ENGINEERING EDGE

Accelerate Innovation with CFD & Thermal Characterization

FloTHERM® & FloTHERM® XT CAD-centric Thermal Simulation Software

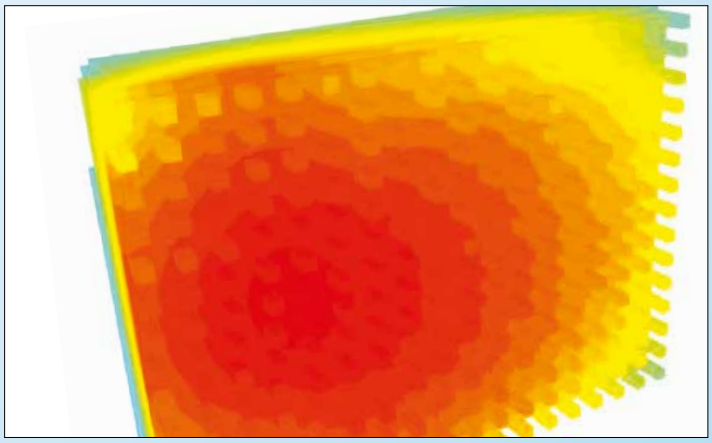

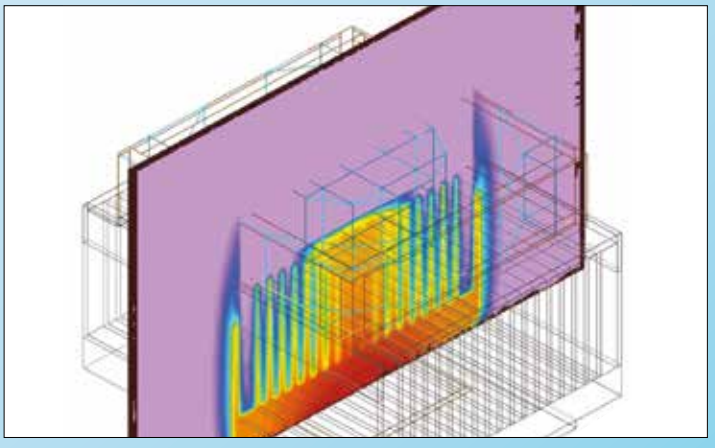

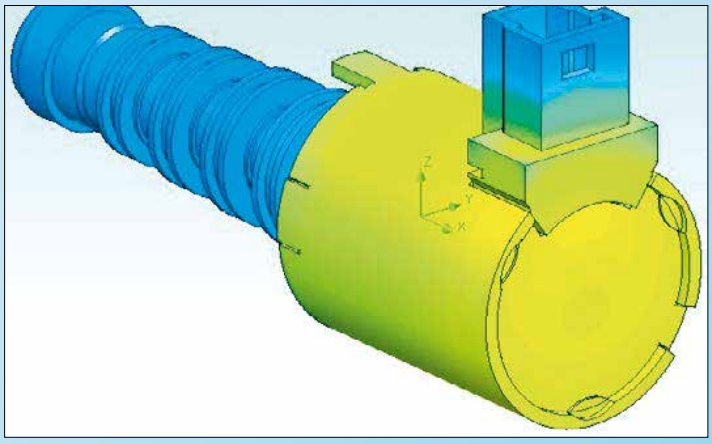

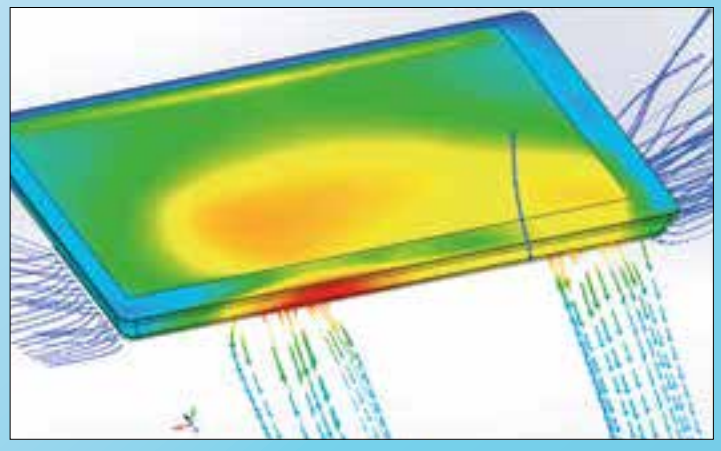

Mentor Graphics

Mechanical Analysis

mentor.com/mechanical

This collection of **Engineering Edge** articles showcases our customers' success with FloTHERM® & FloTHERM® XT CAD-centric Thermal Simulation Software across a number of industries and disciplines including Automotive, Process, Power Generation, Electronic Cooling, Power Electronics, Aerospace, and Military. Test-drive the fully featured **FloEFD, FloTHERM XT,** or **Flowmaster** Virtual Lab.

- Fully operational software, hosted in Mentor Graphics virtual lab environment
- Immediate access from any current PC web browser
- Test at work, home or while traveling
- No IT department approval required
- No special equipment
- Evaluate **FloEFD, Flowmaster,** or **FloTHERM XT** at your own pace

Index

Turn it up!

Thermal Simulation for the Design of Automotive Multimedia Systems

by Dr. Uwe Lautenschlager
Continental Automotive GmbH

Continental's Automotive Group, a system supplier for international automotive OEMs, has three divisions: Chassis & Safety, Powertrain and Interior. The division 'Interior' is itself split into four business units: Instrumentation & Driver HMI; Infotainment & Connectivity; Body & Security and Commercial Vehicles & Aftermarket.

The business unit Infotainment & Connectivity covers Radios, Connected Radio & Entry Navigation, Multimedia Systems, Telematics, Device Connectivity, and Software & Special Solutions, with a particular emphasis on Entertainment, Information, Navigation and Communication.

One of the greatest challenges for automotive multimedia systems is that they have to operate under a wide range of environmental conditions:

- Temperatures ranging from -40°C to +100°C
- Wide range of power supply voltages (-4.5 to 16.0 V)
- Superimposed alternating voltages
- ESD polluted environment (ESD = Electrostatic Discharge up to +/-15 kV)
- Full range of humidity (0% - 100%)
- Chemical influences

And with stringent design criteria:

- Mechanical strength and stability (static / dynamic): Vibration & mechanical shock
- Strong limitations on size and weight
- Strong requirements regarding risk of injury (Head Impact) and product liability

A World Away from Physical Prototyping

Designing systems to meet such requirements requires much more than physical prototyping after the design is complete. Indeed, Continental's simulation vision is to 'Get the Product Right the First Time', and their strategy to achieve this has been through Simulation-Based Design Decisions. As a result, simulation is now very highly integrated into the design process.

Continental undertakes full 3D system-level modeling, with all thermally-relevant and air flow parts included in the simulation. The geometry representation is based on mechanical CAD data, with export/import of all relevant parts and modules from the CAD-System. However, the geometry representation can involve both simplification and idealization. The complexity used depends on the simulation objective: less detail would be used for the analysis of various initial concepts, whereas more details would be included in the model for a mature design.

"Know-How" – Continental's Competitive Edge

Continental has developed considerable in-house 'know-how' associated with building and validating thermal simulation models (FloTHERM™ models) of their complex electronic systems and sub-systems, covering both the physical hardware, and potential cooling solutions.

Hardware-related

- Housing: 1-DIN, 2-DIN, customer-specific
- CD-, HDD-, DVD-Drives (ROM, Video), Changer
- Shielding frames
- Front displays (CCFL-, LED-Backlight)
- Connectors and cables
- Components: Main-CPU, Processors, Memory
- PCBs
- Modules: Power Supply, Drives, GPS, Telephone, Amplifier, Tuner, Display, etc.

Cooling solution-related

- Fan (type, characteristic curve, direction)
- Contact resistances (conductive paste, etc.)
- Vents, filters (free area ratio, pressure loss)
- Heatsinks, heat pipes, heat spreaders
- IC Packages (type, size, thermal model type)
- Thermally conductive gap pads, etc.
- Environmental conditions
- Power dissipations
- Location of sensors for measurement

Uncertainties arising from the various unknowns at each stage in the design process, modeling assumptions and simplifications are mitigated through validation, as the following two examples illustrate.

Continental

Examples:
1. Tuner-Module Modeling

This can be modeled as a single block with homogeneous power dissipation distribution, as a detailed model with all of the internal electrical components and shielding represented, or anywhere in between. In Continental's experience, a simplified model is often sufficient within a complete system analysis. However, conclusions on the internal component temperatures are drawn from a reference simulation as depicted in Figure 1.

Example:
2. Chassis Model Verification

A part of the chassis is used as a heatsink for amplifier cooling. An important aspect of the chassis is the extent to which the brackets on the top and bottom covers are in thermal contact (related to manufacturing tolerances). Verification of modeling assumptions is the key to improving model fidelity. Therefore, a prototype of the product was measured and the results used to compare against the simulation model. Scenario 1 (g1) considers the gap to be fully closed (perfect thermal contact) whereas Scenario 2 (g2) has a gap of 0.3mm, so there is effectively no thermal contact between the covers. By comparison with the measured case temperature the results for Scenario 2 (with gap) seems to be the best representation of the bracket's contact.

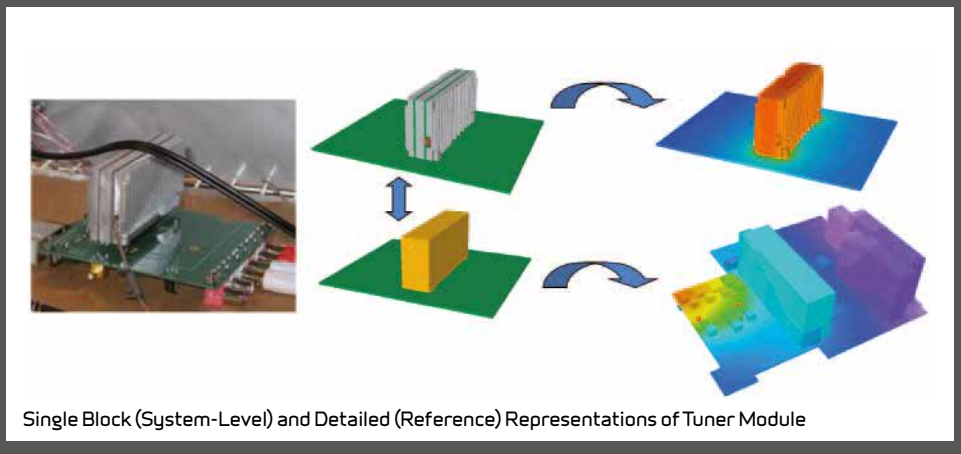

Single Block (System-Level) and Detailed (Reference) Representations of Tuner Module

Multimedia Systems with Audio, Navigation, Internet, Telephone, TV & Rear-Seat Entertainment Functions

"FloTHERM is a key component of our Simulation-Based Design Decisions strategy, ensuring that our thermal design goals are met and we can deliver on Continental's simulation vision of Getting the Product Right the First Time." Dr Uwe Lautenschlager

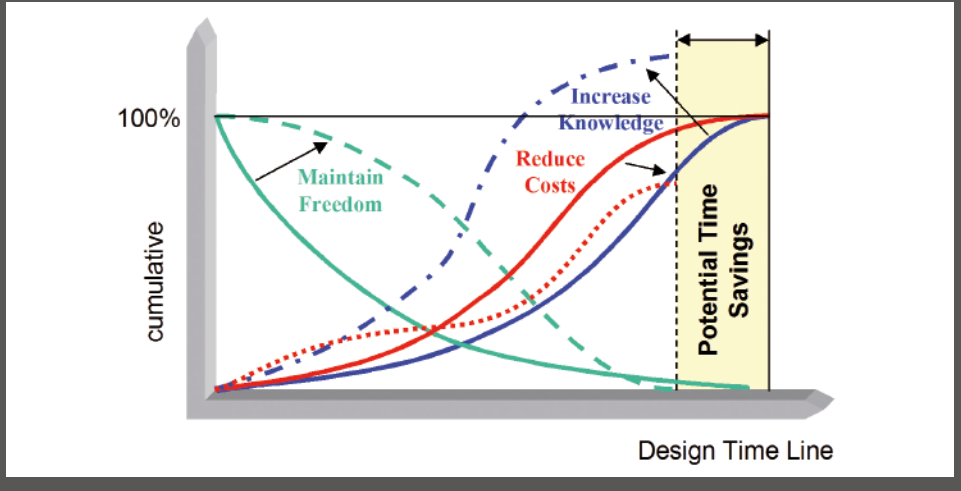

"We selected FloTHERM™ for several reasons but primarily for its robust solution capabilities. FloTHERM's object-associated Cartesian meshing is instantaneous, fully automatic, and most importantly guarantees a mesh that produces accurate simulation results even when geometry changes are made to the base model. This is something we absolutely need for our Multidisciplinary Design Optimization (MDO) activities. Not to mention its modeling and result evaluation capabilities that in total simply made it the best solution for us." Dr Uwe Lautenschlager

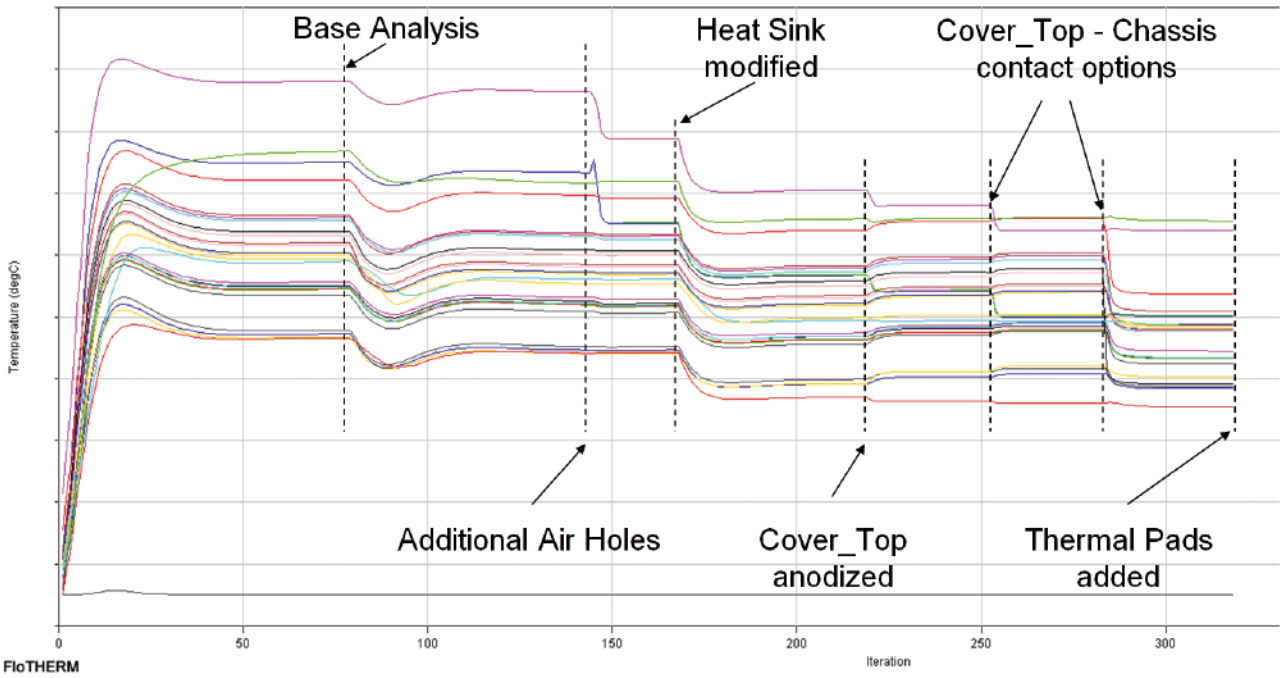

Sequential Concept/Design Improvements from MDO for Multimedia System

Beyond High Fidelity Simulation

Continental's focus on simulation does not stop at building and validating thermal models. Design decisions made to improve the thermal design can impact the mechanical, electrical, and EMC performance of the product. Faced with this problem, a major question facing Continental's designers is: "How can we achieve design flexibility and enable better design decisions before the freedom for such decisions is eliminated?"

The answer to the challenge requires analysis of the product's behavior for all disciplines as well as the identification of independent and coupled system variables. Multidisciplinary Design Optimization (MDO) techniques with simulation, optimization (with discipline-dependent objectives and constraints) and Design-of-Experiments (DOE) & Response Surface Methods are suitable means to solve this design problem. Indeed, DOE tools are the foundation for concept exploration and robust design. Two major aspects of DOE are the planning and statistical analysis of the numerical simulations.

The possible number of discrete and continuous design variables is extremely high. Screening simulations support the selection of the major factors (design drivers) to be included as design variables in the MDO. Even so, the computational effort required can be immense.

Sequential Design Improvements

Changes that improve the overall design, i.e. the compromise between the product's mechanical, electrical, and EMC performance, and additional constraints such as temperature limits, manufacturability or cost allow sequential improvements to be made. For each discipline these improvements can be visualized as improvements against the base design. In the case of the thermal design, these can be visualized as changes in monitor point temperatures that match the locations chosen for sensors that will be used to instrument the physical prototype.

Business Benefits

Modern radio and navigation systems, i.e. multimedia systems, for automotive

applications are highly complex systems with a large variety of mechanical and electrical components and assemblies. The design of these automotive multimedia systems has to fulfil requirements from mechanical stability and thermal management to electromagnetic compliance and optical homogeneity and therefore requires the interaction of mechanical, electrical and software engineers. Interdisciplinary knowledge is essential.

Continental has to react quickly to changing customer requirements. Products have to be developed with regards to customer confidence and quality, costs and design time. MDO has been found to be a suitable means for finding better design solutions in a multidisciplinary environment. Simulation supports knowledge generation and a deeper understanding of product behavior at lower cost, in shorter time and with increased product flexibility, leading to increased customer confidence in our products.

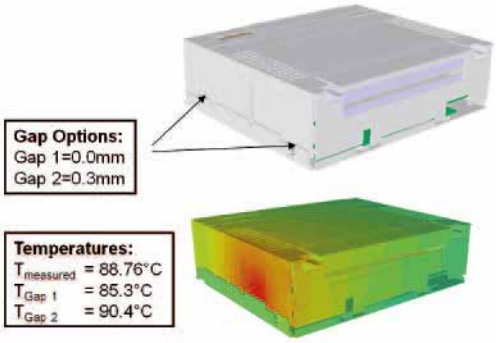

Verification of Case Thermal Contact Modeling Assumption

Structure Functions – The Bridge Between Thermal Measurement and Thermal Simulation

TRADITIONALLY thermal measurements of electronic components have been done using thermocouples. However, thermocouples come with critical disadvantages, (such as contact thermal resistance between target and probe which makes the measured result very unstable, and the tendency of the thermocouple to conduct heat away from the surface being measured) making them near useless for measuring the surface temperature of plastic packages.

Increasingly there is a need to validate thermal models during product design to ensure that the product will perform as designed, by confirming material properties and the thickness of thermal interface material (TIM) layers. A major issue with thermocouples is that they simply cannot measure the temperature of the internal structure. Yet by design, the dominant heat flow path is from the junction, through many materials and material interfaces before passing into a PCB or heatsink, whose temperature can be conveniently measured. Even then, a thermocouple only provides a single temperature value. Thermocouples are therefore a 'blunt instrument' when it comes to thermal design verification.

In this article, we are going to introduce the thermal structural analysis method, which is based on transient thermal measurements, that allow the thermal behavior of the system, including heat spreading, to be characterized as a distributed Resistance-Capacitance (RC) network. By measuring the transient response of junction, we can easily observe the heat spreading path inside package, board, TIM and heat-sink, etc.

The Challenge of Thermal Analysis in the Real World

The best way to study thermal structure is to take a look at the isothermal distribution or heat flux distribution along the heat-spread path. However, in the real world it is impossible to take a picture of heat distribution inside any solid object. The only way to inspect heat flux distribution vitually is by using a software simulator such as FloTHERM™ used in this article.

According to the theory [1], thermal systems are distributed RC systems, which can be modeled by thermal resistance Rth and thermal capacitance Cth. To evaluate a RC system, the most common way is to measure transient response under a step power excitation.

Consider the experiment setup in Fig 1. Ideal heat insulation material prevents heat from escaping to Y and Z direction, the cold plate at the right side of X axis provides an ideal thermal boundary condition. In this setup heat flux will be constrained to X axis which can be considered as one-dimentional heat spreading path starting from the heat source on the left side to the cold plate on the right side along X axis.

Thermal property Rth and Cth on the heat spreading path determines step power response of the system, theoretically we can evaluate the thermal structure by measuring the thermal transient response in an electrical test method as standardized in JEDEC JESD 51-1 in 1995.
In the experiment, we place three kinds of flag material in the middle of heat path.
1. Same as pure copper. (Cu50W)
2. Doubled specific heat against pure copper (Cu50W_2xCth)
3. Halved thermal conductivity against copper. (Cu50W_2xRth)

Figure 2 plots step power responses and Structure Function. In temperature response view, variation caused by different flag material can be seen, however it is not clear enough while in the Structure Function view the structural information can be clearly identified as shown in Figure 3.

Case Study Closer to a Real PCB Board Application

In a real-world application such as a package mounted on PCB board, heat spreads not only vertically but also horizontally as shown in Figure 4.

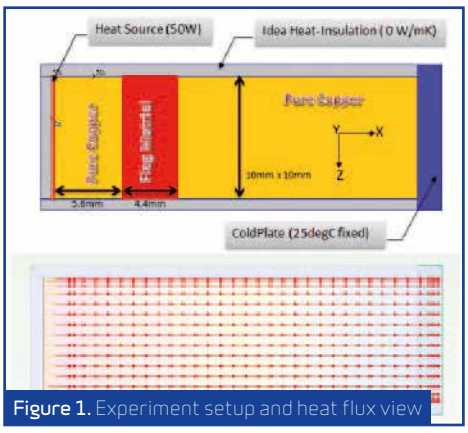

Figure 1. Experiment setup and heat flux view

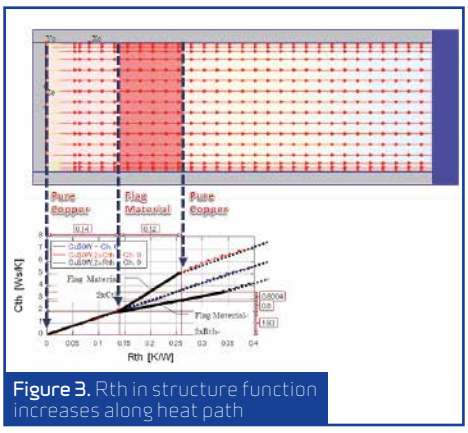

Figure 3. Rth in structure function increases along heat path

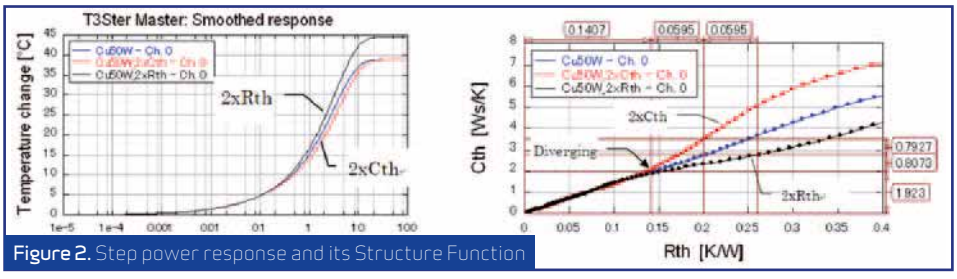

Figure 2. Step power response and its Structure Function

Mentor Graphics
Mechanical Analysis

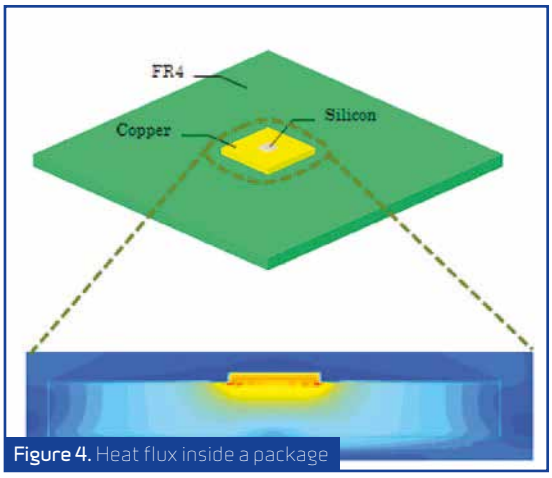

Figure 4. Heat flux inside a package

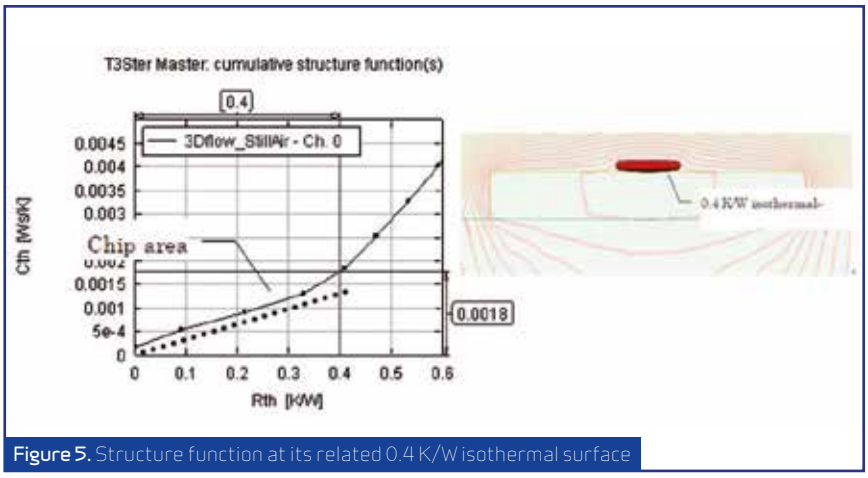

Figure 5. Structure function at its related 0.4 K/W isothermal surface

The heat source (silicon chip) is attached to a metal (copper) substrate and then attached to the FR4 board. Every material is built as a cuboid block and contact thermal resistances are not included for simplicity.

In the Structure Function, a straight line section from 0 - 0.4 K/W can be seen at the beginning. This straight line comes from the nearly 1D heat flow inside the chip as shown in Figure 5. This is because air outside the silicon chip has relatively huge thermal resistance compared to silicon, so that heat is forced to go through the thickness of the chip as discussed in the previous section. After 0.4K/W, structure function curve goes up exponentially which is caused by the 3D heat spreading in the metal substrate as shown in Figure 6.

For the same reason, Structure Function from 0.8K/W - 1.2 K/W also indicates heat spreading in the copper block and the Structure Function curve shows an increasing slope. After 1.2K/W we observe a decrease in the slope of the curve. This is caused by the physical boundary of metal substrate.

Conclusions

Traditionally when doing thermal analysis, thermal models built in CFD simulation software contain many thousands of pieces of data. The challenge for the user is how to verify the correctness of the model. As Structure Functions can be obtained from both experiment and simulation, we are now able to verify package thermal models against the data for real packages by comparing their structure functions. If there is any mismatch we can easily identify and resolve the problem, thereby increasing the fidelity of any board or system-level models in which the package model is used [2]. As Structure Functions track the heat flow path from the die junction to the ultimate ambient, the technique can also be applied to board and system-level models in late design to qualify electronics products before they go into full production.

Note

This article is edited from the paper published in the Transactions on The Japan Institute of Electronics Packaging, Vol. 5, No. 1 (Dec. 2012).

References

[1] V. Székely and T. Van Bien, "Fine structure of heat flow path in semiconductor devices: a measurement and identification method," Solid- State Electron., vol. 31, pp. 1363-1368, 1988.

[2] Andras Vass-Varnai, Robin Bornoff, etc., "Thermal Simulations and Measurements – a Combined Approach for Package Characterization". ICEP 2011 Japan.

[3] Yafei Iuo, "Structure function based thermal resistance & thermal capacitance measurement for electronic system". CPMT Symposium Japan, 2010 IEEE.

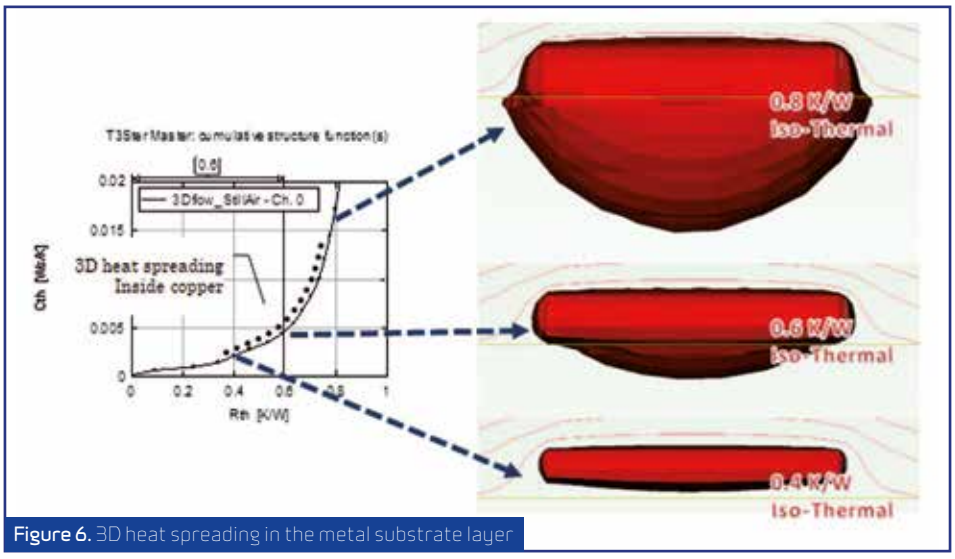

Figure 6. 3D heat spreading in the metal substrate layer

A Tablet for Everything

FloTHERM XT™, FloTHERM®, & T3Ster® Cure Thermal Headaches in Tablet Computers

By Guy Wagner and William Maltz,
Electronic Cooling Solutions Inc.

Mentor Graphics
Mechanical Analysis

It is human nature to follow the path of least resistance. To that end, it has never been easier or quicker to get what you want, when you want it. The demand for this instant gratification has not only resulted in an explosion of online shopping and drive-thru conveniences such as coffee shops and fast food chains, but has also dramatically changed how we work, play, and learn. No longer do we visit the library to get a book, buy newspaper or wait for the evening news to find out what is happening in the world. We read the news, finish an urgent work report, complete the next level of Angry Birds and find out when the next train is due and all from the same little device: The Tablet

The growth of the tablet market is reflected in the continued decline of the PC market, with consumers generally choosing to replace their laptops with tablets rather than purchasing both.[1] Cheaper prices and a wide range of form factors and sizes has resulted in manufacturers designing devices that are slimmer, sleeker and more compact than ever before.

The thermal design of next generation handheld devices must address both comfortable surface touch temperatures and maximum temperature limitations of internal critical components while also meeting aggressive industrial design requirements. This article discusses the challenges in meeting these requirements in tablet designs.

Thermal models of tablets were created using FloTHERM XT to help understand the maximum allowable power dissipation under various operating conditions. The models were also used to conduct parametric studies to determine the best way to move heat from the internal components out to the case of the tablet where it can be dissipated.

Maximum Power Dissipation

Handheld devices are increasingly capable of running applications that used to require laptop and desktop computers. The requirement that these devices provide better performance with a smaller form factor presents significant challenges, especially when one considers that passive cooling is also a requirement. Several studies have focused on the cooling challenges of hand-held devices; Brown et al, Lee et al, Mongia et al, Huh et al, and Gurrum et al. [2-6]

The maximum possible power dissipation by natural convection and radiation has been calculated for this study and is shown in Figure 1 (overleaf). With a 25°C ambient condition at sea level, the maximum total power dissipation was calculated with a requirement that the surface temperature not exceed a touch temperature of 41°C. This is the maximum aluminum enclosure comfort touch temperature as presented by Berhe.[7]

It can be seen that the theoretical maximum total power dissipation is 13.9 watts when the device is suspended vertically in midair with conduction and radiation occurring from all surfaces as shown in Figure 2. When the device is horizontal, the maximum dissipation falls to 13.1 watts. When the device is placed on a horizontal adiabatic surface, heat transfer occurs from the sides and front surface only and the maximum power dissipation is reduced to 7.9 watts as shown in Figure 3. This condition occurs when the user places the tablet on a blanket or pillow effectively blocking heat transfer from the back surface. These values establish bounds for the maximum amount of heat that can be dissipated by the tablet for different orientations in still air.

In order to calculate the total power dissipation, the following assumptions were made: a typical tablet size of 180 mm by

240 mm and an ideal condition of uniform surface temperature. With a conservative assumption of a surface emissivity of 0.8, the radiant heat transfer accounts for more than half of the total power dissipation at an ambient temperature of 25°C.

In order to achieve a 41°C touch temperature, design parameters need to be considered carefully. It is important to design the tablet to be as isothermal as possible to maximize the amount of heat transfer to the surroundings. The reason for this is so that all surface areas of the tablet are as far above ambient temperature as possible to maximize heat transfer without exceeding the maximum allowable touch temperature. When a surface is no longer available for heat transfer, such as when the tablet is placed on a blanket, the amount of power that can be dissipated while staying under the maximum touch temperature drops significantly since heat transfer is effectively blocked from the back surface.

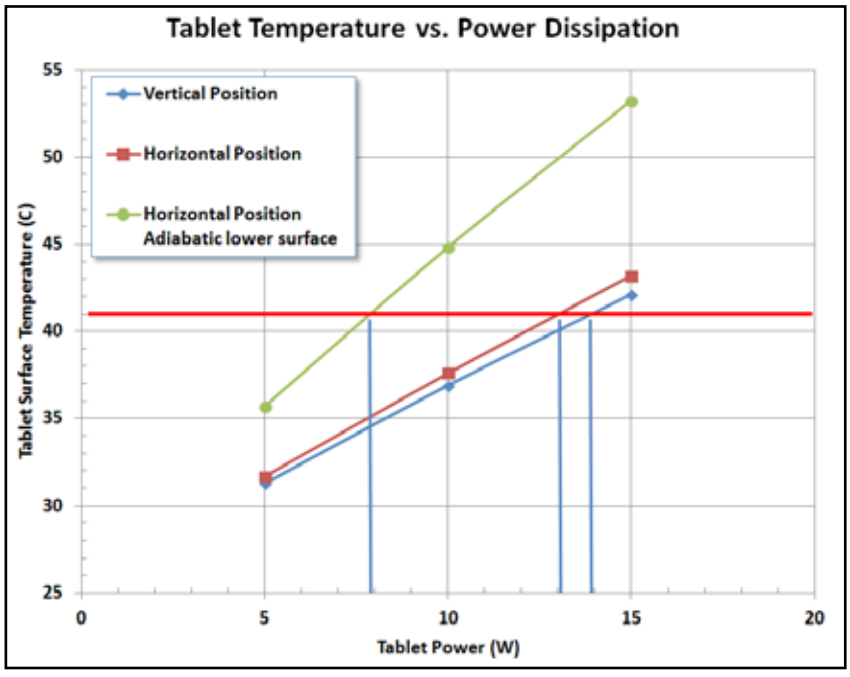

Figure 1. Total power dissipation removed by passive cooling of a tablet

The Value of Numerical Models

In order to analyze the impact of different thermal management techniques, a detailed computational fluid dynamics (CFD) thermal model is constructed using FloTHERM. Since the actual thermal characterization data of the main processor in the tablet may not be known, the thermal characteristics of the processor can be measured with a high degree of accuracy using Mentor Graphics' T3Ster® to determine the thermal resistance from the processor IC to the case and the PCB. This allows accurate capture of heat flow from the top and bottom of the processor. The thermal model of the processor that was generated using T3Ster can be directly dropped into the tablet thermal model in FloTHERM. The results of a detailed thermal model of the interior of a tablet are shown in Figure 4.

Notice how the automatic adaptive meshing in FloTHERM XT follows the small surface details of the components to accurately capture convective and radiant heat transfer from these surfaces.

Hot-Spot Temperature Reduction

Since the goal is to keep the touch temperature at or below 41°C, determining the effect of high conductivity heat spreaders will have a major impact on the design. Some parametric studies were run to determine the effect of making the back side of the case with materials of varying thermal conductivity. The results are shown in Figure 5 and summarized graphically in Figure 6. This study assumes the same power dissipation for each simulation. The only parameter that is being changed is the thermal conductivity of the case. As a reference point, typical thermal conductivity of most plastics is in the range

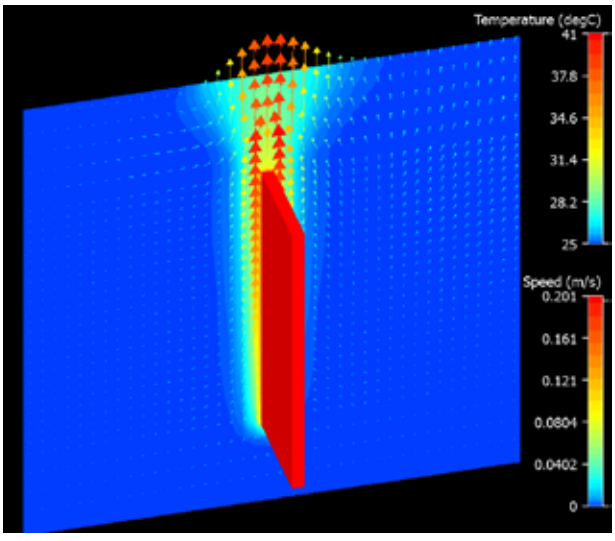

Figure 2. 41°C Isothermal tablet in vertical position

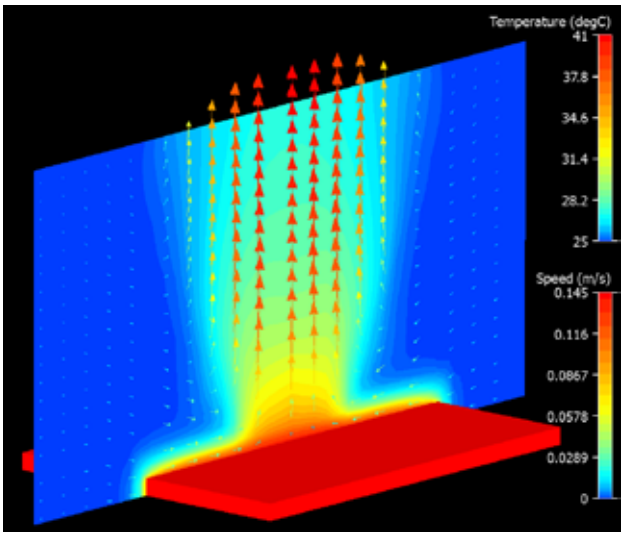

Figure 3. 41°C Isothermal tablet in horizontal position on an adiabatic surface

Mentor Graphics
Mechanical Analysis

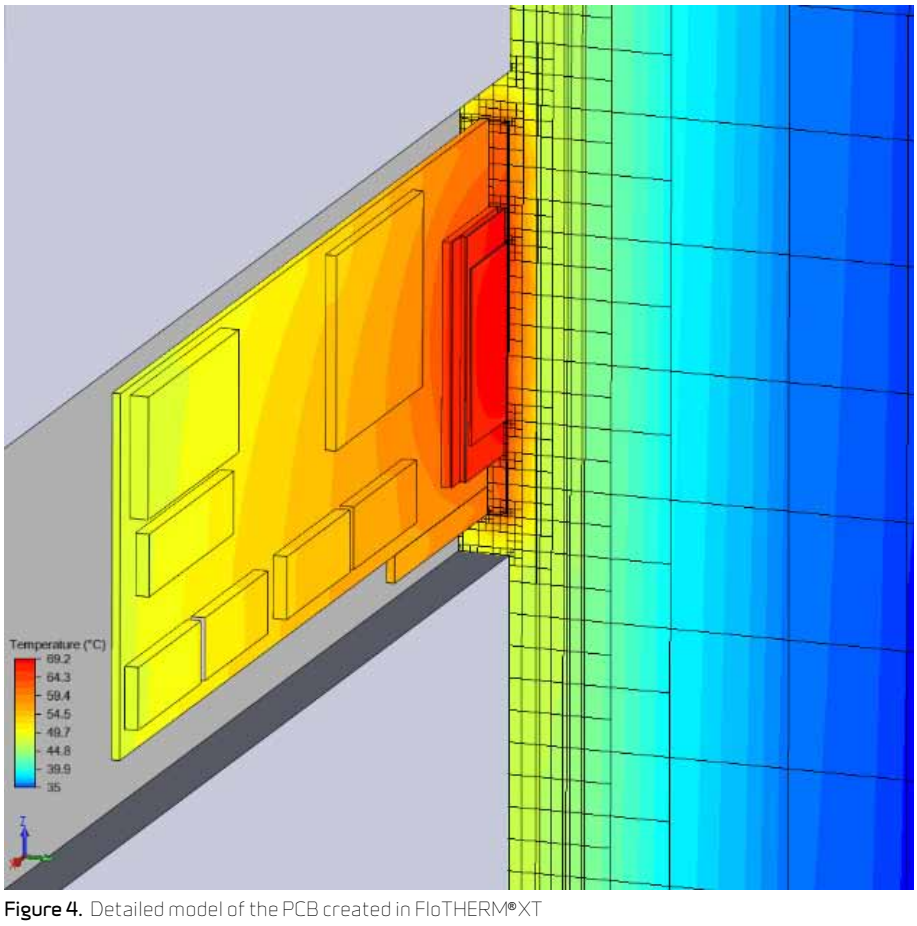

Figure 4. Detailed model of the PCB created in FloTHERM®XT

of 0.2 W/mK while aluminum approaches 200 W/mK. This makes aluminum about 1000 times as thermally conductive as plastic. Hot-spot temperature reduction can be achieved by either providing a high conductivity heat spreader inside the case of the tablet or by making the case itself out of high conductivity material. One must keep in mind that the maximum touch temperature is also a function of the conductivity of the case. As the conductivity goes down the maximum comfortable touch temperature goes up.

As an example, if the case is made of plastic with a thermal conductivity in the range of 0.2 W/mK, the case temperature that the user senses appears to be cooler than that of an alumunum case since the low thermal conductivity of the plastic results in less heat being conducted between the case and the skin. Since the surface area of the case is large in relation to the thickness of the plastic, heat transfer to the air is not reduced significantly over that of an aluminum case. This of course assumes that the heat is spread on the inside of the plastic case using a high-conductivity aluminum plate or a graphite sheet.

Deriving a Thermal Model of a Processor

When building a thermal model of a tablet, the thermal characteristics of the processor are not always known with a high degree of accuracy. It is also true that data sheets from the suppliers of thermal interface materials may not accurately reflect the thermal resistance of the interface material and the wetting properties of the material between the processor chip or lid and the heat spreader.

To overcome this limitation and get an accurate thermal model of the processor, T3Ster was used to determine the thermal resistance from the processor IC to the lid or heat spreader and the PCB. T3Ster is able to do a dynamic thermal characterization of the thermal resistance paths of a packaged semiconductor device.

The transient temperature response of the die is recorded as a function of a step input in power to the die and a structure function is derived from the transient temperature response that characterizes the thermal resistance of all the materials in the thermal path. Figure 7 shows the structure function that was derived for a processor using T3Ster. Note that the thermal resistance from junction to case is measured at 0.23 K/W using this technique.

This thermal model of the processor

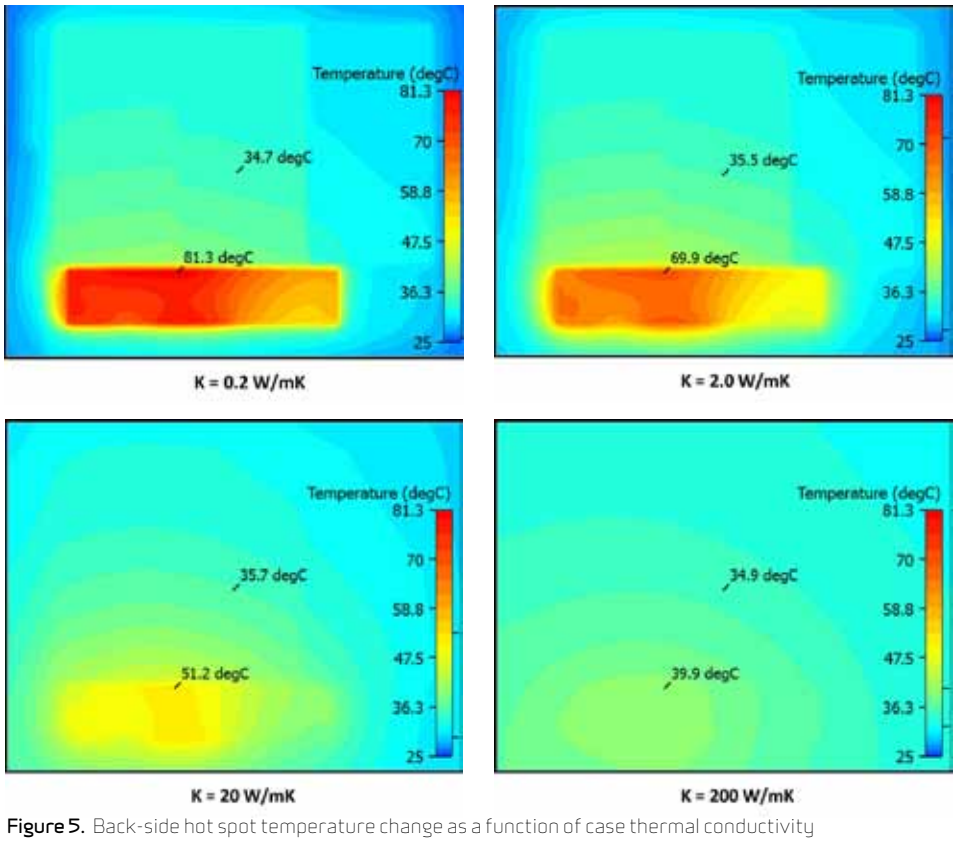

Figure 5. Back-side hot spot temperature change as a function of case thermal conductivity

package is then put back in to the CFD simulation and numerical experiments can be run to determine the change in processor junction temperature as other elements in the thermal path such as heat spreader materials and dimensions, air gap thickness between the heat spreader and back case and case materials are changed.

Summary

The maximum power dissipation of the internal components is not only governed by the size of the tablet but is a strong function of how well that heat is spread internally to reduce hot-spot temperatures. Few engineers realize the importance played by radiation in dissipating the heat from the exposed surfaces of a tablet. It is not until precise calculations are made that the importance of radiation is realized in the thermal design of the tablet. If the emissivities of the various surfaces are high, over half of the heat transfer to the surroundings is due to radiation. Overall heat transfer is maximized by reducing hot spot temperatures and spreading the heat so that all surfaces are effectively providing maximum heat transfer through convection and radiation.

Figure 8. Lab setup with IR camera, power measurement and temperature logging equipment

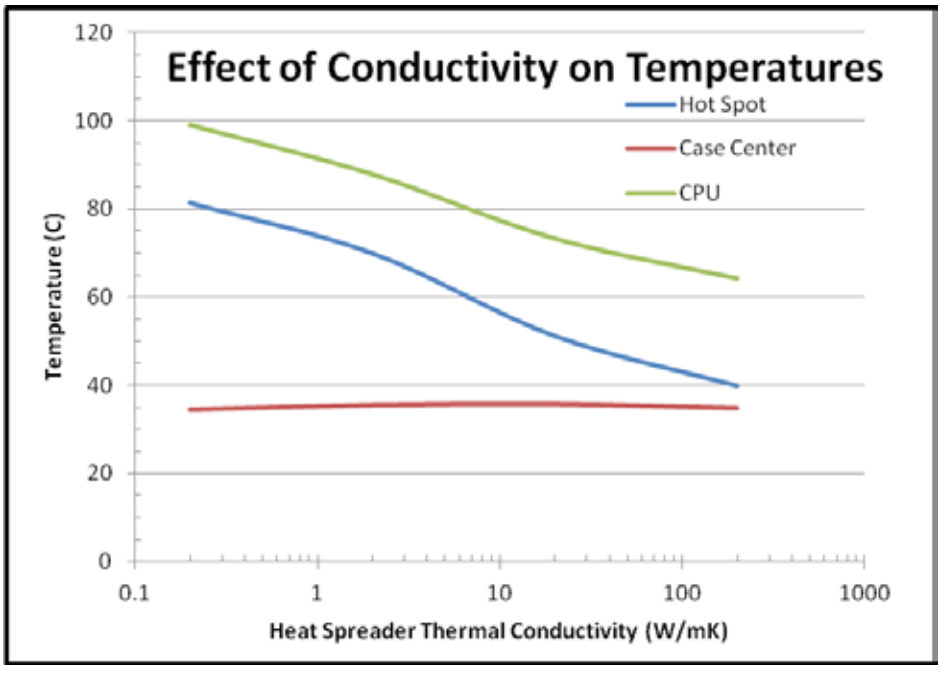

Figure 6. Effect of thermal conductivity of the case on temperatures

In summary, building an accurate thermal model of the tablet allows the designer to rapidly test the effect of design and material changes without incurring the high cost and schedule delays of testing prototypes. A thermal model allows the thermal design engineer to investigate far more alternatives than building prototypes. This results in a highly engineered tablet design that better meets the expectations of the user while providing an edge over the competition. High quality thermal models speed time to market and lower

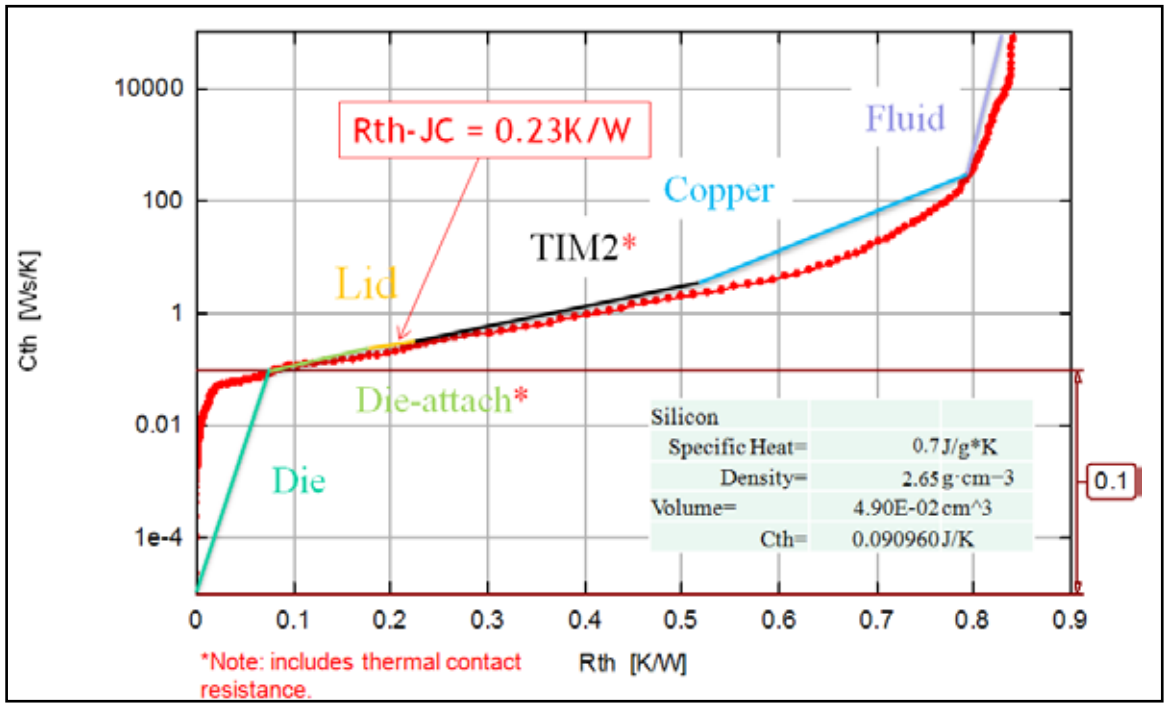

Figure 7. The cumulative structure function measured for the processor and lid using T3Ster®

development costs. With the accuracy of the latest simulation software, the intermediate step of building and testing thermal prototypes can be reduced or eliminated. The only need is final thermal verification of production prototype samples.

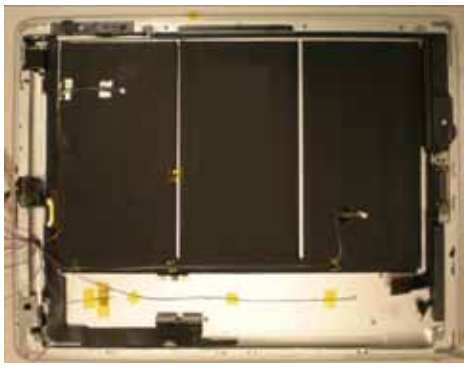

Figure 9. Tablet interior with thermocouples attached

References

[1] Preliminary data from the International Data Corporation (IDC) Worldwide Quarterly Tablet Tracker

[2] Brown, L., Seshadri, H., Cool Hand Linux® - Handheld Thermal Extensions, Proceedings of the Linux Symposium, Vol. 1, pp 75 – 80, 2007

[3] Gurrum, S.P., Edwards, D.R., Marchand-Golder, T., Akiyama, J., Yokoya, S., Drouard, J.F., Dahan, F., Generic Thermal Analysis for Phone and Tablet Systems, Proceedings of IEEE Electronic Components and Technology Conference, 2012

[4] Huh, Y., Future Direction of Power Management in Mobile Devices, IEEE Asian Solid-State Circuits Conference, 2011

[5] Lee, J., Gerlach, D.W., Joshi, Y.K., Parametric Thermal Modeling of Heat Transfer in Handheld Electronic Devices, Proceedings of the 11th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, I-THERM, pp 604-609, 2008

[6] Mongia, R., Bhattacharya, A., Pokharna, H., Skin Cooling and Other Challenges in Future Mobile Form Factor Computing Devices, Microelectronics Journal, Vol. 39, pp 992 – 1000, 2008

[7] Berhe, M.K., Ergonomic Temperature Limits for Handheld Electronic Devices, Proceedings of ASME InterPACK'07, Paper No. IPACK2007-33873

State of the Art Automotive Thermal Design by DENSO

Automotive

ECUs for vehicles and their components are becoming more complicated and intricate to produce energy-efficient and environmentally friendly cars. Successful thermal design is critical for manufacturers

By John Parry, Mentor Graphics

DENSO Corporation is a leading automotive supplier that designs and manufactures advanced vehicle control technology, systems, and components for major automotive manufacturers all over the world. Founded in 1949, DENSO is headquartered in Kariya, Japan, operates in 35 countries, and employs around 120,000 people worldwide. DENSO's Electronic Systems Business Group provides engine, transmission, and power management Electronic Control Units (ECUs) as well as semiconductor sensors, integrated circuits, and power modules.

I met up with Takuya Shinoda, Project Assistant Manager of the Technology Planning Department, Electronics Engineering Division 2, to discuss how DENSO are using thermal simulation to dramatically reduce their design time and cost. In this role, he is responsible for the thermal design of ECUs. Shinoda has the rare quality of understanding both mechanical and electrical disciplines. As he puts it, "Thermal design bridges both the mechanical and electrical disciplines. Thermal management is mainly a mechanical issue, but heat is generated in silicon, so it is necessary to also understand electronics to do thermal design correctly."

Design Challenges

ECUs for vehicles and their components are becoming more complicated and intricate to produce energy-efficient and environmentally friendly cars. Successful thermal design is critical for manufacturers. The junction temperature of the Integrated Circuits (ICs) or Field Effect Transistors (FETs) that drive such a vehicle system must fall within a guaranteed temperature range.

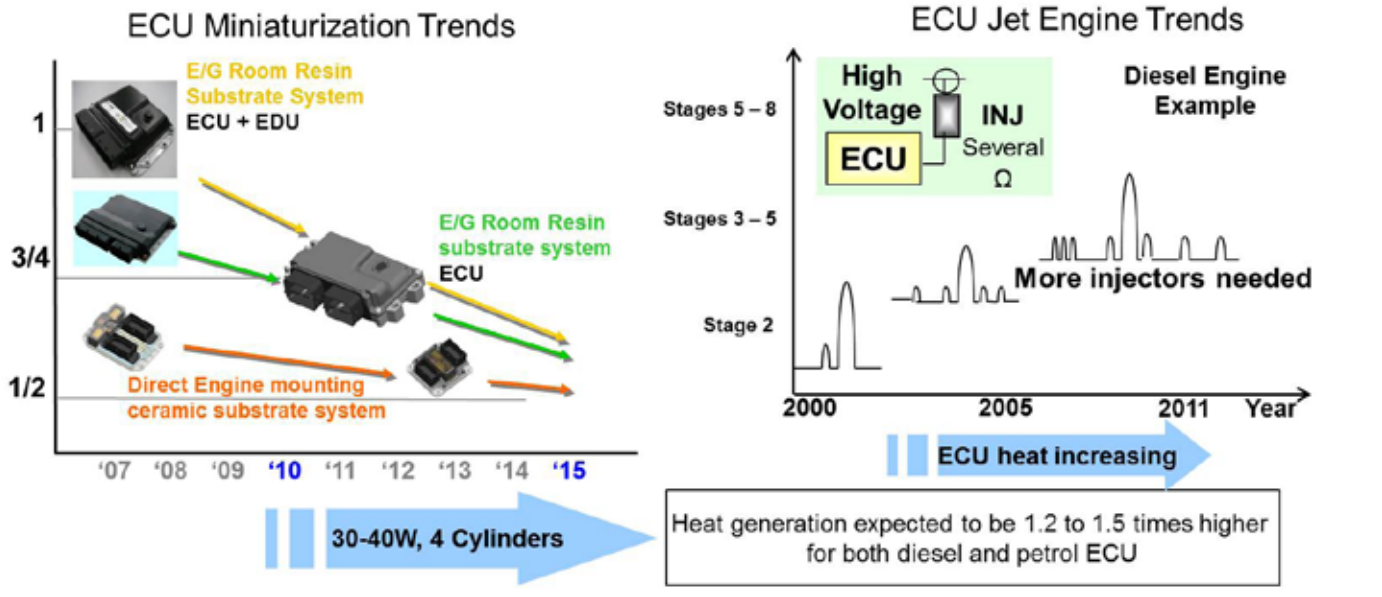

Figure 1. Downsizing drives technology for heat dissipation

Mentor Graphics — Mechanical Analysis

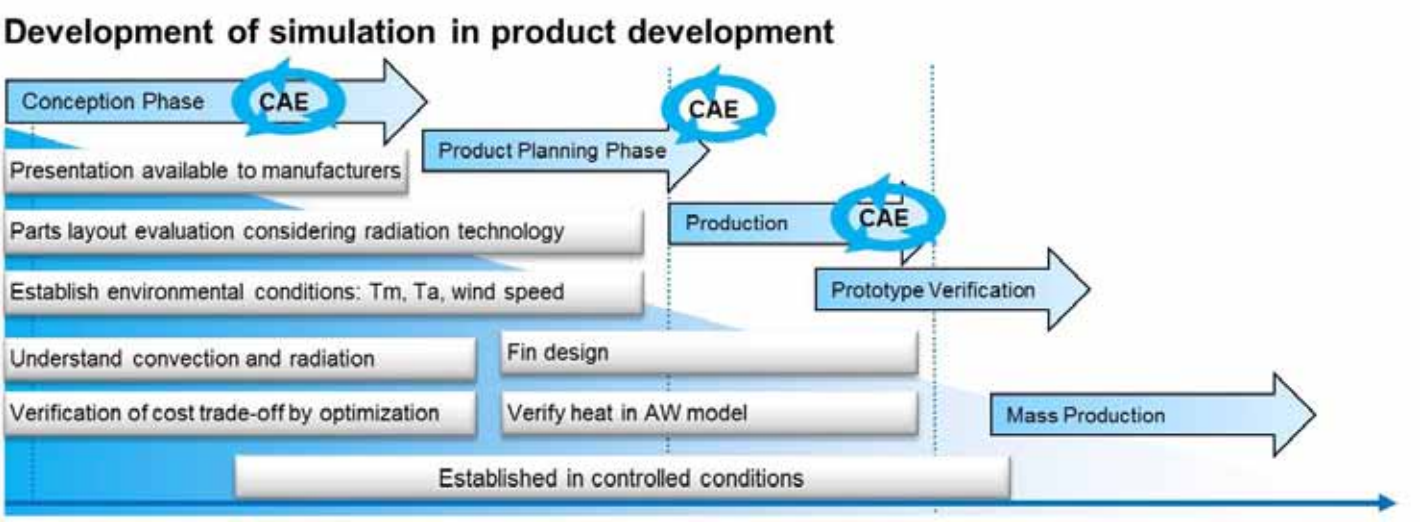

Figure 2. Technological innovation in thermal management

As it has been impossible to directly measure the junction temperature, engineers used to predict electronic elements' junction temperature based on the assumption of the measured surface temperature and set a wider design margin. In order to cope with the current aggressive price competition, it is important to secure quality, optimize the design margin, and achieve overall cost-effectiveness.

Simulation Saves Design Time and Cost

Shinoda first started looking for a thermal design tool in 2006, having seen fluid flow visualization of a circuit board for the first time at an exhibition. DENSO selected FloTHERM® and FloTHERM® PCB for their thermal design flow over other tools through a rigorous benchmarking exercise. Before DENSO started using thermal simulation, a physical ECU prototype had to be created for temperature measurement at the early stage. As a result 2-3 weeks of preparation was needed to perform a day's testing, which might have to be repeated several times before the product was finalized.

Since 2006, DENSO have consistently increased the use of simulation to reduce the time and cost spent on physical prototyping. By 2009 the ratio of simulation to physical prototyping was 20:80, this was increased to 50:50 by 2010. By 2012 the ratio stood at 70:30. This change has resulted in a 50% reduction in both the duration and costs associated with thermal design in less than six years. DENSO plans to go further, aiming to increase the ratio to over 90:10 by 2015.

Benefitting from Centralized Thermal Expertise

DENSO has successfully introduced thermal analysis into the manufacturing process by centralizing their thermal technology and then making this expertise available throughout the company. By listening to the needs of the various design departments, Shinoda's thermal group is able to quickly improve the quality of design by improving the efficiency of the heat removal using thermal simulation backed up by experiments.
In the thermal design, designers usually focus on changing the case form factors to improve heat dissipation. The best results can be achieved by sharing the case design and electronic design among members of the thermal design team. DENSO had decided to use existing component models (Figure 4). The mechanical team created a smaller case, the circuit design team redesigned 10 to 20 % of the circuit according to new specifications, and the measurement team took the temperature for thermal analysis. As a result of this collaboration, DENSO were able to create a working thermal model for the product in two days. Engineers from each team contributed to this effort.

Characterization Supporting Simulation Accuracy

Besides the aforementioned desire to move away from physical prototyping, measurements have an important role to play in DENSO's thermal design process. To support their thermal simulations, DENSO uses Mentor Graphics' T3Ster® to characterize ECU components and thermal interface resistances in situ. The accuracy of T3Ster data has enabled DENSO to increase the accuracy of their thermal simulations and given them the confidence to place such a heavy reliance on simulation. Measurement of T3ster data is taken back into FloTHERM to improve junction temperature prediction accuracy during design to ensure that junction temperatures never exceed their allowed limits. This is quite a tall order, and requires a very high

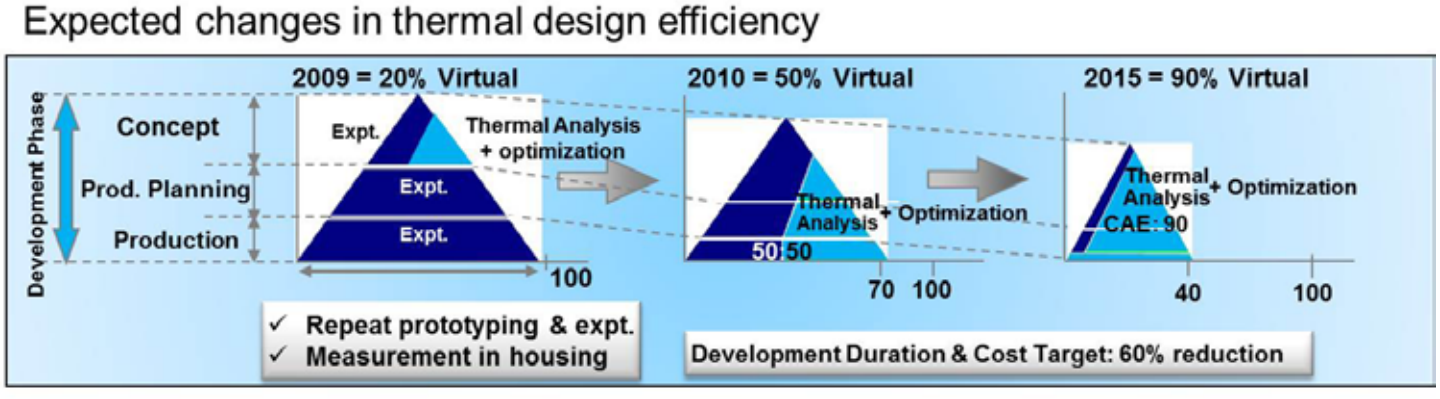

Figure 3. Technological innovation in thermal management

"Our PCB designers use FloTHERM PCB, which has a user-friendly user interface and connection with FloTHERM PACK and we get excellent support from IDAJ and KOZO KEIKAKU ENGINEERING who distribute Mentor Graphics thermal analysis tools"

Takuya Shinoda, DENSO Corporation

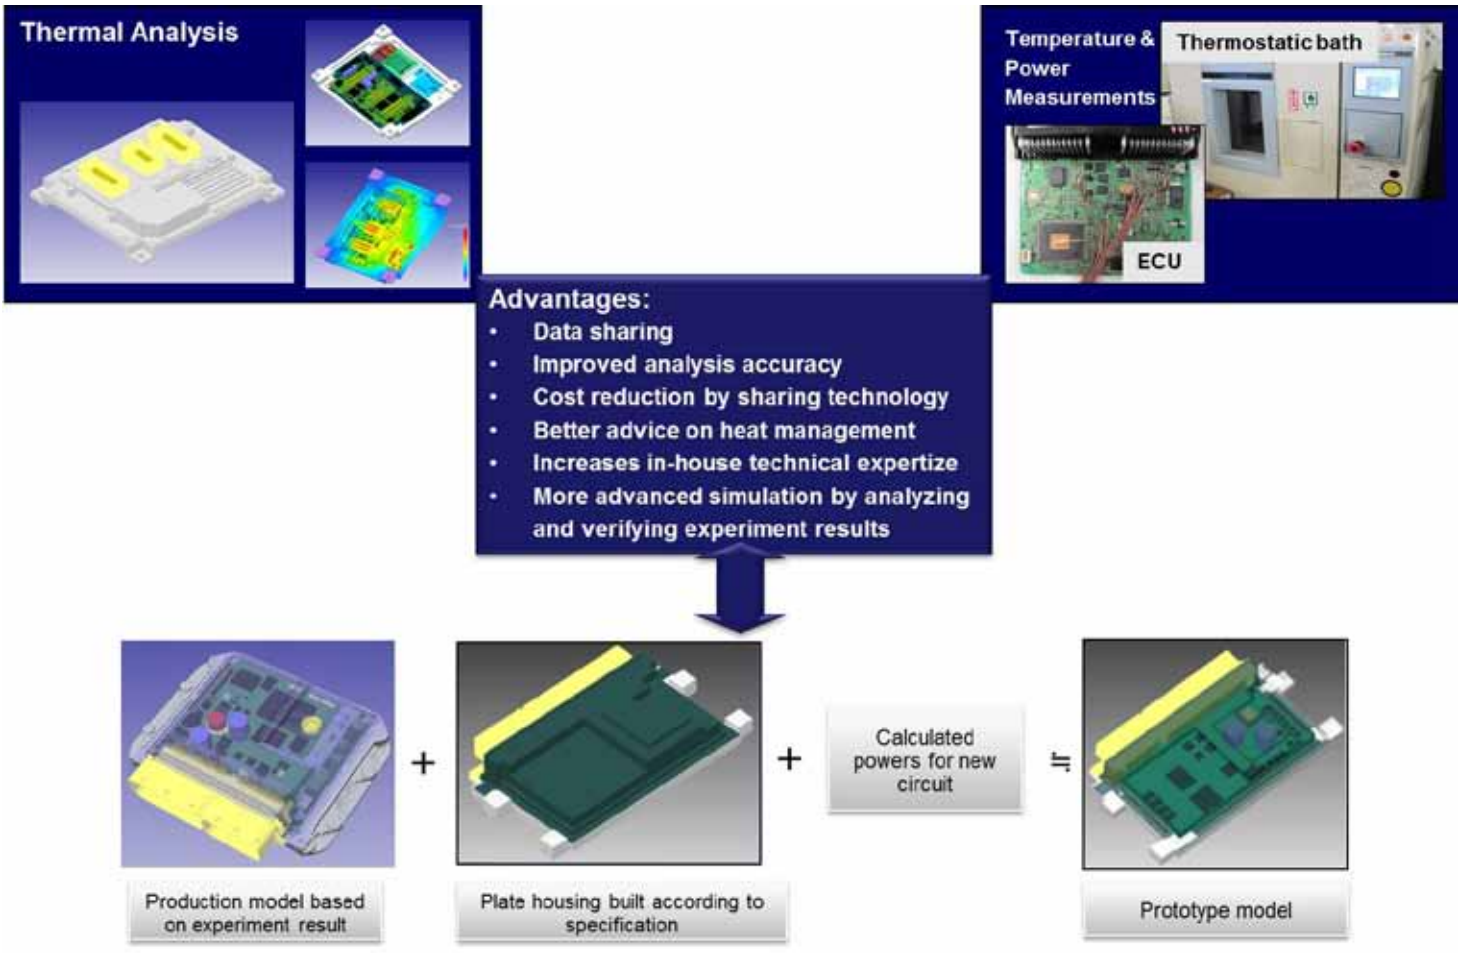

Figure 4. Advantages of centralized thermal technology - Simulation & Measurement

level of confidence in the simulation models. Today, agreement on junction temperatures rise is to within 10% of experiment, and DENSO aim to increase the agreement further, to be within 5% by 2015.

"JEDEC JESD51-14 standard was issued in 2010. It has far exceeded accuracy and repeatability compared to the steady-state measurement that conforms to older standards. T3Ster is the only product available in the market that complies with this new standard, enabling the accurate estimation of thermal resistance and junction temperature. Also, using the structure function, a unique feature of T3Ster, a simple and accurate element model can be generated from the measured data" explained Shinoda.

DENSO has found it indispensable to take accurate measurements from the electronic elements, to improve simulation accuracy, and hence eliminate excess thermal margin from the design.

The FloTHERM suite of products, including FloTHERM PCB and FloTHERM PACK, has become a major toolset and used across the whole of DENSO's thermal design process. Backed up by high accuracy package thermal models, material property data and interface resistance values obtained through measurements using T3Ster, Mentor Graphics' thermal solutions have helped DENSO to achieve over 90% virtualization in their thermal design, and a reduction of more than 50% in both development time and cost, with further savings expected for the future.

With thanks to Takuya Shinoda, Project Assistant Manager at DENSO Corporation
www.globaldenso.com

How small can you go?

A detailed CFD study to optimize the design of multiple compact heatsinks in the mobile communication module of radio access stations

By Seung-Hyun Jeong, Hangang Universtiy

Mobile cellular phone subscriptions hit the six billion mark worldwide in 2011 [1], resulting in an increased demand for base stations. This in turn drives the need to build even smaller, more compact stations which are unobtrusive. Under these circumstances more compact heatsinks become critical to keeping component temperatures below threshold levels for optimal electronics performance. A 2010 study [2] conducted by engineers at the Hanyang University in Seoul addressed this issue using the market leading FloTHERM™ electronics cooling simulation tool from Mentor Graphics.

It was used in association with a new process integration and design optimization (PIDO) tool, PIAnO™, from PIDOTECH Inc. in Korea [3]. Their detailed CFD study looked to optimize the design of multiple compact heatsinks in the mobile communication module of these radio access stations (Figure 1) which enable subscribers' wireless internet connections while on the move.

The compact heatsinks integrate into the front and rear housing of the radio access station. This makes them thermally stable over a wide range of ambient temperatures and they operate with only natural air convection currents to cool the electronics. Natural convection is preferable because of its silent operating mode as opposed to the alternative of fan cooling. The challenge posed by the researchers in this multivariable design project was to find the optimal heights, thicknesses and base thicknesses of the heat sinks, as well as the gap between the sun-shield and the heat sink in the front of the station (Figure 2).

The ultimate goal is to minimize the volume of the system while satisfying the constraints on the junction temperatures of 12 key components in the system (Figure 3).

FloTHERM® CFD simulations (Figure 4) obtained the temperature distribution in the unit for extreme operating conditions. PIAnO was employed to execute an array of CFD simulations needed for a Design of Experiment and to automate the procedure for the multivariable design scenario. In order to obtain an approximate optimal design solution, the thermal analysis results were obtained at 54 experimental sampling

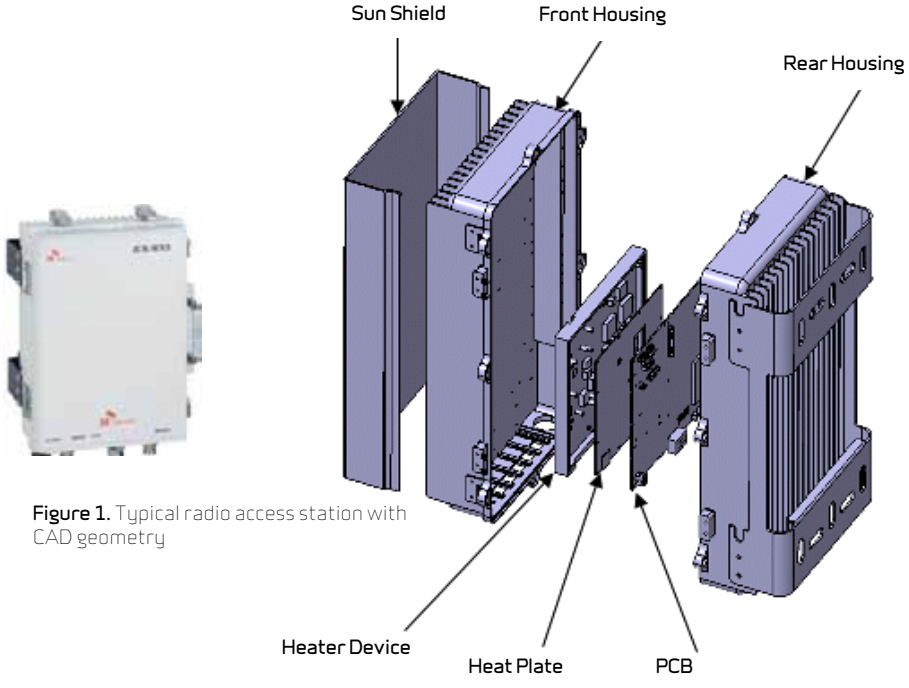

Figure 1. Typical radio access station with CAD geometry

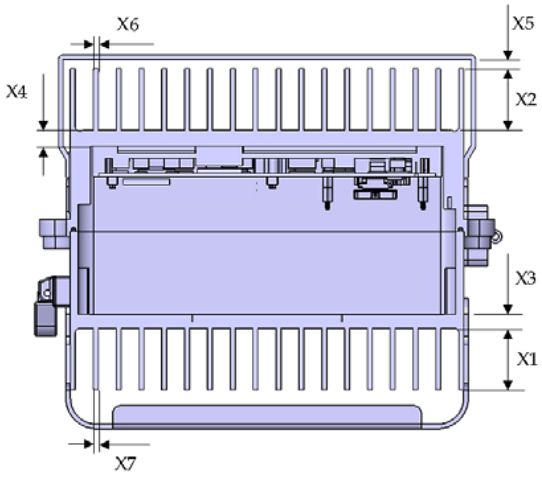

Figure 2. Typical radio access station heat sink variables

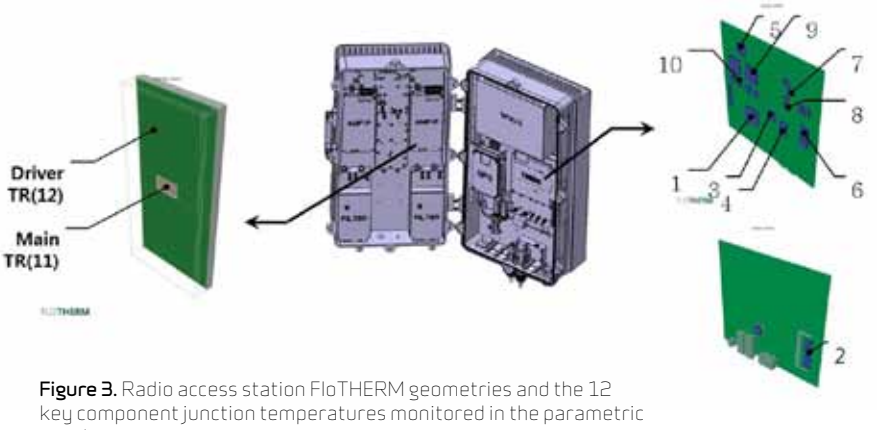

Figure 3. Radio access station FloTHERM geometries and the 12 key component junction temperatures monitored in the parametric simulations

points. These points were specified by an orthogonal array L54(21 x 325) and then full quadratic polynomial regression models were built to approximate temperatures at the 12 important locations. Design optimization was performed using the approximate models.

This design optimization study with PIAnO and FloTHERM resulted in a reduction of volume of the base station unit by 41.9 % while satisfying all the design constraints (Figure 5). The approximate optimum temperature values were also measured and found to be almost the same as those obtained by the FloTHERM simulation at the optimal design point, confirming the validity of this design approach.

References:

[1] "Measuring the Information Society 2012". International Telecommunications Union (ITU)

[2] "Optimal Design of a Heat Sink for Mobile Telecommunication" by Seung-Hyun Jeong, Hyun-Su Jeong, Yongbin Lee and Dong-Hoon Choi, 36th China-Japan-Korea Joint Symposium on Optimization of Structural and Mechanical Systems, June 22-25, 2010, Kyoto, Japan

For more information on PIDOTECH Inc. and PIAnO:
www.pidotech.com/en/product/PIAnO

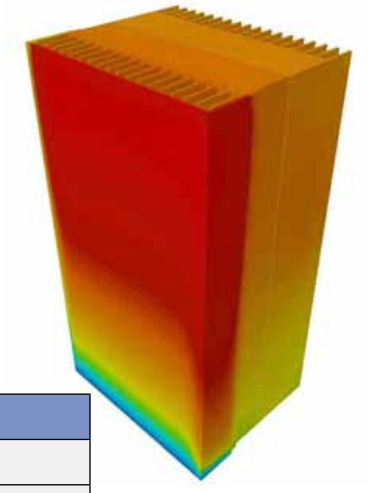

Boundary Conditions	
Ambient Temperature:	50°C
Altitude:	1,800m
Solar Load Heat Flux	753W/m²
Related Standards	Telcordia GR-63-CORE

Figure 4. Boundary conditions & typical FloTHERM calculation of the base station box

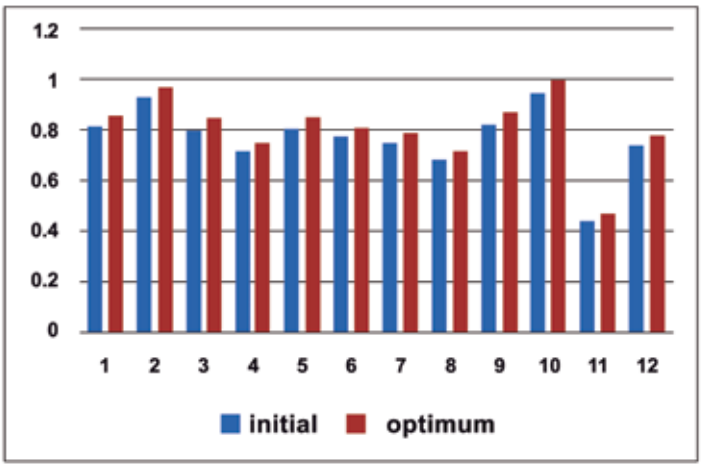

Figure 5. Difference between initial and optimal design of the base station box for the 12 key junction temperatures

Curtiss-Wright & FloTHERM™: From COTS to Custom Deployment

Curtiss-Wright is a provider of rugged, commercial off-the-shelf (COTS) electronic modules and integrated systems for defense and aerospace applications.

By Andrea Schott, Curtiss-Wright Controls, Defense Solutions

Having both a mechanical design and thermal design background has helped me appreciate the many different aspects of taking a product from concept to production in a short amount of time. Unless you are the sole engineer on the project, doing both thermal and mechanical design, a close working relationship between thermal engineering and mechanical design is critical to ensure on-time, low cost product delivery.

Curtiss-Wright provides rugged, commercial off-the-shelf (COTS) electronic modules and integrated systems for defense and aerospace applications. Our highly engineered solutions, ranging from open standards-based modules to fully optimized systems solutions, are deployed in a wide range of demanding applications, including C4ISR, unmanned systems, mission computing, fire control, turret stabilization, data recording and storage solutions.

As part of the Defense Solutions business unit of Curtiss-Wright Controls, the Littleton facility, Massachusetts addresses a niche market by providing quick turn-around custom electronics enclosures to our customers. Typically, this involves very low volume (quantities between 1 and 15 units), and production delivery in around 20 weeks, which leaves no time for prototyping or testing. We rely completely on engineering experience and thermal simulation to meet our requirements. FloTHERM® allows us to iterate multiple scenarios to optimize our systems for not only thermal performance, but also weight reduction, noise reduction, cost and schedule.

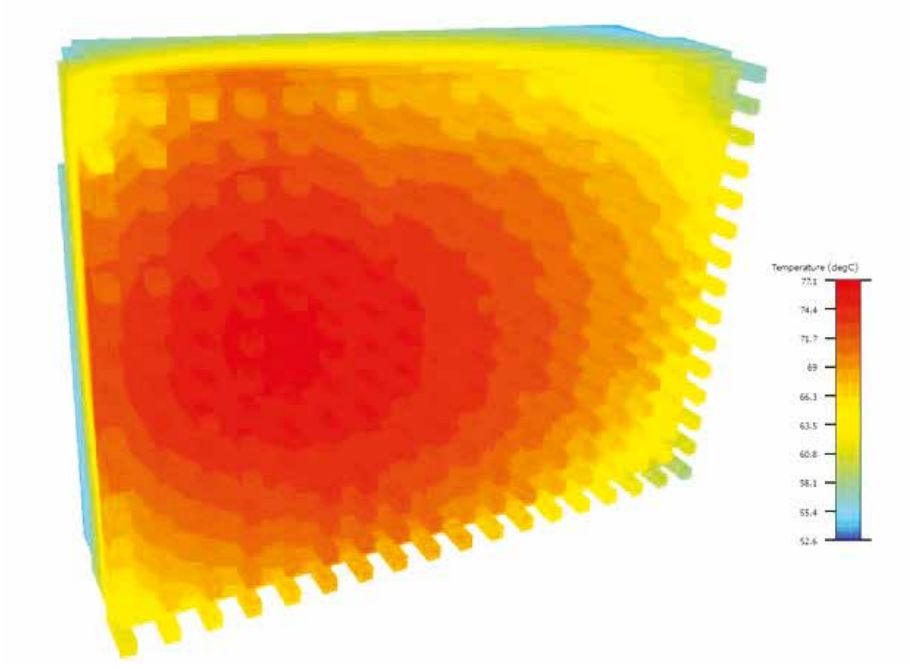

Figure 1. High density, conduction cooled power supply with supplemental air cooling via pin fin heatsink

As a result of FloTHERM's versatility we have the ability to offer our customers a variety of solutions. To further define the operations of Curtiss-Wright's Littleton facility, our activities often involve supplying a metal enclosure (typically brazed aluminum),backplane and power supply all designed to meet specific customer specifications. Our customers populate the enclosure with their own suite (or payload) of electronics. We are usually provided very little information on the design or end-function of the payload. What keeps this process from being straight-forward is the fact the Curtiss-Wright is typically only provided with specifications for the subsystem's overall sizes, power levels, ambient conditions and the temperature requirement for the electronics card mounting. From this limited amount of system detail it is our task to design a solution capable of meeting the required temperature in all environmental conditions at the given power levels. Another challenge that our design team often faces is that in many cases enclosures designed at our facility and are sold to the end customer by a third-party for whom the Curtiss-Wright designed enclosure is essentially a component, not a product level complete system. Because our customers are usually responsible for all verification testing, we rarely receive feedback about results except for the very rare case in which a problem emerges.

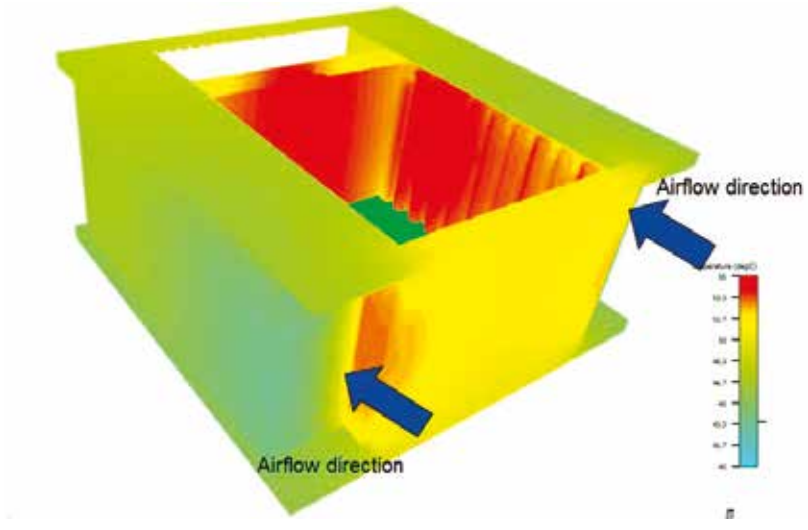

Figure 2. 1300W Custom Electronics Enclosure. The payload is cooled by conduction, Enclosure is air cooled through side walls

As these enclosures are primarily used in military applications, the environmental conditions can vary greatly and are often extreme. New products are often retrofits for older existing equipment and the new higher powered enclosures must be cooled by existing cooling systems. The challenge for our design team, including thermal and mechanical engineers working together, is to meet all of the customer's requirements in a very short timeframe.

The ambient conditions in which the resulting system must perform are harsh, and the power levels required are typically high. Hence thermal simulation is critical for our business, especially since delivery schedules are often very tight.

In order to expedite the design process, Curtiss-Wright has established a consistent method of tracking all thermal design information. This includes tracking the initial customer requirements all the way to documenting the 'as built' configuration. By using a custom spreadsheet template containing the relevant design information we are able to track and maintain information from project to project in a consistent manner. The spreadsheet template contains the initial customer requirements that are used in the quoting process. Since thermal design is one of the highest risk factors in almost every design we undertake, we closely analyze the project's thermal requirements even before we are awarded the job. In many cases we perform some preliminary level of simulation work before the contract has been formally awarded. This initial amount of simulation ensures our customers that we are able to solve their particular design problem. The thermal spreadsheet template also provides our Applications Department a starting point for new designs and saves time in the quoting process.

Geometry used in very early simulation work may turn out to be quite different from what mechanical requirements will later dictate. Frequently, the enclosure space is dictated by mechanical constraints that are not fully defined at the quoting stage. This is because numerous system features such as I/O connections, cabling and air plenum allotment may not yet be determined at this early stage of development. After the results of a preliminary thermal simulation indicate that the customer requirements can be met, the mechanical design process begins. There are often several iterations back and forth between the design engineering and thermal engineering teams to reach a final solution. One design aspect that demands this level of attention is fin optimization. FloTHERM makes it quick and easy to optimize for pitch, thickness, number of fins or base thickness. Each product we design is fully customized so there is very little opportunity for design reuse. For example, the cooling wall (our heatsink) geometry is designed for each particular application to ensure the best design at the lowest cost for each product. While the final solution may be similar to the starting point in the thermal design, it is never exactly the same. Although the time and cost saved by eliminating early prototypes, testing, evaluation and redesign is hard to measure, when delivery schedules are as tight any and all time savings are crucial to our customer's success.

Another key to successful design is a team that works extremely well together. Curtiss-Wright's Littleton facility is staffed with a group of talented mechanical engineers who all work in sync to meet the end goal, which is to deliver a high quality product to our customers every time.

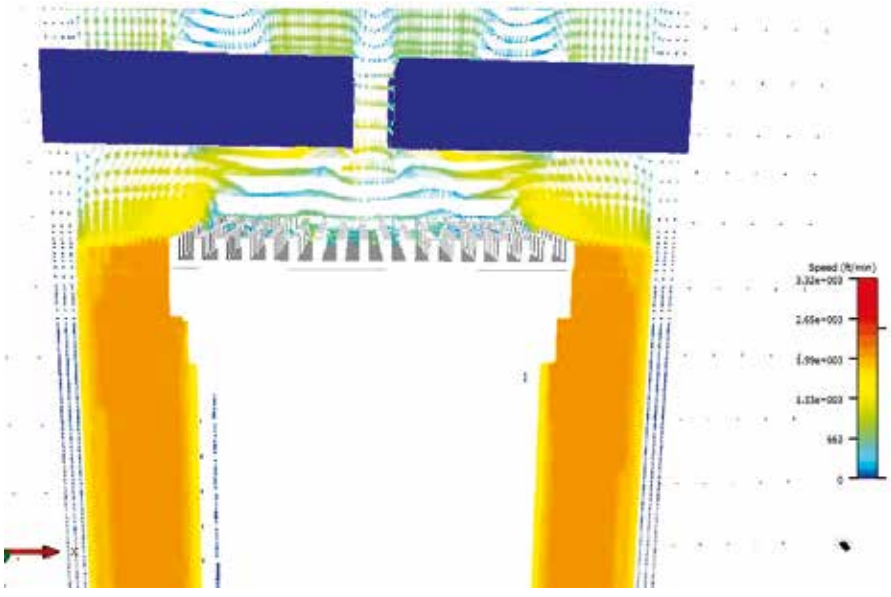

Figure 3. Air flow profile of 1300W enclosure, showing power supply fins and system fans

About Curtiss-Wright Corporation

Curtiss-Wright Controls Defense Solutions (CWCDS) is a long established technology leader in the development of rugged electronic modules and systems for defense applications. CWCDS serves as a technology and integration partner to its customers, providing a full range of advanced, highly engineered solutions from modular open systems approaches to fully custom optimized solutions. Our unmatched capabilities and product breadth span from industry standard based COTS modules to complete electronic subsystems. The company's modules and systems are currently deployed in a wide range of demanding defense & aerospace applications including C4ISR systems, unmanned subsystems, mission computing, fire control, turret stabilization, and recording & storage solutions. Additionally, the company's broad engineering capabilities combine systems, software, electrical, and mechanical design expertise with comprehensive program management and a broad range of life-cycle support services.
For more information visit www.cwcdefense.com.

