




POWER ELECTRONICS

ISSUE IV 2026

 AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

POWERELECTRONICSMAGAZINE.NET

THE FUTURE OF ENERGY



WÜRTH ELEKTRONIK MORE THAN YOU EXPECT

YOU'VE GOT BETTER OPTIONS THAN THAT

Check out our Thermal Management Solutions



**WE meet @
PCIM Europe**
Hall 6 - 342


Thermal management is crucial for developing durable and efficient designs. Our gap filling, heat spreading, and hybrid solutions deliver optimal thermal management solutions, ensuring your design achieves maximum durability and efficiency. Paired with our expert services and custom solutions, we provide the perfect fit for your application. Ready to keep it cool?

www.we-online.com/thermalmanagement

#THERMAL

Highlights

- Gap filling, heat spreading and hybrid solutions
- Customized components for optimal fit
- Advisory services from technical experts



Wide-Bandgap Power Electronics comes of age

➤ The latest developments across the power electronics landscape underline a simple but powerful truth: wide-bandgap (WBG) technologies are no longer emerging - they are maturing into the backbone of next-generation energy and computing systems. This issue's news stories collectively highlight a sector moving with remarkable coherence, where advances in materials, measurement, manufacturing, and system integration are converging to unlock real-world impact.

At the foundation of this progress lies a deeper understanding of materials themselves. The new physics-based framework for detecting defects in SiC and GaN interfaces is more than an incremental improvement - it addresses a long-standing blind spot in device characterisation. By removing uncertainty from defect measurement, researchers can now "see" what had previously been obscured, particularly at the critical band-edge regions that dictate device performance. This signals a broader shift: as WBG devices push theoretical limits, precision metrology becomes just as important as material innovation. Without accurate insight into defects, efficiency gains remain fragile and reliability uncertain.

At the same time, manufacturability and scalability are accelerating rapidly. Bosch's third-generation SiC MOSFETs and Rohm's fifth-generation devices demonstrate how architectural refinements are translating into tangible benefits - lower on-resistance, improved thermal behaviour, and enhanced ruggedness. These are not isolated laboratory achievements; they are production-ready technologies engineered for cost efficiency and volume deployment, supported by moves such as 200 mm wafer transitions. In parallel, Renesas' expansion of GaN production capacity reinforces the industry's confidence that demand will continue to surge across automotive, industrial, and infrastructure segments.

Equally significant is the growing localisation and diversification of supply chains. Cyient's launch of India's first GaN power IC family, combined with its partnership with Navitas, highlights a strategic shift towards regional ecosystems for semiconductor production. Similarly, SweGaN's expanding order book across global markets points to strong commercial traction for advanced materials. These developments reflect a broader industry imperative:

resilience. As power electronics become central to electrification and digital infrastructure, supply assurance is no longer optional - it is strategic.

Application-driven innovation is another defining theme. Nowhere is this more evident than in electric vehicles, where Onsemi's deepening collaboration with Geely illustrates how semiconductor companies are moving upstream into system co-design. Higher-voltage architectures, such as emerging 900 V platforms, are redefining efficiency and performance expectations. Beyond mobility, WBG technologies are transforming renewable energy systems, data centres, robotics, and even space applications - where Infineon's radiation-hardened GaN devices have proven their reliability in the demanding environment of lunar missions.

Meanwhile, the horizon continues to expand with disruptive research. Breakthroughs in ultrawide-bandgap materials such as Ga₂O₃ and diamond hint at a future that extends beyond today's SiC and GaN dominance. From megawatt-class Ga₂O₃ modules to ultrafast diamond switches, these technologies push the boundaries of voltage, current, and switching speed. Though still in early stages, they offer a glimpse of what the next era of power electronics could deliver.

Finally, innovation at the system level remains critical. The development of piezoelectric-based DC-DC converters for data centres shows that even mature architectures are being reimaged to meet the escalating demands of AI-driven computing. Efficiency is no longer just a device-level metric - it is a system-wide imperative.

Taken together, these stories reveal an industry in transition from possibility to deployment. The WBG revolution is not just about better semiconductors; it is about enabling a more electrified, efficient, and sustainable world. As measurement improves, manufacturing scales, and applications diversify, power electronics is stepping firmly into its role as a cornerstone of modern technology.




Contents

Cover Feature

The future of energy

Intelligent power semiconductor computing and the rise of in-power transistor AI.



16 Engineered substrates enhance high-voltage GaN power electronics

Substrates with a ceramic AlN core and a silicon top surface provide a promising platform for the production of high-voltage, fast-switching GaN devices and ICs

22 Minding the ultra-wide bandgap – gearing up for Gallium Oxide

In an interview with PE, Guido Ueberreiter, Vice President Semiconductor Strategy, VON ARDENNE, explains why Gallium Oxide (Ga_2O_3) is attracting strong interest as an ultra-wide bandgap (UWBG) semiconductor with the potential to redefine power electronics beyond the limits of SiC and GaN.

26 Developing scalable traction inverters for compact electromobility solutions

The broad spectrum of electromobility, ranging from the conveyance of individuals to the transportation of commercial goods, requires efficient, reliable and affordable forms of drivetrains.



30 IDTechEX: Wide bandgap momentum meets silicon's staying power

As SiC and GaN gain ground in EVs and data centres, silicon's proven reliability continues to anchor wind power – but for how long? Rebecca Pool talks to Matthew Fall, IDTechEX, to find out more.

32 800VDC power delivery architecture and advantages for next-generation AI & GPU data centers

A comprehensive comparison with traditional AC-based power systems.

36 From conversion tax to DC-native: Re-engineering data centre power for the AI Era

The five-stage maturity model for data centres to transition from AC to DC-native infrastructure and the role played by Solid-State Transformers in defining the next generation of AI factories.

40 WE-MXGI: High-efficiency power inductor for demanding DC/DC converters

The WE MXGI series was developed for use in compact, high-power-density DC/DC applications in which high currents, low losses, and stable performance are required even at high switching frequencies.

42 How to select a shunt resistor for high bandwidth current measurements

Selecting the right shunt resistor is critical to achieving accurate, high bandwidth current measurements.

48 Driving efficiency through system-level optimization in next-generation power electronic systems

Future efficiency gains in power electronics require holistic system optimization across electrical, thermal, mechanical, and control domains, not isolated components.

51 Japan's power semi giants move toward landmark consolidation

Can Mitsubishi Electric, Rohm and Toshiba's potential merger reshape Japan's position in global power electronics?

54 Advancing SiC substrate and epitaxial technologies for next-generation power devices

Silicon carbide is rapidly reshaping next-generation power electronics, but the real performance gains depend on advances in substrate and epitaxial material quality that improve yield, efficiency and reliability at scale.

58 HfO₂ for better GaN transistors

Introducing a HfO₂ gate dielectric improves the vertical GaN transistor, boosting its drain current density and its breakdown voltage

NEWS

06 Cyient launches India's first GaN power IC family

07 A new way to measure WBG device defects

08 Renesas expands GaN production with Aixtron tools

09 SweGaN wins commercial orders worth €2.3 million

10 Bosch launches 3rd generation SiC MOSFETs

11 Novel DC-DC design could reduce energy consumption



Editor

Phil Alsop
phil.alsop@angelbc.com

News Editor

Christine Evans-Pughe
christine.evans-pughe@angelbc.com

Technology Editor

Rebecca Pool

Contributing Technical Editor

Richard Stevenson
richard.stevenson@angelbc.com
+44 (0)1923 690215

Sales & Marketing Manager

Shehzad Munshi
shehzad.munshi@angelbc.com
+44 (0)1923 690215

Design & Production Manager

Mitch Gaynor
mitch.gaynor@angelbc.com
+44 (0)1923 690214

Graphic Design & Multimedia Assistant

Harvey Watkins
harvey.watkins@angelbc.com

Sales and Product Manager

James Cheriton
james.cheriton@angelbc.com
+44 (0)2476 718970

Publisher

Jackie Cannon
jackie.cannon@angelbc.com
+44 (0)1923 690205

CEO

Sukhi Bhadal
sukhi.bhadal@angelbc.com
+44 (0)2476 718970

CTO

Scott Adams
scott.adams@angelbc.com
+44 (0)2476 718970

Power Electronics International Magazine is published six times a year on a controlled circulation basis. Non-qualifying individuals can subscribe at: £105.00/e158 pa (UK & Europe), £138.00 pa (air mail), \$198 pa (USA). Cover price £4.50. All information herein is believed to be correct at time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher.

Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd © Copyright 2026. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing. © Copyright 2026.

We strive for accuracy in all we publish; readers and contributors are encouraged to contact us if they recognise an error or omission. Once a magazine edition is published [online, in print or both], we do not update previously published articles to align old company names, branding, marketing efforts, taglines, mission statements or other promotional verbiage, images, or logos to newly created or updated names, images, typographic renderings, logos (or similar) when such references/images were accurately stated, rendered or displayed at the time of the original publication.

When companies change their names or the images/text used to represent the company, we invite organizations to provide Angel Business Communications with a news release detailing their new business objectives and/or other changes that could impact how customers/prospects might recognise the company, contact the organisation, or engage with them for future commercial enterprise.

Published by Angel Business Communications Ltd

6 Bow Court,
Fletchworth Gate,
Burnsall Road, Coventry CV5 6SP, UK.
T: +44 (0)2476 718 970
E: info@angelbc.com
W: angelbc.com



Cyient launches India's first GaN power IC family

650V DPAK GaN devices for edge AI computing and e-mobility are based on Navitas GaN technology

Indian company Cyient Semiconductors has announced the launch of seven new GaN power devices for the Indian market, developed using Navitas Semiconductor's GaN technology.

The DPAK devices combine drive, control, and protection functions, with integrated EMI management and current sensing.

This is Cyient's first commercial GaN product family and hailed as a major milestone in advancing India's domestic power semiconductor ecosystem. The new portfolio is designed to address the rapidly growing demand for high-efficiency, high-power-density solutions across AI data centres, telecommunications, consumer fast charging, industrial power systems, and e-mobility platforms.

Building on the collaboration with Navitas announced in December 2025, the partnership enables customers in India to access commercially available

GaN power solutions with enhanced local support, supply assurance, and alignment with emerging domestic sourcing initiatives.

Under the agreement, Cyient Semiconductors will license Navitas's GaN technology for use in India, accelerating the adoption of high-performance GaN solutions across a broad range of markets. In addition, Cyient Semiconductors will serve as a second source for select Navitas GaN devices already in mass production and strengthening supply chain resilience.


Over time, the licensing agreement with Navitas Semiconductor is expected to enable the domestic manufacturing of GaN power devices in India.

Suman Narayan, CEO, Cyient Semiconductors, said: "With this launch, Cyient Semiconductors is introducing its first family of GaN power ICs, marking our entry into high-performance power semiconductors with a strong roadmap

for expansion. Built on Navitas' proven GaN platform, these highly integrated devices are designed to accelerate adoption and support next-generation power applications. This represents the foundation of a broader GaN portfolio that will address the growing power and efficiency demands of AI infrastructure, industrial systems, consumer power, and e-mobility applications"

Chris Alexandre, president and CEO of Navitas Semiconductor, said, "India is a key market in Navitas's high growth, high power strategy with Navitas 2.0. This launch furthers our vision of a robust local supply chain and manufacturing in India for the governments "Make in India" initiatives. The partnership with Cyient delivers a strong local support infrastructure for our customers as we pivot our India strategy to focus on GaN based product enablement and customer success"

Cyient Semiconductors expect to begin sampling the first set of GaN power products by June 2026.



A new way to measure WBG device defects

US team develops physics-based framework that more accurately detects hidden defects in SiC and GaN

Researchers at Sandia National Laboratories and Auburn University in the US have developed a new method to more accurately detect atomic-scale defects in electronic materials, an advance that could be particularly useful for wide-bandgap (WBG) semiconductors such as SiC and GaN.

The study ‘The completed High-Low method for interface state analysis in MOS capacitors’, recently accepted for publication in the Journal of Applied Physics, addresses a longstanding challenge in understanding what happens at the critical boundary where a semiconductor meets an insulating layer.

At this interface, microscopic defects can trap electrical charge and quietly reduce device performance, even when the device otherwise appears to function normally. These defects can limit efficiency, increase electrical losses, and reduce the performance of advanced semiconductor devices.

Scientists commonly study these defects by comparing how a device responds to slow and fast electrical signals. However, the technique depends on knowing a key device property, the insulator capacitance, with very high accuracy. Even tiny errors can produce misleading results, sometimes making it appear that far more defects exist than are actually present.

The problem is like tuning a radio. If the frequency is slightly off, the sound becomes distorted. If it is far off, only noise is heard. Only at the correct setting does the signal come through clearly. In these measurements, researchers are effectively trying to “listen” to the defects, with the assumed capacitance acting as the tuning dial. However, unlike a radio listener who easily recognises a clear broadcast, researchers do not inherently know what the true defect signal should look like. If this tuning

parameter is not set exactly, the resulting data cannot be trusted.

Instead of relying on estimates, the researchers developed a physics-based framework that identifies the correct device conditions automatically by enforcing a fundamental electrostatic constraint. In simple terms, all the voltages inside the device must add up consistently according to the laws of physics. By enforcing this condition, the method removes guesswork and allows researchers to accurately measure defects in regions where the technique previously failed, particularly near the semiconductor band edge, where defects can have the strongest impact on device behaviour.

“This work resolves an otherwise fundamental limitation in one of the most widely used techniques for studying semiconductor interfaces,” said Brian D. Rummel, senior member of technical staff at Sandia National Laboratories and lead author of the study. “By introducing a physically consistent framework, we can now extract information that was previously obscured by measurement uncertainty.”

The advance is particularly important for SiC and GaN, widely used in high-efficiency and high-power electronics. These materials’ performance is often limited by defects at material interfaces. By providing a more accurate way to measure these defects, the new method gives researchers a clearer understanding of what is happening inside the device and how to improve it, according to the team.

“The new analytical framework will help researchers more accurately measure defects in transistor materials,” said Sarit Dhar, professor of physics at Auburn University and co-author of the study. “We are excited to apply this method to measure a variety of

technologically relevant and exploratory interfaces in our future work.”

“Defects at semiconductor interfaces play a major role in determining the performance and reliability of power electronic devices,” said Robert J. Kaplar, senior scientist and manager of the semiconductor materials and device physics group at Sandia National Laboratories. “More accurate characterisation tools help us better understand these materials and ultimately improve next-generation electronic technologies.”

At its core, the work improves a widely used semiconductor measurement technique by ensuring that it remains fully consistent with fundamental physical principles. The framework helps researchers separate meaningful signals from noise and obtain a clearer picture of the microscopic defects that influence device performance.

As demand grows for faster, more efficient, and more reliable electronics, advances such as this one could help guide the development of next-generation semiconductor technologies.

Instead of relying on estimates, the researchers developed a physics-based framework that identifies the correct device conditions automatically by enforcing a fundamental electrostatic constraint.

Renesas expands GaN production with Aixtron tools

GaN ramp-up aims to meet surging demand across power electronics applications

Aixtron has supplied Renesas with multiple Planetary G5+C MOCVD systems to expand its GaN production in response to surging demand across critical power electronics applications.

Since the acquisition of Transphorm in June 2024, Renesas has been accelerating the deployment of GaN technology in e-mobility and automotive, advanced Internet of Things (IoT) solutions, fast-charging infrastructure, AI data centre power

architectures, as well as renewable energy and industrial sectors to enable efficient and compact power conversion at scale.

“GaN is one of the fastest-growing segments in our industry and is becoming a key growth driver for our Power Business,” said Rohan Samsi, VP, GaN business division, power product group at Renesas. “We are excited to build on the proven Planetary platform, originally introduced at Transphorm, and to seamlessly

expand our production capacity with additional Aixtron GaN Planetary systems.”

“Renesas’ strategy to ramp up GaN production in high-volume manufacturing sends a clear signal of accelerated momentum in the wide-bandgap sector. As a long-standing technology partner, we support this expansion with our proven Planetary MOCVD production solutions,” said Nicolas Müsgens, director of product management GaN at Aixtron.

Onsemi and Geely expand EV collaboration

Onsemi and Geely Auto Group have expanded their collaboration to accelerate next-generation EV development.

The partnership includes deeper integration of Onsemi’s EliteSiC technologies across vehicles built on Geely’s SEA-S, the Super Hybrid variant of Geely’s Sustainable Experience Architecture.

These technologies enable higher-voltage 900V architectures that improve efficiency, extend driving range, and reduce charging times to

deliver a faster, more reliable, and more convenient driving experience for customers worldwide.

As part of the announcement, Geely is showcasing its SEA-S–based Super Electric Power (SEP) System that incorporates Onsemi’s EliteSiC power technologies in its electric drive systems.

By enabling higher power density, Onsemi says its technology allows more performance to be packed into a smaller, lighter system, contributing to improved vehicle dynamics and interior

space. At the same time, enhanced thermal performance helps maintain consistent output during demanding driving conditions while supporting long-term reliability. Together, these capabilities translate into a more responsive, efficient, and refined driving experience.

“The EV market is entering a new phase, driven by higher-voltage architectures and increasing system complexity. Our expanded engagement with Geely reflects how automakers and semiconductor partners are working more closely and earlier to shape critical vehicle design decisions. Together, we are enabling EVs that deliver higher efficiency and performance,” said Hassane El-Khoury, president and CEO, Onsemi

Gan JiaYue, CEO, Geely Auto Group commented: “As we advance our electrification strategy, closer coordination with technology partners is increasingly critical. Integrating Onsemi technologies across our SEA-S–based EV platforms reflects this deeper collaboration in action.”



SweGaN wins commercial orders worth €2.3m

GaN-on-SiC epi wafer orders span applications from power devices for data centres to RF devices for defence, aerospace, and telecoms

SweGaN, a European semiconductor manufacturer specialising in GaN-on-SiC epitaxial wafers, today announced that it has secured multiple new commercial framework agreements from customers across Europe, Asia, and the United States.

The orders received during the first four months of 2026 represent a total contract value of approximately €2.3m (SEK 25m) and will contribute to revenue over the next coming 6 to 18 months.

The new orders span a broad range of applications, including power devices for data centres and RF devices for defence, aerospace, and telecommunications. They reflect the growing interest and demand for SweGaN's differentiated material technology and innovation capabilities.

"We are pleased to see the strong growth and commercial traction across a diverse global customer and application base," said Jr-Tai Chen, CEO of SweGaN. "This indicates the recognition the market is showing to the value our material performance and tailored solutions can bring to their products and these design-in and design-win orders will propel our abilities to invest in capacity expansion and continued product innovation for the future", Chen added.

In conjunction with this update, SweGaN also announces the appointment of Ulf Nolemo as chief commercial officer (CCO), effective May 1st, 2026. In this role, Nolemo will be responsible for leading SweGaN's global commercial strategy, including sales & marketing activities, business

development, and strategic customer partnerships.

Nolemo brings 30 years of experience from the semiconductor industry with a strong track record in scaling commercial operations and driving innovation across international markets. Ulf has previously held leading positions at industry leaders such as Ericsson, Wolfspeed, Macom, NXP, and Infineon.

"The appointment of a chief commercial officer marks an important step as we enter our next phase of growth," said Jr-Tai Chen, CEO of SweGaN. "The industry experience and commercial leadership brought by our newly appointed CCO will be critical as global demand for our GaN epitaxial solutions continues to expand.



Bosch launches 3rd generation SiC MOSFETs

Taking SiC semiconductors to the next level of performance, robustness, and heat dissipation

Bosch has announced its 3rd generation SiC MOSFETs, based on the company’s dual-channel trench SiC architecture with lower on-resistance, better ruggedness, and smaller chip sizes for broader EV adoption.

A 20 percent reduction in specific on-resistance, around 10 percent higher short-circuit withstand capability, and a 40 percent thinner die enable clear efficiency gains and more cost-effective power module designs, according to the company.

At the same time, the reduced die size has a direct impact on production efficiency and system-level costs, further supported by the transition to 200 mm wafer manufacturing, which increases chip output per wafer.

While Bosch’s SiC chips are well-established in the demanding environment of high-performance electric vehicles, the company says these new features change the economic equation around SiC chip uses, enabling compact, efficient inverter and power module designs across a wider range of EVs from premium sports cars to mass-market models.

Gen 3 SiC MOSFETs offer measurable improvements in three key areas for automotive applications: electrical performance, thermal performance, and ruggedness.

In terms of thermal performance, reducing die thickness by 40 percent directly enhances heat dissipation and thermal conductivity. This supports higher power density and more efficient thermal management in high-performance semiconductor systems.

For ruggedness, the short-circuit withstand capability increases by about 10 percent, supported by the optimised two-zone JFET region and enhanced

gate oxide design. The design also improves robustness against parasitic turn-on at high switching speeds.

Additionally, the intrinsic body diode is optimised for soft recovery behaviour across the full automotive temperature range of -40°C to 200°C, reducing electrical stress during switching events and supporting stable operation under demanding conditions.


Three specific advancements in Gen 3 SiC-MOSFET architecture are shown in the diagram above including the trench (green), the new shield implant below the trench (blue), and the two-zone JFET region (yellow).

First, additional p-type shielding region introduced directly below the trench, provides full protection against off-state electric fields. This contributes to improved long-term gate oxide reliability under typical operating conditions. The existing hard mask of the trench etching process was reused for the new trench shield implant, eliminating the need for an additional lithographic process step. Second, two-zone structure in the

JFET region beneath the trench is such that each channel now corresponds to its own JFET region, providing significantly finer control over electric field distribution during both normal operation and fault events. This improves the critical trade-off between on-resistance (RonA) and short-circuit withstand time (SCWT), expanding the available design space for robust, high-performance SiC power semiconductors.

Third, the die thickness in Gen 3 is reduced by 40 percent to only 100 µm. This reduces material usage and places the active semiconductor layer closer to the heat sink, accelerating heat dissipation.

Anne Bedacht, head of product management for power semiconductors at Bosch commented: “The Gen 3 upgrades are noteworthy for two reasons: their combined effect and how our engineers integrated them into well-established production processes. Our design-for-manufacturability approach ensures that Gen 3 is production-ready and cost-effective from the outset.”



Novel DC-DC design could reduce energy consumption

US team says piezoelectric tech could be more efficient way to convert high voltages to lower levels in data centres


In an effort to meet the rising energy demands of data centres, engineers at the University of California San Diego have developed a new device that could improve how graphics processing units (GPUs) convert and manage power.

They say the technology demonstrates a more efficient way to converting high voltages into lower levels required by computing hardware. In lab tests, a prototype chip performed the type of voltage conversion used in modern data centres with high efficiency.

The advance ‘A Hybrid Piezoelectric Resonator-based DC-DC Converter’, published in *Nature Communications*, could lead to the development of smaller, more energy-efficient systems for advanced computing.

Data centres often distribute power at 48 volts, while processors in GPUs need much lower voltages, typically between 1 and 5 volts. However, converting between these levels efficiently, and within limited space, has become increasingly difficult as computing demands grow.

