Texas Instruments

ISSUE 1, 2025

Analog Design JOURNAL

Table of contents

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

Avishek Pal
Power Switches

Rakesh Panguloori
Power Switches

Introduction

With advancements in artificial intelligence (AI) and machine learning, enterprise servers have become extremely power-hungry as they simultaneously process a large amount of data and storage. The steady-state power rating of each server motherboard has gone up to 5kW or 6kW, in contrast to 1kW or 2 kW for general servers. The form factor remains the same, however, which imposes system design challenges given the increased power density. The load amplitude, slew rate and frequency of transient loads on AI servers have increased three to four times compared to general servers.

Figure 1 shows a typical power distribution in a 48V rack server where the input is protected by the hot-swap circuit - and then distributed to all downstream system loads.

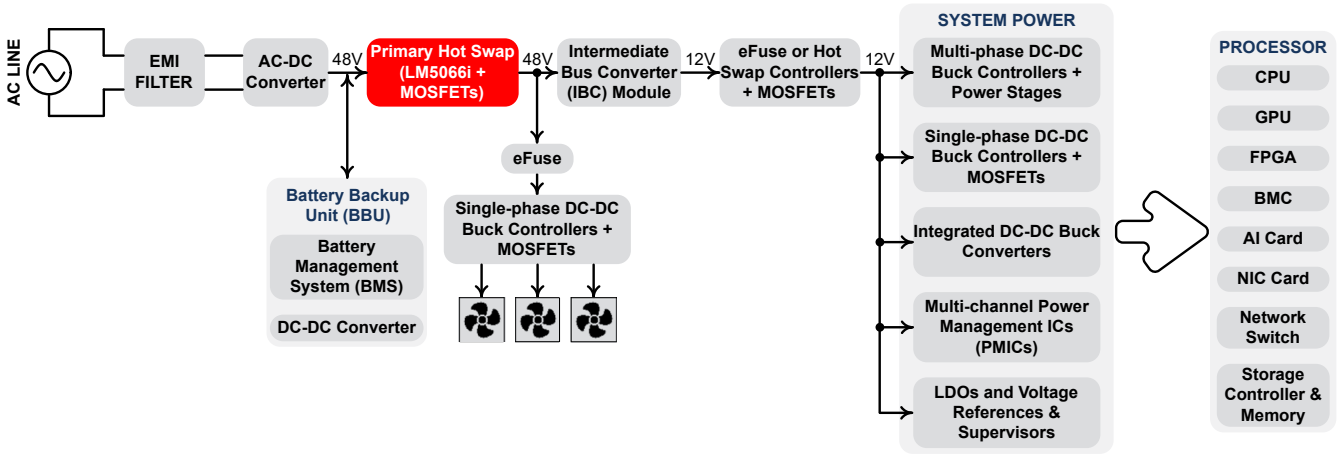

Figure 1. *Typical block diagram of a 48V rack server power distribution.*

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

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

Table 1. *Typical system specifications.*

Challenges in designing a hot-swap circuit for a 48V AI server

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

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

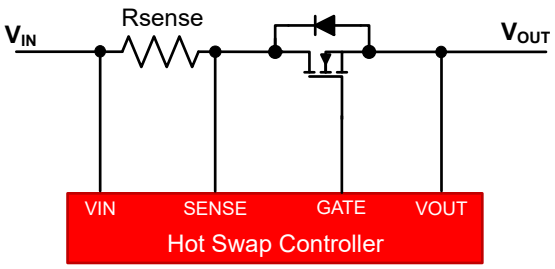

Figure 2. *Traditional power-limiting hot-swap circuit.*

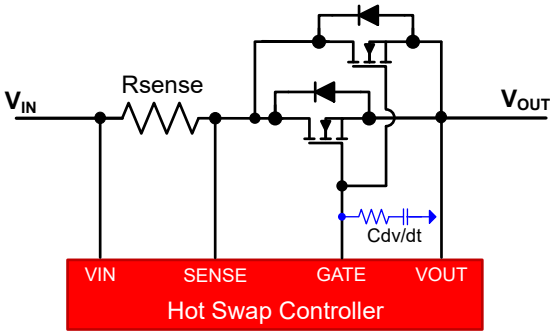

Figure 3. *Hot-swap circuit with GATE slew-rate control.*

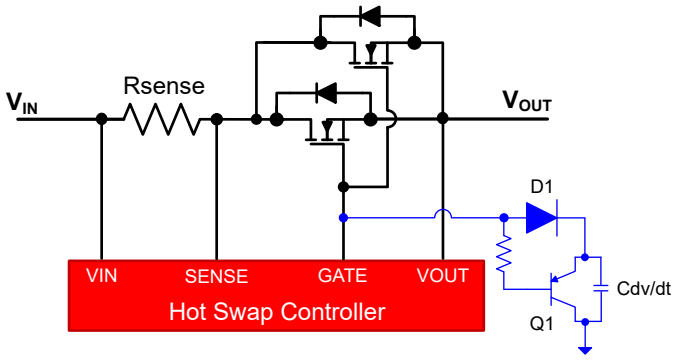

Figure 4. *Hot-swap circuit with local discharge path for Cdv/dt.*

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

In the output-voltage slew-rate control method, capacitor $C_{dv/dt}$ placed across GATE-GND (see **Figure 3**) limits the slew-rate of the gate and output voltages, which limits the inrush current. MOSFETs can handle more energy when the power dissipation in them is reduced and spread over longer durations. Therefore, as the output capacitance increases, you need a higher $C_{dv/dt}$

to reduce both the inrush current and power dissipation in the MOSFET during startup.

A higher Cdv/dt interferes with the turnoff process, however, the hot-swap controller has limited pulldown strength. This necessitates a local P-channel N-channel P-channel (PNP)-based discharge circuit for Cdv/dt, as shown in **Figure 4**. During startup, Cdv/dt controls slew-rate in the same way, but during a turnoff event, the Q1 PNP transistor activates and discharges $C_{dv/dt}$ locally. Diode D1 blocks the discharge of Cdv/dt into the GATE pin, which reduces the stress on the GATE pin and also ensures proper operation of the controller.

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

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

With the increase in load current, more MOSFETs need to be parallel to limit the maximum steady-state MOSFET junction temperature to a safe value (100°C to 125°C). For example, to support a steady-state load current of 150A at an ambient temperature of 70°C, eight Texas Instruments (TI) CSD19536KTT MOSFETs need to be in parallel to limit the steady-state MOSFET junction temperature to 100°C. Paralleled MOSFETs help thermally, but increase the effective capacitance on the GATE pin of the hot-swap controller and impact the turnoff response.

During an output short-circuit, the MOSFETs need to turn off fast enough to prevent further buildup of fault current and avoid damage to the MOSFETs, input power supply, or printed circuit board (PCB). The gate pull-down strength of the TI LM5066I hot-swap controller is limited to 160mA, which is not enough to turn off all eight MOSFETs completely during a short-circuit event, as shown in **Figure 5**.

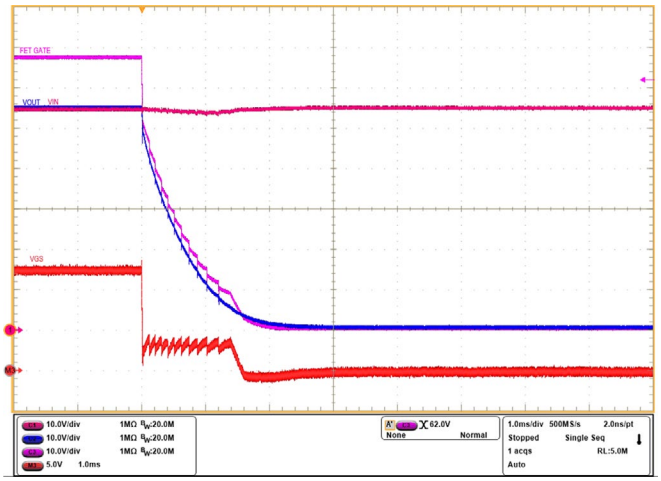

Figure 5. *Short-circuit response of the LM5066I controller with eight MOSFETs.*

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

Although the local PNP-based discharge circuit for Cdv/dt helps reliably turn-off the MOSFETs during an output short-circuit event, it causes a false GATE turn-off in the presence of high-frequency, high slew-rate load transients. During load step-up, the MOSFET source node drops because of the finite input and output impedances of the hot-swap circuit. The voltage drop at the source node gets coupled to the MOSFET gate node through the C_{GS} capacitance of the MOSFET and causes the gate node to drop as well. The MOSFET source node recovers during load step-down. The gate node cannot recover completely to its previous level, because of the limited gate current (20µA typical) of the LM5066I hot-swap controller. As a result, the hot-swap controller gate continues to drop further in the subsequent load transient cycles developing the base-emitter voltage for Q1. Finally, PNP bipolar junction transistor Q1 turns on, and falsely shuts down the system. **Figure 6** illustrates the whole process, while **Figure 7** shows the corresponding test result.

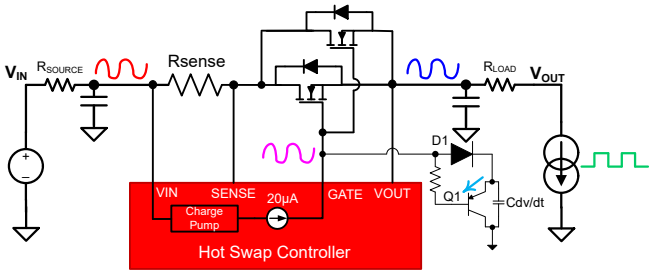

Figure 6. *Illustration of a hot-swap circuit for a dynamic load.*

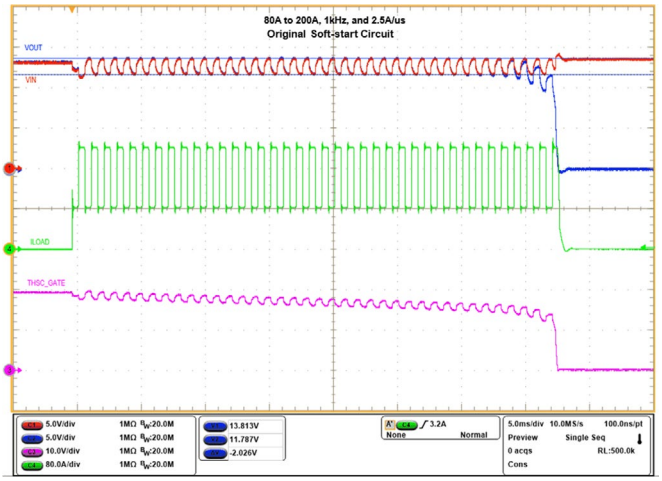

Figure 7. *Response of a hot-swap circuit to a dynamic load.*

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

Generally, parallel MOSFETs are more prone to parasitic oscillations than a single MOSFET in the linear region of operation. This is because of the presence of parasitic stray package inductances and capacitances on the drain, source and gate nodes, which form a resonant tank circuit resembling a Colpitts oscillator. Unlike switching regulators with a gate-drive strength of >2A, hot-swap controllers with a lower gate-drive strength (20µA) limit the inrush current during start-up by operating the MOSFETs in the linear region. As a result, the parallel combination of hot-swap MOSFETs is highly susceptible, with more chance of generating sustained oscillations. This phenomenon causes the violation of the MOSFET SOA during a power-into-short fault, leading to MOSFET damage.

Proposed circuit enhancements

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

Improving the turn-off response

In the proposed solution shown in **Figure 8**, introducing an external fast pull-down circuit using - PNP transistor (Q_{PD} and R_{PD}) will boost up the turn-off speed. During an output short-circuit event, the gate pull-down current of 160mA creates a substantial voltage drop across the R_{PD} resistor and enables fast pull-down of the PNP transistor (Q_{PD}). This in turn shorts the gate-to-source of all parallel MOSFETs, turning off the MOSFETs immediately to quickly disconnect the power path. **Figure 9** shows the experimental result for a short-circuit event with a fast pull-down circuit.

Overcoming false turn-off for dynamic loads

In this solution, the hot-swap gate node is decoupled from the MOSFET gate terminal by placing the D_{SS} diode between them, again shown in **Figure 8**. This modification helps eliminate the reflection of output voltage ripple to the hot-swap controller GATE node and avoids false turn-on of the soft-start PNP transistor, Qss. Changing the position of the diode does not impact controller behavior during start-up nor any of the fault events. As shown in the test result (see **Figure 10**), the system operates continuously even for large load steps from 20A to 120A at a 1kHz frequency.

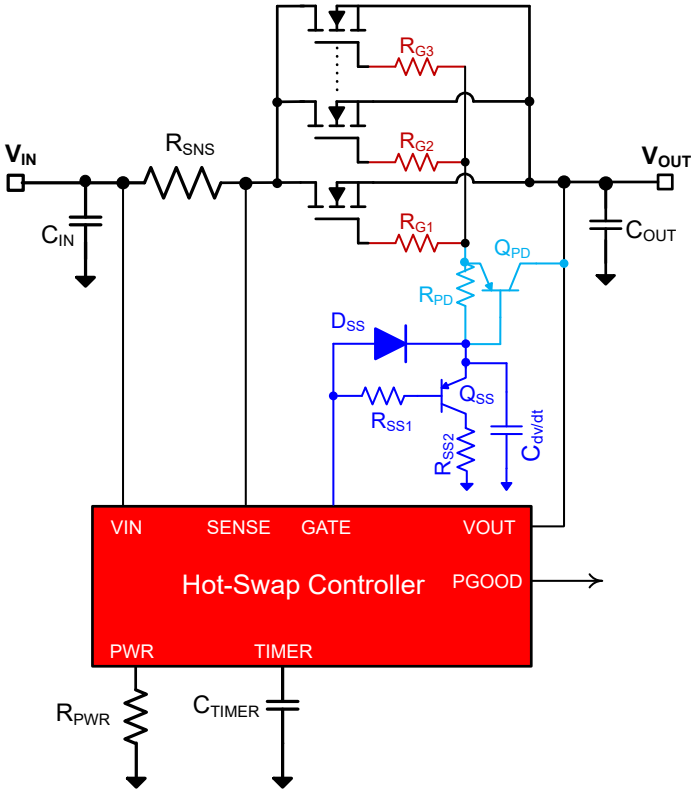

Figure 8. *Proposed hot-swap circuit configuration.*

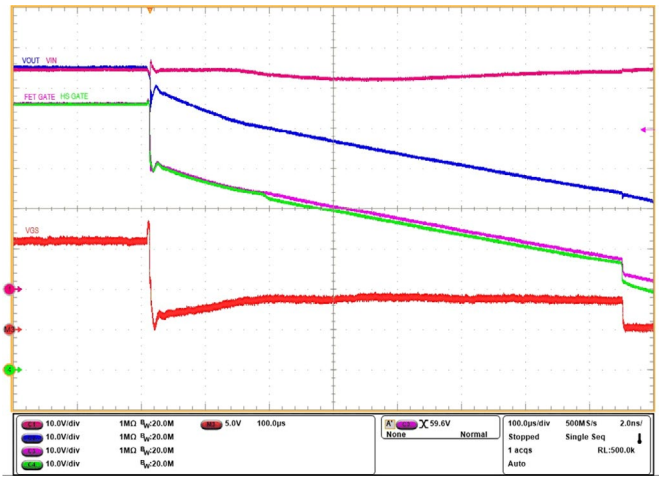

Figure 9. *Output short-circuit response with fast pull-down circuit.*

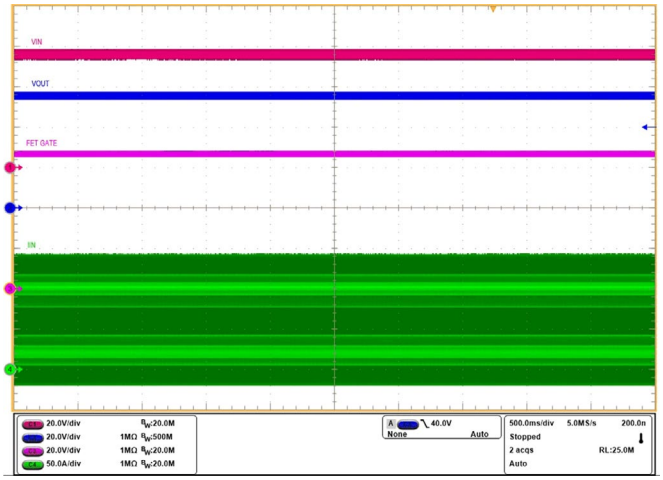

Figure 10. *Load transient performance for steps from 20A to 120A to 20A at a 1kHz frequency.*

Damping parasitic oscillations

Adding a damping resistor (R_{G1}, R_{G2}, R_{G3}) in series with the gate of each MOSFET can eliminate the parasitic oscillations in the system. Usually, we recommend a 10Ω 0603 package resistance, but based on the parasitics, a low value around 1Ω may also help. We suggest testing on your PCB and deciding the value of the damping resistor.

