WHITE PAPER
FLEX POWER MODULES

Peak Power Load Management in DC/DC Converters

DC/DC converters can be specified with limited peak power ratings to reduce costs and save space. In order to avoid stress, DC/DCs must have secure over-temperature immunity. This white paper discusses some existing approaches and a new proposal to achieve effective protection.

Table of Contents

Many computing tasks have a high ratio between average and peak activity. An extreme example might be a battery-powered IoT node that idles with microamps of current draw, increasing to hundreds of milliamps during a wireless transmit/receive cycle. At the other end of the scale, servers in AI and hyperscale computing in data centers could now consume more than 30 kW per rack during peak loading, according to industry body AFCOM (1), although with processing demands arriving through 'the cloud', actual load on processors and GPUs is hard to predict. In telecom applications, peak power demand is also emerging as an increasingly important consideration, driven by Time Division Duplex operation (TDD). If the server hardware and its cooling systems are configured to handle the possible peak load continuously, with no bottlenecking and with acceptable temperature rise, there would inevitably be underutilization for some of the time, with its associated costs. The effect has been long-appreciated and mechanisms have been put in place to use servers more efficiently. For example, 'server virtualization' allows workload to be spread more evenly across data center hardware, avoiding individual server under- or over-utilization. In contrast, a new approach is to segregate servers into application types. For example, high precision calculations required in scientific research might require compute-intensive 64-bit, double-precision floating point calculations, while 16-bit, half-precision can be adequate for applications such as image processing with deep neural networks, according to NVIDIA (2). In this case, if the applications are routed to servers with their power supplies and cooling systems configured and scaled to match, efficiency and utilization are improved. For a solution at the hardware level, in 2008 Intel® introduced their 'Turbo Boost' technology, which allowed processors to be 'over-clocked' above their rated frequency, as long as power, current and temperature specification limits were not exceeded. The technique can be dynamic, responding to demand and goes hand-in-hand with dynamic voltage scaling (DVS) where the supply voltage to ICs is raised to allow higher clock speed for higher throughput or lowered during periods of reduced activity, with a corresponding decrease in supply current and power dissipated.

Rating of power supplies and cooling for peak loads can allow cost and size savings

Although constant, controlled loading of servers could enable them to operate more efficiently in the computational sense, there are system advantages in allowing high peak to average operating modes. If the power supply and cooling arrangements can also be scaled for the average load but with a peak capability, they can potentially be smaller, lighter and lower cost. Critical parts are the DC/DC converters that run from a system bus at typically 12 V or increasingly 48 V and which in turn provide power rails for 'Point of Load' (PoL) converters. These output the final supply voltages required by the server processor ICs, often sub-1 V.

As with processors, GPUs and FPGAs etc., there is a thermal limit to peak power operation in the design of a DC/DC converter and this in turn depends on thermal resistance to the cooling system, starting ambient temperature, allowable temperature rise, peak load value, and its duration and repetition rate. These parameters need definition in a system and manufacturer Flex Power Modules (3), for example, identifies these ratings for their DC/DC converters:

- **'Thermal Design Power' (TDP)**, the continuous power available for a given cooling regime to keep a specified thermal hotspot below typically 125°C. This point may or may not be accessible for a temperature probe as it is typically a MOSFET solder joint. The design of the converter is such that this also corresponds to an acceptable temperature for the PCB. This is also 125°C to give a 5°C margin to its 130°C rating. This is applicable to all products and the actual TDP value will depend on the application airflow and its temperature. When operated within its TDP, a converter will be accurately characterized for reliability and lifetime through exhaustive analysis and testing.

- **'Peak Power' (PP)**, a higher peak, but lower average rating than TDP for a specified limited time, typically less than one second, with a cool-down period. This could be around 10 seconds, with the same cooling regime as TDP. E.g., Flex Power Modules products BMR491, BMR492, BMR350, BMR351. PP rating is not specified over around one second, as this approaches the thermal equilibrium time and TDP rating. Peak power may be taken repeatedly and Flex Power Modules for example characterizes their parts with no less than ten seconds between each event. Reliability and lifetime are inevitably reduced when a DC/DC is operated repeatedly to its peak power level, with exact reductions difficult to evaluate, depending on repetition rate, peak load level, starting temperature and more factors. Repeated temperature swings also cause mechanical stress to components but DC/DC product qualification tests give a confidence level that this stress is not excessive.

