flex.
Power Modules

Peak Power and Transient Peak Power

Abstract

This Application Note describes the basic principles of Peak Power and Transient Peak Power, which are definitions used by Flex Power Modules. For clarification we will provide some examples of different product families where these definitions are used. This AN applies for the following products :

- **BMR458**
- **BMR480**
- **BMR491**
- **BMR492**
- **BMR350**
- **BMR351**
- **BMR352**
- **BMR313**
- **BMR314**
- **BMR316**

Contents

Background

Many processing tasks have a high ratio between average and peak activity, which therefore affects the power consumption. There are numerous different scenarios ranging from IoT nodes that idle at microamps of current, increasing to hundreds of milliamps during a wireless transmit/receive cycle, to power hungry servers used in AI which consume more than 30kW per rack during peak loads.

If the server hardware and its cooling systems are configured to handle the possible peak load continuously, with no bottlenecks and with an acceptable temperature rise, there would inevitably be underutilization for much of the time, with its associated costs.

Terminology

As with processors such as GPUs (Graphical Processor Unit) and FPGAs (Field Programmable Gate Array) etc., there is a thermal limit to the peak power operation in the design of a DC/DC converter, and this in turn depends on several factors including the thermal resistance of the cooling system, the initial ambient temperature, the allowable temperature rise and the required peak load value including its duration and repetition rate. These parameters need to have a clear definition within a system, and manufacturers like Flex Power Modules identify these ratings for DC/DC converters as follows:

- **Thermal Design Power' (TDP**) is the continuous power available for a given cooling regime to keep a specified thermal hotspot below a maximum temperature, typically 125°C. This particular 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 max rating. This is applicable to all products and the actual TDP value may be limited by the application airflow and its temperature. When operating within its TDP, a converter will be accurately characterized for reliability and lifetime through exhaustive analysis and testing.

- **'Peak Power' (PP)** is defined as a higher peak value for a specified limited time (typically less than one second) that yields a lower average rating than TDP, and includes a cool-down period. This could be around 10 seconds, with the same cooling regime as TDP. Product examples include the [BMR491](#), [BMR492](#), [BMR350](#), [BMR351](#) and [BMR352](#). Most applications use PP for shorter bursts, but the maximum time is set with the cooling properties of the DC/DC converter in mind. In many cases PP can be applied for up to 1 sec without any significant local hot-spots (for instance a MOSFET junction running hot due to excessive power loss). Peak power may be taken repeatedly, but in that case several shorter events must be accumulated in effect to one overall event when calculating the total time. Longer periods exceeding the maximum length need appropriate cooling time between events. In this type of situation temperatures are allowed to exceed 125°C, since usually it is the MOSFETs that contribute the most to the temperature rise, and these are usually rated at 150°C. Even starting with a maximum hotspot temperature of 125°C, drawing the full PP amplitude for the maximum allowed time will still not cause the MOSFET to reach a 150°C junction temperature.

Figure 1: BMR352

- '**Transient Peak** Power' (TPP) is 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. Examples include Flex Power Modules products such as [BMR313](#), [BMR314](#) and [BMR316](#). TPP is most relevant to very power dense applications, where thermal time constants are short, often utilizing unregulated IBC converters.

Figure 2: BMR316

Power loss

From a user perspective the output power is the interesting number. However, from an understanding of what drives the limitations for peak power, it is the power loss in the DC/DC converter itself that is more important.

The power dissipation in the load is linear with the peak demand, so if the load increases by 50% the heat dissipation will also be 50% larger.

For the DC/DC it is a different scenario since losses are proportional to I^2, so when the load increases by 50% the power dissipation increases by an additional 125%, and so twice the load will result in 4 times higher losses. This puts extreme demands on the thermal design when the amplitude of the peak increases. It also explains why the thermal resistance plays such an important role.

