




POWER ELECTRONICS

ISSUE IX 2025

 AN ANGEL BUSINESS COMMUNICATIONS PUBLICATION

POWERELECTRONICSMAGAZINE.NET


ADDRESSING THE PROTECTION CHALLENGES OF 48V AI SERVERS USING HOT-SWAP CONTROLLERS



AIXTRON

THE NEW G10 SERIES

Your Productivity Solution for All Advanced Epitaxy Materials



G10-SiC

- ▶ 9x150 mm (6") or 6x200 mm (8")
- ▶ New hardware & process surpassing Single Wafer Reactor uniformities
- ▶ +50% productivity per fab area with large batch technology & small footprint

End Markets/Products:

EV inverters & charging infrastructure

G10-GaN

- ▶ 8x150mm (6") or 5x200 mm (8")
- ▶ 1st fully automated compact GaN MOCVD cluster designed 100% for Si Power fabs
- ▶ Novel hardware solution for unmatched barrier uniformities and device yields

End Markets/Products:

Power Electronics & Wireless communication



G10-AsP

- ▶ 8x150 mm (6") or 5x200mm (8")
- ▶ 1st fully automated AsP MOCVD system enabling 10x lower defect density
- ▶ Unmatched wavelength uniformity on all wafer sizes

End Markets/Products:

Micro LED, Optical Data communication, 3D-sensing & LiDAR



Research community delivers some seasonal successes

THE PAST SEVERAL WEEKS have delivered some significant research breakthroughs across the wide-bandgap landscape — a reminder, perhaps, that power electronics is not merely an industry of incremental advancement but one continually reinventing its scientific foundations. Something of a genuine research renaissance, spanning materials, devices, packaging, metrology, and manufacturing infrastructure?

Nowhere is this more evident than in gallium nitride, where a collaboration led by north carolina state university has achieved a nine-orders-of-magnitude reduction in the specific on-resistance of magnesium-implanted vertical p–n diodes. By boldly revisiting one of gan’s thorniest issues — realising high-quality p-type contacts after implantation — the team demonstrated performance on par with epitaxial devices. Their clever use of magnesium deposition and annealing to eliminate the parasitic schottky barrier is not just a process optimisation; it reopens the viability of selective-area doping for advanced vertical gan architectures. For designers of vertical mosfets, jbs diodes and next-generation high-voltage gan devices, this work is strategically enabling.

Imec’s record-breaking >800 v gan hemt on shin-etsu’s 300 mm qst substrate is another milestone that signals gan’s steady march into large-diameter, cmos-compatible manufacturing. By solving the historic warpage and cracking problems that plagued gan-on-silicon scaling, qst may become a cornerstone substrate technology — arriving just in time as ai data centres and electrified mobility demand cheaper, higher-performance gan power devices.

Sic research, meanwhile, is undergoing its own transformation. The university of osaka’s two-step hydrogen/argon anneal — notably without nitrogen — is a decisive step forward in resolving the long-standing mobility–reliability trade-off at the sic/sio₂ interface. Achieving high mobility while suppressing threshold-voltage drift is precisely the combination needed to support 600–1200 v mosfets in ev and renewable-energy inverters. In tandem, nrel’s ulis

module demonstrates that sic innovation is not confined to materials: packaging engineering can still rewrite the performance envelope. With parasitic inductance cut by nearly an order of magnitude and a design that is both lightweight and health-monitoring-capable, ulis hints at a new breed of intelligent power modules for aviation, microgrids and defence.

If gan and sic dominate the headlines, gallium oxide and aluminium nitride are proving that the materials frontier remains wide open. Nagoya university’s demonstration of the world’s first ga₂o₃ p–n diodes, using industry-standard implantation and annealing steps, marks a pivotal advance for this ultra-wide-bandgap contender. Likewise, new movpe-grown si-doped aln films with tunable properties for high-temperature schottky devices highlight how research is expanding the palette of viable power-semiconductor materials.

Even fundamental heat transport science is being rewritten. The university of houston’s discovery that ultra-pure boron arsenide can exceed 2,100 w/mk, potentially outperforming diamond, could redefine thermal management for power electronics and high-power computing alike — a rare case where theory is being forced to catch up with experiment. Crucially, this surge in research is being matched by investment in infrastructure and standards. The University of Arkansas’ MUSiC facility gives the U.S. its first open-access SiC MPW fab — a powerful lever for accelerating prototyping and training. The TRACE-Ga initiative signals strategic recognition of gallium’s geopolitical importance. And the ITRI–NPL partnership reflects a global push toward metrology leadership in compound semiconductors.

Wishing all our readers a happy and relaxing festive season and here’s to the power electronics industry enjoying a successful 2026.



COVER STORY

Addressing the protection challenges of 48V AI servers using hot-swap controllers

Enabling the design of a reliable input protection solution for a 48V AI server



26

14 MIT engineers develop a magnetic transistor for more energy-efficient electronics

A new device concept opens the door to compact, high-performance transistors with built-in memory

16 Small material, big impact

Dr Conor O'Brien, Senior Technology Analyst at IDTechEx, explores graphene in electronic applications

18 Powering the AI era: how Infineon Technologies' SiC solutions are shaping the future of data centre power supplies

As artificial intelligence (AI) continues to redefine the frontiers of computational performance, the supporting infrastructure particularly the power conversion within data centers must evolve rapidly to meet new performance and density benchmarks



18

32 The rise of 800V High Voltage DC: Can collaboration drive NVIDIA's AI infrastructure vision?

NVIDIA's transition to an 800 VDC architecture for AI data centres isn't just a technological shift – it's a call for industry-wide alignment on power delivery, semiconductors and standards

36 Ferric brings power efficiency to the heart of AI chips

As power-hungry AI processors strain energy systems, Ferric's integrated voltage regulators promise more efficient power delivery

40 Efficient debugging workflow in chip circuit verification - A targeted simulation approach

Chip circuit verification is a vital to the semiconductor industry, ensuring that integrated circuits meet design specifications and quality standards

44 The criticality of performance per watt optimisation for AI chip development

AI workloads create an urgent need for compute and power efficiency improvements from the architecture stage forward

48 Taking SiC switches to new highs

How APC Electronics and Luminus aim to reshape high-power, high-frequency design with next-generation silicon carbide technology

52 From chargers to data centres: power GaN market set for rapid sixfold expansion by 2030


Power GaN 2025, powered by Yole Group, provides a detailed look at the market growth, technology trends, and strategic shifts shaping the gallium nitride power device industry

54 New CEO to lead Wise Integration's global growth

Ghislain Kaiser, the recently appointed Chief Executive Officer of Wise Integration, explains how his main task is to transition the company from a CEA Leti spinout into a pioneering force in GaN and digital power management innovation with strong growth potential

58 SMT Inductors - Coils and ferrites - selection made easy

How do you choose the right inductor for your application? This article compares the differences between three inductive SMT components: ceramic inductors, SMT ferrites, and wire-wound ferrite inductors



NEWS

06 SemiQ launches Gen3 1200V modules

07 STMicroelectronics unveils e-fuse automotive smart switch

08 VisIC announces \$26M investment backed by Hyundai

09 X-FAB speeds time-to-market for SiC MOSFET solutions

10 European DC-Power projects to boost collaboration

11 Wolfspeed receives nearly \$700m in tax refunds

12 Wise, Powernet and KEC to co-develop AI power supplies



PEI POWER ELECTRONICS INTERNATIONAL

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

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 the time of going to press. The publisher does not accept responsibility for any errors and omissions. The views expressed in this publication are not necessarily those of the publisher.

Every effort has been made to obtain copyright permission for the material contained in this publication. Angel Business Communications Ltd will be happy to acknowledge any copyright oversights in a subsequent issue of the publication. Angel Business Communications Ltd © Copyright 2025. All rights reserved. Contents may not be reproduced in whole or part without the written consent of the publishers. The paper used within this magazine is produced by chain of custody certified manufacturers, guaranteeing sustainable sourcing. © Copyright 2025.

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

When companies change their names or the images/text used to represent the company, we invite organizations to provide Angel Business Communications with a news release detailing their new business 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

Angel 
BUSINESS COMMUNICATIONS

SemiQ launches Gen3 1200V modules

Family includes 608A half-bridge module with 2.4mΩ RDSon and best-in-class thermal resistance

SEMIQ INC, a developer of SiC solutions, has announced an expansion of its third-generation QSiC MOSFET product line, including devices with what it claims are industry-leading current density and thermal resistance.

Seven devices have been launched, including high-current S3 half-bridge, B2T1 six-pack and B3 full-bridge packages.

This expansion addresses the growing demand for ultra-efficient conversion in high-power systems, and features devices engineered for current capabilities of up to 608A and a junction-to-case thermal resistance of 0.07°C/W (in the 62 mm standard S3 half-bridge format).

The six-pack modules integrate the three-phase power stage into a compact housing and have an RDSon range of 19.5 to 82mΩ. They are designed to optimize layout and minimise parasitics in motor drives and advanced AC-DC converters.

The full-bridge modules deliver high current capabilities of up to 120A and an ultra-low on-resistance down to



8.6mΩ. This combination, coupled with a low thermal resistance of 0.28°C/W, maximises power density and efficiency in single-phase inverters and high-voltage DC-DC systems.

All parts are screened using wafer-level gate-oxide burn-in tests to guarantee the gate oxide quality. They are also breakdown voltage tested to over

1350 V. Modules using these third-generation chips operate at lower gate voltages than previous generations as a result of the 18 V/-4.5 V gate voltage of the third-generation chips.

SemiQ's Gen3 technology reduces both RONsp and turn off energy losses (EOFF) by up to 30 percent versus previous generations, according to the company.

Mitsubishi adds to high-voltage IGBT range

MITSUBISHI ELECTRIC launching new standard-isolation (6.0kVrms) and high-isolation (10.2kVrms) modules in its 4.5kV/1,200A XB Series of high-voltage HVIGBTs .


These new high-capacity power semiconductors achieve high moisture resistance for more efficient and reliable inverters used in large industrial equipment, such as railcars, operating in diverse environments including outdoors.

Mitsubishi Electric will exhibit the new modules at the 40th Nepcon Japan R&D and Manufacturing show in Tokyo

from January 21 to 23, 2026, as well as other exhibitions in North America, Europe, China, India and additional locations.

The new modules use IGBT elements that incorporate Mitsubishi Electric's proprietary relaxed field of cathode (RFC) diode and carrier-stored trench-gate bipolar transistor (CSTBT) structure.

New structures for electric field



relaxation and surface charge control enabled Mitsubishi Electric to reduce the chip's termination region size by about 30 percent while also achieving about 20 times greater moisture resistance than existing products.

In addition, the module reduces total switching loss by approximately 5% compared to previous models, and reverse-recovery safe-operating area (RRSOA) tolerance is about 2.5 times greater than that of compared to previous models.

STMicroelectronics unveils e-fuse automotive smart switch electrification strategy

New addition to STi2Fuse family reacts within 100µs


STMICROELECTRONICS' VNF1248F automotive e-fuse MOSFET controller combines ST's fast-acting, advanced e-fuse features for harness protection and enhanced flexibility, enabling power saving and advanced performance in automotive functional-safety systems.

A new addition to the STi2Fuse family, the VNF1248F reacts within 100µs, which is faster than a conventional wire fuse and ensures flexible and robust protection to avoid fault propagation inside the vehicle.

The VNF1248F integrates the new capacitive charging mode (CCM) functionality to ensure proper driving of large capacitive loads with high inrush current. In addition, an enhanced Standby-ON mode with current capability up to 600mA and current consumption lower than 75µA enhances the vehicle's efficiency when in park mode, helping increase autonomy. An optional external supply pin for logic reduces power consumption by 0.4W in 48V systems and battery-undervoltage shutdown compatible with the automotive LV124 standard ensures system stability.

The chip facilitates reaching high safety-integrity levels (ASIL) in ISO 26262 functional-safety applications thanks comprehensive dedicated features like advanced fault detection and reaction, fail-safe mode and limp-home mode. There are also built-in self tests for automatic diagnosis and a dedicated pin for direct hardware control of the external MOSFET gate in case of a microcontroller fault

Suitable for 12V, 24V, and 48V boardnets, the VNF1248F handles power distribution in zonal vehicle electrical architectures and general



fuse and relay replacement. Other uses include as an ECU main switch and active supply for always-on circuitry in parking mode.

Joining ST's STi2Fuse family of smart switches for wire harness protection and dynamic power distribution, the VNF1248F has an SPI port to interact with a host microcontroller. Configuring the device using SPI saves conventional external programming components and allows flexible control of settings including hard short-circuit (HSC) latch-off, current-vs-time latch-off in fuse emulation, and MOSFET-desaturation shutdown. With non-volatile memory (NVM) also on-chip, the VNF1248F can additionally store an immutable default configuration to ensure consistent behavior in failsafe or locked states.

The associated EV-VNF1248F evaluation board, pre-assembled and ready to connect directly to the load, power supply, and microcontroller, simplifies integration of ST's intelligent fuse protection into prototype circuitry. A software package, STSW-EV-VNF1248F, is also available and provides a graphical user interface and control firmware to configure and monitor the VNF1248F using an EV-SPC582B microcontroller board.

The VNF1248F MOSFET controller is in production now in a 5mm x 5mm 32-lead QFN32 package, from \$2.63 for orders of 1000 pieces.

Infineon GaN powers Enphase solar microinverters

ENPHASE ENERGY, a leading supplier of microinverter-based solar and battery systems, is using Infineon's CoolGaN bi-directional switch (BDS) technology for its next-generation of solar microinverters.

According to Infineon, its GaN technology enables significant enhancements in power output, energy efficiency, and system reliability for Enphase's IQ9 Series Microinverters. For the new IQ9N-3P Commercial Microinverter, it will also help simplify design complexity and lower installation and balance of system costs.

CoolGaN bi-directional switch technology replaces two or four unidirectional switches with a single bi-directional switch. This allows power flow in both directions, and enables significant

cost savings, while enabling the design of smaller, more efficient power devices, according to Infineon. The BDS technology has also demonstrated a reduction in power loss, outperforming conventional silicon switches by 68 percent and GaN unidirectional switches by 42 percent. By achieving higher power output and efficiency, Enphase's new microinverters can now service not only the residential solar market, but also a much larger portion of the commercial market.

"Utilising Infineon's CoolGaN bi-directional switch technology allows us to service a much larger segment of the commercial market with our IQ9 Series Microinverters," said Ron Swenson, SVP of operations at Enphase Energy.



VisIC announces \$26 million investment backed by Hyundai


Hyundai and Kia join as strategic investors with a commitment to integrate GaN tech into mass-production EV platforms

GaN POWER chip company VisIC Technologies has announced the successful second closing of its Round B funding, securing \$26million.

The round was led by a global semiconductor leader, with car companies Hyundai and Kia (HKMC) joining as a strategic investor.

The lead investor's focus on advancing critical semiconductor technologies complements VisIC's proprietary D³GaN platform, designed to deliver efficiency, scalability, and reliability for automotive drivetrains. HKMC's participation underscores its commitment to integrating GaN technology into mass-production EV platforms.

VisIC says the new capital will accelerate its roadmap, including: optimisation, qualification, and release of Gen3 750V GaN dice and power modules; development of Gen4 1350V GaN technology, supporting the full spectrum of EV designs; stabilisation of the supply chain and ramp-up of GaN product delivery for EV traction inverters; and expansion into emerging 800V data centre power requirements,



leveraging the same advanced GaN platform.

Tamara Baksht, CEO of VisIC Technologies said: "This investment marks a major milestone for VisIC and the global EV industry. Our D³GaN technology is redefining power electronics for electric vehicles, and the support of our strategic partners accelerates our mission to deliver high-efficiency, scalable solutions for the

next generation of mobility."

Hyundai Motor Company and Kia (HKMC) said: "Hyundai Motor Company and Kia are committed to advancing sustainable mobility. Partnering with VisIC enables us to integrate cutting-edge GaN power technologies into our EV platforms, enhancing efficiency, reliability, and performance as we shape the future of electric transportation."

Company's SiC technology to help advance Toyota's electrification strategy

WOLFSPEED has announced that its automotive SiC MOSFETs will power onboard charger systems for Toyota's battery electric vehicles (BEVs).

"Toyota is known for its uncompromising approach to quality and reliability, and we're honoured to be supporting their next wave of electrification," said Robert Feurle, CEO,

"Wolfspeed's US.-based supply chain and domestic SiC manufacturing

footprint ensure the stability and continuity they need to achieve their electrification goals."

"Our work with Toyota is built upon years of trust in engineering expertise, supply reliability, as well as a shared obsession with quality," said Cengiz Balkas, Wolfspeed's chief business officer, "This reinforces our role in driving electrification with SiC technology that delivers performance, efficiency and safety."



X-FAB speeds time-to-market for SiC MOSFET solutions


XbloX platform gives easy access to a standardised SiC process technologies

SPECIALTY foundry X-FAB is offering easy access to a standardised yet flexible set of SiC process technologies that accelerate the development of advanced power devices.

From rapid prototyping to full production, the modular and scalable XbloX platform is said to help SiC device developers to expedite engineering assessments and technology release, with production starts achieved up to nine months faster than traditional methods.

The XbloX WBG discrete foundry model has two major benefits, according to the company. Firstly, X-FAB takes on process development activities with the introduction of a Process Installation Kit (PIK), where design and implant recipes provide key differentiators.

Secondly, the use of XbloX ensures wafer manufacturing at X-FAB becomes a highly scalable activity in line with application requirements, differing considerably from the less scalable production provided by a traditional foundry model for customer-specific SiC technologies.



“Thanks to a PIK, qualified SiC process development modules, and an automated onboarding process, customers need do little more than access our global hotline for support on block selection and deployment,” explains Brian Throneberry, business director SiC Foundry at X-FAB.

“We have robust rules in place to help guide design, mask tooling, engagement, and so on. Once the selection is finalised, XbloX automatically generates the process flow, which subsequently integrates quality systems, business functions, and commercial aspects for the customer. It’s a highly expedited way of providing customers with the flexibility required to build custom SiC MOSFET technologies at an accelerated rate.”

SK Keyfoundry ramps up SiC activity

SK KEYFOUNDRY, an 8-inch pure-play foundry in Korea, is accelerating the development of SiC-based power semiconductor technology with the acquisition of SK Powertech, a key SiC player.

“The acquisition of SK Powertech, a specialist in SiC, marked a pivotal step for SK Keyfoundry in securing its own distinctive technological edge in the compound semiconductor field,” said Derek D. Lee, CEO of SK Keyfoundry. “By combining the core development capabilities of both companies and

launching high-efficiency SiC power semiconductor process technologies and products, SK Keyfoundry aims to establish differentiated technological leadership in the rapidly growing global market for high-voltage and high-efficiency compound semiconductor applications.”

Building on this foundation, SK Keyfoundry aims to provide SiC MOSFET 1200V process technologies by the end of 2025 and launching SiC-based power semiconductor foundry business in the first half of 2026.

Coherent 300mm SiC platform addresses AI needs

COHERENT has announced that its next-generation 300mm SiC platform will address the increasing higher power density, faster switching, and thermal efficiency demands in AI data centre infrastructure.

“AI is transforming the thermal-management landscape in data centres, and SiC is emerging as one of the foundational materials enabling this scalability,” said Gary Ruland, SVPt and general manager at Coherent.

“Our 300mm platform, which we plan to ramp in high volumes, delivers new levels of thermal efficiency that translate directly into faster, more power-efficient AI data centres.”



The platform’s conductive SiC substrates provide low resistivity, low defect density, and high homogeneity, enabling low-dissipation, high-frequency, and good thermal stability.

In AI and data infrastructure, these properties boost energy efficiency and thermal performance in next-generation data centre systems.

In addition, the company says the technology brings benefits for AR/VR devices, enabling thinner and more efficient waveguides for AR smart glasses and VR headsets, improving reliability in compact immersive display modules.

European DC-Power projects to boost collaboration

THEUS, DC-POWER, InterSCADA, FLAGCHIP, PROSECCO and SAFEPOWER look towards building a coherent European DC and hybrid AC/DC innovation ecosystem

EARLIER this month, seven European projects in DC and hybrid AC/DC grid technologies held a joint coordination meeting to strengthen cooperation and identify synergies.

The participating projects which include THEUS, DC-POWER, InterSCADA, FLAGCHIP, PROSECCO and SAFEPOWER, aim to create a more modern, flexible, and efficient electricity system, capable of supporting increased renewable integration, sector coupling, and the EU's long-term climate and energy security goals.

During the meeting, each project presented its objectives, activities, and expected outcomes. A dedicated discussion session was held to explore common challenges, shared research priorities, and opportunities for joint actions across technology development, demonstration activities, and dissemination efforts.

This first collective exchange represents an important step toward building a more coherent and collaborative European DC and hybrid AC/DC innovation ecosystem. By aligning efforts and sharing expertise, the projects aim to accelerate progress toward more efficient, resilient, and sustainable electricity systems.

"The DC-Power project has developed a solution for the cost-efficient distribution of electricity for MW scale appliances, such as datacenters, electrolysers, fast charging stations of electric trucks and cars.

Key is the use of medium voltage DC grids which has been found to be an optimal solution for electro intensive industry", says Jens Merten, the coordinator who is with CEA, a large scale technology developer in France.



Carlos Pérez Montero, THEUS Coordinator: "Our synergy call showed that the challenges we face in hybrid AC/DC grids are shared across projects, and that by aligning our methods and sharing use cases we can accelerate the deployment of next-generation grid solutions across Europe".

Eduardo Garcia Martinez, FLAGCHIP Coordinator: "By sharing our objectives, innovations, and approaches, we gained a clearer collective picture of where the field is heading."

Dirk Van Hertem, PROSECCO Coordinator: "Hybrid AC/DC grids introduce new protection, control, and system engineering challenges that no single project can solve alone. This exchange allows us to align our approaches, share insights, and accelerate solutions that will support a resilient and interoperable European grid."

About the projects:

- **DC-POWER:** Medium-voltage DC grid solutions for industrial and utility applications. dcpower.tech

- **InterSCADA:** Open-source SCADA platform for hybrid AC/DC monitoring and control. interscada.eu

- **THEUS:** Tools and methodologies for planning and operating hybrid AC/DC networks across all voltage levels. theus-project.eu

- **FLAGCHIP:** High-efficiency power electronics using wide-bandgap and ultra-wide-bandgap semiconductors. flagchip-project.eu

- **PROSECCO:** Enhanced protection, congestion management, and system engineering for HVDC and hybrid grids. prosecco-project.be

- **SAFEPOWER:** Next-generation MVDC converters with advanced semiconductors and AI-enabled condition monitoring. safepowerproject.eu

Wolfspeed receives nearly \$700m in tax refunds


Cash injection gives company more financial flexibility as it ramps its 200mm SiC line

WOLFSPEED has received \$698.6m in cash tax refunds from the US Internal Revenue Service (IRS) from the Advanced Manufacturing Investment Credit (AMIC) under Section 48D of the Internal Revenue Code. The company says this refund represents a significant step in getting back the approximately \$1 billion of Section 48D cash tax refunds. In fiscal 2025, the company received \$186.5m in cash tax refunds related to its 2023 and 2024 federal tax filings.