- **'Transient Peak Power' (TPP)** rating, an even higher peak power rating typically for <<100 ms with local power dissipation and heat stored in the thermal capacitance of the interface to the cooling system, which itself has a lesser effect. E.g., Flex Power Modules products BMR313, BMR314. TPP is most relevant to very power dense applications, where thermal time constants are short, often utilizing unregulated converters. Again, effects on reliability and lifetime are difficult to quantify but a combination of simulation and measurements can give an indication.

The progressively higher currents during peak and transient-peak power events cause exponential dissipation increase due to the resistive characteristic of semiconductors utilized, because the 'square law' applies. For example, a transient increase in loading from 2.6 kW to 3 kW is only 15% more, but the proportional increase in current through the on-resistances of semiconductors could produce nearly a third more 'I^2R' dissipation. Worse, on-resistance increases with temperature for MOSFETs, so the dissipation is higher still. This makes effective thermal protection a vital consideration during peak loading.

Figure 1 shows the typical relationships between the different power ratings. DC/DCs with peak power capability have been available for some time, for example the Flex Power Modules BMR458 introduced in 2017, but the industry now demands higher peak/average ratios, forcing system designers to focus more on this performance metric in power converters.

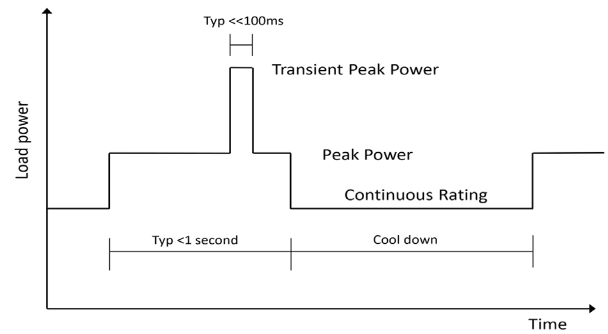

Figure 1: Continuous thermal, peak power and transient peak power limits for DC/DC converters

Design limitations to DC/DC converter peak power capability

DC/DCs can be designed with or without isolation and with or without regulation. The type chosen will often depend on system configuration, for example whether there is already required isolation upstream or downstream and whether downstream PoLs can accept a wide, varying input from an unregulated bus converter. The choice however also affects the converter design topology and implementation and its inherent peak and transient power capability. For example, if we compare state-of-the-art designs from Flex Power Modules, with similar input and output voltages, their regulated, non-isolated BMR350 achieves 1300 W thermal design power/1700 W peak power in a standard quarter brick size. This represents 66 W/cm^3, whereas the recently introduced BMR313, as a 1 kW continuous, un-regulated, non-isolated part, achieves nearly 15x peak power per cubic centimeter rating of the BMR350 at 942 W/cm^3, in its much smaller footprint.

There are various reasons why the different conversion topologies suit provision of peak power to a greater or lesser extent. For example, regulated converters such as the Flex BMR350 with pulse width modulation require output inductors which store and release energy each switching cycle to maintain a DC output and have a limiting current beyond which magnetic saturation occurs, components are stressed and functionality is lost. The effect is on a cycle-by-cycle basis **(Figure 2)** rather than cumulative, such as a thermal limit and if the inductor is designed to not saturate at a high peak load, it would be impractically large.

