Powering and Sensing AI: From Cloud Infrastructure to Physical Autonomy

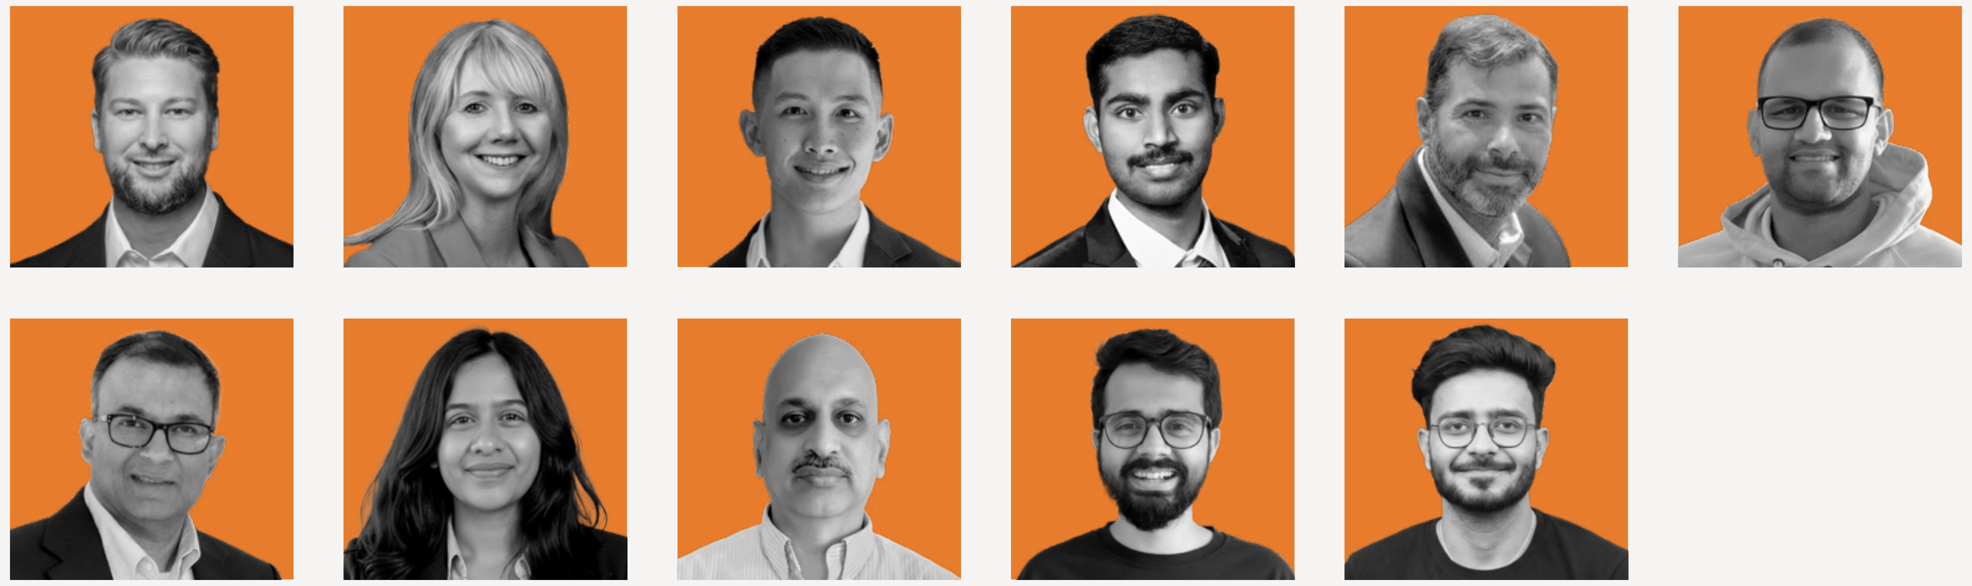

Mighty Guides

Sponsored by

onsemi

MOUSER ELECTRONICS

Table of Contents

Meet Our Experts

Today's AI systems—spanning cloud infrastructure, autonomous machines, and industrial automation—are increasingly defined by the demands of efficient power delivery and high-fidelity sensing. We interviewed leading experts about the shift from viewing AI as a purely software-driven domain to recognizing it as a full-stack engineering challenge grounded in physical constraints. This eBook explores how engineers are addressing challenges in energy efficiency, power density, real-time perception, and system reliability through advanced power architectures, wide-bandgap semiconductors, and intelligent sensing technologies that enable scalable, real-world AI deployment.