Traditional step-down converters, for instance, lose efficiency and struggle to deliver enough current when the gap between input and output voltage is large. Most step-down converters rely on magnetic components such as inductors, which, while effective, are approaching their physical performance limits and are growing difficult to scale further. “We’ve gotten so good at designing inductive converters that there’s not really much room left to improve them to meet future needs,” said study senior author Patrick Mercier, professor in the Department of Electrical and Computer Engineering at the UC San Diego Jacobs School of Engineering.



To address this challenge, Mercier and members of his research group, including study first author Jae-Young Ko, a PhD student at UC San Diego, explored a promising alternative: piezoelectric resonators, which store and transfer energy through mechanical vibrations. Piezoelectric-based converters could potentially be smaller, more energy dense, more efficient and easier to manufacture at scale. “They have a lot of room to grow and have the potential to deliver better performance than anything that’s come before them,” Mercier said.

However, early versions of piezoelectric-based converters have struggled to maintain efficiency and deliver enough power when handling large voltage differences.

In this study, the team developed an improved step-down converter that combines a piezoelectric resonator with small, commercially available capacitors arranged in a strategic way. This new circuit design allows the converter to handle larger voltage conversions more effectively. The team implemented the design in a prototype chip. In tests, it converted 48 V down to 4.8 V — a level commonly required in data centres — with a peak efficiency of 96.2 percent. The chip also delivered about four

times more output current than earlier piezoelectric-based designs.

There are several advantages with this hybrid circuit design: it creates multiple pathways for power to flow; reduces wasted energy; and eases the workload on the resonator. As a result, it boosts both efficiency and power delivery with only a small increase in size.

Although the technology is still in its early stages, the researchers say it represents an important step toward overcoming the limitations of today’s power converters. Future work will focus on improving materials, circuit design and packaging. Because piezoelectric resonators physically vibrate, they cannot be soldered onto circuit boards using conventional approaches, and will require different strategies to integrate them into electronic systems, Mercier explained.

“Piezoelectric-based converters aren’t quite ready to replace existing power converter technologies yet,” Mercier added. “But they offer a trajectory for improvement. We need to continue to improve on multiple areas — materials, circuits and packaging — to make this technology ready for data centre applications.”

The future of energy

Intelligent power semiconductor computing and the rise of in-power transistor AI.

BY DANILO PAU, IEEE FELLOW, INTELLIGENT POWER SEMICONDUCTOR COMPUTING,
STMICROELECTRONICS

The silent revolution in power electronics

IN THE MODERN industrial landscape, power electronics serves as the “heartbeat” of technological progress. From the traction inverters driving electric vehicles (EVs) to the massive grid injection systems of AI server farms, power transistors are the foundational backbone. However, as we push these systems toward higher efficiency and greater power density, they encounter significant physical and reliability challenges. The traditional approach—relying on external microcontrollers and rigid physics-based models—is reaching its limits. We are now entering the era of Intelligent Power Semiconductor Computing, a paradigm shift where the transistor itself becomes capable of sensing, learning, and predicting its own operational future.

The critical role of power transistors across sectors

Power transistors are no longer just simple switches; they are sophisticated components integrated into diverse, high-stakes environments. According to recent research from STMicroelectronics, their roles can be categorized into several critical sectors:

- **EV Powertrain:** In this sector, power transistors act as the traction inverters for electric motors and manage DC-DC conversions for battery chargers and Battery Management Systems (BMS).
- **Energy Converters:** They provide the backbone for Uninterruptible Power Supplies (UPS) and network protection, ensuring stable power delivery to server supplies.
- **Aerospace:** In aviation, these semiconductors control propellers, actuators, and manage isolated power distribution systems.
- **AI Farms:** As artificial intelligence demands

massive energy, power transistors manage grid injection and panel optimization to keep data centers running efficiently.

Despite their ubiquity, these components face mounting challenges that threaten system reliability.

The reliability wall: challenges in modern semiconductors

As industry transitions toward Wide Bandgap (WBG) materials like Silicon Carbide (SiC) and Gallium Nitride (GaN), the performance ceilings are rising, but so are the complexities. These materials allow for higher switching frequencies and better breakdown voltages, yet they introduce new failure modes.

The most significant factor leading to the failure of power semiconductors is heat. Thermal management is a constant struggle, with “self-heating” and “thermal runaway” posing catastrophic risks to the entire system. Furthermore, scaling these transistors for better integration introduces secondary physical effects that are difficult to model using traditional methods.

The intelligence dilemma: physics-based vs. traditional AI


To prevent catastrophic failures, industry has traditionally used two modeling approaches:

- **Physics-Based Models:** These offer high interpretability but are computationally demanding and rigid. They often fail to match the time-varying occurrences of real-world operations.
- **Legacy AI Models (RNNs/LSTMs):** While highly accurate and adaptable, Recurrent Neural Networks (RNNs) are too “heavy” for deep edge deployment. They require too many data to feed the learning process, very long learning times.

Current Convolutional Neural Processing Units (NPU) are not capable to accelerate them. Furthermore these workloads are not quantization friendly. All in one these methods fails to meet tight latency requirements of power electronics, which often operate in the microsecond (μ s) or hundreds of nanosecond (ns) domain.

This creates a gap: we need the intelligence of AI but to fit within the extreme constraints of a power transistor package.

➤ A high-detail view of a power electronic circuit board with prominent copper coils and semiconductor components, representing the heart of modern power systems.



Tiny perceptual on line learning AI: intelligence at the nano-edge

The solution lies in Tiny Perceptual On Line Learning AI. The goal is to move beyond external monitoring and enable the power transistor to “predict its own operative conditions” in real-time. And including the capability to learn and deploy at scale, transistor per transistor as their single story over operative time can be its own experience. This requires a radical rethinking of AI architecture to fit within extreme hardware constraints:

- **Power Consumption:** μW (microwatt) envelope.
- **Clock Speed:** 10 to 40 MHz.
- **Memory:** Total SRAM between 40 and 88 KiB.

Reimagining the neuron: the RBF-NN breakthrough


To meet these constraints, researchers have moved away from fixed, deep neural networks toward Radial Basis Function Neural Networks (RBF-NN). This architecture offers three transformative advantages:

- **Dynamic Topology**
Unlike standard AI, the RBF-NN utilizes a dynamic hidden layer. It can auto-allocate new neurons when the prediction error exceeds a certain threshold and de-allocate inactive ones to minimize memory footprint. A typical configuration might include 10 fixed input neurons, a hidden layer that scales up to 50 neurons, and a single output neuron.
- **Gradient-Free Learning**
Perhaps the most significant breakthrough is the removal of the Backpropagation (BP) algorithm. Traditional AI learning is “gradient-heavy,” requiring massive RAM to store activation derivatives. The RBF-NN uses Gaussian activations characterized by specific centers and radii, allowing for zero-gradient computation. This drastically accelerates learning and permanently removes the memory overhead associated with traditional training.
- **Physics-Informed (PI) Integration**
To ensure the AI doesn’t produce “hallucinated” or physically impossible predictions, a Physics-Informed (PI) variant was developed. The PI-RBF-NN integrates physical degradation constraints directly into the loss function. This enforces monotonic decreases and boundary consistency, guaranteeing that predictions for Remaining Useful Life (RUL) remain stable and grounded in physical reality.

Validating the vision: The data behind the hype

The effectiveness of this approach has been validated using the NASA Ames Research Center Prognostics Center of Excellence dataset. Researchers tested discrete IGBTs (Insulated-Gate Bipolar Transistors) under thermal overstress. The results were staggering:

- **Accuracy:** The RBF-NN achieved an average R^2 of 0.9867, outperforming traditional RNNs (0.9836).



➤ A diagram showing the multiple ST power transistors, mounted on a board.

- **Efficiency Delta:** While a standard RNN required 7,381 parameters and over 68 MB of training RAM on an Intel i7, the RBF-NN required only 561 parameters (7.6% of the weight) and fit entirely within < 21 KiB of RAM.

From concept to silicon: the STRED DSP

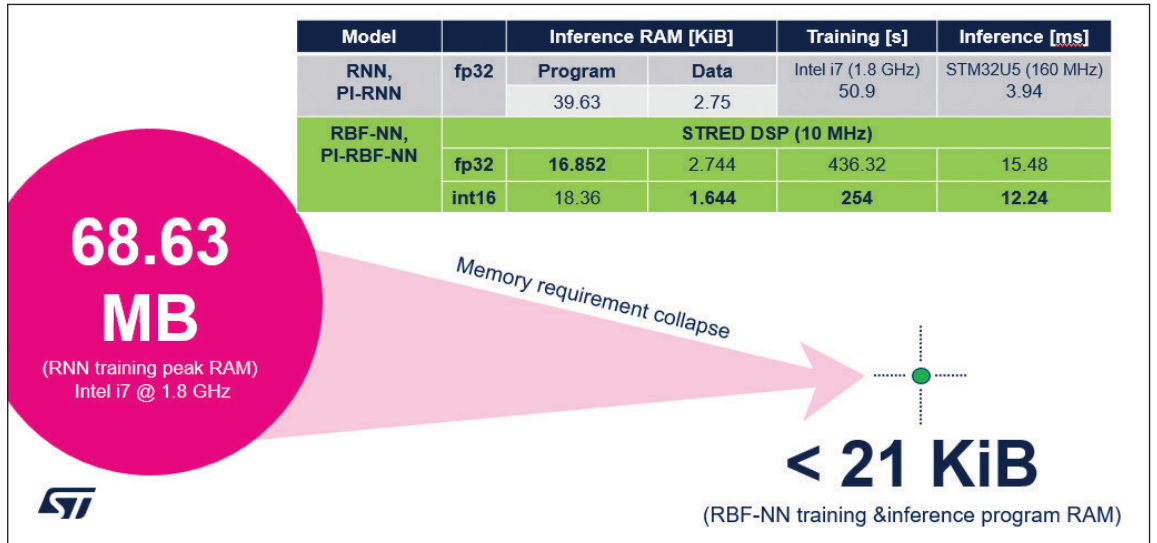
The hardware realization of this vision is the STRED DSP. This STMicroelectronics home-made micro-processor is so tiny in power and silicon area that it can be physically bonded inside a standard discrete power transistor package. This is the ultimate vision of System-in-package (SiP): a transistor that no longer relies on an external processor but autonomously manages its own lifecycle.

The roadmap: toward agentic AI power control


The integration of AI into power electronics is following a clear evolutionary path:

- **In-Transistor AI Computing:** Real-time lifecycle prediction at the component level.
- **In-Power Module AI Computing:** Coordinated intelligence across multiple components, handling aging of discrete and their varying tolerances.
- **Agentic AI Power Control:** Autonomous systems capable of microsecond to nano-latency critical

> A comparison table showing the drastic reduction in parameters and memory usage between traditional RNNs and the new RBF-NN architecture.



> A cross-section of a "System-in-package" showing a DSP chip physically bonded within a transistor housing, representing autonomous intelligence.



decision-making to optimize power distribution (especially in AI farms), performance and prevent failure before it happens.

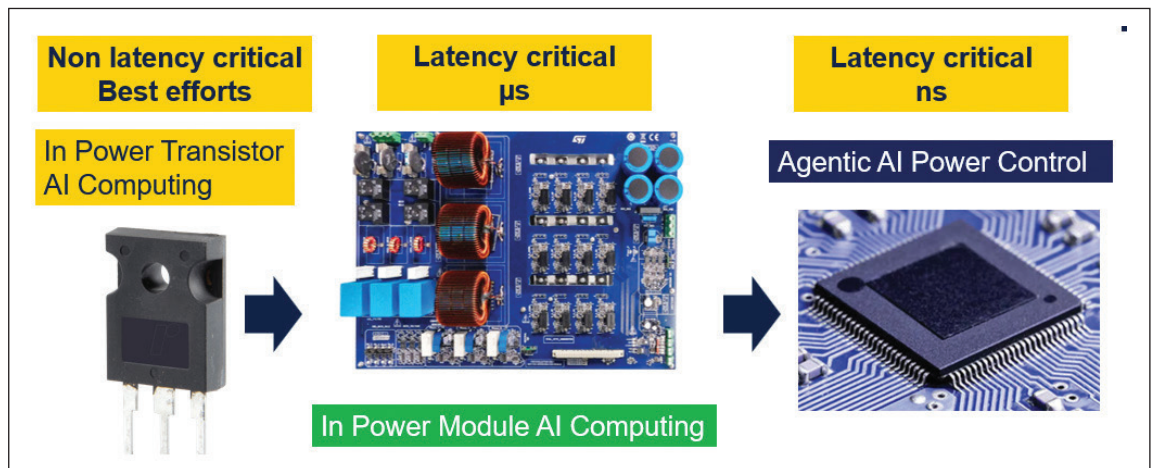
Conclusion: the autonomous future

The marriage of power electronics and Tiny Perceptual AI is more than just a technical curiosity; it is a necessity for a sustainable, electrified future. By enabling power transistors to "learn" and "predict" without the burden of heavy computation or external dependencies, we are creating a more resilient industrial infrastructure. From the cars we drive to the grids that power our AI, the next generation of energy management will be silent, autonomous, and incredibly computationally clever.

REFERENCES

> S. Tognocchi, M. Marcon and D. P. Pau, "Integrated IGBT with Tiny Learning of Radial Basis Networks for Remaining Useful Lifetime Estimation," 2026 IEEE Applied Power Electronics Conference and Exposition (APEC), San Antonio, TX, USA, 2026, pp. 3635-3641, doi: 10.1109/APEC51134.2026.11517072. <https://ieeexplore.ieee.org/vdocument/11517072>

> A roadmap diagram showing the progression from simple transistors to agentic AI power control, marked by decreasing latency requirements.



HITACHI




Next Gen LV LinPak

Meet our PCIM 2026 highlight:

Next Gen LV LinPak. Generation 2 LV LinPak with a state-of-the-art 3.3 kV trench chipset, even lower stray inductance, ideal gate-unit connection, current-sensing features and easy paralleling. It enables compact, high-efficiency, robust converter designs for renewable energy systems and grid-connected power converters. Ready for silicon IGBT and SiC. Proven standard for wind, PV, BESS and H₂ electrolysis.

Visit us 9-11 June in Nuremberg: Hall 9, Booth 302.



hitachienergy.com/semiconductors

Engineered substrates enhance high-voltage GaN power electronics

Substrates with a ceramic AlN core and a silicon top surface provide a promising platform for the production of high-voltage, fast-switching GaN devices and ICs

BY KAREN GEENS AND BENOIT BAKEROOT FROM IMEC

THINK of an application for SiC power electronics and you'll probably think of electric vehicles (EVs). After all, it's the battery-powered automobile that's driving the growth in sales of SiC MOSFETs, a trend that's forecast to continue throughout this decade and beyond. Employed in the traction inverter, these transistors tend to operate in systems operating at 400 V – but higher voltages, particularly 800 V, will be introduced over the coming years, a shift that moves towards the sweet spot of SiC. Commercial products associated with this class of transistor span several hundred volts to a few kilovolts, with the well-established 1.2 kV MOSFET offering ideal headroom for EVs operating with 800 V power systems.

However, SiC is not the only wide-bandgap semiconductor with the attributes needed to make a compelling case for deployment in EVs. There's also GaN, which many are overlooking today for this application. This attractive alternative is already enjoying tremendous success at lower voltages, and is now the key component in fast chargers. Increasing the blocking voltage to 1 kV or beyond remains a challenge for fabrication of devices on large-area substrates. A possible route to success is turning to engineered substrates, an approach we are pioneering at the European microelectronics research centre, imec. Alternatives such as bulk GaN exist, though they come with smaller wafer diameters and significantly higher cost.

The GaN HEMT has a number of strengths, including a high critical-electric-field and a high electron-mobility, the latter realised in combination with very high density for the two-dimensional electron (2DEG) that's created at hetero-interfaces with other III-nitrides, such as AlGaN. In recent years, there has been significant improvement in the leading forms of this device, which include the enhancement-mode



(E-mode) lateral p -GaN gate HEMT and the cascode depletion-mode (D-mode) HEMT.

To realise high power levels in D-mode devices in a cascode configuration, engineers tend to direct their efforts at lowering the 2DEG channel resistance and increasing the electron channel mobility, steps that result in a very low on-resistance.

In comparison, the E-mode p -GaN gate HEMT has a higher 2DEG channel sheet resistance and thus a higher on-state resistance. However, this weakness must be weighed against a device architecture that allows a higher level of monolithic integration, thanks to the E-mode transistor ensuring a safe stand-alone operation and a simpler gate driving circuitry in many applications.

In addition to these advantages, the p -GaN gate HEMT stands out for its maturity and performance. That's why we have selected this particular device architecture for the first demonstrator for a 1200 V GaN switch, while exploring the potential of vertical GaN structures for future high-voltage applications.

Engineered foundations

To produce a switch that operates at such a high voltage demands a thicker epitaxial layer, which presents challenges in manufacturing. This is where QST, often referred to as just QST, comes into play.

QST's commercially available, engineered substrates feature a ceramic AlN core that combines a high mechanical strength with close matching of the coefficient of thermal expansion of GaN and AlGaIn layers. In addition to these features, the seed layer for the epitaxial growth of the GaN and Al(Ga)In layers remains well-known silicon, with a (111) crystal orientation.

This engineered platform provides a great foundation for the growth of thick, crack-free GaN stacks, significantly reducing the risk of wafer breakage, a common issue when using silicon substrates for growing thicker stacks targeting 1200 V, especially at larger diameters, such as 200 mm.

Using these proprietary substrates, we have fabricated lateral GaN power HEMTs with 9 μm -thick AlGaIn buffer stacks that deliver excellent electrical performance. These devices have a hard breakdown voltage of over 1800 V, and a specific on-resistance of around just 5.8 $\text{m}\Omega\text{ cm}^2$.

Another asset of QST is its scalability. Applicable to 300 mm wafers, this technology is a major step forward for cost-effective, high-volume production. Migrating to larger wafers allows more devices per batch, helping to drive down manufacturing costs and accelerate the adoption of GaN in high-voltage markets.

Vertical or lateral?

One key question when pursuing higher voltages, such as 1200 V or more, is this: which architecture, vertical or lateral, will offer a superior performance? With lateral GaN devices, all connections – that's

the gate, the source, and the drain – are made to the front side. With this geometry, as voltage targets rise, the gate-to-drain distance has to increase to maintain safe operation. In addition, it's critical to lengthen the field plates – they are employed to distribute the high-voltage in the off-state, with source-connected field plates used to optimise the electric-field distribution (see Figure 1 (a)). Due to these requirements, the introduction of higher blocking voltages requires a larger footprint, and leads to a higher specific on-state resistance and larger capacitance, limiting efficiency and scalability.

To address these issues, a number of researchers, including ourselves, are investigating vertical GaN architectures. Unlike lateral devices, vertical structures distribute the high off-state voltage throughout the thickness of the epitaxial stack, rather than across its surface. Thanks to this, the device footprint remains compact, regardless of voltage rating – a noteworthy asset for high-voltage devices.

A key challenge associated with the development of vertical GaN devices is realising low defect densities in the epilayers. In GaN-based epitaxial stacks grown on foreign substrates – including silicon, employed in all commercially available GaN devices – the most common defects are one-dimensional (1D) dislocation lines perpendicular to the surface. In lateral devices, it appears that these dislocation lines do not have a significant impact on device operation and reliability. However, that's not the case for vertical devices, where current flow is inherently from bottom to top, and 1D defects interfere with vertical current flow. Due to this limitation, it's preferable to produce vertical devices using homoepitaxial growth on GaN substrates. But GaN substrates are expensive and limited in size, restricting commercial scalability. One possible avenue for further research involves investigations to understand the full extent of these 1D defects on the operation of vertical GaN devices, as well as determining what 1D defect level may be acceptable regarding device performance and reliability.

Using QST substrates for vertical device development, we have grown thick, high-quality n -type GaN drift layers on 200 mm wafers, paving the way for manufacturable vertical GaN devices.



➤ Figure 1: Cross-section STEM of (a) the p -GaN gate architecture used in lateral devices and (b) the trench gate MOS architecture used in vertical devices.

We have produced epitaxial stacks with a thickness exceeding 11 μm on this engineered foundation, enabling robust device performance.



We have processed these epiwafers into semi-vertical trench gate MOSFETs. When these devices are in the on-state there's an electron inversion channel in a p -type GaN layer, similar to traditional MOSFETs, albeit at the sidewall of the trench MOS gate. During device operation, electrons flow from the source n^+ -GaN layer through the electron inversion channel towards the deeper drift layer, before they are transported into a buried n^+ layer at the bottom of the structure. We use a drain metal contact from the frontside towards the buried n^+ -GaN layer to collect these electrons.

Vertical GaN devices are emerging as a promising direction for future high-voltage power electronics, attracting significant research interest worldwide. Although still in an exploratory stage, this technology holds potential to overcome key limitations of conventional lateral architectures.

GaN ICs on chip

Today, most GaN-based power systems are built from multiple discrete chips, assembled on a printed circuit board. While this modular approach is effective, it introduces parasitic inductances – they are unwanted electrical effects that can degrade performance, especially in fast-switching applications.

A promising route to truly harnessing the speed and efficiency of GaN is monolithic integration – that's the bringing together of multiple power devices and



➤ Figure 3: Image of fabricated GaN-ICs, along with test structures, on 200 mm QST substrate.

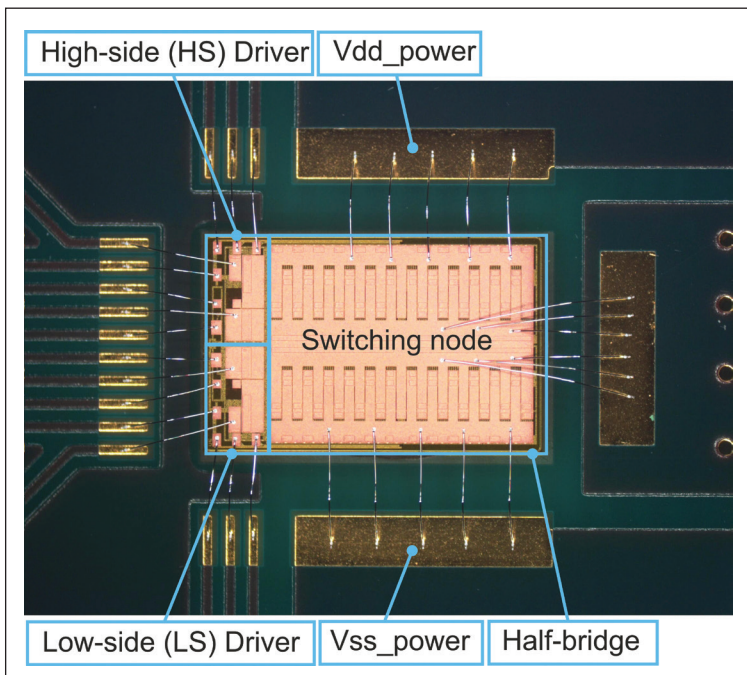
control circuits on a single chip. GaN power ICs are already finding their way into power supplies, motor drives, and converters, offering compact, high-performance solutions.