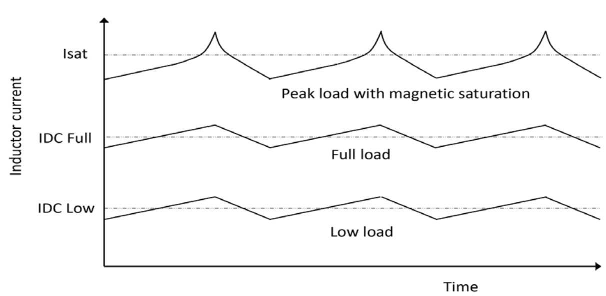

Figure 2: Output inductors in regulated DC/DCs saturate at excess current levels

Unregulated converters that work at near 100% duty cycle, with little or no requirement for storage inductance can avoid the problem if a varying output is acceptable and for a fixed down-conversion voltage ratio, a transformer is normally included. This is not however a significant limit to peak power as load current does not directly cause saturation of transformers in the 'forward' converter topologies typically used.

Another topology used is the 'LLC' converter which is a resonant type which is also switched at near 100% duty cycle. If regulation is required, the switching frequency is varied so that the converter operates up and down the slope of the 'tank' resonance curve, effectively varying the voltage gain of the power stage. If regulation is not required, the circuit can operate at fixed frequency, exactly at resonance. At this point, gain is independent of load current and efficiency is highest, on the border between zero-voltage and zero-current switching. For down-conversion, a transformer then scales the voltage in the desired ratio and this component can be relatively simple if input-output isolation is not required, meaning agency-rated insulation and separations are therefore not necessary. Again, there is no output choke to saturate at peak load, and the resonant inductor is typically anyway small. This is often formed from transformer leakage inductance, which by definition is not coupling into the transformer core and therefore cannot saturate it. For these reasons, the unregulated, LLC converter peak power capability can be very high, set practically by thermal limits, and is the topology adopted in the Flex Power Modules BMR313 mentioned.

Given converter topologies that are not limited in peak power by saturation of magnetic components, the remaining constraints are voltage drops across devices and interconnections, and thermal, either short or long-term. If a system is designed requiring no regulation from the converter with its output inherently varying with input and load, small voltage changes due to transient voltage drops with peak currents can typically be ignored. However, thermal effects need careful consideration.

Modelling thermal performance

The simple model for heat transfer from a semiconductor junction through thermal interfaces to ambient is shown in **Figure 3**. Thermal resistances are shown as 'R' and capacitances as 'C' which affect the resulting temperatures with transient loading conditions. As the die in a semiconductor is small, the value C_{THJ-C} is the smallest of the thermal capacitances by a large margin so a dissipation step at the device can cause a rapid temperature increase locally before the downstream heat path 'catches up'. The 'macro' model of Figure 3 has weaknesses however – it does not reflect coupling of heat between different devices at the local heatsink and local ambient levels. 'Ambient' is not well defined anyway and can vary between physical locations when modelling multiple heat sources together. Also, the model might rely on thermal resistances obtained from device manufacturers, which will be specified under different conditions. An accurate model therefore can only be derived from an actual sample DC/DC converter with empirical data taken in a specified end-product configuration. Flex Power Modules has characterized the complex thermal resistances and capacitances within their products, resulting in a 'micro' thermal model which produces a version of Figure 3 with multiple parallel thermal paths for the various devices, with thermal resistances and capacitances coupling across at each stage. As the junction temperature cannot be sensed directly, these comprehensive models give a more accurate T_J figure. This enables Flex Power Modules to define peak and continuous power capability given defined maximum hot-spot temperatures, to ensure there is no undue stress on the semiconductor junctions.