David Rogelberg
Publisher,
Mighty Guides Inc.
david@mightyguides.com
(516) 234 2969

Brandon Becker
Product Marketing Manager SiC JFETs, onsemi

Felicity Carson
Chief Marketing Officer, onsemi

Ee Kin Chin
Senior Machine Learning Engineer, Meta

Naveen Dodda
Applied Machine Learning Engineer, Google

Alessandro Maggioni
Sr. Manager, Regional Marketing EMEA, AMG, onsemi

Pranav Maheshwari
Staff Machine Learning Engineer, Tesla

Ganesh Narayanaswamy
Senior Business Marketing Manager, onsemi

Aarati Noronha
Machine Learning Engineer, Amazon AGI

Prasad Paruchuri
Technical Marketing Team Leader, Power Solutions Group, onsemi

Shreyash Rawat
Machine Learning Engineer, Meta

Vipul Tiwari
Senior Solutions Architect for Agentic AI and FinOps Center of Excellence, Amazon

Foreword

By **Felicity Carson**, Chief Marketing Officer, onsemi

Artificial intelligence (AI) is often framed as a breakthrough in software—defined by models, algorithms, and data. But as AI moves from abstraction into real-world deployment, it becomes clear that its true limits are physical. Every inference, decision, and autonomous action depends on two foundational capabilities: efficient power and high-fidelity sensing.

Today's AI inflection point is driven not only by exponential growth in compute, but by the strain that growth places on infrastructure. Data centers are confronting unprecedented power densities, while robots, vehicles, and autonomous machines must perceive and respond to the physical world under tight energy, latency, and safety constraints. In both environments, inefficiency at the hardware layer compounds quickly, hence driving higher cost, greater energy consumption, and reduced reliability.

This eBook explores why power delivery and sensing are no longer supporting functions of AI, but strategic enablers. From high-voltage DC architectures and wide-bandgap semiconductors in the cloud to intelligent image sensors, analog front ends, and edge integration in physical systems, the quality of the underlying hardware determines what AI can safely and sustainably achieve.

At onsemi, we focus on making AI possible in the real world—by powering it efficiently, sensing it intelligently, and integrating both into scalable, production-ready platforms. As intelligence moves from the cloud into machines that interact with people and environments, solving for power and perception will define the next era of AI.

onsemi delivers intelligent power and sensing technologies that enable electrification, energy efficiency, safety, and automation across automotive, industrial, and AI data center end-markets. With a highly differentiated and innovative product portfolio, onsemi helps customers solve complex challenges to achieve higher efficiency, improved performance, and lower system cost, while supporting a safer, cleaner, and more energy-efficient world. The company is part of the S&P 500® index. Learn more at www.onsemi.com.

Introduction

People often talk about AI as a software revolution, with the attention heavily placed on models, algorithms, and training data. However, that framing misses that every inference, robotic action, and real-time decision depends on power and sensing.

The industry has spent years optimizing the model layer, but AI's infrastructure has become the latest bottleneck. Data centers are straining under the power demands of graphics processing unit-dense (GPU-dense) AI workloads, and physical AI systems need to perceive and respond to the world in real time.

While not obvious, digital AI and physical AI belong in the same category. Cloud-based intelligence trains and informs physical systems, and physical AI systems must then act on that intelligence locally, with limited power and in unpredictable environments. What happens in the cloud has a direct bearing on what is possible with physical AI.

Ultimately, the quality of the underlying sensing and power infrastructure determines whether AI delivers on its promise or falls short. Low-quality sensor data forces processors to work harder and, in turn, increases power draw and latency. Similarly, inefficient power delivery can introduce noise, cause brownouts, and undermine system reliability. Together, these failures can have a compounding effect, in which a degraded sensor signal feeding an underpowered processor produces slow, inaccurate, or dangerous decisions. It's for those reasons that edge AI requires solving for power and perception at every level of the stack.

In this eBook, we'll explore the growing energy demands of AI infrastructure and the sensing technologies behind physical autonomy, and how onsemi is responding to these challenges with its EliteSiC Silicon Carbide, gallium nitride (GaN), Hyperlux image sensors, and Treo Analog and Mixed Signal Platform technologies.

POWERING THE AI BUILDOUT

AI workloads have fundamentally changed the power requirements of the data center.

For example, rack-level power has grown significantly from roughly 3.3 kW just a few years ago to 12 kW today, and next-generation AI racks are targeting 30, 120, and 240 kW per rack. At the extreme, by 2028, Open Compute Project requirements point toward 1 MW rack-level deployments. These power levels pose an entirely different class of infrastructure challenge.

The move from 12 to 48 V and now to 800VDC bus architectures is a direct response to these demands. At 12 V, conduction losses across a high-current bus become unsustainable at scale. Moving to 48 V minimizes those losses through decreased current at equivalent power, and ±400 V or 800 VDC reduces them further at higher power. For the designer, the goal is to eliminate unnecessary power conversion stages. Going directly from ±400 to 12 V output, for example, removes an entire intermediate bus converter step. Some customers are exploring output voltages as low as 6 V to further streamline the conversion path.

In data centers and other industrial applications, these higher power densities create serious thermal and electrical stress points throughout the system.

