

# Don't Be Misled by Power Device Specs

By Roger Stout, Senior Research Scientist, ON Semiconductor, Technology Development, Advanced Packaging, Phoenix

Understanding that maximum device power is situation dependent can go a long way in explaining why semiconductor data-sheet specifications for power can be misleading.

More than once, I've fielded this question from applications engineers on the telephone: "I just found a data sheet for the XYZ device, and it has two different maximum power values listed. Which one is correct?" Or, perhaps designers have seen one of those device sample boards that have a dozen or more package types mounted on the front, with a table on the back showing two columns with two different maximum device power values for each package style. How can this be?

A couple of years ago at a major industry conference, at the first break in a half-day seminar I was presenting (after I'd already covered the basics), someone came up to me and pleaded, "But for a given package, shouldn't there just be some intrinsic amount of power dissipation that results in that maximum junction temperature?"


Speaking as a thermal expert, the answer to the question is yes and no. Unfortunately, this answer probably doesn't help the designer much. However, a mental model rooted in electrical fundamentals will help clarify these concepts.

## An Electrical Analogy


For starters, let's take a 10- $\Omega$  resistor and pass 1 A of current through it (Fig. 1). Now, we may ask, "What's the voltage of this resistor?" Actually, such a question

doesn't make any sense. For the question to make sense, what needs to be asked is, "What's the voltage drop across the resistor?"

Here's the problem. Voltage, as such, implies a reference. With no reference specified, it might reasonably be assumed that the voltage at one terminal be compared with the voltage at the other. Then the answer is clear and simple: 10 V.



**Fig. 1.** Given the values shown for this resistor, a question about absolute voltage is meaningless. Even if we know the resistance and current values, all we know is the voltage drop. The two ends can be known relative to each other, but not with respect to the outside world unless we make a specific reference to the outside world.

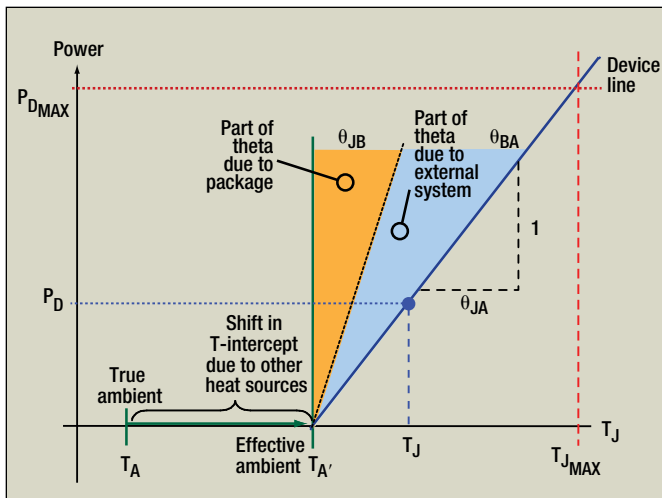


**Fig. 2.** The value of voltage  $V_A$  in this circuit cannot be determined without more information, even if current I1 is known. The value of current I2 must also be known to figure out what the value of  $V_A$  is.

It should be intuitive that the actual voltage at one end or the other could really be anything; this resistor would act the same in any circuit and, as long as it was conducting 1 A, the voltage at the input end would always be 10 V relative to the output end. If the voltage at the output end were determined by other things in the circuit, it might be ground (that is, a systemwide 0-V reference), or it might be something else entirely. For example, there's no reason the input end of the resistor couldn't be ground, and then the output end would be -10 V.

Let's suppose we have three connected resistors with two current sources (Fig. 2). If I1 is 1 A, can we tell what the value of  $V_A$  is? How about this variation: What if we don't allow  $V_A$  to be higher than 10 V, what would the maximum current of I1 be?

Obviously, we can't answer either of those questions without more information. Believe it or



**Fig. 3.** For a given device, the solid blue line, which represents all the possible operating points, can never become any more vertical than allowed by the thermal characteristics of the device's package.

not, most semiconductor package thermal systems boil down to a thermal circuit analogous to Fig. 2; yet most designers who ask for maximum power dissipation think of Fig. 1 (with the output end grounded).

