

How Hot is Too Hot?

By:

Craig Hillman, Ron Wunderlich, and Greg Caswell

DfR Solutions, LLC

301-474-0607

chillman@dfrsolutions.com, rwunderlich@dfrsolutions.com, gcaswell@dfrsolutions.com

Abstract

Power supply engineers have access to more powerful thermal simulation tools than ever before. Finite element analysis and computational fluid dynamics can provide highly accurate predictions for even very complex thermal management solutions. However, these new capabilities fail to answer the most critical question: How hot is too hot?

Introduction

The power supply is the core of all electronic equipment. It often has to inexpensively manage high levels of power and voltage in relatively tight spaces. To meet these requirements, power supply designers need to call upon an impressive level of creativity and skill.

But, creativity relies on knowledge. This is especially true in power supply design, where solutions to electrical noise, timing, and efficiency are born from expertise and experience. Unfortunately, the feedback loop on thermal solutions is not always so direct. While amazingly powerful thermal tools can predict the distribution of junction, case, and ambient temperatures in minute detail, it is often difficult to know what the temperatures should be as opposed to knowing what the temperatures will be.

The classic approach for determining appropriate component temperatures is based on a combination of datasheet information and derating strategies. As an example, if a Power MOSFET is rated to 125C and internal derating guidelines specify a 20C margin, the selected thermal solution must ensure a temperature no higher than 105C. Beyond the typical confusion over which 125C (Junction temperature? Case temperature? Lead temperature?) and if thermal derating should be absolute (e.g., 20C margin) or relative (e.g., 80% of the rated temperature), this methodology is popular because it is so straightforward.

Increasingly, the electronics industry is realizing that classic derating is no longer satisfactory. Its broad assumptions, which are not based on actual failure modes and degradation mechanisms, can result in overly conservative and expensive designs or products with insufficient reliability. Either case loses customers and shrinks market share. A more effective approach is one that takes the results of thermal modeling or measurement and inserts these results into design rules or predictive tools that are based on reliability physics.

This article will identify the most temperature sensitive components in power supply design, discuss how temperature induces degradation in these components, and what existing knowledge is available to allow the designer to make physics-based decisions on how hot is too hot.

Components At Risk

The derating approach was always a questionable practice, but had some legitimacy in older electronics because solid state mechanisms typically took decades, if not hundreds of years, to evolve and induce any significant number of failures. So, derating was more about function (parametric drift, etc.) than reliability. Reliability has become a greater concern as the need for more functionality in less space has resulted in such fine feature sizes that degradation mechanisms are now evolving in a matter of years, or even months, even when designers adhere to classic derating guidelines.

The component technology of greatest concern in regards to temperature and reliability in power supply design are

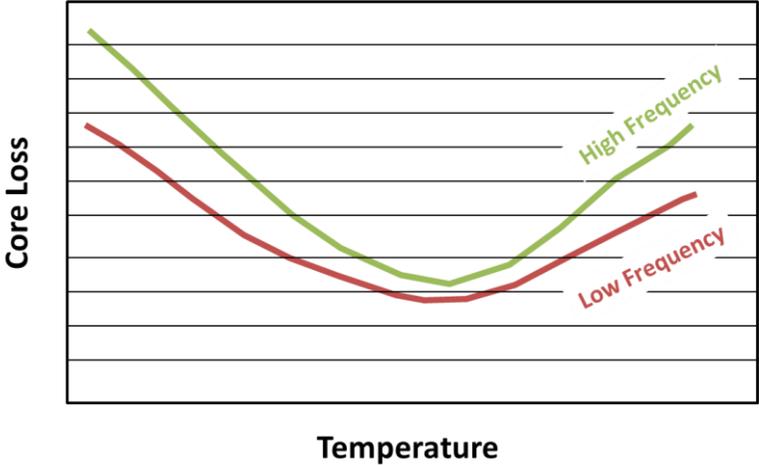
- Magnetics (Transformers / Chokes)
- Optocouplers / Light Emitting Diodes (LEDs)
- Capacitors (Electrolytic / Ceramic / Film)
- Integrated Circuits
- Solder Joints

Magnetics (Transformers / Chokes)

