The Challenging Task of Thermal Management

Martin Schulz, Infineon Technologies, Germany, martin.schulz@infineon.com

Abstract

The constant increase of power densities inside of power electronic components is an omnipresent trend that in future is going to continue. One of the consequences related to this trend is a constant increase of junction temperatures in power electronic devices. With this, the thermal management becomes a more important design criteria. The present paper deals with several aspects of thermal management, focusing on the often neglected influence of building a proper interface between power module and heat sink. Measured data of an ongoing research is presented concentrating on applying and testing adequate thermal interface materials.

1 Introduction

From the first days of IGBT-Module technology to the present, chip sizes for a given current carrying capability of silicon have constantly been decreased. Vice versa, the current density inside the silicon die and therefore the power modules has grown. Smaller chips also allowed the integration of more silicon into one device, often doubling the current a given housing could handle. Though the losses per amp in every new chip generation have also decreased, the higher power densities lead to increased junction temperatures as well. While in the year 2000 a single IGBT 1200V/1800A had a footprint of 130mm x 140mm [1] current ratings of up to 900A are available today in a package 62mm x 106mm [2]. Regarding the module, this is an increase of current density from 6.6A/cm^2 to 13,7 A/cm^2, a factor of about 2. However, the beneficial increase in performance is often wasted by underestimating the thermal design. From this point of view it becomes obvious, that the thermal transfer from the chip's junction to the ambient gets more important.

2 Simplified thermal model

To evaluate the options in increasing the thermal performance of a given power electronic component, a simplified model as depicted in figure 1 becomes helpful.

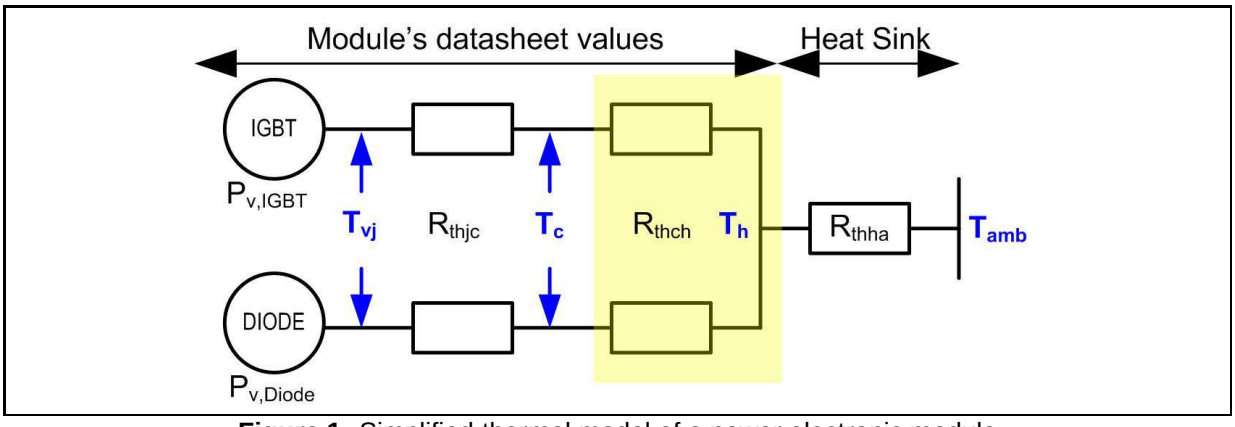

Figure 1: Simplified thermal model of a power electronic module

The model includes the two sources for heat, the IGBT and the diode die. The junction as a source of the heat reaches a given junction temperature T_{vj}. The heat flow crosses the junction to case path R_{thjc} that is defined by the internal module construction. A later section describes ongoing developments that will enhance the thermal performance in this area as well.
Up to this point, the module's construction is responsible for the thermal transfer and therefore defines the thermal performance.
The shaded box in figure 1 introduces the thermal path from module case to the heat sink R_{thch} that in the datasheet is considered to consist of a certain amount of thermally active compound.

