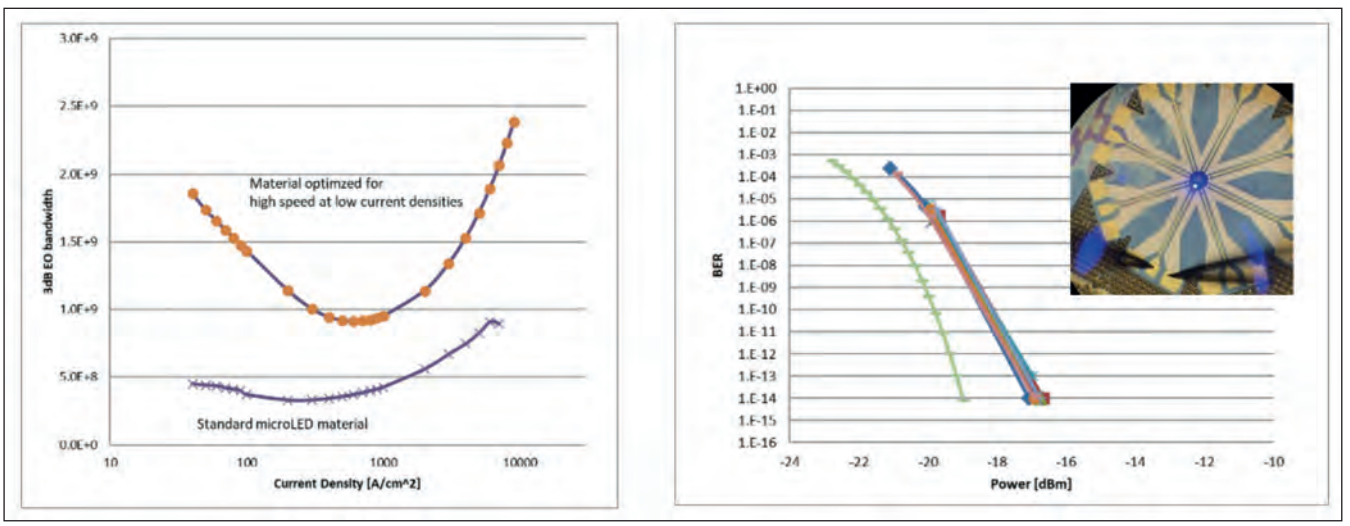

➤ Figure 4. CROMEs have a higher 3dB bandwidth than their conventional cousins, and can obtain high modulation speeds, even at low current densities. On the right, eight channels modulated at 1.25 Gbit/s show no error floors. The inset photo is a probe test of a single element of the eight-channel array.

are integrated with amplifiers, virtually all capacitance is eliminated. We estimate that such links can deliver an energy efficiency below 100 fJ/bit and deliver a multi-Tbit/s throughput at a density in excess of 10 Tbit s^{-1} mm^{-2}.

A block diagram of our 'LightBundle' approach, together with details of a single lane, is shown in Figure 2. Each lane consists of an optical transmitter, a fibre core (or waveguide), and a receiver. The CROME transmitters are powered by simple drive circuitry, driven between 20 μA and 500 μA and optically coupled to a waveguide. Each light emitter has a diameter between 1 μm and 10 μm and may be modulated up to around 10 Gbit/s. The receiver consists of a silicon photodetector, optically coupled to a waveguide and monolithically integrated with CMOS transimpedance and limiting amplifiers.

There has already been considerable interest in high-speed visible wavelength LEDs for free-space optical interconnects and 'LiFi', where LEDs are used to generate ambient lighting and transmit data. Such LEDs have a range of criteria to fulfil, needing to be both very efficient and fast. But in a chip-to-chip application considerations are markedly different, and one can trade off quantum efficiency for speed. In fact, the benefits of doing so are tremendous. By optimizing the device's doping, quantum well structure, device design, and other features, CROMEs can be fast enough for very high-density, high-performance IC interconnects. We have demonstrated CROME-based links with wide open eyes at even 10 Gbit/s (see Figure 3), a speed previously accessible only with high-speed lasers or modulators.

Other factors that increase speed at very low current densities are coulombic enhancement, microcavity effects, and non-radiative recombination. All these help the CROMEs achieve high modulation speeds.

Figure 4 shows some CROME experimental results. At very low current densities the 3dB bandwidth reaches about 2 GHz (or around 4 Gbit/s for on-off modulation). Good link performance is realised at drive currents of a few tens of microamps. This allows the power consumption per bit to be far lower than

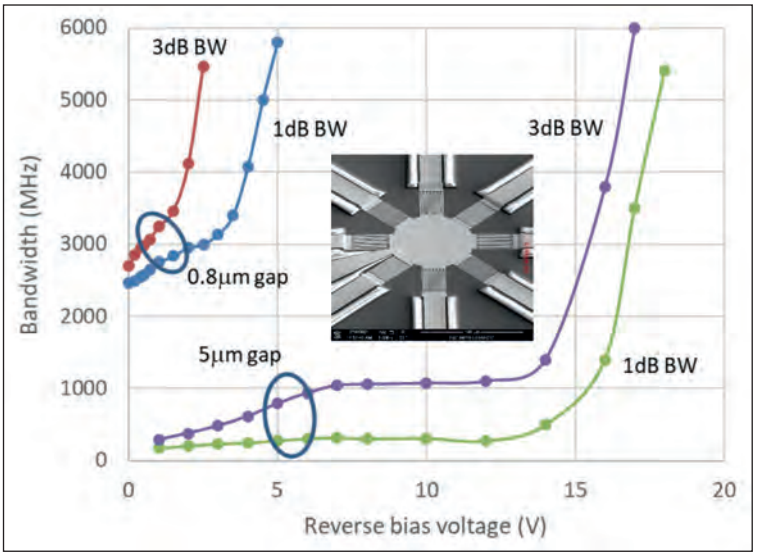

➤ Figure 5. The frequency response of eight-element CMOS-compatible detectors – data is limited by the 6 GHz bandwidth limit of the network analyzer used.

that for lasers. Note that even very small VCSELs, the most energy-efficient lasers, typically have a threshold current of a milliamp.

Using blue light for data transmission is incredibly advantageous at the receiver. Silicon is a nearly ideal material for detecting blue light, with an absorption length of just 0.2 μm. This enables CMOS-compatible fabrication of receivers with integrated photodiodes. These photodiodes with very low parasitic capacitance (less than 10 fF) allow use of innovative, simplified transimpedance amplifier designs, enabling receiver power dissipation of below 50 fJ/bit. Such a low power is not possible with 'typical receivers', which are hampered by a much higher capacitance and greater complexity.

Such a detector structure can be fabricated by using the source and drain implant/diffusions for the CMOS transistors to make lateral *p-i-n* diodes. Our eight-element array, used for the multi-channel measurements of Figure 4, shows a bandwidth of over 6 GHz (see Figure 5). The shadowing by the

Using blue light for data transmission is incredibly advantageous at the receiver. Silicon is a nearly ideal material for detecting blue light, with an absorption length of just 0.2 μm. This enables CMOS-compatible fabrication of receivers with integrated photodiodes

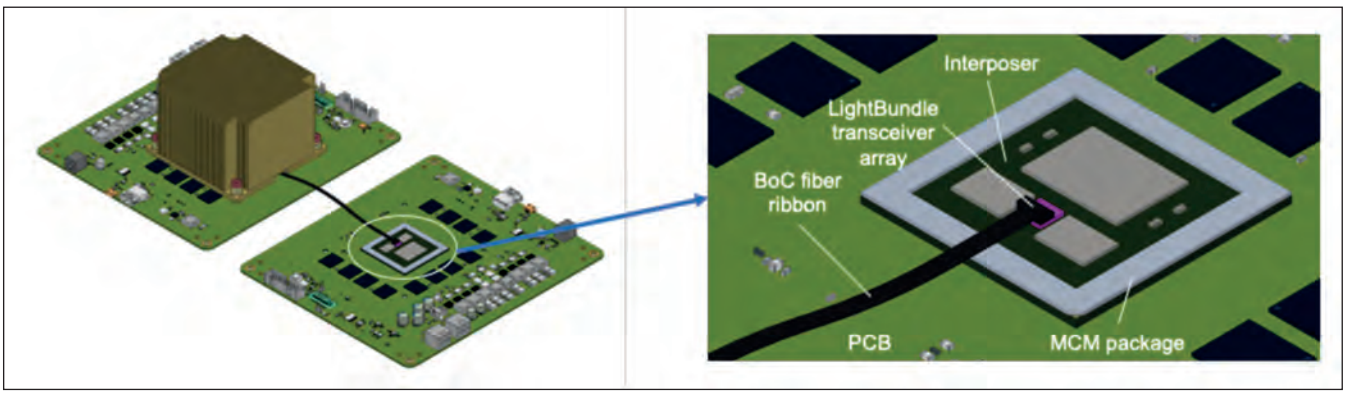

➤ Figure 6. A fibre-based architecture, showing two MCMs on different PCBs connected by a ribbon fibre interconnect. More advanced versions support interfaces directly to the surface of a complex IC.

fingers on the detector and a thin SOI substrate limited quantum efficiency to about 50 percent, but quantum efficiencies of more than 90 percent should be achievable in optimized devices.

Perfecting the package

We use highly multimode fibres and waveguides to realise efficient optical coupling to the sources and detectors. This relaxes alignment tolerances. One of the merits of this approach is that it allows us to select between a variety of useful packaging architectures, each with a different range of benefits.

Our fibre-based links are optimized for longer interconnects, spanning distances from 10 cm to 10 m. These links transfer data between chips and multi-chip modules at the board, shelf and rack levels. Transmitters are connected to receivers using high core-count or imaging fibre (see Figure 6).

A typical imaging fibre has hundreds or thousands of cores, each with a diameter ranging from 2 μm to 20 μm. Manufacture involves standard 'stack and draw' optical fibre fabrication techniques. Optical channels can be sent through this fiber on a square grid with 20 μm centre-to-centre spacing, with each carrying 4 Gbit/s. This gives an areal interconnect bandwidth density of 10 Tbit s^{-1} mm^{-2}. Sending 256 of these channels through a fiber provides a total throughput of 1 Tbit/s. As such a fibre is less than 500 μm in diameter, multi-fibre ribbons and cables capable of carrying more than 10 Tbit/s are very compact and flexible. That's in dramatic contrast to bulky DAC twinax.

Since all our processes are fully CMOS compatible, interconnects can also be made directly from the surface of a large complex IC. This direct 'optical pin-out' enables ultra-dense, ultra-low-power optical interconnects directly from anywhere on an IC. Operating with an interconnect density of 10 Tbit s^{-1} mm^{-2} and an energy per bit of just 100 fJ/bit, these links could provide outputs and inputs for a 100 Tbit/s (bidirectional) switch IC from a 20 mm^2 footprint, and a power consumption of just 20 W.

For distances of just 1 mm up to 10 cm, data can be routed through lithographically formed waveguides. An array of multimode waveguides, formed from SiO_2, SiN or polymers, can be fabricated on a planar substrate. Deploying these 'optically-enhanced' interposers greatly increases practical interconnect distances while decreasing power consumption. The use of lithographic registration enables small sources, which have a typical diameter below 2 μm, to be efficiently coupled into optical waveguides that are only slightly wider than the CROME diameter.

Compelling credentials

The strengths of our CROME interconnect technology go well beyond performance. They also fulfil all the

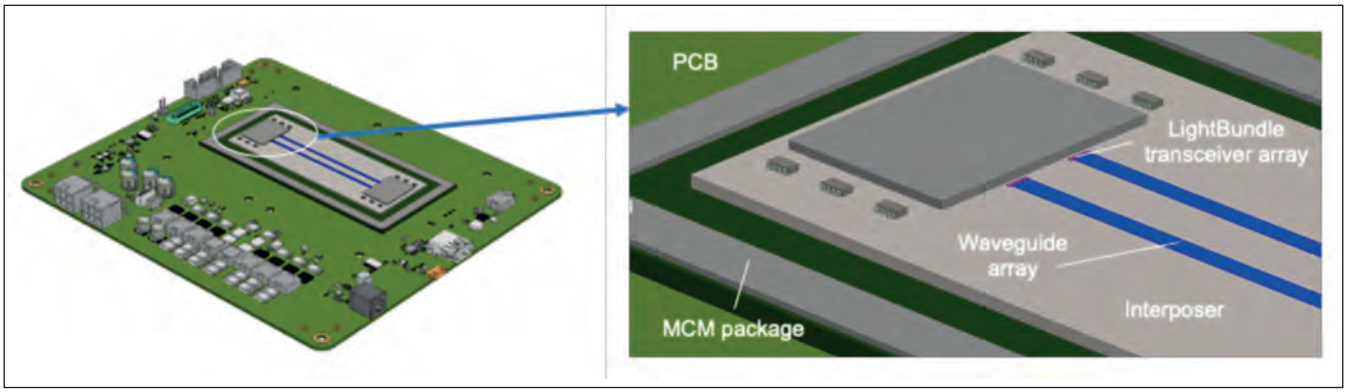

➤ Figure 7. Potential waveguide architecture on an 'optically-enhanced' interposer.

other requirements for practical high-volume products: they have a high reliability, they deliver excellent high-temperature operation, they are low in cost, easy to package, and they are compatible with existing high-volume manufacturing and test capabilities.

Our links' optical interfaces are formed from a variety of technologies that are used in very high volumes in different applications. From a manufacturing perspective, CROME arrays can be thought of as very small GaN microLED displays. As regular readers know, GaN is manufactured in huge volumes for solid-state lighting and power devices and is becoming increasingly important for displays. Production of this device is underpinned by a massive ecosystem that supports high volume, low-cost manufacturing. In general, GaN is a far more reliable and robust material system than other III-Vs, such as GaAs and InP, due to its excellent high-temperature performance and insensitivity to defects. Unlike devices made from other III-Vs, those built from GaN can operate at high temperatures, such as 150 °C, with extremely low failure-in-time rates.

We are able to draw on existing high-volume manufacturing processes used in the lighting and display industries for the mass transfer of thousands to millions of CROMEs from a sapphire source wafer to a target silicon CMOS driver wafer containing transceiver circuitry. Compared to display requirements, the data interconnect is very undemanding: there are fewer than tens of thousands of lanes, and the application is insensitive to colour and brightness variations, both critical concerns for display makers. Another positive for us is that some redundancy can be built in with spare channels, which is also done in wide slow electrical busses. Our transceiver circuitry and our photodiodes are manufactured using standard CMOS. We can use older process nodes, thanks to the modest link speeds. For coupling we can use polymer micro-optics, similar to those in smartphone cameras.

By leveraging all these high-volume technologies, we estimate that the costs of these Tbit/s wide parallel GaN-based links are one to two orders of magnitude lower than the cost of other high speed serial optical link technologies, such as pluggable optical transceivers, board mounted optics, or even silicon photonics co-packaged optics aimed for datacenter and computing environments. Although our technology is multimode and limited to 10 m reaches, it clearly promises to have an important role to play in advancing the performance of computing systems.

OPPORTUNITIES FOR CROME-BASED LINKS

CLOUD-SCALE data centres and high-performance computing are dominated by interconnects. To ease the data bottleneck, a great deal of attention and investment has been directed at the co-packaging of silicon photonics with ASICs, GPUs and memory.

CROME-based interconnects offer a complementary technology to silicon photonics. While silicon photonics may provide an excellent solution for interconnects longer than 10 m, CROME-based links offer a far more dense, lower-cost solution at smaller length scales.

CROME-based interconnects could provide welcome solutions to various problems. As switch ASICs in routers and switches move to more than 100 Tbit/s total throughput, CROME-based interconnects could move the data from the switch ASIC to the chassis front plate, where standard ethernet transceivers could be deployed, and down the rack to the NIC cards in the servers. In the move towards disaggregated infrastructure, CROME-based links could provide low latency, low power connections to shared memory or connect processor clusters. These high-density links increase connectivity for both learning and inference in machine-learning and deep neural networks.

Looking further ahead, there are opportunities for optical interfaces to be implemented directly on large complex ICs, where they will provide the best power, latency, density, and cost. In decades gone by, the fundamental density and power limitations of electrical interconnects compelled the migration to SerDes-based input/output (I/O) for high-speed interfaces. Now these interfaces consume significant IC power and real-estate, and low power optical interfaces hold much promise. Introducing them could unlock the door to a new generation of ICs incorporating 'lightweight' high-performance IOs with greatly improved density, latency, and reach while consuming far less IC real estate and power.

In the latter part of the twentieth century, numerous groups in academia and research labs looked to photonics to revolutionize computing, pointing out that optical beams enable very high connectivity.

In this vein, a team at AT&T Bell Laboratories demonstrated simple optical computers using quantum well modulators with gratings and holograms. But despite massive funding and huge interest from DARPA and other resources, the technology never matured – perhaps because the components and the applications needed to make these breakthroughs did not exist. Today that's not the case. The technology and the need are now both there, and perhaps the dreams of the visionaries can be realized with these new CROME devices.

FURTHER READING

- D. A. B. Miller J. Light. Technol. **35** 346 (2017)
- B. Pezeshki *et al*. "High speed microLEDs for visible wavelength data communication," Proceedings of SPIE, vol 11706 Light-emitting devices, materials, and applications XXV, 117060N, 2021.
- S. Rajbhandari *et al*. Semicond. Sci. Technol. **32** 023001 (2017)

CS shares race ahead

Stocks in the CS industry soar to well above pre-pandemic levels

BY RICHARD STEVENSON

WHILE THE PANDEMIC has taken its toll on all of us, it's been hardest on the poor. Those with bigger salaries are less likely to have lost their jobs, and if they've had a chance to invest any spare cash, they are probably enjoying a good return.

Canny investors will have considerably strengthened their financial position over the last year or so by buying shares in technology companies. As a glance at our annual Shareprice Leaderboard shows, over the 12 months up to the end of April 2020 the tech-heavy NASDAQ climbed by just over 60 percent, while some companies in our industry have fared even better (see Shareprice Leaderboard on p. 21). Topping our list is Cree, thanks to a valuation that has leapt more than three-fold. Substrate maker AXT is not far behind, and even third-placed II-VI has delivered a return over this timeframe of around 140 percent.

It's worth noting that all these impressive performances are not a simply a result of nosedives at the start of the pandemic followed by recoveries throughout last summer to regain the lost ground.

➤ Construction of Cree's Mohawk Valley facility is well underway. This fab is the centrepiece to Cree's massive expansion plans.

Compare the share prices at the very beginning of last year to where they are at the end of April 2021, and you'll see that all but one of the companies we cover are enjoying an increase in valuation, with five doubling their worth over that timeframe (see the second Shareprice Leaderboard, on p. 23, that considers the period from the beginning of 2020 to late April of this year).

Champion Cree

Cree's soaring share price reflects the success it has had in transforming its business. It has now sold-off its LED chips and lighting businesses, which brought declining returns over the last decade or so, and has turned its attention to making SiC substrates, SiC power devices and GaN-on-SiC RF products. Marketed through the moniker Wolfspeed, sales of these products are ramping, and could continue to do so throughout this next decade and beyond as the company executes a massive capacity expansion programme to produce higher volumes more efficiently.

The company still has a long way to go along this road. Sales of its revamped portfolio are growing, with revenue for the third fiscal quarter of 2021, ending on 28 March, netting $137.3 million. This is above guidance, an 8 percent increase sequentially, and up 21 percent compared with the equivalent fiscal quarter of the previous year. However, due to a combination of Covid-19 safety measures and a temporary increase in higher factory costs, gross margins are heading in the wrong direction, having fallen 0.4 percent sequentially, to 35 percent. This will have contributed to the loss for the latest fiscal quarter, which is $24.7 million, according to non-GAAP accounting.

All figures should improve substantially over the coming years. Gross margin is tipped to climb towards 50 percent, thanks to the launch of the 200 mm Mohawk Valley Fab that will swell shipments. Construction of this fab accounts for much of the $550 million that Cree is investing in capital expansion in fiscal 2021. Helping to pay for this is the sale of the LED business that recently netted around $500 million.

During a third fiscal quarter earnings conference call, held on 28 April 2021, Cree's CEO Gregg Lowe

Compound Semiconductor Shareprice Leaderboard

Rank	Company	Ticker	Share value, April 27, 2020 ($)	Share value, April 27, 2021 ($)	% appreciation	Change in Rank
1	Cree	CREE	38.58	119.76	210.4	+14
2	AXT	AXTI	4.14	11.61	180.4	+11
3	II-VI	IIVI	31.10	74.56	139.7	+9
4	Veeco	VECO	10.24	23.29	127.4	+6
5	Qorvo	QRVO	90.93	200.36	120.3	-2
6	Aixtron (Frankfurt)	AIX	11.04*	22.89*	107.3	+3
7	Skyworks	SWKS	97.80	201.18	105.7	-2
8	IPG Photonics	IPGP	122.44	234.06	91.2	+6
9	IQE (London)	IQE	51.6*	88.05*	70.6	+7
10	Infinera	INFN	6.01	9.80	63.1	-9
11	NASDAQ composite	IXIC	8717.98	14170.91	62.5	-5
12	Emcore	EMKR	4.14	6.69	61.6	-4
13	WIN Semiconductor (Taipei)	3015.TWO	9.00*	13.30*	47.8	-9
14	Riber (Paris)	RIB	1.38*	1.82*	31.9	-3
15	Rubicon	RBCN	8.17	10.54	29.0	-8
16	Lumentum	LITE	79.24	94.90	19.8	-14

* * Converted to dollars using the exchange rates on 27 April of 1 EURO = 1.2079 USD, 1 GBP = 1.3910 USD and 1 TWD = 0.035839

➤ It's been a great 12 months for Compound Semiconductor stocks, with a rise in valuation across the board.

offered some insight into the latest trends in the markets where the company's products serve. "We are seeing more demand in our core automotive and RF markets, as well as additional interest in new areas across energy, industrials, and aerospace and defence," revealed Lowe. Diversity of deployment is impressive, including devices designed into forklifts, 3D laser printers, air conditioners, motor drives, robotic arms, flying vehicles and a beer truck.

Lowe added that the company's device pipeline now stands at more than $10 billion, with new opportunities being identified at a rapid pace. In just the last quarter, the company's sales team secured more than $580 million in design-ins; and over the last five quarters this figure totals $2.5 billion.

Over the next few years, Cree is set to benefit from the $2 trillion infrastructure plan recently unveiled by the US administration. Within the plan, $100 billion is dedicated to increasing broadband access, with a special emphasis on 5G infrastructure. Remarking on this, Lowe commented: "This development, combined with strong sales of 5G smartphones during the pandemic, underscores how 5G is continuing to gain momentum and offers a global opportunity in the years ahead."

Also included in the infrastructure plan is significant funding for electric vehicles, including sales rebates, tax credits and charging stations. "We anticipate this will have a significant impact on the adoption of electric vehicles," reasoned Lowe. "We are now seeing US automakers make big commitments to ramp their EV efforts. For instance, General Motors and LG Chem recently announced plans to invest $2.3 billion to build a battery cell plant to support the automaker's efforts to expand its electric vehicles."

