Managing Temperature Differences Between IGBT Modules

By John Parry, Industry Manager, Mentor Graphics

Guy Diemunsch's interest in thermal design started while he was preparing his PhD in Physics. He taught how to build comfortable, low energy homes and buildings in the University of Franche-Comté located in Besancon in the 1980's. In 1994 he joined Hewlett Packard to manage the thermal design of HP's range of professional PCs and workstations, where the challenge was to make their operation silent for the European market. In 2002 Guy joined Schneider Electric to optimize the thermal design of high-end UPS (MGE-UPS). Cost reduction was a major objective and this got Guy involved in power electronics. Five years later Guy was invited to join Aavid Thermalloy, a supplier to Schneider Electric, to extend Aavid Thermalloy's business in Power Electronics. Guy joined Electronic Cooling Solutions Inc. at the beginning of 2013.

I met Guy Diemunsch at THERMINIC in Berlin September 2014 we talked about some work he was doing on cooling IGBTs for high power inverters & converters used in renewable energy applications (wind turbines & photovoltaic power plants), drives and electric networks.

Back in 1994, Guy first came across the challenge of minimizing the temperature difference between different components when he was working on a computer cooling problem for Hewlett Packard. Twenty years on Guy was now faced with the same challenge, this time for a power electronics application cooling IGBTs as a project for a customer of Electronic Cooling Solutions.

Their initial design gave an unacceptable temperature variation between three IGBT modules. To ensure the efficiency of the system was not impacted it was necessary to hold the temperature of the modules to within 2°C of each other, otherwise the operational performance of the modules would be too different.

Faced with this challenge, a choice would need to be made between air, water, or phase change cooling. The best option depends on how well the chosen solution allows the temperature level and uniformity to be controlled while managing the mass flow and preheating of the cooling fluid.

Phase change cooling is a great solution to reduce the temperature difference between components; however, this solution is often the most expensive when compared with air cooling with a heatsink or water cooling with a Liquid Cold Plate (LCP). Therefore Guy chose to focus his attention on using air or water cooling to control temperature.

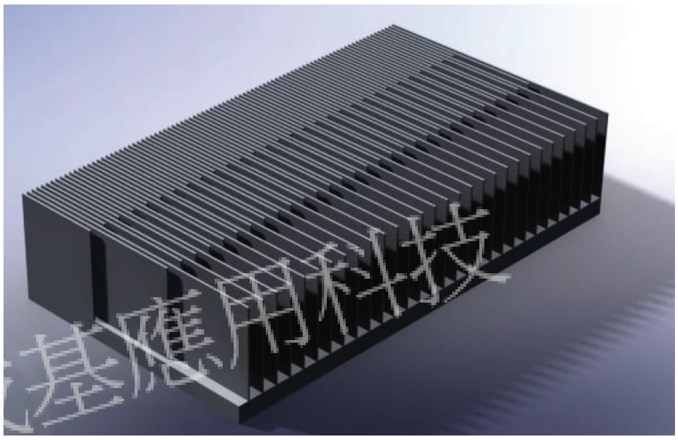

Figure 1. Preliminary Air Cooling (Heatsink) Design

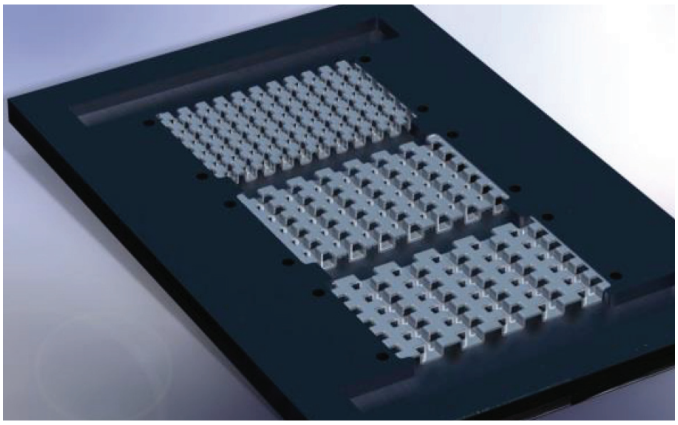

Figure 2. Preliminary Water Cooling (LCP) Design

The most obvious solution is to arrange to bring the same amount of cooling fluid at the same temperature to each identical component. There are at least two such examples in liquid cooling applied to Power Electronics: the ShowerPower® solution from Danfoss Silicon Power and the parallel cooling patent from Schneider Electric.

