MOSFET Safe Operating Area and Hot Swap Circuits

Dan Eddleman

"Is this MOSFET's SOA (safe operating area) adequate for my application?" This is the most frequently asked question by designers of hot swap circuits. When evaluating a MOSFET's suitability for a specific application, it helps to have an intuitive understanding of SOA. After a short review of the SOA specifications found in MOSFET data sheets, this article presents a simple explanation of SOA, with a focus on MOSFET thermal behavior. With this understanding, designers can confidently use SOAtherm, a tool included with LTspice, to accurately evaluate MOSFET SOA in their circuit simulations.

WHAT IS SOA?

Every MOSFET data sheet includes an SOA plot, which describes the maximum time a MOSFET can be exposed to a specific voltage and current. Figure 1 shows the SOA plot from NXP Semiconductor's data sheet for the PSMN1R5-30BLE 30V $1.5m\Omega$ N-channel MOSFET. Consider the condition where 10V at 100A is applied to the MOSFET. Looking at the corresponding point on the SOA plot, we see that it falls between the 1ms and 10ms lines. The plot indicates that you can apply 10V and 100A for at least 1ms without damaging the MOSFET if the case (tab) is held at 25° C.



Derating for higher case temperatures (above 25°C) is discussed in the "Transient Thermal Impedance" section below. For events longer than 10ms, see the "Beyond 10ms" sidebar.



Figure 1. Safe operating area; maximum allowable time for a pulse at a fixed voltage and current



Figure 2. Transient thermal impedance from junction to mounting base as a function of pulse duration

STEADY STATE

Before delving deeper into transient SOA events, it is helpful to step back and review the more familiar steady state (DC) limits. MOSFET data sheets specify the maximum silicon temperature (typically 150°C or 175°C) as well as θ_{JC} and θ_{JA} , the thermal resistance from silicon (junction) to the bottom of the package (case), and the thermal resistance from silicon (junction) to the environment (ambient), respectively. (Sometimes $R_{th(JC)}$ and $R_{th(JA)}$ are used as alternative names for θ_{IC} and θ_{IA} .)

From the PSMN1R5-30BLE data sheet, $\theta_{JA} = 50^{\circ}$ C/W and $\theta_{JC} = 0.3^{\circ}$ C/W. θ_{JA} is used to calculate the temperature rise from ambient to the MOSFET's silicon die with a PC board configuration specified in the data sheet. In the PSMN1R5-30BLE data sheet, θ_{JA} is specified with a minimum footprint on FR4 PC board. Assuming your PC board is identical to the one the MOSFET manufacturer used to specify θ_{JA} , the silicon die temperature is:

 $T_{JUNCTION} = T_{AMBIENT} + \Theta_{JA} [°C/W] \bullet Power[W]$

For example, with a θ_{JA} of 50°C/W and an ambient temperature of 75°C, the die temperature will be 125°C when 1W is dissipated by the MOSFET. For most modern MOSFETs with exposed metal tabs, θ_{JA} is primarily determined by the PC board layout rather than the MOSFET itself (although the exposed pad shape and size play a role). Because θ_{JA} is highly dependent on the PC board layout and airflow, the manufacturer's specified θ_{IA} is only suitable for rough estimates.

 θ_{JC} is often a more useful metric, as it describes the MOSFET behavior without the influence of the PC board layout. To determine silicon temperature, use the following:

 $T_{JUNCTION} = T_{CASE} + \theta_{JC}[^{\circ}C/W] \cdot Power[W]$

With 1W of power dissipation, the silicon temperature is only 0.3° C above the case temperature. When using this formula, the case temperature (T_{CASE}) must be determined by physical measurement or through a thermal simulation of the PC board. Obviously, the PC board layout, airflow, and heat sinks are critical factors when calculating the steady state conditions.

TRANSIENT THERMAL IMPEDANCE

Most MOSFET data sheets also include a transient thermal impedance plot. The "single pulse" transient thermal impedance ($Z_{th(JC)}$) is the temperature rise produced by a time-limited power pulse. The longest time point on the transient thermal impedance plot always matches the θ_{JC} specification, because θ_{JC} is, by definition, the steady state (infinite time) thermal impedance. Figure 2 shows the transient thermal impedance from the PSMN1R5-30BLE data sheet. For the purposes here, only the "single pulse" curve is important.