“

By tailoring an integrated circuit (IC) specifically for a task, application-specific integrated circuits (ASICs) enable high-intensity computations in everyday devices through low-cost production and reduced thermal issues. While often underestimated due to their rigid nature, these components are essential for successful AI scaling.”

Naveen Dodda

Applied Machine Learning Engineer, Google

> ***The transition to 800VDC AI infrastructure depends on SiC JFETs, which provide the backbone for solid-state protection, seamless high-voltage hot swap, and high-efficiency DC-DC power conversion."***

Brandon Becker

Product Marketing Manager SiC JFETs, onsemi

For example, factory automation equipment, industrial robots, and autonomous mobile platforms all face the same fundamental tradeoff: as power density increases, so does heat, and thermal management becomes as much of a design constraint as efficiency. In a data center rack, this means liquid cooling and chillers. In a robot or drone, where heat dissipation options are far more constrained and form factor is fixed, the problem is more acute. Without highly efficient conversion at every stage, cooling requirements escalate quickly. For most, the priority should be to reduce power loss before adding cooling, because efficient conversion is cheaper and more effective than managing waste heat after the fact.

Wide-bandgap (WBG) semiconductor materials such as gallium nitride (GaN) are central to solving these challenges. For example, lateral GaN 650 V devices suit the intermediate bus converter (IBC) function, as they can operate at the megahertz-level switching frequencies required for high-voltage IBC applications. Similarly, SiC cascode-JFETs can operate in the megahertz ranges needed for IBC while also supporting the 1200 V processes needed to operate within 800 V buses.

Vertical GaN devices, while currently under development, can extend this behavior even further by offering lower reverse recovery charge (Qrr), lower Coss, lower junction-to-case thermal resistance, and operation up to 2 MHz.

Overall, grid-to-processor efficiency has direct implications for the total cost of ownership in AI-driven applications. Consider the typical example of a high-voltage 800 V architecture in hyperscale data centers. Here, designers can perform grid conversion to the DC bus via a rectifier or solid-state transformer, which then passes down the DC bus to the GPU. Inefficiency at any stage in this power tree can multiply across the entire fleet, where even a 1-2% loss in efficiency can translate into megawatts of additional losses in a large data center. In practical terms, that non-ideal behavior results in more CO_2 emissions, greater cooling demand, higher electricity costs, and additional stress on grid infrastructure.

The same logic applies to other industrial applications, such as factory automation systems, drones, and autonomous robots. Where systems draw from batteries or local power supplies, every percentage point of conversion efficiency directly determines how long the system runs, how much it weighs, and how much it costs to operate.

onsemi supports AI data center power by:

- Delivering EliteSiC MOSFETs and JFET-based solutions for high-efficiency AC-DC conversion at rack-level power densities up to and beyond 100 kW.
- Advancing vertical GaN technology that reduces switching losses, enables 2 MHz operation, and provides lower thermal resistance than existing lateral GaN or SiC alternatives.

The core tradeoff is between power density and thermal headroom—you can pack more compute into a smaller form factor, but heat removal doesn't scale the same way. The shift is toward heterogeneous designs: offloading workloads to purpose-built accelerators within a watt budget."

Shreyash Rawat

Machine Learning Engineer, Meta

"As AI scales, power is emerging as a first-order constraint alongside compute. Larger models deliver higher capability but at significantly higher energy and heat costs. This is driving a shift toward optimizing AI efficiency per watt rather than raw compute alone."

Ee Kin Chin

Senior Machine Learning Engineer,
Meta

- Offering 1200 V SiC cascode-JFET solutions for IBC that can operate in the megahertz frequency range.
- Providing reference designs for 100 kW three-phase AC-DC supplies, 12 kW single-phase solutions, and high-voltage IBC designs.
- Offering the Treo platform as a monolithic point-of-load solution for the next generation of 6V-input, sub-1V-output multiphase power delivery to AI processors.

"If you have high power losses, then you need to add more liquid cooling technologies and need to spend more on chillers and other thermal management tools. For end applications, you need high efficiency and high density—reliability is also important in cloud data centers."

