Texas Instruments

ISSUE 2, 2026

Analog Design JOURNAL



Navigating the power delivery demands of next-generation optical modules

Beyond the limit: How smart audio amplifiers dynamically boost microspeaker sound

Using matched N-JFETs for lead resistance cancellation in three-wire RTD measurements

Achieving high power density in a 48V buck converter with integrated GaN FETs and dual-side cooling

Redundant power supplies in automotive applications – priority switching done right

Table of contents



02

Navigating the power delivery demands of next-generation optical modules

Driving next-generation DSPs at sub-1V core voltages requires unprecedented levels of power density. Explore how MagPack™ technology addresses the critical trade-offs between footprint, current density and EMI in high-speed optical modules.

07

Beyond the limit: How smart audio amplifiers dynamically boost microspeaker sound

Mechanical constraints in microspeakers limit loudness and bass extension. Discover how smart audio amplifiers use real-time IV sensing and DSP protection to deliver 8dB-to-10dB higher loudness without compromising hardware reliability.

16

Using matched N-JFETs for lead resistance cancellation in three-wire RTD measurements

Uncompensated lead resistance causes significant RTD measurement errors. Discover how matched N-channel JFETs enable ratiometric measurements to cancel lead resistance, providing a high-precision, low-cost sensing solution.

23

Achieving high power density in a 48V buck converter with integrated GaN FETs and dual-side cooling

Manage the high thermal density of integrated GaN FETs by leveraging dual-side cooling and PCB thermal spreading. This article examines a 48V-to-12V four-phase synchronous buck converter designed for maximum power density and reliability.

30

Redundant power supplies in automotive applications – priority switching done right

Parasitic elements can cause instability during priority switching in redundant automotive power architectures. Learn how to implement external hysteresis and modified feedback dividers to eliminate switchover oscillations and ensure reliable, always-on power delivery to ECUs.

Navigating the power delivery demands of next-generation optical modules

Thomas Hoffman
Marketing engineer, switching regulators

Introduction

Higher data rates do not come with more board space, particularly in high-speed pluggable optical modules. Networks are moving from 400Gbps and 800Gbps toward 1.6Tbps and beyond, while standardized form factors such as quad small form-factor double density and octal small form-factor pluggable remain tightly constrained. As processing power and current requirements grow, power management components must handle heat within the same fixed dimensions.

At the heart of optical modules, digital signal processors (DSPs) convert electrical signals to optical and back again at multi-terabit speeds. Modern DSPs operate at core voltages below 1V while drawing substantial current, and those demands are increasing with each new generation of optical module, making it harder to deliver higher current in less space without adding losses, heat or electromagnetic noise.

One approach – integrating the magnetics directly in the power package – meets multiple design challenges at once. Power modules with integrated magnetics packaging technology use a metallic mold compound as part of the package structure. As a result, the power stage is no longer an afterthought: it has become part of the overall system architecture.

System architecture and power requirements

An optical module brings together a wide range of functions, including DSPs, clock and data recovery circuitry, gearboxes, laser drivers, photodiodes, transimpedance amplifiers, power management and control electronics. Among these, the DSP core supply is often the most demanding rail, with core voltages ranging from 0.4V to 1.0V and currents reaching tens of amperes. These loads are addressed by both high-current (>10A, red circle) and lower-current (2A/4A/6A, blue circle) power modules (**Figure 1**).

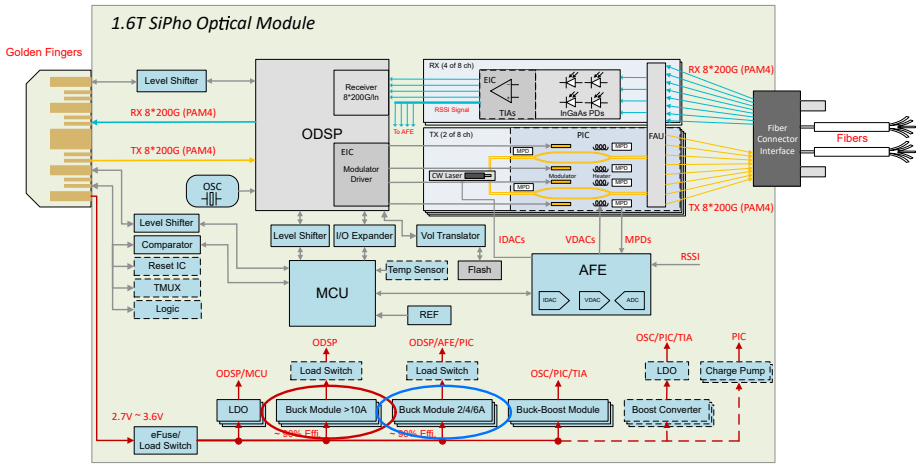


Figure 1. *System block diagram of a 1.6Tbps optical module*

At these levels, a fraction of a watt may be negligible in a larger system, but not inside a compact optical module, which typically has restricted airflow and several heat-generating components packed closely together.

One way to reduce power consumption is to adjust the DSP supply voltage to match the operating conditions. Some systems use a fixed voltage selected during characterization, while others support dynamic voltage scaling (DVS). A DSP may operate at a higher voltage when the signal quality is poor or the temperature is high, and at a lower voltage when conditions allow. The impact can be significant: reducing a 15A core rail from 0.9V to 0.5V lowers the load power by 6W, a substantial reduction when the total power budget of a module is only a few tens of watts.

DVS also means that the regulator may need more than simple voltage conversion. An I^2C interface, for example, allows the system controller to adjust the output voltage or operating mode during operation. Control flexibility alone, however, does not address the design challenge of fitting the power stage within the available space.

Why the inductor needs attention

In a conventional step-down DC/DC converter, the regulator integrated circuit (IC) and power inductor are separate components. The inductor often represents one of the largest elements in the power stage. It takes up board area and adds height, and thus requires careful placement. Furthermore, the connection between the switching node and the inductor also carries fast, high-current waveforms. Managing the resulting losses and electromagnetic interference (EMI) emissions is straightforward on a larger printed circuit board (PCB), but presents a significant challenge inside an optical module.

Inside an optical module, every few square millimeters matter. Multiple supply rails may already be competing for space with high-speed signal routing and sensitive analog circuitry. While integrating the inductor into a power module reduces the external component count, the physical implementation is paramount. Conventional integrated modules may place the inductor beside or above the silicon, which can force compromises in footprint, package height, or heat flow.

Incorporating the magnetics directly into the package improves this integration. Power modules leveraging TI's MagPack™ technology use a metallic mold material as part of the package structure around the semiconductor die and the switching circuitry. For system designers, keeping the switching node integrated inside the package eliminates the need to select and route an external inductor, simplifying the overall power module design.

The 6A **TPSM82866C** module measures 2.3mm by 3mm by 1mm and delivers up to 6A of continuous output current. This complete solution footprint occupies only 6.9mm^2 of PCB space, achieving nearly 1A per 1mm^2 – a 50% solution size reduction compared to previous generations.

To meet the increasing power and thermal demands of optical modules, TI's next-generation 50A power module offers a significant increase in current capability. While the 30A **TPSM8287B30** achieves a density of 450mA/mm^2 and serves a wide range of needs, the **TPSM8287C50** delivers 50A of current (590mA/mm^2) within a similar footprint (**Figure 2**). This enables designers to increase total current delivery by 30% without significantly expanding the power stage's footprint, meeting the current density requirements for 1.6Tbps and beyond.

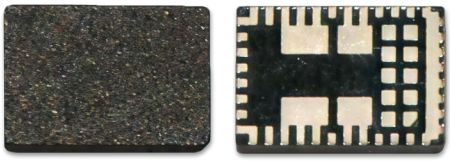


Figure 2. *TI's TPSM8287C50 power module delivers up to 50A of continuous output current in a 4.5mm-by-6.7mm-by-2.85mm footprint*

Small size is useful only if you can remove the heat

Putting more current into a smaller area raises heat density. A compact converter only helps if efficiency and thermal performance keep pace with the reduction in size. Optimizing the magnetic element in an integrated package together with the silicon helps reduce DC and AC losses, and the package itself can help spread heat away from the die.

Both advantages are important in an optical module. DSPs, lasers, and analog circuitry already contribute to the thermal load, and the fixed form factor limits cooling options. If the power stage runs too hot, your options are rather limited: reduce performance, improve the mechanical heat path, or accept less operating margin at elevated temperatures.

A current rating alone does not reveal whether a converter will work well in an optical module. Designers also need to look at efficiency at the actual operating point, temperature rise, and safe operating area (SOA). The question is not simply whether a converter can supply the current, but whether it can do so inside the real thermal environment of the module. **Figure 3** shows the SOA curve for the TI 30A TPSM8287B30 power module with a 5.0V input voltage.

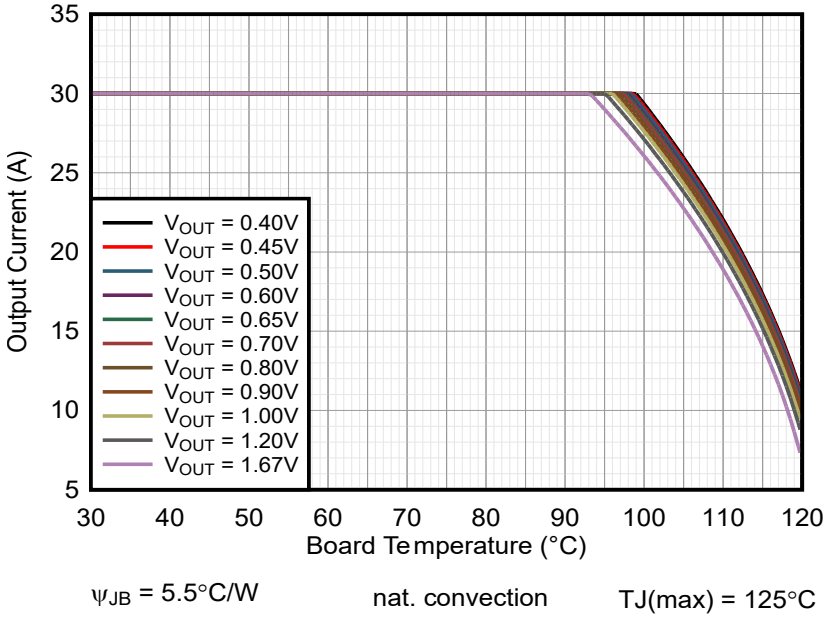


Figure 3. *SOA curve of the TPSM8287B30 power module at V_{IN} = 5.0V*

Considering EMI at the same time

The same density that creates thermal problems also introduces potential EMI concerns. Optical modules place high-current switching power supplies close to sensitive transmit and receive circuitry. In a discrete buck converter, the path between the switching node and the external inductor is a primary source of noise; at high currents and switching frequencies, even a small loop area can generate significant unwanted emissions. While adding shielding, filtering, or moving the power supply farther away from sensitive circuitry can mitigate this, each approach consumes additional board space or component count.

Integrated magnetics packaging addresses these EMI challenges by enclosing the die, inductor and switching node within a single shielded package, drastically reducing the internal switching loop area. **Figure 4** demonstrates that a 6A power module with MagPack technology maintains radiated emission levels well below both CISPR (International Special Committee on Radio Interference) 32 Class A and Class B (3m) limits across the 30MHz to 1GHz range.