One of the key advantages that comes from forming a GaN IC through the on-chip integration of multiple GaN power HEMTs and accompanying control circuitry is a dramatic reduction in parasitic inductance. This enables faster, more-efficient switching. At the heart of these ICs are lateral p -GaN HEMT power arrays, serving as core switching elements.

To further enhance performance and reliability, we are combining a GaN-on-SOI (silicon-on-insulator) technology with trench isolation to ensure galvanic isolation between switches on the same chip. This design helps prevent electrical interference and enables stable operation at high voltages.

One of our critical innovations is the use of a local deep contact to the top silicon layer in the SOI substrate, which is connected (shorted on chip) to the source. This prevents back-gating effects – a condition that threatens to disrupt switch performance, especially in integrated half-bridge configurations, where the high-side switch might otherwise fail.

It's worth noting that QST substrates, which share similarities with SOI in terms of structure and



➤ Figure 2: Example of 650 V GaN-on-QST monolithic IC, depicting a symmetrical 140 mm half-bridge with integrated driver and level shifters.

insulation, also support this technology platform. We have successfully demonstrated monolithically integrated GaN ICs on both SOI and QST substrates, with voltage ratings up to 650 V.

These integrated circuits are not limited to just power switches, and can also include gate drivers, control and protection logic, and even sensing functions, such as temperature monitoring – all on the same chip. This level of integration paves the way for smarter, more compact, and more efficient power systems, ideal for everything from consumer electronics to industrial automation.

◉ *The work on the fabrication of GaN ICs both on SOI and QST received funding from the ECSEL Joint Undertaking (JU) HiEFFICIENT project under grant agreement no. 101007281. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Austria, Spain, Belgium, Germany, Slovakia, Italy, Netherlands, and Slovenia. This work also received funding from the ALL2GaN Project for the development of lateral HEMT devices on QST for 1200 V (Grant Agreement No 101111890), and has been supported by the Chips Joint Undertaking and its members including the top-up funding by Austria, Belgium, Czech Republic, Denmark, Germany, Greece, Netherlands, Norway,*

FURTHER READING

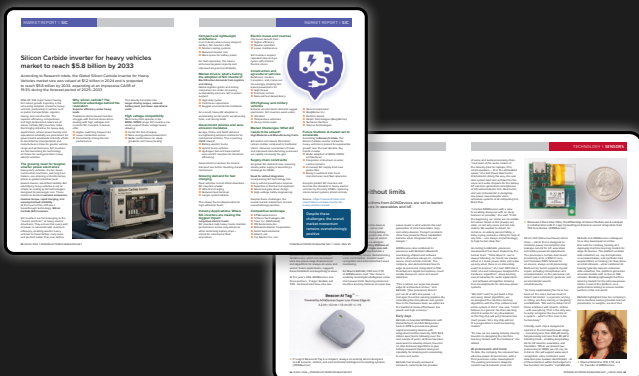
- A. Vohra *et al.* "Epitaxial buffer structures grown on 200 mm engineering substrates for 1200 V E-mode HEMT application", *Appl. Phys. Lett.* **120** 261902 (2022)
- S. Kumar *et al.* "1.2 kV Enhancement-Mode p-GaN Gate HEMTs on 200 mm Engineered Substrates," in *IEEE Electron Dev. Lett.* **45** 657 (2024)
- G. Baratella *et al.* "1200-V Dual-Gate p-GaN Bidirectional Switches on 200-mm Engineered Substrates," *IEEE Trans. Electron Devices* **72** 6574 (2025)
- K. Geens *et al.* "Route Toward Commercially Manufacturable Vertical GaN Devices," *IEEE Trans. Electron Devices* **71** 1488 (2024)
- U. Chatterjee *et al.* "A fully integrated half-bridge driver circuit in All-GaN GAN-IC technology", *Solid-State Electronics* **207** 108707 (2023)

Slovakia, Spain, Sweden and Switzerland. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the national granting authorities. Neither the European Union nor the national granting authorities can be held responsible for them.

PEI POWER ELECTRONICS INTERNATIONAL

BOOK YOUR REPRINT TODAY!

A reprint from Power Electronics International amplifies your editorial exposure and provides authoritative third-party validation. It's a powerful tool for sales and marketing, helping you showcase innovation to engineers, designers, and decision-makers with independent credibility. Use it across events, pitches, newsletters, and social channels to reinforce trust, extend visibility, and strengthen your position in the power electronics market.



Contact: Shehzad Munshi
shehzad.munshi@angelbc.com

20
27

Sheraton Brussels
Airport Hotel

Brussels, Belgium

12-14 April 2027



ANGELTECH

**INTERNATIONAL
CONFERENCE**

REGISTER NOW

Power Electronics International Conference 2027

Building on four hugely successful events, PE International returns for its 5th year. Chaired by Aleksei Cherkasov from Infineon, the conference brings together executives, end users, and technology experts for two days of insight, discussion, and collaboration on the technologies driving the global power electronics industry.

The 2027 programme builds on the core themes from 2026; covering automotive, renewables, data centres, and wide bandgap innovation, while placing a stronger emphasis on roadmaps and bridging the gap between research and manufacturing.

peinternational.net



SCAN ME

80+ EXHIBITORS

800+ ATTENDEES

4 CO-LOCATED CONFERENCES

120+ PRESENTATIONS





20
27

The European Event for The Semiconductor Supply Chain

REGISTER NOW

Conference Themes

- Wide Bandgap Dominance: From Adoption to Optimisation
- Electrification at Scale: EV, Charging and High Voltage Architectures
- AI, Data Centres and the Power Density Challenge
- Powering the Energy Transition: Grid, Storage and Energy Systems Innovation
- Integration, Packaging and Manufacturing Transformation

Why Attend?


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

THE WHOLE SEMICONDUCTOR SUPPLY CHAIN UNDER ONE ROOF



PREVIOUS SPEAKERS & SPONSORS





Minding the ultra-wide bandgap – gearing up for Gallium Oxide

In an interview with PE Magazine, Guido Ueberreiter, Vice President Semiconductor Strategy, VON ARDENNE, explains why Gallium Oxide (Ga_2O_3) is attracting strong interest as an ultra-wide bandgap (UWBG) semiconductor with the potential to redefine power electronics beyond the limits of SiC and GaN.

As AI infrastructure, electrified transport and renewable energy systems drive demand for higher-voltage, higher-efficiency power conversion, gallium oxide is emerging as a serious contender in the next generation of ultra-wide bandgap semiconductors. Once viewed as a long-term research material, gallium oxide is now moving rapidly toward industrial deployment, promising new levels of performance, scalability and cost efficiency in applications beyond the reach of today's silicon carbide and gallium nitride technologies.

For much of the last decade, the power electronics industry has been dominated by discussions around silicon carbide and gallium nitride. Both materials have transformed expectations surrounding efficiency, switching speeds and power density, enabling major advances across electric vehicles, renewable energy systems, industrial automation and data centre infrastructure. Yet even as silicon carbide and gallium nitride continue

their expansion into mainstream markets, attention is already shifting toward what many now see as the next major breakthrough material for high-voltage power electronics: gallium oxide.

Until recently, gallium oxide remained largely confined to research laboratories and academic discussions. Although engineers and scientists recognised its impressive electrical properties many years ago, there was widespread scepticism about whether the material could ever be manufactured economically at scale. Today, however, that perception is changing rapidly. A combination of mounting market demand, improvements in manufacturing pathways and the unprecedented rise of AI-driven power consumption is accelerating interest in gallium oxide across the semiconductor industry.

What makes the timing particularly significant is that the wider power electronics landscape itself is evolving.

Electrification trends have already pushed operating voltages higher across multiple industries. Electric vehicles have moved from 400V to 800V architectures in pursuit of faster charging and improved efficiency, while renewable energy systems increasingly require more advanced power conversion technologies capable of handling greater loads with minimal losses. At the same time, hyperscale AI infrastructure is creating entirely new power delivery challenges inside modern data centres, where operators are now discussing the possibility of bringing hundreds of volts — and eventually even kilovolt-level power distribution — directly onto server motherboards.

These trends are forcing the industry to look beyond the capabilities of existing semiconductor materials. While silicon carbide and gallium nitride remain highly effective in many applications, the move toward ultra-high-voltage power conversion is creating demand for materials capable of operating more

efficiently at significantly higher electric fields. This is where gallium oxide begins to stand apart.

Gallium oxide belongs to the category of ultra-wide bandgap semiconductors. Its bandgap is approximately 4.9 electron volts, substantially wider than the roughly 3.2 to 3.3 electron volt bandgaps associated with silicon carbide and gallium nitride. Although this may appear to be a modest numerical difference, the implications for device performance are considerable. A wider bandgap enables semiconductor devices to tolerate much higher electric fields before breakdown occurs, allowing engineers to design smaller, faster and more efficient devices capable of operating at far higher voltages.

The practical significance of this becomes clearer when examining the changing nature of modern power architectures. Historically, many power conversion systems operated at relatively modest voltages, particularly in battery-driven and distributed electronic systems. However, as industries seek greater efficiency and lower transmission losses, higher voltages are being pushed progressively closer to end-use applications. Data centres are perhaps the clearest example of this transformation. Just a few years ago, server-level power architectures centred largely around 48V distribution. Today, some of the world's largest AI infrastructure providers are discussing 800V distribution schemes inside computing systems, with long-term ambitions extending even further.

Higher voltages reduce current requirements and improve overall efficiency, but they also place enormous demands on power semiconductors. Devices must not only tolerate these voltages safely, but also maintain compact form factors and minimise thermal losses. Gallium oxide's electrical properties position it as a potentially transformative solution for these emerging requirements.

Importantly, gallium oxide is not simply viewed as a replacement for existing technologies. Instead, many experts believe it could open entirely new application domains for semiconductor-based power conversion. One area

receiving particular attention is grid infrastructure. Traditional electrical grids still rely heavily on large oil-filled transformers and mechanical power conversion systems that are expensive, maintenance-intensive and physically cumbersome. As the energy transition accelerates, there is growing interest in replacing portions of this legacy infrastructure with solid-state transformers and advanced semiconductor-based conversion systems capable of operating directly at much higher voltages.

This is precisely the type of environment where gallium oxide could offer major advantages. Its ability to operate efficiently in kilovolt-level applications could enable entirely new approaches to power conversion across grid infrastructure, renewable energy systems and industrial power distribution. In electric vehicles, meanwhile, the ongoing shift toward higher-voltage architectures may eventually create further opportunities for gallium oxide devices, particularly as manufacturers continue pursuing improvements in charging speed, efficiency and power density.

Yet one of the most significant reasons for growing excitement around gallium oxide lies not only in its electrical performance, but also in its potential manufacturability. Historically, many promising semiconductor materials have failed to achieve widespread commercial success because they proved too difficult or expensive to produce at scale. Silicon carbide itself continues to face substantial manufacturing challenges, particularly in the production of large-diameter wafers with sufficiently low defect densities for high-volume device fabrication.

Gallium oxide, by contrast, offers a potentially more scalable production pathway. One of its key advantages is that it can be grown using melt-based crystal growth methods that resemble traditional silicon manufacturing techniques more closely than silicon carbide production processes do. This compatibility with established semiconductor infrastructure


could prove decisive in accelerating commercial adoption.

The semiconductor industry has spent decades building highly optimised manufacturing ecosystems around large-format wafers, particularly 200mm and 300mm platforms. Fabrication equipment, lithography systems, deposition tools and process technologies have all evolved around these standards. Gallium oxide's compatibility with such infrastructure means the industry may be able to leverage much of its existing manufacturing base rather than requiring entirely new production ecosystems.

This could dramatically reduce both cost and complexity. Larger wafers enable more devices to be manufactured per substrate, improving equipment utilisation and lowering the cost per chip. In silicon carbide manufacturing, achieving defect-free 300mm wafers remains a major technical challenge. Gallium oxide, however, appears to offer a more straightforward route toward large-diameter substrates, potentially giving it a significant long-term economic advantage.

The existence of mature manufacturing equipment is also highly important. Much of the infrastructure required to process gallium oxide devices already exists within semiconductor fabrication facilities. Physical vapour deposition, chemical vapour deposition, lithography and related technologies are already widely deployed across advanced semiconductor manufacturing. If

gallium oxide substrates can be produced reliably at scale, many elements of the industrial ecosystem are effectively ready to support them.



Physical vapour deposition technologies, in particular, are expected to play a major role in unlocking the full potential of gallium oxide devices. Although the material offers exceptional electrical properties, it also presents one important challenge: relatively poor thermal conductivity. Compared with silicon carbide, gallium oxide conducts heat far less efficiently, creating potential concerns around heat dissipation in high-power applications.

Thermal management has become one of the defining engineering challenges across the semiconductor industry, extending far beyond power electronics alone. AI processors, GPUs and high-bandwidth memory systems are all generating unprecedented thermal loads, forcing engineers to develop increasingly sophisticated cooling and packaging strategies. This broader industry effort may ultimately benefit gallium oxide significantly.

Experts point out that the semiconductor sector already possesses many of the tools required to address thermal limitations. Advanced packaging technologies, thermally conductive metallisation layers, wafer thinning techniques and engineered interfaces are all areas where substantial progress is being made. Physical vapour deposition processes are expected to contribute heavily in this area by enabling highly controlled interfaces, advanced contact structures, passivation layers and thermal management coatings.

Interestingly, this is creating a convergence between previously separate areas of semiconductor development. Historically, the world of advanced logic and memory packaging evolved largely independently from power semiconductor manufacturing. Today, however, many of the same packaging concepts developed for high-performance computing are becoming increasingly relevant to next-generation power electronics.

This convergence is one of the most striking features of the current semiconductor landscape. As industries push simultaneously toward higher performance, greater efficiency and increased miniaturisation, technologies developed in one sector are rapidly finding applications in others. The challenges facing AI processors and

advanced power semiconductors increasingly overlap, particularly in the areas of thermal management, packaging density and power delivery.

For many industry observers, this broader technological convergence is helping accelerate gallium oxide's path toward commercialisation. Rather than requiring entirely new inventions, the material may benefit from innovations already occurring elsewhere in the semiconductor ecosystem.

That acceleration is also reflected in changing industry perceptions. Only a short time ago, many semiconductor manufacturers viewed gallium oxide as a technology that remained at least a decade away from practical deployment. Today, however, the conversation has shifted noticeably. Companies are actively evaluating how ultra-wide bandgap semiconductors might fit into future product strategies, driven largely by the growing demands of AI infrastructure, grid modernisation and electrified transportation.

As a result, gallium oxide now appears to be transitioning from a purely academic topic into the early stages of industrial deployment. Initial commercial devices could emerge within the next few years, while broader high-volume manufacturing may follow later in the decade if substrate scaling and packaging challenges continue progressing successfully.

Despite the excitement, few experts expect gallium oxide to replace silicon carbide or gallium nitride entirely. Instead, the future power electronics landscape is likely to become increasingly diversified, with multiple semiconductor materials serving different operating ranges and application domains

Despite the excitement, few experts expect gallium oxide to replace silicon carbide or gallium nitride entirely. Instead, the future power electronics landscape is likely to become increasingly diversified, with multiple semiconductor materials serving different operating ranges and application domains. Silicon carbide and gallium nitride will continue playing vital roles in many voltage classes, while gallium oxide is expected to establish itself primarily in the highest-voltage applications where its unique properties offer the greatest advantages.

In some areas, particularly at the upper end of the voltage spectrum, gallium oxide may eventually displace portions of the silicon carbide market, especially if its manufacturing economics prove superior at scale. However, the more likely scenario is one of coexistence rather than outright disruption. Future power architectures may increasingly incorporate multiple semiconductor materials across different stages of energy conversion, from the electrical grid all the way down to processors and edge devices.

What is becoming increasingly clear is that semiconductor innovation cycles themselves are accelerating dramatically. Technologies once expected to require decades to commercialise are now advancing far more quickly, driven by the strategic importance of semiconductors across nearly every major industry. AI, electrification and the global energy transition are all intensifying pressure for faster innovation and deployment.

Gallium oxide is emerging at exactly the moment when the industry needs new solutions for ultra-high-voltage power conversion. While significant engineering challenges remain, the combination of compelling electrical performance, scalable manufacturing potential and strong market demand is pushing the material rapidly toward commercial reality.

The next chapter in power electronics may therefore extend beyond silicon carbide and gallium nitride sooner than many expected. Increasingly, gallium oxide appears poised to become one of the defining semiconductor materials of the high-voltage era.



Thin film technologies **reducing energy consumption** in a power-hungry world

From more efficient powertrain management to reduced energy consumption for vehicle displays and lighting, Evatec thin film tools and production process know-how help automobile manufacturers reduce energy consumption extending battery life and range for the latest generation of electric vehicles. Evatec expertise in the manufacture of Wide Band Gap Power Devices reduces losses, while its mini- and micro-LED capabilities enable bright, efficient lighting and display technologies that drive down energy consumption in a power-hungry world.



Find out more about how Evatec thin film technology solutions are shaping our future

How can Evatec help you with your next thin film production challenge?

Learn more about Evatec front- and backside process know-how for the deposition of **Si, SiC** and **GaN** based power devices.



Find out more about the deposition of **metals, TCOs** and **dielectrics** in **optoelectronics** applications including **LED**.



Developing scalable traction inverters for compact electromobility solutions

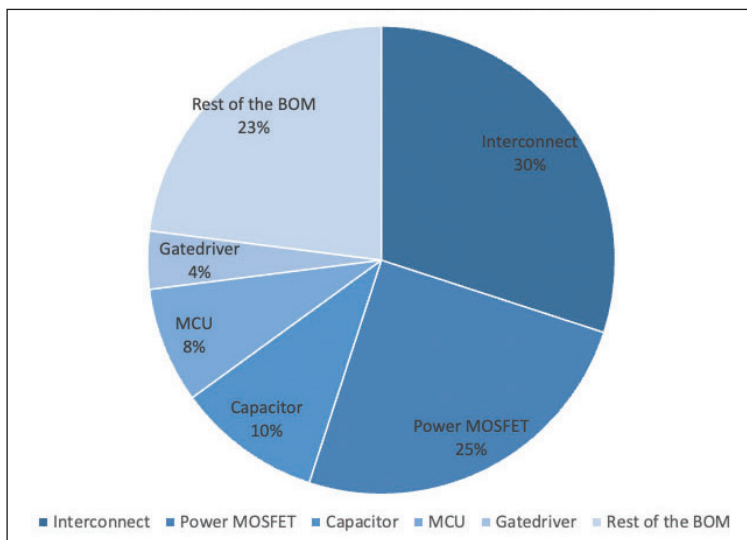
The broad spectrum of electromobility, ranging from the conveyance of individuals to the transportation of commercial goods, requires efficient, reliable and affordable forms of drivetrains.

BY PHILIPP MAI, VICE PRESIDENT ENGINEERING EMEA – AUTOMOTIVE & TRANSPORTATION, ARROW ELECTRONICS

THE TERM ‘electromobility’ is commonly associated with the electrification of automobiles. However, the concept of e-mobility extends far beyond this market to cover various contemporary forms of mobility. Examples include electrically-powered scooters, motorbikes and small utility vehicles. While e-scooters contribute to local emissions reductions, mainly in urban areas, e-motorbikes and electric small utility vehicles are increasingly being used as sustainable alternatives for wider leisure, transport and logistics purposes.

This broad spectrum of electromobility, ranging from the conveyance of individuals to the transportation of commercial goods, requires efficient, reliable and affordable forms of drivetrains. The traction inverter is a prime example. This key component transforms the direct current (DC) supplied by a battery into alternating current (AC) that can be used by the motor, as well as adjusting speed and engine torque to meet user requirements.

➤ Fig 1. Example of cost breakdown



Power variations

While battery-powered automobiles require high-voltage battery packs (400 V/800 V) to ensure sufficient performance and range, low-voltage batteries in the range of 36 V to 96 V have become standard for other forms of e-mobility systems. These diverse power requirements demand different technological solutions from the corresponding traction inverters.

In each case, the traction inverters must be designed to conform to the relevant technical framework conditions and the specification profiles. During development, particular emphasis is placed on:

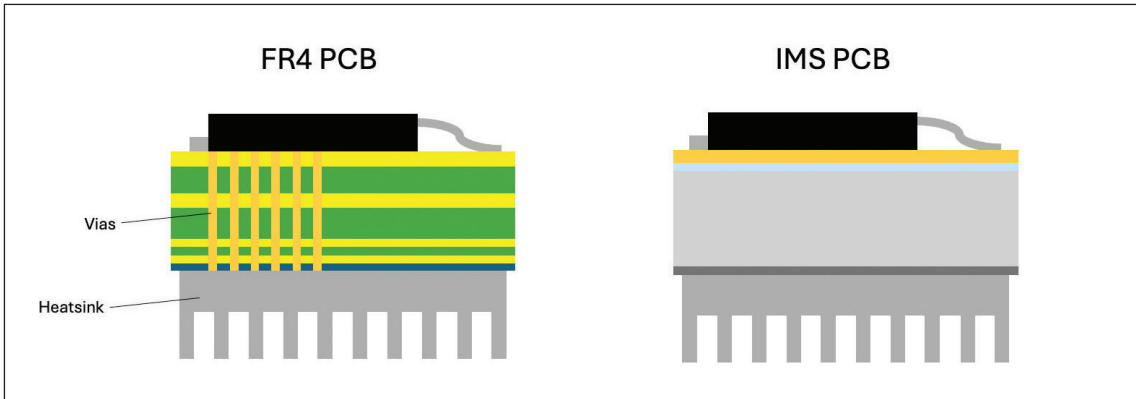
- Compact construction
- High efficiency
- Cost-effectiveness
- Robustness and reliability
- Functional safety

To meet these requirements, it is advisable to start with a basic concept that is scalable to the required power (typically 3 to 5 kW) and voltage range (36 to 96 V). Particular attention should be paid to the circuitry design, because it enables adaptation to the varying use profiles and determines the cost of the inverter. (Fig 1.)

MOSFETs and power modules

The use of MOSFETs in SMD housings is one tried and tested concept for the design of final circuitry in low-voltage traction inverters. This approach allows the heat generated to be transferred via the circuit board to a cooling element or the housing for dissipation. Thermal vias with aluminium base PCBs have emerged as the most cost-effective solution for this application. (Fig 2.)

However, a disadvantage of this concept is the excessive heating of the PCB. This can make it



➤ Fig 2. SMD Concept

impossible to position the necessary filter capacitors directly in connection with the PCB, as the increased heat input would significantly reduce their service life. As a result, the corresponding designs are frequently more complex and less compact.

TSC (top-side cooled) MOSFETs have subsequently become the solution of choice in high-power applications. (Fig 3.)

This method offers the distinct advantage that MOSFETs can be directly connected to a cooling element, providing effective discharge of heat. In addition, a more cost-effective PCB can be employed. However, increasing performance demands mean that parallel connection of the MOSFETs becomes essential, which in turn results in additional complexity of designs (thermal management, layout, eMV) and makes compact construction more difficult.

Although both these methods are scalable, their respective disadvantages limit their use within contemporary low-voltage traction inverters, especially when trying to meet challenging load and performance requirements. A viable alternative can be found in the type of power modules used in e-vehicles.