Following receipt of these funds, Wolfspeed's cash balance is approximately \$1.5b, positioning the company with more financial flexibility as it ramps its 200mm SiC line.

"This substantial cash infusion further strengthens our liquidity position at a critical phase in Wolfspeed's strategic evolution," said Wolfspeed CFO, Gregor Van Issum. "It provides us with the financial agility to support long-term growth, manage our capital structure responsibly, and continue driving innovation across the SiC value chain for our customers."

Over the past several years, Wolfspeed has made significant investments to establish a vertically integrated, US-based, and highly resilient supply chain



for SiC materials and power devices. Section 48D support has helped accelerate the company's transition from 150mm (6-inch) to 200mm (8-inch) wafer technology, a critical next-generation platform for increasing efficiency and expanding global adoption of SiC.

As per the agreement with the senior secured lenders, the company plans to allocate \$192.2m of the refund toward retiring approximately \$175m of outstanding debt, with the remaining funds for general corporate purposes.

Over the past several years, Wolfspeed has made significant investments to establish a vertically integrated, US-based, and highly resilient supply chain for SiC materials and power devices

PEI **POWER ELECTRONICS INTERNATIONAL**
DESIGN • MATERIALS • PACKAGING • PERFORMANCE



- Based around a hot topic for your company, this 60-minute recorded, moderated zoom roundtable would be a platform for debate and discussion
- Moderated by editor, Phil Alsop, this can include 3 speakers
- Questions prepared and shared in advance

Cost: €5995

Contact: Jackie Cannon
jackie.cannon@angelbc.com

Wise, Powernet and KEC to co-develop AI power supplies

French digital GaN pioneer and Korean power supply specialists join forces as Korea accelerates AI data centre build-out

FRENCH GaN company Wise Integration has signed a memorandum of understanding (MoU) with Korean power firms Powernet and KEC to co-develop next-generation switched-mode power supply (SMPS) solutions for AI server applications in South Korea.

Under the agreement, Wise Integration will supply its GaN power devices, digital-control expertise and technical support. Powernet Technologies will lead development of new SMPS designs using Wise's WiseGan and WiseWare technologies. KEC will manage backend manufacturing, including module integration and system-in-package production tailored to the thermal and reliability demands of AI-server racks.

In addition to accelerating the design and development of competitive AI-server power supply solutions and creating business opportunities in Korea's AI server market, the project aims to shorten the solutions' time-to-market using WiseGan and WiseWare technologies. The collaboration builds on an earlier partnership between Wise Integration and Powernet, launched to serve OEMs that require compact, digitally controlled power-supply systems for faster, smaller and more energy-efficient electronic equipment.

AI servers draw extensive power, generate intense heat, and depend on SMPS designs capable of converting high-voltage input (e.g., 400 V) to

stable 48 V rails with minimal loss. GaN devices with digital control are uniquely suited for that task: they enable higher switching frequencies, greater efficiency, and more precise management of fast, high-current load transients.

The South Korean government's investments in AI-dedicated data centres includes high-performance GPU clusters and digital infrastructure that supports demand for more efficient, compact, and scalable power-conversion systems.

"Korea is moving quickly to build the next generation of AI data centres, and power architecture is a critical piece of that effort. Working with Powernet and KEC lets us bring GaN-based digital control into server-grade designs at scale—delivering the efficiency, thermal performance, and responsiveness that modern AI hardware depends on," said Ghislain Kaiser, CEO of Wise Integration.

Beyond the MoU, Wise has been widening the foundation that supports strategic partnerships. As Korea accelerates AI-ready data centre development, the demands on power-conversion systems intensify — higher switching frequencies, tighter thermal budgets and more aggressive efficiency targets. Wise has been expanding its GaN + digital-control portfolio to meet those constraints, ensuring the technologies Korea aims to deploy at scale already have a maturing, production-validated base behind them.



Wolfspeed to power Toyota BEVs

Wolfspeed has announced that its automotive SiC MOSFETs will power onboard charger systems for Toyota's battery electric vehicles (BEVs).

"Toyota is known for its uncompromising approach to quality and reliability, and we're honoured to be supporting their next wave of electrification," said Robert Feurle, CEO.



"Wolfspeed's US.-based supply chain and domestic SiC manufacturing footprint ensure the stability and continuity they need to achieve their electrification goals."


"Our work with Toyota is built upon years of trust in engineering expertise, supply reliability, as well as a shared obsession with quality," said Cengiz Balkas, Wolfspeed's chief business officer, "This reinforces our role in driving electrification with SiC technology that delivers performance, efficiency and safety."

For example, the company recently launched WiseWare 1.0 for totem-pole PFC and LLC topologies, its newest fully digital controller. While aimed today at gaming, displays and industrial systems, it shares the same architecture—high-frequency GaN operation, compact form factor, digitally managed efficiency—that naturally scales into the server-class designs targeted under the Korea-focused collaboration.

WÜRTH ELEKTRONIK MORE THAN YOU EXPECT

ULTRA LOW LOSSES

WE-MXGI



e eiS05



With the WE-MXGI Würth Elektronik offers the newest molded power inductor series. It combines an innovative iron alloy material that provides high permeability for lowest R_{DC} values combined with an optimized wire geometry.

Ready to Design-In? Take advantage of personal technical support and free samples ex-stock.


www.we-online.com/WE-MXGI

Highlights

- Extremely high power density
- Ultra low R_{DC} values and AC losses
- Magnetically shielded
- Optimized for high switching frequencies beyond 1 MHz

#UltraLowLosses

MIT engineers develop a magnetic transistor for more energy-efficient electronics



A new device concept opens the door to compact, high-performance transistors with built-in memory.

BY ADAM ZEWE, MIT NEWS OFFICE

TRANSISTORS, the building blocks of modern electronics, are typically made of silicon. Because it's a semiconductor, this material can control the flow of electricity in a circuit. But silicon has fundamental physical limits that restrict how compact and energy-efficient a transistor can be.

MIT researchers have now replaced silicon with a magnetic semiconductor, creating a magnetic transistor that could enable smaller, faster, and more energy-efficient circuits. The material's magnetism strongly influences its electronic behavior, leading to more efficient control of the flow of electricity. The team used a novel magnetic material and an optimization process that reduces the material's defects, which boosts the transistor's performance.

The material's unique magnetic properties also allow for transistors with built-in memory, which would simplify circuit design and unlock new applications for high-performance electronics.

"People have known about magnets for thousands of years, but there are very limited ways to incorporate magnetism into electronics.

We have shown a new way to efficiently utilize magnetism that opens up a lot of possibilities for future applications and research," says Chung-Tao Chou, an MIT graduate student in the departments of Electrical Engineering and Computer Science (EECS) and Physics, and co-lead author of a paper on this advance.

Chou is joined on the paper by co-lead author Eugene Park, a graduate student in the Department of Materials Science and Engineering (DMSE); Julian Klein, a DMSE research scientist; Josep Ingla-Aynes, a postdoc in the MIT Plasma Science and Fusion Center; Jagadeesh S. Moodera, a senior research scientist in the Department of Physics; and senior authors Frances Ross, TDK Professor in DMSE; and Luqiao Liu, an associate professor in EECS, and a member of the Research Laboratory of Electronics; as well as

others at the University of Chemistry and Technology in Prague. The paper appears in *Physical Review Letters*.

Overcoming the limits

In an electronic device, silicon semiconductor transistors act like tiny light switches that turn a circuit on and off, or amplify weak signals in a communication system. They do this using a small input voltage. But a fundamental physical limit of silicon semiconductors prevents a transistor from operating below a certain voltage, which hinders its energy efficiency.

To make more efficient electronics, researchers have spent decades working toward magnetic transistors that utilize electron spin to control the flow of electricity. Electron spin is a fundamental property that enables electrons to behave like tiny magnets. So far, scientists have mostly been limited to using certain magnetic materials. These lack the favorable electronic properties of

semiconductors, constraining device performance.

“In this work, we combine magnetism and semiconductor physics to realize useful spintronic devices,” Liu says. The researchers replace the silicon in the surface layer of a transistor with chromium sulfur bromide, a two-dimensional material that acts as a magnetic semiconductor.

Due to the material’s structure, researchers can switch between two magnetic states very cleanly. This makes it ideal for use in a transistor that smoothly switches between “on” and “off.”

“One of the biggest challenges we faced was finding the right material. We tried many other materials that didn’t work,” Chou says.

They discovered that changing these magnetic states modifies the material’s electronic properties, enabling low-energy operation. And unlike many other 2D materials, chromium sulfur bromide remains stable in air.

To make a transistor, the researchers pattern electrodes onto a silicon substrate, then carefully align and transfer the 2D material on top. They use tape to pick up a tiny piece of material, only a few tens of nanometers thick, and place it onto the substrate. “A lot of researchers will use solvents or glue to do the transfer, but transistors require a very clean surface. We eliminate all those risks by simplifying this step,” Chou says.

“People have known about magnets for thousands of years, but there are very limited ways to incorporate magnetism into electronics. We have shown a new way to efficiently utilize magnetism that opens up a lot of possibilities for future applications and research,”

Leveraging magnetism

This lack of contamination enables their device to outperform existing magnetic transistors. Most others can only create a weak magnetic effect, changing the flow of current by a few percent or less. Their new transistor can switch or amplify the electric current by a factor of 10.

They use an external magnetic field to change the magnetic state of the material, switching the transistor using significantly less energy than would usually be required.

The material also allows them to control the magnetic states with electric current. This is important because engineers cannot apply magnetic fields to individual transistors in an electronic device. They need to control each one electrically.

The material’s magnetic properties could also enable transistors with built-in memory, simplifying the design of logic or memory circuits.

A typical memory device has a magnetic cell to store information and

a transistor to read it out. Their method can combine both into one magnetic transistor.

“Now, not only are transistors turning on and off, they are also remembering information. And because we can switch the transistor with greater magnitude, the signal is much stronger so we can read out the information faster, and in a much more reliable way,” Liu says.

Building on this demonstration, the researchers plan to further study the use of electrical current to control the device. They are also working to make their method scalable so they can fabricate arrays of transistors.


This research was supported, in part, by the Semiconductor Research Corporation, the U.S. Defense Advanced Research Projects Agency (DARPA), the U.S. National Science Foundation (NSF), the U.S. Department of Energy, the U.S. Army Research Office, and the Czech Ministry of Education, Youth, and Sports. The work was partially carried out at the MIT.nano facilities.

VACUUM SOLDERING SYSTEM

UNCOMPROMISING QUALITY,
EFFICIENT DESIGN

VADU core

- Standardized, efficient system design
- Individually expandable to solve complex soldering tasks
- Soldering with preforms or solder pastes
- Proven and tested in series production for 25 years



PiNK®



productronica

18. – 21. November
Trade Fair Munich
Booth A4.255

PiNK GmbH Thermosysteme . 97877 Wertheim/Germany . info@pink.de . www.pink.de

Small material, big impact

Dr Conor O'Brien, Senior Technology Analyst at IDTechEx, explores graphene in electronic applications

GRAPHENE was initially celebrated as a sci-fi material destined to transform electronics, quantum devices, and futuristic sensors. While expectations soared, commercial reality proved more modest. Rather than enabling transparent phones or levitating trains, graphene first succeeded in practical roles such as enhancing polymer composites in tires, coatings, and sporting goods through improved strength and conductivity. Now, with advances in manufacturing and renewed interest from major electronics and energy companies, the question resurfaces: are we finally approaching the breakthrough era once promised?

IDTechEx has covered the graphene market since 2012, providing a unique perspective on technological progress, commercialization strategies, and the long-term market outlook for the material. A new report has been released, "Graphene & 2D Materials

2026-2036: Technologies, Markets, Players", which includes granular 10-year graphene market forecasts, based on profiles of 90+ key players and leverages extensive in-depth coverage of many end-use markets for graphene.


Graphene is increasingly used in consumer electronics for thermal management, leveraging its exceptional heat conductivity to improve device cooling and reliability. Beyond passive materials, its high surface area and conductivity also make it a strong candidate for energy storage applications, particularly in next-generation batteries and supercapacitors, where faster charging and longer lifetimes are sought.

Material requirements

Pristine graphene produced via chemical vapor deposition (CVD) offers exceptional electronic performance but is costly and complex to scale. As

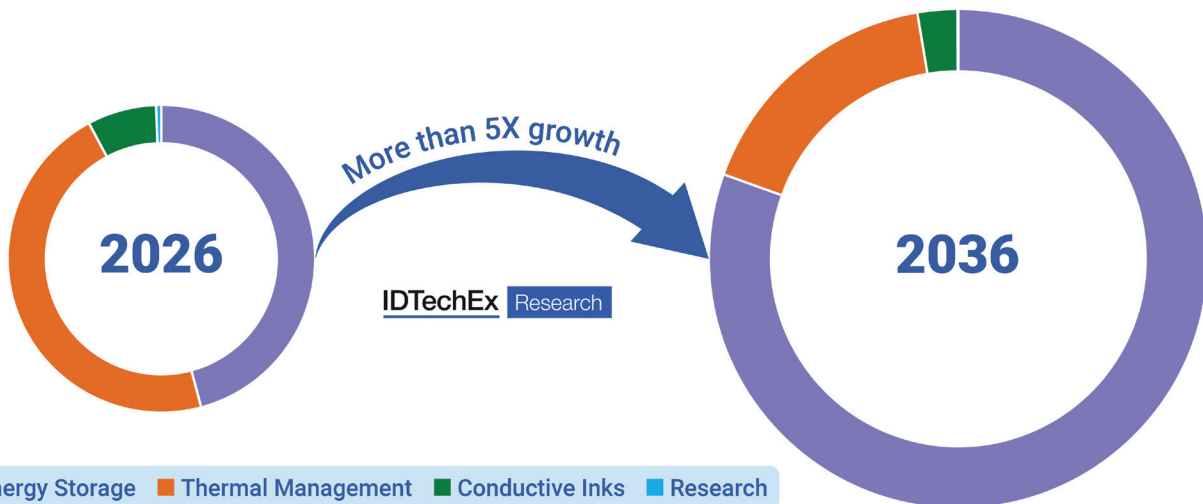
a result, it's primarily reserved for high-end sensing applications. Most other uses such as thermal management films or battery additives can rely on lower-grade, bulk graphene materials without requiring perfect atomic structure.

IDTechEx believes that thermal management applications is a key early area for graphene adoption. There is increasing demand for high-performance heat spreaders, and the graphene properties and morphology are well-suited. The thermal conductivity of carbon materials tend to be higher than that seen in metals, with graphene oxide (GO) reporting a thermal conductivity in excess of 3,000 W/m-K. For anisotropic materials such as graphene, achieving through-plane alignment is key for a thermally efficient system, capable of minimising material usage and enabling cost savings. Multiple routes exist based on the material to be



Small Material, Big Impact: Graphene in Electronic Applications

The **electronic applications of graphene** have received a lot of hype and bordered on science fiction over the past two decades. Tangible success is now being achieved in **thermal management**, while the **next generation of energy storage** devices could be unlocked by graphene. The sector is set to grow by a factor **more than 5X** over the coming decade.



aligned, including mechanical methods and dielectrophoresis. This need for alignment to enhance thermal conductivity is mirrored in applications leveraging the electrical conductivity of graphene, such as when using graphene dispersions as conductive inks.

Overview of the growing importance of graphene in electronic applications over the coming decade. Source: IDTechEx market report “Graphene & 2D Materials 2026-2036: Technologies, Markets, Players”.

China leads the way for consumer electronics

For some years the centre of gravity for graphene has been shifting towards China. This trend has shown up in several indicators including investment levels, patent filing trends, academic publication and in the number and size (at least nominal size) of graphene companies. Key players include The Sixth Element, Leadernano and SCF Nanotech.

The use of graphene for thermal management has increasingly been observed across the leading Chinese smartphone manufacturers as higher performances are demanded. This varies from acting as very localised heat spreader, a larger heat spreader or replacing a vapor chamber. A cooling

system, such as vapor chamber (VC) technology, manages a phone’s heat by using evaporation and condensation to dissipate warmth from internal components. This prevents overheating, enhances performance, and extends battery life.

Huawei has taken this concept further with the Mate X6 by incorporating graphene sheets. Other leading smartphone players including Realme, ZTE, Xiaomi and OnePlus have reported use of graphene in various heat spreading technologies.

Other consumer electronics to have adopted graphene heat spreaders include gaming-grade laptops and monitors with Acer, ViewSonic and Samsung Display active in this area. Beyond heat spreaders, adoption of graphene has also been seen in headphones with players reporting that coating audio drivers with graphene can reduce vibrations and allow powerful bass with enhanced mid- and high-range frequencies.

Household names have released graphene-enhanced headsets, including Logitech, and LG, with graphene typically featuring in high-end headsets. Philips bucked this trend when the mid-range H8000E was launched in January 2025, featuring graphene-coated 40 mm drivers.

Will graphene achieve success in energy storage?

The term “graphene battery” is widely misused. In reality, graphene is more likely to be one of several conductive additives selected based on cost, morphology, and performance, potentially used at the electrode, current collector, or pack level. Today, it is rarely found in commercial lithium-ion cells, with limited success mainly in high C-rate consumer devices. In contrast, carbon nanotubes are already established in current cathode formulations with well-established supply chains for lithium-ion batteries.

Looking ahead, graphene’s best opportunity within energy storage lies in enabling next-generation technologies such as silicon anodes, which are expected to scale commercially. However, it is only one contender in an intensely competitive, well-funded landscape where many alternatives also offer effective conductive networks. If graphene is the enabling technology that unlocks silicon anodes, then strong and significant growth will be seen for this material.

For more details on the graphene market, including granular 10-year forecasts for 18 different application areas, see the IDTechEx market report “Graphene & 2D Materials 2026-2036: Technologies, Markets, Players”.

Powering the AI era: how Infineon Technologies' SiC solutions are shaping the future of data centre power supplies

As artificial intelligence (AI) continues to redefine the frontiers of computational performance, the supporting infrastructure particularly the power conversion within data centers must evolve rapidly to meet new performance and density benchmarks

BY DEEPAK VEEREDDY, DIRECTOR OF TECHNICAL MARKETING, INFINEON TECHNOLOGIES AND ALEKSEI CHERKASOV, SENIOR MANAGER PRODUCT MARKETING, INFINEON TECHNOLOGIES

WITH data center power levels rising exponentially due to the proliferation of GPU- and TPU-based processing systems, it is crucial to use new semiconductor materials like silicon carbide (SiC) and gallium nitride (GaN) to complement a well-established silicon (Si) infrastructure.

Today Infineon Technologies is leading the charge to future-proof power conversion infrastructure by deploying its latest generations of all three


technologies making the most of each to support the extreme demands of AI-driven data center architectures.

Each technology has its own strengths and hence is positioned for power supply design as follows [1]:

Silicon is:

- Perfect for legacy or conventional designs
- Optimized cost / performance solutions
- Suitable for traditional topologies up to 150 kHz switching





➤ Figure 1. Infineon Technologies' PSU reference designs span across a wide range of data center needs.

Silicon carbide is:

- Perfect for half-bridge (HB) hard-switching topologies like totem-pole PFC requiring high voltages ($\geq 400V$) and high temperature operating conditions
- Suitable for resonant topologies like LLC up to 500 kHz

Gallium nitride is:

- Perfect for soft-switching topologies such as LLC, for very high frequency (≥ 500 kHz) operation
- Best in efficiency x power density FOM

As we consider the diverse needs of various topologies, it becomes apparent that a single technology is not sufficient to meet all requirements. This is where the value of a hybrid approach comes into play, combining the strengths of Si, SiC, and GaN devices in a single system to create a tailored solution.

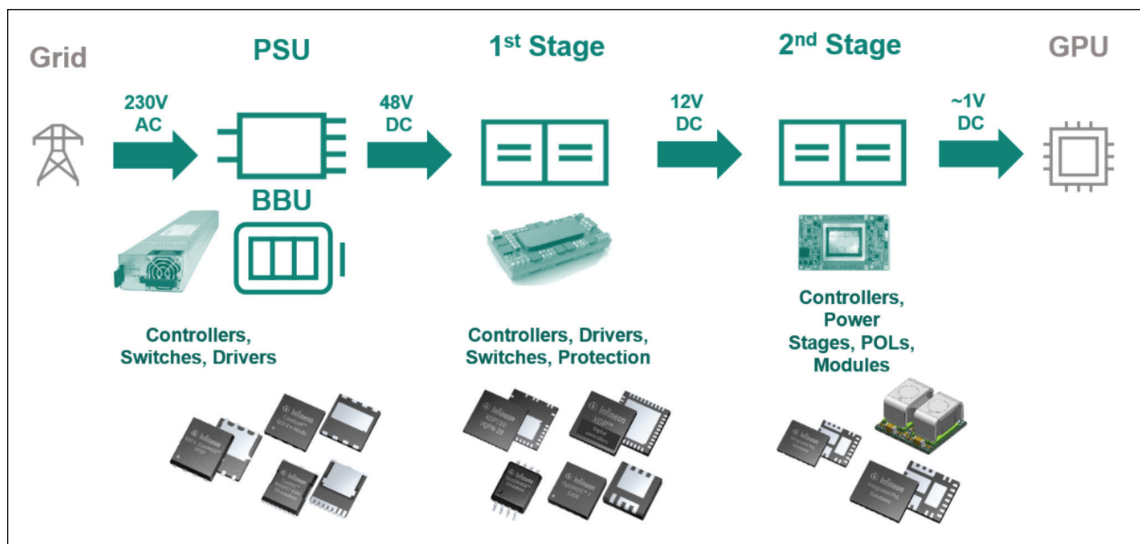
By leveraging this hybrid approach, engineers can achieve rapid scaling and optimal performance, making it an attractive option for developing modern power supply unit (PSU) designs for AI data centers. Having access to a broad portfolio of technologies

and expertise, such as that offered by Infineon Technologies, is crucial to create innovative and efficient PSU designs that meet the unique demands of AI data centers.

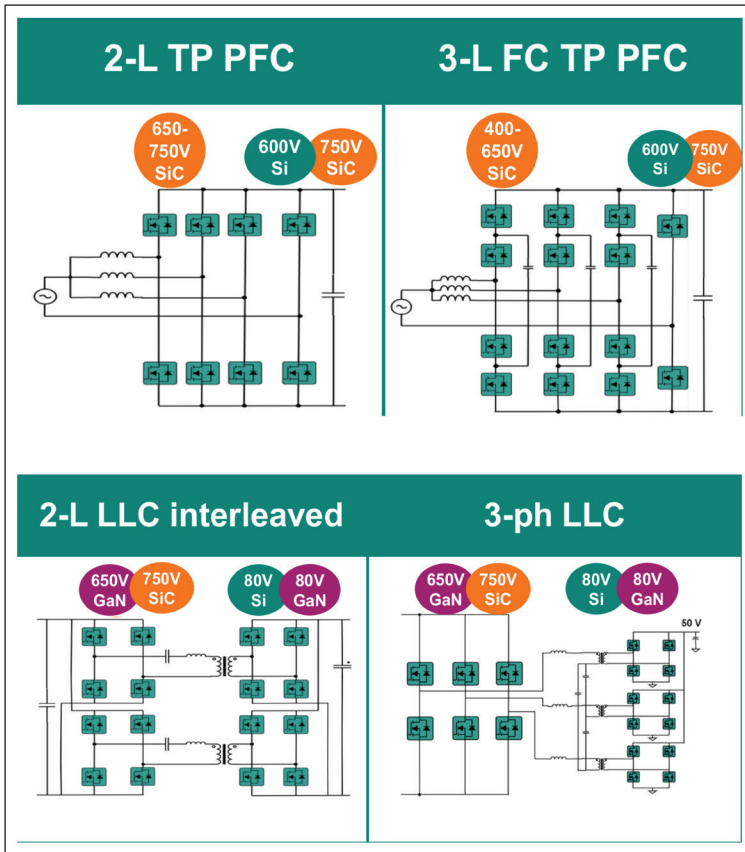
Infineon's AI centric SMPS demo solutions roadmap (Figure 1) reflects such approach that is engineered to factor in the rugged realities of 24/7 data center operation, offering full turn-key solutions. They are a blueprint for achieving the highest efficiency and density targets.