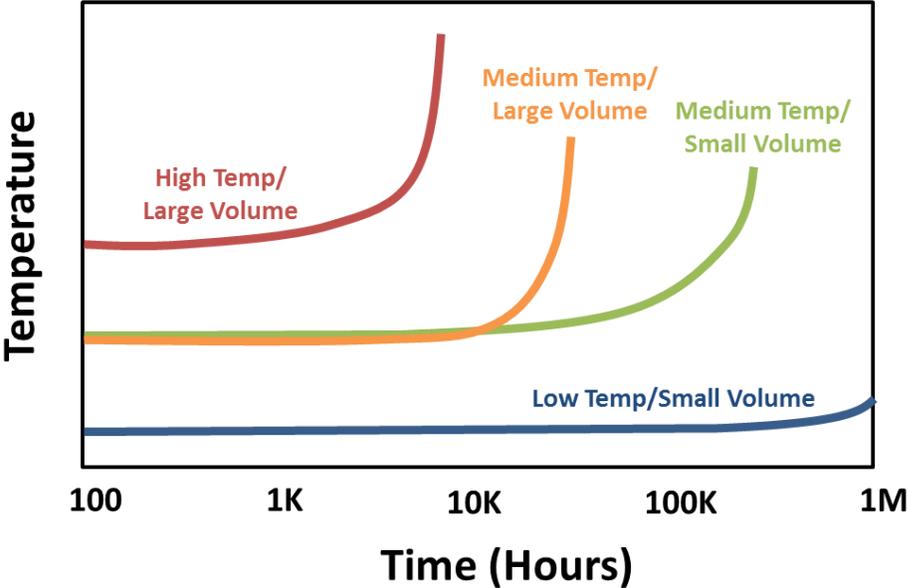
Magnetics, such as transformers and chokes, get top billing as they are the component technology often least considered when concerns about temperature arise during design reviews and component stress analysis. Since transformers are also typically custom, many do not even come with a temperature rating. So, how to determine how hot is too hot for magnetics? There are three key issues of concern.

First, the saturation current in ferrite material has a soft saturation curve which can tend to obscure when a material starts to saturate. Never the less, this will vary with temperature. For example, B_{sat} could be 3000 gauss at room temperature, but at 100C it can drop down to 2000 gauss. Saturating the magnetic material will not damage a magnetic, but it will appear to be shorted to the electronic circuit and can cause the circuitry to fail. Debugging this scenario is quite difficult because the transformer or inductor can seem ok at room temperature.

Second, designers are sometimes under the mistaken impression that maximum temperature rating is equal to the Curie temperature (this can be between 100C to 300C). However, core loss usually reaches a minimum at temperatures between 50C to 100C, which is typically below the Curie temperature. Depending on the ferrite design, structure and cooling, the magnetic can go into a thermal runaway if the core temperature is on the right side of this minimum. This because for some magnetic designs, the core loss will increase over time, resulting in a higher temperature, which causes an even higher core loss, and so on and so on. This increase in Core Loss is due to Thermal Aging.



Thermal aging is primarily a concern for powder iron cores, which are lower cost and sometimes more appropriate than ferrite cores. Powder iron cores are mixture of extremely small particles of iron oxide or other magnetic material mixed with an organic binding agent. It is well known that long-term exposure at elevated temperatures can induce 'Thermal Aging' of the binding agents. There are several variables that affect the rate of Thermal Aging of magnetic components, including the core material, the peak AC flux density, the operating frequency, the core geometry (larger cores age faster), copper loss and core temperature. As Thermal Aging progresses, the eddy current loss, which is a critical component of the core loss, becomes significantly higher. Increasing core loss, eventually results in higher core temperatures and failure of the magnetic component.



There are several things that a designer can do to mitigate the higher temperature effects on magnetics. For example, the device could utilize a low loss core material. Also, higher frequency with total turns

reduction while maintaining turns ratio or using Litz wire for reduced coil heating can be done. Or, varnish applied with vacuum evacuation (also allows for washing) can also facilitate improved thermal performance.

Light Emitting Diodes / Optocouplers

Light emitting diodes (LEDs) are typically incorporated into power supplies as indicator lights or as one half of optocoupler technology. LEDs have a natural lifespan that ends in a wear-out mechanism. Defects within the active region can spur nucleation and dislocation growth and are particularly affected by temperature and current. Lifetime of LEDs is typically described by the following functions

$$t_2 = t_0 \left(\frac{I_2}{I_0} \right)^{-n} \cdot t = c e^{Ea/kT}$$

where I is forward current, n is forward current exponent, Ea is activation energy, k is Boltzmann constant, and T is temperature in Kelvin. For LED technology, the forward current exponent is typically 1.5 and the activation energy is 0.7 eV.