Prasad Paruchuri

Technical Marketing Team Leader,
Power Solutions Group,
onsemi

NTBL023N065M3S Silicon Carbide (SiC) MOSFETs

Learn More

UF3N120007K4S 1200V JFET N-Channel Transistor

Learn More

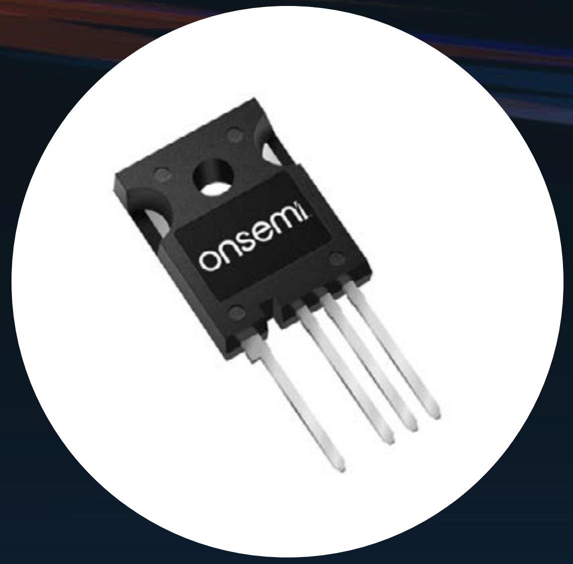

UG4SC 750V 8.4mΩ Combo-FETs

Learn More

Key Points

- **Rack-level AI power requirements are scaling from 12 kW today toward 1 MW by 2028.**
- **Transitioning to 800 V DC bus architecture reduces conduction losses and can eliminate entire conversion stages.**
- **A 1-2% improvement in power conversion efficiency translates to megawatts of savings at hyperscale, with meaningful reductions in CO_2 emissions, cooling infrastructure, and grid demand.**
- **SiC, lateral GaN, and emerging vertical GaN each offer distinct benefits in the AI power tree.**

Chapter 2

THE SENSING FOUNDATION OF PHYSICAL AI

While vision is the most prominent sensing modality in physical AI, it is far from the only one. A robot that can see but cannot feel has no way to detect the force it's applying—it can damage parts, injure humans, or fail at precision tasks. In practice, the true full sensing stack for physical AI includes force and torque, tactile and pressure feedback, proximity, depth, motor current, and more, where each modality contributes information that the others can't supply.

Force and torque sensing at joints and end-effectors is a prerequisite for compliant motion, collision detection, and high-precision assembly.

> *A poor quality sensor output will force the processing engine to do a lot more work. Cost increases, power increases, and everything goes downhill from there. When sensing outputs are high quality, you really get the processing down to the bare minimum needed to make an effective output."*

Ganesh Narayanaswamy

Senior Business Marketing Manager, onsemi

> ***"Sensing quality decides the ceiling of your eventual performance. Compensating for lower-quality sensing requires bigger and more sophisticated models, which can also be prone to hallucinations in unexpected ways."***

Pranav Maheshwari

Staff Machine Learning Engineer, Tesla

Without this feedback, a robotic arm has no way to modulate its grip or respond to unexpected resistance. Additionally, the speed and precision of this feedback loop directly determine the robot's ability to interact safely with its environment. Tactile sensing, enabled by pressure analog front-ends, adds awareness of contact, slip, and pressure for grasping and manipulation tasks, while ultrasonic sensing can provide proximity detection without line-of-sight requirements.

Similarly, motor current and load sensing can provide real-time insight into interaction forces for adaptive control loops that go beyond what vision alone can offer. This feedback enables robots to respond intelligently to physical contact rather than only to what they can observe. In environments where robots operate near humans, such capabilities can save lives.

Regardless of modality, sensor quality has a direct and compounding effect on the computational cost of AI training and inference. Poor-quality sensor data forces the processing engine to do more work to extract usable information, demanding higher-end silicon, more memory bandwidth, and more power to compensate for data that should have been better at the source. The downstream effects include higher system cost, increased latency, and, in some cases, unsafe decisions.

In addition to higher-quality sensors, another way to reduce the burden on central compute is to use intelligent edge-processing sensors. Rather than streaming raw data to a processor and expecting it to manage bandwidth, denoise to increase fidelity, and handle redundancy to optimize compute, intelligent sensors can locally perform subsampling (which minimizes bandwidth), denoise, and enhance dynamic range to deliver high-fidelity outputs.

Reducing data movement is one of the most effective ways to decrease latency and power consumption in the sensing pipeline. Protocols like 10BASE-T1S with Remote Control Protocol (RCP) go further by removing local microcontroller units (MCUs) from sensor nodes entirely, reducing bill of materials (BOM) complexity and software update overhead.

Finally, sensor fusion ties the entire sensing stack together. Where physical AI systems rely on vision, depth, proximity, pressure, motion, and thermal sensors, the processing engine must fuse these otherwise disparate inputs into a unified model of the environment.

However, designers must implement sensor fusion in a way that provides safety through redundancy, such that, if one sensor fails, the rest of the system can continue to operate. New devices with redundancy coordination (e.g., detecting coordinator failure in a T1S chain and automatically switching to a backup) and certified for higher safety integrity levels (SIL) are a meaningful step toward bringing safety right to the edge.

"Bad input data pushes complexity downstream. You end up adding more processing or heavier models just to stabilize outputs. Good sensing simplifies the entire pipeline and reduces both compute cost and latency."