Power modules offer several advantages in the context of contemporary low-voltage traction inverters. Features such as component integration and efficient thermal management enable a compact inverter design and reduce development costs. In addition, power modules are more reliable because they have lower failure rates than more discrete solutions. Their design also makes scalability easier, allowing developers to adapt circuitry more rapidly and cheaply to the various power and voltage classes. An example of such a module is illustrated in Fig 4.

Making good contacts

In addition to the design for the final circuitry, the contacts required for high currents represent a challenge for developers. In the past, simple screw contacts and ring cable lugs were often employed. Although cost-effective, this method no longer conforms to current reliability and safety requirements.


A modern contact system must exhibit minimal contact resistance and offset tolerances and provide a reliable contact despite the extensive shock, vibration and thermal cycles to which the system will be exposed. The contact system must therefore be based on a locking design.

The SW1 COEUR by Molex is one such example. (Fig 5.)


Despite its high current-carrying capacity, this contact system has a compact design, allowing it to be readily installed in the field or in workshops. There are three different contact system sizes (6.00, 8.00 and 11.00 mm), facilitating scalability and cost-effective adaptation to various circuitry designs requiring currents of up to 300 A.

A matter of control

Using the correct control card is essential in ensuring the scalability of a low-voltage traction inverter - and the microcontroller plays a key role. Solutions are widely available for providing




➤ Fig 3. TSC Concept



➤ Fig 4. Vishay module

➤ Fig 5.
The SW1
COEUR



sufficient power to process the complex algorithms that control the motor (such as FOC - field-oriented control) and suitable peripherals (e.g. precise PWM generation or ADC). However, they are not always cost-effective and offer only limited options regarding functional safety. This is a concern because the functional safety of the microcontroller and the system as a whole is becoming increasingly important in gaining access to international markets for compact e-mobility solutions.

When choosing a suitable microcontroller, hardware features are no longer the most crucial factor. Increasingly, it is the software that determines the complexity and costs of traction inverter development. For example, the S32K3 family by NXP offers a microcontroller platform suitable for low-voltage traction. These units come with certified software packages for use in the field, such as Real-Time (RTD) and Safety Peripheral Drivers (SPD), ASIL D, an AMMCL (Advanced Math and Motor Control Library), SAF (Safety Software Framework), and a Structural Core Self-Test (SCST) Library. This significantly reduces the necessary design input.

Isolation aspects

Galvanic isolation is a further consideration in designing and scaling compact systems for e-mobility. In circumstances where low-voltage traction inverters are exposed to moisture, dust and soiling, the isolated driver circuit design not only provides efficient control of a MOSFET but also supports the reliability and galvanic isolation of the system as a whole.

For high-voltage traction inverters, special gate drivers are used. These are considered oversized for use with systems operating at voltages below 96 V. One example of a potentially suitable gate

driver is the Si8285CD-AS by Skyworks. In addition to sufficient capacity and isolation, it offers suitable integrated safety features (UVLO, DESAT), together with a dedicated pin for automatic feature activation.


Achieving a complete, scalable design

The Arrow reference design, developed by Arrow Electronics in collaboration with its partner Vishay, represents a complete, scalable low-voltage traction inverter for compact e-mobility solutions. The control board and circuitry are strictly separated to provide for modularity.

The emphasis is on a scalable, final circuitry design based on an innovative, top-side cooled power module by Vishay. This module design allows for flexible and effective adaptation to various performance classes from 3 kW to 15 kW and voltage classes up to 96 V – all within a single form factor. To optimise the form factor, Vishay SMD hybrid capacitors have been attached to the circuitry system and directly integrated into the cooling concept of the power modules. This innovative combination of functionality and flexibility opens up a wide range of possible applications, setting new standards in the design of final circuitry for use in compact e-mobility solutions.

The control board has been developed with a clear focus on reliability and safety. The isolation provided (1 Kvdc) ensures the system can be flexibly used across various voltage classes, making it versatile and robust. Various aspects relating to functional safety were considered during development, including the comprehensive use of safety PMICs from the FS23 family by NXP. These are characterised by a wealth of features that provide for a wide range of diagnostic operations and thus ensure system integrity. Moreover, the gate drivers are monitored by a diagnostic output that ensures early detection of malfunctions. The phase voltage and current of the final circuitry design are measured and monitored by isolated amplifiers from the Si89xx series by Skyworks.

This combination of integrated safety features and reliable design concepts provides the highest level of performance. It also complies with the most stringent requirements for safety-critical low-voltage traction inverters for use in e-mobility applications, from electrically-powered scooters to commercial vehicles.




➤ Vishay
control board



www.EVGroup.com

HIGH-VACUUM BONDING



EVG® ComBond®

Oxide-free, covalent wafer bonding
High-vacuum process environment
Room-temperature process



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

IDTechEX: Wide bandgap momentum meets silicon's staying power

As SiC and GaN gain ground in EVs and data centres, silicon's proven reliability continues to anchor wind power – but for how long? Rebecca Pool talks to Matthew Fall, IDTechEX, to find out more.

WITH THE power electronics market on course to more than double in the next decade, the industry is entering an era of transition. Analyst firm, IDTechEx Research, forecasts market growth from US\$25.5 billion this year to US\$ 65.2 billion in 2036 – representing a CAGR of 10% largely driven by electric vehicle and data centre applications.

As OEMs push for ever-greater efficiency, reliability and power density, both silicon carbide and gallium nitride will capture market share from silicon. Yet despite the clear momentum behind wide bandgap semiconductors, the industry mainstay is not stepping aside yet.

“[SiC and GaN] technologies have the potential to revolutionise the power electronics industry, enabling high-voltage operation and new power architectures, such as the 800 V e-powertrain and 800 VDC data centers,” says IDTechEx Technology Analyst, Matthew Fall. “But silicon's proven reliability is slowing the adoption of these wide bandgap technologies in wind energy – a smaller segment of the overall market but an important contrast to EVs and data centres.”

As Fall highlights, wind generation operates at very high voltages and power, in harsh environments – think extreme temperature cycling, high humidity and strong vibrations – and component failure is very costly. Given this, component reliability is paramount and the risk of switching to newer technologies remains a major constraint.

“A lot of the key industry players in this sector are reluctant to switch away from silicon – you've got at least 50 years of proof of concept with silicon compared to around eight years for silicon carbide,” he says.

Also, given the sheer size of the wind turbine nacelle, where space and weight constraints are less critical, the advantages of wide bandgap semiconductors become less compelling. “Miniaturisation is less relevant and while silicon carbide has certainly proven its reliability in electric vehicles, reliability concerns are so much higher,” adds Fall, “So we're seeing that the transition to silicon carbide in wind generation is still at an earlier stage.”

Still, late last year, Wolfspeed partnered

with Hopewind to integrate its 2.3kV LM Pack module into the China-based renewable energy tech firm's 950 VAC wind power converter. The all-SiC power module promises to simplify overall system design and provide higher efficiency, power density and reliability. As Wolfspeed noted at the time, the partnership marks a critical milestone in the ongoing evolution of the wind power industry, helping to pave the way for cleaner, more efficient energy solutions on a global scale.

But while this early integration of SiC in wind power is significant – is it a clear industry shift or more of a cautious first step? SiC MOSFETs have been researched and prototyped in wind power electronics for well over a decade with manufacturers introducing advanced packaging technologies and improved gate driver designs to ensure stable operation within the challenging conditions of the wind turbine nacelles. On the grid-side, customized switching is also being investigated to help reduce electromagnetic interference in long cable runs typical in wind farms. Even so, Fall is certain that silicon is safe, at least for now.



➤ Offshore wind turbines: silicon's proven reliability means manufacturers are reluctant to adopt silicon carbide – for now.

“I’ve spoken to some large industry players in wind power, and they don’t even have a transition to silicon carbide on their roadmaps yet,” he says. “We know that GaN and SiC are already making their way into solar [generation] but for wind power, this will be a slower, steady adoption – silicon will have a significant place here in the next ten years.”

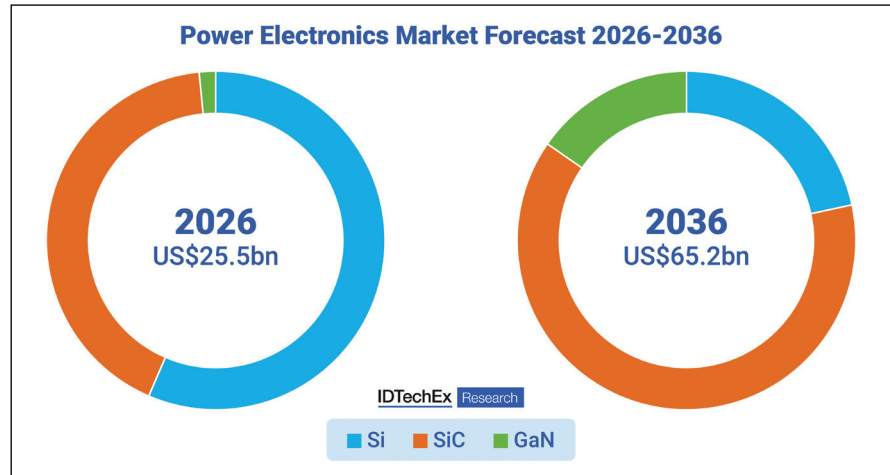
Beyond turbines

Electric vehicles are undoubtedly the dominant market for power electronics with silicon IGBTs and other silicon components having been the power device of choice in traction inverters, onboard chargers and DC-DC converters for the past two decades. But as Fall highlights, SiC MOSFETs now account for a growing revenue share of these applications and are expected to dominate traction inverters, onboard chargers and DC-DC converters by 2036. The high-temperature operation, faster switching speeds and smaller form factor offered by these devices, can deliver improved efficiency as well as weight and volume savings that ultimately support increased EV range and performance.

“What we’re now seeing is that SiC has reached this critical inflection point in the EV space – and is no longer seen as emerging,” he says. “The challenge for SiC is now to convince manufacturers that its benefits are worth [the switch from silicon] for more than EV performance models, as right now its cost premium makes it a harder sell for budget vehicles.”

Still, the harder sell may not last a lot longer. Fall points to aggressive competition between Chinese SiC manufacturers dramatically reducing the price of SiC wafers. Factor in the rising quality of these wafers coupled with more and more eight-inch substrates reaching the market, and it’s easy to see why revenues are consistently growing.

And it’s not just silicon carbide. Relatively early work on GaN devices in EV traction inverters from the likes of Cambridge GaN Devices (CGD) and VisIC Technologies shows promise. Late last year, CGD won Hyundai’s ‘Open Innovation Challenge for Sporty Driving’ with its one-chip ICeGaN power transistor and also partnered with Global Foundries on production.



➤ Silicon carbide is forecast to take the majority share of the power electronics industry by 2036. (IDTechEX)

Around the same time, VisIC raised some US\$26 million to advance its D³GaN platform for electric vehicles, with the funding round led by Hyundai.

According to Fall, the development of automotive GaN depends on proof of its long-term reliability as well as its ability to operate at high voltages for 800V EV power architectures – but actual tests in cars are expected in the next couple of years. “We don’t expect the EV market to completely dry up for silicon though,” he says. “The cost trade-off will remain for your budget EV that’s used for driving around cities at lower speeds.”

Data centres follow electric vehicles in terms of market share for power electronics, and offer significant opportunities – but less so for silicon. While silicon has long formed the backbone of data-centre electronics, today’s large language models and other neural networks draw substantial power when processing complex AI and machine-learning workloads – and the semiconductor’s physical limitations are becoming more and more apparent. As the industry moves from AC power delivery towards 800 VDC architectures, the pressure to improve efficiency, reduce losses and shrink power systems is also intensifying. To support compute-intensive workloads, power supply unit manufacturers are increasingly turning to wide bandgap semiconductors with their lower on-resistance, higher switching speeds, reduced losses – and decreasing costs.

“The space for processing data centre workloads is becoming more and

more a premium... and there’s a strong driver for miniaturisation here,” says Fall. “In the transition to 800 VDC, I think a winning technology will be GaN, because of the need for miniaturisation and ultra-fast switching.”

Fall highlights how organisations such as Infineon and imec are also rapidly developing 12 inch GaN-on-silicon wafers to boost economies of scale. “We believe that GaN will overtake SiC in terms of cost, and could well approach the cost of silicon in the long-term - beyond ten years,” he says.

Looking to the future, Fall also emphasises how all power electronics applications are interconnected, with innovation in one sector driving innovation in a different sector. “Switching to SiC in the wind sector would never have made a compelling argument without the history of significant research in electric vehicles,” he says. “Many OEMs still have a certain amount of tunnel vision and focus on their market. But it’s the developments in other markets that can direct research and development.”

And amidst the undeniable market erosion for silicon, Fall is confident that the industry incumbent will remain, at least within the planning horizons of today’s OEMs. “Wide bandgap semiconductors fundamentally exist to complement silicon in your more challenging environments, whether that’s high power at high voltages or ultra-fast switching frequencies,” he asserts. “We’re unlikely to see a time when silicon is completely replaced.”

800VDC power delivery architecture and advantages for next-generation AI & GPU data centers

A comprehensive comparison with traditional AC-based power systems.

PART 1

BY SEAN BURKE, CEO OF ENTELENT

MODERN data centers are undergoing a fundamental transformation driven by the explosive growth of artificial intelligence and high-performance GPU computing. The power demands associated with these systems are increasing at a rate far beyond the trajectory of traditional enterprise IT. Whereas historical rack power densities ranged from 3 to 12 kW, today's AI training clusters routinely exceed 50 kW per rack and are rapidly pushing toward 80 kW, 100 kW, and even 150 kW in advanced deployments. This rise is not incremental but exponential, driven by the escalating computational requirements of large-scale neural networks, massive memory bandwidth needs, and the shift toward tightly integrated multi-GPU architectures.

In this environment, power delivery and power density are defining constraints. High-performance GPU

servers can draw several kilowatts per node, and a single GPU module may spike hundreds of amperes during intense computational bursts. As these loads aggregate across dozens of servers within a rack, the limitations of conventional AC power distribution become painfully apparent. Thermal dissipation, conductor ampacity, and conversion inefficiencies introduce bottlenecks that make traditional architectures increasingly impractical. Many facilities now find themselves limited not by the availability of utility power, but by the physical and electrical constraints of delivering that power safely and efficiently to the rack.

These challenges are magnified by the multi-stage nature of AC power systems. At each step—from medium-voltage utility delivery to transformers, UPS systems, branch distribution, rack-level PDUs, and finally server PSUs—AC

must be transformed or rectified, with each conversion introducing losses and heat.

This stacked inefficiency compounds rapidly at high densities, resulting in significant waste, elevated cooling requirements, and escalating operational costs. As data centers expand to accommodate thousands of high-density AI racks, the cumulative penalties of the AC model threaten both scalability and sustainability.

Against this backdrop, High-Voltage Direct Current distribution is emerging as a clear alternative. HVDC architectures, particularly those centered around an 800VDC backbone, address the fundamental limitations of AC by reducing current, minimizing conversion stages, and enabling dramatically higher efficiency. By delivering power at high voltage



and converting it to 50VDC at the rack—where modern servers already operate—HVDC aligns power delivery with the actual electrical needs of contemporary compute systems. The result is a streamlined, highly efficient energy pathway optimized for the next generation of AI-driven workloads.

The acceleration of AI power density

The rapid adoption of GPU-accelerated computing has fundamentally altered the power profile of the modern data center. AI servers now incorporate dozens of high-performance GPUs tightly coupled to CPUs, memory systems, and high-bandwidth interconnects. These configurations routinely draw 2–3 kW per GPU, and full server nodes often exceed 16–32 kW of continuous load. When aggregated across a rack, these systems demand unprecedented power densities that strain traditional electrical infrastructure.

Unlike legacy compute, AI workloads generate highly dynamic load patterns, with rapid transitions between idle and full-power states. These transitions create steep current spikes that challenge voltage regulation and thermal management. As racks push past the 50 kW threshold, the AC distribution model encounters both mechanical and electrical limitations: cable bulk increases, termination points overheat, and multiple AC conversion stages struggle to deliver consistent and efficient power. The power delivery system, rather than the compute hardware, increasingly becomes the primary bottleneck for AI capacity growth.

The case for HVDC

High-Voltage Direct Current provides a fundamentally more efficient and scalable electrical architecture for AI-dense environments. By transmitting power at 800VDC, current is drastically reduced compared to 240VAC distribution, lowering resistive losses and enabling smaller, lighter, and more thermally manageable conductors. Equally important, HVDC eliminates the need for multiple AC conversion stages by leveraging a single high-efficiency rectification step before distributing power throughout the facility.

At the rack, compact 800VDC-to-50VDC converters deliver power directly into servers designed around

modern 48/50/54VDC backplanes. This architecture removes the traditional AC power supply from each server, eliminating one of the least efficient components in the IT power chain. With fewer conversions, higher efficiency, reduced copper consumption, and improved reliability, HVDC naturally supports the scale and density of future AI clusters.

The technology itself is mature: electric vehicles, solar inverters, industrial drives, and commercial battery systems already rely on high-voltage DC architectures. Data centers now face the same power delivery challenges those industries confronted years ago, and they are poised to adopt the same solution at scale.

Traditional AC power path architecture

Traditional AC distribution relies on large transformers, UPS systems, PDUs, and per-server AC power supplies. These components introduce conversion losses and heat. Traditional high-current 240AC feeders also have the higher, unnecessary cost from oversized copper wire and conduit.

For decades, traditional AC-based power distribution has formed the backbone of data center electrical infrastructure. This architecture matured during an era when rack power requirements were modest—typically between 3 and 12 kW per rack—and when compute loads were dominated by CPUs, storage appliances, and networking systems operating well within the thermal and electrical constraints of low-density IT environments. As a result, AC systems evolved around a multi-stage power conversion chain that was considered efficient enough for the times, even though each stage introduced losses, heat, and operational complexity.

Thermal limitations represent another growing challenge. AC systems require high currents to deliver sufficient power to AI racks, and these currents generate substantial localized heating in conduits, cable trays, and distribution panels. As power densities increase, operators often find themselves constrained not by available wattage from the grid, but by the physical ability of AC distribution to safely carry and dissipate the heat generated by high currents. These thermal bottlenecks

The technology itself is mature: electric vehicles, solar inverters, industrial drives, and commercial battery systems already rely on high-voltage DC architectures. Data centers now face the same power delivery challenges those industries confronted years ago, and they are poised to adopt the same solution at scale

place a natural ceiling on how far AC architectures can scale without major redesigns or costly reinforcements.

In summary, the traditional AC power architecture remains functional for legacy compute but struggles to accommodate the massive electrical demands of next generation GPU clusters. Its multi-stage conversion path, high current requirements, thermal limitations, and reliance on bulky copper conductors introduce inefficiencies and scalability barriers that become increasingly unsustainable as data centers transition into the AI era. These limitations set the stage for the emergence of high voltage DC distribution as a more capable and forward looking alternative.

Typical AC power chain

In a conventional data center, medium-voltage utility power is stepped down by large transformers before passing into UPS systems designed to ensure power quality and continuity. These UPS units commonly employ double conversion architectures that rectify AC to DC and then invert it back to AC, incurring significant energy losses in the process.

From the UPS, power is routed through low-voltage AC distribution panels and into rack-level power distribution units (PDUs), which divide the incoming feed into multiple branch circuits. Inside each server chassis, the AC supply must then be rectified once more - this

time by the server's internal power supply - to create low-voltage DC suitable for CPU, GPU, memory, storage, and motherboard VRMs.

A standard modern data center uses the following flow:

- Utility Medium-Voltage AC (13.2–34.5 kV)
- MV ⇄ LV transformer
- UPS (double-conversion or line-interactive)
- Low-voltage AC distribution (208/240/277/480 VAC)
- Rack PDUs
- Server PSUs (AC ⇄ 12V or AC ⇄ 48V)
- Onboard VRMs powering CPUs/GPUs

Every stage introduces losses. As power climbs above 50 kW per rack, the AC model becomes physically cumbersome and inefficient.

Efficiency losses

The compounding effect of multiple conversion stages further erodes system efficiency. Transformers, UPS systems, PDUs, and server power supplies each have their own loss profiles. Even with modern designs, UPS systems typically operate in the mid-90% efficiency range while server PSUs often achieve 92–95% efficiency under ideal conditions. When all stages are combined, the usable end-to-end efficiency may drop to as low as 78–85%. This means that a substantial portion of incoming utility power is dissipated as heat before it ever reaches the compute hardware - an untenable inefficiency in large-scale AI infrastructure where electrical demand rises exponentially.

Thermal limitations

While effective for legacy compute, this AC-centric power path is increasingly misaligned with the demands of modern AI and GPU workloads. As rack power densities rise into the 40 kW to 120 kW range, the inherent inefficiencies of the AC model become more pronounced. At 208-240V, delivering high power requires high current levels (200-400 A per rack), which demands thick copper conductors and, in many cases, multiple parallel cable runs. These heavy cables generate significant heat due to I^2R losses, placing stress on both conductors and terminations while driving up cooling requirements. Mechanical constraints emerge as well: routing large conductors through

Every stage introduces losses. As power climbs above 50 kW per rack, the AC model becomes physically cumbersome and inefficient.

overhead busways or underfloor systems becomes difficult, labor-intensive, and costly.

Higher currents (e.g., 200–400 A per rack) cause:

- Excessive I^2R losses
- Increased heat within conduits
- Requirement for parallel conductors
- Difficulty in using flexible whips or overhead busways
- Increased mechanical stress at terminations

Ultimately, AC won't work for the power densities demanded by AI demand.

The 800VDC rack power architecture

800VDC distribution reduces current by ~70% compared to 240V AC, lowering resistive losses and voltage drop. Rack level HVDC (High-Voltage DC) power supply converters achieve 97–98% efficiency and eliminate the AC conversion stages which is achieved outside of the building by a high-efficiency utility-scale inverter. The transition to an 800VDC distribution architecture represents a fundamental shift in how power is delivered within next-generation data centers. Traditional AC systems struggle to keep pace with the rapidly increasing rack-level power densities demanded by AI and GPU infrastructure. By contrast, an 800VDC backbone dramatically reduces current levels for a given power capacity, which in turn lowers resistive losses, minimizes copper mass, and allows much longer power distribution runs without the voltage-drop penalties seen in AC systems. This approach mirrors the power architectures that have already proven successful in electric vehicles, utility-scale solar installations, and industrial automation systems.

Overview

At the core of the architecture is a centralized rectification stage that

converts medium-voltage AC into a stable 800VDC bus. From this point forward, the power remains in DC form as it is distributed throughout the facility. Each rack contains a high-efficiency HVDC converter, typically achieving 97–98% efficiency, that steps down the 800VDC to a native 50VDC level suitable for modern AI servers. This reduces the number of conversion stages and effectively eliminates the server-level AC power supply entirely. In addition to lowering losses, this architectural simplicity enhances overall reliability by reducing the number of power conversion components susceptible to thermal and electrical stress.