At this point, the chosen thermal interface material, the application process, layer thicknesses and further parameters of the chosen material define the conditions for the thermal behavior. Though it seems trivia to use and apply this material especially careful investigation of the thermal interface material (TIM) needs to be done.
Besides the obvious ability to conduct thermal energy flow, the material needs to be long-term stable. In power electronic designs, 10 to 20 years of expected lifetime are common while traction or wind mill application being even more demanding.

3 Testing TIM

To serve the needs in the application, a thermal interface material needs to undergo an extensive qualification procedure to prove that it will

- Provide an appropriate thermal flow
- Withstand the pumping effect due to thermal mechanical movement
- Operate from the first turn-on cycle
- Operate during the predicted lifetime of the application

The tests done include all the reliability tests mandatory for power electronic modules. This way it is ensured that the thermal interface will survive all stresses that a power module is designed to handle. Furthermore, particular tests are added or derived from the standard procedures to examine the TIM's capabilities.

3.1 Cycle life Tests

A dedicated test bench shown in figure 2 was designed to get comparable results and different test candidates were stressed for several months.

Figure 2: Test bench dedicated to conduct active thermal cycling featuring space to test six candidates at a time

Pump-out of thermally active material due to thermal electrical movement reduces the amount of material between module and baseplate over the lifetime, diminishing the thermal qualities of the interface. Additionally, the effect increases the possibility that channels form inside the TIM-Layer, thus increasing the speed of ageing of the material [3]. Pump-out can best be observed in active heating cycles. In this test, modules with applied TIM are mounted on heat sinks and are actively heated by current flowing through the IGBT. A complete cycle consists of an on-time of about one minute heating up the device and a following off-time of about two minutes for cooling down. The times were chosen to get an appropriate temperature swing and at the same time allow for a stable maximum to be reached preventing thermal runaway. Though testing this way takes a long time it resembles the thermal load as close to the real world as possible.
As seen in figure 2 the tested modules were blackened allowing to record the chip temperatures using a thermographic camera. Doing so on a daily base provides an insight into ageing of the candidates under test. As long as no failure mechanism is triggered, a fair comparison is based on the averaged maximum temperature reached during the test. The maximum temperature per module was recorded and an average on the data was calculated. The result gathered in a 100.000 cycle test run is displayed in figure 3:

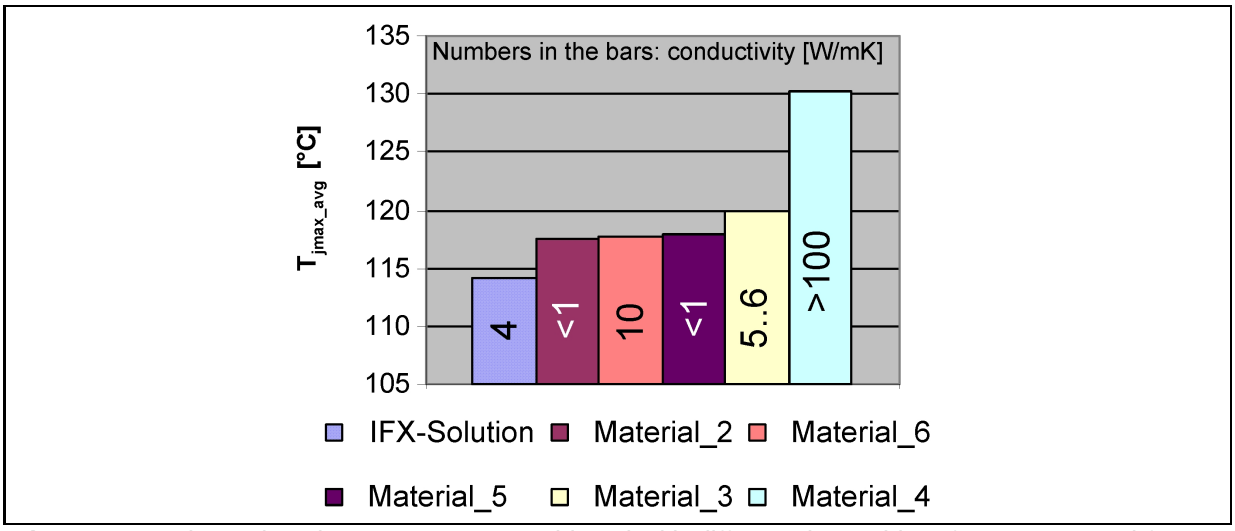

Figure 3: Maximum junction temperatures achieved with different thermal interfaces, averaged from a 100k cycle run.