Thermal resistance and capacitance

A typical 'brick' type DC/DC converter (a quarter brick for instance) has a thermal time constant in the region of 5-10 seconds, which means that it takes at least this time to reach a thermal equilibrium. When connected to an application board and heatsink, more thermal capacitance (mass) is added, and the time constant will be even longer. Bear in mind that thermal equilibrium may occur at a significant over-temperature, and hence the maximum time must be reduced.

The heat generated will be a very localized hot spot that needs to be dissipated to larger areas for cooling. When running PP the thermal gradient across the product increases. In most cases the thermal resistance is so low that even the hot spot can get decent cooling during this situation. Theoretically, it might even be possible to keep PP for longer times if the ambient temperature is low enough.

In transient situations however, where the power loss is so great that the thermal resistance is not sufficient, this will limit the PP duration dramatically. When the time gets very short it can itself be referred to as a transient and referred to as a "Transient Peak Power" (TPP) event.

During a TPP event most of the heat energy is absorbed locally to the heat-source, such as the silicon-die, component lead-frame etc. When the load is subsequently lowered it can be dissipated to the cooling system.

TPP is a thermal trap situation. Exactly where and when this happens depends on several parameters. Lower initial temperatures create both a possibility for more heat absorption and a greater heat flow from the component due to higher Δt to the cooling mechanism surrounding it. Therefore, the TPP time varies considerably with both initial temperature and load.

It is almost impossible to measure the temperature in the MOSFET in this situation, mostly because things happen so fast (it is difficult to monitor temperature accurately in timeframes of 10msec, especially inside a MOSFET).

Models

In order to calculate TPP time, we create simulations of what happens close to the limiting components. For PP, simulations are also made but they are expanded to cover thermal resistance in the complete DC/DC converter. In this way both heat-trap situations and steady-state flow will be observed. Steady-state performance is confirmed by actual measurements, and the models of the components are verified by the suppliers.

The most difficult parameters to get correct are often the load and the ambient temperature, both of which are set by the application. These are critical parameters for correct calculation of TPP duration.

Reliability

This application note will only briefly mention the reliability topic. Current as such is not a large contributor to ageing, but heat is a great contributor. Running maximum peak duration and maximum amplitude will cause a local temperature swing, whereas very short transients are merely a "thermal noise". Temperature swing is also an ageing accelerator. Therefore 10 instances of 100 msec pulses with cooling time in between causes less stress than a single 1 sec pulse with the same amplitude.

To guarantee high reliability, many standard tests are performed, and usually the converter reaches steady state at each step in the test cycle - tests that toggle between off and TDP power for instance. For peak power conditions, other methods must be used which include time limitations before steady-state is reached.

BMR491

As an example, the BMR491 underwent the following test sequence:
The DC/DC converter was set to operate at 90% of TDP as a baseload (90% of 1540W). A peak was then applied for 1 second at the maximum power (2450W). Then the converter needs to cool at the 90% level for 9 seconds. The thermal reference point was kept just below the maximum temperature for "T_{P1}". The test was stopped after >300,000 cycles, which takes 7 weeks to perform. Statistical models can then be applied for re-calculation in less demanding situations.


Figure 3: Peak cycle test

Protection mechanisms

Creating an overload protection is a fine line. It must provide a decent protection from failures, but it should not accidently generate a shut-down of the system. The details below show how specific products handle this in different ways. They all have their differences and reasons behind the selection that has been made.

BMR313 / BMR314 /(Peak Power & Transient Peak Power)

The unregulated BMR313/BMR314 and BMR316 allow very high TPP levels. The control circuit has advanced algorithms for calculating the die temperature in the GaN FETs (BMR313) or Silicon MOSFETs (BMR314 and BMR316). When the temperature gets high enough, the converter sends a "Throttle down signal" via the Alert pin to the load.

When this signal is sent, the load must be reduced with highest priority to avoid a shutdown, which will occur if the signal is neglected.

It is not necessary to monitor this signal, but if the full TPP potential shall be utilized it is recommended.