While a hybrid approach combining Si, SiC, and GaN devices is essential for meeting the diverse needs of AI data centers, silicon carbide CoolSiC™ MOSFETs are a critical component in this ecosystem, and in this article, we will discuss their specific applications and benefits in modern power supply unit (PSU) designs.

The AI revolution is reshaping the design criteria of data center power systems. Unlike conventional server environments that were CPU-centric and operated in relatively predictable load conditions, AI servers operate in highly dynamic states with rapid current surges and higher sustained power demands.



➤ Figure 2. Server power flow and Infineon Technologies' comprehensive product portfolio to address each stage from grid to core in 48V bus architecture.



► Figure 3. Common PFC and LLC topologies addressed with SiC MOSFETs.

In fact, racks that once consumed under 10 kW are now pushing beyond 300 kW per rack and even above 500 kW in specialized training clusters. SiC based semiconductor solutions are the undisputed choice for such high power rack designs. And here comes the question – why silicon carbide?

SiC devices are built on a wide bandgap semiconductor material with a bandgap of 3.26 eV—more than three times that of silicon. This physical characteristic provides the following intrinsic advantages for SiC based MOSFET:

- High breakdown field strength allowing for thinner drift regions and thereby reducing

specific on-resistance

- High thermal conductivity (3-5x better than Si) improving heat dissipation
- Low switching losses allowing high frequency operation and reducing magnetics size

The low reverse recovery charge (Q_{rr}) of the intrinsic body diode of SiC MOSFET is a decisive differentiator in fast commutation legs, particularly in bridgeless totem pole PFC. The combination of low gate charge (Q_g) and reduced output charge (Q_{OSS}) broadens zero voltage switching (ZVS) windows in resonant topologies like LLC, enabling operation in the 200–250 kHz range without disproportionate penalties in magnetics or EMI.


From reliability standpoint, SiC MOSFETs exhibit robust avalanche and short circuit capability and a relatively lower positive temperature coefficient of R_{DS(on)} than Silicon MOSFETs. These features simplify current sharing mechanisms in parallel configurations common to multi kilowatt shelves. All these device level benefits align directly with the system level objectives of AI power design: higher efficiency across the load range, smaller and lighter passives, predictable thermal behavior under elevated coolant temperatures, and resilience to abnormal events without compromising on quality and reliability.

From an economic standpoint, SiC enables:

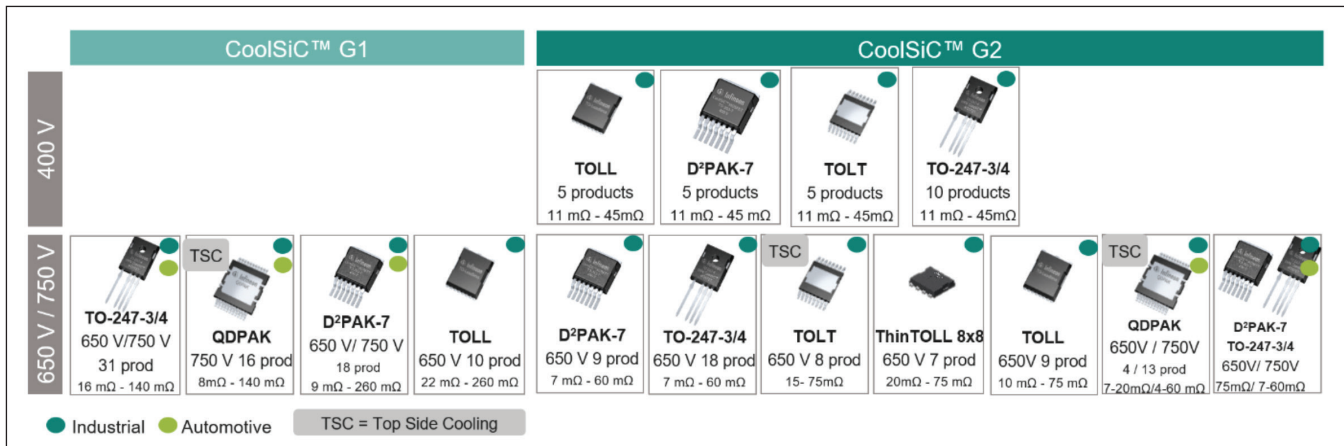
- Smaller power stages and fewer passive components, reducing BOM and footprint
- Higher reliability, minimizing service costs and extending product life
- Superior performance under extreme conditions

The growth of SiC in AI is very well supported by Environmental, Social, and Governance regulations as well as performance and total cost of ownership metrics:

- Hyperscalers like Google, Meta, and Microsoft have set aggressive energy efficiency targets for AI workloads
- Open Compute Project (OCP) standards, such as ORv3, require modular, high-efficiency PSUs
- EU and U.S. energy regulations (e.g. Ecodesign Directive, DOE Level VI) demand stringent



► Figure 4. 5.5 kW three-phase interleaved LLC converter: reference design and efficiency curve.



► Figure 5. Infineon Technologies' CoolSiC™ MOSFET 400-750V product families.

- efficiency curves across load ranges
- Edge data centers and modular deployments favor compact, high-density designs achievable only with wide bandgap (WBG) semiconductors

This explosive growth poses an immediate and long-term engineering challenge: how to efficiently convert and distribute power from high-voltage AC grid inputs to ultra-low voltage rails (often under 1V) that feed modern AI processors. Moreover, in future, these conversions may happen within space-constrained, thermally stressed environments like liquid-cooled environments.

To meet the current demands of 48V architecture (Figure 2), Infineon Technologies has comprehensive offerings that addresses every node of the power path – from the power grid interface all the way to the processor core. The building blocks of this ecosystem include power switches, gate drivers, controllers, power stages, points of load, and power modules.

With PSUs being at the frontier of power conversion flow in servers, Infineon Technologies' CoolSiC™ MOSFET families provide the building blocks to support most advanced AI data center architectures.

PSU manufacturers have been widely adopting SiC in Totem Pole Power Factor Correction (TP PFC) designs, where SiC's fast recovery possible by low Q_{rr} characteristics eliminate the need for slow, lossy boost diodes. Compared to traditional bridged PFC, TP PFC offers 1–1.5% higher peak efficiency in continuous conduction mode (CCM) operation.

The TP PFC stage of the PSU design could involve a 2-level or a 3-level flying capacitor (FC) topology (Figure 3) with or without interleaved approach. 2-level TP PFC is relatively simpler to implement than a 3-level FC design. However, the 3-level FC TP PFC boosts efficiency and density further compared to 2-level approach [2].

In ~400input systems, LLC converters (Figure 3) using SiC MOSFET can operate at up to 500 kHz frequency leading to minimal system loss for the following reasons.

The conduction losses of primary FET are the dominant portion of its overall loss. SiC MOSFET's relatively lower $R_{ds(on)}$ temperature co-efficient results in relatively lower conduction losses at high temperature operation.

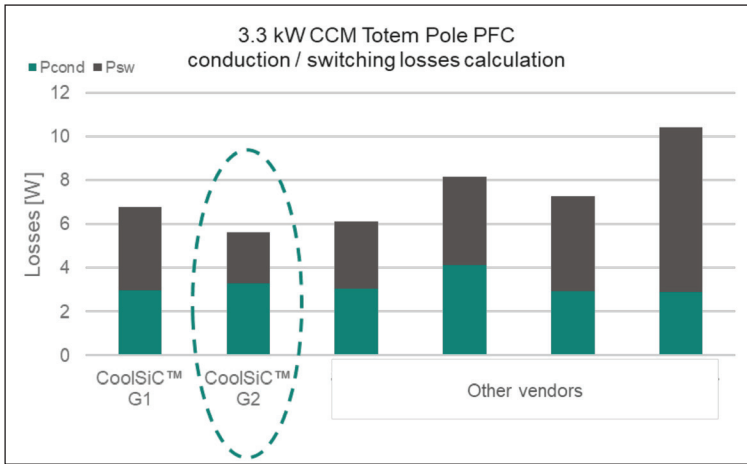
The second dominant loss in the primary FET of LLC is the turn-off loss. SiC MOSFET, being a fast-switching device, enables lower turn-off losses. In addition to the above, the gate driving loss contribution of SiC MOSFET is low due to its lower gate charge compared to Si MOSFETs.

For instance, Infineon Technologies' 5.5 kW LLC reference design [3] using Gen 2 650V CoolSiC™ MOSFETs in TOLT package achieves almost 99% peak efficiency (Figure 4).

Today Infineon Technologies offers one of the most comprehensive SiC MOSFET offerings in the market and recently has expanded the product families with Generation 2 across 400 V, 650 V, and 750 V voltage classes (Figure 5).

Thus, 400 V CoolSiC™ MOSFET family shows its advantages the best in the multilevel topologies [4]. In 3L Flying Capacitor Continuous Conduction Mode

The AI revolution is reshaping the design criteria of data center power systems. Unlike conventional server environments that were CPU-centric and operated in relatively predictable load conditions, AI servers operate in highly dynamic states with rapid current surges and higher sustained power demands. In fact, racks that once consumed under 10 kW are now pushing beyond 300 kW per rack and even above 500 kW in specialized training clusters



➤ Figure 6. 650 V CoolSiC™ MOSFET Gen 2 conduction/switching losses in 3.3 kW CCM TP PFC.

(CCM) Totem Pole topology 400 V SiC provides:

- Lower PFC ripple current leading to lower Root Mean Square (RMS) currents and Total Harmonic Distortion (THD) compared to Triangular Current Mode (TCM) mode.
- Easier control, sensing and electromagnetic interference (EMI) filter design with fixed-frequency operation compared to variable-frequency operation in TCM mode
- Effective frequency on the 3L-PFC inductor is double that of the device switching frequency. Due to the multi-level voltage output of the HF-leg, volt-seconds on the inductor is further reduced. This leads to a potential 4x shrink in PFC-inductor size
- HF-leg switches block half the DC-link voltage, leading to lower switching losses and lower dv/dt swings for a reduced EMI emissions enabling a reduction in the EMI filter size

650 V CoolSiC™ MOSFET Gen 2 devices (when compared to Gen1) are engineered for higher switching speeds, lower gate charge, reduced reverse recovery charge and output capacitance/charge values [5] — making them ideal for demanding hard-switching topologies (Figure

6). These enhancements directly translate into improved efficiency and thermal behavior.

For hardware design manufacturers, the benefit lies in reduced cooling requirements, greater design freedom, and the ability to hit higher power density targets without compromising on system reliability.

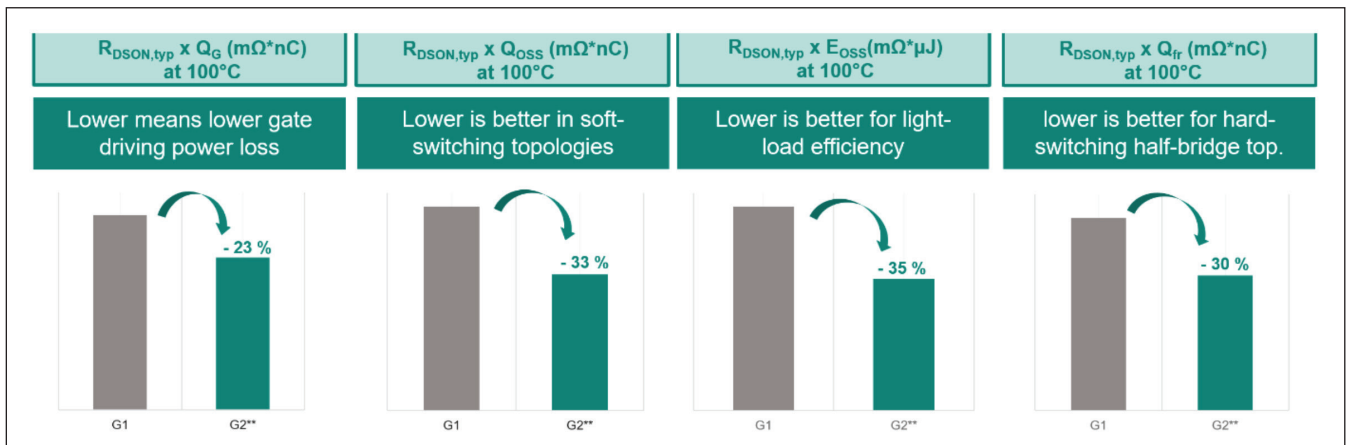
In short, the CoolSiC™ 650V Gen 2 family fits well in the building blocks to support today's most advanced AI data center power supply architectures encompassing both hard-switching (PFC) and soft-switching (LLC) topologies as described before. Infineon Technologies' 750V CoolSiC™ MOSFET portfolio is engineered for higher voltage and harsher operating conditions, making it ideal for both industrial and high-end server PSU applications [6]. Gen 1 MOSFETs offer excellent switching performance and ruggedness, while Gen 2 pushes the performance frontier even further.

With RDS(on) values as low as 4 mΩ, these MOSFETs deliver ultra-low conduction losses, essential for efficient high-power systems. Available in through-hole (THD), surface mount device (SMD), as well as SMD top-side cooled (TSC) like Q-DPAK package, this portfolio offers design flexibility.

Gen 2 devices have industry leading lowest specific on-resistance, lower gate charge (QG), output capacitance (QOSS), and output capacitance stored energy (EOSS), translating into faster switching and reduced switching losses. These characteristics are critical for high-frequency, soft-switching topologies where fast transition speeds and minimal dead time improve both efficiency and system density.

Moreover, these devices are being adopted in everything from electric vehicle fast chargers to AI data center PSUs. Their robustness, efficiency, and thermal characteristics are redefining what's possible at the 8 – 30 kW PSU level.

One of the key features that brings CoolSiC™ Gen 2 on top is the .XT [7] interconnection technology offered across all voltage classes, which marks



➤ Figure 7. 750 V CoolSiC™ MOSFET Gen 2 figures of merit vs 750 V Gen 1.

From reliability standpoint, SiC MOSFETs exhibit robust avalanche and short circuit capability and a relatively lower positive temperature coefficient of RDS(on) than Silicon MOSFETs

a significant leap in packaging and thermal performance.

Traditional solder-based interfaces exhibit limitations in thermal resistance and mechanical robustness under high current cycling. The .XT uses a diffusion solder process to create a highly thermally conductive bond between the die and the lead frame. This reduces RthJC (thermal resistance junction-to-case) by up to 25%—significantly enhancing thermal performance and allowing for more compact and reliable system designs.

This advancement is particularly critical in AI data centers where power density is climbing up and thermal margins are shrinking. The .XT also improves reliability under power cycling, reducing the risk of joint fatigue or thermal cracking over time. For system designers, this means longer operational life, less frequent maintenance, and fewer cooling constraints. The powerful combination of CoolSiC™ MOSFETs and .XT technology enables PSUs that are not just smaller and cooler, but also significantly more robust—exactly what’s needed in mission-critical AI infrastructure.

Nowadays environmental impact of power efficiency improvements is often overlooked. However, deploying 650V CoolSiC™ Gen 2 in the PFC stages of all global data centers would save approximately 2.9 terawatt-hours of electricity over five years [1]. That equates to 1.35 million metric tons of CO₂ emissions avoided—the same as removing nearly 300,000 cars from the road.

These calculations are not based on optimistic projections; they are grounded in real operating profiles, mission loads, and validated efficiency gains from production-ready reference boards of Infineon Technologies. As global electricity prices tend to rise with power hungry data centers, and carbon regulation intensifies, these savings become more than a “nice to have”—they are operational imperatives. Today Infineon Technologies helps customers meet Environmental, Social and Corporate Governance targets while also enhancing their bottom line.

Through energy-efficient CoolSiC™ MOSFETs, we contribute to a cleaner, greener digital infrastructure. Sustainability and performance no longer need to be tradeoffs — they go hand in hand. The data centers are no longer a passive utility — it is the active brain of the digital economy. As AI accelerates toward exascale compute, the infrastructure behind it must evolve. Silicon carbide MOSFETs provide the building blocks for that evolution. With unmatched performance in efficiency, density, and reliability, 400V, 650V, and 750V SiC devices are now essential for AI server power conversion stages of the existing 48V and the upcoming 800V or ±400V bus architectures.

Infineon Technologies’ full stack SiC strategy, supported by industry partnerships and deep application expertise, ensures that power electronics keep pace with computing demands. In this new era of power-hungry, thermally constrained, and environmentally conscious compute, the message is clear: AI cannot run without SiC MOSFETs.

FURTHER READING

- [1] A. Cherkasov; “SiC-based power solutions for AI: Enabling next-generation server power supplies with unparalleled efficiency and reliability”; Power Electronics International, 2025
- [2] 3.3 kW high-frequency, high-density PSU with three-level flying capacitor PFC for server and data center applications, Application note, 2025 - <https://www.infineon.com/assets/row/public/documents/24/42/infineon-reference-design-ref-3k3w-3lfc-psu-applicationnotes-en.pdf>
- [3] 5.5 kW three-phase interleaved LLC converter with 650 V CoolSiC™ TOLT, Application note, 2025 - <https://www.infineon.com/assets/row/public/documents/24/42/infineon-evaluation-board-eval-5k5w-3ph-llc-sic2-applicationnotes-en.pdf>
- [4] CoolSiC™ MOSFET 400 V and 440 V G2 - <https://www.infineon.com/content/dam/infineon/row/restricted/documents/24/296/disti/infineon-coolsic-mosfet-400-v-g2-customerconnector-en.pptx>
- [5] CoolSiC™ MOSFET discrete 650 V G2 Industrial, 2025 - <https://www.infineon.com/content/dam/infineon/row/restricted/documents/24/156/customer/infineon-coolsic-mosfet-650v-g2-productpresentation-en.pptx>
- [6] CoolSiC™ 750 V G2 Industrial MOSFET, Application note, 2025 - <https://www.infineon.com/assets/row/public/documents/24/42/infineon-coolsic-750v-g2-industrial-mosfet-applicationnotes-en.pdf>
- [7] .XT interconnection technology - <https://www.infineon.com/technology/xt>



INTERNATIONAL CONFERENCE

ANGELTECH

20-22 APRIL 2026
BRUSSELS
BELGIUM

CONNECTING, EDUCATING & INSPIRING THE POWER ELECTRONICS INDUSTRY

REGISTER NOW: [PEINTERNATIONAL.NET/REGISTER](https://peinternational.net/register)

Power Electronics International 2026 returns for its 4th edition as part of the leading series of senior-level chip level conferences for the global power electronics industry, bringing together executives, end users and technology experts for two days of insight, discussion and collaboration. The conference will focus on four key themes:

- Putting the Automotive Industry in Pole Position
- Ensuring a Renewable, Sustainable Future
- Delivering AI-driven Data Centre Densification
- Engineering Innovation with WBG Devices, Circuits and Materials

Co-located with **CS International**, **PIC International** and **Advanced Packaging International**, the event features four focused conference programmes and a shared exhibition hall.

With over **800 attendees** and more than **80 exhibitors**, the event also includes two evening networking drinks receptions, providing valuable opportunities for conversation, collaboration and business development.

POWER ELECTRONICS SPONSORS

PLATINUM



STARTUP



SUPPORTING PARTNER



OPENING RECEPTION




REGISTRATION



MEDIA PARTNERS



2026 SPEAKERS INCLUDE



ORGANISERS



Sponsorship / Exhibition / Speaker OPPORTUNITIES AVAILABLE!

Sponsorship, exhibition, and speaker opportunities at PE International, enable and empower you to educate and penetrate the global power electronics industry using multiple marketing tools.

Demand is high, with sponsorship, exhibition and speaking slots already being secured for 2026, so to avoid disappointment, please get in touch today.



10% off registration offer by scanning the code



Key Themes for 2026

Putting the Automotive Industry in Pole Position

What are the most effective ways to improve reliability, efficiency, and fast-charging for electric vehicles while managing thermal and wide bandgap challenges?

Ensuring a Renewable, Sustainable Future

Can greater efficiency in power conversion and better integration of renewables lead to a more stable and low-carbon energy grid?

Delivering AI-driven Data Centre Densification

What solutions will deliver reliable, high-density power for AI data centres as space and thermal demands grow?

Engineering Innovation with WBG Devices, Circuits and Materials

Will advances in SiC and GaN technologies finally break through technical and supply chain barriers for wide bandgap adoption?



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



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



Addressing the protection challenges of 48V AI servers using hot-swap controllers

Enabling the design of a reliable input protection solution for a 48V AI server.

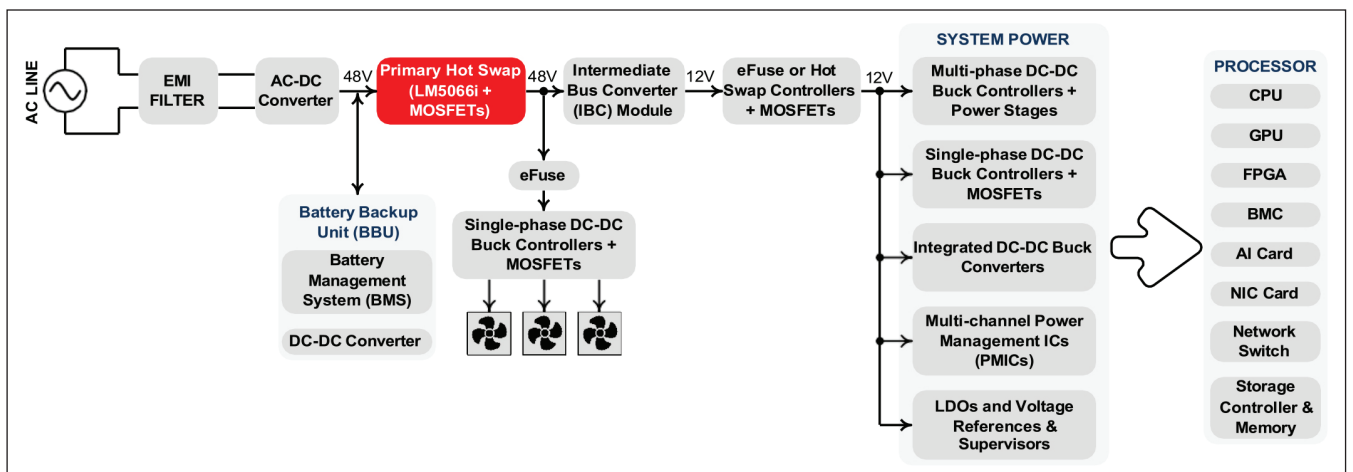
BY AVISHEK PAL, POWER SWITCHES AT TEXAS INSTRUMENTS AND RAKESH PANGULOORI, POWER SWITCHES AT TEXAS INSTRUMENTS

WITH ADVANCEMENTS in artificial intelligence (AI) and machine learning, enterprise servers have become extremely power-hungry as they simultaneously process a large amount of data and storage. The steady-state power rating of each server

motherboard has gone up to 5kW or 6kW, in contrast to 1kW or 2 kW for general servers. The form factor remains the same, however, which imposes system design challenges given the increased power density. The load amplitude, slew rate and frequency

of transient loads on AI servers have increased three to four times compared to general servers.