The difference between the best material in the test and the worst candidate turns out to be 18K in average. Due to the exponential nature of the degradation effects, 20K lower temperature swing increases the device's lifetime by a factor 10. A further result coming from the test was, that the temperature reached and the value for thermal conductivity stated in the according datasheets had no correlation at all; the according values for thermal conductivity in $[W/(mK)]$ are printed into the bar diagram as a reference. The reason can be found in the procedure to determine that particular value. Here conditions are defined that dramatically vary from those below a power module making the datasheet value for thermal conductivity an inadequate parameter for choosing a proper material.

3.2 High Temperature storing

The previous test, done under conditions close to the real application, does not feature accelerating factors. Though the change observed for the preferred solution is a very low one, additional tests with higher stress levels were chosen to get a more comfortable statement regarding long-term stability.
For this High Temperature Storing or HTS Test, modules with applied TIM were mounted on common heat sinks. The first step was to determine the initial condition. Therefore, power cycling was done for several cycles until a steady maximum temperature on the chip was recorded. Afterwards the modules were stored in dry heat at 125^oC for one week, 168 hours respectively. Cycling and temperature recording was repeated afterwards. This procedure was repeated until 1000 hours of HTS were completed.

The results for three different candidates in the test can be seen in figure 4:

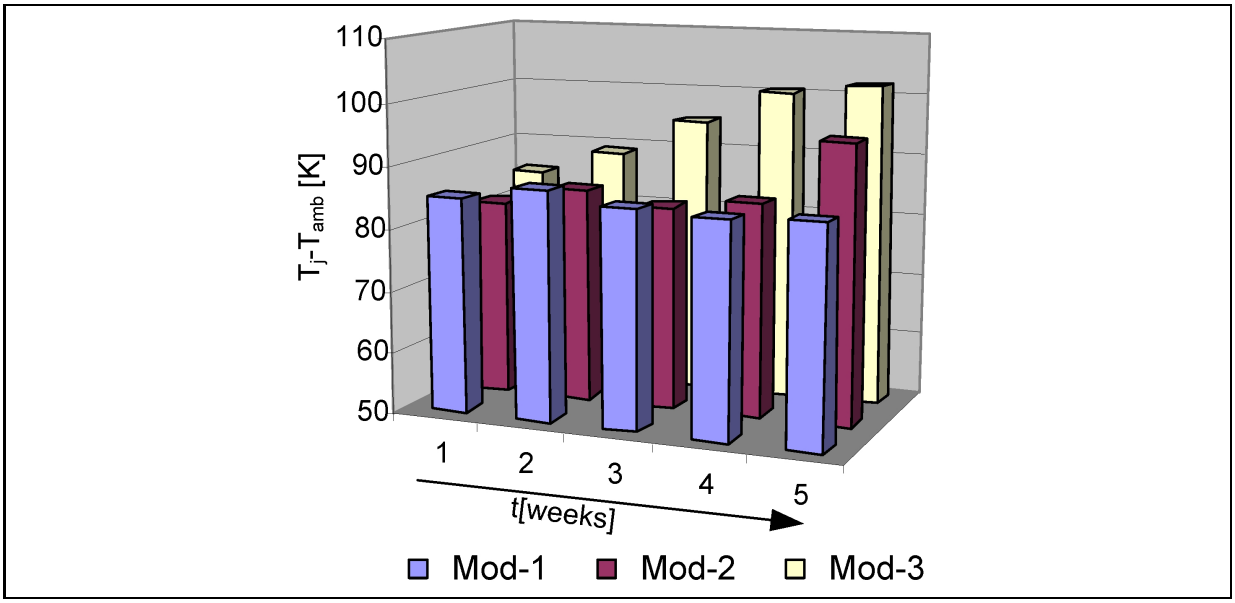

Figure 4: Results of a 1000 hour high temperature storing test showing three different types of ageing