CURTISS WRIGHT Controls Defense Solutions

Mentor Graphics Mechanical Analysis

Making Computers Cool by Design

NMB Minebea use thermal simulation to predict design flaws

By Dr.-Ing. Anton Breier,
NMB Minebea GmbH

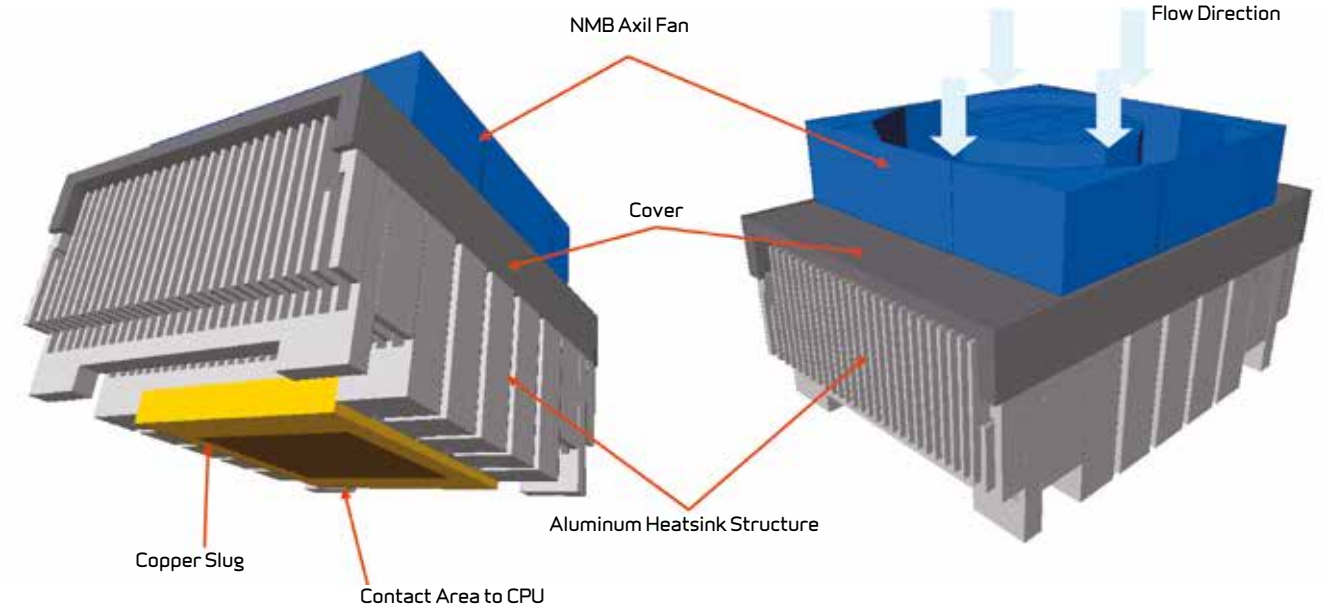

Figure 1. FloTHERM Simulation Model of CPU

CPU cooling is a critical aspect of a functioning computer system, and for this reason the need for forced-air cooling is a significant factor that should be determined at an early stage in system design. Good airflow to heat-generating components, and adequate space and power for the cooling fan are a critical design requirement for any forced convection system. One of the first steps in system design should be to estimate the required airflow. This will depend on the heat generated within the enclosure, and the maximum permitted air temperature rise.

The question becomes, how can we best determine the optimal design for computer cooling, and in particular, for CPU cooling? The answer from fan experts NMB Minebea is to use thermal simulation. Using FloTHERM® from Mentor Graphics, NMB Minebea can anticipate design flaws, and evaluate the thermal behavior of the various components that will ultimately become an important part of the final cooling design. By way of example, here's a quick view of how thermal simulation can assist in predicting and improving the airflow of an average CPU application in order to produce optimal CPU cooling within a desktop computer. The structural design of a typical CPU will look similar to Figure 1.

In this example, the axial fan sits on top of the cover of the heatsink structure. Under the heatsink structure you will find the thermal slug (shown in yellow) and the contact area to the CPU surface. Between the heatsink fin tips and the cover it is necessary to provide a gap in order to allow air to exhaust from the fan.

The FloTHERM results show the distribution of the surface temperature on the heatsink and CPU for this particular heatsink-fan (or 'fansink') combination. As expected the highest temperature is seen directly at the contact area between the CPU and heatsink.

Heat is transported inside the heatsink material by conduction, and from the surfaces to the air by convection. The temperature profile within the heatsink, and the resulting CPU temperature, depend on both the conduction and convection within the assembly. The heatsink has to be designed to deliver the best performance for the chosen fan in order to optimize the entire fansink design. Many factors, including the material, fin design, air velocity and surface treatment all influence the thermal performance of a heatsink.

"Increasing performance gives rise to problems related to equipment cooling. During the development phase, thermal simulations provide us with crucial information about airflow distribution as well as both air and component temperatures. The use of Computational Fluid Dynamics Software not only eliminates the need for thermal redesigns, but also facilitates shorter development times and optimized equipment cooling. We offer our customers complete system solutions, combining fans, heatsinks and power supplies with layout and dimensioning, optimized precisely to the customer's equipment."

Dr. Anton Breier, Deputy General Manager, NMB-Minebea-GmbH

As the cooling air from the fan does not provide a uniform flow pattern, there can be considerable temperature variation within the heatsink. The insight that FloTHERM provides indicates exactly what measures should be taken for optimal cooling. In general, if high temperature gradients are observed within a heatsink then conduction should be improved. This can be done by choosing another material (alloy) with a better thermal conductivity, or by increasing the local cross-sectional area to improve the heat conduction.

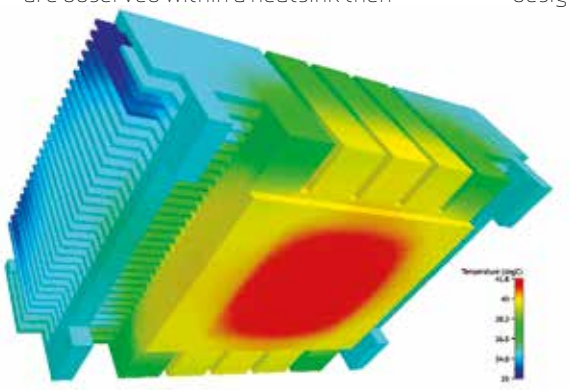

Figure 2. CPU Surface Temperatures

Looking at the flow vector field in a center cross-section, above left, FloTHERM reveals a zone below the fan motor with almost no air flow. This behavior is seen for all axial fans used in this design of heatsink. As a consequence of the stagnant air below the fan motor, very little heat is removed from the heatsink fins in this area and so almost no cooling occurs. The fins heat the stagnant air close to the fin temperature as shown on the right hand side of the graphic above.

At the system level, obstructions in the airflow path increase the static pressure drop within the enclosure, reducing the air flow, so airflow obstructions should be minimized. Obstructions in the form of baffles are sometimes necessary to direct the airflow over the components that need cooling. This can lead to unforeseen consequences, such as regions of flow recirculation can occur behind the baffles leading to unexpected hot spots. This is why system-level simulators ,like FloTHERM, are critical for good equipment thermal design. The final system design should show continuous airflow through all parts of the enclosure for optimum thermal management of all heat generating components.

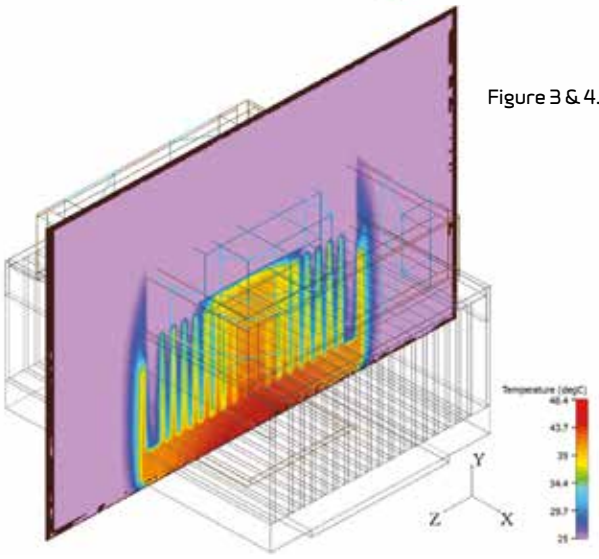

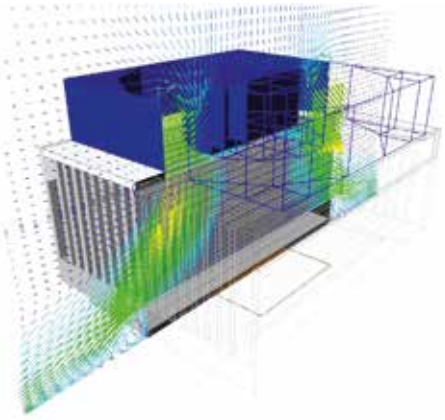

Figure 3 & 4. CPU Cooler Analysis

To assist customers, NMB offers Thermal Management Consultancy as part of their design services. This has been proven to facilitate the design process of many of NMB's customers by providing key information and solutions for their thermal cooling applications.

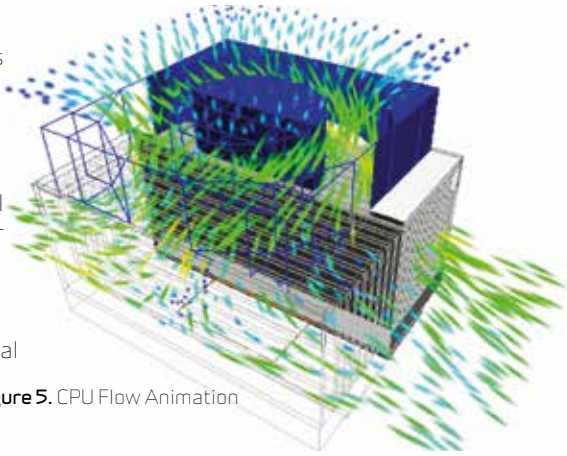

Figure 5. CPU Flow Animation

The Latest Philips TVs make LEDs Dance

'Dancing' above refers to dimming and boosting of LEDs in time with the picture

By Wendy Luiten, Philips

In recent years there has been an explosion in LED-LCD displays, and a proliferation of LED TVs in people's homes. They offer a great picture, great styling with a thin form factor, increased functionality like 3D and internet access, and great value for money, but cooling the LEDs, while steadily increasing power densities in ever thinner product enclosures poses a big challenge for designers

In an LED, die (junction) temperature affects both performance and lifetime as LED lumen efficacy is lowered and the color temperature shifted. Heat degrades the epoxy lens, and both the absolute temperature and the temperature distribution over the LCD screen can lead to performance and lifetime issues. Non-uniform LED temperatures lead to unwanted spatial non-uniform lighting in terms of color uniformity and brightness. Effective heat spreading is therefore a key goal of the thermal design.

The temperature of the LCD has to be just right. Too hot (typically 60 - 70°C) and optical materials age , too cold and the switching is sluggish, causing fast changing pictures to display blurred - very performance degrading in modern high resolution 3D TVs! At a system level, the temperature the user experiences when touching the TV is a key safety issue.

While it is entirely feasible to do electrical and optical tests in a standalone setting

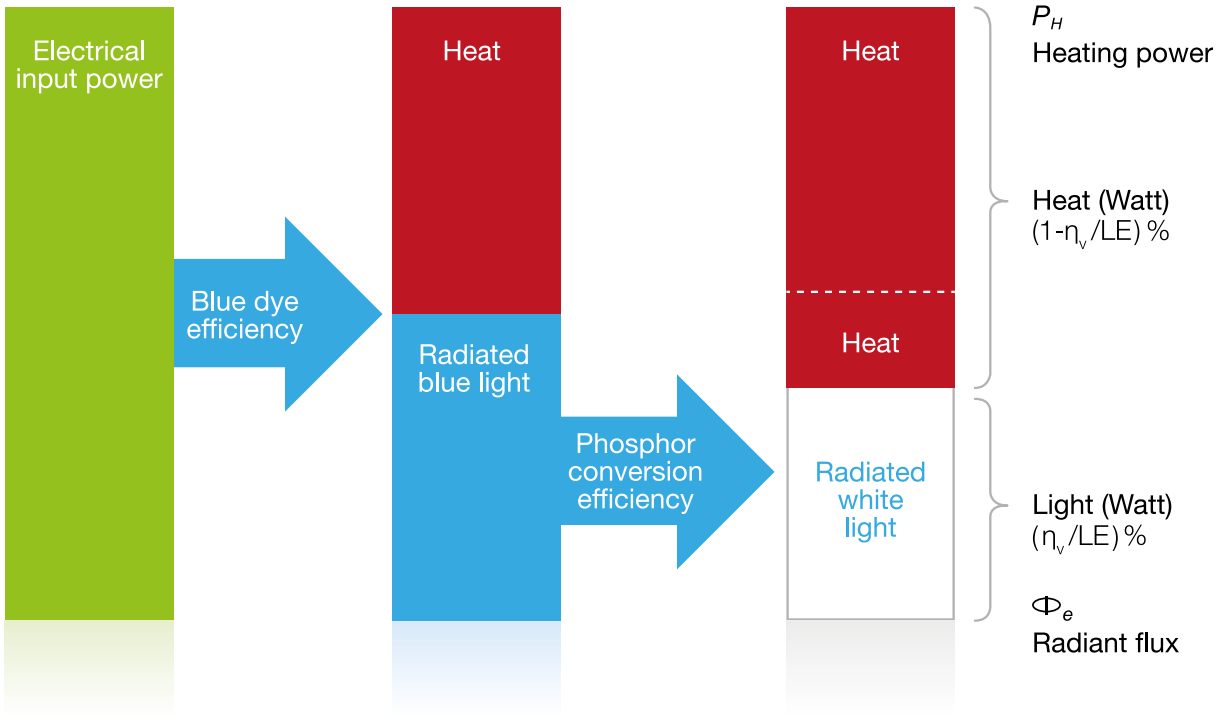

Figure 1. Power Balance of a White LED

and directly translate the results to the system situation, for thermal this is not the case. The presence of the set back cover, which causes an additional thermal resistance from the display back to the room temperature, and second, the direct thermal interaction between the display and boards illustrates the difficulties encountered in co-design.

Philips have been pioneers in the TV industry since the 1940s, so the chances are the first TV you ever saw was a Philips. In fact, it was probably the first TV your grandparents ever saw. Philips has used FloTHERM extensively in the design of TVs since the software was first released in the early 1990s.

Today TPVision, the maker of today's Philips TVs, uses FloTHERM to look at all aspects of the thermal design of LED LCD TVs from the LED package, module, up to the full TV in different environments, starting from conceptual design through to the final product, to optimize all aspects of the thermal design. Factors considered include evaluating different design architectures, such as direct lit display vs. edge lit displays, and optimizing the cost benefit of different cooling solutions.

While LEDs are used to light the display, there is a world of difference between using LEDs for display backlighting and using LEDs in a lighting product. In lighting applications LEDs are typically used in a steady state manner, with timescales in the order of hours or longer. In a TV two different timescales prevail. Heating and steady state behavior of the set as a whole is governed by the average power consumption, and this has a long timescale similar to domestic lighting applications. However, the moment-to-moment changing of the video content happens at a much higher frequency, and this creates an additional much shorter timescale.

A further complicating factor is the thin film transistor (TFT) panel that covers the LCD-LED display. Light emitted from the TV is a combination of the light emitted by

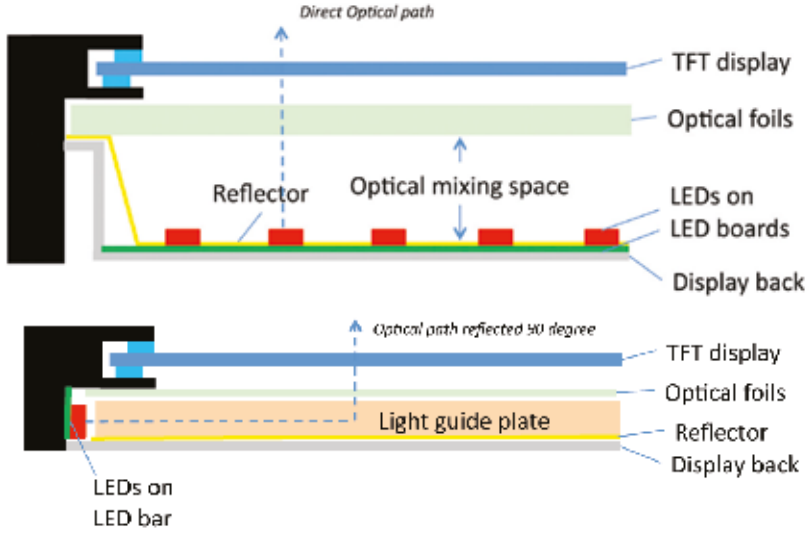

Figure 2. Principle of Direct Lit (top) and Edge-lit (bottom) LCD-LED Display

the backlight and the state of the pixels in the TFT panel. If a TFT pixel is open, light goes through to reach the viewer. If a TFT pixel is closed, light is blocked. In display with a static light emitting backlight, pixels are predominantly open in a bright image, and pixels are predominantly in the 'closed' state if the image is dark. However, from a picture quality and energy consumption point of view, a static light output from a backlight is not optimal.

To improve the picture quality (deeper black) and reduce energy consumption it is common to dim the LEDs in dark scenes. A refinement is to use 2D (or local) dimming, where the LEDs are dimmed not only in time, but also depending on location, providing a further improvement both in picture quality (higher contrast over the screen area) and in energy consumption. However, it is possible to go further still.

The timescale over which the area of a TV lit by a single LED changes is one to two orders of magnitude smaller than typical thermal time constants for display LED packages. As well as dimming LEDs during dark scenes it is also possible to boost their light output for short periods of time. LED dimming and boosting scenarios, while not primarily intended as thermal control measures, are very beneficial to the thermal management of LED-LCD TV sets as the associated LED temperature is more highly correlated to the average LED power, which is much lower than the peak. The result is exceptional picture quality.

Philips TV takes a similar approach for the ambilight feature: the LED temperature is determined by the average LED power dissipation, and large instantaneous LED peak powers can be allowed to increase the immersive experience. Tight thermal management algorithms are deployed to prevent LED boosting from adversely affecting lifetime and reliability of the display.

FloTHERM® simulations were performed on a stand-alone direct lit display, cooled by a heat transfer coefficient typical of natural convection including radiation on the front and on the back. Figure 4 illustrates an important trade-off between the number of LEDs and thermal issues.

The calculated temperature field compares well to the measured temperatures on the front of a direct lit TV set, in Figure 6
In the direct lit TV, the effect of hot air rising is visible, as higher temperatures at the top of the display. Also, the positions of the three boards are visible as locally higher screen temperatures. The infrared picture

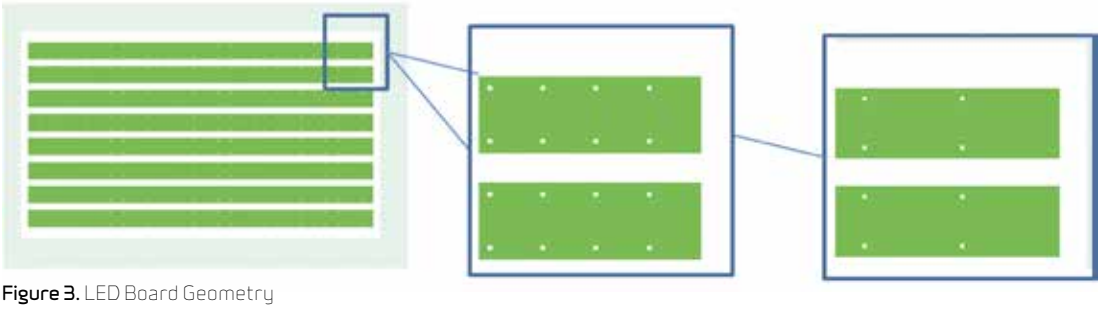

Figure 3. LED Board Geometry

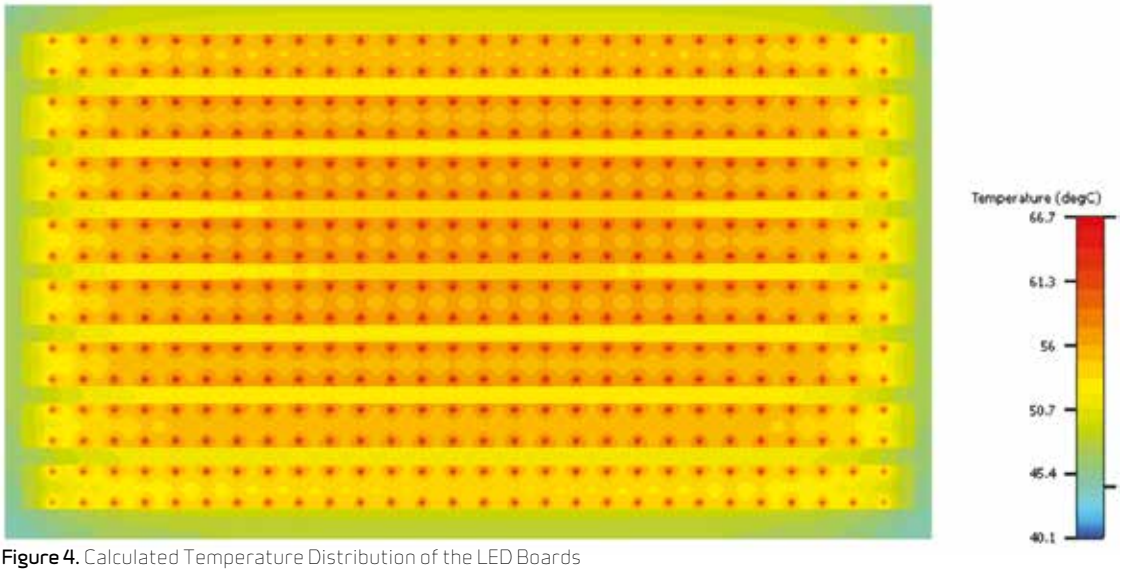

Figure 4. Calculated Temperature Distribution of the LED Boards

"FloTHERM has long been an important tool to our LED-LCD TV thermal design, and is routinely used at module (display) level and at system (TV) level as the temperature distribution in the LED-LCD display is a system-level issue due to the strong thermal interactions. FloTHERM helps us select and optimize the thermal solution so we have confidence from a very early design stage."
G.A. Luiten, Philips Research, Eindhoven, Netherlands

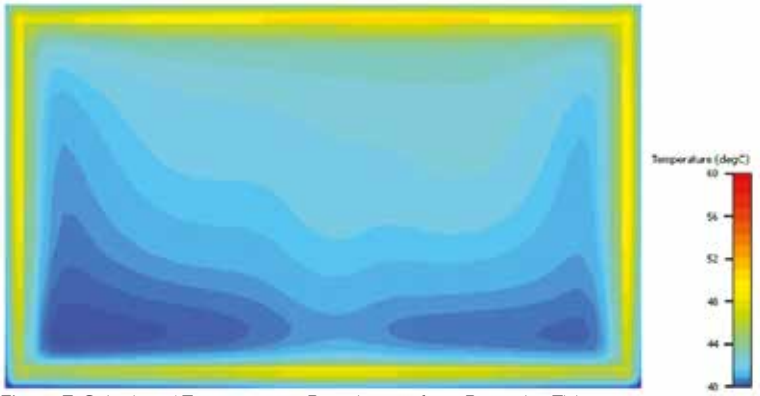

Figure 5. Calculated Temperature Distribution for a Direct Lit TV

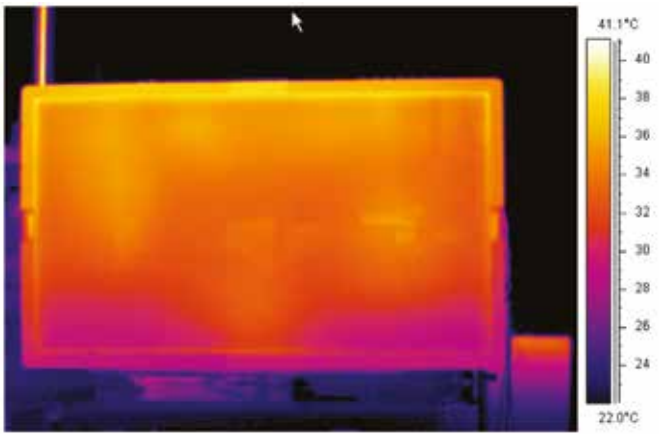

Figure 6. Infrared Picture of a Direct Lit TV

confirms that screen temperatures are well below the aging limit of approximately 60 °C (in 35 °C ambient) and that the temperature difference over the screen is around 8 °C Figure 17 shows the infrared images of the front of a side lit TV, equipped with internal graphite heat spreaders. Comparison of the infrared images with the simulation results shows good agreement. The measurement confirms that the screen is critical with respect to the aging criterion in the zone directly adjacent to the LED bars, and that there is roughly 20 °C temperature difference between the high temperatures at the side and the temperatures in the center.

Acknowledgement
With thanks to TP Vision Innovation Lab, Eindhoven and TP Vision TV Development, Bruges, formerly known as Philips TV

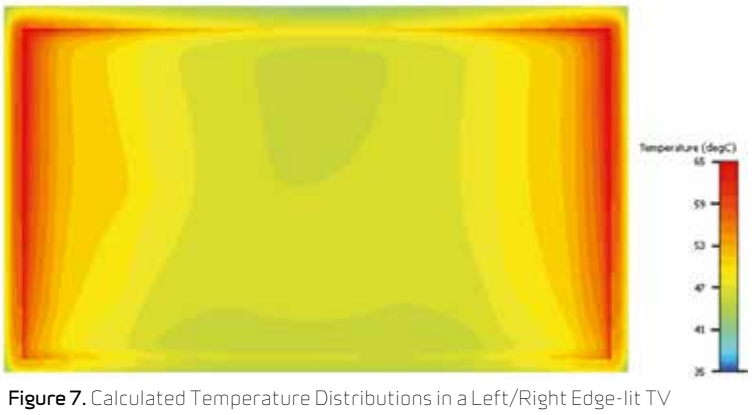

Figure 7. Calculated Temperature Distributions in a Left/Right Edge-lit TV

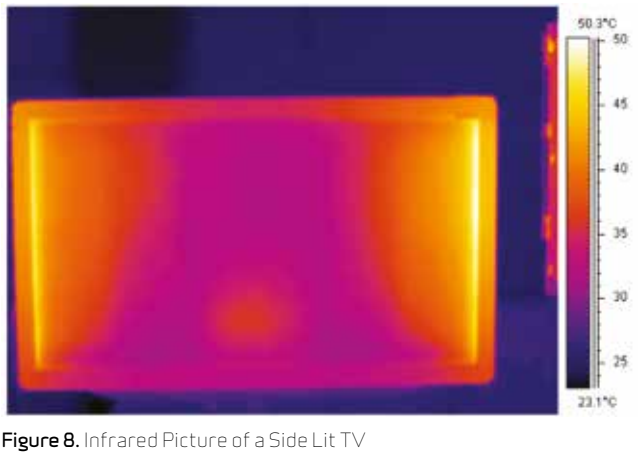

Figure 8. Infrared Picture of a Side Lit TV

Divide & Conquer

Power Savings from Server to Datacenter

A study by The Department of Engineering at the University of Taiwan to reduce energy costs and improve efficiency in datacenters

By John Parry
Electronics Industry Manager,
Mentor Graphics

Increasingly organizations are choosing to host their data requirements in large purpose-built, energy-hungry datacenters. Datacenters house many racks and a large number of servers where a significant amount of heat is generated from the IT equipment. To remove the heat from equipment so that the electrical components can operate normally, cooling systems must be introduced to provide adequate cooling. Fans, Computer Room Air Handlers (CRAH), and chillers, consume 35%-45% of the total power budget. A big concern for datacenter operators is reducing energy bills and one approach is to improve cooling efficiency.

Figure 1 shows one datacenter room with a single CRAH and 20 racks; each rack is 42U high, and so can house 42 1U-servers. The datacenter includes a raised floor to provide cold air from the CRAH for the IT equipment and a dropped ceiling to draw hot air from the IT equipment into the CRAH. This is a typical datacenter arrangement in which some of the hot and cold air mixes, increasing the room and server inlet temperature. To provide a suitable ambient environment for IT equipment, lowering the temperature may require a higher cooling capability from the CRAH, which will waste more power.

Divided Zone Partitioning

A recent study by The Department of Engineering at the National Taiwan University explored the idea of a divided zone approach to cooling efficiency. A divided zone partition works by concentrating airflow for key components to avoid airflow bypass and controlling different individual zones independently with the aid of a Fan Speed Control (FSC) for the system.

Mixed airflow challenges can be overcome by cooling airflow path management to improve cooling efficiency and power saving [1]. Some datacenters implement hot and cold aisle containment [2, 3], while Zhou et al. [4, 5] propose adaptive vent tiles that can vary their opening for air flow adjustment. Most hot and cold aisle containment systems encompass and seal off the racks in the same row of a datacenter, but providing adequate cooling performance when the loading of one or more racks is much lower than the others remains an issue.

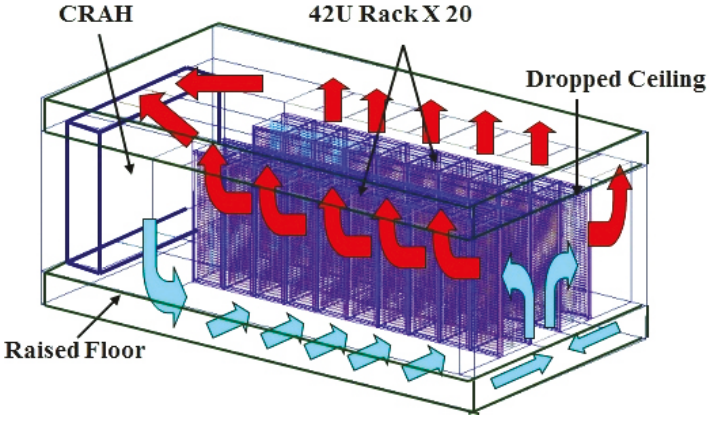

Figure 1. Datacenter configuration and airflow direction

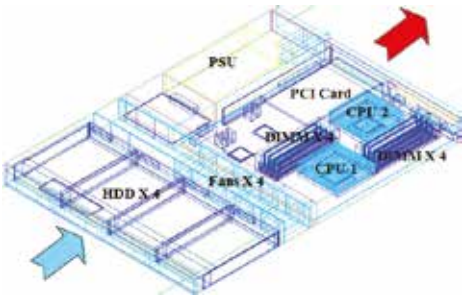

Figure 2. Configuration of the 1U server CFD model and airflow direction

partition shows a significant power saving for IT equipment from the server level up to the datacenter level.

As Simple as One, Two, Three with FloTHERM®

The CFD analysis for the building simulation model is performed with FloTHERM 3D CFD software using a structured Cartesian grid that can be localized and nested to minimize solve times and enable multi-scale modeling for accurate results.

Server-Level Power Savings from a Divided Zone Partition

Figure 2 shows a standard 1U height server CFD model with four fans. Three fans are directed at the CPUs and DIMMs, and one fan towards the PCI card. The PSU includes its own fan, which is located at the rear. We find the thermal solution first to tune the suitable fan curve performance to pass the component thermal specifications at 35°C system ambient. While all the equipment temperatures meet the system thermal requirements, it cannot be assumed that this is the optimum design. Figure 4 shows the airflow distribution. Some airflow does not follow the desired path, providing an opportunity for power savings.

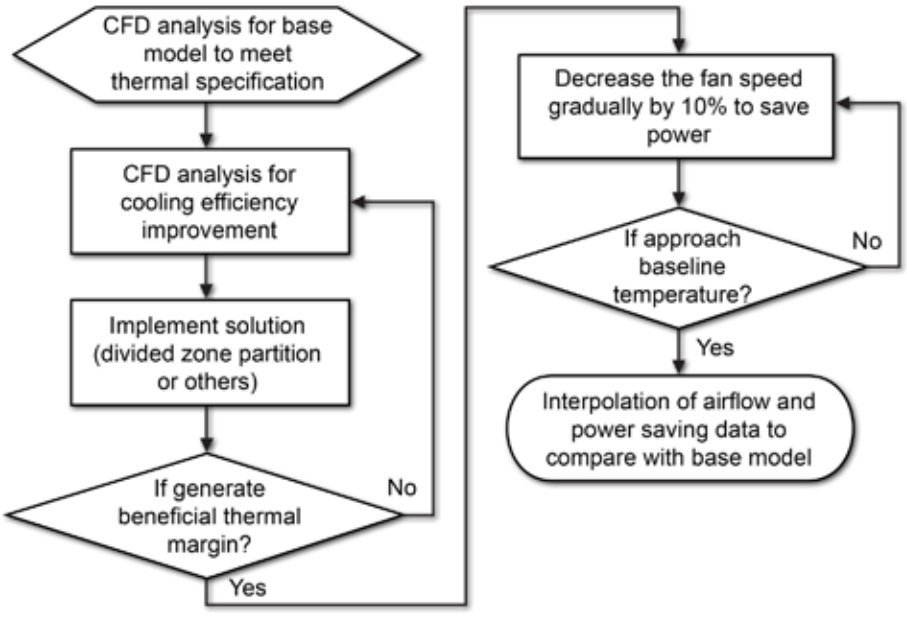

Figure 3. Flow chart of a divided zone analysis for power savings

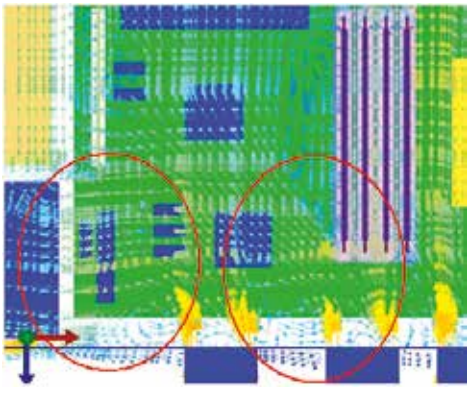

Figure 4. Server system airflow bypass illustration

Servers of the type shown in Figure 2 consisting of Hard Disk Drives (HDDs), Central Processing Units (CPUs), and Dual In-Line Memory Modules (DIMMs), etc., are the main type of IT equipment and can be the target for major power savings in a datacenter, where research [6, 7, 8] into server liquid cooling have shown improvements in cooling efficiency.

In this study, the divided zone method is developed to improve the cooling efficiency for both a server and a datacenter. The effect of a divided zone on airflow management and fan power savings under normal conditions and during a component load change were investigated in detail. Additionally, the utilization of a divided zone

To ensure the design is workable, a fully loaded system is considered at 35°C ambient according to the ASHRAE maximum allowable temperature of the A2 class [9]. First, the thermal solution for the fan and heatsink to meet the fully loaded system requirements that can satisfy all component and device thermal specifications is found for the base model. Second, design optimization is performed to determine whether the divided zone method would improve the cooling. Finally, the solution is analysed to calculate the resulting power savings. The flow chart is shown in Figure 3.

The divided zone partitions were implemented to determine their optimum position for airflow management (Figure 5 overleaf). The resulting effect was that components that generate higher temperatures, such as the CPUs, receive better cooling. The resulting CPU temperature margin allows fan speeds to be reduced, achieving considerable power savings.

The divided zone partition not only saves server fan power directly, but also decreases the system airflow rate requirement, as shown in Table 1. The airflow rate savings can reduce the CRAH blower load. For this case, the divided zone partition can help decrease the server system airflow rate from 59.5 CFM to 51.3 CFM, a 13.8% reduction in the system airflow rate that can further improve datacenter cooling efficiency.

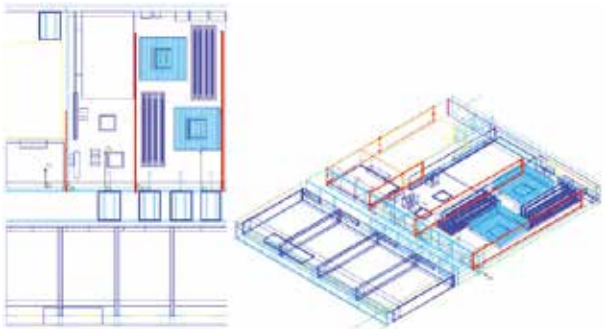

Figure 5. Divided zone partition in the server system

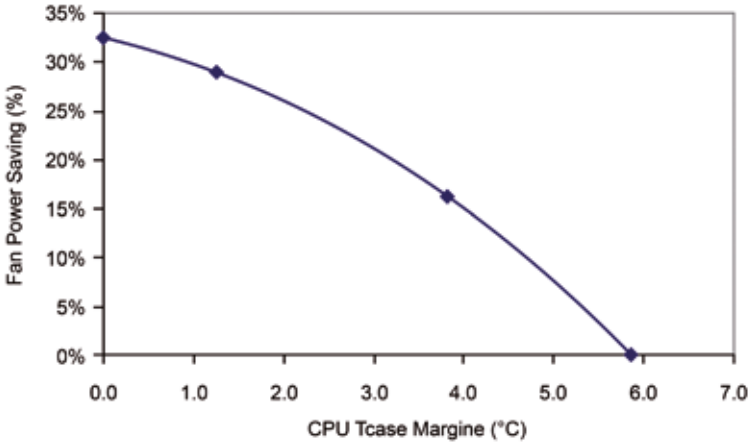

Figure 6. Fan power savings rate vs. the CPU temperature margin of the divided zone partition system

Case	Divided Zone Partition	Fan Speed	CPU2 Tcase (°C)	Fan Power (W)	Fan Power Saving (%)	System Airflow Rate (CFM)	Airflow Rate Saving (%)
1	**No**	**100%**	**79.4**	**66.0**	-	**59.5**	-
2	Yes	100%	73.5	66.0	-	59.9	-
3	Yes	90%	75.5	55.2	16.4%	56.1	5.7%
4	Yes	80%	78.1	46.8	29.1%	52.5	11.7%
5	Yes	70%	82.0	39.9	39.5%	48.7	18.1%
6	**Yes**	**76.7%**	**79.4**	**44.5**	**32.6%**	**51.3**	**13.8%**

Table 1. Simulation result data of the divided zone partition in the server compared with the original model

Server-Level Power Savings with Regional Load Change Condition

The components in a server system are not always at full load, and the load will not be constant, so there will typically be a thermal sensor in the system to detect component temperature variations caused by ambient temperature or load changes. The fan speed can be modulated according to the thermal sensor data by a controller chip in the server. Figure 7 shows the thermal sensor locations. In addition to the ambient sensor placed in the front of the server, there are component sensors placed in CPUs, DIMMs, etc.

The CPUs and DIMMs are controlled by the three fans shown in the right region, the PCI card and chips are controlled by the one fan shown in the middle region, and the PSU is controlled by its own fan in the left region. In this case, the load in the right region was changed to decrease the CPU power from 95 W to 76 W, a reduction of 20%. The CPU was decreased to 71.9°C, providing a temperature margin corresponding to a fan power saving of 31.4%, as shown in Table 2. When the divided zone partition is added, the power savings can be improved to 46.8%. The divided zone partition is not only useful for typical conditions but also for FSC where the local fan is already independently controlled by temperature in a specific area.

Divided Zone Partition Implemented in a Datacenter

The configuration in Figure 1 follows a typical datacenter layout. Earlier studies have shown that air containment separating hot and cold air can be very effective [2, 3]. In this study a simulation model was created for 20 racks with a total of 252kW of power consumption. Small gaps on each side of the rack were included to emulate a true IT equipment environment. Temperatures were set to an ambient of 27°C, the ASHRAE recommended temperature for this class of datacenter [9].

The hot air from the outlet of the racks and

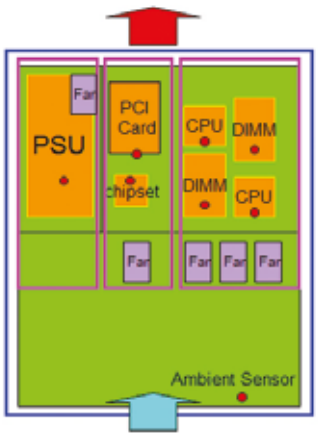

Figure 7. Thermal sensors in the server

the cold air provided to the inlet of the racks are not separated completely, which allows mixing between the hot and cold areas, as shown in Figure 8. To avoid this, some advanced datacenters are implementing air containment, as shown in Figure 9.

At the server level, the component loading is not consistent over a long period of time. At the datacenter level, the rack also experiences different loadings. Although a datacenter with air containment can improve cooling performance, different loadings of the racks in the same row still cause problems. Assume that one of the racks has a zero loading state and that there is no driving fan in the rack. With hot air containment , Figure 10 shows a backflow of hot air from the rear of other racks through the zero-loaded rack to the cold air area. The temperature of the zero-power rack, shown in Figure 11, is seen to be higher than that of other racks.

The partition material could be flexible and transparent plastic sheet for easy access and low cost, with the benefit that every rack can operate independently from the others. From the FloTHERM results, the divided zone has a significant influence on the improvement of the cooling performance of the datacenter under different loading conditions. The study showed that a rack filled with servers with a lower load and correspondingly lower server fan speed resulted in inadequate rack airflow when contained alongside racks containing high-load, high-airflow servers.

The divided zone partition can improve the situation by increasing the airflow rate in the separated region. This can provide better cooling performance for the specific rack and prevent the rack from receiving inadequate airflow, which would lead to an increase in server fan speed and power consumption as the FSC function worked to meet the server's thermal specifications.

Conclusions

A divided zone method has been successfully developed to improve the cooling efficiency for a datacenter. The performance was simulated and investigated under different operating conditions. The major findings are that the divided zone partition can avoid airflow bypass to gain power savings. The partition can save 32.6% of the total fan power consumption and reduce the server airflow rate by 13.8%, reducing the CRAH blower load. For a specific load change case in the server, the FSC function can save 31.4% of the fan power consumption when the CPU load decreases from 95 W to 76 W. Power savings can be enhanced from 31.4% to 46.8% by implementing the divided zone partition with the FSC function. For an advanced datacenter design, the air containment system can avoid the mixing of hot air and cold air to improve cooling efficiency. However, different rack operating load in the same containment region remains an issue. For a 30% load rack case, the implementation of a divided zone partition in the air containment system can improve the airflow rate by 39% for the fan operating in a server case. The divided zone partition shows a significant power savings for IT equipment from the server level to the datacenter level, making it a good choice for datacenter refits to reclaim lost capacity due to cooling.

References

1. M. Green, S. Karajgikar, P. Vozza, N. Gmitter, and D. Dyer, "Achieving Energy Efficient Data Centers Using Cooling Path Management Coupled with ASHRAE Standards2, Proc of 28th IEEE SEMI-THERM Symposium, pp. 288-292, 2012.
2. J. Niemann, K. Brown, and V. Avelar," Impact of Hot and Cold Aisle Containment on Data Center Temperature and Efficiency", White Paper 135 from Schneider Electric - Data Center Science Center.
3. J. Niemann, K. Brown, and V. Avelar, "Hot-Aisle vs. Cold-Aisle Containment on Data Centers", White Paper 135 from Schneider Electric - Data Center Science Center.
4. R. Zhou, C. Bash, Z. Wang, A. McReynolds, T. Christian, and T. Cader, "Data center cooling efficiency improvement through localized and optimized cooling resources delivery", ASME 2012 International Mechanical Engineering Congress & Exposition, 2012.
5. R. Zhou, Z. Wang, C. E. Bash, C. Hoover, R. Shih, A. McReynolds, R. Sharma, and N. Kumari, "A holistic and optimal approach for data center cooling management", 2011 American Control Conference , 2011.
6. M. Iyengar et al, "Server Liquid Cooling with Chiller-less Data Center Design to Enable Significant Energy Savings", Proc of 28th IEEE SEMI-THERM Symposium, pp. 212- 223, 2012.
7. P. R. Parida, "Experimental Investigation of Water Cooled Server Microprocessors and Memory Devices in an Energy Efficient Chiller-less Data Center", Proc of 28th IEEE SEMI-THERM Symposium, pp. 224-231, 2012.
8. M.M. Ohadi, S.V. Dessiatoun, K. Choo, and M. Pecht, "A Comparison Analysis of Air, Liquid, and Two-Phase Cooling of Data Center", Proc of 28th IEEE SEMITHERM Symposium, pp. 58-63, 2012.
9. ASHRAE TC 9.9, "2011 Thermal Guidelines for DataProcessing Environments - Expanded Data Center Classes and Usage Guidance", ASHRAE, 2011.

Divided Zone Partition	Fan Speed	T CPU2 Tcase (°C)	Fan Power (W)	System Airflow Rate (CFM)	Fan Power Saving (%)
No	100%	71.9	66.0	59.7	-
No	77.8%	79.4	45.3	51.7	31.4%
Yes	60%	79.4	35.1	51.3	46.8%

Table 2. Power savings from the divided zone partition for CPU load change situation

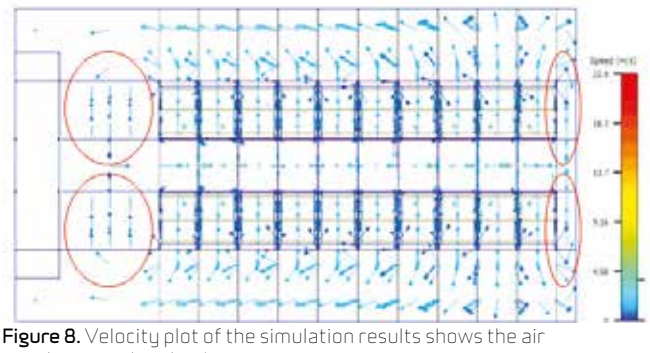

Figure 8. Velocity plot of the simulation results shows the air circulation within the datacenter

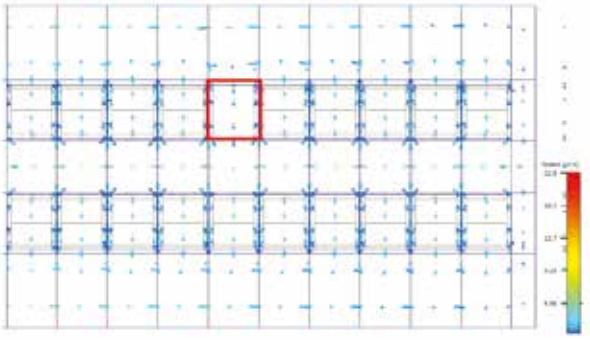

Figure 10. Velocity plot of the simulation result shows that the airflow flows back to the cold air area from the zero loading rack

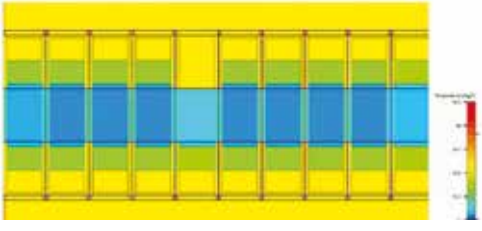

Figure 11. The temperature of the zero power rack is higher than that of other racks

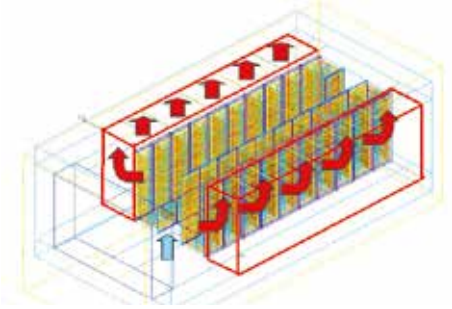

Figure 9. Hot aisle containment within the datacenter

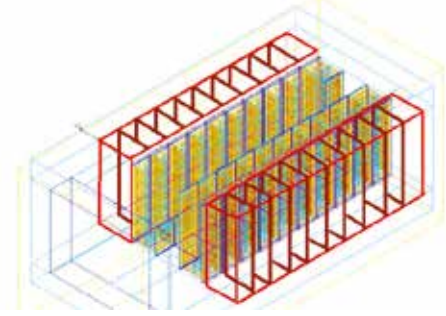

Figure 12. Divided zone partition within the datacenter

Thermal Efficiency: Facebook's Datacenter Server Design

By John Parry, Industry Manager, Mentor Graphics

Large-scale datacenters consume megawatts in power and cost hundreds of millions of dollars to equip. Reducing the energy and cost footprint of servers can therefore have substantial impact.

Web, Grid, and Cloud Servers in particular can be hard to optimize, since they are expected to operate under a wide range of workloads. For its first datacenter in Prineville, Oregon, Facebook set out to significantly improve its power efficiency, cost, reliability, serviceability, and environmental footprint. To this end, many dimensions of the datacenter and servers were redesigned, using a holistic approach. This article is abstracted from the Facebook paper "High-efficiency Server Design," which was presented at the 2011 ACM Conference on Supercomputing, and focuses on this server design, combining aspects of power, motherboard, thermal, and mechanical design. In this article we have looked at the thermal aspects in isolation. In the full paper, Facebook calculated and confirmed experimentally that its custom-designed servers can reduce power consumption across the entire load spectrum while at the same time lower acquisition and maintenance costs. The design does not reduce the servers' performance or portability, which would otherwise limit its applicability. Importantly, the server design has been made available to the open source community via the Open Compute Project, a rapidly growing community of engineers around the world whose mission is to design and enable

the delivery of the most efficient server, storage and datacenter hardware designs for scalable computing.In the past decade, we have witnessed a fundamental change in personal computing. Many of the modern computer uses such as networking and communicating; searching; creating and consuming media; shopping; and gaming—increasingly rely on remote servers for their execution.

The computation and storage burdens of these applications has largely shifted from personal computers to the datacenters of service providers such as Amazon, Facebook, Google, and Microsoft. These providers can thus offer higher-quality and larger-scale services, such as the ability to search virtually the entire internet in a fraction of a second. It also lets providers benefit from the economies of scale and increase the efficiency of their services.

As one of these service providers, Facebook leased datacenters and filled them with commodity servers. This choice made sense at small to medium scale, while the relative energy cost is still small and the relative cost of customization outweighs the potential benefits. As the Facebook site grew to become one of the world's largest, with a corresponding growth in computational requirements, they started exploring alternative, more efficient designs for both servers and datacenters.

Thermal Design

The goal of server thermal design is to cool down the hot components to their operating temperatures with a minimal expenditure of energy and component cost. The typical mechanism used to cool servers at the datacenter level is to cool air at large scale and push it through the servers using their internal fans. The cool air picks up heat from the server components, exits from the server outlet, and is then pushed back to the atmosphere or chilled and recirculated.

More efficient cooling is achieved with air containment in aisles, with the front (or inlet), side of the server facing the "cold aisle" and the back facing the "hot aisle." Yet another technique to improve cooling efficiency is to create an air-pressure differential between the aisles using large datacenter fans. In this case the specific design goal was to be able to cool the upcoming datacenter without chilling the outside air almost year round by allowing effective server cooling even with relatively high inlet air temperature and humidity. To achieve this goal, a more effective design was needed for heat transfer than currently used in the commodity servers.

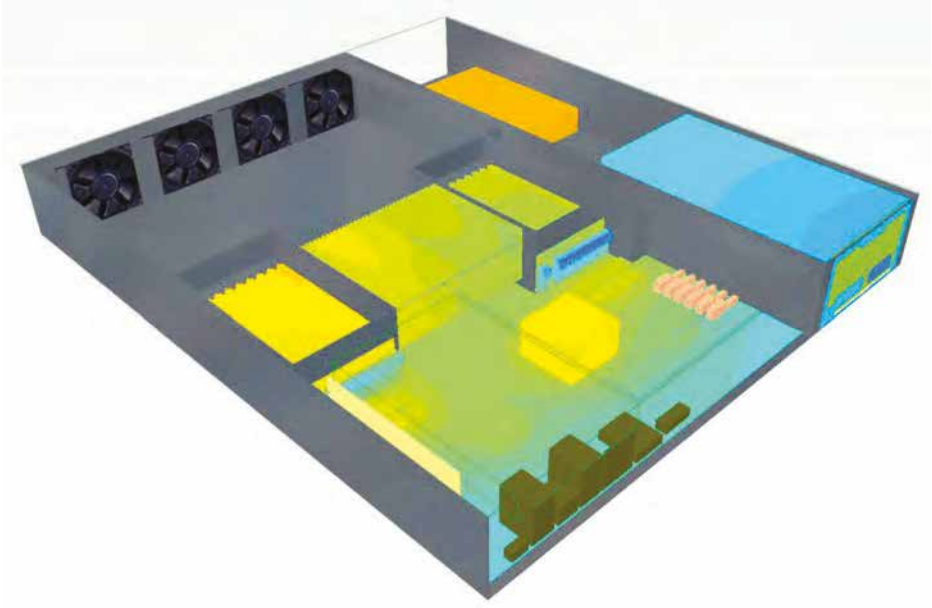

Figure 1. FloTHERM isometric view of thermal design shows chassis, motherboard (with dual processors and memory slots side-by-side), fans, and the hard-disk drive (HDD) behind the PSU. The temperature range here assumes an inlet temperature of 27°C. The air duct on top is elided for visualization purposes.

Improving airflow through the server is a key element here: when internal server components impede airflow, more cooling energy is expended (for example, by faster fans, cooler inlet air, or higher air pressure). One technique by which improved airflow is achieved in the chassis is to widen the motherboard and spread the hot components side by side, not behind each other. The hottest components—processors and memory—were moved to receive the coldest air first, by locating them closer to the air inlet than in the typical back-mounted motherboard.

Another modified dimension was the server height: given a relatively constant rack height (for servicing purposes), a taller server reduces cooling energy but also the rack's computational density. Calculations found that the optimal server height to maximize the compute-capacity per cooling-energy ratio to be the uncommon 1.5U height with large-surface-area heat sinks. This height also allows for an air duct that sits on top of the motherboard and "surgically" directs airflow to the thermal components in parallel heat tracks, reducing leaks and air recirculation inside the chassis. Obstructions to airflow are kept to a minimum, decreasing the number of fans required to push the air out (Figure 1).

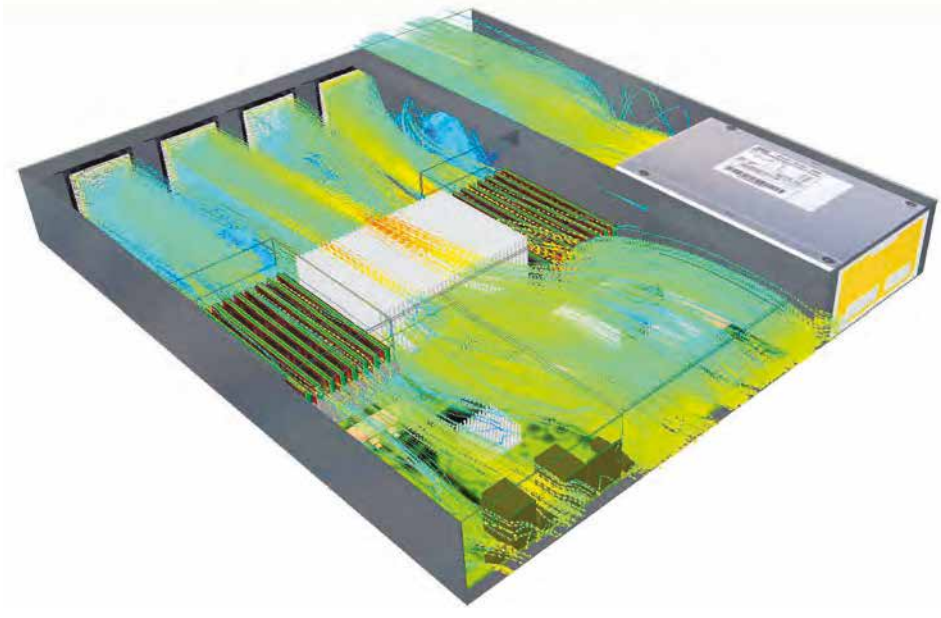

Figure 2. FloTHERM CFD simulation of airflow speed at minimum continuous fan speed

And since the high-efficiency PSU generates less than 20W of waste heat under load, the HDD remains well within specified temperature operating range even behind the PSU. Contrast this with typical server designs that locate the HDD in the front of the chassis to meet its cooling requirements. Also reduced is the amount of airflow required through the system to keep it cool—up to half the volume flowrate compared to standard 1U servers, for the same inlet to-outlet temperature difference (Figure 2).

This low requirement, combined with smart fan-speed controllers, results in fans that spin at their minimum continuous speed nearly year-round, depending on ambient temperature and workload.

An additional advantage of this low speed, continuous operation is a longer expected fan lifetime compared to the typical fan's start-stop cycles, leading to overall improved server reliability. It also naturally translates to lower power and operating costs for server cooling—approximately 1% of the total server power—compared to the more typical 10% in commodity servers. Somewhat surprisingly, even the CAPEX of the server's cooling components alone is about 40~60% lower than a typical server, depending on OEM component pricing. The two main reasons for this improvement are the use of thinner fans (owing to the reduced airflow) and simpler heatsinks without a heat pipe (owing to the larger surface area). Closing the cycle, these efficiency gains carry forward to the datacenter level as well. The server is capable of working reliably at air inlet temperatures of 35°C and a relative humidity of 90%, exceeding the most liberal ASHRAE recommendations for datacenter equipment. In practice, this allows Facebook's datacenter to be cooled almost exclusively on free (outside) air, relying on infrequent evaporative cooling instead of chillers only on particularly hot days.

Methodology

Facebook have evaluated the power, thermal and performance properties of a prototype of the new design against two commodity servers. Both commodity servers are a common off-the-shelf product from two major OEMs, with dual Xeon X5650 processors, 12GB DDR3 ECC memory, on-board Gigabit Ethernet, and a single 250G SATA HDD in a 1U standard configuration. The first server, "Commodity A," is widely deployed in the leased datacenters for Facebook's main Web application. The second server, "Commodity B," is a three-year-old model that was updated to accept the latest generation processors. To ensure a fair comparison, the exact same CPUs, DIMMs, and HDD unit are used in turn, moving them from server to server. The only differing components between the three servers were therefore the chassis, motherboard, fans, power supply, and power source (208V ac/277V ac).

Thermal Efficiency

Thermal efficiency is another important element of the total cost of ownership (TCO), both in terms of cooling energy in the server (fan energy) and in the datacenter. The thermal design is based on a spread and unpopulated board placed in a 1.5U pitch open chassis, and employs four high-efficiency custom 60 × 25mm axial fans. In contrast, the commodity servers use a thermally shadowed, densely populated 1U chassis with six off-the-shelf 40×25mm fans. To evaluate the thermal efficiency, each server was placed in a specially-built airflow chamber that can isolate and measure the airflow through the server, expressed in cubic-feet per-minute (CFM). The measured CFM value was also confirmed analytically by measuring the server's AC power and air temperature difference between inlet and outlet. The servers are loaded with an artificial load resembling Facebook's production power load (around 200W, with leakage power at less than 10W), while maintaining the constraint that all components remain within their operating thermal specifications. The results for the prototype (Figure 3) show a significant improvement. For a typical 7.5MW datacenter, this reduced airflow translates to a reduction of approximately 8~12% of the cooling OPEX. More importantly, it enables free air cooling to be used for the datacenter.

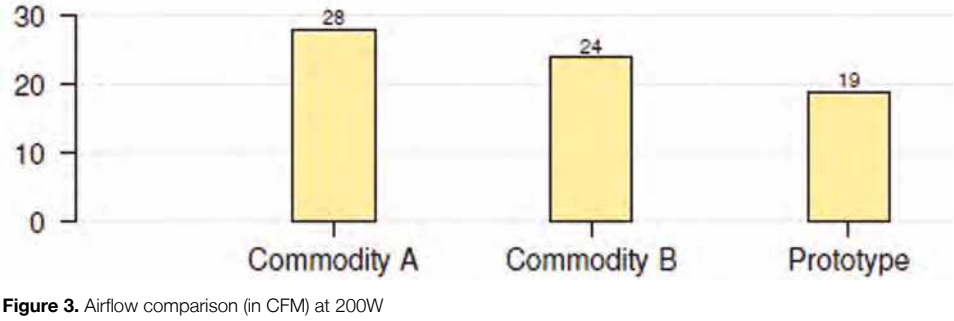

Figure 3. Airflow comparison (in CFM) at 200W

Conclusions

This new server design measurably reduces TCO without reducing performance. The customized server design can:

1. Reduce operating and cooling power (e.g. efficient power conversions, higher-quality power characteristics, fewer components, thinner and slower fans, improved airflow).
2. Lower the acquisition cost and server weight (e.g. fewer and simpler components, lower density, fewer expansion options).
3. Cut costs on supporting infrastructure (e.g. no centralized UPS, no PDUs, no chillers).
4. Increase overall reliability (e.g. fewer and simpler components, distributed and redundant batteries, smooth normal / backup transitions, staggered HDD startup, slower fans).
5. Improve serviceability (e.g. all-front service access, simpler cable management, no extraneous plastics or covers).

At large scale, this design translates to substantial savings. Facebook calculate that over a three year period, these servers alone will deliver at least 19% more throughput, cost approximately 10% less, and use several tons less raw materials to build than a comparable datacenter of the same power budget, populated with commodity servers. When matched with a corresponding datacenter design (including all aspects of cooling, power distribution, backup power, and rack design), the power savings grow to 38% and the cost savings to 24%, with a corresponding power usage effectiveness (PUE) of ≈ 1.07.

Reference:

[1]. Eitan Frachtenberg, Ali Heydari, Harry Li, Amir Michael, Jacob Na, Avery Nisbet, Pierluigi Sarti, Facebook. High-Efficiency Server Design, 2011

Getting Heat Out

Cooling the Next Generation of Embedded Military Computing Products

By Darryl McKenney, VP, Engineering Services, Mercury Systems, Inc,

Mercury Systems is a publicly listed company based in Chelmsford, MA, USA and is a leading supplier of commercially developed, open sensor and Big Data processing systems for critical commercial, defense and intelligence applications. We design and build end-to-end, open-sensor processing subsystems. Our product set spans the entire ISR (Intelligence, Surveillance, and Reconnaissance) sensor processing chain, from acquisition to dissemination, helping customers address a broad range of sensor processing. Mercury Systems have worked on over 300 programs, including Aegis, Patriot, SEWIP, Gorgon Stare and Predator/Reaper.

If we examine typical military electronics CPU Modules and Mezzanines over the last few decades (Figure 1), what is very clear is that their Power levels have increased dramatically. The VPX, formerly known as VITA 46, is an ANSI standard (ANSI/VITA 46.0-2007) that provides VMEbus-based (Versa Module Europa bus) systems for CPUs with support for switched fabrics over a new high speed connector. It was defined by the VITA (VME International Trade Association) working group, that includes Mercury Systems, and it has been designed specifically with defense applications in mind, with an enhanced module standard that enables applications and platforms with superior performance. Basically, all CPU boxes in ISR applications must comply with this standard and its successor VITA 48.

We are finding that devices such as microprocessors and FPGAs (field-programmable gate arrays) have been running ever faster while their size has been constantly shrinking, which obviously has increased heat densities and threatened product reliability. But after nearly a decade of honing our Design for Reliability (DfR) thrust we have produced new design processes and implemented new procedures such that we have reduced the number of engineered change orders by over an order of magnitude. This demand for higher and higher functionalities in defense electronics has led to conflicting demands for more heat management, more sensitive signals, shorter design cycles, higher test coverage and all within ever tighter defense budgets. To add to all this, our products have to be highly reliable with years of operational run time in a wide range of harsh environments. You can imagine the challenges this poses for test engineers, signal-integrity engineers and mechanical engineers when it comes to designing new PCBs and enclosures. Many of today's high powered modules cannot be cooled using legacy cooling approaches. The bottomline is that in our business heat is the primary enemy of module reliability.

At Mercury we offer three types of products to our customers, Air-cooled (A/C) Modules, Conduction-cooled (C/C) Modules and what we call Air Flow-By (AFB) Modules.

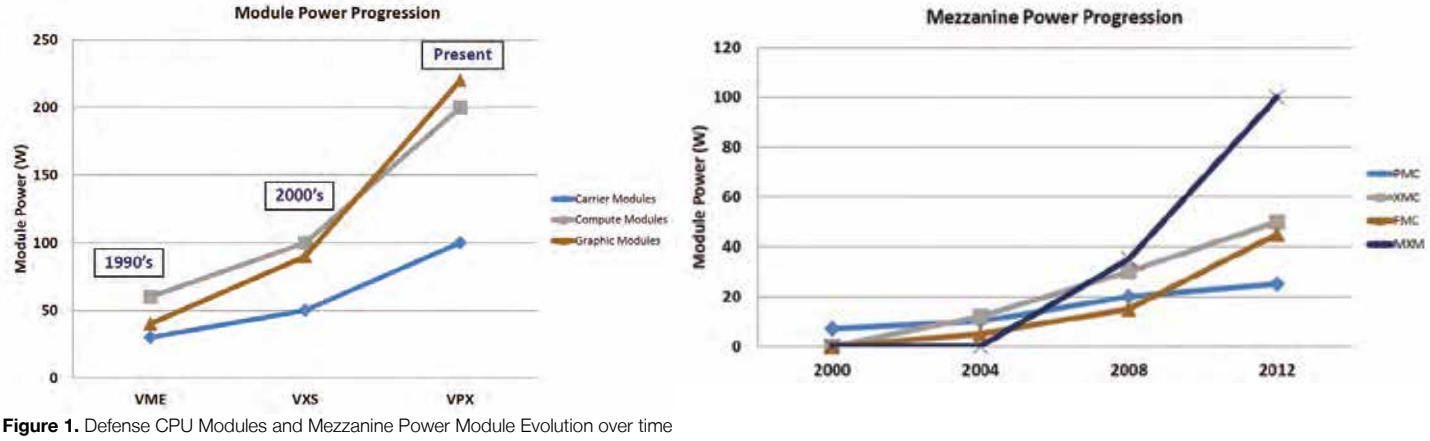

Figure 1. Defense CPU Modules and Mezzanine Power Module Evolution over time

In all cases we go through detailed design and testing processes to design the units for our customers. Our evaluation of each technique's cooling efficiency is highlighted in Figure 2. For our thermal CFD simulation needs we use Mentor Graphics' FloTHERM product which helps expedite our design process.

Air-cooling provides easy access to module debug connectors, front panel I/O and mezzanine modules. This combination simplifies system development and configurability while the system is in its greatest state of flux and requirements are not all identified. A major drawback is that air-cooled modules are not typically designed to be deployed in rugged environments. Conduction-cooling has been the preferred method of cooling for deployed systems for many years.

The modules are designed to handle the rugged shock and vibration levels, while the systems seal the modules away from harmful elements. A major challenge with conduction cooling is that it is heavier than air-cooled and thermally challenged with higher power modules. Air Flow-By – a new cooling technology designed by Mercury – delivers the best of both worlds. It provides the efficient point source cooling of an air-cooled module with the rugged deployment capabilities of conduction-cooling.

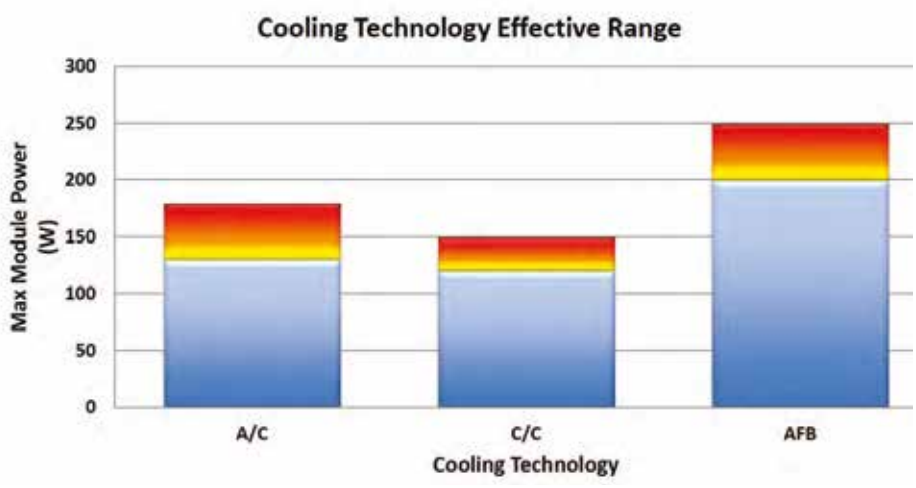

Figure 2. Air-Cooled, Conduction-Cooled and Air Flow-By Cooling Technology Ranges

To give a simple example of how we apply FloTHERM to one of our XMC-Air-Cooled products (Figure 3), we employed a standards based approach to bring heat from the mezzanine modules to the carrier module's heatsink. We discovered via CFD simulation (Figure 4) that we could do this by adding "hooks" for a thermal bridge between the carrier module heatsink and the mezzanine module heatsink. The net effect was a thermal solution that was compliant to standards and allowed for a wide range of mezzanine modules to be placed on a host while limiting any potential changes to a single component. We discovered with FloTHERM that we could get a 5°C Processor thermal reduction - half an Order of Magnitude. This leads to significant impact on mean time between failures (MTBF) too.

In summary, our new thermal-management solutions are capable of dissipating tremendous amounts of thermal energy, while still meeting the same or smaller size, weight and power requirements for the overall solution. By understanding the thermal profile for each specific component that makes up a system using FloTHERM, we created innovations in the mass transfer of thermal energy that work at the individual component, module and subsystem level.

References:

[1]. "We Need More Power Scotty! Getting the Heat Out: Innovations for Cooling the Next Generation of Embedded Computing Electronics" by Dan Coolidge & Darryl McKenney, Embedded Technology Trends, Long Beach, CA, January 21-23, 2013
[2]. December 2012, "Mercury Computer Systems Announces Cold Plate Technology for Embedded and OpenVPX Technologies": http://www.coolingzone.com/index.php?read=162&onmag=true&type=press#sthash.wssfl7Hp.dpuf
[3]. Test & Measurement World, Nov 2008, "Quality by Negotiation", pp32-35

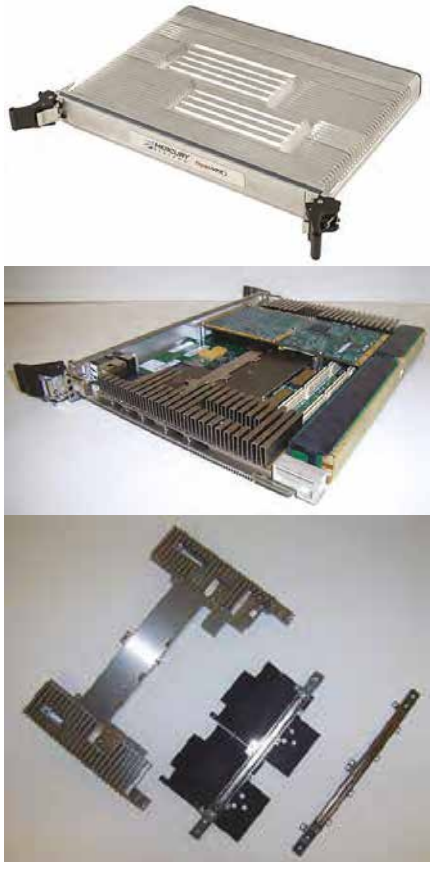

Figure 3. Typical Mercury Systems Integrated XMC Air-Cooled Thermal Solution showing details of Thermal Bridge Hooks

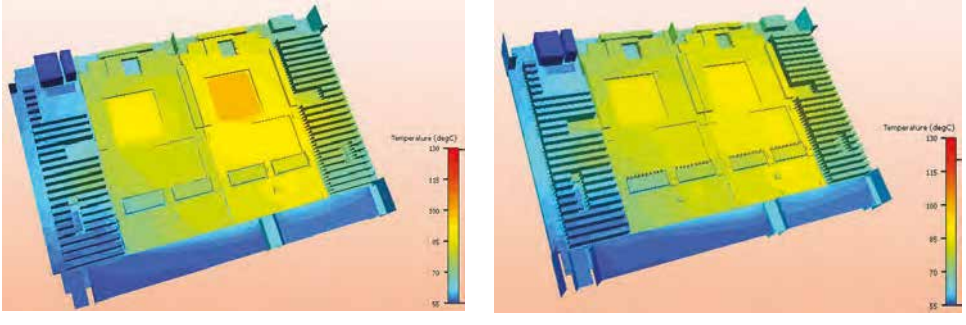

Figure 4. FloTHERM thermal analysis (L) without Integrated thermal bridge (R) with Integrated thermal bridge

Seiko Epson Corporation - Empowering Engineers since 1989

By Mr. Hiroshi Abe, Mr. Naoki Ishibashi, Mr. Shigeki Kikuchi, & Mr. Fumio Yuzawa, Seiko Epson Corporation;
Ms. Hideko Murano, Syuzaiya;
Ms. Hiromi Sugihara, Kozo Keikaku Engineering Inc.

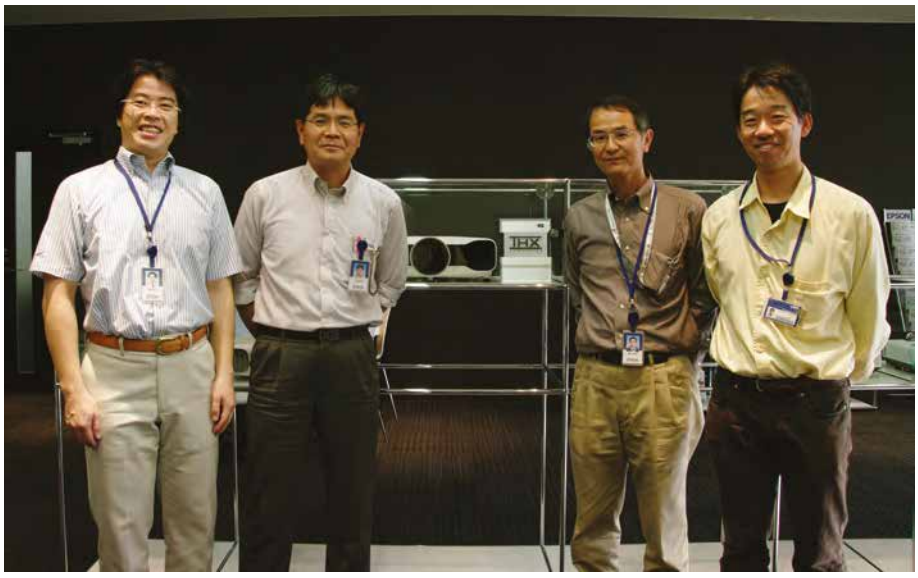

Figure 1. Seiko Epson Corporation in Toyoshina,, VI Planning & Design Department

The Seiko Epson Corporation in Toyoshina, Japan, is home to the VI Planning & Design Department. In March 1989, the Visual Instruments Operations Division was established, with liquid crystal televisions and liquid crystal projectors as its core products.

Epson used its liquid crystal projectors to develop an entirely new market that the company continues to lead: data projectors as multimedia presentation tools. It is here that they developed the technology for the world's first compact, full-color liquid crystal video projector. The first Epson branded projector, the VPJ-700, was built here with revolutionary technology that allowed the projector to use liquid crystal panels instead of a traditional 3-gun CRT to present a picture, thus showing the world a brand new application for liquid crystal displays. With the release of the VPJ-700, Epson became the company to watch in the visual instruments field, where it combined liquid crystal panel technology with optics to develop new products.

Epson's 17 year period with top market share for the Japanese projector market has been achieved by empowering engineers and having an efficient design lifecycle. Today, projectors are used for business, education, and for home entertainment. With each application there is a need for these products to be compact, light-weight, portable and most importantly durable. The challenge with all these attributes is of course thermal.

Heat sources in projectors, like power supplies and lamps, result in high temperatures inside the projector housing. As smaller, more compact portable projectors are continuously being developed the first consideration for the designers at Epson is always making heat sources smaller. By the nature of the design and materials used, the units tend to retain heat in their body unless it is vented into the air or through other parts. Radiant heat transfer is the most important consideration in the development of projectors.

The Thermal Challenge

Previously projector models did not require intense brightness and therefore didn't radiate high temperatures. With only a few models available on the market at the time, Epson could afford a lengthy development cycle. If any problems arose during testing they were able to simply redesign the units. When LCD technology was introduced, the development period halved as there was now a demand for projectors that were brighter, had more functionality, and were smaller.

As with many miniaturized devices, the path for heat radiation is limited, nevertheless air cooling is required within the housing. As well as this, development timescales are sometimes underestimated and can prove costly if they overrun.

Empowering Analysis Engineers

Thermal analysis simulation started at Epson in the 1990's by a team of in-house analysts. This team analyzed all of the products in SEIKO EPSON, including projectors and other electronic devices. Despite the new challenges faced in thermal analysis there was still a requirement to reduce development times and costs. The solutions tended to be focused on each type of device Epson manufactured in order to solve the more complex problems for physics analysis. Hence a new team of dedicated analysts was established to analyze projectors in 2007. It soon became apparent that with the need to speed up development time it was difficult to complete development within the developing phase. Meanwhile designers found it frustrating that they had to wait for the result of analysis by analysts, thereby squeezing the time they had in the design cycle to design and change geometries as necessary. Consequently, SEIKO EPSON started to use FloEFD for Creo in 2009 in order to empower designers to analyze their own designs and to speed up productivity.

Why FloEFD for Creo was the right choice for Epson

According to Mr Hiroshi Abe, "The most important consideration in selecting an analysis software tool was that all team members could use it regardless of their level of ability. We

	**VPJ-700** (1989) World-first 3CCD LCD projector Brightness: 100lm Weight: 7.6kg Pixel number: (320×320) Price: US$ 4,670
	**ELP-3000** (1994) World-first data projector Brightness: 250lm Weight: 7.7kg Pixel number: VGA (640 ×480) Price: US$ 8,650
	**ELP-730** (2002) Brightness: 2000lm Weight: 1.9kg Pixel number: XGA (1024×768) Price: US$ 3,400
	**EB-1775W** (2010) Brightness: 3000lm Weight: 1.7kg Pixel number: WXGA (1280×768) Price: US$ 970 ～ 1,750

Figure 2. LCD Projector history.

evaluated the following three criteria:
1.The people who don't have much experience of analysis can use it easily. In particular, meshing, as this is one of the most difficult processes. FloEFD's automatic meshing enables you to just set a specific area of a model. As for workflow, we only needed to select "yes" or "no" by using the wizard and then we can also learn what we should set in the analysis process by habit.
2. It was important that the tool integrated with Pro/ENGINEER. We didn't want to have to create another model for analysis and being CAD-embedded we could validate various analysis models repeatedly. We also wouldn't have any difficulty in switching between processes (from design to analysis).
3. A comprehensive database. FloEFD has a world-standard database. Especially, we are able to use other databases in the Mentor Graphics suite of products, such as FloTHERM. It has real benefits for users."

Improving Electronics Cooling Simulation

Mr Naoki Ishibashi at Epson notes that often the challenge when a company purchases a new tool is the adoption of its use within established teams with established processes. When Seiko Epson employed FloEFD the take up was swift. The product's intuitive environment was a contributing factor that quickly saw the number of license requests spread. "Actually, I didn't have confidence in the growth of users in the early days. After we tried one license as a test, there were many people who wanted to use it because a good reputation from other users had already spread. Now, we have six licenses. Sometimes I tried to say to some people, 'this software is really comfortable. You can use it anytime. Give it a try', and then the number of users increased. I didn't force them to use it at all."

By adopting FloEFD, Seiko Epson designers were able to affect designs as they developed, with the confidence that any discrepancies in results would be picked up by the analysis team during testing. FloEFD users found new ways to solve difficulties with their newly acquired skill set in analysis. Something the analysts weren't able to deliver as they serviced all the products in the company. Simulation is an essential part of the product development cycle at Seiko Epson so the usability by engineers of any tool is crucial.

"Computational Fluid Dynamics (CFD) is difficult for me even though I've been experienced in analysis for 20 years! However, the first time I tried to use FloEFD, I was amazed by its simplicity," said Mr. Fumio Yuzawa.

"Typical so-called 'high end CFD' software uses really complex meshing techniques. FloEFD requires you only select the resolution level between one and eight. However appropriate the mesh resolution levels are, an adequate number of cells is entrusted to the operator so we have accumulated know-how. If we have a lot of time and a high spec machine, is it better to use high level mesh? I think that's not necessarily so."
"One day, I had some difficulties with the analysis of fan air cooling. The difficulty was solved by using FloEFD for Creo. I didn't consider creating a one-to-one relationship between the characteristic of P-Q curve and the cooling system. I finally succeeded after trying to analyze the fan of air cooling." Mr. Fumio Yuzawa.

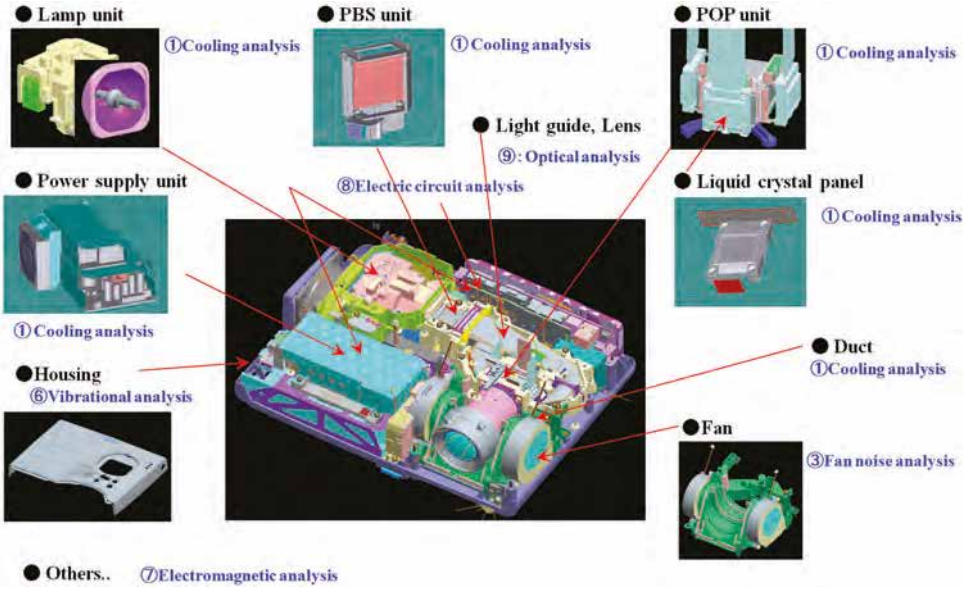

Figure 3. Analysis required for the improvement of design quality

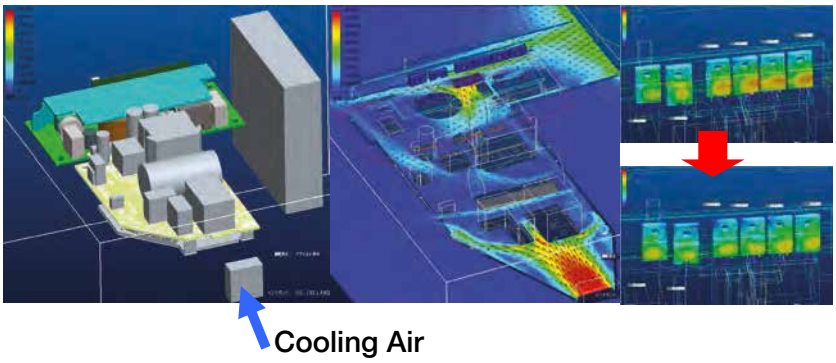

Figure 4. Verification of cooling airflow after changing a geometry path.

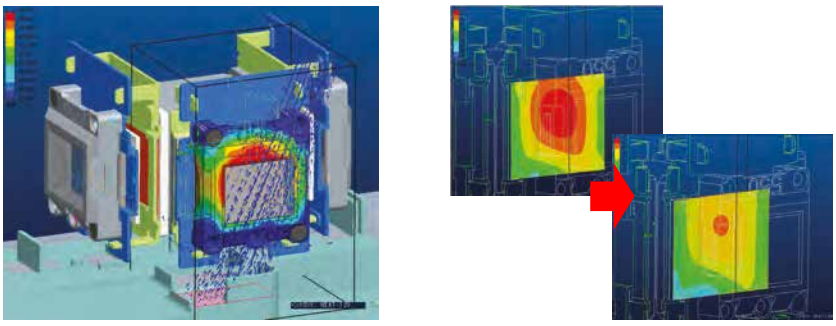

Figure 5. Design study based on a deflecting panel's temperature.

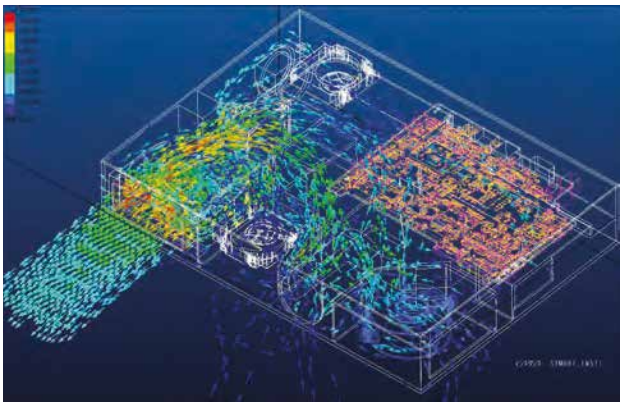

Figure 6. Cooling airflow verification inside of the whole enclosure.

The Future brings its own Challenges

"We want to tackle the problem of projector noise. Our projectors are spread worldwide, not only in Japan. There are some countries with high temperatures, humidity and elevation. We have to design to accommodate different environments as these variables cause the projector to be louder, which is inconvenient for the customer" said Mr. Shigeki Kikuchi. "The difficulty lies in the ability to attach larger fans into modern compact projectors. We use five or six small fans for cooling. A countermeasure for noise is absolutely needed so I joined this team as a sound analyst. We started to measure unpleasant noise over the recent few years and reflect these results in our products."

FloEFD was able to assist in the challenges the team faced with design of semiconductors in projectors. Semiconductors reach high temperatures with natural air cooling and are typically designed under the 60% attainment of industry standards. Designer knowledge and experience will not achieve this level of attainment easily, so FloEFD proved to be the key to solving this challenge. "As well as ease of use and versatility, FloEFD is accurate" said Mr. Fumio Yuzawa, "while some complex parts still require specialist analysis knowledge there are others, such as enclosure air cooling, that can by-pass testing altogether as we are extremely confident in the accuracy of the results we are getting. FloEFD is essential in our daily work. In many cases, we can't predict the results of analysis but FloEFD leads us to correct results automatically."

The world of technology is constantly evolving but Seiko Epson is at the forefront of projector innovation. The company is in the privileged position of being the pioneers as well as the leaders of their market space. Epson's projectors are widely utilized in offices, schools, retailers, museums, movie theaters, and living rooms. Most recently in 2011, Epson developed the Moverio BT-100. This revolutionary product is the world's first* standalone see-through mobile viewer that allows users to enjoy the big-screen experience anywhere and at any time. Moverio signaled Epson's intention to create a new visual communications culture. Going forward, Epson aims to leverage its original projection technology to create more original products that will deliver a big screen experience to all kinds of customers worldwide.

* According to Epson research as of November 25, 2011. The world's first civilian-use see-through mobile viewer that allows users to view video contents without being connected to another device.

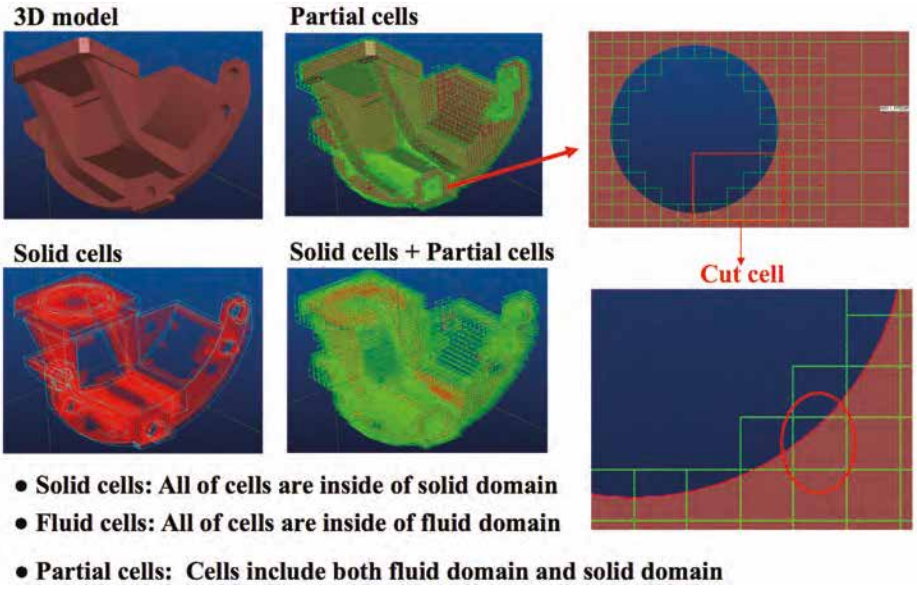

Figure 7. Optimization of mesh calculation.

Comparing Tablet Natural Convection Cooling Efficiency

Consumers are rapidly switching from desktop computers and laptop computers to tablets and smart phones for their computing, gaming and communication needs.

By Guy R. Wagner & William Maltz, Electronic Cooling Solutions

Handheld devices are increasingly capable of running applications that used to require high performance PC systems. With a smaller form factor, this presents significant challenges, especially when one considers that passive cooling is almost an absolute requirement.

With the rapidly increasing performance of tablets and smart phones, the result is increased power consumption leading to devices that are uncomfortably hot to hold. This is especially true after watching videos or playing games on these devices has become very popular. These types of operations are both CPU and graphics intensive which involve much higher power dissipation than viewing a relatively static screen. The thermal design of the next generation handheld tablet device and smart phone will need to address both a comfortable surface touch temperature and maximum temperature limitations of critical internal components while also meeting aggressive industrial design requirements.

The limits of cooling for handheld devices explored in this study are based on both testing and simulation under various conditions, and provides a method for evaluating the quality of the cooling options for these devices. Factors affecting the maximum possible power dissipation are: surface area and emissivity; outer shell materials; thermal interface materials; heat spreader properties; and air gaps. The limiting factor in the thermal design of these devices is generally the touch temperature of external surfaces. There have been studies that address the maximum allowable touch temperature of a handheld device. This article discusses the effect of maximizing the heat spreading within these devices which aids in keeping the touch temperatures of these devices within comfortable limits.

Performance Expectations

Handheld devices are increasingly capable of running applications that used to require their more powerful predecessors. The user expects these devices to provide similar performance to the notebook and desktop computers that were commonly used in the past. This presents significant challenges, especially when one considers that passive cooling is almost a requirement in these small form factor devices. Several studies have focused on the cooling challenges of hand-held devices; Brown et al [2], Lee et al [5], Mongia et al [6], Huh et al [4], and Gurrum et al [3].

Maximum Power Dissipation

The maximum possible power dissipation by natural convection and radiation has been calculated for this study using FloTHERM® CFD simulations and is shown in Figure 1.