Figure 1 shows a typical power distribution in a 48V rack server where the input is protected by the hot-swap



➤ Figure 1. Typical block diagram of a 48V rack server power distribution.

circuit - and then distributed to all downstream system loads.

In this article, we'll discuss various challenges that AI-based processors bring into 48V server designs, along with design guidelines and important tips and tricks for the design and layout to achieve a reliable hot-swap solution for the system specifications outlined in Table 1.

Challenges in designing a hot-swap circuit for a 48V AI server
It's interesting to look at how hot-swap circuit configurations have evolved over the years. A hot-swap solution consists of three main components: a N-channel metal-oxide semiconductor field-effect transistor (MOSFET) that serves as the main power control switch; a sense resistor that measures the current; and the hot-swap controller, which includes a current-sense amplifier completing the loop to control the MOSFET's pass current.

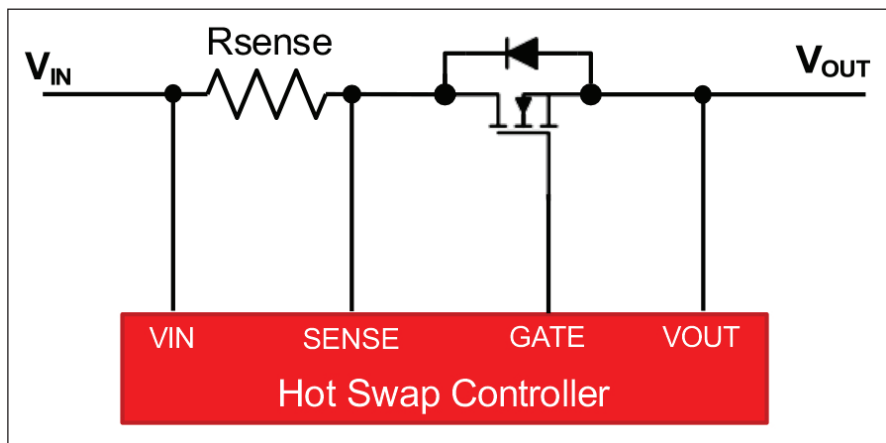
As shown in Figure 2, you can use a single MOSFET-based hot-swap solution for low-power designs. Fundamentally, the hot-swap controller comes with current- and power-limiting functionalities to limit the inrush and fault currents while ensuring the MOSFET's safe operating area (SOA). These functionalities are good enough to design low-power (<500W) hot-swap solutions.

With the increase in digital load, the system needs a higher output capacitance (>470µF), requiring parallel MOSFETs to support steady-state current and the adoption of output-voltage slew-rate control [1] to keep the MOSFET within its SOA.

In the output-voltage slew-rate control method, capacitor C_{dv/dt} placed across GATE-GND (see Figure 3) limits the slew-rate of the gate and output voltages, which limits the inrush current. MOSFETs can handle more energy

Design Parameter	Value
Input voltage range	40V to 60V
Output capacitance	4.2mF
Steady-state thermal design power rating	6kW
Transient power rating	8kW at 400µs
Transient load profile	15% to 100% of the transient power rating with a 10% duty cycle
Load slew rate	>2A/µs
Frequency of transient load	>1kHz

➤ Table 1. Typical system specifications.



➤ Figure 2. Traditional power-limiting hot-swap circuit.

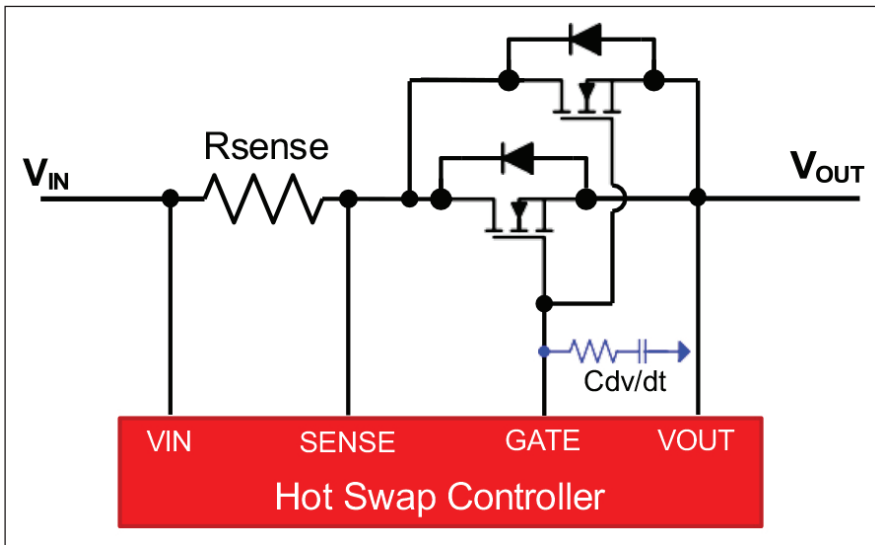
when the power dissipation in them is reduced and spread over longer durations. Therefore, as the output capacitance increases, you need a higher C_{dv/dt} to reduce both the inrush current and power dissipation in the MOSFET during startup.

A higher C_{dv/dt} interferes with the turnoff process, however, the hot-swap controller has limited pulldown strength. This necessitates a local P-channel N-channel P-channel (PNP)-based discharge circuit for C_{dv/dt}, as shown in Figure 4. During startup, C_{dv/dt}

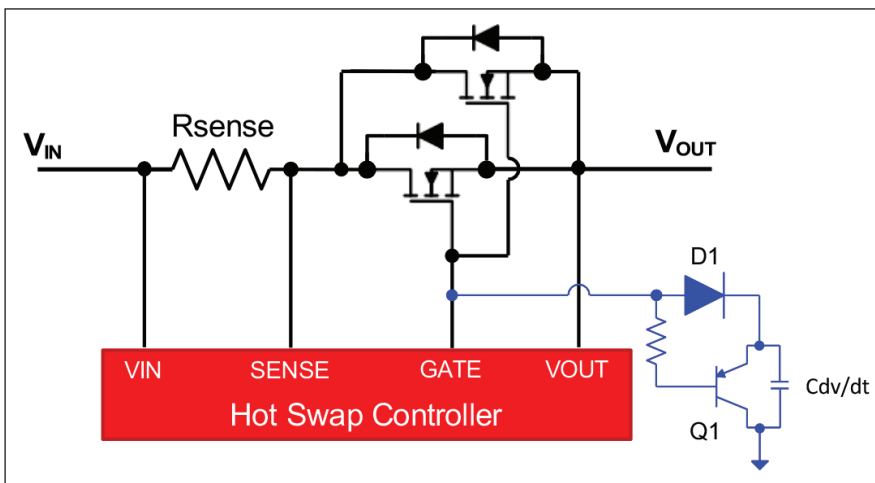
controls slew-rate in the same way, but during a turnoff event, the Q1 PNP transistor activates and discharges C_{dv/dt} locally. Diode D1 blocks the discharge of C_{dv/dt} into the GATE pin, which reduces the stress on the GATE pin and also ensures proper operation of the controller.

In AI-powered graphics processing unit applications, the hot-swap solution has to support currents around 150A and must support high-frequency, high slew-rate load transients, which present three new challenges.


The hot-swap gate node is decoupled from the MOSFET gate terminal by placing the DSS diode between them. This modification helps eliminate the reflection of output voltage ripple to the hot-swap controller GATE node and avoids false turn-on of the soft-start PNP transistor, Q_{ss}.



➤ Figure 3. Hot-swap circuit with GATE slew-rate control.



➤ Figure 4. Hot-swap circuit with local discharge path for Cdv/dt.



➤ Figure 5. Short-circuit response of the LM50661 controller with eight MOSFETs.

Challenge No. 1: Turnoff delay during an output short-circuit

With the increase in load current, more MOSFETs need to be parallel to limit the maximum steady-state MOSFET junction temperature to a safe value (100°C to 125°C).

For example, to support a steady-state load current of 150A at an ambient temperature of 70°C, eight Texas Instruments (TI) CSD19536KTT MOSFETs need to be in parallel to limit the steady-state MOSFET junction temperature to 100°C. Paralleled MOSFETs help thermally, but increase the effective capacitance on the GATE pin of the hot-swap controller and impact the turnoff response.

During an output short-circuit, the MOSFETs need to turn off fast enough to prevent further buildup of fault current and avoid damage to the MOSFETs, input power supply, or printed circuit board (PCB). The gate pull-down strength of the TI LM50661 hot-swap controller is limited to 160mA, which is not enough to turn off all eight MOSFETs completely during a short-circuit event, as shown in Figure 5.

Challenge No. 2: False gate turn-off during a load transient

Although the local PNP-based discharge circuit for Cdv/dt helps reliably turn-off the MOSFETs during an output short-circuit event, it causes a false GATE turn-off in the presence of high-frequency, high slew-rate load transients. During load step-up, the MOSFET source node drops because of the finite input and output impedances of the hot-swap circuit. The voltage drop at the source node gets coupled to the MOSFET gate node through the CGS capacitance of the MOSFET and causes the gate node to drop as well. The MOSFET source node recovers during load step-down. The gate node cannot recover completely to its previous level, because of the limited gate current (20µA typical) of the LM50661 hot-swap controller.

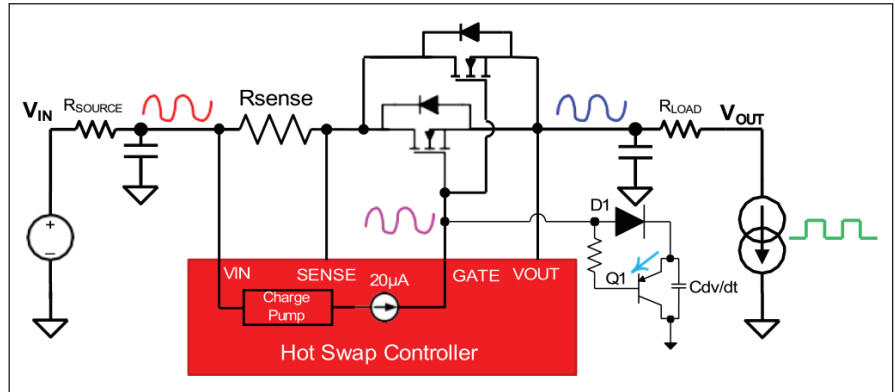
As a result, the hot-swap controller gate continues to drop further in the subsequent load transient cycles developing the base-emitter voltage for Q1. Finally, PNP bipolar junction transistor Q1 turns on, and falsely shuts down the system. Figure 6 illustrates the whole process, while Figure 7 shows the corresponding test result.

Challenge No. 3: Parallel resonance during controlled (slow) turn-on

Generally, parallel MOSFETs are more prone to parasitic oscillations than a single MOSFET in the linear region of operation. This is because of the presence of parasitic stray package inductances and capacitances on the drain, source and gate nodes, which form a resonant tank circuit resembling a Colpitts oscillator. Unlike switching regulators with a gate-drive strength of >2A, hot-swap controllers with a lower gate-drive strength (20µA) limit the inrush current during start-up by operating the MOSFETs in the linear region.

As a result, the parallel combination of hot-swap MOSFETs is highly susceptible, with more chance of generating sustained oscillations. This phenomenon causes the violation of the MOSFET SOA during a power-into-short fault, leading to MOSFET damage.

Proposed circuit enhancements
Let's discuss circuit enhancements to help solve these three challenges.



➤ Figure 6. Illustration of a hot-swap circuit for a dynamic load.

Improving the turn-off response

In the proposed solution shown in Figure 8, introducing an external fast pull-down circuit using a PNP transistor (QPD and RPD) will boost up the turn-off speed. During an output short-circuit event, the gate pull-down current of 160mA creates a substantial voltage drop across the RPD resistor and enables fast pull-down of the PNP transistor (QPD). This in turn shorts the gate-to-source of all parallel MOSFETs,

turning off the MOSFETs immediately to quickly disconnect the power path. Figure 9 shows the experimental result for a short-circuit event with a fast pull-down circuit.

Overcoming false turn-off for dynamic loads

In this solution, the hot-swap gate node is decoupled from the MOSFET gate terminal by placing the DSS diode between them, again shown in Figure



Enabling AI at the Edge

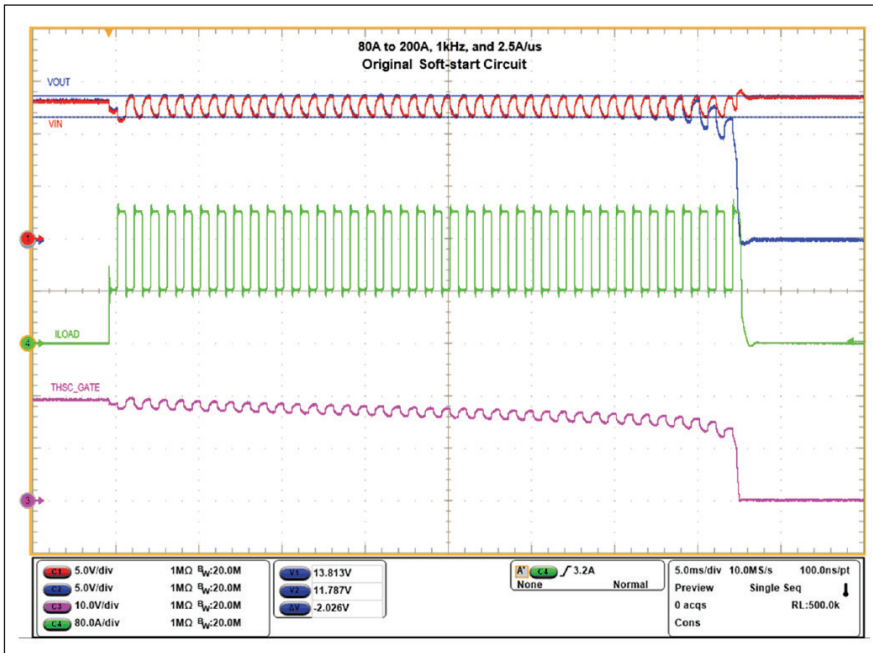
Optimized Hardware for Machine Learning Models

Artificial Intelligence (AI) and Machine Learning (ML) are transforming industries by enabling smarter, more autonomous systems. However, deploying AI at the edge—where real-time decision-making is essential—requires highly efficient, low-power solutions tailored for embedded environments. Microchip Technology is leading the way in this transformation with a comprehensive suite of processors, FPGAs, and development tools that make it easier than ever to bring intelligence to the edge.

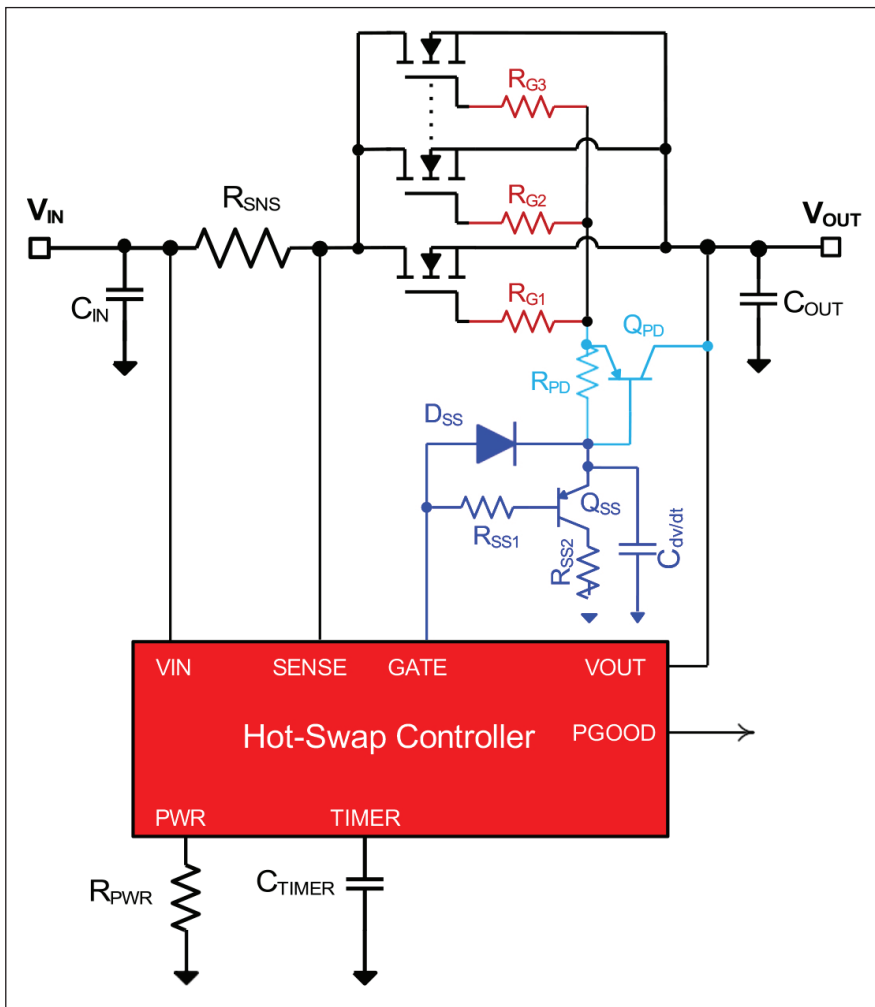
Across critical sectors like **smart factories, medical diagnostics, automotive systems, construction equipment, agriculture, and smart cities**, edge AI is unlocking real-time insights and automating tasks where cloud connectivity is limited or latency is unacceptable. Whether it's predictive maintenance on a factory floor, patient monitoring in a hospital room, object detection on a construction site, or energy-efficient traffic management, Microchip's solutions empower developers to deploy ML where it's needed most.



The Microchip name and logo and the Microchip logo are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries. All other trademarks are the property of their registered owners.
© 2025 Microchip Technology Inc. All rights reserved.
MEC2629A-UK-09-25



➤ Figure 7. Response of a hot-swap circuit to a dynamic load.



➤ Figure 8. Proposed hot-swap circuit configuration.

8. This modification helps eliminate the reflection of output voltage ripple to the hot-swap controller GATE node and avoids false turn-on of the soft-start PNP transistor, Qss.

Changing the position of the diode does not impact controller behavior during start-up nor any of the fault events. As shown in the test result (see Figure 10), the system operates continuously even for large load steps from 20A to 120A at a 1kHz frequency.

Damping parasitic oscillations

Adding a damping resistor (RG1, RG2, RG3) in series with the gate of each MOSFET can eliminate the parasitic oscillations in the system.

Usually, we recommend a 10Ω 0603 package resistance, but based on the parasitics, a low value around 1Ω may also help. We suggest testing on your PCB and deciding the value of the damping resistor.

Design guidelines and component selection

Reference [1] iterates a procedure for designing a hot-swap circuit to protect the system and MOSFETs. We recommend reviewing Reference [1] to become familiar with the design.

Feeding the system specifications shown in Table 1 into the LM50661 design calculator will obtain the values of the current-sense resistor (R_{SNS}), power-limiting resistor (R_{PWR}), fault timer capacitor (C_{TIMER}), soft-start capacitor ($C_{dv/dt}$) and number (N) of selected MOSFETs to parallel. In the 8kW Hot-Swap Reference Design for 48V Artificial Intelligence Servers [2], $R_{SNS} = 330\mu\Omega$, $R_{PWR} = 28.7k\Omega$, $C_{TIMER} = 10nF$, $C_{dv/dt} = 47nF$ and $N = 8$. Looking at Figure 8, select the R_{PD} resistor using $R_{PD} > (1)$

$$R_{PD} > \frac{V_{BE(sat)}}{I_{GATE(CB)}}$$

You will need output Schottky diodes to protect the output pin of the hot-swap controller against a negative where, $V_{BE(sat)}$ is the base-emitter saturation voltage of the QPD PNP transistor and $I_{GATE(CB)}$ is the power-on reset circuit-breaker sink current in the LM50661 hot-swap controller. The 8kW hot-swap reference design uses an R_{PD} value = 20Ω.

C_{dv/dt} discharge circuit

Figure 8 uses a 100V signal diode for DSS. The diode should handle a few tens of milliamperes of forward current. The 8kW hot-swap reference design uses the BAV16W-7-F from Diodes Inc.

You will have to select R_{SS1} , R_{SS2} and Q_{SS} iteratively so that none of the three components become stressed during turn-off. For Q_{SS} , you can select any standard PNP transistor with collector-emitter (V_{CEO}) and collector-base (V_{BCO}) voltages of $>100V_{DC}$ and a continuous collector current of $>200mA$.

Select the values for R_{SS1} and R_{SS2} and their respective power ratings to limit the current flowing through the Q_{SS} transistor to a safe value.

You must use a special high-power resistor for R_{SS2} to manage the transient peak power stress during turn-off. The 8kW hot-swap reference design uses onsemi MMBT5401LT1G for Q_{SS} , with $R_{SS1} = 100\Omega$ and $R_{SS2} = 499\Omega$ (the Vishay RCS0805499RFKEA).

Input transient voltage suppression (TVS) diodes are required to protect against transient overvoltages during

input hot-plug and output short-circuit events. The TI TVS diode recommendation tool can help you obtain the part number (voltage and power ratings) of the TVS diode and the number of TVS diodes to parallel.


The 8kW hot-swap reference design uses three Littelfuse 8.0SMDJ60A TVS diodes. For a deeper analysis into TVS diode selection, see Reference [3]. You will need output Schottky diodes to protect the output pin of the hot-swap controller against a negative transient in the event of an output short-circuit event. The 8kW hot-swap reference design uses three onsemi FSV20100V Schottky diodes.