The diagram clearly shows three different, partially unexpected results:
Modification 1 is the only candidate in the test that performs as mandatory showing no noteworthy change during the whole test.
Modification 2 performs quite well initially but shows a sudden increase in temperature after about 700 hours of testing. This is related to a multitude of effects including separation and drying of the grease; it is less a consequence of thermal mechanical stress. However, as the dominant effects were supposed to be continuous effects it was expected to measure a gradual increase in the chip temperature rather than a sudden change.
Modification 3 also performs well in initial state but here the gradual increase can be observed as expected.
The results also lead to the conclusion that long term stability can not be evaluated by simply measuring an actual value and even stressing the material for several days or a few thousands of cycles can be inconclusive.

4 Applying Thermal Interface Material

The amount of TIM below a module massively influences the resulting thermal resistance so the target has to be to apply as much as necessary but as little as possible. However, there is only a very fine line between just enough material and a too small amount.
To prevent too much material being applied to the module, manual processes should be avoided as far as possible. Using roll-on techniques, brushes or simple scrapers does not lead to a reliable and reproducible result. A first step into the right direction is the application of TIM in silk-screen or stencil screen printing. The stencil can be manufactured with thickness tolerances down to the range of μm and the hole geometry defines the volume of grease applied to a certain area. It is important to know the module's baseplate geometry in deep detail because a homogenous layer will never produce the best result. To achieve the best possible results, two things need to be taken into account:

1.) TIM only has to be applied in the cavities of the base plate. Here the TIM is needed to conduct the heat

2.) TIM should not prevent the forming of a metal-to-metal contact as the setup strongly benefits from this direct contact

The stencil presented in figure 5 gives a good impression how an optimized pattern allows the partial decrease of the amount applied to a certain area:

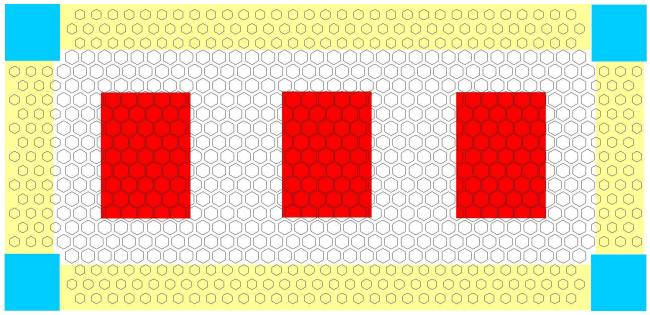

Figure 5: Stencil based on honeycomb shaped holes to locally apply varying amounts of TIM in one single step. Different colors refer to different volumes applied.

In the picture, the corner regions later carry the screws and have no TIM to get the largest area for metal-to-metal contact here. The spaces along the outer edges have low TIM content, the center regions have cavities to be filled with higher amount of TIM while the white regions are in between.
The honeycomb shape also poses a good solution to prevent air being caught during the spreading of TIM. As hinted out in [3] this would lead to accelerated ageing of TIM as well.
As mentioned, the stencil geometry needs to be tuned to the base plate geometry of a given module. This geometry depends on the initial condition of the base plate before the module is assembled, the DCB-material that is used to carry the chips, the soldering process along with the soldering temperature profile and several other parameters. As these vary from manufacturer to manufacturer it becomes clear that a stencil designed to perfectly suit a given module does not necessarily need to be the best solution for all modules with the same footprint. On the contrary, it is safe to assume the opposite making module, stencil and thermal interface material an interconnected system that has to be tested, qualified and used as a unit.
To create this unit, automated printing as shown in figure 6 is recommended. This way, squeegee speed, pressure and angle remain within the tolerances necessary for a reliable process.