Cree could also benefit from a ramp in electric vehicle production in Europe. On a recent trip to Germany, Lowe had very positive conversations with a number of tier one carmakers, who are struggling with supply constraints associated with silicon chips. He offered them hope, telling them that Cree is going into production with the world's largest SiC fab, right in the middle of this supply issue. This great timing, which Lowe openly admits is more down to luck than judgement, bodes well for Cree's future. "[It] is certainly a very nice light at the end of the tunnel for some of these guys as they start placing bets on silicon carbide."

Until the Mohawk fab is up and running, sales are likely to show a moderate increase, rather than a massive ramp. During the next quarter revenue is tipped to rise to between $142 million and $148 million, with the majority of the increase coming from the power and RF device product lines. A small uptick is also anticipated from the materials division.

AXT accelerates

Throughout most of last year, AXT's share price hovered at around $5, before leaping to almost twice that in November, and surging as high as just over $15 in late February. The company's valuation has dropped back a little since then, but sales continue to climb. They leapt from $22.1 million for the second fiscal quarter of 2020 that ended on 30 June 2020 to $25.5 million and $27.0 million in the third and fourth fiscal quarters that followed, and have broken the $30 million barrier in the most recent quarter.

This growth is attributed to an increase in sales to chipmakers serving a broad range of markets, from 5G telecommunications and its related technologies to datacentre connectivity, LED-based sensing and display, healthcare monitoring and consumer devices.

Of the $31.4 million netted in the most recent quarter, $23.4 million came from substrate sales and $8 million from the two joint-ventures: BoYu, a manufacturer high-temperature crucibles and tools for OLEDs; and JinMei, a diversified industrial supplier of high-purity materials.

➤ Earlier this year II-VI has launched double-junction VCSEL arrays, targeting next-generation world-facing 3D sensing applications.

As well as growing revenues, AXT is reducing reliance on its biggest customers. Company CFO Gary Fisher highlighted this welcome move in a call discussing earnings on 28 April, 2021, when he revealed that no customers accounted for 10 percent or more of overall revenue, and the top five contributed just 26 percent to total sales. "Usually, we do have at least one 10 percent customer," added Fisher. "And usually, the top five customers contribute approximately 35 percent to 40 percent of total revenue."

Gross margin is on the rise, reaching 36.8 percent in the latest fiscal quarter, up 2.9 percent sequentially and 10.2 percent year-on-year. Thanks to this, plus a favourable product mix and and increasing sales, net profit hit £3.4 million, up $1.3 million sequentially.

During a recent earnings call, AXT CEO Morris Chang discussed sales for all three types of substrate produced: InP, GaAs and germanium.

Shipments of InP are seeing strong growth. "We saw continued strength from 5G and its related technologies," said Chang, who did not know whether these wafers were for optical connections for 5G equipment, or technologies such as passive optical networks that support 4G and 5G functionality. "But from our perspective, any modernization of telecom infrastructure that utilizes indium phosphide is positive for our business."

Demand for GaAs substrates has recently increased, due to sales to makers of LEDs and wireless products. For this class of substrate, revenue for the first fiscal quarter 2021 exceeded all quarters since the first fiscal quarter 2018. Sales in the coming quarter will be helped by strong demand by LED makers, but hindered by a softening of sales to those making wireless products.

AXT is continuing to invest in both its InP and GaAs technology. It is developing 6-inch InP substrates and has just unveiled 8-inch GaAs substrates for LED applications. Commenting on the latter, Chang remarked: "This is no longer a test tube laboratory program, because we are now shipping wafers according to our customer specifications. This is a tremendous step for AXT, and we are very proud of our team." It is hoped that the large substrate size will help to provide scale and efficiency for very high volume manufacture of VCSELs for 3D sensors and LiDAR, and microLEDs for displays."

For germanium substrates, revenue for the recent fiscal quarter is down slightly. However, growth is expected this year, thanks to a strong satellite market. This should provide a modest increase in the next quarter.

For that quarter – the second fiscal quarter of 2021 – AXT predicts that contributions from germanium, plus growth in InP and GaAs revenues, will propel sales to between $30.5 million and $31.5 million. Profit is expected to be in the range $2.6 million to $3.4 million.

Third for II-VI

II-VI, third on our Leaderboard, has seen its share price follow the same trend as AXT: broadly flat to last November, climbing to peak in February, and now a little below that high point.

This multi-national powerhouse, with products that include GaAs electronic devices, SiC substrates, laser systems and components for optical networks, claimed several records in its last quarterly results, announced on 9 February, 2021. Quarterly sales broke new ground, hitting $787 million, backlog reached a record $1.08 billion, and cash from operations climbed to a new high of $221 million.

The management at II-VI are very pleased with the integration of Finisar into the existing business. Discussing this in a second fiscal quarter earnings call, held on 9 February 2021, company CEO Chuck Mattera revealed that progress on this front is ahead of schedule: "We are now on track to achieve our $150 million total synergy target in 24 months, or 12 months ahead of schedule. And we are now increasing our three-year total synergy target to $200 million." This success is helping to strengthen margins and ensure a strong cash flow.

Details of particular product successes were highlighted in the call by Giovanni Barbarossa, who is Chief Strategy Officer and President of the Compound Semiconductor Segment.

Barbarossa revealed that products for 3D sensing increased more than 140 percent sequentially. "3D sensing growth came from shipments of production volumes of VCSEL arrays for multiple end customers, including for front-facing and world-facing applications, as well as for other consumer electronics, and automotive in cabin sensing," said Barbarossa. According to him, sales of VCSEL arrays should continue to rise as II-VI expands its customer base with additional wins, including in the Android ecosystem and personal computing platforms.
II-VI has devoted much time and effort to building up its 3D sensing products. Back in 2013 it bought Oclaro's GaAs laser business; three years on it purchased the Anadigics GaAs fab and the Epiworks facility; and in the autumn of 2019 it bought the Finisar fab.

"When we acquired Finisar some asked us which gallium arsenide fab we plan on closing," remarked Barbarossa. "Our answer was none because we needed the capacity to gain share and become the market share leader, by offering breakthrough solutions at scale."

As well as sensing, the VCSELs made by II-VI are targeting LIDAR. Barbarossa is confident of success in this sector, arguing that the company has the broadest portfolio of products in the industry. "Unlike our pure play laser competitors, we have an entire vertical integrated portfolio of both active and passive components, made from our engineering materials that are critical for these next generation LiDAR designs."

SiC products for power electronics could also help to grow II-VI revenue. Encouraging signs on this front include a purchaser of II-VI's SiC substrates winning selection by a tier-one automotive manufacturer in Japan.

Another promising portfolio of products are InP-based lasers. Shipments of high-data-rate coherent transceivers are on the rise, adding bandwidth to new and existing networks, and sales are up for 200G and 400G components. "We are also excited to announce that we have just sampled our first 800G transceivers to a large web-scale customer, who has already provided exciting feedback," added Barbarossa.

II-VI is widening its broad portfolio with the acquisition of Coherent, a manufacturer of a broad range of lasers, including CO, CO_2, excimer, solid-state and fibre sources. II-VI emerged victorious in a bidding war with MKS Instruments and Lumentum during the first few months of this year. The stock-and-cash deal, worth around $7 billion, is expected to close at the end of this year.

Results just in for the third fiscal quarter show that the company continues to thrive, generating a revenue of $783 million and a cash flow from operations of $447 million. For the fourth fiscal quarter, sales are expected to net $752-802 million, with earnings of $98-127 million.

Languishing Lumentum

Footing this year's table is laser diode manufacturer Lumentum. It's not that Lumentum has performed poorly – it's just that its peers have fared better. Over the timeframe considered in the annual Shareprice Leaderboard, Lumentum has seen a rise in its valuation by about 20 percent, while based on a pre-Covid reference point of the start of 2020, the increase in share price is around 17 percent.

Over the last year or so Lumentum's share price has not fluctuated wildly. It's nadir over that time has been just below $70, and it has peaked at a little more than $105.

Lumentum has been experiencing mixed fortunes. Sales of it VCSELs for 3D sensing applications have been strong, but shipments of telecom lasers have

Compound Semiconductor Shareprice Leaderboard

Rank	Company	Ticker	Share value, January 2, 2020 ($)	Share value, April 27, 2021 ($)	% appreciation	Change in Rank
1	AXT	AXTI	4.36	11.61	166.3	+12
2	Cree	CREE	46.90	119.76	155.4	+13
3	Aixtron (Frankfurt)	AIX	10.28*	22.89*	122.9	+6
4	Emcore	EMKR	3.05	6.69	119.3	+4
5	II-VI	IIVI	34.26	74.56	117.6	+7
6	Qorvo	QRVO	117.26	200.36	70.9	-3
7	Skyworks	SWKS	122.01	201.18	64.9	-2
8	IPG Photonics	IPGP	146.99	234.06	59.2	+6
9	NASDAQ composite	IXIC	9039.46	14170.91	56.8	-3
10	Veeco	VECO	14.93	23.29	56.0	0
11	IQE (London)	IQE	67.74*	88.05*	30.0	+5
12	Rubicon	RBCN	8.65	10.54	21.8	-5
13	WIN Semiconductor (Taipei)	3015.TWO	10.93*	13.30*	21.6	-9
14	Infinera	INFN	8.13	9.80	20.5	-13
15	Lumentum	LITE	81.25	94.90	16.8	-13
16	Riber (Paris)	RIB	3.225*	2.20*	-31.8	-5

** Converted to dollars using the exchange rates on 27 April of 1 EURO = 1.2079 USD, 1 GBP = 1.3910 USD and 1 TWD = 0.035839

➤ Comparing share prices of compound semiconductor companies to pre-Covid times shows that they have fared far better than just recovering from the nadir caused by the pandemic.

weakened, due to a push-out in 5G deployments. There has also been a sharp decline in sales of commercial lasers, due to weakness within the industrial sector.

While the company's management may also rue their failure to secure Coherent, market reaction to missing out on this deal has actually been favourable. When news broke in March that II-VI had finally won the bidding war, Lumentum's valuation climbed by around 5 percent, probably spurred on by a boost to the company's financial assets – it would now be in line for a termination fee, worth $218 million.

Very recently Lumentum has released its results for its third fiscal quarter of 2021, ending 31 March. Sales missed guidance, coming in at $420 million, down 12 percent sequentially and below the guidance given in February of $425 million to $440 million. This news didn't go down well with the market, with shares tumbling 15 percent.

Company CEO Alan Lowe accounted for unexpectedly low revenue in a quarterly earnings call given on 12 May. He blamed the sales shortfall on China's roll out of 5G, which is going slower than anticipated, due to geopolitical tensions with the US and shortages of electronic components. Continued delays in 5G fronthaul deployments in China have dragged down third quarter revenue for directly modulated lasers to significantly below year-ago levels. There will be no immediate improvement, with sales of this class of laser for the upcoming quarter expected to be down by more than $20 million year-on-year. "At this time, we expect 5G fronthaul deployments could resume this summer," added Lowe. "This timing would drive increased demand for our products towards the middle of fiscal '22, once customers ramp up and burn through existing inventory."

Another sector weakening in the short term is 3D sensing. Due to customer design decisions, the global market for 3D sensing lasers is expected to decline between 20 percent and 25 percent during Lumentum's fiscal 2022. However, during that fiscal year, and the one that follows, laser-based sensing is expected to expand into more applications and markets. Lowe says that this will set the stage for reacceleration of market growth in fiscal 2023.

Although total sales are expected to be sluggish for the next few quarters – revenue for the first half of fiscal 2022 is expected to be down approximately 5 percent relative to the first half of fiscal 2021 – the company is very bullish about its longer-term prospects. Sales for this latest quarter were rich in new and differentiated products that enjoyed double-digit year-on-year growth, and can serve in markets with great prospects for expansion over many years. These products include InP coherent components and modules, next-generation contentionless ROADMs, high-speed electro-absorption modulated lasers and 3D sensing lasers.

Commenting on these opportunities, Lowe remarked: "The computer and machine vision revolutions are in their early days, and we expect 3D sensing and LiDAR capabilities will expand to many more applications in multiple markets." In his view, there are opportunities in augmented and virtual reality, 3D machine vision for industrial applications, frictionless and contactless biometric security and access control, and automotive and delivery vehicle applications. Another sector with much promise is laser-based material processing, which Lowe says is critical to the manufacturing of the devices that enable the digital transformation and transition to 5G wireless, as well as the production of electric vehicles and technologies for energy storage.

With so many great opportunities ahead, it's unlikely that Lumentum will languish at the bottom of our leaderboard for too long. But with strong prospects in the next 12 months for all companies in the compound semiconductor sector, it will not be easy to climb rapidly up the table.

➤ II-VI has just expanded its SiC wafer finishing manufacturing footprint in China to serve the market for electric vehicles and for clean energy applications.

IQE

THE GLOBAL LEADER IN EPIWAFER SUPPLY

Enabling the Best RF, Photonics & Power Semiconductors

- 'One-stop-shop' for III-V & Silicon epiwafer products
- 2" to 6" epitaxy ready substrates & polycrystalline materials
- Multiple manufacturing sites; Europe, Asia & America
- Leveraging 30+ years of compound semiconductor IP
- Custom engineered structures; R&D to production volumes

RF

GaN HEMT
5G Infrastructure
Radar, EW/ECM, & SATCOM

GaAs HBT / pHEMT / BiHEMT
5G Smartphones & Infrastructure
WiFi6 Hotspots, GPS & IoT

SiGe HBT
5G Smartphones
WiFi6, Bluetooth, GPS & IoT

PHOTONICS

InP Laser & Detector
Optical Comms & Data Centres
SW IR Imaging

GaSb Laser & Detector
MW-LW IR imaging
Biometrics

GaAs VCSEL
3D Sensing & LiDAR
Datacoms

POWER

GaN on Si HEMT
GaN on GaN
Electric Vehicle Systems
Power Conversion & Storage

GaAs Multi-Junction Solar
High Efficiency Terrestrial CPV & Space PV

GaN LED & Laser
MicroLED Display & AR/VR
UV Sterilisation

EUROPE: +44 2920 839 400
AMERICAS: +1 508 824 6696
APAC: +886 3 579 8181
ALL INQUIRIES: sales@iqep.com

www.iqep.com

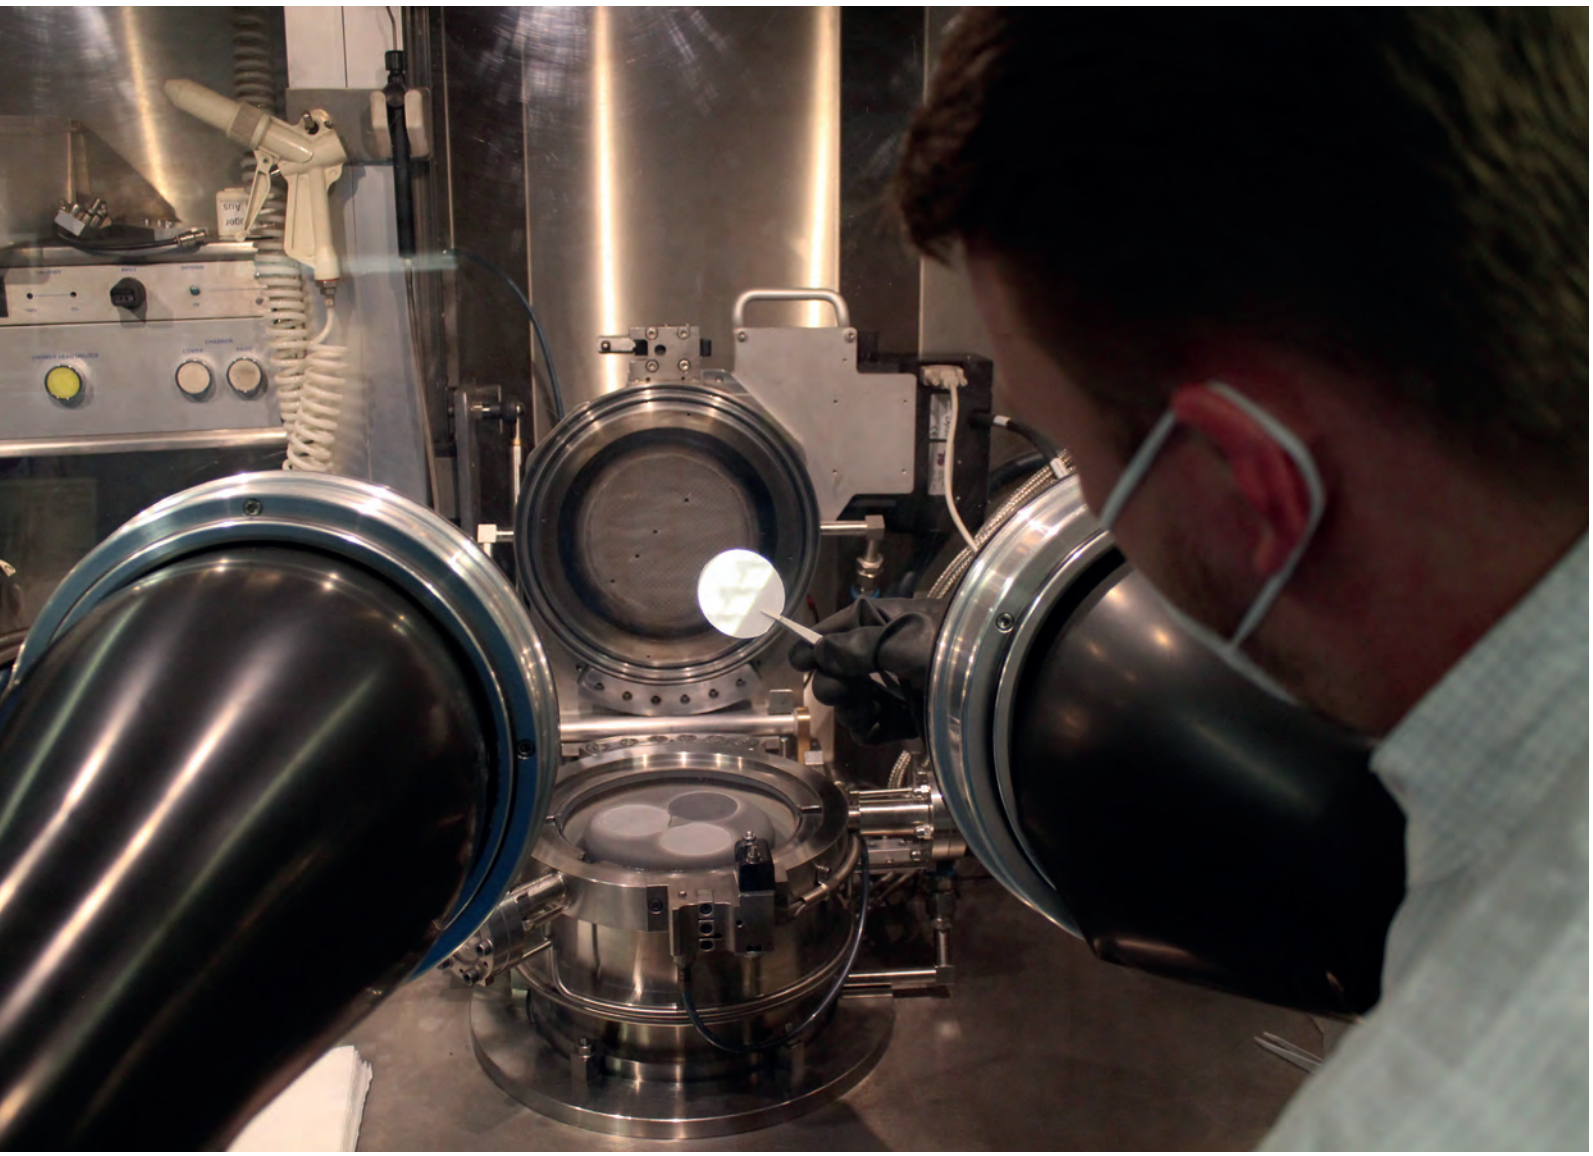

Tackling the transparency conductivity conundrum

Deep-UV LEDs enjoy a hike in efficiency, thanks to the transparency of the MOCVD-grown tunnel junction

BY CHRISTIAN KUHN, TIM WERNICKE AND MICHAEL KNEISSL FROM TU BERLIN

THERE ARE MANY applications requiring a source of deep-UV emission. Radiation in this spectral domain is employed for materials processing, sensing, and disinfection targeting bacteria and viruses, including SARS-CoV-2.

Today, the incumbent source for these tasks is the mercury lamp, which is available in low-pressure or medium-pressure formats. A more promising option, though, is the LED: it is smaller, more robust, and does not require high voltages or employ toxic materials. There is also potential for production costs of this chip to fall, just like they have for its blue and white cousins.

Holding back the deep-UV LED is its low external quantum efficiency, which is limited by poor carrier injection and a low light-extraction efficiency – in short, it suffers from a transparency-conductivity dilemma. Makers of UVC LEDs tend to trade conductivity for transparency by utilising GaN:Mg *p*-layers for the injection of holes. While a good choice for electrical performance, GaN is compromised by its UV light-absorbing nature. Increasing the aluminium mole fraction in the *p*-AlGaN layers combats absorption losses, but at the expense of higher magnesium acceptor ionization energies that lead to a significant increase in the resistivities of the AlGaN:Mg layers, and in turn, the operating voltage. Another downside of AlGaN:Mg is that despite much effort, an ohmic contact is elusive. Forming such a contact is challenging, due to the large work function mismatch between UV-reflective metals and the high electron affinity of *p*-type III-nitrides.

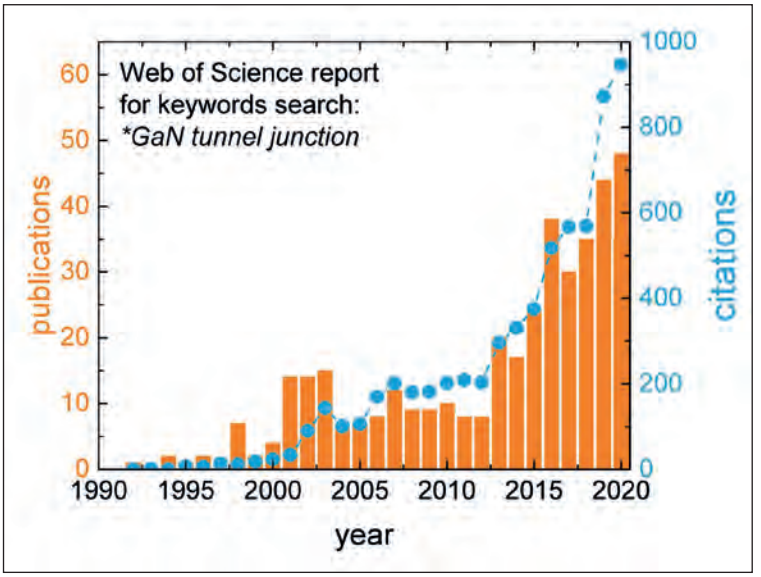

➤ Figure 1. Accelerating research efforts and rapid developments are illustrated by the growing number of scientific publications and citations regarding (In,Al)GaN-based tunnel junctions. This field has been propelled by these main breakthroughs: first, the GaN-based tunnel-junction LEDs, in 2001; and then AlGaN-based tunnel heterojunctions and lateral activation mechanism in 2013, as well as MOCVD-grown deep-UV LEDs in 2019.