Vipul Tiwari

Senior Solutions Architect for Agentic AI and FinOps Center of Excellence, Amazon

> *High-quality sensors increasingly include onboard processing, where filtering, compression, and/or early feature extraction are performed directly on the sensor before data is sent downstream. This can minimize data movement, lower bandwidth, and compute requirements in addition to improving end-to-end latency.”*

Aarati Noronha

Machine Learning Engineer,
Amazon AGI

onsemi supports physical AI sensing by:

- Providing Hyperlux image sensors with built-in Smart ROI, wake-on-motion, and subsampling capabilities that reduce load on AI processors.
- Enabling intelligent edge processing at the sensor level to reduce raw data movement, lower latency, and decrease power and silicon costs.
- Offering analog front-end solutions for force, position, tactile, and proximity sensing that complete the non-vision sensing stack.
- Advancing 10BASE-T1S connectivity with RCP-enabled sensor node architectures that reduce local MCU count and improve sensing chain reliability and redundancy.

> *"Poor sensor data compounds at every layer of the stack. Noisy inputs force models to be larger and more uncertain, increasing compute load and inference latency. Bad sensing doesn't degrade accuracy linearly, it can trigger failure modes that spike power draw and system cost."*

Shreyash Rawat

Machine Learning Engineer, Meta

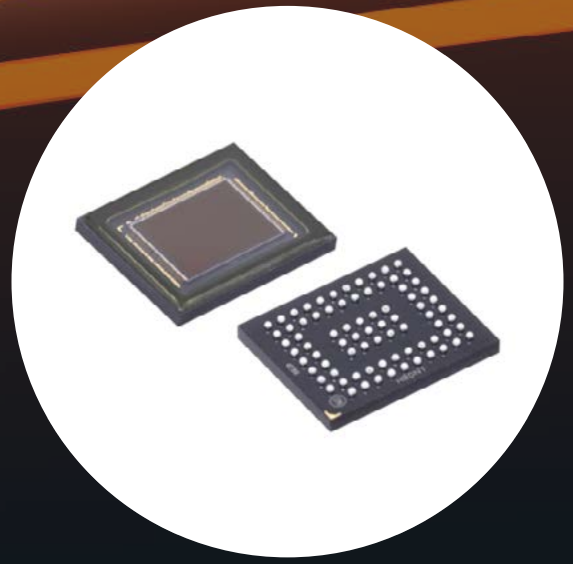

AR2020 Hyperlux™ LP 20 MP Image Sensors

Learn More

NCS32100 Inductive Position Sensor

Learn More

Key Points

- **Physical AI requires sensing well beyond vision for safe, precise interaction with the physical world.**
- **Poor sensor data quality forces processors to overcompensate, increasing silicon requirements, power draw, latency, and system cost while reducing reliability.**
- **Intelligent edge sensors that perform subsampling, dynamic range enhancement, and smart region-of-interest processing reduce the volume of data reaching the processor.**
- **Sensor fusion across multiple modalities provides fault tolerance and keeps physical AI systems safe when individual sensors degrade or fail.**

INTELLIGENCE MOVES INTO THE REAL WORLD

Designers have optimized cloud-based AI models for controlled environments with no physical constraints. But when intelligence moves into the physical world, those assumptions break down entirely.

Real-world systems must contend with variable power budgets, mechanical vibration, thermal stress, unpredictable scenes, and latency requirements measured in milliseconds. Models trained in the cloud have no inherent preparation for these conditions, and the burden falls on the hardware layer.

In sensing, latency is a safety parameter. In the example of a manufacturing or warehouse setting, a robot that detects a human obstacle 100 ms too late may not be able to stop in time. In that context, reducing latency in the sensing pipeline directly expands the safety margin of every physical AI system. Cobots and automated guided vehicles (AGVs) operating near people need deterministic, low-latency sensing to slow, stop, or re-route without relying on hard safety stops that interrupt production. The goal is sensing fast enough to keep the system running safely without exceeding power budgets.

Naturally, physical AI systems need efficiency at every level. Whereas industrial robots can draw from facility power, battery-powered autonomous mobile and humanoid robots must carry their energy source.

“Reducing latency is the critical aspect for safety. Robots are becoming faster, and are interacting more with humans. Having the lowest-latency sensing possible optimizes both safety and efficiency.”

Alessandro Maggioni

Sr. Manager, Regional Marketing EMEA, AMG, onsemi

"Low-latency sensing matters most in closed-loop systems like robot obstacle avoidance or drone control, where even small delays affect safety and control quality. In field robotics, keeping processing on-device improves consistency and reduces dependence on external infrastructure."

Ee Kin Chin

Senior Machine Learning Engineer, Meta

Therefore, every milliwatt saved in the sensing and power management layer extends battery life and allows designers to deploy more capable compute within the same energy budget. That trade-off among sensing quality, computational capability, and energy consumption ultimately defines what a system can actually do in the field.