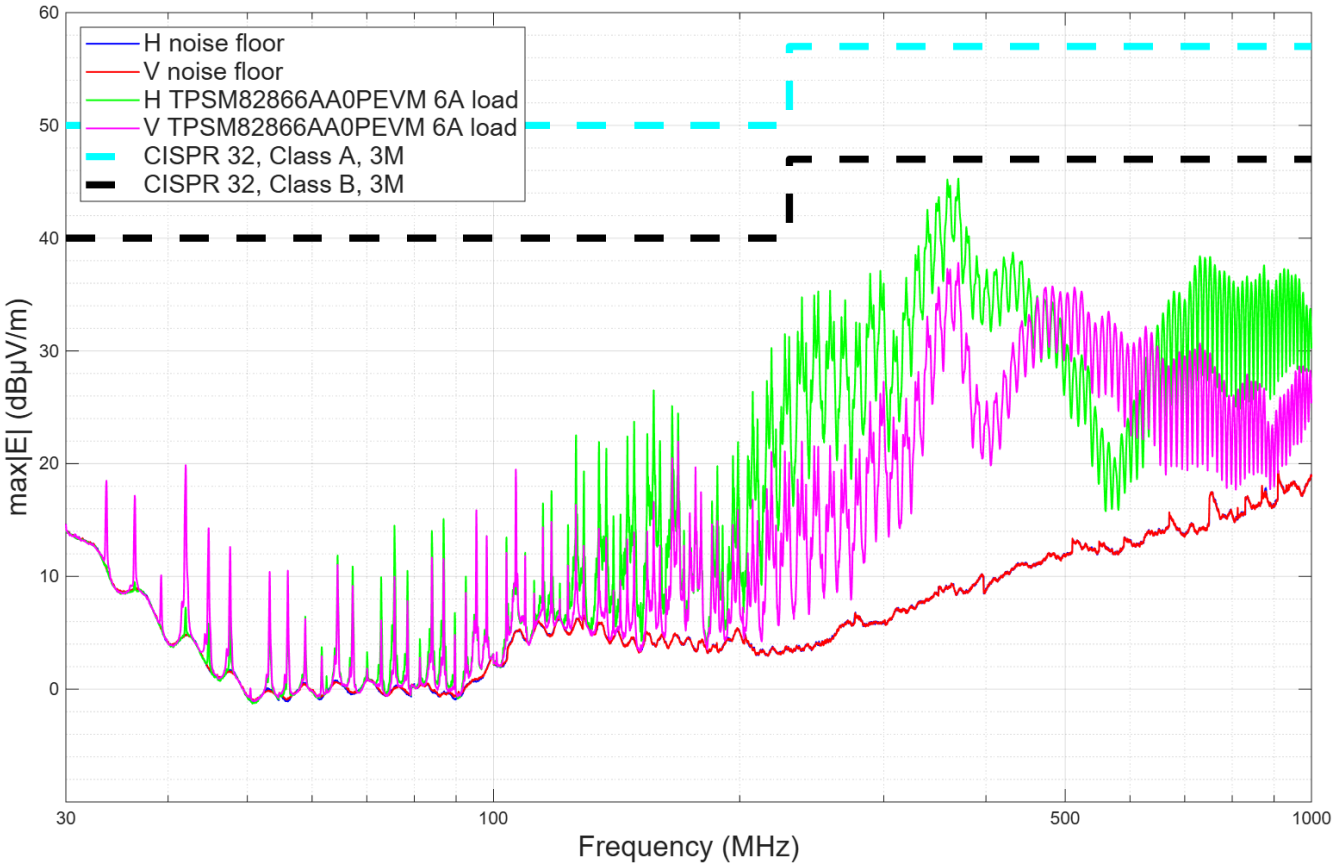


Figure 4. *Radiated EMI performance of a 6A power module with MagPack technology from 30MHz to 1GHz, demonstrating compliance with CISPR32 Class A and Class B limits*

Lower emissions are useful, but the placement flexibility that those emissions create may be just as important. If the regulator can sit closer to the DSP, the high-current path between supply and load becomes shorter, reducing the resistive voltage drop and parasitic impedance. Both factors become increasingly critical for maintaining regulation accuracy as core voltages fall below 1V. EMI affects placement, placement affects power integrity, and both influence how much board area is left for the rest of the module – an example of why design constraints cannot be viewed in isolation.

Conclusion

As optical modules move to higher data rates, power delivery becomes more tightly coupled to the mechanical, thermal and electrical design of the module. Integrated magnetic packaging, such as MagPack technology, helps address these challenges by reducing the footprint of the power stage, shortening switching paths, and improving both thermal and EMI performance.

However, integrated magnetic packaging does not replace the need for thorough design and validation. Current capability, switching frequency, compensation, control strategy, passive components and PCB placement still require careful optimization. Designers must validate the thermal behavior and EMI performance within the final enclosure, rather than relying solely on evaluation board results.

As data rates increase in next-generation optical modules, the primary limiting factor may not be the DSP or the optical components, but rather how much power can be delivered to them, quietly and efficiently, within the available space and thermal budget.

About the author

Thomas Hoffman is a marketing engineer focused on data center and data communication modules. He has more than 30 years of experience in the electronic industry, spanning hardware design, product definition, and product promotion. Thomas holds a master's degree in electrical engineering from the Bergische Universität Gesamthochschule Wuppertal (BUGH) in Wuppertal, Germany.

Additional resources

- Read the application brief, "**Achieving High Thermal Performance in Compact Buck Power Modules**" and application note, "**Understanding SOA Curves to Operate at High Output Currents and Temperature**."
- These technical articles cover MagPack technology: "**MagPack technology: Four Benefits of New Power Modules That Can Help You Pack More Power in Less Space**" and "**Power Module Package Types and Their Benefits**."
- Watch the videos, "**TI's Power Modules Utilize MagPack Packaging Technology for Improved EMI Performance**" and "**Double the Density: Reshaping Power Modules with MagPack Technology**."

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.

MagPack™ is a trademark of Texas Instruments.
All trademarks are the property of their respective owners.

© 2026 Texas Instruments Incorporated

Beyond the limit: How smart audio amplifiers dynamically boost microspeaker sound

Supriyo Palit
Systems engineer, consumer audio amplifiers and haptics

Aditya Sundar
Systems engineer, consumer audio amplifiers and haptics

Introduction

Delivering loud, immersive sound from thin, lightweight devices is challenging because microspeakers face strict loudness ceilings. To prevent blown voice coils or torn diaphragms caused by high voltage, designers traditionally cap peak volume well below a speaker's rated power, restricting achievable sound pressure level (SPL). This limitation also affects larger speakers in Bluetooth devices and soundbars. Smart audio amplifiers address this by fusing a high-precision IV-sense analog-to-digital-converter (ADC) front end with real-time DSP and protection algorithms. By intelligently overdriving speakers for short durations within a safe zone, a smart amplifier allows hardware to reach its true physical limits, delivering 8dB to 10dB higher perceived loudness and deeper bass.

The physics of microspeaker vulnerability

A standard microspeaker consists of a voice coil suspended in a magnetic field, attached to a diaphragm (or cone) and held in place by a suspension system. When an electrical signal passes through the coil, it creates an electromagnetic force that moves the cone, vibrating the air to create sound. See **Figure 1**.

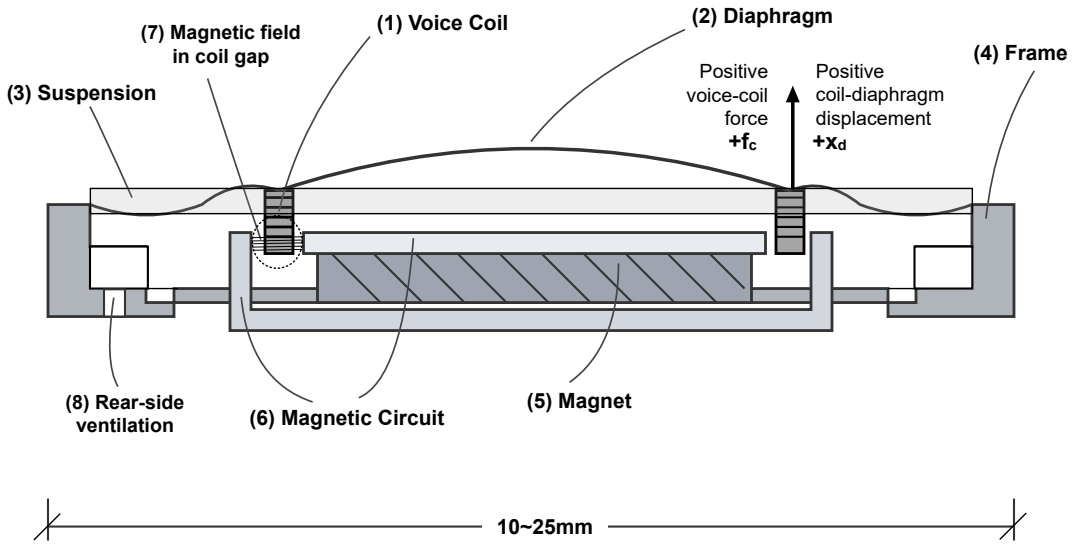


Figure 1. *Microspeaker construction*

When engineers increase audio output in these tiny components to produce louder sound, they encounter two catastrophic failure mechanisms: thermal overload and mechanical over-excursion.

From a thermal perspective, loudspeakers are highly inefficient, where 95% to 99% of input electrical energy becomes heat rather than sound. Thus, continuous high-voltage signals can push the voice coil past its rated temperature (T_{max}), melting the wire insulation and causing electrical shorts. Mechanically, producing bass at high volumes requires large cone excursion; if the drive voltage pushes the cone past its physical travel limit (X_{max}), the suspension can stretch beyond its elastic limit, letting the voice coil strike the magnet's backplate (bottoming out) or tearing the diaphragm from its surround and causing distortion, buzzing or total acoustic failure.

Traditional versus smart approaches to audio design

In a traditional audio design, engineers cap the peak audio power at the speaker's continuous rated power to ensure safety. This static voltage ceiling forces the entire signal down to a defensive floor, resulting in low volume and weak bass.

A smart amplifier constantly monitors the state of the speaker and adapts its output in real time. By understanding exactly how hot the voice coil is, and exactly how far the diaphragm is moving at any given millisecond, a smart amplifier can safely drive speakers with high-voltage peaks that would instantly destroy it under a traditional system.

A smart amplifier algorithm dynamically modulates loudness, allowing for 8dB to 10dB higher peak loudness SPL by ensuring that the speaker's estimated excursion stays within X_{max} and that its estimated voice coil temperature stays within T_{max}. The algorithm only pulls back the energy when the speaker actively approaches a literal breaking point. See **Figure 2**.

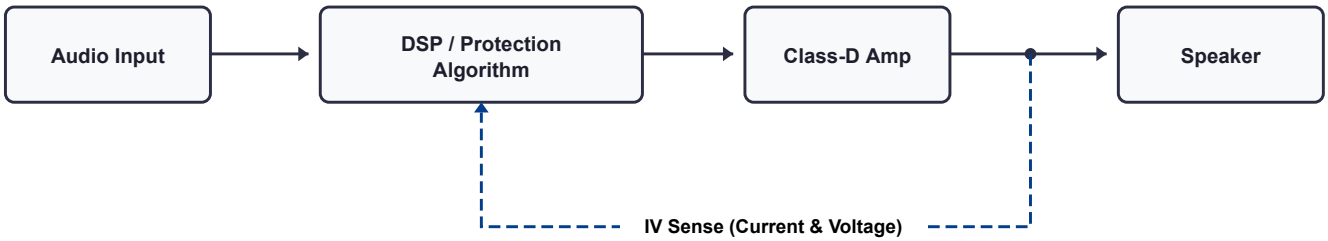


Figure 2. *Smart amplifier system architecture*

Specialized sensing hardware, together with an acoustic software model running on the DSP, form a real-time feedback loop that enables the amplifier to adapt its output on the fly.

Integrated IV-sense circuitry

Smart amplifiers utilize current and voltage (IV)-sense circuitry. After integrated analog-to-digital converters (ADCs) capture the exact voltage applied to the speaker and the exact current drawn by it, the system computes the speaker's electrical impedance ($Z = V / I$) in real time. Dynamic impedance changes – driven by temperature and mechanical movement – provide the raw sensory input for protection algorithms.

The ADC's conversion range is set to ensure that it is possible to sense the full-scale audio signal played through the amplifier. High-precision hardware ensures maximum utilization of the speaker's loudness. Any current or voltage sensing

error is a direct hit to loudness, since the speaker protection algorithm would need that much additional error guard banding to never breach the X_{max} and T_{max} limits.

A typical ±0.35% error results in ±1°C error in the voice coil temperature estimation through IV sensing. TI's **TAS2781**, **TAS2783**, and **TAS2573** ADCs achieve high accuracy in IV sensing, resulting in minimal guard-band requirements for temperature and excursion. A high signal-to-noise ratio (SNR) and dynamic range (DNR) avoid sensing errors and reproduce the audio content through the speaker with high fidelity. See **Figure 3**.