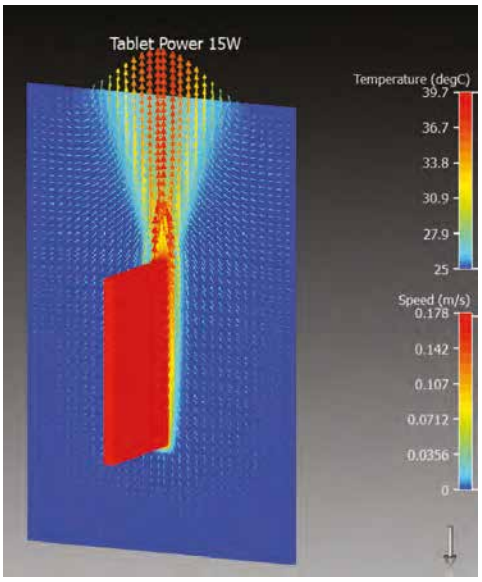

Figure 1. FloTHERM® CFD Model of a 10 Inch Tablet in a Vertical Orientation

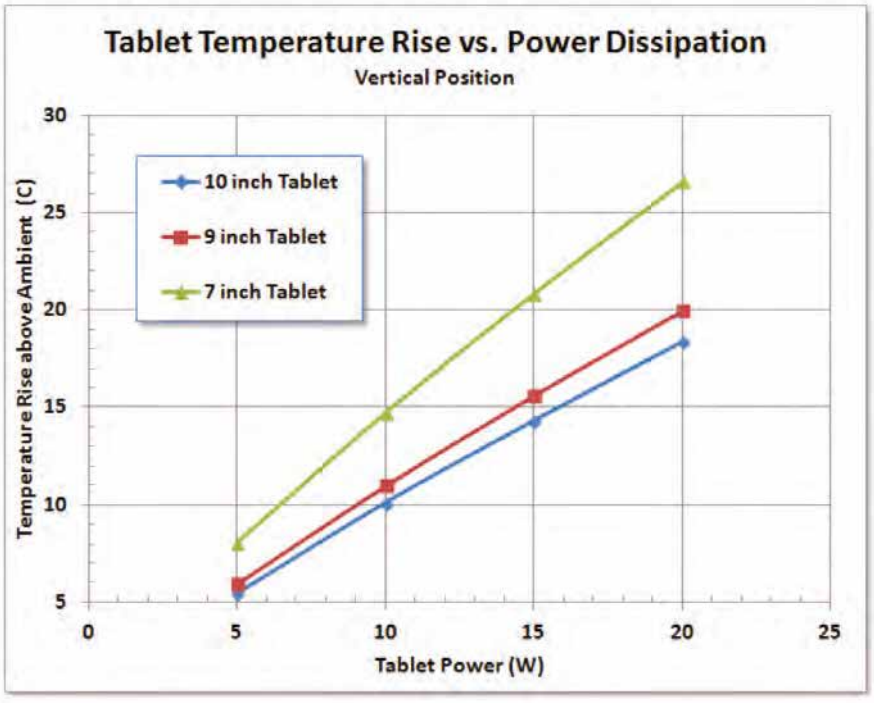

Figure 2. Temperature Rise Above Ambient as a Function of Tablet Power

In a 25°C ambient condition, the maximum total power dissipation is calculated with the requirement that the surface temperature does not exceed a touch temperature of 41°C. This is the maximum aluminum enclosure comfort touch temperature as presented by Berhe [1].

It can be seen that the theoretical maximum total power dissipation is 17.1 watts when a 10 inch tablet is suspended vertically in mid air with conduction and radiation occurring from all surfaces. When the device is placed on a horizontal adiabatic surface, heat transfer occurs from the sides and front surface only. These values establish bounds for the maximum amount of heat that can be dissipated by the device for different orientations in still air.

In order to calculate the total power dissipation for a full-size tablet, the following assumptions were made as inputs to the FloTHERM model with 610k elements: a typical 10 inch tablet size of 180 mm (w) x 240 mm (h) x 10 mm (d) with an ideal condition of uniform surface temperature. The tablet was simulated as a very high thermal conductivity block (k = 10000 W/mK) with uniform internal heat generation to yield an essentially isothermal surface. With an ideal surface emissivity of 1.0, the radiant heat transfer component accounts for a surprising 9.9 watts of the 17.1 watts total power while the convective heat transfer component is 7.2 watts. Since radiation can account for more than half of the total power dissipation at an ambient temperature of 25°C, the use of high emissivity surfaces finishes is very important.

In order to achieve maximum power dissipation, design parameters need to be considered carefully. It is important to design the device to be as isothermal as possible to maximize the amount of heat transfer from all surfaces to the surroundings. If a surface is no longer available for heat transfer, such as when the device is placed on a blanket, the amount of power that can be dissipated is almost cut in half.

The power dissipation at 41°C maximum touch temperature is calculated using CFD for three devices ranging from smart phone size through to mini and full-size tablet, the power dissipation versus surface area is shown in Figure 3, for the device in both the vertical and horizontal positions with an adiabatic lower surface. There is a high probability that the tablet will be used in the horizontal position with a near-adiabatic

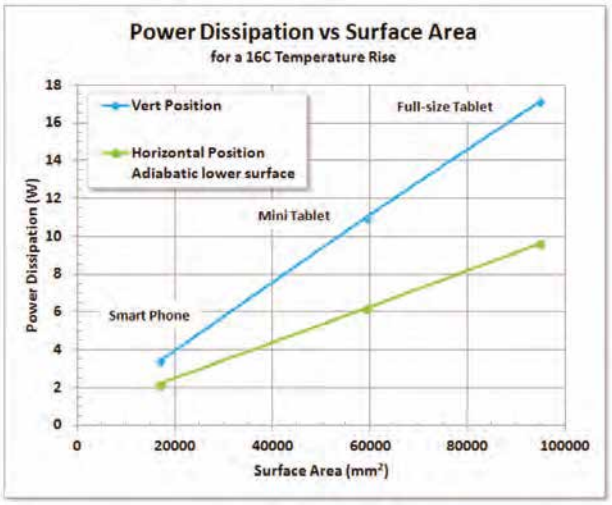

Figure 3. Power Dissipation vs. Surface Area for a Device with a 41°C Skin Temperature

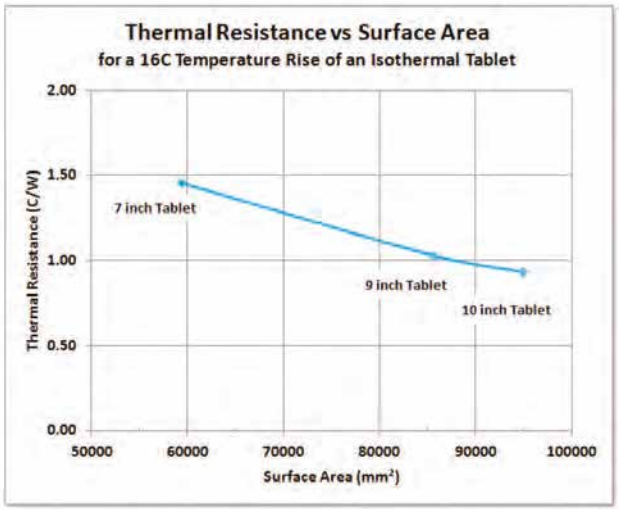

Figure 4. Thermal Resistance vs. Surface Area for a Vertical Isothermal Device with a 41°C Skin Temperature

lower surface while the smart phone will rarely be run in this orientation. From the power and temperature rise computations, it is now possible to calculate the thermal resistance of a vertical, isothermal tablet as a function of the exposed surface area as shown in Figure 4.

Numerical Models

In order to analyze the impact of different thermal management techniques, a detailed Computational Fluid Dynamics (CFD) thermal model was constructed using FloTHERM XT™. Since the thermal characterization data for the main processor in the tablet may not be known, the thermal characteristics of the actual processor can be measured with a high degree of accuracy using Mentor Graphics' T3Ster® hardware to determine the thermal resistance from the processor IC to the case and the PCB; Wagner et al [7]. This allows accurate capture of heat flow from the top and bottom of the processor. The thermal model of the processor can be directly dropped into the tablet CFD model. The adaptive mesh in FloTHERM XT allows the fine features of the internal components of the tablet to be included in the model while keeping the mesh count to a reasonable size as shown in Figures 5 and 6 (overleaf).

With the CFD thermal model, the following questions can be addressed:

1. How much do high-conductivity heat spreaders improve heat dissipation while reducing the touch temperature?
2. What is the best way to move the heat from the heat-producing components to the surfaces of the tablet where it can be safely dissipated?
3. How can air gaps be strategically used in the thermal management process?
4. How important it is to account for radiation in addition to convection?
5. Should the dissipated power be spread at the source or at the surface?

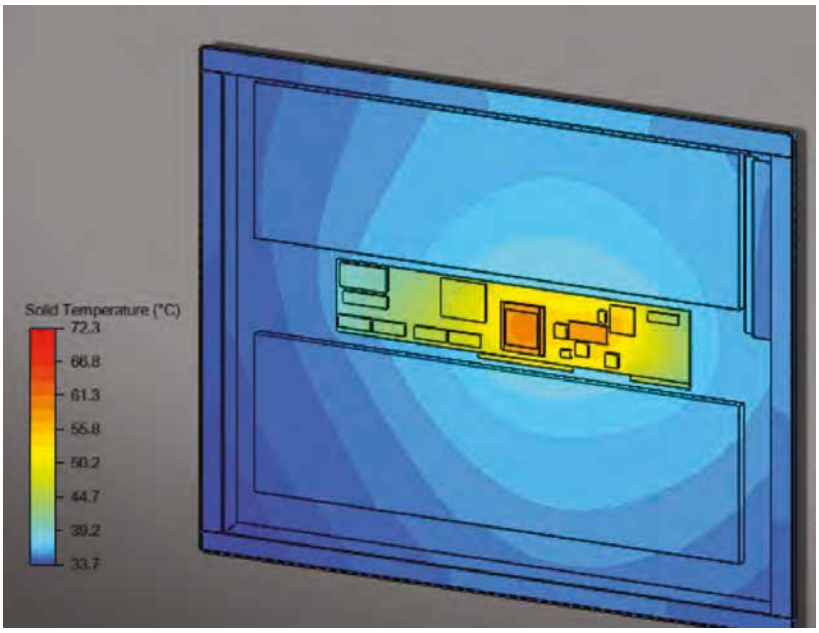

Figure 5. FloTHERM XT Model of the Internal Component Temperatures of a Tablet

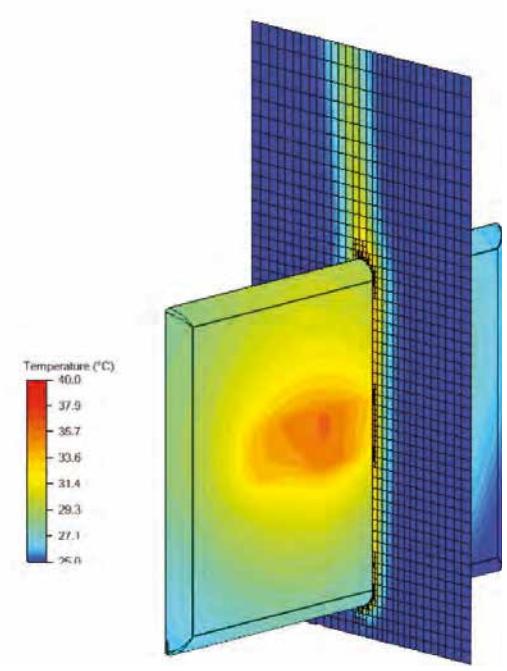

Figure 6. FloTHERM XT Model Showing the Hot Spot and the Natural Convection Airflow around the Back of a Tablet

6. How does designing the tablet with a plastic case compare to an aluminum or magnesium case?

Since we have a goal of keeping touch temperature at or below 41°C, answering question #1 has a large impact on the design.

Temperature uniformity can be achieved by either providing a high conductivity heat spreader inside the case of the tablet or by making the case itself out of a high conductivity material; Wagner et al [8]. One must keep in mind that the maximum touch temperature is a strong function of the conductivity of the heat spreader or case. As the conductivity of the case goes down, the maximum comfortable touch temperature increases. For example, if the case is made of plastic with a thermal conductivity in the range of 0.2 W/mK, the case temperature that the user senses feels lower since the low thermal conductivity of the plastic conducts less heat to the user's skin. The apparent touch temperature decreases as the product of thermal conductivity, density and specific heat (k.ρ.Cp) of the case material decreases. When this product is low, the touch temperature may be increased by approximately 5°C over that of a solid metal case. Since the surface area of the case is large relative to the thickness of the plastic, heat transfer to the air is not reduced significantly over that of an aluminum case. This assumes that the heat is spread across the inside surface of the plastic housing using a high-x-y conductivity aluminum plate or graphite sheet.

Infrared Images

When evaluating the effectiveness of heat spreaders, whether internal or through the use of high conductivity case material, it is very useful to take infrared images of the units under test while they are performing at their maximum computation levels. It has been found that graphic intensive processes also require considerable computation and will exercise the tablet at maximum power. In general, it takes a minimum of 30 minutes for the tablet to reach a steady state temperature. Figure 7 shows three tablets running the same game during thermal testing.

Since the emissivities of the tablet surfaces are not always known, thermocouples are placed at various locations on the front and back of the tablet to read the temperatures at selected locations. The emissivity setting of the image from the IR camera is adjusted until the difference between the thermocouple readings and the IR image is minimized.

Figure 8 consists of infrared images of the back side of four different models of tablet. The tablets were running a game called Riptide GP, while the graphics and computational capabilities of mobile devices were measured. Ambient air temperature at the time of the test was recorded for each tablet to determine the temperature rise of the hot spot. The dark area running up the center of tablet A is from the support that was holding that tablet in the vertical position.

Note how the location of the hot components inside "print through" the case forming a hot spot.

Figure 7. Images of three Different Tablets Running Riptide GP while Temperatures are being Monitored with Thermocouples.

Figure of Merit for the Quality of the Thermal Solution

Since heat spreading is the most important factor for dissipating heat from the outer surface of the tablet and reducing the temperature of the hot spots, the authors propose the following figure of merit to determine the effectiveness of the thermal design of the tablet.

The thermal heat spreading efficiency of the tablet can be defined as the ratio of the ideal thermal resistance of an isothermal tablet divided by the measured or simulated thermal resistance of the actual tablet. The thermal resistance of an isothermal tablet with emissivity equal to 1.0 is the very best that can be achieved on a theoretical basis. The isothermal thermal resistance is calculated by dividing the temperature rise above ambient by the dissipated tablet power for an isothermal tablet. The actual thermal resistance is calculated by dividing the temperature rise of the hot spot above ambient by the dissipated tablet power.

Heat Spreading Efficiency = R_i/R_a

$R_i=\Delta T_i/Q_i$ Thermal resistance of an ideal isothermal tablet

$R_a=\Delta T_a/Q_a$ Thermal resistance of an actual tablet

Where

ΔT_i = Temperature rise above ambient for an ideal isothermal tablet with emissivity = 1.0

Q_i = Power dissipation of the ideal isothermal tablet

ΔT_a = Temperature rise of the hot spot above ambient for the actual tablet

Q_a = Power dissipation of the actual tablet

The table below summarizes the results of the testing for the four tablets and calculates the heat spreading efficiency of the thermal design for each tablet.

Summary

In summary, building an accurate thermal model of the tablet allows the designer to rapidly test the effect of design and material changes without incurring the cost and schedule delays of testing prototypes. This speeds time to market and lowers development costs.

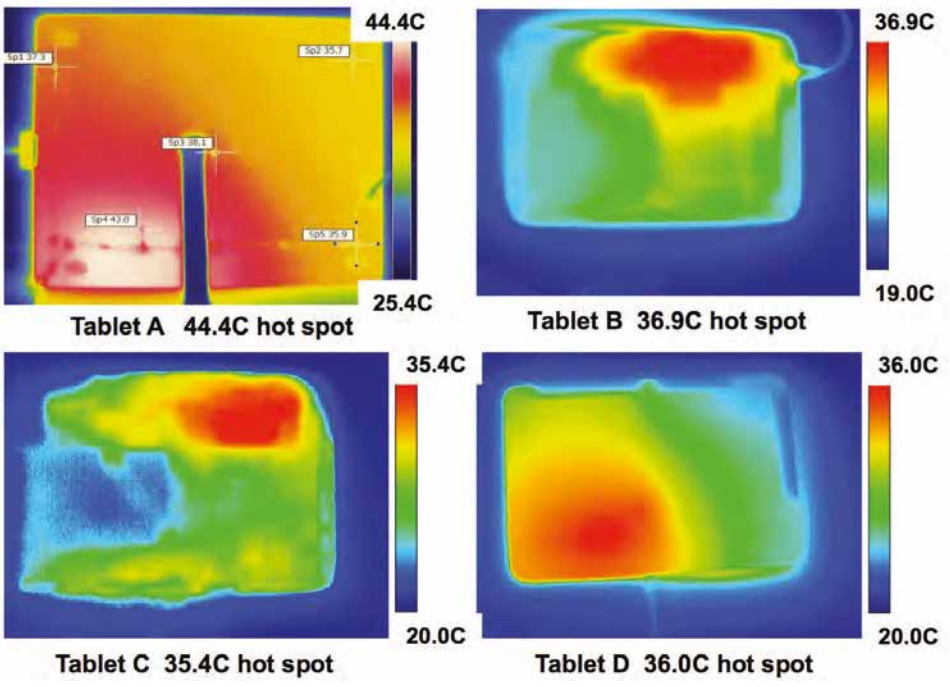

Figure 8. Infrared Images of the Back Side of Four Different Tablets Running Riptide GP

The maximum power dissipation of the internal components is not only governed by the size of the tablet but is a strong function of how well that heat is spread internally to reduce hot-spot temperatures. Few engineers realize the importance radiation plays in dissipating the heat from the exposed surfaces of a tablet. It is not until precise calculations are made that the importance of radiation in the thermal design of tablets really becomes apparent. If the emissivities of the various surfaces are high, over half of the heat transfer is due to radiation.

Introduced is the Heat Spreading Efficiency figure of merit which measures the actual cooling efficiency of a tablet against the theoretical maximum cooling efficiency. The perfect thermal design for a tablet cooled under natural convection would have a Heat Spreading Efficiency of 1.0. However, tablet thickness and weight have to be traded off against efficiency.

Heat Spreading Efficiency

	Tablet			
	A	B	C	D
Ambient Temperature (C)	25.4	19.0	20.0	20.0
Hot Spot Temperature (C)	44.4	36.9	35.4	36.0
Actual Delta T (C)	19.0	17.9	15.4	16.0
Power (W)	8.8	6.9	7.2	5.2
Ra (C/W)	2.16	2.60	2.14	3.05
Delta T iso (C)	16.0	16.0	16.0	16.0
Power-isothermal (W)	17.1	17.1	17.1	15.5
Ri (C/W)	0.94	0.94	0.94	1.03
Heat Spreading Efficiency	**0.43**	**0.36**	**0.44**	**0.34**

References:

[1]. Berhe, M.K., Ergonomic Temperature Limits for Handheld Electronic Devices, Proceedings of ASME InterPACK'07, Paper No. IPACK2007-33873

[2]. Brown, L., Seshadri, H., Cool Hand Linux® - Handheld Thermal Extensions, Proceedings of the Linux Symposium, Vol. 1, pp 75 – 80, 2007

[3]. Gurrum, S.P., Edwards, D.R., Marchand-Golder, T., Akiyama, J., Yokoya, S., Drouard, J.F., Dahan, F., Generic Thermal Analysis for Phone and Tablet Systems, Proceedings of IEEE Electronic Components and Technology Conference, 2012

[4]. Huh, Y., Future Direction of Power Management in Mobile Devices, IEEE Asian Solid-State Circuits Conference, 2011.

[5]. Lee, J., Gerlach, D.W., Joshi, Y.K., Parametric Thermal Modeling of Heat Transfer in Handheld Electronic Devices, Proceedings of the 11th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, I-THERM, pp 604-609, 2008

[6]. Mongia, R., Bhattacharya, A., Pokharna, H., Skin Cooling and Other Challenges in Future Mobile Form Factor Computing Devices, Microelectronics Journal, Vol. 39, pp 992 – 1000, 2008

[7]. Wagner, G.R., Maltz, W., Thermal Management Challenges in the Passive Cooling of Handheld Devices, Proceedings of Therminic 2013, pp 135-140.

[8]. Wagner, G.R., Maltz, W., On the Thermal Management Challenges in Next Generation Handheld Devices, Proceedings of the ASME InterPACK 2013, Paper No. InterPACK2013-73237

Unravelling the Complexities of Automotive Instrument Cluster Design

Visteon Electronics Corporation's use of FloTHERM®XT for Validation and Optimization

By Sam Gustafson, Thermal Analysis Engineer, Visteon Electronics Corporation

Visteon Corporation who recently acquired the automotive electronics business of Johnson Controls Inc. (JCI), has become one of the world's three largest automotive electronics suppliers of instrument clusters and vehicle cockpit electronics. The combined global electronics enterprise has more than $3 billion in annual revenue, with a No. 2 global position in driver information and above-average growth rates for the cockpit electronics segment, supplying nine of the world's 10 largest vehicle manufacturers.

In this article, Sam Gustafson, a Thermal Analyst at Visteon , shares the importance of the thermal requirements and their influence on the PCB and mechanical design of a cluster. In particular, how PCB data from Mentor Graphics' Expedition software can be embedded inside a FloTHERM® XT simulation model to investigate the full thermal behavior of the instrument cluster and optimize it.

Also discussed is a comparison of the PCB layout thermal simulation using Thermography before full enclosure modeling.

In an automotive instrument cluster design, where enclosure shape and internal complexity significantly influences thermal management considerations, engineers focus their attention on areas such as PCB structure/layout and active display dimming to ensure durable performance. Efficient exchange of data between the PCB layout, the Mechanical and Thermal analysis tools therefore becomes key in designing such systems.

As the component most utilized by the driver, the vehicle instrumentation panel greatly impacts the driver experience and therefore consumer satisfaction, Visteon are at the forefront of this technology. The design of instrument panels (Instrumentation Cluster Assemblies) must optimize quality, reduce costs and lead times, and guarantee flawless product launches for their customers.

Thermal integrity becomes a top priority, with heat dissipation the most important consideration for suppliers such as Visteon.

A typical instrumentation cluster (Figure 1 overleaf) consists of two analog gauges on either side of the unit with several LEDs, a significant number of which are bright LEDs. Each system has a Thin Film Transistor (TFT) display (blue section in Figure 1 overleaf), which can only operate 10°C above the maximum ambient air temperature. It is therefore critical to keep this component's temperature tightly controlled. The cluster is then encased into housing to the dashboard that has limited air openings inhibiting ventilation.

When designing an instrument cluster, the most pertinent consideration is finding an effective way to overcome the temperature sensitivities of most of the components in the assembly. For instance, the LED light and color will degrade if the junction temperature becomes too hot for long periods

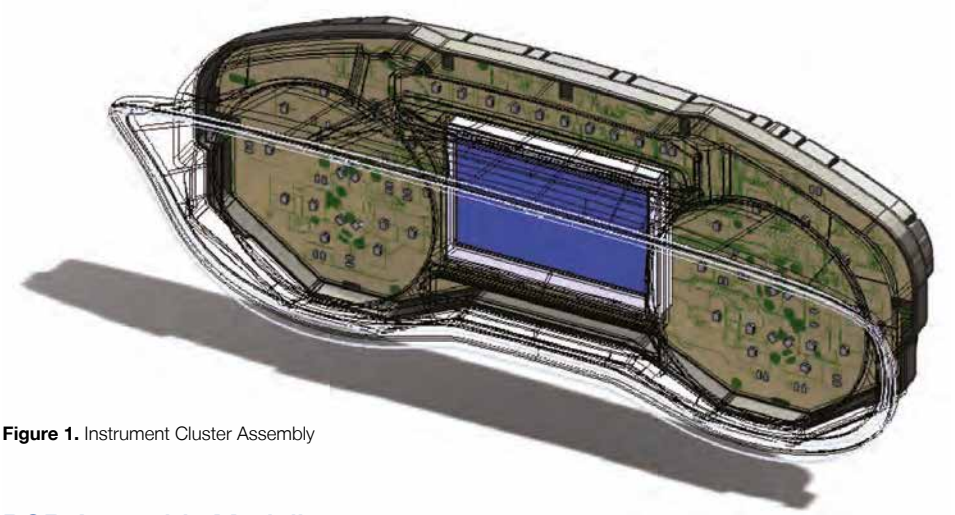

Figure 1. Instrument Cluster Assembly

of time. Another important factor to consider is the higher power dissipation of components such as linear power supplies, reverse protection diodes and microprocessors. With automotive ambient temperature requirements dictating that the instrument cluster remain at an ambient air temperature of 85°C, space is limited for the components to operate below the temperature limit specified by manufacturers.

To address these challenges upfront, various strategies are applied to the design. First, working closely with the Hardware Engineer creating the cluster architecture. The Mechanical Group then ensures the created architecture is thermally acceptable: Does it contain unnecessary high power density areas? Is the generated power being used efficiently throughout the system? The instrument cluster is housed in the dashboard; the primary route for the heat to exit the system is by conduction through the PCB. Working to optimize the heat spread through the PCB will make it an efficient heatsink before other cooling strategies are considered. This is made possible by changing the component layout as well as adjusting the copper content of the PCB. When possible, holes are placed on the back of the housing to vent the cluster. Finally, because of the temperature sensitivity of the TFT display, its brightness may be diminished at high ambient temperatures. This parameter is something that JCI and their customers work closely on together so that the TFT display is expected to operate at 100% power up to a fixed ambient temperature. Once this is reached, the TFT display is then dimmed to preserve its integrity and maximize its performance.

Using FloTHERM XT upfront allows designers to gain a better understanding of these strategies and their most efficient implementation.

PCB Assembly Modeling

With the PCB being used as an optimized heatsink for the component conduction path, it is crucial to model it accurately, especially when considering the copper content of the board. Fortunately, FloTHERM XT seamlessly leverages an existing board layout from Xpedition PCB via a *.cce file. The direct import of this file (Figure 2) with the FloEDA Bridge interface allows for the definition of a thermally comprehensive 3D model of the PCB. This model contains the traces in each layer based on the predefined routing as well as the placement, size and names of each component. Parts can be linked to an existing library of predefined thermal models, ranging from a simple block representation to a more detailed one, along with thermal networks such as a two-resistor or a DELPHI network. The components are then swapped automatically with the most appropriate fit based on the user's definition and match of either the package name or the part number.

Figure 2. PCB Assembly 3D Thermal Representation Front and Rear Views

For this cluster design the component thermal models used were the simple 2R and detailed models. Figure 3 demonstrates a detailed model of the processor sitting on the cluster board. This model was provided directly by the component manufacturer. Once all the components are properly defined, the power dissipation can simultaneously be allocated to each part from the import of a *.csv file containing the power per reference designator. Alternatively, the PCB can be modeled depending on how refined the thermal results need to be. FloTHERM XT can create an equivalent compact model of the board based on the copper content for each of the four layers transferred from Xpedition PCB. An orthotropic in-plane and through plane conductivity is then computed based on this data. This approach works well for the instrument cluster overall but is insufficient for components relying on copper patches to disperse their heat. In this case, the image of the traces can be used to draw

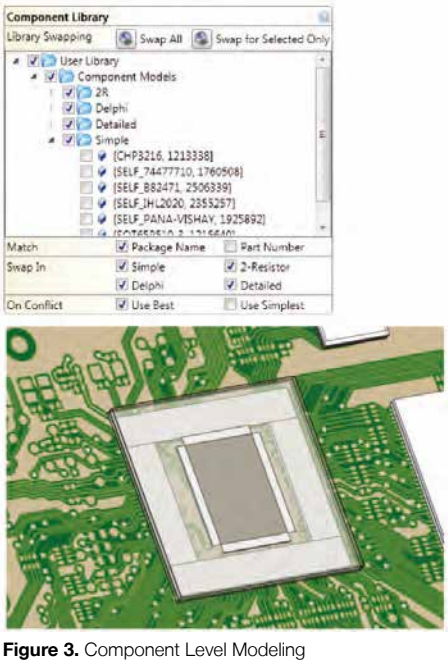

Figure 3. Component Level Modeling

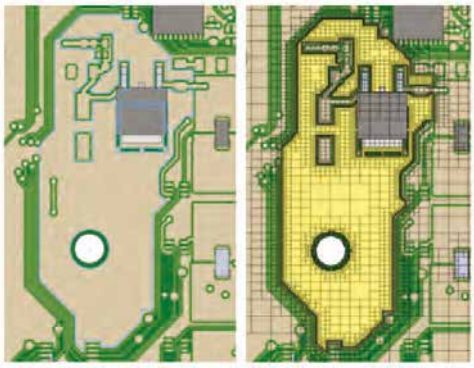

Figure 4. PCB Detailed Thermal Modeling

Figure 5. PCB Experimental and Numerical Set-Up

the outline of the desired copper shape and extrude the corresponding copper patch (Figure 4). Placing the newly created copper pad lower in the FloTHERM XT tree makes it overwrite the PCB area of interest. Thermal resistance values were also defined on the edges of the patch to represent the buffer that the FR4 creates between the detailed shape and the rest of the copper plane. Being able to combine the existing ECAD data with thermal representations of components in FloTHERM XT allows the user to create the thermal analysis model quickly, saving time in the design process.

Comparing Experimental Data and Numerical Results

In order to correlate the numerical results, an experimental set-up of the cluster board without the housing was created. Figure 5 demonstrates how the PCB was mounted standing vertically above a table in a room at 24°C. The PCB was operating with the backlight running at nominal brightness. IR camera images of the cluster were generated during the test.

A first comparison between the experimental data and the numerical results was done in order to validate which PCB modeling technique was the most appropriate for this model. The second technique (Figure 6) gives a component temperature which is within 5% of the measured temperature while the first modeling technique shows a difference of 14%. Therefore, the second modeling technique was chosen as a reference for this particular model.

PCB Model	Total Cell Count	Temperature Rise		Difference	
		Measured	Simulated		
Compact	88,153	24.7 °C	28.2 °C	3.5 °C	14.2%
Detailed Copper Flood	124,525		25.9 °C	1.2 °C	4.9%

Figure 6. Results Comparison for PCB Modeling Approach

Looking deeper, Figure 7, demonstrates that both the measured IR camera images and the simulated images compare very well, with heated areas being highlighted consistently between the two. A comparison component by component shows a maximum discrepancy of 5% with the experimental data, confirming the validity of the thermal analysis model in FloTHERM XT. As a result of this model being validated, the PCB assembly can then be placed in the cluster full assembly to determine its performance in worst case conditions at an air ambient temperature of 85°C.

Figure 7. Results Comparison for PCB Modeling Approach Front and Rear Views

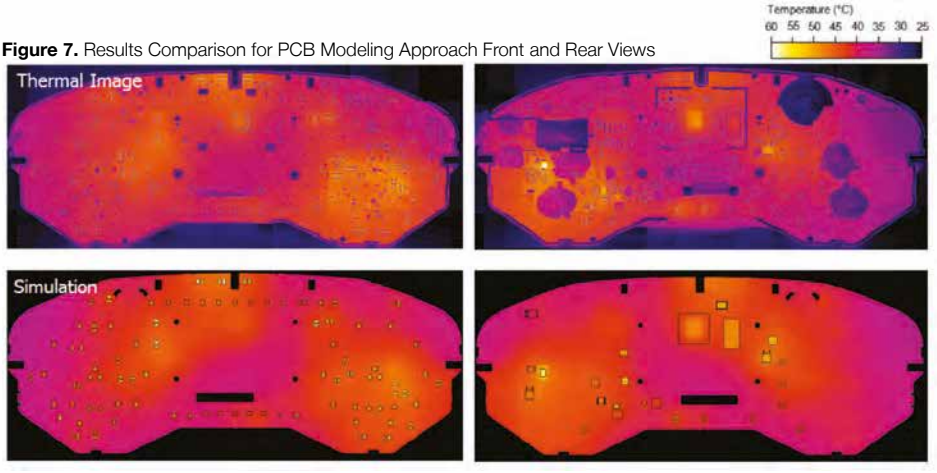

Housing Modeling

The housing (Figures 8.1 & 8.2) was designed in CATIA. It contains features such as screw holes and pointers: these could potentially increase the thermal analysis solution time with no gain in the level of accuracy. Other features, such as the back cover holes were important to keep as they allow some airflow movement inside the cluster. The mesh settings were therefore defined so that the 2mm hole opening would be captured appropriately by a minimum gap size setting applied to a local mesh region. The remainder of the mesh was set-up automatically and was able to capture all the relevant details such as the PCB, its components and the housing as well as the critical heat transfer paths between the PCB and the housing.

There is a fairly high temperature distribution (Figure 9.1) on the components along with recirculation zones in the cluster, highlighting issues with the airflow distribution. This venting pattern does not provide effective airflow movement inside the cluster to cool the PCB efficiently.

Component	Original Vents	Larger Vents
Linear PS	135 °C	132 °C
Reverse Protect Diode	122 °C	118 °C
Microprocessor	122 °C	118 °C
Backlight LED	125 °C	121 °C
Backlight LED	123 °C	115 °C

Figure 9. Thermal Analysis Results

To improve that, the venting pattern was modified in FloTHERM XT to add openings and also make them slightly wider. This design change can be easily implemented using sketches which are then extruded to remove material from the housing. The results of this design modification are displayed in figure 9.2 and emphasize that the airflow distribution was improved, allowing components to operate at a lower temperature. It can also be noted that the PCB temperature decreased accordingly.

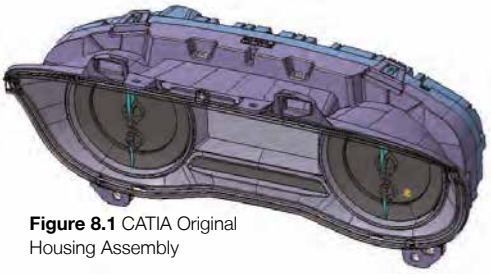

Figure 8.1 CATIA Original Housing Assembly

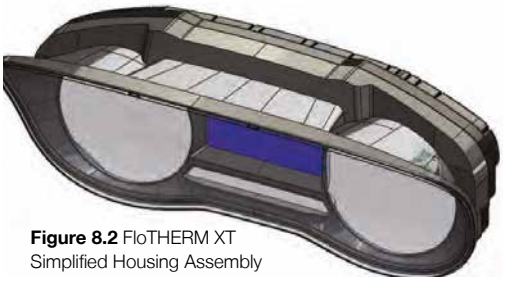

Figure 8.2 FloTHERM XT Simplified Housing Assembly

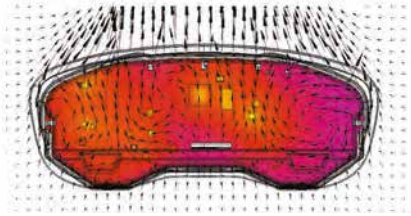

Figure 9.1. Original Venting Pattern

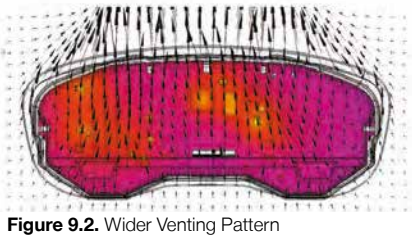

Figure 9.2. Wider Venting Pattern

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.

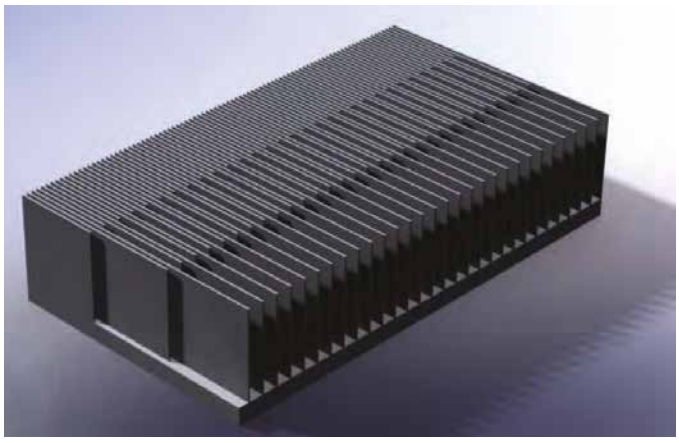

Figure 1. Preliminary Air Cooling (Heatsink) Design

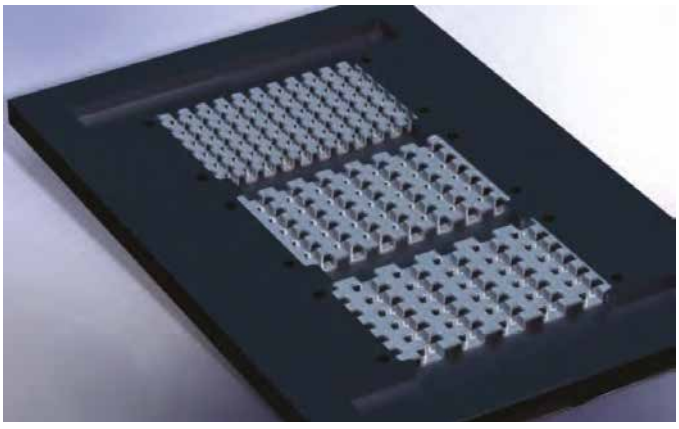

Figure 2. Preliminary Water Cooling (LCP) Design

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.

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.
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.

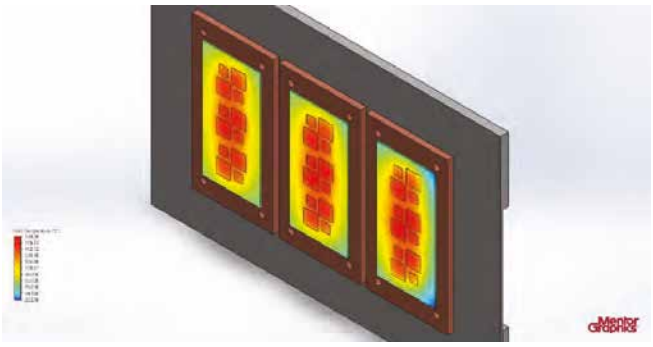

Figure 3. Chips Temperatures (Liquid Flow from Right to Left)

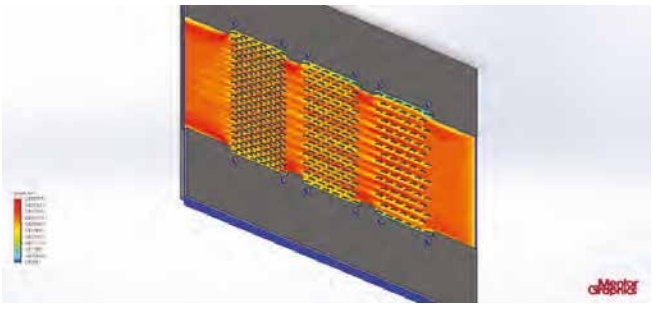

Figure 4. Fluid Speeds

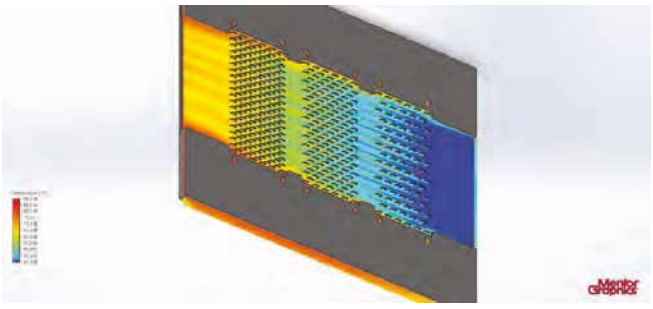

Figure 5. Fluid Temperatures

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.

Heat Pipe Heatsink Design in a Wipro Ltd. 3U High-End Server

Wipro Ltd is a global information technology, consulting and outsourcing company with 145,000 employees serving over 900 clients in 61 countries, headquartered in Bangalore, Karnataka, India. Wipro's Engineering Design Services Group has over 400 experienced professionals with expertise in advanced materials, CAD and CAE tool customization and competence in design, analysis prototyping and testing, covering thermal, vibration, acoustic and shock.

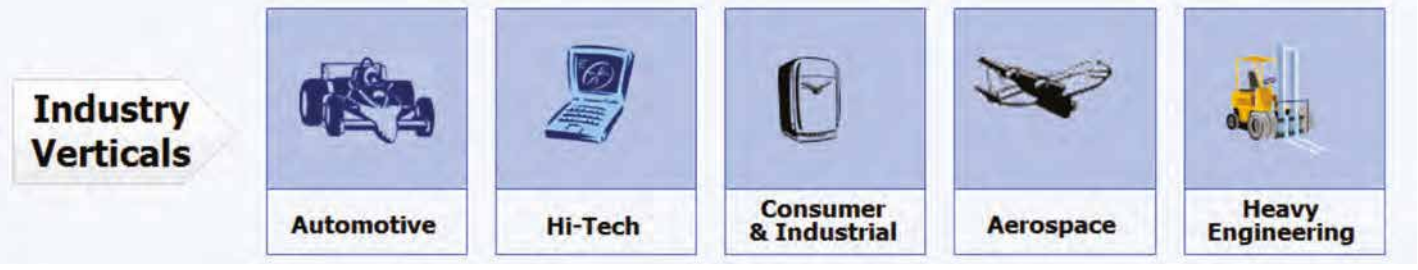

Wipro's team of experts were contracted to provide the thermal design for a 1.5kW, 3U server with a predefined form factor for the enclosure. The system would contain multiple fans, and was required to work when one fan failed (N+1 active-active resilience) to meet operational reliability requirements. Within this specification Wipro were required to perform fan selection and ducting design to optimize the airflow.

Fan make	Rated Input (W)	Airflow (CFM)	Static (Inch H_2O)	Rated speed (rpm)	Rated current (A)	dB(A)
Fan A	23.52	32.349	2.382	15150	1.96	64.4
Fan B	19.20	32.5	2.61	17500	1.6	64

Two fans were chosen for evaluation. When compared, the components performed almost identically in terms of airflow and so as expected were found to perform the same from a thermal perspective. Therefore, the only further considerations in the evaluation to be made were cost, efficiency, and lifetime. All things considered, this led to the selection of Fan B, which produced slightly less noise and consumed less power.

Having selected Fan B, the impact of fan failure on critical component temperatures within the system was studied. Using FloTHERM to fail each fan in turn, to examine the impact on system temperatures, revealed that for all fan failed cases the component temperatures remained acceptable.

A key challenge Wipro faced was selecting a heatsink for the two COMPs that would keep the processor junction temperature under the vendor's specified junction temperature limit of 96°C for power dissipation, 155W per COMP. This involved pairing the heatsink with a duct design that would ensure a uniform airflow through each COMP for any fan failed condition with minimum pressure drop.

Wipro selected two heatsinks for comparison. The first, heatsink A, is fabricated from copper. Wipro performed flowrate vs pressure drop measurements on the heatsink in order to accurately include it in the FloTHERM model, and found that a pressure drop of 0.18 inch of H_2O is required to deliver the minimum airflow of 33.5CFM.

An alternate, heatsink B, fabricated from aluminum was also measured and its performance in the system simulated. Despite the cost and weight advantages of using aluminium, heatsink B was found not to be suitable for the application. Although it resulted in just over a 3°C increase in temperature, this almost halved the thermal margin for the design from 6.8°C to 3.6°C.

For the IC critical component, which was offset from the main airflow path, a heat pipe was needed to carry heat away from the chip to keep the IC's junction temperature under the vendor specified junction temperature limit of 105°C at power dissipation 50W. Wipro investigated different designs of chiller block to compare

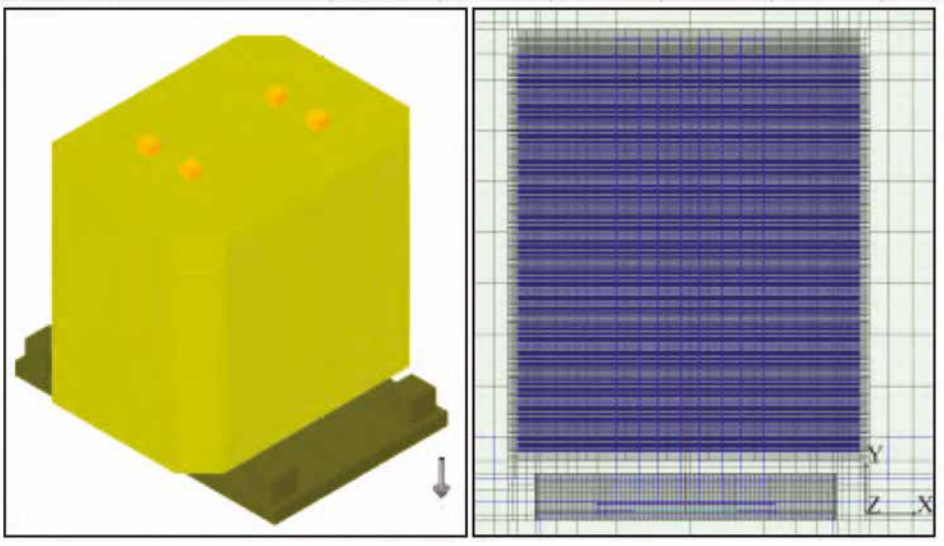

Figure 1. FloTHERM Model of Heatsink A in 40 Million Cell System Model (1.2 Million Cells within Heatsink with 5 Cells per Fin Channel)

2 x 6mm, 3 x 6mm and 2 x 8mm diameter heat pipe configurations, that connected to a finned region directly behind the COMP's heatsinks. Simulation showed that the best arrangement was to use 2 x 180mm long 8mm diameter heat pipes with 46 fins and a 3mm fin spacing, which reduced the junction temperature down to 96°C.

When the prototype was built, Wipro found there was only an ~8% variation in airflow rate from simulation to validation, with the measured airflow being lower, due to extra resistance created from the cables inside the system which were ignored in the simulation. This is well within the design margin, experimental uncertainty and fully met the client's expectations.

Cases	System	Dimms1	Dimms0	CPU1	CPU0	Interconnect chip
Fan:1 failed	136.68	67.76	68.95	49.20	48.49	26.46
Fan:2 failed	137.00	67.54	69.47	49.01	48.84	23.87
Fan:3 failed	137.12	67.38	69.72	48.90	49.07	23.67
Fan:4 failed	137.32	66.33	70.95	48.04	49.95	23.11
Fan:5 failed	137.16	65.78	71.32	47.60	50.25	24.36
Fan:6 failed	136.06	64.84	71.22	47.14	50.15	26.70

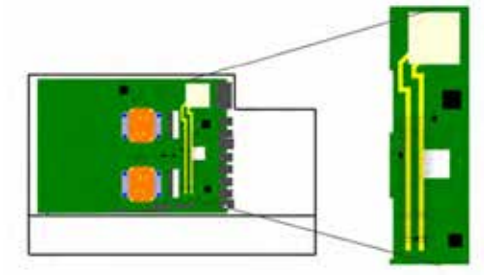

Figure 2. Detail of Heat Pipe Design for IC

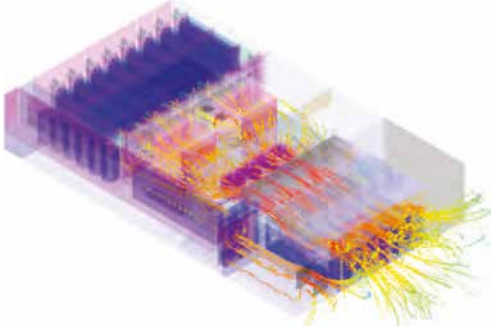

Figure 4. Final Simulation Model

Sl.No	Component Name	TDP (W)	Allowable Tj (°C)	Simulated Tj (°C)	Calculated* Tj (°C)	Simulated airflow , CFM	Validated airflow , CFM
1	Processor, CPU 0	155	96	89.7	92.5	109.5	100.5
2	Processor, CPU 1	155	96	90.2	93.8		
3	Interconnect chip	50	105	99.6	102.9		

Figure 5. Final Model Accuracy Validated Against Experiment (*calculated Tj based on measured case temperature and junction to case thermal resistance).

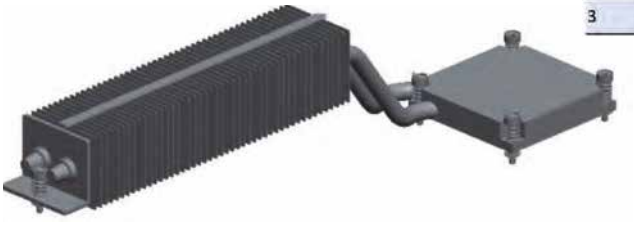

Figure 3. Detail of Heat Pipe Design for IC

By John Parry, Industry Manager, Mentor Graphics

FloTHERM® - Powering Power Supply Development

A case study of Emerson Network Power

By John Parry,
Electronics Industry Vertical
Manager, Mentor Graphics

The development of electronic devices has seen a trend towards their miniaturization and an increase in their functionality. Both of these put high demands on heat dissipation. High MOSFET temperatures can lead to instability, shortened life expectancy, and even lead to the device exploding! High temperatures can also contribute to thermal breakdown in electrolytic capacitors, dramatically shortening their life. While inductors and other magnetic components can withstand high temperatures, peripheral devices can be impacted.

To ensure that the power devices operate effectively within a given temperature range, attention needs to be paid to a product's thermal design. Reliable thermal design can guarantee the life and reliability of the product, and ultimately reduce costs in the long-term.

Emerson Network Power are a $6.2bn company with about 45,000 employees [1]. Thermal Management is a Center of Expertise that has been strengthened by the acquisitions of Liebert, Hiross, Cooligy and Knurr, to critical cooling issues at room, rack or chip level. Emerson take a holistic approach to thermal design, using FloTHERM from the chip package up to the room in which the equipment sits, to produce the best overall thermal design. This endeavor is greatly assisted by FloTHERM's Cartesian meshing supporting localized mesh regions, allowing multiple length scales to be included within a single model.

At the package level, Emerson aim to create a thermal model that is as close to reality as possible to help evaluate the transient response of their power electronics components with accurate junction temperature prediction – the goal of all thermal design studies.

At the PCB level, emphasis is placed on using the correct material property values, capturing the actual copper distribution in the board. Attention is paid to refining the mesh around large and high power components such as, MOSFETs, to accurately capture temperature gradients in both the air and the solid. Thermal radiation contributes to both heat loss and heat exchange between components, so radiation is also included in the calculation, so the surface emissivity of the components and board must be factored in.

At the modular level, the environment is an important factor to capture in the simulation. In the case of a small sub-rack for example, which is cooled by natural convection, the ambient temperature can be 30°C. With stratification in the ambient air affecting the air flow through the enclosure, the simulation model must be extended to include the local environment. As the scale of the model increases, care is needed to preserve the fidelity of the modeling effort.

Comparison with test data becomes tricky, as it is difficult to control the environment, particularly when testing products cooled by natural convection, as small drafts and other effects can affect the measurements.

Integrating the electronics together at the cabinet level adds a new level of complexity to the simulations, as failures or disturbances in the power network can cause the supply voltage to drop. Under these circumstances, the power electronics must be able to function as normal to avoid an impact on downstream equipment.

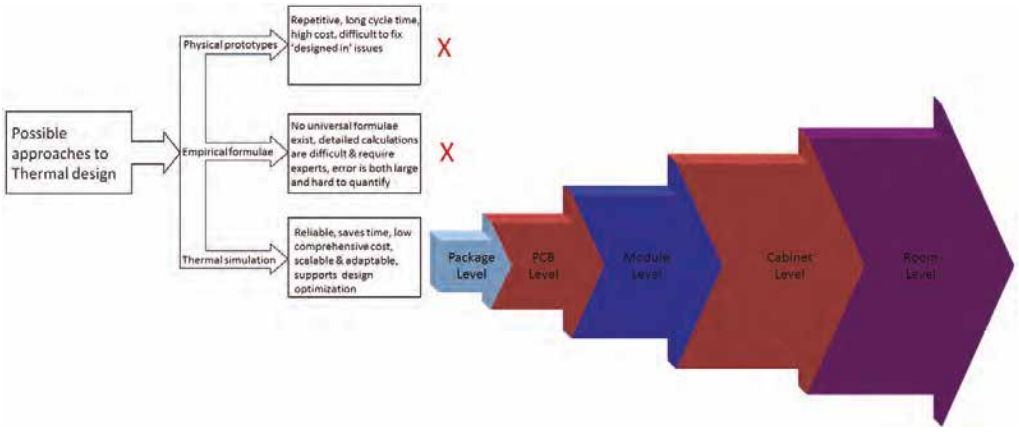

Figure 1. Emerson's exclusive use of thermal simulation at all packaging levels

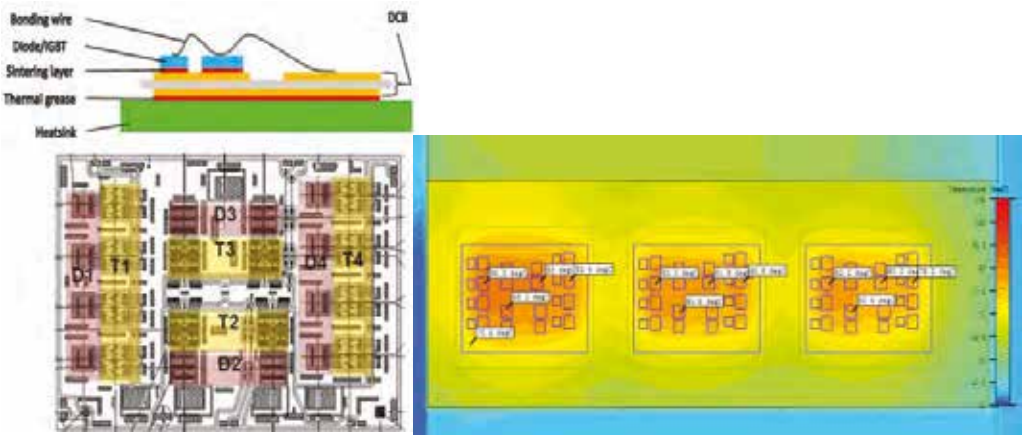

Figure 2. Stackup of Power Module showing die detail, and FloTHERM model

When operating in this manner, the power dissipation within the devices increases, leading to a further rise in temperature within the components in just a few seconds, this must however, remain below the maximum operating temperature for the components.

Emerson Network Power's equipment can be used in multiple configurations and in a variety of environments, so there is considerable scope for equipment-to-equipment thermal interaction at the room level. In some installations, this is prevented by using a room-level cooling strategy, such as cold aisle containment in datacenters.

To make the most efficient use of FloTHERM, Emerson use Mentor webparts [2] to define models of porous materials, centrifugal blowers, thermoelectric coolers (TECs), fans at high altitude, and complex geometry. In this way thermal simulation is repeatedly used to validate ideas during design and improve knowledge sharing with colleagues and other stakeholders.

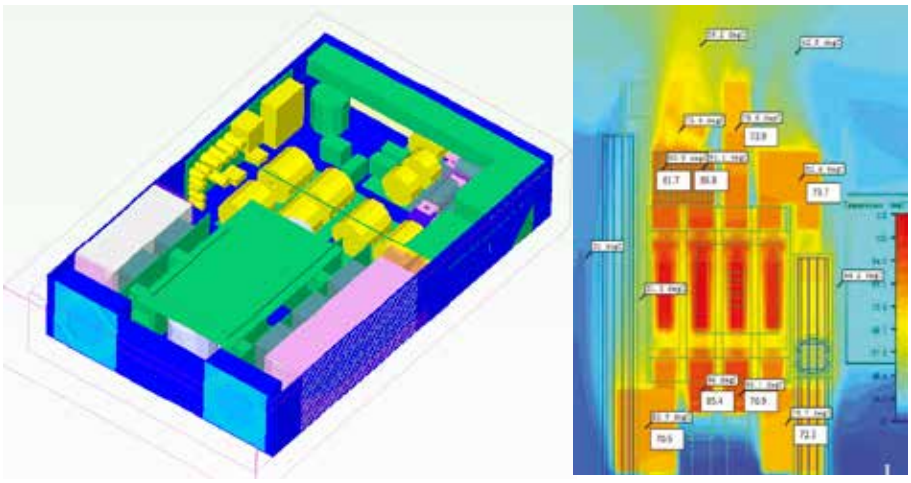

Figure 4. Typical Module and Temperatures within a Sub-Rack

Emerson also make extensive use of thermal measurements, including Mentor's Transient Thermal Tester, T3Ster®, to obtain thermal data for components and heatsink thermal resistances, to increase the fidelity of their thermal simulation work – a practice that is becoming increasingly common in best-in-class companies.

References:

[1] www.emersonnetworkpower.com/documentation/en-US/About/Documents/FactSheetEmersonNetworkPowerFY13.pdf
[2] webparts.mentor.com/flotherm/support/webparts.jsp

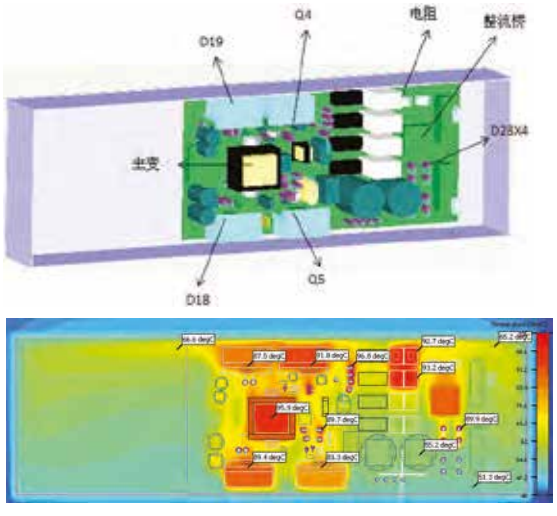

Figure 3. PCB and FloTHERM simulation results

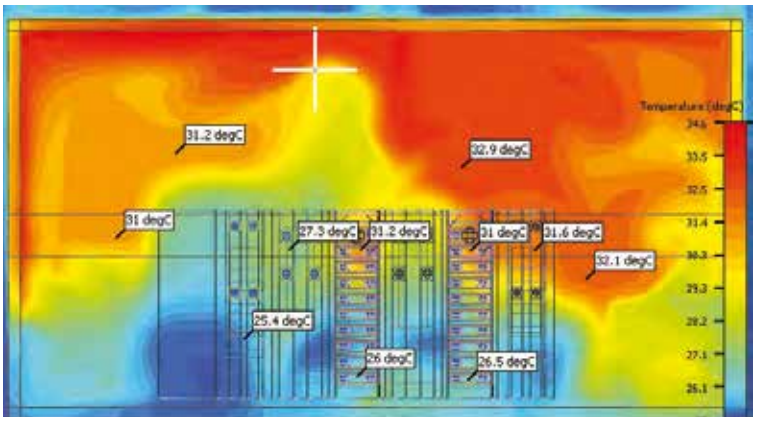

Figure 5. Temperature distribution through several racks showing stratification in the room

Thermal Simulation to Model Airflow and Heat Dissipation for Vehicle ECUs

By Kelly Cordell-Morris, QA/Test Engineer, Mentor Graphics

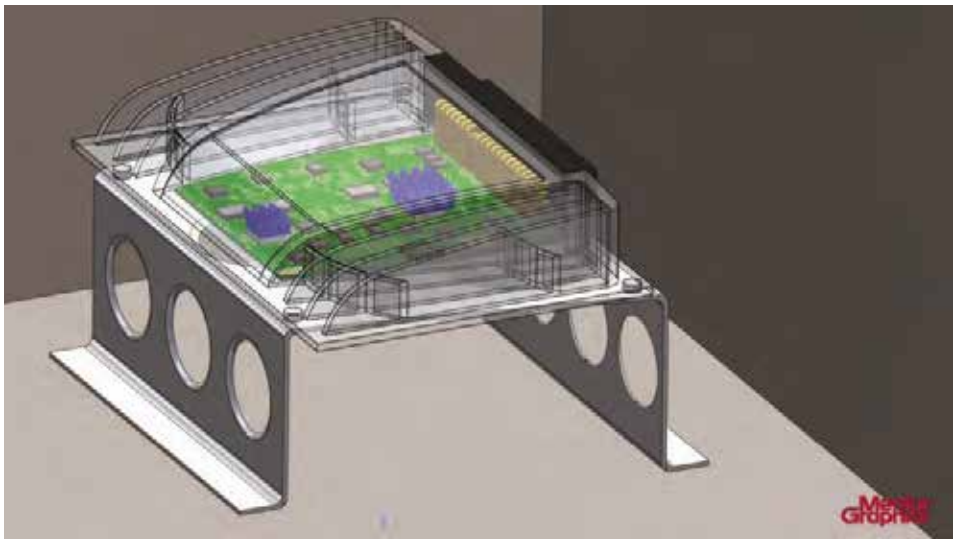

Figure 1. The ECU with the PCB

Electronic control units (ECUs) are a vital part of vehicle operation, controlling engine functions - fuel injection, ignition timing, and idle speed control. The ECU is part of a feedback loop in which data from sensors around the engine is monitored and used to optimize the control outputs. The ECU must be highly reliable in a harsh environment. Components can be exposed to extreme temperatures, and the push for minimal footprint designs increases the thermal issues. It must take into consideration environmental implications and maintenance-free operation, as well as suitability for high volume manufacturing.

Historically, automotive electronics have been cooled via convection by ambient air. Within the engine bay, ambient air temperatures vary significantly. Often, the ECU is tucked away leaving convection around it compromised. Electronics modules in vehicles, and in particular, modules mounted out-of-cab, are often sealed to prevent moisture ingress. Preventing external air flow from circulating around the electronics directly, making forced convection impossible.

Convection cannot be relied on as a heat transfer mechanism in these situations. Conduction, normally to the vehicle structural body, becomes the dominant means by which to cool components. With a conduction-influenced design, the materials used in the ECU are an important part of thermal considerations. In this case, we are comparing materials used in the mounting brackets and their effects on the ECU's internal component average temperatures.

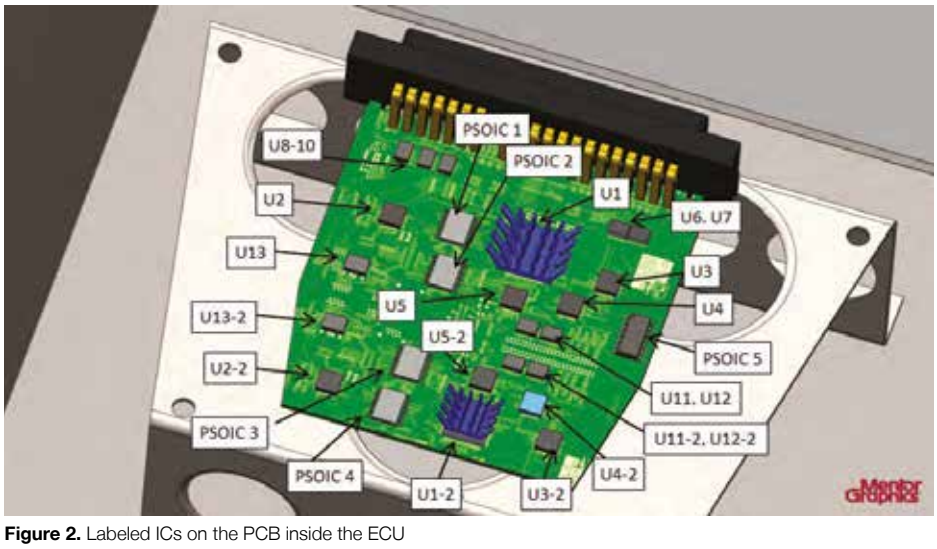

Figure 2. Labeled ICs on the PCB inside the ECU

Modeling the ECU as Mounted in a Vehicle

In the FloTHERM XT model, the ECU was mounted to a metal bracket, mounted to the vehicle chassis (Figure 1). Two thin walls were located close to two sides of the ECU to block airflow around the unit. Total heat dissipation in the unit was 19.8 W (Joules per second).

We modeled two internal ICs as two-resistor (2R) representations to obtain a prediction of junction temperature. Each one dissipated 4.5 W. Pin heatsinks were also mounted on these 2R components. The remaining 24 IC representations were simple conducting cuboids with typical effective conductivity properties applied depending on the type of IC represented. These dissipated between 0.1 to 1 W each. The results give an indication of case temperature (Figure 2).

Results for Unmodified Mild Steel and Aluminum Brackets

With walls on two sides, there was, as expected a compromised airflow, and hotter

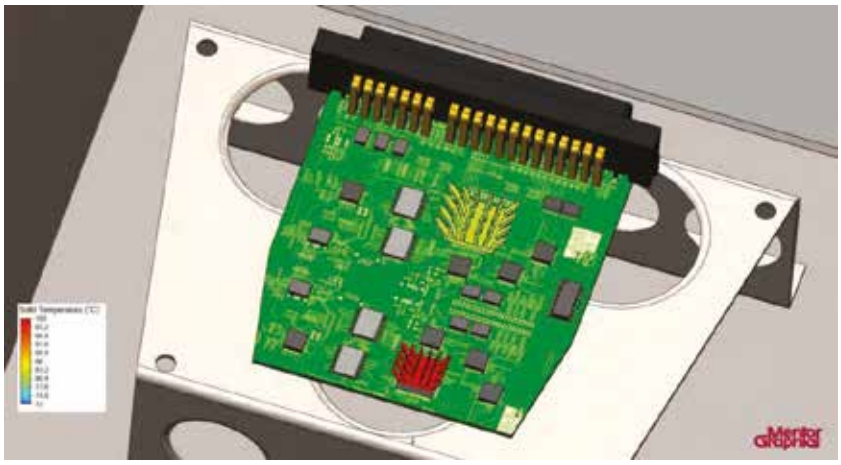

Figure 3. Heatsink surface temperatures for aluminum bracket in FloTHERM XT

enclosure when made of mild steel. Plus a higher airflow velocity resulting from energy lost via convection.

A change of material provided an average improvement of 8°C. PSOIC3 located over a hole in the bracket recorded the smallest change, with convection playing a large part in heat transfer making this IC less sensitive to changes in the material. However, U1 would be more sensitive to change given its location. Including a heatsink gave an alternative heat transfer path to conduction (Figure 3). In conclusion, holes may help airflow; however, on occasion solid mounts may be preferable to maximize conduction when convection is virtually zero.

Results with Solid Bracket

A new configuration for the bracket where the holes were suppressed was created for both models. Another configuration for each model with the heatsinks removed was also solved to determine just how necessary heatsinks may be to the final design (Figures 4-6).

With a fully solid mounting bracket IC temperatures were generally reduced by around 4°C. Clearly, conduction was more efficient with the new design despite convection being compromised by the removal of the bracket cut-outs. Changing the bracket material to aluminum reduced temperatures by a further 7-8°C on average. We observed the greatest temperature drop on PSOIC 1, which may be because of its location close to U1. With lower temperatures on the 2R model and thus lower temperatures for the surrounding PCB, more efficient heat loss from PSOIC 1 to the PCB can be achieved.

Results with No Heatsink

We saw little gain in temperature for the majority of the components. U1 and U1-2 junction temperatures were raised 1.7°C and 1.5°C, respectively, when the heatsinks were removed in the mild steel bracket model. The rise was slightly less when the bracket was aluminum. The greatest change for the 2R models was in case temperature as would be expected with the removal of a heatsink. For the other components on the PCB, PSOICs 1 and 2 were the only ones significantly affected. This is probably because of their proximity to U1, and PSOIC 1 was the most affected because it is closest. Assuming that the small increase in temperature for these few components is allowable, the addition of heatsinks would appear superfluous. Again, using aluminum for the mounting bracket reduced component temperatures by 7-8°C on average.

Conclusion

In cases where airflow around the ECU is compromised, material changes can have significant effects to be considered for the thermal design. Using a more thermally conductive material such as aluminum is a simple change that results in a robust and fairly inexpensive design.

Where convection, either forced or natural, is compromised or impossible, then a simple mounting bracket design can be more effective than a design intended to increase airflow around the ECU casing. This has the added benefit of making manufacturing simpler and thus less costly. As has been shown in this case, heatsinks in a closed system may not be that effective and add unnecessary cost to the final design. Thermal simulation software provides a quick and easy way to test material and thermal changes early on in the design process.

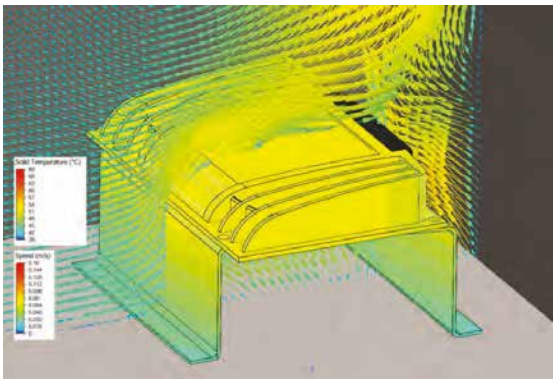

Figure 5. Surface temperatures and air flow vectors for enclosure and aluminum solid mounting bracket with no heatsinks

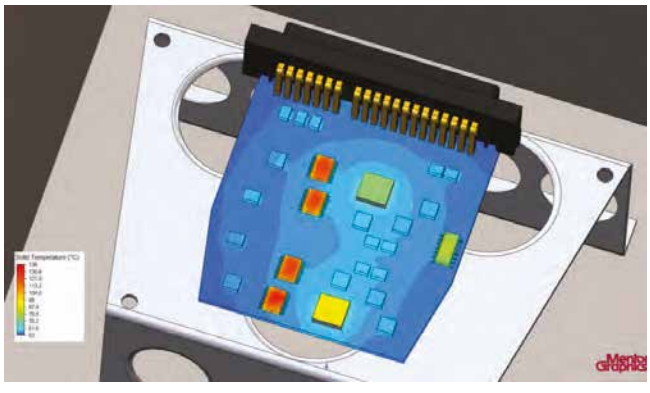

Figure 4. IC and PCB surface temperatures for aluminum solid mounting bracket (heatsinks not shown)

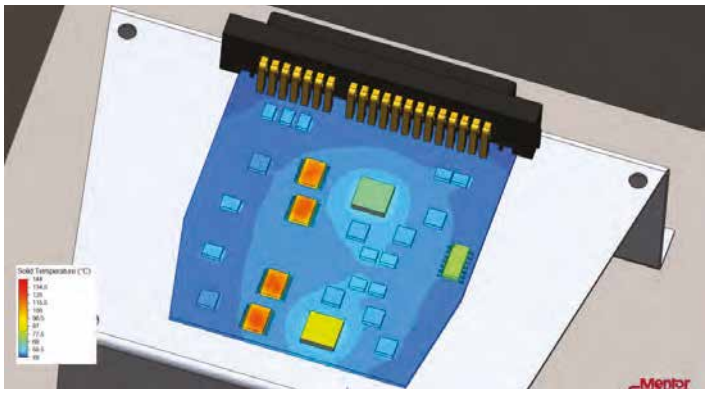

Figure 6. IC and PCB surface temperatures for aluminum solid bracket with no heatsinks

The Three Waves of Commercial CFD

By Ivo Weinhold, User Experience Manager
and John Parry, Mentor Graphics

In recent years, many papers have been published on the history of flow simulation. Many early CFD pioneers like Brian Spalding, David Tatchell, Ferit Boysan and Michael Engelman have talked or written about their memories. This pool of historical facts, technical information and personal impressions give a remarkably consistent description of the way engineering simulation software evolved from academic research codes towards the modern CFD products we know today. Developed and supported on an industrial scale by multinational software companies.

Closely linked to the performance of available computing hardware, this development was, particularly in the early stages, driven primarily by research and development projects for aerospace and defense, but latterly also increasingly by interest from civilian industry. Looking back, three major phases of the development of CFD software for industrial applications can now be recognized:

- The First Wave: The beginnings of commercial CFD software in the 70s and 80s.
- The Second Wave: In the 90s, CFD enters the research and development departments of large industrial enterprises.
- The Third Wave: After the millennium, CFD becomes an indispensable part of the product development process.

The First Wave: The beginnings of commercial CFD software

Since 1958 the codes of the CFD software engineers in the first phase had its roots in the work of the Fluid Dynamics Group T-3 at the Los Alamos National Laboratory (USA), and the research activities under Prof. D. B. Spalding at Imperial College London in the 1960s and 1970s.

In the late 1960s, Concentration, Heat and Momentum (CHAM) Ltd., founded by Spalding, and initially located at Imperial College London, dealt with consulting work. The era of commercial CFD software began in 1974, when CHAM Ltd moved to its own offices in New Malden near London. Initially, the development of customized CFD codes was central to the business activities of CHAM. That became too time-consuming and inefficient, so CHAM decided to develop a general-purpose CFD package for in-house consultancy work, and released this as a commercial product, PHOENICS, in 1981. This may well be regarded as the birth of the CFD software industry (see CHAM Ltd, 2008). Others quickly followed suit. For instance, Fluid Dynamics International (USA) followed in 1982 with FIDAP, a FEM-based CFD package, and in 1983 Creare Inc (USA) released the finite-volume CFD code, Fluent. Computational Dynamics/ADAPCO (UK/USA), co-founded by Prof. David Gosman, another professor at Imperial College London, released StarCD in 1989.

Figure 1. Fluid flow simulation in the 1980s, taken from Hanna & Parry (2011)

The basic technologies behind most of the CFD packages of this era had been created by former employees or guest scientists of the two aforementioned research institutions in London and Los Alamos, or were based on their scientific publications. But there were also other developments of CFD technology: in the 1980s alternative approaches for CFD simulation emerged as part of the military and civilian aviation and space program of the former Soviet Union, largely unnoticed by the Western scientific community due to the political situation. Their technical tasks for CFD simulations were similar to those in the West, but the available computing resource for their solution was much more limited. Conversely, because of the high political priority of these research programs, very extensive experimental data for numerous fluid flow and heat transfer phenomena, especially in the near-wall area, were generated. This situation led to the development of alternative CFD methods, which, building on known methods for

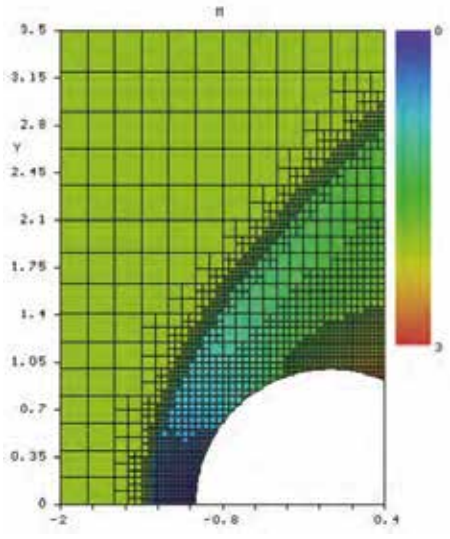

Figure 2. Result plot of Aeroshape-3D (Parry et al., 2012)

Cartesian grids as published in the scientific publications in the West, were based on a combination of numerical, analytical, and empirical data. This innovative approach yielded high-quality simulation results in virtually, arbitrarily complex computational domains while maintaining the low resource requirements and the effectiveness of methods using Cartesian grids. With the gradual economic liberalization in the Soviet Union in the late 1980s, several teams of scientists have commercialized this CFD technology and, since the early 1990s, sold their products and services in Europe and Asia. The best known products of this kind were Aeroshape-3D by Prof. V. N. Gavriliouk and team (Petrowa, 1998 & Alyamovskiy, 2008) and FlowVision by Dr. A. A. Aksenov and team (Aksenov et al., 2003).

From the beginning of the 1990s, the conditions for CFD software and simulations changed quite rapidly. Computer hardware, mathematical methods and physical models all experienced huge performance gains. Numerical methods such as unstructured Finite Volume methods, multi-grid methods, sliding mesh, etc., suitable for complex geometry and optimized for HPC, became commercially available as well as more reliable, more flexible and more broadly applicable physical models. CFD technology became much more feasible, and for the first time, quite realistic model sizes, for real industrial applications were possible. These new capabilities heralded a new phase in the usage of commercial CFD software - entry into the research and development departments of industry across the board.

The Second Wave: CFD enters the Research and Development Departments of the Industry

Using technology typical of the first phase, Flomerics Ltd., founded in 1988 by David Tatchell and Harvey Rosten in Kingston-upon-Thames (UK), played a pioneering role in marketing CFD software developed exclusively for industrial applications with its software package FloTHERM, first released in 1989. Both founders worked for CHAM Ltd in senior positions before leaving to found Flomerics, with the aim of "providing good science to industry" (Tatchell, 2009). FloTHERM was a first paradigm shift in the CFD industry, away from the focus on complex CFD technology, towards the solution of engineering tasks in industry as the central goal. This also meant that from then on engineers working in product development, and not scientists, were the main target users of this type of CFD software. The available CFD technology, computer hardware and operating systems imposed certain limits on such an innovative approach. Therefore Flomerics concentrated initially on only two application areas: electronics cooling (with FloTHERM) and built environment HVAC (with FloVENT). The requirements of engineering-oriented CFD software for these application areas were relatively clearly defined and, more important, also just feasible.

This concept opened up completely new market opportunities, because for the first time engineers in product development without special knowledge of numerical methods and without extensive CFD experience were empowered to employ CFD simulations as a development tool. The solution of a technical engineering task became the center of attention, while the underlying CFD technology was more or less just a means to an end.

Obviously, other CFD providers also recognized the beginning of this paradigm shift and especially the new business opportunities associated with it, responding to this trend with their own product offerings. Overall, huge investments from all CFD software vendors in better user interfaces, robust solvers and reliable physical models could be observed, with the clear objective of entrenching CFD into the research and development departments of large industrial enterprises and thereby attracting a new generation of CFD users.

After establishing CFD as a successful tool for the functional design, verification and optimization of product designs, features, processes and physical effects in large industrial companies in the early 2000s, the reputation of this technology amongst engineers improved significantly. As a result, the demand for CFD simulations showed strong growth, especially in medium-sized and small companies keen to reduce the costs associated with physical prototypes. Another important aspect was the need to integrate CFD simulation into the regular product development process, as these companies usually had as yet no

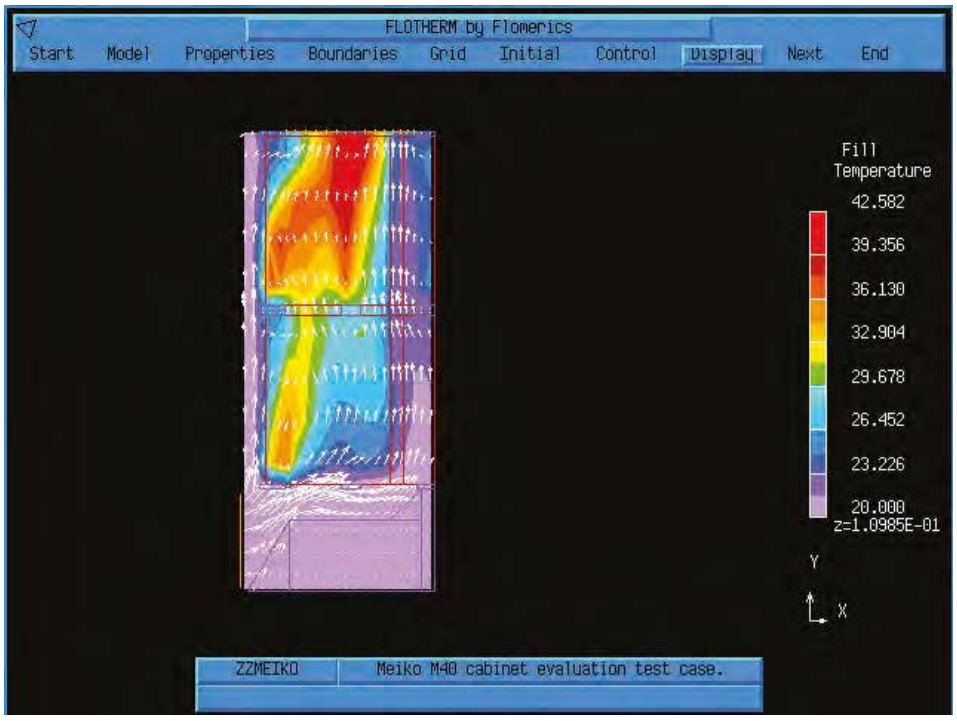

Figure 3. Early version of FloTHERM (Hanna & Parry, 2011)

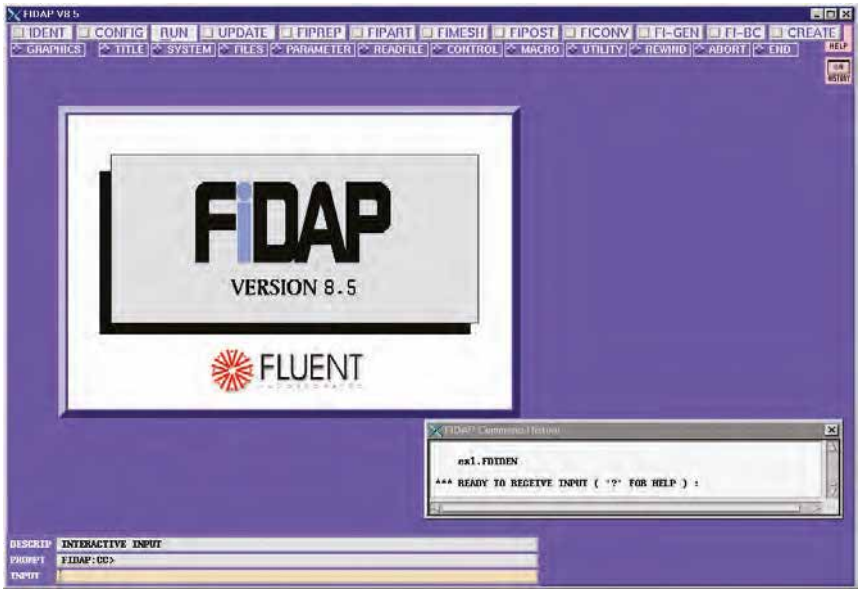

Figure 4. FIDAP User Interface in the late 1990s (University of Delaware, 2007) URL

dedicated simulation departments. This meant that qualified engineers from product development or design groups would perform the simulation themselves. The efficiency with which simulation projects were conducted had to be increased, so CFD results would be available sufficiently in sync with the product design cycles, for the results to help guide proposals for design improvements. In this context the handling of industry-level geometry played a key role. At that time this was already being provided as 3D CAD data which should, of course, be used with as little simplification and modification effort as possible to be useable by the subsequent and preferably fully-automated mesh generation step. The CFD software market responded to these demands with new and improved products – and a third wave of CFD software for industrial product design began and continues to this day.

The Third Wave – CFD Becomes an Essential Element of the Product Design Process

The major CAD and Product Lifecycle Management (PLM) vendors play a key role in this third phase. Since the 1990s, they have been successfully introducing the concept of PLM, which encompasses CAE. As a result, customers have put pressure on commercial CFD software vendors to conform to this concept and to take steps to integrate their products into the major PLM systems. In the 2000s, virtually all CFD software providers upgraded their systems with, at the least, CAD import interfaces. Some have developed bi-directional links with major CAD/PLM systems, and a few have even embedded their CFD technology directly into the 3D CAD systems themselves. New CFD techniques to support these requirements were also developed, partly from scratch, and partly as enhancements of existing technology.

The company NIKA GmbH, founded in 1999 as a German-Russian joint venture, was a typical example of a new commercial CFD software vendor emerging at the start of this third wave. NIKA exclusively developed, based on the above mentioned Aeroshape-3D technology, CAD embedded CFD software, which is now offered as FloEFD for several major 3D CAD systems.

The current third wave has allowed newcomers from other areas the opportunity to enter the CFD market, refreshing it with new technologies. But all have one thing in common: The industrial user, with his or her need for easy-to-use, task-oriented, automated, reliable, efficient and readily-available CFD software as an indispensable tool for digital prototyping is the focus. The result of changing development processes and, as a consequence, the changing role of the simulation engineer. Aspects like process integration, reliability, modeling safety, and reproducibility are becoming the center of attention, and influence purchasing decisions for CFD software. The further development of CFD software based around these requirements will bring exciting new technologies and products to market. A new fourth wave may be expected to follow soon…watch this space.

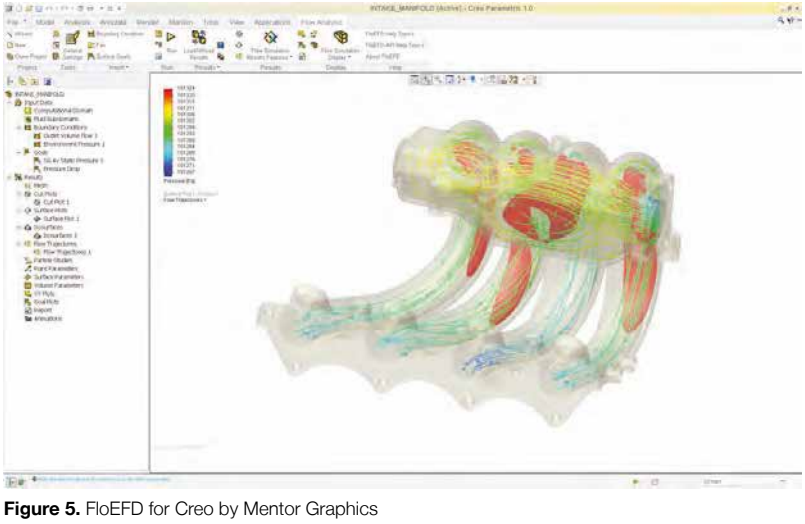

Figure 5. FloEFD for Creo by Mentor Graphics

References:

Aksenov, A. A., Kharchenko, S. A., Konshin, V. N., Pokhilko, V. I. (2003), FlowVision software: numerical simulation of industrial CFD applications on parallel computer systems. In Parallel Computational Fluid Dynamics 2003: Advanced Numerical Methods, Software and Applications, Elsevier, 2004, pp. 401-408

Alyamovskiy, A. A. (2008), SolidWorks 2007/2008. Компьютерное моделирование в инженерной практике, bhv-St. Petersburg, 2008, pp. 467-468

CHAM Ltd (2008), Earlier versions of PHOENICS: -81 to -1.6 A brief history. Available: http://www.cham.co.uk/phoenics/d_polis/d_chron/earlyver.htm

Hanna, K., Parry, J. (2011), Back to the Future: Trends in Commercial CFD, NAFEMS World Congress, Boston (Paper and Presentation Slides)

Parry, J., Kharitonovich, A., Weinhold, I. (2012), FloEFD – History, Technology & Latest Developments, Mentor Graphics, 2012

Petrowa, J. (1998), GUS - Informationstechnologien im CeBIT-Spiegel: Partner gesucht. ComputerWeekly, Volume 8, 1998. http://scripts.online.ru/it/press/cwm/08_98/gus.htm

Tatchell, David (2009), David Tatchell's Blog, Mentor Graphics, 2009. Available: http://blogs.mentor.com/davidtatchell/

University of Delaware (2007), FIDAP/GAMBIT. www.udel.edu/topics/software/special/statmath/fidap/

Electronics Thermal Design with Thales and FloTHERM® XT

A Step Further in Thermal Modeling of Electronic Components

By Eric Monier-Vinard, Thermal Domain Manager,
Thales Corporate Engineering

The goal of electronics thermal design is to accurately predict component junction temperatures to ensure that they are within specification. Easier said than done. Before CFD was used, designers used simple metrics, such as junction-to-case thermal resistance, as a 'thermal model' of the component in calculations by hand, with very wide safety margins to ensure the design was thermally viable. CFD allowed designers to predict the flow of cooling air, and include 3D thermal simulations of the board and components, increasing the need for more accurate component-level modeling.

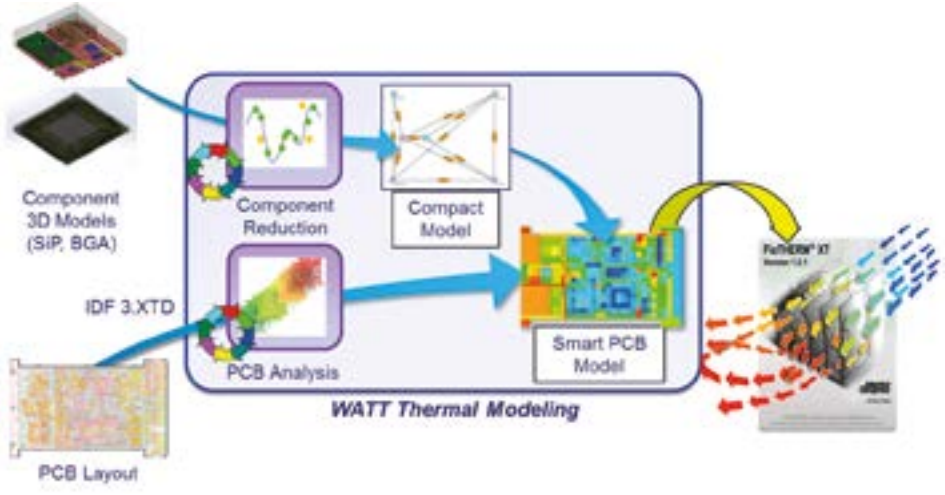

Figure 1. Thales' Thermal Analysis Workbench (WATT)

Various methods were devised in the 1990s. Junction-to-case and junction-to-board thermal metrics were combined to form a 2-resistor model, and the DELPHI Consortium developed multi-resistor models that accounted for multiple heat flow paths in the package, increasing the predictive accuracy further. The most accurate thermal models, which also account for transient effects and are able to handle multi-die packages, are detailed 3D conduction models. The increasing use of miniaturized high-powered devices and High Density Interconnection boards intensifies the coupling effect with neighboring thermally-sensitive components, increasing the need to predict the temperatures of all components accurately.

The Thales Corporate Engineering Thermal Team is responsible for the introduction of new technologies inside the Thales Group and is consequently at the leading edge of thermal research. This is aimed at achieving more accurate simulation results in the shortest possible time to meet the industrial requirements of the divisions they support, including Defense, Aerospace & Space and Security. As a DELPHI consortium partner, the team has continued its own research on the use of reduced order models, created from detailed models, to provide Thales' divisions with the resources they need through Thales' Thermal Analysis Workbench (WATT).

Until now, this effort has been hampered by the inability to incorporate all thermally-relevant details into the detailed model due to the large number of microscopic elements that are present within an electronic component. More than ever, a fine representation of all the details of a small package is today mandatory to avoid an overestimation of the semiconductor temperature.