## The Thermal Domain

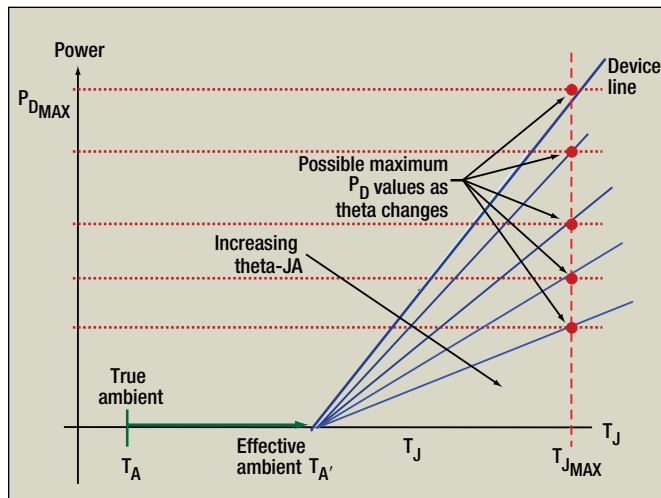
Let's formally switch the electrical analogy back into the thermal domain. First, temperature is like voltage. We're mainly interested in temperature *differences* between points in a system. An important point to consider is thermal ground, otherwise known as ambient. But ambient is probably not zero in a given system. Another important point is the junction temperature. Although we need to get to it relative to ambient, we probably have a nonrelative maximum junction temperature to worry about.

Second, thermal heat is like electrical current. Be very careful here. Because we so often talk about electrical current flowing in a semiconductor device, it's all too easy to forget that when we're specifying amperes of electrical current, we basically haven't said anything at all about how much *heat* is being dissipated. We'll have to do a little electrical analysis to get from electrical current into power, but we have to make sure to use power as the current analog in the thermal model.

Third, thermal resistance must have units of °C/W. Then, when we multiply a thermal current (W) by a resistance, we get a temperature difference in °C. Note again that we're referring to a temperature difference, not a temperature.

## Data-Sheet Specifications

Now, when a data sheet specifies a value of thermal resistance for a package and calls it theta-JA, what have we really learned? What the data sheet says is that, in some specific environment, the total temperature difference between a junction (what we really care about, keeping



**Fig. 4.** The actual maximum power dissipated by a device depends on theta-JA; it decreases with increasing values of theta-JA.

it at or below some significant limit) and ambient — the ubiquitous thermal ground — is this “known” coefficient multiplied by the power we're dissipating in the device.

What the data sheet doesn't tell us is the junction temperature. And even if we know the power dissipation, until we know ambient, we still don't know the junction temperature. (It also may be intuitive that given a certain amount of power, as ambient goes up, so does the junction temperature.) In the mean time, the actual power dissipation is determined by an electrical system design.

What may not be intuitive is that the quoted thermal resistance is also actually determined by the system design. Moreover, it is probably not intuitive that the effective ambient (the temperature the device's junction will experience under its own zero-power conditions) is also determined by the thermal design. Let's look at Fig. 3 for an explanation.

In Fig. 3, the blue device line describes all the possible operating points of the device, where device power and junction temperature are the only variables, with everything else in the system being fixed. Where the device line intersects the horizontal (temperature) axis is the effective ambient. We can see that the effective ambient is shifted to the right of the *true* ambient, which is a “given” in whatever environment the device is operating. The amount of that temperature shift represents the background heating of the device in question by the rest of the components in the system when those components are turned on and the device of interest is turned off.

The slope of the device line is the reciprocal of theta-JA. There's usually some control over that (copper spreader area, airflow), but if nothing else, we must understand that theta-JA is actually comprised of a “package dependent” portion and an “external system” portion. This is illustrated in Fig. 3 by subdividing the wedge between the vertical line through the intercept and the device line. What definitely can't change is the package contribution to that

wedge. In other words, the device line can never get any more vertical than allowed by thermal characteristics of the package itself (let's call this theta-JB, where B represents the board under the device).

A lot of money can be spent trying to reduce the external contribution to the wedge (theta-BA) until it's just a sliver; or, it can be very economical to realize that not much can be done about this — it'll just be what it ends up being. In any event, when all we have from a data sheet is a theta-JA, then all we have is the total wedge. We don't know how much of it is due to the manufacturer's lab test conditions and how much is due to the package itself. Worse, of course, is when we don't know what the lab test conditions were, because then we actually don't even know if we have the correct theta-JA at all.

Now, let's backtrack to Fig. 2, where almost nothing was known, yet we still wanted to know the maximum allowable current in the 10-Ω resistor. If you've analyzed the environment around your system, you may have the correct total theta-JA (which would be the sum of R1 and R2 — and splitting it up means knowing a value for theta-JB, which may or may not be on the data sheet). Fig. 4 illustrates how the actual maximum power that can be dissipated depends on theta-JA.