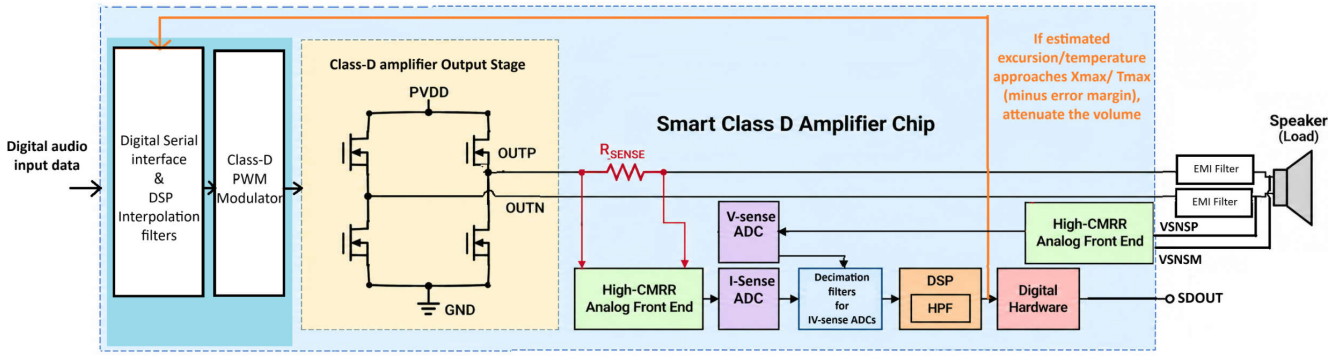


Figure 3. *The co-working of hardware and software in a smart amplifier*

The VSNSP and VSNSN pins of the device sense the voltage across the speaker. These are sense lines that must be connected as close to the speaker terminal as possible to ensure accurate measurement.

Since the voltage across the speaker in a Class-D amplifier is a high-voltage switching node, the IV-sense ADCs require a high common-mode rejection ratio analog front-end circuit to extract the audio content from the switching voltage and to convert the high voltage to a lower voltage for precision sensing. Sigma-delta or successive approximation register ADC circuitry senses this low-voltage audio content to extract precise IV information. The ADC output is typically at a much higher frequency, such as 3MHz, since the data ADC is oversampled to achieve high SNR and DNR. The ADC output data is decimated down to the audio rate (48KSPS) and high pass-filtered to remove the DC content (if any), before being used by the SmartAmp algorithm running in the integrated DSP of the TAS2781, TAS2783, or TAS2573.

Phase-matching the I and V audio channels to high precision – within <±1µs – enables highly precise real-time estimation of the speaker impedance. The IV-sense data can also be monitored by the user and brought out on TDM SDOUT pins or through the Soundwire data interface.

The DSP is the brain of a smart amplifier. It hosts the speaker protection algorithms, running complex mathematical models of the loudspeaker alongside real-time predictive filters. The DSP acts as a gatekeeper, modifying the incoming digital audio stream on the fly based on the telemetry provided by the IV-sense hardware.

For a smart amplifier to work, it must be calibrated to the specific loudspeaker it is driving. Audio engineers use specialized diagnostic tools to characterize the target microspeaker in an acoustic chamber. They measure its thermal resistance (R_E), resonance frequency (f_0), mechanical damping (Q_{ms}) and magnetic strength (Bl). Programming these parameters into the DSP serves as a model for the speaker to create a customized safety envelope. TI's PurePath™ Console 3 audio toolkit provides calibration, modelling, and end-system tuning support.

The temperature protection algorithm

Thermal protection in a smart amplifier relies on the physical law that the electrical resistance of a metal conductor changes predictably with temperature. This is known as the temperature coefficient of resistance (α).

As a voice coil heats up, its R_E increases linearly, and the smart amplifier's IV-sense hardware isolates this resistance from the measured impedance. **Equation 1** solves for the coil's internal temperature:

$$R_{dc}(T) = R_E(T_0)[1 + \alpha(T - T_0)] \quad (1)$$

where R_E (T) is the measured resistance at current time T, R_E (T_0) is the baseline resistance from factory calibration, and α is the coil material's temperature coefficient of resistance (such as copper).

This eliminates the need for bulky, slow-responding physical thermistors.

This measurement is invoked only periodically – a checkpoint reading rather than a continuous stream – for two reasons: it is less trustworthy at low output levels or during low-to-high power transitions (noise-limited signals, settling error and offsets), and running the sense ADC and its analog front end continuously would consume unnecessary power, especially in battery-powered, low-volume listening.

Thermal modeling and control

Because the R_E measurement only arrives periodically, the algorithm needs a way to know the voice coil's temperature continuously, in the gaps between measurements. The solution is a predictive thermal model – a multistage resistor-capacitor thermal network (see **Figure 4**). This model simulates heat transfer from the voice coil to the speaker magnet structure, and finally to the device's chassis and ambient air, and it runs at the full control-loop rate. Each time a new, IV-sense-based R_E reading arrives, re-anchors the model, and corrects any drift accumulated since the last checkpoint.

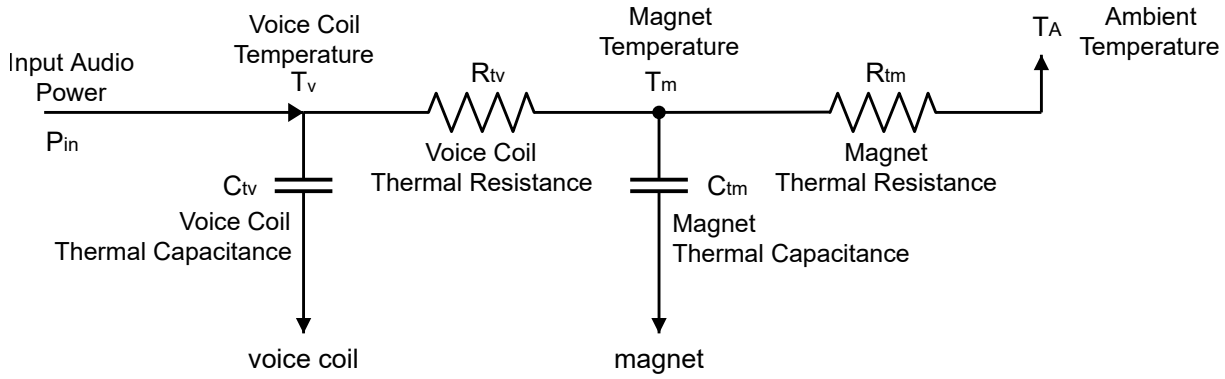


Figure 4. *Loudspeaker predictive thermal model*

By tracking heat dissipation paths in this way, the algorithm can predict how fast the temperature will rise during an upcoming audio peak and how quickly it will cool down during a pause. This estimate feeds a power controller that enforces the safe operating limit of the audio signal continuously and without interruption. As the estimated voice coil temperature approaches the $T_{max}$of the loudspeaker, the controller reduces the permissible limit to protect the loudspeaker from overheating. When the loudspeaker cools down during soft audio segments, the controller will allow higher-power delivery to the loudspeaker.

The thermal protection algorithm controls the audio signal based on the power limit predicted by the thermal controller, applying a smooth, broadband gain reduction that lowers the overall root-mean-square (RMS) power delivered to the speaker so that the voice coil can cool down without abrupt, audible drops in volume. This gain reduction uses slow attack and release times; because human hearing is less sensitive to gradual changes in overall volume, the attenuation blends seamlessly into the background, preserving the perceived dynamics of the audio. See **Figure 5**.

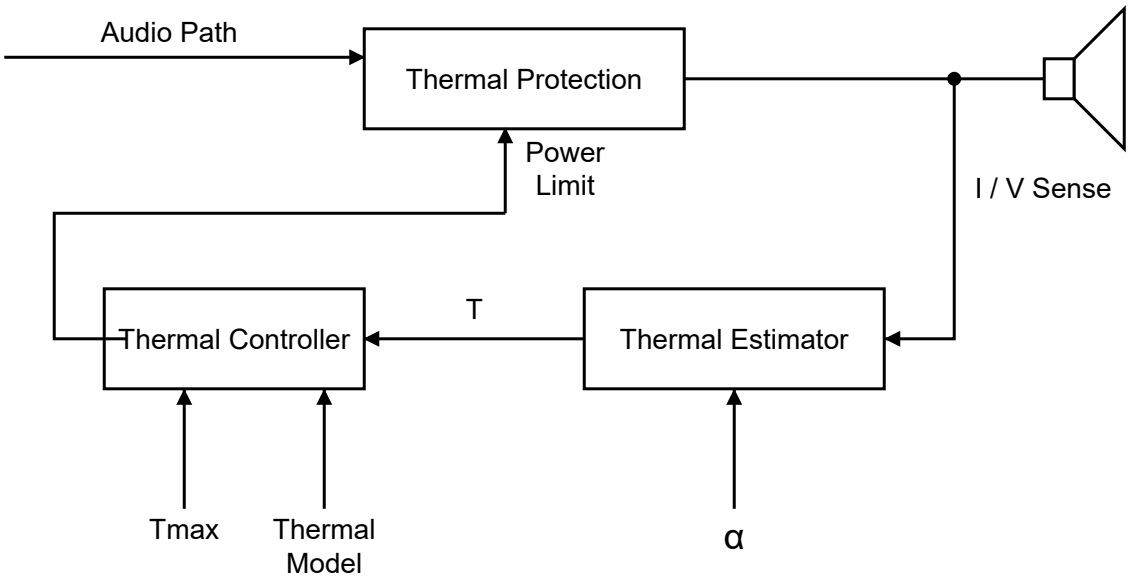


Figure 5. *Closed-loop temperature protection*

In **Figure 6**, the blue curve represents the smart amplifier's estimated temperature, while the red line indicates the T_{max} limit. The amplifier dynamically modulates the audio to ensure that the speaker is always within reliable temperature limits (leaving some headroom for IV-sensing errors).

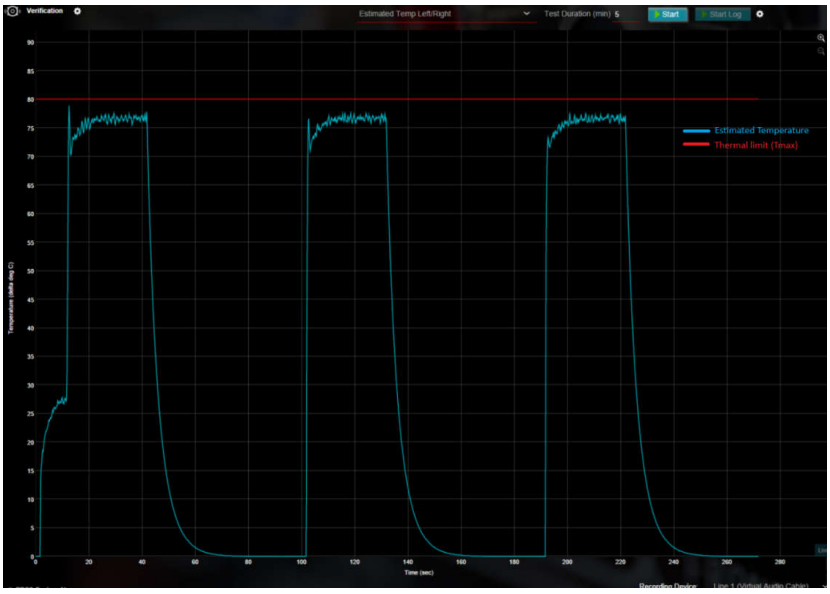


Figure 6. *Smart amplifier temperature estimation vs the T_{max} limit*

The excursion protection algorithm

Unlike thermal protection, which deals with slow-moving thermal masses and RMS power, excursion protection is highly volatile. A speaker cone can travel from rest to maximum over excursion within a single waveform cycle in a fraction of a millisecond.

Because putting a physical laser displacement sensor inside a smartphone to track speaker cone movement is impossible, the smart amplifier relies on a mathematical simulation of the speaker mechanics. This is based on Thiele-Small parameters and linear state-space models of the transducer. See **Figure 7** and **Table 1**.

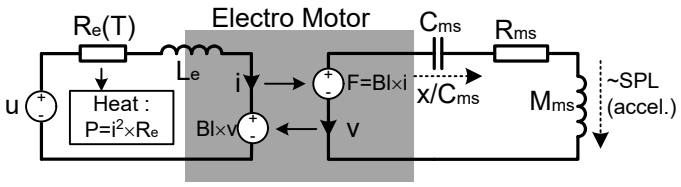


Figure 7. *Electromotive model of a loudspeaker*

Electromotive parameter	Description
R_{ms}	Mechanical damping factor
M_{ms}	Mechanical mass
C_{ms}	Mechanical compliance
Bl	Force factor
R_E	DC resistance of the voice coil
L_E	Inductance of the voice coil
u	Amplifier voltage
i	Voice coil current
v	Velocity of membrane
F	Mechanical force in magnetic field
x	Magnetic excursion

Table 1. *Main parameters in a speaker electromotive model*

The algorithm uses real-time current and voltage data to calculate the back-electromotive force (back-EMF) generated by the speaker's movement. After combining back-EMF data with the “digital twin” speaker model, the DSP runs an observer loop such as a normalized least mean square filter to update the model in real time. The speaker model is used to calculate the instantaneous physical displacement (x) of the cone using an excursion predictor voltage-to-excursion filter (**Equation 2**):

$$Z_{exc}\left(s\right) = \frac{X(s)}{V(s)} = \frac{Bl}{M_{ms}R_E} \times \frac{1}{s^2 + s \times \left(\frac{R_{ms} + \frac{Bl^2}{Re}}{M_{ms}} + \frac{1}{M_{ms}C_{ms}}\right)} \quad (2)$$

Look-ahead limiter

If a low-frequency bass transient (such as a snare drum hit or an explosion in a movie) threatens to push the cone past its X_{max}, the excursion protection algorithm deploys rapid, targeted countermeasures. A brief digital delay buffer (typically a few milliseconds) enables the DSP to see a dangerous peak before it actually reaches the amplifier stage, giving the algorithm time to apply a precise, instantaneous peak attenuation to that specific waveform cycle. See **Figure 8**.