A promising solution to the transparency-conductivity dilemma is to turn to a tunnel junction for hole injection. Twenty years ago, researchers incorporated this junction into visible, GaN-based LEDs for the first time (see Figure 1). This body of work showed that blue and green LEDs, as well as laser diodes, produce excellent performance when MBE provides the growth or re-growth of the tunnel junction.

One well-known strength of MBE is that, in contrast to MOCVD, it is well suited to realising very abrupt doping profiles that are beneficial for a tunnel-junction interface. In addition, MBE is capable of very high magnesium doping levels, such as more than 10^{20} cm^{-3}; and it can avoid hydrogen passivation of magnesium acceptors and magnesium diffusion during overgrowth of the *n*-doped layer. However, despite all these strengths associated with MBE, there is still much merit in trying to develop an MOCVD-based process for producing tunnel junctions. The latter is the dominant manufacturing technology for producing III-nitride LEDs, with processes applied on an industrial scale. What's more, the use of MOCVD allows tunnel-junction LEDs to be grown in a single run, eliminating the need for a two-step growth process.

Using this growth technology, our team at TU Berlin, Germany, has had much success developing deep-

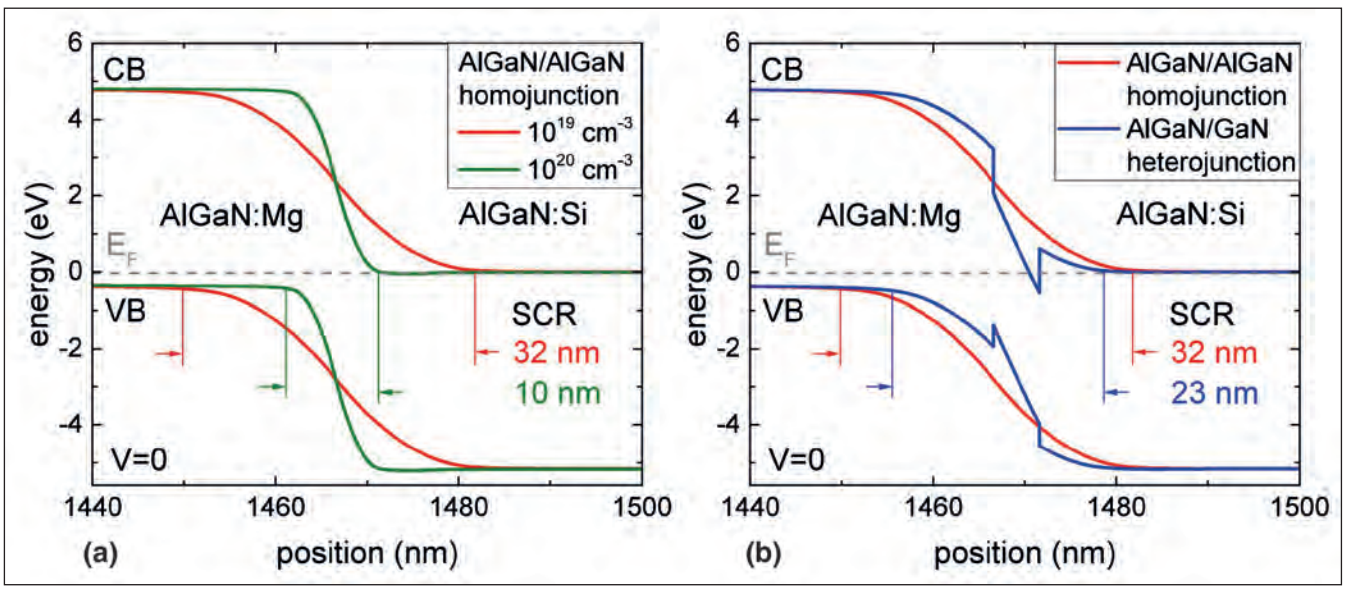

➤ Figure 2. Band diagram simulations of AlGaN-based tunnel junctions at zero bias highlight two mechanisms for reducing the tunnelling barrier width: the heavy doping approach (a) and the heterojunction polarization engineering approach (b). Both leads to thinner space charge regions and higher tunnelling probabilities.

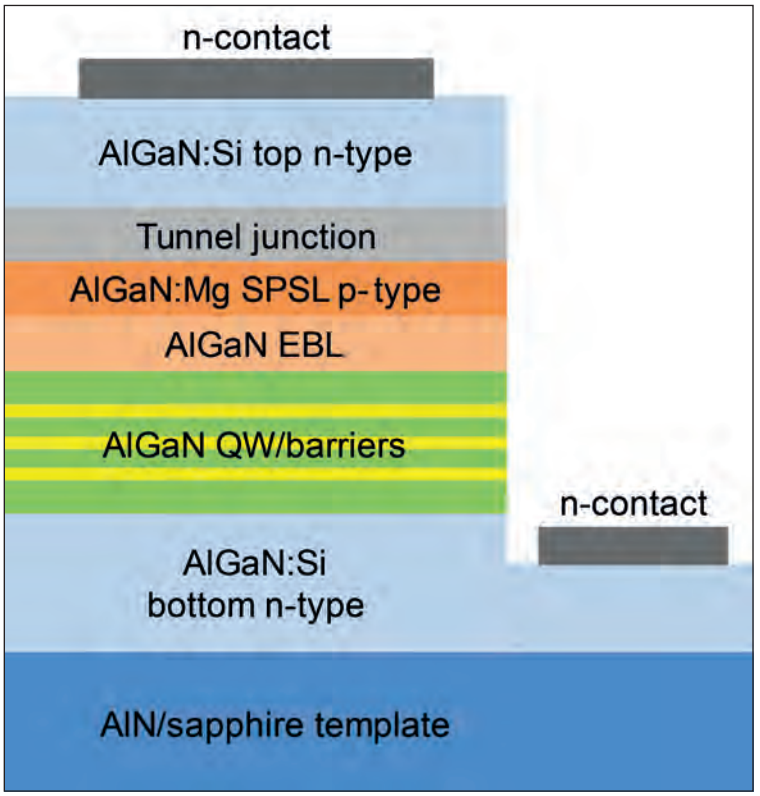

➤ Figure 3. Researchers at TU Berlin have investigated the UVC tunnel-junction LED heterostructure.

UV LEDs that incorporate a tunnel junction. We have broken new ground by demonstrating the first MOCVD-grown AlGaN-based tunnel-junctions, as well as the first AlGaN/AlGaN homojunctions; and we have produced the world's shortest-wavelength LED that has a tunnel junction.

There are several benefits associated with the incorporation of an AlGaN-based tunnel junction in a deep-UV LED. This junction provides an alternative architecture to the conventional absorbing p-side for hole injection into the active region. With this design, which leads to hole injection via interband tunnelling, silicon-doped *n*-AlGaN layers partially replace highly resistive *p*-doped layers, and ohmic *n*-contacts replace rectifying *p*-contacts. This scheme simplifies the contact fabrication process, with the device just requiring *n*-metal contacts to AlGaN:Si. Further gains, stemming from the use of tunnel-junction technology for non-equilibrium hole injection, include more efficient lateral current spreading in highly conductive *n*-AlGaN layers and reduced optical losses in UV LEDs and laser diodes. There is also the opportunity to realise new device concepts, such as: intra-cavity contacts for VCSELs; vertical integration of cascaded laser diodes; and power adjustment by voltage-scaling, rather than current scaling. The latter refinement promises a reduction in efficiency droop.

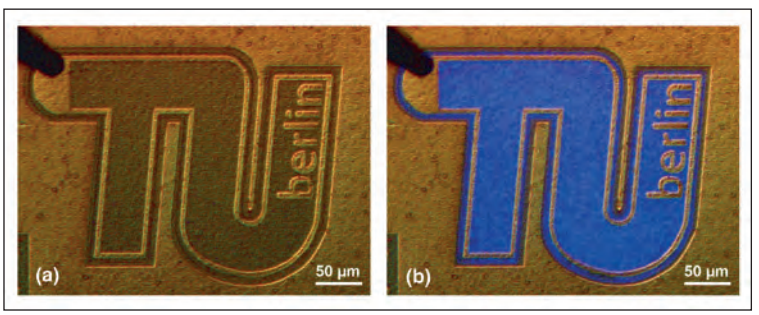

➤ Figure 4. Microscopic images of the processed TJ-LED surface before operation (a) and during operation (b). In image (a), golden top and bottom contacts are visible in bright contrast. The darker area shows the metal-free surface of the AlGaN heterostructure, which is UV transparent down to the multi-quantum-well plane. (b) A few milliamps are supplied to the LED, using the top contact at the left with a black needle attached.

Injection by interband tunnelling

When a tunnel-junction is incorporated into the *p*-side of a deep-UV LED, this leads to non-equilibrium hole injection via tunnelling (see Figures 2 and 3). Operating in forward bias, such a device produces a tunnelling of electrons from the AlGaN:Mg valence band into the AlGaN:Si conduction band. This process leaves a free hole that is injected into the active region, where it takes part in radiative recombination. The efficiency of this tunnelling may be described by a tunnelling probability.

In general, the quantum-mechanical tunnelling probability – and with that the total device resistivity – is governed by the height and the width of the tunnel barrier. The height of the barrier depends on the bandgaps of the materials at the heterojunction, while its width is determined by the space charge region. We control both characteristics with our design. Our investigation has focused on two key elements: maximizing magnesium- and silicon-doping at the tunnel junction; and the introduction of polarization-field-enhanced tunnel-junction heterostructures, which utilise lower bandgap interlayers.

We have simulated the properties of our LEDs, considering an $Al_{0.7}Ga_{0.3}N$ tunnel junction (Figure 2 shows energy band diagrams at zero bias). These calculations included a reference structure with homogeneous and abrupt silicon- and magnesium-doping of 10^{19} cm^{-3} at both sides of the junction – for this structure, the space charge region is 32 nm wide. When we increase the doping levels to 10^{20} cm^{-3}, the width of this region decreases to just 10 nm. Another option for reducing the width of the space charge region is to introduce a polarization-field-enhanced tunnel-junction heterostructure with a GaN interlayer. This modification shortens the width of the space charge region from 32 nm down to 23 nm (see Figure 2 (b)).

Within certain limitations, there is no need to choose between these two options, but rather incorporate both of them. This has a huge pay-off, because reductions in the space charge region deliver an exponential increase in tunnelling probability, and in turn a significant hike in device performance.

We have evaluated different designs, starting with a tunnel-junction LED delivering emission in the UVC (its design is shown in Figure 3). The foundation for this

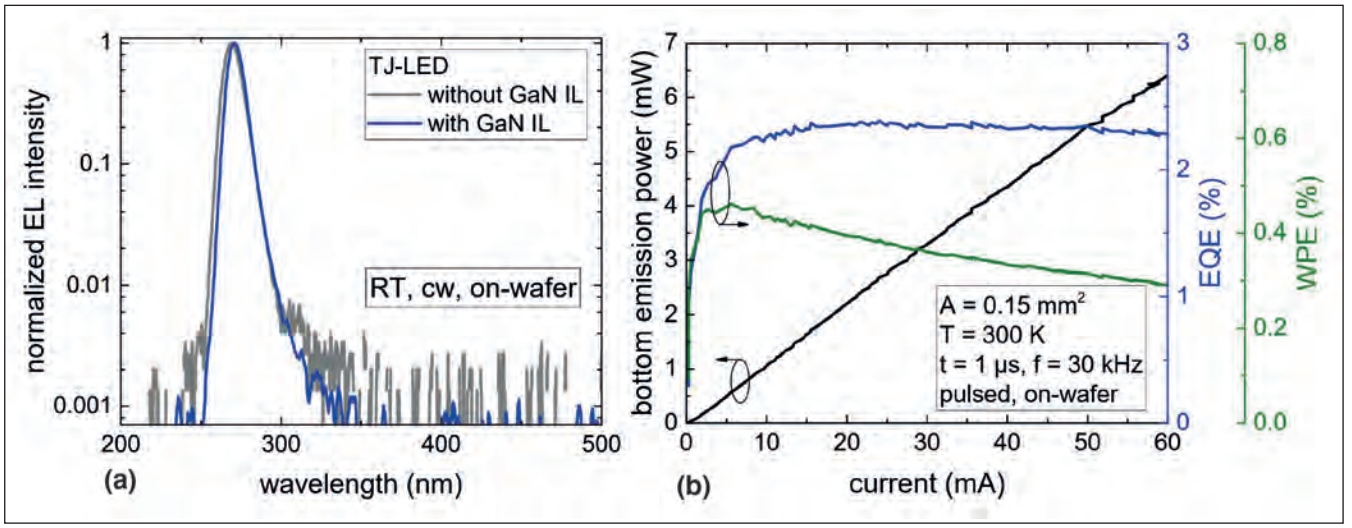

➤ Figure 5. Electroluminescence spectra (a) and emission characteristics (b) of tunnel-junction LEDs emitting at 268 nm.

device, provided by FBH Berlin, is a defect-reduced AlN-on-sapphire template formed by epitaxial lateral overgrowth of a patterned AlN layer. On this base we add, by MOCVD, a heterostructure that includes: an *n*-type AlGaN:Si bottom current-spreading layer; a threefold AlGaN multiple quantum well, separated by AlGaN barriers; and a magnesium-doped AlGaN electron-blocking layer with a wide bandgap. For the *p*-side, we have a thin AlGaN:Mg short-period superlattice, followed by the tunnel junction and a top AlGaN:Si current-spreading layer. The tunnel junction features an 8 nm-thick GaN:Si interlayer sandwiched between a heavily doped AlGaN:Mg superlattice (9×10^{19} cm^{-3}) and an AlGaN:Si layer (5×10^{19} cm^{-3}).

Initial emission

Fabrication of our first MOCVD-grown LEDs that feature tunnel junctions involved an *in-situ* thermal annealing step after the growth of the magnesium-containing layers to activate magnesium acceptors. With this step, acceptor passivation by hydrogen is reduced to negligible levels – according to secondary ion mass spectrometry, the ratio of hydrogen to magnesium atoms is below 10 percent. Device processing is completed with mesa etching, followed by simultaneous deposition of top and bottom V/Al/Ni/Au *n*-metal contacts and then a rapid thermal annealing of these contacts under nitrogen (see Figure 4 (a) for a microscopic image of our processed tunnel-junction LED, with a TU Berlin logo).

Applying a few milliamps to our device produces homogeneous emission from the whole area of the tunnel-junction LED (see Figure 4(b), which captures luminescence from the AlGaN quantum well plane using a 'blue' false colour). Proof of deep-UV emission comes from spectrally resolved measurements. A plot of normalized electroluminescence intensity reveals a single-peak in the UVC, centred at 268 nm, and no significant parasitic luminescence (see Figure 5 (a)). Note that these emission profiles are produced for CW injection at room temperature, for tunnel-junction LEDs with and without a GaN interlayer. Success on both fronts has given us two breakthroughs: the first experimental demonstration of an AlGaN-based tunnel-junction LED made by MOCVD, and the first ever AlGaN/AlGaN tunnel homojunction produced by any means.

As we expected, our variant with the homojunction has a higher tunnelling resistivity, with high voltages producing low currents. The sibling with the heterojunction produces better results. For that device, on-wafer measurements at a 60 mA drive current produce an output power up to 6.6 mW, corresponding to an external quantum efficiency of 2.3 percent and a wall-plug efficiency of 0.42 percent (plots obtained using a pulsed current are shown in Figure 5 (b)). Encouragingly, light-current and current-

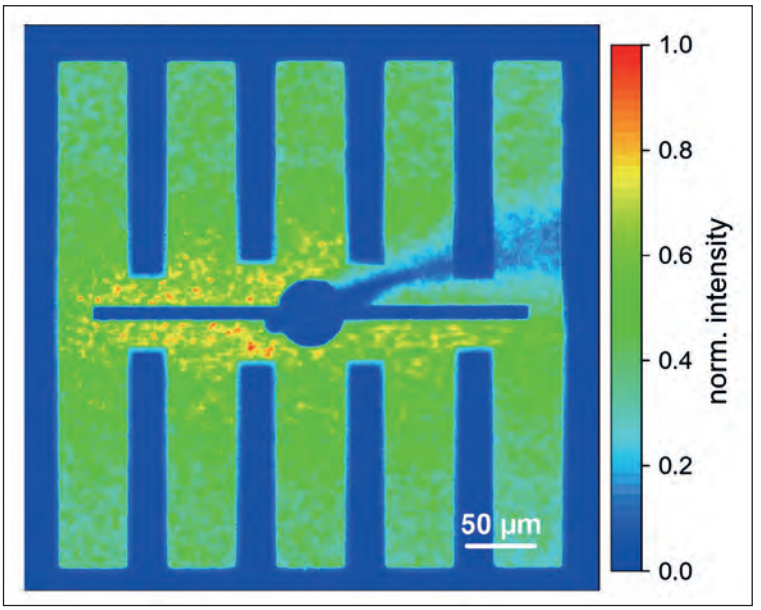

➤ Figure 6. UV-sensitive microscope image of a 268 nm-emitting tunnel-junction LED with a finger-shaped mesa geometry and central stripe contact during operation.

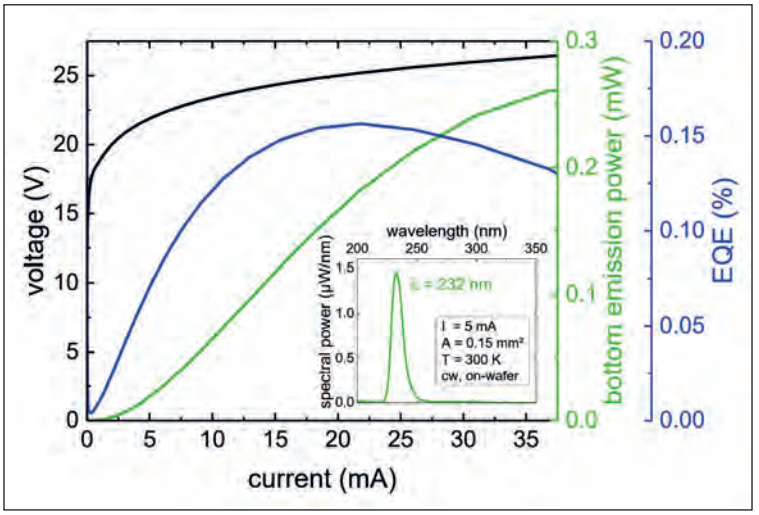

➤ Figure 7. Optical and electrical characteristics and spectral power density of a MOCVD-grown tunnel-junction LED with a peak emission at 232 nm.

voltage characteristics show no indication of carrier leakage or parasitic current paths.

Using a UV-enhanced microscope camera, we have mapped the intensity across the tunnel-junction LEDs. Imaging of our finger-shaped LED mesa, with a thin central metal contact on top, shows emission from the entire mesa area (see Figure 6). The uniform emission highlights homogenous acceptor activation and current injection through the tunnel junction. If our tunnel-junction LEDs were impaired by incomplete activation or issues associated with local injection problems, this would have led to dark areas without any emission.

Close inspection of our intensity image reveals a slight intensity enhancement near the centre compared with the finger edges. The most likely cause of this is limited current spreading in the top AlGaN:Si layer along the 170 μm-long LED-fingers. This can be resolved by optimizing the LED layout or the heterostructure design through the introduction of a thicker top AlGaN:Si layer with higher conductivity.

Deeper into the UV

Our technology is capable of going deeper into the UV than the devices discussed so far, which emit at 268 nm. The fundamental limitation is the bandgap of AlN, which is just above 6 eV – this corresponds to emission at around 200 nm. Success in this domain is not easy, and gets progressively harder with every nanometre further into the deep UV. Reasons for this include: increasing severity of UV absorption; reductions in doping efficiency; limited options for carrier confinement, all issues for such LEDs in general; and lower tunnelling probabilities, resulting primarily from an increase in the tunnel barrier, due to an increase in the AlGaN bandgap and a larger space charge region.

By taking on these challenges, we have managed to expand the spectral range of our MOCVD-grown tunnel-junction LEDs down to 232 nm, a world record for this class of device (see Figure 7). Electroluminescence spectra show a dominant emission peak from the multiple quantum wells, and only marginal parasitic long-wavelength emission from these non-optimized structures.

Plots of current-voltage characteristics highlight one weakness of our device. Its operating voltage is above 15 V, and clearly needs to be improved. Emission power is encouraging, at almost 0.2 mW at 20 mA, helping to deliver a peak external quantum efficiency as high as 0.16 percent. For UV LEDs operating at such ultra-short emission wavelengths, this value – obtained on-wafer – is close to the highest values ever reported. In fact, it is enabled by the highly UV transparent TJ-LED heterostructure that overcomes the transparency-conductivity-dilemma of deep UV LEDs with potential for even more improvements.

The conventional approach to quantifying the efficiency of tunnelling is to use differential device resistivity (this is dV/dj, simply calculated from j-V measurements). So long as the tunnel junction provides the main contribution to device resistivity, this metric offers a solid basis for comparing different tunnel-junction LEDs.

We have compared values for device resistivity versus current density for tunnel-junction LEDs grown by MBE and MOCVD, using data reported by various groups (see Figure 8). Our plot highlights the underlying trend of decreasing resistivity with increasing current density, which occurs over several orders of magnitude and is independent of device design and growth method. One implication is that MOCVD is just as good as MBE for growing tunnel-junction LEDs.

Another benefit of this plot, shown in Figure 8, is that it allows small variations at a fixed current density to compare properties of individual tunnel-junction LEDs. When viewed in this manner, the results highlight the outstandingly low device resistivity of nearly 10^4 Ω cm^2 at 2 kAcm^{-2} for InGaN-based devices produced at Unipress.

Device resistivity appears to be fairly independent of the bandgap of the tunnel-junction and its growth method, but it is connected to the wavelength. A low resistivity is only realised in GaN- and InGaN-based tunnel-junction LEDs emitting beyond 350 nm. When tunnel junctions employ AlGaN they show higher operating voltages as well as higher differential device resistivities due to lower achievable current densities. However, as well as the work by our team, independent efforts at Ohio State University and CRHEA-CNRS show that when tunnel junctions are formed with an aluminium mole fraction way above 50 percent – this corresponds to a bandgap above

4.5 eV – they offering an alternative approach to carrier injection that combines transparency with conductivity.