The 800VDC distribution system includes:

- Central rectifier (MV AC → 800VDC)
- 800VDC backbone distribution to each rack
- Rack-mounted high-efficiency converter (800VDC → 50VDC)
- Direct 50VDC bus feed into servers. This eliminates the server PSUs entirely or reduces them to simplified DC input modules.

Efficiency advantages

- State-of-the-art HVDC converters
- 97–98% efficient at rack level
- Maintain high efficiency across 20–100% load

The end-to-end chain reaches 92–95% efficiency, a significant gain over the 78-85% efficiency of AC-coupled UPS systems.

Reduced current & copper

The reduction in current afforded by 800VDC distribution is one of its defining advantages. For example, delivering 100 kW at 800V requires only 125 amps, compared to more than 400 amps in a 240VAC system. This reduction translates to smaller-gauge conductors, lighter and more flexible cabling, and reduced installation labor. The lower current also decreases heat generation within conduit and raceways, improving both safety and thermal management. These improvements prove to be decisive as rack power levels move beyond 80 kW, where AC distribution approaches its practical operational limits.

At 800VDC:

- 100 kW → 125 A
- 50 kW → 62.5 A

Contrast with 240 VAC:

- 100 kW → 417 A
- 50 kW → 208 A

Safety of HVDC

Equally important, modern HVDC systems have evolved significantly in terms of operational safety. Semiconductor based fast-acting breakers, arc-fault detection algorithms, insulated bus ducting, and touch-safe connectors all contribute to a robust and compliant safety ecosystem. The widespread adoption of 800–1500VDC systems in commercial vehicles and renewable energy systems further demonstrates the maturity and reliability of HVDC technologies. When applied to data centers, these advances provide a strong foundation for a scalable, safe, and efficient power delivery architecture.

The proliferation of 800–1500VDC is already standard in EVs, solar farms, and commercial batteries.

Server-Level power delivery: 50VDC advantage

48/50/54V server backplanes improve VRM efficiency and reduce board-level copper losses. Higher voltage distribution enhances transient response for GPU loads. Eliminating AC PSUs increases reliability and rack density.

Why 50VDC is optimal

As data centers evolve toward increasingly power dense compute environments, server-level power delivery requirements are also undergoing significant transformation. The industry's shift toward 48/50/54VDC input buses reflects a convergence on the optimal voltage range for balancing efficiency, thermal performance, and practical implementation. Traditional 12V architectures struggle to support the extreme currents demanded by modern GPUs and AI accelerators. By contrast, a 50VDC distribution plane substantially reduces current flow across backplanes, connectors, and motherboard copper traces, allowing for simpler mechanical designs and more efficient voltage regulation.

The industry is converging on 48/50/54V buses for reasons, not limited to:

- Lower copper loss than 12V
- Higher efficiency for 1V VRMs

- Better current handling for >20 kW chassis
- Backplane distribution becomes easier
- Standardization across OCP (Open Compute Platform) and hyperscale platforms

GPU load characteristics

One of the key benefits of a 50VDC input bus is its impact on VRM performance.

High-performance GPUs frequently draw hundreds of amperes during rapid load transitions, placing enormous instantaneous demands on the server power subsystem. VRMs operating from 50VDC enjoy wider duty-cycle margins and reduced stress on switching components, enabling superior transient response and lower conduction losses. This is especially important in AI workloads, where GPUs often move abruptly between idle and full-compute states, repeatedly stressing the power delivery system.

AI workloads demand:

- Step loads of hundreds of amperes per GPU
- Low-impedance, low-noise input rails
- Fast transient response
- High VRM current capacity (>1000 A requiring multiphase topologies)

50VDC busbars simplify this, enabling:

- Wider VRM duty cycles
- Lower peak currents
- Reduced heat on motherboards

Eliminating traditional PSUs

Transitioning to a 50VDC rack power environment also eliminates the need

for traditional AC power supplies within each server. AC PSUs are comparatively bulky, add thermal overhead, and introduce an additional conversion stage that reduces overall efficiency. By replacing them with lightweight DC input modules or direct-feed distribution rails, server density can be improved while simultaneously reducing points of failure.

Rack-level HVDC-to-50VDC converters effectively become shared, high-efficiency power shelves that can be operated in redundant configurations, further enhancing system reliability.

AC server PSUs are bulky, expensive, and low-efficiency compared to modern DC shelves.

With 50VDC:

- PSUs are replaced by simple input modules
- One HVDC → 50VDC unit can feed multiple chassis
- Hot-swap DC power shelves support redundancy

This significantly improves server density and reliability.

Overall, the adoption of a 50VDC power bus at the server level is a critical enabler for next-generation AI systems. It harmonizes with OCP standards, improves board-level efficiency, simplifies mechanical and electrical design, and provides a foundation for supporting extremely power dense GPU platforms that will define the coming decade of AI compute.



From conversion tax to DC-native: Re-engineering data centre power for the AI Era

An interview with Dafna Granot, senior manager strategy and innovation at SolarEdge, author of the white paper: Powering the AI revolution: A Maturity Roadmap to Integrated 800V DC, who discusses the five-stage maturity model for data centres to transition from AC to DC-native infrastructure and the role played by Solid-State Transformers in defining the next generation of AI factories.

THE RAPID expansion of artificial intelligence (AI) workloads is forcing a fundamental re-evaluation of how data centres are powered. For decades, alternating current (AC) has served as the foundation of data centre electrical infrastructure, supported by mature standards, well-established supply chains, and familiar operational practices. However, the emergence of GPU-intensive AI systems is exposing critical inefficiencies in these legacy architectures, particularly in the form of repeated power conversions that drain energy and limit scalability.

What was once considered an acceptable level of inefficiency - often described as a “conversion tax” -


has become a central constraint. As data centres strive to deliver ever greater compute performance within finite power envelopes, the industry is beginning a decisive shift toward high-voltage direct current (DC) architectures. In this transition, 800V DC is emerging as the immediate target, while more advanced DC-native systems promise further gains in efficiency, scalability, and economic return.

AI workloads are driving an unprecedented shift in power density

Traditional data centre applications, such as enterprise computing and web services, typically operated within

relatively modest power constraints. Rack densities in the range of 5 to 10kW were once the norm, and infrastructure inefficiencies could be tolerated because the overall demand profile remained manageable.

AI workloads are fundamentally different. They require dense clusters of GPUs operating at high utilisation levels, often in tightly integrated configurations that demand both high performance and low latency. This has led to a dramatic escalation in power density. Where 10kW per rack was typical in the past, modern AI deployments frequently reach 100kW per rack, and projections indicate that this figure could rise to nearly 1MW



per rack within the next generation of systems.

This rapid increase introduces significant challenges. Existing AC infrastructure places a practical upper limit on rack density, often around 200kW. Beyond that point, thermal management, power delivery, and physical constraints become increasingly difficult to overcome. At the same time, the sheer scale of power being consumed means that even small efficiency losses translate into massive amounts of wasted energy.

The rising cost of conversion losses

Legacy data centre designs rely on multiple stages of power conversion. Electricity typically enters the facility as AC from the grid, undergoes voltage transformation, is distributed throughout the facility in AC, and is then converted into DC within the uninterruptible power supply (UPS). The power may subsequently be converted back to AC and ultimately converted again to DC at the server or rack level, where it is actually used by electronic components.

Each of these conversions introduces losses. When combined, they result in overall system efficiencies that generally fall between 84 percent and 91 percent. In practical terms, this means that a significant portion of incoming energy is dissipated as heat rather than being used for computation.

In the past, this inefficiency was largely accepted. Today, however, it represents a critical barrier. In a large-scale facility operating at 100MW, conversion losses alone can account for upwards of 10MW of wasted power. This wasted energy must also be removed through cooling systems, adding further overhead and complexity.

With AI workloads pushing both capacity and performance limits, such inefficiencies are no longer tolerable. Every unit of power that is lost to conversion is one that cannot be used to operate GPUs or deliver value to customers.

Elevating voltage: the case for 800V DC

One of the most effective ways to improve power delivery efficiency is to increase voltage. Higher voltage enables lower current for the same



power level, which in turn reduces resistive losses and minimises heat generation. This principle is particularly important in high-density environments, where thermal constraints are already a major concern.

As a result, the industry is converging on 800V DC as a near-term standard for next-generation data centres. This voltage level offers a balance between performance gains and practical implementation. It allows for significant reductions in current, enabling thinner conductors, reduced material use, and improved overall system efficiency.

Another key advantage of 800V DC is that it is already widely used in other industries. Electric vehicle platforms and solar energy systems have established mature ecosystems around this voltage level, including components, design expertise, and safety standards. This existing infrastructure enables data centre operators to adopt 800V DC without waiting for entirely new supply chains to develop.

While 800V represents a major step forward, it is not the end of the journey. The industry is already discussing a future transition to 1500V DC, which

would offer even greater efficiency gains. However, challenges related to safety, standardisation, and component availability mean that 800V is currently the most practical and immediate solution.

A structured path: the five-stage transition model

The shift from AC to DC-native infrastructure is best understood as a progression through several stages of maturity. At the starting point, identified as Stage 0, most data centres operate using traditional AC distribution. These systems are characterised by multiple conversion stages and relatively low overall efficiency.

The first step toward modernisation, Stage 1, involves introducing a local conversion device, sometimes referred to as a sidecar, near the rack. This approach allows operators to support 800V DC loads without fundamentally altering the rest of the infrastructure. While this enables higher-density compute, it does little to address overall efficiency.

Stage 2 builds on this by replacing AC-based UPS systems with DC-based alternatives. This eliminates



redundant conversions and increases overall efficiency into the low-to-mid 90 percent range. While still a hybrid approach, it represents a meaningful improvement.

The transition becomes more significant at Stage 3, where the architecture begins to take on a truly DC-native form. This is achieved through the introduction of solid-state transformers, which convert AC to DC and regulate voltage in a single, highly efficient device. By simplifying the power chain, this stage enables efficiency levels approaching the high 90 percent range.

At the final stage, Stage 4, the system is fully integrated around a DC-native design. Here, advanced solid-state

transformers incorporate additional functionality, reducing component count and enabling direct connection to medium-voltage grid inputs. This stage represents the highest level of efficiency and the most streamlined architecture.

Solid-State transformers as a key enabler

Central to the evolution toward DC-native data centres is the development of solid-state transformers. Unlike conventional transformers, which perform only voltage conversion, these devices combine multiple functions within a single system. They not only transform voltage but also convert AC to DC and provide power conditioning.

This multifunctional capability allows for a dramatic simplification of the power distribution chain. By reducing the number of conversions required, solid-state transformers improve efficiency and reduce system complexity. They also offer greater flexibility, as their behaviour can be controlled through power electronics rather than fixed electromagnetic designs.

Although the technology is advancing rapidly, solid-state transformers are not yet widely available in commercial deployments. Industry momentum suggests that this will change in the near future, particularly as the rollout of next-generation AI hardware creates demand for more efficient power systems.

The role of silicon carbide in high-efficiency systems

The performance of solid-state transformers and high-voltage DC systems is closely tied to advances in semiconductor technology. In this context, silicon carbide has emerged as a critical enabler.

Silicon carbide devices offer several advantages over traditional silicon components. They can operate at higher voltages, making them well suited for 800V and beyond. They also exhibit lower switching losses and lower conduction losses, both of which contribute to improved efficiency. These characteristics are essential for achieving the high performance required in modern data centre power systems.

Although gallium nitride is another promising wide-bandgap material, its current limitations in high-voltage applications make it less suitable for use in large-scale power distribution. However, continued development may expand its role in the future.

Modular architectures and the benefits of scalability

Another important aspect of the emerging DC-native paradigm is the shift toward modular system design. Instead of relying on large, centralised transformers, operators can deploy smaller, standardised building blocks in the range of 2 to 5MW.

This modular approach offers several advantages. It enables faster

Central to the evolution toward DC-native data centres is the development of solid-state transformers. Unlike conventional transformers, which perform only voltage conversion, these devices combine multiple functions within a single system

deployment, as components can be pre-manufactured and installed as needed. It also improves scalability, allowing capacity to be added incrementally in response to demand. From a cost perspective, mass production reduces capital expenditure and shortens lead times, which are currently a significant challenge for traditional transformer procurement.

Modularity also enhances system resilience. In a distributed architecture, the failure of a single module has minimal impact on overall operation. This contrasts with centralised systems, where a single point of failure can disrupt an entire facility.

Overcoming barriers to adoption

Despite its clear advantages, the transition to DC-native infrastructure presents several challenges. One of the most significant is the need for a mature supply chain capable of supporting high-voltage DC components. While progress is being made, the ecosystem is still developing.

Standards and regulatory frameworks are another area of concern. AC systems benefit from decades of established guidelines, whereas DC systems are still in the process of being standardised. Safety considerations are also different, particularly in relation to fault detection and arc suppression, requiring new approaches and expertise.

Additionally, many data centre operators lack experience with DC

systems, creating a knowledge gap that must be addressed through education and partnerships with specialised vendors. Finally, the transition can be disruptive, particularly at higher stages of maturity where significant infrastructure changes are required.

Economic imperatives drive change

Ultimately, the move toward DC-native architectures is driven by economic considerations. Efficiency gains translate directly into increased compute capacity and revenue. By reducing conversion losses, operators can reclaim power that would otherwise be wasted and redirect it toward productive use.

In large-scale facilities, even modest improvements in efficiency can have a substantial financial impact. The ability to support more GPUs within the same power envelope not only improves performance but also accelerates return on investment.

Designing for the next generation

For organisations planning new data centre builds, the question of future-proofing has become increasingly important. Facilities expected to come online within the next few years must be designed with emerging AI requirements in mind.

In practical terms, this means planning for 800V DC compatibility and considering how infrastructure can evolve toward fully DC-native

architectures. Hybrid approaches may be used during the transition, but the long-term trajectory is clear.

Looking further ahead, the eventual move from 800V to 1500V DC is expected to be less disruptive than the current transition from AC. Many components are likely to be designed with forward compatibility, enabling a smoother evolution.

A new era for power electronics in data centres

The transformation of data centre power infrastructure represents a shift from incremental optimisation to fundamental redesign. As AI workloads continue to drive demand for higher performance and greater efficiency, the limitations of traditional AC systems are becoming increasingly apparent.

High-voltage DC architectures, supported by innovations such as solid-state transformers and silicon carbide semiconductors, offer a compelling path forward. While challenges remain, the direction of travel is unmistakable.

In the era of AI, power electronics has moved from a supporting role to a central position in data centre design. The ability to deliver energy efficiently will not only determine operational costs but also define the limits of computational capability.

The conversion tax is no longer an acceptable compromise. It is a problem that must be solved - and DC-native architectures are emerging as the solution.





WE-MXGI: High-efficiency power inductor for demanding DC/DC Converters

The WE MXGI series was developed for use in compact, high-power-density DC/DC applications in which high currents, low losses, and stable performance are required even at high switching frequencies.

THE INCREASING complexity of electronic systems places growing demands on power conversion stages. In many applications, DC/DC converters must deliver high output currents while operating efficiently within limited space and under constrained thermal conditions. At the same time, higher switching frequencies are used to reduce the size of passive components and enable more compact layouts. These developments directly affect the requirements for power inductors, which influence efficiency, thermal behavior, electromagnetic compatibility, and achievable power density.

Within this context, the WE-MXGI series was developed as a molded SMD power inductor for modern DC/DC converter designs. Rather than focusing on a single design parameter, the series combines material properties, electrical performance, and mechanical robustness. This reflects the reality that power inductors are rarely optimized for one parameter alone but must operate reliably across a wide range of electrical and environmental conditions.

A central feature of the WE-MXGI design is the use of an iron alloy core material together with an optimized winding geometry. The core material influences both DC and AC losses, which are critical factors for efficiency and temperature rise. Low DC resistance reduces conduction losses at higher load currents, while low AC losses help maintain efficiency



as switching frequency increases. By addressing both mechanisms, the WE-MXGI series supports stable operation over a broad range of operating points.

High switching frequencies create additional challenges for inductor design. As frequency rises, conventional core materials may show increasing losses or shifts in inductance behavior, reducing efficiency and complicating converter design. The

iron alloy material used in the WE-MXGI series is intended to maintain stable electrical characteristics at switching frequencies beyond 1 MHz. This supports higher-frequency converter operation and enables more compact power stages with improved dynamic response.

Current handling capability is another key requirement in many power electronics applications. Even when higher switching frequencies reduce



component size, inductors must still carry the required current without saturation or excessive temperature rise. The WE-MXGI series offers saturation currents of up to 28 A, depending on inductance value and package size. With available inductance values from 0.16 μ H to 10 μ H in 4020 and 5030 SMT packages, the series gives designers flexibility in matching inductor characteristics to specific converter requirements.

This range of inductance and current ratings allows targeted tradeoffs between ripple current, transient response, efficiency, and PCB area. In point-of-load converters, for example, lower inductance values may improve transient performance, while higher values may be preferred where ripple reduction or steady-state efficiency is more important. Because multiple values are available within a consistent package concept, designers can make such adjustments without major layout changes.

Electromagnetic interference is another important consideration, especially in densely populated PCBs. Power inductors can be a source of stray magnetic fields that couple into nearby signal lines or sensitive components. The WE-MXGI series uses a magnetically shielded molded construction to reduce stray field emission. This can simplify PCB layout and support compliance with EMC requirements, particularly in systems where several power stages are placed close together.

The molded construction also improves mechanical stability. In applications exposed to vibration, thermal cycling, or assembly-related stress, a robust inductor design helps maintain consistent electrical performance over time. Although mechanical robustness is often secondary to electrical parameters, it is important for long-term reliability, especially in systems with demanding operating conditions.

Thermal behavior is closely linked to both electrical losses and construction. The WE-MXGI series is specified for an operating temperature range from -40°C to $+125^{\circ}\text{C}$, allowing use in a wide variety of environments. Combined with low DC and AC losses, this supports predictable behavior under changing load and ambient



conditions. From a system perspective, lower inductor losses can reduce the need for additional cooling measures and increase flexibility in component placement.

The application range of the WE-MXGI series reflects these characteristics. Typical use cases include DC/DC converters for high-current power supplies, point-of-load converters, and power management circuits for FPGAs and processors. The series is also suited for mainboards and graphics cards, where space constraints and thermal demands are especially relevant.

Beyond computing-related applications, WE-MXGI inductors are also used in battery-powered devices, wireless communication equipment, and power supplies for mobile electronics such as smartphones and tablets. In these systems, efficiency directly affects battery life or thermal performance, while compact size supports higher integration. The wide operating temperature range also extends

suitability to industrial and embedded applications.

From a design methodology perspective, the WE-MXGI series supports scalability. The availability of multiple inductance values within a limited set of package sizes allows designers to adapt power stages to different load requirements while maintaining a consistent layout strategy. This can simplify design updates, product variants, and qualification processes.

Rather than presenting the WE-MXGI as a universal solution, it is more accurate to view it as part of a broader inductor portfolio. Different applications may prioritize minimum footprint, maximum current capability, or specific thermal characteristics. Within the Würth Elektronik portfolio, other series such as WE-MAPI, WE-PD, and WE-XHMI address additional design priorities. Together, these series give designers a structured set of options for selecting power inductors according to technical requirements.

In summary, the WE-MXGI series is a power inductor family designed to meet the combined demands of high switching frequency operation, low losses, compact packaging, and reliable mechanical construction. Its characteristics support a broad range of DC/DC converter applications and help designers balance efficiency, power density, thermal performance, and layout constraints in a controlled manner.

The molded construction also improves mechanical stability. In applications exposed to vibration, thermal cycling, or assembly-related stress, a robust inductor design helps maintain consistent electrical performance over time

<https://www.we-online.com/we-mxgi>

How to select a shunt resistor for high bandwidth current measurements

Selecting the right shunt resistor is critical to achieving accurate, high bandwidth current measurements. From power dissipation and insertion inductance to bandwidth optimisation and compensation techniques, careful shunt selection can dramatically improve signal integrity, dynamic range, and measurement reliability

BY SEAMUS BROKAW, APPLICATION ENGINEER, TEKTRONIX

SHUNT RESISTORS are necessary for many applications requiring higher bandwidth and lower inductance current measurements than a Rogowski or clamp-style current probe can offer. Understanding the resistor's properties and careful selection from the wide variety of available shunt resistors will ensure accurate, wide-bandwidth current measurements that minimize parasitic inductance and maximize Signal-to-Noise Ratio.

The selection process involves several interconnected considerations:

- Target current range and dynamic range requirements
- Power dissipation limits and thermal management
- Resistance value optimization for signal amplitude and burden voltage
- Form factor selection based on space, power, and bandwidth requirements
- Bandwidth limitations due to parasitic inductance
- Test point implementation strategies

Current range identification

Target Current Analysis

The target current RMS value, peak current, and ideal bandwidth are the first place to start exploring shunt options. These factors will all affect the downstream variables like power dissipation of the shunt, burden Voltage, the resistor value, insertion inductance, and compensation strategy.

In the ideal case, a high value resistor with near-zero insertion inductance has infinite power capability, leading to perfect signal:noise ratio, several GHz bandwidth, and no impact to the device-under-test. In reality, compromises must be made for the shunt's power capacity and the DUT's burden Voltage and inductance sensitivity.

For example, identify an input current that averages 8 A RMS with occasional spikes to 12 A and a 40 mA signal-of-interest riding on top of the 8 A. These

values will factor into the shunt selection, as we will see later.

Power limits and thermal considerations

With a known target current, browse for current shunts in values that would lead to a power dissipation below 3 W for larger CVRs and 1 W being a safer limit for SMD devices. This power limit will determine the resistance values that you can select from, with higher resistances providing more signal to measure and a lower noise floor at the tradeoff of more power dissipation and a higher burden voltage dropped across the resistor.

Shunt Value to Measure <i>I</i> _{max} (Amps) (Ω)	Noise Floor at BW (A RMS)	Dynamic Range (dB)	Power Dissipation at 100% duty cycle(W)	<i>I</i> _{max} (A)	Bandwidth (Hz)
81.3E-3	3.1E-3	68.3	5.2E+0	8	120.0E+6
81.3E-3	2.3E-3	70.7	5.2E+0		
62.5E-3	2.6E-3	69.7	4.0E+0		
43.8E-3	2.4E-3	70.5	2.8E+0		
31.3E-3	3.0E-3	68.4	2.0E+0		
22.5E-3	3.1E-3	68.3	1.4E+0		
16.3E-3	3.7E-3	66.7	1.0E+0		
11.3E-3	4.6E-3	64.9	720.0E-3		
7.5E-3	6.9E-3	61.3	480.0E-3		
5.0E-3	10.3E-3	57.8	320.0E-3		


➤ This table shows how different shunt values (Column A) can make the same 8 A measurement with varying Power levels (Column D). Higher resistance values lower the noise floor (Column B) at the cost of more power dissipation in the shunt.

For example, an 8 A RMS signal with a 40 mA signal of interest riding on top of it will require that the resistor be able to handle $P=I^2R$ power dissipation in its body. With a 1 W shunt, the maximum resistance value will be $R = 1 W / (8 A)^2 = 15.6 m\Omega$. Sourcing a 3 W shunt raises this limit to 46.9 m Ω and will generate more signal to measure, raising Signal:Noise ratio and making that 40 mA signal-of-interest come further out of the noise floor.