Indicator lights rarely experience 'failure' because designers typically introduce a low duty cycle (refresh rate simply must be 60 Hz or faster) and a low forward current. Because of the high intensity of LEDs at rated forward current, use of forward currents $1/10^{\text{th}}$ the rated forward current is typically sufficient for all but the brightest environments. With such a high exponent, the derating on the forward current is more than sufficient to extend lifetime out by several decades. In addition, failure definition in LEDs is a 50% reduction in brightness. Even with this reduction in brightness, most users can still perceive the indicator lights and can respond to their signaling.

The real issue with LEDs comes with their use in optocouplers. Low voltage optocouplers are typically rated to 40 to 60mA. While derating guidelines typically allow for up to 80% of the rated forward current, most designers, to optimize CTR, typically drive the LED at 1 to 10 mA. And this approach is a good thing when the optocoupler sees anything close to elevated temperatures.

As an example case study, a solar power conversion company was operating a 50mA optocoupler in their design at 11mA, which was well within derating guidelines, at a 50% duty cycle (solar panels don't produce much power at night). This particular manufacturer's reliability goals were 95% reliability after 15 years.

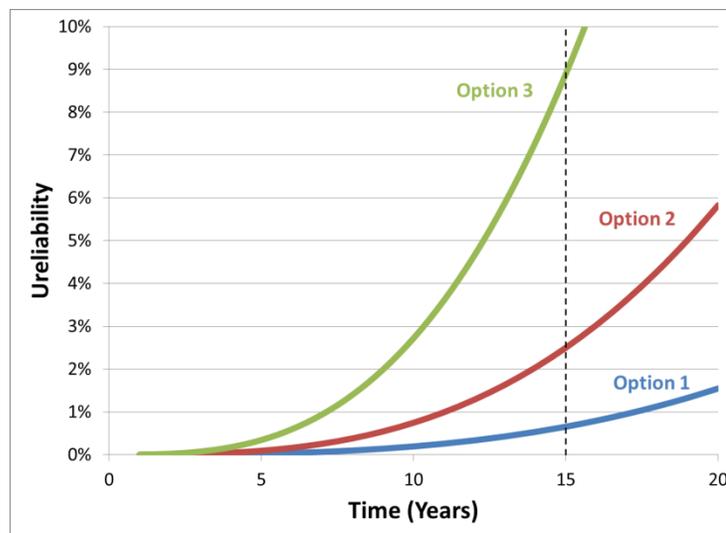
The challenge was in where to locate the optocoupler. The three locations had various advantages and disadvantages, but they had distinct differences in temperature. Option 1 was at ambient, Option 2 was 5C above ambient, and Option 3 was 10C above ambient. Because of this application, ambient was driven by diurnal temperature changes (see table below)

Table 1: Simplified (rounded to the nearest 5°C) diurnal cycle experienced in Phoenix environment

Month	Cycles/Year	Ramp	Dwell	Max. Temp (°C)	Min. Temp. (°C)
Jan.+Feb.+Dec.	90	6 hrs	6 hrs	20	5
March+November	60	6 hrs	6 hrs	25	10
April+October	60	6 hrs	6 hrs	30	15
May+September	60	6 hrs	6 hrs	35	20
June+July+August	90	6 hrs	6 hrs	40	25

To use the equations above, a rated lifetime is necessary. While a surprising number of optocoupler manufacturers do not provide this information, the standard LED in optocouplers is typically specified to have a mean time to failure (MTTF) of 50,000 hours at maximum rated forward current at 25C. It is important that MTTF is not lifetime; MTTF is the time to 63% failure. The more important reliability metric of time to 1% failure is typically far shorter.

Using the relevant equations and inputs, the manufacturer quickly realized that the criticality of keeping the optocoupler as cool as possible (see below). The 20X difference in unreliability would not have been captured using the traditional approaches of derating to determine appropriate temperatures.



Electrolytic Capacitors

If there is any component designers should be aware of in regards to temperature, it would be electrolytic capacitors. The only electronics component that relies on liquid for functional operation, it has long been known that electrolytic capacitors have a limited lifetime due to the gradual evaporation of their liquid electrolyte. This loss of electrolyte leads a gradual decrease in capacitance and an increase in equivalent series resistance (ESR). As a result of this process, in addition to the standard part rating values, such as voltage, current (ripple), and temperature, all electrolytic capacitor manufacturers also provide a rated lifetime.