For these reasons, system-level integration is the next major frontier for physical AI developers. Selecting the best individual components is no longer sufficient. Engineers need to design the interactions among sensing, power, and compute holistically. With highly-integrated off-the-shelf solutions, engineers can take a system that previously required multiple discrete components and significantly reduce form factor, BOM complexity, and power overhead. These systems can include integrated analog front-ends that process sensing data locally and advanced power management that delivers compute resources on demand to remove the idle power overhead that normally plagues always-on edge systems.

onsemi supports real-world AI deployment by:

- Integrating intelligent sensing and efficient power management in products within the Treo platform.
- Offering low-power image sensor architectures that allow designers to deploy more sensors per system within the same thermal and energy envelope.
- Enabling analog front-ends with higher bandwidth that translate directly to lower-latency feedback loops.
- Providing system-level reference designs and integrated platforms that reduce BOM complexity and minimize form factor for battery-powered autonomous machines.

"The idea is to provide system-level integration. Integrating more intelligent power and intelligent sensing into the system means that every milliwatt here and there can help to extend the battery life of the system, achieve more payload, or extend range."

Alessandro Maggioni

Sr. Manager, Regional Marketing EMEA, AMG, onsemi

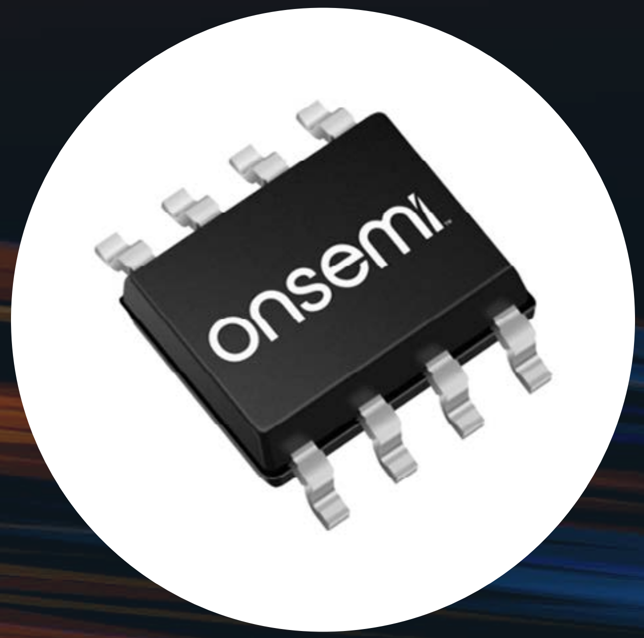

NCV7192 Bridge Signal Conditioning Interface

Learn More

Key Points

- **Real-world AI operates under constraints that cloud-based models are not meant for.**
- **Sensing latency is a safety variable, as milliseconds of additional response time can determine whether a human interaction is safe or dangerous.**
- **System-level integration is what makes high-performance physical AI viable within the strict energy budgets of untethered machines.**
- **Products within onsemi's Treo platform reduce the sensing-to-actuation loop by integrating sensing, power, and edge inference locally.**

Chapter 4

FROM ROBOTICS TO PHYSICAL AUTONOMY

The hardware behind physical autonomy is more mature than the public narrative suggests. Deployable sensing and power solutions already exist. The remaining challenge is less about inventing new hardware and more about integrating it intelligently, scaling it cost-effectively, and writing the software and AI layers that make it coherent. The next three to five years of physical AI progress will largely depend on how well the industry assembles the hardware stack.

A prime example is cobots, one of the most demanding integration challenges in physical AI. A cobot must operate safely in close proximity to humans without halting production whenever a person enters its workspace. That requires multi-layer, real-time sensing, with force and torque at the joints, rotational and angular positioning sensing, motor current monitoring, and tactile feedback all working in concert. Meanwhile, vision adds an outer detection layer, where more cameras per robot means better spatial awareness and earlier hazard detection. Engineers will commonly design humanoids with six to eight cameras because the additional perception quality justifies the cost.

In humanoid robotics, power design is challenging because power requirements can vary dramatically across body segments.

“In a humanoid robot, a hand has five fingers, and each finger requires a motor and dedicated control. That means your control plus FETs for the gate drivers all must be integrated into a single chip.”

Prasad Paruchuri

Technical Marketing Team Leader,
Power Solutions Group,
onsemi

"Better integration matters most because AI must work in harmony and perform multiple tasks across diverse environments. Intelligence is adaptive and fluid, requiring constant feedback and seamless integration between I/O sensors and compute."

Naveen Dodda

Applied Machine Learning Engineer, Google

Hand-level actuation may require less than 500 W, with multiple brushless DC motors per hand that need compact, integrated gate drivers and control. In contrast, hip and knee joints can operate at up to 4 kW, requiring high-efficiency, high-frequency devices and potentially 48- or 60-V bus architectures. Given the extreme variability in power requirements and power densities, integration is the clearest path to space-efficient solutions.