Design guidelines and component selection

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

Feeding the system specifications shown in Table 1 into the **LM5066I design calculator** will obtain the values of the current-sense resistor (R_{SNS}), power-limiting resistor (R_{PWR}), fault timer capacitor (C_{TIMER}), soft-start capacitor ($C_{dv/dt}$) and number (N) of selected MOSFETs to parallel. In the **8kW Hot-Swap Reference Design for 48V Artificial Intelligence Servers** [2], R_{SNS} = 330µΩ, R_{PWR} = 28.7kΩ, C_{TIMER} = 10nF, $C_{dv/dt}$ = 47nF and N = 8.

Looking at **Figure 8**, select the R_{PD} resistor using **Equation 1**:

$$R_{PD} > \frac{V_{BE(sat)}}{I_{GATE(CB)}} \quad (1)$$

where, $V_{BE(sat)}$ is the base-emitter saturation voltage of the Q_{PD} PNP transistor and $I_{GATE(CB)}$ is the power-on reset circuit-breaker sink current in the LM5066I hot-swap controller. The 8kW hot-swap reference design uses an R_{PD} value = 20Ω.

$C_{dv/dt}$ discharge circuit

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

You will have to select R_{SS1}, R_{SS2} and Q_{SS} iteratively so that none of the three components become stressed during turn-off. For QSS, you can select any standard PNP transistor with collector-emitter (V_{CEO}) and collector-base (V_{BEO}) voltages of >100V_{DC} and a continuous collector current of >200mA. Select the values for R_{SS1} and R_{SS2} and their respective power ratings to limit the current flowing through the Q_{SS} transistor to a safe value. You must use a special high-power resistor for R_{SS2} to manage the transient peak power stress during turn-off. The 8kW hot-swap reference design uses onsemi MMBT5401LT1G for Q_{SS}, with R_{SS1} = 100Ω and R_{SS2} = 499Ω (the Vishay RCS0805499RFKEA).

Input transient voltage suppression (TVS) diodes are required to protect against transient overvoltages during input hot-plug and output short-circuit events. The TI **TVS diode recommendation tool** can help you obtain the part number (voltage and power ratings) of the TVS diode and the number of TVS diodes to parallel. The 8kW hot-swap reference design uses three Littelfuse 8.0SMDJ60A TVS diodes. For a deeper analysis into TVS diode selection, see Reference [3].

You will need output Schottky diodes to protect the output pin of the hot-swap controller against a negative transient in the event of an output short-circuit event. The 8kW hot-swap reference design uses three onsemi FSV20100V Schottky diodes.

Conclusion

The emerging 48V AI servers demand significantly more power, both in peak and steady states, than traditional servers. The high-power consumption along with fast and transient dynamics impose challenges in designing front-end protection using a hot-swap controller and parallel MOSFETs. The challenges include fast turn-off of parallel MOSFETs for real faults while avoiding false turn-off for high-frequency transients from the computational load. The proposed solution in this article eliminates the limitations of legacy hot-swap controllers and enables the design of a reliable input protection solution for a 48V AI server.

References

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

Related Websites

- **LM5066I**
- **PMP23496**

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.

© 2025 Texas Instruments Incorporated

How to limit PFC re-rush current

Bosheng Sun
Power Design Services

Introduction

The recently released Modular Hardware System – Common Redundant Power Supply (M-CRPS) specification [1] requires the re-rush current (which is different from the well-known inrush current) needs to be limited when the input voltage resumes after an input brownout or blackout event on the power supply used in data center. Previously, this re-rush current was unspecified, and no special control action for this event existed. In this article, I'll present a low-cost, simple and very effective method to meet the M-CRPS requirement.

Inrush current vs. re-rush current

Power factor correction (PFC) is required for front-end power supplies (PSUs) greater than 75W. PFC forces the input current to follow the input voltage so that the electronics load appears as a resistor. PFC usually has a large output capacitor. Before startup, the PFC output capacitor is fully discharged. Because the PFC structure provides a current path when $V_{AC} > V_{OUT}$, applying the AC voltage will generate a huge current, since the input voltage is essentially applied to the PFC output capacitor directly. This current is called the inrush current.

Placing a thermistor (R_T) with a positive temperature coefficient and a mechanical relay at the PFC input side will limit the inrush current, as shown in **Figure 1**. During PFC power up, the relay is off. The inrush current is limited by R_T to a low value, and the PFC output bulk capacitor (C_{BULK}) charges gradually. Once the output voltage (V_{OUT}) charges to equal the peak value of the AC voltage (V_{AC}), the inrush current drops to 0. Then the relay turns on, with R_T bypassed to reduce power losses during normal operation.

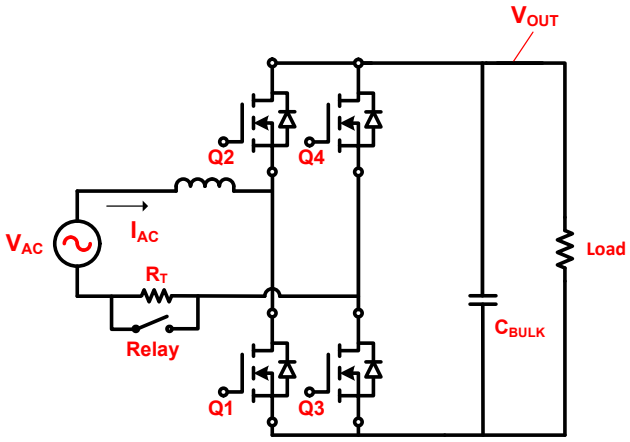

Figure 1. *Using R_T and a relay to limit the PFC inrush current.*

Re-rush current is different; it occurs during normal PFC operation. As shown in **Figure 2**, the AC input voltage suddenly drops out when PFC is operating normally. Since the load is still applied, the PFC V_{OUT} could drop to a lower value. Then when the AC voltage returns, if the AC input voltage is higher than V_{OUT}, there will be an inrush current again. This current is called the re-rush current.

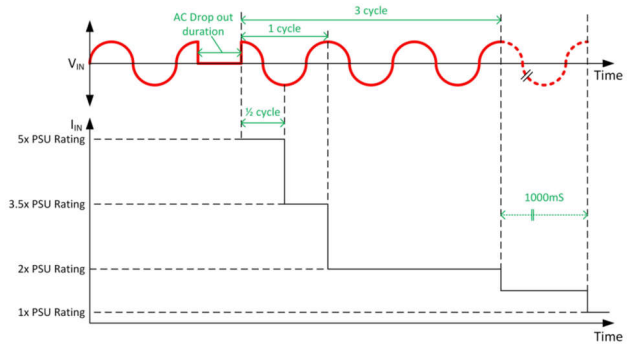

Figure 2. *M-CRPS re-rush current limit and timing.*

Previously, it is solely relies on the power stage components' ability to handle the re-rush current. Test results show that the re-rush current can jump to more than 10 times higher than the PFC-rated maximum input current. Such a high re-rush current can either damage the power supply or reduce its lifetime, which is why

the M-CRPS specification limits the amount of re-rush current after the AC voltage returns. The root-mean-square value of the re-rush current should be less than five times the maximum PSU current rating ($5 \times I_{rated,RMS}$) over a half cycle of input frequency, and less than $3.5 \times I_{rated,RMS}$ over one cycle of input frequency. In addition, the input current should settle to a value $\leq 2 \times I_{rated,RMS}$ within two cycles of the input frequency after applying the AC input.

It gets more complicated when considering PFC pulse-width modulation (PWM) operation during this period. If the PFC is not well controlled, an inappropriate PWM duty cycle may occur when the AC voltage resumes, resulting in another large input current spike that may also exceed the M-CRPS specification.

On the other hand, when the AC voltage resumes, PFC needs to supply enough current to boost the PFC output voltage to its regulation level as soon as possible; otherwise, V_{OUT} will keep dropping because of the heavy load and eventually trip the input undervoltage lockout level of the DC/DC converter. Charging the PFC output capacitor once the AC voltage resumes will require a large input current, either from re-rush when $V_{IN} > V_{OUT}$ or from the PFC control loop when $V_{IN} < V_{OUT}$.

This paper provides a solution to handle this re-rush current so that when the AC voltage comes back from dropout, both the re-rush current (when $V_{IN} > V_{OUT}$) and the non-re-rush current (when $V_{IN} < V_{OUT}$) are well controlled and high enough to rapidly boost V_{OUT} but not exceed the M-CRPS limit specification.

Proposed re-rush current control method

Figure 3 illustrates the proposed low-cost re-rush current control method. There are two differences compared to **Figure 1**. First, R_T has moved from the AC side to the DC side. Second, a metal-oxide semiconductor field-effect transistor (MOSFET), Q_5, has replaced the traditional mechanical relay. The reason to choose a solid-state relay is that you need to rapidly turn the relay on and off, and a mechanical relay is too slow for this purpose. Also, because the MOSFET cannot turn off the AC voltage, it is put on the DC side. The inrush current limit works the same as the traditional method. The first time that the input voltage is applied to the PSU, R_T will limit the inrush current. Once the inrush current passes, Q_5 turns on and R_T is bypassed.

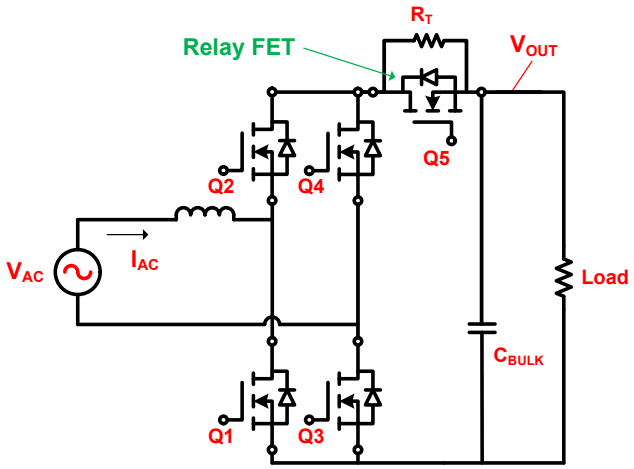

Figure 3. *The proposed re-rush current limit hardware structure.*

Figure 4 illustrates the proposed re-rush current control method. V_{AC} is the PFC input voltage, V_{OUT} is the PFC output voltage, and I_{AC} is the input current. Q_1 and Q_2 are high-frequency switches that work as either a PFC boost switch or a synchronous switch alternatively in each AC half cycle. The AC line drops for a period of 10ms and then comes back at its peak while the PFC operates at full load. This is the worst case for AC voltage dropout.

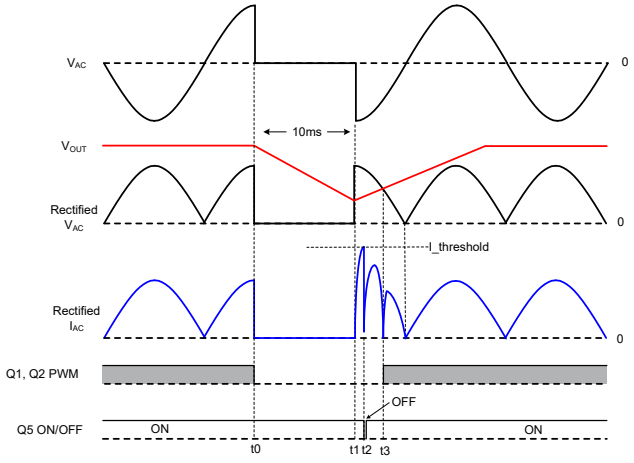

Figure 4. *The proposed AC drop and re-rush current limit control algorithm.*

This is the proposed re-rush current control method:

- At t_0: Upon detection of an AC voltage drop, Q_1 and Q_2 turn off. You must also turn off both the PFC voltage and current loops because if the voltage loop and current loop keep running, their integrators will accumulate. When the AC voltage comes back and PFC turns on, a large PWM duty cycle will then be present, resulting in a large current spike that may damage the power supply.
- Once the current loop is turned off, reset it to 0 and clear its integrator history. If you do not clear the integrator, when the AC voltage comes back and PFC turns on, PFC will turn on with the same PWM duty cycle before the AC voltage drop, and it may not be the appropriate duty cycle. For example, if the AC voltage dropout occurs at zero crossing, the PWM duty cycle is almost 100%. If the AC voltage comes back at the AC peak without a cleared current-loop integrator, an almost 100% duty cycle will occur at the AC peak and generate a large current spike, which could damage the power supply. For the voltage loop, once it turns off, freeze it to keep its internal value. The voltage loop output represents the load and is used for current loop-reference generation; therefore, you want to keep its value so that the load does not change during AC dropout.
- At t_1: The AC voltage returns. Because $V_{AC} > V_{OUT}$, a generated re-rush current will charge the bulk capacitor. Q_1 and Q_2 remain off.
- At t_2: The re-rush current exceeds a programmable threshold and trips a relay Q_5 turnoff event. The re-rush current is then limited by R_T when Q_5 is off, and its magnitude rapidly drops. Relay Q_5 only turns off for a very short period of time (for example, 10µs), then turns on again. Once Q_5 turns on, the re-rush current rises again until it exceeds the threshold. This process repeats until the re-rush current never exceeds the limit again. **Figure 5** shows the flow chart for this process.
- At t_3: $V_{AC} < V_{OUT}$. Now it is time to turn on PFC. Set the voltage loop reference equal to the instantaneous V_{OUT} value at t_3, then turn on the voltage loop. After that, gradually increase the voltage loop reference until it reaches the normal setpoint. For the current loop, first calculate a duty cycle $D = (V_{OUT} - V_{AC})/V_{OUT}$ and inject it into the current loop such that the current loop output starts from the calculated D when the current loop is on. Then turn on the current loop. Finally, turn on Q_1 and Q_2 to allow PFC normal operation.