Conclusion

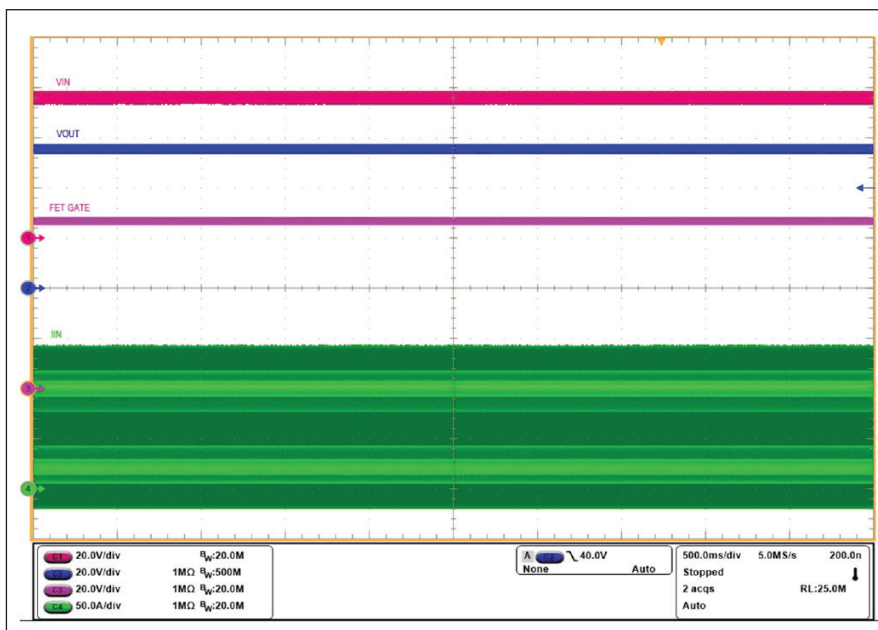
The emerging 48V AI servers demand significantly more power, both in peak and steady states, than traditional servers. The high-power consumption along with fast and transient dynamics impose challenges in designing front-end protection using a hot-swap controller and parallel MOSFETs.

The challenges include fast turn-off of parallel MOSFETs for real faults while avoiding false turn-off for high-frequency transients from the computational load.

The proposed solution in this article eliminates the limitations of legacy hot-swap controllers and enables the design of a reliable input protection solution for a 48V AI server.



➤ Figure 9. Output short-circuit response with fast pull-down circuit.



➤ Figure 10. Load transient performance for steps from 20A to 120A to 20A at a 1kHz frequency.

FURTHER READING

- 1. Rogachev, Artem. “Robust Hot Swap Design.” Texas Instruments application report, literature No. SLVA673A, April 2014.
- 2. Texas Instruments. “8kW Hot-Swap Reference Design for 48V Artificial Intelligence Servers.” Texas Instruments test report, literature No. PMP23496, August 2024.
- 3. Hegarty, Timothy. 2011. “TVS Clamping in Hot-Swap Circuits.” Power Electronics Technology, October 2011.

Related Websites

- LM5066I
- PMP23496

The rise of 800V high voltage DC: can collaboration drive NVIDIA's AI infrastructure vision?

NVIDIA's transition to an 800 VDC architecture for AI data centres is a technological shift that will demand industry-wide alignment on power delivery, semiconductors and standards.

BY REBECCA POOL, TECHNOLOGY EDITOR.

WHEN NVIDIA first announced its 800 V high-voltage direct current architecture for AI servers and data centres in May this year, it wasn't a huge surprise. AI servers, especially those using GPUs and accelerators, were already drawing up to ten times more power per rack than traditional servers, and your 48 and 54 VDC distribution systems were suffering major power losses.

The shift to the higher voltage bus aims to address these issues, raising power efficiency and density in the data centre infrastructure of tomorrow.

As NVIDIA senior power architect, Jared Huntington, and colleagues, said in their blog at the time : "Traditional 54 V in-rack power distribution, designed for kilowatt-scale racks, isn't designed to support the

megawatt-scale racks coming soon to modern AI factories.

[We are] leading the transition to 800 VDC data centre power infrastructure to support 1 MW IT racks and beyond, starting in 2027."

Power distribution in today's data centres involves multiple voltage conversions, introducing inefficiencies



➤ Data centre servers: Infineon is collaborating with NVIDIA on power delivery architecture for future AI server racks. [Infineon]

into the electrical system. But by moving to this new architecture - in which AC grid power can be directly converted to 800 VDC at the data centre - many conversion steps will be eliminated, reducing energy losses during AC/DC and DC/DC transformations. Clearly the 800 VDC shift hinges on industry-wide coordination on voltage ranges, connector interfaces and supporting standards – but from word go, numerous partners pledged support. Early on, Infineon highlighted how it will apply its expertise in silicon, SiC and GaN semiconductors in a collaboration with NVIDIA to ‘create new power delivery standards’ for AI data centres. The Germany semiconductor heavyweight pointed out how the future system architecture will be centralized, making the best possible use of the constraint space in a server rack – increasing the importance of power semiconductor solutions using fewer power conversion stages and allowing upgrades to even higher distribution voltages.

Onsemi, which has long-developed both silicon and SiC semiconductors, and GaN device manufacturer, Innoscience, both trumpeted support for NVIDIA’s transition to 800 VDC power architectures. And Navitas was also quick to announce that its GaN and SiC power semiconductors will support NVIDIA’s Kyber rack-scale systems powering AI-accelerating GPUs, such as ‘Rubin Ultra’. The US firm highlighted that NVIDIA’s 800 VDC architecture is expected to improve end-to-end power efficiency by up to 5%, reduce maintenance costs by 70% and lower cooling costs by directly connecting the HVDC to the IT and compute racks. Other semiconductor providers offering support include Analog Devices, AOS, EPC, MPS, Renesas, Richtek, ROHM, STMicroelectronics and Texas Instruments. Meanwhile additional support has come from power system components partners, such as BizLink, Delta, Flex Power, GE Vernova, Lead Wealth, LiteOn and Megmeet, as well as data centre power system providers, including ABB, Eaton, Mitsubishi Electric and Schneider Electric.

Setting standards


Since the initial reaction to NVIDIA’s plan to lead industry towards to the 800 VDC architecture, even more announcements relating to

next-generation data centres have followed – particularly at the Open Compute Project (OCP) Global Summit, in October. For starters, the OCP - recognised as a crucial forum for establishing open standards for hardware designs and specifications - unveiled its ‘Open Data Center for AI’ strategic initiative, with initial support from Google, Meta and Microsoft. The initiative expands on the OCP’s Open Systems for AI initiative, and aims to develop the necessary standards so the already rapidly-evolving AI data centre infrastructure can be flexibly deployed alongside traditional systems.

According to OCP Chief Innovation Officer, Dr Clifford Grossner, the initiative was a “natural evolution” of work the organisation had started some 18 months ago, knowing that a transition from cloud data centres to AI data centres was coming. “If you look at a typical enterprise, one rack of equipment might consume 20 to 30 kW of power... when we start to move towards architectures with GPUs and building AI clusters, we are already approaching several hundred kilowatts and we can envision at least one megawatt racks in the future,” he explains.

“We also realised that AI workloads aren’t like cloud workloads... and [involve] massively parallel computation so the power draw may go from very small to very large, very quickly,” he adds. “[Factor in] that the AI data centres of the future might draw as much power as a small city, and we knew the time was right to understand the ways that we can bring standardizations here and solve problems in a collaborative fashion.” Tech giants have already contributed several OCP specifications. For example “Diablo” from Google, Meta and Microsoft describes a high-density sidecar rack for powering AI clusters. Meanwhile AMD, Google and Microsoft have co-authored “Hyperscale CPU Reliability, Availability, and Serviceability (RAS) and debug requirements for CPUs”.


Such specifications can only help to accelerate the deployment of AI data centres, but as Grossner emphasises: “We don’t stop at specifications – we want products, and so we’re going to be building the multi-vendor supply chain that can supply this equipment.”



➤ NVLink Fusion enables partners, such as Fujitsu and Qualcomm, to couple custom CPUs with NVIDIA GPUs in a rack-scale architecture to boost AI performance. [NVIDIA]

Many in industry agree the time is ripe for unified industry action on future AI data centre requirements. “With the AI infrastructure market moving very fast, there is a risk of higher costs due to fragmentations,” commented Technology Analyst, Alan Weckel of 650 Group. “It is the right time for an organization like the OCP to be facilitating a community to determine commonalities in data centre facilities and IT infrastructure that can help accelerate the market for future generations of AI cluster deployments and data centre facility builds.”

Weckel’s comments echoed across numerous plenary presentations from tech executives at the OCP Global Summit. Microsoft’s Saurabh Dighe, Corporate VP, Azure Strategic Planning and Architecture, highlighted how colleagues have re-engineered the entire stack from system to silicon to deploy AI, with the design being contributed to OCP. ARM Senior Vice President of Infrastructure, Mohamed Awad, also described how the company has been working with OCP on chiplet system architecture specifications. “It’s unclear to me if we’ll ever get to a truly interoperable chiplet marketplace, but there’s no reason we can’t work together to define a common system architecture,” he said. “This would meaningfully accelerate time to market and innovation... and allow for new dimensions of systems optimization that



► Microchip president and chief executive, Steve Sanghi, with a 3nm wafer: the firm's Gen 6 PCIe switches are manufactured using a 3 nm process. [Microchip]

are too complex for most companies to take on by themselves."

Collaborating on complexity

Plenaries aside, the flurry of products and white papers describing power architecture components and plans from semiconductor players at the OCP Summit signify a meeting of minds on the complexity that is to come, and the will to tackle issues head on. Infineon, EPC, Alpha and Omega Semiconductor, and a host of other industry players announced developments to support the power demands of NVIDIA's 800 VDC architecture.

Texas Instruments debuted a dual-phase smart power stage, a reference design for a 30 kW AI server power supply unit and a GaN intermediate bus converter that can convert energy down from 800 VDC to a lower intermediate bus voltage – all to address the growing power consumption in next-generation data centres. "The transition to an 800 VDC power architecture fosters a renewed discussion about how to design power delivery [within an IT server rack] given the trade-offs in overall conversion efficiency, size and performance," TI stated in its white paper, 'Power delivery trade-offs when preparing for the next wave of AI computing growth'.

In a similar vein, Renesas outlined how GaN semiconductors can benefit the data centres of tomorrow, describing how power devices can be used in new architectures as well as sidecar AC/

DC racks, for efficient, high-frequency power conversion at the higher voltages and power densities. The Japan-based semiconductor supplier also highlighted how its silicon REXFETs can support future designs. "[We are] helping power the future of AI with high-density energy solutions built for scale, supported by our GaN FETs, MOSFETs, controllers and drivers," said Zaher Baidas, Senior Vice President and General Manager of Power at Renesas. "[Such] innovations will deliver performance and efficiency, with the scalability required for future growth."

Power Integrations also outlined how high voltage GaN HEMTs can deliver high power density and efficiency in 800 VDC power architectures. And in another example, Microchip revealed what it calls the industry's first 3 nm PCIe Gen 6 switch, designed to reduce power consumption in high density AI systems when connecting and routing high speed data between critical compute resources such as CPUs, GPUs, AI accelerators and storage devices. Systems level player Flex Power also launched an open platform integrating power, cooling and compute resources into modular designs for next-generation data centres while Eaton unveiled a reference architecture for 800 VDC power systems.

More than hardware

In May this year, NVIDIA also introduced 'NVLink Fusion', a system for partners to integrate their own silicon AI processors

with the firm's GPU ecosystem using NVLink, interconnect technologies designed to connect NVIDIA GPUs and CPUs. Custom silicon and technology partners include MediaTek, Marvell, Alchip, Astera Labs, Global Unichip Corporation, Samsung, Synopsys and Cadence while CPU partners are Fujitsu, Qualcomm Technologies and Intel.

Prior to NVLink Fusion, NVLink technologies could only be used with NVIDIA chips - and while Fusion drives demand for NVIDIA's underlying AI infrastructure amongst hardware developers, it also lowers the barrier to developing and scaling data centre infrastructure. In his recent technical blog, 'Scaling AI Inference Performance and Flexibility with NVIDIA NVLink and NVLink Fusion', Joe DeLaere, NVIDIA senior manager for accelerated computing solutions for data centre, wrote: "Being available as a modular Open Compute Project (OCP) MGX rack solution enables NVLink Fusion integration with any network interface card, data processing unit or scale-out switch, giving customers the flexibility to build what they need."

Importantly, NVIDIA's NVLink Fusion, twinned with recent strategic partnerships, also indicate the firm is positioning itself not just as a straightforward AI infrastructure provider, but moreover, as a collaborative platform for future AI-driven computing. As industry barrels towards the modern AI data centre, this openness will be critical to power semiconductor firms developing the devices that will support AI data centre scaling and the ensuing rises in energy consumption.


And of course, for OCP, openness remains fundamental. In October, this year, the organisation welcomed three data centre powerhouses to its Board of Directors, reflecting the importance of collaboration amongst hyperscalers, semiconductor firms and infrastructure providers.

"We've expanded our board with AMD, ARM and NVIDIA," explains Grossner. "And just as ARM is bringing its chiplet ecosystem to OCP, we're encouraging NVIDIA to do the same around NVLink." "The impact of [AI data centres] is becoming really obvious and real," he adds. "We'd be remiss to not take action and avoid serious issues."



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



Ferric brings power efficiency to the heart of AI chips

As power-hungry AI processors strain energy systems, Ferric's integrated voltage regulators promise more efficient power delivery.

BY REBECCA POOL, TECHNOLOGY EDITOR

AT A TIME when AI accelerators are pushing power consumption to kilowatt levels and beyond, power semiconductor players are joining forces to hone power delivery to these high-performance processors. Earlier this year, US-based custom silicon and SoC firm, Marvell Technology, partnered with half a dozen power tech manufacturers from around the world to integrate voltage regulators directly into its silicon platforms.

In these integrated voltage regulator (IVRs), all components of the switched inductor DC-DC power converter are integrated on-chip, bringing voltage regulation closer to the processor. This vertical power delivery set-up offers

clear benefits: it shortens electrical paths to reduce losses, improves power delivery response, and frees-up all-important board space – ultimately enabling more computer density per rack and easing the energy bottlenecks for hyperscalers.

US-based IVR manufacturer, Ferric, is amongst the power players working with Marvell Technology to prove the devices are well-and-truly ready for future data centres. Ferric chief executive, Noah Sturcken, has been working on these devices for nearly twenty years – and back then, he and colleagues at Columbia University, New York City, had already realised a processor power delivery bottleneck

was looming. As he points out, power demand was on the rise yet processor voltages were decreasing, and traditional power delivery systems with external voltage regulators couldn't efficiently deliver the required low-voltage to processors.

“With funds from Intel and US government, we were developing some of the underlying technologies critical to enabling IVRs at Columbia,” he says. “We were working on [thin film] ferromagnetic materials that could be integrated with semiconductors to miniaturise the entire power converter system and enable dense IVRs that can address this bottleneck – and that's where we are now.”



Early days

Ferric was founded by Sturcken and Professor Ken Shepard, also from Columbia, in 2012, to commercialise integrated DC-DC converters using thin film magnetic power inductors, rather than bulky, wire-wound inductors, for vertical power delivery. Early on, the firm partnered with Tier-1 fabless semiconductor companies to implement CMOS compatible thin film fabrication processes, and by 2017, had joined forces with TSMC to integrate its thin film inductors onto CMOS chips. TSMC licensed the technology, and as Mark Liu, TSMC’s President and Co-Chief Executive has said, integrating Ferric’s inductors with his company’s CMOS back-end-of-line enables efficient, high density on-chip/on-package power conversion at any process node.

Along the way, Ferric has won more than \$80 million in US government funds and venture capital - a mighty \$32 million came late last year - to support manufacturing and production ramp-up. Following the recent Marvell Technology partnership, the company launched its latest IVR in August this year, which specifically targets AI and high-performance processors, delivering 160 A from what is claimed to be the industry’s smallest footprint of 4.2 x 8 x 1 mm (35.5 mm²), with 4.5A/mm² current density. Ferric also says this IVR can provide three times more power per area and more than twenty times the power per volume, compared to competitors’ devices, and could scale to more than 10 kW with 64 IVRs.

“Some incumbent devices have a profile of a couple of millimetres, but we’re achieving what we believe to be the highest current

➤ Ferric’s latest IVR: the Fe1766 is a 16-phase e buck converter with integrated power switches and inductors designed to address power demands of AI processors, datacentre infrastructure and advanced electronics [Ferric].



➤ Placing voltage regulators close to the point of load can save power, board-space and cost. [Ferric].

with a component that’s barely a millimetre thick,” says Sturcken. “A high current density is very important for AI processors - right now the current density on some processors can exceed four amps per square millimetre in hot spots, and this is likely to increase to more than six amps per square millimetre in the future. It’s really important that the current density of converters keeps up with the processor.”

“Our low profile also means you can put our converters on the bottom-side of a motherboard in form-factors that might not provide enough room to accommodate vertical power delivery - that could be a mezzanine card or a PCI Express card,” he adds.

Critically for high performance processors, IVRs also provide granular power management, so the amount of power reaching a portion of a processor can be specifically controlled and rapidly adjusted according to workload. “This is a consequence of the power converters being very small and modular, being closer to the processor and having a faster regulation bandwidth,” highlights Sturcken.

“The power converters can operate ten to 100 times faster than conventional power converters and move the supply voltage around that

“ A high current density is very important for AI processors - right now the current density on some processors can exceed four amps per square millimetre in hot spots, and this is likely to increase to more than six amps per square millimetre in the future. It’s really important that the current density of converters keeps up with the processor. Our low profile also means you can put our converters on the bottom-side of a motherboard in form-factors that might not provide enough room to accommodate vertical power delivery - that could be a mezzanine card or a PCI Express card ”

much faster to provide best energy efficiency.”

Power-hungry processor demands aside, feedback from Ferric’s customers confirms that the timing of its latest IVR, and work with Marvell Technology, is apt. According to Sturcken, within the last 18 months, AI processor developers have been telling Ferric that the conventional power delivery path simply doesn’t work any more. “They could lose hundreds of Watts through inefficiencies while also facing reliability issues from delivering a lot of current through the package ball grid arrays [connecting the processor to the PCB],” he says. “Right now power delivery is the biggest threat to the ongoing development of these AI processors and future AI services.”

So what now for Ferric? Collaboration is clearly the way forward. The company has multiple projects underway with Marvell Technology, in which its power converters are being integrated into different locations within the silicon,

such as the motherboard or the processor package substrate. “[With Marvell], we can assess the different technical and business advantages of each different integration scheme, although we expect all to be used effectively in high volume production,” says Sturcken.

Ferric is also working with other organisations to develop power conversion technologies upstream of its own IVRs, so power delivery systems will be ready for, say, the 800 VDC power architecture of future data centres, as announced by NVIDIA earlier this year. And the firm also continues to work on ever-higher levels of integration, developing modular and flexible set-ups.

For example, a power converter might have terminations on both the top and bottom side to enable the most direct power path from the PCB to a processor. “We expect this and embedding the IVR in the package substrate to become common,”

comments Sturcken. Still, as the Ferric chief executive highlights, such developments will take time to be implemented in commercial applications, as any change raises questions. “Customers might ask, ‘if you take that substrate and embed a bunch of power converters into it, how will you then test it?’ Or, ‘what happens if one of those power converters aren’t assembled correctly, do I then have to throw out the entire system and deal with some pretty expensive economics?’” he says. “These questions create uncertainty that needs to be addressed before the technology is adopted at large-scale.”

However, Sturcken is certain that every high-performance processor will need integrated voltage regulation. “We’ve had this sudden recognition, of yes, this is what is needed for future processors,” he says. “Right now, many organisations still need to go out and find technology partners such as us, so that they can stay competitive – so [IVR adoption] is only a matter of time.”

Its blue sky all the way with Evatec



Find your way in thin film technology for Power Devices with Evatec

From 300mm silicon processing to Wide Band Gap (WBG) applications on SiC or GaN, Evatec is your partner for thin film production in power devices.

Process know-how including secure thin or bonded wafer handling combined with management of wafer stress and temperature keeps you on the right track.

Visit evatecnet.com/markets/power-devices to find out how you can power ahead in your own production or contact your local Evatec Office at evatecnet.com/about-us/sales-service

Evatec – The Thin Film Powerhouse
Advanced Packaging, Semiconductor, Optoelectronics, Photonics
www.evatecnet.com

Thin Film Production Tools and Processes

Silicon, SiC and GaN ↑

Increased yields straight ahead ↑



**SMART
MANUFACTURING** 

**SMART
MOBILITY** 

**SMART
MEDTECH** 

**SMART
DATA-AI** 

**WORKFORCE
DEVELOPMENT**

**Global
ADVOCACY**

Sustainability

Mit MARKET
INTELLIGENCE
TEAM

1000+
STANDARDS

**THOUGHT
LEADERSHIP**

SEMICON

**TECH
COMMUNITIES**

**SEMI
UNIVERSITY**

**Cyber
Security**

Connecting Semiconductors and Electronics

About SEMI:

SEMI connects more than 2,500 member companies and 1.3 million professionals worldwide to advance the technology and business of electronics design and manufacturing. The breadth and depth of our events, programs and services help members accelerate innovation and collaboration on the toughest challenges to growth.



Efficient debugging workflow in chip circuit verification - A targeted simulation approach

Chip circuit verification is a vital to the semiconductor industry, ensuring that integrated circuits meet design specifications and quality standards.

BY SLAVA ZHUCHENYA, PRODUCT ENGINEER SUPPORTING CALIBRE INTERFACE TOOLS IN THE DESIGN-TO-SILICON DIVISION OF SIEMENS DIGITAL INDUSTRIES SOFTWARE.

DURING the verification process, engineers aim to identify and resolve potential design flaws or violations before the chip goes into production. To meet tight tape-out deadlines, verification simulations need to be as fast as possible, as there can be many iterations of these runs.

However, that is only part of the challenge. IC designers need to quickly identify circuit issues, but fixing them requires debug data. Traditional approaches either sacrifice speed by including all debug data in the main verification run or compromise debug capabilities by omitting crucial data to expedite the process. Balancing the need for fast verification with the necessity of detailed debug information presents a dilemma.

Chip circuit verification is a vital to the semiconductor industry, ensuring that integrated circuits meet design specifications and quality standards. During the verification process, engineers aim to identify and resolve potential design flaws or violations before the chip goes into production. To meet tight tape-out deadlines, verification simulations need to be as fast as possible, as there can be many iterations of these runs.

However, that is only part of the challenge. IC designers need to quickly identify circuit issues, but fixing them requires debug data. Traditional approaches either sacrifice speed by

including all debug data in the main verification run or compromise debug capabilities by omitting crucial data to expedite the process. Balancing the need for fast verification with the necessity of detailed debug information presents a dilemma.

Challenges to debugging workflows in IC verification

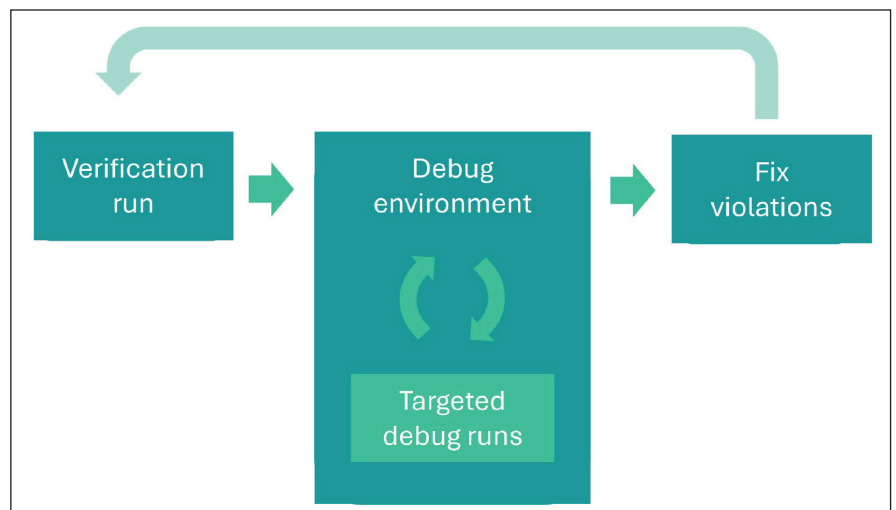
In chip circuit verification, there is often a trade-off between the speed of verification runs and the availability of helpful data to fix the violations. Running a full verification job with all the useful debug data can be time and resource-consuming and lead to an explosion of data volume. The more information there is to process and write to disk, the more it will cost.