The transient thermal impedance plot may be used to calculate the temperature rise for a power pulse of any duration. For example, assume a MOSFET drain-to-source voltage (V_{DS}) of 12V and a drain current (I_D) of 100A. The power dissipated by the MOSFET is 12V • 100A = 1.2kW. If we look at the transient thermal impedance plot at 1ms, the thermal impedance is 0.075°C/W. The silicon junction temperature is:

$$\begin{split} T_{JUNCTION} &= T_{CASE} + Z_{th(JC)} \ [^{\circ}C/W] \bullet Power[W] \\ &= 0.075^{\circ}C/W \bullet 1.2kW + 25^{\circ}C \\ &= 115^{\circ}C \end{split}$$

for a 1ms, 1.2kW pulse with a fixed case temperature of 25°C.

For moderate V_{DS} voltages (below the Spirito region, see sidebar), MOSFET manufacturers generate the SOA plot from the transient thermal impedance plot. In other words, these two plots are alternate expressions of the same information. The SOA plot shows the time it takes for the silicon die to reach its maximum junction temperature (150°C or 175°C) for each V_{DS} voltage and I_D current combination. Be aware that the SOA plot is only valid for a case temperature of 25°C and must be derated for higher case temperatures, including the case

Spirito Effect

Years ago when the maximum current in a hot swap design was less than 10A, it was easy to find a MOSFET to satisfy most applications. Two things have changed in the last decade. First, supply currents have increased significantly with 100A or more becoming common. Second, MOSFET manufacturers have been hard at work improving the resistance specifications of MOSFETs ($R_{DS(ON)}$) when they are fully turned on. Ironically, this has reduced the available SOA at higher drain-to-source voltages in what has been termed the "Spirito effect." Professor Paolo Spirito¹ explained that as MOSFET manufacturers have increased transconductance to improve the on-resistance, there is a greater tendency for the MOSFETs to fail by forming unstable hot spots.

Two primary factors compete to determine if hot spots cause MOSFET failure. One factor is the MOSFET's ability to dissipate power without a rapid increase in temperature. (This is reflected in the transient thermal impedance curve.) The second competing factor is the tendency of MOSFET cells to "run away" by stealing more current from neighboring cells as they get hotter. This second factor is dominated by the temperature coefficient of the MOSFET's threshold voltage, which drops with increasing temperature, causing "current

Spirito Region



crowding" in hotter cells. (MOSFET transconductance falls with increasing temperature due to reduced carrier mobility in the MOSFET's conduction channel. It somewhat counteracts the current crowding effects, but can be safely ignored in this explanation.)

Inside a MOSFET package, there is a silicon die that contains an array of MOSFET cells with their gates,

drains and sources connected in parallel. As some cells become hotter than others, their threshold voltages decrease relative to the cooler cells, causing the hotter cells to conduct more current. If the competing factors cited above reach an unstable condition, certain cells may thermally run away, drawing more and more current until they self-destruct.

The Spirito effect is observed primarily at high $V_{\rm DS}$ voltages since a given change in cell current results in a greater change in power at high $V_{\rm DS}$, resulting in an increased tendency of cells to thermally run away. Similarly, the Spirito effect is most pronounced at lower currents where there is more time for the MOSFET cells to thermally run away. (At higher currents, the average die temperature reaches the 150°C or 175°C before any cells exhibit significant thermal runaway.) For this reason, the high $V_{\rm DS}$ and low $I_{\rm D}$ region of the SOA plot, where the Spirito effect is dominant, is sometimes referred to as the "Spirito region" and is highlighted in the PSMN1R5-30BLE SOA in Figure 3.

Notes

¹ G. Breglio, F. Frisina, A. Magri, and P. Spirito, "Electro-Thermal Instability in Low Voltage Power MOS: Experimental Characterization," IEEE Proceedings ISPSD 1999, Toronto, p233 temperature rise that occurs from the pulse itself. (See "Beyond 10ms" sidebar.)

Knowing that the maximum junction temperature of the PSMN1R5-30BLE is 175°C, and using a case temperature of 25°C, we can calculate the maximum allowable time at 1.2kW. $T_{JUNCTION} - T_{CASE} = Z_{th(JC)}[°C/W] \bullet Power[W]$ 175°C - 25°C = $Z_{th(JC)} \bullet 1.2kW$ $Z_{th(JC)} = 0.125°C/W$

Looking at the transient thermal impedance plot we find that $Z_{th(JC)}$ crosses 0.125°C/W at roughly 2ms, which also matches the SOA plot.

With an understanding of transient thermal impedance plots, we can calculate the allowable time for case temperatures other than 25°C. In the previous 1.2kW example, the allowable time was 2ms with a 25°C case temperature. Now, assume the case temperature is 85°C:

$$\begin{split} T_{JUNCTION} - T_{CASE} &= Z_{th(JC)}[^{\circ}C/W] \bullet Power[W] \\ 175^{\circ}C - 85^{\circ}C &= Z_{th(JC)} \bullet 1.2kW \\ Z_{th(JC)} &= 0.075^{\circ}C/W \end{split}$$

Looking at the transient thermal impedance plot we find that $Z_{\text{th}(IC)}$ crosses 0.075°C/W at 1ms,

significantly less than the 2ms we found for a 25°C case temperature.