Pulsed vs. continuous current



With a signal that is periodic, rather than continuous, significantly higher current levels can be achieved even with a nominally 1 W resistor. The 5 m Ω Wideband Shunt from Tektronix can reach 200 A square or triangle pulse for a 10 μs pulsewidth where it'd normally be limited to 14 A RMS for a continuous signal.

Form Factor Analysis



SMD resistors

SMD Resistors come in a wide variety of form factors and Ohm values, leading them to be a flexible and affordable option for many applications. The



Shunt Type	Resistance Range	Power Rating	Key Applications
SMD Thick/Thin Film	50 $\mu\Omega$ –10 Ω	0.25 W–2 W	Compact designs, pulsed current
Bus Bar/Metal Strip	50 $\mu\Omega$ –10 m Ω	10–100 W	High RMS current, thermal stability
Wideband Shunts	5 m Ω to 5 Ω	1 W	High bandwidth, precision compensation
Ultra Fast Current Shunts	1 m Ω –52 m Ω	3 W	WBG testing on devices sensitive to insertion inductance

tradeoffs of SMD devices is their limited power dissipation capabilities and uncompensated frequency response.

➤ Shunt Type Comparison

Inductance for SMD resistors can be low to medium, often landing in the few nano-Henry range. The form factor has flexible options with the caveat that as the width of the device increases relative to the length, the resistor starts to look more like a distributed transmission line and less like a lumped element. Impedances vary based on the probing position and therefore currents start being distributed unevenly across the SMD resistor.


Bus bar types

Bus bars have a larger size than SMD devices and therefore can dissipate much more heat. This can be a big advantage when in a high-RMS current measurement such as a 100 A continuous DC current with a smaller, 1 A transient of interest riding on top. The resistor needs to dissipate the power created by the full RMS current, not just the smaller signal of interest. These devices also are made of materials that have stable resistance over temperature and can have inductance values in the 10s of nH range, barring the same transimpedance issues that start to appear on large geometry resistors, similar to the larger SMD resistor devices.

Their larger size means the designer often needs to think ahead and design in these resistors to the PCB layout. If your traction inverter or motor drive application already uses current sense devices for over-current protection or short-circuit detection, building in the bus bar shunt and adding an additional test point for a high performance current measurement is much more feasible than an impromptu test point with the same shunt.

Wideband shunts



One shunt designed to work with the Tektronix Isolated Current Probe and MSO oscilloscopes are the Wideband Shunts. This design interfaces with square pins on the board and features frequency compensation, temperature compensation, and high-performance fuses to protect the test equipment. Values range from 5 m Ω to 5 Ω , providing 250 MHz of bandwidth and easy implementation onto the board.



Note for Wide Bandgap devices, the square pin Wideband shunts add several nH of insertion inductance. If placed in the current loop, that could affect the device performance and make measurements on the lab bench differ from simulation or production devices.

Ultrafast current shunts (and other new entries)

New current shunts from test and measurement companies are launching every year to improve on off-the-shelf resistor performance. One is the




Ultrafast Current Shunt (UFCS) from PMK. In testing, the UFCS shows several-hundred MHz bandwidth with a desirable rolloff, indicating it is a frequency-compensated design. It performs well with low insertion inductance and somewhat high 3 W power dissipation capability. The challenges with using the PMK device come from its large form factor. The low insertion inductance requires a wide space on the test board and the shunt is also tall, presenting positioning challenges in crowded enclosures.

This device is well-suited for Wide Bandgap current measurements where the insertion inductance is paramount and compromises can be made to fit the shunt onto the test board. It's cost is more in line with a measurement device than a mass-produced SMD component, so it's also more suited for R&D lab benches than production tests.

Device	Insertion Inductance
Picotest	<100 pH datasheet
PMK Mess- & Kommunikationstechnik GmbH Ultra-Fast Current Shunt UFCS	<150 pH datasheet
Surface Mount Shunt	<1 nH empirical
Shorted Square pins	3 nH empirical
T&M Research Products Coaxial Shunt SDN-414-01	>5 nH (depends on lead length) empirical
Tektronix Wideband Shunts	7 to 8 nH empirical + datasheet
Clamp Probe wire loop	>80 nH empirical
6" twisted pair wire	123 nH empirical

Bandwidth and compensation

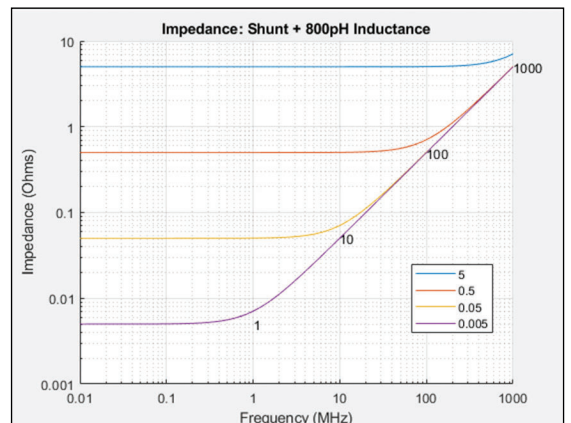
At a first order level, a shunt resistor can be modeled as a single pole RL circuit: a resistor and inductor in series. The inductance is caused by parasitic design of the shunt resistor and has a large effect on the effective bandwidth of the current measurement. This basic model ignores transimpedance and coupling effects which can be significant but are very difficult to characterize and model correctly.



Bandwidth calculation

$$f_c = \frac{R}{2\pi L}$$

Corner frequency is correlated with R and inversely correlated with L. That shows how the resistor value selection and form factor will greatly impact the measurement bandwidth. Here is a simulation showing the effect of resistance value and 800 pH of inductance on the frequency response. The



➤ Uncompensated Frequency Response

corner shifts with increased Resistance.

Compensation network design

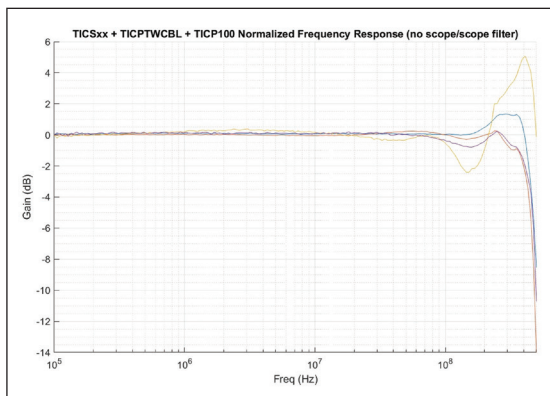
The parasitic inductance can be compensated using an RC network where:

$$R_{comp} * L_{comp} = \frac{L_{parasitic}}{R_{shunt}}$$

After compensation, frequency response can be flattened and extended for an order of magnitude or more! Notice also how the addition of a steep rolloff element can change the deceiving high frequency amplification in the original chart to a more desirable HF rolloff.

One important note is that although the measurement inductance is improved through compensation, the DUT is still experiencing the same insertion inductance caused by the shunt. It is only in the measurement that the effects of the inductance are compensated.

Resistor value selection



➤ Compensated Frequency Response

Design Trade-offs

The resistance value selection involves balancing several competing factors:

- **Power Dissipation:** Higher resistance increases power dissipation in the shunt ($P = I^2 \cdot R$)
- **Burden Voltage:** Higher resistance increases voltage drop across the shunt, lowering final rail Voltage experienced by the DUT. This factor can matter on low Voltage processors that have a 5% Power Rail tolerance. (e.g. ± 40 mV tolerance on a 0.8 V core Voltage power rail)
- **Signal Amplitude:** Higher resistance provides larger signal amplitude, raising signal:noise ratio
- **Measurement Bandwidth:** Higher resistance and lower insertion inductance will both improve measurement bandwidth.

$$f_c = \frac{R}{2\pi L}$$

Selection Methodology

Now familiarized with the tradeoffs associated with different current shunts, select a device based on

what is mechanically possible and then optimize for measurement bandwidth and device insertion inductance.

Calculate Maximum Acceptable Burden Voltage

With good decoupling capacitor selections, the burden Voltage imposed by the shunt will be limited by the average, or RMS, input current.

$$R_{shunt} = \frac{V_{burden(max)}}{I_{RMS}}$$

Example: a 40 mV Voltage budget with 8 A RMS average current allows a 5 mΩ resistor.

Calculate Power Dissipation

The next limiting factor is power dissipation in the shunt.

$$R_{shunt} = \frac{P_{max}}{I^2}$$

Example: The 5 mΩ resistor experiencing 8 A RMS current needs to dissipate 1/3 W heat continuously. That Power value is easy to find in SMD devices and trivial for larger form factor bus bars, wideband shunts, and other dedicated current viewing resistors.

Verify Signal Amplitude and Check Noise Floor

After finding some devices that will survive the current measurement, the next task is to optimize the measurement for dynamic range, bandwidth, and insertion inductance.

Dynamic range is the ratio of the largest Current that can be measured to the noise floor of the measurement device, in dB.

$$DynamicRange(dB) = 20 \cdot \log\left(\frac{I_{max}}{I_{noisefloor}}\right)$$

Example: The same 8 A RMS current has occasional impulses up to 12 A which need to be measured. On a 5 mΩ shunt, that creates a 60 mV measurement. Using a TICP probe in its ± 45 mV range allows this measurement when at least 15 mV of offset is applied. This TICP input range with a 5 mΩ shunt results in a 13.3 mA RMS noise floor.

$$DynamicRange(dB) = 20 \cdot \log\left(\frac{12 A}{13.3 \cdot 10^{-3} A}\right) = 59 dB$$

If this same measurement were performed on a 1 mΩ shunt, the noise floor raises to 66 mA RMS and Dynamic Range falls to 45 dB.

Check Measurement Bandwidth

Resistor value selection affects dynamic range and it also affects another critical measurement specification: bandwidth.

As shown in section Bandwidth and Compensation, both resistance and insertion inductance impact the measurement bandwidth. Compensation helps but the effectiveness is limited by the signal coming off of the shunt in the first place.

$$f_c = \frac{R_{shunt}}{2 \cdot \pi \cdot L_{shunt}}$$

For these reasons, a resistor with the largest value that is possible is always recommended for current shunt measurements. Also for bandwidth, L_{shunt} of the device should be minimized. This can be accomplished through resistor selection, PCB layout, and test point strategies.

SMD resistors by themselves could have several nanoHenries of inductance but try putting four to eight in parallel to drop the inductance. This approach has limits as current stops being distributed evenly at wider form factors, but it does work to lower the insertion inductance.

Newer shunts, designed to optimize insertion inductance, have accomplished more and some designs also include compensation networks to further tune the measurement frequency response.

Shunt Selection Summary

Choose the highest value Resistor that will survive the RMS current while being minimally invasive to the DUT. Then choose a resistor that is designed for lowest insertion inductance.

Test point implementation

Ideally the shunt resistor, compensation network, and test points are all built into the board design. In the second best case, a shunt resistor already exists on the board. And finally, there are cases where no measurement point exists and the entire chain needs to be added.

Implementation Strategies

- [1] **Integrated PCB Design:** Best performance, requires early planning
- [2] **Existing Shunt Modification:** Quick implementation, may degrade signal integrity.
- [3] **Complete Addition:** Maximum flexibility, highest complexity

Connector choices


In addition to the shunt location and compensation, the physical connection to the device will also impact measurements through their sensitivity to radiated emissions. In the examples above, MMCX, square pins, a coaxial cable with MMCX connector, and twisted pair with square pins were shown. These different methods have different levels of radiated emissions sensitivity and care should be taken to choose the best connection that is realistic and possible for each measurement.

Shielded test points using MMCX, SMA, and coaxial cables will perform the best, as the outer braid provides effective shielding to the center conductor. Even better than coaxial is a twinax cable as used with the Tektronix Wideband shunts. Twinax cables put both conductors inside a third, braided shield—protecting both positive and negative conductors


from EMI with the added benefit of true differential loading down at the DUT.

Conclusion


Selecting the appropriate shunt resistor for high bandwidth current measurements requires careful consideration of multiple factors including current range, power dissipation, resistance value, form factor, bandwidth limitations, and test point



➤ [1] On this board, the I_s _Lower testpoint at MMCX connector J3 is designed into the board, complete with compensating pole at C1, R32, and R20.



➤ [2] An existing 0.1 Ω SMD shunt has a MMCX cable soldered across, allowing a TICP current measurement. This implementation does not include a compensation circuit for the parasitic inductance which amplifies high frequency ringing.



➤ [3] After cutting the current carrying trace, a twisted pair wire is added in series with the current up to a newly added shunt. This design is convenient but adds the most insertion inductance into the current measurement path, impacting the device behavior.

implementation. The selection process involves balancing competing requirements to achieve optimal performance for the specific application.

Key takeaways:

- Current range and dynamic range requirements drive initial selection criteria
- Power dissipation limits constrain resistance value selection
- Form factor selection affects bandwidth, power handling, and cost
- Parasitic inductance limits bandwidth, affects the DUT, and may require compensation
- Test point implementation strategy affects measurement accuracy and repeatability

REFERENCES

- Tektronix TICS Shunt Resistor Datasheet
- T&M Research Current Viewing Resistor Specifications
- PMK UFCS Datasheet
- Signal Integrity Journal Articles from Steve Sandler
Picotest Transimpedance



SAVE THE DATE
12-14 APRIL
2027

**Sheraton Brussels
 Airport Hotel**

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

info@csinternational.net
info@picinternational.net
info@peinternational.net
info@angel-tech.net

Driving efficiency through system-level optimization in next-generation power electronic systems

Future efficiency gains in power electronics require holistic system optimization across electrical, thermal, mechanical, and control domains, not isolated components.

BY DR. FABIAN HOHMANN, POWER ELECTRONIC SYSTEMS, SYSTEMS ENGINEERING ONSEMI

GLOBAL megatrends such as climate-change mitigation, electrification of mobility, expansion of data centers, and the integration of renewable energy sources are driving an explosive growth in electrical power demand. Across these domains, power electronic systems serve as the essential interface between generation, storage, and consumption. Consequently, efficiency and power density have become the most critical key performance indicators (KPIs) for next-generation power electronics.

Historical development shows a continuous increase in system-level power density of approximately 10–20 % every two years (see Figure 1). Looking ahead, this trend is expected to accelerate further, driven by the rapidly growing energy demand of hyperscale data centers and the increasing need

for smaller, lighter, and more efficient systems in traction, aerospace, and emerging application domains such as neurotechnology. As a result, many roadmaps now assume a transition from linear to exponential scaling, with an effective doubling of system-level power density approximately every two years.

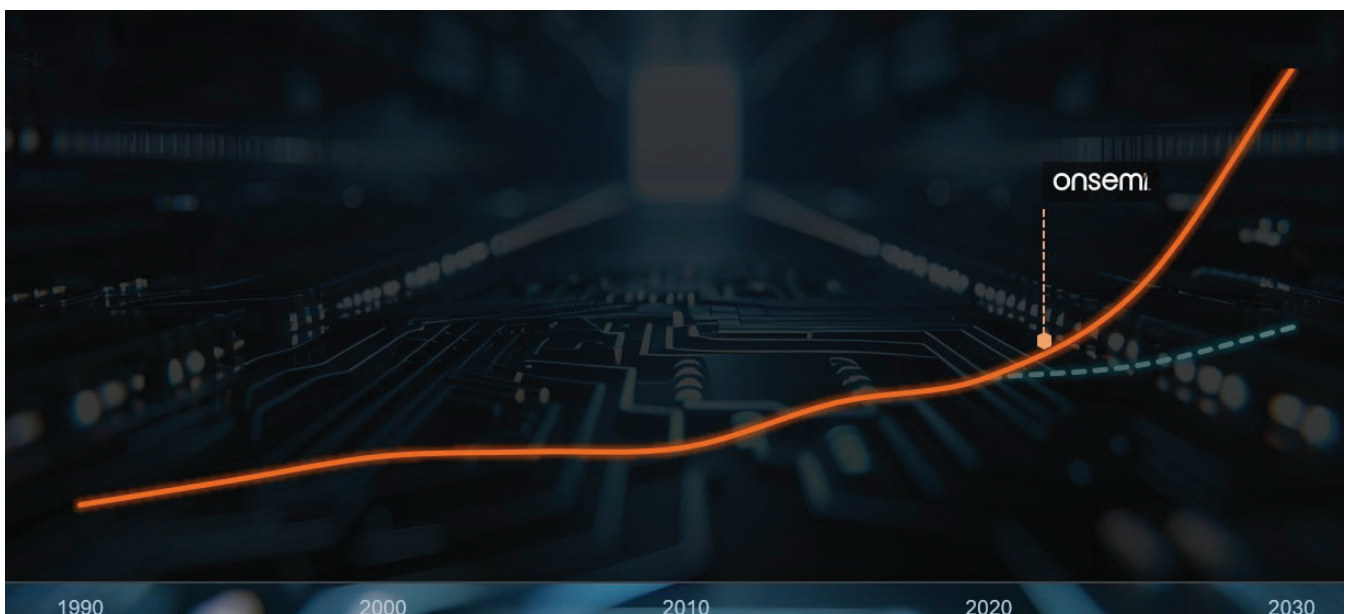
Modern power electronic systems comprise a tightly coupled network of power semiconductors, passive components, cooling solutions, control software, and mechanical integration. Optimizing these elements independently typically leads to local optima but fails to unlock the full performance potential of the overall system. This article outlines a system-level optimization methodology for power electronic systems and illustrates why breaking down traditional silos

between design disciplines is essential to achieve the efficiency gains required for future applications.

Drivers for system-level optimization

Electrification and power-demand growth

Electrified vehicles, heat pumps, and renewable-energy integration significantly increase local and global power demand. In parallel, data centers already consume several gigawatts of electrical power in individual regions, with demand rising sharply due to artificial-intelligence workloads. Even marginal efficiency improvements in conversion stages translate into substantial energy savings, reduced cooling requirements, and lower operational costs when deployed at scale.



➤ Figure 1: Historical and future power density in power electronic systems.

At the same time, regulatory pressure and sustainability targets increasingly demand higher overall system efficiency rather than isolated component improvements. This shifts the optimization focus from individual devices toward holistic energy flow across the system.

Limits of component-level optimization

Wide-bandgap devices such as SiC MOSFETs and GaN HEMTs provide lower switching losses and allow higher operating frequencies compared with silicon devices. Nevertheless, their full potential cannot be realized if parasitic inductances, thermal bottlenecks, or overly conservative control strategies dominate system behavior. For example, reducing transistor switching losses may expose the DC-link capacitor, gate-driver dynamics, or cooling system as new limiting factors.

In such cases, optimization confined to a single component merely shifts losses within the system rather than reducing them globally. As power density increases, these interactions become increasingly critical, reinforcing the need for system-level design methodologies.

From component silos to system thinking

Traditional power-electronic design follows a hierarchical structure: semiconductor selection, packaging, cooling, passives, and control are treated as largely independent design steps. This “silo” approach simplifies project organization and responsibility allocation but inherently restricts achievable system performance.

In contrast, system-level optimization treats the power-electronic converter as an integrated electromechanical system. Electrical waveforms, thermal behavior, mechanical constraints, and control dynamics are evaluated simultaneously. This enables informed trade-offs between competing objectives such as switching speed, electromagnetic interference (EMI), thermal cycling, efficiency at partial load, and long-term reliability. Breaking down functional silos does not imply abandoning domain expertise. Instead, it requires structured interfaces and shared metrics that allow each discipline to contribute toward a common system objective.



➤ Figure 2: System of inverter and electrical machine

Efficiency gains through system optimization

Automotive traction inverters as a case study

In electric vehicles, drivetrain efficiency is strongly determined by the interaction between the traction inverter and the electric machine. Optimizing both jointly yields significantly higher efficiency gains than improving the inverter alone, resulting in a driving range increase that is more than four times greater than what can be achieved through inverter-only optimization. This requires deliberately operating neither component at its individual optimum but instead aligning both toward the global optimum of the combined system.

The optimal operating point emerges from a balanced trade-off between inverter switching behavior and the loss characteristics of the electric machine. A system-level perspective ensures that losses are minimized across both components simultaneously, rather than optimizing them in isolation.

Coordinated electrical and thermal design

While the example in the last caption only considers two subsystems, a complete vehicle offers significantly more degrees of freedom for system-level optimization (see Figure 3). Only by optimizing all relevant subsystems in an integrated manner can the full system performance be realized. Higher switching frequencies reduce passive-component size and improve dynamic performance but simultaneously increase switching losses and thermal stress. Conversely, oversized cooling systems reduce junction temperatures

but add volume, mass, and cost, reducing overall power density.

System-level optimization quantifies these trade-offs by jointly modeling electrical losses and thermal resistances. For example, a power module with Figure 3: Electric vehicles represent a near-ideal thermal coupling to its cooler enables more multidisciplinary system-level optimization problem. aggressive switching strategies without violating junction-temperature limits. However, the cooling system must be designed to handle transient heat flux rather than relying solely on conservative steady-state assumptions. By treating electrical and thermal domains as interdependent rather than sequential, designers can avoid overly conservative design margins while maintaining reliability.

Control and software as system-level enablers

Control algorithms and embedded software increasingly define system efficiency and robustness. Adaptive modulation schemes, dynamic gate-driver control, and operating-point-dependent switching strategies allow converters to maintain high efficiency across wide load ranges.

In a system-optimized design, control software is not an afterthought but a core optimization variable. Software can compensate for hardware limitations, mitigate parasitic effects, and enable operation closer to physical limits when supported by accurate sensing and thermal models. The integration of software into system-level optimization also enables future updates and performance improvements without hardware changes.

DC-Link, magnetics, and passive components

Passive components are frequently the hidden bottleneck in high-power-density systems. DC-link capacitors, bus structures, and magnetic components strongly influence switching behavior, EMI, and reliability. Their electrical performance cannot be separated from mechanical placement and thermal coupling.

System-level optimization explicitly considers the interaction between switching transients, parasitic inductances, and capacitor technology. This enables reduction of voltage overshoot, improved dv/dt control, and lower overall loss. Similarly, optimized magnetic design balances copper loss, core loss, thermal dissipation, and acoustic noise within a unified framework.

Multidisciplinary design and parallel development

One major obstacle to system-level optimization is organizational rather than technical. Electrical, mechanical, and software engineers often operate sequentially, leading to long development cycles and suboptimal integration.

Future power electronic systems demand parallel, multidisciplinary development where design iterations span the entire stack simultaneously. This approach shortens design cycles, exposes system-level trade-offs early, and avoids costly late-stage redesigns. It also enables continuous optimization rather than fixed design freezes between stages.

Role of advanced modeling, data, and validation

Accurate system-level optimization relies on high-fidelity models that capture interactions between components. This includes:

- Electrical models with realistic parasitic elements and switching behavior
- Thermal models linking junction temperature to cooling architecture and ambient conditions
- Control models reflecting real-time constraints, sensor accuracy, and communication delays

Increasingly, data-driven methods and artificial-intelligence techniques support design-space exploration and sensitivity analysis. However, these tools

complement rather than replace physical insight. Meaningful optimization requires identifying which parameters materially affect system performance and which variations are secondary.