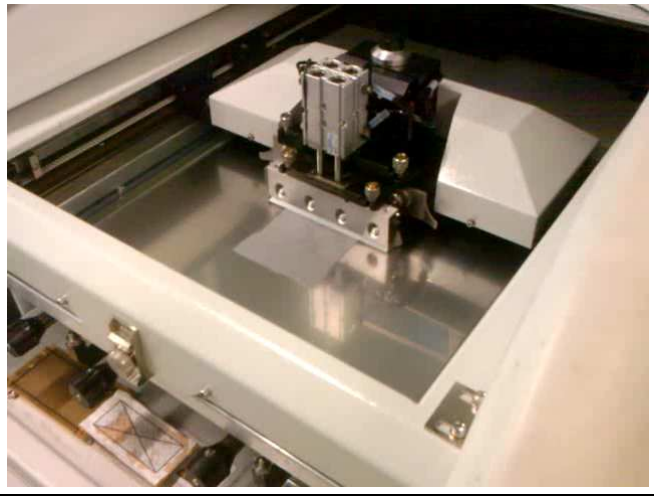

Figure 6: Automated stencil printing in process keeping all relevant parameters within the specified limits

Knowing the unexpected complexity of the topic a logic consequence is, that modules with thermal interface material preapplied to the baseplate will be available to the marked. This way, material quality, process stability and reproducibility are defined by the manufacturer easing the power electronic designers work in thermal management aspects.

5 TIM benefits

It seems to be a lot of effort for something trivia but power semiconductors are today pushed to reach every limit. In this aspect it would be short sighted to underestimate the influence of the thermal interface. Figure 7 hints out the possible benefit:

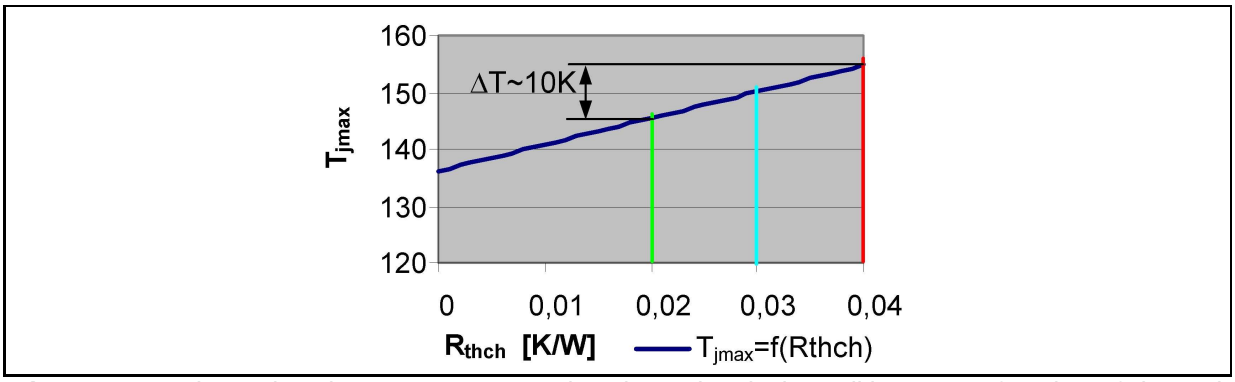

Figure 7: Maximum junction temperature under given electrical conditions as a function of thermal interface resistance R_{thch}

The line at $R_{thch} = 0,03$ represents the typical value stated in the datasheet [5] for an IGBT. While a sophisticated material in conjunction with a well aligned and controlled process of application allows to even exceed this value, less good conditions are often observed in the field. The difference can easily reach up to 10K as denoted in the picture. As the effect is of linear nature this is also true if the overall temperature limit is a lower one. On the other hand, lifetime limiting effects are of exponential nature. Reducing the junction temperature during power cycling by 10K from $\Delta T = 40^oC$ to $\Delta T = 30^oC$ means an increase in power cycling capability by a factor 10 [7]. A second option coming from improved thermal performance would be higher output power at constant temperature levels.