With the continuous increase of the system density, thermal and system engineers are very keen to minimize total flow rate needed for cooling. Therefore in most cases the focus is on optimizing a serial cooling solution to keep this flow rate as low as possible, with the same air being passed over several components. If however, the aim is to keep components at the same temperature, a very high volume of air flow would be required, so that the temperature rise in the air as it flows through the system is very low, such that all components are cooled by air at the same temperature.
As an example, for temperature difference

below 2°C, for three components generating 500 W per IGBT, serial cooling will then require more cooling fluid than parallel cooling. Is it therefore necessary to use parallel cooling? Guy realized the answer is no, because he can increase the cooling performance of the heatsink for each successive IGBT to account for the increase in air temperature passing over it.
To accomplish this from a purely thermal standpoint, each IGBT could have its own heatsink design, and this would be the easiest arrangement to optimize. This solution would increase the Bill of Materials (BoM) and increase assembly costs. However, the most compelling reason for not choosing this approach is the risk of wrong assembly. As the risk of wrong heat sink placement is both high and critical for the system, this is forbidden in all risk analysis. The avoidance of this issue through design would increase the manufacturing and assembly costs substantially.

To solve the problem, Guy created a single heatsink designed to allow the thermal resistance to vary in the flow direction. The constraints imposed by such a design are:

- Cooling fluid must be ducted to keep a constant mass flow between chips,
- Components must be grouped together.

The two preliminary designs are shown, without the ducts visible, in Figures 1 and 2.

The initial reaction of the customer was that the design of these solutions is complex, so they were concerned that it would be necessary to do many design iterations that may not arrive at an acceptable final solution.

Guy was able to assure them that it would be possible to converge on an acceptable design in two or three iterations by applying the following simple process:
Step 1: Define the most critical (i.e. lowest) thermal resistance needed and search for an existing heatsink design that will meet this duty and note the mass flow rate associated with it (A),
Step 2: Define the flow rate (B) needed considering the max temperature of the cooling fluid at the intake.
Step 3: Iterate 1 & 2 until the two mass flow rates (i.e. A & B) are very close.
Step 4: Define each local thermal resistance which is always possible because we have already defined the lowest thermal resistance in Step 1.

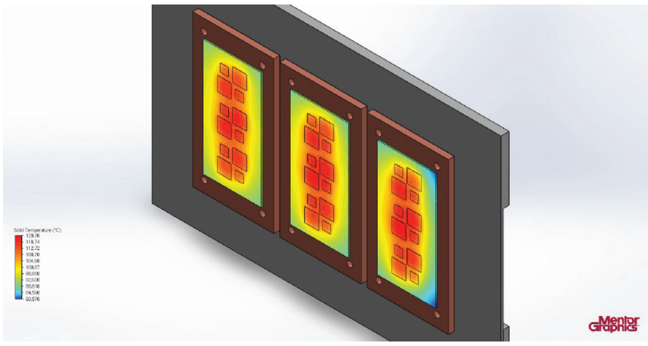

Figure 3. Chips Temperatures (Liquid Flow from Right to Left)