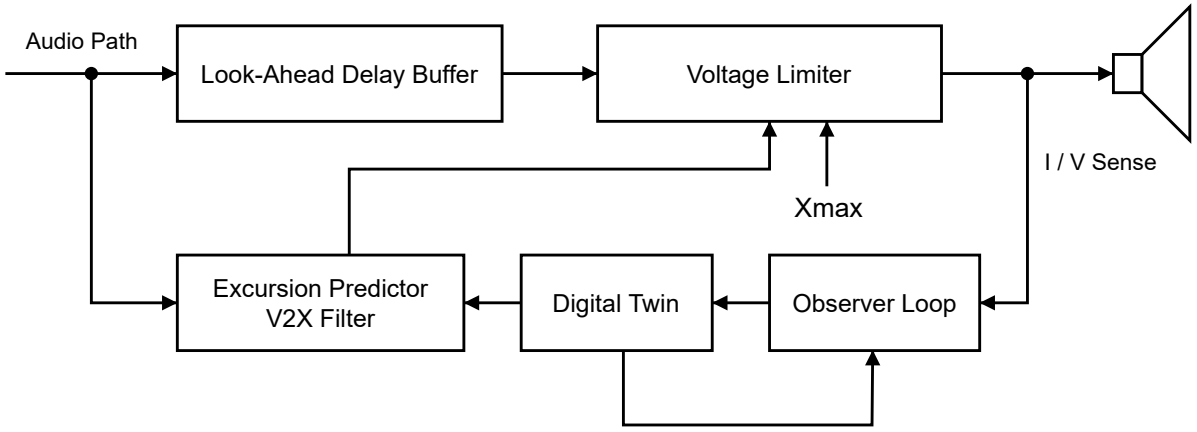


Figure 8. *Closed-loop excursion protection*

In **Figure 9**, the blue curve represents the smart amplifier's estimated excursion, while the red lines indicate the X_{max} limits. The amplifier dynamically modulates the audio to ensure that the speaker is always within reliable excursion limits.

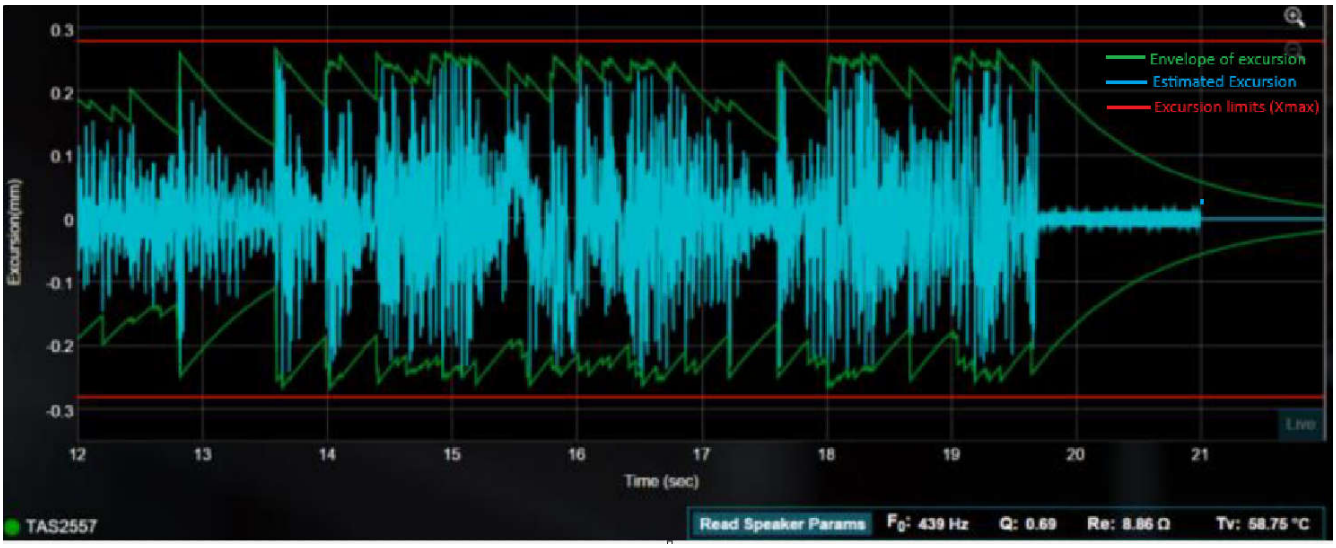


Figure 9. *Real-time excursion tracking relative to the X_{max} limits*

Integrating the audio tuning chain

Unlike conventional audio designs that use the equalizer (EQ) and dynamic range compressors (DRCs) as aggressive limiting tools to ensure safety of the microspeaker, the smart amplifier paradigm enables true artistic tuning of the audio signal chain. For example, tuning engineers can use the parametric EQ to aggressively boost low frequencies (bass extension) and make tiny smartphone speakers sound like a much larger acoustic enclosure. The EQ can also flatten response peaks to optimize vocal clarity.

A smart amplifier also uses multiband DRCs aligned with its protection models, splitting the audio into distinct low, mid, and high bands. Designers can optimize the mid and high bands for maximum perceived loudness, crisp dialogue and wide spatial imaging. The low band is tightly coupled with the excursion model in order to provide the optimum bass response based on the actual physical displacement of the cone.

Smart amplifiers enhance loudness and frequency response by safely maximizing speaker output, providing advantages in hardware performance, design, and reliability. Specific benefits include:

- Extended low-frequency response: Instead of rolling off bass permanently with a static hardware filter, smart amplifiers use an intelligent, sliding bass extension strategy. At low volumes, the bass extends deep into lower frequencies. As the volume increases, the look-ahead limiter activates only as much as needed to prevent over-excursion, retaining maximum possible bass presence across all volume levels.
- Component optimization and space savings: Because software handles protection, product designers can source smaller, thinner speaker modules, freeing up space inside tight device enclosures for larger batteries, advanced camera sensors, or slimmer product profiles.
- Near-zero field return rates: By acting as a reliable digital shield against scorched voice coils and torn diaphragms, smart amplifiers drastically reduce hardware failures and save millions in warranty support costs and field return claims.

Figure 10 shows a comparison of loudness and frequency response, comparing smart amplifiers versus traditional amplifiers.

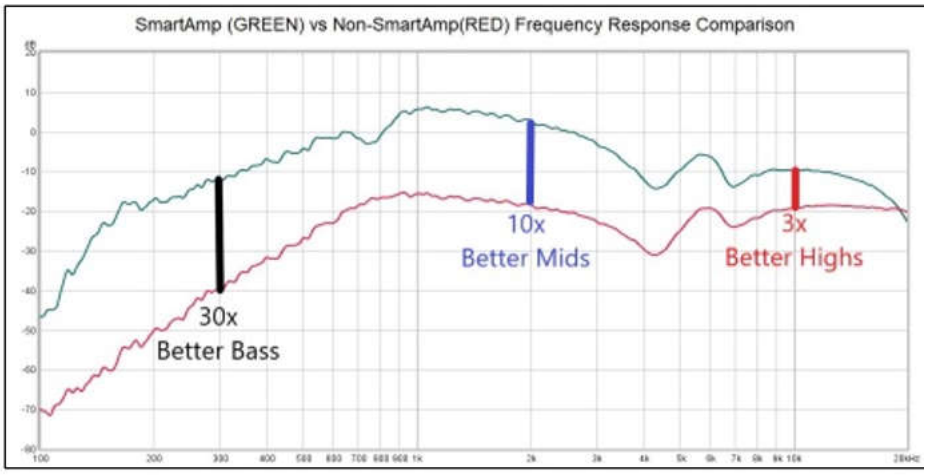


Figure 10. *Comparison of loudness and frequency response of smart amplifiers vs traditional non-smart amplifiers*

Conclusion

By replacing passive safety margins with active, hardware-software co-driven surveillance, smart amplifiers bridge the gap between fragile microacoustics and consumer demand for powerful, high-fidelity sound. Through real-time current and voltage sensing, nonlinear thermal and mechanical modeling, and audio engineering, these systems successfully push microspeakers to their absolute physical boundaries. The result is a generation of smartphones, laptops, and smart home devices that sound remarkably loud, rich, and full, all while maintaining component reliability.

About the authors

Aditya Sundar, a systems engineer for consumer audio amplifiers and haptics products, has worked at Texas Instruments for more than 11 years. He previously served as an analog design manager. Sundar holds a bachelor's degree in electronics and instrumentation engineering from the Birla Institute of Technology & Science, Goa. He specializes in mixed-signal Class-D audio signal chain design and works extensively on mixed-signal boosted Class-D audio amplifiers, piezo drivers and loudspeaker protection technology.

Supriyo Palit is a systems engineer for consumer audio amplifiers and haptics products at Texas Instruments, where he has spent more than 27 of his 30 years in the industry. He holds a postgraduate degree in electrical communications from the Indian Institute of Science in Bangalore. His work focuses on audio signal processing, loudspeaker protection technology, audio power management, haptics and voice algorithms.

Acknowledgments

The authors would like to acknowledge Vallabha Hampiholi, Nilendra Gadre, Michael Kim, Jasjot Chadha, and Mashkurahmed Badebade for their technical contributions during the development of this article.

Additional resources

- Check out this **smart amplifiers presentation from CES 2015**.
- Read the application reports, "**Smart Amp Tuning Guide**," "**Smart Amp Quick Start Guide**" and "**Getting Started with Smart Amp Development**."

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.

PurePath™ is a trademark of Texas Instruments.
All trademarks are the property of their respective owners.

© 2026 Texas Instruments Incorporated

Using matched N-JFETs for lead resistance cancellation in three-wire RTD measurements

John Caldwell
Technologist, linear amplifiers

Three-wire sensing is a common approach to sensor interfaces where the sensor impedance is of similar magnitude to the leads connecting to the sensor. Resistance temperature detectors (RTDs) are a common example of such a sensor. A Pt100 RTD has a resistance of 100Ω at 0°C. Even 1Ω of lead resistance can contribute >2.5°C of error to the measured temperature. In a three-wire measurement, matched current sources are used to ensure that the voltage drop across the sensor connecting leads are equal and therefore can be removed from the measurement.

Figure 1 illustrates a three-wire, ratiometric, RTD measurement circuit that employs a monolithic, dual, N-channel JFET, the **JFE2325**, and two resistors as a simple, low-cost, approach to creating two matched current sources. This solution is also able to operate at relatively low supply voltages. The voltage drop across the JFET current source circuit only needs to be above the pinchoff voltage of the JFET, 1.08V in the case of the JFE2325, to operate properly.

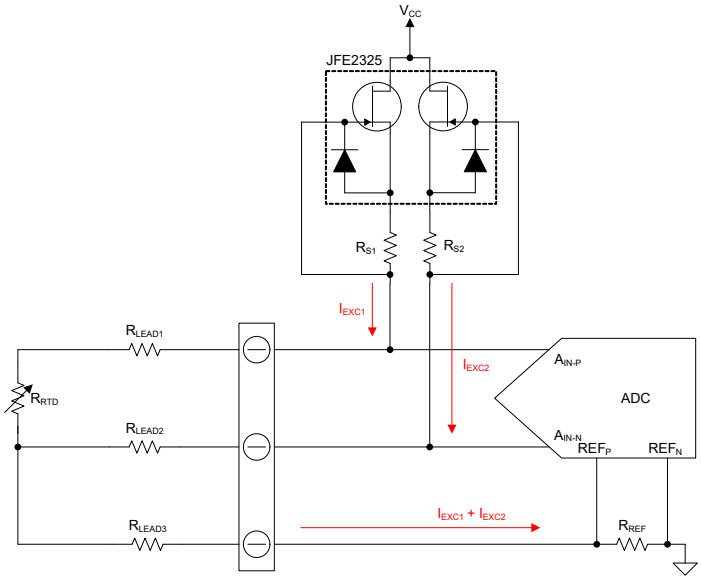


Figure 1. *Three-wire, ratiometric RTD measurement using the JFE2325 dual JFET as matched current sources*

To understand the benefit of the three-wire approach, we can derive the equations for the voltage at each input to the ADC:

$$V_{AIN-P} = I_{EXC1}(R_{LEAD1} + R_{RTD}) + (I_{EXC1} + I_{EXC2})(R_{LEAD3} + R_{REF}) \quad (1)$$

$$V_{AIN-N} = I_{EXC2}(R_{LEAD2}) + (I_{EXC1} + I_{EXC2})(R_{LEAD3} + R_{REF}) \quad (2)$$

If $R_{LEAD1} = R_{LEAD2} = R_{LEAD}$ and $I_{EXC1} = I_{EXC2} = I_{EXC}$, then the differential input voltage sensed by the ADC can be simplified to **Equation 3**:

$$V_{ADC} = V_{AIN-P} - V_{AIN-N} = I_{EXC}(R_{LEAD} + R_{RTD}) + 2I_{EXC}(R_{LEAD3} + R_{REF}) - I_{EXC}R_{LEAD} - 2I_{EXC}(R_{LEAD3} + R_{REF}) = I_{EXC}R_{RTD} \quad (3)$$

The voltage drop across the third lead and the reference resistor are common to both ADC inputs and are therefore removed.