Equally important is experimental validation. System-level optimization must be verified through targeted measurements that confirm predicted loss distributions, thermal behavior, and dynamic performance under real operating conditions.

Innovation beyond incremental improvement

Historical breakthroughs often arise not from incremental refinement but from rethinking fundamental assumptions. The transition from silicon to wide-bandgap semiconductors exemplifies such a paradigm shift. Likewise, future efficiency leaps may result from architectural changes—such as higher DC-link voltages, integrated drive units, advanced cooling concepts, or software-defined power stages—enabled by holistic system design.

Leadership in power electronics therefore requires fostering a culture that encourages cross-disciplinary thinking, questioning established practices, and embracing controlled risk in pursuit of disruptive improvements.

Implications for industry and research

For industry, system-level optimization

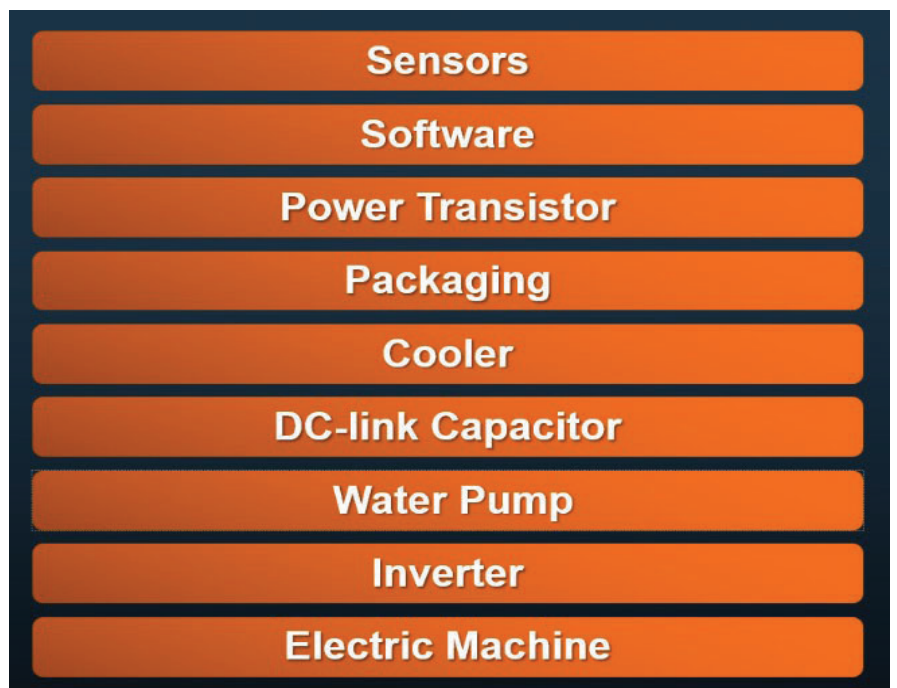
translates directly into competitive advantage through higher efficiency, reduced material usage, improved reliability, and shorter time-to-market. For academia and research institutions, it highlights the need for integrated curricula and research programs that bridge traditional disciplinary boundaries.

Close collaboration between semiconductor manufacturers, system integrators, and end users becomes increasingly important, as many optimization potentials span organizational interfaces rather than individual components.


Conclusion

The next generation of power electronic systems will be defined not by isolated component advances but by the effectiveness of system-level optimization. Efficiency and power-density targets driven by electrification and sustainability goals can only be met by coordinated co-design of semiconductors, passives, cooling systems, and control software.

By removing silo barriers, adopting multidisciplinary development approaches, and leveraging advanced modeling and data-driven validation, the power-electronics community can continue to drive innovation and enable the electrified future.



➤ Figure 3: System of inverter and electrical machine



Japan's power semi giants move toward landmark consolidation

Can Mitsubishi Electric, Rohm and Toshiba's potential merger reshape Japan's position in global power electronics?

BY REBECCA POOL, TECHNOLOGY EDITOR

AS Rohm, Toshiba, and Mitsubishi Electric continue discussions to merge their power semiconductor businesses, industry can only wait and watch for what could be the most significant shake-up in Japan's power semiconductor sector for decades. A consolidation of these industry giants would create the world's second-largest power chip group - trailing only Germany's Infineon - and by pooling resources and manufacturing scale, the combined entity would be set to compete aggressively in the global market.

As the firms signed a memorandum of understanding to 'initiate discussions' earlier this year, Toshiba emphasised - in a company statement - the need to ensure a secure and stable supply of semiconductors in line with Japanese Government's plans. But what lies beneath the company rhetoric?

First and foremost is the need to revitalize Japan's fragmented power

semiconductor market amidst a backdrop of the West's technological strength and mounting price pressure from China. As John Li, Senior Analyst at UK-headquartered Omdia highlights: "The power semiconductor ecosystem in Japan is acknowledged to be quite fragmented compared to the US and Europe, where you have these behemoths such as OnSemi, STMicro and Infineon."

"Collaboration and mergers have been discussed for quite a while now, especially with the cost decreases in wide bandgap semiconductors - particularly silicon carbide - in China," he says. "Japan wants to optimise production efficiency, increase scale of manufacturing and allocate resources to high growth areas - and is probably trying to do this sooner rather than later, and not fall further behind."

Recent analysis from Omdia's 'Competitive Landscaping Tools, March

2026', ranked the combined power semiconductor businesses of Mitsubishi Electric, Rohm and Toshiba as second in the power device market with a 11.3% market share, whilst individually the firms hold 4th, 8th and 9th place, respectively. Infineon currently holds first place with a 24.4% market share. "A merger would also place them ahead of OnSemi and ST Micro," points out Li.

Economic security is also key, and for Japan, this is inextricably entwined with semiconductor resilience. The nation was hit particularly hard by the global supply chain disruptions and semiconductor shortages of 2020 - intensified by the Covid-19 pandemic - prompting Japanese government to label chip production a matter of strategic national importance. As Li's colleague and Omdia Research Director, Paul Pickering, puts it: "We all remember that time when \$50,000 cars were stuck because we couldn't get a \$3 part as

factories had closed – no-one wants to put themselves in that situation again.”

Japan is also keen to future-proof itself against rising overseas competition, especially from China. The Chinese government has channelled huge resources into developing silicon carbide manufacturing, from boules through to devices, whilst doggedly prioritising domestic markets.

Given Japan’s global automotive industry might, and the rising demand for SiC semiconductors in electric vehicles, Japan, as a nation, would also likely prefer a stable supply of the devices without relying on China. “SiC is important for driving domestic strength,” adds Li.

According to Pickering, he and colleagues have also long been tracking China’s power electronics strategy, noting a gradual increase in the share of Chinese suppliers within the Chinese market itself – which has profound implications for the rest of the industry. “If you take a look at Rohm’s semiconductor sales, Japan accounts for a third of sales, less than 10% comes from both EMEA and the Americas, and the remaining [proportion] of sales are in China,” he says. “So I would imagine that Rohm is now thinking that’s nearly a third of my sales that are a prime target for getting chopped away, especially as Chinese

companies issue more of an aggressive posture.”

Pickering also questions if the likes of Infineon, as well as Mitsubishi Electric, Rohm and Toshiba, have concerns over China’s next move. “Could China be coming to Europe and Japan next?” he notes. “[These companies] will all want economies of scale in order to compete with that looming threat from China.”

On the flipside, the Japanese market has tended to favour domestic suppliers where possible, and, as Pickering notes, a merger of key power semiconductor firms would deplete supply sources. “Power semiconductors are commodity products which means customers are used to sourcing form-, fit- and function-compatible devices from multiple suppliers,” he says. “So the likes of Toyota and Nissan might now be thinking, ‘we had three power discrete companies competing for our business and now we only have one – do we now have a sole source situation?’”

“I don’t think Japan-based [companies] would want their supply to come from a single Japanese company, especially if the next competitors are, worst case, Chinese companies,” he adds.


Still, while much of the merger’s strategic focus looks to be centered on wide bandgap technologies, especially with the SiC competitive threat from China,

silicon looks set to remain the backbone of the new power semiconductor business. “Silicon is probably going to remain the major player... [it] does the job and there’s a huge number of suppliers,” says Pickering.

Beyond Japan

The proposed merger from Mitsubishi Electric, Rohm and Toshiba comes at a time of extensive semiconductor investment, and change, in Japan. TSMC recently unveiled plans to mass-produce 3 nm chips in a reported \$17 billion investment in its second Japan-based fab. The plant will eventually churn out some 15,000 12-inch wafers, via its 3 nm process, every month – well-and-truly strengthening Japan’s domestic supply chain. The Taiwan-based chip-making giant launched ‘Japan Advanced Semiconductor Manufacturing’ in December 2021, which now has support from Sony Semiconductor Solutions, Japan-based automotive firms, Denso and Toyota, and the Japanese government – its first plant opened in December 2024.

Meanwhile, in August 2022, a consortium of key Japanese firms, including Toyota, Sony, Denso and Softbank, launched the semiconductor foundry, Rapidus, with significant government subsidies. The organisation will focus on fabricating cutting-edge logic semiconductors, in collaboration with global partners, including IBM and



imec, signifying Japan's clear intention to be a global frontrunner in advanced semiconductors. Additional funds have since come from more than 30 companies, including Canon, Fujitsu, NTT, Honda Motor and Seiko Epson.

According to Pickering, these industry developments accompany tremendous geopolitical change. "People have now realised that it's all very well having a global semiconductor supply chain, where, say, a lot of the fabrication takes place in Taiwan, and the test and assembly in Malaysia, bringing low costs," he says. "But if, for example, you look at the Chinese rumblings in the South China Sea and the threats to Taiwan [that could affect production and shipping routes], people realise that this is also a very fragile system."

The Omdia Research Director points to the US Chips Act, the European Chips Act - now Act 2 - and China's Big Fund, all designed to secure domestic manufacturing and fuel technological sovereignty, especially with rising demands from AI. "This entire regionalization definitely wasn't such a concern five years ago, and I think Japan is seeing the same landscape and thinking, 'we need to do something here,'" he says.

But what will the likes of Infineon, Onsemi and STMicroelectronics make of the stronger competitive force that Mitsubishi, Rohm and Toshiba would bring? Both Pickering and Li believe that all companies are more focused on the rise of China.

"While it's true that it may be concerning for Infineon, for example, that in some of its markets there's this new number two, I think they'll be more concerned about China-based companies," says Pickering.

"Onsemi and ST Micro have a big play in discrete and modules, and are also focused on silicon carbide with their revenues increasing here – but currently, I don't think they'll be too concerned over a merger either," adds Li.

So what now for the power electronics businesses of Mitsubishi Electric, Rohm and Toshiba? Undoubtedly, Rohm's strong SiC presence, Toshiba's expanding wide-bandgap efforts and Mitsubishi's power-module expertise would create an organisation with a broad tech and product range spanning many applications. In a recent presentation Rohm spelt out how the merger could deliver 'significant business synergies' in the burgeoning AI server and data center market,

reduce wafer and assembly costs by consolidating manufacturing sites, and accelerate new product development. "If you can become this sort of 'one-stop shop', that makes you much more competitive on a global scale," says Li.

However, mixed messages have emerged on the merger's current state-of-play. Only a few weeks ago, Denso withdrew its multi-million dollar takeover bid for Rohm to secure a supply of power semiconductors, despite the pair's long-term relationship - leaving industry players speculating the merger would take place sooner rather than later. Mitsubishi Electric's President, Kei Uruma, also said he hoped all parties would reach an agreement as soon as possible when questioned on merger timings at a financial results briefing in April. Yet other industry reports warn of delays in talks, and Rohm CEO, Katsumi Higashi, has since stated: "We want to avoid a situation where too many cooks spoil the broth."

Still, cautious outlooks aside, merger momentum is clearly building. As Li puts it: "There's very many hoops to jump through so I'm not even sure this will all be consolidated by the end of this year – but we already have that memorandum of understanding, so it all looks positive."



Advancing SiC substrate and epitaxial technologies for next-generation power devices

Silicon carbide is rapidly reshaping next-generation power electronics, but the real performance gains depend on advances in substrate and epitaxial material quality that improve yield, efficiency and reliability at scale. As the industry moves toward 200 mm wafers and higher-voltage applications, material engineering is becoming central to the future economics of SiC power devices.

BY GARY RULAND, RAJAN RENGARAJAN, SHANTHI SUBRAMANIAN, QI YANG, ILYA ZWIEBACK
COHERENT CORP

The power electronics inflection point

OVER THE past decade, the SiC power device market has transitioned from a specialty niche to a mainstream enabling technology. The full promise of SiC — lower switching losses, higher operating temperatures, higher breakdown voltages, and smaller form factors¹ — depends on continued advances in substrate and epitaxial quality. The progress in the upstream materials is central to industry success.

Why substrate quality defines device performance

A SiC power device begins its life as a boule of single-crystal 4H-SiC, grown by Physical Vapor Transport (PVT) at temperatures exceeding 2,200°C. The resulting ingot is sliced into substrates that serve as the foundation for epitaxial growth. Every defect present in the substrate — whether a micropipe, threading screw dislocation (TSD), basal plane dislocation (BPD), or stacking fault — has the potential to propagate into the epitaxial layer and

compromise device reliability, yield, and forward voltage stability.

Micropipes, hollow core screw dislocations, are device killers. Once present in SiC at densities of tens per cm², they have been virtually eliminated in commercial substrates through advances in seed preparation, vapor phase control, and thermal gradient engineering. Today's Coherent commercial 150 mm substrates show typical micropipe densities below 0.05 cm⁻². The focus has now shifted to subtler but equally consequential defect classes — BPD and TSD dislocations.

BPDs in Coherent substrates remain the most prevalent structural defects, present at typical densities of 500 cm⁻². During epitaxial growth, some BPDs penetrate from the substrate into the epi in the active device area. Under bias, BPDs dissociate into Shockley Stacking Faults (SSF). As these SSFs grow and expand through the active device area, they act as resistive barriers causing increase in on-resistance and forward voltage and leading to local overheating and device degradation.

In Coherent's production substrates, TSDs are found in typical densities of 300 cm⁻². They generate local regions of high electric field around their cores leading to micro-leakages under reverse bias. The micro-leakages contribute to the overall leakage current of the device and reduce the breakdown voltage.

The transition from 150 mm to 200 mm substrates, now underway across the industry, amplifies these challenges. Larger diameter wafers require tighter control of thermal gradients during growth to avoid lattice curvature, wafer bow and dislocation multiplication. Achieving uniform dislocation density and resistivity across a 200 mm substrate is an



➤ SiC substrate wafer — 150mm SiC Substrate

engineering challenge that requires simultaneous advances in crystal growth hardware, seed preparation, and tight in-growth process control.

Epitaxial growth: where device performance is born

If the substrate defines the defect floor, the epitaxial layer defines device performance. Chemical Vapor Deposition (CVD) of 4H-SiC epitaxial layers on polished substrates enables precise control of doping concentration and layer thickness — parameters that directly determine the breakdown voltage, on-resistance, and switching behavior of the finished device.

For a 1,200 V power MOSFET, the typical drift layer thickness for Coherent epitaxial wafer is ~ 10– μm with a net n-type doping of ~ 1×10^{16} cm⁻³. Typical thickness uniformity is < 3% and doping uniformity < 6%. For 3,300 V devices targeting traction applications, drift layer thicknesses extend to 30–35 μm with correspondingly lower doping levels approaching 3×10^{15} cm⁻³. Maintaining uniformity at these thicknesses while suppressing epitaxial defects, such as in-grown stacking faults and inclusions, requires precise control of growth chemistry, surface preparation, and reactor design. Defect density of < 0.5/cm² are routinely achieved through Epi growth process optimization and improvement in substrate quality.

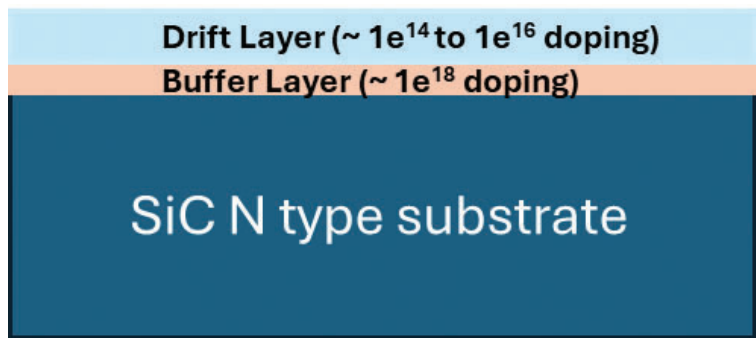
BPD conversion efficiency — the fraction of substrate BPDs that convert to the less-harmful threading edge dislocation (TED) configuration during epitaxial growth — has become a key industry benchmark. Through deliberate miscut angle optimization (typically 4° off-axis toward [11-20]) and controlled off-axis growth conditions, conversion efficiencies above 99% are achievable, dramatically reducing the density of BPDs in the active device layer.

For device manufacturers, these numbers carry direct yield implications. At a hypothetical 1200 V MOSFET die size of 7 mm², an epitaxial defect density of 1 cm⁻² produces a statistical first-pass yield ceiling below 93% before process losses are factored in. Reducing that density to 0.3 cm⁻² pushes the ceiling above 98% — a difference that potentially demarcates profitability at automotive volumes.

The economics and scale-up imperative

Material quality improvements are only meaningful if they can be delivered at the volumes and cost points the market demands. This is where the real competitive differentiation in the SiC supply chain is taking shape.

The automotive sector alone — driven by EV powertrain electrification — is expected to consume hundreds of millions of SiC MOSFETs annually by the end of this decade. Meeting this demand requires not just larger substrates, but higher substrate-to-device yield, and that yield is



predominantly a function of epitaxial defect density with targets now approaching <0.05/cm². This creates a clear strategic imperative: the integration of highthroughput, in-line defect inspection with closed-loop feedback to epitaxial process control. Photoluminescence (PL) mapping, X-ray topography, and automated optical inspection tools capable of classifying defect types across full 150 mm and 200 mm wafers are becoming standard in high-volume manufacturing environments. Machine learning-assisted classification of defect signatures is accelerating the speed at which process deviations are detected and corrected.

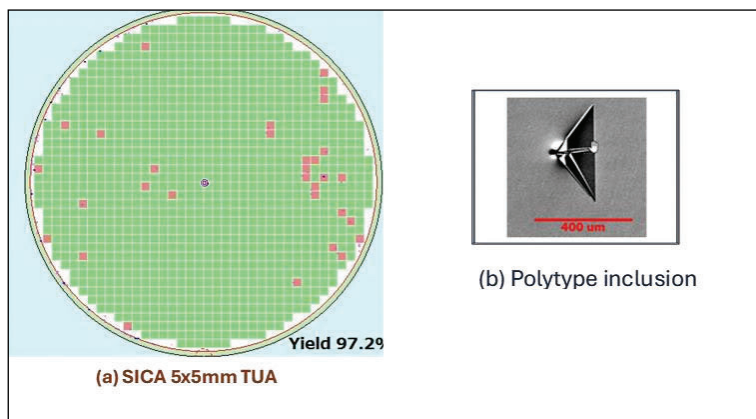
➤ Cross-section schematic of SiC epitaxial layer stack showing substrate, buffer, and drift layer

Simultaneously, the industry is investing in substrate supply security. Vertically integrated supply chains — in which substrate crystal growth, wafering, CMP polishing, epitaxial deposition, and device fabrication are managed within a single organization or through deeply committed supply partnerships — offer advantages in quality feedback loops that are difficult to replicate in fully disaggregated models. The ability to trace a device yield excursion back to a specific boule, growth run, or polishing batch is only possible when data flows freely across the supply chain.

Emerging device structures and their material requirements

Next-generation SiC device architectures are placing even more demanding requirements on substrate and epitaxial quality. SiC trench MOSFETs have been developed to reduce the conduction and switching losses of the traditional planar MOSFET. Trench MOSFET has a lower ON-resistance, which reduces

➤ (a) SICA Defect map: >97% 5mmx5mm Total usable area achieved. (b) SICA image of killer defect : polytype inclusion





► SiC Module

the energy lost as heat and allows for smaller device footprint. A small chip area has low parasitic capacitances, which allows for faster switching.

High-voltage applications above 3,300 V — including solid-state circuit breakers, HVDC conversion, and next-generation traction drives — demand epitaxial layers exceeding 40– 50 μm in thickness with exceedingly tight doping uniformity. At these thicknesses, managing parasitic stress and maintaining step-flow growth morphology across the entire wafer area becomes a significant engineering challenge. Graded buffer layers and strainengineered epitaxial stacks are being explored to manage this challenge without sacrificing the electrical properties of the drift region.

In addition to MOSFETs, there are other SiC devices under development, such as Insulated Gate Bipolar Transistor (IGBT) and Bipolar Junction Transistor (BJT). Fundamentally, operation of both IGBT and BJT is based on injection of the minority carriers. In IGBT, this leads to a very low ON-resistance of the drift region under a high bias up to 20 kV. In BJT, this improves the reliability of the transistor operation under high voltage and at elevated temperatures. SiC BJTs showed the lowest ON-resistance in a broad temperature range. Due to the variety of technological, material and device issues, neither IGBTs nor BJTs entered the commercial SiC device mainstream.

Material requirements for IGBTs and BJTs are different than those applied to MOFETs. IGBT requires precisely doped ($\sim 2 \times 10^{14} \text{ cm}^{-3}$) ultra-thick (up to 200 μm) epilayer having high carrier lifetime (5– 20 μs). The most common n-channel IGBT design requires a *p*-type substrate, which is a challenge. An alternative IGBT design utilizes an *n*-type substrate on which a thick *n*-type drift and *p*-type “cap” epilayers are grown. This is also a challenge. IGBT design may require carbon ion implantation to annihilate carbon vacancies, which are lifetime killers.

BJT design requires precise, heavily doped multiple epilayers with abrupt doping profiles for transistor base, emitter and collector. The lifetime in the BJT

drift layer must be tightly controlled. Both IGBT and BJT are susceptible to bipolar degradation caused by BPDs.

In order to bring IGBT and BJT into the commercial SiC device mainstream, significant progress must be achieved to solve the related material and device design issues.

Looking ahead: the path to 200 mm and beyond

The transition to 200 mm SiC substrates is the defining challenge and opportunity for the materials community over the next three to five years. At 200 mm, the cost per unit area of SiC substrate decreases substantially, and the number of dies per wafer increases by approximately 78% compared to 150 mm — a step change in economics that will be critical to enabling SiC’s penetration into mid-range EV platforms and industrial applications where cost sensitivity is higher.

Achieving this transition at production scale requires overcoming wafer bow and warp issues that emerge at larger diameters, developing CMP processes that maintain surface roughness specifications ($R_a < 0.5 \text{ nm}$) across the full 200 mm area, and qualifying epitaxial reactor platforms capable of delivering the same defect density and uniformity performance achieved on 150 mm wafers. Early production results from leading substrate suppliers are encouraging, but achieving the consistent, high-yield 200 mm substrate supply that device manufacturers will require is still an active engineering effort.