We have several ideas for improving the performance of our devices. Plans include trying to increase the wall-plug efficiency through the introduction of extra-thin AlGaN:Mg *p*-layers that reduce the operating voltage; suppression of donor compensation in the top AlGaN:Si layer, to improve current spreading; and a further reduction in tunnelling resistivity by optimizing doping profiles of the tunnel junction, as well as the interlayer's composition and thickness. Once we have demonstrated high-performance tunnel-junction LEDs at the wafer level, our next step will be chip fabrication. This work will include efforts to improve yield and investigate long-term reliability, which can be easily accomplished by standard techniques used to evaluate UV LEDs.

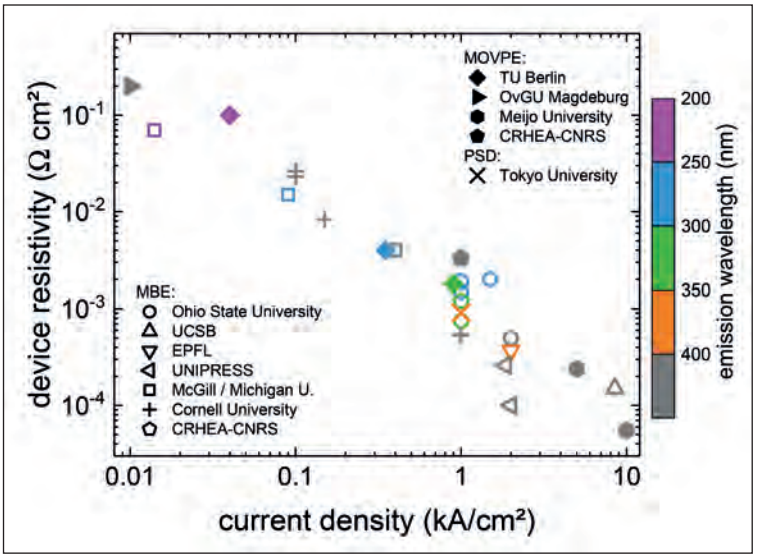

➤ Figure 8. Data from tunnel-junction LEDs reported by different groups highlight the uniform trend of a decreasing differential device resistivity (dV/dj) with increasing current density, independent of the emission wavelength range. The growth method and the material composition of the tunnel junction cause only small deviations from this trend.

FURTHER READING

- C. Kuhn *et al.* Photonics Research **7** B7 (2019)
- F. Mehnke *et al*. Photonics Research **9** 1117 (2021)
- H. Amano *et al.* J. Phys. D: Appl. Phys. **53** 503001 (2020)

How robust is the SiC MOSFET?

Researchers at the International Reliability Physics Symposium report the results of investigations into the impact of defects and gate oxide quality on the performance of the SiC MOSFET

BY RICHARD STEVENSON

THE SiC MOSFET is guaranteed a great future. Thanks to its capability to efficiently control current at high voltages, sales are accelerating in a multitude of applications, including electric vehicles, solid-state circuit breakers, and various types of motors. Multi-billion dollar revenues are sure to follow.

However, exactly how much success the SiC MOSFET will have is not set in stone. Factors weighing heavily on this are price, performance, and reliability. And of these three, reliability is arguably the most critical. That's because many of the adopters of this device will really value robustness, as this allows them to foster a reputation for producing products that never fail.

Efforts within the SiC community continue to take the reliability of the SiC MOSFET to a new level. At the recent *International Reliability Physics Symposium* (IRPS), a forum with a rich history in considering the long-term health of silicon devices, several presentations considered the robustness of the SiC MOSFET.

At this meeting, held on-line due to Covid-19 restrictions, the likes of Infineon, STMicroelectronics and On Semiconductor provided insights into the impact of defects on the lifetime of the device, and offered options for assessing reliability of the gate oxide.

Hidden assassins

Today, most SiC devices are produced using 150 mm SiC substrates. The quality of this foundation has improved a great deal over the last two decades. However, even state-of-the-art substrates grown by physical vapour transport are far from perfect. It is the norm for them to have around ten thousand defects per square centimetre, according to conference speaker Thomas Neyer from On Semiconductor.

Neyer's presentation considered many different forms of defect that occur in SiC MOSFETs. He explained that one common option for exposing them is to treat the material in molten potassium hydroxide, and then examine the wafer under a microscope. But he prefers to classify defects with non-destructive approaches, such as photoluminescence techniques, X-ray topography and imaging with cross-polarised light. Using these techniques to record defects allows the performance of devices to be correlated with the type of imperfection. So powerful is this approach that it is even possible to relate device performance to the actual number of defects in a particular die.

It would be easy to blame all the defects found in epiwafers on imperfections in the boule. That's misleading, however, explained Neyer, who pointed out that they can also be introduced in the slow and costly wafering process flow, involving slicing, grinding and polishing. When substrates are formed from a boule, this creates nanoscale dislocations with open cores that hamper the quality of SiC epilayers. Neyer and his co-workers have categorised the defects found in SiC epilayers into three groups: killer visible defects, which include triangular types of defect, strong topographic defects and carrots; non-killer visible defects, such as obtuse triangles, scratches, pits and V-type defects; and non-killer crystal defects, such as stacking faults, basal plane dislocations, grain boundaries and bar stacking faults (see Figure 1 for more details).

Investigations at On Semiconductor have uncovered the impact of non-killer defects through a study that considered around 3 million devices with a 9 mm^2 die size – they are a combination of Schottky diodes and MOSFETs. According to Neyer, this survey showed that a significant proportion of devices have five or more defects per die. This begs the question: does device performance drop off with more non-killer defects?

The answer is nuanced. For Schottky barrier diodes, an increase in non-killer defects has little impact; but for MOSFETs, the opposite is true – and when a die has more than 10 non-killer defects, this is more of concern than a killer defect. Non-killer defects are to blame for early life rejects, burn-in failures, and outliers in short circuits and avalanche tests.

Neyer and co-workers also discovered that SiC MOSFETs with double-digit numbers of stacking faults have a wider distribution in key device characteristics. Average values for the leakage current and its standard deviation both increase markedly, while values for the breakdown voltage and threshold voltage – these depend on extraction currents – decrease as the number of stacking faults per die increases.

Another investigation by the team from On Semiconductor considered bipolar degradation in

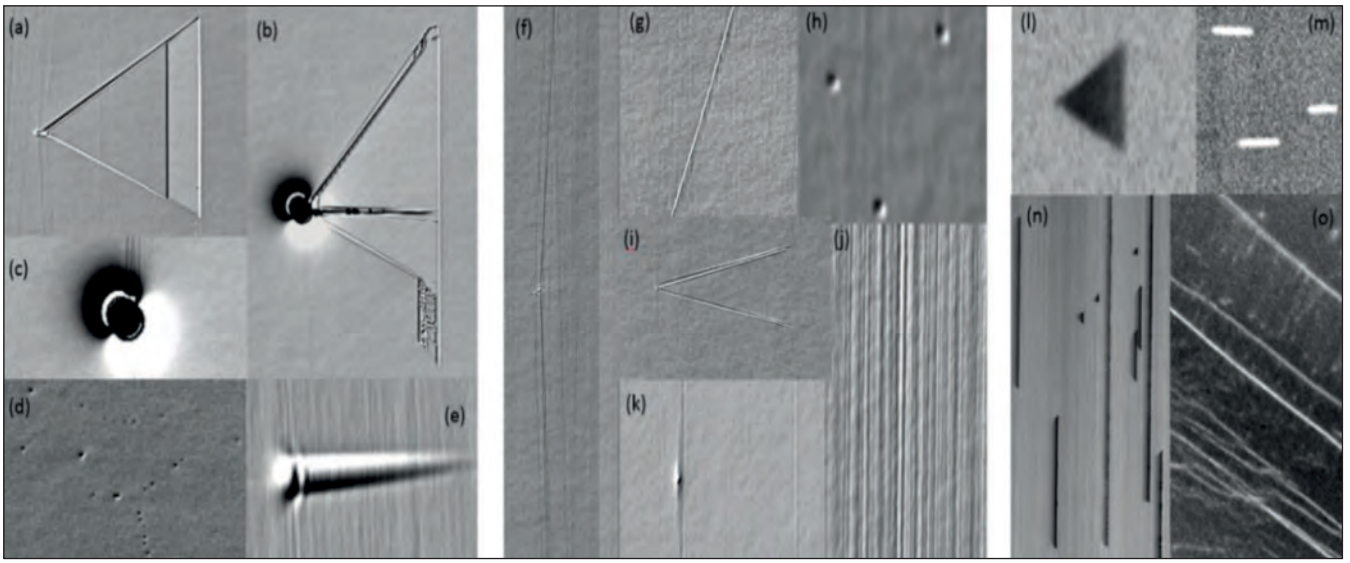

➤ Figure 1. Researchers at On Semiconductor have categorised SiC defects into three classes: killer visible defects (left), non-killer visible defects (middle), and non-killer crystal defects (right). Examples of killer visible defects are (a) triangular defects, (b) particle triangles, (c) particles/downfalls, (d) strong topographic defects and (e) carrots. Non-killer visible defects include (f) obtuse triangles, (g) scratches, (h) pits, (i) V-type defects, (j) roughness/step bunching and (k) small topographic defects. Non-killer crystal defects include (l) stacking faults, (m) basal plane dislocations, (n) bar stacking faults, and (o) grain boundaries.

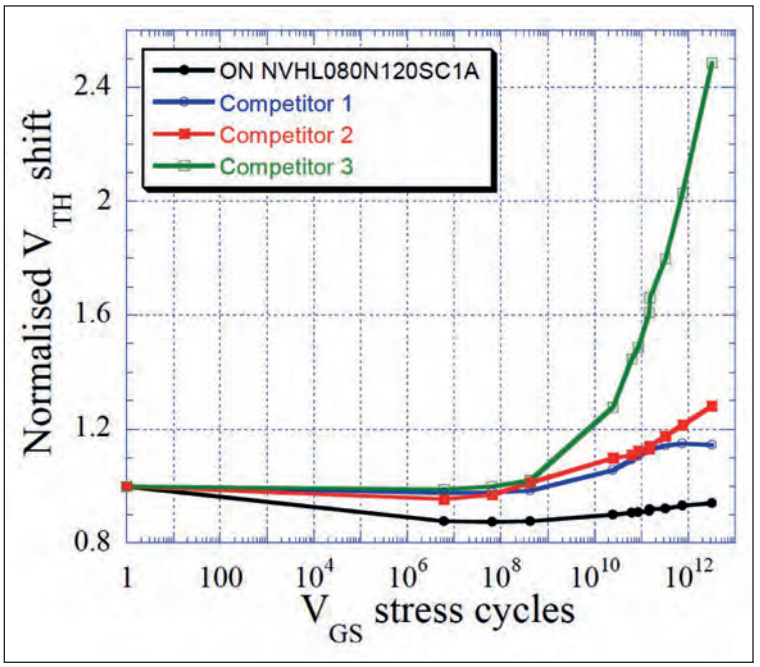

➤ Figure 2. Engineers at On Semiconductor benchmarked their SiC MOSFET against that of three competitors using a gate-source stress at 100 kHz and a 50 percent duty cycle.

the SiC body diode of the MOSFET. This technique, widely employed for evaluating 3.3 kV devices, has been applied by Neyer and co-workers to MOSFETs designed to handle 1200 V and 1700 V. After stressing these devices for between one and three days with a DC drain-source current of 60 A cm^{-2} and a gate-source voltage of -5 V, the engineers found that the greater the basal plane dislocations per die, the greater the shift in on-current. For 1200 V devices, the team recorded an increase in on-current of almost 2 percent for a die with 35 basal plane dislocations, while for the 1700 V equivalent, 30 basal plane dislocations were behind an up-tick in on-current of almost 7 percent.

Encouraging results have come from benchmarking the threshold voltage stability of On Semiconductor's MOSFETs against rival SiC products with planar and trench architectures. Using a gate-source stress at 100 kHz and a 50 percent duty cycle, this device exhibited greater stability than its three competitors used in this study (see Figure 2).

Troublesome triangles

Studies of the role of defects on the MOSFET's breakdown have also been conducted by a team from CNR-IMM, Italy, working in collaboration with STMicroelectronics. Speaking on behalf of this partnership, Patrick Fiorenza from CNR-IMM argued that efforts to understand infant mortality have to begin with wafer-level tests involving thousands of

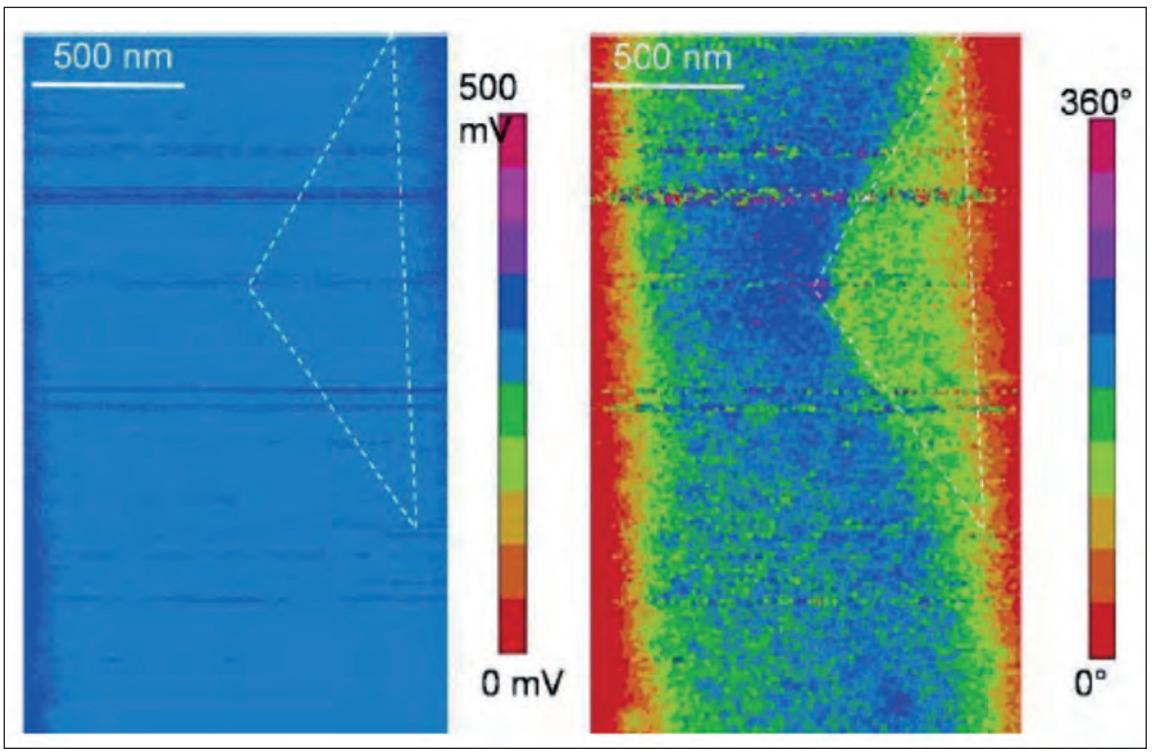

➤ Figure 3. Investigations of threading dislocations by a partnership between CNR-IMM and STMicroelectronics have involved scanning capacitance microscopy measurements. A map of the amplitude of this signal (right) reveals a constant doping concentration; the phase of the signal (left) uncovers local variations in minority carrier concentration, and enables identification of a charge distribution associated with a triangular defect.

devices. He pointed out that it is critical to differentiate between extrinsic breakdown, which happens during the early life of a device, and intrinsic breakdown.

Fiorenza provided an example of a device that had failed instantly. Imaging this device by emission microscopy and scanning electron microscopy revealed a surface pit, which has a hexagonal nature, according to differential atomic force microscopy. When the team delved more deeply into this imperfection with cross-sectional scanning electron microscopy, they found a region with a polytype within the substrate. Weighing up the implications of this finding, Fiorenza concluded: "We have to take care of the fabrication steps, in particular the epitaxial growth of the material."

He added that when working at the buffer level, it is also worthwhile to check the gate current: "This is important to understand if some extrinsic failure can be intercepted before finalisation of the fabrication." At very low electric fields – such as just 4 MV/cm, which ensures no threat of insulator damage – he and his co-workers have found that it is possible to observe gate currents that don't follow the ideal Fowler-Nordheim behaviour. Looking at devices with this attribute in more detail, the team have identified compromised devices that failed high-temperature gate bias tests, due to surface bumps that are seen in atomic force microscopy images.

Devices passing this test were packaged, before undergoing a high-temperature stress test at a 600 V reverse bias. The 2 percent of devices that failed this test, lasting 3 months and involving an elevated temperature of 140 ° C, exhibited a hike in gate current of around seven orders of magnitude. In addition, their characteristics changed dramatically, such as moving to normally-on behaviour.

Imaging the surface of the device with a focused ion beam failed to shed any light on the cause of failure of the MOSFETs. So Fiorenza and co-workers removed the poly-silicon metal gate and the gate oxide, before inspecting the structure once more. This time they discovered triangular defects in the JFET region of the MOSFET. Using a two-beam form of transmission electron microscopy, they found a mixed edge and screw dislocation.

Additional analysis of this stripped back sample, using other forms of probe-based microscopy, enhanced the team's understanding of this imperfection and its consequences. Local current measurements revealed an increase in two orders of magnitude in the conductivity around the threading dislocation, and scanning capacitance microscopy measurements, considering the phase of the signal, revealed local variations in minority carrier concentration and identified a charge distribution associated with the triangular defect (see Figure 3). After drawing on reports in the scientific press, Fiorenza accounted for this observation by reasoning that the threading

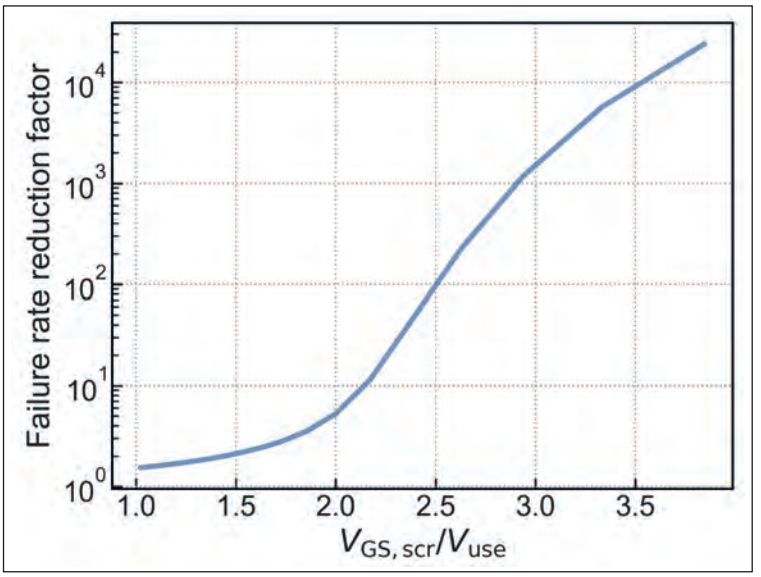

➤ Figure 4. Engineers at Infineon have drawn on previous work (Aichinger and Schmidt, IRPS2020, 3B4_042) that relates a failure rate reduction factor to the ratio of the screening voltage ($V_{GS,scr}$) to the recommended gate use voltage provided by a datasheet (V_{use}). Combining this relationship with rapid stress tests enabled a comparison of the reliability of SiC MOSFETs produced by various vendors.

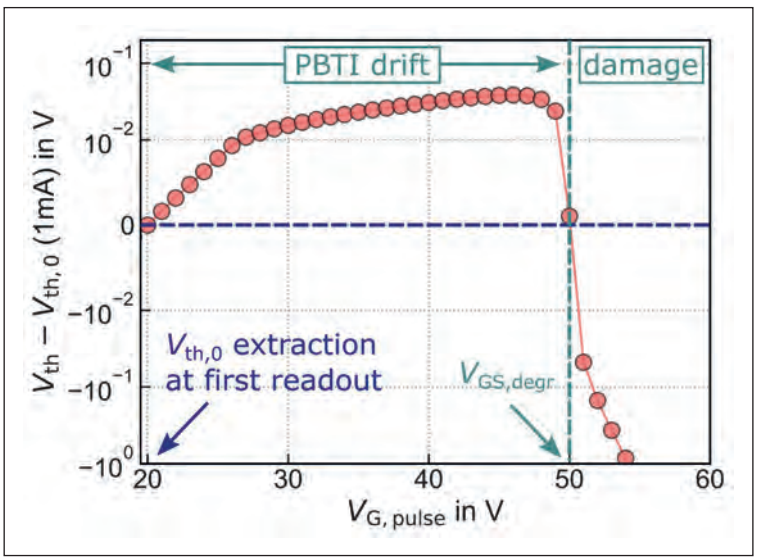

➤ Figure 5. Infineon has developed a new form of stress test for SiC MOSFETs. A series of voltages are applied, each 1 V higher than the previous one, and the threshold voltage is recorded. As the stress voltage incrementally increases, initially the threshold voltage increases as well, due to a positive bias temperature instability. But beyond a certain value the threshold voltage plummets and the device is irreversibly damaged.

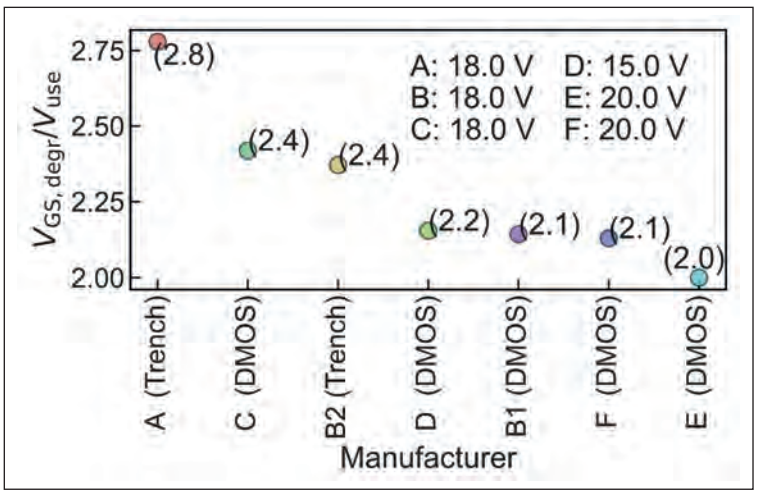

➤ Figure 6. Research at Infineon has uncovered a difference in the failure rate reduction factor of two orders of magnitude between SiC MOSFETs made by different manufacturers. This implies that the most reliable devices have roughly one hundred times fewer failures in the field than the most unreliable devices.

dislocation has an increased hole concentration and a bandgap that is 0.8-1 eV lower than the surrounding SiC. Simulations supported this view and enabled the team to discover that threading dislocations act as quantum wells, increasing hole concentration by 13 orders of magnitude. Operating in reverse bias, these holes are driven through the SiO_2 layer of the MOSFET, accelerating its degradation.