Lastly, because the measurement is ratiometric, the absolute value of the excitation current and its stability over temperature is not as important as the matching between the two current sources. The ADC output only depends on the ratio of the RTD resistance to the reference resistance as shown in **Equation 4**:

$$\text{Output Code} = 2^{n-1} \times \frac{V_{RTD}}{V_{REF}} = 2^{n-1} \times \frac{I_{EXC}R_{RTD}}{2 \times I_{EXC}R_{REF}} = 2^{n-2} \times \frac{R_{RTD}}{R_{REF}} \quad (4)$$

Current source calculations

A very simple JFET current source is formed by placing a resistor in series with the source of an N-JFET and connecting the gate below the resistor as shown in **Figure 2**. As current flows through the source resistor R_S, the resistor raises the voltage at the JFET source with respect to its gate, further pinching off the JFET channel and regulating the drain-to-source current, I_{DS}, through negative feedback.

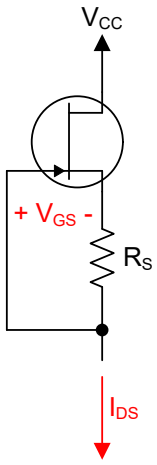


Figure 2. *A simple JFET current source formed by connecting the gate below a source resistor*

The simplest method to determine the output current of the JFET current source is to use the I_{DS} versus gate-to-source voltage (V_{GS}) curve from the JFET datasheet. This curve for the JFE2325 is reproduced in **Figure 3**.

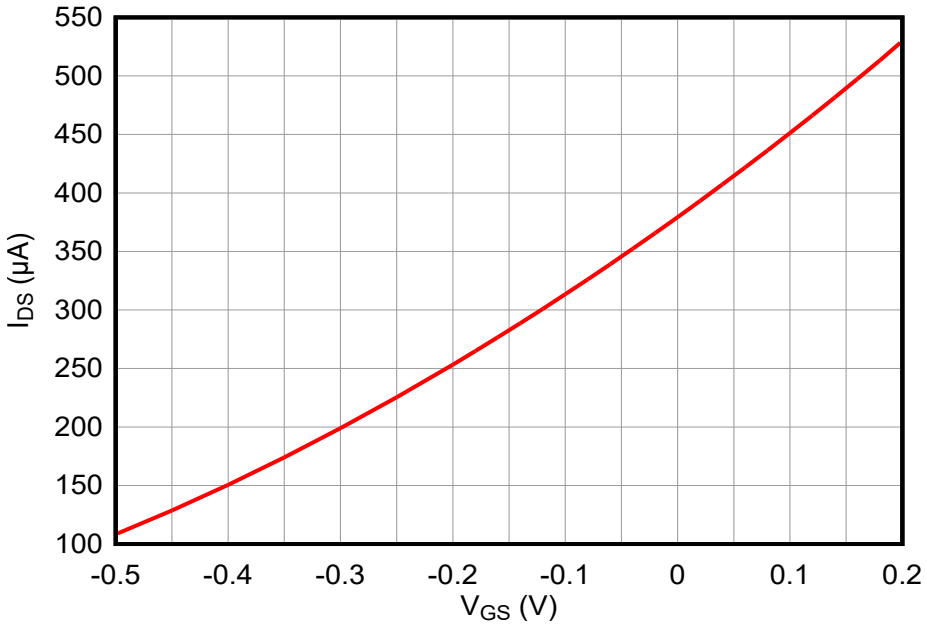


Figure 3. *Drain-to-source current (I_{DS}) vs gate-to-source voltage (V_{GS}) for the JFE2325*

Considering that the voltage drop across R_S is equivalent to the gate-to-source voltage but with opposite polarity, the value for R_S can be calculated for the necessary V_{GS} to produce the desired I_{DS}, shown in **Equation 5**:

$$V_{GS} = -I_{DS}R_{DS} \rightarrow R_S = \frac{-V_{GS}}{I_{DS}} \tag{5}$$

From **Figure 3**, we can see that for a drain current of 200µA, the JFE2325 requires a V_{GS} of -0.3V. Inserting this value into **Equation 5**, we solve for R_S in **Equation 6**:

$$R_S = \frac{-V_{GS}}{I_D} = \frac{-(-0.3V)}{200\mu A} = 1.5k\Omega \tag{6}$$

I_{DS} can be directly calculated if the JFET pinchoff voltage, V_P, and drain-to-source saturation current, I_{DSS}, are known, from **Equation 7**:

$$I_{DS} = I_{DSS}\left(1 - \frac{V_{GS}}{V_P}\right)^2 \tag{7}$$

Inserting **Equation 5** for V_{GS} gives **Equation 8**:

$$I_{DS} = I_{DSS}\left(1 - \frac{-I_{DS}R_S}{V_P}\right)^2 = I_{DSS}\left(1 + \frac{I_{DS}R_S}{V_P}\right)^2 \tag{8}$$

Equation 8 can be expanded and put into the standard quadratic form *($ax^2 + bx + c = 0$)* and solved using the quadratic equation (**Equation 9** and **Equation 10**):

$$\left(\frac{I_{DSS}R_S^2}{V_P^2}\right)I_{DS}^2 + \left(\frac{2I_{DSS}R_S}{V_P} - 1\right)I_{DS} + I_{DSS} = 0 \tag{9}$$

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \tag{10}$$

Where a, b, and c are defined in **Equation 11**, **Equation 12**, and **Equation 13**:

$$a = \left(\frac{I_{DSS} R_S^2}{V_P^2} \right) \tag{11}$$

$$b = \left(\frac{2 I_{DSS} R_S}{V_P} - 1 \right) \tag{12}$$

$$c = I_{DSS} \tag{13}$$

This approach will give two roots, but only one of the two will be a valid value for output current from the current source. Using typical values for I_{DSS} and V_P for the JFE2325, 385µA and -1.079V, gives an I_{DS} of 200.4µA with a 1.5kΩ source resistor.

Estimating mismatch between the two current sources

Calculating mismatch between the two current sources requires investigating the impact of mismatch in the JFET parameters on the resulting output current. Discrete JFETs have extremely wide tolerances on their I_{DSS} and V_P values. However, because the two JFETs in the JFE2325 are built on the same die using modern IC layout techniques, the parameters of the two JFETs closely match each other. **Figure 4** is a production histogram of the I_{DSS} mismatch of the two JFETs in the JFE2325, taken from over 74,000 units. It can be seen from the histogram that the I_{DSS} of the two JFETs matches to 2% within three standard deviations of the mean. The maximum I_{DSS} mismatch, which indicates six standard deviations from the mean, is specified to 5% in the datasheet of the JFE2325.

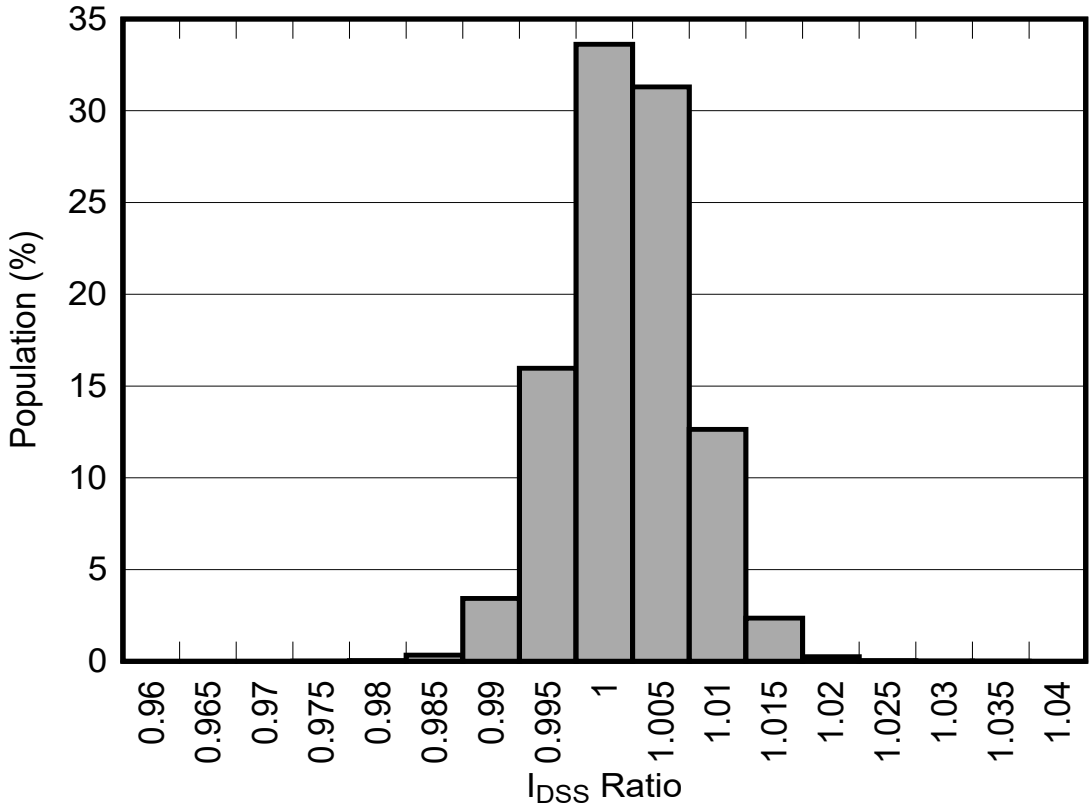


Figure 4. *Production histogram of I_{DSS} mismatch between the two internal JFETs of the JFE2325*

JFET I_{DSS} and V_P are not independent parameters. Rather, the two quantities are correlated such that any increase in I_{DSS} will result in a V_P further from zero. **Figure 5** shows this correlation between I_{DSS} and V_P (called gate-to-source cutoff voltage, V_{GSC}, in the plot) for the JFE2325.

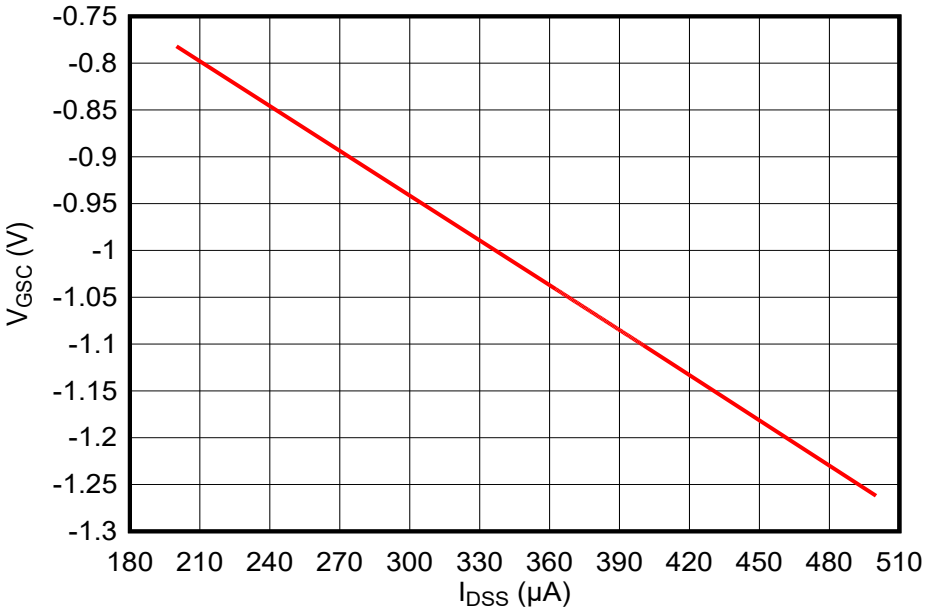


Figure 5. *Pinchoff voltage, or gate-to-source cutoff voltage (V_{GSC}) vs I_{DSS} of the JFE2325*

We can extract a linear equation from this plot to estimate the pinchoff voltage for a JFET given a certain I_{DSS} value, giving **Equation 14**:

$$V_P = -1592.6 \times I_{DSS} - 0.4656V \tag{14}$$

Table 1 shows the values for I_{DSS} and V_P for three possible JFETs in the JFE2325 along with the resulting I_{DS}. JFET1 is a nominal case for an internal JFET. JFET2 has a 2% mismatch in I_{DSS} compared to JFET1, which is equivalent to a unit three standard deviations from the mean. JFET3 has a 5% mismatch in I_{DSS} compared to JFET1, equivalent to a unit six standard deviations from the mean.