So, how do power designers take into consideration electrolytic capacitors and temperature? Most companies extrapolate manufacturer's ratings to actual use environment using a classical Arrhenius equation to develop a conservative prediction of lifetime of electrolytic capacitors (assumes an activation energy of approximately 0.55 eV). The equation is defined as

$$L = L_r \times 2^{\left(\frac{T_r - T_a - \Delta T}{10}\right)}$$

where L_r is the rated lifetime, T_r is the rated temperature, T_a is the ambient temperature, and ΔT is the temperature rise due to ripple current. Does this approach work? The answer, as with most things, is a little bit of yes and a little bit of no.

Unlike most other component technologies, the lifetimes defined by electrolytic capacitors are more conservative than MTTF, with actual failure rates at lifetime between 1 to 5%. Designers have typically applied a 50% derating to the ripple current, which greatly extends the lifetime. And, most designers have also, at least in the past, been conservative with their calculations by assuming a worst-case temperature 24 hours/day, seven days/week.

More recently, in a push to lower cost and reduce dimensions, designers have been forced to become more aggressive on applied ripple current. In some applications, designers are routinely applying ripple currents above the manufacturers' ratings (though, for short durations). In addition, designers can squeeze out more margin by taking into consideration changes in temperature in their life calculations.

Actual lifetime can also vary depending upon the sensitivity of the circuit to change in component parameters. The manufacturers' definition of lifetime is typically a 20% drop in capacitance. By that point, the equivalent series resistance (ESR) could experience an increase of 2X to 5X. Depending on the sensitivity of the circuit, this could induce failure in the product far before 'failure' of the capacitor. When, because of dimensional constraints, designers place electrolytic capacitors adjacent to, or even touching, hot components, the standard lifetime equation may not even apply as non-uniform temperature distributions across the capacitor can induce accelerated degradation and pressure increases that can induce rupture.

Ceramic Capacitors

To meet power supply designers' expectations of improved performance, ceramic capacitor manufacturers have aggressively increased the amount of capacitance available within a given case size (0402, 0603, etc.). This improvement in capacitor technology has required an increase in the number of dielectric layers (>300) and a decrease in the dielectric thickness¹ (see Figure 14). The rate of this decrease has exceeded the reduction in rated voltage, resulting in a significant rise in the electric field across the dielectric.

¹ What is Happening to the Long Term Life of MLCCs?, Clive R. Hendricks, Yongki Min, Tim Lane, Venkat Magadala, CARTS 2010, New Orleans, LA

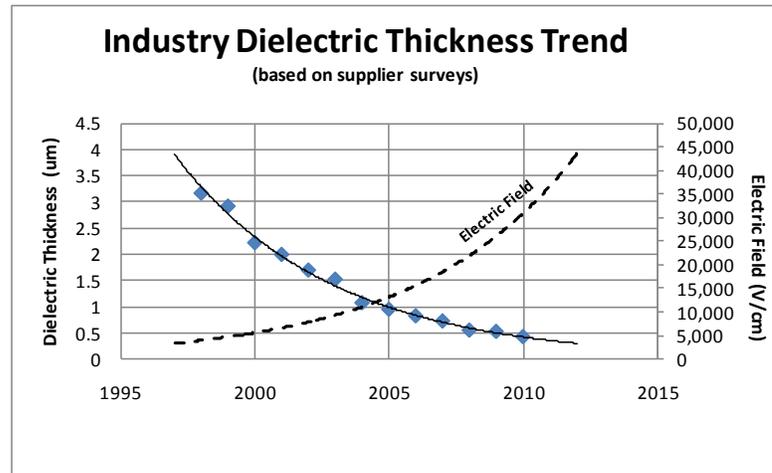


Figure 1

Capacitor processing can introduce vacancies into the BaTiO₃ crystal structure. When ceramic capacitors with high capacitance/volume (C/V) ratios, such as 10uF in an 0603 case size, are placed under a DC electric field, these vacancies will migrate to the cathode/ceramic interface and induce a drop in insulation resistance. This wearout behavior is typically modeled by an Eyring Relationship.