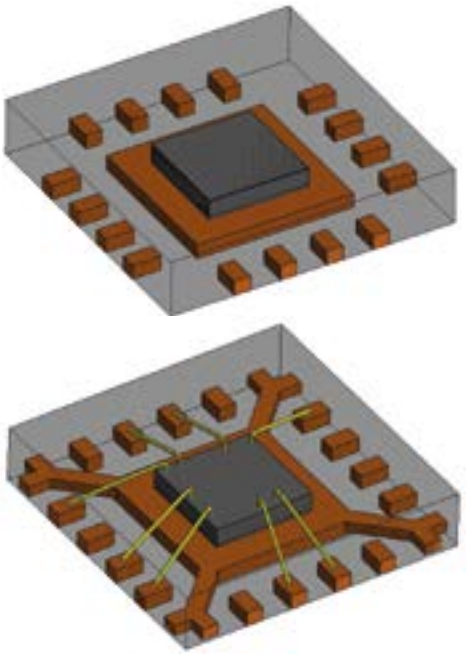

Figure 2. The realistic modeling of a QFN 16 package reduces the temperature prediction by 20%

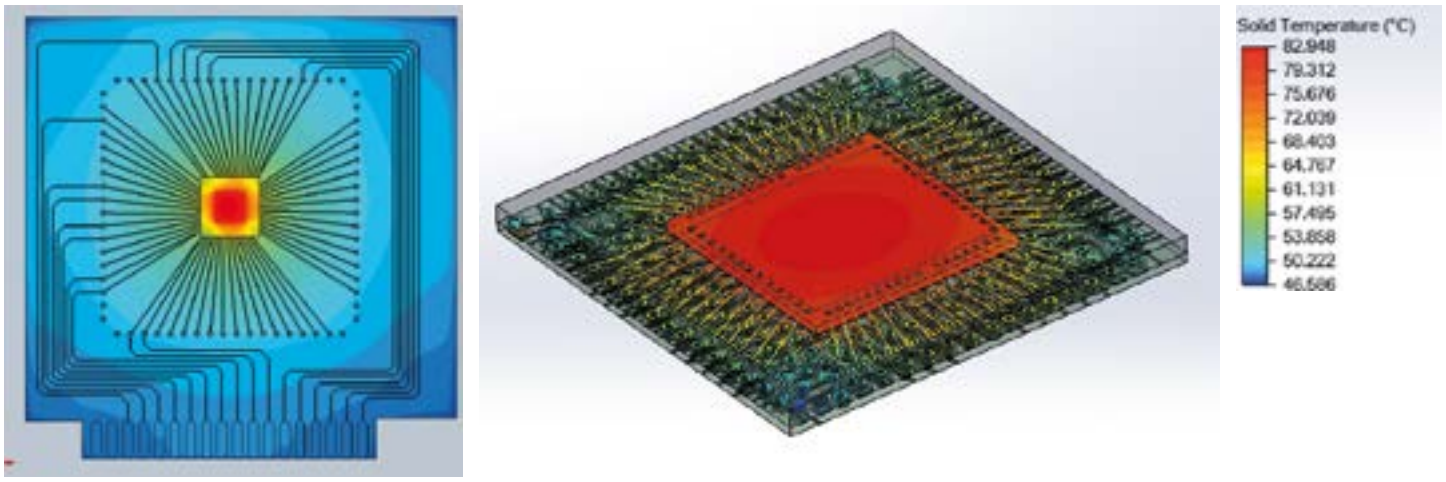

Figure 3. Component on JEDEC 2s2p test board in FloTHERM XT, showing internal detail including bond wires

There has always been a conservative design margin applied at the component level due to the fact that it has either taken too long or simply been impractical to take into account all the geometries inside a component package. For example, the detailed copper traces, copper vias on the substrate as well as the bond wires between the die and the substrate were rarely modeled explicitly, but are known to contribute to the heat spreading. Until now these very small elements were either roughly represented by single parts with averaged thermal properties or simply ignored. Their replacement by single aggregated parts introduces some inaccuracies in the results, while ignoring them leads to a higher calculated temperature and consequently higher margins during the design process.

With FloTHERM™ XT, the Core Thermal Team has been able to take a major step forward, producing, in just a few hours, results for System-in-Package devices or conventional BGA or LGA packages including all geometric elements inside the package.

Velocity	Model	T_{EXP}	T_{CFD}	%E	θja_{EXP}	θja_{CFD}
V = 1 m/s	Model without bond wires	146.2°C	149.4°C	2.1%	39.1°C/W	40.2°C/W
V = 1 m/s	Model with bond wires	146.2°C	146.9°C	<1%	39.1°C/W	39.3°C/W
V = 2 m/s	Model without bond wires	93.0°C	95.0°C	2.8%	23.9°C/W	24.6°C/W
V = 2 m/s	Model with bond wires	93.0°C	92.6°C	0.6%	23.9°C/W	23.8°C/W

Table 1. Thermal performance comparison

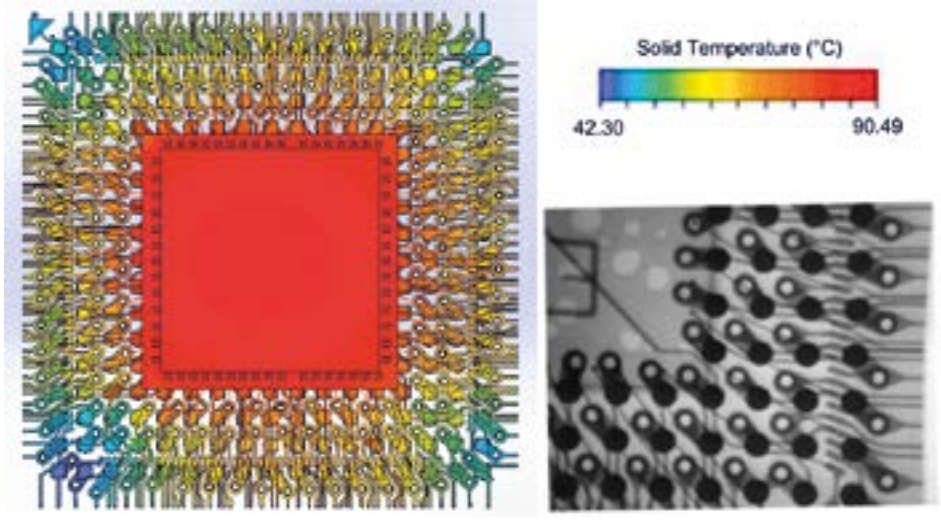

Figure 4. Detail of heat spreading throughout complex copper traces and vias (inset: X-ray showing bond wires)

For instance, the Thales Core Thermal Team has been able to import the complete geometry of a FpBGA 208 package with all its internal details as well as its supported board test vehicle, then set the general boundary conditions in just a few hours.

The meshing strategy of FloTHERM XT, which is based on the local size of the different parts of the model, requires very few user inputs and allows for the creation of an appropriate and easily solvable 1.9-million-cell mesh in less than three hours of computation on a 12-core Intel Xeon processor. The powerful solver needed less than 4.5 hours to reach full convergence on the same processor using 10 Gb of memory. This short computational time has allowed quick comparisons on the influence of the 25µm (1mil) bond wires on the package's thermal performance both in natural convection and in forced convection at different velocities.

The numerical results are very close to the measurements already conducted on this component, and are within 1% for the natural convection case, as shown in Table 1.

Increased simulation accuracy is the only way to break the conservative design margins used in the past.

Respecting these former margins would cost a lot more today than in the past due to the increased power density, so it is essential that cooling systems are

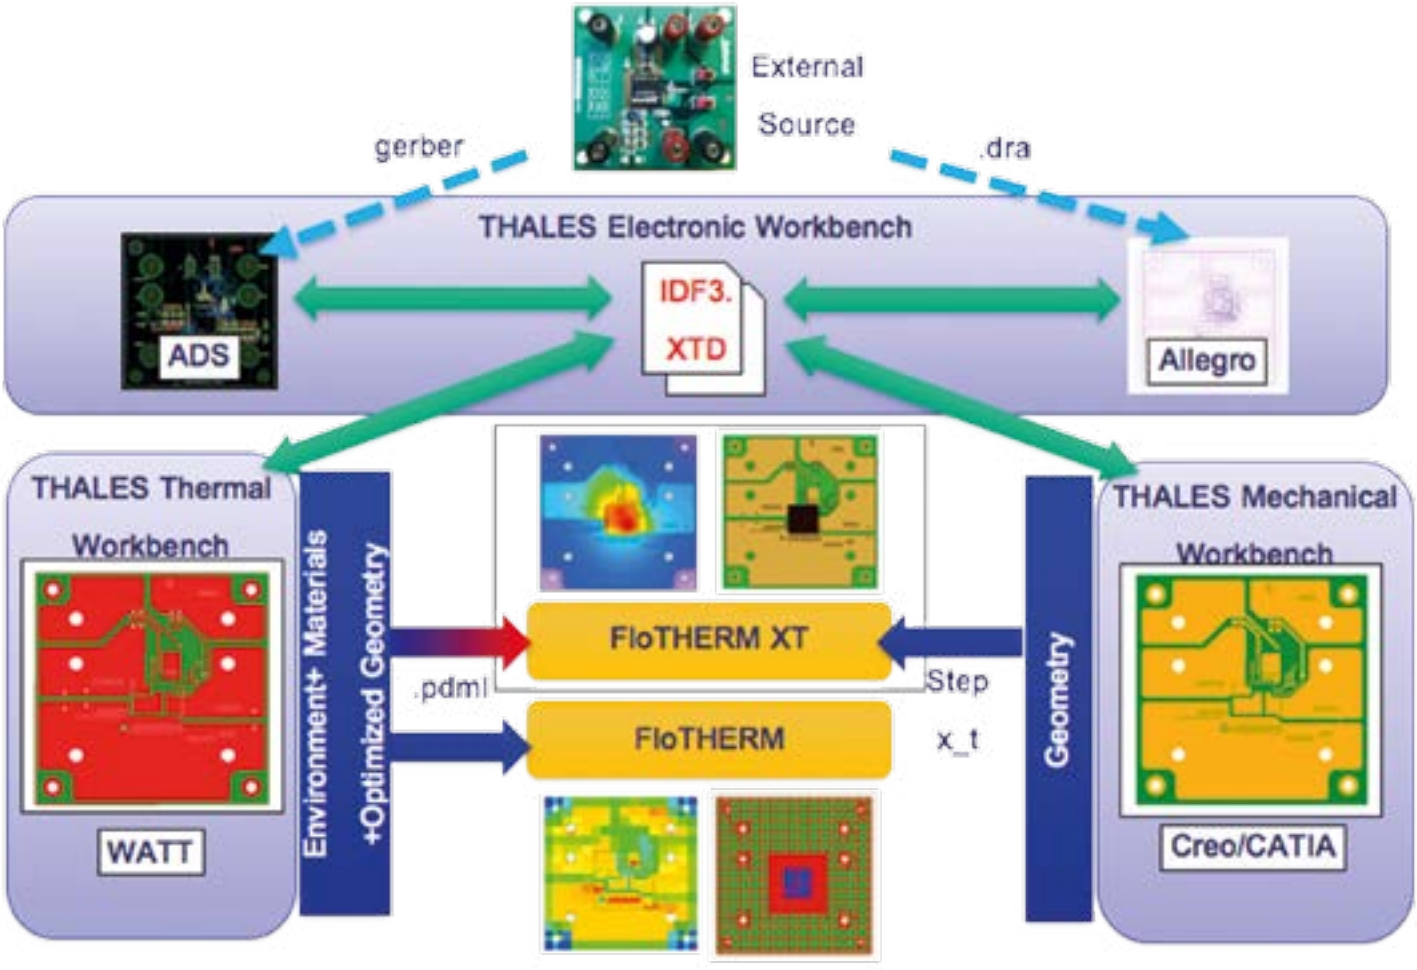

Figure 5. FloTHERM XT is poised to take a major role in Thales' overall Thermal Design Workflow

made as efficient as possible, and for that simulation accuracy at all packaging levels is needed.

Further, a fine representation of FpBGA 208 internal structure permits to better understand the thermal constraints encountered by the PCB interconnect balls, especially at corner locations. Figure 3 highlights a temperature gradient of 48°C for the set of interconnect balls.

If the modeling of the internal structure of the electronic component is crucial to accurately predict the temperature of its chip(s), the layer layout and copper trace design of the electronic board is now essential to efficiently optimize the way the heat is spread throughout its structure. Even there FloTHERM XT can simulate the small and thin elements that make up its composite structure.

This new approach, afforded by FloTHERM XT, means that the conservative design margins of the past can be reviewed, paving the way to accurately predict the thermal behavior of systems at all packaging levels, and particularly at the component level where the highest temperature gradients are located. This will allow Thales to better integrate cooling systems, even in cases where it was impossible with the old conservative margins. And sometimes it helps to comprehend previously misunderstood multiphysics issues.

"The thermal design of electronic component is under increasing control. With FloTHERM XT we can import the complete geometry of a FpBGA 208 package with all its internal details, test board, setup the boundary conditions and solve it in just a few hours. This will allow Thales to better integrate cooling systems, and comprehend previously misunderstood multiphysics issues."

Eric Monier-Vinard, Thermal Domain Manager, Thales Corporate Engineering

Tablet Teardown

Challenges in the Thermal Management of Forced Convection Tablets

By Arun Raghupathy & Bharath Nagendran
Electronic Cooling Solutions Inc.

One of the big changes in computing is the move towards handheld computing, a progression from using laptops and notebooks. Initially the focus was on browsing the web, online shopping, reading e-books etc. Latterly, with ever-increasing internet speeds and reduced access-costs, more often these devices have been used for streaming videos and graphics intensive games.

Though laptops are still widely prevailing, particularly for running CPU and graphics intensive programs, tablets are increasingly used complementarily as a handy digital-assistant.

Getting laptop-like computing performance in a tablet form factor is not an easy task. The ultra-thin form factor of tablets and the densely packed electronics makes forced-air-cooling particularly difficult, yet emerging as inevitable due to performance demands.

Microsoft's Surface Pro, which made its US debut in February 2013, could be converted between a laptop and a tablet (sometimes referred to as a 'laplet'). Despite much praise, Wikipedia notes it received some mixed reviews related to short battery life, bulkiness, excessive heat and fan noise under high load [1]. Yet it was the first product to take mobile computing into a tablet format, running 64-bit Windows 8 Pro, and marketed as "the tablet that can replace your laptop".

With expertise and interest in thermal design, Electronic Cooling Solutions Inc. (ECS) have been investigating tablet thermal design to understand first-hand, the challenges faced by vendors [2]. Hence ECS has undertaken a highly-comprehensive study of the Surface Pro's thermal design including acoustics, airflow and thermal performance.

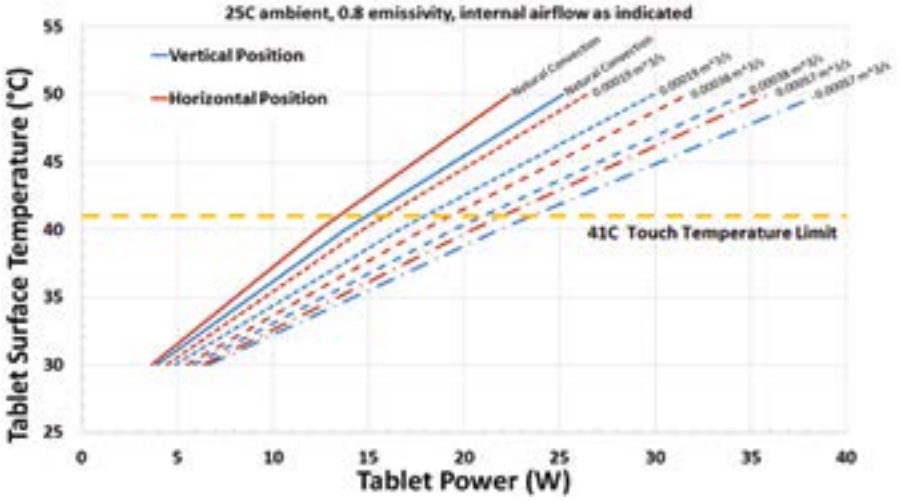

Figure 1. Tablet Power vs. Surface Temperature for 240 x 180 x 10 mm

The study started with an architectural-level thermal analysis to determine the maximum theoretical power dissipation limits for the given form factor, shown in Figure 1, which assumes perfect heat spreading efficiency. As tablets are hand-held devices with touch-sensitive displays, the surface temperature of the tablet is more critical compared to laptops. Guidelines for ergonomic touch temperatures are provided by Berhe [3] for handheld devices with plastic and aluminum surfaces. The form-factor of the device and estimated power dissipation are to be carefully decided at the initial design phase of the product.

All of the available surface area of a tablet needs to be utilised to maximise the heat dissipation. The emissivity of the outer surface is critical for radiative heat loss, which in natural convection tablets can be up to half of the total heat transfer to the ambient. The rear outer cover of the tablet is cast magnesium and electroplated, while the front outer cover is made of plastic.

The tablet was experimentally characterized for airflow and thermal performance. Airflow measurements were completed using an airflow measurement chamber by controlling the blowers at various speeds. Thermal performance was characterized by instrumenting the device with thermocouples internally and externally, and using infrared thermography. Great care was taken to ensure that dismantling and reassembling the device for internal instrumentation did not significantly alter its thermal performance by checking thermocouple and infrared measurements of case temperature before disassembly and after reassembly. The thermocouples pass through a small hole in the case which is sealed to prevent the ingress of air.

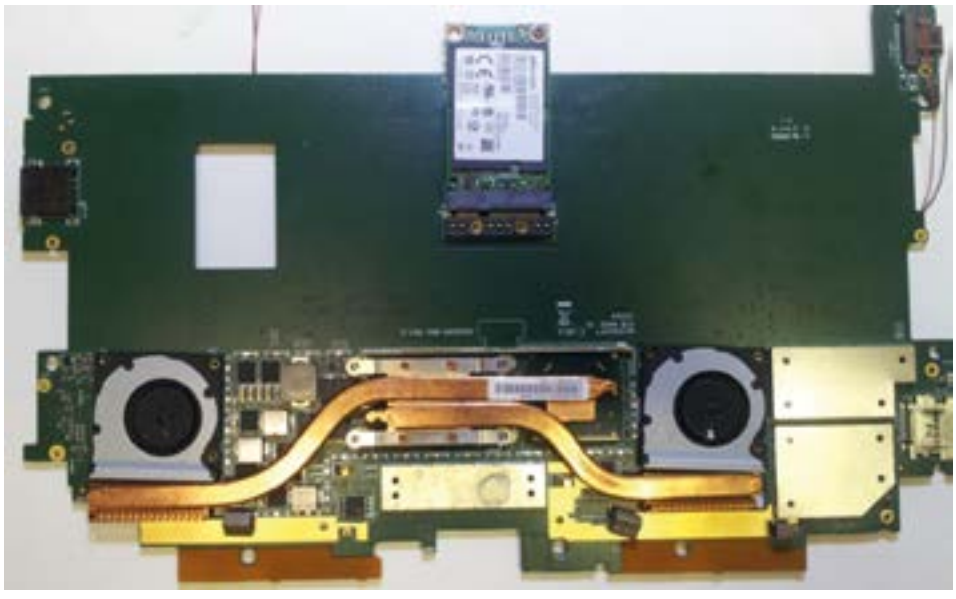

Figure 2. Processor and heat pipe assembly on Surface Pro main board

Thermal measurements were then made with the tablet fully charged and the input power measured while running the Prime95 Torture Test to exercise the CPU. Blowers are located on either side of the main printed circuit board and used to cool heat pipes that are connected to the CPU and GPU, as shown in Figure 2. The blowers have customized angled louvers that match with angled fins in the heat exchangers, which direct the exhaust airflow at about 45 degrees to the plane of the tablet to improve airflow efficiency and reduce aero-acoustic noise. By deflecting the exhaust airflow, entrainment of the exhaust air at inlet vents close to the fan is also minimized.

The tablet in normal operation regulates internal temperature by using a pulsed width modulation circuit to control the blower speed. To measure the overall airflow through

the tablet both blowers were controlled and powered using an external source to control the input voltage. The airflow through each outlet was measured individually at vents identified during disassembly.

Acoustics are extremely important in forced convection hand-held devices, restricting the permissible airflow and limiting the fan speed. Frequencies as well as the sound power level are important for human ergonomics and these parameters were recorded for several blower voltages in a lab that is well-known for its zero background noise. Sound pressure level as experienced by the operator was confirmed to be close to 40dBA.

Having measured the power consumption and the temperatures and dimensions of all critical components within the device, ECS engineers selected Mentor Graphics' FloTHERM® XT to build a thermal model of the entire tablet. Experimental data from the thermal and flow characterization was used to calibrate the model, allowing it to be used to study the heat spreading within the device, the influence of air gaps at various locations, and the impact thermal radiation has on heat transport within the system. The calibrated model in FloTHERM XT was to within 15% of the thermocouple data measured with the system vertical. Sources of this discrepancy include: power budget estimation; difficulty measuring system airflow at low flow rates; difficulty measuring key dimensions accurately, e.g. blower plenum size; and possible disturbances to the internal conduction path due to tear down and instrumentation. The model of the tablet is shown in Figure 3.

As part of the model calibration, ECS engineers used Mentor Graphics' transient thermal tester, T3Ster®, to reverse-engineer a thermal model of the CPU by deriving a Structure Function from the measured temperature vs. time data to reveal details of the thermal path from the junction out to the ambient. From this a Compact Thermal Model (2R model) of the CPU was created and incorporated into the CFD model. This was critical in getting the model to match the measurement data.

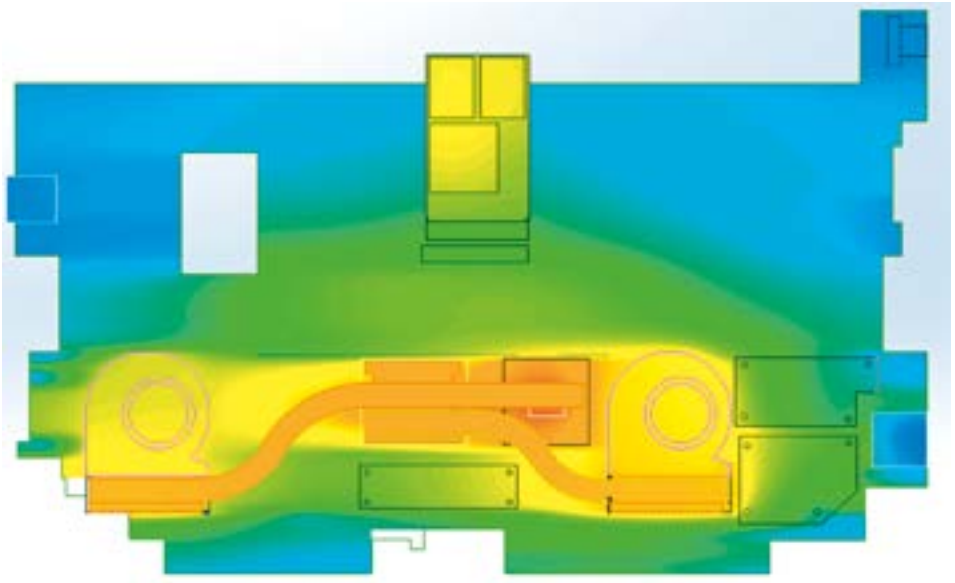

Figure 4. Modeled Surface Temperatures on the Main Board, Blowers, Heatpipes, CPU and GPU.

For mobile devices, where natural convection and thermal radiation are important to the overall cooling, maintaining a fairly uniform surface temperature results in maximum heat transfer to the surroundings. In this tablet under study, about 75% of the total heat generation is concentrated in the CPU and GPU, which occupy less than 1% of the device's area. The tablet also houses two large batteries that dissipate a negligible amount of heat, but need to be maintained at a fairly low temperature compared to the surrounding electronic circuitry. The heat generating components are concentrated in the top of the tablet, with the batteries on the bottom below the main PCB shown in Figure 2, under plastic insulation. The temperature distribution on the main PCB simulated in FloTHERM XT is shown in Figure 4.

The need to keep the tablet's surface close to isothermal, yet isolate components like batteries from the heat, plus the concentration of heat dissipation in a few high power components complicates the design. One key parameter is the thermal conductivity of the case itself. As tablets are partly preferred over laptops due to their light weight, this places a restriction on both the material and thickness that can be used for the base. Low acoustic noise also contributes to the popularity of tablets. Increasing the thermal conductivity of the case, and improving internal heat spreading can reduce the need for forced cooling, making it possible to reduce blower speed, thereby improving both battery life and acoustic performance.

ECS is continuing their work on this tablet, using the calibrated FloTHERM XT model to investigate the potential of using micro-vapor chambers, graphite heat spreaders and phase change materials to further enhance the thermal performance.

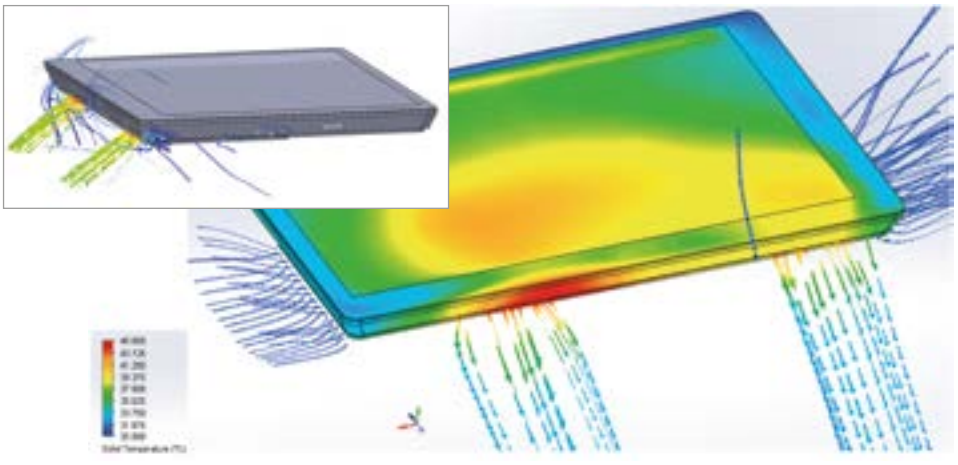

Figure 3. FloTHERM XT Model of the Surface Pro showing flow from the Side Vents.

References:

[1] https://en.wikipedia.org/wiki/Surface_Pro

[2] Wagner, G.R., Maltz, W., On the Thermal Management Challenges in Next Generation Handheld Devices, Proceedings of the ASME InterPACK 2013, Paper No. InterPACK2013-73237.

[3] Berhe, M.K., Ergonomic Temperature Limits for Handheld Electronic Devices, Proceedings of ASME InterPACK'07, Paper No. IPACK2007-33873.

Design Processes and Levels of Thermal Analysis

Design Processes and Levels of Thermal Analysis at United Automotive Electronics Systems Co. Ltd. (UAES)

By Boris Marovic, Product Market Manager, Mentor Graphics

At UAES, a Robert Bosch company, the thermal simulation of products takes an increasing role in the product development process. In a joint venture with Chinese Engine Management Systems Corporation Ltd (CNEMS), the product line includes a large range of automotive powertrain components such as Engine/Electronic Control Units (ECUs), Sensors, Ignition Coils, Inverters and Throttle Valves. The increasing demand in thermal management requires a recognition of the product development process and the introduction of an increased number of simulation steps, in order to ensure product functionality, performance, and reliability.

At UAES' Technical Product Development Center, the CAE Manager Bao Chenyu, and his team ensures the thermal performance of their products is within the limitation of the components and customers' specifications. The application of FloEFD™ and FloTHERM® increases the breadth of the simulation capabilities from early in the design stage, to the start of production (SOP). Both tools are used either separately or in combination, depending on the level of thermal analysis.

Starting at Level 1, UAES uses FloTHERM for any calculation of temperature increase in the chip design. Followed by Level 2, where both FloEFD and FloTHERM are applied depending on the depth of design stage, and need for an electronic cooling specific functionality. The calculation of chip configurations and improvements in cooling measures at level 2 can be very detailed and requires the consideration of the copper traces, vias and the correct thermal modeling of the chips.

Level of Thermal Analysis		Application Area
5	Vehicle Level	Vehicle Thermal Management System
4	Subsystem Level	E-Drive and Battery Thermal Management System
3	Component Level	ECU, PEU and Battery Thermal Analysis
2	PCB Level	Chip Configurations and Improved Cooling Measure
1	Chip Level	ECU, PEU etc. Chip Temperature Calculation

Figure 1. Levels of Thermal Analysis and Application Area

In Level 3, assemblies can often get more complex and require the inclusion of their housings and the environment they work in. Here the thermal analysis applications include ECUs, PEUs (Power Electronic Units), batteries and many more components.

Levels 4 and 5 are executed in system modeling tools.

Back at level 3, and looking more closely at the application in an ignition coil and PEU, we can see that the detail of the simulation

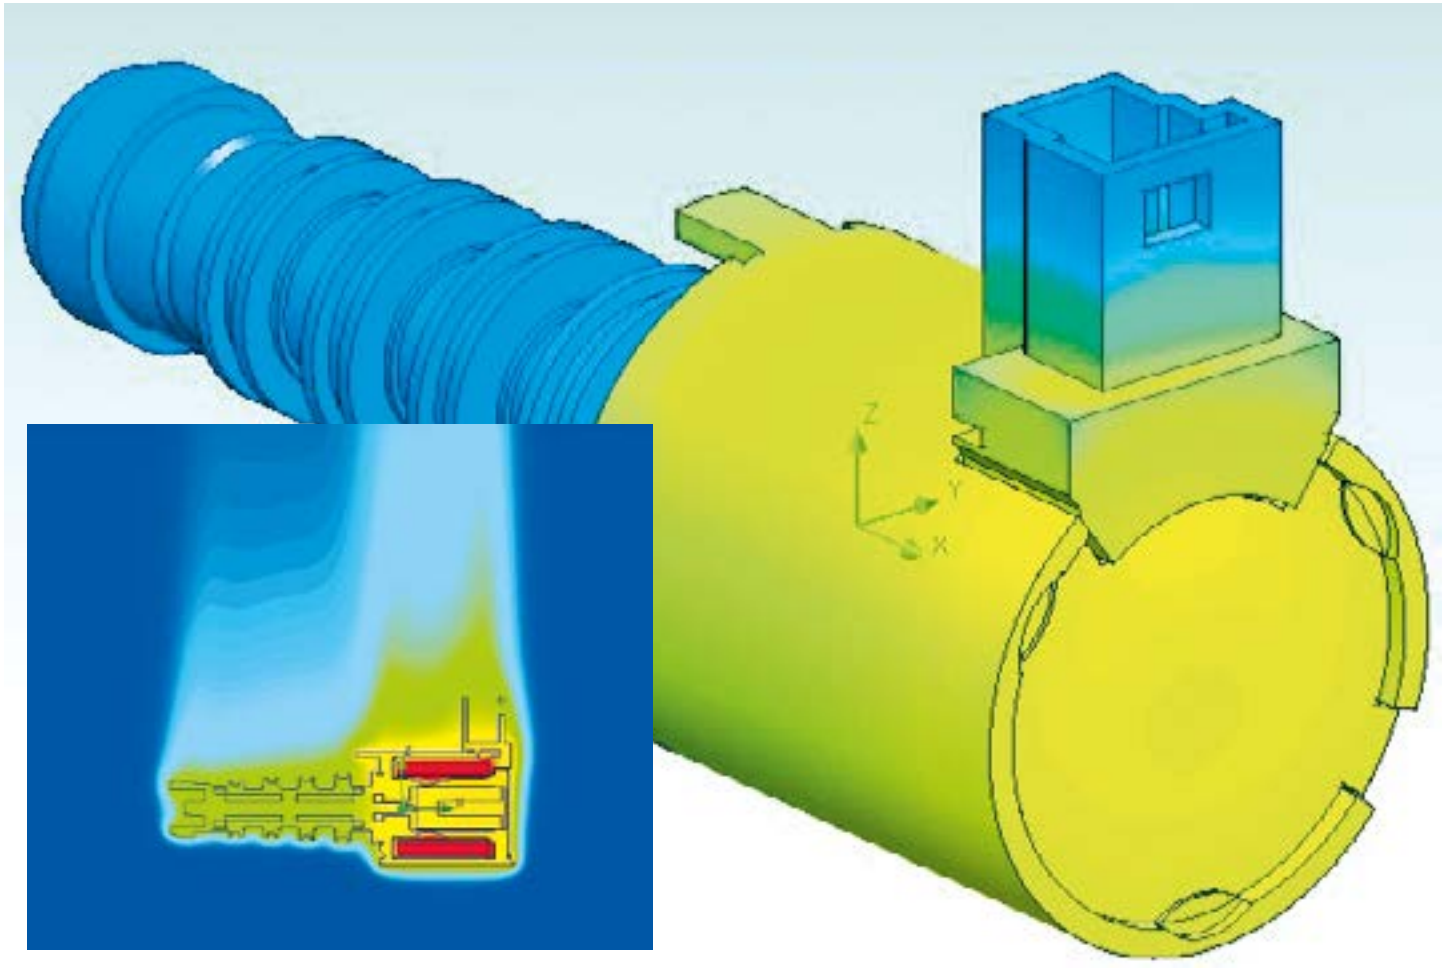

Figure 2. Stepper Motor thermal simulation

model is still high and includes a range of materials from plastic housings to electric connectors and copper coils. With losses applied and under natural convection conditions, these components have to endure the hot underhood environment. (Figures 2 & 3)

As can be seen in Figure 4 the largest proportion of heat loss is coil output, followed by primary coil and core losses. The secondary coil and module both introduce small amounts of heat. All of this has to be considered in order to represent reality as closely as possible.

To enable an improved performance it is necessary to include the arrangement of the semiconductors on the device, the liquid cooling loop in the cold plate and the fins/pins design.

As the geometry can become more complicated to perform, it requires accurate handling by the mesh for the CFD simulation. Often quick calculation tools are used to get a first idea of the concept design; these tools, as well as hand calculations are a source of information that cannot be ignored.

Based on these early concept designs, the first design is created and then analyzed, as

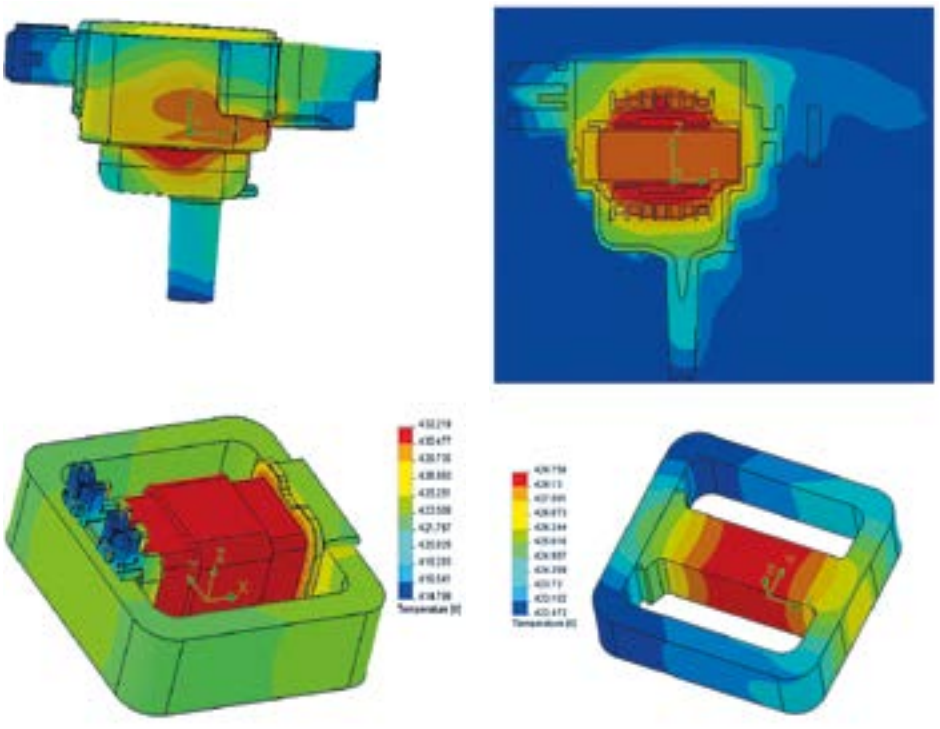

Figure 3. Thermal analysis of the ignition coil with its component losses

concept tools will often only provide a rough and simple idea of the cooling performance. To gain a full 3D flow with thermal behavior, the use of a 3D CFD tool should be employed. With all the tools working together, the next level of detail can be modeled and accurate simulation can be achieved with the necessary information.

It is only when the process and tools are optimized, with the best in class applications, that the efficiency and reduction of costs and time to market can show a significant advantage over the competition.

At UAES, as well as delivering high quality products to their customers, they have defined the process with the combination of tools, to enable them to stay ahead of the competition.

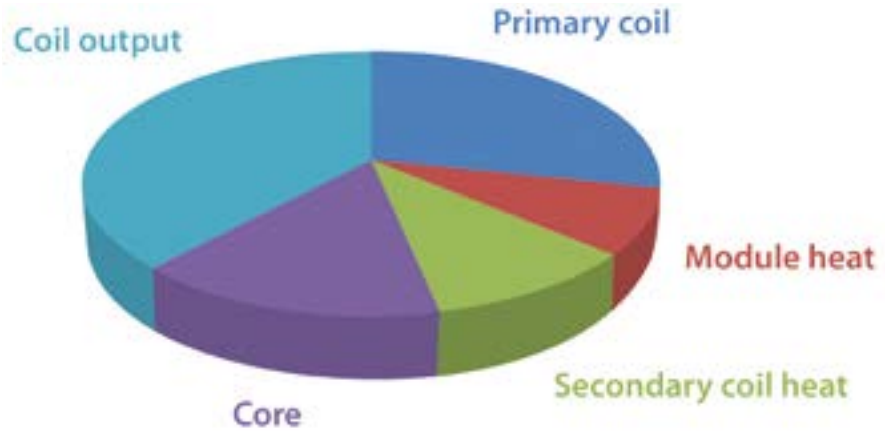

Figure 4. Proportion of Heat Loss

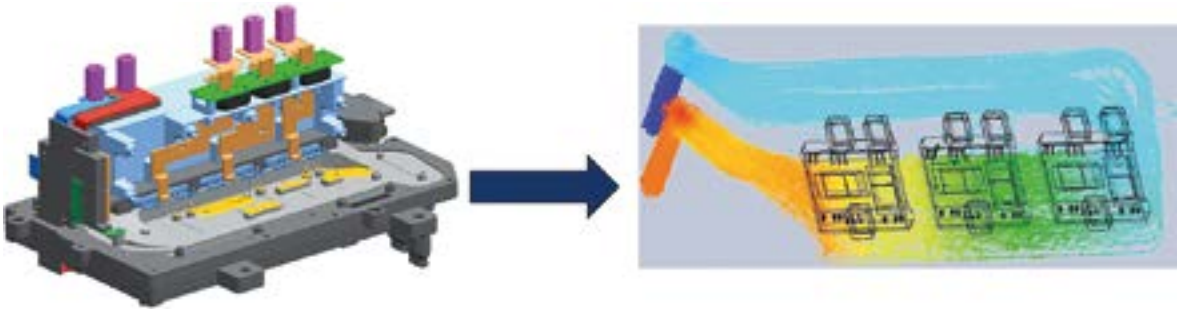

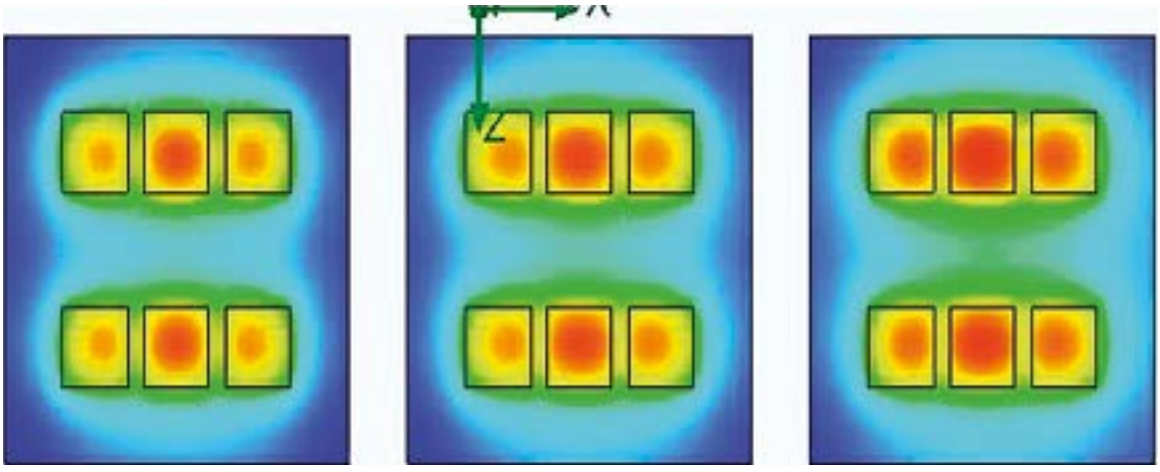

Figure 5. Thermal analysis of a PEU with a liquid cooled cold plate.

Design Guidelines for a Piezoelectric Micro-Blower Fansink

By Fukue Takashi, Thermal Engineering Laboratory (Hirose - Fukue Gr.)
Assistant Professor, National University Corporation Iwate University Faculty of Engineering, Department of Mechanical Systems Engineering

Forced air cooling using a fan is widely used to cool electronic equipment. In recent years, office automation equipment such as printers, and a variety of mobile devices have undergone massive improvements in multi-functionality and performance, coupled with miniaturization. As a result, the thermal design has become more challenging. There is no space to implement an additional fan for cooling. However, forced air cooling is needed to achieve the required cooling performance and prevent overheating. To address this, a number of small air cooling devices have been developed in recent years.

In this study, we investigated the practical use of a small air-cooling device for use within the narrow gaps found in densely packed electronics. The device is an ultra-compact piezoelectric micro-blower that is thin enough to be installed in the gaps between parts and develops a high enough pressure to generate sufficient flow rate. The blower is 20mm x 20mm and just 2mm thick, yet is able to supply 1L/min of air with a jet speed of almost 20m/s. This was investigated in conjunction with a heatsink, optimized using FloTHERM® to get the best thermal performance from the blower.

In recent years, the implementation of micro air movers in narrow spaces has been reported, including, micro fans [1] and piezo fans [2-4]. These devices have dimensions sufficiently small to be considered, but the resulting airflow would be minute, and perhaps adversely affected by the system air flow, limiting the benefits. This study focusses on the use of an ultra-thin piezoelectric micro-blower recently developed by Murata [5] to be 2mm or less (see Figure 1), to cool an attached heatsink. The purpose of the study is to develop design guidelines to maximize the heat transfer from the heatsink structure.

The piezoelectric element attached to the air chamber is vibrated by applying an alternating voltage (nominally 26kHz) to cause expansion and contraction of the air chamber. Air is drawn in during expansion, and forced out as a jet during contraction, entraining air from the flow passage. The study first looked at how

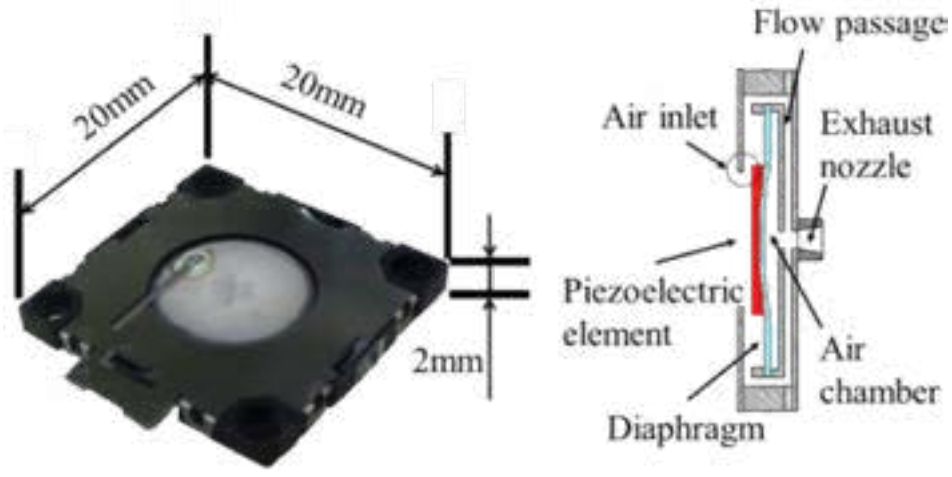

Figure 1. Murata's High-pressure Ultra-thin Microblower.

the flow performance was affected by installing a baffle plate to simulate nearby densely packed electronics. It was found, that provided the distance to the baffle is 1mm or more, there is very little restriction in the flow rate, pointing to the suitability of the device for use with narrow channels.

Having shown that the flow rate does not change even in the presence of very limited available flow space, the next stage was to investigate how the heat transfer enhancement resulting from the jet can be used to greatest advantage by numerical analysis using FloTHERM. This study considered a heatsink that had the same footprint area as the blower, to initially optimize the fin shape. For this a model was constructed as shown in Figure 2, to capture the design of future planned experimental work. The base of the heatsink was uniformly heated, with the base of the heatsink and the heater set into a 150mm x 150mm horizontal acrylic block. An identical acrylic block was set 5mm above, with the blower airflow modeled as a laminar fixed velocity directly above the center of the heatsink. The flow rate was set to 750 mL/min to make allowance for the system pressure drop, with an inlet air temperature of 20°C and a heat source of 1W applied uniformly within the heater.

To investigate the effectiveness of the blower, the cooling performance was also measured with the blower turned off for several heatsink

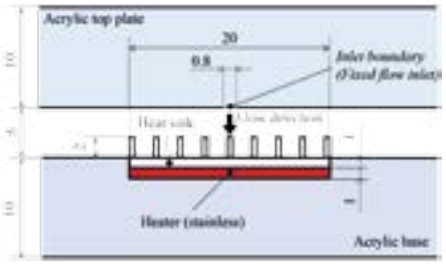

Figure 2. Dimensions (in mm) of the FloTHERM Model

designs, and the assembly cooled through a combination of natural convection, conduction and radiation. This caused an additional increase in the heatsink temperature rise by around 30°C in all cases, showing that although the flow through the blower is small, its effect on cooling the heatsink is large, providing confidence that the full study of heatsink geometries would be worthwhile.

Once the effectiveness of the blower was confirmed, attention turned to studying the influence of the fin shape on the performance of the heatsink. Heatsinks primarily extend the surface area available for cooling, so the hope was that the heat transfer could be increased by switching from an extruded fin heatsink to a pin fin heatsink, which were originally designed for use with impinging flows, and increasing the number of fins. By investigating 10 different heatsink designs in addition to the original extruded fin heatsink, it was found that an in-line arrangement of fins was superior to a staggered

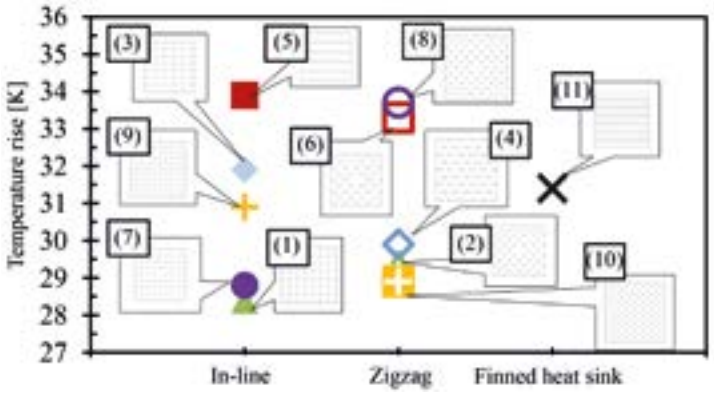

Figure 3. Performance of Different Heatsink Geometries

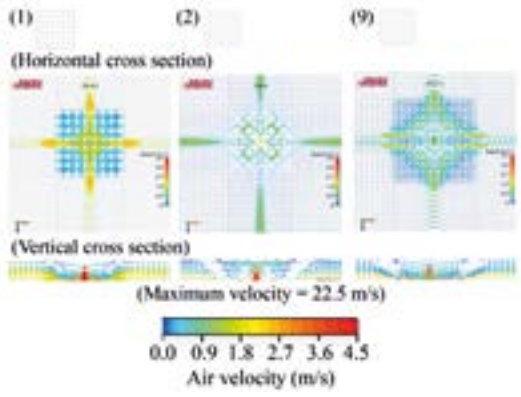

Figure 4. Differences in Flow Field Arising from Different Heatsink Designs

arrangement, as shown in Figure 3, in which heatsink #1 shows the lowest temperature rise above ambient.

It is worth noting that heatsink #9 with the finest fins, each having a cross-sectional dimension of 0.5mm x 0.5mm showed worse performance than heatsink #1 with 1.0mm x 1.0mm fins in the same in-line arrangement. For this reason the flow distribution within the heatsink was then investigated. For this part of the study, heatsinks 1, 2 and 9 were considered. Of these, heatsink #1 gave the highest rate of heat transfer and heatsink #9 the lowest.

From the flow vectors shown in Figure 4, it is evident that heatsink #1 has the highest velocity in the channels between the fins extending from the jet out to the sides of the heatsink, with the flow being ducted in those directions due to the alignment of the pins. The narrower channels in heatsink #9 increase the flow resistance and so act to reduce the flow velocity, causing the flow to spread more uniformly within the fin array. One key difference between heatsink #1 and heatsink #9 is that the latter has a row of pins across the base in line with the centerline of the jet, whereas heatsink #1 has a central gap. The staggered arrangement in heatsink #2 partially breaks up the jets, again reducing the flow velocity and leading to more uniform flow within the finned region.

From this, it was concluded that the main contribution to heat transfer is due to the boundary layer flow forming on the base of the heatsink, and the action of the fins to duct the flow and hence preserve its velocity provides a key to future heatsink designs. To further optimize the design it was decided to investigate how the fin gaps influence the cross flow. For this study the 3mm tip clearance above the heatsink fins shown in Figure 1 was reduced to zero by lowering the top acrylic down to the fin tips. The central fin gap size was varied from 0.5mm to 2.4mm, and the number of fins in each direction set to be six, eight, or ten, with the spacing between

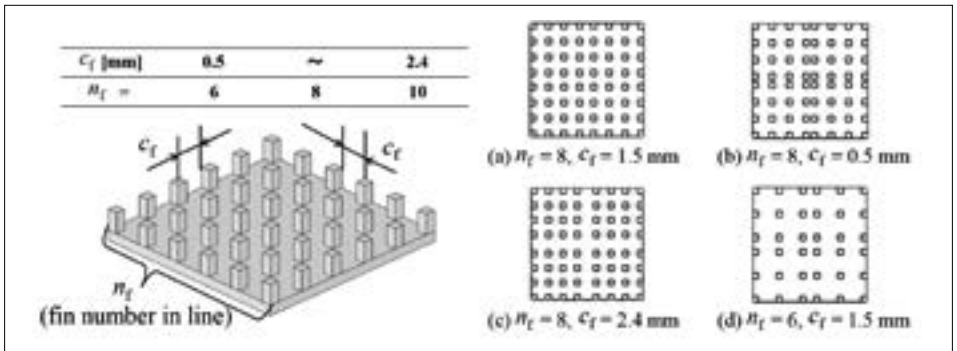

9 fin conditions for optimal fin spacing verification of the jet occurs site

the other fins changed as required. Samples of the designs are shown in Figure 5 along with the results, which show that the key parameter affecting the performance of the heatsink is the size of the central gap, with the results more weakly affected by the spacing between the other fins, with the larger fin spacing resulting from using only six pins in each direction showing the best performance.

By way of conclusion, this work has shown the viability of using a commercially available piezoelectric micro-blower with a customized heatsink design to cool densely packed electronics as found in the latest office automation products and mobile devices. Design guidelines for the heatsink have been developed to maximize the heat transfer from the heatsink by optimizing its design for the impinging flow. Further work is planned to experimentally verify the results of this study. There is also scope to optimize the shape of the fins and their layout beyond the rectangular cross section studied so far to further enhance the flow through the finned region by considering a radial arrangement with circular and elliptical fins.

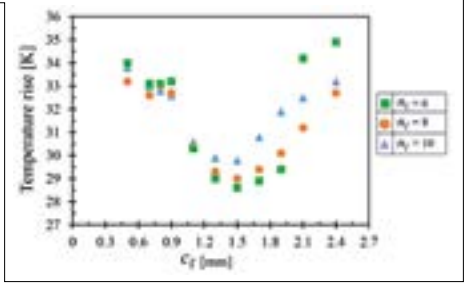

Figure 5. Influence of the Center Gap Spacing on Heatsink Cooling Performance

References

[1] Quin, D. and Grimes, R., "The Effect of Reynolds Number on Microaxial Flow Fan Performance", ASME Journal of Fluids Engineering, Vol. 130 (2008), Paper No., 101101.

[2] Acikalin, T., Wait, S. M., Garimella, S. V. and Raman, A., "Experimental Investigation of the Thermal Performance of Piezoelectric Fans", Heat Transfer Engineering, Vol. 25, No. 1 (2004), 4-14.

[3] Acikalin, T., Garimella, S. V., Raman, A. and Petroski, J., "Characterization and Optimization of the Thermal Performance of Miniature Piezoelectric Fans", International Journal of Heat and Fluid Flow, Vol. 28, No. 4 (2007), 806-820.

[4] Jalilvand, A., Mochizuki, M., Saito, Y., Kawahara, Y., Wuttijumnong, V. and Nguyen, T., "Thinner Thermal Solution Module by Combination of Thin Heat Pipe and Piezo Fan", Proceedings of the ASME InterPACK"11 (2011), Paper No., IPACK2011-52267.

[5] http://www.murata.com/products/micromechatronics/feature/microblower/index.html

Acknowledgement

The author wishes to acknowledge the assistance of Dr. Hirotoshi Terao and Mrs. Tomoko Wauke, ALPS Electric Co., Ltd., Japan

Reading Between the Thermal Lines

New Dielectric Materials for LED-Packages

By Hui Zhang, M.Sc., Dipl.-Phys. Max Wagner, and Prof. Dr.-Ing. habil. Tran Quoc Khanh, Laboratory of Lighting Technology, Technische Universität Darmstadt

FACHGEBIET LICHTTECHNIK

TECHNISCHE UNIVERSITÄT DARMSTADT

One of the main aims in electronic packaging is a good heat transport away from the device, downwards through the package into the board. The LED chip is connected electrically at its top and bottom. That is why there is a need to separate the electric circuit from the metal board. This is achieved by an insulator layer, which is neither electrically nor thermally very conductive. In this study the thermal behavior of new dielectric materials in LED packages are investigated. Furthermore the influence of geometric parameters of the electrical layout has been tested by measurements and simulations.

Figure 1 shows the structure of a high power LED based on a chip-on-board-package. The chip (also called die) not only produces light, but also a certain amount of heat, which should be as low as possible. The majority of the heat from the die will be transferred to the outside by why of conduction. Since the thermal conductivity of molding compounds (e.g. epoxy or silicone) is much smaller than that of the die attach (e.g. silver conductive adhesive), most of the heat will conduct downwards. Parameters such as the thermal conductivity and the geometrical structure of the interface materials have high influences on the heat flow. So a change would lead to different measurement results with the thermal impedance test system T3ster.

Figure 2 shows the structure function of an LED package after the evaluation of the measurement. The structure function presents all the thermal information of the tested LED-package, including

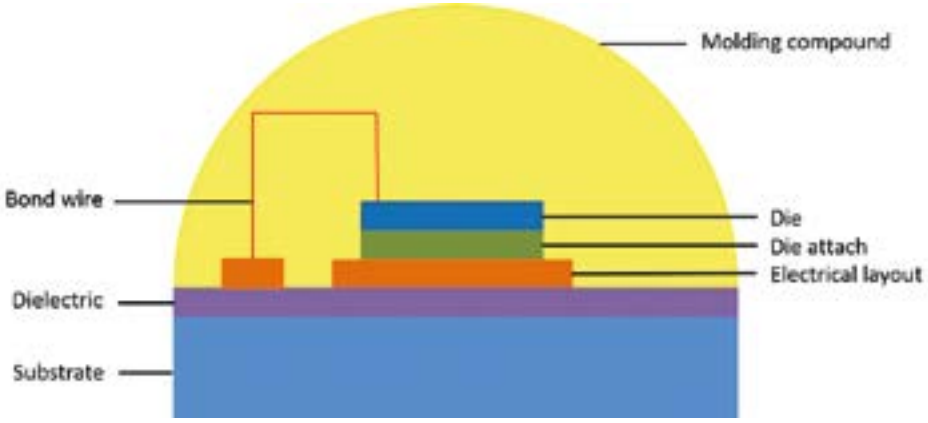

Figure 1. Structure of a high power LED-package

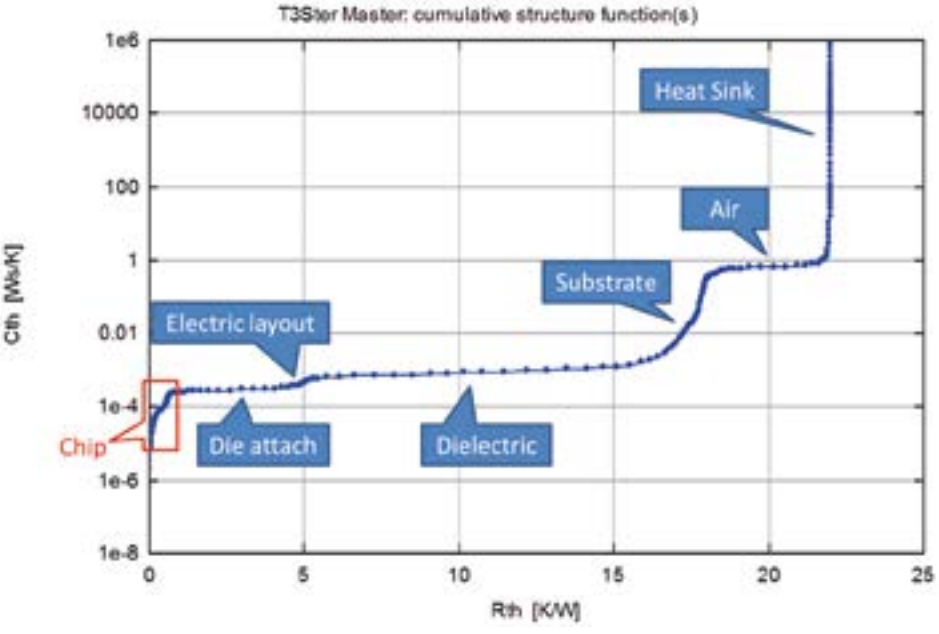

Figure 2. Structure function of an LED-Package

thermal resistance in K/W and thermal capacity in Ws/K. In fact every layer of the thermal capacity represents one kind of a material in the LED-package. The thermal resistances of different materials are of great interests to us, because the electrical components need to be cooled. Of course a lower thermal resistance means a better performance.

Electrical Layout

Considering the electrical connection between the chip and external electrodes, the electric layout plays a very important

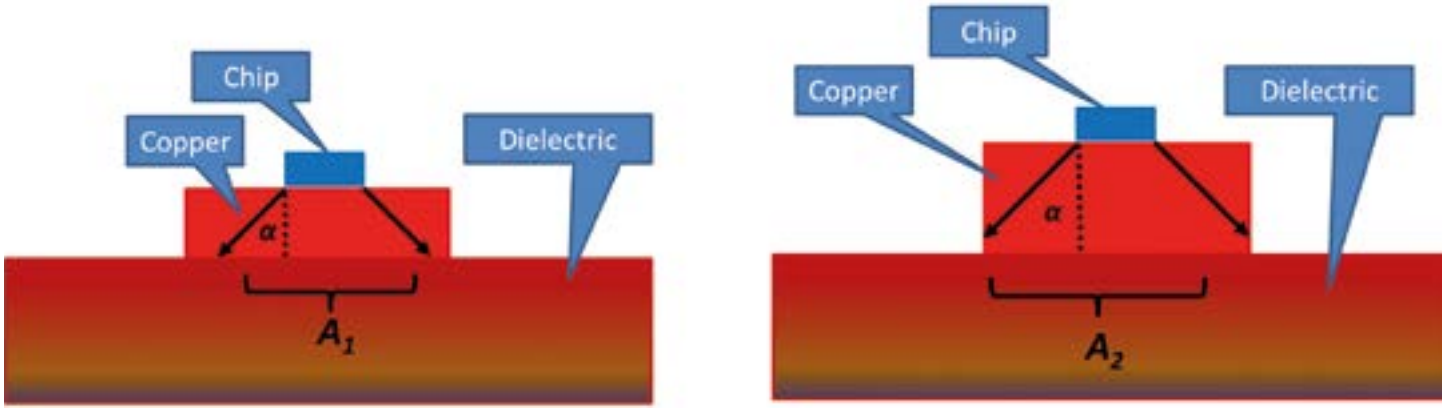

Figure 3. Heat spreading at different thicknesses of electrical layout

role in an LED package. In view of the electric conductivity and cost, copper is the first choice to the manufacturers. That is why different geometrical structures of the electric layout were built and analyzed in a thermal perspective. Beside the size of the surface, the thickness of the layout has been changed in values of 35, 70 and 105 µm. Copper has a high thermal conductivity of λ=385 W/mK. For a copper layout with a surface area of 5 mm x 5 mm and a thickness of 35 µm, the thermal resistance is $R_{th} = \frac{l}{\lambda \cdot A}$ = 0.0036 K/W.

This value is so small that a triplication of the thickness to 105 µm leads to only 0.01 K/W. So, why the concern over thickness? The answer lies in the effect of heat spreading, which is not represented by the formula above, but rather detected in measurement results. Figure 3 shows the approach of the effect by using a refraction model [1].

The die's heat flow reaches the copper layer and will be spread in this material. Compared to optics the refraction at the boundary surface is dependent on the refraction index, which is the conductivity in the thermal picture. The chip's junction to a high conductive material means a high refraction angle α, so that a high spreading effect should appear. By increasing the copper layout's thickness, the surface A, which is flown through by the heat, gets larger. Now, we can calculate the next thermal resistance with the simple formula again. The experimental results and transient thermal simulation with FloTHERM are shown in the following figures 4 and 5.

Both results confirm that the total thermal resistance with a thicker copper layout is lower. That is because heat flows into the dielectric with a higher surface area (Figure 3). So the dielectric itself has a

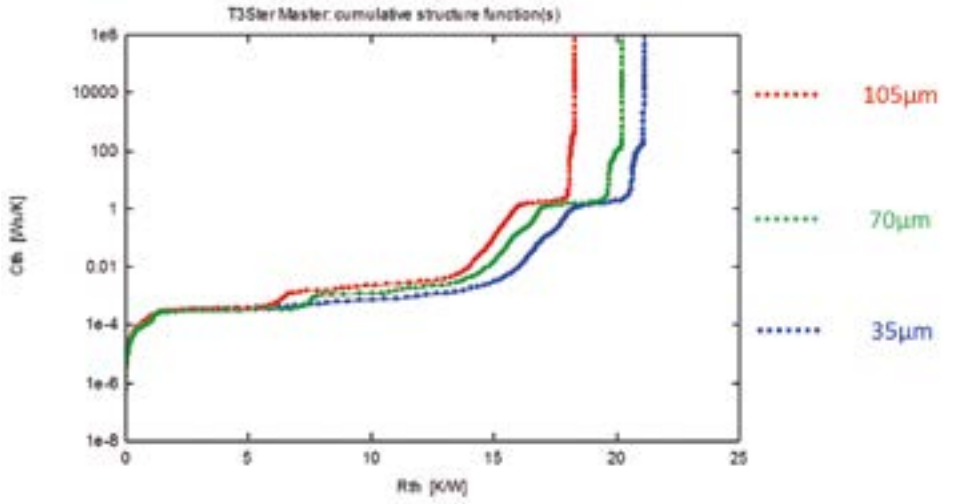

Figure 4. Structure functions (measurement data)

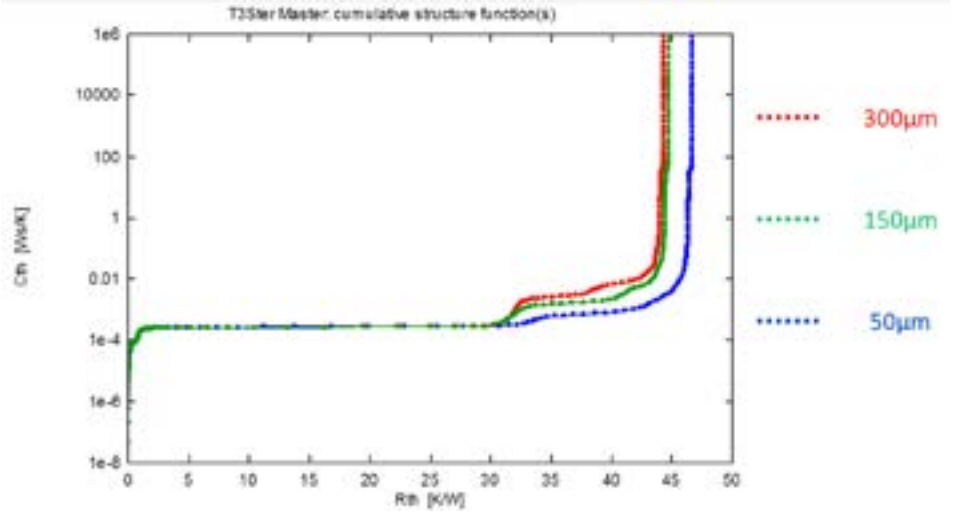

Figure 5. Simulation data (variation of the thickness)

smaller thermal resistance. Since a thicker thickness means more material, a higher thermal capacity should be observed. Since very thin layers mean a long simulation time, the thicknesses of glue and dielectric differ to the test devices. The trend of lower thermal resistance and higher capacity by enlarging the copper layout's thickness appears in measurement and simulation data.

New Dielectric Materials

All used dielectric are based on polymer and/or ceramic materials. The techniques to connect the dielectric layer with the board's substrate and electrical layout

differ. A standard method, is laminating the electrical layout on a ceramic filled polymer by an epoxy adhesive. Since the thermal conductivity of polymer is lower than that of ceramic, a new approach is to use a ceramic layer that consists of nano-crystalline aluminum oxide crystals (Al2O3) [2]. The electrical layout is laminated on this pure ceramic layer, or even contacted directly by a metallization process.

Samples with the same LED package, but different dielectrics are tested with the thermal impedance measurement. A direct comparison of Nanotherm LC and a ceramic paste (Figure 6) as well as Nanotherm DM and a ceramic filled polymer has been made. Figures 8 and 9 show the results of the measurement's structure functions.

Thermal resistances of dielectric and substrate build the main part of the total thermal resistance in the investigated LED-package. The use of a laminated nano-ceramic instead of a ceramic paste reduces the thermal resistance from 40 to 10 K/W (Figure 8). One reason for this is the small thickness of the ceramic layer (10 µm). The direct comparison of pure nano-ceramic with a laminated ceramic filled polymer of a standard PCB also shows a reduction of the dielectric's thermal resistance of about 33%.The conductivity is higher and the thickness of the direct metallized material is smaller, so the consequence must be a reduction in the total thermal resistance.

Conclusion

Not every transient behavior can be covered by a simple one-dimensional calculation. Heat spreading effects in the electrical layout directly affect the measurement curves and must be considered by other models. The change of the PCB's dielectric has the biggest effect on the total thermal resistance. Beside a standard ceramic filled polymer, new dielectric materials are tested. A direct comparison of laminated and direct metallized Nanotherm, and the aging behavior and reliability of both materials are interesting aspects that should be studied in the future.

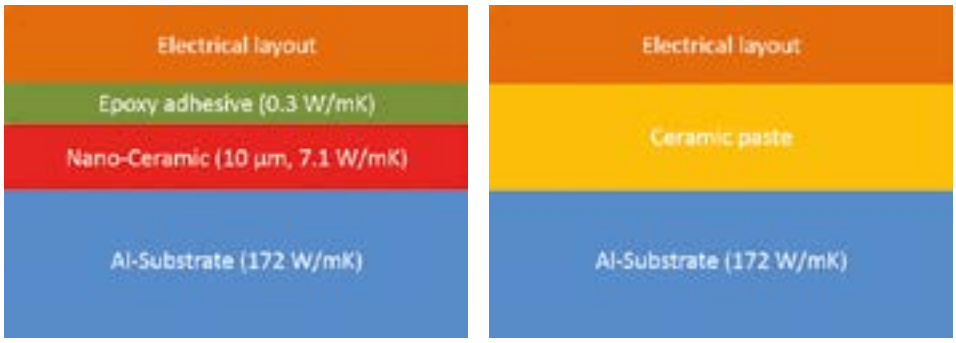

Figure 6. Comparison of Nanotherm LC (Laminated Circuit) and ceramic paste [3]

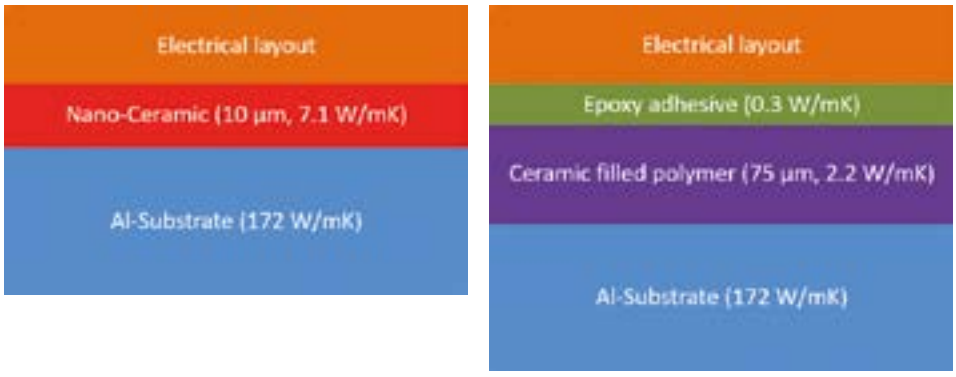

Figure 7. Comparison of Nanotherm DM (Direct metallization) and ceramic filled polymer [3]

References:

[1] David P. Kennedy. "Heat conduction in a homogeneous solid circular cylinder of isotropic media". Product Development Laboratory, Data Systems Div., International Business Machines Corp, Poughkeepsie, NY, 1959.

[2] http://www.camnano.com

[3] Data sheet from Excelitas Technologies and Cambridge Nanotherm.

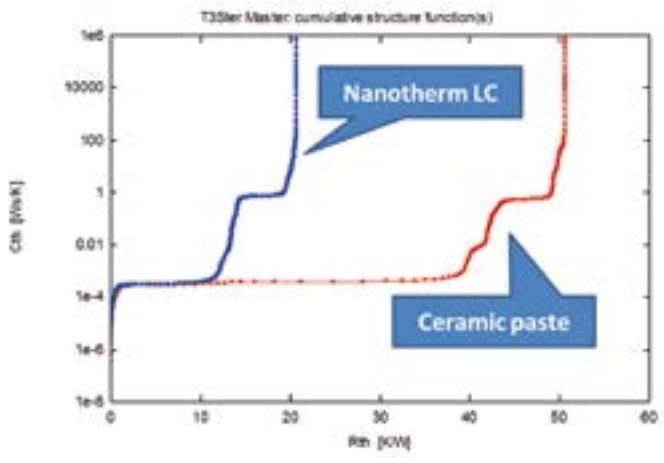

Figure 8. Structure functions of Nanotherm LC and a ceramic paste as dielectric material

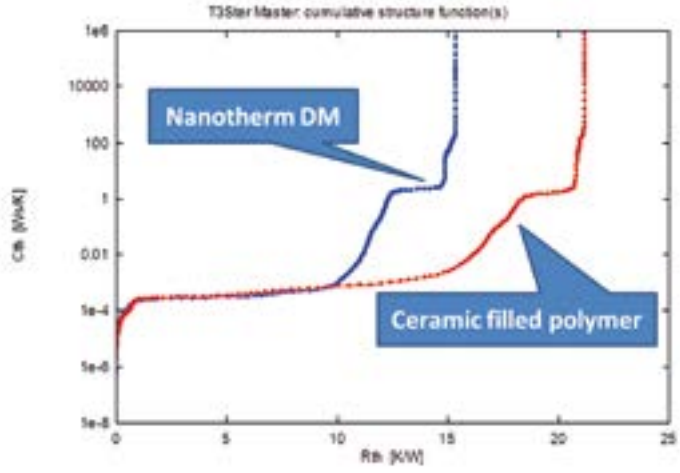

Figure 9. Structure functions of Nanotherm DM

Consumer Electronics Miniaturization:

Thermal Analysis of a Small Outline Package Mounted on a PCB Using Computational Fluid Dynamics

By Robert Day, Senior Application Engineer, Analog Devices;
and Prasad Tota, Application Engineer, Mentor Graphics Corp.

ANALOG DEVICES

The trend towards miniaturization in the consumer electronics industry has driven package component sizes down to the design-rule level of early technologies. Crucial in integrated circuit (IC) package technology is that it must deliver higher lead counts, reduced lead pitch, minimum footprint area, and significant volume reduction. As a result, this has led to semiconductor manufacturers developing the small outline package (SOP), surface-mount memory packaging.

SOP packages consume one-third to one-half of the volume of earlier packaging alternatives and are a logical choice for the small form factor of handheld electronics, portable communication devices, laptop and notebook PCs, disk drives, and other applications. Power SOP (PSOP) packages, when combined with a heat spreading thermal mass (copper slug), make the resulting dimensions an ideal good choice for office automation, industrial controls, networking, and consumer applications that generate internal heat and are exposed to stressful temperature conditions.

To simplify board layout PSOPs can be placed much closer together and to other components as they are designed with their leads located on the long side of the package, leaving two sides of the package open. The open sides of the package can be used to route traces under the component, conserving board layers.

Thermal power density increases when IC packages are downsized, driving the need for heat-transfer path from the die to the external ambient to be optimized to allow for maximum possible power dissipation at the die while ensuring the die temperature is under the maximum allowable value.

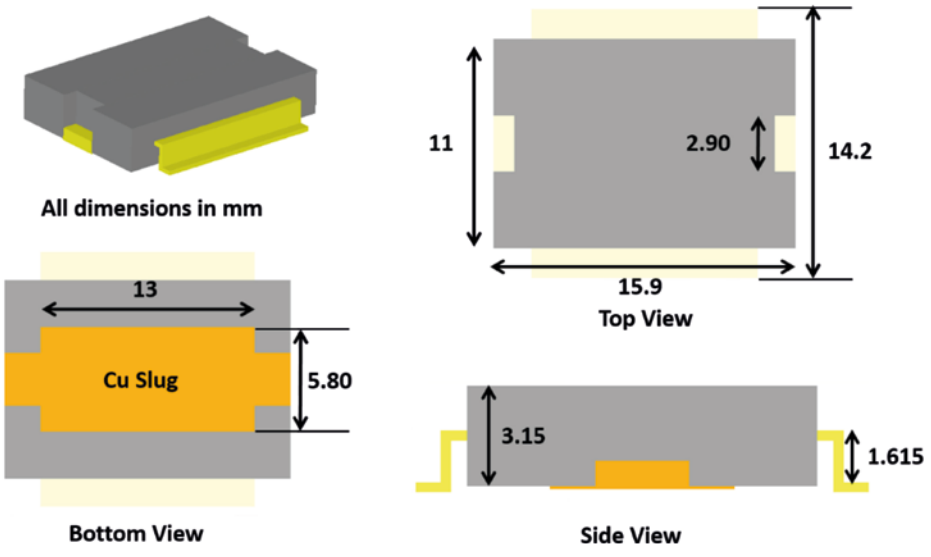

Figure 1. The PSOP dimensions in millimeters, with the copper slug on the bottom.

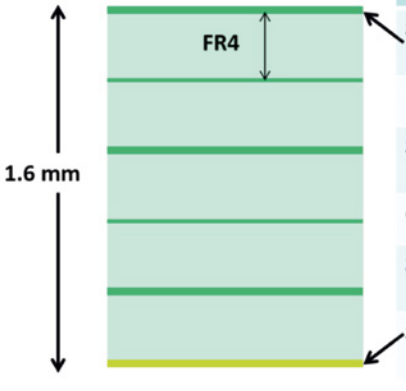

Layer	% Cu	t (mm)	k (W/mK)
Signal (Primary)	20	0.034	77.42
Power	90	0.017	346.53
Signal	20	0.034	77.42
Ground	90	0.017	346.53
Signal	20	0.034	77.42
Signal (Secondary)	100	0.034	385

Table 1. Board stack-up and percent of copper coverage.

PSOPs undergo tests for reliability under various stress conditions at the manufacturer, and it would be time-consuming and expensive to physically test or design test boards to test a package in all its possible applications and configurations. This is where Computational Fluid Dynamics (CFD) software is useful as it can simulate and estimate the junction temperature (T_j) of the IC when attached to the PCB under various conditions. FloTHERM from Mentor Graphics enables a mechanical or electrical

engineer and/or IC designer to quickly see the effect of design changes from a thermal management perspective both qualitatively and quantitatively.

Analog Devices used FloTHERM to perform a computational thermal analysis of a High Speed, High Voltage, 1A Output Drive Amplifier, the ADA4870-1 PSOP mounted on a PCB [1]. Specifically, the goal was to identify the maximum power that could be dissipated on the die active area while keeping the T_J at less than 150°C. Analog Devices studied various environments to estimate this maximum power, for example, changing the board area, adding thermal vias, and attaching a heatsink.

Depending on the direction of the formed leads, the package can be surface-mounted on the board either slug down or slug-up, (Figure 2). In a slug-down configuration, the component is surface-mounted on the primary side of the board where the copper slug is soldered to the top side of the board. In a slug-up configuration, the leads are soldered to the primary side of the board. For the experiment, Analog Devices used a slug-down configuration; first with no heatsink, and then with a heatsink attached to the secondary side of the board with thermal grease between the board and the heatsink base.

For the CFD simulation, the test board used was a six-layer board, with dimensions of 59 x 61 mm with the assumption that the copper coverage for each of the conducting layers was smeared uniformly within the layer's volume. Based on this, the thermal conductivity (k) of each layer was calculated as a volume average based on the percent of copper coverage within an individual layer (Table 1).

To accurately predict the value of the junction temperature, it is recommended to discretely model each of the conducting layers with orthotropic conductivity for the entire thickness of the board. Modeling the layers discretely, rather than with a lumped model, captures the effect of heat spreading within the board more accurately for various heat-transfer paths.

Thermal Simulation without a Heatsink

The first set of simulations were conducted to study the thermal behavior of the PSOP mounted on the primary side of the board where the copper slug was soldered to the board, keeping the board horizontal with respect to gravity in an ambient temperature of 85°C.

To emulate real working conditions, heat was applied to two-thirds of the top of the die.

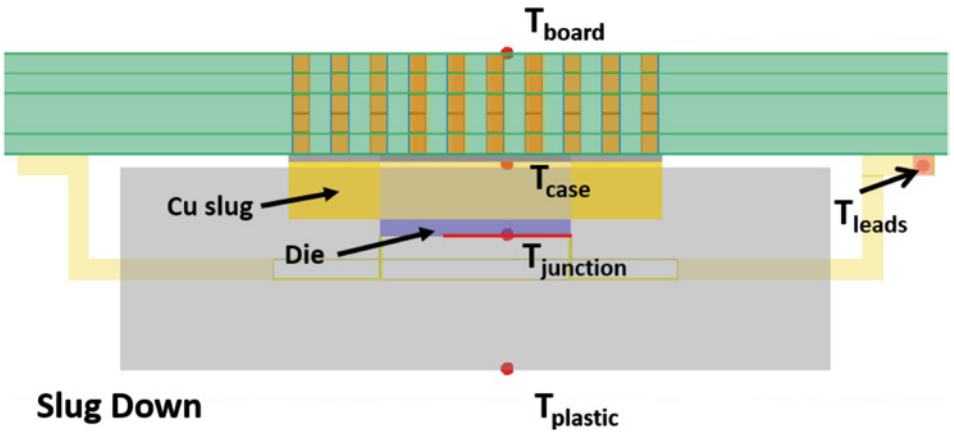

Figure 2. Temperature measurement locations.

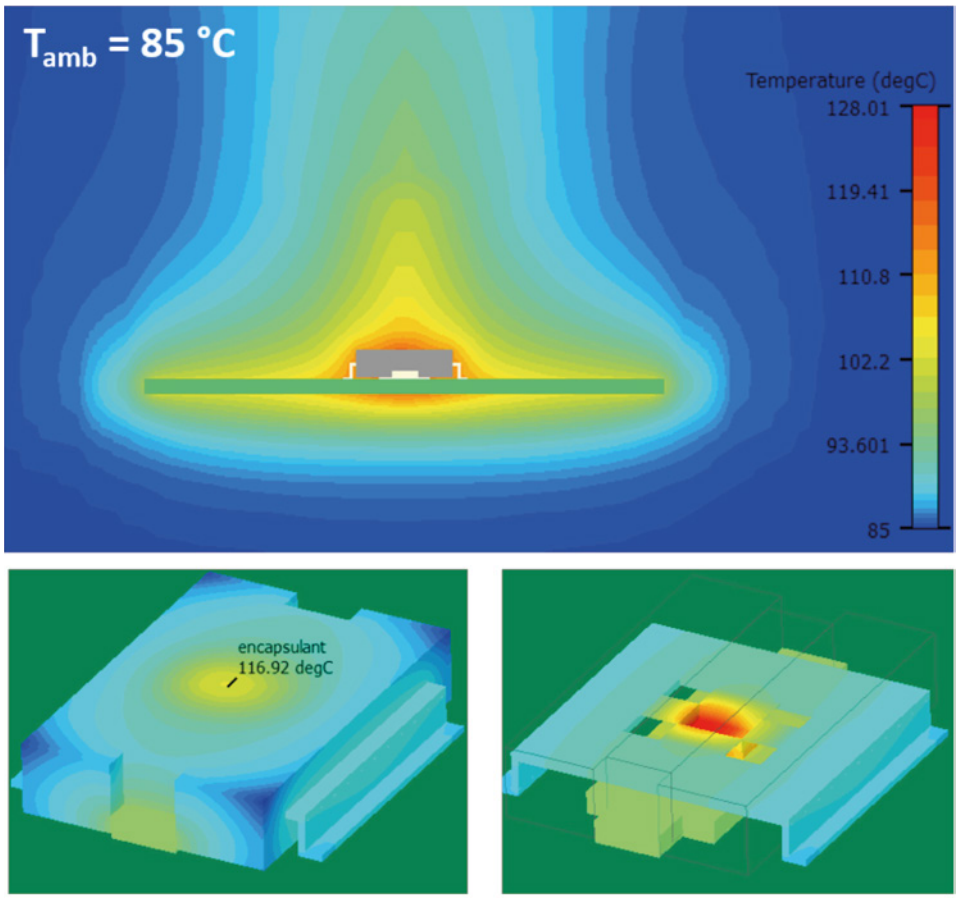

Figure 3. Temperature plots for the package in still air at 85 °C.

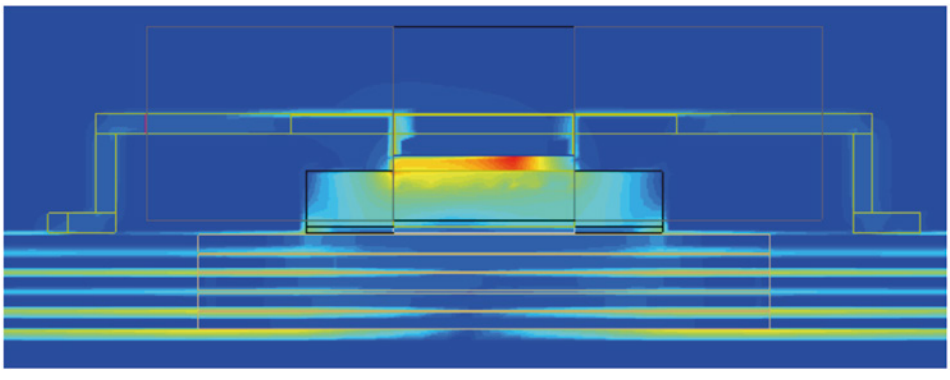

Figure 4. Heat-flux plots for a plane cutting through the package.

The junction temperature (T_j) was measured in the simulation at the geometric centroid of this area, and case temperature (T_c) was measured at a point in the copper slug just above the soldered interface (Figure 2). It is also possible to monitor the temperature of the leads, plastic surface, or any given position to validate the computational results with available test data.

Thermal vias were added under the slug to provide a more conductive path from the copper slug into the board. The vias were placed directly under the copper slug as the numerical investigations revealed a small advantage of adding vias beyond the slug area. This also helps lower board manufacturing costs.

Two possible scenarios for thermal vias were investigated where:

1. Inner layers were isolated; and
2. Inner layers were stitched together.

Stitching the inner layers lowers the junction temperature as a fraction of the heat entering the slug can spread in inner layers; however, including the inner layers raises the core body temperature of the board. Depending on the application, the inner layers could be isolated or used for thermal management. In this study, the secondary side of the board was completely covered with copper.

Figure 3 shows the temperature plots for the package in still air at 85°C and thermal power P = 2W with die-attach material of k = 1.6 W/mK [watts per meter kelvin]). The die-attach was replaced with a more conductive material, k = 50 W/mK, which significantly reduced the junction-to-case thermal resistance (θ_{jc}) of the package from 6.61°C/W (celsius per Watt) to 1.12°C/W.

Thermal Simulation with a Heatsink

A heatsink was soldered to the back side of the board to increase the power dissipation through the package, using thermal grease between the board and heatsink. Adding the heatsink significantly reduced the junction-to-ambient thermal resistance (θ_{ja}) from 16°C/W to 5.73°C/W. Heat-flux plots for a plane cutting through the package show the heat spreading over a larger surface area hence reducing the junction temperature for a given value of thermal power (Figure 4).

Table 2 shows the results for maximum power (Pmax) allowed in the slug-down configuration in still air with and without a heatsink for the two die-attach materials.

Slug-Down Configuration: Still Air at 85 °C				
	Die Attach	**θ_{jc} (C/W)**	**θ_{ja} (C/W)**	**P_{max}**
Without heatsink	Ablebond	6.61	21	3.11
Without heatsink	Cookson	1.12	15.95	4.10
With heatsink	Ablebond	6.78	10.63	6.11
With heatsink	Cookson	1.11	5.73	11.34

Table 2. Thermal resistance for different die-attach materials.

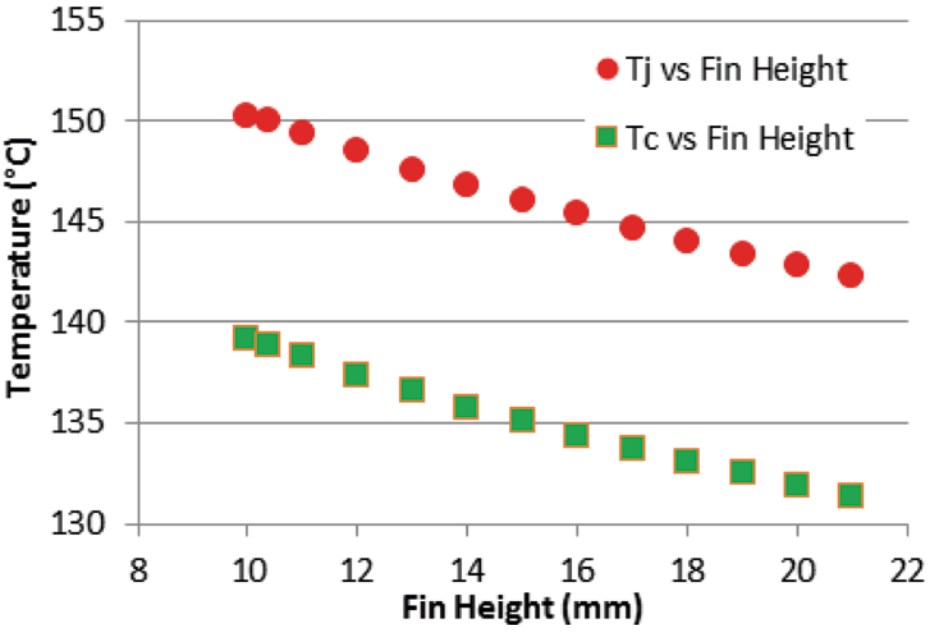

Figure 5. Junction temperature (T_j) and case temperature (T_c) for different heatsink fin heights.

Using the results, the focus of the next study was to use a more conductive die-attach material (Cookson) to find the shortest heatsink sufficient to dissipate 10W of heat at the die. FloTHERM's parametric study capability enabled the team to quickly set up and solve for different scenarios [3]. The variable parameter in this case was the heatsink fin height. The results in Figure 5 show junction temperature (T_j) represented by circles and case temperature (T_c) by squares. It was found that a heatsink with fin height of 10.36mm is sufficient to dissipate 10W.

A further investigation to find Pmax that could be dissipated if there were tighter constraints on the size of board and heatsink was conducted, thereby reducing the size of both to 30 x 30mm. As well as this the team also studied the effect of different fin heights on junction-to-ambient thermal resistance, θ_{ja} (Table 3).

With forced airflow, the junction-to-ambient thermal resistance could be further reduced,

Board and Heatsink Base: 30 x 30 mm		
Fin Height (mm)	**θ_{ja} (C/W)**	**P_{max} (W)**
21	11.82	5.50
15	12.98	5.01
10	14.48	4.49
5	17.12	3.80

Table 3. Thermal resistance vs fin height in still-air environment.

allowing higher powers to be dissipated and T_j to be kept under 150°C. Figure 6 shows the package simulation in a forced-air environment. Table 4 shows the results for heatsink optimization in forced air. Note that, with forced airflow of 2 m/s, the package could dissipate over 20W of heat for a fin height of 21mm and 17W with fins just 10mm high.

A similar parametric study was done for the smaller heatsink with a base of 30x30mm for different fin heights in forced air (Table 5). The smaller heatsink with 10mm high fins (lighter weight) offered the same performance as a larger heatsink with 5mm fin height.

Several parameters affect the thermal conductivity of the board in the region of the vias [4]. Creating a test board for every possible thermal via configuration and testing in a lab is practically infeasible. FloTHERM can be used to perform sensitivity studies of thermal performance to various via parameters, such as the pitch, plating thickness, and fill material (Figure 6). Such computational studies reduce the number of prototypes needed for testing or validation.

In a CFD program, it is computationally intensive to model each and every via discretely, so a lumped approach was used, the region of vias was replaced with a block of orthotropic conductivity that had in-plane conductivity (k_{xy}) and through-plane conductivity (k_z). The board-import tool in FloTHERM was used to calculate the k_{xy} and k_z of this via block, but values could have been calculated analytically [2, 5].

Thermal vias with an outer diameter of 0.3mm were studied. Figure 7 shows the sensitivity of thermal conductivity of via block to pitch and plating thickness (t). The dielectric material used in this calculation was FR4 (k = 0.3 W/mK), and the fill material was pure copper (k = 385 W/mK).

Thermal simulations were conducted for PSOP in still air, based on the conductivity values of the via cuboid (Figure 8). The results show that when plating thickness t is 75µm or higher, even sparsely populated vias are sufficient. However, at low plating thickness, 25µm or lower, the vias need to be populated densely to ensure the component does not experience thermal failure.

Validating Simulation Results

Laboratory experiments were conducted to validate the CFD model results. The IC inside the PSOP package is capable of dissipating 10 Watts of power and has an integrated temperature monitor. The relationship of the voltage at monitor-to-die temperature is not an absolute temperature indicator. However, the change in voltage versus temperature is a reliable indicator of relative changes in die temperature. Calibrating the temperature-monitor voltage verses temperature function was the first step in understanding die temperature used to determine thermal resistance.