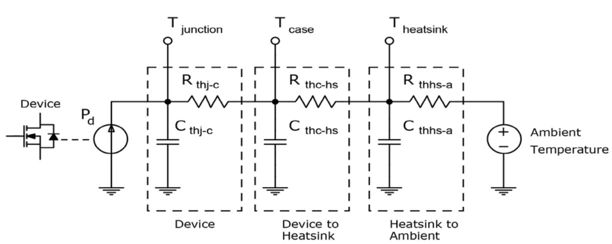

Figure 3: A simplistic thermal model for a semiconductor device with heatsinking

Thermal effects with high peak power loading of DC/DC converters

The maximum continuous load on a DC/DC converter might take the junction temperature of power components, typically MOSFETs, up to a working figure of 125°C. However, a peak power (PP) event is allowed to take this value higher, to the maximum device junction temperature of 150°C or 175°C, $T_{JUNCTION}$ in Figure 3. The time it takes to reach this limiting value is set by the transient thermal impedance of the MOSFET and as said, is usually less than one second. The load cannot be guaranteed to limit its peak power duration in all scenarios, so the DC/DC will include over-temperature protection, typically monitoring a 'hot-spot' with a known thermal resistance to the semiconductor junctions. This would be at position T_{CASE} in Figure 3. A full-function DC/DC will generate a warning signal when the hot-spot exceeds a threshold equivalent to say 142°C junction temperature. This signal can then be used to indicate to the load that it should terminate the peak power demand. If this does not happen, a second hot-spot threshold causes the DC/DC to safely shut down.

Effectiveness of protection measures depends on the DC/DC type. A converter with no peak rating and a current limit a little over its continuous rating can be effectively protected by simple hot-spot over-temperature monitoring. If a converter has a thermal rating of 1k W and a peak rating of, say, 1.6 kW, a ratio of just 1.6:1, the time for junctions to reach maximum levels are relatively long. This means that a temperature sensor can still 'track' the junction temperature relatively accurately and provide protection. This corresponds to the thermal capacitances in Figure 3 'charging' to close to their final value during the load transient. **Figure 4** shows an example of the junction and sensor temperature rises in this converter with its load stepped from 1 kW to 1.6 kW, along with its warning and shut-down thresholds.

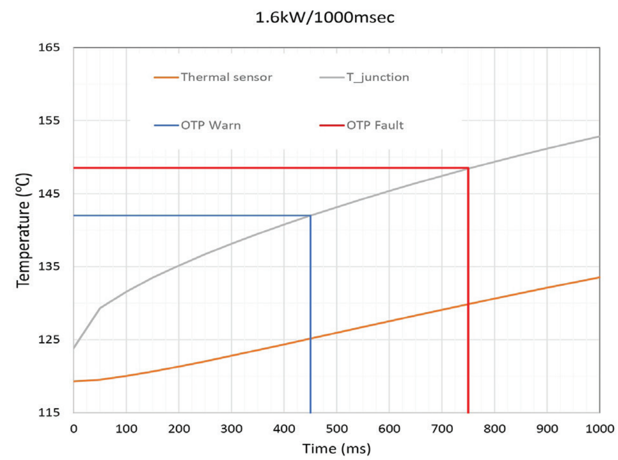

Figure 4: Typical warning and shutdown thresholds for a DC/DC converter during a peak power event

However, if a DC/DC is designed with a very high transient-peak to thermal power rating, such as the BMR313 at 3 kW/1 kW (3:1), rate of temperature rise is much higher. This means that the thermal time constant of the over-temperature sensor and its thermal resistance to the MOSFET junction prevents a warning signal and subsequent shutdown occurring before the junction temperature reaches its absolute maximum value, risking device stress and failure. This corresponds to the thermal capacitances in Figure 3 not reaching full 'charge' during the load transient and delaying heat transfer through the chain. **Figure 5** shows this scenario with the temperature sensor rising at around the same rate for both a step load to 2 kW and 3 kW while junction temperature increases at a higher rate, to a potentially damaging level, due to the delaying effect of C_{THC-HS} in Figure 3. Note that it is assumed that the critical semiconductor junction starts at 125°C. At lower initial temperature, the time before temperature limits are exceeded can be much longer and real-life margins can be much higher. This is, of course, true of all DC/DC types.

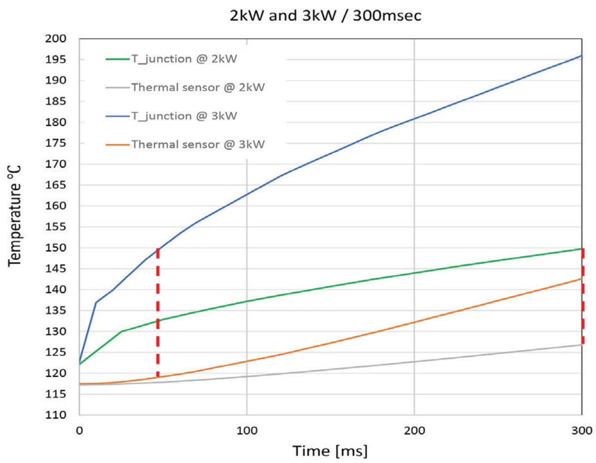

Figure 5: Over temperature protection can be ineffective with high peak power transients