Yet, even if we know both R1 and R2 individually, until we figure out R3 and I2, we still can't figure out  $V_A$  or the




Fig. 5. Even if the slope (theta-JA) is known and fixed, actual maximum power for a device depends on the T-intercept.

maximum current that would go with any particular limit to  $V_A$ . Fig. 5 illustrates how the actual maximum power that can be dissipated depends on the T-intercept, even when the slope (theta-JA) is known and fixed. Think of I2 as being all the other significant heat sources in the system, and R3 as the proportionality between those other sources and how much up the T-axis the effective ambient slides in response to those other heat sources. Clearly it's quite



# MAGNA-POWER ELECTRONICS

**FEATURES:**

- Many models: 5 to 2500 Vdc, 3 to 6000 Adc
- Series and parallel master/slave operation
- High dielectric withstand: 2500 Vac
- 50 Hz to 400 Hz input frequency
- All user interface circuitry referenced to earth ground
- OVT and OCT shutdown standard
- Automatic V/I crossover
- RS232 interface with SCPI commands
- Optional IEEE-488, RS485, and Ethernet programming
- Front panel potentiometers for stepless rotary control
- Front panel calibration
- User friendly controls and indicators
- Remote Interface Software with self-teaching features
- Drivers: Certified LabWindows/CVI and LabVIEW for GPIB, Serial, and TCP/IP communications
- Load dependent air or optional water cooling
- High power factor
- CE Mark



**INDUSTRY LEADING  
AC/DC POWER SUPPLIES**

**⚡ 3.3 kW to 750 kW AC to DC Product Line**

*RACKMOUNT SERIES*

- PQ Series III: 3.3 to 10 kW
- TS Series II: 15 kW to 45 kW

**⚡ Used around the world for:**

- Battery charging
- Power waveform generation
- Alternate power source emulation
- Capacitor charging and forming

*CABINET SERIES*

- MS Series II: 30 kW to 75 kW
- MT Series III: 100 kW to 750 kW

- Focusing magnets
- Electroplating
- Electrodeionization
- Numerous custom applications

**www.magna-power.com**

39 Royal Road, Flemington, NJ 08822  
Tel: 908-237-2200 Fax: 908-237-2201


possible that even a perfectly vertical device line (one with a zero theta-JA) can still have a maximum power allowed of 0 W, if the effective ambient slides all the way up to  $T_{J,MAX}$ .

**A Specific Example**

Fig. 6 is a specific example of the opening dual-maximum power data-sheet scenario. The data sheet gave both theta-JA and theta-JB values, so the implication is that we know both R1 and R2 in the Fig. 2 analogy; and since we're talking about ideal data-sheet systems, we don't have to know R3 because I2 is zero.

The data sheet also says that when using theta-JA, the ambient was 25°C, with a maximum power of 1.25 W; whereas when using theta-JB, the board temperature was 75°C and maximum power was 3 W. Is there a contradiction? Not really. A contradiction arises if we implicitly believe the board temperature is the same in both scenarios. The data sheet doesn't say it is, and if we figure it out for the theta-JA scenario, we get a value of about 119°C.

Look at it this way: Any time we need to dissipate 3 W in this device, we're going to have to hold the board tem-



**Fig. 6.** In many a device data sheet, two different maximum power values are given, one based on total theta-JA (left) and the other based on theta-JB (right). If they differ, it's because they're based on mutually incompatible assumptions as to the board temperature.

perature down to 75°C or less. We won't be able to do that with a theta-JA of 100°C/W and an ambient of 25°C. We might be able to do it with a lower theta-JA (which might mean adding some copper-spreader area). Indeed, recall that device/package sample board I mentioned, the one having two columns of maximum power data on the flip side? That's how those values were derived: two different theta-JAs were chosen based on two different copper spreaders for each device.

For whatever reason, the manufacturer didn't want to do it that way on the data sheet. The company chose to give a true "package" thermal value (theta-JB) and some sort of realistic board temperature that permits more power than the theta-JA scenario, but nothing as good as taking the board all the way down to ambient. Taking the board temperature all the way down to ambient might give a great marketing value, but no customer is going to be able to build a 0°C/W thermal system to wrap around this device.

It might be argued that this really would be the maximum power value, but then it could be said that if we're going to build a 0°C/W system, then why not instead design for a 15°C ambient? Again, there is no standard environment that is meaningful to every possible end user of a package or device; what makes a realistic maximum power value for one is totally outrageous for another.

**Situation Dependent**

Maximum power is really situation dependent. A particular package can have a steep or shallow device line, and certainly the steeper the better. But knowing this slope alone doesn't answer the maximum power question. A poor thermal system design can leave us with a maximum power of zero, no matter what package is chosen.

Undeniably, there are differences between packages, and some packages can dissipate more power than others in the same environment. But a designer can't look just at the maximum power rating of a device on a data sheet to know if he or she can really use the device in a given situation.

**PETech**