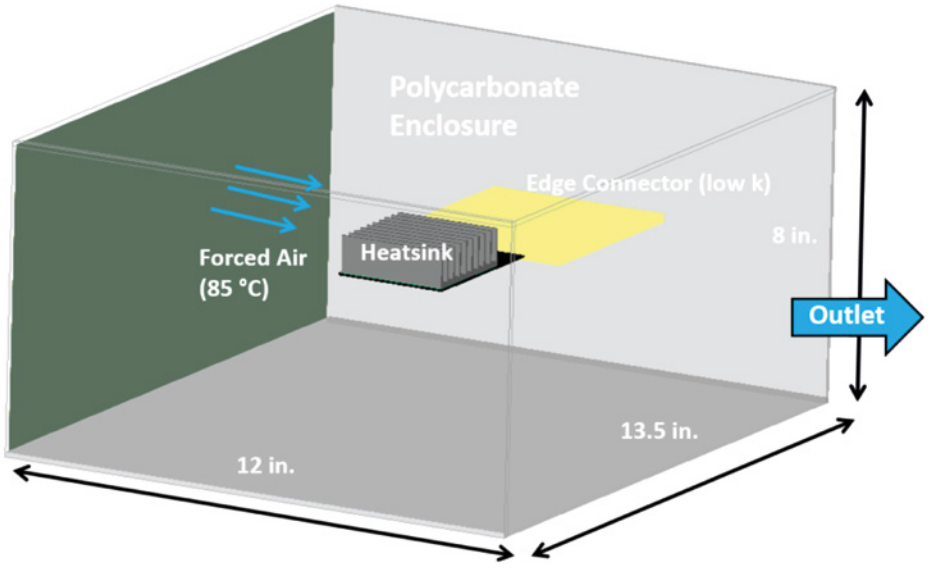

Figure 6. Package with heatsink in a forced-air environment

Forced Air, Heatsink Base 61 x 59 mm				
	1 m/s		2 m/s	
	θ_{ja}	P_{max} (W)	θ_{ja}	P_{max} (W)
21 mm	3.59	18.1	3.18	20.4
15 mm	3.95	16.5	3.42	19.0
10 mm	4.46	14.6	3.8	17.1
5 mm	5.36	12.1	4.49	14.5

Table 4. Thermal resistance versus fin height in forced air. θ_{ja}: junction-to-ambient thermal resistance, P_{max}: maximum power.

Forced Air, Heatsink Base 30 x 30 mm				
	1 m/s		2 m/s	
	θ_{ja}	P_{max} (W)	θ_{ja}	P_{max} (W)
21 mm	4.4	14.8	3.62	18.0
15 mm	4.85	13.4	3.95	16.5
10 mm	4.46	11.9	4.42	14.7
5 mm	6.48	10.0	5.3	12.3

Table 5. Thermal resistance and maximum power for forced air. θ_{ja}: junction-to-ambient thermal resistance, P_{max}: maximum power.

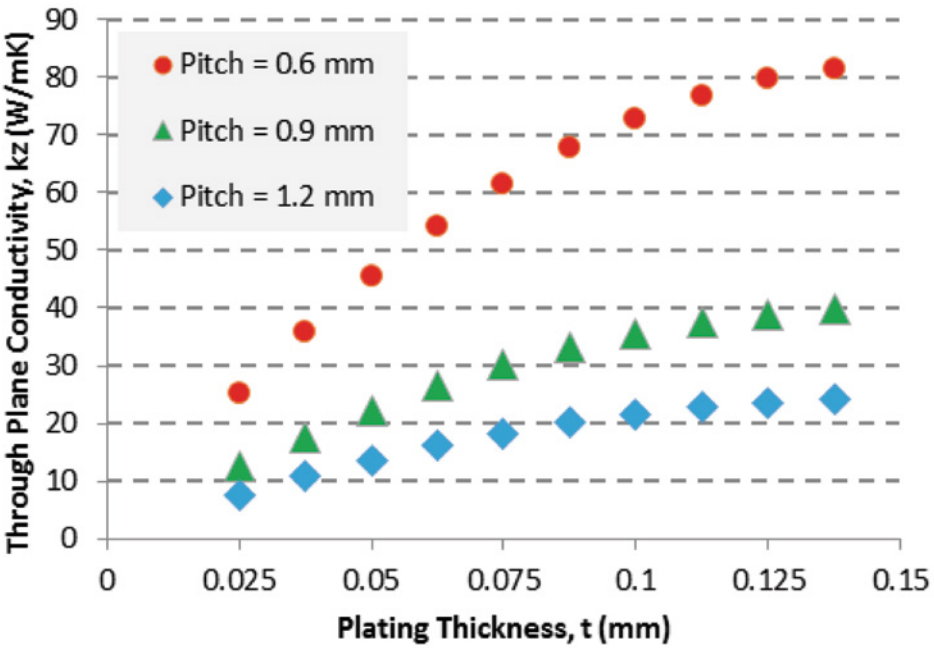

Figure 7. Sensitivity to via pitch and plating thickness. k_z: in-plane conductivity

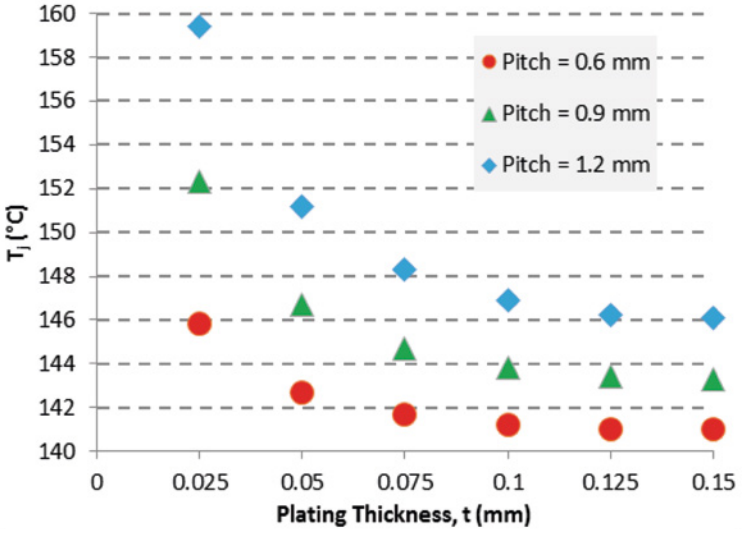

Figure 8. Junction-to-ambient thermal resistance (θ_{ja}) to via pitch and plating thickness in still air.

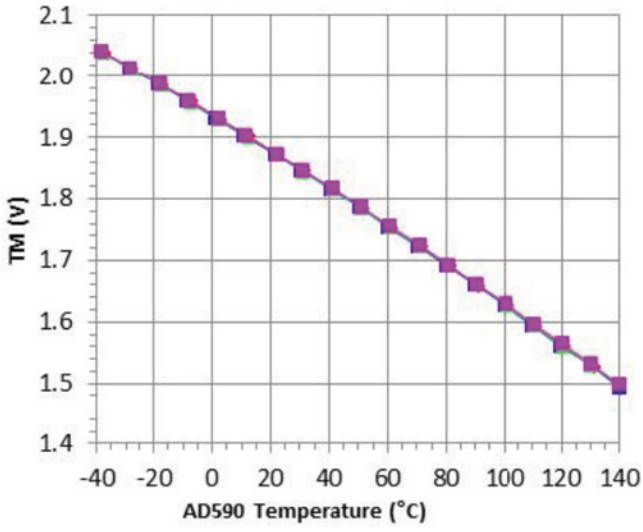

Figure 9. Temperature monitor (TM) volts versus sensor temperature.

The PCB used in the lab was FR4-grade with six layers of copper and exposed copper planes, onto which the ADA4870-1 PSOP package was soldered and heatsinks were mounted. Copper-filled thermal vias were used to conduct heat from the IC side to the bottom of the board where a precise temperature sensor was soldered directly below the thermal slug of the PSOP package onto the back side of the PCB. A heatsink was bolted to the back side that straddled the sensor using silicon grease as a thermal interface material between the heatsink and the PCB.

The PSOP assembly was placed into a still-air chamber using automated instruments and power supplies and allowed to soak overnight without any power applied. The ADA4870-1 IC and temperature sensor were then both turned on and measurements of the PSOP temperature-monitor voltage and sensor-trimmed PTAT (power sub-threshold proportional to absolute temperature) current were made immediately. The temperature-monitor voltage measurement was related to the absolute temperature indicated by the temperature sensor. This process was repeated at several temperatures to develop a calibration of the temperature-monitor voltage to absolute temperature (Figure 9).

Using a linear fit to the curve (T [°C] = TM [V] – 1.93/0.003), the voltage was converted to temperature. Additional steady-state tests were done to reveal the practical limits of power dissipation (maximum power) as a function of the applied heatsink. As shown in Table 6, large heatsinks are necessary when operating at the limits of power dissipation for the tested IC. It was calculated the junction-to-ambient thermal resistance (θ_{ja}) from the measured data by the following relationships at steady state: $\theta_{ja} = \Delta$ TM (V) – 1.93 (V) – 0.003 V/°C Δ Power (W) = °C/W.

The results showed the FloTHERM CFD simulation to be in good agreement with the lab test results with a heatsink mounted, where the dominant heat-transfer path is from the die into the heatsink. There is a higher difference for simulations with no heatsink, where an appreciable fraction of the total heat travels through bond wires and leads into the top layer of the PCB. This difference can be attributed to assumptions in simulation made in modeling the leads and bond wires in the simulation.

Conclusion

With these experiments, Analog Devices found that FloTHERM is a complimentary tool to laboratory testing, enabling quick parametric and design optimization studies in the thermal design. Such data is useful for studying electronics in harsh environments with increasing demands on power. The next step would be to analyze the transient behavior of the package and thermal characterization using structure functions generated by hardware testing, such as the Mentor Graphics T3Ster. A transient thermal simulation validated by test data would go a long way in simulating the transient response of a package for various powering conditions and reduce the number of laboratory tests needed.

Package Mounted in Slug-Down Configuration				
	Test Data		CFD Data	
Test Case	θ_{ja}	P_{max} (W)	θ_{ja}	P_{max} (W)
25 °C no heatsink	12	10.42	16	7.81
25 °C w/ VHS-45	7	17.86	8.87	14.1
85 °C no heatsink	12	5.33	16	4.1
85 °C w/ VHS-45	7	9.14	7.81	8.35
85 °C w/ VHS-95	6.2	10	5.73	11.34

Table 6. Thermal testing versus simulation results. θ_{ja}: junction-to-ambient thermal resistance, P_{max}: maximum power.

References

[1] Analog Devices, High Speed, High Voltage, 1.A Output Drive Amplifier ADA4870, http://www.analog.com/media/en/technical-documentation/data-sheets/ADA4870.pdf

[2] Bornoff, Robin, Blackmore, Byron, Parry, John, "Heatsink Design Optimization using the Thermal Bottleneck Concept," Proceedings of 28th IEEE SEMI-THERM Symposium, San Jose, CA, March 2011, pp.76-80.

[3] Li R.S., "Optimization of thermal via design parameters based on an analytical thermal resistance model," Thermal and Thermomechanical Phenomena in Electronic Systems, 1998. ITHERM 1998, pp 475-480.

[4] Incropera, F., Dewitt, D., et al., Fundamentals of Heat and Mass Transfer, John Wiley and Sons (New York, 1993), pp. 65-67.

A Lesson Learned

Rockwell Collins Improve simulation processes for Commercial Aircraft Avionics

By Mike Croegaert, Industry Vertical Manager, Mentor Graphics

Rockwell Collins

Rockwell Collins is a leading manufacturer of aircraft avionics systems for both commercial and military markets. They have a staff of highly experienced thermal analysts that utilize FloTHERM® Electronics Thermal Analysis Software for upfront simulation to predict the thermal performance of these products early in the design process and make design decisions around thermal management. Some of the analysts have over 20 years' experience using FloTHERM, so when for a particular product, the results of thermal testing were significantly different than the results of their analysis, there was a great deal of surprise. Even after updating the FloTHERM model to better match the final design, the results still did not correlate in a non-conservative way to the test data to one key test scenario. This caused them to kick off a lessons learned exercise to better understand what was causing the discrepancies.

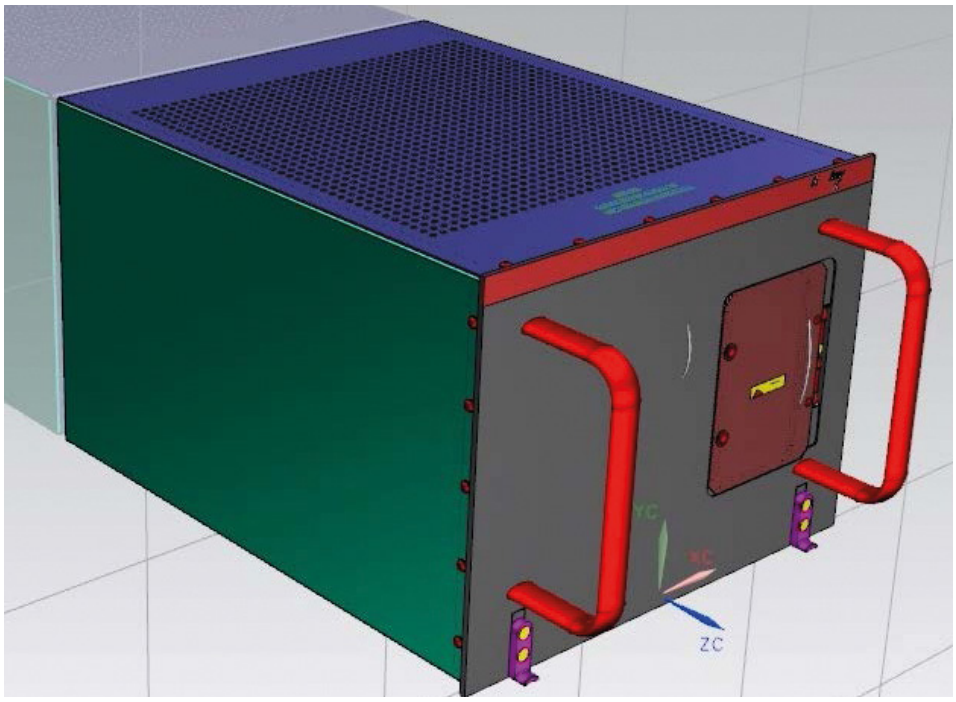

Figure 1. Chassis Model Mechanical Overview

The product in question is the data processing element of a cockpit display system for a new, large commercial aircraft. The product is forced-air cooled; designed to meet Aeronautical Radio, Incorporated (ARINC) Standard number 600. It comprises a top-level chassis or Line Replaceable Unit (LRU,) that dissipates approximately 100W with several subsidiary LRUs or modules inserted into it. The system had a requirement to operate for 180 minutes after the loss of the aircraft supplied cooling air; termed a Loss of Cooling or LoC scenario. It was this scenario where the CFD analysis failed to correlate to test.

In this particular case, the preliminary thermal analysis included an up-front Computational Fluid Dynamics (CFD) analysis using preliminary mechanical and electrical design information to model the thermal situation inside the unit using FloTHERM. The results of this analysis were utilized to establish an initial thermal design strategy for the chassis, which included heatsink design and airflow management. The thermal design plan included a subsequent thermal survey on a fully instrumented early engineering unit, developed to account for the results of this initial thermal modeling. Both the thermal modeling efforts and the thermal survey testing addressed three operating environments: Normal Flight Operating (NFO), Normal Ground Operating (NGO), and Loss of Cooling (LoC). The Loss of Cooling environment required stabilization under Normal Flight conditions followed by operation with no forced-air cooling for 180 minutes. This environment largely drove the design of the system as the COTS components were very near to their upper engineering temperature limits. The custom heatsinks implemented in the unit were optimized for best performance across the various environments using the CFD tool.

During the LoC test portion of the thermal survey, the unit suffered functional failures and many of the temperature predictions were as much as 20°C below the corresponding test data. These discrepancies between analysis and testing gave rise to late design modifications. A quick review of the thermal model indicated that the model was constructed fairly well and seemed to be reasonably representative of the final configuration of the product. There were

some areas where the model fell short, such as where component parameters weren't available, as the part had not yet been fully designed, so their power was spread over the Printed Wiring Board's (PWB's) surface. In general, the model was built to the usual standards. Correcting the obvious few small shortcomings did not completely rectify the errors that were seen in the result.

In order to maximize the efficiency and knowledge benefit of the exercise, the original team of engineers that performed the thermal analysis and heatsink optimization was pulled together. The investigation was run as a small engineering project. The goals defined for the study were to try to understand where the initial modeling effort had fallen short, find, and then document the requisite changes in modeling approach to improve the prediction accuracy of future modeling efforts for a chassis of this type.

The first task undertaken in the review was to revisit the initial thermal model used to evaluate the thermal situation which drove the heatsink and airflow metering strategy for the chassis. The model was updated to match the geometry and component thermal details as they were tested in the thermal survey without significant changes to the modeling assumptions used in its construction. Two specific sets of test data were chosen to pursue correlation that then drove, by necessity, two separate CFD models. The two tests chosen were identified as the most representative of the chassis final configuration with only small, known exceptions that could be modeled separately for each (e.g. presence or absence of heatsinks added in the given test.). The goal for this effort was not so much to accurately model the final configuration of the chassis as it exited the testing but, rather, to get to a correlated model that made engineering sense and that matched each set of thermal test results for each of the two operational configurations.

This chain of events was fortuitous because, as the correlation effort progressed, it became clear that the effort would require two quite dissimilar models in order to get correlated results for each operational situation. The LoC model ended up being different from the NFO model in ways that exceeded just the differences in unit configuration between the two test scenarios.

From these tests several Lessons Learned were obtained. The two models that came out of this effort uncovered a number of nuances to the modeling of this type of chassis and environment that the team was not aware of at the outset. The lessons learned will facilitate modeling efforts on future programs with similar

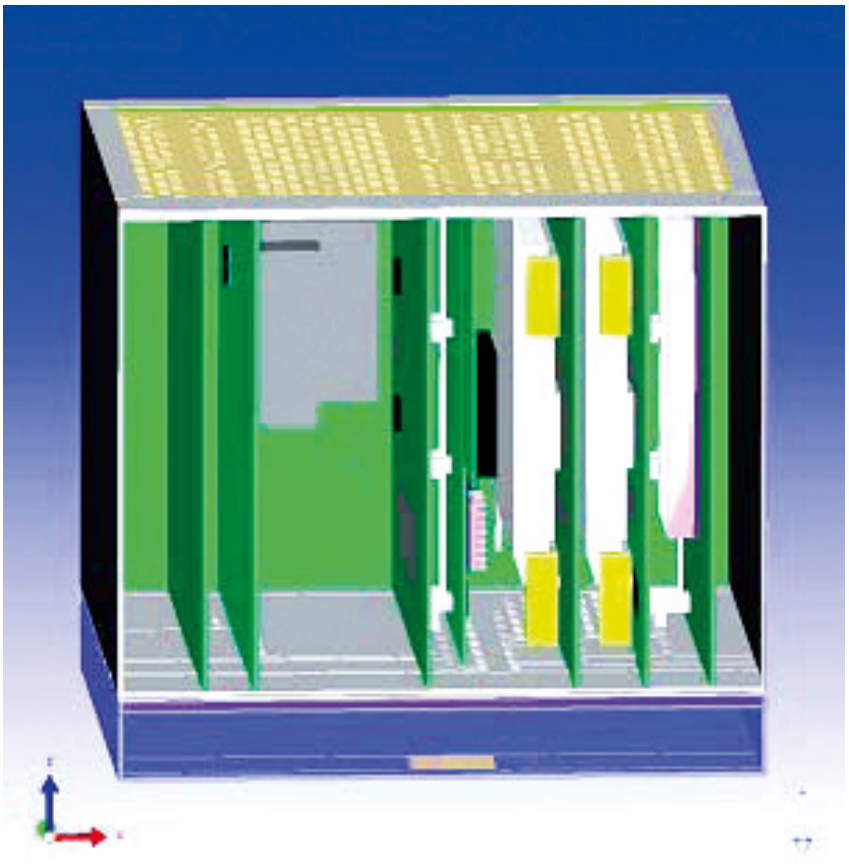

Figure 2. Final NFO CFD Model

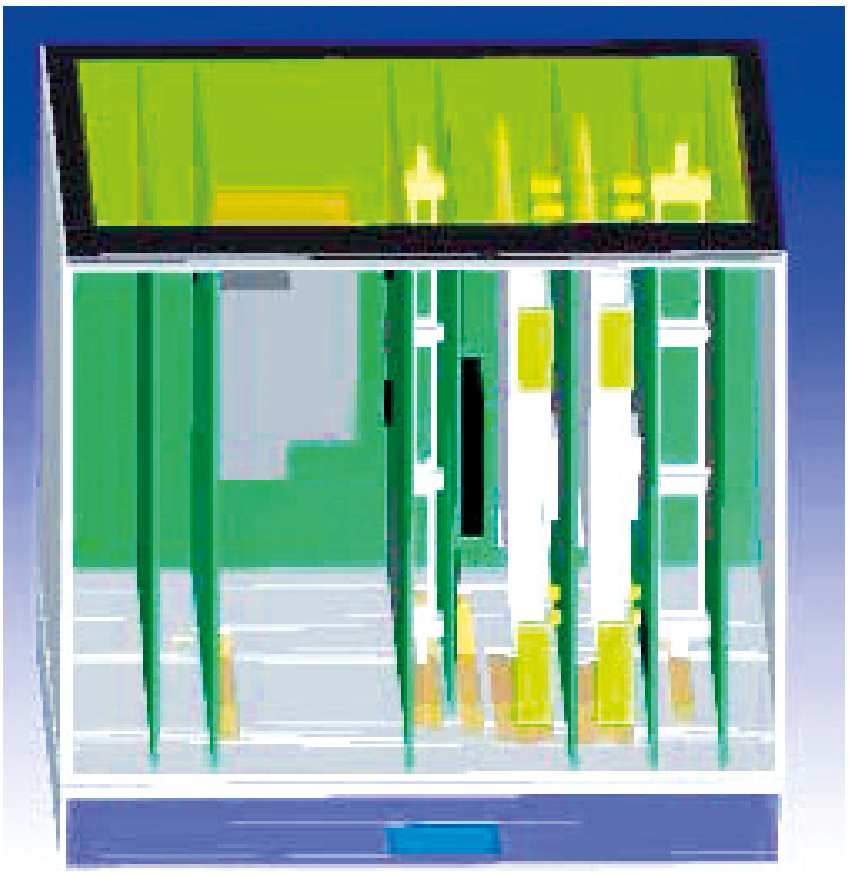

Figure 3. Final LoC CFD Model

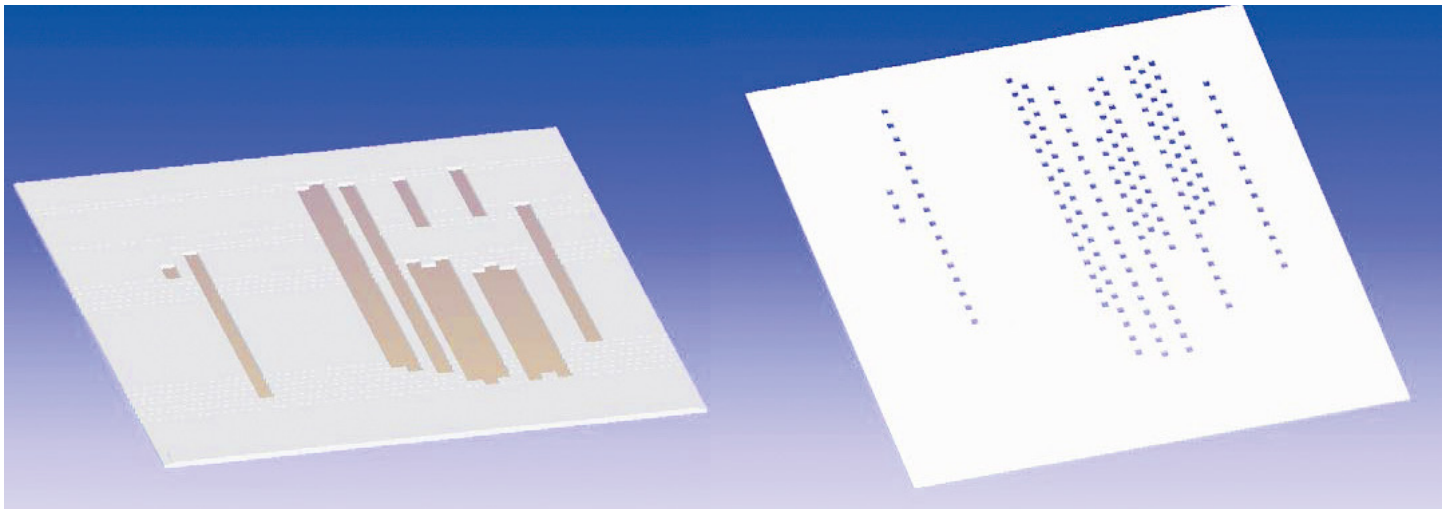

Figure 4. Final NFO (left) and LoC (right) Metering Plates Comparison

chassis designs. Here are some of the more significant findings:

- Both scenarios required refinements of the modeling approach to the inlet conditions for the chassis:
  1. For the NFO case, the original model had utilized correctly sized openings with perforated sheet components with percentage open parameters set to agree with the expected metering plate design. A fixed flow was then imposed on the openings that would provide the required mass flow per the system design. This resulted in a nearly pure vertical flow through the chassis. During the follow-up investigation, the temperatures could not be made to correlate across the entire chassis with this configuration. Two modeling changes were required to fix this issue. The first was to add a detailed model of the plenum used in the test setup. This accurately modeled the airflow within the plenum and introduced lateral and fore to aft flow variations that allowed the model to correlate better. Also for the NFO case, the rows of metering plate holes were modeled as long thin perforated sheet strips, which allowed faster model convergence, but the percentage open had to be adjusted downward to account for the interaction between the inner and outer chassis perforations. See Figure 2.
  2. For the LoC case, the inlet plenum also had to be modeled in detail. Further, getting the mass flow drawn into the chassis by natural convection required that it be monitored and controlled in the simulation. A fixed resistance simulating the test chamber inlet ducting was added and adjusted to match the very low inlet mass flow measured during the LoC tests. While using long thin, perforated sheet strips for the inlet worked well under force air conditions, for the LoC case, this approach did not allow for accurate correlation of the two models. In this case, each metering plate inlet orifice had to be modeled individually, as the velocity profiles across the rows of orifices were not uniform. See Figure 3 and Figure 4.
- The exhaust configuration for both chassis was modeled initially using perforated plate components in FloTHERM. This was found to also not accurately model the exhaust conditions for the LoC case. Ultimately for LoC, the best results were achieved when the chassis top was also modeled as a grid of small orifices below the previous perforated sheet component.
- The LoC model is a steady state model, thus, it produces the temperatures at infinite time. The temperatures used to correlate the model had to be adjusted upward from those measured in the 180 minute LoC test. This was possible to do analytically as the test data was exponential in the last several minutes of the test and a high confidence prediction of the temperatures at infinite time was easy to make. This was a small detail but the error associated with not making this adjustment was greater than the desired 2°C error for predicted temperatures on the hottest components.
- On average, a general component's power dissipation was overestimated under NFO conditions by 20 to 40%. The NFO model, thus, generally overestimated component temperature rises.
- The non-linear thermal behavior versus temperature of several components resulted in their correlated power dissipations being significantly higher than those found in the correlated NFO model. This demonstrated that having a correlated NFO model, which is then run without airflow to simulate the LoC case, would severely underestimate component temperature rises of all these components.
- In general, the initial power dissipation estimates used to construct the original CFD model ended up matching the correlated power out of the LoC test data. It was found, however, that the final correlated power supply component power dissipations averaged approximately 50% higher than the original estimates. This was attributed to the increased system power required to drive the components that were exhibiting non-linear power increases with temperature.
- The initial model was missing several components because the data for them was not available and some turned out to be key to the heat generation. Some of these components ended up driving specific thermal decisions later, during the appraisal tests. Key point here is to have as many components modeled as early as possible in the process.

This Lessons-Learned project uncovered a number of facets of the original analysis work that go beyond a simply flawed analysis approach. Several of the usual assumptions for this type of CFD modeling proved to be inadequate and/or incorrect. As a side benefit of this effort, a procedure for quickly and reliably correlating a large complex thermal model to measured thermal data was developed and refined. The results presented here are applied on and will improve the results of all follow up development projects.

A Study of Electrolytic Capacitor Thermal Conductivity, Behavior & Measurement

By Zhigang NA,
ThinkPad Development Lab, Lenovo

Electrolytic capacitors are widely used in electric circuits, and their durability is an important contributor for the entire lifespan of an electric device. Usually, each supplier would have their own lifetime calculation method. For example:

$$L_x = L_0 \times 2^{\frac{T_0 - T_1}{10}} \times 2^{\frac{\Delta T_0 - \Delta T_1}{10}} \quad [1]$$

Where: L_x = Expected life at temperature T_1 °C

L_0 = Guaranteed life at temperature T_0 °C

T_0 = Maximum operating temperature (°C)

T_1 = Actual operating temperature (°C)

ΔT_0 = Internal temp raise under maximum ripple current (°C)

ΔT_1 = Actual internal temp raise (°C)

According to Eq.1, a 10°C temperature raise (either ambient temperature or internal temperature) will degrade the lifetime of the capacitor by 50%.
In order to devise an adequate cooling solution to prevent the electrolytic capacitor from overheating or even burning, the thermal designer needs to completely understand the component's thermal characteristics.

Due to the constraints of the capacitor corking principals and measurement conditions, it is very difficult to heat a capacitor with an accurately known power. It is also challenging to accurately measure the capacitor internal temperature. Computational Fluid Dynamics (CFD) simulation is a major asset for this type of study. When coupled with real sample tests, CFD can be used to verify key results to ensure the overall accuracy of the study.

Heat Exchange of a Capacitor on PCB Heat Exchange Model

When a capacitor is mounted to a PCB, the PCB acts as a heatsink. From a heat transfer point of view, heat is exchanged between the capacitor, PCB, and the ambient air. The heat transfer modes include conduction, convection, and radiation.

Figure 1 (overleaf) illustrates the heat transfer mechanisms. A thermal resistance network model can also be used to represent this. Since this study was focused on a forced convection system, the effect of heat radiation is ignored because it has very little affect on heat transfer due to the relatively low temperature of the capacitor.

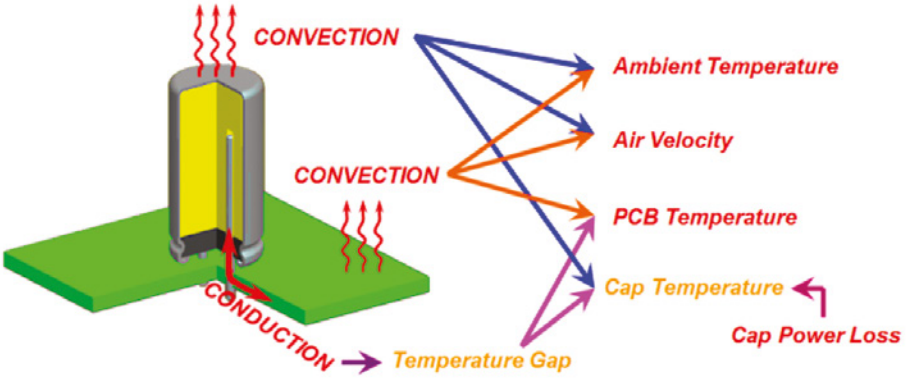

Figure1. Heat exchange modes of a Capacitor on PCB

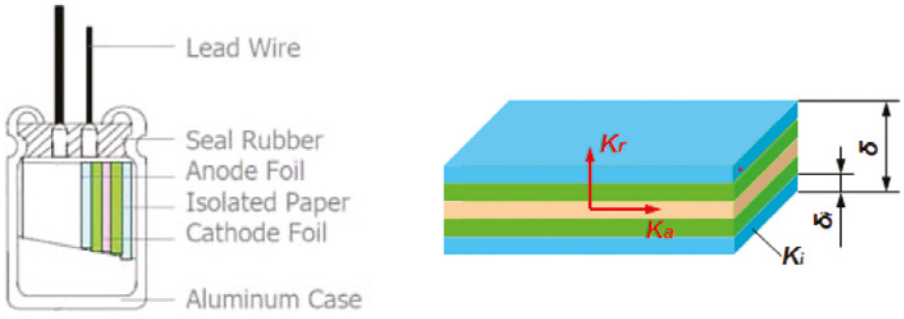

(a) Capacitor structure [1] (b) Simplified winding structure

Figure 2. Capacitor structure and simplified winding structure

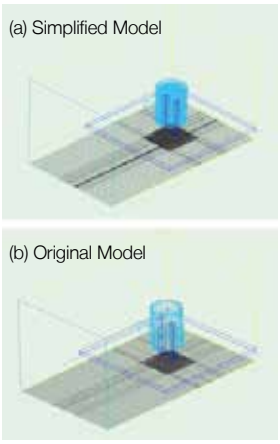

Figure 3. Grid of simplified model and original model

	Simplified Model	Original Model
Cell Quality	141,584	1,790,246
Max Aspect Ratio	7.33	35.10
Number of Iterations	350	750
Residual/ Convergence	1 / Convergent	>10 / Divergent
Solving Duration	13m:55s	58m:21s

Table 1 Difference between simplified and original models

Heat Transfer Boundary Conditions

From Figure 1, the ambient temperature; air velocity; and PCB temperature impact at least one heat transfer mode in this system, and so they are all boundary conditions for heat exchange of the capacitor.

Since the capacitor is a heat source, generating a certain amount of heat, the capacitor's power loss is also a boundary condition. Meanwhile, the PCB can be treated as a heatsink in the system, as it has much bigger thermal mass than the capacitor. The impact caused to the final result by this treatment can be ignored.

Modeling of a Capacitor Internal Structure of Electrolytic Capacitor

Figure 2(a) shows the internal structure of an electrolytic capacitor. In an actual capacitor, the Anode/Cathode Foil and Isolated Paper are wound together to form many layers.

Conductivity Equation of the Winding Structure

By using FloTHERM® Electronics Thermal Simulation software, the thermal designer can set up a capacitor model following the actual structure, but this kind of model is not always recommended, since it won't make the simulation more accurate. Instead, this kind of model increases both the grid density and cell count. A larger grid will result in a longer solve time.

To avoid these issues, the winding structure can be simplified while still retaining the model's accuracy. For this winding structure, if the layers were unwound, the internal structure can be simplified to a stacked structure as shown in Figure 2(b). Based on this simplified structure, the conductivity of the internal winding layer can be calculated by:

$$K_r = \frac{\delta}{\sum_{i=1}^{n} \frac{\delta_i}{K_i}} \quad ; \quad K_a = \frac{\sum_{i=1}^{n} K_i \cdot \delta_i}{\delta}$$

Eq.2 refers to the effective conductivity of multiple objects combined in series and in parallel. In Eq.2, K_r is the conductivity at radial direction, and K_a at axial direction. Obviously, the internal winding structure is anisotropic in terms of conductivity.

If the Anode Foil and Cathode Foil are made with Aluminum (K=180W/m•K), and the Isolation Paper is a typical material which K=0.035W/m•K, then K_r=0.08W/m•K, and K_a= 90.02 W/m•K. In case of a different foil material, such as Tantalum, the capacitor's conductivity can be calculated accordingly.

Comparison of a Simplified Model and an Original Model

The simplified model is much better for solving than the original. The differences are illustrated in Figure 3, which also shows the grid of both models. Table.1 confirms the simulation parameters comparison, it is clear to see that the original model has a longer solving time and eventually becomes divergent.

CFD Model of a Capacitor

With the calculated conductivity of the internal winding structure, a capacitor with a PCB CAE model can be set up as shown in Figure 4. This model is used in the following study.

Capacitor Cooling Simulation

Based on the study earlier, the capacitor's power loss, PCB temperature, air velocity, and ambient temperature all impact the capacitor temperature. The following study verifies how each boundary condition impacts the capacitor temperature. The initial conditions are set to: power loss = 0.3W, PCB temp = 80ºC, air velocity = 1m/s, ambient temp = 45ºC.

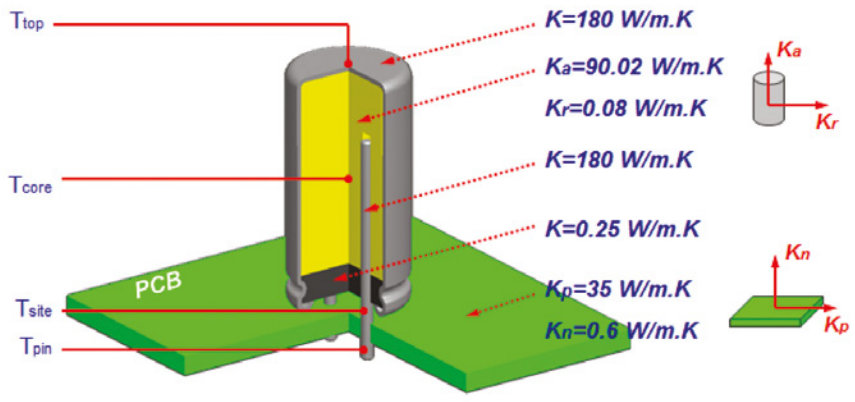

Figure 4. CFD model of a capacitor

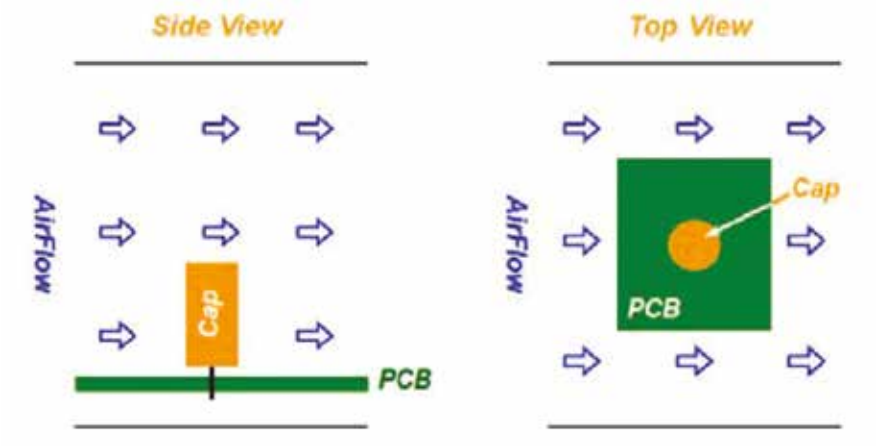

Figure 5. CFD simulation scenario

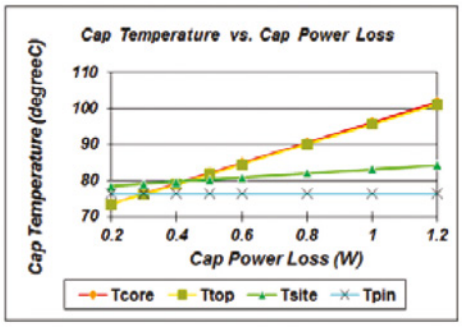

(a) Trend of temperature

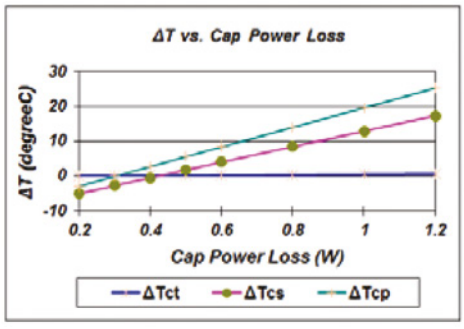

(b) Trend of temperature difference

Figure 6. Variable capacitor power loss

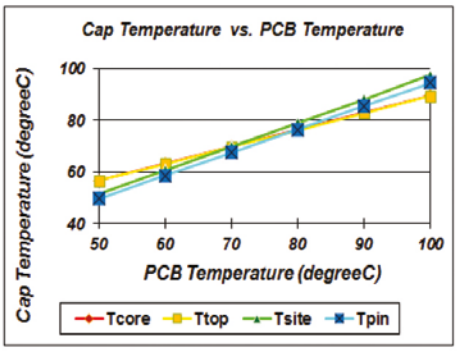

(a) Trend of temperature

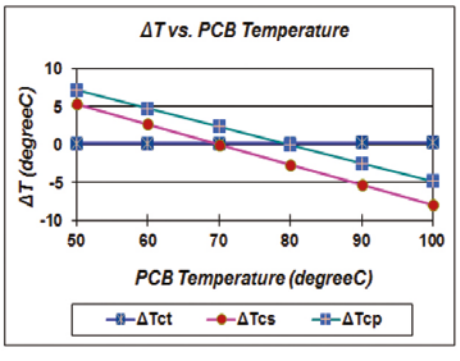

(b) Trend of temperature difference

Figure 7. Variable PCB temperature

In total, four scenarios were studied. In each scenario, three of these four conditions are held constant, while the other is variable so as to show how this condition impacts the capacitor temperature.

Figure 5 shows the solution domain for this study, a DIP (Dual In-line Package) type capacitor with a piece of PCB is placed in a wind tunnel, air flow in the wind tunnel is perpendicular to axis of the capacitor.

As a heat conductor and also heat source, temperature distribution on the capacitor body is not uniform, so the temperature of multiple points on the capacitor are monitored in the study, as follows: T_{top}, T_{core}, T_{side}, T_{pin} (Figure 4). T_{core} is the internal temperature of the capacitor so it is one of the key parameters for the capacitor lifetime evaluation, but T_{core} could not be measured in a real system. So T_{top}, T_{side}, T_{pin} are monitored, and temperature differences between core and top (ΔT_{ct}), core and side (ΔT_{cs}), and core and pin (ΔT_{cp}) were studied.

Variation in Capacitor Power Loss

The power loss was to vary from 0.2W to 1.2W, and the temperature trend of each monitor point was noted.

Figure 6(a) shows temperature trends of each point, T_{core} increases in accordance with the increase in power loss, but T_{pin} is not impacted by the power loss increase at all. T_{side} shows a slight change but keeps within a small range (<5ºC), T_{top} has an obvious increase and the trend is very similar to that of T_{core}.

Figure 6(b) shows the temperature difference trend between the core and other points. It results in ΔT_{ct} only has very slight change (<1ºC), while ΔT_{cs} and ΔT_{cp} have obvious change.

Variation in PCB Temperature

The PCB temperature was set to vary from 50ºC to 100ºC, and then the temperature trend of each monitor point was verified.

Figure 7(a) shows the temperature trend of each point, it appears all four points have obvious increases corresponding with the PCB temperature increase. This means the PCB temperature heavily impacts the capacitor's lifetime, directly conducting heat into the capacitor.

Figure 7(b) shows the temperature difference trend between core and other points. It results in Δ_{Tct} having a very slight change (<1ºC) while ΔT_{cs} and ΔT_{cp} have obvious decrease with the PCB temperature increase.

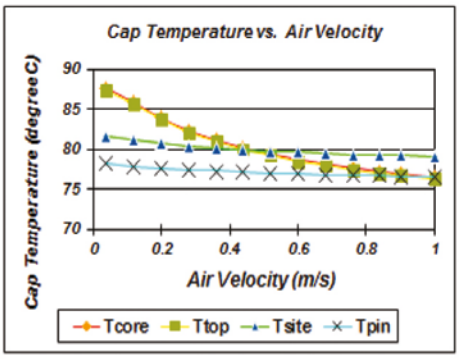

(a) Trend of temperature

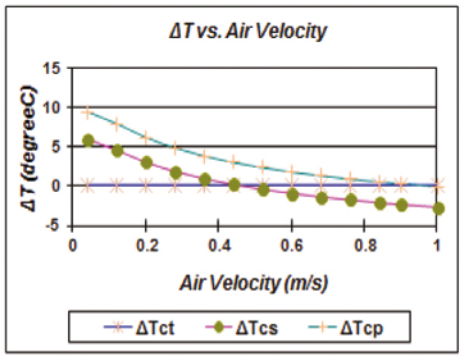

(b) Trend of temperature difference

Figure 8. Variable air velocity

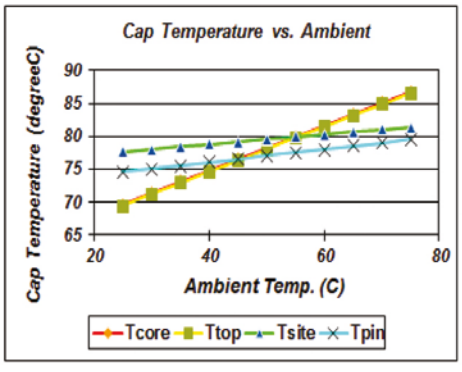

(a) Trend of temperature

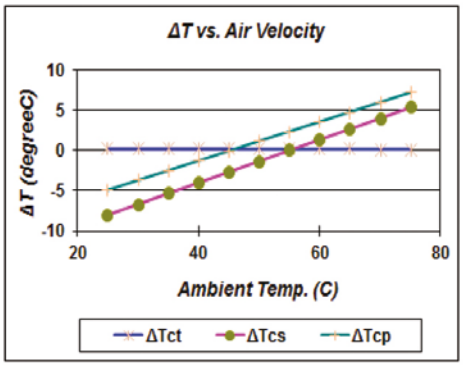

(b) Trend of temperature difference

Figure 9. Variable air velocity

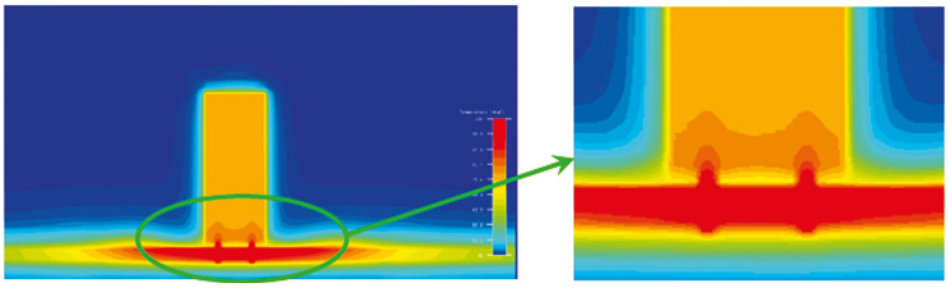

Figure 10. Capacitor temperature field in the FloTHERM simulation

Variation in Air Velocity

Air velocity was set to vary from 0.05m/s to 1m/s, and then the temperature trend of each monitor point was verified.

Figure 8(a) shows the temperature trend of each point, it appears T_{top} and T_{core} decreased in accordance with the air flow velocity increase. While T_{pin} and T_{site} slightly decreased.

Figure 9(b) shows the temperature difference trend between core and other points. It results in ΔT_{ct} just slightly changing (<1°C), while ΔT_{cs} and ΔT_{cp} have obvious decrease with air velocity decrease.

Variable Ambient Temperature

Ambient temperature was set to increase from 25°C to 75°C, and then the temperature trend of each monitor point was verified.

Figure 9(a) shows the temperature trend of each point, it appears ambient temperature impacts the temperature of all points.

Figure 9(b) shows the temperature difference trend between core and other points. It results in ΔT_{ct} also slightly changing (<1 °C) only, while ΔT_{cs} and ΔT_{cp} have obvious increase with ambient temperature increase.

Temperature Measurement Point Study

In a real system, only the outside surface temperature of a capacitor can be measured, while, internal temperature is required for lifetime evaluation. So a proper measurement point which has a small deviation from internal temperature needs to be found.

Traditionally, some capacitor manufacturers recommend measuring pin temperature (T_{pin} in Figure 4) for a DIP type capacitor, as the pin is a high thermal conductor and is in contact with the capacitor internally. However, according to this study, the temperature difference (ΔT_{cp}) is not constant, so the pin temperature should not be used to reflect internal temperature. Figure 9 shows a capacitor's temperature field in the FloTHERM simulation. In this case the PCB temperature is higher, so the pin temperature (T_{pin}) will be also higher than internal temperature (T_{core}).

Figure 10 Capacitor temperature field in the FloTHERM simulation

According to the study, the temperature at the top of the capacitor case (T_{top}) is almost constant when boundary conditions change, so the top of the case is the best measurement point in the case where the airflow pattern is same as shown in Figure 4.

Conclusion

This study developed a simplified capacitor model for use in a CFD simulation. This simplification can improve grid density and quality in the simulation model, and thus improve the accuracy of the simulation. This study also identified all boundary conditions that impact the capacitor's cooling, and then verified how each boundary condition impacts the capacitor temperature. Referring to this study, the thermal designer can improve the capacitor cooling solution by optimizing boundary conditions.

Finally, the top case temperature (T_{top}) was determined as the best point to reflect the capacitor's internal temperature (T_{core}). Across the range of boundary conditions tested, the temperature difference between top and internal is constant and only around 1°C, so the system designer can easily convert the top case temperature to an internal temperature.

References

[1] GDDL, Cap lifecycle calculation template

Fujifilm Frontloads Camera Design with FloTHERM®

By Kazuya Mayumi, Optical Device and Electronic Video Product Development Center, Fujifilm Holdings Corporation

FUJIFILM

Our department focuses on the development of cameras with high performance and high quality, to preserve the culture of photography and to develop imaging technology. The team I belong to develops the elemental technology of such cameras and lenses and supports the business by running simulations. I have been involved in thermal design by predicting temperature with simulation tools ever since the first "X-100" of Fujifilm's X digital camera series. Now simulation technologies are essential for product development today.

Modern digital cameras, like Fujifilm's X-Pro 2 are developed to dissipate the minimum amount of heat as form factors shrink. Excessive heat can increase the size and weight of the cooling solution, preventing miniaturization. It is associated with the concept of poor design by the consumer, as the design can appear to have "regressed" back to an earlier bulkier form factor. It also requires the use of higher-cost parts, adding cost to the product, further reducing its competitiveness.

Heat generated within the camera has to be managed for many diverse reasons, such as the tendency for elevated temperatures to reduce image quality through noise in the Complementary Metal–Oxide–Semiconductor (CMOS) circuitry. The temperature of the chassis must also be limited to avoid the risk of a low temperature burn for the user. For cost reasons, sheet metal parts within the camera can't be replaced with alternatives that spread the heat better if a problem is found in late design or prototyping.

For these reasons considerable care must be taken prior to prototyping. Fujifilm do this by undertaking the following work from the start of development:

- Present an optimal solution for the structure, as an outcome of thermal simulations on several different design variations,
- Predict the maximum temperature without using any correlations, and
- Complete the study by defining what parts, including their size and number, are needed for the heat dissipating structure.

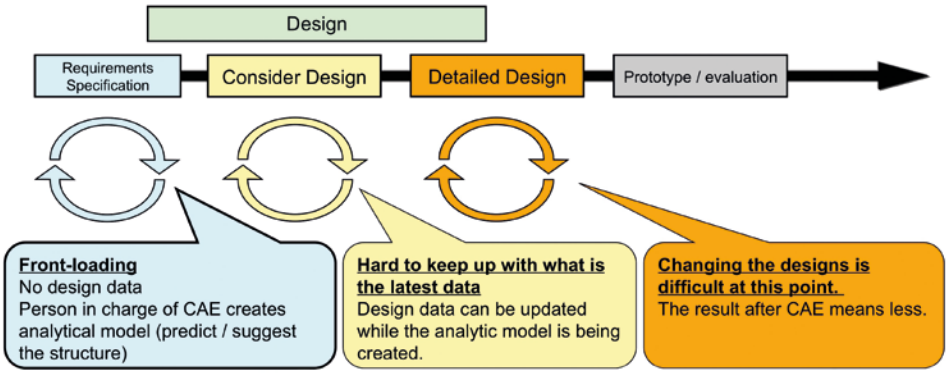

Figure 1. Fujifilm Camera Design Process

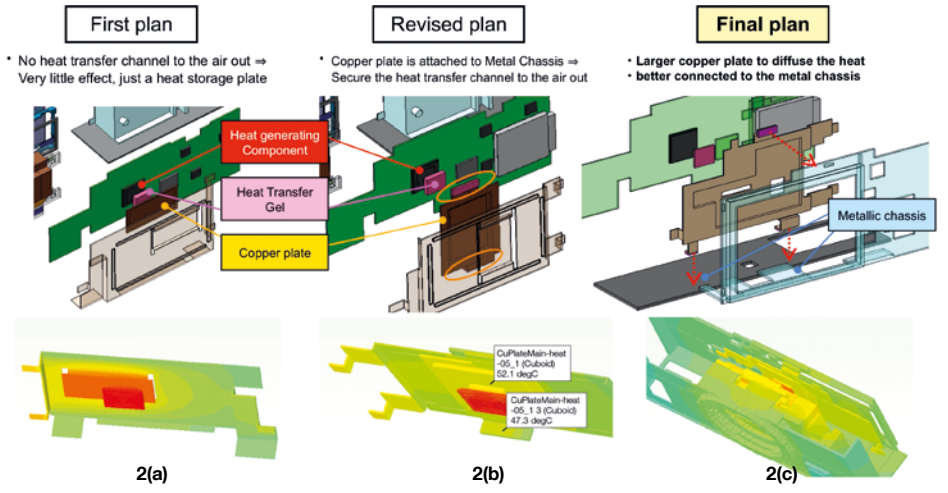

Figure 2. Initial, Intermediate and Final Frontloaded cooling solutions (pre-CAD)

The process is shown in figure 1: It starts with a simple analysis model built in Mentor Graphics' FloTHERM Electronic Thermal Analysis software. It uses geometry simplified from the CAD model – the exact details of the process Fujifilm considers proprietary.

"We use FloTHERM to frontload our thermal design. FloTHERM enables us to explore design alternatives early on, design the optimal solution for removing the heat, and provides very good accuracy against the final test results we get for chip, component, and touch temperature on the case."

Kazuya Mayumi, Optical Device and Electronic Video Product Development Center, Fujifilm Holdings Corporation

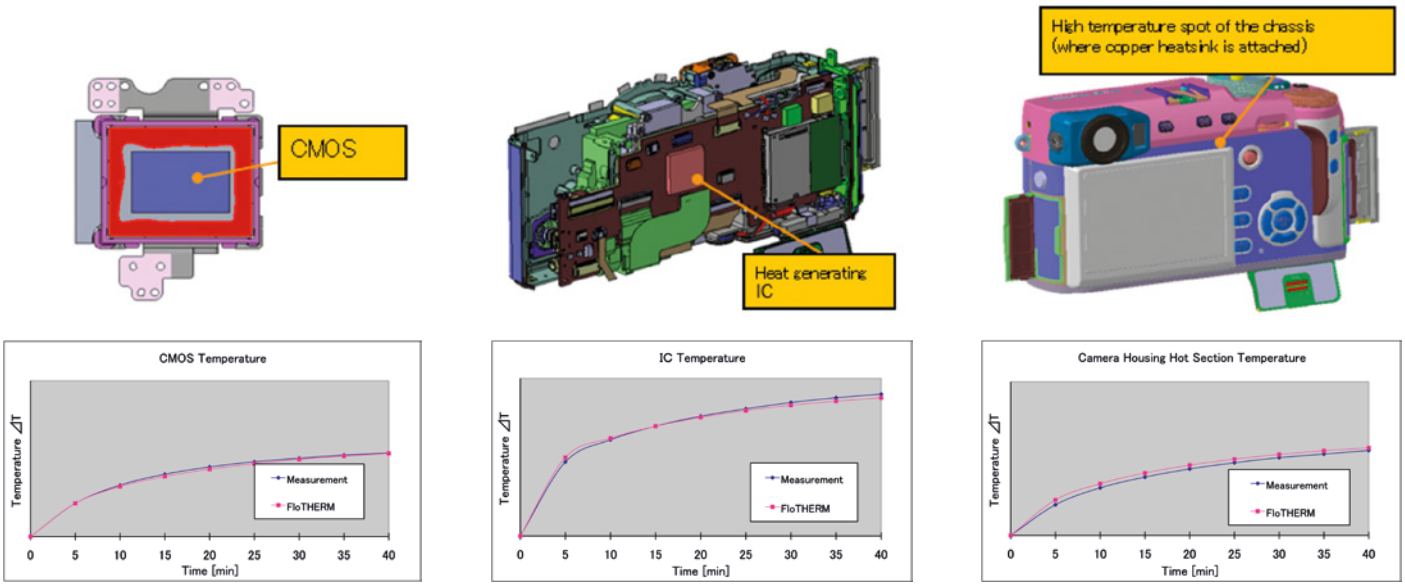

Figure 3. FloTHERM model agreement with test results at different packaging levels.

By keeping this analytical model updated, the effects of changes can be fed back into the design process immediately, frontloading simulation in the development process.

As a case in point, Fujifilm needed to improve thermal performance, introducing a copper heatsink and thermal gel into the product for the first time. In the initial design, shown in figure 2a, along with the FloTHERM result, the copper heatsink just acted as a heat storage device, as there was no heat flow channel to get the heat out. The revised design shown in figure 2b, attached the copper heatsink to the chassis, ensuring that the heat could be transmitted to the air. However, the temperatures were still too high. A final design iteration, shown in figure 2c, increased the size of the copper plate to better diffuse the heat, and additional connections were added to improve heat transport to the chassis.

Using this approach, Fujifilm are able to get highly accurate temperature predictions during the early part of the design process, and use the results of improvements made to the design using FloTHERM to update the CAD model. The predictive accuracy of this analysis model is checked towards the end of the development process once a physical prototype becomes available. The detailed CAD geometry available at this stage is shown in figure 3 for the CMOS chip generating the heat, the component housing the chip, and the hot spot on the camera housing, resulting from this heat source, together with the comparison between measurement on the prototype and the FloTHERM model, which shows excellent agreement at all of these packaging levels.

Using CAD data directly in FloTHERM gives results which are very similar to those shown in figure 3, for the model based on a simplified geometry for the analysis. The import of the CAD geometry is also acceptably fast for use in the design, however this approach is not favored for the simple reason that CAD geometry is not available at the start of the process, so this approach would have the effect of delaying the CFD work until later in the design, when the business benefits of using CFD are less. Whereas the requirements for the cooling solution must be addressed and different options investigated at the outset.

Critical to the design effort is the ability to predict accurate temperature values, not just trends or relative differences. This is because of the thermal limits both on the component's operating temperature, and also the maximum permissible touch temperature of the case. Consequently, it is the responsibility of the CAE engineer to investigate any differences between the simulation result and test data. One source of this discrepancy can be the power values used for the heat sources. It is quite common for electrical engineers to provide power values during design that are conservative, i.e. they over-estimate the likely actual power by say 20% to ensure that the cooling solution will be adequate. However, this makes it impossible to match simulation results to test results, so Fujifilm now ensure that the power estimates that the electrical design team share with the thermal design team are as accurate as possible.

More broadly, Fujifilm consider the ideal role for a CAE engineer is to remain in the background. However, they should be proactive, working to reduce the workload of the designer, and thereby speed the design. CAE engineers should present the results of design optimization and an improved design back to the designer, rather than just reporting back the results of the design iteration they were asked to simulate. Only by optimizing the design process in this way is it possible to perfect the design of the product.

Robert Bosch India Drives ECU Temperatures Down

BOSCH
Invented for life

By Ritwik Alok Pattnayak & Dr. Laxmidhar Biswal, PhD, FIE, CEng,
Robert Bosch Engineering and Business Solutions, India

Modern cars can contain upwards of 100 Electronic Control Units (ECUs) so the reliability of these units is critical to the reliability of the vehicle. ECUs can be found throughout a vehicle, often in hot, harsh environments, yet component temperatures must be kept within acceptable limits even when the ambient temperature is high, making optimizing their thermal performance a key element of the design challenge.

For this, Robert Bosch Engineering and Business Solutions needed a fast, accurate and robust simulation tool they could depend on, to explore different design options quickly and effectively, including the transient response to a time-varying power load, so they turned to Mentor Graphics' FloTHERM.

A product recently developed by the team was an ECU with various functions within the vehicle which result in a different thermal load on the unit. The ECU consists of PCB, components, thermal interface material, and metal housing. There are both active components (e.g. transistors, diodes etc.) and passive components (e.g. resistors, inductors, capacitor etc.) on the board in the ECU that generate heat during the operation. There must be a good conductive path (Figure 1) for heat to flow from components to the housing. Automotive components often contain a thermal slug, to spread the heat internally within the component. These can either be between the die and the top of the component (slug up), to assist heat removal from the top of the component, or between the die and the bottom of the component (slug down), to assist heat transfer into the PCB. Heat entering the housing by either route is dissipated to the external ambient by convection and radiation. The top cover is anodized to improve radiative heat transfer to the environment. Thermal simulation is used to predict the temperature of components and determine their thermal criticality.

The FloTHERM model of the ECU is shown in figure 2, with the housing acting as a

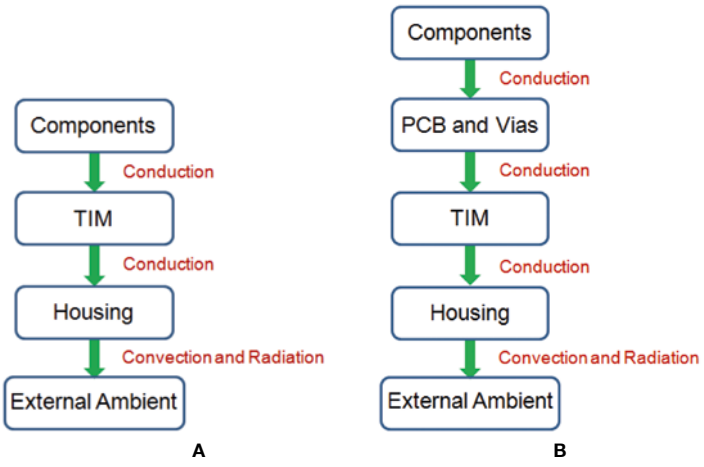

Figure 1. Heat Transfer Paths for (a) Slug-up and (b) Slug-down components

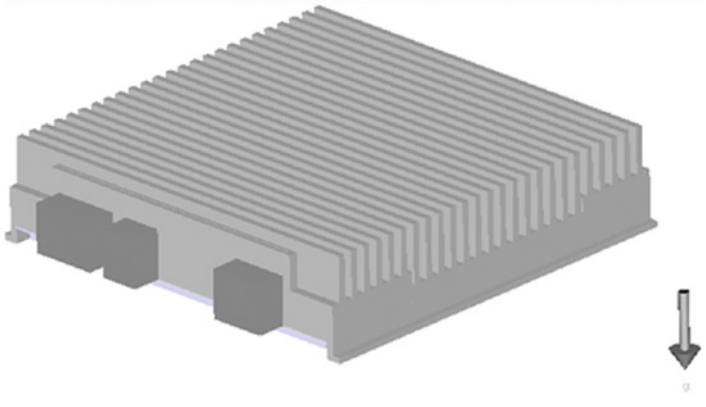

Figure 2. FloTHERM Model of ECU showing Finned Housing

Mentor Graphics
Mechanical Analysis

heatsink to the surrounding air. The ECU is designed to be mounted either inside or outside of the vehicle, and the design has to operate under the worst case ambient condition.

The 3D thermal simulation in FloTHERM includes conduction, convection and radiation. The active components (transistors, diodes, etc.) are modeled in detail (with die, die-attach, die-pad/slug etc.), whereas the passive components that generate heat (resistors, inductors and capacitors) are modeled as lumped cuboids, unless the detailed model is available. The PCB is a multi-layer board constructed from copper and FR4. Each copper layer is modeled as a lumped cuboid with orthotropic thermal conductivity based on the percentage of copper coverage separated by FR4 layers. Appropriate surface emissivity is assigned to the components, PCB and housing to address radiative heat transfer. For the transient simulation the load profile vs. time is needed (Figure 3), where 100% thermal load corresponds to 43.0W. The FloTHERM simulation for the unit in natural convection showed that the maximum component temperature reached 162°C, well beyond the acceptable limits specified in vendor datasheets.

In the baseline design, the PCB does not have a contact with the bottom housing. Hence, heat transfer from the PCB to bottom of the housing is hampered by the absence of a proper thermal conduction path. Convection and radiation from the components and PCB to the housing is not as effective as direct thermal conduction. Bridging the gap between the bottom of the PCB and the housing without requiring a structural change is a challenge, making the thermal design harder.

To address this, Bosch's engineers, Ritwik Pattnayak and Dr. Laxmidhar Biswal, decided to make two changes to the design. The first was to add a metal heat spreader on the underside of the PCB, with a thin layer of thermal interface material in between, to spread the heat and pull down the maximum component temperature. A second change was to add a much thicker section of thermal interface material between the metal heat spreader and the bottom to provide more effective conduction heat transfer.

With the first improvement, the maximum junction temperature of the hottest component is 133°C (Figure 5a), a reduction of 31°C compared to the baseline design. The second modification reduced the maximum junction temperature of the component by a further 5°C to 127°C (Fig. 5b), giving

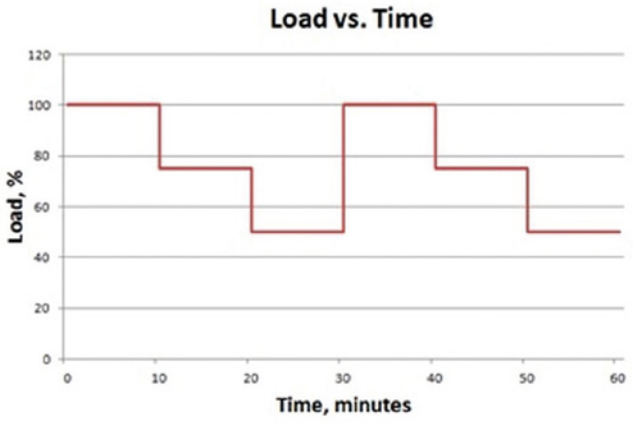

Figure 3. Power Load Profile for the ECU

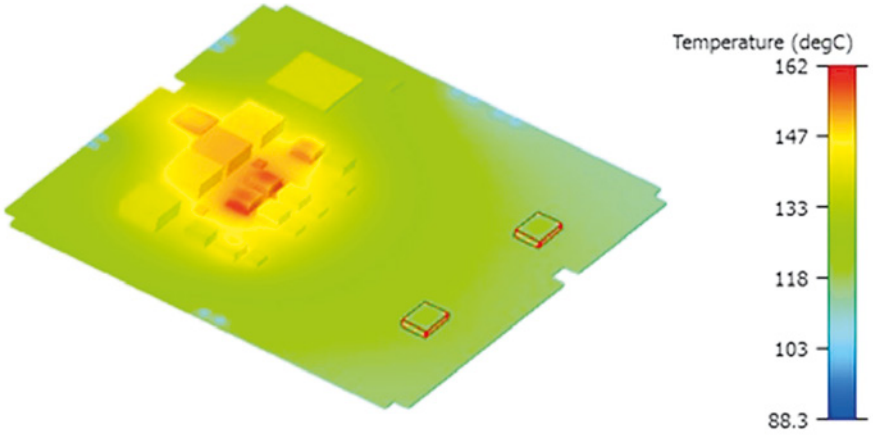

Figure 4. Temperatures of Components in Baseline model.

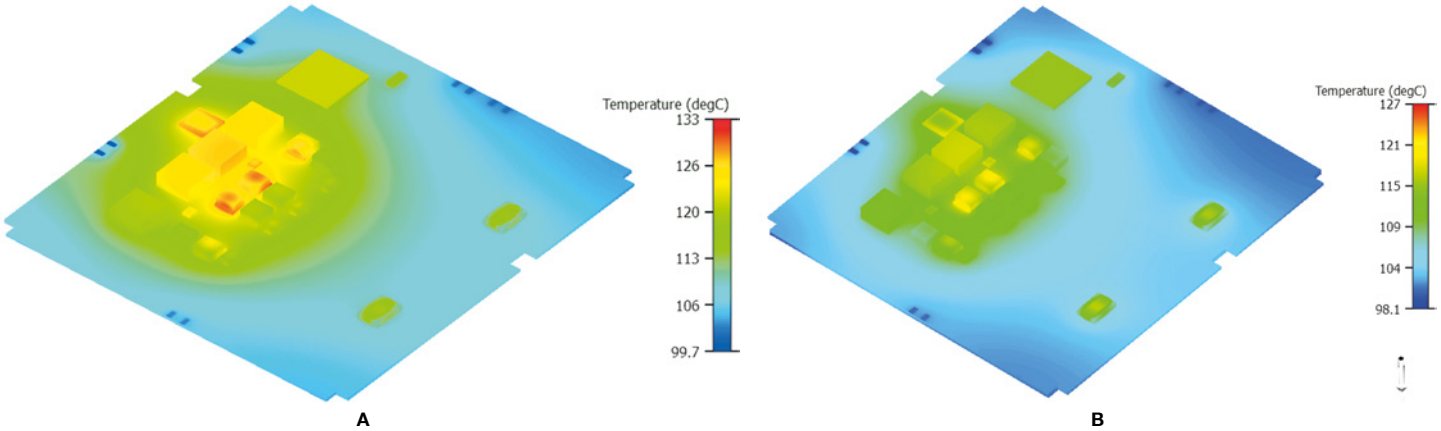

Figure 5. Temperature of Components and PCB for optimization # 1 & optimization # 2.

a total reduction of 35°C is achieved in comparison to the baseline design.

To account for the variation in load power for the system, the sensitivity of the response of individual components to different power levels was tested. Two components were chosen for this, D1600 and V1907. The D1600 component is a slug up design and interfaced with the top finned housing via thermal interface material. Because of lower thermal impedance of the thermal interface material and finned top housing, junction-to-ambient thermal impedance is relatively low, below 5K/W. V1907 is a slug-down component and interfaced with bottom housing via PCB, Thermal Vias, TIM, Heat Spreader and TIM. Due to several layers of materials between the component, bottom housing and associated thermal impedance, junction-to-ambient thermal impedance of 'V1907' is relatively very high in comparison to the slug-up component 'D1600'. This latter component was simulated in FloTHERM to obtain its transient response at 100%, 75%, and 50% power loads. Plotting these temperature responses normalized the applied power to give a graph of dynamic thermal impedance, Zth, the data almost mirror each other as shown in figure 6.

As the data line almost mirror one another, engineers at Bosch were able to derive an analytical model for dynamics thermal resistance (Zth,j-a) of all heat dissipating components from the FloTHERM simulation. This model could then be used to predict the temperature of the components for different power dissipations and external ambient, without having to run a full 3D thermal simulation for each scenario.

The results of the design optimizations were checked with a prototype using a thermal chamber, shown in figure 7. The experiment was conducted for the specified ambient, and at an air speed of 0.6ms-1, with the velocity of the airflow over the ECU measured with an anemometer. Thermocouples were placed on the components of interest to measure their temperature.

The comparison of the simulated temperatures against the experimental measurements showed a good correlation, confirming the robustness and reliability of the FloTHERM model. The outcome of this work was that the design optimizations in FloTHERM overcame the thermal bottleneck and reduced component temperatures by

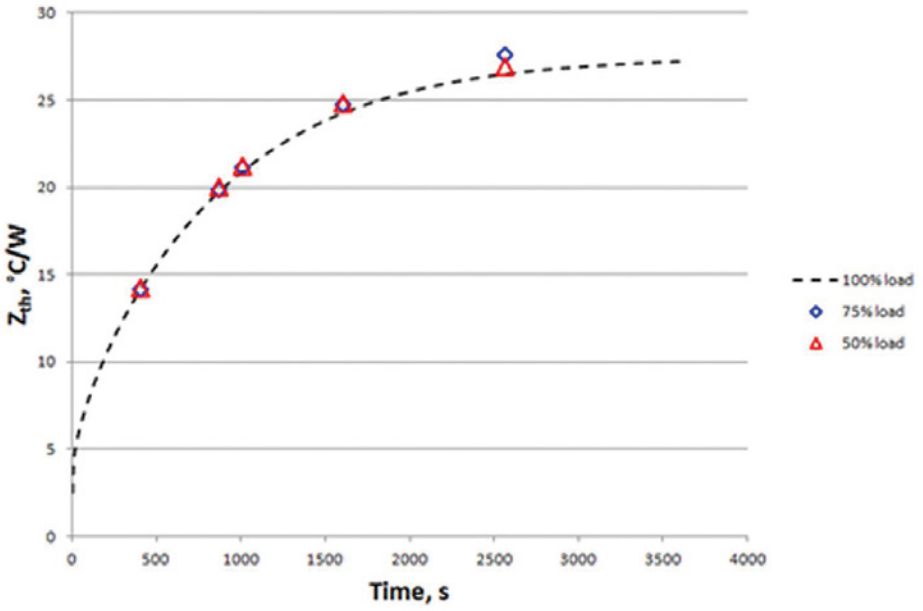

Figure 6. Comparison of dynamic thermal impedance (Zth) of 'V1907' for different loads

Figure 7. Thermocoupled ECU and Anemometer in the Thermal Chamber

35°C from the baseline design, resulting in an acceptable operating temperature of components. Using analytical models derived from FloTHERM results, Robert Bosch Engineering and Business Solutions have been able to predict junction temperature of component at any load for a given design and time duration or operational time duration of the ECU for the desired junction temperature. This has delivered a significant saving in time.

Acknowledgement:

This work "Thermal design and analysis for high power automotive electronic product" by Ritwik Alok Pattnayak & Dr. Laxmidhar Biswal, PhD, FIE, CEng (Robert Bosch Engineering and Business Solutions, India), was first published at NAFEMS' World Congress in San Diego, 21-24 June 2015, www.nafems.org

Liquid Cooling Technology

Leveraging FloTHERM to design a Compact Liquid Cooling System for High Power Microelectronic Devices

By Gongyue Tang, Yong Han, Boon Long Lau, Xiaowu Zhang, and Daniel Min Woo Rhee, Institute of Microelectronics, Agency for Science, Technology and Research, Singapore

Air cooling can provide a simple, low cost, effective, and reliable cooling solution for microelectronic devices. However, with the increase in heat flux dissipation, current air cooling technology is not sufficient for new high power devices. When the heat flux goes beyond 100 W/cm2, air cooling methods become inadequate for most applications. Therefore liquid cooling technology for microelectronic devices with high power chips is required.

There are two major modes of liquid cooling technology; single-phase cooling and two-phase cooling.

Considering the higher pressure drop and complexity of a two-phase liquid cooling system, utilizing the single-phase liquid cooling technology for high-heat-flux microprocessors is an attractive option. For a single-phase liquid cooling technology, both microchannel and microjet heat sinks can dissipate high heat fluxes found in high-power electronic devices. Compared with the impinging microjets, microchannel cooling has a lower averaged heat transfer coefficient but the coolant in microchannels can exchange energy with a larger effective surface area with multiple walls within each of the channels. Combining these approaches into a hybrid microcooler would be the ideal approach.

The proposed liquid cooling system includes three major components:

1. A silicon-based hybrid microcooler with multiple drainage microtrenches (MDMTs);
2. A customized compact liquid-to liquid heat exchanger; and
3. A commercial micropump.

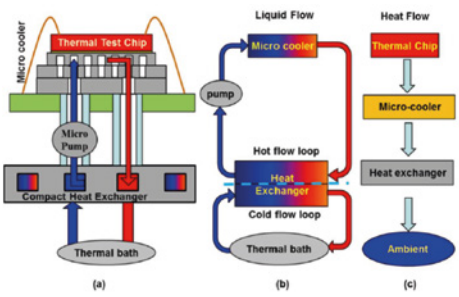

Figure 1. Schematic of the compact liquid cooling system. (a) Concept design of the system. (b) Liquid flow loops in the system. (c) Heat flow path in the system.

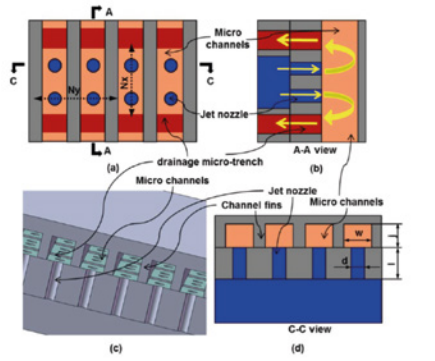

Figure 2. Schematic of the proposed hybrid microcooler with MDMTs design. (a) Pattern of the microchannels, jet nozzles, and MDMTs. (b) Side view of the A—A cross section. (c) 3-D isometric view. (d) C—C cross-sectional view.

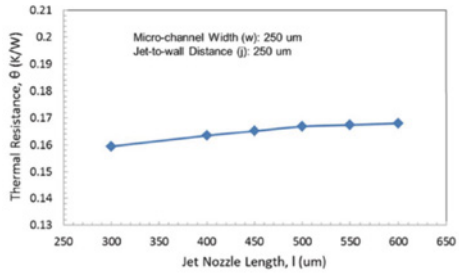

Figure 3. Variation of junction to microcooler thermal resistance with jet nozzle length.

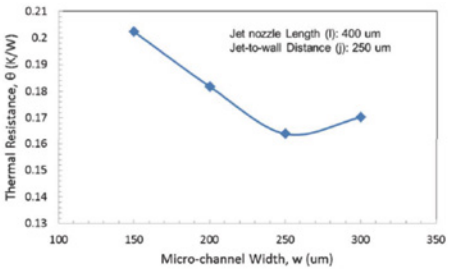

Figure 4. Variation of junction to microcooler thermal resistance with microchannel width.

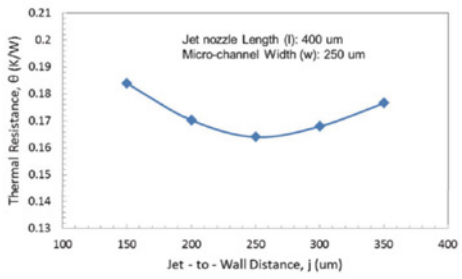

Figure 5. Variation of junction to microcooler thermal resistance with JTW distance.

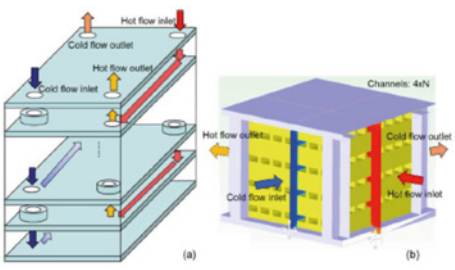

Figure 6. Schematic configurations of the heat exchangers investigated in this paper for (a) counter-flow configuration of commercial heat exchanger A (Hex-A) and (b) cross-flow configuration of customized heat exchanger B (Hex-B).

A compact liquid cooling system for a microelectronic chip is shown in figure 1(a). It mainly consists of a microcooler to remove the heat from the high performance chip, a commercial micropump to drive the liquid flow, and transport the heat from the chip to an external heat exchanger, where the heat is transferred to the secondary fluid and then rejected to the ambient.

To maximize the performance of the system, the following design criteria should be considered in the design of the proposed liquid cooling system:

1. Avoid large pressure drop the flow rate in the loop at chip cooling side should be maintained at a low level;
2. The heat exchanger should be compact with high heat exchange efficiency, or heat transfer density; and
3. It must maintain the thermal resistance and pressure drop of the heat exchanger at a level obviously lower than that of the microcooler.

A silicon-based microcooler, combining the merits of both microchannels and jet impingement, has been developed to dissipate the heat flux for the high-power IC chip. Figure 2 shows the conceptual design of the hybrid silicon microcooler for the pattern of the microchannels, jet nozzles, and MDMTs. As shown in figure 2(a), there are two nozzles between each of the two trenches in the proposed MDMTs design. The jet flow from the microjet nozzle impinges on the top wall, is constrained to flow along the microchannel, and then exits through the nearby drainage microtrenches.

For this work, the CFD modeling and simulation is conducted using FloTHERM® for the microcooler design optimization. For the chip of size 7 × 7 mm^2, the microjet array is designed to cover an area of 8 × 8 mm^2 for the cooling. The jet nozzle diameter is fixed at 100 µm. The nozzle length, microchannel width, and jet to wall distance are varied in the modeling to evaluate their effects on the thermal performance of the microcooler. Other parameters include nine drainage trenches with a width of 150 µm, 16 nozzles along the microchannel direction and 21 jet nozzles along the drainage trench direction. The coolant in the microcooler is water with an inlet temperature of 25°C. A heat source of 175 W is set for the thermal chip, and a gauge pressure of 20 kPa is set at inlet of the microcooler. A mesh of one million cells is generated, and a grid independence study is conducted. The mesh number is sufficient to obtain accuracy within around 1% for

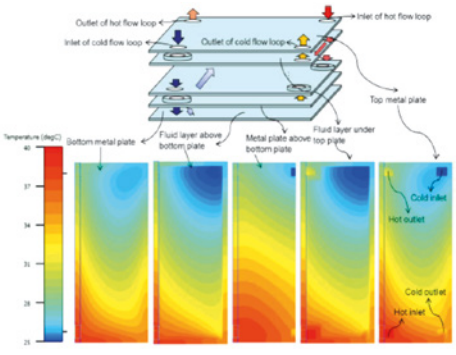

Figure 7. Simulation results of the temperature profiles in the commercial heat exchanger A.

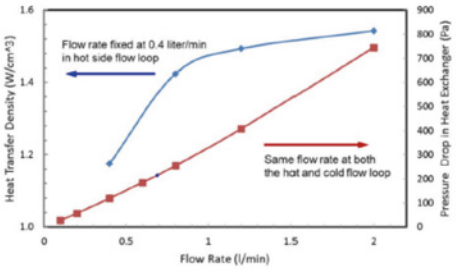

Figure 8. Variation of the heat transfer density, and pressure drop with the rate in Hex-A at the specified conditions.

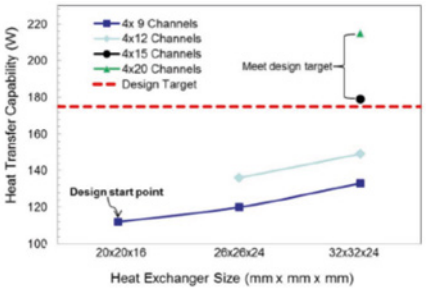

Figure 9. Variation of the heat transfer capability with the channel number's and footprint size of Hex-B at the specified conditions,

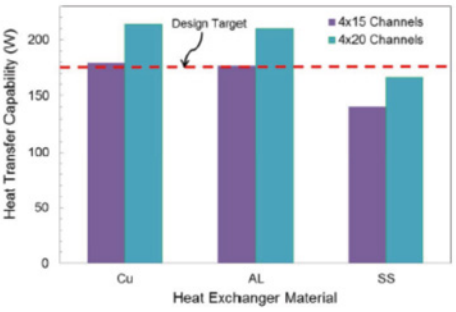

Figure 10. Effect of the material on the heat transfer capability of Hex-B at specified conditions.

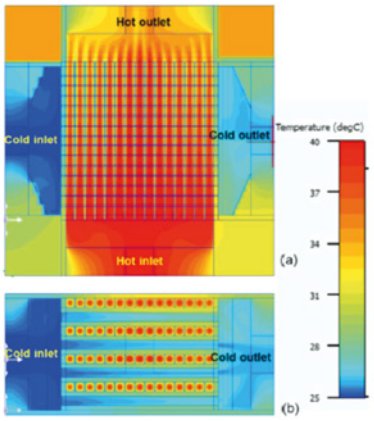

Figure 11. Simulation results of temperature profile for Hex-B with the final design at the specified condition. (a) Top plane view. (b) Side plane view.

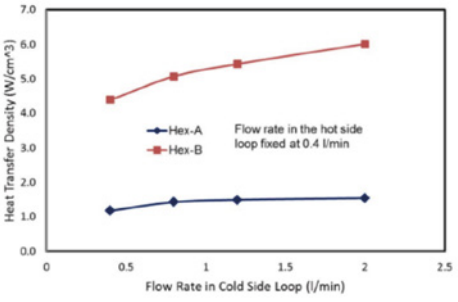

Figure 12. Comparison of heat transfer density for Hex-A and Hex-B at the specified conditions.

hydraulic and thermal performance. The effects of the nozzle length, microchannel width, and JTW distance on the thermal performance of the microcooler are, respectively, presented in figures 3–5.

A liquid–to-liquid heat exchanger was chosen because it has the advantage of higher heat exchange efficiency, smaller size, and centralized management of the secondary liquid flow through facility cooling, such as a coolant distributed unit.

Two liquid-to-liquid heat exchanger configurations are investigated in the present case. The first heat exchanger (Hex-A) is commercially available with a counter-flow configurations, and it is used as a benchmark to design and optimize the second heat exchanger (Hex-B). Hex-B is a customized compact heat exchanger, which represents the cross-flow configuration, consisting of mini-channel layers alternatively stacked for both the hot and cold fluid flows. Hex-A is with the counter-flow configuration, as shown in figure 6(a); the hot fluid and cold fluid layers are stacked alternatively, separated with metal plates. There are a total of 11 stainless steel (SS) layers, including the top and bottom

cover plates with a thickness of 1.1 mm for each layer.

These metal layers form ten layers of liquid channels, which include five layers for the cold liquid and five layers for the hot liquid. The overall size of Hex-A is of 204 mm long by 74 mm wide by 25 mm high. Hex-B is designed with cross-flow configuration, as shown in figure 6(b), in which the hot fluid and cold fluids are stacked, with the two types of flow channels arranged at an angle of 90°.u. Hex-B is targeted to have similar performance as Hex-A while have smaller footprint size than that of the Hex-A. FloTHERM is utilized to simulate the thermal and flow profiles in the heat exchangers.

The inlet flow rate is varied from 0.4 to 2 L/ min with a fixed temperature of 25°C for the cold flow loop of the heat exchangers, whereas the flow rate is fixed to 0.4 L/min for the hot flow loop. The inlet flow temperature in the hot flow loop of the heat exchanger is fixed to 40°C for case study.

Due to the relatively small Reynolds number, the laminar flow is assumed. A mesh of 530 K is generated and the grid independence study is conducted. The fields of flow, pressure, and temperature are obtained through the modeling and simulation; as such the liquid temperature at different locations can be captured through the obtained temperature profiles. The performance of Hex-A is used to benchmark the performance of optimized Hex-B. Figure 7 shows the simulation results of the temperature profile in Hex-A. In this case, the flow rate is 0.4 L/min in both the cold flow loop and the hot flow loop; the inlet fluid temperature is of 25°C in cold flow loop and 40°C in the hot flow loop of the heat exchanger.

The heat transfer density is extracted to evaluate the thermal effectiveness of the heat exchangers. The variation of the heat transfer density with liquid flow rate for Hex-A is shown in figure 8. The flow rate in the hot flow loop is fixed at 0.4 L/min. The inlet fluid temperature is 25°C for the cold flow loop and 40°C for the hot flow loop. The heat transfer density is sensitive to flow rate when the flow rate is low (e.g., less than 1 L/min), and less sensitive to the flow rate when the flow rate is high. The pressure drop in the heat exchanger is also important to the cooling system. The variation of the pressure drop with the flow rate for Hex-A is also shown in figure 8. The flow rates in the hot flow loop and cold flow loop remain the same. As expected, the pressure drop in the heat exchanger increases with an increase in the flow rate.

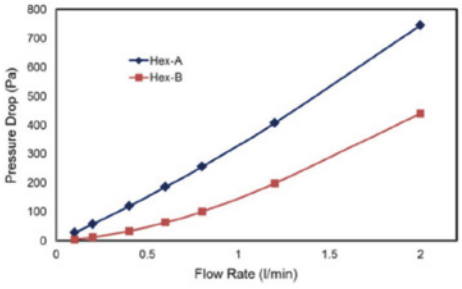

Figure 13. Comparison of pressure drop for Hex-A and Hex -B at the specified conditions.

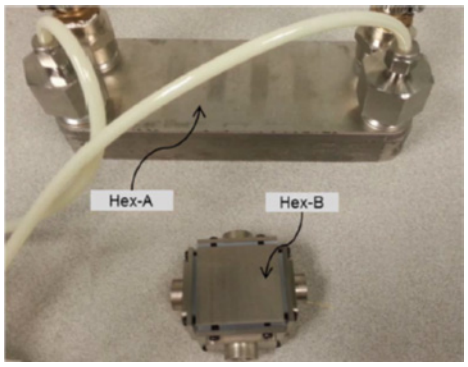

Figure 14. Comparison of the footprint size between the commercial Hex-A and the customized Hex -B

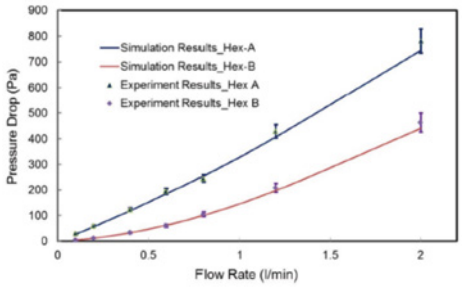

Figure 15. Pressure drop in commercial heat exchanger (Hex-A) and customized compact heat. exchanger (Hex-B).

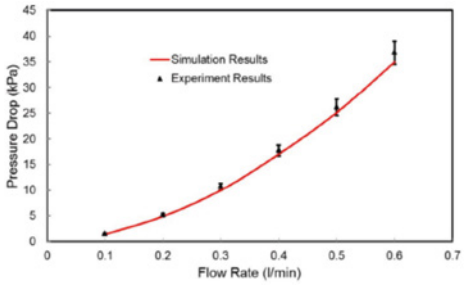

Figure 16. Pressure drop in Microcooler

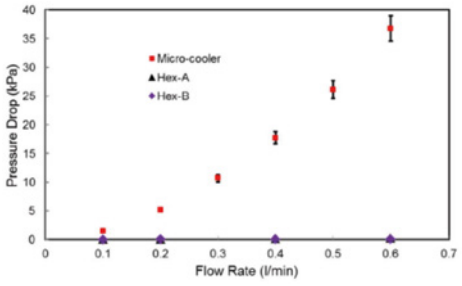

Figure 17. Comparison of the measured pressure drop in microcooler and heat exchangers.

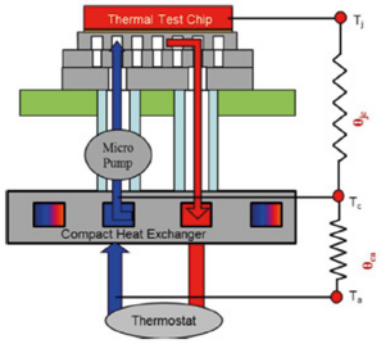

Figure 18. Schematic of the system thermal resistance analysis.

To design the compact Hex-B, the simulation model is first built with 4 × 9 channels in each fluid side with a footprint size of 20 × 20 mm as the design start point. The 6 × 6 cells are assigned in each channel cross section to ensure the computational accuracy, and the grid independent solution is obtained with a mesh number of 1.2 million cells. The simulation results are shown in figure 9. Hex-B with initial configuration, only provides about 110 W heat transfer capability under the specified conditions, which is too low for system design target. Therefore, the design is further optimized to increase the power dissipation capability.

The optimization for Hex-B includes: an increase in channel numbers, an increase in the footprint size and a different the material with higher thermal conductivity. Figure 9 shows the effect of channel numbers and footprint size on thermal performance of the Hex-B. The thermal performance increases as the channel number and footprint size increase. For the heat exchangers with a footprint area of 32 × 32 mm and a channel number of ≥15, their heat transfer capability exceeds the design power of 175W, meeting the design target. The effect of heat exchanger material type is shown in Figure 10. The SS, with a thermal conductivity of 16 W/mK, would have unfavorable thermal performance below the design target, which is not preferred material for Hex-B. Additionally, the aluminum heat exchanger has similar thermal performance as that for the copper heat exchanger. Hence, the aluminum heat exchanger is suggested for practical fabrication due to the light weight.

Based on the above analysis, the optimized heat exchanger has a footprint area of 32 × 32 mm with a channel number of 15 in each row for both fluids, and is made of aluminum. Figure 11 shows the simulation results of the temperature profile for Hex-B with the final design.

To assess the effectiveness of the design of Hex-B, the heat transfer density and the pressure drop of Hex-B are compared with those of Hex-A. Figure 12 presents the comparison of the heat transfer density between Hex-A and Hex-B at the specified conditions. Figure 13 shows the comparison of the pressure drop in Hex-A and Hex-B at the specified conditions.

Compared with the Hex-A, the much higher heat transfer density and the much lower pressure shown for Hex-B, affirm the high effectiveness of this compact heat exchanger. The optimized Hex-B is made of aluminum and includes dual flow channels for both hot fluid and cold fluid. Silicone rubber was used as a sealing ring. The weight of the newly fabricated heat exchanger is 0.23kg with all the connecting flanges, which is much less than that for the Hex-A, which weighs 1.20 kg. Additionally, the footprint area of the newly fabricated heat exchanger is much smaller than that of Hex-A, as shown in figure 14. The optimized compact Hex-B has larger heat transfer area ratio and effectiveness.