Conversely, excluding key data to accelerate verification iterations may hinder the ability to effectively debug results later.

The time savings in one stage may become a penalty in a later stage. The challenge lies in finding a middle ground that allows for fast verification while enabling efficient debugging without compromising either, as both of these are a crucial part of chip design and verification.

A solution for balancing verification speed with detailed debug

A better solution involves integrating targeted debug-focused simulations into the debugging stage itself instead



➤ Figure 1. A targeted debug-focused simulation integrated into the debugging stage.

of an all-or-nothing approach (figure 1). Not every result needs detailed debug data. At least not right away. Some results might be waived later or could be fixed as a result of fixing some larger underlying problems that can eliminate many violations. Instead of choosing between running a full verification job fast or running it with all the necessary debug data, why not give users the ability to decide later when their needs become clearer?

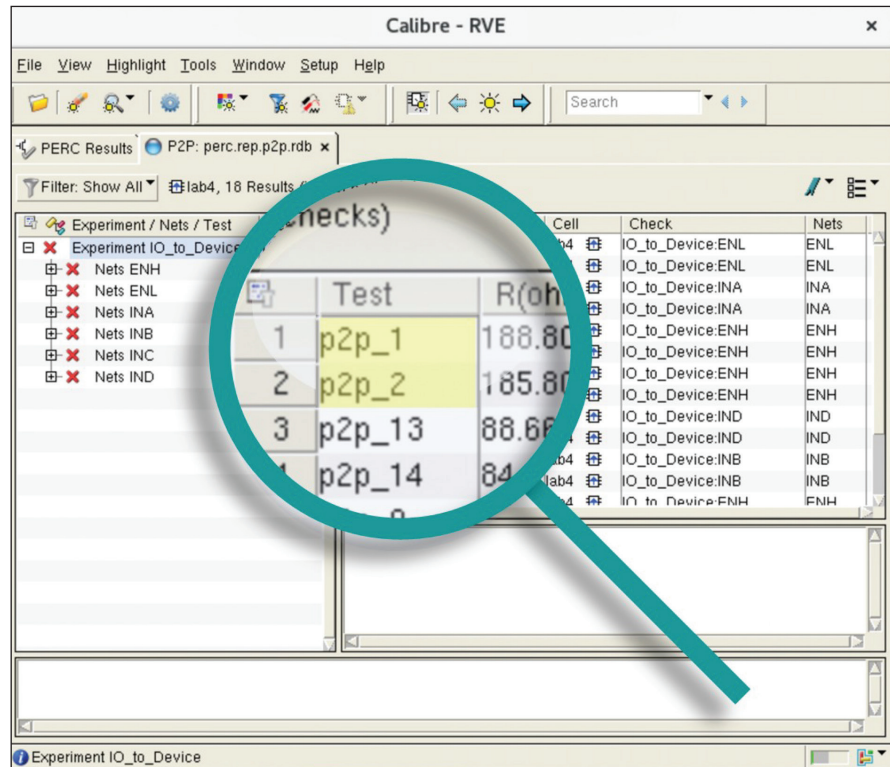
A designer first wants to see what rule violations they have in their simplest form and then decide how they need to act for each result. This approach allows designers to obtain debug data for specific results right from within the debug environment (figure 2).

Once it is clear that more information is needed, the debug tool should let the user run smaller targeted data-rich simulations just on the interesting results. This way, the full verification remains fast, and getting debug data is achievable without the overhead of a massive verification job.

Calibre PERC point-to-point example

Calibre PERC's point-to-point (P2P) simulation is a circuit reliability verification solution that checks for the presence and impact of parasitic resistance in an IC layout. Parasitic resistance in chip design refers to the unintentional resistance that arises due to the physical layout and interconnections of the components on the chip.

This resistance can have a significant impact on circuit performance and reliability when it is manufactured. Calibre PERC P2P simulations calculate the total effective parasitic resistance between two or more points in the IC layout, to check if the resistance does not exceed allowable thresholds that ensure a chip will work as designed.



Test	R(ohm)	Cell	Check	Nets
1	p2p_1	188.80	IO_to_Device:ENL	ENL
2	p2p_2	185.80	IO_to_Device:ENL	ENL
3	p2p_13	88.66	IO_to_Device:INA	INA
4	p2p_14	84.00	IO_to_Device:ENH	ENH

➤ Figure 2. Calibre RVE debug environment offers the data needed for debug.

These simulations can be complex and time-consuming to run but are a crucial for circuit reliability. As with most IC verification flows, simulation runtime is key to successful and timely chip tape-outs. In particular, the more iterations that are required, the more the full verification runtime starts to impact the tape-out schedule. For this reason, P2P simulations report minimal information on the violations, like the effective resistance value and the nets involved. In some cases, a designer might know the layout well enough, and a violation could be simple enough that the fix could be obvious. However, more often than not, additional debug data is needed to fully understand the source of the violation.

Debugging point-to-point violations requires analyzing chip layout interconnect and determining what

routing practices might be causing a higher-than-expected or allowed total resistance value. Doing this with only the net and the effective resistance between some points on the layout can be a significant challenge.

Designers may need more information to determine what part of the net layout is contributing the most to the high parasitic resistance. Making guesses and then rerunning the verification simulation can quickly become very expensive.

Detailed data about the interconnect polygons reveals what is happening between the points and can save time by limiting the number of verification reruns that need to happen. Resistance contribution and current density data for individual layout polygons give designers the data they need to better

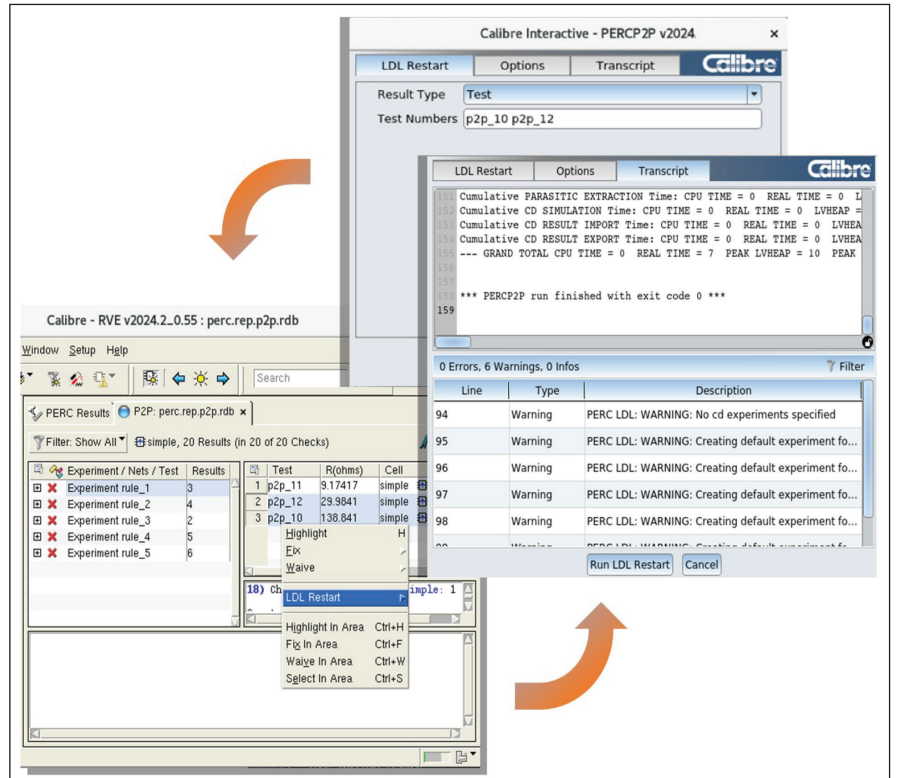
Calibre PERC's point-to-point (P2P) simulation is a circuit reliability verification solution that checks for the presence and impact of parasitic resistance in an IC layout. Parasitic resistance in chip design refers to the unintentional resistance that arises due to the physical layout and interconnections of the components on the chip

debug these results. However, including all that polygon data for every violation would then significantly increase runtime and database sizes for those runs. How do you find a good middle-ground approach?

Calibre RVE PERC now allows designers to run Calibre PERC P2P Debug simulations on select P2P results directly from the debug GUI. This means that designers can get debug data for specific P2P results on demand without leaving the debug environment. They can run the full PERC P2P flow on their design without needing to make any predictions on what debug data they'll need. Once it is clear that more information is needed, RVE lets you run smaller targeted P2P simulations that gather the polygon data just for interesting results (figure 3).

The added database contains information like current density, resistance calculations, and percentage of resistance contribution of the individual polygon. This data significantly increases designers' productivity during the debug process without compromising the runtime of the full Calibre PERC P2P simulation.

The improved workflow begins with the standard execution of the main P2P verification run, optimized for speed by excluding non-essential debug data. Upon identifying potential violations in the results, engineers transition to the debug environment and stay there even if they need to make additional runs to extract more debug data for select violations. This is a much more efficient, productive and intuitive flow.



➤ Figure 3. Run Calibre PERC P2P Debug simulations on select P2P results directly from the debug GUI.

Conclusion

Efficient debugging in chip circuit verification requires a balanced approach that optimizes verification speed without compromising debug capabilities. By integrating targeted simulations into the debugging environment itself, engineers can obtain additional debug data selectively and iteratively, minimizing the need for full verification reruns and setup efforts. This approach maximizes productivity and accelerates time-to-market while

ensuring the reliability and functionality of integrated circuits.

Calibre PERC P2P is one such circuit verification simulation that benefits from this improved debug flow. Calibre RVE PERC streamlines the debug process by enabling targeted P2P simulations within the debug environment, providing detailed debug data without compromising full verification runtime, significantly boosting productivity for designers.

Dedicated webinars for the power electronics (PEI) industry

Using our 30+ years' experience in B2B vertical technical markets, and as the publisher of PEI Magazine, we offer effective webinars, ZOOM interview and virtual events. We help you get your message, to your desired audience, by marketing to over 53,000 PEI professionals.


In addition to organising and managing your webinar, we can also market your webinar to our specialist databases.

Reach Educate Influence

- Brand Awareness
- Lead Generation
- Thought Leadership



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



AP INTERNATIONAL CONFERENCE

ANGELTECH

20 - 22 April 2026

Sheraton Brussels Airport Hotel,
Brussels, Belgium

CONNECTING THE ADVANCED PACKAGING INDUSTRY

The 1st Advanced Packaging International builds on the momentum driving the next era of semiconductor innovation. Over two days, industry leaders and academic experts will deliver around 40 presentations across four key themes:

- Heterogeneous Integration
- Getting the Best Out of Incumbent Solutions
- Advancing Thermal Management
- Optimising Packaging Architectures

Conference attendees will gain insights into the materials, methods, and technologies reshaping chip-level packaging and system integration.

From cutting-edge thermal solutions to the reimagination of wirebond and leadframe formats, the event provides a comprehensive look at how packaging is enabling system-level innovation.

The two-day conference also includes expert panels and a dedicated poster session, creating multiple avenues for peer-to-peer engagement and technical exchange across the packaging ecosystem - from researchers and R&D teams to product engineers and system architects.

Advanced Packaging International is part of AngelTech, which delivers a portfolio of insightful, informative, and highly valued chip-level conferences.

Bringing together an Innovation Summit and four co-located conferences with more than 120 presentations, 700+ delegates, and over 80 exhibitors, AngelTech is the premier global event covering compound semiconductor, photonic integrated circuit, power electronic, and advanced packaging technologies. With significant overlap between the conferences, attendees benefit from exposure to full supply chains, customer bases, and partner ecosystems.

SPECIAL Event Discount:

Register before 31 January 2026 and get 10% OFF!

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

10% OFF




To find out more about our sponsor and speaker opportunities contact us today on:
+44 (0)2476 718970 or email: info@angel-tech.net <https://ap-international.net>

CO-LOCATED WITH:



ORGANISED BY





The criticality of performance per watt optimisation for AI chip development

AI workloads create an urgent need for compute and power efficiency improvements from the architecture stage forward.

BY GODWIN MABEN, LOW-POWER ARCHITECT AND SCIENTIST AT SYNOPSIS

CHIP DEVELOPERS are seeing an urgent rise in demand for compute processing capability driven by AI workloads. This increase in compute requirements drives a corresponding increase in the demand for power consumption.

For example, a ChatGPT query requires nearly 10 times as much power, on average, as a Google search [1].

Power has traditionally been treated as a secondary constraint, with performance taking precedence during development. It is no longer feasible to leave power optimization until the end of the design cycle. The performance per watt metric is now of critical importance for AI chips and chiplets

and must be addressed throughout the development process. Hyperscalers now often revise their metrics to be “tokens/watt.” [2]

Architecting for power efficiency at the earliest possible stage of the chip design process will maximize power-saving potential. Using shift-left methodologies, 30% to 50% power savings can be achieved at the software and hardware architecture phase, compared with single-digit power savings during implementation and signoff stage [3].

An end-to-end silicon to systems solution enables designers to optimize power early in the design cycle while achieving performance goals.

Top Challenges in AI Chip Efficiency

AI chip developers face four main challenges:

- Power efficiency and thermal management
- Memory bandwidth and data movement
- Architecture analysis
- Optimizing hardware and software for the representative workload

Power efficiency and thermal management

AI and other demanding applications are driving the use of multi-die systems and semiconductor devices with multiple homogeneous or heterogeneous dies within a single package. This enables the rapid development of tailored silicon solutions for high-performance computing applications [4].

Heat dissipation is one of the main challenges in designing a multi-die AI system, creating thermal limitations. A well-planned architecture following an iterative process can alleviate thermal stress by exploring options at the front end to avoid getting locked into a partitioning structure that could eventually turn out to be sub-optimal from a power perspective.

System architecture teams can use modeling tools to abstract out pieces of a chip into models for performance/

power analysis and finalize power tradeoffs before the design is locked into its partitions. By mapping a workload onto a multi-die system, the design team can determine the activity per processing element and per communication path. Modeling the hardware and software together is key to generating a robust and thermally efficient design, with scalable software across the die.

Memory bandwidth and data movement

AI applications thrive on high memory bandwidth, fast throughput, and low latency. The growth of bandwidth has not kept up with the growth of compute. For chip designers, overcoming the initial challenge of “The Memory Wall” - the gap between processor speed and memory bandwidth - in AI/ML chip design is paramount. For AI chips, one of the leading causes of power consumption is data movement - even more than compute - and high-speed die-to-die communication needed to pass large data sets between dies within a chip.

Developers must analyze data movement early and identify solutions to optimize memory power and minimize data movement to achieve the highest performance per watt.³ Solutions include high bandwidth memory, analog computing, custom compute units, compute in memory, resistive RAM structure, and algorithmic

solutions like sparse algorithms, to eliminate unnecessary data movement, to eliminate unnecessary data movement.

Design teams should focus on identifying the right memory architecture to minimize data movement, analyzing these architectural changes at an early stage and implementing them.

Architecture analysis

Before the start of the power design cycle, it’s critical to architect a power analysis flow that aims to analyze power as early as System Architecture Stage. Synopsys Platform Architect™, a performance and power analysis tool, enables accurate simulation of system-level function in SystemC, providing crucial early insights into power-performance tradeoffs before hardware descriptions in Verilog have been written.

Platform Architect helps system designers explore and optimize the hardware-software partitioning and the configuration of the System-on-Chip (SoC) infrastructure, focusing on global interconnect and memory subsystem to achieve the right system performance, power and cost. This process helps in deciding the efficient macro architecture of the system, including, but not limited to, technologies such as Dynamic Voltage and Frequency Scaling (DVFS), Power Gating, Network-

on-Chip (NOC) traffic, etc. Using transaction-level simulation, Platform Architect reduces design time by predicting and optimizing architecture KPIs.

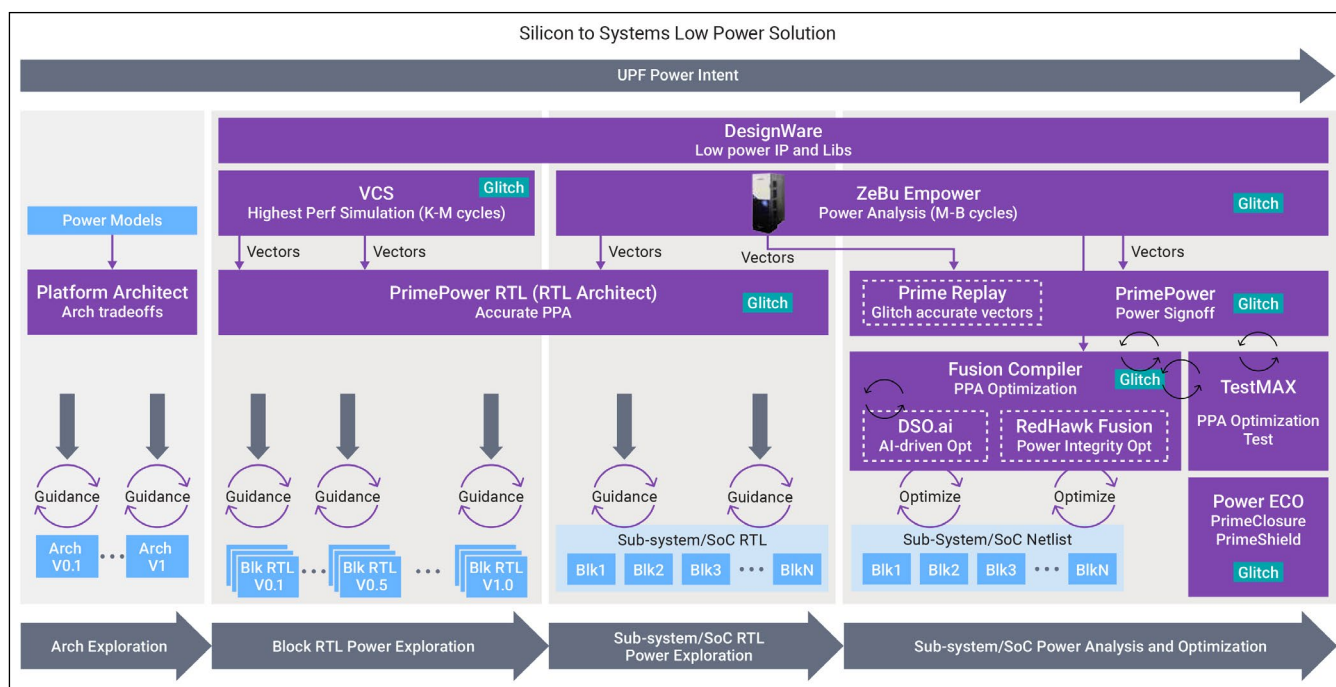
Optimizing hardware and software for the workload

Optimizing both the hardware and software for the specific workload is critical; therefore, developers must model, simulate, emulate and prototype chip performance prior to hardware returning from the fab.

Leveraging early architecture analysis and performance validation in emulation will result in efficient hardware/software partitioning and customized hardware/software for a very specific workload, such as Instagram workload-specific Application-Specific Integrated Circuits (ASICs).

Tools such as Synopsys ZeBu® Empower can be used for workload profiling across multiple vectors to identify the right window and workload for power analysis and optimization. An efficient combination of major benchmark workloads, such as Idle/Sustained/Inference/Training, will be the desired collateral for this analysis and optimization.

Synopsys Solutions for High-Performance, Power-Efficient AI Chips The challenges faced by chip developers illustrate the need for power



optimization throughout the design flow, with a strong emphasis on shift-left methodologies. The potential for power savings diminishes significantly by the implementation phase if not addressed throughout the process. Synopsys' end-to-end power solutions allow for power analysis and optimization at every stage of development—architecture, emulation, functional prototyping, design and verification, implementation and test, engineering change order (ECO), and signoff.

Workload profiling

Key workload profiling and identifying the right workload and window are essential. Synopsys PrimePower™ RTL, along with functional workloads from verification and simulation tools (VCS) or profiled workloads from ZeBu Empower, enables logic designers to design a power-efficient RTL. The Synopsys VCS is the primary verification solution used by the world's top semiconductor companies. VCS provides the industry's highest performance simulation and constraint solver engines, allowing users to easily speed up high-activity, long-cycle tests by allocating more cores at runtime.

ZeBu Empower analyzes billions of cycle workloads from ZeBu emulation and identifies optimal smaller windows for power analysis and optimization.

ZeBu Empower generates workloads for analysis and optimization, revealing opportunities for dynamic and leakage power early in development, reducing the risks of power bugs and missed SoC power goals. ZeBu Empower enables multiple iterations per day with actionable power profiling in the context of the full design and its software workload, optimizing average power, peak, IR drop, wasted clock pin power and others.

Power analysis

Once RTL has been implemented, the PrimePower product family enables accurate power analysis for block-level and full-chip designs, beginning with RTL through the different stages of implementation, leading to power signoff.

Running power analysis on a flat design is challenging due to memory and runtime requirements. Power analysis and optimization must be performed hierarchically and must have a robust methodology in place.