Speeding screening

Those needing to assess the reliability of the SiC MOSFET will welcome the introduction of a relatively straightforward, speedy technique developed by Infineon Technologies and announced at IPRS. This approach focuses on the weakness of the gate oxide, which is the key failure mechanism.

Speaking on behalf of Infineon, Judith Berens detailed this powerful technique after laying the groundwork – the distinction between the intrinsic and extrinsic branch of gate oxide reliability. “The intrinsic branch only plays a role towards the end of the life, and is similar for SiC and silicon MOSFETs,” explained Berens. “Extrinisics, however, might lead to early failures in the field, and for this reason need to be avoided.”

Failures due to extrinsics are a significant cause for concern, given that their prevalence is higher in SiC MOSFETs than in silicon devices. One promising option for uncovering them is to measure device performance at different voltages. “By gate-voltage screening, you mainly sort out devices with extrinsics, and reduce the field failure probability,” said Berens.

The extent of failure reduction realised by screening depends on the ratio of the screening voltage to the operating voltage – Berens refers to this as the failure reduction factor. Note that this factor exceeds a thousand when the screening voltage is three times the value of the recommended operating voltage (see Figure 4).

Of course, there are limitations on how high a voltage can be used for screening. This method of evaluation must avoid permanent damage to ‘good’ chips, accomplished by avoiding degradation caused by the tunnelling of carriers and impact ionisation. Both phenomena can lead to a negative threshold voltage shift. Berens championed a new measurement procedure that avoids using too high a stress voltage. Instead, a series of voltages are applied, each 1 V higher than the previous. After each pulse, which should be as short as possible to reduce the chances of degradation, the threshold voltage is recorded.

To illustrate the value of this approach, Berens presented an example of what one could expect when carrying out this measurement. To begin with, as the stress voltage incrementally increased, so did the threshold voltage, due to a positive bias temperature instability. But beyond a certain value the threshold voltage plummeted, reaching a point of no return and causing the device to be irreversibly damaged (see Figure 5).

Berens explained that determining the likelihood of field-rate failure requires the value for the stress voltage for the onset of irreversible degradation, along with the intended operating voltage for the MOSFET. The ratio of these two determines the failure reduction factor and gives an insight into device reliability.

Demonstrating how this works in practice, Berens compared data for a range of commercial MOSFETs, including trench and DMOS designs. Pulsed stress tests uncovered a 20 V range in voltages required to reach the onset of irreversible damage, with values depending on the thickness of the oxide. This is understandable, as thicker oxides are subjected to a weaker electric field strength that reduces the chances of impact ionisation.

For the next step in this analysis, Berens and co-workers accounted for the operating voltage of the various devices. With this factor included, the differences between manufacturers narrowed significantly; and the new, normalised ratio allowed the failure rate reduction factor to be determined by reading this value from the graph shown in Figure 4. Note that due to the high degree of non-linearity of this graph, small variations in the normalised ratio produce a difference in the failure rate reduction factor by two orders of magnitude (see Figure 6). This implies that the most reliable devices have roughly one hundred times fewer failures in the field than the most unreliable devices.

Berens says that one of the advantages of this method is the speed with which it determines the likelihood of

failure – it takes less than a minute per test run. The technique is also straightforward, with no special test equipment required. "And last but not least, we didn't need any special knowledge, only publicly available data sheet values."

The problem of pits

SemiQ, a producer of SiC power devices based in California, has also been investigating the screening of SiC MOSFETs using the gate voltage. Speaking on behalf of the company, Yongju Zheng detailed measurements on 1200 V, 80 mΩ SiC planar depletion MOSFETs with a 50 nm-thick gate oxide. This investigation considered a variety of voltages and voltage ramp rates, with measurements taken at room temperature and 130 °C.

Zheng and co-workers have pursued a two-step process. In their study, they began by ramping from 40 V to 50 V, which screened out about 4 percent of devices from the sample size of 289. "These are defined as weak gate oxides, as the median voltage of our devices is about 53-54 volts," said Zheng. The second step involved ramping from 40 V and 70 V. The upper end ensured breakdown and enabled a failure distribution to be obtained. Failure is defined when gate leakage hits 9 mA.

The engineers found that the failure voltage is higher when the ramp rate is faster. "This could be due to longer stress time, leading to more charging at the interface of silicon carbide and silicon dioxide," speculated Zheng. Note that for the two temperatures used in this study, no notable difference in failure voltage is found.

Zheng and her colleagues have considered whether large epi pits could be behind the early failure of some SiC MOSFETs. "It turned out that these early failures show very strong correlation to large pit defects, that are defined as pit defects larger than 100 micron-squared." This correlation is so strong that 80 percent of these large pits caused early failure (see Figure 7).

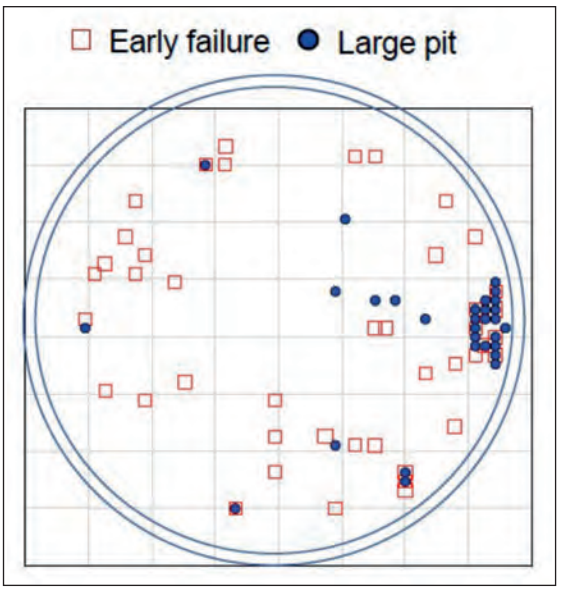

➤ Figure 7. A study at SemiQ has shown that when pit defects are larger than 100 μm², 80 percent of them cause early failure of a SiC MOSFET.

According to Zheng, the presence of these pits may alter the epitaxial growth conditions, and lead to local oxide thinning that would result in electric field crowding – both factors could account for early failures.

One of the downsides of using a high screening voltage is that it causes a negative shift in the threshold voltage, which could impair reliability, especially when operating the MOSFET at high temperatures under reverse bias. Measurements by the team from SemiQ have shown that ramping from 40 V to 50 V produces a negative shift in the threshold voltage of 0.85 V, taking this device close to normally-on behaviour.

"Based on the study in this work, the screening voltage is suggested to be below 38 volts for a 50 nanometre gate oxide in production," said Zheng, adding that such a value should be low enough to detect the extrinisic failure rate, while not significantly

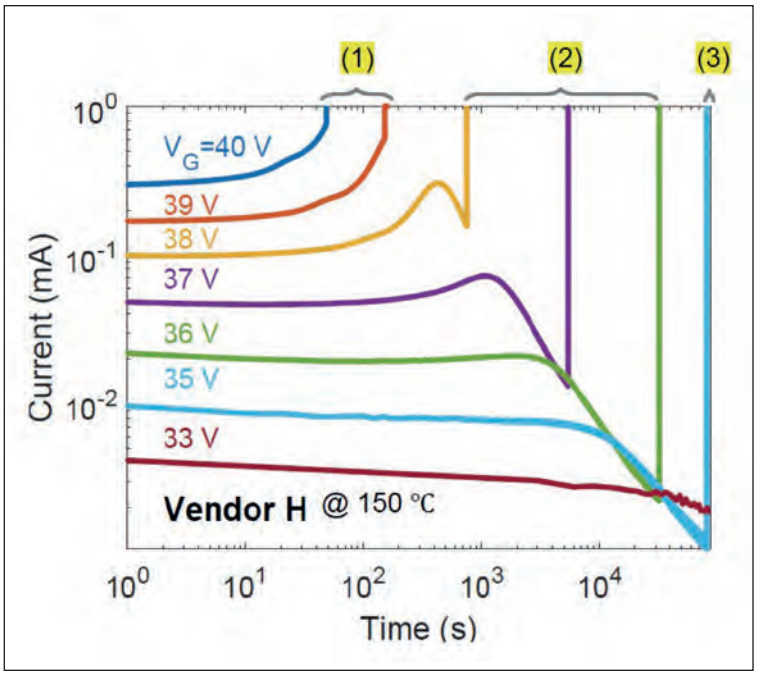

➤ Figure 8. Researchers have identified three different regimes for SiC MOSFETs subjected to gate-stress tests. When the gate voltage is 39 V or more, the gate current accelerates; for voltages of 38 V to 35 V, initially the current varies, before a rapid hike; and for a voltages of 33 V, current falls steadily.

lowering the threshold voltage of the MOSFET, or shortening the lifetime of its gate oxide.

Looking into leakage

A thorough study into the leakage current of 1.2 kV commercial SiC power MOSFETs has been conducted by researchers at Ohio State University, working in partnership with engineers at Alpha and Omega Semiconductor.

Highlighting the findings at IPRS, Shengnan Zhu from Ohio State University explained that one of the challenges with this type of study is that the device makers do not disclose the thickness of their gate oxides. So Zhu and her co-workers have had to estimate this from the value for the breakdown voltage, obtained by measuring the gate leakage current while ramping the gate voltage.

This approach indicated that the oxide in 1.2 kV power MOSFETs varies from 39 nm to 46 nm, assuming a dielectric breakdown field in the range 10 MV/cm to 11.5 MV/cm. The team also recorded variations in the gate current leakage over a 24 hour period for a range of gate voltages, while maintaining the device at an elevated temperature of 150 °C. Plotting the data revealed three regimes (see Figure 8). For gate voltages of 39 V or more, the gate current accelerates; for voltages of 38 V to 35 V, there are variations in the current, before a rapid hike; and for a voltages of 33 V or less, the current falls steadily.

Zhu offered a detailed explanation for this range of behaviour. She argued that for gate voltages of 39 V or more, corresponding to an oxide field of at least 9.8 MV/cm, hole trapping dominates. This takes place due to Fowler-Nordheim tunnelling of electrons into the device – these carriers trigger impact ionisation and generate electron-hole pairs, with holes driven to the semiconductor-oxide interface, where they are joined by additional holes resulting from anode hole injection. A high hole density at this boundary reduces the width of the barrier for tunnelling of electrons and leads to an increased injection of this carrier, which generates yet more holes through impact ionisation. With positive feedback at play, there is an acceleration of the gate leakage current.

For voltages of 38 V to 35 V, the oxide field is slightly lower, at 9.5-9.0 MV/cm. In this case, there is also hole trapping that enhances the electric field near the interface. However, this time subsequent electron trapping in the oxide leads to a reduction in the strength of the local field, a widening of the tunnel barrier, and ultimately a reduction in leakage current. Dominance of hole trapping, followed by electron trapping, results in a rise and then a fall in the gate leakage current. When investigating this in more detail, the team found that their conjecture is supported by variations in threshold voltage.

When the gate voltage is 35 V or below, the field across the gate oxide is restricted to no more than 8.8 MV/cm. For devices operating in this regime, electron trapping takes place, relaxing the electric field. This increases the barrier width, and in turn reduces Fowler-Nordheim tunnelling. While this takes place, the threshold voltage increases.

The insights provide by Zhu and the other speakers at IPRS showcase the progress being made to understand the reliability of the SiC power MOSFET, and differences in the devices of many manufacturers.

As insights feed in to device development and production, robustness of this transistor should increase, driving its deployment in ever more applications.

A high hole density at this boundary reduces in the width of the barrier for tunnelling of electrons and leads to an increased injection of this carrier, which generates yet more holes through impact ionisation. With positive feedback at play, there is an acceleration of the gate leakage current

2D transistors look to extend the logic roadmap

Development of WS_2 2D transistors in a 300 mm CMOS fab provides a promising pathway for scaling the transistor

BY INGE ASSELBERGHS AND IULIANA RADU FROM IMEC

THE ERA OF 'happy scaling', driven by Moore's Law and played out by the semiconductor community, has faced severe challenges since 2005. Up until then progress came relatively easily, with shrinking of the silicon transistors delivering multiple benefits, including a lower power consumption, increased performance, and a reduction in the fabrication cost per transistor. A new, more powerful generation of technology launched roughly every two years – denoted by a new node size – to double the number of transistors packed into an identical-sized chip.

The first sign that the benefits of scaling would not go on forever appeared when the node-to-node performance improvements at a fixed power – referred to as Dennard scaling – started slowing down. Applying the brakes were short-channel effects. Leakage currents started shooting up, even when not applying a voltage to the gate, due to significant reductions in gate length and a shortening of the conduction channel. Scaling also caused source and drain contacts to begin to have a dramatic impact on the channel region.

To compensate for this slow-down in device improvement and allow integrated circuits to continue to advance, much effort has been directed at developing new transistor technologies. Many directions have been pursued, including new channel materials, new transistor architectures and co-optimisation of the chip's design and its technology.

A significant architectural improvement has been the replacement of the planar MOSFET with the FinFET – the latter is now the incumbent design in mainstream chip production processes. In a FinFET, the channel takes the form of a very thin fin, positioned between source and drain terminals. Wrapping around this three-dimensional channel is a gate that provides control from three sides and combats short-channel effects. It is an architecture that has been crucial to scaling over the last decade, but will fail to provide enough electrostatic control beyond the 5 nm or 3 nm node.

Building on this approach is the vertically stacked, gate-all-around nanosheet transistor. It provides superior channel control, thanks to a gate that fully wraps around and in between the channels. At imec of Leuven, Belgium, this architecture is on our roadmap, followed by the forksheet transistor – that is a design that uses a forked gate structure to control nanosheets, and allows a far tighter *n*-to-*p* spacing than that possible with FinFET and nanosheet devices. Another candidate for well into the future is a stack of multiple channels, which could effectively extend scaling with silicon or SiGe semiconductors.

Running in parallel to this evolution, the research community has been investigating gains provided by enhancing carrier mobility. Options for the channel include strain engineering and migrating to high-mobility semiconductors, such as germanium and the III-Vs. Note that there is much interest in InGaAs: as it transports charge much faster than silicon, it promises faster speeds and lower power consumption. Innovative processes have already been developed to incorporate foreign channels with silicon CMOS, using conventional manufacturing techniques. However, just like silicon, it is unlikely that these high-mobility materials will enable sufficient gate length scaling at the very advanced nodes that lie ahead. It seems that III/V-material-based devices will find a more natural adoption in high-frequency applications – as required for (beyond-) 5G applications – that allow their high-speed properties to reach their full potential.

The promise of 2D semiconductors

To realise further gate length reduction, thinner semiconductor channels are needed to keep short channel effects under control. Migrating to thinner channels restricts the pathway for current to flow, and

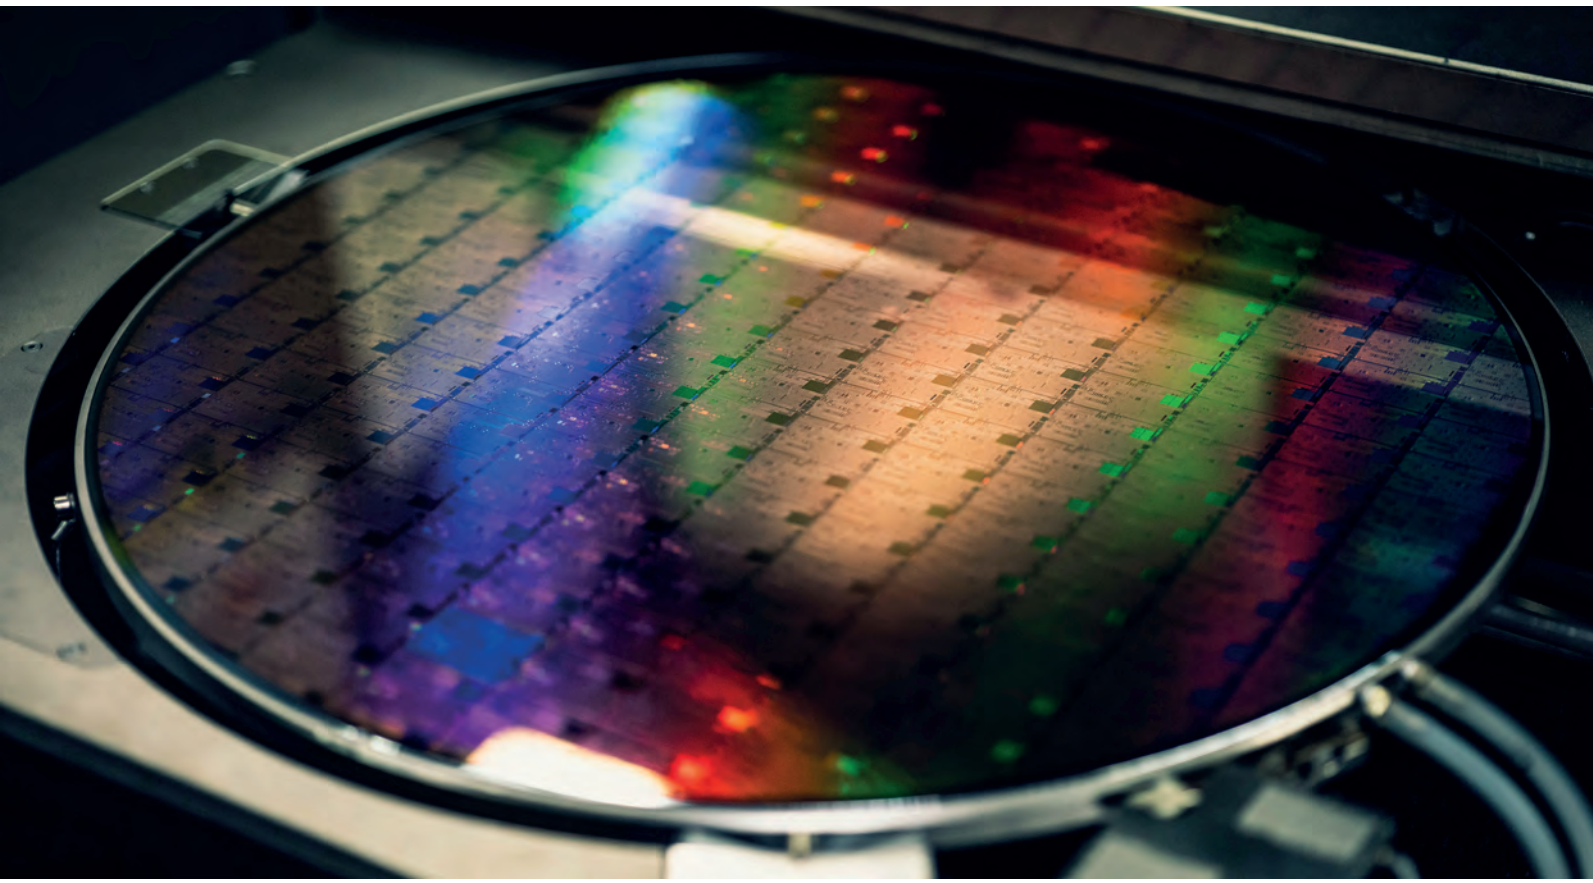

this limits the opportunity for charge carriers to leak when the device is turned off.

Offering much promise in this regard is a class of materials known as two-dimensional semiconductors. They include transition metal dichalcogenides, such as WS_2 and MoS_2. In these semiconductors atoms are arranged in layered crystals, with a single layer thickness of typically just 7 Å – small enough to make these materials a great choice for ensuring very thin channels. Results from theoretical studies highlight their potential, indicating that they maintain a relatively high carrier mobility, independent of channel thickness. It's an attribute that should enable engineers to scale gate lengths below 10 nm without having to worry about short-channel effects.

Working within the Design-Technology Co-Optimization framework, our team at imec has recently highlighted the potential for transistors with a 2D semiconductor channel to further extend the logic scaling roadmap. We anticipate that these 2D-FETs, which will most probably find their insertion point in a stacked-nanosheet-like architecture, will extend the roadmap by providing at least two technology generations. Our circuit-level evaluation of power, performance and area at a node with 36 nm gate pitch revealed that transition metal dichalcogenides in a stacked 2D-nanosheet configuration outperform silicon-based counterparts while having a reduced footprint. Note that this model employed realistic assumptions, drawing on as much experimental data as possible.

The global effort at developing transition metal dichalcogenides has led to the exploration of a variety of materials, and the identification of some of the main challenges for improving device performance. To date, most work has involved semiconductor channels made of MoS_2. Devices based on this material are the most mature, with the best experimental values for mobility getting close to the theoretical value of 200 $cm^2 V^{-1} s^{-1}$.

Recently, promising results have been reported for WS_2-based FETs. According to theoretical work, these devices have the potential to deliver an even higher performance than their MoS_2 cousins. Experimental results are also encouraging. For example, back in 2019, research reported by a TSMC-led collaboration showed that electrical characteristics, such as the on/off ratio and the sub-threshold swing, are comparable to the best recently published values for MoS_2 *n*-FETs. In its turn, imec has demonstrated functional 2D-FETs with a 30 nm-long channel just 1-2 monolayers thick.

With this class of devices a dual-gated device structure improves electrostatic control. Unlike traditional FETs, which just have a gate at the top, dual-gated siblings have a top and a bottom gate – when connected, this increases electrostatic control over the channel. Measurements on our 2D-FETs with connected top and back gates reveal that they outperform single-gated counterparts in drive current (I_{on}), transconductance and sub-threshold swing, three key metrics for evaluating short-channel effects. Another encouraging aspect of the dual-gate structure

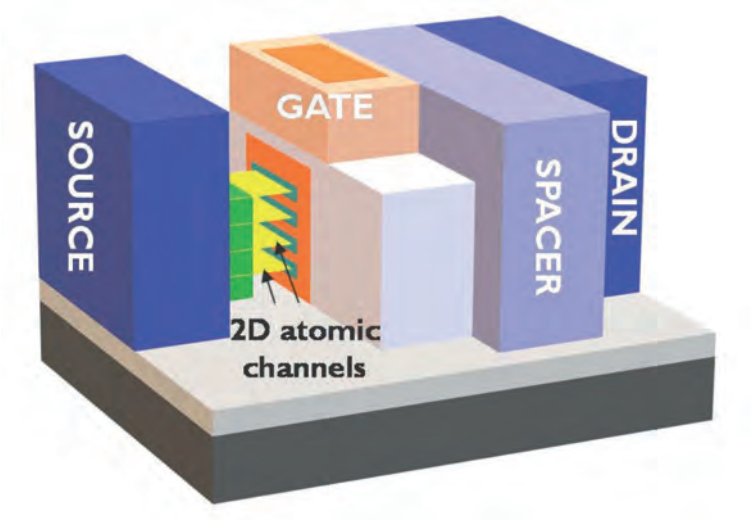

➤ Figure 1. Design-Technology Co-Optimization (DTCO) analysis suggests that stacked 2D materials with side contacts are needed to compete with silicon nanosheets.

is that it shows promise for CMOS operation. While these results suggest that a great future lies ahead for the WS_2-based FET, there is much work still to do, given that the devices we've just described were fabricated on relatively small coupons with patches of synthetic transition metal dichalcogenide material. To build on these hero results from lab-based devices and prepare a pathway for commercial success, an approach must be found that enables their adoption in a 300 mm integration flow.