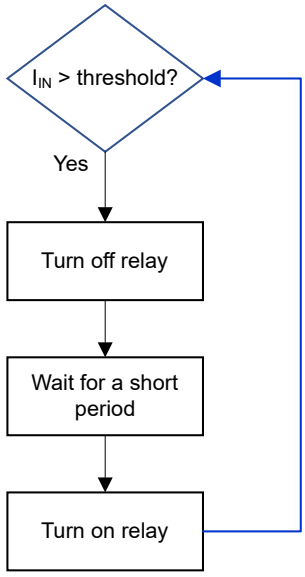

Figure 5. *Flowchart of the proposed relay on and off control algorithm.*

This process repeats until V_{OUT} exceeds V_{AC}.

Test result

I tested the proposed method on a 3.6kW totem-pole bridgeless PFC [2]. **Figure 6** shows that when the AC voltage drops out, it comes back in 10ms at its peak. Channel 1 (blue) is the PFC input current waveform (I_{IN}), while channel 2 (turquoise) is the relay on and off control signal. **Figure 7** is the zoom in at relay on and off. During the AC drop period, relay Q_5 remains on. C_{BULK} continuously delivers the stored energy to the load and V_{OUT} drops. After the AC voltage resumes, because the relay is on and $V_{AC} > V_{OUT}$, the re-rush current quickly rises. Once the re-rush current reaches

a predefined current limit threshold (40A in this example), the relay turns off, and the re-rush current reduces to a very low value because of R_T. The relay only remains off for 10µs, then turns on again. The re-rush current rises once again. This entire process allows the limiting of re-rush current within the M-CRPS specification while still supplying substantial current to rapidly charge C_{BULK}. The waveform also shows that the non-re-rush current, where $V_{AC} < V_{OUT}$, is well controlled without large current spikes.

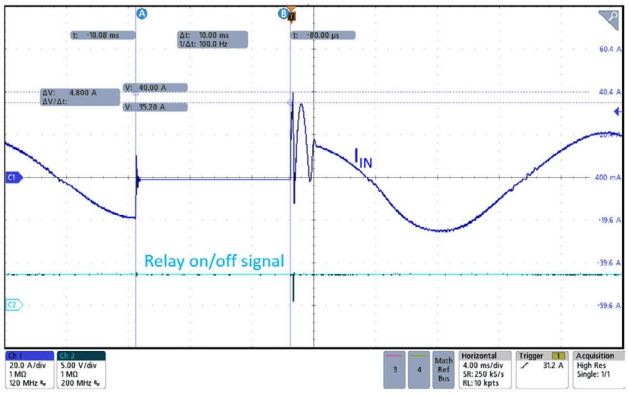

Figure 6. *Re-rush current control when the AC voltage resumes after dropout.*

Figure 7 shows the second re-rush current rising up with a limited slope, which occurs because the PFC input impedance – including the electromagnetic interference filter impedance and printed circuit board trace impedance – limit the current rising slope. The magnitude of the second re-rush current does not exceed the 40A threshold in this example; therefore, the relay turns off only once. If the second re-rush current also exceeds the threshold, the relay will turn off again.

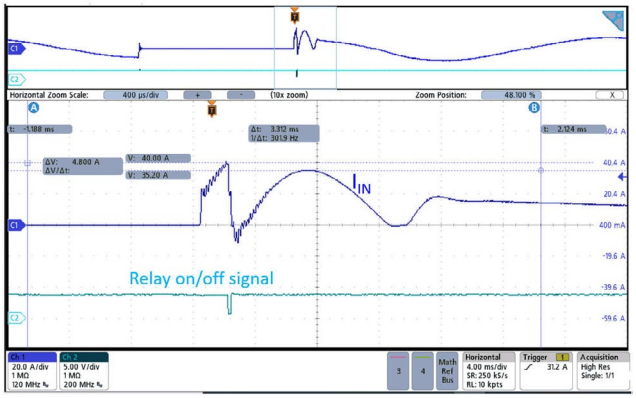

Figure 7. *Zoomed-in Figure 6 at the relay on and off instant.*

Conclusion

The power supply used in data centers requires the re-rush current when the AC voltage resumes from dropout should not exceed a limitation defined in M-CRPS specification. By replacing the traditional mechanical relay with a solid-state relay, and rapidly turn off/on the relay when re-rush current exceeds a programmable threshold, the re-rush current can be well controlled to not exceed the M-CRPS limit specification but high enough to rapidly boost V_{OUT}. Moreover, this firmware-based method leverages the existing R_T, resulting in a low-cost and very effective re-rush current control solution.

References

1. **Modular Hardware System – Common Redundant Power Supply (M-CRPS) Base Specification**. Version 1.05 RC5. Open Compute Project: Austin, Texas, Sep. 25, 2024.
2. Texas Instruments. n.d. **3.6kW Single-Phase Totem-Pole Bridgeless PFC Reference Design with E-Meter Functionality**. Texas Instruments reference design No. PMP23338. Accessed March 24, 2025.

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.

© 2025 Texas Instruments Incorporated

Transient protection design for power switches to achieve robust, reliable power-path protection

Rakesh Panguloori
Applications manager

Kunal Goel
Applications engineer

Introduction

Component integration in modern electronic systems enhances performance by increasing functionality. Most of these systems use sensitive and expensive electronic devices (field-programmable gate arrays, application-specific integrated circuits and microprocessors) that need protection.

Traditional protection solutions such as fuses, positive temperature coefficient resistors, diodes and discrete circuits (including fuses, metal-oxide semiconductor field-effect transistors and diodes) are inaccurate, slower to respond, and lack configurability and repeatability. Thus, active circuit protection using eFuse and hot-swap solutions have begun to replace discrete front-end protection circuits in many applications [1], [2].

Active-circuit-protection eFuses often need additional protection, however, to protect them from transient events. The most common transient events include hot plugging, abrupt current interruption, power surges, hard switching and reverse voltages.

Any of these transient events put electrical overstress on the device, leading to failure. In this article, we'll discuss electrical overstress (EOS) and the design process for transient protection components in a 200A eFuse enterprise server application, including placement and printed circuit board (PCB) layout considerations.

Understanding EOS

The Industry Council on ESD **defines electrical overstress** [3] as "when a maximum limit for either the voltage across, the current through, or the power dissipated in the device is exceeded and causes immediate damage or malfunction, or latent damage resulting in an unpredictable reduction of its lifetime." Of these conditions, it is the overvoltage that can open unintended current paths such as forward or reverse breakdown of diodes or oxides reaching their breakdown voltage within integrated circuits (ICs). Once an overvoltage opens an unintended current path, the resulting currents can cause damage that includes the melting of silicon; the fusing of metal interconnects; thermal damage to packaging material; and the fusing of bond-wires, leading to electrically induced physical damage (EIPD).

It is possible to relate EOS to the absolute maximum ratings specific to the voltage ratings of a device:

- Region A: The safe operating area.
- Region B: No guarantee on device functionality or parameter specifications. Although physical damage is not expected, extended operation may have reliability issues.
- Region C: Beyond the absolute maximum ratings, there is a severe degradation in device lifetime, and a risk of latent failure.
- Region D: Expected to suffer immediate physical damage.

As **Figure 1** illustrates, you should expect problems when the device operates beyond the absolute maximum rating. That's why suppressing transient overvoltages beyond the absolute maximum rating requires protection.

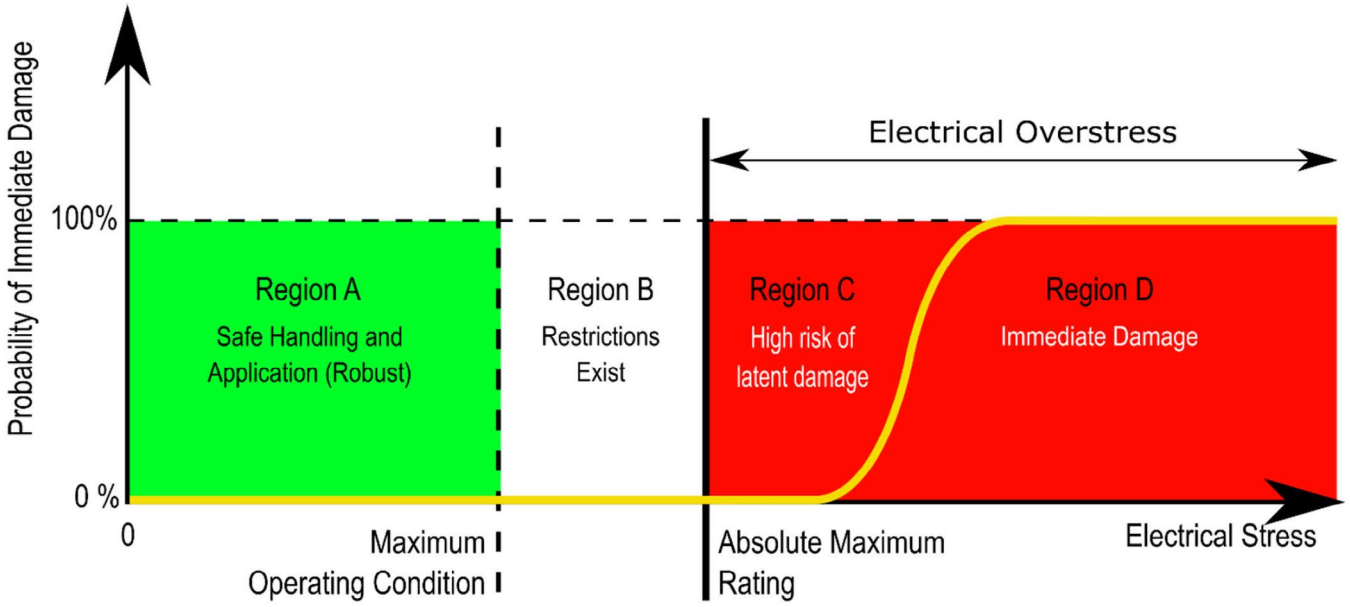

Figure 1. *Interpreting absolute maximum ratings to EOS.*

Enterprise server system example

eFuses are widely used in rack server modules at the front end for input protection and to enable hot-swapping functionality. **Figure 2** shows the typical power distribution architecture of a rack server, where the input comes from a 12V backplane and is then distributed from the eFuse to all downstream loads. The power path, which involves the backplane, PCB traces and interfacing connectors, introduces parasitic inductance (L) that then creates unintended transient voltages during fault events.

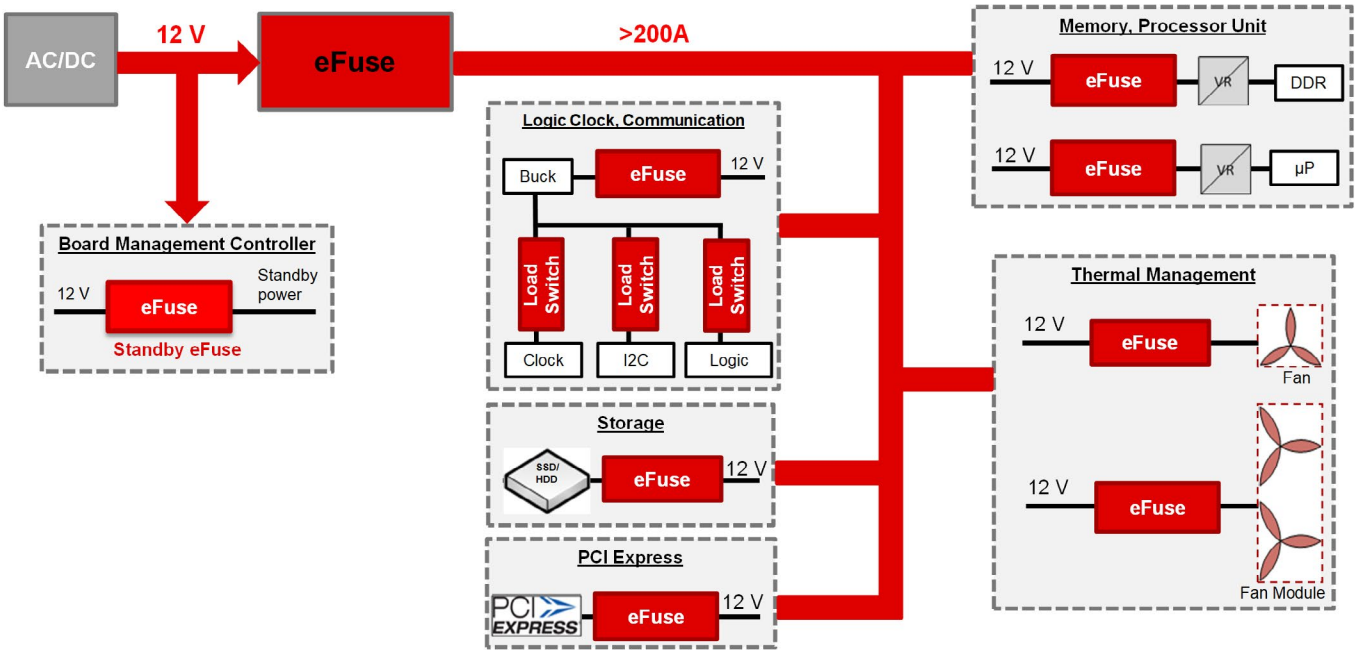

Figure 2. *A typical block diagram of the power distribution of a 12V rack server.*

Let's quantitatively analyze the impact of L on the eFuse, as illustrated in **Figure 3**. In the case of an output short circuit, the eFuse interrupts a large amount of current instantaneously from approximately 200A (overcurrent) to 0A (shutoff for protection) within 1µs, resulting in a large current transient (di/dt), as shown in **Equation 1**:

$$\mathrm{di/dt} = \left(0\mathrm{A} - 200\mathrm{A}\right)/1\mu\mathrm{s} = -2 \times 10^{8}\mathrm{A/s} \quad (1)$$

This current will be trapped as energy in the parasitic inductance and produce a surge, expressed by **Equation 2**:

$$\mathrm{VL} = \mathrm{L} \times \mathrm{di/dt} = 100\mathrm{nH} \ -2 \times 10^{8}\mathrm{A/s} = -20\mathrm{V} \quad (2)$$

That –20V surge will be in series with the 12V input power supply and will effectively create a positive voltage spike of 32V, exceeding the 20V VIN absolute maximum rating of the TPS25984B eFuse from Texas Instruments (TI). Similarly, the output inductance creates a negative voltage spike on the output.