Further, multi-die design and chiplets require many more connections than traditional monolithic chips, and this increased interconnect density creates power distribution challenges, requiring some advanced routing capabilities. Performing power integrity across

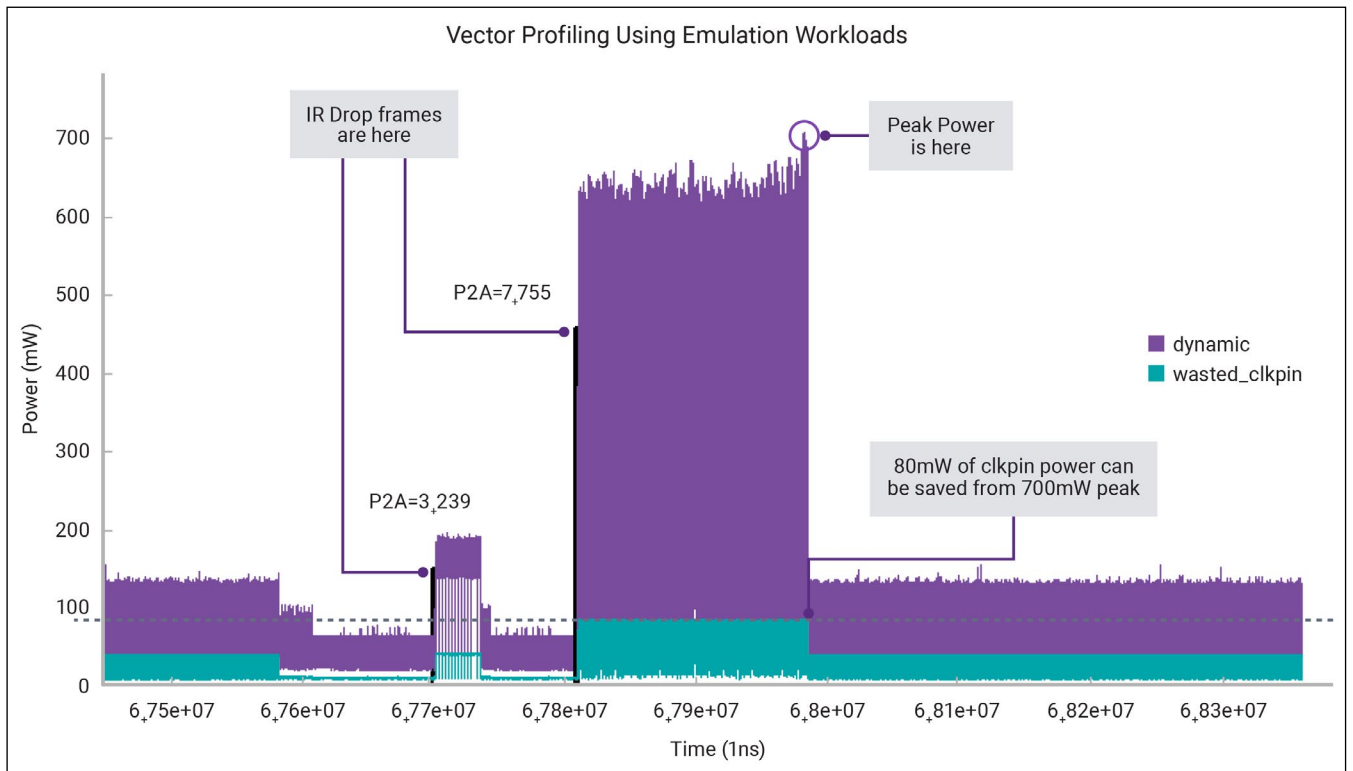
heterogeneous components is more challenging in chiplets due to the complex geometry and relationship between power and temperature.

During implementation and signoff, PrimePower provides accurate gate-level power analysis reports for SoC designers, enabling timely optimization and power target achievement through various metrics such as average, peak, glitch, clock network, dynamic and leakage power and multi-voltage power. The workload for this activity can be derived either from simulation or emulation.

As designs become more complex, designers need a tool that pinpoints the major power sinks while suggesting modifications with the highest return on investment. Synopsys' PrimePower product family provides accurate power analysis for block-level and full-chip designs, improving power efficiency through each stage, shortening the design cycle.

Developers can leverage PrimePower RTL to assist in:

- Power estimation
- Power profiling and distribution
- Identifying specific RTL lines to designers for modification of the architecture and re-analysis



Fixing the glitch

Synopsys technologies address the problem of glitches, where unnecessary signal transitions in a combinational circuit drain power, leading to a significant contributor of dynamic power.

At lower geometry nodes, due to an imbalance in delay between gate and net, most data path-sensitive designs will exhibit a higher percentage of glitches. The majority of glitch sources can be optimized at the architecture level, and final glitches caused due to imbalanced paths can be optimized during the ECO cycle.

Solutions to suppress glitches are needed at every step of the cycle.

- PrimePower RTL computes and identifies sources of glitch early in design cycles, which enables designers to rearchitect to reduce glitch. It can also point to the RTL source line of code generating the highest level of glitch. Optimizing glitch power for a tile can lead to high power savings at the SoC level
- The PrimePower solution offers delay-/glitch-aware vector generation using RTL simulation. The product can generate a Switching Activity Interchange Format (SAIF) with glitch annotation (Inertial Glitch and Transport Glitch) and a delay-aware SAIF from an RTL simulation on any given netlist
- The SAIF or Fast Signal Database (FSDB) can be used during implementation with Synopsys Fusion Compiler™ or during ECO with PrimeClosure to perform glitch-aware optimizations and reduce glitches following the timing phase
- Finally, PrimePower gate-level power analysis and golden power signoff perform glitch power analysis using timing-aware simulation correlating closely to SPICE power numbers

Optimisation

During Place and Route Developers can optimize their design for power during synthesis and the place-and-route stage of the design cycle using technologies such as clock gating, power aware placement, clock tree synthesis, logic restructuring and many more power optimization technologies in Fusion Compiler.

Voltage plays a dominant role in dynamic power consumption via CV2. Reductions in voltage, even if minor, can lead to a quadratic reduction in power. Fusion Compiler's VoltOpt feature enables voltage-based analysis and optimization, and it can be used in the implementation phase to reveal the minimum voltage required to maintain the same timing/area. Once Vmin is identified, designers can redesign their logic to operate at this lower voltage, significantly reducing power. Tools such as Synopsys PrimeShield™ can be used for voltage slack/robustness analysis to ensure that Vmin qualified by VoltOpt is optimal.

Timing/Power ECO using PrimeClosure will be the last piece of the puzzle in squeezing last nW of power before tapeout.

A signoff-level power analysis tool, such as PrimePower, is key to determining any available power recovery at the end of the cycle. This provides developers with a view of the power they will see in the final silicon.

Power optimisation through Silicon Lifecycle Management (SLM)

The performance of a silicon chip does not remain constant over its operating life. Aging effects in the silicon structures and factors in the system operating environment change the performance characteristics of

the device over time. Silicon Lifecycle Management (SLM) provides a data collection, analysis and control environment to monitor these effects and implement corrective actions in the field at a unit device level. The result is a far more stable and secure performance of the silicon device and overall system over time. Synopsys' integrated SLM family of products is built on a foundation of enriched in-chip observability, analytics and integrated automation and improves silicon health and operational metrics at every phase of the device lifecycle.

Through monitoring, deep insights are obtained from silicon to system, enabling meaningful data collection at every opportunity for continuous analysis and actionable feedback.

Summary

The escalating demands of AI workloads require a fundamental shift in AI chip development - one that prioritizes power efficiency from the inception of the architectural design. Power modeling is a key power analysis and optimization challenge for multi-die design. The limitations imposed by thermal constraints in multi-die systems, the energy drain of extensive data movement, and the imperative for workload-specific hardware and software optimization underscore the critical need for early and continuous power analysis and mitigation.

Synopsys' comprehensive suite of tools and the tangible results achieved by innovators like SiMa.ai demonstrate that a proactive, shift-left methodology, coupled with robust power analysis and optimization throughout the design life cycle, is not merely advantageous but essential for achieving the high-performance, energy-conscious AI chips necessary in this new era of artificial intelligence.

FURTHER READING

- [1] AI is poised to drive 160% increase in data center power demand. Goldman Sachs. May 14, 2024. <https://www.goldmansachs.com/insights/articles/AI-poised-to-drive-160-increase-in-power-demand>
- [2] Microsoft/Synopsys at DVCON Keynote 2025. <https://dvcon.org/program/2025-keynote-presentations>
- [3] Ruby, W. How Early Power Analysis Drives Energy-Efficient RISC-V Designs. Synopsys. July 18, 2024. <https://www.synopsys.com/blogs/chip-design/risc-v-power-efficiency.html>
- [4] Synopsys editorial staff. Designing Thermal Management Solutions for Multi-Die Systems. Synopsys. May 22, 2023. <https://www.synopsys.com/blogs/chip-design/thermal-management-solutions-multi-die-systems.html>

Taking SiC switches to new highs



How APC Electronics and Luminus aim to reshape high-power, high-frequency design with next-generation silicon carbide technology.

BY TOM JORY (LEFT), VICE PRESIDENT OF POWER SEMICONDUCTORS AT LUMINUS, AND DUMITRU SDRULLA (RIGHT), CTO AT APC-E

SILICON CARBIDE (SiC) has long been recognised as one of the most transformative technologies in modern power electronics, promising higher efficiency, better thermal behaviour and superior ruggedness compared with silicon incumbents. Yet within the rapidly growing SiC ecosystem, some companies are pushing the boundaries far beyond today's mainstream implementations - towards higher power, higher frequency and more integrated, application-specific solutions.

Two such organisations are APC Electronics (APC-E) and Luminus, who recently announced an exclusive worldwide sales partnership. Their combined capabilities - APC-E's deep power semiconductor design expertise and Luminus's global commercial footprint and wide-ranging power-technology background - form the basis of a new approach to SiC switches and integrated driver solutions.

In this conversation, Tom Jory, Vice President of Power Semiconductors at Luminus, and Dumitru Sdrulla, CTO at APC-E, explain how APC-E's high-performance SiC MOSFET and Schottky diode portfolio emerged, how these devices outperform established competitors, and how the companies intend to bring 'unprecedented' switching frequencies and power densities to applications ranging from data centres to megawatt-scale charging infrastructure.

A design powerhouse 'born in Bend'

APC Electronics may not yet be a

household name, but its origins reveal a team steeped in semiconductor heritage. The company's founders, representing more than a century of collective experience, established APC-E in Bend, Oregon, when Microchip Technology closed its fab in the region. Rather than relocate, the engineers stayed in Bend and formed a specialist design house focused on high-speed SiC MOSFETs, Schottky barrier diodes (SBDs) and high-performance gate-driver ICs.

"We really liked Bend and wanted to stay here," recalls Sdrulla. "So, we started a fabless chip design service, creating high-speed silicon carbide MOSFETs and gate drivers for customers worldwide, with an emphasis on fast switching, high efficiency and high reliability."

APC-E works with foundries around the world, including local facilities such as the former Sycamore fab, while leveraging international partners for volume and specialty manufacturing. This global footprint gives the team flexibility in device structures construction/design, process flows and cost optimisation - advantages that become clear when examining its product lineup.

A portfolio built for performance

APC-E began with a family of SiC Schottky barrier diodes, intentionally optimised for low conduction losses and minimal recovery charge - critical parameters for efficiency and thermal performance in fast-switching converters. The company now offers SBDs rated at 650, 1200, 1700 and

even 2000 V, positioning them for industrial power conversion, renewable-energy systems and high-voltage infrastructure.

The MOSFET range is equally diverse. Initial releases include 650 and 1200 V SiC MOSFETs for both industrial and automotive applications, with upcoming ceramic-isolated variants in TO-247 three-lead and four-lead packages. Ceramic isolation enables more efficient thermal integration and simplified assembly, key advantages for modern power-module architects.

But the raw specifications only tell part of the story. APC-E's devices are engineered to maximise several figures of merit that correlate directly with system-level efficiency.

"From an electrical-performance point of view, the diodes have the lowest figure of merit we've benchmarked," says Sdrulla. "If you take forward voltage multiplied by recovery charge, and then divide that by surge current, we're still significantly ahead of competing devices."

For MOSFETs, the company emphasises $RDS(on) \times$ total switching charge (Q_{sw}) as a primary metric, with particular attention to the MOSFET's hot $RDS(on)$ - a performance dimension frequently overlooked by the market.

"By design, our MOSFETs maintain lower $RDS(on)$ at elevated temperatures compared with the competition," Sdrulla explains. "And when you look at time-to-failure metrics from long-duration TDDDB evaluations, we're able

to define an integrated figure of merit that incorporates reliability as well as performance. In that category, too, our devices stand at the top.”

Of equal importance is the company’s focus on cost-effective, streamlined manufacturing flows. High performance is only half the equation; wide adoption requires predictable, scalable and competitive pricing, something APC-E believes it has achieved through careful process optimisation.

SiC, GaN and silicon: understanding the technology landscape

With silicon carbide now firmly established in EVs, renewable energy and industrial drives, there remains ongoing debate around the roles of SiC versus GaN versus advanced silicon superjunction devices. APC-E approaches the question pragmatically. “GaN has very significant advantages in lower-voltage, lower-power applications,” Sdrulla acknowledges.

“The electron mobility in GaN is excellent for fast switching in the sub-650-volt space. But silicon carbide, because of its vertical structure and high breakdown capability, is the technology of choice for tens of kilowatts through tens of megawatts. It is fundamentally suited for high voltage, high power and high reliability.” APC-E’s technology roadmap aligns with guidelines from organisations like PowerAmerica, emphasising efficiency, reliability, cost and ecosystem readiness.

Target sectors include:

- Automotive EV and hybrid platforms
- Heavy-duty electric transport
- Photovoltaic inverters and energy-storage systems
- Data centres and AI-driven computing infrastructure
- High-voltage switchgear and circuit-breaker systems

Each sector brings unique requirements. In photovoltaics, for example, the company’s 1,200-V MOSFETs and SBDs are well suited to multi-kilowatt string inverters and large solar farm installations. In data centres - an industry undergoing rapid expansion due to AI training and inference workloads - APC-E sees particularly strong demand for 750-V SiC MOSFETs, positioned to displace

silicon superjunction MOSFETs that dominate PFC and DC-DC conversion today.

The migration from silicon to SiC in the data centre market is already underway, as server power density increases and operators pursue aggressive efficiency and sustainability targets. Here, APC-E’s high-frequency capabilities create a notable differentiator.

Pushing SiC to megahertz frequencies


Perhaps the most groundbreaking aspect of APC-E’s work is its development of SiC RF MOSFETs capable of operating not at tens of kilohertz, as is standard in most SiC applications, but at tens of megahertz. “Across the industry, silicon carbide MOSFETs typically operate in the kilohertz range,” Jory says. “We’re working on devices that run in the megahertz range. This takes things to an entirely new level.”

Such switching speeds unlock dramatic reductions in the size of magnetics and capacitors, two of the most costly and bulky elements in any power-conversion system. Faster switching means smaller passives, lower stored energy and significantly higher power density.

APC-E has validated discrete RF MOSFET designs at 50 MHz, 80 MHz and even 100 MHz. The company has also demonstrated a power module combining an RF SiC MOSFET with a chain of custom ICs, producing more than 4 kW at 54 MHz.

This is not merely a technology showcase, APC-E views high-frequency SiC as a critical enabler for future wireless charging, including inductive and capacitive transfer systems for electric vehicles and even heavy transportation.

“Very high power at very high frequency allows you to shrink system size dramatically,” says Sdrulla. “This opens the door to new architectures for wireless charging of cars, trucks or industrial vehicles. We are already in discussions with interested parties.” The implications extend well beyond transportation. High-frequency SiC



could reshape resonant power conversion, medical systems, industrial heating, plasma generation and any application where efficiency and compactness are paramount.

Reliability, ruggedness and the value of ‘Made in the USA’

Despite its ambitions in cutting-edge performance, APC-E emphasises reliability above all else. The company designs its MOSFETs and diodes to withstand harsh operating environments, with a focus on short-circuit robustness, unclamped inductive switching (UIS) survivability, high surge-current capability, stable threshold voltages over lifetime and enhanced avalanche energy ratings.

“Ruggedness and reliability by design are non-negotiable,” Sdrulla stresses. “We look very carefully at harsh-environment behaviour - high current, high voltage, short-circuit events, inductive transients. Our MOSFETs and diodes must do excel in all these areas.” APC-E is also committed to growing U.S.-based manufacturing, following broader industry and government trends encouraging domestic semiconductor supply chains.

“We are extremely interested, and focused, on having mass production capability in U.S. foundries,” Sdrulla adds. “The design is ready for volume manufacturing, and we want to support a Made-in-USA ecosystem.”

The Luminus partnership: technical depth meets global reach

For APC-E to scale its technology and engage customers worldwide, a commercial partner with technical credibility and established distribution was essential. The choice of Luminus was not arbitrary; it grew from an existing engineering relationship.

Luminus, founded in 2002 as an MIT spin-out, built its reputation in high-power LED technology, engineering chips capable of handling ten times the current of standard devices. This background in high-current packaging, thermal management and rugged

design proved directly applicable to SiC.

Luminus was acquired in 2013 by Sanan Optoelectronics, which has since become one of the world's largest LED chip manufacturers, and a major investor in silicon carbide. Sanan established a vertically integrated SiC megafab, now one of the top three global producers.

"APC-E was hired by Sanan to design silicon carbide MOSFETs, and that's how we got to know them," Jory explains. "APC-E wanted to launch its own line of power semiconductor products but needed a partner to focus on sales and marketing. It was a natural fit."

Under the partnership, APC-E remains focused on engineering, while Luminus becomes the exclusive global sales and marketing organisation for APC-E's SiC portfolio.

But the relationship is more than a sales channel. It creates a feedback

loop between customers, Luminus's application engineers, and APC-E's design team.

"We need the voice of the customer," Sdrulla emphasises. "Luminus brings that to us. Their seasoned marketing and sales team connects us to end users, so we can understand system-level needs and implement those insights directly into our designs."

A shared roadmap to breakthrough products

Both companies view the partnership as a platform for long-term innovation. APC-E will continue to advance its diode and MOSFET technologies, while the teams collaborate on integrated solutions that combine SiC devices, high-speed drivers and advanced packaging.

"We've identified directions of development that bring the knowledge and innovation of both groups into unique, differentiating products," says Sdrulla.

"There is already a patent application in place. The combination of Luminus's expertise and our SiC technology will make a meaningful difference in the field."

Jory agrees: "Both Luminus and APC are driven by innovation and by delivering components that enable customers to achieve higher performance - higher power levels, faster switching, higher efficiencies, and lower system costs. It's really a win-win for us and for the customers."

Industries from EVs to green energy to AI-driven data centres are hungry for more efficient, more compact, more reliable power solutions. With demand rising and system architectures evolving, SiC technology is entering its next chapter.

And if APC Electronics and Luminus succeed in pushing silicon carbide to megahertz frequencies and multi-kilowatt power levels, that next chapter may look dramatically different from today's playbook.

APC Electronics
Engineered in Bend, Oregon

Industry-Leading SiC Efficiency

- Low $R_{ds(on)}$ MOSFETs
- Best-in-class V_f in SBDs
- Lower capacitance

Superior Reliability & Ruggedness

- Higher breakdown voltage
- Higher surge capability
- 650V, 1200V, 1700V and 2000V options

Optimized Thermal Performance

- Lower junction-to-case resistance ($R_{\theta jc}$)
- Ceramic Isolated TO-247-3L MOSFETs

Brought to you by:
LUMINUS
www.luminus.com



**SAVE THE
DATE
20-22
APRIL
2026**

**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



From chargers to data centres:

Power GaN market set for rapid sixfold expansion by 2030

Power GaN 2025, powered by Yole Group, provides a detailed look at the market growth, technology trends, and strategic shifts shaping the gallium nitride power device industry.

BY ROY DAGHER TECHNOLOGY & MARKET ANALYST, COMPOUND SEMICONDUCTORS AT YOLE GROUP

Yole Group has released its new edition of the Power GaN report, Power GaN 2025. This annual bestseller provides an in-depth analysis of the global power GaN market trajectory, covering market values and volumes, supply chain dynamics, and technology innovations from 2020 to 2030.

The report's objective is to deliver a comprehensive understanding of GaN's growing role in the power electronics industry, from consumer applications to industrial, automotive, and data center systems, and to identify emerging players, competitive strategies, and manufacturing capacities that will shape the market over the next decade.

Power GaN is transitioning from promise to production reality. At Yole Group, we see an acceleration

across all end markets. Its efficiency, compactness, and performance advantages make it a key technology for the next decade of power electronics.

GaN has emerged as one of the most disruptive semiconductor technologies of the decade, with a market projected to reach about \$3 billion by 2030.

Consumer applications, particularly fast chargers, have been the early adopters, driving volume growth and ecosystem maturity. By 2030, the consumer and mobile segment is projected to constitute over 50% of the total power GaN device market... Read the related article: The power GaN race: market growth, consolidation, and new entrants.

Data centers: the golden road for GaN

The explosion in AI computing and data traffic is transforming data center power architectures. The growing number of servers and communication systems is driving electricity demand and CO₂ emissions, intensifying the need for higher-efficiency power conversion.

GaN technology offers a compact, efficient, and thermally superior alternative to silicon-based systems. As detailed in the Power GaN 2025 report, GaN is particularly well-suited for PSUs above 3 kW, delivering improved form factors, reduced heat losses, and lower operational costs.

In 2025, NVIDIA's new data center architecture announcement catalyzed a wave of collaborations with leading

power semiconductor manufacturers, including Texas Instruments, Navitas, Infineon Technologies, Innoscience, and onsemi. The aim is to integrate GaN devices into 800V HVDC power systems. These partnerships mark the beginning of large-scale GaN deployment, with Yole Group anticipating first commercial rollouts around 2027. See also the related Monthly Billet, “Nvidia defines power electronics future,” and more information about the associated report, “Power Electronics for Data Centers 2025.”

Together with telecom, these segments are expected to generate more than \$380 million in GaN revenues by 2030, making data centers one of the most promising growth pillars of the power GaN market.

“Data centers represent a turning point for GaN,” adds Roy Dagher from Yole Group. “The combination of AI, electrification, and sustainability goals makes GaN indispensable for next-generation server and telecom power systems.”

Other segments are gaining traction, with Enphase Energy’s first GaN-based microinverter and Changan Automobile’s first GaN-based OBC marking significant milestones that strengthen confidence in GaN

technology and pave the way for wider penetration across the power electronics market.

Following the power GaN revolution...

As the semiconductor industry strives for higher efficiency and sustainability, GaN stands at the core of the global energy and digital transition. From fast chargers to data centers and electric vehicles, power GaN is reshaping the way energy is converted and managed. At the same time, the market is shifting toward IDM-driven business models – where vertical integration offers tighter control over technology and supply – as foundries remain essential to the industry’s growth. Established foundries are expanding their GaN capacity, and new entrants are emerging to serve fabless companies and IDMs seeking additional sourcing flexibility.

This coexistence between integrated and foundry-based ecosystems is reinforcing the resilience and scalability of the global Power GaN supply chain.

With the release of Power GaN Transistor Comparison 2025, Yole Group delivers an in-depth comparison of the leading GaN power devices, revealing how design choices and manufacturing strategies define competitive advantage in this rapidly evolving market.