	I_{DSS}	V_P	I_{DS} (R_S = 1.5kΩ
JFET1 (Nominal)	385.0µA	-1.079V	200.4µA
JFET2 (2% Mismatch)	377.3µA	-1.066V	197.1µA
JFET3 (5% Mismatch)	365.8µA	-1.048V	192.2µA

Table 1. *Possible parameters for mismatched JFETs in the JFE2325*

The resulting I_{DS} for a current source built from each JFET is calculated by solving **Equation 9** with the respective values for I_{DSS} and V_P. The mismatch in current sources is calculated in **Equation 15** and **Equation 16**:

$$\text{Mismatch}(\pm 3\sigma,\ \%) = 100 \times \frac{|200.4\mu A - 197.1\mu A|}{200.4\mu A} = 1.647\% \tag{15}$$

$$\text{Mismatch}(\pm 6\sigma,\ \%) = 100 \times \frac{|200.4\mu A - 192.2\mu A|}{200.4\mu A} = 4.092\% \tag{16}$$

It can be instructional to also view the I_{DS} versus V_{GS} plot with a load line added for the source resistor R_S as shown in **Figure 6**. The I_{DS} value will be where the curve for the JFET and the load line for R_S intersect.

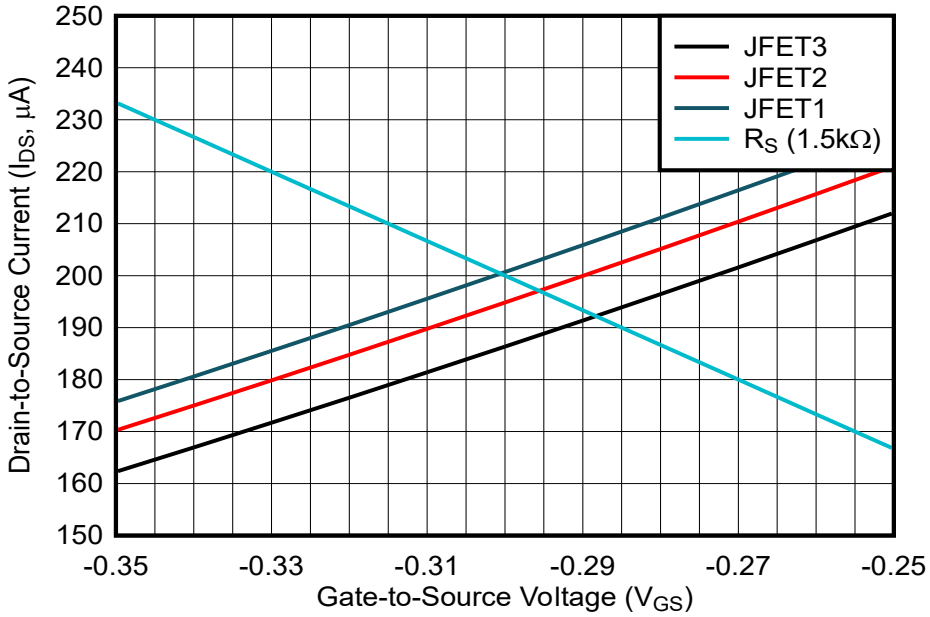


Figure 6. *I_{DS} vs V_{GS} plot for the three JFETs given in Table 1 with a 1.5kΩ load line*

Notice that the lower I_{DSS} devices, JFET2 and JFET3, intersect the R_S load line at V_{GS} values closer to zero, effectively pinching off the JFET channel less. Conversely, if the JFET I_{DSS} were increased, R_S would force the V_{GS} further from 0V and pinch off the JFET channel more. This is why a 2% mismatch in I_{DSS} values of the JFETs does not result in a 2% mismatch in output current from the current sources. Increasing R_S will reduce the mismatch between the two current sources, but with the downside of decreasing the output current of the current source.

Conclusion

Three-wire configurations are a common and effective way to remove the errors introduced by lead resistance to RTD measurements. However, three-wire systems require a pair of well-matched current sources to effectively remove errors. A simple method to create these matched current sources is to use a monolithic pair of N-JFETs such as the JFE2325. The JFE2325 includes two, well-matched, JFET devices on the same die inside of an extremely small 0.8mm × 1mm package. Because the two JFETs share the same silicon area, their parameters match each other far better than discrete single JFETs. Furthermore, adding a source resistor to the JFETs introduces negative feedback which reduces variation in the output current due to changes in the JFET parameters. This approach delivers two current sources with exceptional matching at a very low cost, small solution size, and able to operate over a wide range of power supply voltages.

References

1. Wu, Joseph, "**A Basic Guide to RTD Measurements**."
2. Texas Instruments, **JFE2325 Dual, Low-Power, N-Channel JFET for Electret Microphones**.

About the author

John Caldwell is a technologist for the linear amplifiers BU where he works across all product lines to define the next generation of high performance signal conditioning amplifiers. John has more than 15 years of experience in linear amplifier development and precision, low noise, circuit design. John received his B.S. and M.S. in electrical engineering from Virginia Tech, holds 3 U.S. patents, and has published more than 30 articles.

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.

All trademarks are the property of their respective owners.

© 2026 Texas Instruments Incorporated

Achieving high power density in a 48V buck converter with integrated GaN FETs and dual-side cooling

Matthew Bowers
Systems engineer, power design services

Introduction

The continued adoption of 48V power distribution architectures in automotive, industrial, and data center systems has increased the demand for compact, high-efficiency DC/DC converters capable of supporting high output currents while minimizing solution size. Integrated gallium nitride (GaN) power converters address these requirements by reducing switching losses and package parasitics, which enables increased power density compared to conventional silicon-based solutions.

As converter power density increases, thermal management becomes a critical design consideration. This article examines a 48V-to-12V, four-phase synchronous buck converter using integrated GaN power integrated circuits (ICs) and demonstrates how designers can leverage dual-side cooling alongside printed circuit board (PCB) thermal spreading to manage power dissipation at high load operation.

Integrated GaN for high power density conversion

Silicon metal-oxide semiconductor field-effect transistor (MOSFET) technology has served as the foundation of power converter design due to its availability and proven performance. However, as converter power levels and power density requirements continue to increase, conventional silicon-based implementations face growing challenges associated with switching losses and package parasitics. Integrated GaN power ICs address these limitations by combining high-speed GaN power FETs, gate drivers, and control circuitry within a single package, reducing parasitic inductance while simplifying converter implementation. The higher switching capability of GaN also enables operation at increased switching frequencies, which can reduce the required size of magnetic and capacitive components while enabling faster control-loop response.

Reducing the parasitic inductance between the GaN power FETs and integrated gate driver also improves control of the switching transitions, reducing unwanted ringing and associated switching losses. This improved switching behavior can help reduce electromagnetic interference (EMI), an important consideration for high-density power converters operating at high slew rates [1].

The four-phase converter design presented in this article uses the TI **LMG708B0** integrated GaN power multi-chip module (MCM) IC to support the efficiency and power-density requirements of 48V-to-12V step-down conversion.

Thermal challenges and dual-side cooling

Integrated GaN power ICs can improve converter efficiency and reduce solution size, but concentrating power dissipation within a small package increases thermal density and heat flux. As a result, achieving the full performance potential of integrated GaN depends not only on minimizing electrical losses but also on providing effective pathways for transferring heat from the device to the surrounding ambient environment. In a conventional PCB-cooled implementation, heat generated by the device must conduct through the package into the PCB, where copper planes and thermal vias spread the heat toward the surrounding surfaces and to the local ambient environment. As a general rule of thumb, approximately 1in^2 of cooling area per watt of dissipated power may be required for natural-convection cooling. Maximizing the thermal spreading capability of the PCB is therefore a critical part of managing the device temperature in a high-power-density converter.

The addition of a dual-side cooling structure provides a parallel thermal path from the LMG708B0 package directly to the external cooling interface. This additional path allows heat to be transferred through both the PCB and the dual-side cooling structure, reducing the overall junction-to-ambient thermal resistance. The top-side cooling approach therefore builds upon the PCB thermal path rather than fully replacing it, providing additional heat-removal capability while maintaining the compact footprint of the integrated GaN power ICs. **Figure 1** is a simplified representation of the package heat-flow path.

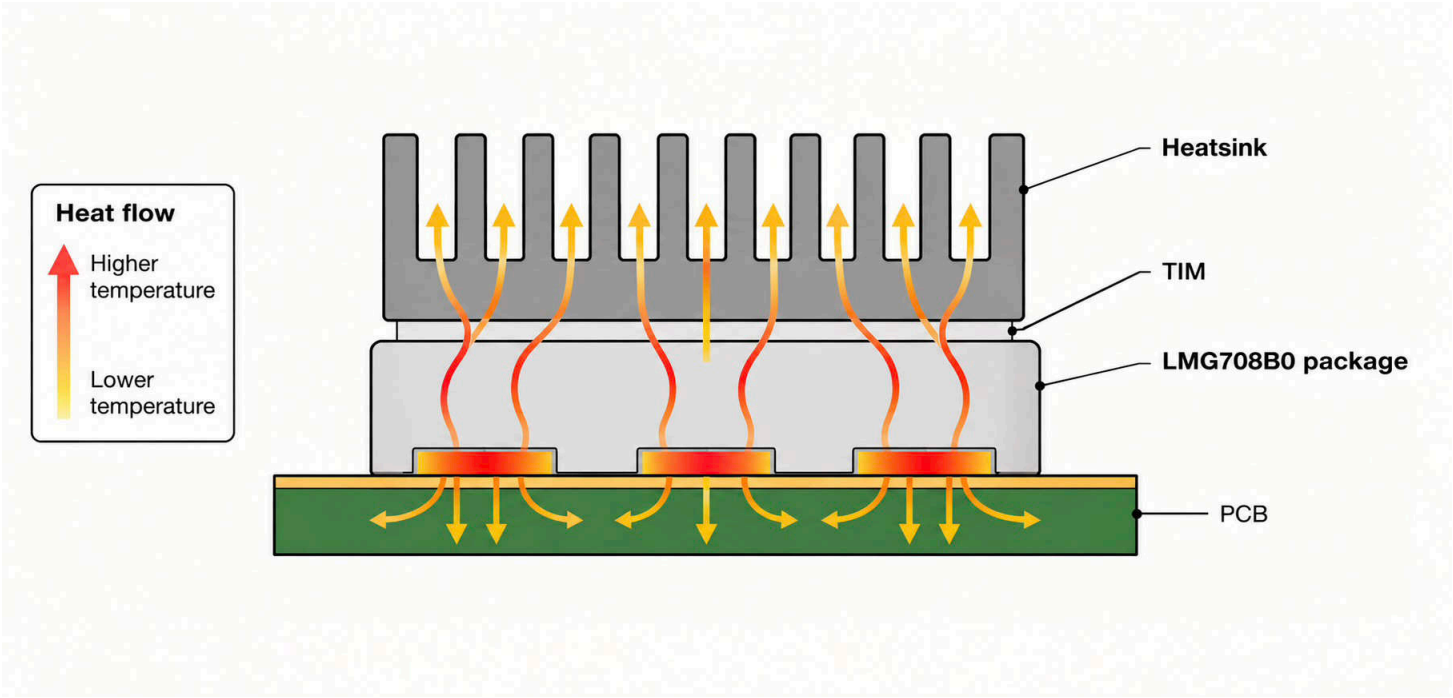


Figure 1. *Simplified heat-flow paths through the LMG708B0 integrated GaN power multichip module package*

GaN converter IC design and implementation

The **48V$_{IN}$, 960W Four-Phase Buck Converter with Integrated GaN Reference Design** was used to evaluate the thermal-focused design considerations. Distributing the load across four interleaved integrated GaN phases reduces the current and power dissipation in each phase while simultaneously supporting high output-current capability and compact solution size that 48V power conversion requires. **Table 1** lists the overall design specifications.