To prevent this, a transient voltage suppressor (TVS) diode will clamp the voltages on the positive side, while a low-forward-voltage freewheeling Schottky diode will clamp the voltages on the negative side. Careful selection of these components is necessary to ensure reliable system protection.

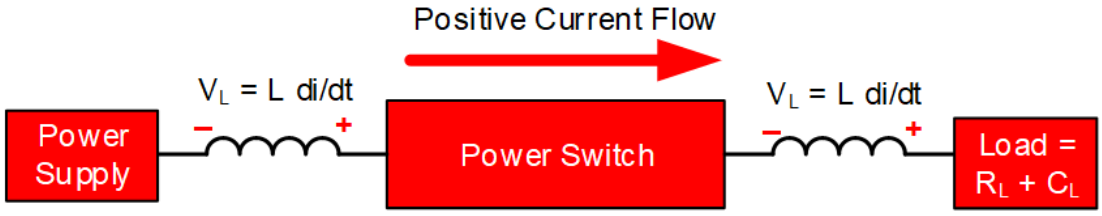

Figure 3. *Inductive kickback voltages caused by an abrupt interruption of fault current in a power switch.*

TVS diode selection

A TVS diode is designed to protect electronic components from voltage spikes. The TVS diode begins to operate once the voltage on the diode exceeds the avalanche breakdown potential. **Figure 4** is a graph of a TVS diode's current-voltage curve.

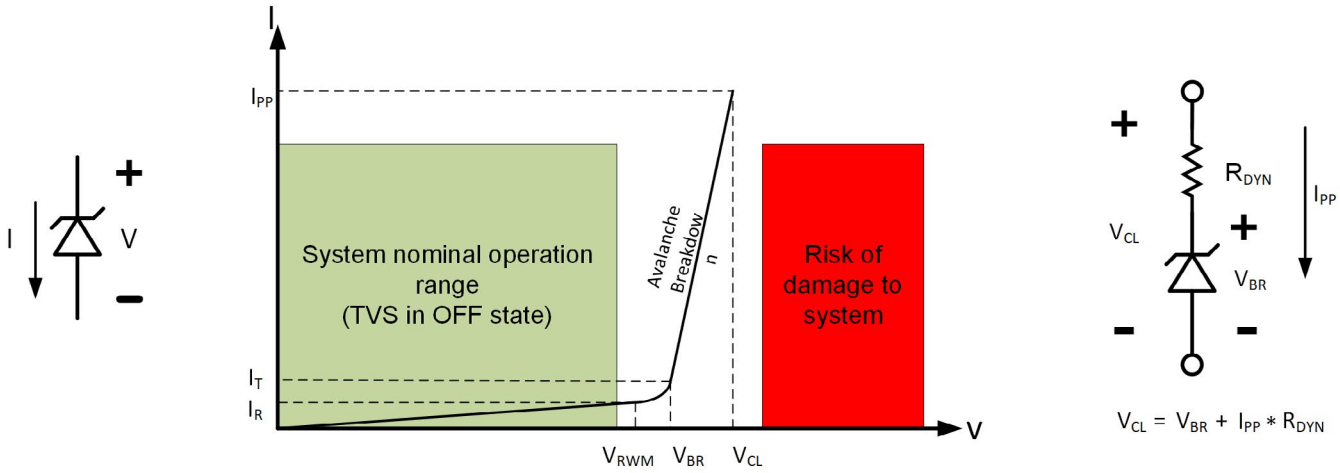

Figure 4. *TVS diode characteristics.*

As you can see in **Figure 4**, the final clamping voltage is a function of the current, which the TVS diode shunts, and the dynamic resistance (R_d) of the TVS diode. Again, the dynamic resistance is a function of diode package size and the time duration (t_P) that the TVS diode shunts the current.

For example, the SMAJ diode (13.52mm^2) has a higher R_d than the SMBJ diode (19.44mm^2), so the SMAJ diode leads to a higher clamping voltage at a given shunt current.

Use these values of R_d to calculate the clamping voltage, these values are available from the TVS diode manufacturer's data sheet.

For $t_p \leq 20\mu s$:

$$R_{d(t_P)} = R_{D(8/20\mu s)} \tag{3}$$

For $20\mu s < t_p < 1ms$:

$$R_{d(t_P)} = \frac{R_{D(10/1{,}000\mu s)} - R_{D(8/20\mu s)}}{980} \left[t_P - 20\mu s\right] + R_{d(8/20\mu s)} \tag{4}$$

For $t_p \geq 1ms$:

$$R_{d(t_P)} = R_{D(10/1{,}000\mu s)} \tag{5}$$

This multiparameter dependency leads to a challenging iterative design process. To ease design, TI released an online tool for TVS selection [5]. **Figure 5** illustrates the design methodology as a flow chart, while **Table 1** lists the typical specifications of a rack server.

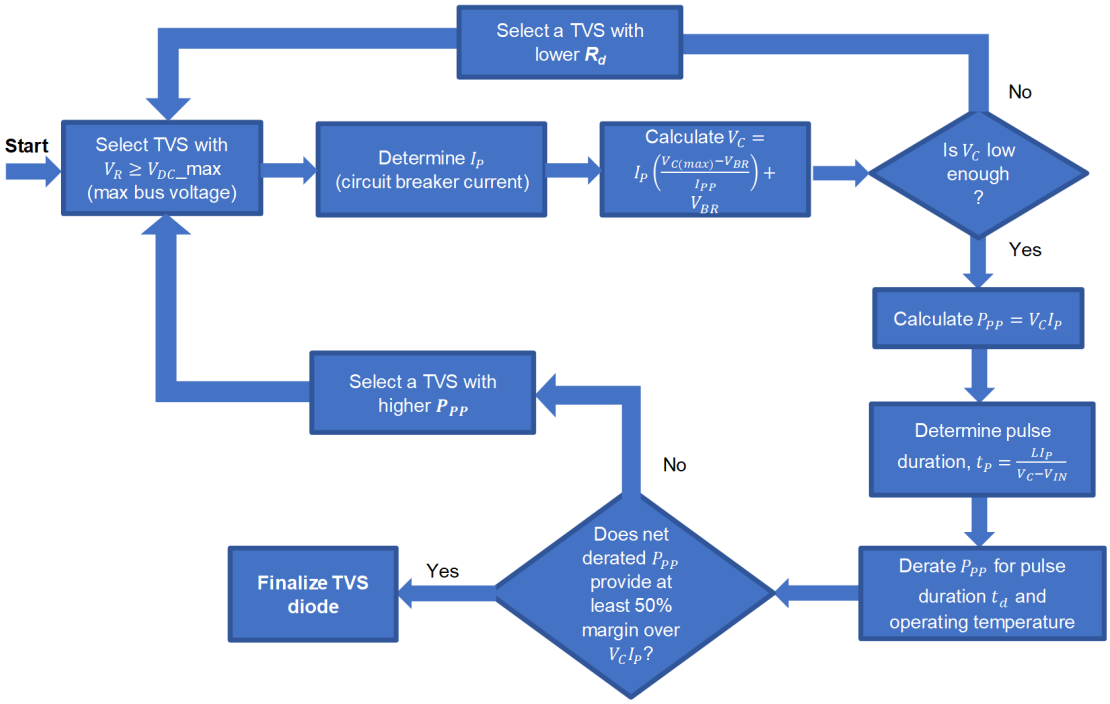

Figure 5. *Flow chart for TVS diode selection.*

Parameter	Value
Nominal operating voltage (V_{IN})	12V
Maximum operating voltage (V_{DC}_max)	13.2V
Circuit breaker current (I_P)	200A
Parasitic inductance (L)	100nH
Maximum tolerable voltage ($V_{C(max)}$)	20V
Maximum operating temperature	75°C

Table 1. *Typical system specifications.*

Design steps

When designing an input protection for the rack server, select the supporting component values for the TPS25984B eFuse from its data sheet and then follow the following design steps for TVS selection. First, select a unidirectional TVS with a reverse standoff voltage equal to or greater than V_{DC}_max. We chose the Littlefuse SMDJ12A diode [4] as a starting point. Next, determine the I_P, which is nothing but the circuit breaker current. Then calculate clamping voltage. Because R_d is a function of t_P, use **Equation 6** to find t_P:

$$t_p = \frac{LI_p}{V_{c(max)} - V_{IN}} = \frac{100nH \times 200A}{20V - 12V} = 2.5\mu s \quad (6)$$

For pulse widths below 20µs, you can approximate the dynamic resistance to that at an 8/20µs test pulse. From the SMDJ12A data sheet, our calculations were:

$$V_{BR(max)} = 14.7V \quad (7)$$

$$V_{C(max)} \text{at} I_{PP}(8/20\mu s) = 25.71V$$

$$I_{PP}(8/20\mu s) = 754A$$

Therefore:

$$R_{d(8/20\mu s)} = \frac{V_{C(max)} - V_{BR(max)}}{I_{PP}} = \frac{25.71 - 14.7}{754} = 14.6m\Omega \quad (8)$$

Now, using the R_d of 14.6mΩ, calculate the clamping voltage:

$$V_C = V_{BR(max)} + I_P \ R_d = 14.7V + 200A \times 14.6m\Omega = 17.6V \quad (9)$$

Because the clamping voltage is less than maximum tolerable voltage, $V_{C(max)}$ (the 20V absolute maximum rating of TPS25984B eFuse), you can proceed further with the SMDJ12A; otherwise, you will have to consider a TVS diode with a lower R_d, or parallel TVS diodes.

Calculate the peak power using:

$$P_{PP} = V_C \ I_P = 17.6V \times 200A = 3.52kW \quad (10)$$

Because the SMDJ12A supports a peak power of 60kW for 2.5µs (see **Figure 6**), you can proceed further.

Now, derate the power rating with temperature by using **Figure 6**. The maximum power support at 75° is:

$$P_{PP} \times \text{derating\_factor} = 0.8 \times 60kW = 48kW \quad (11)$$

Because 48kW > 3.52kW and V_C < 20V, the SMDJ12A is a good choice for this application.

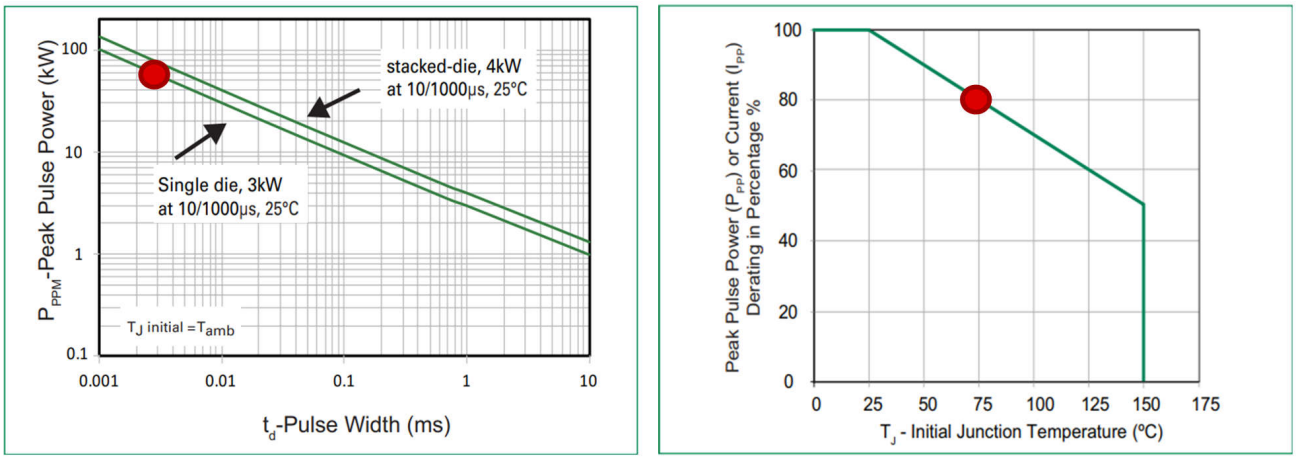

Figure 6. *Peak pulse power rating (left) and peak pulse power derating curve (right).*

Figure 7 shows the clamping performance of the SMDJ12A on the TPS25984B system.

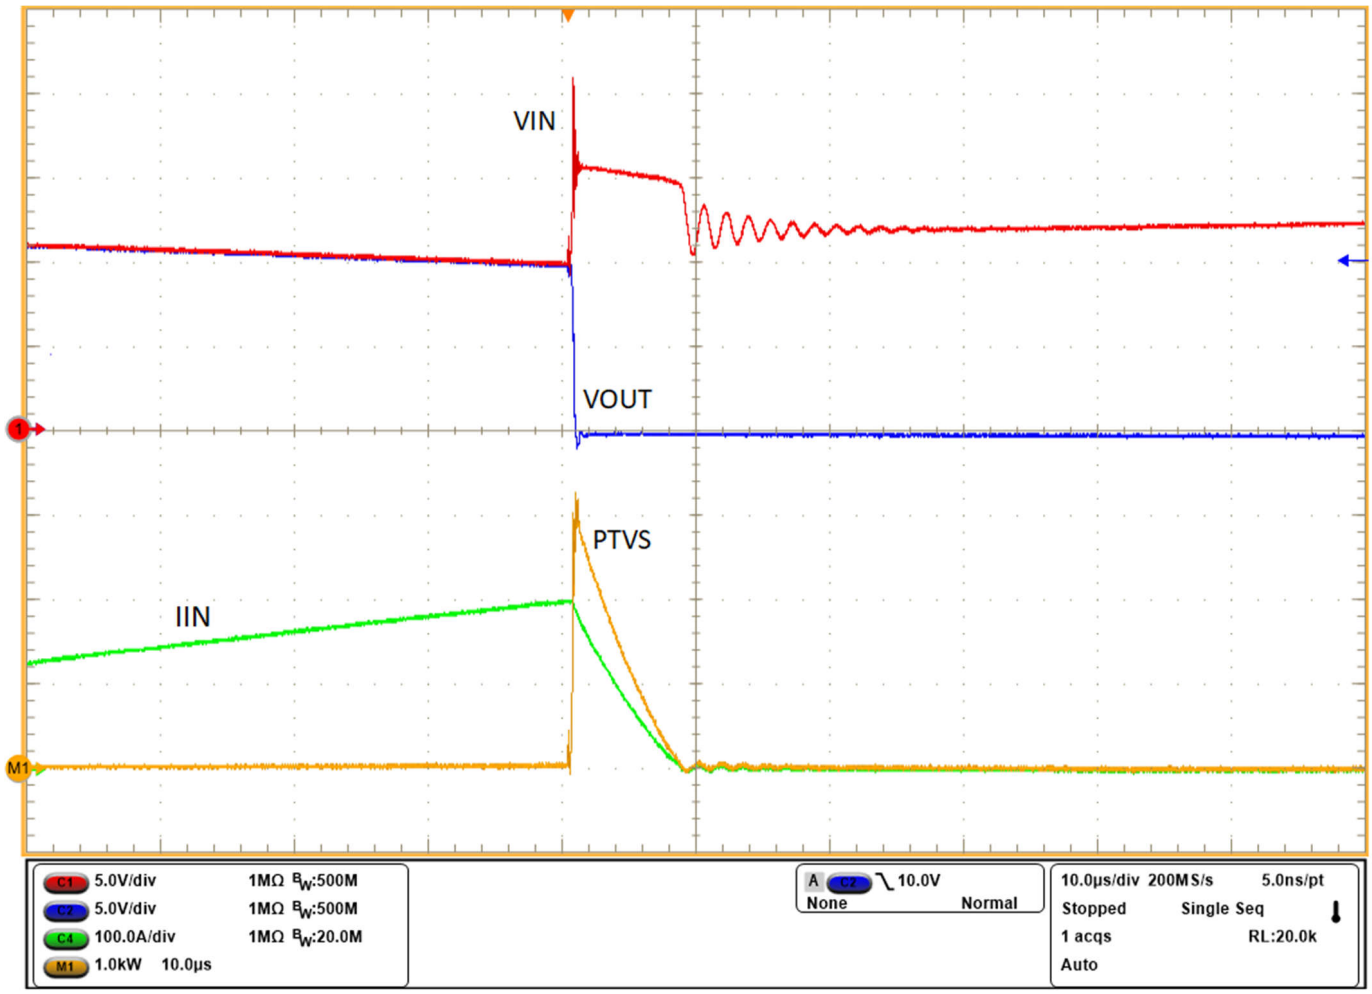

Figure 7. *Transient protection with the SMDJ12A diode at the input of the TPS25984B eFuse.*

Output Schottky diode selection

Figure 8 shows the sequence of events that could happen internal to the IC when taking the OUT pin below ground. The parasitic PN junction diode starts conducting, which injects free electrons into the substrate. These free electrons interfere with other control units that could reset the IC or cause a latch-up event. A large current conduction through the parasitic PN junction diode could cause EOS and lead to EIPD.