Key takeaways:

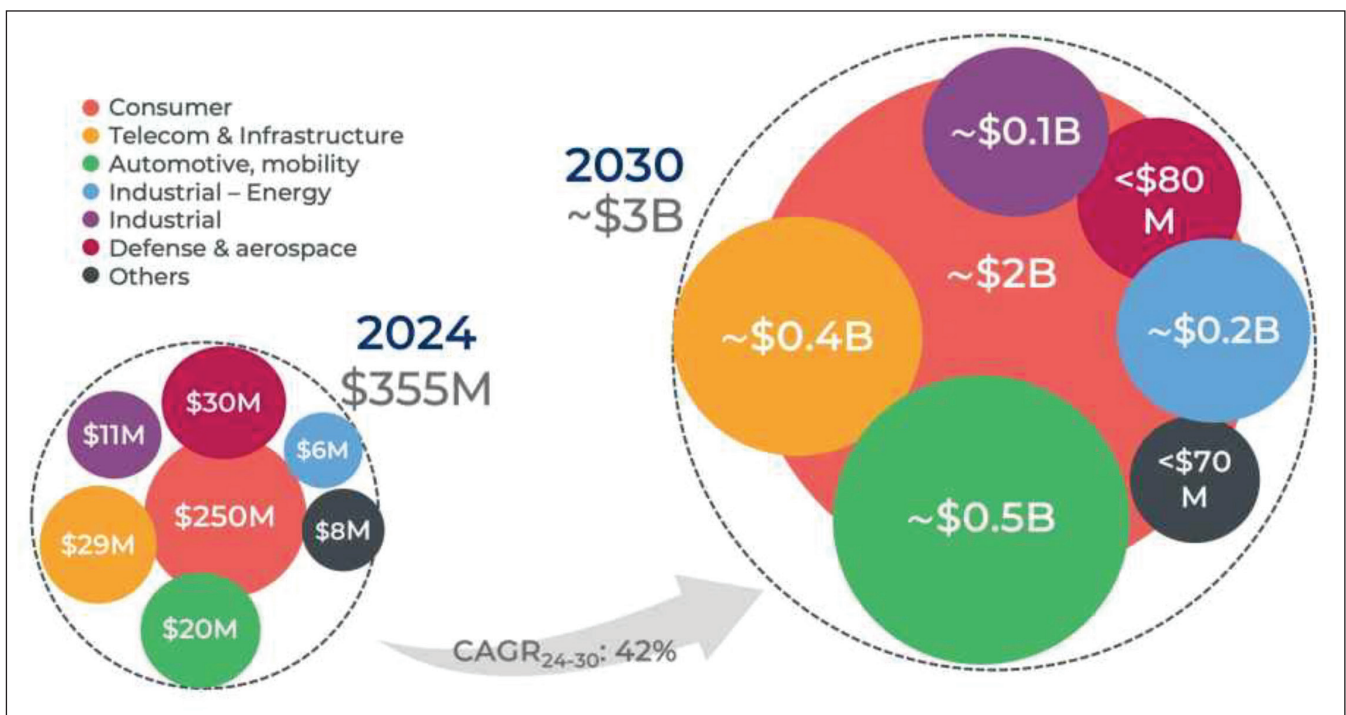
A booming \$3 billion market by 2030: the power GaN device market will have an impressive 42% CAGR from 2024 to 2030.

By market segment: Consumer electronics leading adoption: power GaN continues to dominate the fast-charger segment... Emerging momentum in automotive and mobility: despite short-term delays linked to the xEV market slowdown, this segment is expected to grow at a 73% CAGR between 2024 and 2030.

Data centers are accelerating the next wave: the AI boom is intensifying demand for high-efficiency power systems.

GaN is emerging as a key enabler for data centers. Together with telecom, this segment will have a 53% CAGR over this period.

Power GaN ecosystem: rapid consolidation amid intensifying competition and restructuring across the value chain.



➤ GaN power device market in \$m - Split by end-markets between 2024 - 2030.

New CEO to lead Wise Integration's global growth



Ghislain Kaiser, the recently appointed Chief Executive Officer of Wise Integration, explains how his main task is to transition the company from a CEA Leti spinout into a pioneering force in GaN and digital power management innovation with strong growth potential.

PEI: Please can we start with a little bit of an overview of your career to date, the roles you've had and the work you've undertaken in those?

GK: I began my career at ST Microelectronics at the late '90s, where I had a number of technical and leadership roles, spending time on test and product engineering and product engineering, design, and lead architect for different business units, set-up box, and the wireless business units.

So in 2006, actually, that was linked to my job at ST Microelectronics, where I was in charge of optimising the power consumption and the heat dissipation for wireless applications. I wanted to have better tools in order to do this job at the system level. And because I wasn't able to find anything that was very suitable for this job, I decided to

jump into the startup world and create my own startup to develop this new generation of system simulation tools.

So, I co-founded DOCEA Power, a French electronic design automation startup focused on full-chip power and thermal modelling. At DOCEA, our goal was to address, as I said, the increasing challenge of power consumption and thermal management in ICN platform design. And I served as CEO and led the company to become a recognised leader in this domain, resulting eventually in the acquisition by Intel in 2015.

This adventure actually also took me to relocate in the US, so from France to the US, and more specifically in the Silicon Valley, where we had our ecosystem, not only the main customers, but also the analysts,

partners, and all the electronic design automation ecosystem.

I stayed in the US for about nine years. First, as DOCEA CEO, and then I moved to the Intel headquarters in Santa Clara in 2015. I stayed at Intel for about a decade, holding senior director positions, and most recently leading global engineering and customer enablement programmes worldwide, focusing on system-level power, thermal, and performance challenges using simulation across the Intel main market segments. This included consumer electronics, laptops, desktops, notebooks, data centres, and, of course, AI, because now everyone talks about AI. Intel as well, of course, as a key leader in this domain, and AI from edge devices to network to cloud. And then I joined WISE.

PEI: What were the factors which made you want to join WISE?

GK: What attracted me to WISE was a unique convergence of deep tech innovation, strong foundational work, and a talented team. WISE has developed truly differentiated technologies in gallium nitrate power devices, and also a very innovative way of controlling that with digital. That's our differentiator.

We have a true R&D-driven culture, coupled with a clear mission to disrupt power electronics with smarter, more power-efficient systems, which is perfectly aligned with my own vision. I saw a great opportunity to help scale the company globally, not only in terms



of commercial reach, but also in terms of technological leadership.

PEI: *The company started life as a CEA-Leti spin-off. Your main task is to develop it into a fully commercial GAN technology, pioneering organisation?*

GK: Absolutely. We started from strong foundational work within CEA-Leti, a very well-known research centre in Europe, especially for semiconductors. And now, WISE is entering into a new phase in its journey. Having built this strong technology foundation since spinning out of CEA-Leti in 2020, the next step is to scale commercially, expand globally, because, as you know, we play on an international playground, and we want to establish ourselves as a leader in digitalisation of power converters. There are many things that are moving from the analogue world to digitalisation, with all the benefits of having digital control, having digital features, which leads to a more scalable application and reduce the time to market or development cycles.

PEI: *In order to help develop the company, you'll be looking at some of the key rapidly expanding markets, which include AI and the data centre and also the electric vehicle market. It would be good to have your thoughts as to the opportunities there are there in some of these rapidly developing markets?*

GK: We had our first successes with WISE in the consumer market, starting with applications below 500 watts for chargers and TVs, notebook chargers, smartphone chargers. Now that we are improving our roadmaps, that we are able to address higher power levels, we are expanding into AI data centres and AI servers - it's definitely a growing market. There are a lot of opportunities these days because we have some worldwide key leaders like NVIDIA that are pushing the transition from silicon to wide bandgap devices like gallium nitrate and silicon carbide. And we are at the right time to be a leader in this transformation. So that's exactly why we are positioning ourselves on our roadmaps to address this emerging AI data centre market.

But there is also the automotive innovation market that is going through this same transformation. These two markets are definitely our future for

growth opportunities. We know that those markets are very challenging, and we will not go alone. We will go with partners, in order to scale from an industrial point of view to address these very demanding markets.

And maybe last, but not least, this is also all the experience that I obtained at Intel in the data centre.

PEI: *In terms of the technology offering from WISE Integration, what are the USPs you have developed so far that make the solutions attractive to the market? And are you able to outline the roadmap as to how the technology may be developed - new solutions over time? What you can share on the technology side?*

GK: On the technology side, we have built two main platforms. One is WiseGan, a family of high efficiency, gallium nitrate power transistors. And the second one is WiseWare, our proprietary digital control solution using software, firmware, running on a microcontroller. This is the digital path, the digital control that I was talking about.

Together, these two solutions form a unique gallium nitride plus digital combination that enables high frequency operation for gallium nitride.

What attracted me to WISE was a unique convergence of deep tech innovation, strong foundational work, and a talented team. WISE has developed truly differentiated technologies in gallium nitrate power devices, and also a very innovative way of controlling that with digital

Switching the transistor at high frequency, a higher frequency than we have today with the existing solutions. We also use the digital control to enable what we call the soft switching, zero voltage switching, which are key features to develop a more compact and more power-efficient design while reducing the overall cost of these power converters applications.

I already gave a few examples of these - notebooks, smartphone chargers, AI server power supplies, rack power shelves in data centres, and also onboard chargers for electrical vehicles. All of these applications will benefit from either more compact capabilities or higher power efficiency (or both), meaning that we keep the volume the same, but we can put more power into the same volume.

A clear example of a benefit for an electrical vehicle is that if you have more power efficiency, it means that you can charge your vehicle with more power, so reducing the overall power charging timing.

So, these are all the benefits that we bring with this combination of GaN, power switch, smart power device and our power control using software running on an MCU.

PEI: *In terms of the future, is it further refining those solutions, or are there other solutions that are on the drawing board or in development? You possibly don't want to reveal too much, but just anything you can tell us as to what else you think the company is looking at?*

GK: Let me talk a little bit about our roadmap. It's driven by three customer benefits. One is improving compactness and power efficiency for power converters. To see how we can reduce the existing components, how we can remove some analogue parts on the existing power converters in order to improve the compactness.

We have some innovative patented new topologies that help reduce this volume for the same power or increase the power efficiency for the same volume.

Second is reducing overall cost for power converters. It's not only having more power in the same volume or more compact application, it's also doing all of that, this benefit, while reducing



the overall cost. And third is enabling scalability and additional high-value features, thanks to what we call the software-defined approach.

Because as you have understood, with our digital approach, we have a microcontroller which is running software. And with software, you have more flexibility than hardware to add features and also reduce the cycle to develop differentiated products - because of this flexibility of the software. All of this is made possible through our proprietary digital control, our system expertise for integrated solutions, and innovative architectures or topologies.

I can give you a few examples by way of illustration. We are developing system in package solutions that combine a discrete gallium nitrate power switch with custom analogue digital ASICs for smarter integration. And as we integrate more components into this system in package, we also reduce the bill of materials.

On the control side, our next-gen WiseWare controllers will support higher power interleaved multi-phase architectures to support higher power requirements and will embed intelligent algorithms for adaptive control and fault detection. And fault detection is one of the digital features I mentioned earlier, that can be very useful in a data centres

In order to address these new markets and to support these growth opportunities, we are expanding our global infrastructure. We have a design centre in Canada, subsidiaries in Asia - Hong Kong, South Korea, Taiwan, because there are a lot of customers, opportunities and partners in Asia in the power electronics domain. We are also building strategic partnerships, and we are strengthening our internal talent to support our growth objectives

for preventive maintenance. I can imagine that in the future we will add some AI algorithm in order to manage this type of new requirements in data centres or in other applications.

PEI: *Alongside the technology focus which you've outlined, are there plans to develop the company in terms of partnerships, whether it's existing ones expanding or forming new ones and/or by adding personnel. Please give us, if you can, an idea of how the infrastructure of the company is going to be built out?*

GK: In order to address these new markets and to support these growth opportunities, we are expanding our global infrastructure. We have a design centre in Canada, Sales and Technical support subsidiaries in Asia - Hong Kong, South Korea, Taiwan, because there are a lot of customers, opportunities and partners in Asia in the power electronics domain.

We are also building strategic partnerships, and we are strengthening our internal talent to support our growth objectives. As an example, we have a strong partnership with well-known research centres in Europe, like imec and CEA-Leti. We have also developed strong partnerships with all the ecosystem players, including IDMs, OEMs and distributors.

PEI: *In terms of what you've outlined for the future, do you see any particular challenges that you have to address, or is it mainly just some great opportunities for the company going forward?*

GK: There are always challenges, unfortunately, and especially these days, for a startup that plays in an international semiconductor playground. As you know, we have seen in the past five years that semiconductors have become a high-stakes topic subject to geopolitics. Competition is global with competition with startups, with established companies all over the world.

All of this is part of the challenges for any semiconductor startup these days. Our opportunity, though, is to leverage our agility and differentiated technology to create trust, because as we are a newcomer, we need to create trust, and position ourselves on the path to becoming a leader in the power electronics domain. We think that we have a strong, very solid technology and differentiated products that will help us overcome these barriers and get on the path to growth.

PEI: *In terms of how you measure the success of the plans that you have in place, do you have any specific objectives and timelines for achieving them, or is it a more relaxed approach that you know what you want to do, but there's no particular pressures to achieve certain things by a certain time? What does success look like over the next couple of years?*

GK: In the next two years, my main objectives with Wise will be, one, to establish the company as a leader in digital control for wide bandgap power electronics. If you notice here, I didn't say just gallium nitrate, but I said wide bandgap because our digital control is not specific to gallium nitrate. We can also use digital control for silicon carbide. We don't think that we are limited to one specific material, but we have opportunities to expand and address different stages of power levels, power requirements, thanks to our digital control approach. Two is to drive innovation to increase customer value and expand into key markets - data centres, AI, and automotive. That's definitely a second, big objective. And this will require scalability. It will require partnerships in order to grow and address these very demanding markets and customers with the right integration.

Third is that, in order to do this, we need to build a resilient and high-performing organisation. We already have a very strong, talented team. We will continue to expand this team in Europe, in Asia, and in North America, because you know that at the end of the day, what matters and makes the difference is not only the technology, but it's the team behind the technology.

And fourth is our objective to focus on sustainable growth and customer success. We have to put in place a solid discipline in order to have flawless execution and to establish strong relationships and to develop quality and differentiated products to ensure this

growth and customer success. These are the four objectives that I have in mind for the next couple of years.

PEI: *Any final thoughts as to your new role at Wise Integration and/or how you see the company's place in the power electronics industry into the future?*

GK: I'm excited to lead Wise Integration at a pivotal moment. The industry is moving toward more intelligence and energy-conscious systems. The gallium nitrate and digital combination that we put in our products is exactly what is needed in order to address these industrial challenges.

For me, this road is more than leadership. It's about powering the innovation of tomorrow while contributing to the global energy transition. Why the innovation of tomorrow? Because you know that with AI and all the AI workloads that we need to support in the future, we will need to have massive computers available. And all of that, the GPUs and CPUs, will consume a tonne of power, and this requires innovative solutions like the product developed by Wise, in order to have sustainable growth.

It will be the same with electrical vehicles, this is also part of this energy conscious transition. We are very happy at Wise to contribute to a better future. This is what's driving our team, why people wake up every morning with this motivation, because they feel that at Wise they contribute to something that is big and absolutely required, very urgent for the future world.



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

SAVE THE DATE

20-22 APRIL 2026

Sheraton Brussels Airport Hotel

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

SMT inductors with coils and ferrites

How do you choose the right inductor for your application? This article compares the differences between three inductive SMT components: ceramic inductors, SMT ferrites, and wire-wound ferrite inductors. Their electrical properties make them suitable for different applications, and their physical parasitic properties can be used to your advantage in circuit design.

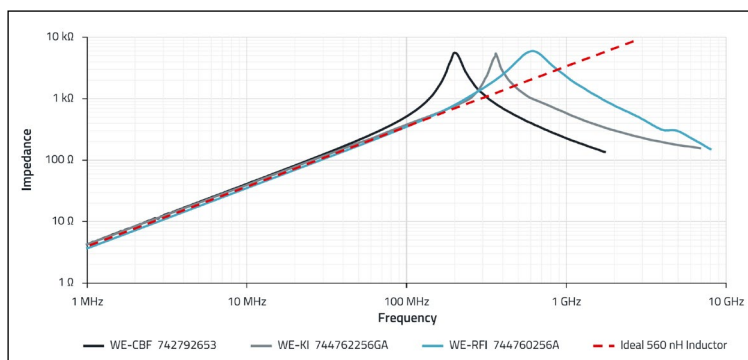
BY DR. HEINZ ZENKNER, WÜRTH ELEKTRONIK EISOS

INDUCTIVE COMPONENTS are available in numerous forms. SMT versions are very popular, as they are easy to mount on circuit boards. However, not all inductors are the same, as this article demonstrates using three representatives of SMT inductors: ceramic inductance, SMT ferrite, and wire-wound ferrite inductance.

Figure 1 shows the impedance curves of the three inductor types in comparison:

- **SMT ferrite:** WE-CBF [1].
- **SMT inductor with ceramic core (“air coil”):** WE-KI [2].
- **Wire windings on ferrite core (“wire-wound ferrite”):** WE-RFI [3].

The components were selected so that they have a similar impedance curve in the range below their impedance peaks. The differences in the impedances can be seen in the area of the peaks, the SMT ferrite has its maximum at the lowest frequency, while the wire-wound ferrite at the highest. The ceramic inductor shows the steepest rise and fall in the area of the impedance maximum and therefore the highest quality factor Q.



➤ Figure 1: Comparison of the impedances of SMT ferrite, ceramic/air inductance and wire-wound ferrite.

Electrical parameters in comparison

For inductors without a ferrite core (WE-KI), the inductance is given as a value in the data sheets, in this case 560 nH. Although the inductance value is specified at a particular measuring frequency, the value below the resonant frequency is almost constant against frequency.

On closer inspection, the reactance of the inductance does not increase linearly with frequency. Two effects increase the impedance of the air-core coil: the increase in reactance due to Lenz’s law and the increase in resistance due to the skin effect.

In the purely inductive circuit, the coil is connected directly to the AC supply voltage. As the voltage rises and falls with frequency, the self-induced counter-electromagnetic force (EMF) in the coil also rises and falls as a function of this change. This self-induced back EMF is directly proportional to the rate of change of the current through the coil (Lenz’s law) and therefore increases with frequency. Consequently, the reactance of the inductance also increases with frequency; this relationship is proportional.

A further increase in impedance is caused by the skin effect. At low frequencies, a conductor uses its entire cross-sectional area as a transport medium for charge carriers. If the frequency increases, an increased magnetic field in the direction of the center of the conductor represents an impedance for the charge carriers, causing the current density in the center of the conductor to decrease and the current density at the edge of the conductor to increase. This increased current density near the edge of the conductor is known as the skin effect. The effect increases with frequency and also occurs with all other inductors (with a ferrite core).

The resonant frequency of the inductance without a ferrite core is primarily caused by the parasitic capacitance between the individual windings. Whenever two conductors are arranged in close proximity but separated by a dielectric and there is a voltage difference between them, a capacitor is formed.

The chain of these winding capacitances is connected in parallel to the winding inductance and thus forms a parallel resonant circuit. In addition, there is a parasitic capacitance between the connections (solder pads), which is parallel to the winding capacitance. This results in a parasitic total capacitance in parallel with the winding as an equivalent circuit. The equivalent circuit is shown in Figure 2 on the left.

For inductors with a ferrite core (WE-RFI) and SMT ferrites (WE-CBF), the data sheet does not specify an inductance value, but rather an impedance at a measuring frequency. It can also be seen that the SMT ferrite has the highest tolerance, whereas the inductance without ferrite has the lowest tolerance.

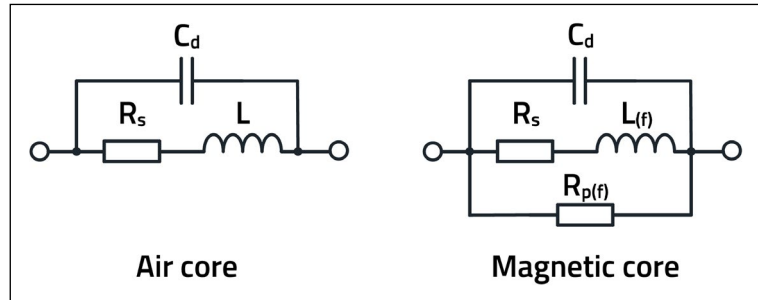
Since the WE-KI inductor does not have a ferrite core, it requires more turns of wire for the same impedance than the components with ferrite, which explains why the WE-KI also has the highest wire resistance R_{DC} . A Q factor, i.e. a quality factor, is specified for both the WE-KI and the WE-RFI, but not for the WE-CBF. The Q factor is a measure of the dissipative property of an inductor. Inductors with a high Q factor have low losses and a narrower impedance curve. Inductors with a low Q factor, on the other hand, have higher losses and a broader impedance curve.

The magnitude of the maximum impedance of the inductor is related to the quality factor Q. Low-loss inductors with a high Q factor have a very high maximum impedance, while an inductor with high losses has a lower maximum impedance. By changing the way an inductor is wound or the core materials used, the impedance maximum and the frequency range of the impedance maximum can be adjusted.

Magnetic core materials

The WE-RFI inductor and the WE-CBF SMT ferrite use ferrite materials as the core material. The diagrams each show three different curves, R - as resistive (ohmic) resistance, X_L - as reactance (inductive) and Z - as the value of the component's impedance. It is important to understand these diagrams in order to successfully use inductors with ferrite material in circuits.

In many radio frequency applications where large inductance values are required in a small space, inductors with an "air core" cannot be used due to their size. The design of the inductor can be smaller if the air core is replaced by a core material with a higher magnetic permeability ($\mu_r > 1$).



➤ Figure 2: Equivalent circuit of an inductor without magnetic core (left) and with magnetic core (right).

If the size remains the same, the inductance value is maintained despite the reduced number of windings.

This allows several advantages to be achieved:

- **Smaller size** - due to the smaller number of turns required for a given inductance.
- **Increased Q** - fewer turns mean less wire resistance.
- **Adjusting the impedance of the inductance over frequency** - through targeted selection/mixing of the core material.

However, the use of magnetic cores presents some major problems and care must be taken to ensure that the core material chosen is the right one for the job. These problems can be overcome if care is taken during the development stage to ensure that the inductors are selected correctly for their intended use. For this purpose, the impedance diagrams with the three different impedance curves R, X_L and Z are required. The behavior of the curves as a function of the frequency depends strongly on the magnetic properties of the core material.

The AppNote ANP129 [4] show "Inductive SMT Components in Comparison – the Wire Makes the Difference" provides in-depth information of the most important parameters of the three inductor types and their preferred fields of application.


FURTHER READING

- [1] SMT ferrite WE-CBF from Würth Elektronik: <https://www.we-online.com/en/components/products/WE-CBF>
- [2] SMT inductor with ceramic core from Würth Elektronik: <https://www.we-online.com/en/components/products/WE-KI>
- [3] Wire-wound SMT bead ferrite from Würth Elektronik: https://www.we-online.com/en/components/products/WE-RFI_FERRITE_BEAD
- [4] Zenkner, H.: Inductors, SMT ferrites and wire-wound SMT ferrites - The wire makes the difference. AppNote ANP129 from Würth Elektronik: <https://www.we-online.com/ANP129>

AIXTRON

THE NEW G10 SERIES

Your Productivity Solution for All Advanced Epitaxy Materials



G10-SiC

- ▶ 9x150 mm (6") or 6x200 mm (8")
- ▶ New hardware & process surpassing Single Wafer Reactor uniformities
- ▶ +50% productivity per fab area with large batch technology & small footprint

End Markets/Products:

EV inverters & charging infrastructure

G10-GaN

- ▶ 8x150mm (6") or 5x200 mm (8")
- ▶ 1st fully automated compact GaN MOCVD cluster designed 100% for Si Power fabs
- ▶ Novel hardware solution for unmatched barrier uniformities and device yields

End Markets/Products:

Power Electronics & Wireless communication



G10-AsP

- ▶ 8x150 mm (6") or 5x200 mm (8")
- ▶ 1st fully automated AsP MOCVD system enabling 10x lower defect density
- ▶ Unmatched wavelength uniformity on all wafer sizes

End Markets/Products:

Micro LED, Optical Data communication, 3D-sensing & LiDAR