Another major objective of the design of the liquid cooling system is to reduce the system pressure drop so as to achieve the compactness of the system by selecting the micropump with low pumping power and small size. The variation of the simulated and measured pressure drops versus the flow rate in the heat exchangers and microcooler are presented in figures 15 and 16, respectively. This shows that the experimentally measured pressure drop agrees reasonably well with the simulated value, and the simulated results are slightly lower (5%~10%) than the measured values. The pressure drop in the newly designed and fabricated compact Hex-B is about half of the pressure drop in the Hex-A. It can also be seen that the microcooler with microjet nozzle and microchannels is the main contributor to the pressure drop of the full system, accounting for around 97~99% of the system pressure drop. The proportions of system pressure drop in the microcooler and heat exchanger are shown in figure 17. This indicates that the hydraulic design of this system is efficient, and the pressure drop in the heat exchanger has been reduced to a reasonable level.

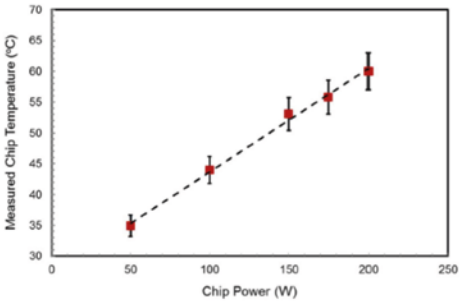

Figure 19. Measured chip temperature through the IR camera for different chip powers.

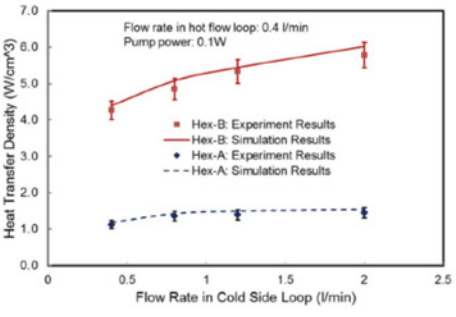

Figure 21. Measured thermal transfer density for commercial heat transfer A and compact Hex -B (a power of 175 W).

The thermal resistance is extracted to evaluate the thermal performance of the full system. The schematic 1-D thermal network of the system is shown in Figure 18. It is seen that the junction to ambient thermal resistance of the system (θ_{ja}) consists of two parts; the first part is the thermal resistance from the junction to microcooler (θ_{jc}), which represents the thermal performance of the microcooler and the other is the thermal resistance from the microcooler to the ambient (θ_{ca}), which represents the thermal performance of the heat exchanger.

The design target for this silicon microcooler system is to have a heat dissipation capability of 350 W/cm^2, which corresponds to a heat power of 175 W on a 7 × 7 mm chip. Assuming the system works in an environment of 40°C, and the allowable junction temperature is 85°C. The design target of the junction to ambient thermal resistance for this system is about 0.25°C/W.

In the tests, the power applied to the thermal test chip varies from 50 to 200 W. The chip temperature under different chip powers is measured, and the results are shown in figure. 19. The flow rate remains 0.4 L/ min in the hot loop and 2 L/min in the cool loop during the test. It can be seen that the chip temperature linearly increases with the increase in the power applied to the chip, while, the overall thermal resistance of the

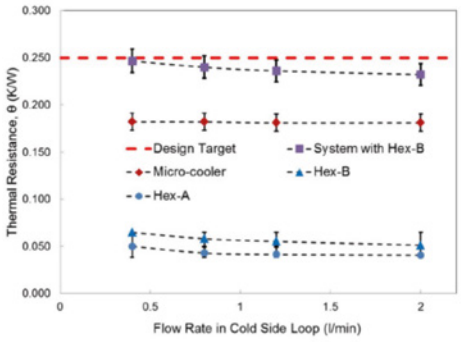

Figure 20. Measured thermal resistance of the heat exchangers, microcooler, and the full system assembled with compact Hex -B (a power of 175 W).

system, including the thermal resistance of the microcooler and the thermal resistance of the heat exchanger, has been obtained. The results are shown in figure 20. The thermal resistance of the heat exchangers decreases when the flow rate in the cool flow loop increases, as such causing the decrease in the overall system thermal resistance with the increase of the flow rate in the cold flow loop. Taking a close examination to figure 20, it shows that the microcooler is the main contributor to the overall system thermal resistance, accounting for around 80~90% of the total thermal resistance, while the heat exchanger only accounts for about 10~20% of the overall system thermal resistance. Additionally, the thermal resistance of Hex-B is slightly larger than that for the Hex-A. While Hex-B is much smaller and lighter than Hex-A. Furthermore, the pressure drop in Hex-B is only half of the pressure drop in Hex-A. Furthermore, as shown in figure 21, the heat transfer density of Hex-B is much higher than that for Hex-A. Hence, it can be concluded that the design of Hex-B is more effective than the design of the Hex-A for the applications in this paper.

Furthermore, the developed compact cooling system is able to meet the design target of 175-W heat dissipation capability with a pumping power of 0.1 W.

Reference:

IEEE TRANSACTIONS ON COMPONENTS, PACKAGING AND MANUFACTURING TECHNOLOGY, VOL. 6, NO. 5, MAY 2016

Development of a Compact and Efficient Liquid

Cooling System With Silicon Microcooler for

High-Power Microelectronic Devices

Gongyue Tang, Yong Han, Boon Long Lau, Xiaowu Zhang, Senior Member, IEEE, and Daniel Min Woo Rhee

Huawei Delivers Outstandingly Accurate Models

By Yake Fang, Senior Engineer, Huawei Technologies Co., Ltd.

Packaging high-performance multi-core IC devices used in communication applications is a key challenge for both manufacturers and system integrators. Traditionally a System-in-Package (SiP) has been taken, with chips mounted side-by-side, allowing differing semiconductor technologies to be mixed. More recently stacking silicon die has become more commonly used, as volumes have increased, making the System-on-Chip (SoC) design approach, which has much higher upfront non-recoverable expenditure, the cheapest option overall. These multi-heat source devices require very careful thermal design frontloaded in the design process, due to the very high cost of changes during detailed design.

The increased functionality and power dissipation has made the thermal design of communication device applications crucial in today's increasingly high computing power communication device industry. These products present a very challenging thermal environment for our SoC packages, with the need to design to very tight thermal tolerances to meet size, weight and form factor goals.

Huawei always pays attention to thermal challenges in product design and uses thermal analysis throughout the whole product R&D phase. We have the ambitious program to develop an internal thermal design flow for multi-core SoC and SiP devices that would enable us to use the highest possible quality thermal data and models to facilitate our own design efforts.

We have used FloTHERM™ for almost 15 years, so we have a rich depth of experience and so are able to create detailed thermal models of our SoC and SiP products. However, we needed a way to test that these packages achieved the 'as-designed' thermal performance, and to test packages from other suppliers rather than rely on their thermal data, which is often not sufficiently accurate or suitable for design. As it is based on the thermal transient measurement principle, giving the highest available fidelity, we turned to Mentor Graphics' T3Ster® hardware to perform a series of thermal tests so our thermal engineers could better understand the thermal impedance characteristics of these high-performance products, and reduce time and resources needed to achieve a good design.

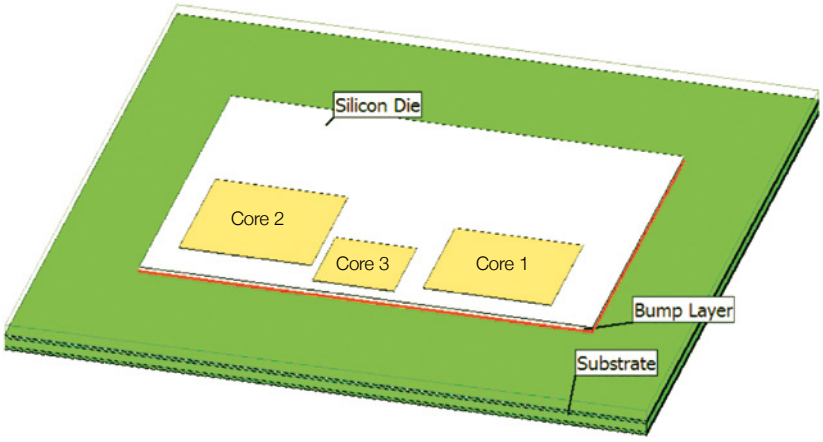

Figure 1. Overview of SoC device construction

As a test vehicle for this new design workflow we chose to investigate a 3-core SoC device typical of those now used in communication device applications and assembled in a Package-on-Package (PoP) format that allows package stacking. The FloTHERM model is shown in figure 1.

As soon as the first working packages were produced, they were measured using Mentor Graphics T3Ster hardware, which records the temperature response of the part to a power step from just a few microseconds until steady-state is reached, with an accuracy of ±0.01°C.

Having creating the detailed numerical model of the package in FloTHERM we performed a transient simulation to get a baseline result for this, as yet, uncalibrated model, which we built using the best available information during the design. The temperature vs. time graphs we got from FloTHERM and T3Ster seemed to match quite well.

Using a model of the package in the same thermal environment in FloTHERM allows the temperature vs. time response from the FloTHERM simulation to also be converted into a Structure Function, and then the two Structure Functions compared. This is a plot of the cumulative thermal capacitance vs. the cumulative thermal resistance that the heat experiences as it leaves the package. If the model exactly matches reality for the package, then the two Structure Function curves should match perfectly, as the heat flow paths from the die junction out to the environment should be identical.

We then post-processed this temperature vs. time data to create a Structure Function graph for the T3Ster and FloTHERM results, and compared these. We were not expecting an exact match, as there are always uncertainties surrounding material thermal properties. The thickness and uniformity of bond lines,

"To design successful communication devices, Huawei first performs thermal simulations in FloTHERM, and then calibrates the detailed thermal model based on accurate thermal measurement data obtained from T3Ster. In parallel with this Huawei optimizes the design through different material selection to decrease thermal resistance, confirmed using T3Ster's transient thermal measurement. This process is much faster and requires fewer resources than we needed before, and provides Huawei with a highly-accurate thermal model to use in our system-level modeling, helping us get products to market faster than our competitors."

Fang Yake, Senior Engineer, Huawei Technologies Co. Ltd.

and contact resistances within the structure where one material is not in perfect thermal contact with its neighbor materials. However, we were somewhat surprised at how big the differences were when we compared Structure Functions, as shown in figure 2. The Structure Function comparison is a very powerful tool for showing up differences between the model and the experimental result, and where those differences are in the heat flow path.

Figure 2 shows the Structure Functions for the initial FloTHERM simulation (black) vs. the T3Ster measurement data (blue), shown here when Core 1 is powered. The two Structure Functions are clearly quite different. At the far right of the plot the thermal capacitance becomes infinite, as heat reaches the environment. For the FloTHERM model this occurs at a lower thermal resistance than found in practice, so the model is underestimating the temperature rise. Some of this difference may be due to deficiencies in the way the cold plate is modeled. However, the differences seen on the left of the graph are within the package model itself, and close to the heat source. Simulation vs. test results, when cores two and three are powered, both showed similar discrepancies.

To modify the model manually to bring it in line with the experimental results would be very time consuming and error prone as it would be a guess-and-correct process that is impractical to undertake a fast-paced design environment like ours. However, part of our reason for choosing T3Ster is that its results can be read into FloTHERM's Command Center tool, which allows automatic calibration of multiple user-selected parameters to adjust the model to match the measurement data. This calibration gives a very good fit, something that would be extremely difficult and time consuming to do manually, and so results in a model that has very high accuracy for us to use in system design.

To do this automatic calibration, we selected the time range corresponding to heating

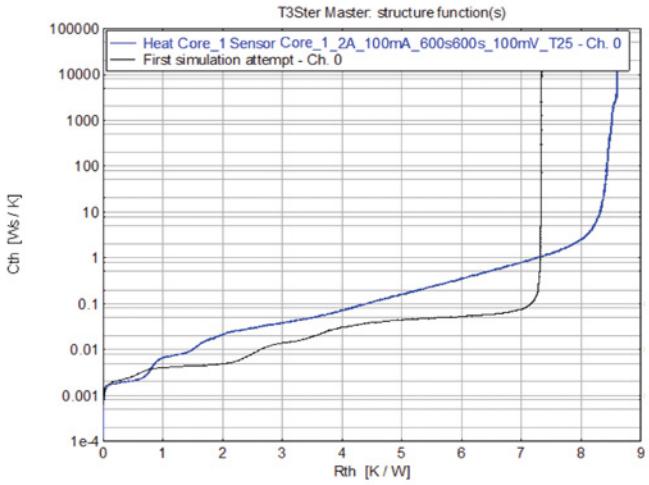

Figure 2. Initial simulated curve (black) vs.curve measured with T3Ster (blue)

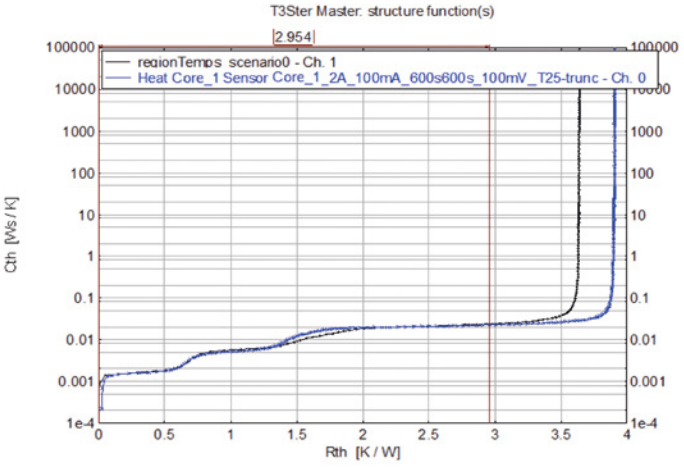

Figure 3. Final Calibration Result Showing Excellent Match

within the package, ignoring the part of the curve that represents heating of the board and the cold plate, so truncating the transient calculation to 0.15 s. The optimization requires a list of parameters, the exact values of which are uncertain, and a possible range for their thermal conductivity values. The parameters selected were the C4 bumps, with a conductivity ranging from 40 W/mK to 70 W/mK; the underfill (0.2 W/mK to 0.7 W/mK; the package substrate (10 W/mK to 20 W/mK); and the solder mask (0.5 W/mK to 3 W/mK).

The result, shown in figure 3, gives an excellent match within the thermal resistance range up to the junction-to-board thermal resistance (RthJB). This has proved to us the effectiveness of FloTHERM's auto calibration technology coupled with T3Ster, and the calibrated model can be used for board-level and system-level thermal analysis with a very high degree of confidence in the fidelity of the results.

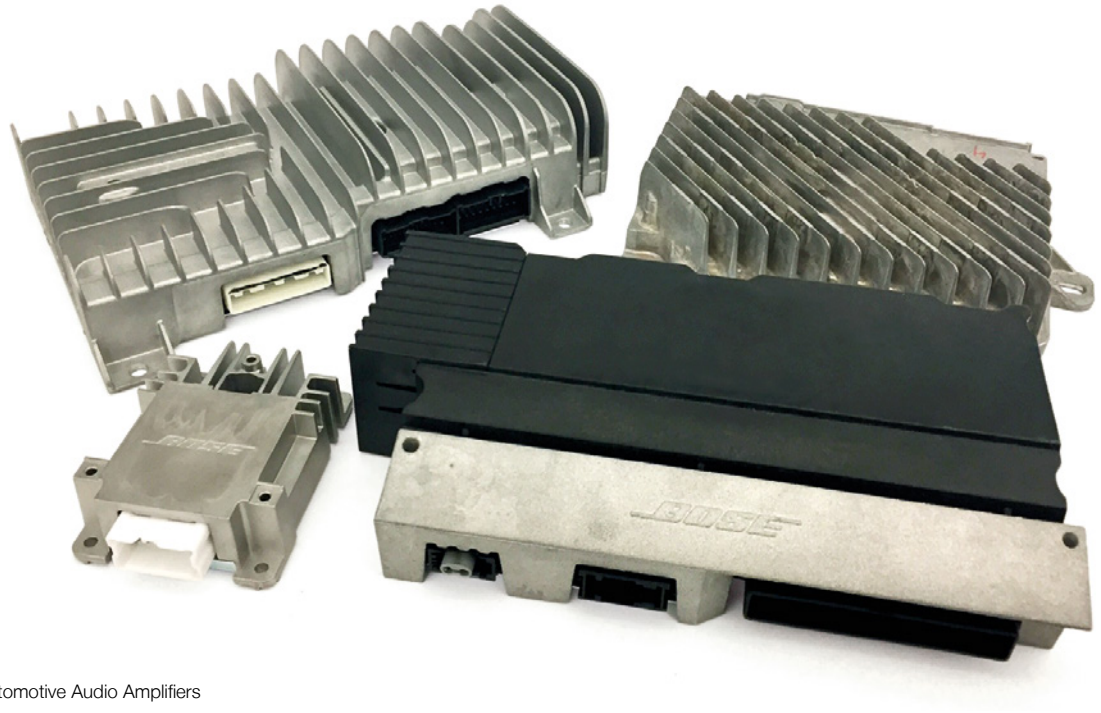

Figure 2. BOSE Automotive Audio Amplifiers

Bose Automotive Systems

Finding new ways to make our Products Cool!

By Brad Subat, Mechanical Engineer, Bose Automotive Systems

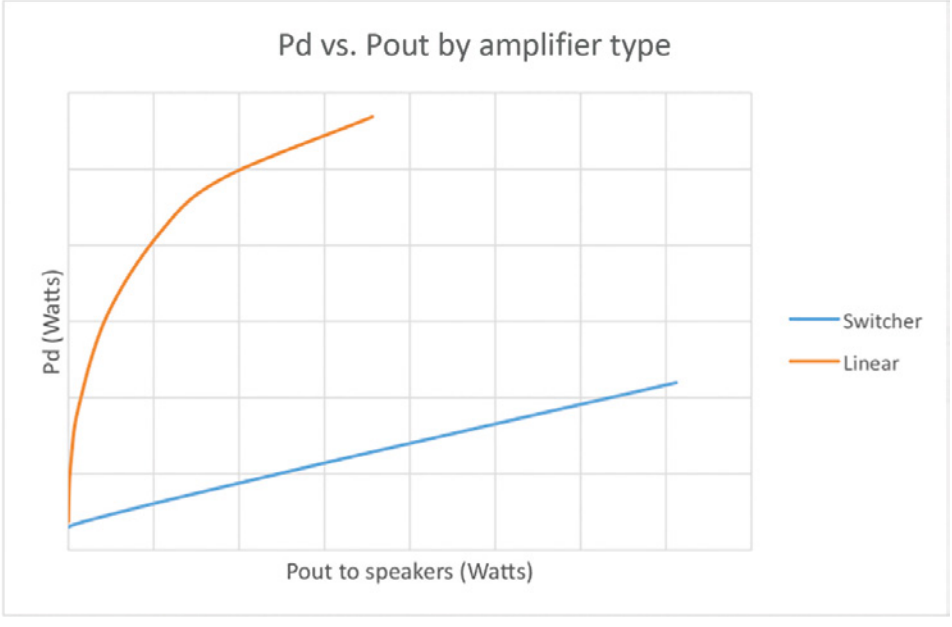

Figure 1. Power Dissipation vs. Amplifier Output Power

Like many industries these days, it is vital to maintain class leading technology and deliver products to market in less time. For Bose, one instrumental tool has been Mentor Graphics thermal analysis software, first FloTHERM and more recently FloEFD. We switched to FloEFD since it handles our more complex geometry and has the benefit of being embedded in our CAD tool (NX). While these tools provide great insight into the thermal design of products, we enter a new project with a bias to use previous designs which brings us to this article.

More than a year ago, I was intrigued by the first Mentor Graphics blog (of what became a several part series) on "Organically Grown 3D Printable Heatsinks". So, while the example was nothing like our actual products, I could see that this technology (once matured) could help us

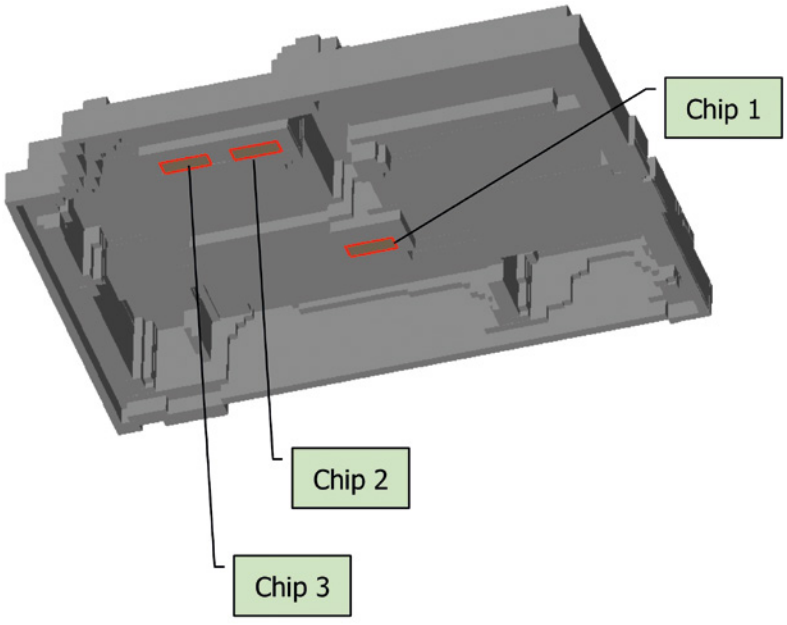

Figure 3. Power Amplifier Heat Source Locations

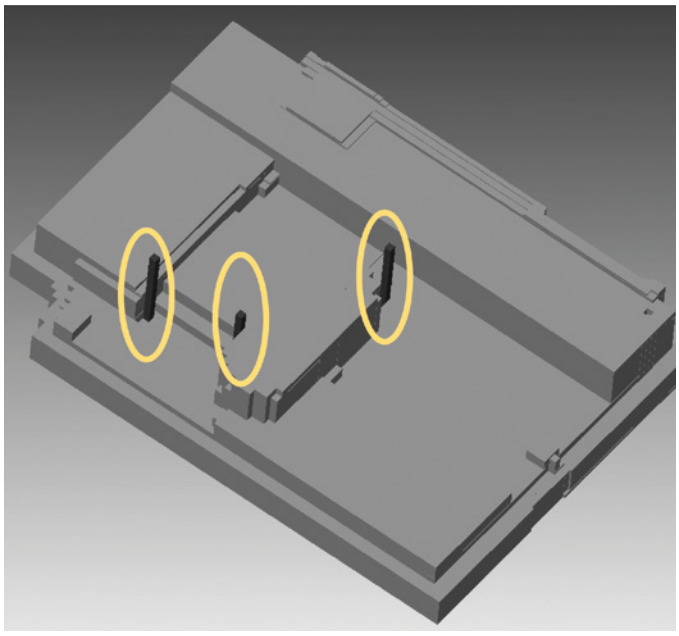

Figure 4. Starting Rod Locations

meet our goals in future years. At Bose, we believe we have "Better products through research", so while a tool might be effective today, we are drawn to developing better tools for the future. I then reached out to Mentor Graphics to propose we start an industrial research collaboration where this automated "growth" approach could be applied to something closer to our products. They would focus their algorithm development on our design space, while Bose would provide feedback and testing at critical progress points.

Designing products for the automotive audio market has many constraints and competing goals. In respect to our amplifiers, we design them to work within a wide temperature range, with high product reliability and nearly perfect initial quality. In the thermal realm, there are a few main components that dominate our designs; power amplifiers, voltage regulators and digital signal processors. Traditionally, the largest power dissipating (Pd) devices are power amplifiers. There are two main technologies that are used in the automotive industry, linear (Class AB) and switching (Class D), figure 1 Linear power amplifiers are less efficient but historically have been more cost effective. In recent years, the industry has focused on improving the cost effectiveness of switching amplifiers and the cost gap has narrowed significantly. The graph below shows the power dissipation differences versus power out for these types of power amplifiers. On both axes, there can be endless discussions on what

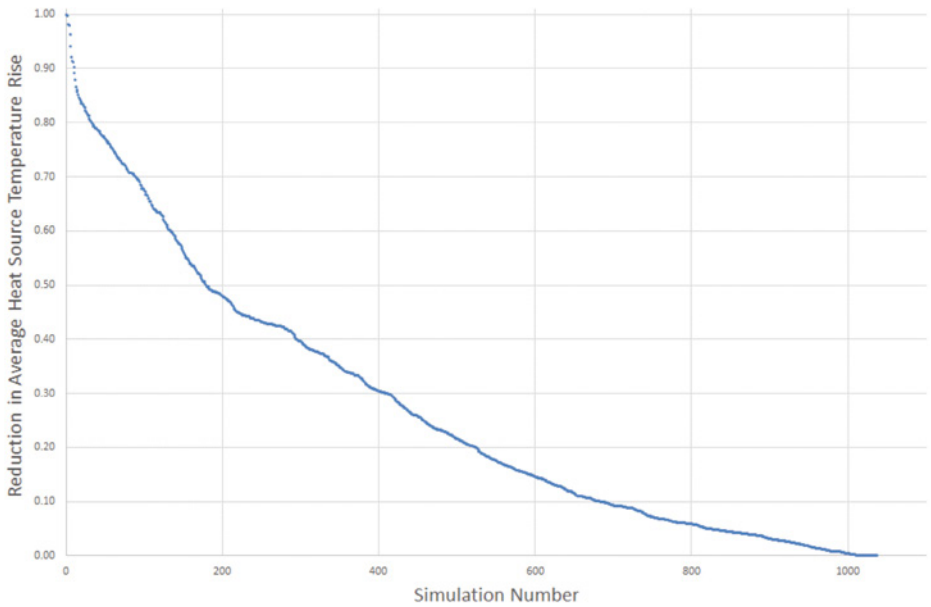

Figure 5. Reduction in Average Power Amplifier Temperature During Growth Process

variables and conditions should be used to produce this data, so, focus on the relative differences.

For the purpose of this initial investigation, the simplified model only had linear power amplifiers as components (figure 3) and a uniform heat source was spread over the simplified PCB.

Bose's current design methodology for our amplifiers is to use die cast housings (both aluminum and magnesium) for five, of the six, sides of the enclosure. The blank canvas of this growth study was a minimum nominal wall shell around a typical size PCB. We included a few thick sections for heat spreading and added three starting rod locations (figure 4) from which the algorithm builds.

In automotive, part re-use is important, so car companies typically choose to have us design a housing that can work in several orientations. This allows them to have design freedom from vehicle to vehicle in regards to the amplifier location and orientation. During a standard development

many orientations are considered. To simplify this initial study, only the fins up axis orientation was considered for "growth" since it usually is the hottest orientation.

Figure 5 shows how the heat source temperature drops, as the heatsink grows. The sharper the initial drop, the greater the effectiveness of the added material. As the growth continues, the rate of temperature reduction slows. If the slope becomes very small, this indicates that you are having to add a lot of material for a small temperature reduction. An ideal thermal design would still have a large slope once the thermal goals have been reached.

Figure 6 shows a few steps of the growth.

With the simulations in and showing favorable results, it was time to do actual thermal testing to see if this innovative simulation technique would correlate well. Before machining our test housings, the grown design was adapted in our CAD tool to account for high level die casting design rules:

- **Fin spacing:** Tool steel is the inverse of your part, so between fins there needs to be enough tool steel to provide cooling and strength during the die-casting process. If pushed too far, the standing steel can break during production requiring major repair or early decommissioning of that cavity.
- **Draft:** Like plastic injection parts, die-casting needs a taper so it can be pressed out of the tool.
- **One area not addressed, was mold fill.** In high volume, automotive die casting process use an extremely short high pressure process to inject molten metal into the tool that can cool or trap air in a complex geometry.

As seen in figure 7, these results are very encouraging for this additive technology, as the unaided design used less material and has a similar thermal result. For thoroughness, we also tested the machined housings to confirm our selected orientation was indeed the hottest.

As with any research, it leads too many other areas to explore. So as Mentor Graphics was working on growing a heatsink, they decided to work on an optimization algorithm in parallel, which was subtractive by removing material from the design. This algorithm was applied to the "grown" heatsink tested above. The graph figure 8 to the right shows the increase in

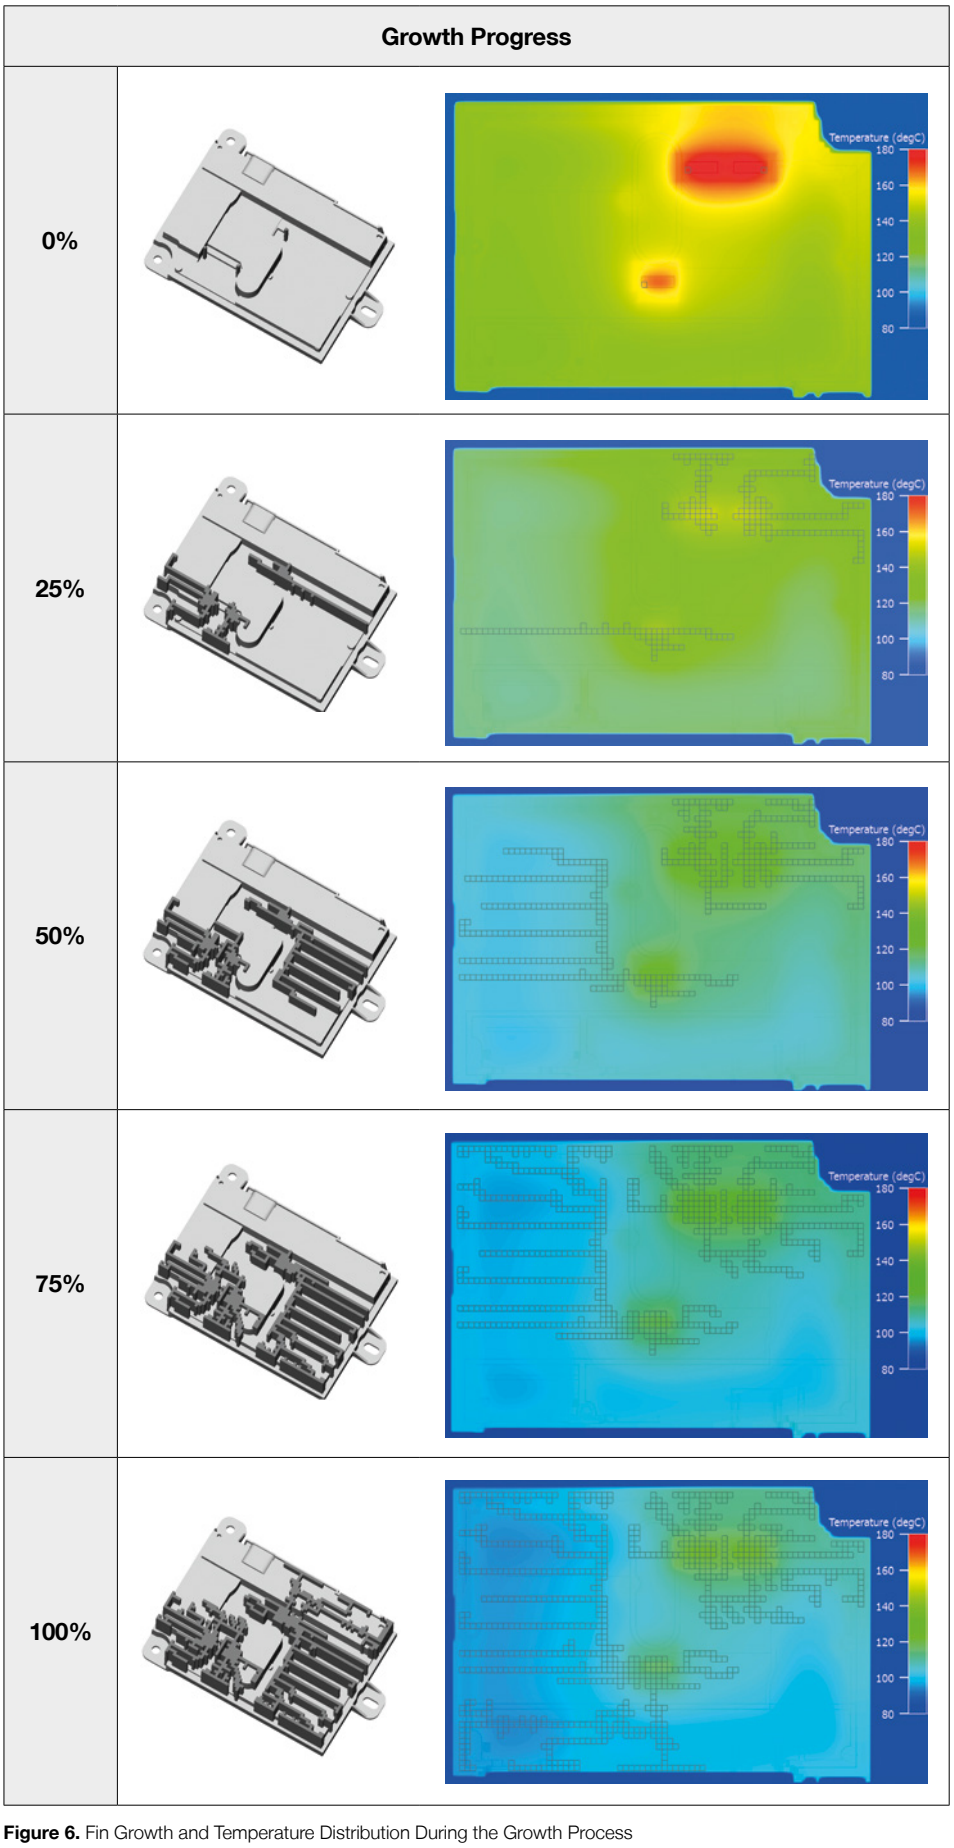

Figure 6. Fin Growth and Temperature Distribution During the Growth Process

"What was seen as an interesting idea over a year ago by Mentor Graphics and Bose for a autonomously grown heatsink, has now proven to be possible technology for the future."

Brad Subat, Mechanical Engineer, Bose Automotive Systems

Predictions and results:					
	Heatsink volume (mm^3)	**Fin only surface area (mm^2)**	**PA1 (Cº)**	**PA2 (Cº)**	**PA3 (Cº)**
Existing design (Prediction)	268,475	100,904	119	124	117
Grown (Prediction)	236,833	76,921	120	122	114
Delta	12% less	24% less	1 Hotter	2 Cooler	3 Cooler

Figure 7. Existing Fin Design vs. Grown Fin Comparisons

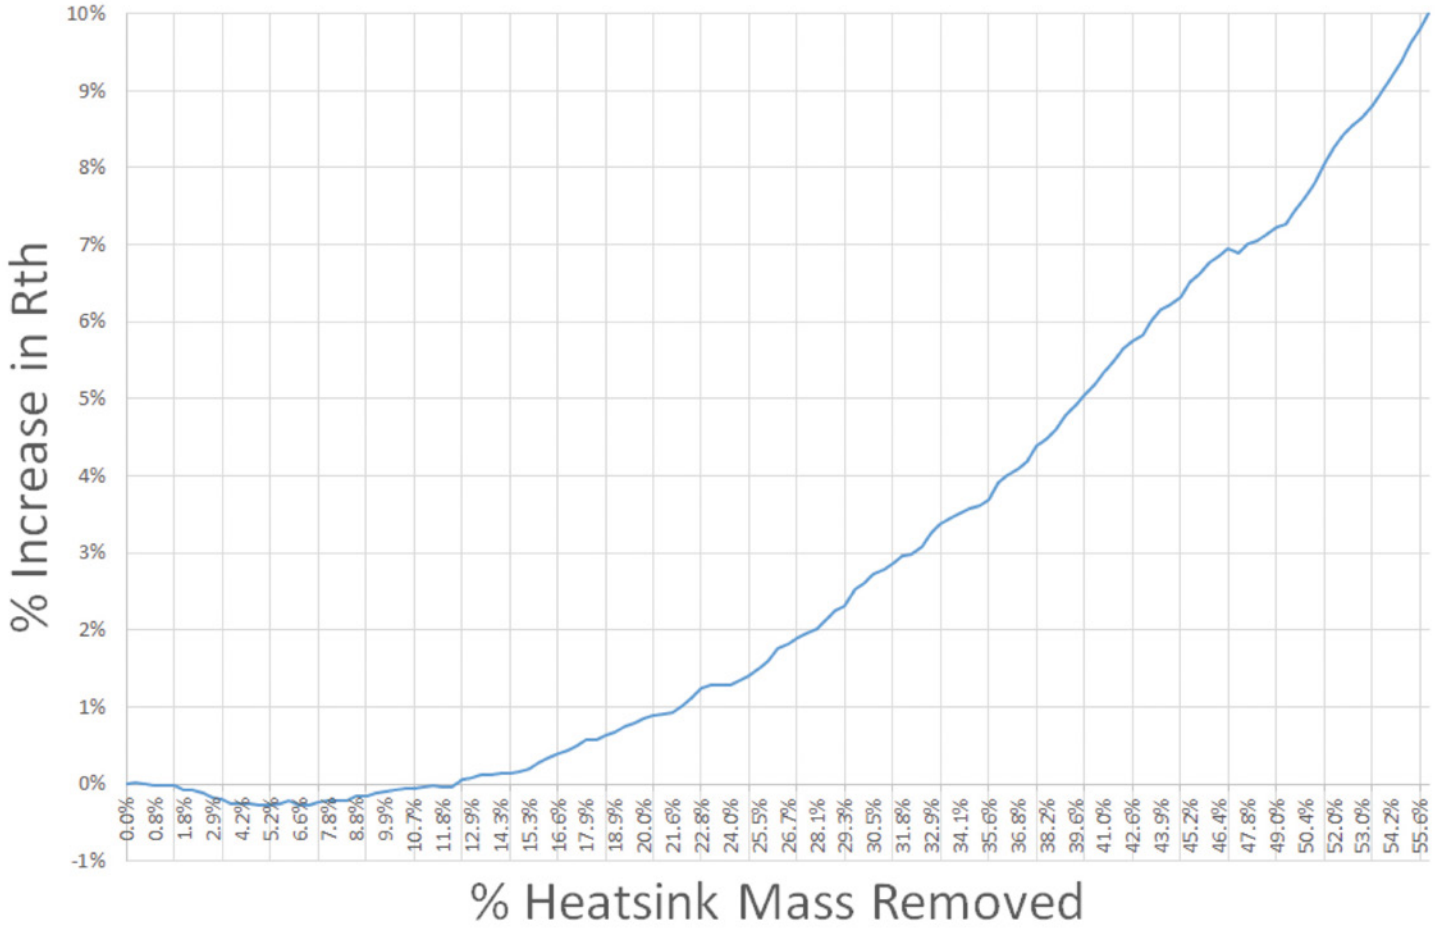

Figure 8. Increase in Thermal Resistance vs. Heatsink Mass Removed

the housings thermal resistance versus the amount on heatsink fin material removed. In this specific design the initial removal of heatsink material is predicted to actually lower the thermal resistance (cooler). Figure 9 shows the optimization of the fins and predicted temperatures at several stages.

We decided to test at 22% since there was a small leveling in the rise of thermal resistance. (Figure 10)

Again, the results are very encouraging for this subtractive algorithm, as heatsink material can be removed for minimal temperature rise. So if you have a product with thermal margin and are looking to optimize it, this technique might be immediately interesting for you.

A special thank you to the Mentor Graphic team this past year for their hard work, vision and passion for this collaborative project. What was seen as an interesting idea over a year ago by Mentor Graphics and Bose for a autonomously grown heatsink, has now proven to be possible technology for the future. I look forward to the next steps in refinements from Mentor Graphics to see if we can establish the confidence needed to add to our product development tools in the future.

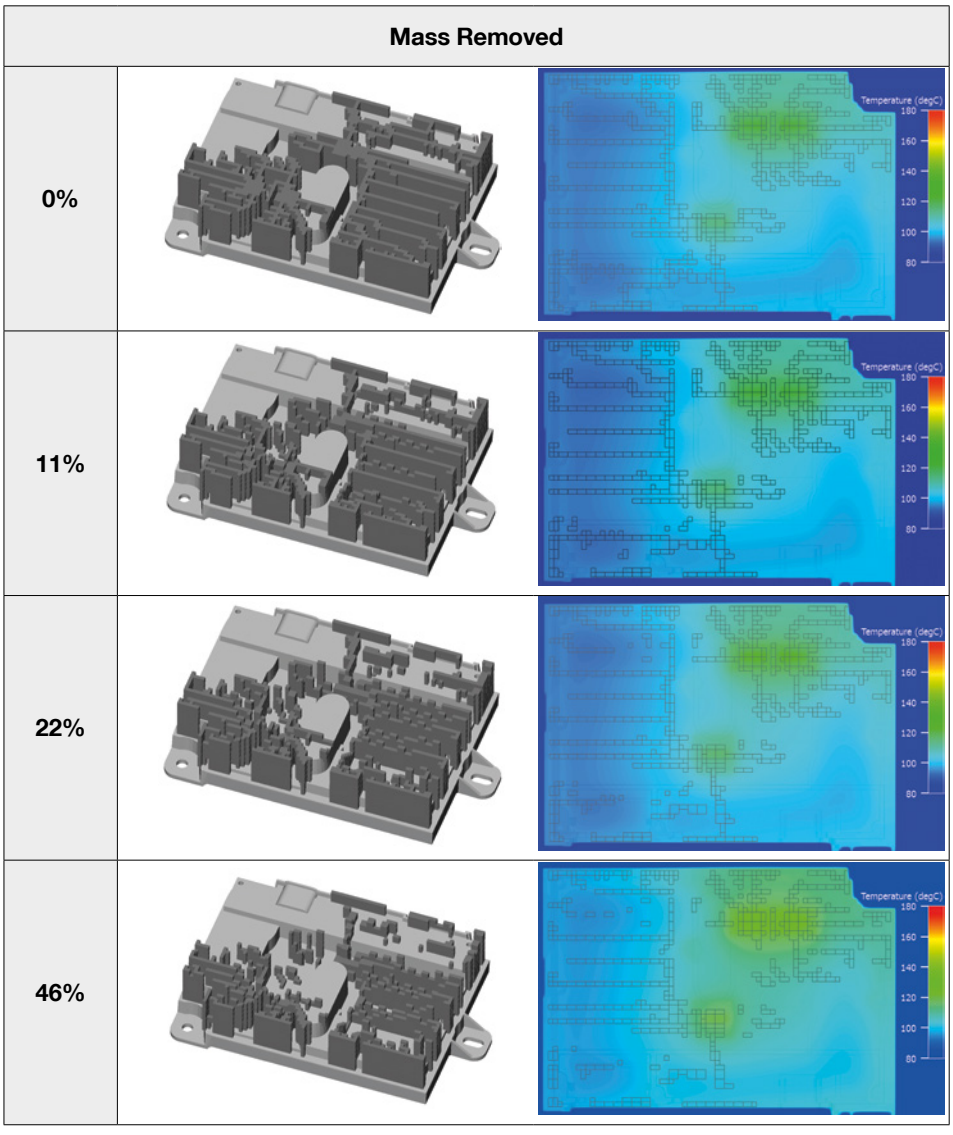

Figure 9. Fin Geometry and Temperature Distribution (Hardly Changing) During the Fin Mass Reduction Process

	Heatsink volume	Fin only volume	PA1 (C°)	PA2 (C°)	PA3 (C°)
Existing design (Prediction)	268,475	100,904	119	124	117
Optimized (Prediction)	219,114	62,268	122	123	117
Delta	18% less	38 % less	3 Hotter	1 Cooler	No Change

Figure 10. Grown Fin Design vs. Mass Reduced Fin Design Comparisons'

The growth methodology involves a series of sequential FloTHERM simulations. The pin with the largest thermal Bottleneck number is identified and another pin added to its coolest side. The process is then repeated. If the addition of a pin causes an increase in the thermal resistance then that pin is removed and a pin addition to the next largest thermal Bottleneck location is attempted instead.

The subtraction methodology starts with a tessellated discretized definition of the heatsink. An initial simulation identifies the portion with the lowest thermal Bottleneck, that portion is removed, the model re-simulated and the change in thermal resistance noted. The process is repeated, building up a graph relating the change in thermal resistance to the decrease in heatsink mass.

Both methods are implemented using an Excel front end, VBA and FloTHERM's FloXML and FloSCRIPT technologies. The processes are fully automated, not requiring any manual intervention due to the stability and robustness of FloTHERM's meshing and solver.

A Unique Design to Generate UAV Electrical Power in Flight

By Guy Wagner, Electronic Cooling Solutions and Travis Mikjaniec, Mentor Graphics

As smaller UAVs are designed with more sensors and communications technology for longer missions, the additional electrical power to run them drives the need to generate onboard electric power. One way to create onboard electrical power would be to harness the remaining 80% "waste energy" produced by the two-stroke engine.

A team of engineers from Electronic Cooling Solutions and Ambient Micro designed, built, and tested an exhaust-heat thermoelectric generator (EHTEG) that can be incorporated into a UAV design to harvest and convert this waste energy into electrical power in flight [2] .

There are several sources of energy loss in a small engine all in the form of heat. The main sources are the heat rejection from the cylinder and cylinder head, loss from friction, and the heat of the exhaust stream. Though the heat from the exhaust system is the only truly usable heat source.

The EHTEG had several requirements:

- Mechanically robust and integrate into the aircraft without compromising flight safety,
- Extract the required heat without impairing engine performance,
- Provide the largest possible temperature differential across the thermoelectric modules while operating within the maximum temperature limits, and
- Designed with minimal weight and aerodynamic drag.

Figure 1. UAV with the EHTEG attached on top

Designing the Interior and Exterior TEGs

Heat exchangers on the inside of the muffler absorb heat from the exhaust as it flows through. The heat passes through exchangers to 2inch. square TEGs mounted on the outside of the UAV and finally passes through another row of heat exchangers to the open air. As the TEGs are exposed to the temperature difference between the hot inside exhaust air and the cool outside air, they generate electric current. The greater the temperature difference, the more current is generated.

The team modeled the thermal design of the system using Mentor Graphics' FloTHERM® Computational Fluid Dynamics (CFD) modeling software [1]. They simulated airflow on the outside (cool air) and the hot exhaust inside to estimate the temperature difference, which enabled them to optimize the internal and external fins of the heat exchanger and the number and location of the TEGs.

They built engineering models of several likely EHTEG configurations and ran them on a test stand using the same engine and propeller that is used in a MLB Company Bat4 UAV. The models were validated for a range of operating parameters that simulate flight conditions

The engineers used FloTHERM software that models conduction, convection, and radiation as well as the fluid flow of both the external cooling air and the hot internal

exhaust gas. They created virtual models of the exhaust system and performed thermal analysis and test design modifications quickly and easily before building any physical prototypes. The results of the CFD models correlated well with those obtained on the engineering test bed.

The engineers used the internal volume and length for the muffler recommended by the manufacturer to make the system act as an efficient expansion chamber exhaust system. These were used to develop the first half-symmetry CFD models that would determine the number of TEGs needed to optimize the electrical output with minimal weight. The model is symmetrical, so building a half-symmetry model reduced the number of elements down to 1.04 million cells with no loss in accuracy.

The fin parameters of the internal and external heat exchangers were kept relatively constant while varying the location of the heat exchangers and the placement of the TEGs. The goal was to maximize the temperature differential across each TEG to extract the most heat energy from the 455°C exhaust gas.

13 configurations of heat exchangers and TEGs were modeled in the first optimization study. They tabulated the power output from each configuration and chose the best configuration.

After optimizing the placement of the TEGs, it was found that the central fins on the interior heat exchangers disrupted the exhaust flow and the hot exhaust gas wasn't reaching the front end of the muffler. But if they removed the center fins, the exhaust gas would not channel down the center and the exhaust pulse would not reach the front end of the muffler. So instead of removing the central fins, they placed them in a semicircular pattern. (Figure 3) This configuration kept the exhaust pulse moving through the center of the muffler and it was able to curl back symmetrically as the hot gas flowed back along the outsides and through the fins.

The interior heat exchangers decreased in temperature as the hot exhaust gas flowed from the front through the fins to the rear of the muffler and then out of the vertical exhaust pipe.

The FloTHERM simulation was able to show the exhaust gas flow pattern for the heat exchanger with the center fins removed. (Figure 4) The flow pattern is disrupted

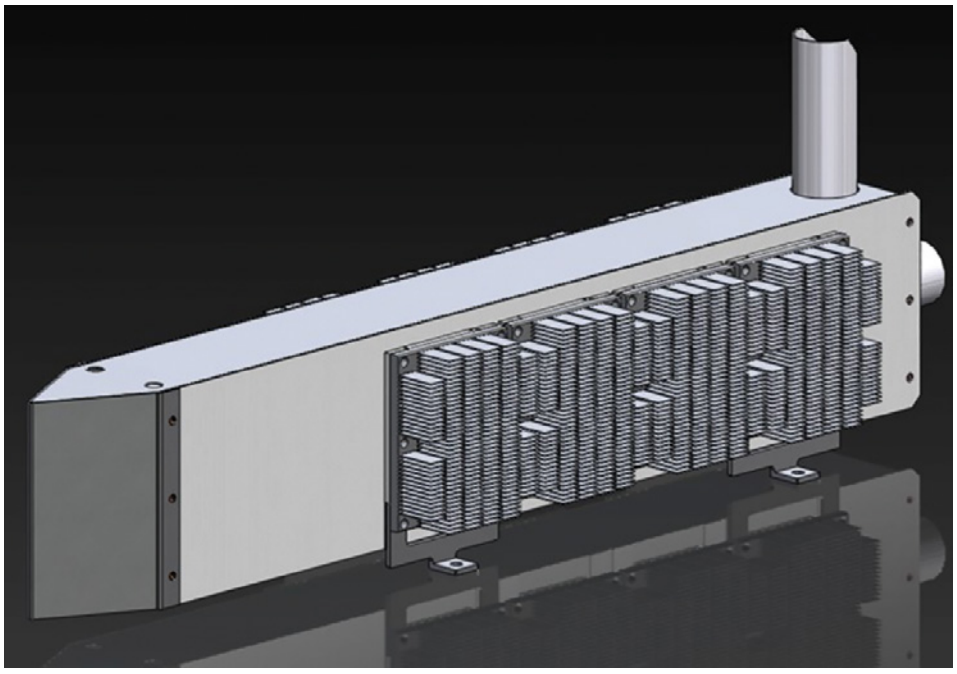

Figure 2. The positions of the external heat exchangers that resulted in the highest average power generation for the system

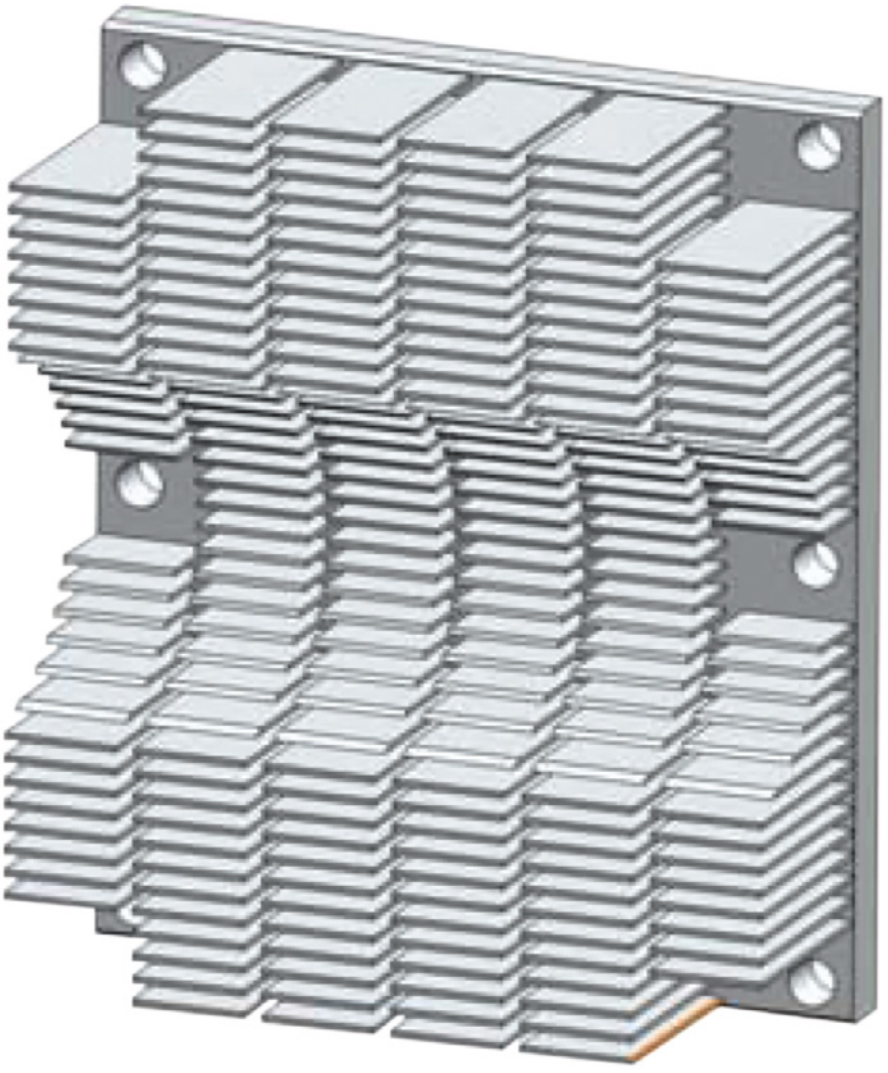

Figure 3. The interior heat exchanger with the fins in a semicircular shape in the center

before the exhaust stream reaches the front end of the muffler.

The simulations showed that when the semicircular-shaped fins are present, the flow pattern retains its shape all the way to the front of the muffler. (Figure 5) This allowed them to optimize fin placement.

Typically, the outside heat exchangers on TEGs are placed in a line. This causes an issue because the units toward the rear of the external airflow receive more preheated air than the heat exchangers that are upstream. This reduces the delta-T across the heat exchanger and the power output.

Therefore, nine configurations were analyzed to determine the optimum fin parameters for the external heat exchangers. The results were plotted against the total power generated by the TEGs. All of the custom heatsink designs analyzed out-performed the stock heatsinks for total power generation, but due to time and budget constraints, stock heatsinks were used for the first flight test model. This resulted in a reduced power output 11W compared to the custom heatsinks.

The outside TEGs were arranged in four columns by two rows on each side of the muffler. The TEGs maintained a cool side temperature below 58.3°C with external air at 18 °C and 22.3 m/s velocity, while keeping the hot side temperature below the maximum allowable temperature of 225°C. (Figure 6) The outside heat exchanger temperature ranges from 26.1°C on the leading edge of the front fins to 58.1°C on the base plate next to the hottest TEG.

Some of the first simulations demonstrated that the heat loss through all other surfaces of the muffler had to be minimized to maximize heat flow through the TEGs for maximum power generation. Mineral wool insulation was used on all the exposed surfaces to minimize the heat loss.

The power output of the EHTEG system was modeled by summing the power contribution of each pair of TEG modules. By first calculating the hot side and cold side temperatures of each TEG pair, these values could then be used to compute the open circuit voltage. The harvesting and power conditioning circuitry matches the equivalent series resistance for maximum power transfer; thus, the voltage of the load resistance is exactly half the open circuit voltage. This data defines the power harvested per TEG pair.

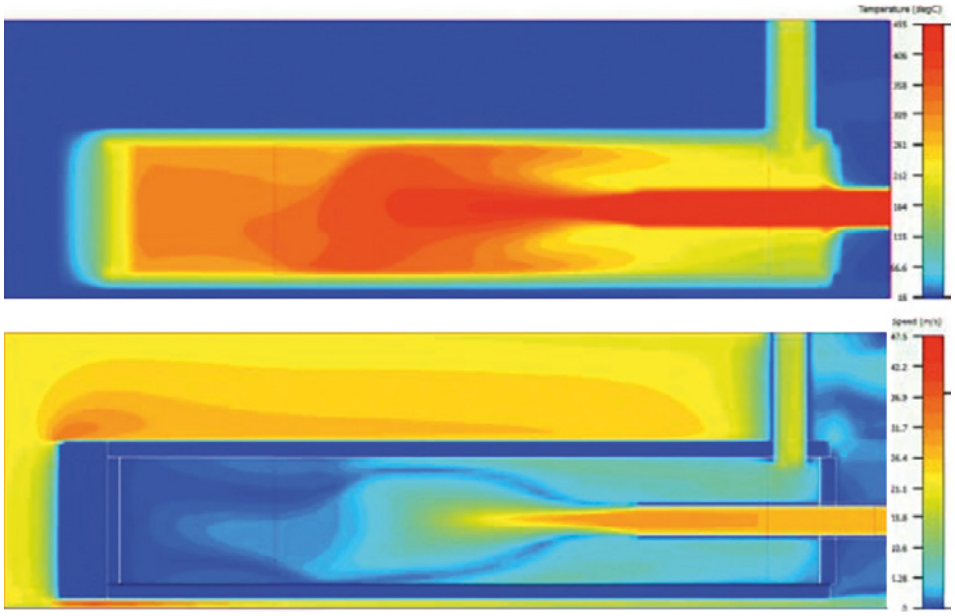

Figure 4. Thermal simulation shows the exhaust gas flow pattern without fins in the center

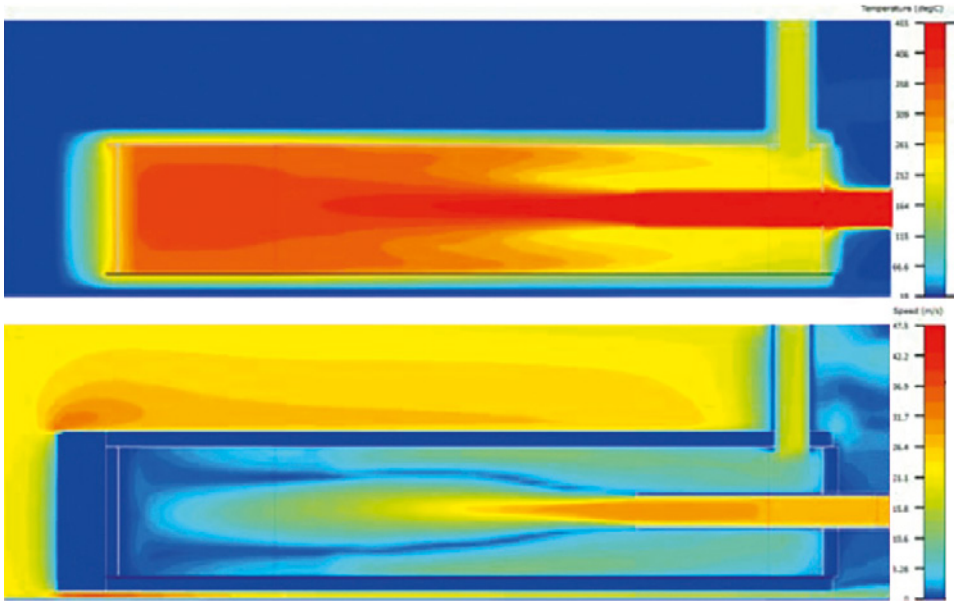

Figure 5. Thermal simulation shows the exhaust gas flow pattern with fins placed in a semicircle in the center

Mechanical Design

The mechanical design of the EHTEG had two main requirements:

- Able to support the thermal components and allow adequate flow of cold air around the external heatsinks, and
- Be lightweight with minimal frontal area to reduce aerodynamic drag.

The heatsinks were fastened from the outer (cold side) to the inner (hot side) forgings. The TEG modules were sandwiched between the outer heatsinks and the EHTEG sidewalls. They used a high-temperature thermal interface material at each interface in the thermal path to maximize heat transfer from the inner heatsinks through the thermoelectric generator modules and the outer heatsinks. Figure 8 shows an exploded view of the EHTEG.

Converting to Usable Electricity

The thermal environment of each thermoelectric module was slightly different because of its location on the EHTEG, so each module's output was also different from its neighbors. The input interface modules received the output from a pair of

modules and converted the input voltage to 12 V. The input module also automatically adjusted its input impedance to match the source impedance, thus operating at the maximum power transfer point.

The outputs of the input modules were combined and fed to a single 12-V bus regulator that provided a regulated 12-V output to external loads. An electronic load was included in the power conditioning electronics for testing purposes. The electronic load automatically adjusts its resistance to extract the maximum power available from the EHTEG system. The data I/O board provided voltage levels proportional to selected voltage and current levels for input to the onboard data-logging system.

In Summary

The team used FloTHERM to create virtual models of the exhaust system, analyzing various design configurations quickly before building any physical prototypes. And the results of the CFD models correlated well with those obtained on the engineering test bed. Since the TEGs are actually thermoelectric coolers run in reverse, their efficiency is only around 5%; that is, 5% of the heat energy flowing through is turned into electricity. If this efficiency rate can be doubled, the technology could be used in many practical and profitable applications. New commercial opportunities are spurring interest in thermoelectric power generation. The design techniques described here could be used to develop much higher power output thermal energy harvesting power systems.

References

[1] FloTHERM, http://www.mentor.com/products/mechanical/products/flotherm

[2] J. Langley, M. Taylor, G. Wagner, and S. Morris, "Thermoelectric Energy Harvesting from Small Aircraft Engines," SAE International, 2009.

[3] Marco Nuti, Emissions from Two-Stroke Engines, Society of Automotive Engineers, Inc., Chapters 7 and 9.

[4] Combustion Products Applet: http://www.engr.colostate.edu/~allan/thermo/page9/Combustion/combustion.html, Allan T. Kirkpatrick, Colorado State University, Fort Collins.

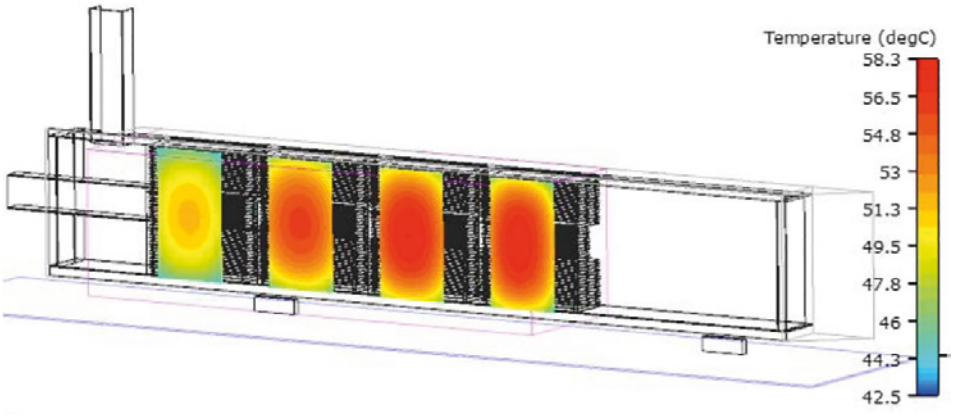

Figure 6. Cool-side temperatures of the outside TEGs.

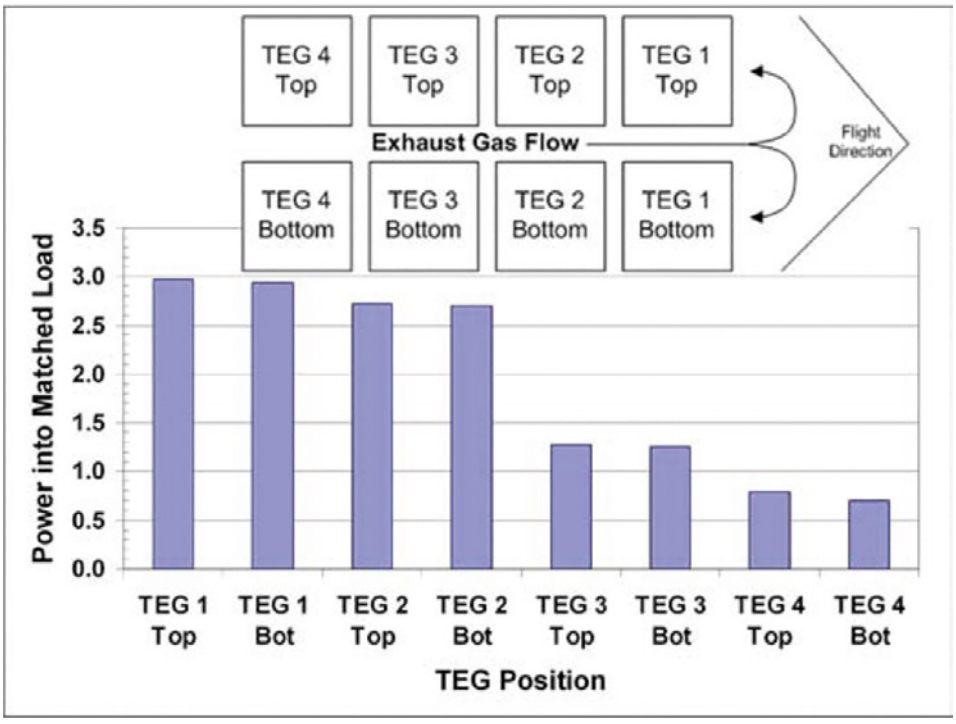

Figure 7. Single-side power output (W) for the flight configuration

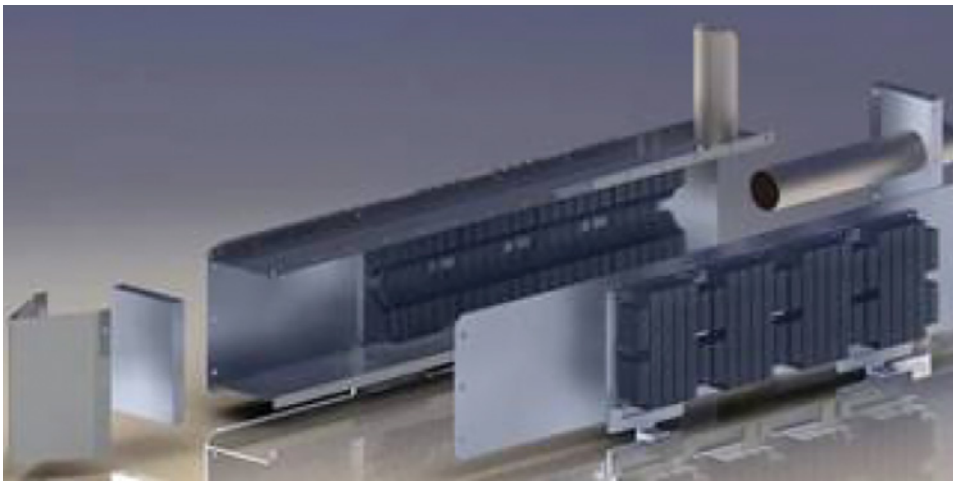

Figure 8. The inside and outside of the EHTEG

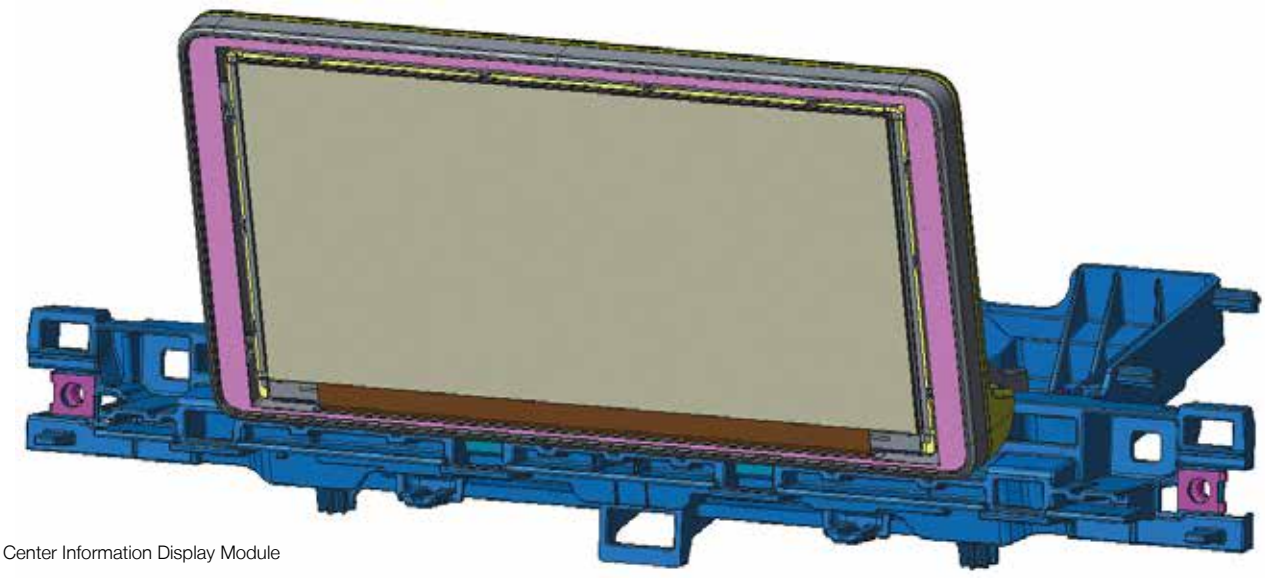

Figure 1. Center Information Display Module

Thermal Design Approach for Automotive Display Integration

By Clemens J.M. Lasance, Philips Research Emeritus, Consultant SomelikeitCool

Innolux Corporation is a world leading manufacturer of TFT-LCD displays, and supplies customers that include many of the world's leading information and consumer electronics manufacturers. Innolux Corporation employs more than 68,000 employees worldwide and is a leader in the global optoelectronics industry. Innolux actively recruit and train R&D talent to consolidate its prominent status in the industry. Innolux is striving to realize the ultimate in visual infotainment. Their unique intelligent management platform has not only enhanced the company's management capabilities, but also provides customers with prompt, accurate delivery information.

Engineers at Innolux Technology Europe B.V. were confronted with the challenge of how to solve expected thermal problems with the integration of a display-centered system into a high-end automobile. The thermal design philosophy behind the activities to be taken is the subject of this article.

A thermal model was already available based on the code CST Microwave which had an embedded thermal solver, with detailed MCAD data as input. However, there were some doubts concerning its accuracy for this application due to the fact that the code does not solve the Navier-Stokes equations. Furthermore, its radiation treatment was not well-documented and caused further concern. Hence, as a starting point for the numerical analysis the CFD code FloTHERM from Mentor Graphics would be used.

The main objective was to provide preliminary models based on the available (but not sufficiently accurate) input data to study the feasibility of the proposed complete system as a sound starting point for subsequent discussions regarding design decisions with the customer. Tests on an available prototype were performed early on to provide a first-order comparison with the numerical models. In a later stage, experimental calibration by means of dedicated tests, such as thermal resistances and thermal conductivities and component thermal data, were required due to the inaccurate thermal data delivered. One key challenge was that the system was being designed to operate in an environment where the ambient could be as high as 70°C, while the testing had to be carried out at room temperature.

The numerical models were developed as a tool for optimization: sensitivity analysis of gap pads, wall thicknesses and thermal conductivities, PCB thermal data, boundary conditions, coupling of elements to housing, decoupling LEDs from light guide, properties of the display stack, etc. These models were also to compare with dedicated tests to determine several unknowns, specifically the real-life boundary conditions and thermal interface resistances.

To get an idea of the CFD output, figure 2 depicts three screenshots for the three cross sections through the hottest point. The module is mounted vertically, showing the internal airflow, for a total dissipation of 18.5W.

Results were made available to the customer for several system test conditions, heat transfer coefficients on the external surfaces to represent natural convection, with and without radiation, various choices of gap pad, PCB data, thermal conductivity changes, etc. including component temperatures, which should be interpreted as local solder temperatures. An additional temperature rise caused by the internal thermal resistances of the components should be added, but this is only appropriate when reliable component thermal data is available.

A comparison was made between the existing model created using CST Microwave and Mentor Graphics' FloTHERM. After a series

of trials, a reasonable match was obtained between the two codes, for identical input data and boundary conditions. However, there were a number of reasons for not continuing to proceed with CST:

- The CST model does not solve the Navier-Stokes equations, hence one never knows the consequences except by comparison to a code that does. Better to continue using the CFD code. Knowing the errors arising from the assumption, it was possible to continue in conduction-only mode to speed up the process of sensitivity analysis;
- How CST solves the internal radiation transfer is unclear, and subject to doubt; and
- On the issue of importing CAD files: the author is not in favor of using brute force when trying to get insight in thermal problems. While modern tools make it easy to import CAD files, the not-so-easy next step is to get rid of all mechanical and electrical details that are useless for thermal modeling. While the speed of modern computers may allow millions of cells, the consequences are twofold: the layout becomes very complex, hindering insight in what is going on, and the CPU times go through the roof, hindering fast optimization.

Comparison with Experiments

A number of experiments were available for an existing prototype. While the author does not recommend to use general results for checking the accuracy of the numerical model, in this case, because the results are already available, a comparison gave some idea of trends, and could be used to extract some average boundary conditions in the lab. After calibration of the boundary conditions, the model and the experiments matched to within 2°C.

Under normal operating conditions the SOT solder temperature was above the recommended 100°C, the rest of the components appeared to be acceptable. The importance of radiation was tested, and found to be making a significant contribution to the cooling. Fixing a strip of gap pad material between the LEDs, bracket, housing and carrier had limited effect, however introducing a gap pad below between the bottom and main PCBs reduced the temperature by up to 6°C and hence is considered effective. The results, particularly the SOT temperature, was found to be sensitive to the assumed PCB thermal conductivity. Due to uncertainties in the model data, a gap pad was introduced.

As in many other cases encountered, we were faced with a specified maximum "ambient" for a module or component. This is impossible to design to, as the temperature within the system varies throughout the system so there is no one temperature a component sees as its own ambient. Another challenge is that the customer cannot control the design of the environment, in this case a car dashboard, into which this product fits, so defining the boundary conditions for Innolux's system is the responsibility of their customer, not Innolux.

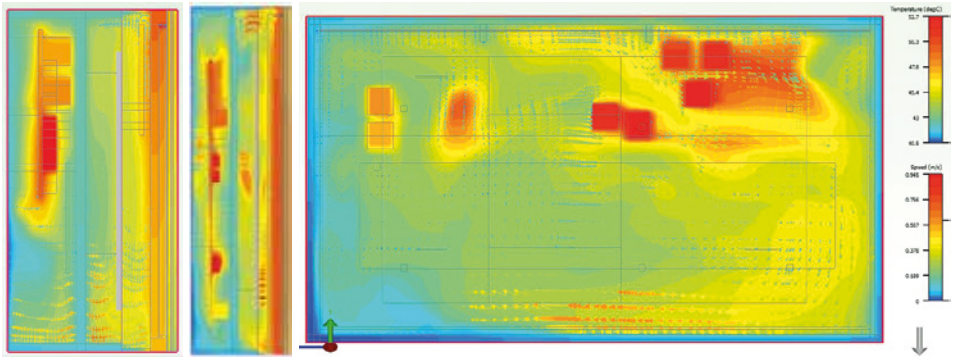

Figure 2. Typical results with flow, total dissipation 18.5W

It would be much better for the supplier to define a unique point at the outside of their component or module where a given temperature should not be exceeded. Such a spec is independent of the application and so avoids this issue.

Tests with a Dummy System

Tests with a dummy system should ideally focus on just a few issues, the most important being the real-life external boundary conditions for the complex display. Because of the complexity of the layout it does not make sense to model every detail, because it makes the fitting considerably more difficult. Unfortunately, this principle is often forgotten. Determining the boundary conditions that govern real-life applications is always a challenge, because to get sufficient accuracy one should simulate a very big part of the environment wherein the device is operational, including the radiators, the windows, the people etc. noted above. To avoid these problems, it is highly recommended to impose a heat transfer coefficient, including radiation, to be calibrated with real-life measurements.

Gap Pads

It is clear from the analyses that using gap pads or gap fillers improves the thermal performance. Be aware that the thermal data supplied by the manufacturers should be treated with caution, and gap pad data are not an exception unfortunately. The test method usually applies far too high pressure, resulting in too optimistic values. So, if critical, perform a dedicated test, ideally using Mentor Graphics' DynTIM system within which the pressure can be accurately controlled.

Recommendations

How to retrieve the right thermal data for all critical components, including the maximum junction temperature, how this is to be determined, and exactly now useful thermal data such as Rjc and Rjb are generated, requires investigation and discussion. It is highly recommended to ask the vendors of all critical components for more accurate models. For many components even basic but useful thermal resistances such as Rjb and/or Rjc are not available. As a customer, and especially if you build high-end or high reliability systems, you are entitled to get the right data, and especially when you will be held responsible for system failures!

Conclusions

The main conclusion is that, provided some measures are taken, the estimated temperature specs (as far as they are known), could be met, given the current dissipations and assumed boundary conditions. However, it could be foreseen that under certain circumstances these conditions would be worse, to such an extent that natural convection only cannot prevent overheating of the display at maximum ambient temperatures, and hence a forced convection solution should be explored. However, with the available numerical models of the system, it should be easy to check the consequences of such design changes.

In the opinion of the author, there is no other way to reach these conclusions within a realistic timeframe except by using a CFD-code combined with dedicated tests to confirm uncertain parameters required for the modeling. The biggest hurdle to improving the speed and accuracy of the design process is the lack of accurate data for the boundary conditions for the system, coupled with the widely-spread but incorrect habit of prescribing maximum ambient temperatures for the components.

Unique Challenges

with Autonomous Vehicle Systems Design and Integration

By Puneet Sinha, Automotive Manager, Mentor, a Siemens Business

Stories about autonomous vehicles are regular fare in the tech news cycle and usually include forecasts about the eventual ascendancy of self-driving cars. It is projected that by 2035, 25% of all cars will have partial or full autonomy [1]. In the short few years since the Google concept car, auto manufacturers new and old have announced their plans [2] for commercializing autonomous vehicles starting as early as 2021. That's disruption! So what will it take to deliver on this future? Without a doubt, artificial intelligence with robust machine algorithms to comprehend road and traffic conditions and make appropriate decisions is the most critical. However, it will take a lot more than artificial intelligence to build commercially-viable autonomous vehicles.

To deliver, the automotive supply chain will reshape very drastically in the coming years, where necessary component-level technology (sensors/fusion box/new electronics) will be driven by the new entrants (technology/electronics startups or large organizations) but the responsibility of vehicle integration will continue to be with auto OEMs. This fast evolving supply chain coupled with a paradigm shift in desired vehicle functionalities, pose unique challenges to autonomous vehicle engineering. In this article, I highlight the key challenges and Mentor's solutions to address them.

Sensor Design Exploration

Design goals for lidars, radars and cameras, three of the most critical sensors for autonomous vehicles, are largely centered on size and cost reduction without sacrificing high resolution and the high range necessary to support various levels of vehicle autonomy. Additionally, these sensors when integrated in vehicles, must function reliably in an automotive environment and in all-weather conditions. Desired small form factors and functionality consolidation for signal processing in the sensors may cause significant heat build-up that may be detrimental to performance and/or reliability of sensors. This may deter sensor size (and cost) reduction efforts. Thermally-conscious designs for sensor electronics as well as for

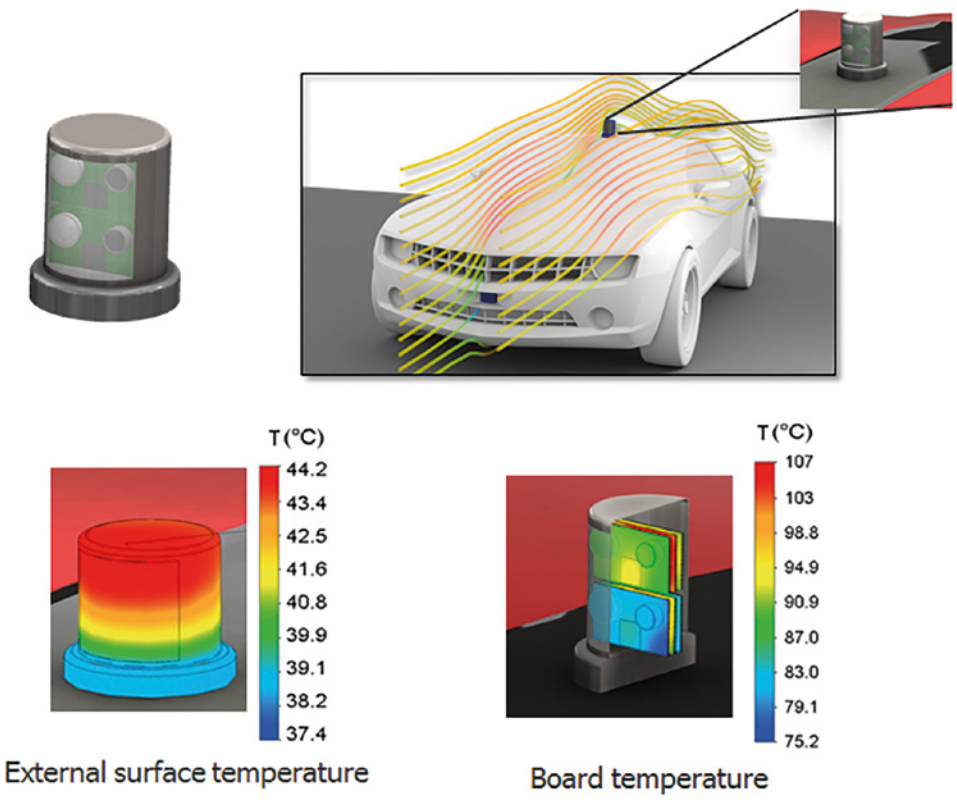

Figure 1. FloTHERM XT CAD-embedded thermal simulation of a rotating lidar mounted on a vehicle. Vehicle is moving at 10m/s and the ambient temperature is 25ºC.

their enclosures, while taking into account their vehicle integration locations, is critical to achieve the desired size (and cost) reduction while achieving high resolution and range goals are met. Mentor's EDA-centric, CAD-embedded softwares are ideally positioned to address these challenges from the early design stages. Figure 1 exhibits thermal simulation of autonomous vehicle lidar, accounting for its vehicle integration location using FloTHERM XT. In this Engineering Edge, there is also an article on IR camera thermal management design exploration based on the Masters' thesis of Hugo Falk from KTH Industrial Engineering and Management Institute. Reliable estimation of thermal behavior for lidars (solid-state lidars and/or mechanical rotating lidars) allow hardware engineers to achieve the desired size (and therefore cost) reduction without jeopardizing sensor life, due to undesired hotspots, in the automotive environment. Lidar integration in headlights, pursued by many companies, may pose additional challenges to lidar thermal footprint, as well as may impact lidar performance with headlight fogging or icing. FloEFD™, with its design-centric headlight design simulation capabilities, is helping companies to account for such vehicle integration issues. Additionally, heat build-up in sensors, especially in cameras, can negatively impact output quality, posing challenges to building reliable a 360º view around the vehicle. Sensor thermal simulation, therefore, taking its vehicle integration location and real world driving scenarios, impacts autonomous vehicle virtual testing, verification and validation.

Sensor Fusion Box Reliability and Power Consumption

Most of today's autonomous test vehicles have a trunk full of computers to ensure data from multiple vision and non-vision sensors can be efficiently fused together to create an accurate 360° view around a vehicle. However, for a commercial ready vehicle, desired processing efficiency must accompany size reduction for the fusion box for easy integration in a vehicle. This poses a significant challenge to sensor fusion electronics board design, as well as box enclosure design to ensure robust active cooling design considerations are taken into account. Our electronic cooling software FloTHERM™, a de facto standard for electronic cooling for the last 30 years, and its seamless integration with T3ster®, is ideally suited to address shrinking geometry related design challenges for sensor fusion box. With an electric powertrain, sensor fusion box will likely be powered by high voltage battery, through a DC-DC converter, and that can

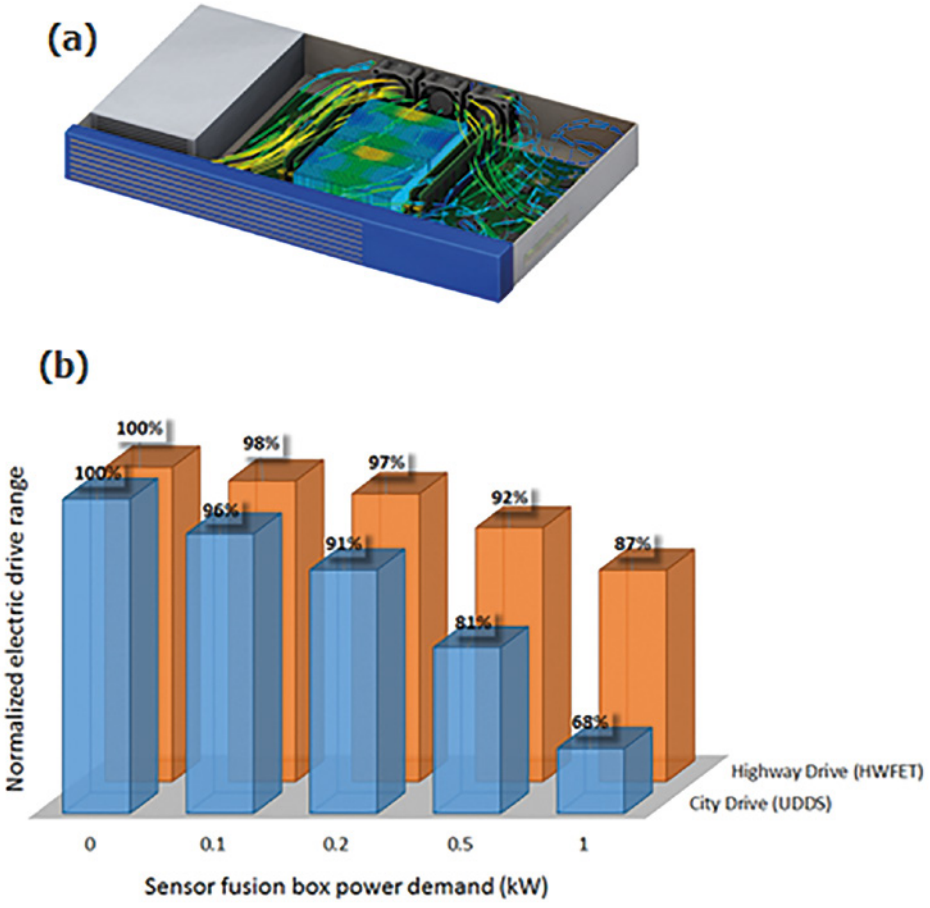

Figure 2. (a) CAD-embedded electronics cooling simulation of sensor fusion box using FloTHERM XT (b) FloMASTER autonomous vehicle system simulation to evaluate sensor fusion box power demand on electric drive range. Sensor fusion box is powered by the high voltage battery via DC-DC converter, in this example

impact electric drive range. This makes power consumption of the sensor fusion box a key criterion for vehicle integration. However, there is no publicly mentioned target for sensor fusion box maximum power consumption. FloMASTER™ system simulation can evaluate the impact of the sensor fusion box power demand impact on battery drive range and can empower suppliers and auto OEMs alike to develop and integrate a sensor fusion box that can deliver the required intelligence without impacting on vehicle performance/ range. For instance, Figure 2 shows a 250W (NVIDIA PX2) or higher power consumption sensor fusion box can reduce electric drive range by 10% or more, especially in city drive. This analysis further highlights the advantage of Mentor's DRS360 sensor fusion platform, which consumes no more than 100W- for commercially-viable autonomous vehicles.

Vehicle Integration and E-Powertrain Implications

Electric powertrain is indispensable for autonomous vehicles as it offers a) higher fuel efficiency and reduced CO_2 emissions, b) an easier platform to support drive-by-wire

"FloTHERM, a de facto standard for electronics cooling for the last 30 years, and its seamless integration with T3Ster®, is ideally suited to address shrinking geometry related design challenges for sensor fusion boxes.

systems needed for vehicle autonomy, and c) as battery prices keep dropping sharply, an attractive proposition of lower cost of ownership and maintenance, especially for fleet owners in a ride-sharing ecosystem. However, integrating vehicle autonomy with electrification will not be simple additive manufacturing. Vehicle autonomy poses additional challenges as well as unique opportunities to optimize electric powertrain size and energy management. For instance, the tremendous increase of vehicle electronics can impact electric drive range, especially in city drive. Whereas, autonomous vehicles are expected to be driven in a pre-determined way, especially Level 4 and Level 5 autonomous vehicles, that eliminate the need to account for 90th percentile driver (with highly aggressive driving pattern) and its impact on electric powertrain sizing and operation. Additionally, in a ride-sharing city-centric usage of autonomous vehicle, power demand to support cabin comfort is expected to challenge the available drive range and may warrant a complete redesign of cabin comfort. For these and many other aspects of autonomous electric vehicle, Mentor's frontloading e-powertrain simulation capabilities in FloMASTER and FloEFD allow users to define, design and evaluate electric powertrain architectures in the early phase of design. To learn more about Mentor's e-powertrain component to system simulation and reliability characterization offering, refer to our very recent white paper [3].

Vehicle Safety and In-Cabin Experience

With the increase in vehicle autonomy, safety-critical functions of steering and braking will depend on electronic control units (ECU). One of the biggest challenges for ECU design engineers is to manage power (and hence thermal) load on electronics, especially in harsh automotive operation. It has been demonstrated that thermal problems with ECUs invariably lead to electronic failures, which for autonomous vehicles will have severe safety implications. Mentor's Mechanical Analysis Division has been empowering major suppliers all over the world in developing reliable ECUs [4]. Additionally, as vehicles evolve towards Level 5 autonomy, passengers' expectations from in-cabin experience are expected to go through a paradigm shift. Infotainment systems are very likely to evolve towards thin large screens that may be integrated in the vehicle interior in ways that haven't been done to date. To meet these challenges, ECU and infotainment systems design engineers are exploring electronics consolidation. This can pose unique electronics cooling challenges for such systems. Mentor's T3Ster products are one-of-a-kind tools for non-destructive reliability characterization of PCB/semiconductors that allow suppliers to accurately quantify their product's life in real world drive cycle conditions. This coupled with FloTHERM XT/FloEFD electronic cooling software, benefit users to significantly compress time to design and develop ECUs for the drive-by-wire units and new infotainment systems.

As mentioned before, autonomous vehicles warrant technology integration from two vastly different verticals: electronics/technology and automotive. Historically, the two industries have very different product lifecycle requirements and product development trajectories. Simulations are, therefore, expected to play a critical role to connect these two industry verticals from component level design exploration to vehicle integration to vehicle-level verification and validation. Mentor's Mechanical Analysis Division frontloading, design-centric software, as shown through some of the examples in this article, are well-suited to address the challenges for autonomous vehicle engineering. Various customer stories in this edition of Engineering Edge and in the previous ones showcase the improvement in design efficiency our softwares are bringing to the customers in automotive and electronics industries all over the world. Extrapolating these benefits to autonomous vehicle engineering promises to bring a significant reduction in time and cost for design and vehicle integration for autonomous vehicle hardware.