Equation 1
$$\frac{t_1}{t_2} = \left(\frac{V_2}{V_1} \right)^n \exp \frac{E_a}{K_B} \left(\frac{1}{T_1} - \frac{1}{T_2} \right)$$

where t is time to failure, V is voltage, n is a voltage exponent, E_a is activation energy, K_B is Boltzmann’s constant, and T is temperature (in K). While this equation is currently the standard for life prediction, variations in the constants prescribed by various organizations suggest underlying behavior and attributes that are not currently being captured. Table 2 lists examples of some of the variations in parameters identified by DfR Solutions.

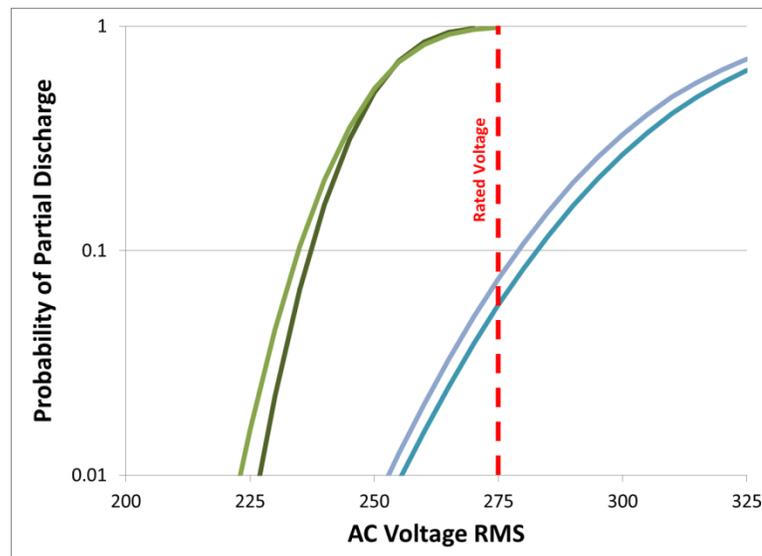
Table 2

Organization	Voltage Exponent, n	Activation Energy, E _a (eV)	Comments
DfR	2.5	0.9	Based on case studies with clients
Panasonic	3	0.31	Roughly equivalent to 2X / 15C
Murata	3	0.57	Roughly equivalent to 2X / 8C

So, how important is it to keep this tiny, ubiquitous component cool? A combination of accelerated testing by DfR Solutions and use of the Murata constants in the table above determined that 0603 / 10uF / 6.3V / X5R capacitors operated at 40C and 3.3VDC could see 2% failure after 10 years. It may not sound like much, but count up all the capacitors in your design and it starts to add up to a real problem.

Film Capacitors

Film capacitors can fail by one of two failure mechanisms, both of which are sensitive to case temperatures. The first is partial discharge (also known as corona). This is the self-healing mechanism of film capacitors, but excessive levels can destroy large areas of the polyester dielectric, resulting in a drop in measured capacitance. Partial discharge will typically initiate at voltages below the rated voltage (see graph below for several 275V film capacitors), but the occurrence drops off dramatically with lower applied voltages.



The other failure mechanism is embrittlement of the dielectric material. Most of the common dielectric materials undergo a slow aging process by which they become brittle and are more susceptible to cracking. The higher the temperature, the more the process is accelerated.

Unfortunately, there are no good formulas that clearly separate the effect of the two distinct failure mechanisms on film capacitor lifetime. Instead, the typical approach for predicting lifetime is to extrapolate from the standard IEC 60384-16 endurance test (85C and 1.25Vr for 2000 hours) using the following equation.

$$\left(\frac{Voltage_{test}}{Voltage_{field}} \right)^7 \times 2^{\left(\frac{Temp_{test} - Temp_{field}}{10} \right)} \times Test\ time$$

Of all the capacitor technologies, the lifetime of film capacitors is by far the most sensitive to variations in voltage. Because of this behavior, expert designers are typically willing to allow film capacitors to get a

little hotter than electrolytics or ceramics because sufficient voltage derating can extend lifetime sufficiently for most applications.

IC Wearout

Because complex integrated circuits within their designs may face wearout or even failure within the period of useful life, it is necessary to investigate the effects of use and environmental conditions on these components. The main concern is that submicron process technologies drive device wearout into the regions of useful life well before wearout was initially anticipated to occur.

Development of these critical components such as microprocessors and microcontrollers have conformed to Moore's Law, where the number of transistors on a die doubles approximately every two years. This trend has been successfully followed over the last two decades through reduction in transistor sizes creating faster, smaller ICs with greatly reduced power dissipation. Although this is great news for developers of high performance equipment, a crucial, yet underlying reliability risk has emerged. Semiconductor failure mechanisms which are far worse at these minute feature sizes (tens of nanometers) result in shorter device lifetimes and unanticipated early device wearout, particularly as a function of increased thermal stresses. See Figure 4.