To address the issue of providing more accurate protection thresholds, while maximizing the transient power time available for a high peak-to-thermal rated converter, a different approach can be taken. A junction temperature can be predicted from the starting temperature, output current value and duration. These can be correlated when the electrical and mechanical arrangement of the DC/DC converter is accurately simulated in multiphysics software. This can be a relatively simple calculation, as the peak temperature from a high transient peak load is largely affected by just the local MOSFET die thermal characteristics –the complex thermal interactions further down the heat path can be ignored. The method can be effective for peak power loading with low duty cycle, but when the repetition rate is higher, the slower 'cooling down' period may be insufficient to ensure that a subsequent peak load does not cause the junction maximum temperature to be exceeded. This 'ratcheting effect' can be seen in **Figure 6** where a single short load pulse is effectively just 'noise' in the system whereas a string of pulses can quickly add up towards a stress level with the temperature sensor 'lagging' behind. Effective protection should ideally detect this effect.

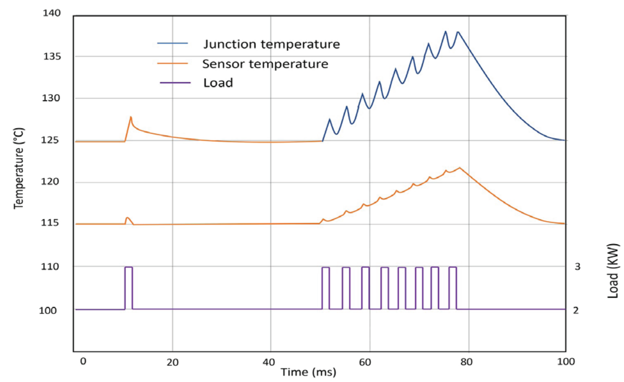

Figure 6: A single load pulse can produce negligible temperature rise whereas close strings of pulses are additive.

Monitoring of the hot-spot temperature to provide an alarm is less accurate due to the multiple, complex thermal time constants in play during cooling, from junction to package, package to DC/DC module and DC/DC to external cooling system at its starting temperature. **Figure 7** shows an example cooling down characteristic with multiple thermal time constants indicated for a single pulse load event.

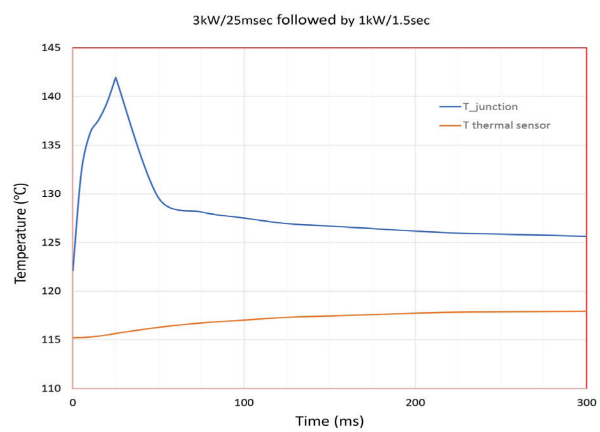

Figure 7: Cooling down of a MOSFET die after a transient peak power event showing multiple thermal time constants

Further intelligence can be incorporated into a DC/DC converter to provide more comprehensive protection, both for rapid local monitoring and shutdown, and for signaling to the load. For example, in the Flex BMR313 converter, rated 1 kW continuous, 3 kW maximum, multiple thresholds for over-temperature and timed over-current are included. At currents equivalent to 3.4 kW, a fast hardware shutdown circuit with configurable filtering operates within a few microseconds. At a current designated OC_FAULT equivalent to 2.6 kW, a slower hardware shutdown operates within one to two milliseconds. At a current level OC_WARN equivalent to 1.8 kW, an ALERT pin is asserted to request the load to reduce power draw.