References:

[1] https://www.bcg.com/expertise/industries/automotive/autonomous-vehicle-adoption-study.aspx

[2] http://www.businessinsider.com/companies-making-driverless-cars-by-2020-2017-1/

[3] "E-Vehicle Thermal Management Powertrain Simulation", a Mentor White Paper (Link)

[4] "ECU Thermal Reliability Becomes Mission Critical", a Mentor White Paper (Link)

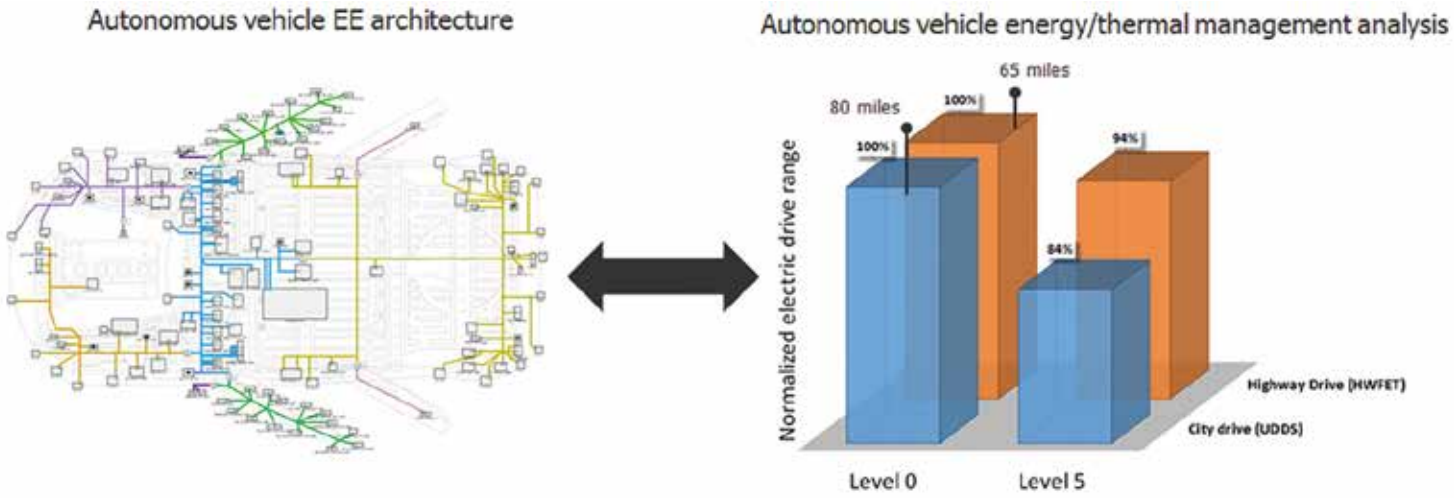

Figure 3. Impact of Level 5 autonomous vehicle system integration on electric drive range. In this example, Level 5 autonomous electric vehicle (80 miles electric range with 80kW front motor drive unit) is simulated where vehicle EE architecture (simulated using Mentor's Captial software) includes 30 sensors, DRS360 centralized sensor fusion and ECUs for steering and braking. Electric drive range is simulated at 25°C ambient conditions using FloMASTER based e-powertrain simulation framework.

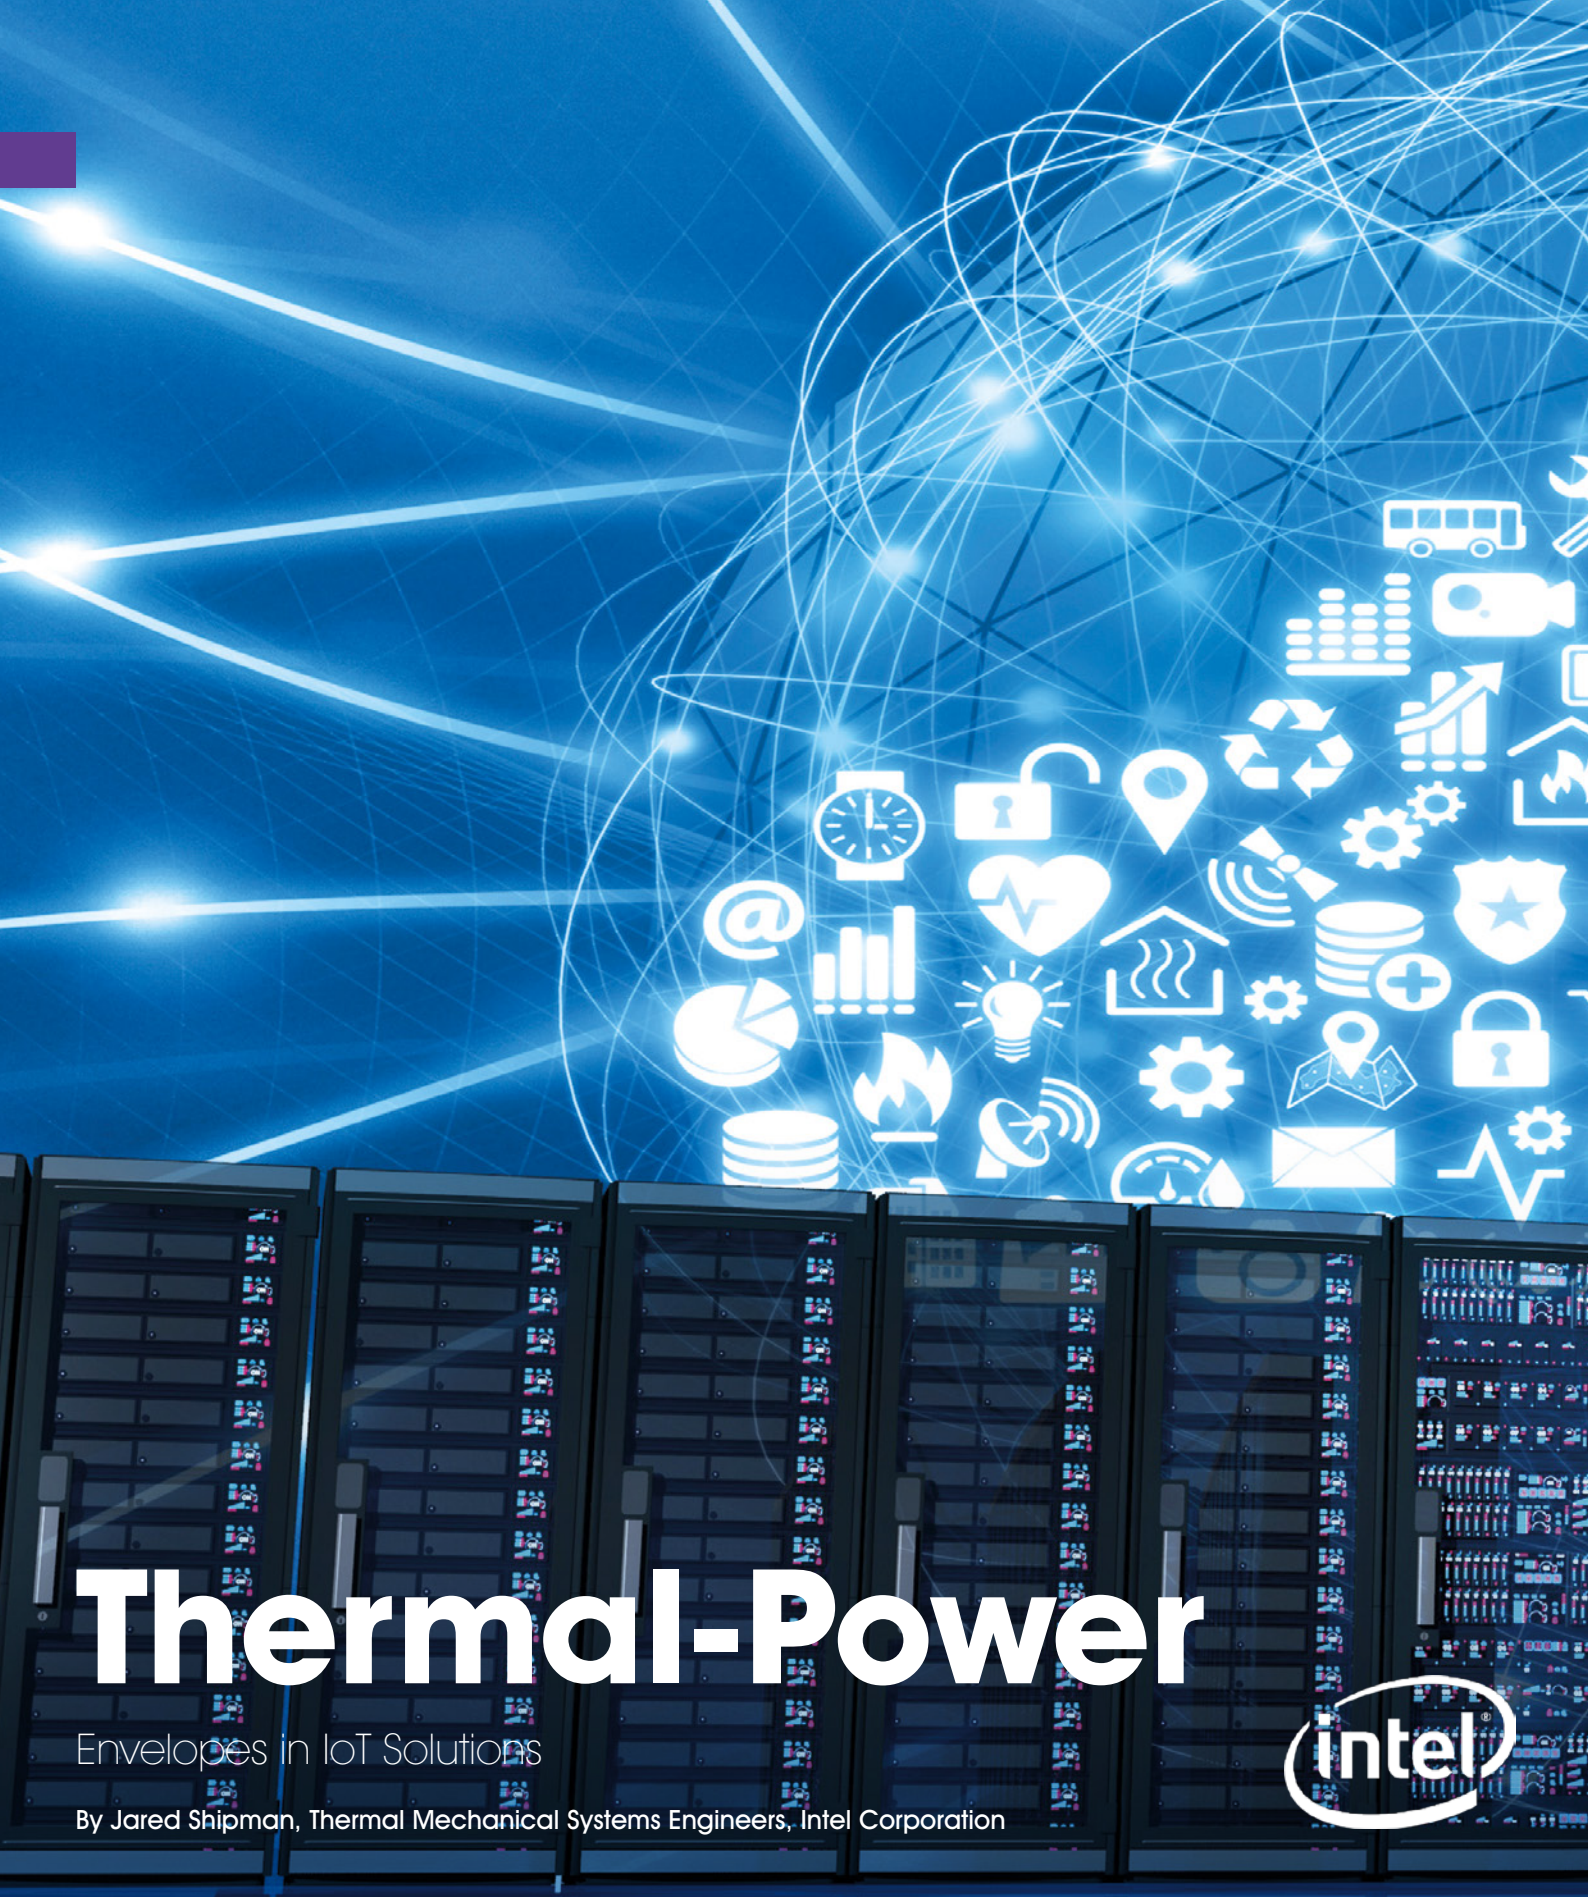

Thermal-Power

Envelopes in IoT Solutions

By Jared Shipman, Thermal Mechanical Systems Engineers, Intel Corporation

Intel's Internet of Things Group (IOTG) comprises a variety of divisions Retail Solutions, Smart Home/Building, Industrial Solutions and Transportation to name a few. Each division has its own challenges when it comes to designing solutions for their customers. The Thermal/Mechanical Systems Engineering team within IOTG is there to ensure that the solutions are designed with adequate thermal capabilities for the Intel processor within them.


10101
01001
11010
10101
01001
11010

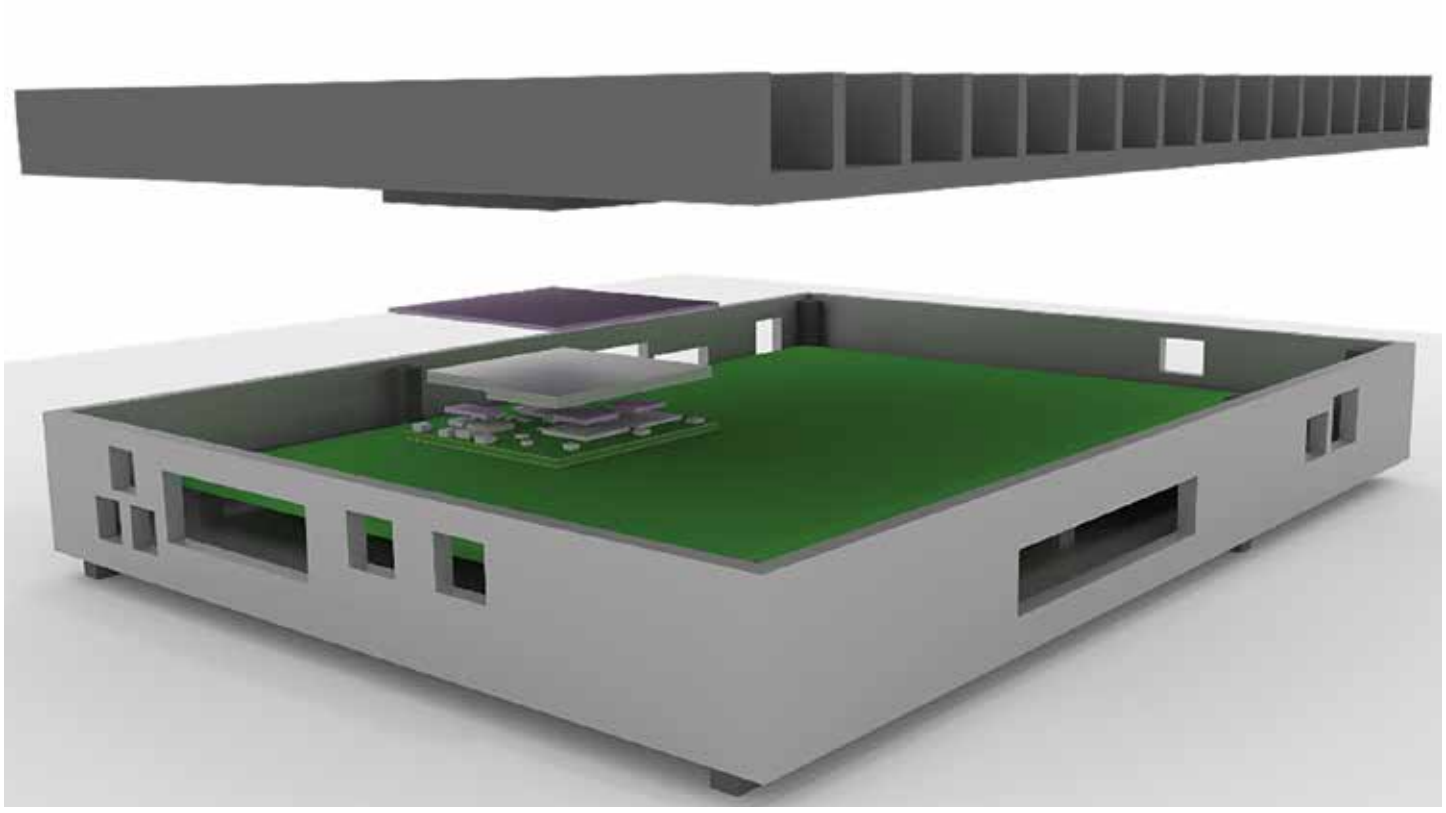

Figure 1. Rendering of a reference solution for an IoT module

Challenge:

As the Internet of Things continues to expand with more and more devices being connected to the Cloud the demand for smaller form factor devices with higher compute power is increasing. Not only do customers want better computing, the devices also need to be able to communicate via cellular or WiFi to send the data back to the customers' servers. With higher compute power comes an increase in junction temperature which creates a challenge for the thermal engineer. Adding to the small form factors is the spectrum of end customer use cases and operating environments. For example, a device in a storage container could be placed in the scorching Arizona sun or the frigid cold of Alaska and these devices need to operate flawlessly in both conditions. One of the tasks as a thermal engineer within the Internet Of Things Group at Intel is to aid our customers in their thermal solution by designing and validating a reference design.

Solution:

With so many end customer use cases it is impossible to validate them all so one way we have decided to solve this is by classifying the use cases into two broad categories of Application Only Processing and Application plus Communication Processing. This allows us to focus in on the most thermally significant use cases for each category and design and validate to those. We also want to give our customers a method for understanding the thermal capabilities of our reference design, without it being too confusing, to the point that someone not in the thermal field could understand. This inevitably led us to the creation of the Thermal-Power Envelopes for each classification of use case.

Benefits:

With FloTHERM's simulation environment we were able to create and optimize the reference design with a high level of confidence that the design will perform as it did in the simulations. To verify this we built a prototype of the design shown in Figure 1 and tested it in an environmental chamber at different temperatures.

The test results showed that the simulations were accurate within two degrees. This is expected given the current limitations of measuring technology today. Validating the simulation results means that our model is now correlated to real world data which opens the door for our customers to use these models as a first step in their own

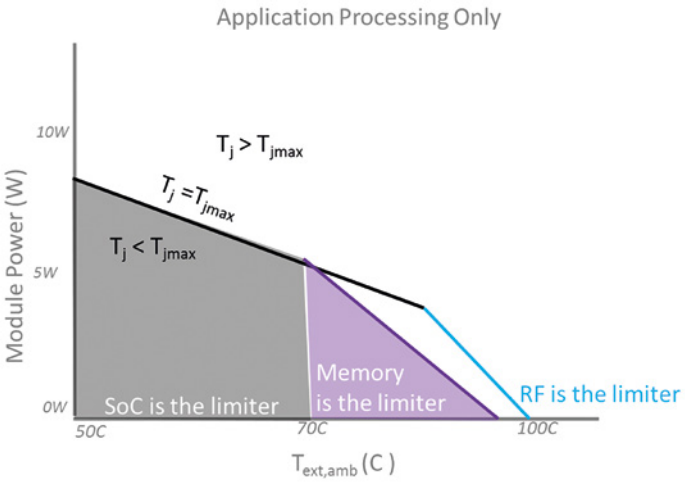

Figure 2. Thermal Cooling Capability Envelope for AP only work load.

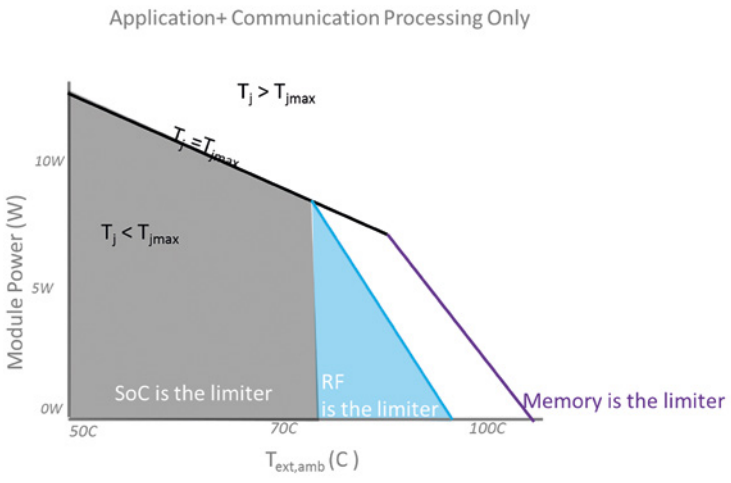

Figure 3. Thermal Cooling Capability Envelope for AP+CP work load.

thermal solutions saving them prototyping and testing costs. The correlated models are now used to generate the Thermal-Power Envelopes for this system. These Thermal-Power Envelopes are system dependent which means that a change in the thermal solution whether that be heatsink size, TIM selection, or some other variable will result in a shift in the curve.

These curves provide the customers with a fast and simple way to analyze and compare the thermal solution of a system. The charts also help the customers to understand the power and performance limitations of the system by outlining how much power the system can consume before reaching a component's specification limit at a given external ambient air temperature.

Conclusion:

The challenges that lay before today's engineers in a world that demands instant and accurate analysis of complex computations at the edge and in the palm of their hands are far from solved. None of which is more important than ensuring these devices can operate in the harshest environments. This means a properly designed thermal solution. With FloTHERM's simulations environment we are able to create, optimize, verify and validate a reference solution with minimum impact to prototyping and testing costs. FloTHERM's results were within two degrees of accuracy leading to the creation of the system's Thermal-Power Envelopes that provide our customers with a tool that they can evaluate and compare the performance of a system's thermal solution in a fast and simple way.

"The challenges that lay before today's engineers in a world that demands instant and accurate analysis of complex computations at the edge and in the palm of their hands are far from solved. None of which is more important than ensuring these devices can operate in the harshest environments.

IR-Camera **Thermal Management**

By Atsushi Ishii, Director Sensor and System Functions, FLIR Systems
and Hugo Ljunggren Falk, KTH Royal Institute of Technology

FLIR Systems develops thermal imaging cameras and components for a wide variety of commercial and government applications. Depending on the end use, a product design might be driven by economical or dimensional factors while other applications are constrained by extreme environmental operating conditions. Whether the product is designed for firefighting purposes or to monitor a datacenter, FLIR continually explores opportunities to shorten their product development cycle.

Many development cycles are shortened by integrating simulation early in the design process. In addition to the time saved, the cost of developing a product can decrease substantially by reducing the number of prototypes built and tested. In terms of sustainable energy engineering, by "streamlining" the development cycle by adding simulation, less physical resources are used to create prototypes which means a reduction in environmental footprint.

To explore the benefits of introducing simulation early in the design process, the FLIR AX8 stationary camera was used in a case study with FloTHERM™ XT. The AX8 camera has been developed to monitor apparatus in industrial sensor networks, such as telecom electrical boxes or refrigeration units in supermarkets. A rendering of the camera is provided in Figure 1.

The case study was considered in two phases, Alpha and Beta. Alpha represents the earliest part of the design process, where decisions such as board layout are considered, and Beta represents the fully detailed analysis.

The simulation during the Alpha phase only includes the PCB and is shown in Figure 2. With FloTHERM XT, when the PCB design changes, the board can simply be swapped out through the FloEDA Bridge. The initial conditions and values set will be kept for the new board. Also, parametric studies can be used to set different configurations which can range from initial conditions to changes in geometry and is useful when comparing different solutions against each other.

The Beta phase, which represents the detailed analysis, is then considered. Simulation during the Beta phase would include any thermal relevant components including the mechanical assemblies and enclosure. The simulation setup for the Beta phase is shown in Figure 3.

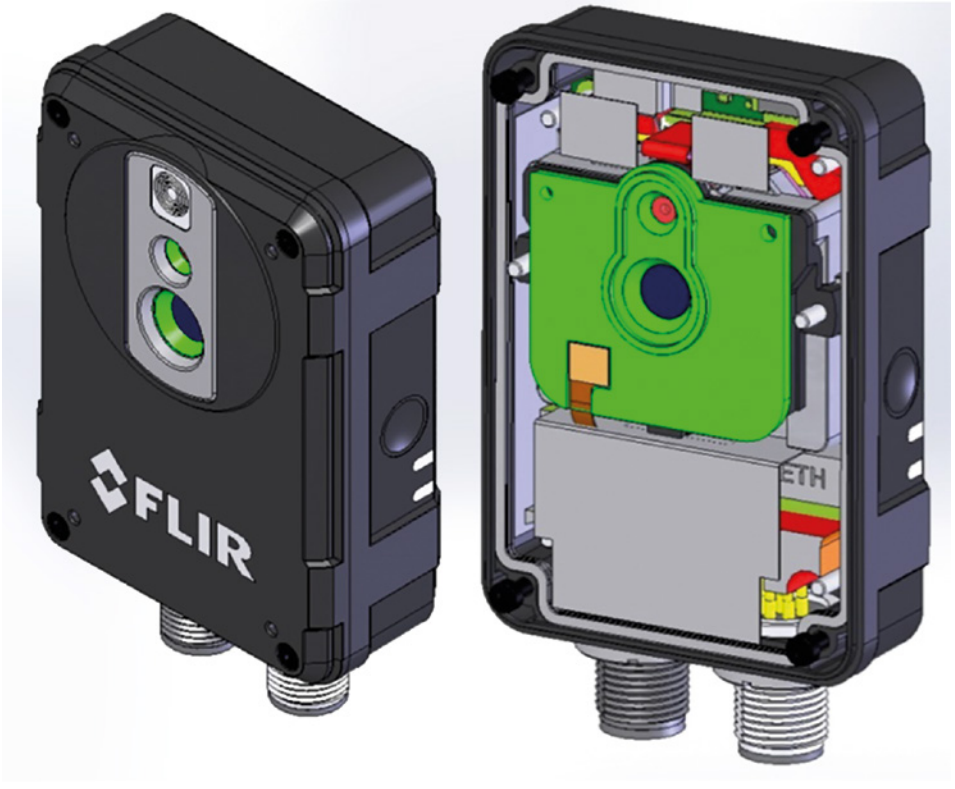

Figure 1. FLIR AX8 with and without cover

To assess the integrity of the simulations they were compared to temperature measurements of the camera in a loaded condition. Measurements during the Alpha phase were conducted using a FLIR T640 thermal imaging camera mounted on a tripod. Thermocouples were used to measure temperatures during the Beta phase.

Results Comparison: Alpha phase

The Alpha phase was used to determine if correlation between measurement and analysis could be achieved with only the PCB, in a lab environment. With correlation the analysis could be used to explore PCB design choices with respect to layout and overall PCB footprint. Figures 4 and 5 show the IR images captured with the FLIR T640 for the FLIR AX8 printed circuit board. Figures 6 and 7 show the FloTHERM XT results of the PCB in a natural convection environment. Table 1 shows a component temperature comparison of the IR image and simulation results.

The Alpha phase simulation results compare well with the bench top IR camera measurement. Any proposed board level design changes could be explored with confidence without building and testing numerous prototypes.

Result Comparison: Beta phase

The Beta phase was used to determine if correlation between measurement and analysis could be achieved when considering the full assembly. With correlation, the analysis could be used to explore thermal design choices such as materials, thermally conductive gap fillers, and venting. In addition the analysis could be used to predict performance under any environmental or usage condition. Figure 8 shows a solid temperature cross section of the simulation results and identifies the dominant heat transfer paths.

The Beta phase simulation results compare well with the bench top thermocouple measurements. This simulation model would allow a thermal design team to explore design alternatives and consider operating environments that would be too expensive or time consuming to otherwise consider.

Summary

Whether a product design is constrained by extreme operating environments, cost, or form factor, introducing simulation early in the process will shorten their product development cycle. Validated simulation processes reduce the cost of design through time savings and number of prototypes.

This article is summarized from the Master of Science Thesis: Thermal Management in An IR-Camera. Hugo Ljunggren Falk, KTH School of Industrial Engineering and Management, Stockholm

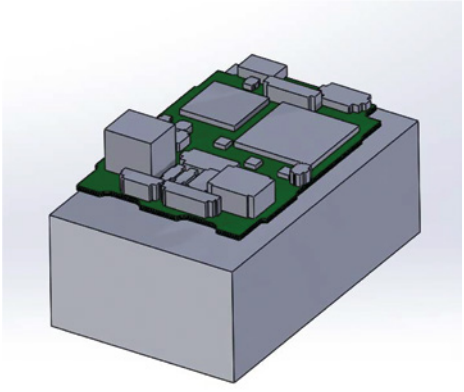

Figure 2. Alpha simulation setup

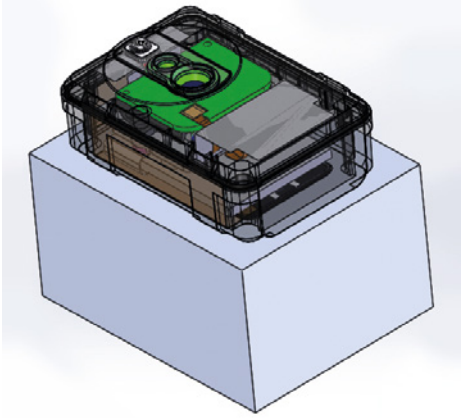

Figure 3. Beta simulation setup

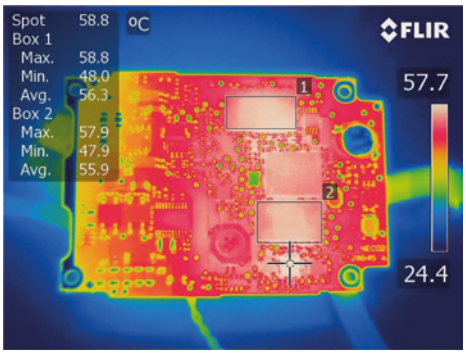

Figure 4. Alpha phase PCB top surface IR image

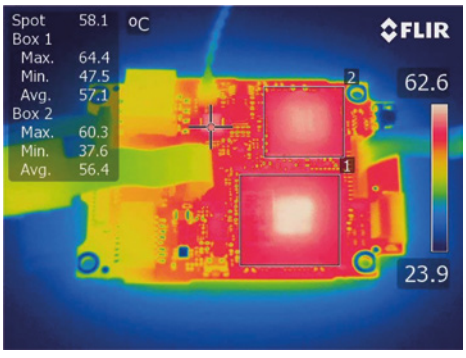

Figure 5. Alpha phase PCB bottom surface IR image

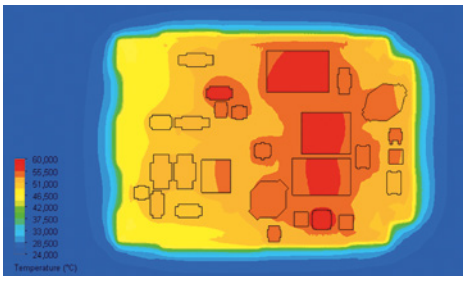

Figure 6. Alpha phase top view simulation results

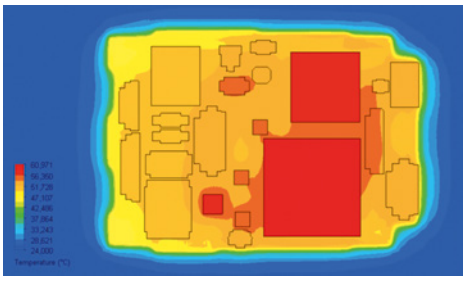

Figure 7. Alpha phase bottom view simulation results

	Temperature [°C]		
Probe	**Measurement**	**Simulation**	**+/-**
CPU	57.1	61.0	3.9
FPGA	56.4	57.7	1.3
Rectifier for PoE	58.1	55.8	-2.3
FPGA memory	55.9	55.5	-0.4
CPU memory	56.3	55.7	-0.6
Power mgmt. unit	58.8	57.0	-1.8

Table 1. Alpha phase results comparison

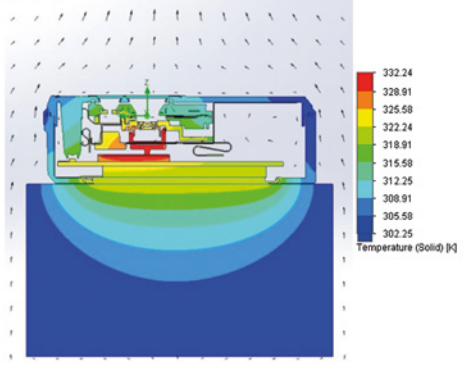

Figure 8. Beta solid temperature distribution simulation results

	Temperature [°C]		
Probe	**Measurement**	**Simulation**	**+/-**
Aluminum back plate	52.5	49.1	-3.4
CPU	58.9	58.7	-0.2
FPGA memory	55.4	49.1	-6.3
FPGA	57.1	58.1	1.0
Visual Cam	48.5	52.4	3.9
IR Sensor	47.2	48.2	1.0

Table 2. Beta phase results comparison

INTERVIEW

Wendy Luiten

Q. You have been in the electronics cooling field for over 20 years. How has electronics cooling changed over that time?
A. If we are really counting, I guess it is closer to over 25 years, although I remember way back in the early nineties that my then boss thought it was not worthwhile for me to learn flotherm – famous last words. I think the main change by far is the use of CFD and especially the GUI. I remember having to number my own nodes and integration points in the FEM code I used when I first started my career, and colouring in the printed number fields that were delivered on huge reams of folded paper – unimaginable in these days.

Q. What is your background in thermal design?
A. Worked in Lighting, TV, and Healthcare. I actually majored in heat and fluid flow, and started at Phillips Manufacturing Technology lab – the applied split-off of Philips Research- as a thermal and thermo-mechanical specialist so when I moved to the division of consumer electronics, getting started at electronics cooling was a natural progression. I started out in TVs and that grew to consumer electronics products in general. After a couple or restructurings and the sell-off of the CE divisions I found myself back at Philips Research and working on Lighting and Personal Health products.

Q. What are the common electronics cooling challenges, and what are the differences?
A. I have predominantly worked in consumer goods – so end-customer focus, speed of development, cost and manufacturability are very major drivers common to all three fields of Consumer electronics, Lighting and healthcare. The big differences I saw were in development speed and in platform use – the art of developing a set of common parts to cover not one product but a whole range of products. In general, the markets with more competition tend to be more challenging – but that is part of the fun.

Q. Can you tell us about Wendy Luiten Consulting and what your company does?
A. I do thermal design and simulations for clients and I provide training. Wendy Luiten Consulting is my next big adventure- I took voluntary leave from Philips research to get back in touch with the business again. In the past few years I have also branched out in Design for Six Sigma – I am an Innovation Master-BlackBelt-in-progress. What that actually means is that I look into the entire chain from discovering to designing-in and delivering value in electronics products. WLC develops and delivers training and advises in the combined field of thermal design and design for six sigma, which are surprisingly interlinked if you look into it.

Q. What are the challenges your customers are facing today?
A. For most people, thermal design is the ultimate performance limiter, and it is hard to do really well because so many different disciplines and so many system levels are involved and, in the end, you are working with invisible entities. The majority of electronics products are air cooled- which means the main component is not even on the bill of materials, and all kinds of enabling design features tend to get overlooked or not documented, and then people tend to re-invent the same wheel many many times.

Q. Which project have you done that you are most proud of and why?
A. I have to answer with two projects here – THE most exciting thermal project by far was the thermal design and architecture of the world's first ever fan less flat TV monitor. I wrote up that story in the paper 'Cooling of a Flat Tv monitor', which won the Semitherm Best paper award in 2002. More recently I am very proud of the work that I did on the Philips DfSS Green Belt and Black Belt training curriculum, expanding, re-writing and re-structuring the material to deliver present day product development skills.

Q. When you teach electronics cooling, what key points you want people to get?
A. Oh, there are so many key points! I think the most important one is that your thermal design is made or broken in a few vital early architecture choices – it is absolutely essential to figure out what they are – and then guard these choices through the entire design flow. And these can easily be choices that the EEs and MEs overlook as non-thermal related, surface sizes, gap sizes, material choices – that sort of thing.

Q. What do companies get wrong when it comes to thermal design?
A. There are still a lot of companies out there to which thermal malfunction just happens. They are unaware that there is such a thing as thermal design. One step up are the companies that use thermal design, and even CFD, but with a way of working that is straight from the hardware dominated nineties – they make an electrical design and a mechanical design, and when these are fully finished and detailed they submit the CAD and CAE files to CFD calculations. That really is just substituting the hardware prototype with a CFD run, and not profiting from the vastly larger possibilities that using CFD in concurrent architecture brings you. It is like when the car was replacing the horse – you have the engine outside of the vehicle like a horse was. But since then we have found out that integrating the engine into the vehicle has big advantages.

Q. How have simulation tools changed the way thermal design is done?
A. I would say the impact is vast, and on two points. The first is of course the speed of development – you do not have to wait for hardware prototyping and repeated hardware measurements. The second is related to the speed: insight in how your design functions. Computer simulations enable a thorough investigation of your solution space that would be inaccurate with analytical models and just impossible with traditional hardware and in addition they can calculate the effects of limit cases that would be impossible to build. This enables to really find the optimal thermal design in terms of total system desirability. Of course, there is also a risk – a CFD code in the hands of an unskilled user can make lots of untrue predictions, but that is really the risk of any tool at all in the hands of an unskilled user.

Q. What is your philosophy about how simulation tools should be used across the product creation process?
A. My philosophy is that they should be used from the earliest phases of concept design to discover the main parameters and guide the vital choices. I typically do the first CFD right after the back-of-the-envelope hand calculations – and this has the advantage of having like a double check both on the hand calculations and the earliest CFD model – just to be sure that I did not forget to model an important effect. Then, as the design grows, the CFD model evolves in complexity as well and can be used to monitor the effect of changes in the subsequent steps.

Q. What would you like electronics CFD be able to do that it can't do today.
A) In general: Maybe make the whole process more fool-proof from a thermal concept point of view. Like I said, a CFD tool in the hands of an unskilled user is a dangerous thing. Maybe with warning messages like 'I see that you have switched off radiation but there is no fan or external flow involved – are you aware that radiation is very significant in natural convection cooling?' or ' I detected a significant heat flow over an interface between two parts modelled as ideally conducting – are you sure?' or even 'I detect a significant temperature difference across an interface that you modelled as non-conducting – are you sure?'. I can imagine you could use Robins B and Sc numbers for something like this.

B) For me personally: Adapt the command center so I can run my DOEs more easily. The ability to switch groups of variables would already be a HUGE help, as well as (in flotherm XT) the ability to switch geometry on or off. The fact that the geometry is in a different scenario table in flotherm XT makes it very complicated to do a DOE in architecture phase. In addition, calculating the effect would be comparatively easy to do in flotherm itself.

C) In design and optimize phase making the transfer function and doing the montecarlo and the tolerance analysis would also be mathematically easy to do from my point of view. Obviously, you already have a random generator on board for the monte carlo because you do ray tracing, and I can imagine that the capability of doing the thermal tolerance analysis would be a feature that could be unique to flotherm and also add considerable value to the end user.

Fast Design of Medical Ultrasound Probes with FloTHERM

By Kazuya Motoki, Hitachi, Ltd., Japan

HITACHI
Inspire the Next

Ultrasound systems (Figure 1.a) for diagnostic images have been around for several decades. They have advanced in the last twenty years to include much more sophisticated electronics that deals with high quality scan images on much larger screens than ever before but in smaller, more compact systems that draw much higher power consumptions to manage.

Hitachi supplies ultrasound scanner systems and a wide range of probes (Figure 1.b). Each probe includes piezoelectric transducers, as an imaging sensor, which generates heat during scan. Recently, more of the electronics associated with the probe have been used inside the probe to improve image quality. As the performance of probes increase, heat generation is becoming a bigger issue compared to our previous designs.

International safety standards (IEC 60601-2-37) exist that manufacturers must comply with in terms of the thermal limits for ultrasound probes (Table 1). These temperature limits must not be exceeded while the equipment is powered continuously for up to 30 minutes so as to prevent burns. As well as this we need to allow for real world probe situations like on and off usage, variations in ambient conditions, and body temperature variations in patients. Prediction of the transient surface temperature rise, rather than the equilibrium temperature, is a key measure we use for assessing probes and hence any CFD code must be able to do a transient thermal simulation correctly. We chose the FloTHERM thermal analysis tool a few years ago because it was faster at getting simulation results than our existing tool, easy to mesh, had a good interface with our MCAD software, and it was easy to pass our CFD predictions to our FEA tool.

We typically have two to three weeks to influence the design of a probe. Our typical workflow before we used FloTHERM is shown in Figure 2 (a). We tended to finalize on the product design after several rounds of prototyping but we found difficulties with our then CFD tool and it was difficult to know what the problems were in the final assembled product. Figure 2 (b) shows the process workflow we have iterated to over the last few years with FloTHERM, because of its capabilities, to shift design

(a)

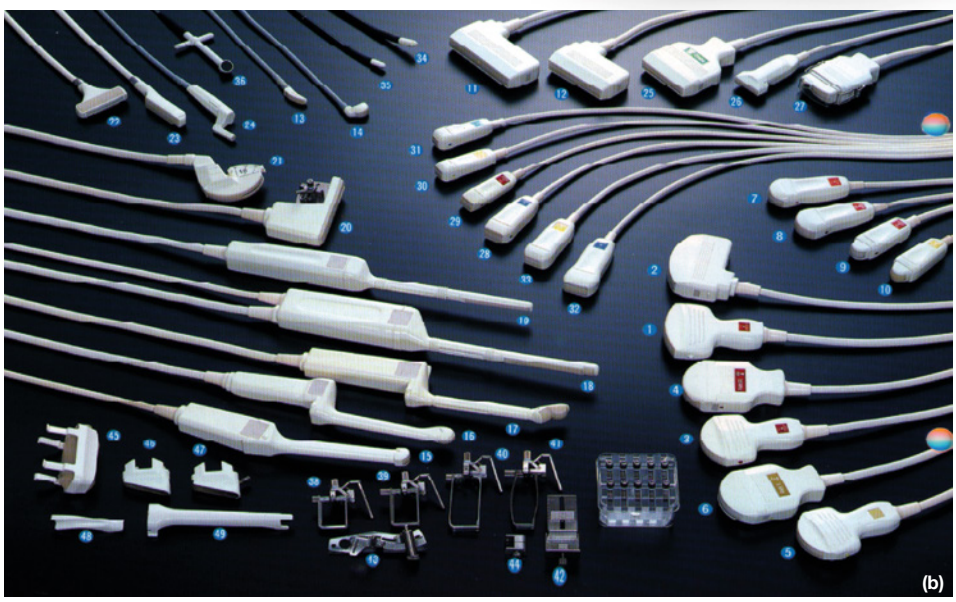

(b)

Figure 1. (a) Typical Hitachi Ultrasound Scanning System and (b) a selection of Hitachi Ultrasound Probes

to the left and an earlier stage that has allowed us to have many fewer prototypes and a cost saving of X2. In this instance, the result achieved was because we were able to review the thermal analysis predictions versus actual measurements early on, identify and rectify problems for each assembly process, and then move to the next part of the process.

From a product design perspective, we also want a CFD tool that has high analysis precision and low prediction errors. To achieve this we first made sure our probe constituent material properties were correct – such as thermal conductivity, emissivity, specific heat and density. To ensure high precision, we carried out rigorous test measurements on all of our supplier materials (Table 2) and found that certain materials (for example, Materials B that was composites) had big variations (in red font) than those shown in manufacturer datasheets.

The next thing we did in this thermal simulation process was to improve our transient thermal analysis predictions by first simulating a simple probe inside FloTHERM (Figure 3) at steady state and then we did a more detailed probe model (with each component part accurately modeled) reviewing each process transiently. Here, the temperature dependence of the ceramic heater should be taken into account via measurement of its resistance at a fixed voltage (Figure 4) when the heater was used in the simple mock-up in experiment. These heater temperature characteristics needed to be input into FloTHERM and controlled to maintain a constant power during the simulation process, otherwise the CFD simulation would overestimate or underestimate the actual measurement results. Through these processes, we found that using Cartesian meshes inside FloTHERM gave the most accurate CFD predictions (Figure 5) for the probe assembly when compared with actual measurements, especially when we used measured material properties in the detailed probe model (Figure 6). We also used FloTHERM's excellent Command Center capability for multiple early Design of Experiment analysis. In summary, we found that our simulations are very accurate, ten times faster than before, and we have a design process that fits our manufacturing needs and we use FloTHERM 100% of the time on our four core machine. Indeed, FloTHERM has helped us to invent several patents for our probes.

Test state	Body surface applications	Invasive applications
Tissue-mimicking material	Initial temperature: 33°C or more **T < 43°c**	Initial temperature: 37°C or more **T < 43°c**
Left in air	Initial temperature: 23°C ± 3°C **T < 50°C**	

Table 1. Temperature standards for ultrasound probes

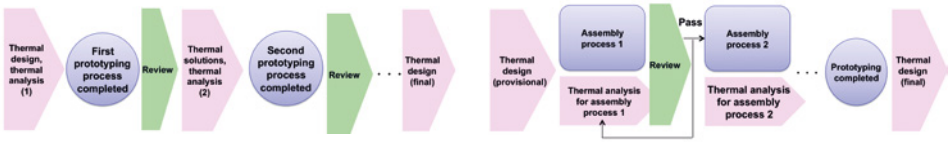

Figure 2. Hitachi CFD simulation workflow before FloTHERM (a) and today (b)

Material	Thermal conductivity [W/mK]	
	Manufacturer Catalog value	Measured value
Material A	0.3	0.3
Material B	15	5
Plastic A	0.2	0.2
Adhesive A	0.2	0.2

Material	Emissivity	
	Manufacturer catalog value	Measured value
Material A	0.86	0.96
Material B	0.8	0.15 (Gold-plated surface)
Material C	0.9	0.96

Material	Specific heat [J/kg]	
	Manufacturer catalog value	Measured value
Material 1	900	850
Material 2	1,900	1,500
Material 3	1,000	650

Table 2. Ultrasound probe constituent material thermal property measurement results

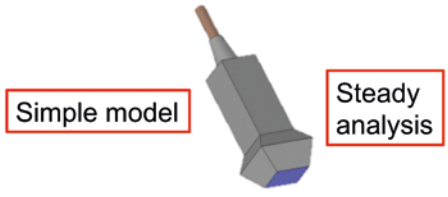

Figure 3. Simple FloTHERM model geometry of a Hitachi probe

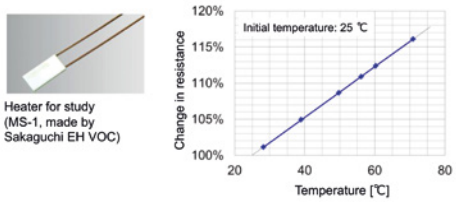

Figure 4. Ceramic heater characteristics

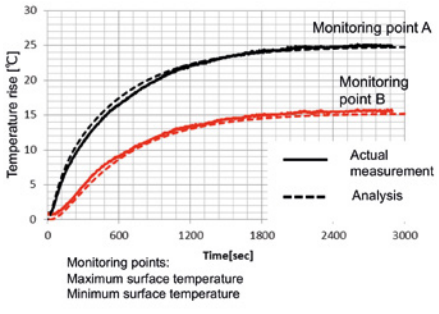

Figure 5. Transient thermal measurements versus FloTHERM predictions of a simple probe

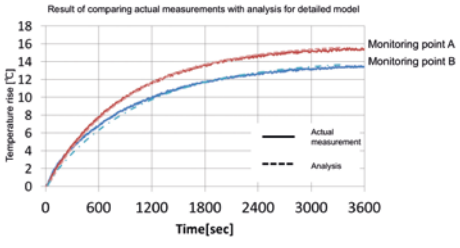

Figure 6. Transient thermal measurements versus FloTHERM predictions of a detailed model probe

Peristaltic Pump Electronics Thermal Design with FloTHERM XT

By Michael Clements, CEng MIET, Electronics Engineer, Watson-Marlow Fluid Technology Group

Watson-Marlow Fluid Technology Group produce a range of chemical metering pumps. Their Qdos pumps are specified to replace diaphragm and other designs of metering pump as a result of their precision flow rates and significantly lower maintenance costs when handling abrasive or corrosive fluids. The success of these pumps and their technology makes them the ideal choice for rugged applications.

One such application for Qdos pumps is metering chemical coatings on to seeds, with the pump mounted on the seed planter towed behind a tractor. Seed treatment is now common practice for agriculture and horticulture industries, where the thick fluid coating may contain growth promoters, inert carriers or fertilizers, as well as antimicrobial or antifungal treatments.

The requirement to deploy Qdos pumps externally has introduced some new design complexity. As all Watson-Marlow pumps are mains powered, the mains powered 48V 200W internal Switched Mode Power Supply Unit (SMPSU) had to be replaced by a 12/24V DC input power supply.

The original power supply was constrained by the footprint of the mains SMPSU, with an initial heatsink design intended to ensure that the power supply components would be kept cool enough at the output power and input currents required for a target 90% efficiency, giving a thermal load of 20W. The pump housing is non-optimal for cooling due to its construction of glass-loaded plastic and also contains a brushless DC motor that can vary between 70°C and 90°C surface temperature at a maximum ambient of 40°C. The initial heatsink design looked reasonable, but could not be evaluated without an expensive machined prototype being manufactured.

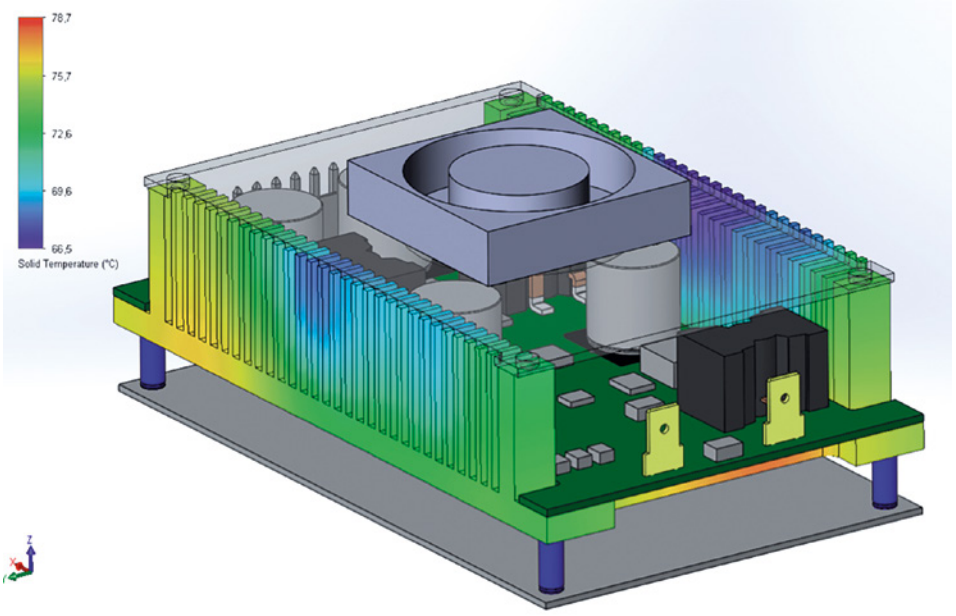

Figure 1. Original Watson-Marlow design as evaluated by IC Blue

Previous experience of FloTHERM, coupled with use of PADS PCB design software suggested that FloTHERM XT for PADS would be the best tool for evaluating heatsink designs.

Working with the Mentor business partner, IC Blue, Watson-Marlow evaluated the original design using FloTHERM XT for

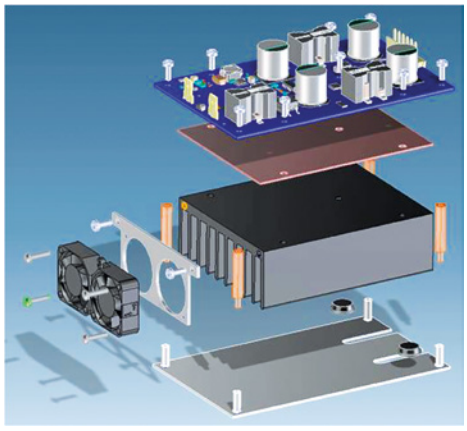

Figure 2. Dual Fan Extruded Heatsink Design showing Complexity of PCB Imported from PADS

PADS, discovering that the heatsink did not work well, with excessive temperatures shown on the components, as well as being prohibitively expensive to machine, requiring a large number of fins.

Based on the success of the evaluation work by IC Blue, Watson-Marlow were able to make a business case to purchase FloTHERM XT for PADS to be used by their own development team. The purchasing decision was used to great effect, enabling a completely redesigned heatsink utilising dual 30mm fans and a standard heatsink extrusion, with the DC power supply meeting the cost and size target, as well as matching the original mains SMPSU, shown in Figure 2.

Throughout this development work, Watson-Marlow were able to take component placement, board design and power data, directly from PADS (Figure 3) into FloTHERM XT in a smooth and efficient manner. As a result, FloTHERM XT simulations were fully synchronized with the PCB layout and circuit design, without thermal risks. Watson-Marlow also use PTC Creo for their mechanical design, and were able to read native Creo geometry directly into FloTHERM XT for PADS, allowing other designs of heatsink to be incorporated directly and used without modification or further simplification.

With this workflow, Watson-Marlow used FloTHERM XT for PADS to undertake many further investigations, re-designing the heatsink to a form that can be cast to save costs over the extruded part. They were also able to optimize air flow paths through the cast heatsink to achieve the equivalent cooling performance of the above dual fan design, but using a single 30mm fan resulting in further cost savings. Two of the many variants studied are shown in Figure 4, indicating how complex the interaction is between the flow exhausting the fan and passing along the heatsink fins.

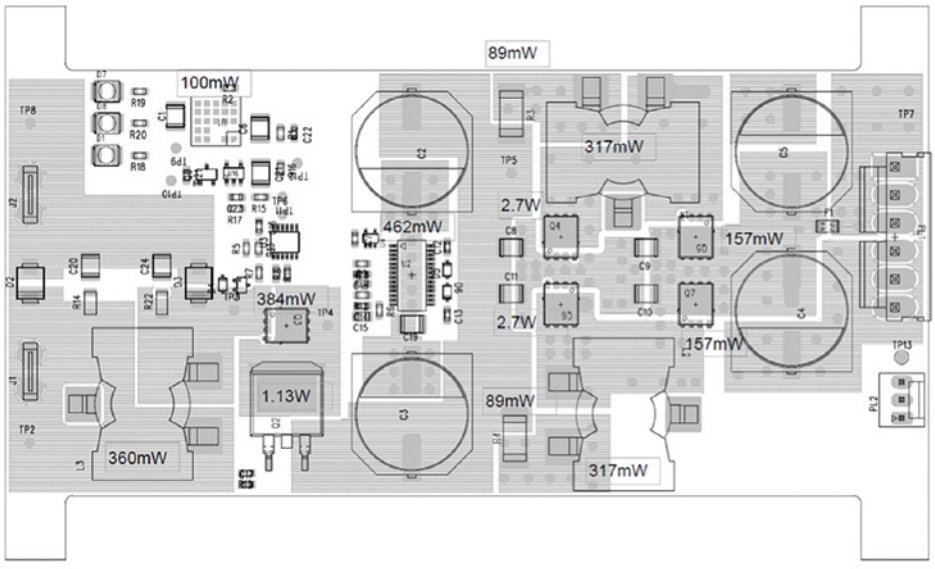

Figure 3. PCB component layout drawing from PADS

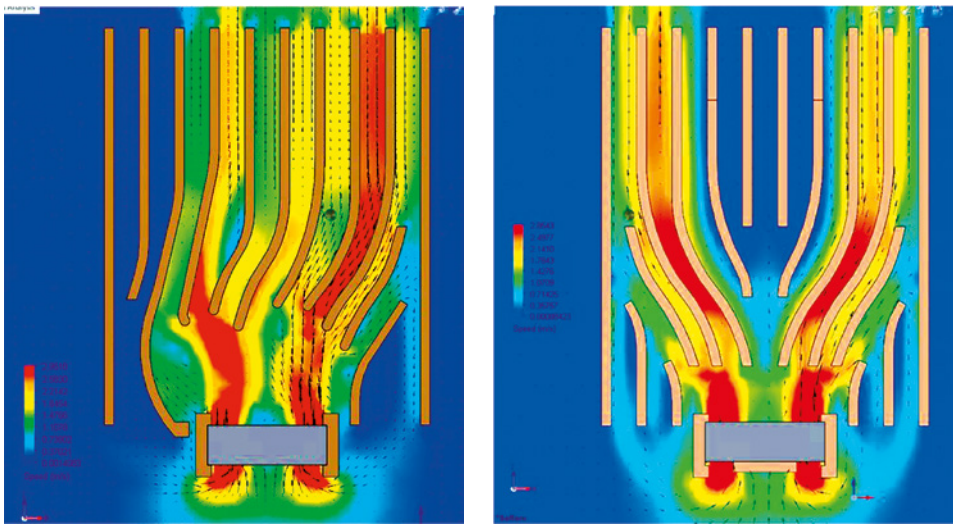

Figure 4. Preliminary Designs of Cast Heatsink Evaluated by Watson Marlow

The success of this project, which has enabled Watson-Marlow to produce an efficient and effectively powered DC dosing pump, has also opened up new markets, for example, in solar powered pumps for remote applications where mains power is not available. Examples of this application are for waste water treatment plants in remote locations or deployed directly into the PLC-BUS for use in factory automation applications.

Design for **Six Sigma in Electronics Cooling**

By Wendy Luiten, Principal, Wendy Luiten Consulting

Temperature problems are well known in the high-tech industry. Everyone knows of cases where overheated products stopping working and in the best case scenario, resume their function, but only after an extensive period of cooling down.

In past years, electronics cooling was very much about cooling consumer electronics, mainly TVs, computers and networking devices. Nowadays we see the focus shifting towards smartphones, tablets, smart watches and other wearable electronics but also towards all other transitional fields in our society. Energy savings? Cooling of LED lighting is a hot item. Energy transition? Cooling of power electronics, for example in automotive electronics, for solar cells and inverters, or in fast chargers for electric cars and electric buses. Also, the trend towards more and more data communication is enabled by cooling of electronics. One can think about cooling of datacenters, servers and telecom equipment for 5G, and the Internet-of-Things

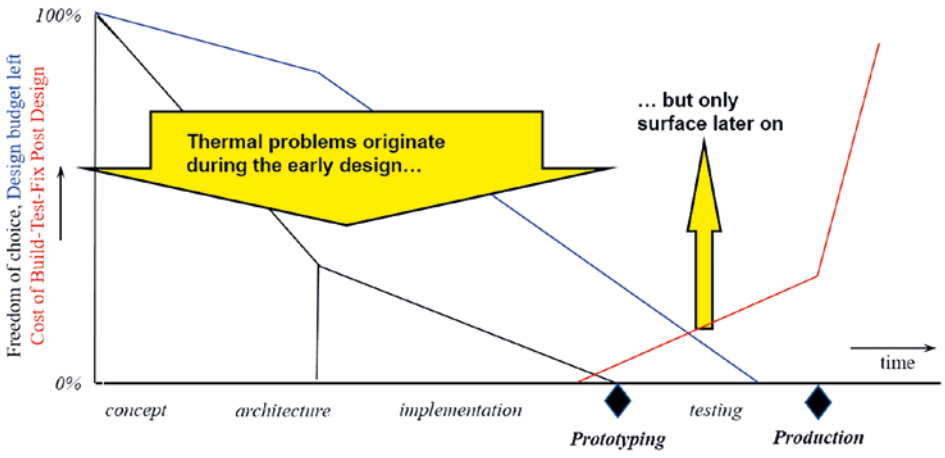

Figure 1. Typical development trajectory

Everywhere energy is stored or transformed, part of it is released as heat. Higher operating temperatures impact reliability, lifetime and safety, hence embedded control algorithms are increasingly used, dialing back performance or switching off once a measured temperature exceeds a threshold value. A good thermal design will keep your product cool and its performance up.
For many product developers, the temperature of their product is something that just happens, and they discover only at the stage that the total product has been realized in hardware and software and is tested for the first time. At that point it is far too late for a cost-effective solution. In addition, problems are difficult to solve, solutions often comprise multiple parts, and are difficult to realize, requiring multiple prototypes and may also involve further fixes when production starts (Figure 1). In the worst case scenario, problems are found in the field with product reliability being impaired due to thermal reasons.

Why are thermal problems so difficult to solve? Simply put, it is because heat flows are so elusive. A typical product has multiple, often interconnected heat flow paths each consisting of multiple steps, and each step represents a thermal resistance. A high source temperature is the result of a high heat dissipation, a high thermal resistance or a combination of both. The heat dissipation is a direct consequence of the functional performance, and usually this cannot be lowered without a performance penalty. This leaves low thermal resistances and short paths as the preferred option to control temperature.

In the ideal case, each heat path from each source to the environment is formed through

a chain consisting of a low number of low thermal resistances. But this is not easy to realize. Part of the thermal resistances is either air-related or infra-red radiation related, which makes them invisible to the human eye. While air might be the primary cooling medium, it does not appear on the Bill-Of-Materials (BOM) for the product, and so far not tracked through the normal change request procedures. Changes that affect the air flow are unseen, and so also likely to recur in the next design iteration, and indeed in the next product development.

Design for Six Sigma (DfSS) is a design philosophy aimed at improving the success of innovation processes. The method is very well suited to the thermal field. Thermal design starts with identification of the product requirements (Define Phase), and flows down to how this translates to the thermal requirements, usage conditions, magnitude and location of heat sources, environment, and the location of temperature critical components (Identify Phase). A good thermal design then provides a robust solution to heat removal, in close collaboration with the mechanical design and electrical design flows (Design Phase) Identification of input parameters, exploration of the solution space and making a conscious design choice are key concepts.

In the very early design phases thermal concept design can be done analytically, using hand calculations and estimations. This has the advantage of additionally identifying the key input parameters that influence the thermal behavior of the total product, but requires experience and engineering judgement. For more complex cases computer simulations are increasingly used, both in the architecture phase and in the implementation phase. For air-cooled products, use of CFD is the highly preferred option, as in this approach both the heat transfer coefficient and the temperature of the cooling air are calculated as a function of the air flow, rather than assumed to be a generic value to a constant air temperature.

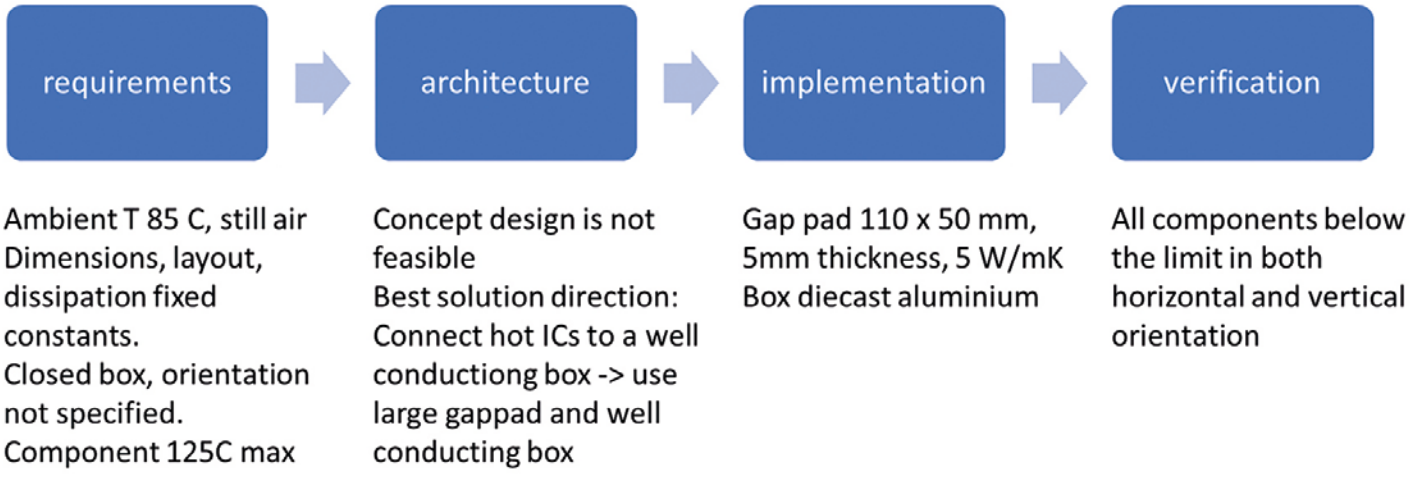

Figure 2. Example of the thermal design process of an automotive electronic box

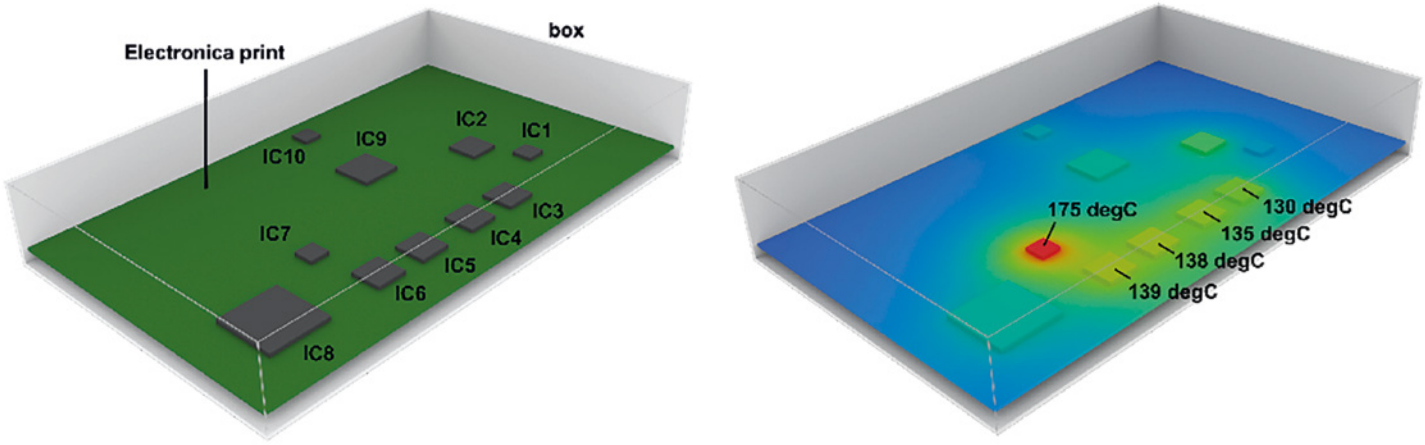

Figure 3. Left: Geometry, Right: Simulated temperature field with the first concept

	0	1	2	3	4	5	6	7	8	9	10	11
Gravity Direction	Negative Z	Negative X	Negative Z	Negative Z	Negative Z	Negative Z	Negative Z	Negative Z	Negative Z	Negative Z	Negative Z	Negative X
box material : Conductivity (W/(m K))	0.2	0.2	0.2	0.2	2	8	15	15	15	130	130	130
gappad for IC7 : Activated	No	No	No	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No
large gappad for hot area : Activated	No	No	No	No	No	No	No	No	Yes	No	Yes	Yes
board : Conductivity (W/(m K))	10	10	10	10	10	10	10	20	10	10	10	10
Heat Sink for IC7 : Activated	No	No	Yes	No	No	No	No	No	No	No	No	No
IC7 : Temperature (degC)	175.6	174.6	148.2	163	154.7	144.2	138.2	129.4	137.1	122.1	121.3	121
IC3 : Temperature (degC)	130.8	128.2	130.6	130.8	130.2	129.8	129.6	122.5	112.6	128.8	103.4	103
IC4 : Temperature (degC)	135.7	133.1	135	135.6	134.7	134.2	133.8	126	115.2	132.4	104.3	103.9
IC5 : Temperature (degC)	138.7	136.4	136.7	138.2	136.7	135.4	134.4	126.9	117.3	131.5	104.9	104.6
IC6 : Temperature (degC)	139.8	138.7	135.9	138.8	136.4	133.9	132.3	125	117.8	127.5	105	104.7

Figure 4. Scenario table with all virtual experiments

In fan-less products the magnitude and the topology of the air flow field and the heating of the cooling air flow are non-trivial and can have very significant impact on the temperature behavior of a product.

In many cases the true power of computer simulation is not in number crunching the detailed mechanical and electrical CAD design as a final check just before production. Rather the true added value is in the use of a series of numerical experiments in the early architecture phase. Choosing the right architecture from the start can free up design space to pursue the most desirable product from the earliest stage. Using computer simulations one can virtually explore the solution space and choose the most appropriate solution direction without incurring the large cost in time and financial resources that would be needed to pursue a similar goal through testing hardware. In the author's experience it would be very common to perform between 10 and 40 computer simulations to finally get to the most optimal architecture for the design. After choosing the most appropriate architecture – tailored to the specific product requirements and usage conditions – detailed design can then follow, often helped by additional computer simulations on the CAD data.

Figure 2 shows the thermal design process of an automotive electronics box as an example. The concept design starts from a closed plastic box containing a Printed Circuit Board (PCB). The thermal requirements are that the box is placed inside a non-ventilated environment of 85ºC with a maximum allowed component temperature of 125ºC. In the architecture phase, CFD simulations start from this concept design, and an aggressive estimate for the powers to ensure that design is robust enough to deal with the anticipated worst case power consumption.

Figure 3 shows the geometry (left). Note that this concept simulation is performed with only a very rough mechanical/electrical model, and without using mechanical or electronic CAD files or data. Rather, these architecture-stage simulations precede the detailed mechanical and layout CAD design which take place after the architecture is chosen. The results of the simulations are shown in the calculated temperature field in Figure 3 (right) show that the proposed concept design is not thermally feasible. Multiple ICs are above the 125°C temperature limit, and the hottest component being 50ºC higher.

In this product, the key parameters for the thermal resistances consist of the dimensions and material properties of the box, the layout and heat dissipation of the PCB, and the mounting of the PCB inside the box. The size of the box, the fact that the box needs to be closed, the layout of the PCB and the heat dissipation are fixed. Parameters that can be changed can all be investigated in the architecture phase of the thermal design. These are:

- The material of the box, especially the thermal conductivity (normal plastic k=0.2 W/mK, thermal plastic – low budget k=2 W/mK, electrically insulating k=8 W/mK or electrically non-insulating k=15 W/mK, or die-cast aluminum k=130 W/mK) are all discrete options
- Thermal management products: using a heatsink and/or a gap pad – a thermally conductive solid material bridging the air gap between the PCB and the box can be investigated, with
  - **a)** Location limited to the hottest component, IC7
  - **b)** Distributed over the entire hot zone
- Thermal conductivity in the printed circuit board itself (layout and construction related, e.g. number of layers, buried power and ground planes).

Another aspect of DfSS is de-risking the potential influences of causes of variation. As the orientation of the box is not prescribed in the product requirements, the box must be simulated both in horizontal and in vertical orientation, since its mounting is under the control of the end user company.

> "I find FloTHERM's Command Centre indispensable for running the many scenarios that I need to make sure that I do not just have a working design, but the best working design possible for my circumstances.

Figure 4 shows FloTHERM's scenario table with all the calculated cases and corresponding results. In the scenario table, each column represents a virtual experiment. In total, the table shows 12 different virtual experiments. In the top blue part of the table the chosen key design parameters are shown. The bottom, orange, part shows the corresponding calculated temperatures for the key ICs on the PCB.

Scenario 0 is the original concept design. In this design IC7 is 50°C above spec and IC3 to IC6 are also above spec. Scenario 1 shows the results for the vertical orientation. It shows that the horizontal orientation can be considered worst case, and we continue subsequent scenarios with horizontal orientation. In scenario 2, a heatsink is used on IC7, and this is not a viable solution. Scenario 3 – 6 virtually explore the use of a gap pad in conjunction with a closed box of increasing thermal conductivity. The results show that a small gap pad to a plastic box does not work, also not if the box is made of thermal plastic and also not if the PCB itself is better conducting (scenario 7). Scenario 8 shows that also a large gap pad to a thermal plastic box does not solve the thermal problem – clearly a metal box is needed.

Scenario 9 shows that a local gap pad on IC7 to an aluminum die cast box solves the problem for IC7, but IC3 is still above specification. Finally, scenario 10 shows that to apply a gap pad over the hot zone in conjunction with a die-cast aluminum box is a feasible solution, and scenario 11 shows that this solution is also robust with respect to different orientations: in vertical orientation, this solution also fills all requirements.

Figure 5 shows the calculated temperature and flow fields for scenario 10, the final solution, in horizontal orientation. The layout of the board is unchanged, and applying a gap pad of sufficient size bridging the air gap between PCB and box in conjunction with an aluminum die-cast box, a thermal design is realized that will keep all temperatures within the specified boundaries irrespective of the orientation of the final product.

The automotive box example illustrates the importance of a good thermal design and a methodical exploration of the solution space.

Using a heatsink on IC7 was not a solution because the heatsink lowered a resistance in a path that contains a very large second and third thermal resistance: the heat transfer of the air inside the box to the wall of the box, and the heat transfer of the box to the environment. In the chosen solution, multiple resistances in the same heat path are lowered through a strategic choice of thermal input parameters.

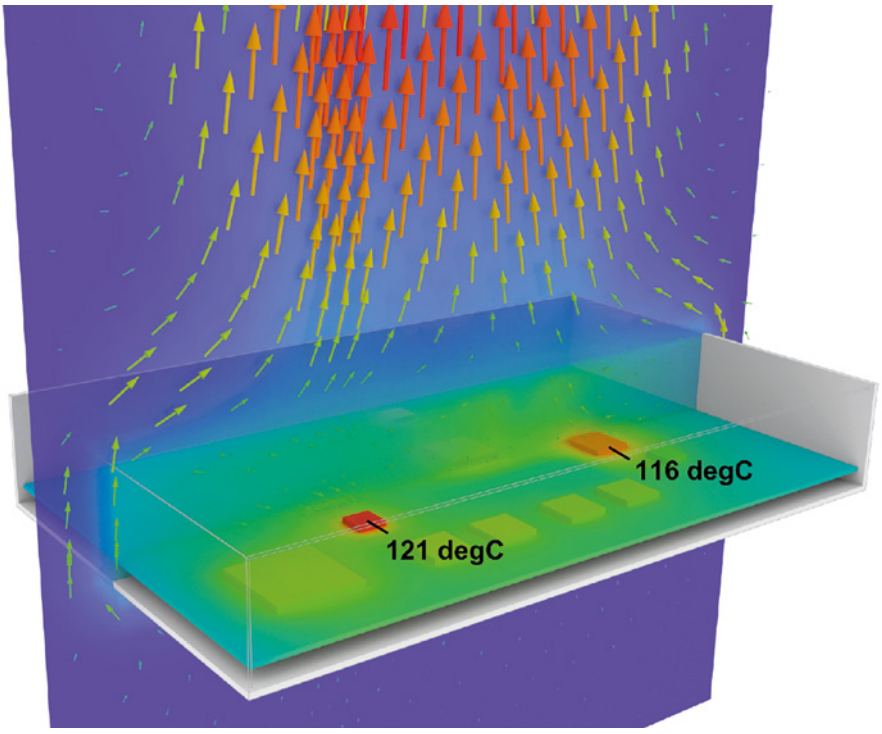

Figure 5. Surface temperature, air temperature and flow field for the final chosen architecture

In the case of the simple automotive box example, purchasing can now proceed to source a supplier for a die-cast box and the gap pad, while in parallel the mechanical developer can proceed to implement the detailed CAD design for a die-cast box. Thermal design in the pre-CAD phases was sufficient to make informed thermal design choices and lower the risk of wasted time and project resources by detailing an unfeasible design to a significant degree. Rather, having de-risked the project from the outset, as more becomes known about the component placement, board layout, component powers etc. the detailed design can explore ways in which the cost can be reduced. In this example, a smaller gap pad may be possible, or a cheaper material with a lower thermal conductivity might perform adequately. Exploration of the solution space is one of the pillars of Thermal Design for Six Sigma as it allows thermal solutions to be found that potentially free up design space to increase product desirability as well. As an example, avoiding the use of the heatsink potentially enables the box to be flatter and reduce its volume claim. It is not limited to the architectural design phase and can be used to very good effect from both engineering and financial standpoints throughout the development.

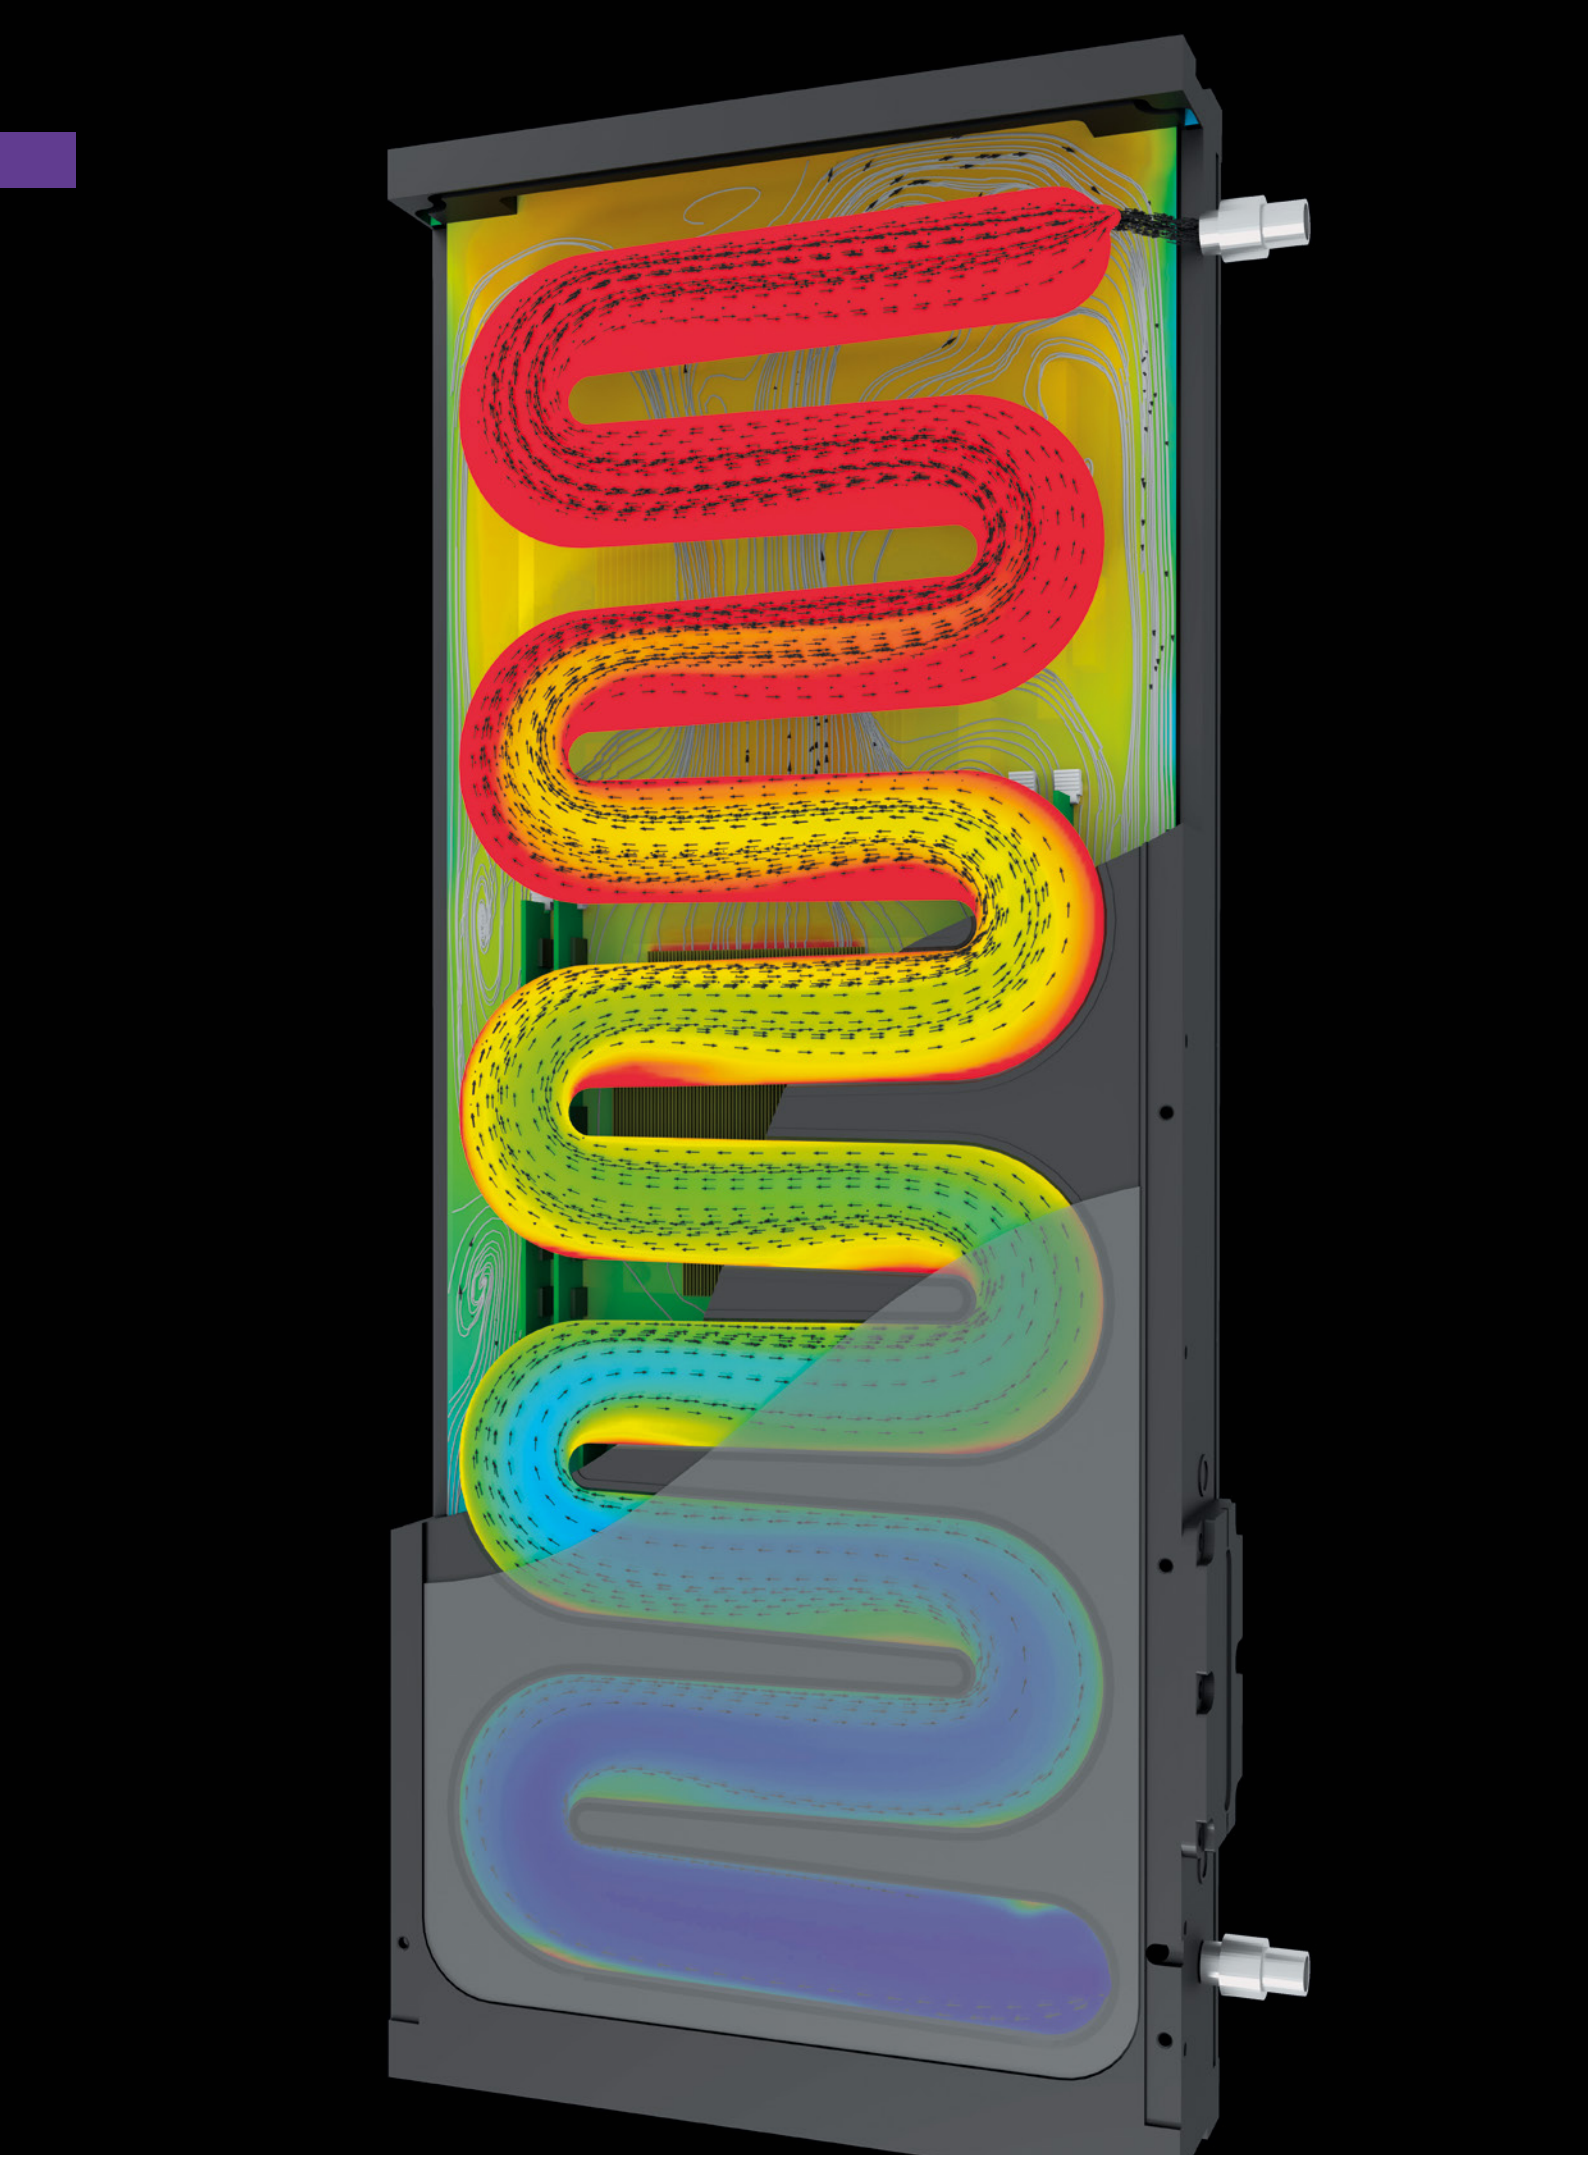

Liquid Cooled Computing

Enabling The Digital Enterprise In The Advent Of Industry 4.0

By Jon Halestrap, Business Development Director, ICEOTOPE, UK

One of the core tenets of mass manufacturing is centralization. It's an idea pretty much as old as the Industrial Revolution itself. Putting everything in one place – raw materials, workforce, machines, post-production – enabled scale and scale meant efficiency.

The advent of Industry 4.0 has disrupted these long established paradigms. Smart factories are challenging the accepted thinking underpinning volume manufacturing.

- **Innovative manufacturing techniques** – 3D printing has undermined the relationship between volume and cost efficiency. The significant capital expenditure of investing in new tooling meant that any part had to be produced in high volumes to be justifiable.
- **Increasing automation** – factories of the future will be ever less reliant on a local workforce to staff the production line. Factories will be part of a distributed network and contain machines that are more capable of intelligent operation independent of human interaction.
- **Autonomous vehicles** – automated machines will shift the logistics of manufacturing. Rather than relying on a centralized distribution network, manufacturers will use autonomous vehicles, robots and drones to move parts between factories or deliver products directly to retail and even end-users.
- **The "Digital Twin"** – from initial designs through to simulation, materials and production, the digital twin concept will accompany every machine in order to create, test and build products. Taking a digital approach for critical parts of production will open new ways of productivity and efficiency.

This fundamental shift from centralized manufacturing to a flexible network of smart factories is supported by micro and edge data centers. Just as factory hardware no longer has to sit under a single roof, IT systems will also be spread across a network.

Datacenters and high performance processing systems will sit throughout a manufacturer's network, perhaps on the factory floor itself. Through wireless technology, computers will be in constant communication with other smart factories in the network, the cloud, and a range of connected devices.

Digitalization Puts Strain On Existing Infrastructure

Despite these exciting advances in automation and communications, let's not forget that manufacturing is still about making products – it's a physical business. Machining parts can generate huge amounts of noise, heat, and dust which combine to create a very challenging environment for IT systems.

As systems will be 'edge' deployed as part of a decentralized network they could well exist in remote locations where local on-hand IT support is not guaranteed. Therefore it is paramount that smart factory devices are easy to manage, secure, and resilient to the environmental challenges of the factory floor.

As new technologies such as the Internet of Things (IoT) are driving volumes of data, along with substantial investment at the Edge of Network, it's clear that cost effective and efficient cooling technology is needed.

Cooling is essential to high powered electronics and should not burden your IT investment. Traditional air-based cooling equipment cannot cope with ever-increasing

heat loads that hotter processors and applications demand. Sticking with the legacy approaches will only lead to larger footprints and increasing cost and complexity, with no competitive benefit.

Roughly 25% of datacenter unplanned outages are caused by weather, water, heat or air-conditioning related issues. Combine this with air cooling equipment being bulky, power-hungry and costly, sufficient cooling methods are crucial for business operation and continuity.

Redefining The Cooling Landscape For Industry 4.0

Liquid cooling has been around for many years but, until now, has been regarded as a niche technology with many compromises. However hotter processors and new technologies are driving the need for faster speeds, better flexibility with minimal disruption and downtime.

ICEOTOPE, developers of immersion liquid cooling technology, enables IT infrastructure to operate seamlessly in any environment. The patented technology revolves around immersing and protecting high powered electronics in a specially engineered coolant.

Thanks to Iceotope's total immersion cooling technology, expenses such as chillers, computer room air handling (CRAH) equipment, raised floors or ducting are no longer necessary. The result is an advanced product portfolio ranging from tower to server-level to suit a range of computational demands and environments.

EdgeStation™ - Fan-Free, Resilient Workstation

EdgeStation adds robust, portable powerful computing resource to the digital enterprise. Bringing critical applications closer to the manufacturing operation reduces latency, therefore improving your time-to-market.

Without the industrial drone of fans and pumps, you can put a high performance machine to work side-by-side with your team. Removing the fans ensures a sealed system which protects critical IT components from hot, harsh or contaminated environments delivering reliable performance.