Design parameters	Value
Input voltage	48V (36V to 60V)
Output voltage	12V
Output current	80A (electrical design and thermal design current)
Topology	Four-phase synchronous buck
Switching frequency	400kHz (each phase)
Power stage	Four LMG708B0 devices
Reference design number	PMP23595

Table 1. *Buck converter design specifications*

To evaluate the resulting temperature distribution and heat flow through the package, cooling structure, and PCB a thermal simulation was conducted. The thermal simulation was created using Ansys Icepak® thermal modeling software tool. **Figure 2** shows the steady-state thermal simulation of the four-phase buck converter at the full 960W output at an ambient temperature of 25℃ with no airflow, and 6.29W of power dissipation assigned to each of the four LMG708B0 devices. The power dissipation was calculated by obtaining the total power losses of the reference design at the full 960W output under the 25℃ ambient temperature condition and subtracting the losses generated from the surrounding passives, leaving only the losses generated by the LMG708B0 devices. The device power dissipation corresponded to 25.16W of total power dissipation across the four phases. The simulated package temperatures are approximately 108℃ across the four phases, still well below the recommended case temperature of the LMG708B0 device, which is 135℃.

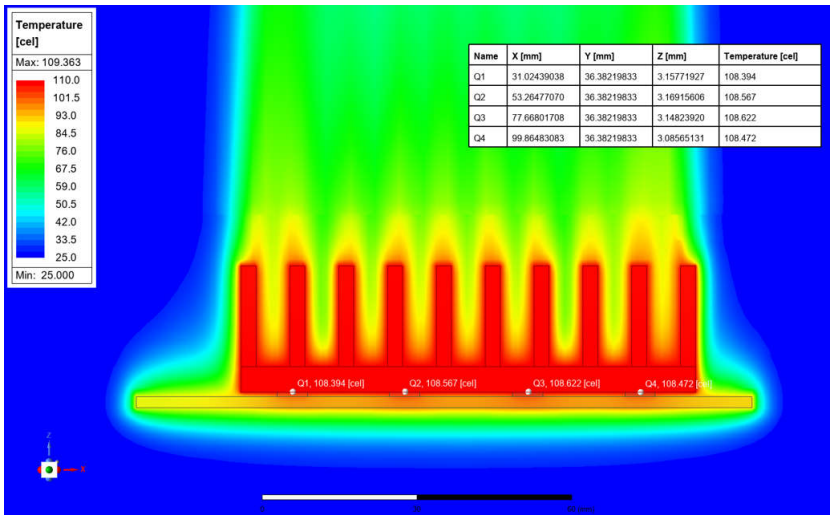


Figure 2. *Thermal simulation of the four integrated GaN power stages with the dual-side cooling structure*

The simulation also illustrates the thermal paths from the integrated GaN devices through both the dual-side cooling structure and the PCB. However, the simulation represents an approximation of the final hardware implementation. The modeled heat-sink geometry differs slightly from the final heat sink and does not account for the thermal interface material (TIM) used between the LMG708B0 packages and the heat sink. The TIM introduces additional case-to-sink thermal impedance that can affect the overall thermal performance of the dual-side cooling path. Therefore, the simulated temperatures provide a thermal reference for comparison with the measured converter performance discussed later in the article.

Experimental evaluation

The operating load range for evaluating converter performance was at a 48V input and 12V output. **Figure 3** shows the measured efficiency and corresponding power losses as a function of output current. Presenting both quantities provides a more complete view of converter performance. Efficiency indicates the effectiveness of power conversion, whereas the power-loss curve represents the thermal load that the converter must dissipate. To maximize the effect of the dual-side cooling package, a heat sink was attached during data capture with no additional airflow used. Details about the heat sink and the test setup are discussed later in the article.

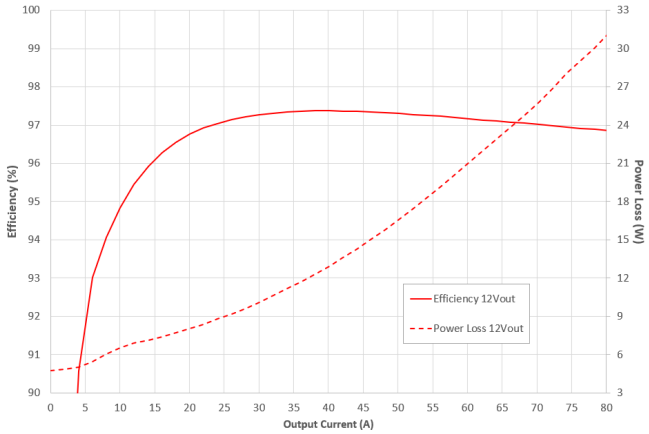


Figure 3. *Measured efficiency and power loss vs output current*

At lighter loads, relatively fixed losses associated with inductor core loss, gate drivers, and other control-related losses represent a larger percentage of the processed power. As the output current increases, these fixed losses become less significant relative to the output power and efficiency improves, reaching a maximum near 40A. Beyond this point, current-dependent losses become increasingly dominant because resistive conduction losses increase approximately with the square of the current (I^2R). Conduction losses within the power ICs, inductors, and high current PCB copper paths therefore increase rapidly with load current, resulting in a gradual reduction in efficiency as the converter approaches full load.

Although the efficiency decreases by only 0.5% between the peak efficiency and the maximum load, the corresponding power losses increase by 18.2W. This increase in absolute power losses represents a significant thermal load within the compact converter footprint and reinforces the importance of providing effective heat-removal paths for the power stages and surrounding components.

The measured power-loss data also provides context for the thermal simulation in **Figure 2**. The simulation assigned 6.29W of dissipation to each LMG708B0 IC, corresponding to 25.16W across the four phases at the 80A full-load operating point. At the measured 80A operating point, the total converter power losses are 31.2W. The difference between these values represents losses distributed among other converter components, mainly inductors and current sense shunt resistors, and PCB current paths, while also reflecting the assumptions used in the thermal model. Nevertheless, the comparison indicates that a substantial portion of the total thermal load is associated with the four integrated GaN power ICs, supporting the use of dedicated dual-side cooling to supplement heat spreading through the PCB.

Thermal performance

To characterize both the package-level temperature distribution and the effectiveness of the dual-side cooling approach, thermal performance was evaluated under two steady-state operating conditions. The first measurement, shown in **Figure 4**, is a 50A output current with the heat sink removed, enabling the thermal camera to directly observe the four LMG708B0 packages and surrounding PCB. At this operating point, the four-phase implementation delivers 600W of output power while dissipating approximately 16.6W of total measured power losses. The maximum observed temperature was approximately 97℃, with the highest temperatures concentrated around the power-stage region.

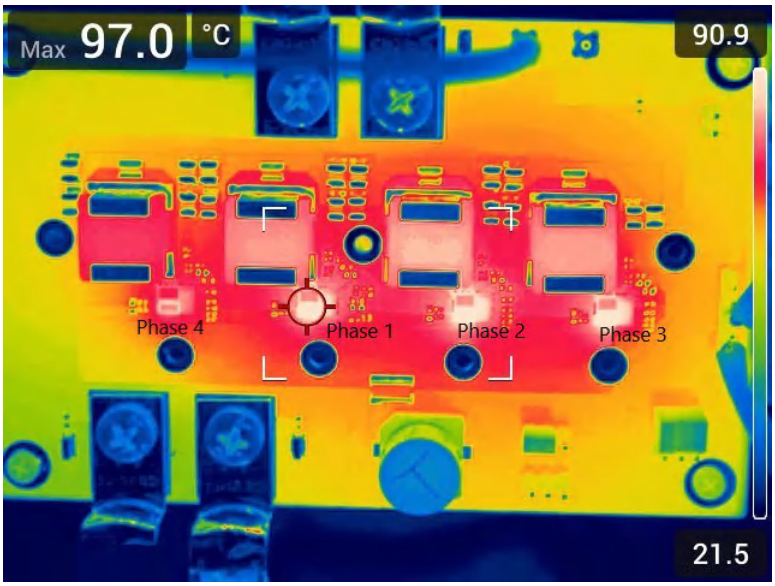


Figure 4. *Infrared thermal image at 50A load without heat sink*

Subsequently, the converter was evaluated at the full 80A load condition with a heat sink installed to maximize the effect of the dual-side cooling structure. The heat sink is made from aluminum, with a base size of 10.72 centimeters (4.22 inches) by 2.70 centimeters (1.06 inches) with a height of 1.61 centimeters (0.63 inches). Additionally, the heat sink has seven 2.54 centimeters (1.00 inch) high fins that extend from the base. The heat sink is designed to have a simple yet effective way of extracting heat away from the IC through the dual-side cooling package. Most applications, notably in the automotive sector, will have either a larger heat sink, such as the housing of the electronic control unit (ECU), or a more optimized cooling structure that will enable better performance. The TIM used between the heat sink and the LMG708B0 devices was the T-WORK9000, which is 1mm thick and has a thermal conductivity of 20W/m-K. At this operating point, the converter delivers 960W, corresponding to approximately 31.2W of total converter power losses. With the heat sink covering the integrated GaN packages, the thermal image shown in **Figure 5** primarily illustrates the temperature distribution across the external cooling structure and surrounding converter. The thermal image was taken after letting the converter maintain steady-state operation at 960W for approximately 20 minutes to reach thermal equilibrium.

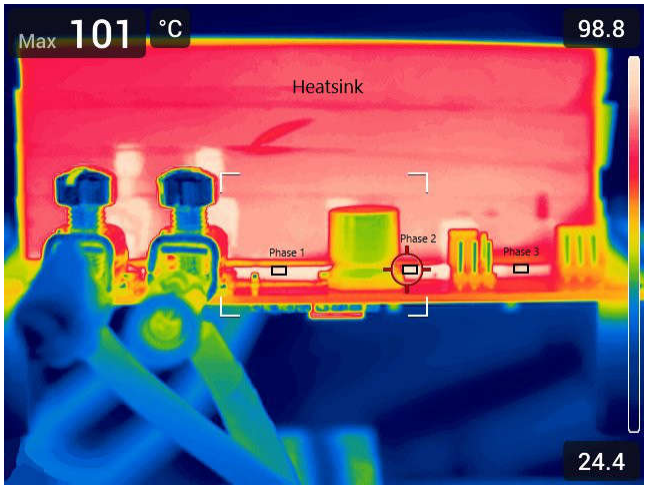


Figure 5. *Infrared thermal image at 80A load with heat sink installed*

Together, the two measurements demonstrate the complementary thermal paths incorporated into the converter design. The 50A measurement in **Figure 4** provides a direct observation of the power-stage and PCB temperature distribution, while the 80A measurement in **Figure 5** evaluates the complete cooling implementation under the maximum operating condition and illustrates the effectiveness of the dual-side cooling structure. These results complement the thermal simulation presented in **Figure 2**, which modeled the 80A full-load condition. **Figure 5** presents a glimpse of the power stage region showing each IC to have an approximate temperature of 101℃. A more accurate case temperature was taken using thermal coupling with the wire leads placed on top of the primary LMG708B0 device but underneath the TIM and heat sink. The case temperature of the primary LMG708B0 device was 105.4℃. **Figure 6** shows the setup used for measuring the case temperature of the primary IC.

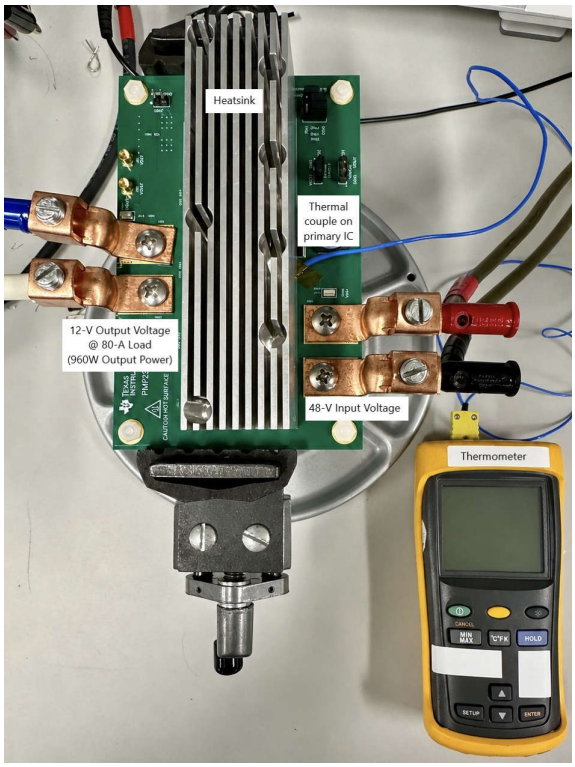


Figure 6. *Test setup for thermal evaluation*

These results coincide with the thermal simulation results to show the significance of the dual-side cooling structure. Moreover, the use of a larger or more reliable cooling system or adding a fan to provide airflow would achieve lower temperatures and thus enable even better system performance.