It is possible to prevent these issues by either reducing the peak negative voltage at the OUT pin or by limiting the current through the OUT pin. Adding an output capacitor close to the OUT pin will absorb some of the energy from the negative voltage spike and control the slew rate to limit the peak negative voltage. Adding a low-forward-voltage Schottky diode at the OUT pin provides an alternate path for current and limits the current through the IC.

Effective clamping requires a combination of capacitors and Schottky diodes. While a higher-output capacitor is helpful, use these guidelines when selecting a Schottky diode:

- The DC blocking voltage must be more than the maximum input operating voltage.
- The non-repetitive peak forward surge current of the selected diode must be higher than the I_P.
- The forward voltage drop at the I_P must be within the absolute maximum rating of the OUT pin (which is –1V for the TPS25984B).

We used two SBR10U45SP5 [6] diodes from Diodes Incorporated in parallel in this application.

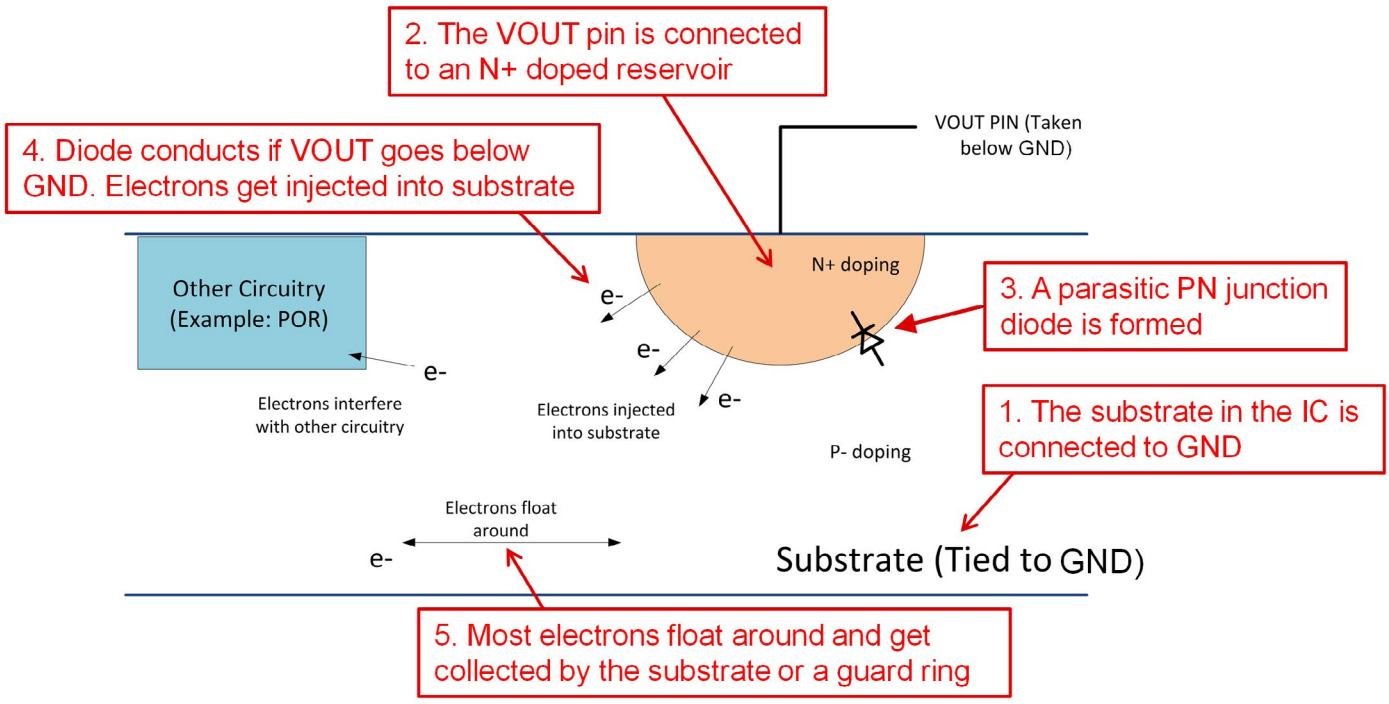

Figure 8. *Graphical illustration showing the consequences inside the IC when taking the output below ground.*

Figure 9 shows the output clamping performance with and without Schottky diodes in the TPS25984B solution.

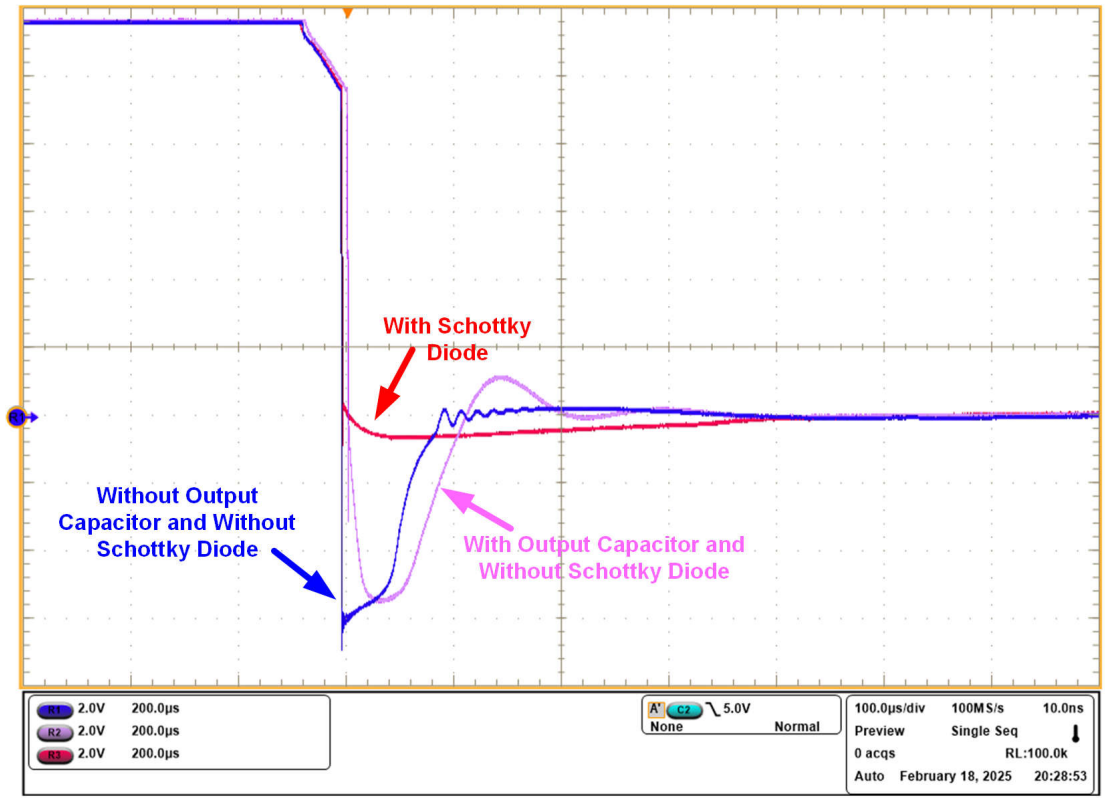

Figure 9. *Transient protection at the output of the eFuse.*

When dealing with high-current hot-swap solutions, secondary protection (shown in **Figure 10**) can minimize the Schottky diode requirement at the output. As you can see, D1 will absorb most of the energy from the negative voltage transient. Adding a small-value resistor (R1) such as 47Ω and a diode (D2) such as an SS13 will significantly limit the remaining energy.

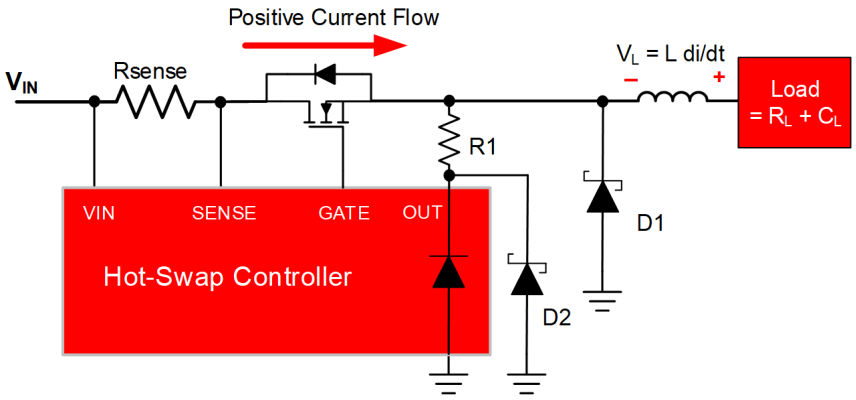

Figure 10. *Secondary protection in high-current hot-swap solutions.*

Placement and PCB layout considerations

You must place protection devices such as TVS diodes, decoupling capacitors and Schottky diodes physically close to the device they are intended to protect. The distributed inductance limits the bandwidth of the effectiveness of the shunt elements, such as decoupling capacitors and TVS diodes. It restricts the surge current flow and also leads to larger transient voltage spikes during clamping, as shown in **Figure 11**. Therefore, the layout should be such that these shunt elements have the least series impedance. When routing, use short traces and multiple vias to reduce inductance.

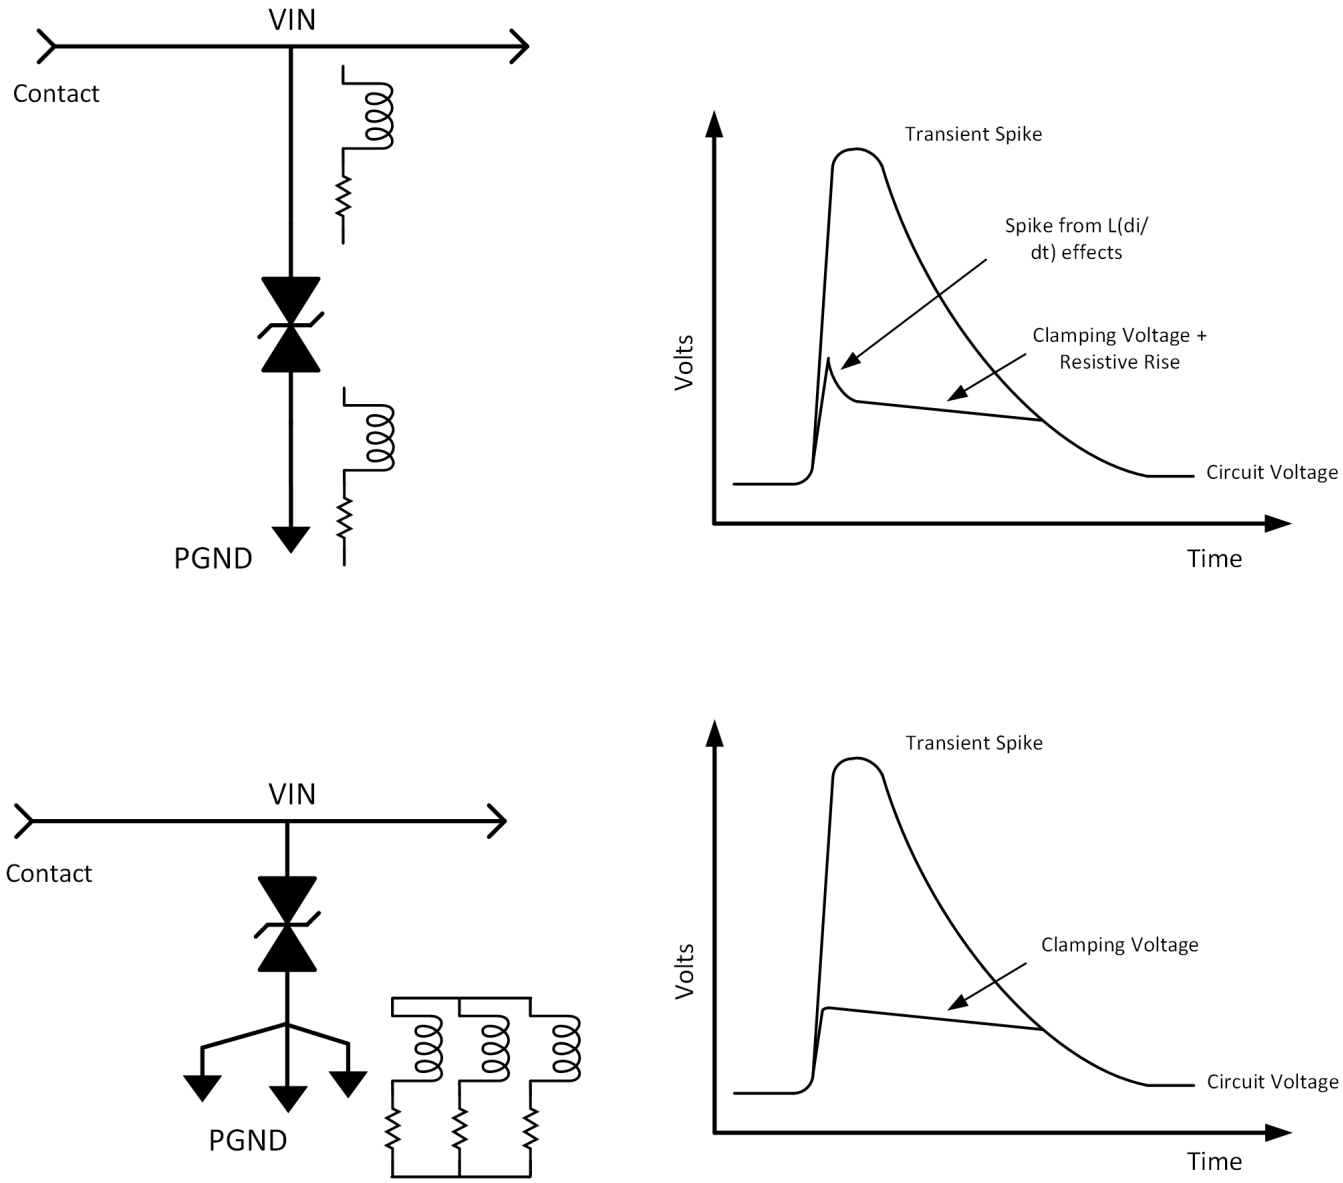

Figure 11. *Impact of TVS clamping performance with respect to PCB layout.*

Conclusion

Discrete front-end protection circuits are being replaced with active circuit protection devices such as eFuses to improve performance. However, an eFuse often needs transient protection to prevent its absolute maximum spec violation. The component selection guidelines and layout considerations discussed in this article can help you design a solution to ensure reliable power-path protection.