A similar field, autonomous mobile robots (AMRs), is maturing quickly. Early AMRs depended on fixed maps and controlled environments where obstacles didn't move. Today, AMRs can handle unstructured environments with dynamic obstacles, human co-workers, unmapped areas, and unpredictable terrain. Real-time sensing, low-latency edge AI, and adaptive control loops make this possible, as the robot continuously perceives and reinterprets its environment rather than executing a pre-planned path.

Still, the move from a flat warehouse floor to uneven outdoor terrain fundamentally changes the power profile of a mobile robot. On a flat surface, power demand is relatively steady and predictable.

In dynamic environments, variable slopes and constant torque corrections can drive higher average power and much larger short-duration power spikes. In turn, actuators draw higher peak and transient currents, while sensing and compute loads increase to maintain control. Future power systems, therefore, must handle this dynamic range without sacrificing efficiency at either extreme.

Ultimately, mass deployment of physical AI systems will require hardware that is smaller, cheaper, and more integrated. Today's systems often rely on multiple discrete components for functions that must be consolidated, but the path to deploying humanoids and autonomous systems at scale runs through integration.

onsemi supports physical autonomy by:

- Providing integrated motor drive solutions that fit the form factor and power constraints of humanoid joints and robotic end-effectors.
- Delivering the sensing technology for cobots, including image sensors, analog force/torque front-ends, and proximity sensing.
- Offering high-efficiency GaN and silicon power devices in integrated half-bridge configurations.
- Advancing its EliteSiC and imaging portfolios in parallel with the automotive-grade and industrial-grade requirements of autonomous vehicles.

“

In collaborative robots, safety comes from real-time, multi-layer sensing rather than stopping motion. Force and torque sensing, high-accuracy position sensing, motor current monitoring, and tactile feedback continuously detect human contact or proximity. The robot dynamically reduces torque or changes trajectory in milliseconds.”

Alessandro Maggioni

Sr. Manager, Regional Marketing EMEA, AMG, onsemi

"Real-world robots like self-driving cars operate on a fraction of the compute used in cloud-based solutions, with much tighter latency constraints. You need to co-design your network architecture and edge inference accelerator for maximum throughput and efficiency."

Pranav Maheshwari

Staff Machine Learning Engineer, Tesla

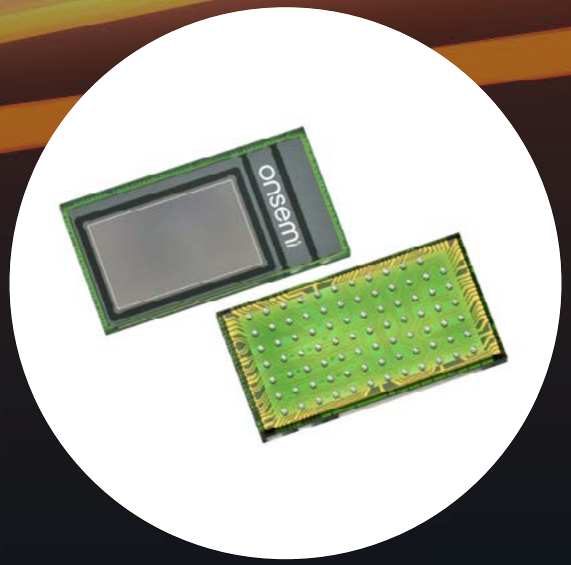

AR0235 Digital Image Sensors

Learn More

NCN26010 Industrial Ethernet Transceiver

Learn More

Key Points

- **The next challenge for physical autonomy is intelligent integration and the development of coherent, scalable software and AI layers.**
- **Cobots need multi-layer real-time sensing to comply safely with human proximity rather than defaulting to hard stops that interrupt production.**
- **Power requirements in a humanoid vary by more than an order of magnitude between body segments.**
- **Moving AMRs from structured to unstructured environments requires edge AI, adaptive control, and dynamic power systems.**

Conclusion

The systems of 2030 will look markedly different from today's. We see a future with more sensors per robot, higher compute density, lower power per inference, and tighter integration between sensing and actuation. At the same time, the AI layer will become more capable as the hardware layer becomes more efficient, but only if engineers can solve the power and sensing problems outlined in this eBook. Fortunately, onsemi has developed a roadmap for making AI physically possible and energy-efficient, from the high-voltage DC bus powering the data center rack to the integrated motor drive built into a robotic finger.

Learn More About Our Experts

Brandon Becker

Product Marketing Manager SiC JFETs, onsemi

Brandon Becker is an accomplished Product Manager with over a decade of experience in the semiconductor industry, specializing in wide-bandgap technologies (SiC and GaN), power MOSFETs, and digital interface solutions. Since 2014, Brandon has played a pivotal role in driving product strategy, business planning, and market expansion across multiple technology domains.