A 300 mm platform

Several years ago, our team started working towards 300 mm integration of the family of transition metal dichalcogenide 2D semiconductors. This created a unique 300 mm test vehicle for 2D-FETs, allowing fabrication of functioning devices with gate lengths down to 18 nm. We have used the flow to study the impact of various processing conditions, such as channel deposition technology and gate stack formation. Based on these insights, we are developing improved process steps to enhance device performance.

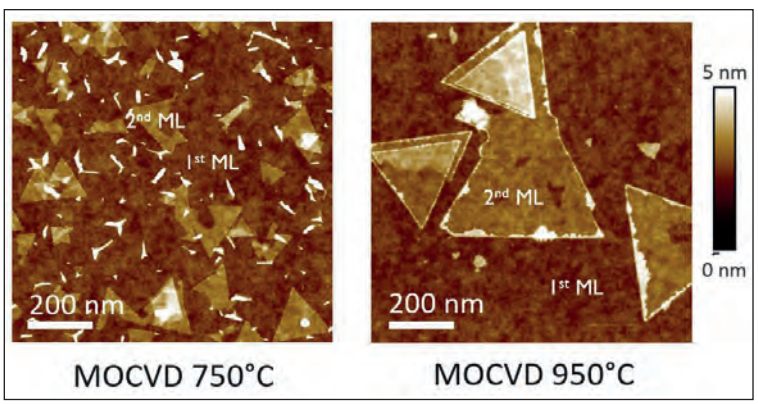

➤ Figure 2. Atomic force microscopy of films of WS_2 grown by MOCVD at (left) 750 °C and (right) 950 °C (as presented at the 2020 IEDM conference).

Our first foray into the formation of integrated transistors on the 300 mm platform pointed to high-temperature MOCVD as the best option for depositing high-quality channels, critical to realising high-performance devices. This growth technology offers thickness control of the 2D semiconductor down to a single monolayer. However, there are small multi-layer spots over the full 300 mm wafer.

We have found that the growth temperature impacts material quality. This conclusion came from the growth and characterisation of a layer of WS_2 deposited at 750 °C and 950 °C, using $W(CO)_6$ and H_2S sources. Using atomic force microscopy, we compared films grown on 300 mm silicon/SiO_2 substrates, produced using a growth time of 170 minutes, and found that a higher deposition temperature increased the size of the WS_2 grains. The higher temperature also improved crystallinity and reduced the defectivity of the WS_2 layer, according to photoluminescence and Raman measurements. It is possible that larger crystal sizes could be obtained at lower temperatures using different precursors or other deposition techniques – this is a topic of further research.

One of the challenges associated with producing WS_2 is that it is not easy to deposit insulating materials on top of the 2D surface and form a gate dielectric. An intrinsically passivated process is to blame. When traditional ALD processes are adopted, relying on nucleation of surface dangling bonds, growth only occurs at defect sites. Since the MOCVD-grown WS_2 has relatively few defect sites, this hinders direct oxide deposition.

To tackle this issue, we are investigating novel approaches for oxide deposition. They include making use of a nucleation layer of silicon seeds, deposited by a molecular beam. Another problem we shall have to consider is the low adhesion of the WS_2 to most oxides, resulting from self-passivation. This poses challenges for typical patterning schemes using hard masks.

Results on our devices formed on 300 mm wafers reveal that their performance is an order of magnitude lower than it is for reported lab devices. For example, on-current is typically just 10 μA/μm. To understand why these devices are inferior, we have considered the integrated flow. Our characterisation of dual-gated devices with source/drain side contacts suggests that channel material crystallinity is the biggest challenge to improving device performance. Success requires further breakthroughs in material growth and processing. We also have room for improvement in the processes used to form gates, dope material and add contacts.

As 2D semiconductor-based FETs are most likely to be introduced in stacked nanosheet-like architectures, we have used the 300 mm-compatible flow to identify

and overcome challenges associated with building these advanced device architectures. The creation of a stacked nanosheet requires formation of a superlattice structure, containing alternating layers of a channel material and either silicon or SiGe. Once this is formed, nanosheets could be released by selectively etching away the silicon or SiGe layers. Looking further ahead, there may come a time when 2D semiconductors are integrated in a complementary FET-like architecture, using *n*-type 2D-FETs on top of *p*-type 2D-FETs.

While MOCVD is the preferred technique for depositing high-quality 2D semiconductor channels, the high temperatures that are used threaten to exceed the thermal budget. Options to prevent this from happening include introducing different precursors and switching to alternative deposition technologies. There is also the more radical, complex approach of using a transfer process to move the 2D channel to a pre-patterned 300 mm silicon wafer.

Evaluating variability

Efforts at imec have not been limited to just developing and integrating 2D-based transistors. We have also undertaken the first-ever variability study of a large set of nanoscale lab-based 2D-FET devices, using transistors with a channel width of 115 nm and lengths of 100 nm and below. This investigation considered various sources of variability – including the thickness of the 2D-channel; the presence of bilayer islands, such as grains; and the 2D growth template – and the respective impact on electrical performance, with a focus on the sub-threshold regime.

Within this study, there have been simulations, along with the construction of devices that have a median sub-threshold slope of 80 mV/dec and maximum on-current in excess of 100 μA/μm. This research uncovered a strongly reduced sub-threshold slope and threshold-voltage variability when thinning the 2D material from three monolayers to one. This is an encouraging result, indicating that very thin channels are needed for further transistor scaling. For atom-thick channels, this work shows that the intrinsic variability is low, and comparable with silicon FinFETs. To make further progress in driving down device variability, so that it is suitable for future nodes, there needs to be better control of key process steps, like cleaning and contacts.

Worldwide progress with 2D-FETs is positioning this class of transistor as a prime candidate for extending the logic device scaling roadmap. At imec, work by our team and our colleagues has started to lay the groundwork for introducing 2D semiconductors into a 300 mm integration flow – a key requirement for industrial adoption. We have already taken significant steps by improving device performance and developing a fundamental understanding of this form of FET.

➤ *This work is the result of a collaborative effort of a broad imec team working on exploratory logic.*

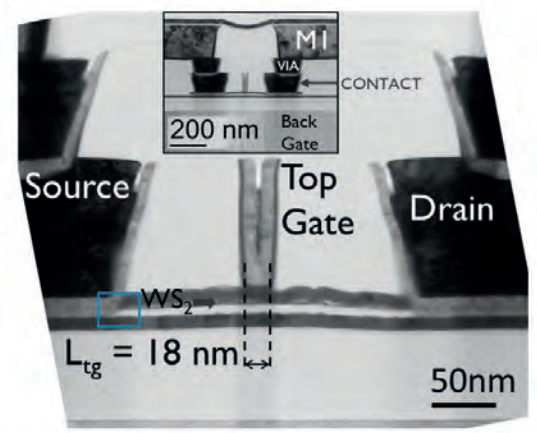

➤ Figure 3. Transmission electron microscopy (TEM) image of a 2D device fabricated with 300 mm processes (as presented at the 2020 IEDM conference).

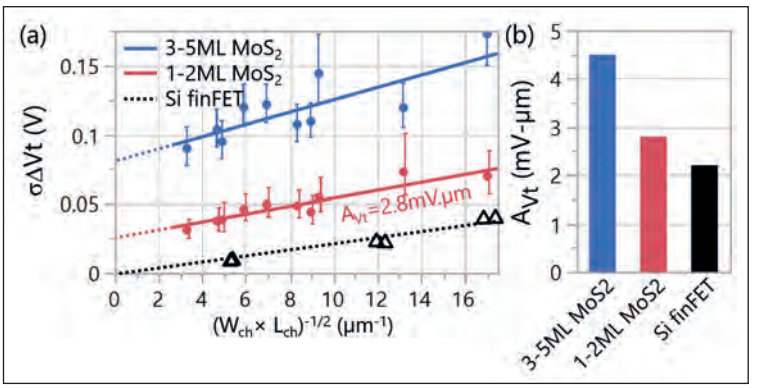

➤ Figure 4. 1-2 monolayer MoS_2 FETs (with an equivalent oxide thickness (EOT) of 2.6nm) have higher threshold-voltage variability, but their slope approaches the silicon FinFET reference (EOT=0.8nm) (as presented at the 2020 IEDM conference).

FURTHER READING

- Z. Ahmed *et al.* 'Introducing 2D-FETs in device scaling roadmap using DTCO' IEDM 2020.
- K.K.H. Smithe *et al.* 'Intrinsic electrical transport and performance projections of synthetic monolayer MoS_2 devices'2D Materials **4** 011009 (2017)
- C-C Cheng *et al.* 'First demonstration of 40nm channel length top-gate WS_2 pFET using channel area-selective CVD growth directly on SiO_x/Si substrate' 2019 Symposium on VLSI Technology Digest of Technical Papers.
- D. Lin *et al.* 'Dual gate synthetic WS_2 MOSFETs with 120μS/μm Gm 2.7μF/cm² capacitance and ambipolar channel', IEDM 2020.
- I. Asselberghs *et al.* 'Wafer-scale integration of double gated WS_2-transistors in 300mm Si CMOS fab', IEDM 2020.
- Q. Smets *et al.* 'Sources of variability in scaled MoS_2 FETs', IEDM 2020.

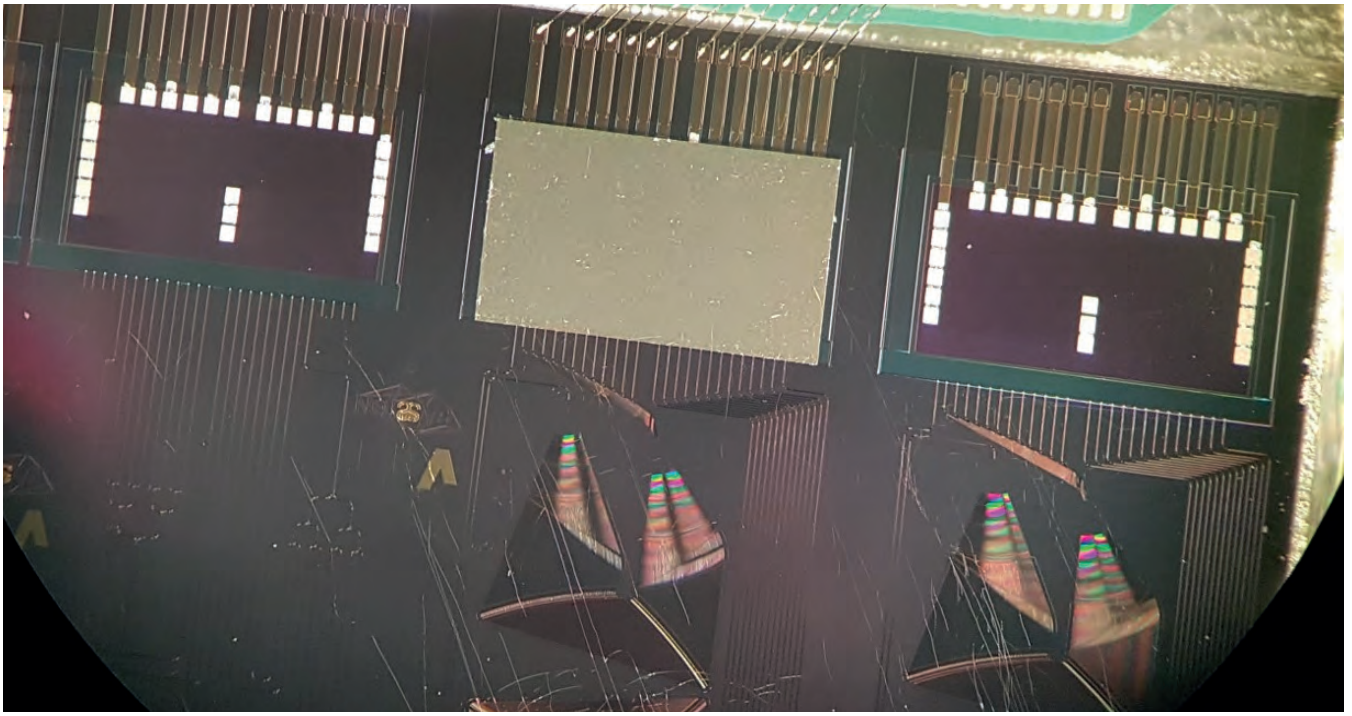

Superior switching

Data centres and optical networks could benefit from a new generation of wavelength blockers, formed by uniting the strengths of InP and silicon photonics

BY NETSANET TESSEMA, KRISTIF PRIFTI, AREF RASOULZADEHZALI, YU WANG, RIPALTA STABILE AND NICOLA CALABRETTA FROM EINDHOVEN UNIVERSITY OF TECHNOLOGY AND GIOVANNI DELROSSO, SRIVATHSA BHAT AND TIMO AALTO FROM VTT RESEARCH INSTITUTE

WHILE MANY OF YOU KNOW quite a bit about LEDs, lasers and solar cells, it's possible that you have never heard of the wavelength blocker. Yet despite its low profile it is produced in high volumes, playing a key role in today's optical networks, as it serves as the smallest unit of a wavelength-selective switch.

The deployment of wavelengths blockers is on an upward trajectory, driven by technologies such as 5G, Cloud computing and the Internet of Things. In all these applications, it is essential to efficiently utilize the available bandwidth within the optical fibre network.

In order to understand the role of the wavelength blocker, you must first get to grips with the basics of dense wavelength division multiplexing (DWDM), a technology based on bringing together optical signals at different wavelengths and transmitting all this aggregated information through an optical fibre.

A key component in these networks is the optical add/drop multiplexer (OADM). It taps (that is, drops) information at some wavelengths, while sending it on at other wavelengths. Another attribute of the OADM is that it can add new wavelengths onto an optical fibre, allowing additional signals to be transported to other destinations.

The first generation of OADMs connected spectral demultiplexers and multiplexers by a fibre patch cord. The downside of this arrangement is that it did not have any capability to be reconfigured – that required manual intervention to change the wavelength dropped or added. This weakness

stemmed from employing a fixed wavelength for port mapping in the spectral multiplexing/demultiplexing devices, and resulted in the need to physically disconnect fibre patch cords during reconfiguration. Such a state of affairs is clearly undesirable. It is far better to avoid any manual intervention – and that motivated the migration to a new form of OADM that has dynamic reconfigurability.

This is where the wavelength selective switch (WSS) has a major role to play. Without any manual intervention, this form of switch can forward an incoming wavelength at the input port to one or more output ports of an OADM. This transforms the OADM into a Reconfigurable variant, known as a ROADM.

It is this more sophisticated successor that lies at the backbone of today's high-capacity optical network, enabling dynamic transport of traffic in a manner that leads to efficient use of bandwidth and network resources. Note that the WSS enables independent switching of any wavelength to any port, without restrictions imposed by the switching of other wavelengths.

An additional role played by the WSS is to provide dynamically reconfigurable connectivity between the switches at the top of the racks within a data centre network. Helping to fulfil this task is a software defined network, which automates provisioning of network resources according to current traffic demands.

Marrying materials

Lying at the heart of a WSS are a number wavelength blockers (WBLs), each with one input port and one output port. Several research groups have demonstrated WBLs using silicon photonics, but they are compromised by a poor on/off switching ratio and a high insertion loss, attributed to a large port count. As this platform does not have the capability to include a gain section that compensates for on-chip power losses, it is not an ideal candidate for implementing a WBL.

Another material for making a WBL is InP. Its strength is that it can compensate for on-chip losses by incorporating semiconductor optical amplifiers (SOAs), which also provide a high on/off switching ratio. But the limited wafer size of InP hampers scaling to a large wavelength and port count.

The best way forward is to adopt a hybrid approach that addresses all the limitations just described. For this architecture, silicon photonics is used for the passive parts – that is, the demultiplexer and multiplexer – and InP is employed to realise gain, forming an SOA that provides on/off switching and amplification. With this combination, scaling up the port and wavelength count draws on the merits of silicon photonics, while InP provides switching with a high on/off ratio and amplification.

Our team at Eindhoven University of Technology, The Netherlands, working in partnership with engineers

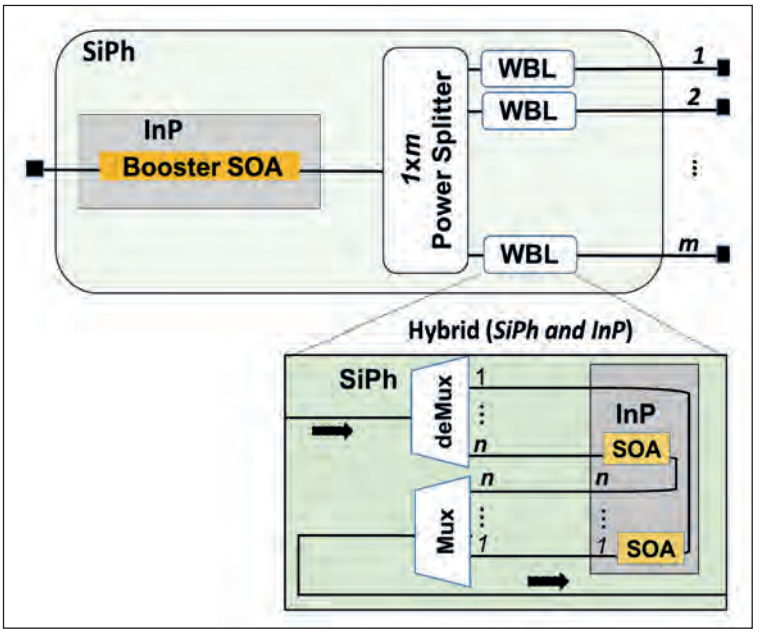

➤ Figure 1. Architecture of hybrid integrated wavelength selective switch (WSS). A hybrid WSS can be developed with one input port and m output ports by integrating *m* hybrid WBL modules. The WBL has 1 to *n* channel mapping at the demultiplexer and *n* to 1 channel mapping at the multiplexer. There are *n* SOA switching gates, formed from InP.

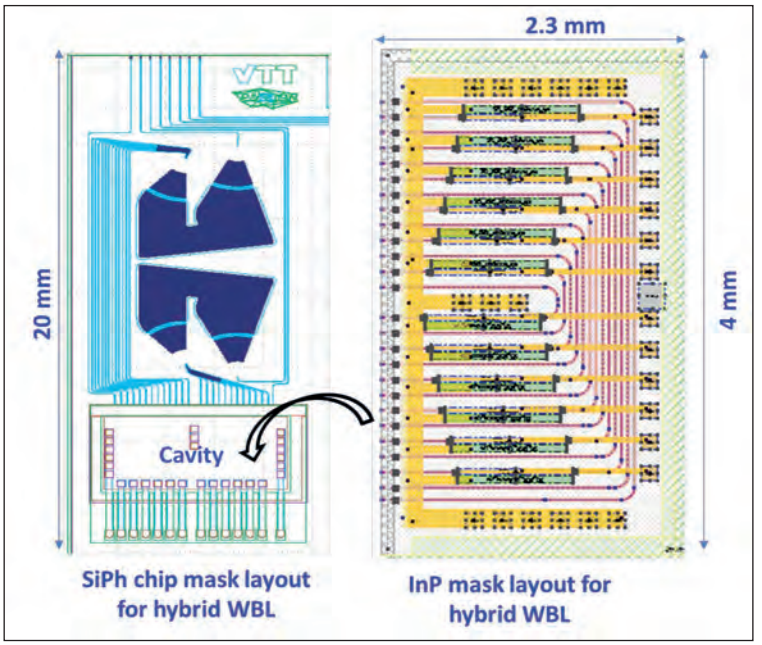

➤ Figure 2. The mask layout of silicon photonics (SiPh) (left) and InP (right) PIC design for the implementation of a hybrid WBL module with *n* =12 wavelength channels (10 functional and 2 test channels). Channel spacing is 0.8 nm (100 GHz).

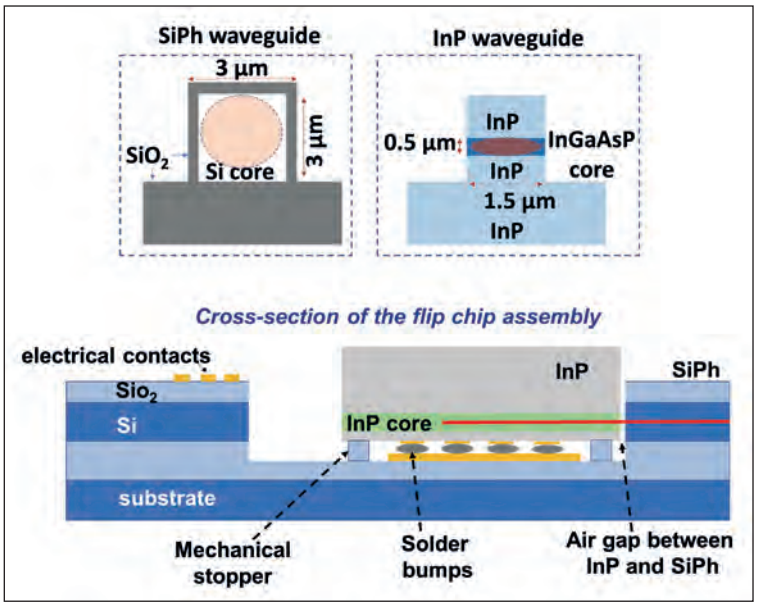

➤ Figure 3. Cross-sectional diagrams of SiPh/InP waveguides (top) and a flip-chip assembly process (bottom).

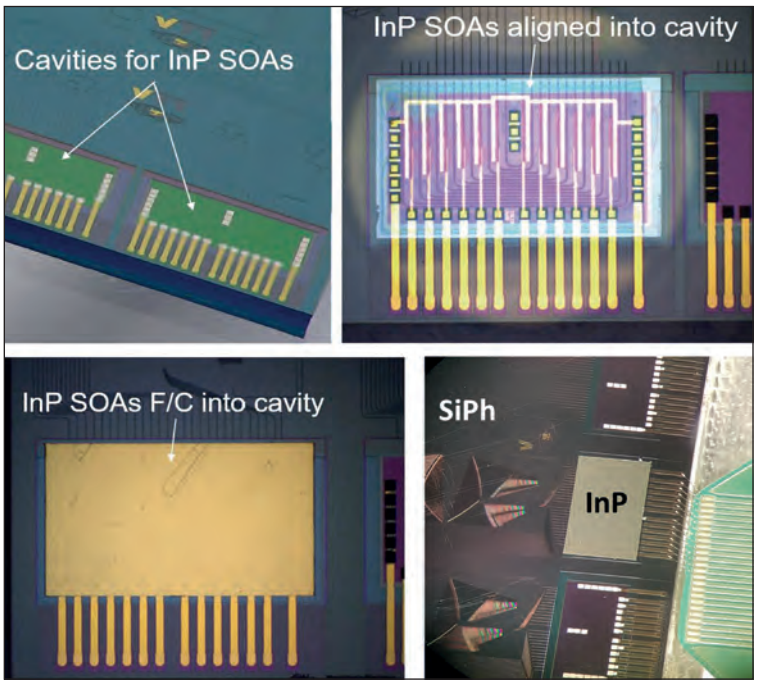

➤ Figure 4. Images of a: fabricated SiPh chip with cavities (top left); SOA chip aligned into the cavity during flip-chip bonding (FCB) (top right); InP chip placed into a SiPh cavity (bottom left); and a hybrid WBL after FCB (bottom right).

at VTT Research Institute, Finland, has broken new ground by fabricating the first ever hybrid WBL. Collaborating through an EU project named PASSION – its full name is Photonic technologies for prograMmable transmission and switching modular systems based on Scalable Spectrum/space aggregation for future agIle high capacity metrO Networks – we are taking on the ambitious goal of developing a software-defined network featuring transceiver and switching node technology for high capacity metropolitan optical networks. The goal is to reach a capacity as high as 8 Tbit/s per second per polarization, using 160 wavelength channels within the C-band, with transmission routed down a single fibre.