Beyond 200 mm, the long-term roadmap envisions continued improvements in dislocation density, the broader adoption of nitrogen and aluminum co-doping strategies to improve substrate conductivity uniformity, and deeper integration of AI-driven process control across the entire substrate and epitaxial manufacturing sequence. The physics of SiC crystal growth will continue to set limits, but the engineering systems surrounding those physical processes have enormous room to improve.

Conclusion

Silicon carbide is no longer an emerging technology — it is the power electronics platform of the energy transition. The performance, reliability, and cost trajectory of SiC power devices is inseparably linked to the quality of the substrate and epitaxial materials from which they are built. The progress achieved over the past decade in reducing micropipe density, improving BPD conversion efficiency, and delivering uniform, low-defect epitaxial layers at 150 mm scale has been remarkable. The challenge ahead — scaling those achievements to 200 mm, reducing defect densities further, and enabling next-generation device architectures — is equally ambitious.

The organizations that lead this materials challenge

will define the competitive landscape of SiC power electronics for the decade ahead. For the industry as a whole, the continued investment in substrate and epitaxial technology is not a cost of doing business — it is laying the cornerstones that support the most important power electronics revolution of this generation.

Gary Ruland is a technology executive at Coherent. With a PhD in physical chemistry and decades of experience in advanced materials, he is a recognized thought leader in SiC substrate and epitaxial technology for power electronics applications.

Rajan Rengarajan is a Senior Director at Coherent, leading SiC crystal growth and wafering R&D. With a PhD in crystal growth and deep expertise in developing advanced ferroelectric and wide-bandgap materials, his focus has been on SiC materials and technologies for the past two decades.

Shanthi Subramanian is a Senior Director at Coherent, leading SiC metrology and epitaxy R&D. She holds a PhD in materials science and engineering, brings more than 25 years of experience, driving innovation across crystal

growth, process development, and metrology. Qi Yang is a Senior Manager at Coherent, leading SiC data systems. With a PhD in experimental condensed matter physics, his work at Coherent brings software, data, and AI technologies to SiC manufacturing.

Ilya Zwieback is a Senior Director at Coherent. With a PhD in semiconductor materials and devices, and deep experience in crystal growth, epitaxy, and devices of a wide range of semiconductors, he has been a key contributor to SiC technology for the past two decades.

FURTHER READING

- [1] Form Factor is in essence the package. Individual SiC and Si devices use the same packages as Si, for instance, individual TO-220, TO-247, as well as surface packages, etc. It is only the SiC modules and systems have much smaller packages (form factors) due to the much higher switching frequencies, much lower heat load and much smaller passive components.

PEI POWER ELECTRONICS INTERNATIONAL


BOOK YOUR REPRINT TODAY!

A reprint from Power Electronics International amplifies your editorial exposure and provides authoritative third-party validation. It's a powerful tool for sales and marketing, helping you showcase innovation to engineers, designers, and decision-makers with independent credibility. Use it across events, pitches, newsletters, and social channels to reinforce trust, extend visibility, and strengthen your position in the power electronics market.



Contact: Shehzad Munshi
shehzad.munshi@angelbc.com





HfO₂ for better GaN transistors

Introducing a HfO₂ gate dielectric improves the vertical GaN transistor, boosting its drain current density and its breakdown voltage

BY ENRICO BRUSATERRA, ELDAD BAHAT TREIDEL, PALLABI PAUL, MARTIN DAMIAN CUALLO, FRANK BRUNNER, INA OSTERMAY AND OLIVER HILT FROM THE FERDINAND-BRAUN-INSTITUT, BERLIN

THE LATERAL GaN HEMT is without doubt a great success, generating significant sales that will continue to climb to billions of dollars per annum by the end of this decade. But despite the ramping revenues, it's not the best geometry for this higher class of power device. The vertical architecture has the upper hand on a number of fronts, promising a higher breakdown voltage, a higher on-state current, a reduced resistance and a lower thermal impedance for a given chip size.

In addition to all these advantages, the vertical device has another valuable asset – it provides a normally-off behaviour, with a positive threshold voltage required to turn the transistor on. This characteristic is a mandatory requirement for safe, high-power operation.

Developers of vertical transistors tend to use gates that are either based on metal-oxide-semiconductors or junctions. Of these two, it's only a variant of the former, the inversion-type MOSFET, that provides a sufficiently high positive threshold voltage required for fail-safe normally-off operation,

alongside a low gate leakage current necessary for stable operation in a noisy power-electronic environment.

Within the family of vertical GaN MOSFETs, those with a trench-gate architecture are the most common. A key merit of this design is that it allows the *p*-type GaN channel-layer to be formed by epitaxial growth, thereby avoiding the need for ion implantation of magnesium, followed by successive activation during the fabrication process. Such activation is not ideal, demanding lengthy annealing above 1300°C, a temperature so high that it already decomposes the GaN crystal. What's more, as the processing of the trench MOSFET does not require particularly high-resolution lithography, these transistors can be produced with a higher yield and a lower cost than other architectures, such as FinFETs.

One of the requirements for the gate module is that it has an electrical connection to all three epitaxial layers: drift, channel and source. To fulfil this requirement, a trench is etched into the epitaxial

stack prior to deposition of the gate dielectric and the gate metal.

When operating the trench MOSFET, electrons in the *p*-GaN layer, induced by a positive gate bias, form the transistor inversion-channel. During device processing, trench-etching damage may occur, slowing electron transport. Causes include interface roughness and the formation of defect states at the interface of the GaN and the gate dielectric at the trench side-walls. Due to this, channel electron mobility can be quite low.

An undesirable consequence of this state of affairs is that it's hard for trench MOSFETs to achieve channel mobilities and switching speeds as high as those in planar gate devices.


For engineers employing modern GaN-based semiconductor technology, there are two primary options for the gate dielectric of the trench MOSFET. For production purposes, SiO₂ is preferred, due to the higher throughput of low-pressure CVD machines; and for research and development, Al₂O₃ is the standard, as high-quality thin layers can be deposited with atomic layer deposition, and this oxide is a well-established high-*k* dielectric in the semiconductor industry.

To improve the performance of the trench MOSFET, our team at Ferdinand-Braun-Institut, Berlin, is pioneering a new gate dielectric, consisting of a stack formed from HfO₂ and Al₂O₃. This pairing is designed to combine a high relative dielectric constant with a low interface trap density and the realisation of a small threshold voltage shift, as typically observed for pure Al₂O₃ layers. A high dielectric constant is very desirable, as it increases inversion charges in the *p*-type channel, as well as helping manage the electric field inside the gate trenches.

Selecting HfO₂ is not a surprising choice, given that this oxide has gained a growing popularity for GaN-based transistors (see Figure 1). Merits of HfO₂ include a very high relative dielectric constant of around 20, while still maintaining a high enough bandgap – it is around 5.4 eV – and, more importantly, a good band offset to both the GaN conduction (1.1 eV) and the valence (1.6 eV) band. HfO₂ is also a tried and tested material in the semiconductor industry, having been developed for silicon-based applications in DRAMs since 2007.

To ensure that we can work with this oxide in our standard trench MOSFET process line, we tested the material properties of HfO₂, and its stability to other process steps. This investigation determined that HfO₂ is stable up to 350°C, shows signs of decomposition at 500°C, and starts to crystallise at 550°C.

Based on these findings, we know the extent of the temperature budget available for all the process




➤ Figure 1. Publication trend of the last 20 years on HfO₂ in general, and in combination with GaN.

steps for transistor fabrication after oxide deposition. Another factor influencing oxide stability is the surface quality of GaN. We discovered that ammonia plasma-treated surfaces are more stable overall. In addition, plasma cleaning helps reduce oxygen impurities at the GaN surface, leading to improved electrical performance for the dielectric.

We have compared our new dual material dielectric with our standard gate technology that's based on just Al₂O₃. Our stacked oxide consists of 60 percent hafnium and 40 percent aluminium, a combination that provides a relative dielectric constant above 15 while maintaining a low interface trap density associated with the Al₂O₃ component of below 4 x 10¹² cm⁻², and a high breakdown field provided by HfO₂ of 6.6 MV cm⁻¹. Another encouraging finding is that our combined oxide is more than twice as stable to voltage-induced stress as either just Al₂O₃ or just HfO₂. This robustness is especially beneficial for transistors, as exposing the gate to voltage stress may induce a shift in threshold voltage and ultimately instability during operation.

Our newly developed dielectric stack has been implemented in our standard trench-MOSFET process line. We have produced devices on 2-inch ammonothermally-grown bulk GaN substrates provided by the Institute of High-Pressure Physics of



➤ Figure 2. (left) An image of a multi-cell trench-MOSFET cross-section taken using focused ion beam (FIB). (right) A device schematic diagram.

► Figure 3. Photo of a packaged trench MOSFET with 1800 nm gate periphery processed by the Ferdinand-Braun-Institut (FBH).




the Polish Academy of Science (for more details, see the papers listed in 'Further Reading'). The design of this trench MOSFET and details related to epitaxy are provided in Figure 2.

Electrical measurements on our trench MOSFETs with the novel stacked gate dielectric reveal that compared to that with a standard Al_2O_3 gate dielectric, forward current is up to three times higher, and there is a substantial reduction in the threshold-voltage shift induced by positive gate-bias stress. Another strength is the much higher channel mobility of $11.1 \text{ cm}^2 \text{ V}^{-1} \text{ s}^{-1}$. Thanks to this, on-state resistance falls by a factor of three, despite no apparent difference in accumulation capacitance or threshold voltage.

We explain these phenomena by considering that, for the same gate capacitance, the two different dielectrics are exposed to very different electric fields, given the different oxide thicknesses of 25 nm and 50 nm needed to maintain the same input capacitance. A lower electric field inevitably reduces trapping effects in the oxide. The diminished electric field that reduces interface trapping accounts for the superiority of our trench MOSFET with a gate that features both HfO_2 and Al_2O_3 .

Incorporating HfO_2 into the gate stack increases the breakdown voltage of the gate modules. With HfO_2 , a 50 nm-thick dielectric can handle up to 250 V, while the standard Al_2O_3 -based dielectric only withstands 50 V over 25 nm. Our measurements have determined that the breakdown voltage of transistors that incorporate HfO_2 are 410 V. Testing our trench MOSFETs with the new gate dielectric under high-voltage switching conditions produces promising results (a photo one of these



► Figure 4. Dynamic switching characteristic of a trench MOSFET with HfO_2 gate dielectric. After 10 ms off-state stress with 20–360 V drain bias, the transistor is turned on by switching the gate bias from 0 V to 15 V. The on-state drain current is monitored for 10 ms at 10 V drain bias.

transistors, which have a large gate periphery of 1800 nm that's designed to deliver high currents, is shown in Figure 3). When stressing these devices under a high voltage in their off-state for 10 ms, we monitored the drain current after turn-on for another 10 ms (results are shown in Figure 4). Operating at an off-state drain bias of 360 V, on-state resistance degraded by only 50 percent, increasing from $16.7 \text{ m}\Omega \text{ cm}^2$ to $25.6 \text{ m}\Omega \text{ cm}^2$. Note that the correlation between the dynamic on-state resistance and off-state drain voltage is not typical degradation due to hot electrons (this is common in lateral GaN HEMTs, where current flows close to the surface and is sensitive to the presence of localised charges and defects).

Following in-depth analysis, based on the degradation of the on-state resistance over time, we concluded that for low stress biases, the non-depleted part of the drift region interacts with the gate trench, showing a wide range of traps and defects. This situation differs for high stress biases. In this case, the portion of the drift region close to the trench is fully depleted, making interaction less likely. This suggests that the degradation of on-state resistance is related to the gate trench itself, and suggests that improving the dielectric interface to the GaN inside the trench can advance the switching performance of the trench MOSFET.

With interest in GaN power electronics increasing, and efforts at launching commercial vertical GaN transistors underway by both start-ups and established chipmakers, our efforts provide a timely development on how to improve the performance of these most promising of devices.

FURTHER READING

- E. Brusaterra et al., "Vertical GaN Trench MOSFETs With $\text{HfO}_2 / \text{Al}_2\text{O}_3$ Layered Gate Dielectric," in *IEEE Trans. Semicond. Manuf* **38** 750 (2025)
- P. Paul et al. "Plasma enhanced atomic layer deposition of HfO_2 – A potential gate dielectric for GaN-based devices," in *J. Vac. Sci. Technol. A* **43** 042406 (2025)

Advancing Ga₂O₃ diodes with copper anodes

Copper anodes are enabling Schottky diodes to deliver an impressive performance while offering robust thermal reliability

ENGINEERS from China are claiming to have delivered a significant advancement in the performance of β -Ga₂O₃ power rectifiers through the introduction of copper anodes.

According to the spokesman for the team, Xiaohui Wang from Southern University of Science and Technology, by integrating a Cu₂O/Ga₂O₃ heterojunction with a low work-function copper anode, this collaboration has produced a diode that combines a turn-on voltage of less than a volt with a breakdown voltage beyond 2 kV. The result is a leading figure-of-merit for power devices.


Another feature of this device, produced by a partnership involving researchers at Southern University of Science and Technology, Peng Chen Laboratory, The University of Hong Kong and Shenzhen Polytechnic University, is its enhanced field control. It is said that the highly doped p -Cu₂O layer and the Cu₂O-based junction termination extension effectively redistribute the electric field and suppress peak electric field crowding. According to Wang, these breakthroughs have played a key role in establishing a new material platform and design strategy beyond conventional NiO/Ga₂O₃ heterojunction barrier Schottky diode structures.

Wang and co-workers are not the first to try and improve NiO/Ga₂O₃ by turning to Cu₂O, an alternative that helps to excel on two key fronts: a high blocking voltage and a low forward voltage. Previous success has been reported by engineers at NCT, producing an ampere-class large-area diode that's optimised for high current density and switching performance.

"Our work introduces a low work-function copper anode with an *in-situ* platinum layer to form an optimised p -contact to copper oxide," explains Wang, who points out that the device from NCT, featuring nickel contacts and produced using a lift-off trench-filling Cu₂O process, has significantly higher values for on-voltage and specific on-resistance.

Another strength of the device produced by the team from China is a result of its optimised planar p -Cu₂O ring, which suppresses electric-field crowding and shifts the peak electric field into the drift region.

Efforts by Wang and co-workers have not been restricted to device design, and include the fabrication process: "Our device benefits from rigorous interface engineering, including copper oxide with a high hole concentration, an *in-situ* platinum contact and surface preparation. These processes lead to lower



leakage, a more uniform electric field distribution, and superior thermal/electrical reliability."

Devices have been produced by processing epiwafers, featuring a 10 μm-thick silicon-doped drift layer on a heavily tin-doped native substrate, into Schottky barrier diodes and heterojunction barrier Schottky diodes (see Figure for details). Characterisation of the latter determined a turn-on voltage of 0.83 V, a breakdown voltage of 2345 V, and a figure-of-merit for power of 1.22 GW cm⁻².

Efforts have also been directed at evaluating long-term reliability. Operating under a 200 V reverse-bias stress at 425 K for 10,000 s produced negligible degradation in dynamic on-resistance and on-voltage. "Combined with a trap-related mechanism analysis, these results offer critical experimental evidence for the practical deployment of gallium oxide devices in high-temperature, high-voltage applications," argues Wang.

One of the team's goals is to improve heterojunction quality through a combination of controlled oxidation, refined surface treatment, and alternative deposition strategies.

"These efforts are expected to suppress interface trap density, enhance carrier transport uniformity, and further reduce leakage and barrier inhomogeneity," says Wang.

There are also plans to: increase the breakdown voltage while maintaining the low specific on-resistance of 4.5 mΩ cm², a goal that will be pursued by introducing either multi-zone junction termination extension, field plates, or hybrid termination structures; and to explore the on-chip integration of the team's diodes, by using surge-protection devices, field-control components, and high-voltage rectifiers.

➤ Schottky barrier diodes (SBDs) and heterojunction barrier Schottky diodes (JBSs) benefit from a highly doped p -Cu₂O layer and a Cu₂O-based junction termination extension.

REFERENCE

➤ X. Wang *et al.* Appl. Phys. Lett. 127 212102 (2025)

Improving MBE for GaN power devices

Dividing the growth of a channel into a dozen layers boosts the blocking voltage of MBE-grown GaN HEMTs

A TEAM from China has strengthened the case for manufacturing GaN HEMTs by plasma-assisted MBE by producing simple devices that can block more than 2.5 kV.

While MOCVD dominates the manufacture of GaN HEMTs for power electronics, there are a number of benefits associated with plasma-assisted MBE.

One of the strengths of this class of epitaxy is that it is capable of producing an insulating buffer without the need for doping. Avoiding doping is a significant asset, as it eliminates current collapse when switching GaN HEMTs between their on and off states.


Another advantage of plasma-assisted MBE is the suppression of background impurities and defects, thanks to growth under ultra-high vacuum.

Exploiting these assets, the team from Massphoton, Suzhou Powerhouse Electronics Technology, and Suzhou Institute of Nano-Tech and Nano-Bionics, have produced GaN HEMTs that feature a 70 nm-thick channel produced with a novel growth scheme.

To form this device, the team divides the growth of the channel into 12 layers and employs a gallium droplet depletion treatment. This approach is said to enhance the quality of the GaN channel, due to fewer defect states and a lower pore defect density.

Fabrication of the HEMTs began by loading a 3-inch sapphire substrate with a 2° off-cut into a plasma-assisted MBE chamber and depositing a 1000 nm-thick AlN buffer at 920°C, followed by growth of the 70 nm-thick GaN channel, two monolayers of AlN and a 28 nm-thick layer of $\text{Al}_{0.22}\text{Ga}_{0.78}\text{N}$, all at 815°C. After growing this III-N stack, the team added, *in situ*, a 10 nm-thick SiN passivation layer. In addition, the team produced a control device with a slightly thicker channel (see Figure for details).

Processing these epiwafers into devices began with the deposition of a 30 nm-thick SiN layer by low-pressure CVD, followed by nitrogen implantation to realise mesa isolation, and the addition and annealing of source and drain ohmic contacts. After the growth of a 150 nm-thick layer of SiN, gate windows with a recess depth of 120 nm were defined, to leave 30 nm of SiN as the gate insulator. Deposition of Ti/Al/Ti followed, creating a gate metal and field plate, prior to the addition of interconnects and deposition of a 2 µm-thick SiN top passivation layer.



➤ A superior channel improves the HEMT's blocking voltage, leakage current and on-off ratio.

Electrical measurements using the transmission line method determined a contact resistance of 3.6 Ω mm and a sheet resistance for the wafer of 491 Ω sq⁻¹. Meanwhile, Hall measurements provided values for the two-dimensional electron gas density of 1.07×10^{13} cm⁻² and the field-effect mobility of 884 cm² V⁻¹ s⁻¹, suggesting a sheet resistance of 660 Ω sq⁻¹. The engineers attribute a discrepancy between values for sheet resistance obtained from different measurements to a buffer layer below the GaN channel that's not sufficiently insulating and leads to parasitic conduction.

It's claimed that improving the growth conditions of the buffer layer would reduce the sheet resistance and enable a thicker channel layer.

Transfer measurements of the GaN HEMT featuring modified channel growth show current saturation in the output characteristics. This characteristic, not observed in a control device with a conventional channel, is attributed to the superior channel produced with subdivided layers.

Additional benefits associated with the superior channel include an increase in the on-off ratio to 10⁷ and a reduction in the drain leakage current at a drain bias of 0.5 V to 1 nA mm⁻¹ – in both cases, there is an order-of magnitude improvement over the HEMT with the conventional channel.

Both forms of HEMT have a relatively high on-resistance, due to the thin GaN channel thickness.

According to the team, increasing the channel thickness will help to reduce the on-resistance, which can also fall by introducing a graded layer and improving the growth process.

The control device breaks down at 1000 V, while the HEMT with the superior channel can withstand around 2500 V. The team speculates that the far higher value stems from the superior insulating properties of the GaN channel.

REFERENCE

➤ Z. Yang *et al.* Appl. Phys. Express 18 114001 (2025)

INDUSTRY SERVICES DIRECTORY

To promote your products and services contact:

Shehzad Munshi T: +44 (0)1923 69015 E: shehzad.munshi@angelbc.com

AIXTRON

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



Brought to you by:
LUMINUS

Email: apc-e@luminus.com
www.luminus.com/products/APC-SiC-MOSFETs

evatec
process systems

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

GT Greene Tweed

SEALS | CONNECTORS | COMPONENTS

gtweed.com

HITACHI
Inspire the Next

Hitachi Energy

www.hitachienergy.com

WINSOURCE
ELECTRONICS

Tel: + 86-755-83957316
Email: service@win-source.net
www.win-source.net



WÜRTH ELEKTRONIK
MORE THAN YOU EXPECT

Tel: +49 7942 945-0
Email: eiSos@we-online.com
www.we-online.com/en/components/products

Showcase your presence in PEI for 12 months

Promoting your products or your brand through a Corporate Partnership Program in is an economical way to generate interest and drive prospects. Your message will reach over 53,000 professionals worldwide through 3 different mediums of magazine, website and newsletter which creates maximum visibility.

With the largest and most diverse global audience of qualified buyers served by Power Electronics International, secure a corporate partnership and buyers guide entry and receive 12 months of promotion and exposure to the power electronic industry across the magazine portfolio including:

- **Newsletter:** One sponsored newsletter : A sponsor's message (up to 100 words) A 728 x 90 banner displayed at the centre of the alert
- **Magazine:** Listing in 4 issues: A position on the Corporate Partners page for company logo, contact details and web address 300 x 150 pixels.
- **Website:** A button banner 160 x 54 pixels for 12 months plus sponsor logo: 24/7 and 365 days (jpg or .gif; max. File size: 15k click-thru URL). <https://POWEELECTRONICSMAGAZINE.NET/home>
- All Corporate Partners are included in online buyers guide <https://powerelectronicismagazine.net/buyers-guide/companies>

WEBINARS

Specialists with 30 year+ pedigree and in-depth knowledge in overlapping sectors



Expertise: Moderators, Markets, 30 Years + Pedigree

Reach: Specialist vertical databases

Branding: Message delivery to high level influencers via various in house established magazines, websites, events and social media

For more information contact:

Jackie Cannon T: 01923 690205 E: jackie@angelwebinar.co.uk W: www.angelwebinar.co.uk
T: +44 (0)2476 718 970 E: info@angelbc.com W: www.angelbc.com

Angel 
BUSINESS COMMUNICATIONS

WÜRTH ELEKTRONIK MORE THAN YOU EXPECT

YOU'VE GOT BETTER OPTIONS THAN THAT

Check out our Thermal Management Solutions



**WE meet @
PCIM Europe**
Hall 6 - 342

Thermal management is crucial for developing durable and efficient designs. Our gap filling, heat spreading, and hybrid solutions deliver optimal thermal management solutions, ensuring your design achieves maximum durability and efficiency. Paired with our expert services and custom solutions, we provide the perfect fit for your application. Ready to keep it cool?

www.we-online.com/thermalmanagement

#THERMAL

Highlights

- Gap filling, heat spreading and hybrid solutions
- Customized components for optimal fit
- Advisory services from technical experts