Felicity Carson

Chief Marketing Officer, onsemi

Felicity Carson is Senior Vice President and Chief Marketing Officer at onsemi, leading Corporate Marketing, Customer Experience, and ESG & Sustainability. Since joining in 2021, she has driven onsemi's brand and demand transformation, including establishing a Global Demand Center to strengthen pipeline generation and marketing effectiveness. A seasoned B2B marketing executive, Felicity brings deep experience across technology industries and is currently leading AI adoption across onsemi's marketing operations.

Ee Kin Chin

Senior Machine Learning Engineer, Meta

Ee Kin Chin is a Senior Machine Learning Engineer at Meta, where he works on improving Facebook's news feed recommendation system. He previously spent nearly seven years at DataRobot, leading engineering teams and developing features across AutoML, explainability, multimodal clustering, anomaly detection, and generative AI. He is also a Kaggle Competitions Master and author of The Deep Learning Architect's Handbook.

Naveen Dodda

Applied Machine Learning Engineer, Google

Naveen Dodda is an Applied ML Engineer at Google, focused on improving edge acceleration. He develops ISA-specific kernels for Coral NPU (a RISC-V-based neural processor) and SW simulations. Previously, he worked on Coral Edge TPU, deploying production CV and NLP solutions for the medical and automotive sectors through optimized pipelines that streamline model development and real-time deployment.

Learn More
About Our Experts

Alessandro Maggioni

Sr. Manager, Regional Marketing EMEA, AMG, onsemi

Alessandro Maggioni is the Senior Regional Marketing Manager, responsible for EMEA, for the Analog and Mixed Signal Group (AMG) at onsemi. In the semiconductor industry for over 23 years, he has worked in different domains, including automotive and Industrial. He joined onsemi in 2017. His area of expertise is in Technical, mainly Analog and Power, and Strategic Marketing.

Pranav Maheshwari

Staff Machine Learning Engineer, Tesla

Pranav Maheshwari is a Staff Machine Learning Engineer at Tesla, where he focuses on AI and autonomous systems. His background spans motion planning, probabilistic robotics, and autonomous vehicle technologies, with previous engineering roles at Luminar Technologies and robotics research experience at Bristol Robotics Lab. He holds an M.Sc. in Robotic Systems Development from Carnegie Mellon University and a B.Tech. in Engineering Physics from Delhi Technological University.

Ganesh Narayanaswamy

Senior Business Marketing Manager, onsemi

Ganesh Narayanaswamy leads the image sensor product line and business management in onsemi's Industrial & Commercial Sensing Division, driving vision solutions for commercial and industrial markets worldwide. Previously, he led Industrial Vision and Electric Drives businesses at Xilinx, and networking ASIC product lines at Avago and STMicroelectronics. Ganesh holds an MBA from Santa Clara University and an MS in Electrical Engineering from Mississippi State University.

Aarati Noronha

Machine Learning Engineer, Amazon AGI

Aarati Noronha is a Machine Learning Engineer at Amazon AGI, working on reinforcement learning and agentic capabilities for the Nova family of foundation models. Previously at Zordi, she worked on fine-tuning VLA models and edge deployment. She earned her master's from Carnegie Mellon University, where her research at the Bot Intelligence Group focused on promoting safer human-AI interactions.

Learn More About Our Experts

Prasad Paruchuri

Technical Marketing Team Leader, Power Solutions Group, onsemi

Prasad Paruchuri is the Technical Marketing Team Lead of the Industrial PSG Business Unit at onsemi, where he is driving innovative EliteSiC family products for AI Cloud Power conversion and Industrial applications. Prasad brings over 28 years of semiconductor experience in areas that span Application Development, System Solutions Marketing, and Strategic Marketing of New Technology @ STMicroelectronics (15 Years) and onsemi (13 Years). Prior to that, Prasad worked as a Power Supply Design Engineer at ASTEC Custom Power (Advanced Energy).

Shreyash Rawat

Machine Learning Engineer, Meta

Shreyash Rawat received his M.S. in Information Systems with Specialization in ML from Carnegie Mellon University (2022) and his B.Tech. in Computer Science from Vellore Institute of Technology (2021). His research interests include natural language processing, representation learning, and labor market analytics. He has held machine learning engineering roles at Meta, Insight Global, and early-stage startups, where his work has spanned large-scale occupational taxonomy alignment, recommendation systems, and sparse neural network architectures.

Vipul Tiwari

Senior Solutions Architect for Agentic AI and FinOps Center of Excellence, Amazon

Vipul Tiwari is a Senior Solutions Architect for Agentic AI and FinOps Center of Excellence at Amazon. He brings deep experience in intelligent automation, RPA, low-code development, and enterprise-scale solution design across financial services, telecom, IT, and services. A longtime Blue Prism MVP, community leader, speaker, mentor, and hackathon judge, he has helped advance automation practices through both technical delivery and community engagement.