The motivation for our development of our hybrid WBL is that it can play a crucial part in a scalable and modular WSS for a high capacity switching node within a metropolitan optical network (see Figure 1). This technology is based on a broadcast-and-select scheme, with a WDM signal broadcast by a power splitter and selected by WBL modules at the output ports. To compensate for loss associated with a power splitter, signals can be boosted by an SOA, formed in InP. At each output port of the WSS, a WBL is used to pass or block a wavelength channel by turning gate SOAs on or off respectively. If the port count of the WSS needs to increase, a new WBL switching module is added.

Demonstrating devices

We have demonstrated the proof-of-concept of our hybrid WBL by designing and fabricating a pair of chips on silicon photonic and InP platforms (see Figure 2 for the mask layouts of two 1 x 12 arrayed waveguide gratings formed using silicon photonics and an array of 12 SOA gates on InP). By placing SOA switching gates in a U-shape, we allow input/output straight waveguides to be accessed from the same side. Flip-chip bonding unites the silicon photonics and InP chips, and realizes the hybrid WBL. As well as housing de-multiplexer and multiplexer arrayed waveguide gratings, the silicon photonics chip has a cavity, designed to accommodate the InP chip during flip-chip bonding assembly.

In our hybrid chip, waveguides are formed in InP and silicon (see Figure 3). Those formed in silicon, based on VTT's 3 µm by 3 µm waveguide technology, provide a polarization-insensitive performance that doubles the transmission capacity per wavelength channel – data can be simultaneously transmitted in transverse electric and the transverse magnetic polarizations. The InP waveguide features a 0.5 µm-thick InGaAsP core. Due to the difference in the size of this core and that in silicon, to ensure that hybrid coupling losses are below 3 dB when the waveguides are brought together, their air-gap must be less than 2 µm, and the alignment accuracy must be below ± 0.5 µm in the vertical and lateral directions.

Our flip-chip bonding process involves gold-coated contact pads on the InP chip and tin-coated contact pads on the silicon photonics chip. These two entities are united by thermal compression at an elevated

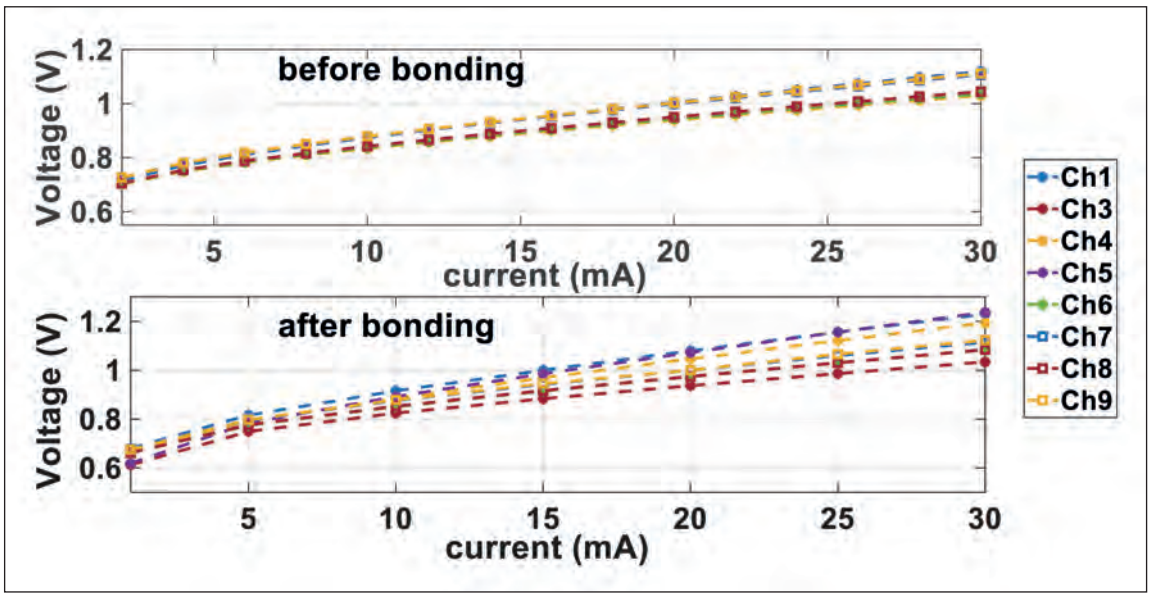

➤ Figure 5. Current-voltage curves before and after flip-chip bonding.

temperature of 220°C. Once bonded, the patterned stopper on the silicon photonics chip provides mechanical support to the InP chip. Use of lithography allows us to vertically align the top contact areas of the silicon photonics mechanical stoppers and the recessed areas in the InP chip (see Figure 4 for images of the: silicon photonics chip with cavities after fabrication; the InP chip aligned to the silicon photonics cavity; and the hybrid WBL switch after assembly).

Flip-chip bonding provides us with a placement accuracy below 0.5 μm in the vertical and lateral directions. However, a limitation of the visual system that we use for bonding is that it prevents us from precisely controlling the airgap between the two chip facets, resulting in an air gap of the order 8 μm; this leads to excess coupling losses between silicon photonics and InP waveguides.

Channel characteristics

We have characterized our hybrid WBL electrically and optically. Of the ten wavelength channels, eight are working – they are channels Ch1, Ch3, …, Ch9. The most likely reason that two channels are not working is optical damage at the waveguide facets of the SOAs, occurring during the flip-chip bonding process.

Plotting the current-voltage of the SOA gate switches before and after integration shows that the variation in resistance is only 4.5 Ω (see Figure 5). Transmission characteristics of the eight wavelength channels reveal that the typical on-off ratio is around 30 dB (see Figure 6).

After optimally tuning the current for each gate SOA to minimize fibre-to-fibre loss and the effect of back

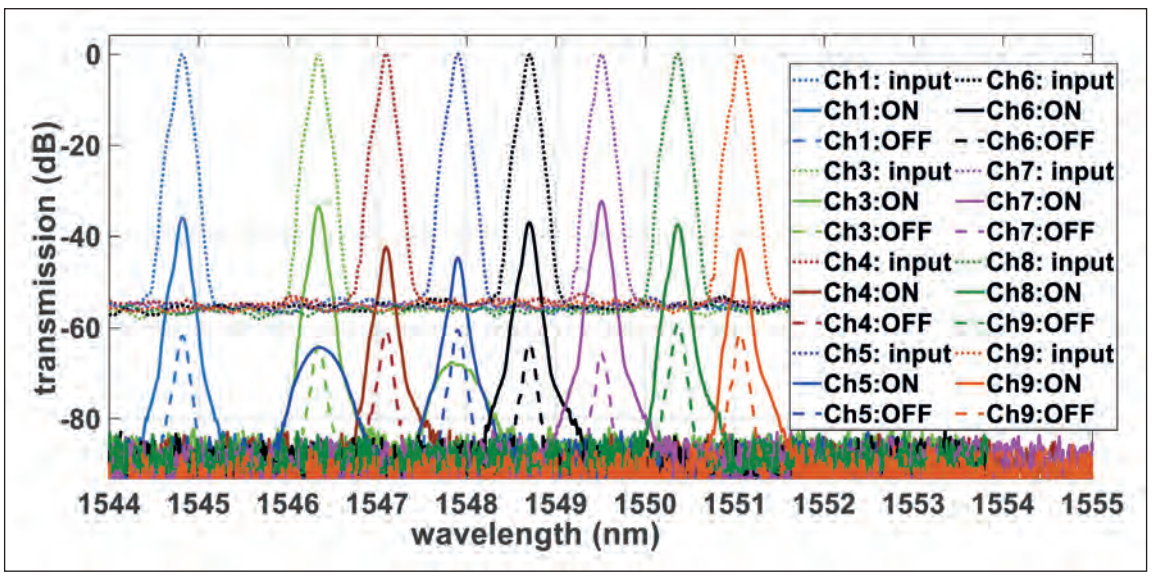

➤ Figure 6. Transmission characteristics of eight wavelength channels when SOA gates are turned On and Off for an input optical power of 0 dBm. Input optical signal (dotted line), output optical signal of the hybrid WBL switch ON (solid line) /Off state (dashed line).

Channel number	1	3	4	5	6	7	8	9
Fiber-to-fiber loss (dB)	36	33.4	42.3	44.8	37.0	32.2	37.5	42.8
SOA current (mA)	40	80	80	80	90	40	80	28

Table 1. Fibre-to-fibre loss of each channel and the corresponding current of the SOA gate.

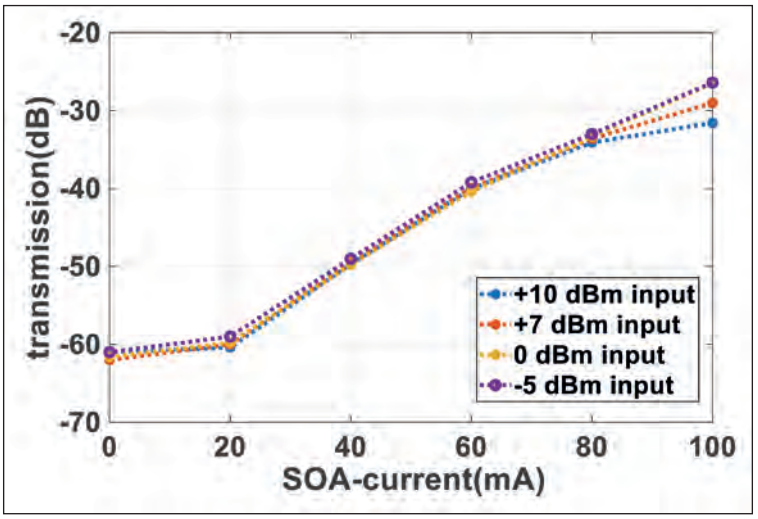

Figure 7. Fiber-to-fiber transmission and gain characteristics of a SOA (Ch3). The fibre-to-fibre loss is of the order of 33 dB at 80 mA for all input power values. The loss decreases to 26 dB at a SOA current of 100 mA; for input power of 0 dBm and -5 dBm. Effects of a SOA saturation on the fibre-to-fibre loss are slightly observed at 100 mA for input powers of 7 dBm and 10 dBm, resulting a fibre-to-fibre losses of 29 dB and 31.57 dB, respectively.

reflection on the hybrid WBL, fibre-to-fibre losses for Ch1, Ch3, Ch6, Ch7, Ch8 range from 32.2 dB to 37.5 dB. Such high losses stem from the large air gap, which needs to be below 2 μm in the horizontal direction to keep hybrid coupling losses below 3 dB. To help address this loss, we plan to move to on-chip integrated spot-size converters, because they relax the tight airgap tolerance. Another improvement that we will be making is to introduce angled waveguides – they alleviate the impact of back reflection and fully employ the gain of SOAs.

Investigating the gain characteristics of the gate SOA for Ch3, for varying current and varying input power, uncovered a significant variation in fibre-to-fibre loss. It varies from 60 dB when the SOA has no drive current and is in an absorption state, to 26 dB, due to gain provided by the SOA at 100 mA (see Figure 7, and also Table 1, which summarizes the fibre-to-fibre loss of each channel and the corresponding current of the SOA gate).

To evaluate data transmission, we have employed a tuneable laser source to generate the wavelengths corresponding to the channels of the WBL, and used 10 Gbit/s NRZ 2^7-1 and 20 Gbit/s NRZ 2^{31}-1 PRBS patterns (see Figure 8). Using 10 Gbit/s signals, measurements of the bit error rate at various received powers produced clearly opened eye diagrams for Ch1, Ch3, Ch6, Ch7, and Ch8 (see Figure 9). Power penalties range from 0.3 dB to 1.5 dB.

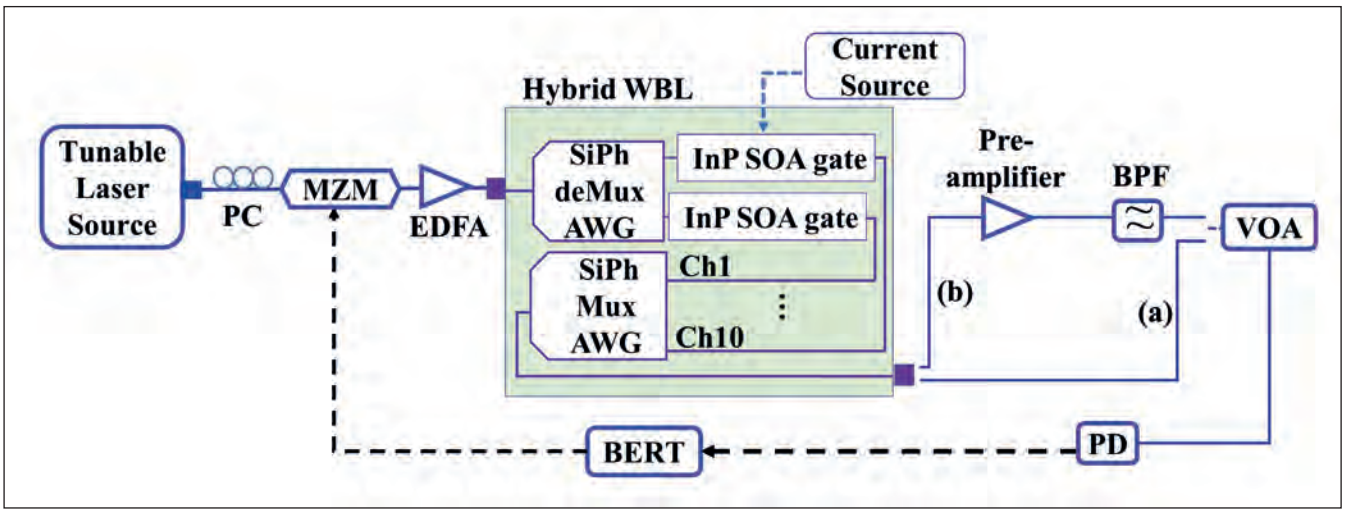

Figure 8. Experimental set-up for data transmission at (a) 10 Gbit/s and (b) 20 Gbit/s. Light is coupled in and out of the hybrid WBL switch with lensed fibre. To compensate for losses in this switch, input optical modulated data is boosted by 17 dBm. For the 20 Gbit/s signal, after light is transmitted through the WBL, it is boosted by a second preamplifier, so that it is sufficient for the sensitivity of the receiver. Inserting a bandpass filter with a 3 dB bandwidth of 200 GHz removes the out-of-band noise of the pre-amplifier. For each channel the current of the SOA is optimised to minimize power penalties.

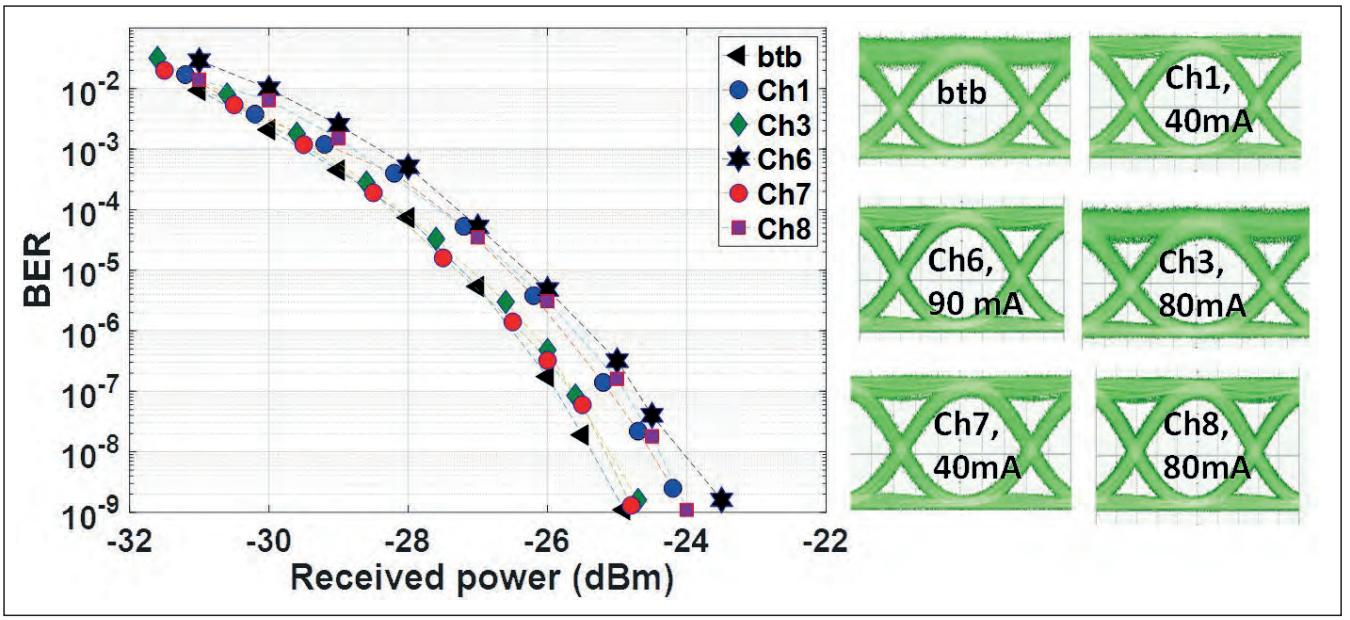

Figure 9. The bit error rate (BER) as a function of received power and eye diagram at 10 Gbit/s.

When data transmission capacity increased to 20 Gbit/s, we still observed clear eye diagrams and obtained error-free performance (see Figure 10). Power penalties ranged from 0.7 dB in the case of Ch3 to 2.8 dB in the case of Ch8 (see Table 2 for the power penalty at a bit error rate of 1 x 10^{-9}).

These results, and those at 10 Gbit/s, showcase the suitability of our modular hybrid WBL for deployment in a high-capacity switching node. For this application, capacity may be scaled by adding new modules in a pay-as-you-grow manner.

The road ahead

Success to date has enabled us to demonstrate a proof-of-concept hybrid WBL module that is based on silicon photonic passives and InP actives. For this we used a flip-chip bonding process to produce just four samples, undoubtedly a small number for such an assembly process. We are now laying the foundations for trimming hybrid coupling loss between the silicon photonic and InP waveguides, and maximizing the yield of the flip-chip bonding assembly process.

Addressing the low loss in our module will unlock the door to a high degree of integration, involving a large port count and wavelength count WSS. In turn, this will equip our device for deployment in current metropolitan optical networks and data centre networks.

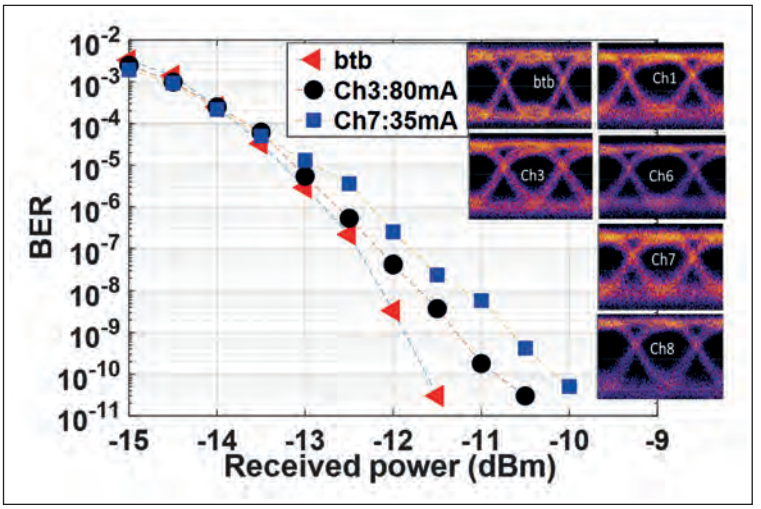

Figure 10. The bit error rate (BER) as a function of received power and eye diagram at 20 Gbit/s.

For both of these applications, our WBL benefits from the large degree of integration associated with silicon photonics. We can scale the number of wavelength channels per WBL up to 40 and enjoy a wide operating range, thanks to the optical transparency of the silicon photonics VTT waveguide technology through O and L bands. Another strength of our technology is that the

Channel	Ch1	Ch3	Ch6	Ch7	Ch8
Penalty (dB) at 1E-9	1.3	0.7	2	1.4	2.8
Current (mA) of SOA	35	80	90	35	60

Table 2. Power penalty at a bit error rate of 1 x 10^{-9}.

SOA switching gates span a substantial spectra range, and can be optimised for O, E, S, C and L bands. This provides multi-band WSS functionality, a major asset in high-capacity metro and data centre networks.

One of the major goals of the PASSION project is to develop technolgy for a software-defined network. This features programmable metro nodes, referred to as Level-1 nodes, designed to meet the agile capacity demands of the metropolitan network (see Figure 11 for an illustration of the metro-core architecture and the details of how this can aid these networks). As the network grows, new WSS modules can be added in the drop and add directions.

Our WSS could also have a role to play in data centre networks, which provide full connectivity to thousands of servers hosting various applications, such as cloud computing, big data and IoT. In these networks, optical switches can help to support high bandwidth, and are a promising candidate to replace the electronic switches that struggle in this regard. A move to optical switches can also trim power consumption in these networks by eliminating electrical-optical-electrical conversion, which is essential when using electrical switches.

In data centre networks, WSS are used to create a reconfigurable connectivity between the top-of-the-rack switches. The technology also enables the introduction of optical top-of-the-rack switches, which create dynamic connectivity within several servers in a single rack. Using our hybrid WBL, switching can take place in few nanoseconds, fulfilling the speed requirements of packet switching needed for broadband operation within the datacentre.

While there is still much to do to improve and refine our hybrid chip, by uniting the best attributes of silicon photonics with the strengths of InP, this device clearly has much to offer.