References

1. Panguloori, Rakesh. 2016. "**Basics of eFuses**." Texas Instruments application report, literature No. SLVA862A, December 2016.
2. Rogachev, Artem. 2014. "**Robust Hot Swap Design**." Texas Instruments application report, literature No. SLVA673A, April 2014.
3. The Industry Council on ESD. "**White Paper 4: Understanding Electrical Overstress – EOS**." August 2016.
4. "**SMDJ12A TVS 3KW 12V 5%UNI DO-214AB TR13 RoHS**." Littlefuse data sheet. 2025.
5. Texas Instruments. n.d. **TVS diode recommendation tool**. June 26, 2023.
6. "**SBR10U45SP5 10A SBR super barrier rectifier**." Diodes Inc. data sheet, literature No. DS31371, January 2024.

Related Websites

- **TPS25984B**
- **LM5066I**

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.

© 2025 Texas Instruments Incorporated

Optimizing signal chain cost and accuracy for isolated current sensing in energy metering

Dr. Ralph Oberhuber
Precision analog-to-digital converters

Introduction

As we move toward a more sustainable energy grid with sources such as solar and wind power, the need arises for energy metering devices to obtain detailed insights into individual power consumption in order to identify areas for improvement, optimize usage, and reduce costs. The types of electronic end equipment that require an energy metering subsystem include smart electrical meters, electric vehicles (EVs) charging stations, power supplies and power distribution units, smart appliances, street lightning, and building automation components. The sheer number of these products drives the need for the lowest possible cost for energy metering solutions, while regional metering standards such as American National Standards Institute C12 (in the U.S.) or Measuring Instruments Directive (in Europe) impose strict accuracy and safety requirements [1], [2].

Figure 1 illustrates a typical signal chain inside an energy metering application, showing only one phase for simplicity. An analog-to-digital converter (ADC) simultaneously measures and digitizes voltage and current for each phase. Subsequently, digital signal processing extracts metrology parameters such as active and reactive power and energy, line-to-line voltages, fundamental power and energy, and harmonics [3].

The fundamental buildings blocks of the signal chain are:

- Line voltage sensing front end (A in Figure 1).
- Current measurement sensor (B).
- Front-end and signal conditioner between the current sensor and the ADC (C).
- ADC (D).
- Digital signal processing hardware (E).
- Galvanic isolation (F).

While the line voltage sensing front end is in most cases implemented using simple resistor-dividers [3], various options exist for the choice of every other building block. For each of these signal-chain components, performance, size and cost trade-offs exist. This article focuses on the performance vs. cost trade-offs for the current measurement sensor and signal conditioning, as well as the ADC.

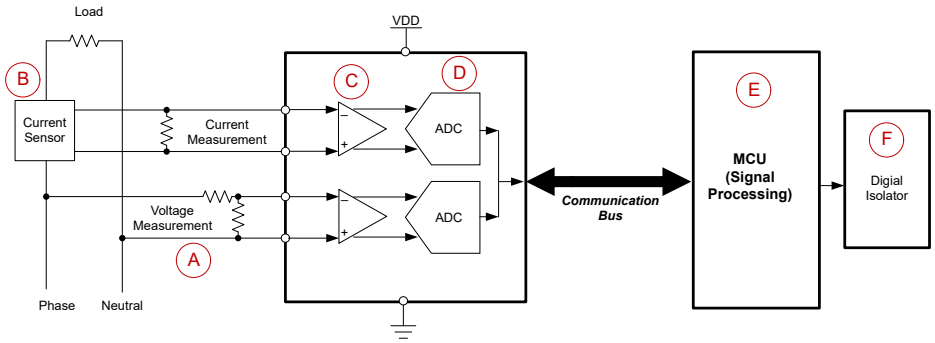

Figure 1. *Energy metering subsystem signal chain.*

Comparing current-sensor elements used for energy metering

Table 1 summarizes the performance benefits and challenges vs. cost for three current-sensing technologies used in energy metering applications. Current transformers are the most popular sensors given their wide dynamic range, durability and low insertion impedance (noninvasive current measurement) [4]. Their cost is most likely the highest of all current-sensing technologies, however. Shunt resistors (shunts) are very attractive given their magnetic immunity, smaller size and reduced cost, but lack isolation and provide less accuracy at higher currents because of thermal self-heating [5].

Rogowski coils are an interesting alternative to the other two sensors and are the lowest-cost option, especially when considering printed circuit board (PCB) coils vs. bulk Rogowski coils [6].

Sensor Type	Current Transformer	Rogowski Coil	Shunt
Block diagram	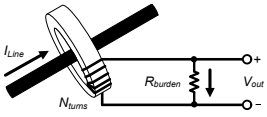	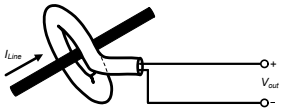	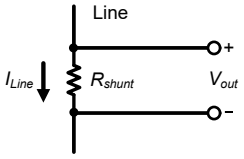
Transfer function	$V_{out} = \frac{I_{line} R_{burden}}{N_{turns}}$	$V_{out} \propto \frac{dI_{line}}{dt}$	$V_{out} = I_{line} R_{shunt}$
Benefits	• Provides isolation • High linearity • High durability • High accuracy	• Provides isolation • Saturation not possible • Small form factor (PCB) • High linearity • Fast response time • Low power-loss	• Not isolated • Antimagnetic • Small form factor
Challenges	• Saturation possible • Power loss in burden resistors • Phase calibration required • Size and weight	• Integration required • Cannot measure DC • PCB coils can have lower sensitivity	• Resistance may vary when self-heated • Less reliability for overload signals
Cost	≥$0.21 (no shield) ≥$0.26 (with shield)	<$0.10 for some PCB coils ≥$0.10 for bulk coils	≥$0.10 (manganese)

Table 1. *Current-sensor comparison.*

Because of their low cost and flexibility of installation, PCB Rogowski coils are attractive for low-cost energy metering applications [7], [8]. Let's analyze the benefits and challenges of a PCB Rogowski-based metering design and how to optimize the signal chain for the lowest possible cost while complying with regional metering standards.

Sensitivity of PCB Rogowski coils used in energy metering vs. ADC noise performance

PCB Rogowski coil sensitivity is typically specified in microvolts per ampere and depends on the geometry (number of turns, coil dimension); core material (if any); current frequency; and environmental factors (temperature, humidity, external magnetic fields) [9]. Common sensitivities range from tens to a few hundred microvolts per ampere [9].

It's a common accuracy requirement for residential electricity meters to measure a 250mA root-mean-square (RMS) phase current with 2% accuracy [1]. For example, with a 200µV/A Rogowski coil, the signal at the input of the ADC is only 200µV/A × 0.250A = 50µV for this phase current. The required ADC performance (that is, the noise which determines the effective resolution) to measure this signal with 2% accuracy is as low as 0.02 × 200µV/A × 0.250A = 1µV, as defined by **Equation 1**:

$$V_{nADC} = tol \times k \times I_{phase-rms} \quad (1)$$

where, V_{nADC} is the required noise level of the ADC, tol is the specified measurement accuracy in percentage for a given phase current, $I_{phase-rms}$ in amperes, and k is the sensitivity constant of the Rogowski coil in microvolts per ampere.

Therefore, in this example, the total noise of the ADC (the quantization noise plus white noise) needs to be lower than 1µV.

Comparing the 1µV ADC noise requirement to the specification of a precision ADC such as the Texas Instruments (TI) **ADS131M08** [10], it is clear that achieving the intended performance level may require additional averaging of the ADC samples. **Table 2** illustrates this averaging, and also shows the total ADC noise in RMS microvolts for various gain settings and data rates as defined by the oversampling ratio (OSR). With gain of 1 and a sample rate of 4kSPS (OSR = 1,024), the ADC noise is approximately 5µVrms. As noise improves at a factor of √2 for doubling the time of averaging, achieving the requirement of <1µV of ADC noise requires a time period of ≥16ms. This is acceptable for most energy metering systems, which commonly require an update rate of 20ms [1]. This type of averaging may be practically implemented with a combination of ADC internal oversampling using the delta-sigma ADC's internal oversampling ratio (OSR) feature and external post-averaging.

Another option suggested by **Table 2** is to select a higher gain for the programmable gain amplifier (PGA) internal to the ADC, as it reduces the noise referred to the input [10]. Alternatively, you could precondition the signal with an external gain stage before it arrives at the ADC. An external gain stage increases the cost of the signal chain significantly, however.

Averaging Time (ms)	OSR	Data Rate (kSPS)	Noise (µVrms), Gain 1	Noise (µVrms), Gain 128
16	65,392	0.0625	0.95	0.07
8	32,696	0.125	1.34	0.10
4	16,384	0.25	1.90	0.42
2	8,192	0.5	2.39	0.57
1	4,096	1	3.38	0.77
0.5	2,048	2	4.25	1.00
0.25	1,024	4	5.35	1.20
0.125	512	8	7.56	1.69
0.0625	256	16	10.68	2.40

Table 2. *ADC noise performance vs. speed, averaging time and OSR.*

Sensitivity analysis of an ADC signal chain for Rogowski coil-based current sensors

The primary concern for systems using a low-cost PCB Rogowski current sensor is that the signal amplitude at the sensor output is typically very small – in most cases only a few microvolts. You must design the signal chain carefully to meet the accuracy requirements driven by the metering standard. The signal conditioning of such a small signal must include significant differential gain, either by selecting a high-resolution ADC with internal gain or by cascading an external gain stage between the sensor and the ADC. Adding an external gain stage is

often detrimental, as it increases total cost; therefore, it makes more sense to quantify which solutions require external gain stages and when you can avoid them.

Table 3 introduces three different Rogowski coils in order to analyze the effectiveness of an external gain stage:

- Coil A is a PCB Rogowski coil based on the High Accuracy AC Current Measurement Reference Design Using PCB Rogowski Coil Sensor [**11**], with a sensitivity of approximately 20µV/A.
- Coil B is another proprietary Rogowski coil with a sensitivity of approximately 100µV/A.
- Coil C is a commercially available bulk Rogowski coil (Pulse PA3209NL) [**12**] with a sensitivity of approximately 500µV/A.

Number	Type of Coil	Source	Sensitivity (µV/A)	Cost
A	PCB	TI reference design	20	Low
B	PCB	Proprietary	100	Medium
C	Bulk	Pulse PA3209NL	500	High

Table 3. *Rogowski coils characterized during signal-chain analysis.*

Figure 2 illustrates the measurement setup for the sensitivity analysis. The output of each individual Rogowski coil, shown in **Table 3**, connects to a signal-conditioning interface board where you can select or bypass the TI INA188-based gain stage [**13**] with four jumpers. The gain-defining resistor R_G (see **Figure 2**) is 390Ω, yielding an optional external gain of 128.

The output of the instrumentation amplifier (INA) interface board connects to the phase 1 current input of the Three-Phase Current Transformer E-Meter Reference Design with Standalone ADC [**3**]. This reference design does include burden resistors R37 and R38, which are required only when connecting to a current transformer and were physically removed for this analysis. The ADC on the e-meter reference design is the TI **ADS131M08**, a high-precision, eight-channel, simultaneous-sampling delta-sigma ADC with internal gain options ranging from 1 to 128.

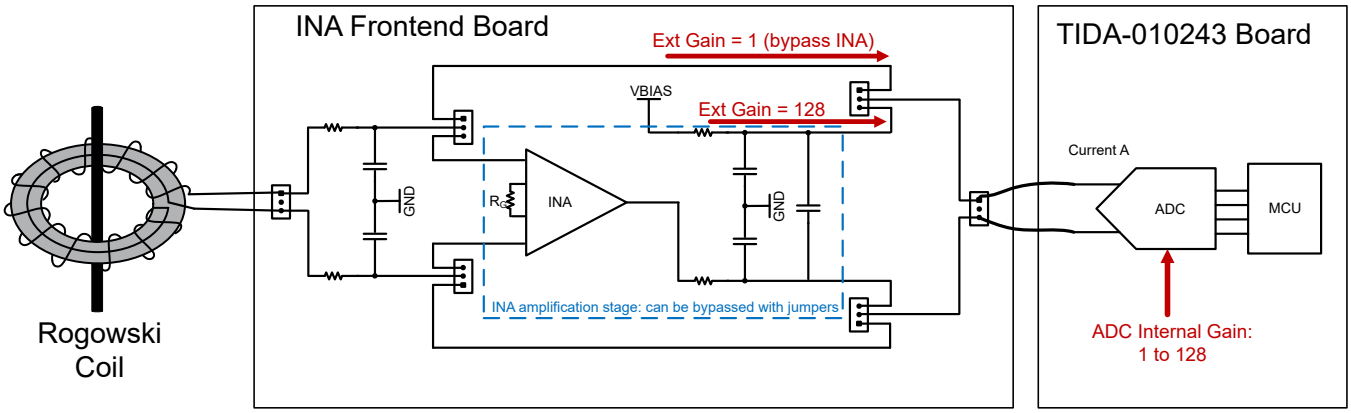

Figure 2. *Measurement setup for the signal-chain analysis.*

Figure 3 and **Figure 4** show the measured current accuracy for a 50Hz line current from 100mA to 10A using MTE's PTS3.3C source generator and reference meter. Using the same test procedure as in [3], samples of current and energy are averaged over a 20ms time period. We implement the Rogowoski signal integration in the digital domain following the procedure outlined in [11]. The alternative would be analog active integration as shown in [14], however this technique is ignored for our analysis as the two methods usually yield similar results.