When a current between OC_WARN and OC_FAULT is detected, two counters are started. The ALERT counter, which immediately asserts the ALERT pin for a configured T_{WARN} time, and the FAULT counter which will be incremented for every sample in the OC_WARN to OC_FAULT region.

If the FAULT counter reaches half of the configured T_{WARN} time, the converter will turn off. The FAULT counter will be reset after the configured T_{WARN} time and the ALERT will be de-asserted.

Below the OC_WARN level the converter has peak power protection by over-temperature sensing, with the ALERT pin asserted if the temperature is above the TEMP WARN level. Above the TEMP FAULT level, the converter will turn off. **Figure 8** shows the levels and timing.

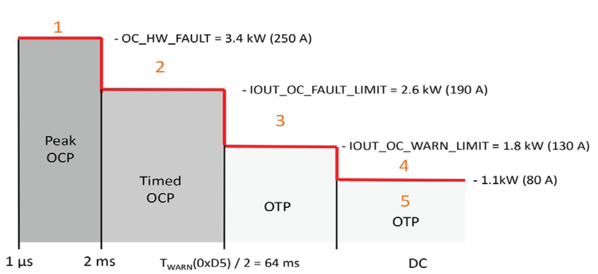

Figure 8: The Flex BMR313 continuous, peak and transient peak power protection scheme

As a worst case, the BMR313 configurable filtering and timing can be set assuming the converter is already running at its maximum MOSFET continuous junction temperature of 120°C, but longer current pulses can be allowed if the starting temperature is lower. However, this will also generate larger temperature swings, risking high mechanical stress to the MOSFET package and must be evaluated carefully to maintain reliability. It is generally accepted that wide and fast temperature swings are more stressful than high repetition rates of smaller swings, for example. **Figure 9** indicates actual transient peak power (TPP) pulse duration allowed for the BMR313 for peak power up to 3 kW for different sensor starting temperatures.

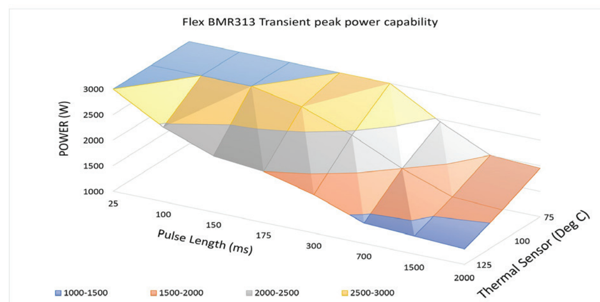

Figure 9: Transient peak power capability of the Flex BMR313

Another example of peak power protection is employed in the 1.3 kW Flex BMR350 DC/DC converter. In this design, current is monitored and filtered such that a high load transient yields a signal which is a good analog of junction temperature. The filter employed uses the 'Exponential Moving Average' (EMA) method, where instantaneous current is added to a fixed proportion of the previous evaluated EMA value **(Equation 1)**. In the equation, α is a coefficient controlling the weight of older measurements in the new value, larger values of α corresponding to faster discarding of older measurement weighting.

$$EMA_N = (1-\alpha) . EMA_{N-1} + \alpha . MEASUREMENT \quad \text{(Equ. 1)}$$

When EMA_N reaches a limit value, an overcurrent fault is registered. In the case of the BMR350, an

ALERT signal is not generated as the part follows an industry standard pin-out which does not have this provision. **Figure 10** shows the comparison of a Simple Moving Average (SMA) and EMA filter response for a step input, the EMA plot resembling a first order filter response, corresponding to the heating of a semiconductor junction. As a regulated part, the peak power is limited by magnetics to 1.7 kW so if the load is configured never to exceed this value, this simple over-temperature protection is adequate.