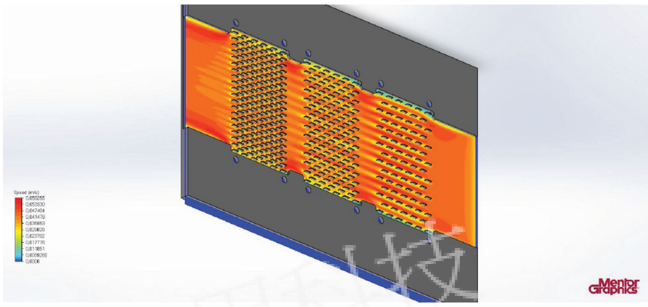

Figure 4. Fluid Speeds

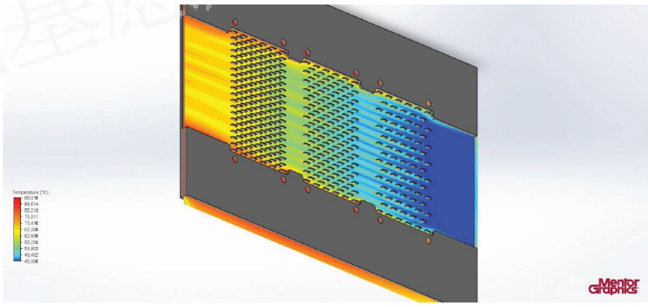

Figure 5. Fluid Temperatures

Step 5: Validate the solution using simulation.

The solution that Guy arrived at is shown in Figures 3, 4, and 5 where he used FloTHERM XT to simulate the current heat transfer in the cooling channels. The total power is 1.5 kW (500 W per IGBT module). The heatsink is cooled by water with 30% of Glycol. The fluid intake temperature is set to the maximum of 45°C.
With 15 g/s of fluid the pressure drop of 22 Pa (0.09 inch H2O) and an average fluid speed of 0.03 m/s, low enough to ensure there is no erosion even with Aluminum. The total temperature elevation of the fluid is 30°C. The three thermal resistances (from the right to the left) are 0.12, 0.10 & 0.08 °C/W. This detailed analysis showed that within each module the average difference between chips in the direction of the flow is 1.6 °C. This arises from the fins being uniform below each IGBT, while as fluid passes through each finned region the boundary layer thickens. The worst case is for the chips in the top and bottom corners on the right in the above plots where the temperature difference is up to 2.9°C. This boundary layer effect was corrected in a subsequent design refinement.

We also see fluid bypasses between fins for the first IGBT. During the design process fin spacing was optimized in laminar & turbulent flow regimes. In fact if the gap is too large the fluid in the middle of the fins is not heated. If the gap is too small, the pressure drop is increased without a corresponding improvement in heat transfer. Another observation Guy made was that within the IGBT modules the temperature difference between the central ICs and those at the edge were up to 7.5°C. This is related to the layout and the module and the packaging design, neither of which were under Guy's control. Again, this can be corrected through refinements to the LCP fin design. However, this improvement will require nonstandard folded fins manufacturing, which might not be cost effective.

In conclusion, with care, serial cooling can be used even when it is necessary to respect tight design criteria for temperature differences between components and chips. Using a simple design process it is possible to meet the design goal for temperature control while using a low flow rate with a correspondingly low pressure drop. This maximizes the energy efficiency of the cooling solution.

This article was originally published in Engineering Edge Vol. 3 Iss. 2

Mentor Graphics
Mechanical Analysis

This article originally appeared in Engineering Edge Vol. 3 Iss. 2

Download the latest issue:
www.mentor.com/mechanical/products/engineering-edge

ENGINEERING EDGE

Accelerate Innovation with CFD & Thermal Characterization

You might also be interested in...

Yamaha Motor Company Ltd: An Evaluation of Delamination of Power Modules using the MicReD T3Ster®

Saipem S.p.A: Moves FloEFD Offshore

NAMICS Corporation: Use the MicReD T3Ster® to Develop High Thermal Conductivity Sintered Silver Paste

Visteon Electronics Corporation: Unravel the Complexities of Automotive Instrument Cluster Designs

Download the latest issue:
mentor.com/mechanical/products/engineering-edge