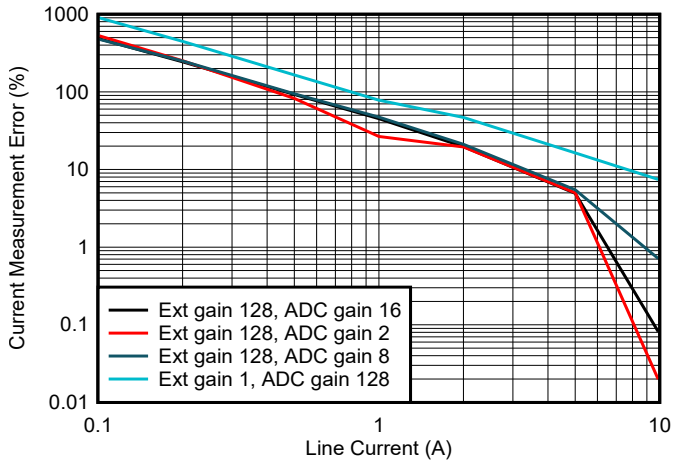

Figure 3. *Measured current accuracy for the 20µV/A coil with different gain settings.*

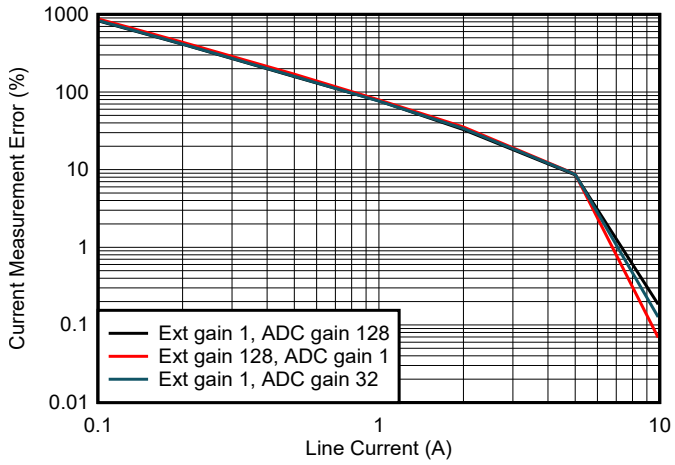

Figure 4. *Measured current accuracy for the 100µV/A coil with different gain settings.*

For a PCB coil with very low sensitivity (for example, 20µV/A), there is significant improvement when using an external gain of 128 by cascading the INA stage (see **Figure 3**). The internal PGA gain (even with a gain of 128) of the ADS131M08 alone does not sufficiently raise the small input signal above the quantization noise level, as explained previously.

When using PCB coils with a sensitivity of ≥100µV/A (see **Figure 4**), the selection of internal vs. external gain results in a comparable error, indicating that the sensor output amplitude is now well above the quantization noise level for the relevant phase current range. The absolute value of the resulting error is higher than acceptable for some revenue-grade energy metering systems, which target ≤0.5% accuracy. This increased error is the result of a simplified calibration procedure applied in this setup: a single-point (gain) calibration. In a typical metering design, applying up to three calibration steps (offset calibration, gain calibration and phase calibration) can further reduce the absolute error.

Figure 5 and **Figure 6** illustrate the dependency of the measurement error on the sensitivity of the Rogowski coil for the three different coils listed in **Table 3**.

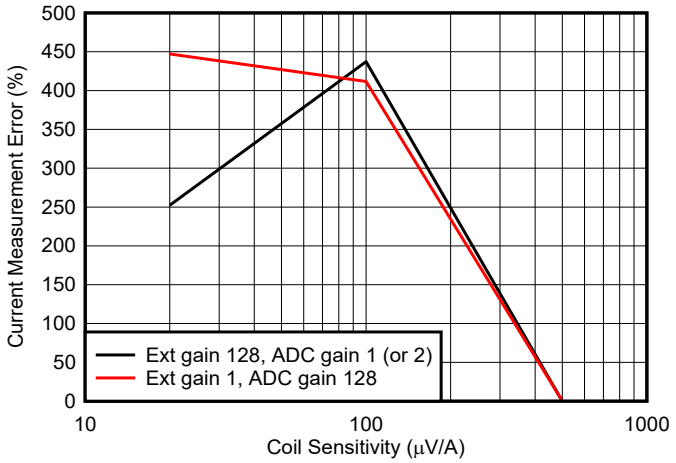

Figure 5. *Measured current accuracy for three coils at a line current of 200mA.*

Both in the case of small phase current (200mA, **Figure 5**) and mid-level phase current (5A, **Figure 6**), the 20µV/A Rogowski coil achieves a significant improvement (a smaller error) by employing the external gain stage. As expected, all errors scale to smaller values when detecting the larger line-current value (5A, **Figure 6**). For

the 100µV/A and 500µV/A Rogowski coils, applying an external gain of 128 vs. using the internal ADC gain results in comparable accuracy.

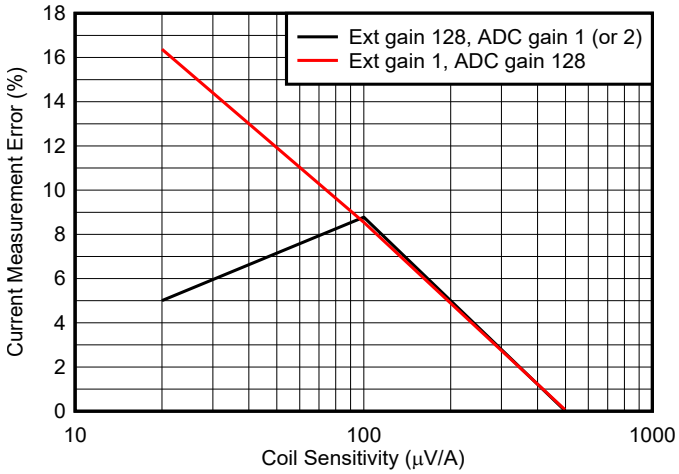

Figure 6. *Measured current accuracy for three coils at a line current of 5A.*

Conclusion

For ≥100µV/A Rogowski coils, the external gain stage is not necessary, resulting in a reduced cost for an energy metering signal-chain solution. For <100µV/A Rogowski coils, external gain may be needed to meet residential electricity metering accuracy when using the TI **ADS131M08** or a comparable ADC. Or alternatively, a lower-noise, higher cost ADC may be considered to avoid the additional circuitry.

References

1. Kelechava, Brad. **ANSI C12.20-2015 – Electricity Meters – 0.1, 0.2, and 0.5 Accuracy Classes**. American National Standards Institute blog, May 8, 2017.
2. Directive 2004/22/EC of the European Parliament and of the Council of 31 March 2004 on Measuring Instruments (**https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CONSLEG:2004L0022:20130101:en:PDF**). Document 32004L0022. European Union: Brussels, Belgium. March 31, 2004.
3. Texas Instruments. n.d. **Three-Phase Current Transformer E-Meter Reference Design with Standalone ADC**. Texas Instruments reference design No. TIDA-010243. Accessed March 14, 2025.
4. Blue Jay. n.d. **What Are Advantages and Disadvantages of Current Transformer**? Accessed Jan. 21, 2025.
5. Maniar, Krunal. **Comparing Shunt- and Hall-Based Isolated Current-Sensing Solutions in HEV/EV**. Texas Instruments application brief, literature No. SBAA293C, June 2018.
6. Institute of Electrical and Electronics Engineers (IEEE) Power System Relaying Committee (PSRC). **Practical Aspects of Rogowski Coil Applications to Relaying**. IEEE PSRC Special Report, September 2010.
7. Aim Dynamics. **Rogowski Coils vs. Current Transformers**. Aim Dynamics blog post, July 24, 2023.
8. Brooks, Chris. **Why Choose Rogowski Coils**? Setra blog post, Feb.13, 2017.
9. Salomon, Victor. **Two-Phase Rogowski Coil Based Electricity Meter Analog Front-End Circuit**. Texas Instruments Analog Engineer's Circuit, literature No. SBAA385A, July 2020.
10. Texas Instruments. n.d. **ADS131M08 24-bit, 32-kSPS, 8-channel, simultaneous-sampling, delta-sigma ADC**. Accessed Jan. 21, 2025.
11. Texas Instruments. n.d. **High Accuracy AC Current Measurement Reference Design Using PCB Rogowski Coil Sensor**. Texas Instruments reference design No. TIDA-01063. Accessed March 14, 2025.
12. Pulse Electronics. n.d. **PA3209NL** Sidewinder-Current Sensor. Accessed Mar. 19, 2025
13. Texas Instruments. n.d. **INA188 36-V, Zero-Drift, Rail-to-Rail-Out Instrumentation Amplifier**. Accessed Jan. 21, 2025.
14. Texas Instruments n.d. **Active Integrator for Rogowski Coil Reference Design With Improved Accuracy for Relays and Breakers**, Texas Instruments reference design No. TIDA-00777. Accessed March 19, 2025.

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.

© 2025 Texas Instruments Incorporated

Sampling around Nyquist holes in high-speed converters

Luke Allen
Application Engineer
High-Speed Converters

Rob Reeder
Application Engineer
High-Speed Converters

Chase Wood
Application Engineer
High-Speed Converters

Frequency planning is a paramount part of any frequency-based application. Making sure that the frequency or frequencies of interest are valid, within band, and not reduced to the point where you lose the spurious dynamic range should be part of any frequency development strategy. In this paper, we will cover high-speed analog-to-digital converter (ADC) Nyquist rules as they apply to super Nyquist sampling. We will also cover how to use decimation to aid in your frequency planning and how to guard band against falling into a frequency "hole" during the design and development stage.

Nyquist rules

Super Nyquist sampling, intermediate frequency (IF) sampling and subsampling are popular in many frequency-based applications that employ software-defined radio (SDR) or radar-like receiver architectures (see **Figure 1**).

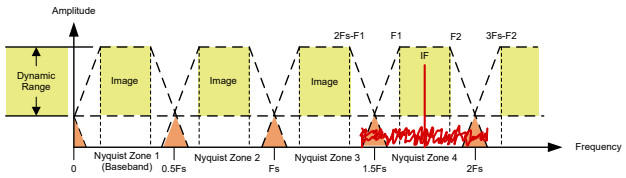

Figure 1. *Example of a super Nyquist sample vs. sampling at baseband (first Nyquist).*

There are two main reasons to plan frequencies outside of baseband (first Nyquist). The first reason is to gain the relaxation constraints put on the antialiasing filter design (AAF) (see **Figure 2**). Initially, the general filter rolloff needs to be much steeper when designing a baseband filter versus a filter design for a higher Nyquist zone. A steeper filter rolloff leads to a more complex filter where passive components become cumbersome. It's simple physics; you cannot purchase a 100µH inductor in a 0201 size. Therefore, when employing a higher Nyquist zone and possibly a higher sampling rate, the trade-offs and requirements for the rolloff in the stopband region are more relaxed, resulting in fewer components and smaller component sizes.

The second reason to use the high-frequency subsampling technique is to relax the radio-frequency (RF) receiver signal chain in front of the ADC. Assuming that the ADC can support the bandwidth requirements beyond the first Nyquist, which is almost always the case, relaxing the receiver signal chain could eliminate one or even two mix-down stages in the RF signal chain, resulting in even fewer components, less noise and less complexity.

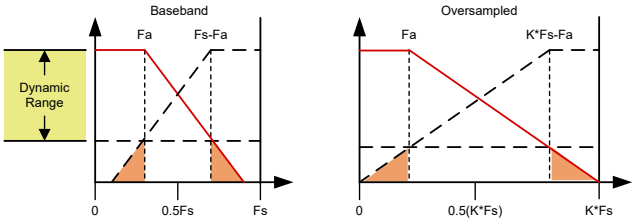

Figure 2. *Dynamic range vs. AAF stop-band attenuation.*

For example, **Figure 3** shows the Texas Instruments (TI) ADC3669 sampling an intermediate frequency of 800MHz relative to a 500 MSPS sampling frequency (Fs). Essentially, the signal is in the fourth Nyquist zone. The image or alias of the interest frequency reflects back to the first Nyquist zone appearing as a 200MHz signal. Most fast Fourier transform (FFT) analyzers, such as **High-Speed Data Converter Pro**, only plot an FFT of the first Nyquist zone, or 0Fs to 0.5Fs. Therefore, if the frequency of interest is above 0.5Fs, an image reflects down to the first Nyquist zone or baseband. This can make things confusing if spurious tones are in the band of interest as well.

So how does an ADC sample above 0.5Fs and still hold true to the Nyquist criteria? The Nyquist rule states that a signal must be sampled at a rate equal to or greater than twice its bandwidth in order to preserve all of the signal's information (see **Equation 1**):

$$\mathrm{Fs} > 2 > \mathrm{FBW} \qquad (1)$$

where Fs is the sample frequency and FBW is the maximum frequency of interest.

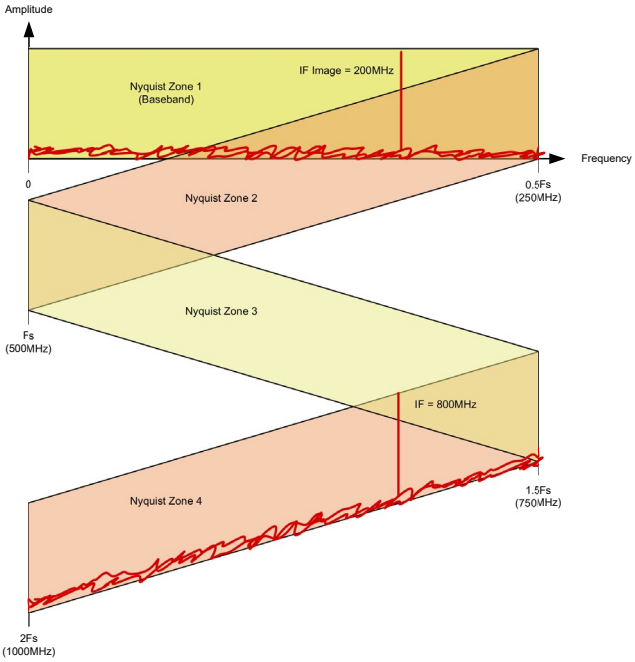

Figure 3. *ADC3669 example, where Fs = 500MSPS and the intermediate frequency = 800MHz.*

The key to holding the Nyquist rule to be true is the location of the frequencies of interest. As long as the signals do not overlap and stay within a single Nyquist zone, the Nyquist criteria still holds true. The only thing that has changed is the location of the first Nyquist zone to a higher one. IF sampling is becoming very popular because of these trade-offs.

What is process gain?