Conclusion

High-power-density converters require consideration of both electrical and thermal performance throughout the development process. While integrated GaN power ICs enable improvements in efficiency and power density, realizing these benefits depends on careful converter architecture, PCB layout, and thermal management.

The $48V_{IN}$-to-$12V_{OUT}$ four-phase synchronous buck converter design presented in this article demonstrates how integrating dual-side cooling within the overall converter design supports high efficiency while providing an effective path for managing power dissipation at high load. Considering electrical performance, power losses, and thermal behavior together enables compact and reliable integrated GaN power converters for high-current 48V power-conversion applications.

About the author

Matthew Bowers is a systems engineer in TI's Power Design Services team, focused on developing power solutions for automotive applications. He received his bachelor's degree in electrical engineering from Texas Tech University in 2023.

References

1. Texas Instruments. March 2016. "**Optimizing GaN performance with an integrated driver**." Texas Instruments white paper No. SLYY085. Accessed September 22, 2026.

Additional resources

- For more insight into the optimization of 48V multiphase applications, read the article "**Power Tips #151: Improving Efficiency in 48V-Input Multiphase Buck Converters with GaN**."

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.

Icepak® is a registered trademark of ANSYS, Inc.
All trademarks are the property of their respective owners.

© 2026 Texas Instruments Incorporated

Redundant power supplies in automotive applications – priority switching done right

Björn Görner
Systems engineer, power design services

Introduction

Safety-critical systems in automotive applications often depend on redundant power supplies to achieve the required Automotive Safety Integrity Level. Systems designed specifically for autonomous driving must work under all circumstances, at least to a degree where the vehicle can achieve a safe state. For redundant power supplies, designers typically use one of two implementations: ORing or priority switching. While both implementations are nearly identical, there are some details that can cause issues if not properly addressed.

The differences between ORing and priority switching

An ORing application connects two or more inputs with diodes, or ideal diodes or ORing controllers if the system requirements call for higher power. **Figure 1** shows the typical implementation of such an ORing scheme. This implementation delivers whatever input voltage is the highest to the output. Equal input voltages are all loaded relatively equally.

The ORing approach is easy to implement and a good fit for systems that offer equally usable supplies. Because both power supplies will load simultaneously when their voltage levels are close to each other, the transition from one supply to the other happens somewhat smoothly. The use of back-to-back field-effect transistors (FETs) is optional in this implementation, as long as there is no need to disconnect one or both power sources.

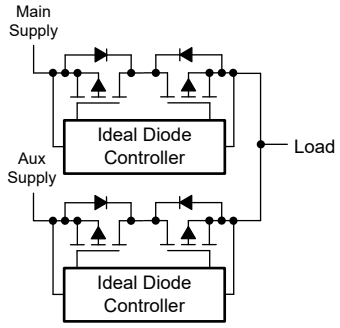


Figure 1. *Input ORing*

For an electric vehicle, which has a high-voltage traction battery with a powerful step-down converter and a small 12V or 48V auxiliary battery, the standard ORing approach may not be suitable. This is because of the low capacity of the auxiliary battery and charging circuit, which are both designed for emergency operation only. In such cases, the priority switching approach achieves redundancy without unnecessarily loading the small auxiliary battery.

As shown in **Figure 2**, priority switching monitors the main power supply through the ideal diode controller of the auxiliary supply and only switches over to the auxiliary supply if the main supply drops below a certain threshold. The ideal diode controller must support an overvoltage protection feature, which in this case is used to detect an undervoltage event on the main power supply. The LM74900-Q1 device family would be an example of ideal diode controllers that have an overvoltage protection feature. The use of back-to-back FETs in this application is mandatory, in order to avoid current draw from the disabled source.

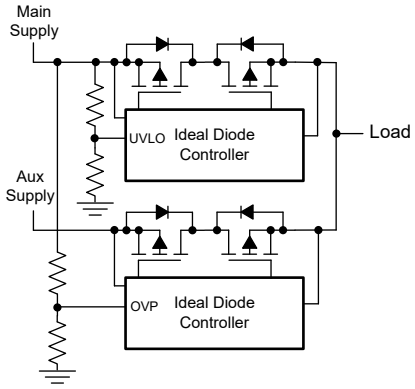


Figure 2. *Priority switching*

Priority switching details

In general, the circuits used for ORing and priority switching are not particularly different from one another, but the behavior of the priority switching implementation requires some attention to detail.

While the ORing solution enables both power paths if they are somewhat equal, the priority switching approach keeps the auxiliary power path disabled until the voltage of the main power path drops below a certain threshold. Reaching this threshold disables the main power path and enables the auxiliary power path, causing a sudden load shift from one supply to the other and generating a sudden voltage step on both supplies. Because of this sudden load shift, the voltage on the main power supply will often spike for a short period before completely collapsing, although this depends on the failure causing the initial drop.

Because the auxiliary supply switches based on the voltage on the main supply, a voltage spike can cause the system to switch back to the primary supply, which immediately lowers the primary voltage and causes the system to revert to the auxiliary supply. This results in uncontrolled oscillation between the two power supplies and creates a large ripple on the output, which may disrupt system functionality. The reason behind the voltage shift on each supply is the sum of all parasitic resistances and inductances within the power path, such as printed circuit board traces, wiring, and the MOSFET drain-to-source on-resistance ($R_{DS(on)}$).

Figure 3 shows an example of such behavior. During a disconnect of the primary power supply, a short burst occurs right when the input voltage drops below the switchover threshold. Channel 1 in **Figure 3** shows the voltage on the auxiliary supply, while Channel 3 shows the output voltage. The auxiliary supply voltage in this scenario has a lower value than the primary supply voltage, with the switchover threshold set right above the voltage level of the auxiliary supply.

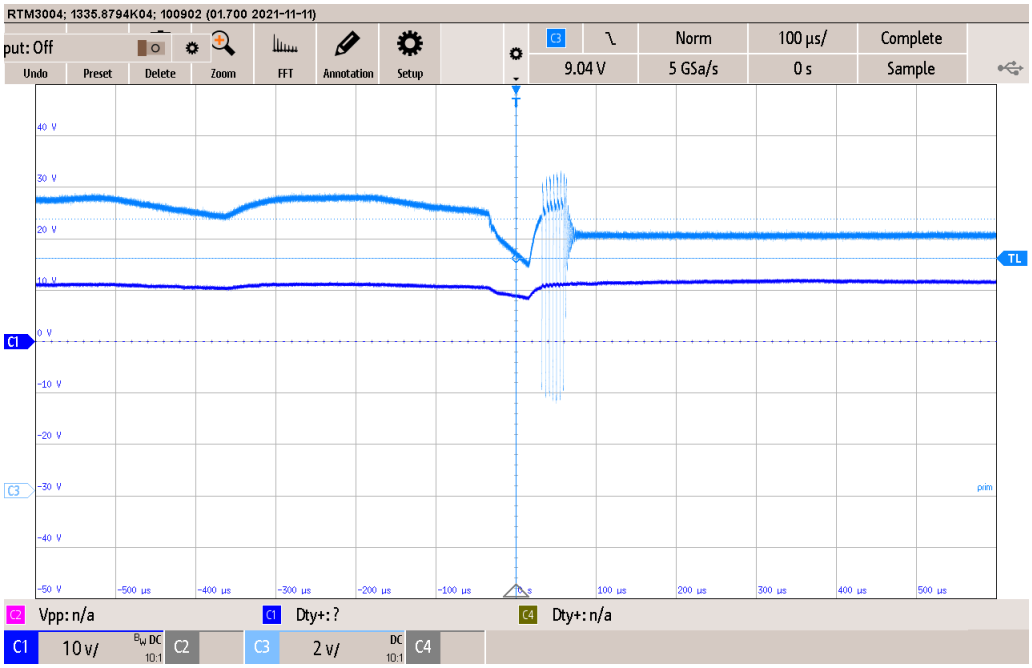


Figure 3. *Oscillation during input switchover in a priority switching system*

Depending on what problem causes the main supply to drop, this oscillation may even be continuous; for example, if the cable for the main supply is partially damaged and causing a huge voltage drop instead of a complete disconnect. In many cases the hysteresis built into the ideal diode controllers, together with a sufficiently high input capacitance, can suppress this behavior. But in some cases the hysteresis may not be enough to avoid unnecessarily large input capacitors.

Although it is not possible to change the internal hysteresis, on some devices there is a simple solution to increase it externally. Many devices such as the LM74910H-Q1 offer a fault output that you can use to add external hysteresis with minimal component count. On the LM74910H-Q1, the fault pin is an open-drain output that pulls low for overvoltage conditions and other faults. This output can also control an external transistor, which alters the feedback divider for the overvoltage input. **Figure 4** illustrates a typical implementation.

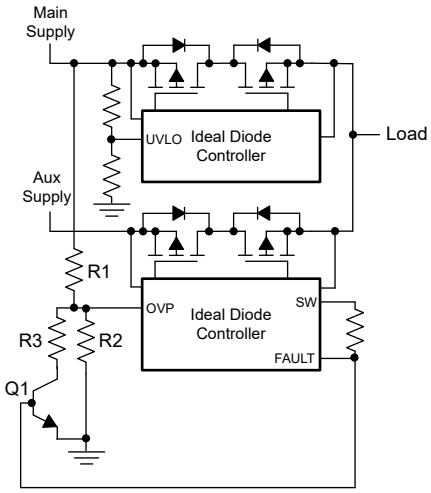


Figure 4. *External hysteresis*

Since the fault output is an open-drain output, you will need a bias resistor to turn on the bipolar junction transistor. This resistor connects to the software pin of the integrated circuit, which offers a switchable low-current path to the input that becomes disabled when the system shut downs to minimize quiescent current.

As long as the primary power supply is above the switchover threshold, the auxiliary controller is in an overvoltage fault state, disabling the power switches and pulling the fault output low. Pulling down the fault output disables transistor Q1; therefore, the feedback divider for the overvoltage input only consists of R1 and R2. This results in an overvoltage feedback voltage defined by the ratio of both resistors and the switchover threshold, as expressed by **Equation 1**:

$$V_{ov_falling} = V_{sense_falling} \frac{R1}{R2} \quad (1)$$

As soon as the primary supply voltage drops below this overvoltage threshold, the fault output no longer pulls low and the transistor switches on, connecting R3 in parallel to R2 and reducing the voltage appearing at the overvoltage input. **Equation 2** now gives a new value for the rising threshold:

$$V_{ov_rising} = V_{sense_rising} \frac{R1}{\frac{R2 \; R3}{R2 + R3}} \quad (2)$$

Changing the value of R3 easily adjusts the resulting hysteresis, offering great design flexibility.

Conclusion

The choice between an ORing or priority switching redundant power solution depends on your system design. While both designs offer the same function, the ORing implementation is less critical when it comes to actual hardware. Parasitic effects such as board inductance, board resistance, the $R_{DS(on)}$ of the power transistors and cable resistance can negatively affect the switchover behavior of a priority switching implementation. In some instances, the internal hysteresis of the controllers is enough to counteract the influence of the parasitic elements, but not always.

The circuit described in this article shows a simple workaround for adding external hysteresis with minimal component count. The only requirement is a controller, which offers a suitable signal for switching between the two voltage levels. This controller helps reduce the capacitance on the supply inputs required for a clean switchover, resulting in a stable system and cost savings.

About the author

As part of the Power Design Services team, Björn Görner has eight years of experience designing custom power solutions and reference designs, with his main focus on automotive applications. He is a regular presenter at TI's Power Supply Design Seminars, which Texas Instruments offers biannually to educate and support engineers interested in power conversion.

All trademarks are the property of their respective owners.

© 2026 Texas Instruments Incorporated

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, regulatory or other requirements.

These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you fully indemnify TI and its representatives against any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI’s products are provided subject to TI’s Terms of Sale, TI’s General Quality Guidelines, or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI’s provision of these resources does not expand or otherwise alter TI’s applicable warranties or warranty disclaimers for TI products. Unless TI explicitly designates a product as custom or customer-specified, TI products are standard, catalog, general purpose devices.

TI objects to and rejects any additional or different terms you may propose.

Copyright © 2026, Texas Instruments Incorporated

Last updated 10/2025