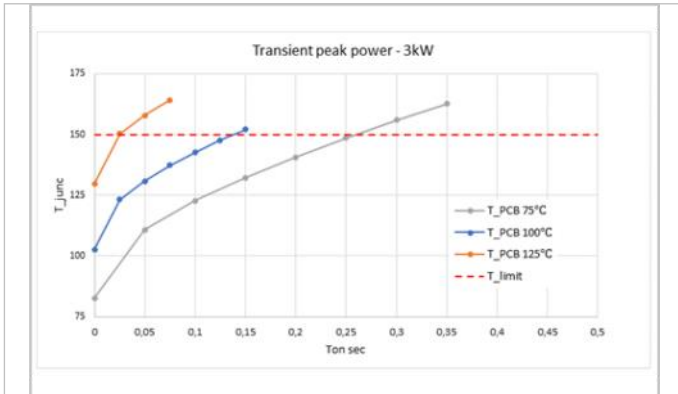This shows how the converter interacts with the load in an advanced way.

Figure 4: Transient peak power

BMR491 / BMR492 / BMR480 / BMR458 / (Peak Power)

The regulated IBCs with Peak Power have a conservative maximum power compared to the latest unregulated devices. The slow reacting thermal protection will protect the device since the internal thermal gradients are not allowed to become too large even at peak power. The internal thermal resistance has a low value in relation to the power loss.

The 1 second length should still be respected, and the converter will not actively shutdown or limit the peak power if the T_{P1} temperature is within its limits.

BMR350 / BMR351 / BMR352 (Peak Power)

The regulated IBCs BMR350, BMR351 and BMR352 have a so-called average current based protection. A rolling average of the output current is stored in a register, and when this register exceeds approximately 115% of IO(-TDP) the converter enters into a protection mode and shuts down.

When this happens, the product has been operated in the PP region for an excessive time.

This prevents the device from being operated close to the over current protection (which is set

above Peak Power, sometimes referred to as C-B-C, Cycle-by-Cycle) for extended periods - which causes local hot spots.

The benefit with average current protection is that both the OCP (C-B-C) and the OTP limits can be set at less restrictive levels since this type of protection does not allow a hot spot to be very dominant. It prevents a situation where the thermal gradient across the converter becomes very large due to high overload.

This type of protection does not operate with a fixed time - the higher the peak loads the quicker the protection activates. It also takes earlier readings into account, balancing just below the threshold which will cause the protection to activate quickly if full PP is requested. A long period of low output current will allow extended PP duration.

Summary

This application note has attempted to explain both the technical limitations of peak power delivery within our devices, and how different products perform this functionality in practice.

The trend however is very clear - average power (TDP) and maximum power (PP/TPP) requirements must be clearly differentiated. A key requirement going forward is how the converter should communicate with the load, and then how it should calculate the stress. The methods used in BMR 313/314/316are very attractive, but they require loads that can also accept the DC/DC communication to be completely effective. This is simply not practical (yet) when there can typically be one IBC powering dozens of different load types.

For this reason there will continue to be different ways employed to handle the peak power demands, and Flex Power Modules will continue to offer alternative solutions across a number of different platforms.

flex. Power Modules

Flex Power Modules, a business line of Flex, is a leading manufacturer and solution provider of scalable DC/DC power converters primarily serving the data processing, communications, industrial and transportation markets. Offering a wide range of both isolated and non-isolated solutions, its digitally-enabled DC/DC converters include PMBus compatibility supported by the powerful Flex Power Designer.

EMEA (Headquarters) | Torshamnsgatan 28 A, 16440 Kista, Sweden

APAC | 33 Fuhua Road, Jiading District, Shanghai, China 201818

Americas | 6201 America Center Drive, San Jose, CA 95002, USA

✉ pm.info@flex.com   💬 flexpowermodules.com/wechat

🌐 flexpowermodules.com   in linkedin.com/showcase/flex-power-modules

🌐 flexpowerdesigner.com

▶ youtube.com/flexintl