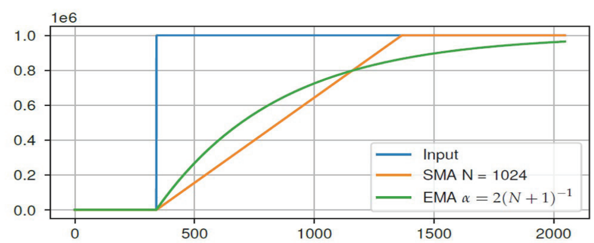

Figure 10: Comparison of Exponential and Simple Moving Average filtering of a DC/DC output current signal

The next generation of peak power protection methods

We have considered the peak power protection methods for two Flex Power Modules products and both have merits:

The BMR 313 implementation can alert the host system that it needs to reduce its power consumption, but the time in this mode is fixed, and does not adjust based on the actual current level. This implementation allows the host system to know the time from when the ALERT pin is asserted to the time the unit will shut down. In contrast, the BMR 350 can more accurately recreate the junction temperature, but no ALERT signal is available This implementation can more easily adjust the time allowed for a peak power demand based on the actual current level, but a problem is how to alert the host system when to reduce power. There is therefore a possibility that the end user is taking a higher load than they thought, close to the DC/DC limit, and a small transient extra load could cause a nuisance shutdown. There is therefore a practical trade-off between high transient peak power capability and risk of loss of power to the load.

To improve performance, a novel method being investigated by Flex Power Modules is to implement two separate Exponential Moving Average calculations, one representing the DC/DC unit average temperature and the other the MOSFET die temperature. This will allow the modelling of the cool-down characteristic of the DC/DC since longer peak power events will heat up the entire part. An ALERT signal is set to an 'active impedance' mode Z, when the DC/DC is operating in the transient peak power area and then pulled down to zero at a pre-defined time before shutdown, based on system requirements. The time is dependent on the inferred MOSFET junction temperature and the maximum allowed peak power. The technique could be called Safe Operating ARea protection (SOAR) and although it requires precise and high speed monitoring, this is facilitated by the latest digital DC/DC converter technologies. **Figure 11** shows a scenario with the monitoring responding to two incremental load steps, die temperature rise and the corresponding ALERT pin levels. The second step to the transient peak power maximum of 3 kW triggers a faster moving average generating an ALERT and eventual shutdown before the inferred maximum junction temperature is exceeded.

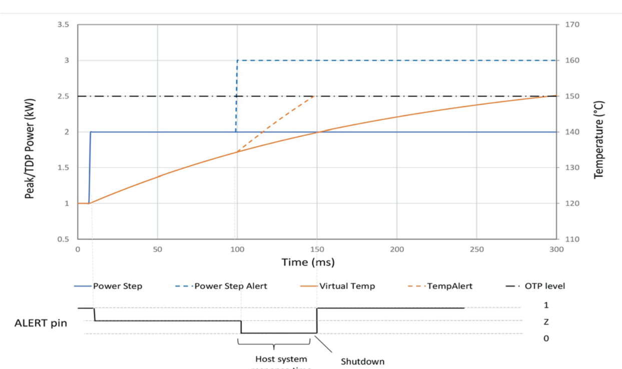

Figure 11: Dual EMA response to peak power monitoring proposed by Flex Power Modules

Conclusion

Ability to provide peak and higher peak transient load current is a valuable attribute of DC/DC converters to maximize power density and minimize cost in modern data centers. However, high loads cause rapid heating in compact applications and protection must be provided in the DC/DC converters against short- and long-term overloads that might take critical device junction temperatures over maximum specifications. Various methodologies have been described as implemented in Flex Power Modules products with their relative advantages, along with an approach for the future which promises more accurate protection, to enable parts to run closer to their rated limits without degrading reliability.

References

(1) https://www.networkworld.com/article/3454626/8-ways-to-prepare-your-data-center-for-ai-s-power-draw.html

(2) https://developer.nvidia.com/blog/mixed-precision-training-deep-neural-networks/

(3) www.flexpowermodules.com

Connect with us

pm.info@flex.com

flexpowermodules.com

flexpowerdesigner.com

linkedin.com/showcase/flex-power-modules

twitter.com/flexpowermodule

flexpowermodules.com/wechat

youtube.com/flexintl