Because thermal behavior is linear, we can use the transient thermal impedance plot to determine the temperature rise for any power shape. While it is possible to do this calculation using convolution, it is easier to model thermal behavior in an electrical circuit simulator such as SPICE. In particular, the SOAtherm tool in LTspice can be used to model the MOSFET thermal behavior.

SOATHERM THERMAL MODELING IN LTspice PREDICTS MAXIMUM MOSFET DIE TEMPERATURE

A designer armed only with MOSFET data sheet SOA plots faces a difficult challenge in predicting a MOSFET's suitability for a hot swap design. Fortunately, MOSFET thermal behavior (and SOA) can be modeled in circuit simulators such as LTspice.

The SOAtherm symbol included in LTspice includes a collection of MOSFET thermal models that simplify the task of predicting MOSFET maximum die temperature over time, even in the Spirito region. The thermal model reports the temperature of the hottest point on the MOSFET die without influencing the electrical behavior of the MOSFET model.

For better or worse, the SOAtherm models are based on the MOSFET manufacturers' data sheets, and as such are only as accurate as the manufacturers' data itself. With that in mind, design with plenty of margin since the SOA curves provided by MOSFET manufacturers are usually typical numbers without sufficient derating to account for part-to-part variation.

Using SOAtherm

To use SOAtherm, place the SOAtherm-NMOS symbol on top of a MOSFET in an LTspice simulation (Figure 3). The voltages at the Tc and Tj pins of the SOAtherm-NMOS symbol indicate the case temperature and silicon junction temperature respectively. (Refer to the SOAtherm-NMOS tutorial for more information about using this model, including how to adjust the ambient temperature settings and other parameters.)

After running the simulation, the silicon and case temperature can be observed in the waveform viewer (Figure 4). In the waveform shown here, the MOSFET silicon junction temperature rises from 25°C to

Beyond 10ms

For most MOSFETs, the case temperature does not rise significantly during transient events lasting less than 10ms, because it takes time for the heat to move through the MOSFET silicon and copper. At roughly 10ms, the heat begins to reach the PCB.

If the MOSFET's copper tab is small, the temperature of the MOSFET begins to rise faster as the heat reaches the PCB. For packages where the copper tab is larger (i.e., D2PAK packages) the heat begins to move outward into portions of the copper tab that are still cool. As a result, packages with more copper perform better in high SOA applications (hot swap designs, linear amplifiers, etc.) than MOSFETs with less copper, even if their transient thermal impedance and SOA plots appear similar.

Think of the copper as a reservoir that helps to limit the MOSFET temperature rise during events in the 10ms–10s timeframe. The transient thermal impedance plot and the SOA plot are often deceptive, because they are created by assuming the case temperature is fixed at 25°C by an impossibly perfect heat sink.

The figure here shows the simulated thermal characteristics of a Power-SO8 package and a D2PAK package soldered to a PCB with a 1oz copper plane on the top layer. The figure also includes the thermal impedance curves of the type found in MOSFET data sheets where the case temperature is fixed.

At 1ms, the heat is concentrated within the silicon die. The D2PAK silicon is cooler with a thermal impedance of 0.075°C/W compared to 0.14°C/W for the Power-S08, primarily due to the larger silicon die in the D2PAK. At 10ms, the heat begins to reach the bottom of the copper tab, and the temperatures start to diverge. At 100ms, the Power-S08 die has a temperature rise of 4.2° C/W, while the temperature rise of the D2PAK is only 0.6°C/W. Clearly, the extra copper in the D2PAK saves the day.









72°C. The case temperature rises from 25°C to 35°C. (A 1V rise on the Tc or Tj pin is equivalent to a 1°C temperature rise.)

Remember to Simulate Special Circumstances

There are several important special circumstances that should not be overlooked when using SOAtherm to determine if MOSFET SOA limits may be exceeded.

- Input supply step. For example, the SOA requirements of a -48V telecom application, where the input supply may quickly step from -36V to -72V can require a MOSFET with significant SOA capabilities. When supplies are pre-regulated, or are well controlled to eliminate such steps, the SOA requirements are reduced.
- Start-up into a load. The downstream circuitry may turn on and draw current before the supply is fully ramped-up, or a component such as a capacitor may fail in a resistive short. Simulating a resistive load at the output can indicate when a MOSFET might unexpectedly be subjected to a condition requiring significant SOA.
- A short circuit at the output occurring during otherwise normal operation. You never know when the user is going to drop a paperclip into the

chassis, and a hot swap circuit (or fuse) prevents a call to the fire department.