6 Improvement due to new technologies

From figure 1 it is obvious that other thermal resistances also have a potential for improvement. Along with reliability improvements coming with new interconnection technologies [4] the thermal resistance from the chip junction to the module's case is lowered as well. Replacing the soft soldering used today by diffusion soldering leads to more robust interconnection layers. Diffusion soldered joints also feature very low bond line thicknesses and additionally the materials used provide higher thermal conductivities. Both these properties improve the thermal situation for the silicon dies. It can be seen in figure 8 that the layer between the chip and the carrier substrate is reduced in thickness by about 90%.

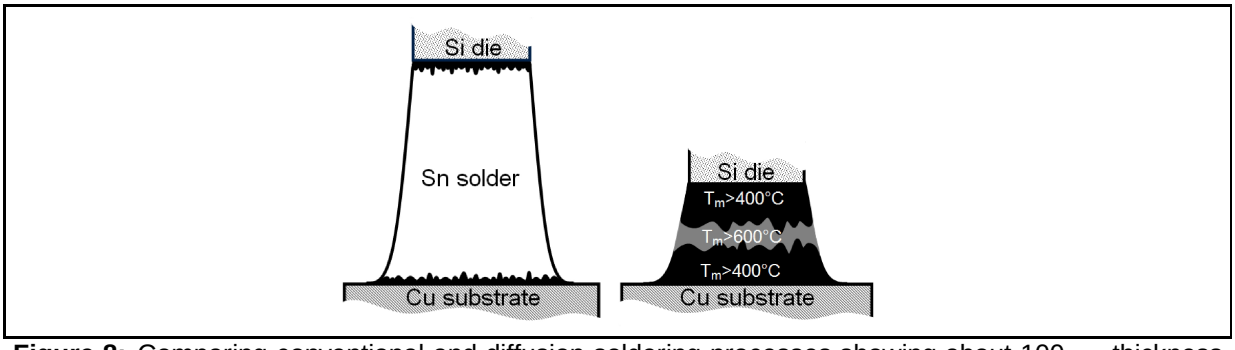

Figure 8: Comparing conventional and diffusion soldering processes showing about 100μm thickness for conventional but far less for diffusion soldering

Further improvement for the thermal behavior results from the use of substrates with thinner ceramics and thicker copper layers. As described in [6] a reduction of 10% is possible for the resistivity from case to heat sink R_{thch} due to this. Along with the diffusion soldering and a proper thermal interface, this will further support designers in managing the challenging task of thermal management.

Conclusion

The marked driven demand for higher power density designs inevitably leads to higher temperatures at the silicon level. The interconnection of power semiconductor and heat sink often forms a limiting bottleneck. Carefully choosing and properly applying thermal interface material is mandatory to get the best possible performance. Thus, power module, TIM and the process of application have to be considered as an associated, closely defined subsystem in power electronic designs. New interconnection technologies in future designs will be mandatory improvements in thermal aspects.

References

[1] Infineon Technologies,
Datasheet of the FZ1800R12KF4, 1998
see also www.infineon.com

[2] Infineon Technologies,
Datasheet of the FZ900R12KE4
see also www.infineon.com

[3] Ijeoma M. Nebe and Claudius Feger
Drainage-Induced Dry-Out of Thermal Grease IEEE Transaction on advanced packaging, VOL31, No3, August 2008

[4] Carsten Guth et al.
New assembly and interconnects beyond sintering methods PCIM 2010 Nuremberg, Germany in May 2010

[5] Infineon Technologies,
Datasheet of the FF450R12ME4
see also www.infineon.com

[6] R. Ott et al.
New superior assembly technologies for modules with highest power densities,
PCIM, Nuremberg, May 2010

[7] Infineon Technologies
Power Cycling Capabilities of IGBT4,
see also www.infineon.com