Four failure mechanisms drive this issue. These mechanisms are Electromigration (EM), Time Dependent Dielectric Breakdown (TDDB), Hot Carrier Injection (HCI) and Negative Bias Temperature Instability (NBTI). Each of these failure mechanisms is driven by a combination of temperature, voltage, current, and frequency.

The ability to analyze and understand the impact that thermal effects have on these failure mechanisms and device reliability is necessary to mitigate risk of system degradation, which will affect measurements being taken by the system, and even cause early failure of that system. Physics-of-Failure (PoF) knowledge and an accurate mathematical approach which utilizes semiconductor formulae, industry accepted failure mechanism models, and device functionality can access reliability of those integrated circuits vital to system reliability.

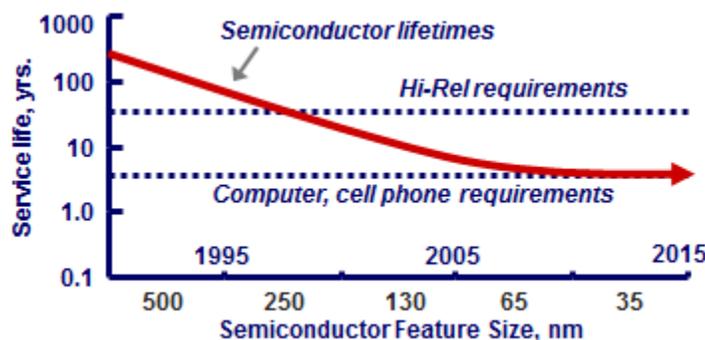


Figure 4 – Service Life versus Gate Geometry

Solder Joints

Solder joints, also known as interconnects, provide electrical, thermal, and mechanical connections between electronic components (passive, discrete, and integrated) and the substrate or board to which it

is attached. Solder joints can be a first level (die to substrate) or second level (component package to printed board) connection.

During changes in temperature, the component and printed board will expand or contract by dissimilar amounts due to differences in the coefficient of thermal expansion (CTE). This difference in expansion or contraction will place the second-level solder joint under a shear load. This load, or stress, is typically far below the strength of the solder joint. However, repeated exposure to temperature changes, such as power on/off or diurnal cycles, can introduce damage into the bulk solder. With each additional temperature cycle this damage accumulates, leading to crack propagation and eventual failure of the solder joint.

The failure of solder joints due to thermo-mechanical fatigue is one of the primary wearout mechanisms in electronic products, primarily because inappropriate design, material selection, and use environments can result in relatively short times to failure.

What are the drivers for Thermo-Mechanical Solder Joint Fatigue? Thermo-mechanical solder joint fatigue is influenced by maximum temperature, minimum temperature, dwell time at maximum temperature, component design (size, number of I/O, etc.), component material properties (CTE, elastic modulus, etc.), solder joint geometry (size and shape), solder joint material (SnPb, SAC305, etc.), printed board thickness, and printed board in-plane material properties (CTE, elastic modulus).

Finding a Solution

While the issues are clear, power supply designers have had a challenging time resolving these issues due to the lack of effective tools. As discussed earlier, the majority of tools are either focused on predicting temperatures (but, not their consequences) or are so simplistic, such as derating or MTBF, they miss the real temperature-driven risks.

A more viable approach is to utilize reliability-based physics, where degradation behavior is predicted and tradeoff analyses are performed using validated algorithms that use environmental, material, and architectural information to provide accurate guidance and prediction on the performance of the power supply. This approach, while a best practice, has typically been road-blocked by the lack of relevant tools. DfR Solutions has recently released Sherlock Automated Design Analysis™ software to help rectify this gap in design tool capability.

Sherlock uses physics of failure (PoF) to help the power supply engineer clearly understand when hot is too hot. Standard design information (ODB, Gerber, Net List, etc.) is combined with comprehensive embedded databases to provide the inputs necessary to perform these complex calculations. Streamlined software architecture ensures that thousands of these calculations are completed and results displayed within a matter of minutes. Most important, this type of analysis can now be performed by the design team well before any prototyping activity.

By using Sherlock and being part of the modeling and simulation revolution, power supply engineers can now be sure that THEIR hot will never be TOO hot.