EXAMPLE USING THE LTC4226 WIDE OPERATING RANGE DUAL HOT SWAP CONTROLLER

The LTC4226 is a dual hot swap controller that drives external N-channel MOSFETs in applications with supply voltages as high as 44V.

In the circuit in Figure 5, the LTC4226 provides current limit and circuit breaker features for a 12V supply and a 5V supply. The circuit breaker timer is configured with the capacitors connected to the FTMR1 and FTMR2 pins. When the voltage across either sense resistor is between 50mV and 86mV, the corresponding capacitor at FTMR1 or FTMR2 is ramped up with a 2μ A current.

Because the current limit is not engaged until the sense resistor voltage reaches 86mV, the power dissipation in the MOSFET is negligible as long as the current remains below $86mV/5m\Omega = 17.2A$. When the current exceeds that level, current limit is engaged and the FTMR1 or FTMR2 pin ramps up with 20µA. The appropriate channel's MOSFET is turned off when the corresponding FTMR pin reaches 1.23V, setting a maximum time before the MOSFET is shut off. In this example, 100nF capacitors configure a 6.2ms current limit timeout for both channels.

With the LTC4226, the worst-case MOSFET power dissipation occurs when the output is shorted to ground. As a result, determining the required SOA is straightforward. (With hot swap controllers that feature current foldback or power limiting, more effort is required to determine the worst-case loading condition.) Referring to the SOA plot in Figure 1 for the PSMN1R5-30BLE, it can be seen that 6.2ms is well inside the SOA limit at 17.2A and 12V. An SOAtherm simulation confirms that the total junction temperature rise is less than 50°C. The same simulation shows a negligible case temperature rise of roughly 5°C, which would be expected from this rather large D2PAK package during a short 6ms event.

The 5V supply in this example application uses a powerPAK-SO8 package for the MOSFET, which is smaller than the D2PAK used for the 12V supply. A smaller package may be used for the 5V supply because the worst-case power dissipation of the 5V supply's MOSFET is $17.2A \cdot 5V = 86W$ versus the $17.2A \cdot 12V = 206W$ worst-case dissipation of the 12V supply's MOSFET. An SOAtherm simulation of this circuit



Figure 5. LTC4226 hot swap controller protects a 12V and a 5V supply. Both supplies provide 9A steady state current and up to 17.2A during transients

predicts a junction temperature rise of 40°C, including a case temperature rise of 30°C. The larger case temperature rise is explained by the smaller size of the powerPAK-SO8 package (and correspondingly less copper) relative to the D2PAK used for the 12V supply.

The above calculations and simulation help to verify a circuit design and MOSFET selection, but the ultimate test must be done in the lab with an assembled circuit. Because the worst-case SOA requirement of the LTC4226 occurs with an output short, the lab test is as simple as quickly applying the input power supply with a grounded output. A good technique is to hot-plug the LTC4226 circuit into a live power supply to simulate an actual hot swap event. Alternatively, an output short can be applied while the input supply is fully powered. To determine if the circuit has extra margin, swap the timer capacitor to a larger value and test again.

The circuit in Figure 6 shows a technique for using two parallel MOSFETs when a single MOSFET may not satisfy the SOA requirements of an application. In general, it is not advisable to use parallel MOSFETs to increase the SOA capability of a circuit. Mismatches between the MOSFETs, especially mismatches in the threshold voltages, may result in one MOSFET thermally running away and conducting all of the current. Nevertheless, the circuit in Figure 6 safely uses parallel MOSFETs by implementing independent current limit in each channel, preventing either MOSFET from running away.

Additionally, the cross-coupled PNPs, Q1 and Q2, only allow the circuit breaker timers to activate when both MOSFETs are conducting their full current. Without the cross-coupled PNPs, one channel's circuit breaker timer could activate if it was delivering a greater share of the load current.

CONCLUSION

As the power levels required in hot swap applications have increased, so have concerns regarding MOSFET safe operating area. Frequently, the most challenging aspect of designing a high power hot swap circuit is determining whether a specific MOSFET is capable of supporting the application. At a minimum, the circuit designer must be comfortable interpreting MOSFET SOA plots. As power levels increase and approach the limits of existing MOSFET technology, an understanding of transient thermal impedance plots and the ability to simulate this behavior in SPICE circuit simulations are invaluable tools in the hot swap circuit designer's arsenal.