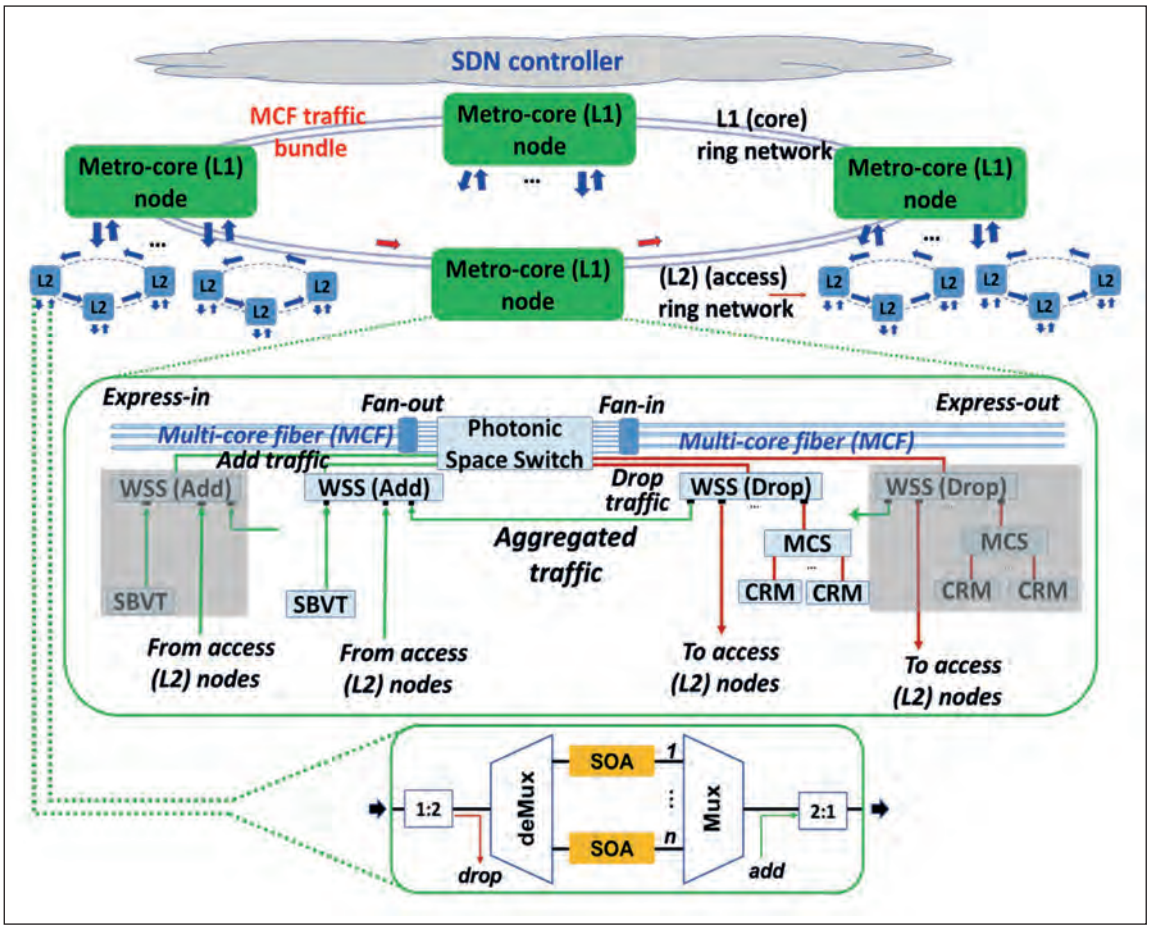

➤ Figure 11. In this network, the L1 node is a high-capacity ROADM node, equipped with colour-less, direction-less, contention-less capability. This node can transparently switch space-division multiplexed systems via Fan-out and Fan-in of multi-core fibre bundles. The photonic space switch handles traffic flow in the spatial domain in the express in/out path. Incoming express-in traffic is either sent forward directly as by-pass traffic, or directed to the disaggregate switch in the drop path as drop traffic. Operating in a similar manner, this switch adds traffic in the express-out path. The WSS (drop) acts as a disaggregating switch, while the WSS (add) provides an aggregating switch. Working together, they offer dynamic reconfigurability. When traffic has a common destination it is bundled and sent forward in the express-out path so that, at the next L1 node, express ports transparently direct the traffic. The L1 node also enables adding/dropping of traffic from/to the access Level-2 (L2) network, which has L2 nodes with 2-degree ROADMs based on a WBL and 1:2 splitter and 2:1 combiner. In addition, there are local add/drop functionalities at the L1 node. They are embedded at the WSS: in the add directions via sliceable bandwidth variable transmitters (SBVTs); and in the drop direction via multi-cast switches (MCS) and coherent receiver modules (CRM).

Exposing the origins of substrate conduction within GaN-on-silicon HEMTs

Experiments uncover thermal donors and diffused group III elements as the causes of substrate conduction in GaN-on-silicon HEMTs

ONE OF THE BIGGEST questions puzzling the nitride community concerns the origin of substrate conduction in GaN HEMTs. This mysterious malady saps away RF power, and understanding its cause could provide a crucial step in closing the performance gap between GaN-on-silicon HEMTs and those with a SiC foundation.

Shedding new insights on this matter is a UK collaboration between researchers at the University of Cambridge and Cardiff University.

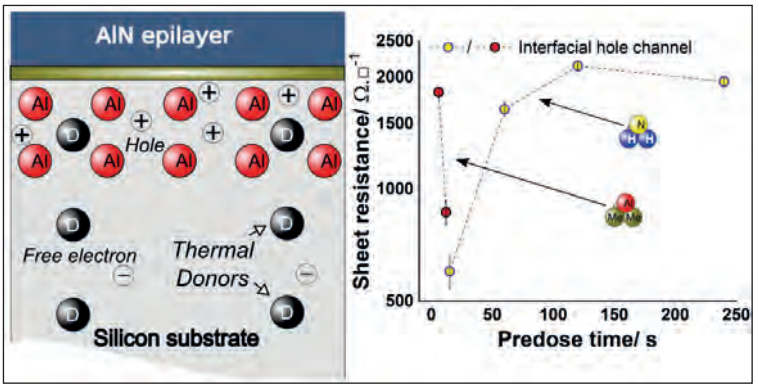

➤ Substrate conduction comes from an *n*-type layer that is distributed throughout the substrate and has a purely thermal origin, and a *p*-type layer that is localised on the substrate side of the AlN-silicon interface and induced by group-III diffusion of the metal-organic precursor.

Saptarsi Ghosh from the University of Cambridge told *Compound Semiconductor* that he and his co-workers have shown, for the first time, that the source of unintentional substrate conductivity for GaN-on-silicon stems from not one but two separate origins. There is an intrinsic contribution from thermal donors, as well as an extrinsic factor at play, group-III diffusion.

The team have shown that by carefully controlling the MOCVD growth process they can reduce the impact of this diffusion on device performance.

To fathom the causes of substrate conduction in RF GaN-on-silicon HEMTs the team carried out many experiments over about 18 months. "When we saw results that could not be explained by what was known, we had to plan new experiments to dig deeper."

Investigations involved scrutinising a portfolio of home-grown samples, all produced on high-resistivity Czochralski-grown substrates. Float-zone high-resistivity wafers were ruled out for this study, due to their higher cost and propensity for plastic deformation.

Two reference samples were included in the investigation – an unprocessed substrate, and another annealed at 1080 °C for 30 minutes. The team also employed this annealing step for the fabrication of six more samples, each with a 250 nm-thick AlN layer. Prior to the growth of this epilayers, the researchers treated the surface of the substrate with either an ammonia pre-dose for 15 s, 60 s, 120 s or 240 s, or a trimethylaluminium pre-dose for 6 s or 12 s.

Eddy current measurements revealed that annealing reduced the sheet-resistance from beyond the limit of the measurement system to 596 Ω/square, a fall of two orders of magnitude. By repeating these measurements after back-side thinning to eliminate any backside contamination, and probing the material with secondary-ion mass spectrometry and depth-dependent Hall measurements, the team pinpointed the cause of unintended conductivity: electrons distributed throughout the thickness of bulk silicon.

The team suspects that the most likely cause of this bulk conductivity is the generation of one or more thermally induced donors. Oxygen participates are viewed as a plausible candidate, due to the high-levels of intrinsic oxygen that are present in silicon substrates grown by the Czochralski process.

Additional experiments by the team identified the second cause of substrate conduction: group-III acceptors that result in a *p*-type conductive layer at the AlN-silicon interface. The duration of the pre-dose with either ammonia or trimethylaluminium impacts sheet resistance. With a judicious choice, values as high as around 2000 Ω/square are possible.

The elemental diffusion that causes *p*-type conductivity may occur before, during or after the nucleation of AlN. Efficacy of the pre-dose is thought to be associated with inhibiting diffusion during the various stages of this nucleation process.

Ghosh and co-workers are planning to carry out electromagnetic simulations to correlate substrate conductivity with RF loss."In parallel, we are planning growth experiments to figure out in which part of the aluminium growth step most of the diffusion occurs."

REFERENCE

➤ S. Ghosh *et. al*. ACS Appl. Electron. Mater. **3** 813 (2021)

Powering underwater vehicles

Spacing thicker fingers closer together increase the conversion efficiency of GaInP power converters

RESEARCHERS at Chiba Institute of Technology, Japan, are claiming to have set a new benchmark for the performance of GaInP power converters. Their latest devices deliver a power conversion efficiency of 43 percent under illumination from a 638 nm laser at a power density of 17 W cm^{-2}.

These GaInP devices are promising candidates for photo-receivers in optical wireless transmission systems powering autonomous underwater vehicles. For this application, lasers should emit in the range 400 nm to 700 nm, as this is where transmittance in sea water is high. That requirement rules out GaAs power converters; while they have an efficiency of over 80 percent between 825 nm to 850 nm, this falls to below 30 percent at wavelengths of 700 nm or less.

Last year the team reported the results of their first GaInP power converters incorporating a distributed Bragg reflector (DBR). Introduction of this reflector boosted power conversion efficiency to 43 percent, but only for irradiation densities below 3 W cm^{-2}.

Significantly higher densities are needed for autonomous underwater vehicles. The team reasons that for a vehicle with a power consumption of 350 W, the laser's power needs to be more than 1458 W, assuming an optical loss of 40 percent and a power conversion efficiency of 40 percent. And in that case, if the convertor cell has a 10 cm diameter, the laser intensity power density will need to be more than 11.1 W cm^{-2}.

To develop a new generation of devices that work better in this regime, the team has investigated power converters that combine a DBR with better surface electrodes, which trim lateral spreading resistance.

The engineers began by fabricating four GaInP power converter cells, all with dimensions of 2.4 mm by 2.4 mm (see Figure for details of the designs). This quartet included devices with and without a DBR that featured 12 μm-wide Ti/Pt/Au finger electrodes in a check pattern at a 485 μm pitch; and a pair of variants, featuring the same metal stack for the finger electrodes, but this time with an 11 μm width and a 115 μm pitch.

Measurements of the cells with a wider pitch showed that for a power density of around 0.1 W cm^{-2}, the introduction of the DBR increased the conversion efficiency via a rise in open-circuit voltage. But at higher power densities these benefits were offset by Joule heat loss, due to increases in electrical and thermal resistance associated with the DBR.

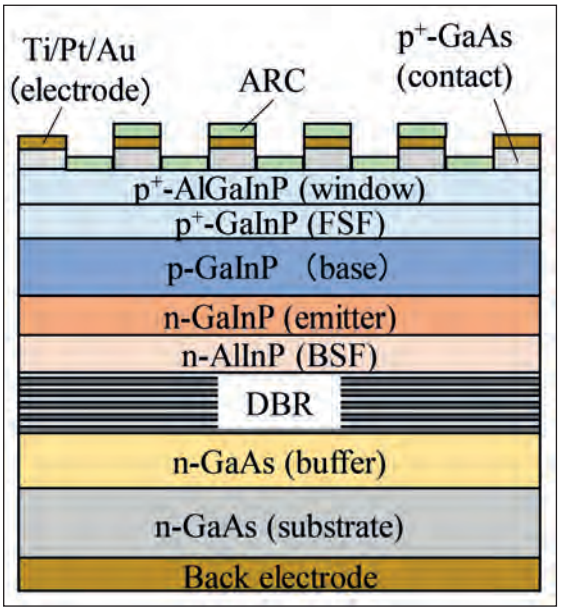

➤ Studies show that reducing the finger pitch boosts the power conversion efficiency.

Narrowing the pitch led to a significant improvement in the fill-factor and efficiency at high incident power densities. To reduce the lateral spreading resistance even further, the engineers increased the thickness of the AlGaInP window layer from 30 nm to 125 nm, and widened the front surface-field layer from 100 nm to 300 nm. With these additional refinements, devices with the 115 μm pitch produced an efficiency in excess of 43 percent from 2.8 W cm^{-2} to 17.1 W cm^{-2}.

Although the fingers in these electrodes have a resistivity of just 4-6 x 10^{-8} Ω m, the resistance of one finger is as high as 9.6 Ω, assuming a 1.0 mm length, a width of 11 μm and a thickness of 475 nm. Consequently, for the 12 parallel finger electrodes in the illuminated area, the total resistance is 0.8 Ω. To trim this resistance, the team increased the thickness of the fingers to 900 nm. This led to conversion efficiencies of 46.0 percent and approximately 45 percent under irradiation densities of 1.1 W cm^{-2} and 8.8 W cm^{-2}, respectively.

Team spokesman Shiro Uchida says that even better results could come from an even smaller pitch, using thicker electrodes and adding more silver and gold into the electrode stack – this could reduce resistance.

REFERENCE
➤ Y. Komuro *et al.* Appl. Phys. Express **14** 041004 (2021)

AlN substrates show much promise for GaN HEMTs

Free-standing AlN substrates have the potential to produce record-breaking X-band HEMTs

ENGINEERS at Fujitsu are claiming to have broken new ground for the performance of the RF GaN-based HEMTs operating in the X-band. Their device, made on an AlN substrate, delivers a power density of more than 15 W mm^{-1}, and far higher values could follow with the introduction of field plates and optimisation of device dimensions.

According to the team, the other reports of GaN-based HEMTs on AlN substrates have used AlGaN channels to improve higher-temperature performance. However, with that design alloy scattering in the ternary diminishes electron mobility and holds back performance. To avoid these issues, the team from Fujitsu employs a GaN channel.

Working with an AlN substrate promises to offer the best of both worlds, as this has a higher conductivity than GaN and avoids the lattice mismatch issues of SiC. But AlN substrates are pricey and limited in size.

Spokesman for the team, Shiro Ozaki, told *Compound Semiconductor* that he hopes that increased use of these devices for ultraviolet laser diodes and LEDs will help to address these issues. He expect that this will ultimately lead to the availability of affordable, 3-inch diameter substrates.

Ozaki and co-workers have produced a range of GaN HEMTs with a 0.25 μm gate length, a 50 μm gate width and a gate-to-drain length of 3 μm. One of these transistors has been produced on a 1-inch diameter AlN substrate with a threading dislocation density below 10^3 cm^{-2}, and another five variants have been fabricated on SiC: one, like that on AlN, has a channel thickness of 200 nm, to trim the drain-leakage current; and the other four have 1000 nm-thick channels and different AlGaN barrier compositions, to evaluate the impact of carrier density on DC characteristics.

Measurements of the drain current as a function of the drive voltage, which is the difference between the gate-to-source voltage and the threshold voltage, uncovered a leakage current path for HEMTs that have a thick channel. Devices with the thinner, 200 nm-thick channel supressed this leakage. The team also found that the AlGaN buffer with a high aluminium composition, employed in the device grown on the AlN substrate, sharply reduced the pinch-off state and had a lower leakage current than the similar device made on SiC.

Simulations show that the HEMT on the AlN foundation has a high carrier density, despite its aluminium-rich back-barrier. And, according to these calculations, when the channel thickness is reduced, the electric field in this $Al_{0.3}Ga_{0.7}N$ barrier diminishes, supressing electron tunnelling into this layer and quashing the leakage current.

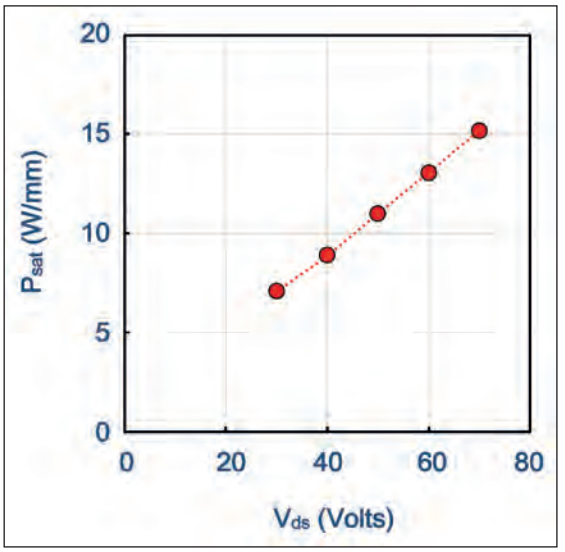

➤ The X-band output power of the HEMT made on an AlN substrate increases with operating voltage. Introducing a field plate should enable an increase in the voltage and the output power.

To ensure excellent output power characteristics, HEMTs must produce a high drain current at a high blocking voltage. Measurements of the five devices made on SiC show that gains on one of these fronts are compromised by losses on the other, while the HEMT on AlN provides an unparalleled performance.

The team's GaN-on-AlN HEMT, which has a total gate width of 1 mm, produces a peak power-added efficiency of 49.1 percent when operating at 70 V, using 10 μs pulses at a 1 percent duty cycle. Associated output power is 41.7 dB, equivalent to 14.7 W mm^{-1}, and gain is 9.6 dBm.

For devices grown on SiC, a power density of 30.6 W mm^{-1} has been reported in the X-band. "Our target is higher than that," says Ozaki.

To meet that goal, Ozaki and co-workers intend to increase the operating voltage for their HEMTs to more than 100 V. "To realise this, field plate optimisation will be a key technology."

REFERENCE

➤ S. Ozaki *et al*. Appl. Phys. Express **14** 041004 (2021)

INDUSTRY SERVICES DIRECTORY

To promote your products and services contact:
Jessica Harrison
T: +44 (0)2476 718970
E: jessica.harrison@angelbc.com

AIXTRON
Tel: +49 2407 9030 0
Fax: +49 2407 9030 40
Email: info@aixtron.com
www.aixtron.com

APPLIED ENERGY SYSTEMS
Tel: +001 610-647-8744
info@appliedenergysystems.com
www.appliedenergysystems.com

APPLIED MATERIALS®
make possible
Tel: +1-408-727-5555
www.appliedmaterials.com

ARM PURIFICATION
BY APPLIED ENERGY SYSTEMS
Tel: 1 610-647-8744
purification@appliedenergysystems.com
www.armpurification.com

BASF We create chemistry | Chemetall expect more
Perfecting Surface Finish
Tel: +1-408-387-5340
Fax: +1-408-809-2883
Email: micro.sales@basf.com
www.chemetallprecisionmicrochemical.com

BRUKER
Tel: +44 (0)191 332 4700
Fax: +44 (0)191 332 4800
sales.semi@bruker.com
www.bruker.com

CS CLEAN SOLUTIONS
Tel No: +49 (89) 96 24 00-0
Fax No: +49 (89) 96 24 00-122
Email: sales@csclean.com
www.csclean.com

VIGO SYSTEM Epitaxy Division
ENT.
VIGO SYSTEM
Tel: +48 538 248 661
Email: ent@vigo.com.pl
www.ent-epitaxy.com

evatec
process systems
Tel: + 41 81 403 8000
Fax: + 41 81 403 8001
www.evatecnet.com

GATAN AMETEK®
info-gatan@ametek.com
www.gatan.com

GT ADVANCED TECHNOLOGIES
Tel: +1 (603) 883 5200
Email: sicinfo@gtat.com
Website: www.gtat.com

HORIBA
Scientific
Tel: +1 732 494 8660
joanne.lowy@horiba.com
www.horiba.com/scientific

INNOVION
FOUNDRY ION IMPLANTATION
Tel: +1 408 501 9100
Fax: +1 408 501 9110
Email: sales@innovioncorp.com
www.innovioncorp.com

III/V-Reclaim GaAs InP
Tel: +49 8728 911 093
Fax: +49 8728 911 156
Email: sales@35reclaim.com
www.35reclaim.de

Non Contact Measurements
Resistance, Resistivity
Wafer Thickness
KITEC microelectronic technologie
Tel: +49 (0)8122 86676
Fax: +49 (0)8122 86652
Email:info@kitec.de
www.kitec.de

kSA
k-Space Associates, Inc.
Tel: +1 734-426-7977
Email: requestinfo@k-space.com
www.k-space.com

MASIMO
semiconductor
Tel: +1 603 595 8900
Email: sales@masimosemi.com
www.masimosemiconductor.com

MBE
KOMPONENTEN | DR. EBERL
Tel: +49 7033 6937-0
info@mbe-components.com
www.mbe-components.com

nel®
Tel: + 1 203 949 8697
Email: info@nelhydrogen.com
www.nelhydrogen.com

Pegasus
Tel: +44 (0)1244 521869
Email: info@pegasuschemicals.com
www.pegasuschemicals.com

Sanan Integrated Circuit Co., Ltd.
Tel: +1 949 241 9501
Email: ray@sanan-ic.com
www.sanan-ic.com/en

Showcase your presence to the global Compound Semiconductor industry for 12 months

COMPOUND SEMICONDUCTOR MAGAZINE is the primary information resource for professionals working with compound semiconductor materials and devices.

Each issue of the magazine is distributed in print and digitally to over 60,000 professionals worldwide.

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

- 12 months of promotion in eight issues of the print and digital versions of Compound Semiconductor featuring your company logo and company details
- 12 months of promotion and exposure on the Compound Semiconductor website **www.compoundsemiconductor.net** featuring a Buyers Guide entry and hyperlinked company logo
- Sponsorship of one e-news alerts sent out to over 60,000 industry professionals with text and a standard web banner.

To find out more, contact: **Jessica Harrison**, Sales Executive
E: jessica.harrison@angelbc.com
T: +44 (0)2476 718970

200mm Process Has Arrived

Introducing the Temescal Systems UEFC-6100

Welcome to your production future. Our UEFC-6100 is the first compound semiconductor production tool designed for 200mm lift-off processing. The world's leading manufacturers are welcoming our Auratus™ enhanced line of Temescal electron beam deposition systems, and you won't want to be left behind. We vaporize metals more efficiently. We deliver uniformity. We return unmatched savings. The results are overwhelming.

With a line of tools optimized for a range of production environments, whether 150mm or 200mm, small-batch or high-volume, this invasion could be coming your way.

Find out more before you take us to your leader. Contact Ferrotec at 1-925-371-4170 or visit temescal.ferrotec.com/200mm.

UEFC-6100
25 Wafers (200mm)

UEFC-5700
42 Wafers (150mm)
18 Wafers (200mm)

UEFC-4900
25 Wafers (150mm)

FerroTec