Earlier forms of high-speed signal chain line ups involved ADCs that used to consume watts of power per single channel, and FPGAs (field-programmable gate array) which captured, filtered and processed all converter data into some useful format. Most designers would use an approach called *process gain*. This approach not only helped in frequency planning by eliminating unwanted spurious and noise; it also enabled the ability to "gain" dynamic range in terms of the signal-to-noise ratio (SNR) by limiting the bandwidth processed within the Nyquist zone. Adding the process gain correction factor to the standard SNR equation results in **Equation 2**:

$$\mathrm{SNR} = 6.02 \times \mathrm{N} + 1.76\mathrm{dB} + 10 \times \log_{10}\left(\frac{\mathrm{Fs}}{(2 \times \mathrm{BW})}\right) \qquad (2)$$

where N is the number of ADC bits, Fs is the ADC sampling frequency and BW is the bandwidth of interest within the Nyquist zone.

With smaller process nodes deployed into both ADC and digital-to-analog converter technology, much of the standard FPGA digital features now reside within the ADC. Some examples include digital downconverters (DDCs), numerically controlled oscillators (NCOs) and frequency hopping. These features significantly help offload FPGA processing, making it possible to use its internal resources elsewhere.

Why frequency plan?

Proper frequency planning is an important step when designing systems using ADCs. Frequency planning ensures efficient utilization of the ADC's dynamic range and minimizes unwanted spurious signals, which are crucial for high-performance applications such as SDR systems or high-density RF signal chains.

One essential aspect of frequency planning is optimizing the ADC's dynamic range. Each additional signal in the analog domain consumes part of the ADC's available input signal power budget, reducing the total dynamic range capability. Proper frequency planning ensures that the ADC fully exercises its capabilities by strategically placing input frequencies to maximize the usable dynamic range. This approach involves considering where unwanted spurs and harmonics will appear in the sampled band, ensuring that there is no overlap between the signals of interest and spurious components.

Another critical aspect is managing the inherent spurious signals generated by ADCs, such as harmonics and interleaving artifacts. A frequency plan is successful when these spurious contributions remain out of the intended band, especially in systems that do not employ digital filtering after the data is sampled. Frequency planning also helps minimize the impact of clocking-related spurs, such as those that appear from clocking devices lacking channel isolation, which are particularly problematic in large element systems that use a high-density clock distribution solution.

Interleaving spur management is also a consideration in systems containing an interleaved ADC, where multiple ADC cores sample a common input signal orthogonally, effectively doubling the sampling rate and Nyquist region by a factor of two. This interleaving introduces spurious tones at Fs/2-Fin, however. Additionally, in systems that use higher interleaving factors, this same Fs/2-Fin spur is modulated once again by the "new" Fs/2-Fin. This results in the new Fin being comprised of the interleaved Fs/2-Fin spur, meaning that the spur count introduced is much higher than a single interleaving factor. Frequency planning enables you to mitigate these spurs by leveraging analog filtering or (preferably) digital decimation filters, attenuating them significantly without having to design a complex analog signal chain. While this approach reduces the instantaneous bandwidth to a selected region, it ensures better dynamic range and cleaner signal performance.

Common pitfalls in frequency planning

Despite its benefits, it is possible to have an improper frequency plan, which can lead to issues that degrade the ADC's performance. One common challenge is that the Nyquist zone overlaps. Poorly planned input signals may fall at the boundaries of Nyquist zones, leading to aliasing effects that reduce system performance. To prevent this, signals must be allocated within appropriate frequency bands to maintain spectral integrity with the Nyquist zone under consideration.

Clock spur contamination is another frequent issue, especially in systems using low-quality clocking devices or suboptimal clock distribution. These spurious signals, modulated onto the ADC spectrum, can severely impact sensitive applications by providing a known offset spur. Careful design of the clocking infrastructure, including the use of higher-quality clocking solutions, helps to mitigate these effects. Another possible method is to digitally filter the data with a high-order band rejection

filter at this offset frequency – although if implemented incorrectly, any wanted signal will be removed along with the spur.

Another challenge to overcome is the correction of tightly modulated third-order intermodulation distortion spurs. These spurs will almost always fall within the passband and are often the spur limiting the spurious-free dynamic range. In the event of a very high decimation factor, it might be possible that these tones land within the attenuation band. This is unlikely for most multitone systems, however, which inherently require a larger instantaneous bandwidth than single-tone systems and thus cannot incorporate such a large decimation filter.

Finally, you must navigate the trade-off between bandwidth and dynamic range. While decimation can suppress spurious components and harmonics, it comes at the cost of reduced instantaneous bandwidth. Balancing these trade-offs is essential to achieving the best performance for specific application requirements.

Advantages of proper frequency planning using decimation

Effective frequency planning delivers several benefits to enhance receiver system design. One advantage is improved spur suppression. Digital decimation filters effectively attenuate spurs, often achieving suppression levels around –85dBFS, which lead to cleaner signal performance and better utilization of the ADC's dynamic range for the intended signal rather than out-of-band spurious noise.

Another benefit is the reduction of data throughput from the ADC. By reducing the ADC's output data rate through decimation, you can interface the ADC with lower-speed, smaller and more cost-effective FPGAs. This reduction in transmitted data not only simplifies hardware requirements but also enables systems to operate in dual band or quad band, thus enabling the sampling of multiple RF bands simultaneously.

The ability of a system to be fully reconfigured in software alone is another significant advantage of using decimation on ADCs. You can plan the hardware interface between the ADC and FPGA to support the maximum data rate expected of the system, which would facilitate the ability to operate many other systems at lower data rates or more narrow bandwidths. Software-reconfigurable systems are particularly valuable in applications that require deployment into multiple scenarios.

Resource savings are also a notable outcome of effective frequency planning. By requiring fewer output lanes, whether high-speed serial data lanes or low-voltage differential signaling pairs, you can conserve valuable pins on both the ADC and FPGA, increasing the utilization factor. This is especially important in high-channel systems with printed circuit board area and power constraints.

Theoretical example: Frequency planning with decimation

Consider an interleaved ADC sampling a common RF input signal at Fs. The interleaving process introduces a spur at Fs/2-Fin, which can interfere with the intended signal. Applying a decimation-by-2 filter, as shown in **Figure 4**, makes it possible to attenuate this spur to levels within the decimation filter's rejection limits. Additionally, the decimation process reduces the ADC's output data rate, enabling cost-effective FPGA interfacing and simplifying downstream processing. Additionally, the wideband noise reduction introduces a 3dB process gain as a result of the N (that is Noise) in SNR being cut in half, while the S (that is Signal) stays the same.

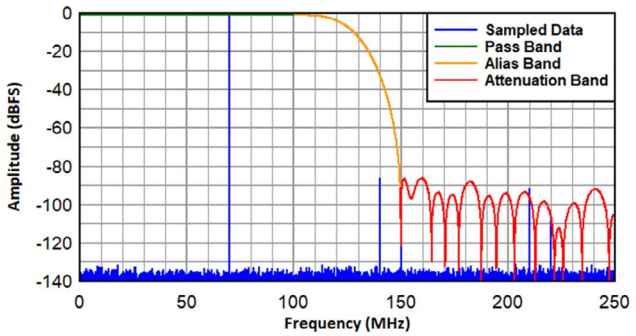

Figure 4. *Decimation filter response of theoretical data at 500MSPS with 70MHz FIN (with decimation factor of 2).*

Real World Examples: Frequency planning with decimation

The use of decimation on an ADC simplifies the frequency planning process because reducing the sampling rate effectively narrows the bandwidth of interest. Think of decimation as selectively focusing in on a narrower part of your spectrum. By focusing on a narrower band, more of the unwanted harmonics or spurs end up falling outside of the passband of interest, and in turn are filtered out. The following examples use the **ADC3669** to demonstrate the difference that decimation makes when frequency planning. **Figure 5** shows a traditional spectrum capture when the ADC is not performing decimation using an FFT size of 16384 points. You can see that the unwanted harmonics are in band and are negatively affecting performance.

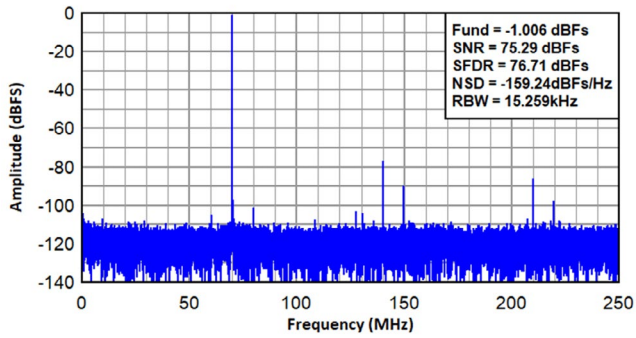

Figure 5. *Real Spectrum captured by the ADC3669 at 500MSPS with near 70MHz FIN (No Decimation).*

It could be that these harmonics are additive noise contributed by the ADC, or some external analog frequencies. **Figure 6** shows an example when the ADC is in real decimation mode, with a decimation factor of 2. You can see that the unwanted harmonic spurs now fall out of band and get filtered out by the decimation filter. Note that there is an additional +3dB improvement due to process gain.

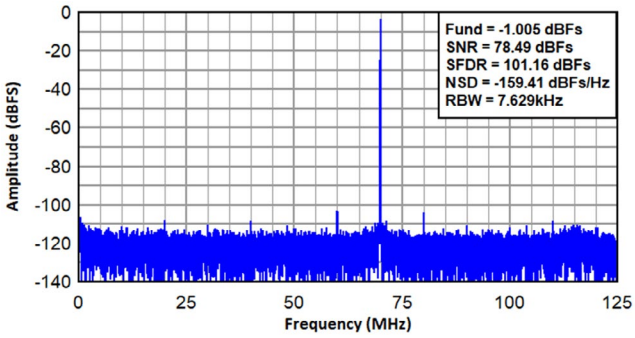

Figure 6. *Real Spectrum captured by the ADC3669 at 500MSPS with near 70MHz FIN (with decimation factor of 2).*

Additionally, the resolution bandwidth of the FFT actually reduces by a factor of two as well, since we maintain the same number of points for the FFT computation. This helps to resolve analog frequencies into closer bins. Up to this point, we have only talked about real decimation, which simply filters the data without any frequency shift. Real decimation is great if your signal of interest falls somewhere less than Fs/4 each time you decimate. But what if you want to decimate a signal that falls outside of this range? The signal of interest is often not centered at zero frequency (baseband), but rather at some intermediate frequency. This is where complex decimation comes in to play. ADCs with newer digital features, such as the ADC3669, incorporate an NCO mixer in the complex DDC stage. Mixing the signal of interest with an NCO frequency shifts the signal to baseband before decimation, enabling you to take advantage of the benefits of decimation for a signal anywhere within the device's bandwidth.

Figure 7 shows a capture of the ADC3669 in complex decimation mode with a decimation factor of 64 resulting in an effective sampling bandwidth of 7.8125MHz when the FFT is calculated using 8192 points. The input frequency is 70MHz and the NCO frequency is 71MHz. When the signal gets mixed with the NCO frequency,

the signal shifts to baseband, resulting in a tone at approximately -1MHz.

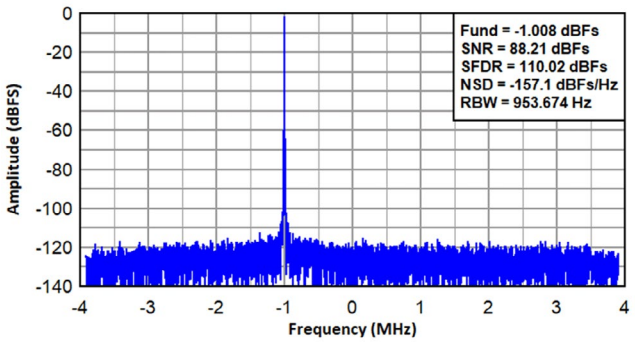

Figure 7. *Complex Spectrum captured by the ADC3669 at 500MSPS with decimation factor of 64 (FIN=70MHz, Fnco=71MHz).*

The ADC3669 can capture a narrow band with a decimation factor of up to 32768, which is beneficial for applications with high-density RF bands or tight channel spacing. Decimating by such a high factor allows you to zoom into your signal of interest, filtering out virtually everything else. The range of decimation factors offered by modern ADCs such as the ADC3669 enables more flexibility when frequency planning, as it is much easier to filter out unwanted spurs. **Figure 8** shows a capture with a decimation factor of 16384 calculated using 8192 FFT points, resulting in a resolution bandwidth of 3.726Hz. Even if your spurs are within thousands of hertz of the fundamental, you can easily filter them out with a high decimation rate.

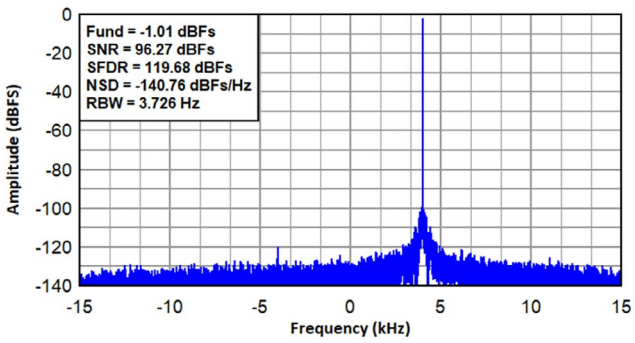

Figure 8. *Complex Spectrum captured by the ADC3669 at 500MSPS with decimation factor of 16384 (FIN=70MHz, Fnco=69.996MHz).*

As the NCO frequency is 4kHz lower than the input signal, the down converted signal appears at a positive frequency offset. While operating in this decimation mode and at 500MSPS, this ADC can sample signals within a 30.517kHz range around the programmable NCO frequency.

Conclusion

Frequency planning is an essential aspect of ADC-based system design, tackling challenges such as spur management, dynamic range optimization and AAF design, as well as efficient data handling. By implementing thoughtful frequency planning upfront, you can avoid common pitfalls such as Nyquist zone overlaps and clock spur contamination, while benefiting from advantages such as improved spurious suppression and dynamic range, a reduction in the ADC's digital interface or data rate, and FPGA resource savings. Carefully balancing these trade-offs and leveraging features such as ADC decimation make it possible to achieve high-performance, software reconfigurable receiver systems for a range of applications, while avoiding your next sampling hole.

Related Websites

- Discover techniques for fast frequency hopping in RF sampling data converters by checking out **SLYT861: Exploring fast frequency hopping in RF sampling data converters** by Chase Wood, featured in the Analog Design Journal, Oct 2024.
- Understand the relationship between sampling and data rate with **TIPL4701: Sampling vs. data rate, decimation (DDC) and interpolation (DUC) in high-speed data converters** by Jim Seton, published in Aug 2017.
- Gain insights into higher-bandwidth decimation examples by reading **Analyzing High-Bandwidth Spectrum Clusters** by Chase Wood, published in Embedded Computing Design, May 2024.
- Optimize your RF-sampling frequency planning using the **RF-Sampling Frequency Planner, Analog Filter, and DDC Calculator**. For detailed specifications, explore the **TI ADC3669 data sheet**.

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.

© 2025 Texas Instruments Incorporated

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATA SHEETS), 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 will 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 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.

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

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2025, Texas Instruments Incorporated