ICEOTOPE has worked with the Advanced Manufacturing Research Centre (AMRC) to solve the problem of dust particles from its carbon fiber facility. This caused continuous IT equipment failure and unscheduled downtime. The AMRC deployed an EdgeStation™ system, which fully encloses and protects

Figure 1. ICEOTOPE EdgeStationTM fan-free resilient workstation

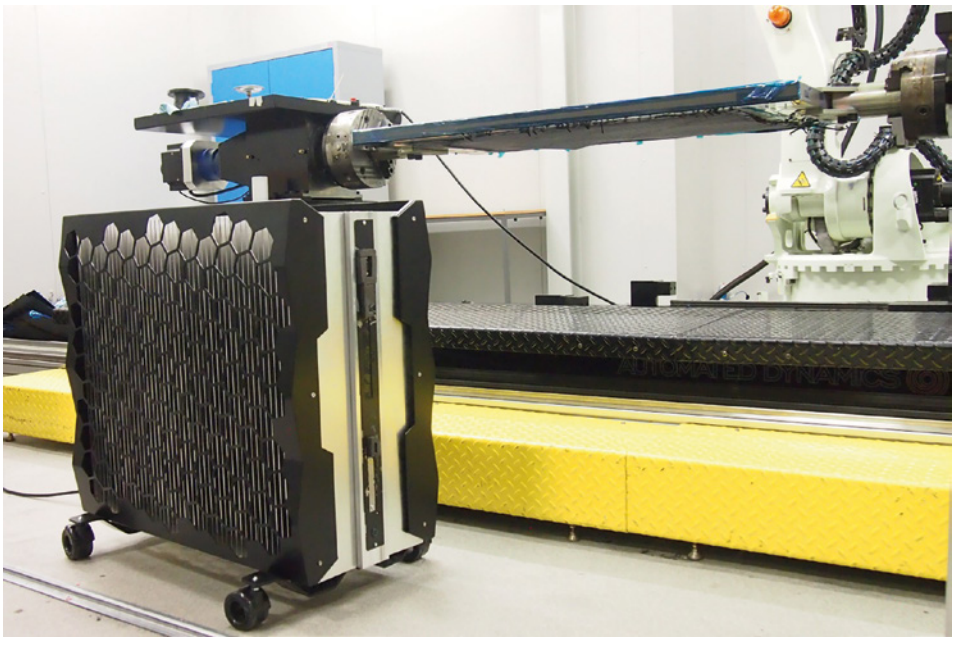

Figure 2. EdgeStationTM deployment at the AMRC carbon fiber laboratory

high-powered electronics to sustain performance and therefore increase product lifespan.

EdgeRack - Simplified Edge Data Center Deployment

EdgeRack is the modern data center to deliver more power in significantly less space so your business can scale one rack-at-a-time.

ICEOTOPE uses a two-stage cooling process – the first stage being a specially engineered liquid to cool all components with the second stage transferring the heat away using a coolant loop. Consider that EdgeRack saves up to 80% energy consumption, and you can reuse the waste heat for district heating, liquid cooled datacenters become a sustainably smart business model.

Summary of ICEOTOPE liquid cooling:

- **Double-digit capex savings** – Iceotope's technology requires next-to-no additional infrastructure leading to significant capex reductions for datacenters..
- **Significant floor space reduction** – using anywhere between 50% - 75% less floor space, savings for customers are substantial. This leads this need to defer or obviate major capex in extending or building new datacenter space.
- **Full integration** – Iceotope can retrofit their servers into existing infrastructure meaning you can get immediate opex benefits without having to redesign your entire datacenter, alter your supply chain or retrain your staff.
- **Consumes less energy** – without the need for power-hungry computer room air handling equipment and zero fans inside the servers, Iceotope's technology can reduce energy bills by up to 80%. Iceotope also allows the recapture and reuse of waste heat leading to an improved corporate risk and social responsibility strategy.

CFD Is Vital For Product Development

The combination of network modeling and CFD are vital to our engineering design processes and greatly reduce the cost and time for development. From prototypes, to performance optimization and design approval, CFD modeling is used for a range of features within the immersion cooling system from natural convection of IT components to forced convection of pipe networks and heat exchangers.

Liquid is 1,000 times more effective at transferring heat than air, however, using liquid comes with challenging physical behaviors. To overcome this, ICEOTOPE uses CFD software to analyze the natural convection inside our server blades, providing details of flow distribution.

To deliver significant energy and cost savings to our customers, we also use CFD to characterize hydraulic resistance in our manifolds to achieve sufficient coolant flow while limiting the operating pressure.

Harnessing the power of digitalization, ICEOTOPE can now utilize its CFD design simulations to create their own "digital twin" to validate and optimize designs.

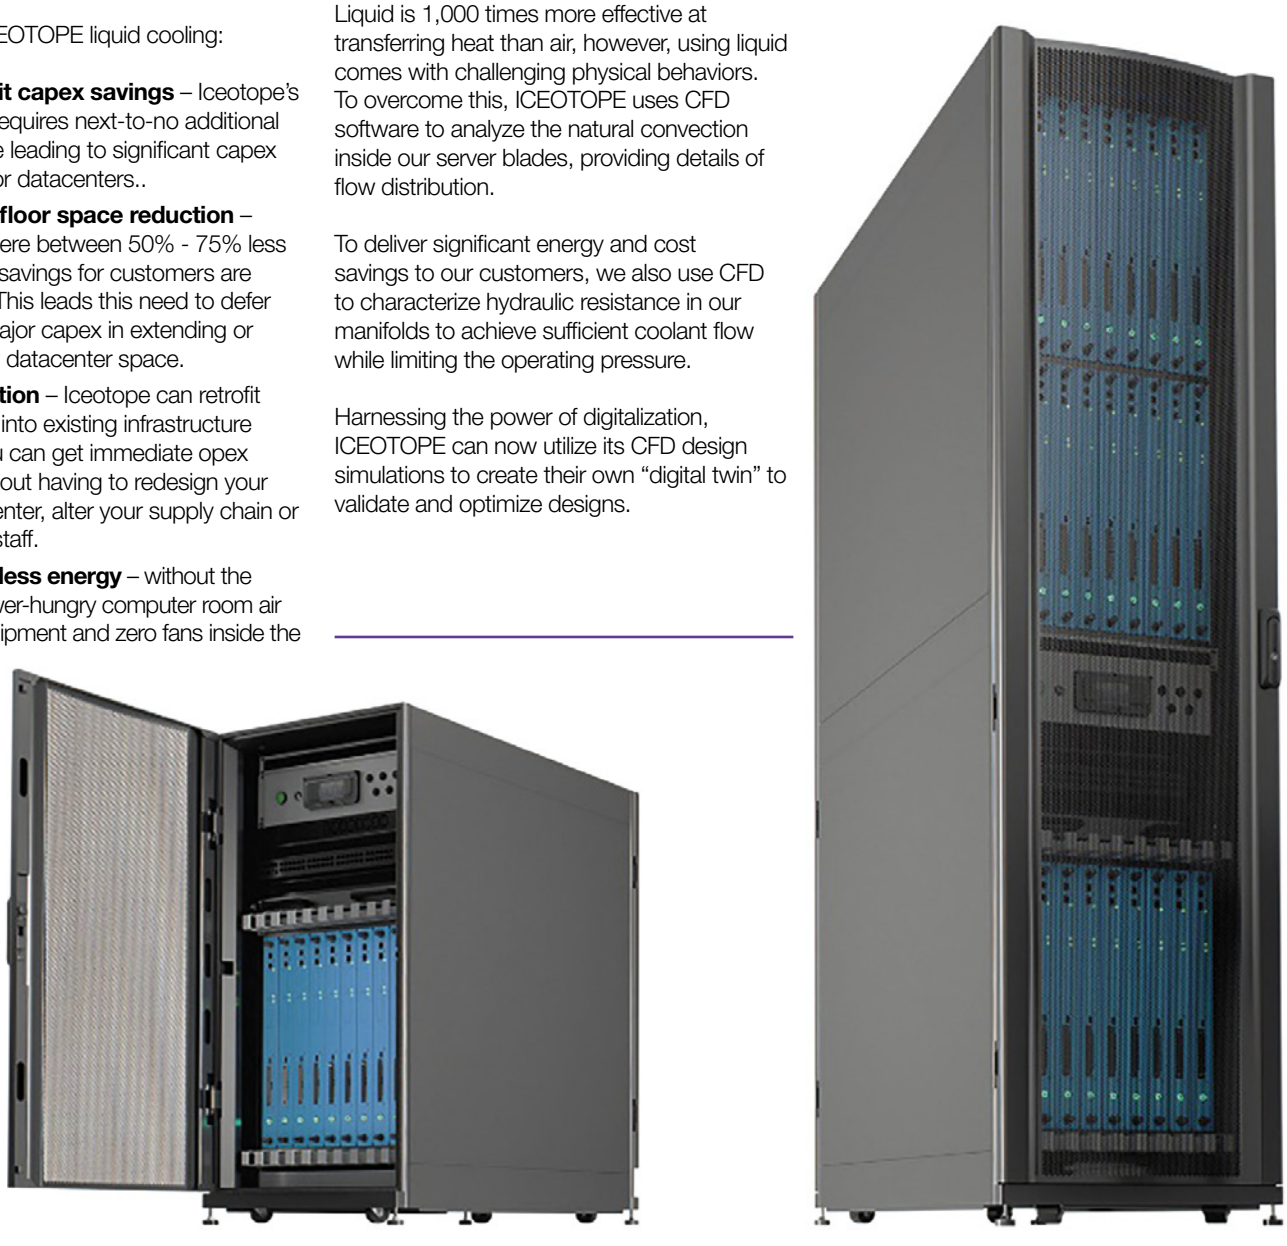

Figure 3. From one to three chassis, ICEOTOPE delivers business scalability

Ask The **GSS Expert**

FloTHERM Advanced Training from GSS

The FloTHERM Advanced Training course is designed to follow the introductory class. Some of the topics, such as advanced gridding and troubleshooting, advanced component level modeling, handling advanced applications in the command center, and import/export options are extensions of material discussed in the introductory class. Other topics, such as transient analysis and joule heating, are presented for the first time in the advanced session.

The transient analysis section begins with a discussion of the reasons for choosing a time dependent simulation as opposed to steady state. Then the modeling options available in FloTHERM are presented along with the steps required for set up. The hands on lab session provides an opportunity to work through a transient analysis from set up to post processing during the class.

The joule heating session is designed similarly to the transient analysis one in that the lecture discusses the theory, and then continues through the set up to post processing in FloTHERM. There is a hands on tutorial for this subject as well.

The advanced gridding and troubleshooting presentation begins with a quick review of topics from the introductory class and then moves on to advanced techniques. There is typically much discussion and demonstration throughout this portion of the class.

Advanced component level modeling also begins with a quick review of the topics covered in the introductory class. This is followed by a more in depth discussion about resistor networks as compared with resistor capacitor networks, how to acquire them and when to use them.

The session covering advanced applications using the Command Center offers a more applied look at the automated parametric analysis and optimization capabilities available in this module. This session includes demonstration as well as a hands on lab session.

> "The course very well showcased V12.0. As well, the various small tips and best practices mentioned by both the instructor and other users in the class is certainly invaluable information.
>
> **Jeffrey Auclair, Hardware Engineer, Dell EMC**

Finally, the lecture on import and export options available in FloTHERM begins with a quick review of mechanical CAD and EDA import capabilities and then progresses to the import of power maps, scripting options as well as bridges to other analysis tools.

Contact Barbara Hazard at barbara_hazard@mentor.com with any questions regarding registration or the topics covered.

INTERVIEW

Takuya Shinoda, Denso, Japan

DENSO
Crafting the Core

Q. Tell our readers about Denso Corporation?
A. Denso is a supplier of advanced automotive technology, systems and components for automakers worldwide. We are headquartered in the city of Kariya, Aichi Prefecture, Japan and last fiscal year we turned over $40Bn with approximately 150,000 employees globally. I work in R&D in the Substrate Hardware Development Department specializing in thermal analysis technology for next-generation Electronic Control Units (ECUs) development. I am therefore tasked with continuously improving our ECU thermal performance and seeking out size, weight, and cost reductions.

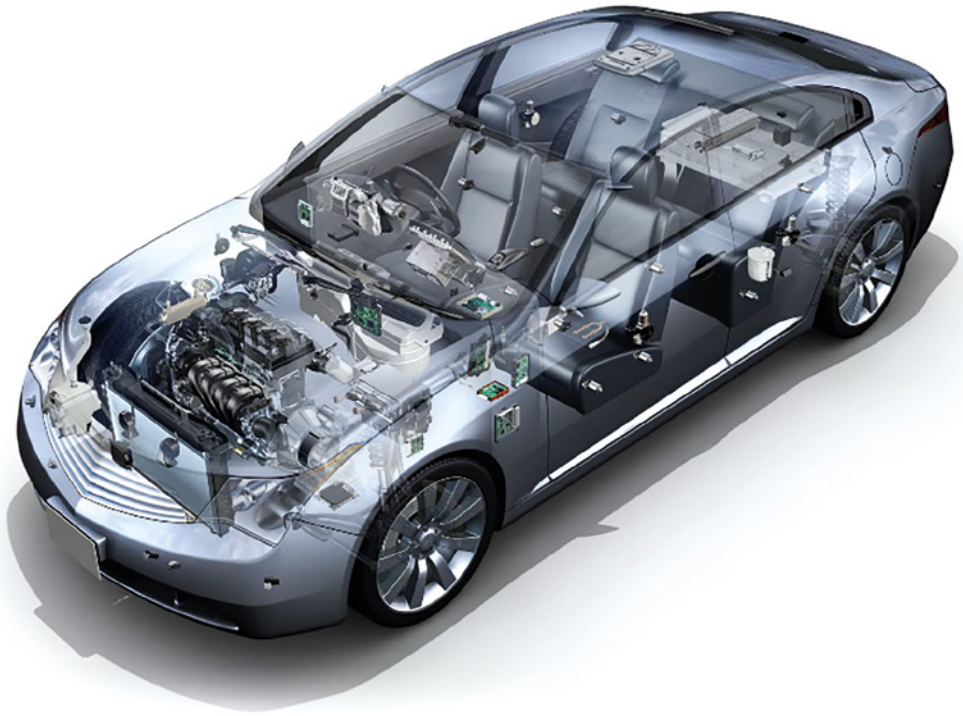

Q. What is your background and your current role as a thermal expert in Denso?
A. This year marks my 31st year inside Denso having joined as an Electronics Engineer in 1987. We use CFD analysis based techniques to improve our products performance - to make our ECUs more reliable, robust, cheap, safe and secure. In 2006 I decided to bring the FloTHERM CFD product from Mentor into our department in order to frontload the thermal design of ECUs. Frontloading CFD with FloTHERM makes our simulations faster and quicker than alternatives. By 2015 we had managed to compress our thermal verification cycles by 62%. Since then, this effort has led to a 78% cost reduction in designing ECUs over a 10 year period.

Q. What are the big trends in automotive ECUs today and in the foreseeable future?
A. We see an inexorable rise in ECUs inside cars now with most having 50-100 on board today. While electrification of vehicles (including EV/HEV) and the emergence of autonomous cars have a big impact on ECU design, we see motors and ECUs associated with those significant emerging trends. Generally we observe that we need to take into account the wider ECU thermal environment in our design processes. I don't see liquid cooled ECUs as a feasible option as engine ECU in the near future because of reliability and cost issues. Smaller and smaller ECUs are coming and higher density of electronics in each ECU. We therefore need to gain deeper insight into components in these complex environments. On the other hand, there is also a need for combining the functions of multiple individual ECUs into a centralized ECU. However, these will have much more challenging thermal design factors to take into account. And a CPU with its higher clock frequencies than what's currently employed will require protection from electromagnetics for instance.

Q. What is your view of trends in thermal simulation & test for ECU design?
A. Now accurate CFD simulations are mission critical to realize frontloading design practice. At the same time, the actual measurement is becoming more important each year because the actual measurement helps us to reach the required accuracy. We bought our first T3Ster in 2012 so that we can use measured thermal data from the PCBs to improve our FloTHERM simulation results. We now have very high accuracy thermal predictions. Five years ago we were using FloTHERM for steady state thermal analysis. We also had to do transient thermal simulations to accurately measure Tj in car drive cycles (actual driving cycles). Therefore, it has become more important for

us to apply transient thermal measurement into actual measurement. Today, we use transient thermal simulation to analyse ECU thermal performance in accordance with car drive cycles. I view myself as a thermal analysis 'frontloader' who does CFD simulation with T3Ster data. T3Ster will be indispensable for creating an accurate FloTHERM analysis.

Q. What are your views on component verification in ECU thermal design?
A. Another challenge I see for assembly makers to frontload their ECU design simulation efforts is a lack of standardized analysis models. We are usually provided models and datasheets from IC makers when they supply us components. There is no standardization of such component models. IC thermal Resistance model (Rth Model) is standardized as DELPHI model based on JEDEC standard. To create a DELPHI model, IC manufacturers need to convert a thermal detail model into a thermal resistance network. However, there is no standard for detailed models. Thus, the accuracy of DELPHI model relies on each IC manufacturer. Also, there is no global standard for Rth-Cth models to be applied to transient thermal analysis.

Today Denso uses T3Ster to characterize IC components coming into Denso so we have accurate values of their thermal resistance and thermal capacitance. I think we should be able to be provided standardized IC thermal datasheets from manufacturers based on T3Ster-FloTHERM models in future.

Q. Tell us about the challenges in getting enough thermal analysis engineers?
A. As the demand for frontloading approach is growing inside Denso, we adopt the tools such as FloTHERM to do CFD simulation. However, a big challenge we and other companies are facing, is a shortage of thermal analysis engineers to use the tools we have. This is partly because thermal analysis of electronics at university is not part of the curriculum. This needs to be addressed collaboratively in future if the annual 20% growth in CFD is going to have enough practitioners of the discipline to satisfy demand from companies who need to use the tools.

> By 2015 we had managed to compress our thermal verification cycles by 62%. Since then, this effort has led to a 78% cost reduction in designing ECUs over a 10 year period.

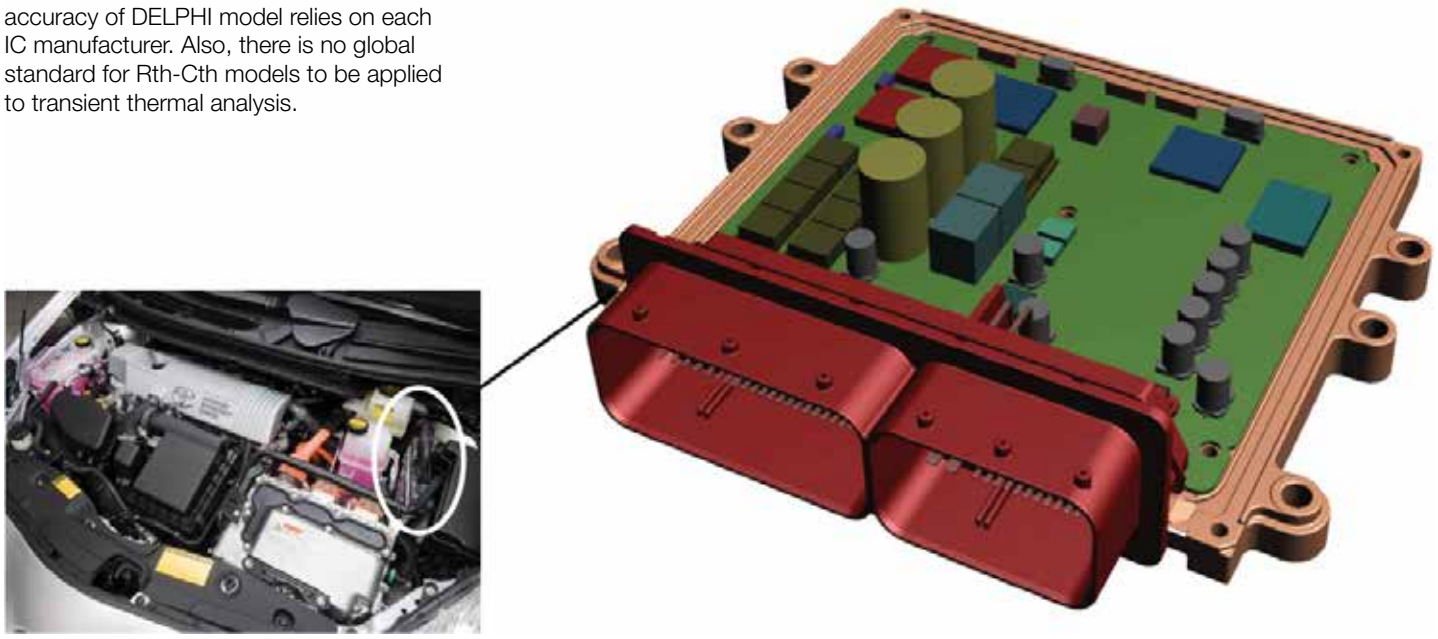

Figure 1. Typical Automotive Engine compartment and Denso ECU geometry

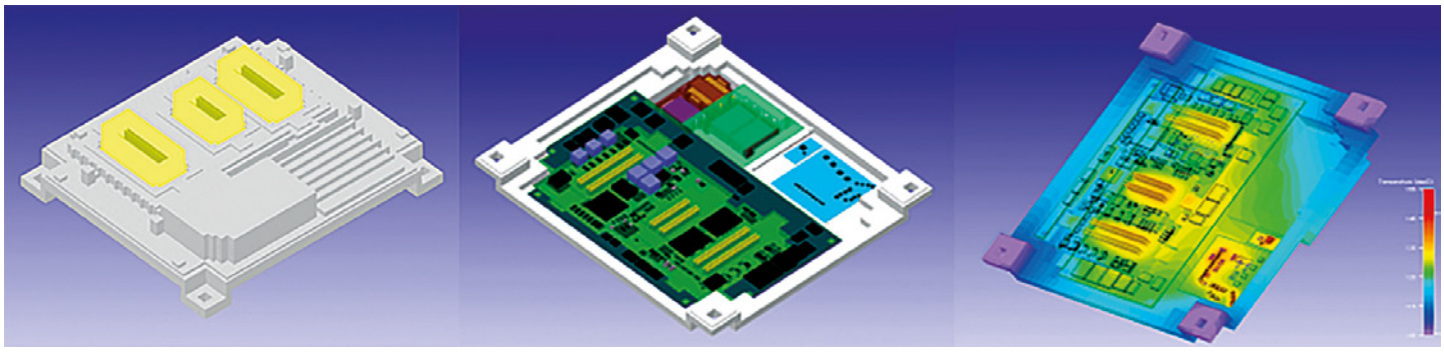

Three Thermal Simulation & Test Innovations for Electronics Equipment Design

By Takuya Shinoda, Denso Corporation, Japan

DENSO
Crafting the Core

Over the last 50 years the automotive industry has been going through huge shifts in electrification (Figure 1). And, of course, we are seeing the electric motor starting to replace the mechanical engine with the advent of Electric & Hybrid Vehicles. Not surprisingly, the demand for ever higher density in a confined electronics, leads to a greater and greater need to resolve the resultant thermal issues.

Thermal simulation technology for automotive electronic equipment design is shifting from steady state to transient analysis. Here, I explain detailed Data Network Resistance Capacitance (DNRC) transient models, which are based on Mentor's T3Ster® tester data. In this article, three innovations to methods for creating DNRC models useful in electronics thermal design are outlined:

1. Placement of measurement locations in models so that the gap between actual measurement and analysis can be identified
2. Preliminary verification of component Tj (junction temperature) value by CFD thermal analysis
3. Creation of compact models from manufacturer datasheets and Tj calculations – what we call a DSRC, Datasheet Resistance Capacitance model.

If we consider a typical automotive engine control unit (ECU) in Figure 2; it can be noted that a typical engine underhood compartment in a car hits relatively high temperatures of approx. 105°C and since electronic parts typically have an upper practical operating limit of which threshold must not be surpassed otherwise electronic component quality cannot be guaranteed.

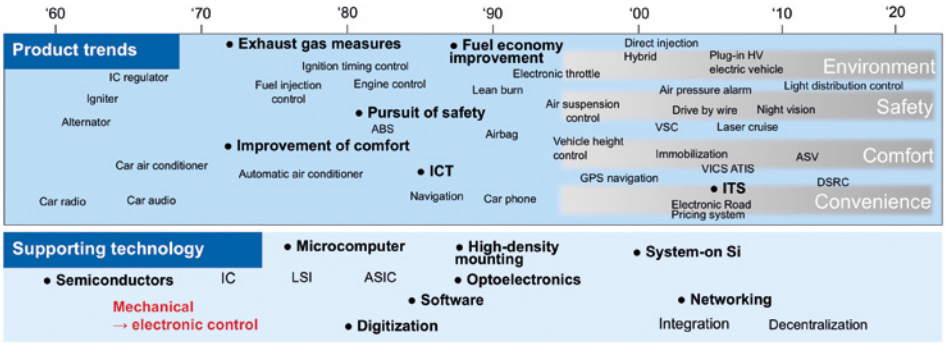

Figure 1. Progression of car electronics product adoption and underlying trend over the last 50 years

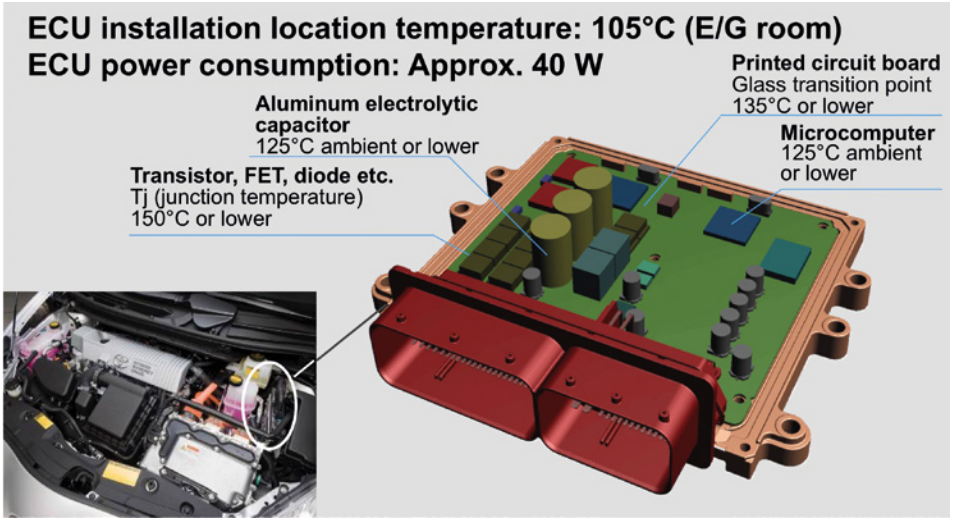

Figure 2. Typical ECU and engine use case and thermal limits

Mentor Graphics
Mechanical Analysis

Moreover, we are seeing that with the advent of more vehicle electrification and with autonomous vehicles on the horizon (Figure 3), there is a need to thermally simulate transient heating effects in ECUs more and more, especially while the advent of Advanced Driver-Assistance Systems (ADAS) increases the number of high-frequency ECUs. Hence, it is necessary today to be able to estimate instantaneous heat generation inside ECUs. Transient thermal analysis needs by JEITA in 2013 were identified to be:

- Tj response during LED lighting operation,
- Instant transient analysis of Transistor, FET, IC and other semiconductors; with measurement of junction temperatures,
- Hot spots during SoC (System on Chip) operation,
- Tj response from engine OFF (dead soak), and
- Business need for the verification of accuracy between measurements and thermal analysis.

The target times for transient thermal analyses depends on whether we are dealing with a chip, a package or an assembly (like an ECU) as illustrated by Figure 4.

From the early days of the electronics cooling simulation market there has been some sort of approximation of chips and packages in terms of their thermal performance and the details inside the chip. DELPHI compact thermal models appeared in the 1990s and quickly became the standard way of modeling packages for CFD simulation that also allows for manufacturers to protect their intellectual property. However, the models do only satisfy transient analysis in terms of the accuracy, because the DELPHI compact thermal models are solely represented by thermal resistance values. Figure 5 depicts Denso's way of looking at all types of analysis models that could be employed to solve our thermal simulations. It is the DNRC and DSRC models that we want to investigate in this study because of their inherent high accuracy and ability to be distributed into the semiconductor supply chain. Based on non-destructive electronics thermal transient testing methodology, DNRC is modeled by 'structure functions, featured by Mentor T3Ster tester (Reference 1). While the form of DSRC is the same as that of DNRC, DSRC model is based on the transient response values in IC component datasheet.

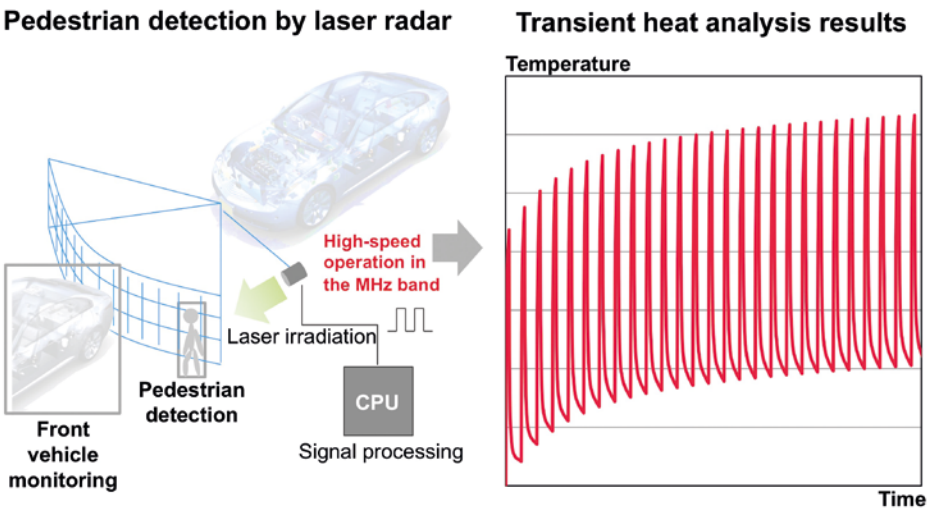

Figure 3. Escalation of transient heat issues with ADAS technology

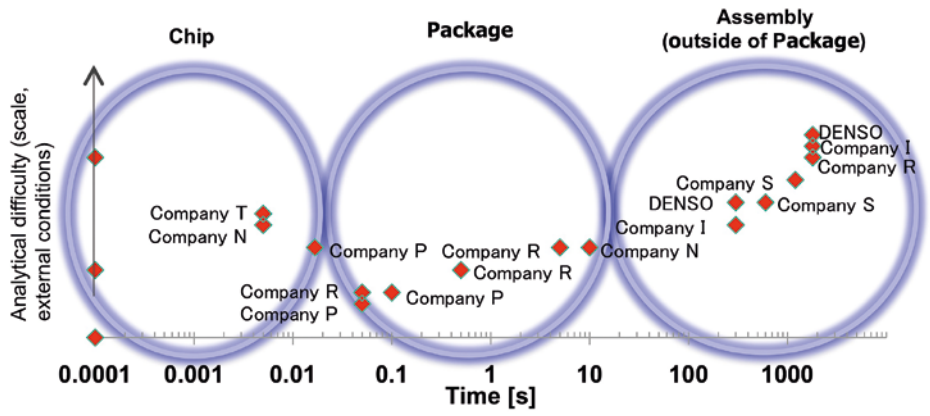

Figure 4. Target time for transient thermal analysis for Chips to Assemblies

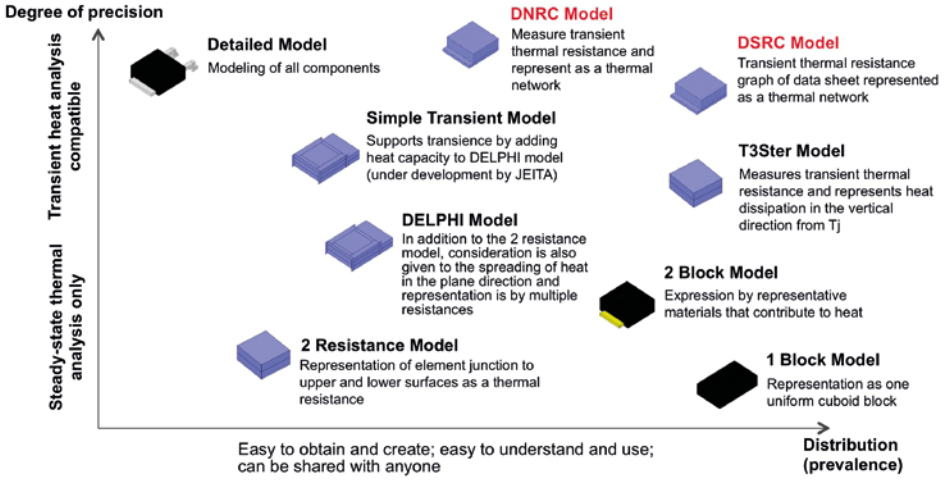

Figure 5. Development of package element models that meet ECU transient thermal analysis needs

This article is not going to go into the details of the T3Ster RC approach for measuring electronic element constituent performance, and how it connects to CFD simulation thermal analysis tools like FloTHERM – see Reference 2 for further details. However, an outline thermal simulation & test workflow is shown in Figure 6 that results in a useful

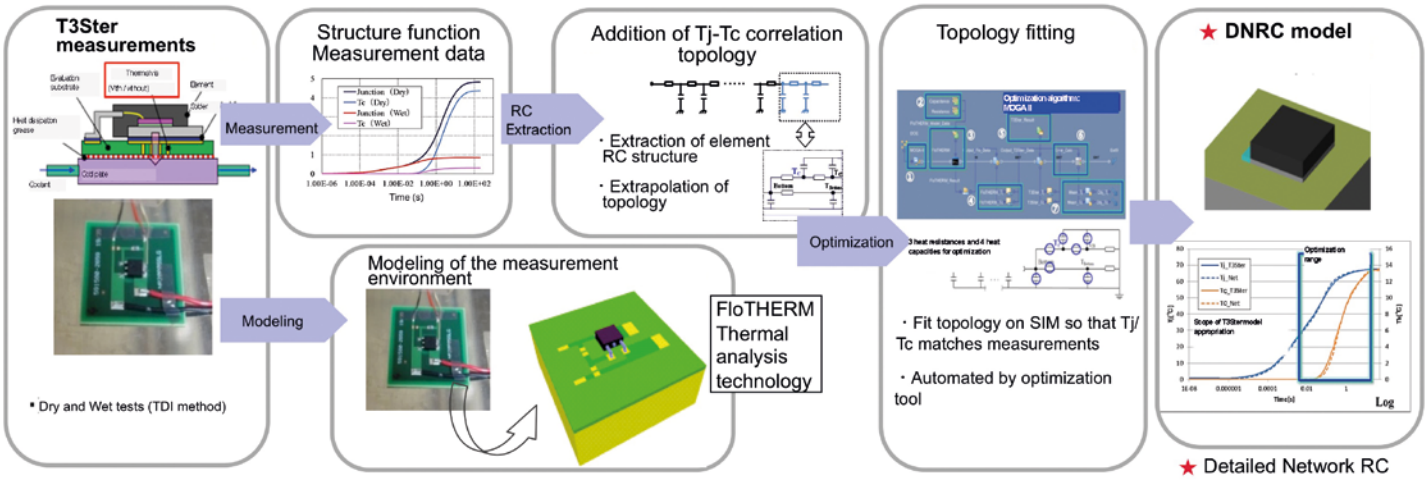

Figure 6. Outline of RC model creation process for a DNRC model by T3Ster measurements & FloTHERM simulations

DNRC model. Here, I emphasize two advantages as below.

- The approach allows for reference point thermal measurement of components in their mounted state on a PCB
- The measured data can be fed into CFD simulation software such as FloTHERM™

Indeed, the beauty of this approach is that it is possible to create a model even if information such as internal dimensions, specific heat, performance etc. of the component is not readily available. The final DNRC is very high precision, because the model is derived by measuring the thermal resistance and thermal capacitance values. Moreover, because DNRC is a thermal network model and has the Tc node existing for verification by measurement with a thermocouple, the values of the model can be compared with the actual temperature measured at Tc for reconciliation. Crucially, this modeling approach can be used by a component manufacturer who does not know the internals of the elements involved. We have completed a validation study (Figure 7) for a typical electronics component package and heatsink array which calculated the relative accuracy of the different types of models outlined in Figure 5. Using the 'Sum of Squares' error estimation method it can be seen that the most accurate approach (the lowest number in the table) is the DNRC approach.

If we now apply this approach to the measurement of transient thermal resistance of a PCB, it is then desirable to measure both the semiconductor component and the PCB simultaneously. We did this with an experiment involving two test circuits with different thermal resistance and

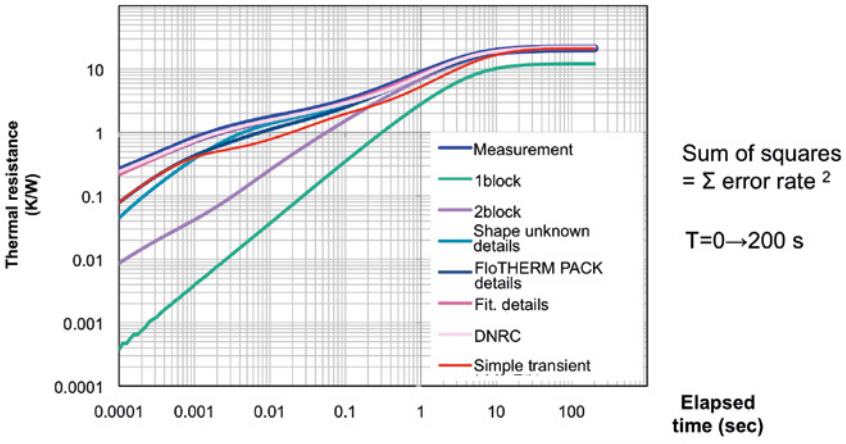

1block	2block	Shape unknown details	FloTHERM PACK details	Fit. details	DNRC	Simple transient
121.0	86.2	49.0	41.5	12.7	0.3	50.3

Figure 7. Results of element model accuracy verification Predictions versus Measurements

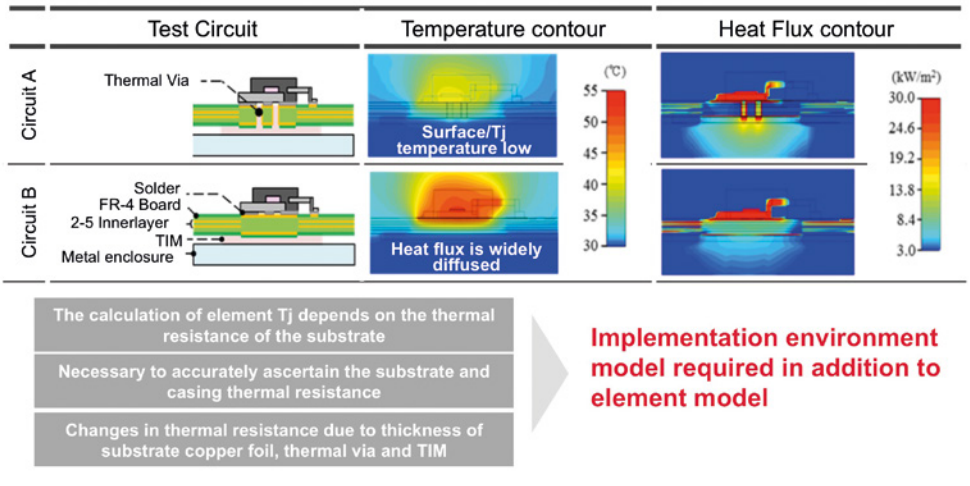

Figure 8. FloTHERM simulation verification of element temperature and heat flux distribution for a circuit with and without a thermal via

capacitance values – one with a through silicon via and one without (Figure 8). The choice of the through Si via meant that the thermal resistance of the PCB without it is overwhelmingly higher than that of the element as a whole – the element being 3K/W with the substrate at 10-20 K/W. Hence, these simulations prove that an environment model in which the element model is located matters and is required for accurate results. There needs to be measurements with T3Ster of the transient thermal resistance of the element in mounted state on the PCB so as to produce thermal resistances equivalent to those in the actual product. This can be seen in Figure 9 where T3Ster confirms the thermal resistance at each point of the two circuits with and without vias being considered in the study. We even examined slight knicks (cracks) in the two structures and saw that structure functions of inside and outside the semiconductor could be discerned correctly.

If we use T3Ster measurement result to calibrate a detailed thermal model for FloTHERM, we devised a seven step workflow process (Figure 10) to accurately obtain the physical property values of any PCB. It involves using FloTHERM coupled with powerful design space optimization tools in order to search for the optimal solution to curve fit onto the actual T3Ster measurements. This in turn reveals accurate physical property values in terms of the thermal conductivity and specific heat of the actual substrate layers. This data can then be used as shown in Figure 11 to evaluate the two circuits we described earlier. And we were able to simulate 1,000 designs in FloTHERM and fit the results to the structure functions under actual measurement. However, the thermal conductivity in copper foil shows 485W/mk, which is not a realistic value, assuming the thermal conductivity for pure copper is as low as 398w/mk. Finally, we were able to do microscopic cross-section measurements (Figure 12) of the PCB substrate with different layers of copper foil thicknesses as a cross-validation of this T3Ster and FloTHERM approach. It revealed that the copper foil is 1.3 times thicker than the design data. When it is converted to the thermal conductivity, it becomes 1.3 times. We got remarkably good agreement and it shows our approach is valid.

The biggest issue facing manufacturers and assembly makers is in measuring actual Tj values (Figure 13). Device manufacturers usually guarantee their electronic components to less than 150°C but how

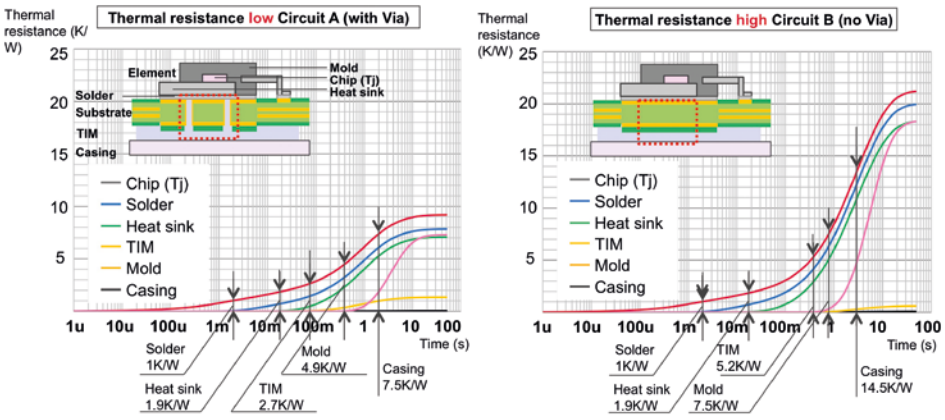

Figure 9. Temperature rise start time and thermal resistance of component for a circuit with and without a thermal via using T3Ster

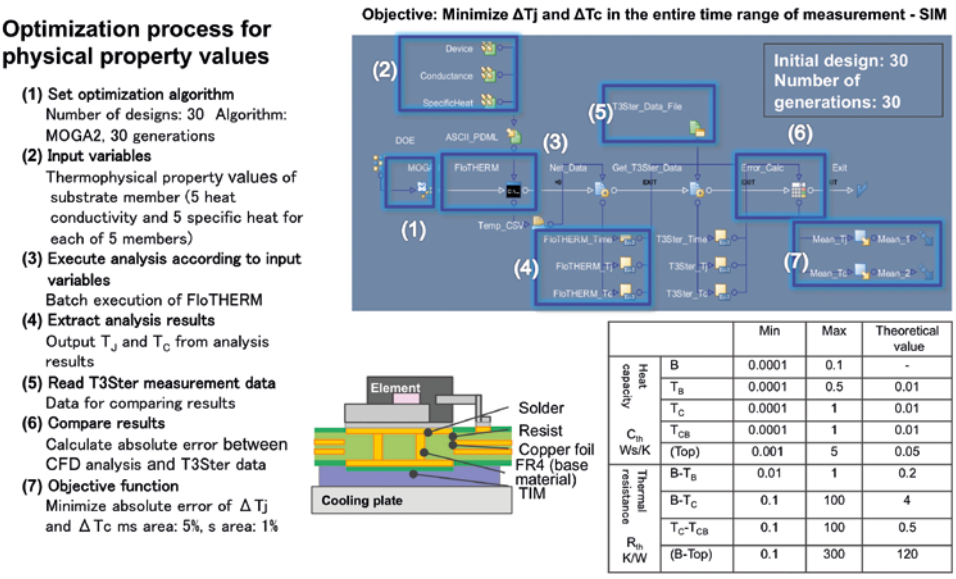

Figure 10. Method of determining physical property values of printed circuit boards (PCBs) using FloTHERM and T3Ster

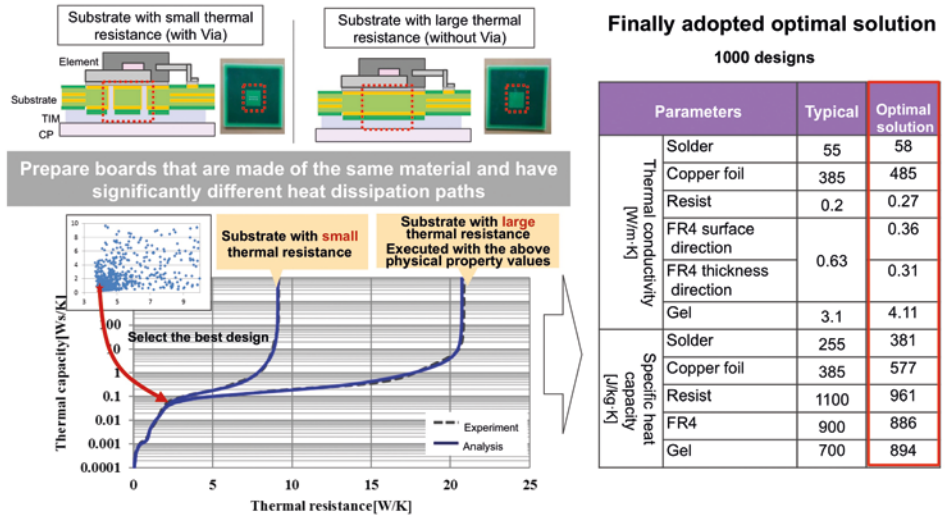

Figure 11. Optimal circuit thermal solution selected from multiple designs

to measure Tj accurately and how to see instantaneous changes in temperature as the device operates? Usually the Tj can be measured by a thermocouple and derived from an electric power waveform analysis and transient thermal graph in datasheet (Figure 13).

If we look at our first proposed innovation, that is, the temperature of the part for coupling with thermal analysis is known. Figure 14 shows this for Tc and Tj values with the various model types outlined in Figure 5. With detailed model and DNRC model, the Tc node enables a comparison with experimental results to ascertain the error of the CFD thermal analysis.

Our second innovation, preliminary verification of the Tj value by FloTHERM thermal analysis is shown in Figure 15. It can be seen that only a detailed DNRC model with a Tc node derived from FloTHERM and T3Ster works well and internal package information is simulated accurately such that a transient Tj and Tc can be estimated in a short period of time.

Our third innovation is that it is now possible to create models based on manufacturer's datasheets and calculate actual Tj values. This is shown in Figure 16 where it can be seen that a DNRC model captures the actual transient thermal resistance in the component's mounted state thus leading to an ultra-high precision model that conforms to actual operating conditions. On the other hand, a DSRC model is based on a transient thermal resistance graph. Hence, the value in the datasheet and the DSRC model based on the datasheet are both guaranteed by the manufacturer. Table 1 shows a relative comparison between the recommended usage of DRNC and DSRC models. As noted in the table, a DSRC approach is a great bridge between analysis models and component manufacturer's guaranteed property values in their datasheet. Moreover, as the DSRC matches the manufacturer's datasheet value, it can be guaranteed for product design. If the data applicable in product design can be converted into a model for analysis, both component manufacturer and assembly maker can easily exchange the models, hence both welcome the approach. Another benefit to this approach is relatively lower costs compared with other measurement techniques that have high measurement and labor costs.

Finally, we carried out a validation test in a standard transient thermal analysis JEDEC still air chamber of a component on a PCB

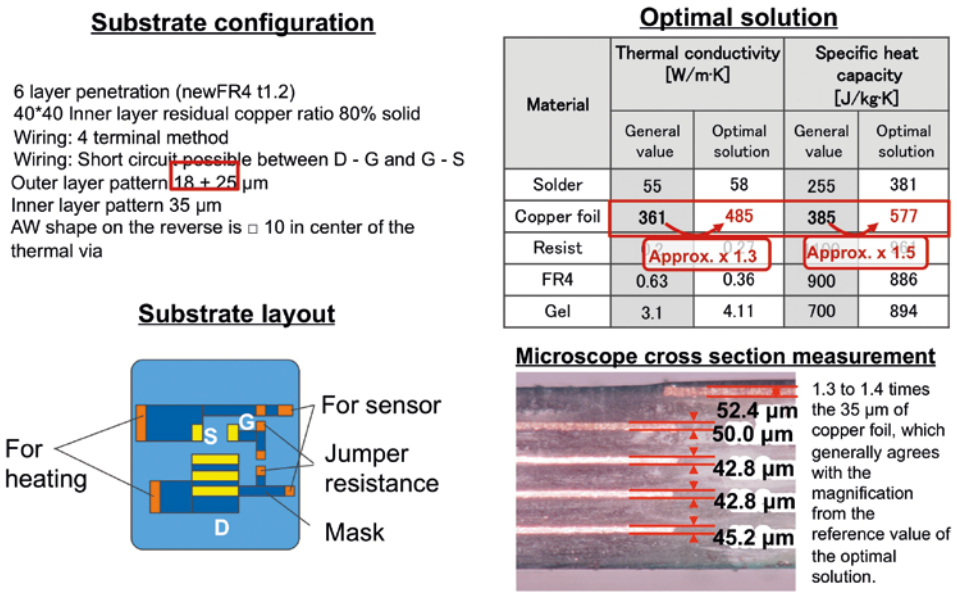

Figure 12. Validation of optimal solution with a structural survey of real substrate samples

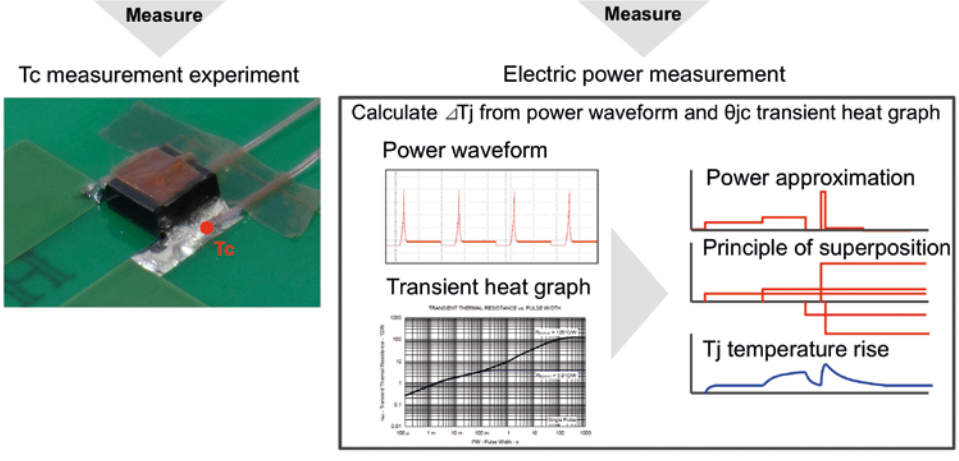

Figure 13. Issues in finding Tj by thermal design for a circuit

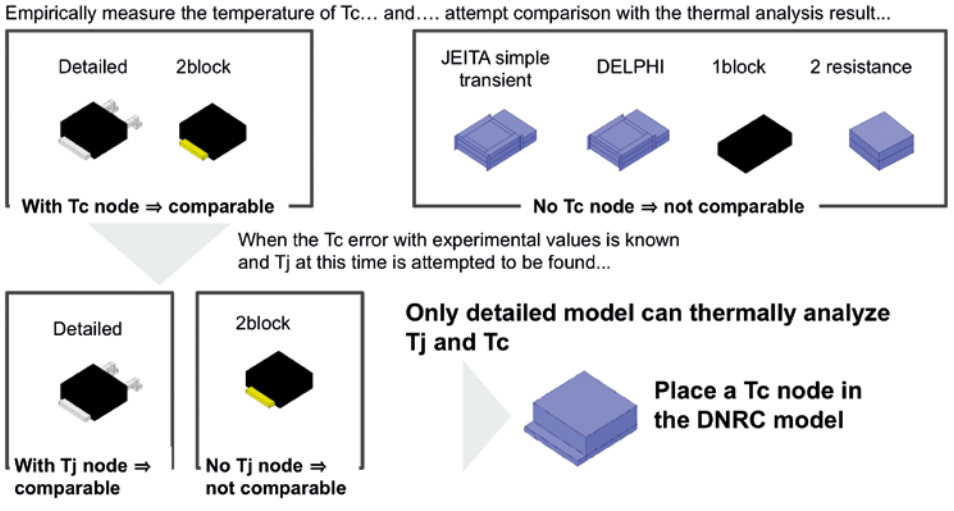

Figure 14. Innovation 1 - Temperature of the part for coupling with CFD analysis is known

	DNRC model	DSRC model
Design phase of model use	Verification of prototype	Concept only
Product status	Actual product assembly	Evaluation environment defined b standards
Transient thermal resistance	Measurement ⟷ This is important	Manufacturer's guaranteed datasheet value

Table 1. Comparison of DNRC models and DSRC models for thermal analysis

(see Figure 17) and compared it with the T3Ster and FloTHERM approach outlined above. Error rates were very low and good agreement in the DSRC models to the manufacturer's datasheet values were observed. Hence, DSRC model creation based on manufacturer's datasheet is possible and improved accuracy can be achieved by these in-built substrate models.

Summary

We have shown that by using a coupling of Mentor's T3Ster transient thermal testing hardware and FloTHERM thermal analysis CFD software it is possible to measure thermal resistance of PCBs accurately and to create thermal simulation models of semiconductor component by using DNRC models for unsteady state analysis. The DNRC modeling approach outlined here allows for the calculation of Tj values with a high level of accuracy in the CFD analysis technology. Comparative verification with measurements by Tc nodes are shown to be possible such that the approach is useful for verification experiments. Furthermore, DSRC based on datasheet value can be used for thermal design according to the manufacturer's datasheet values.

References

[1] Rencz, M. & Székely, V.,"Measuring partial thermal resistances in a heatflow path". IEEE TRANSACTIONS ON COMPONENTS AND PACKAGING. 25 (4): 547–553, 2002.

[2] Bornoff, R. & Vass-Varnai, A., "A detailed IC package numerical model calibration methodology", 29th SEMI-THERM Conference, pp.65-70, 2013.

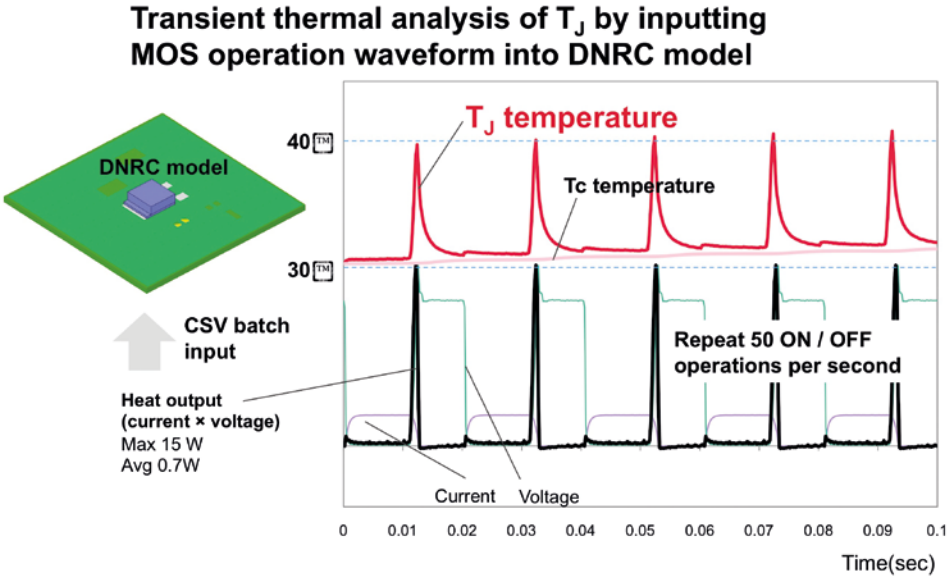

Figure 15. Innovation 2 - Preliminary verification of Tj value possible by FloTHERM thermal analysis

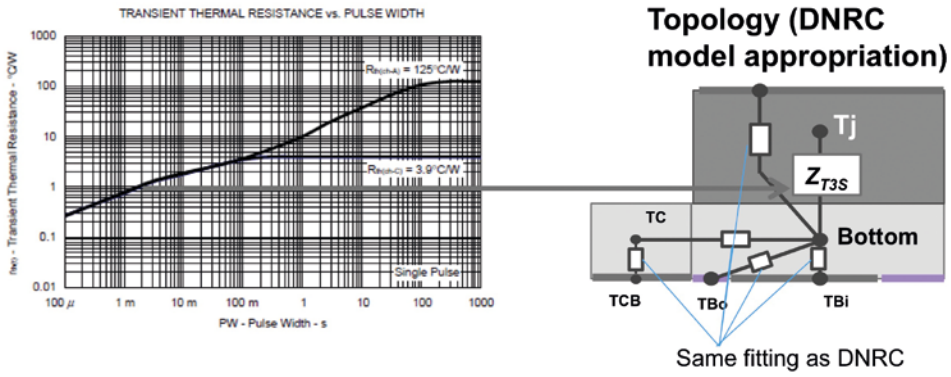

Figure 16. Innovation 3 - Possible creation of models from manufacturer datasheets and calculate Tj

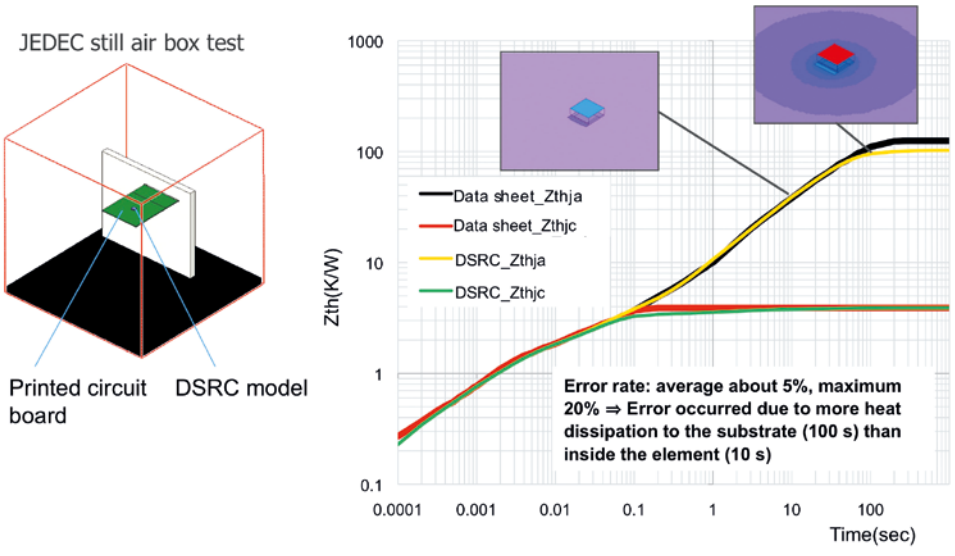

Figure 17. Experimental Validation - Transient thermal analysis in JEDEC environment according to the measurement environment of a component's datasheet value

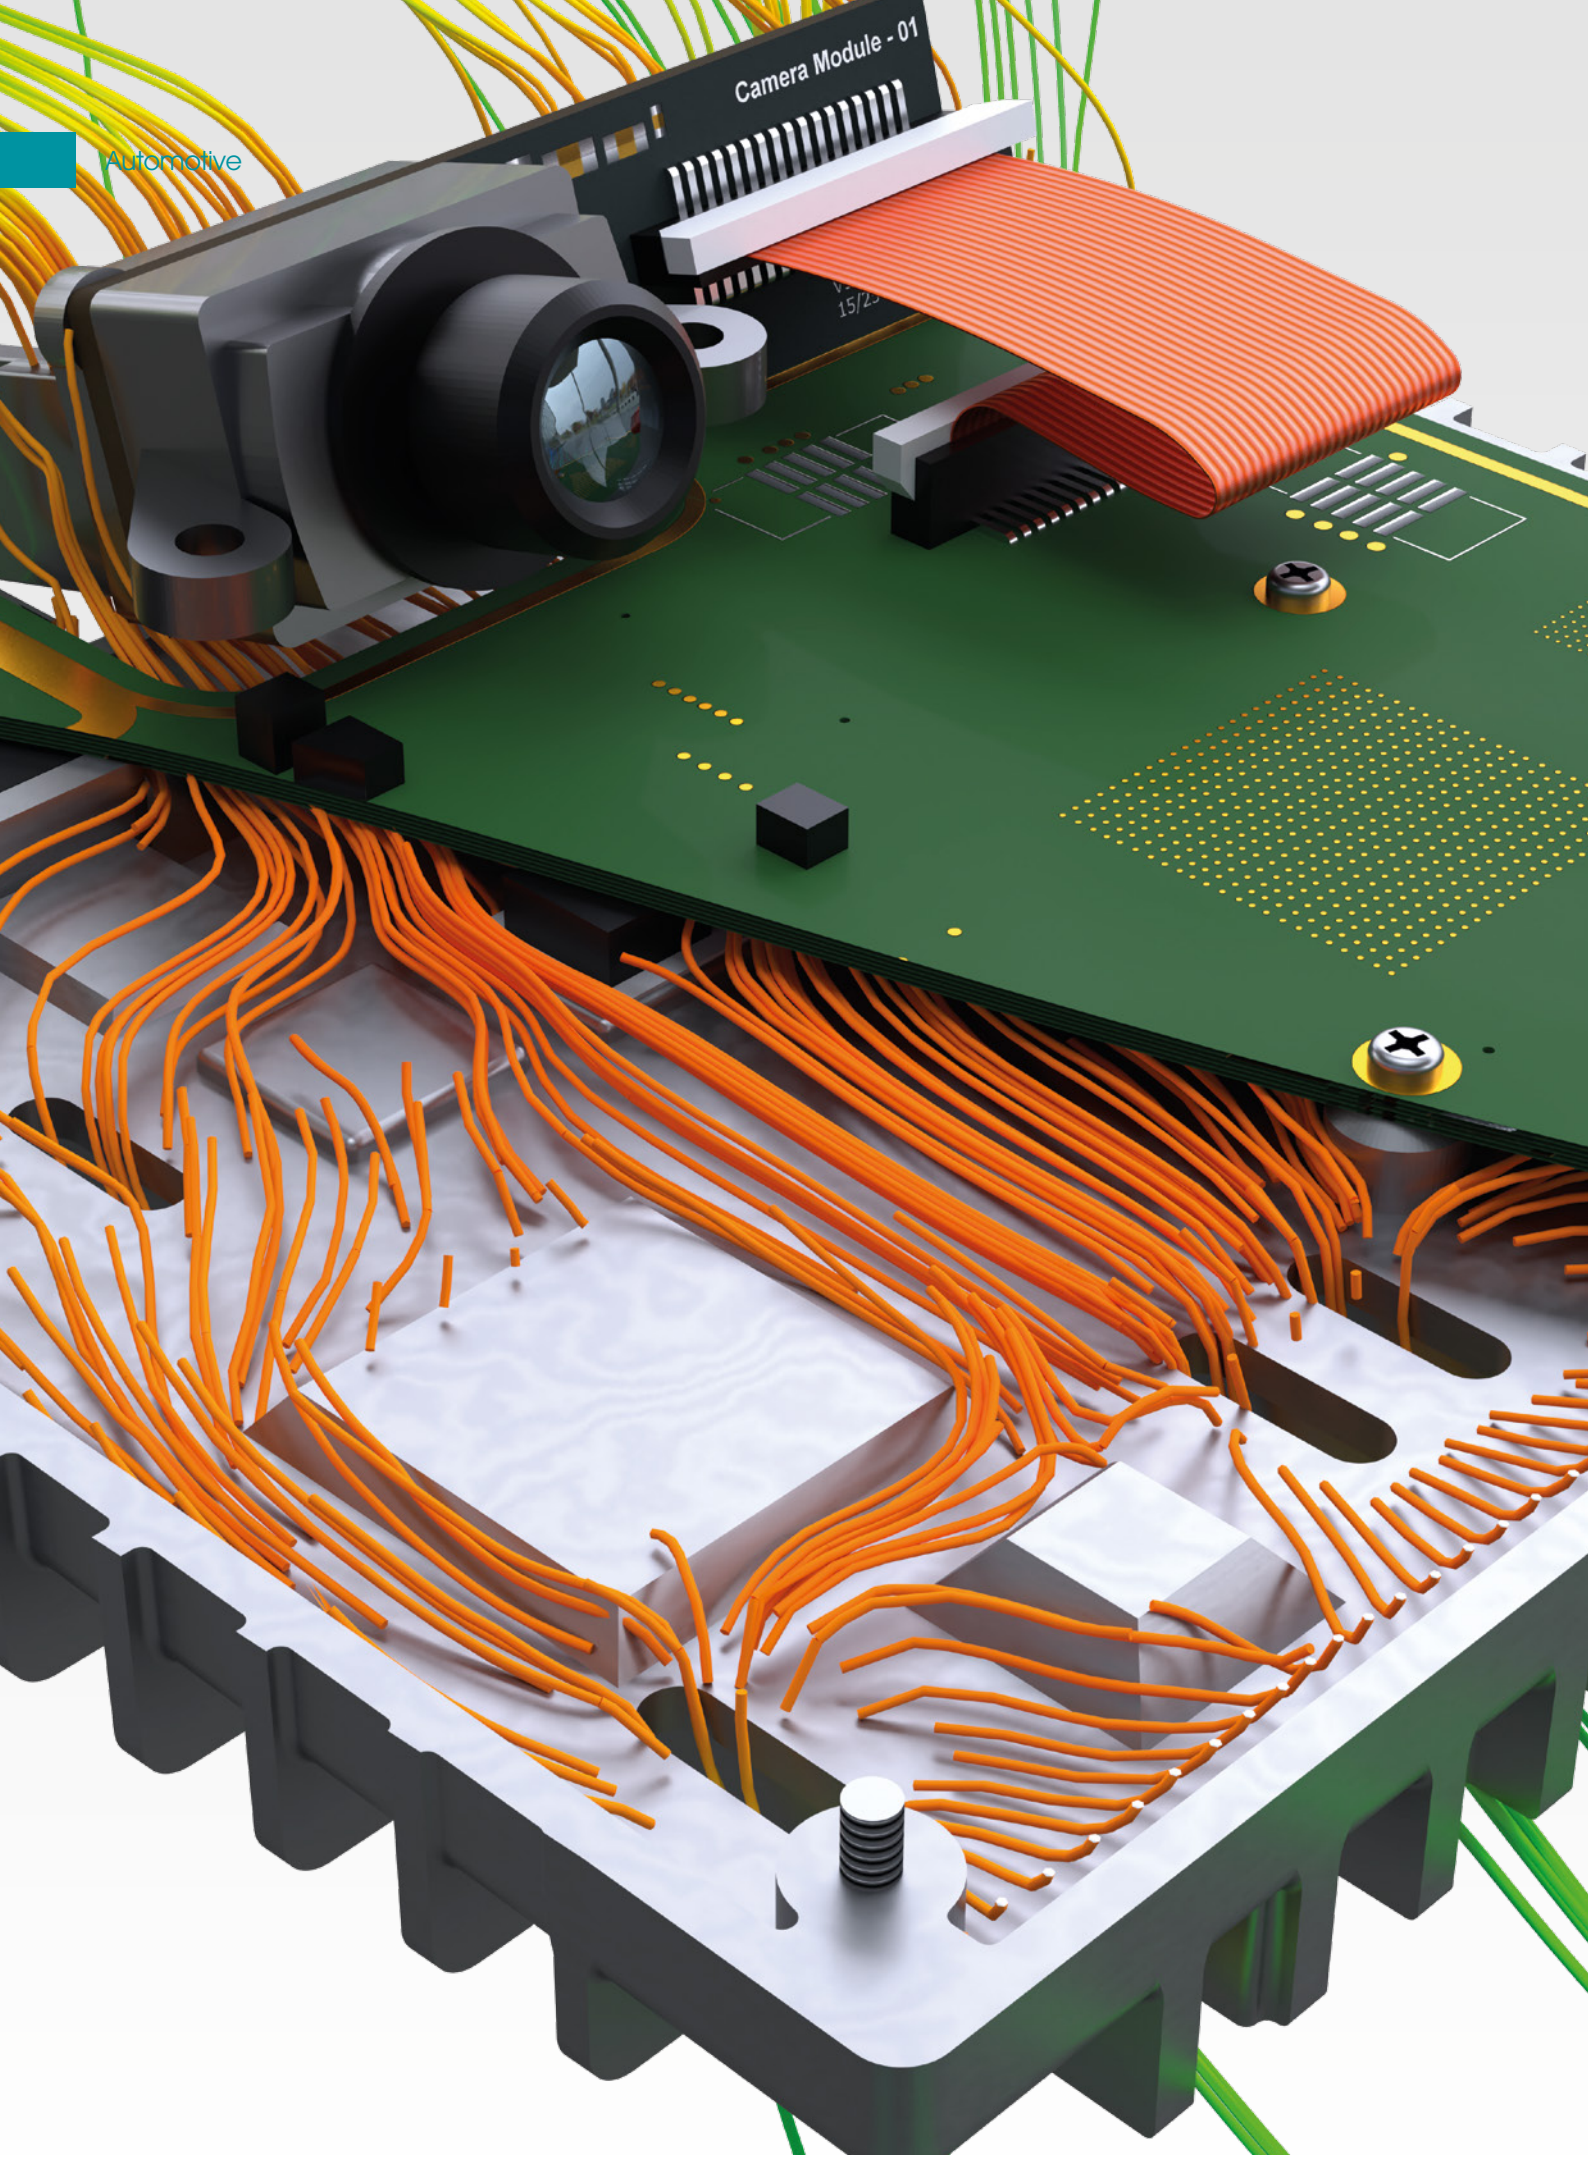
Camera Module - 01

Thermal Analysis of an ADAS Camera in FloTHERM XT

By Vladimir Kirichenko, Student Intern, Mentor, a Siemens Business

As the world continues to evolve, autonomous vehicles suddenly have stopped being seen as something from the future. In fact, both world leading companies and new start-ups are getting involved in the automotive world with the aim of not only providing simple driving assist features, but enabling full control of the vehicle to drive its passengers to their destination of choice. While humans mainly rely on their eyes to navigate, autonomous vehicles have to depend on cameras and other sensors to analyze the outside world to familiarize themselves with their surroundings.

Inaccurate data from these devices could have catastrophic consequences and cost lives of passengers and pedestrians. One of the factors that affect image quality is the thermal performance of the optical cameras, resulting in a challenging task of keeping the temperatures at sensible levels.

Challenges

As can be seen from Figure 1, the temperature of the image sensor chip plays a large role in image quality, as higher temperatures dramatically raise noise levels in the Complementary Metal–Oxide–Semiconductor (CMOS) circuitry. Consequently machine vision algorithms may struggle to distinguish objects from the background. To make autonomous vehicles safer and more trustworthy, the camera complexity has increased, and may also incorporate radar, lidar or ultrasonic sensors. Combining such power-intensive devices in one package leads to problems with heat dissipation, therefore serious consideration should be given to virtual thermal analysis during iteration of the design of the cameras used for autonomous driving to develop a final product that will give satisfactory and reliable operation.

Modeling

To analyze thermal dissipation of the automotive cameras in the current market, Mentor kindly provided to me one of the top of the range ADAS cameras available on the market today to disassemble and examine. Based on the research and the physical device, I created a model, using Siemens NX 10.0, as that was the CAD software I had the most experience with, allowing me to spring into action as soon as possible. It was also a great opportunity to do my primary job role, which was to test FloTHERM XT 3.2 using different import formats to ensure every feature worked correctly. An exploded view of the model I built is shown in Figure 2.

The camera is housed in a plastic shell, apparently for purely aesthetic reasons with cut-outs to provide airflow. The casing of the camera was modeled in a way that the bottom of the casing acts like a heatsink, dissipating most of the heat from major image-processing components in a convective manner. To encourage airflow inside the metal casing, multiple cut-outs have been made to the bottom and top of the casing. This design allowed for the air to circulate inside of the casing, removing the heat from the rest of the components. Any cables and connectors that could restrict the airflow were modeled for increased accuracy of the simulation.

Simulation

Once complete, the model was easily transferred to FloTHERM XT, which meant there was only one part missing before I could perform a thermal analysis on the camera - the main PCB. It was created using FloEDA Bridge, which allowed for a

complex shape of the board and detailed manipulation of necessary parameters. However, if I had the EDA file, it could have been easily imported. Overall, the camera was modeled with eight thermally-significant components assumed to have a total heat dissipation of 6W. The main SoC was dissipating 2.5W of power at the maximum load and has been modeled in detail. All of the power values were assumed based on the research of vision computing SoCs available on the market and engineering judgement. To speed up the simulation, the remaining components were modeled as simple cuboids with an appropriate material attached. For the same reason, compact modeling level of the PCB was chosen instead of detailed or explicit.

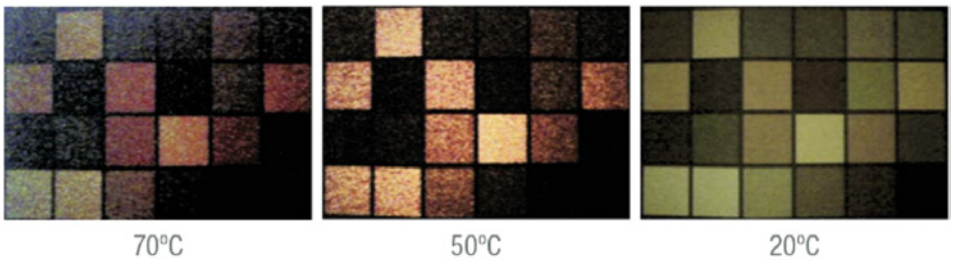

Figure 1. High dynamic range camera image sensor noise under low light for different ambient temperatures (Ref. 1)

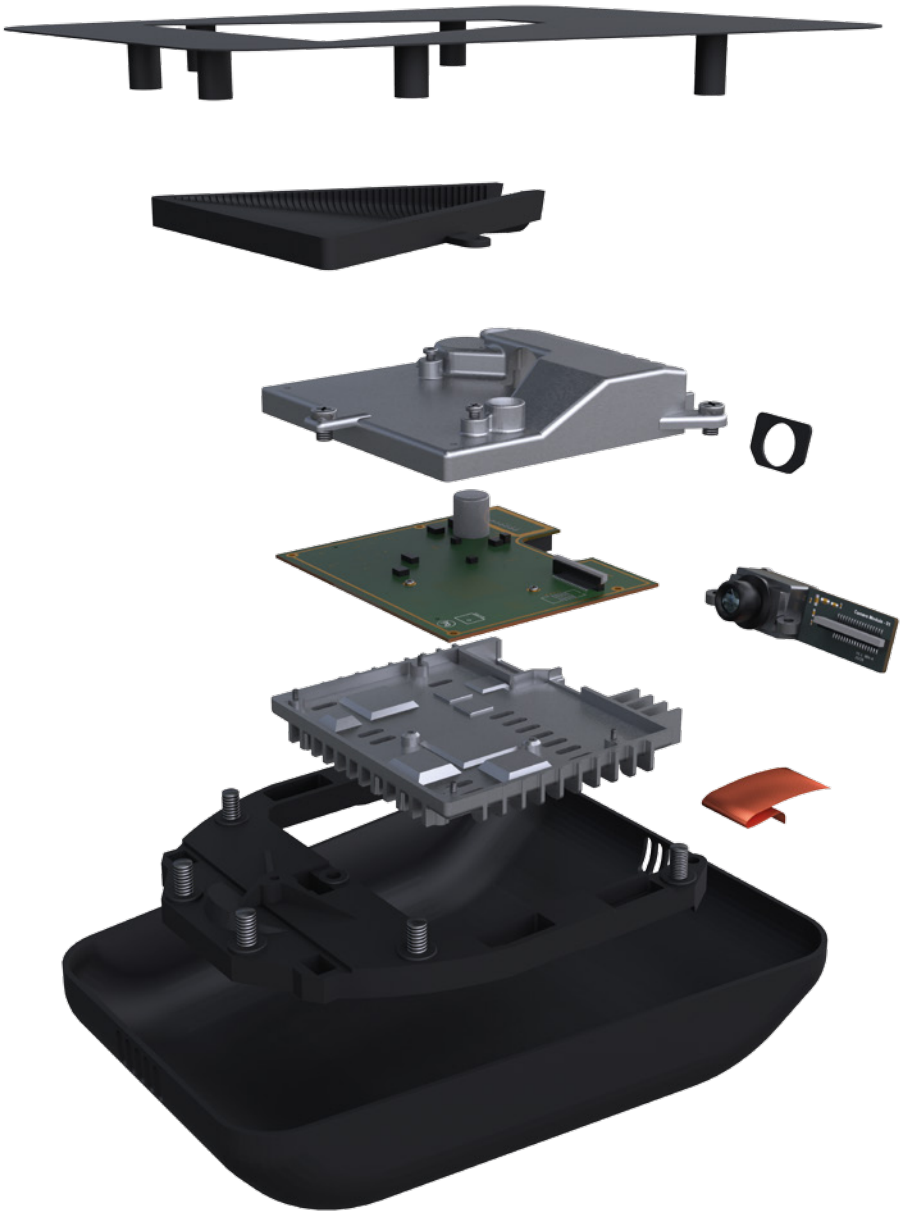

Figure 2. Exploded view of ADAS camera model

Cases

In order to perform a thorough thermal analysis, three different simulations were run. One of the places where autonomous driving technology becomes very useful is the highway, as most of the time the trip becomes monotonous and some drivers can lose concentration. Therefore, to simulate the worst case conditions and analyze a case where, if the limit of 85°C was exceeded when the vehicle is in fully autonomous mode and the camera is functioning at full power, a steady state simulation was created. The vehicle was assumed to be driven at 50 mph on a warm sunny day, consequently, 1000 W/m2 of solar radiation was applied at a right angle to the windscreen. To make the simulation more realistic, 20% of the radiation has been absorbed by the windscreen by using a thermal planar source, 60% was applied as a simulated solar radiation and the remaining 20% of the solar radiation was assumed to be reflected and therefore ignored. Further to this, the heat from the camera has been removed solely by natural convection without the presence of forced airflow inside of the cabin. The absolute maximum temperature on the main SoC reached 82°C, while the air inside of the vehicle was simulated at 22°C and the outside temperature levelled at 35°C. From these results, it becomes apparent that although these are absolute extreme conditions and can be considered as a torture test, the camera components are close to their limits.

Problems

Having tested the thermal limits of camera components, two transient simulations were performed to simulate a more realistic scenario. The main problem with the current model was the excess level of detail. Therefore, to improve simulation times, I have made several adjustments. The main SoC was simplified to a simple cuboid with the same thermal properties and a heatsink for the camera module was replaced by a Heatsink Smart Part available in FloTHERM XT 3.2. This led to a dramatic decrease in the number of cells, allowing for significantly reduced simulation times. Even though the mesh number has decreased significantly, I

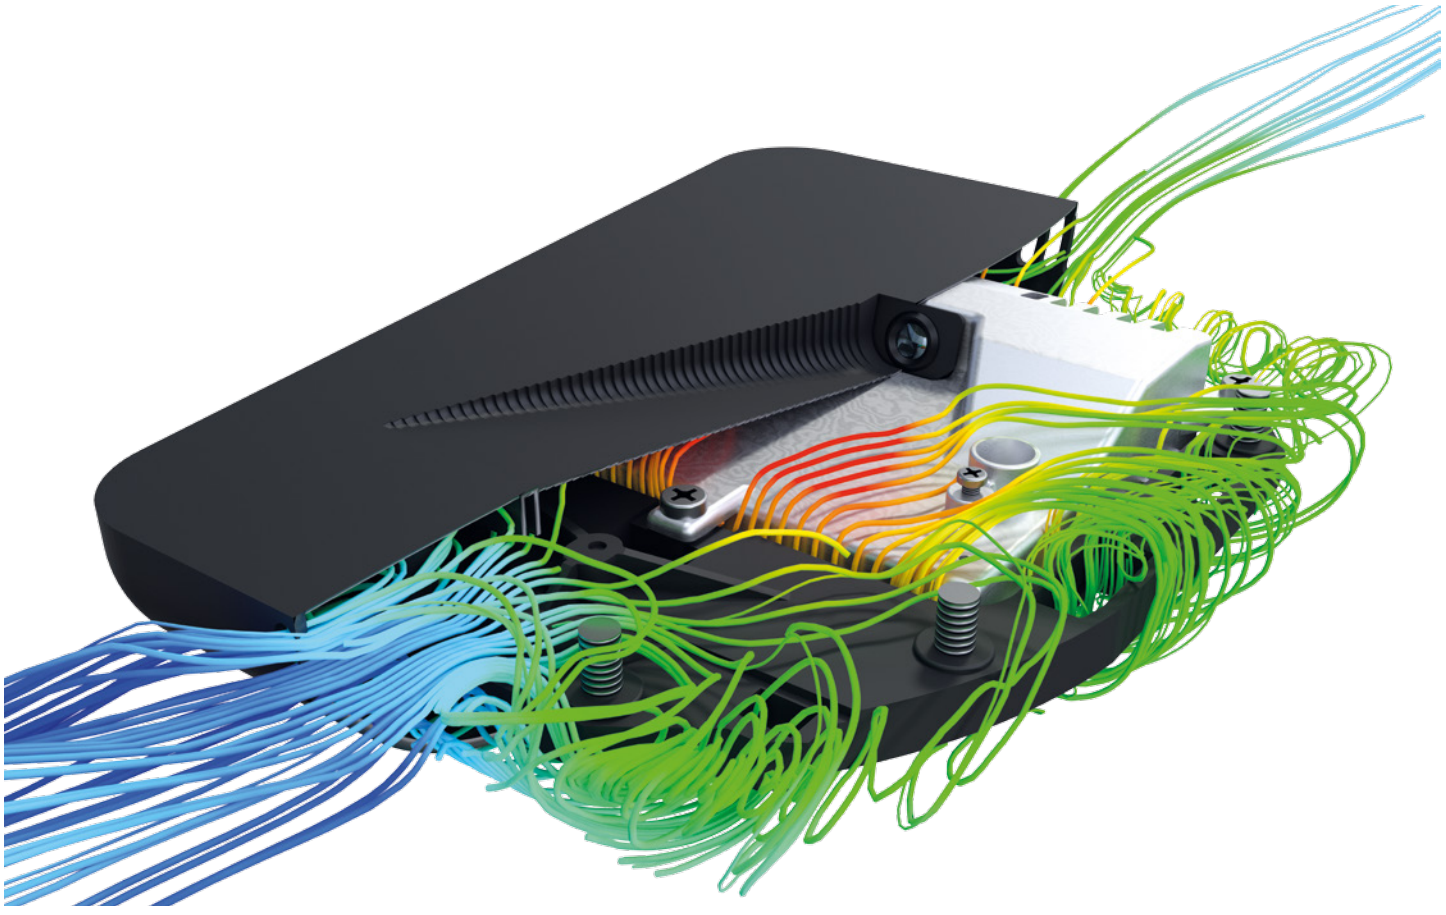

Figure 3. Predicted airflow inside the camera housing

was still unsatisfied with the amount of time it was going to take to perform the thermal analysis. To further optimize the simulation, manually-variable time steps were used.

To further limit the time needed for the simulations, the overall calculation domain was restricted to a box 180x190x135 mm, which worked perfectly for the simulations in still air. However, after an introduction of wind the problems started to appear. As the moving air created turbulence, this solution domain size was not large enough to resolve the flow. Therefore I increased the domain size by approximately 2.5 times.

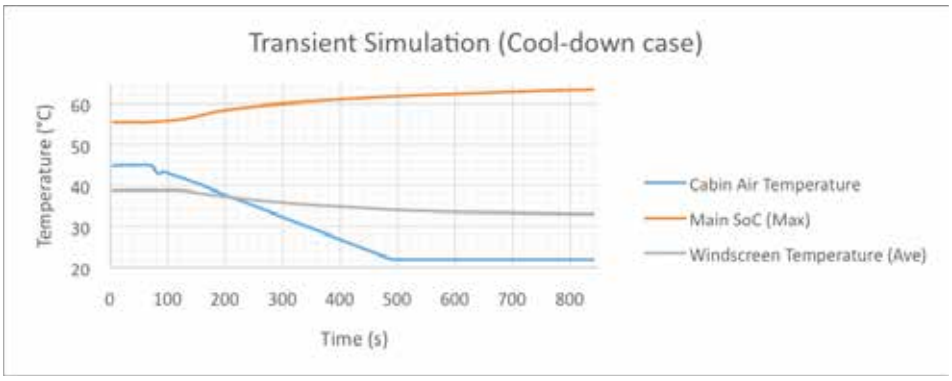

Figure 4. Key temperatures during transient cool-down simulation

Transient results

Continuing with the theme of testing the thermal limits of the camera, one of the transient simulations was performed as a cool-down case. I assumed the vehicle was parked under the sun at 30°C with a solar radiation of 750 W/m^2, 50% of which was applied as radiation, 10% being absorbed by the windscreen and another 40% was reflected, and after that it is driven at 30 mph in a city area.

To ensure the correct starting point of the transient simulation, I have performed a steady state simulation with static air first. This allowed the interior of the vehicle to warm up due to solar radiation and stabilize temperatures of the camera components. These results determined the starting point for the cool-down transient case. The temperatures inside of the cabin stabilized at 45.0°C and the temperature of the main SoC in the camera stabilized at 55.4°C.

The overall time of the transient simulation was 14 minutes, from the first to the eighth minute the AC was turned on to cool the

cabin air down to 22°C and keep it at this temperature until the end of the simulation. From the second minute the vehicle accelerated from 0 to 30 mph in 10 seconds and then continued at this speed. The camera itself was powered at 25% while the vehicle was stationary and the AC was on. As soon as the vehicle started moving, the power was increased to 75%, which was assumed to be sufficient to provide advanced driver-assistance systems. From Figure 4 it is apparent that the maximum temperature of the main SoC of the camera stabilizes at 64°C, which is an acceptable result, as in reality the solar radiation would tend to vary due to the surroundings and the power consumption of the camera unit would be altered by traffic conditions.

Turning to a warm-up case, it was assumed that the vehicle was parked outside at -10°C and there is no solar radiation. Similarly to the cool-down case, a steady state simulation was performed first and the vehicle was assumed to be driven at 30 mph. As the changes were not as dramatic as with warm-up case, the simulation had to be run for much longer for the temperature of the main SoC to stabilize.

Eventually, main SoC temperature levelled out at 19.5°C, which puts the camera assembly at a very comfortable level in terms of working temperatures. In this case the temperature of the cabin air had to be raised to a comfortable temperature for the driver and passengers.

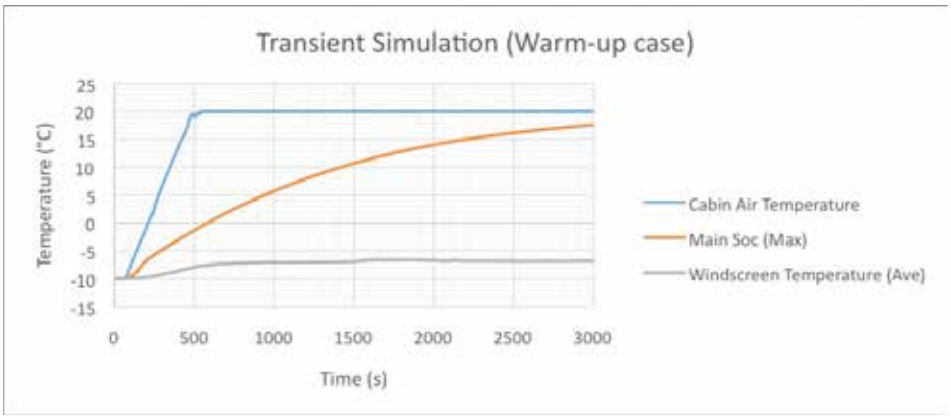

Figure 5. Key Temperatures During Transient Warm-Up Simulation

Power	"Climate Control" temperature	speed
0-1 mins: 0 %	0-1 mins: not set	0-2 mins: 0
1-2 mins: 0-25%	1-8 mins: from cabin ambient to 22 degC	2-2:10 mins: 0-30 mph
2-3 mins: 25-75 %		
3-13:50 mins: 75%	8-13:50 mins: 22 degC	2:10-13:50 mins: 30 mph

Power	"Climate Control" temperature	speed
0-2 mins: 0 %	0-1 mins: not set	0-4 mins: 0
1-2 mins: 0-25%	3-8 mins: from cabin ambient to 22 degC	4-4:10 mins: 0-30 mph
2-3 mins: 25-75 %		
3-50 mins: 75%	8-50 mins: 22 degC	4:10-50 mins: 30 mph

Conclusion

With a rapid growth of the automotive industry, there will be increasingly higher demand for cameras and radars that are required to work for longer periods of time without failures in varying weather conditions. As this article has shown, thermal analysis is extremely important to ensure that the design of the Advanced Driver Assistance Systems (ADAS) camera is effective at removing excessive heat from the internal components.

To improve the design of the device, it will be worth considering several materials before production, as well as, simulating the whole camera assembly with the casing and inside of the vehicle's interior. This would help to identify critical locations for the cut outs in the casing to optimize the airflow. It is important to mention, that while regular processors in a PC or a laptop can be thermal-throttled to reduce the temperatures, this approach is highly undesirable as this would slowdown the processor frequency and reduce the performance, which consequently reduces the speed of image-processing and may affect how efficiently and safely an autonomous vehicle reacts to the road conditions.

Reference:

[1] "Comparing Ethernet and SerDes in ADAS Applications", Dave Lewis, Texas Instruments. http://artsdocbox.com/Television/66316313-Comparing-ethernet-and-serdes-in-adas-applications.html (downloaded 07 March 2018)

Thermal Design Leading the Charge

By Ferdinand Sluijs, Technology Manager, NXP Semiconductors – Smart Power

Quite a lot has been written about the benefits of moving the thermal design activity higher up the product development workflow to cut down the amount of rework needed later on, eliminate physical prototyping for thermal reasons, etc. That said, thermal considerations are normally indicated as advice, or constraints on the main electrical and mechanical design flows.

In the case of the power adapters that NXP Semiconductors' Smart Power Division are developing thermal considerations have really come to the fore, driven by the trends we see in mobile power adaptors, which are getting smaller and have a higher output power. Thermal limits are now constraining all aspects of the design.

A recent project involved fitting a 25W charger into the smaller casing previously used for an 18W charger, some 26% smaller, while still meeting the thermal limits which constrain the casing temperature to a maximum of 50°C averaged over an area not exceeding 2cm x 2cm, against an ambient temperature of 25°C. Increasing the size of the charger, while thermally desirable, would make the charger inconvenient to use. Airflow through the charger, while again thermally desirable was ruled out for safety concerns due to the mains voltage inside.

The first question is "what is possible?" This can be answered by finding out how much power dissipated within the device gives this maximum case temperature condition. By building a simple block representation of the adapter consisting of just two blocks for the body and two for the pins, with a uniform internal power distribution, the power dissipation can be increased until the temperature limit is reached. The adapter casing, and the thermal model of this the ideal case with no hot spots is shown in Figure 1.

This simulation showed that in theory, the adapter could dissipate maximal 2.7W. Given that the output power needs to be 25W, the minimum efficiency, η needs to be:

$$\eta = \frac{25W}{25W + 2.7W} = 90.3\%$$

Providing a performance constraint on the electronic design.

The next step is to model the initial design of the new adapter to see the surface temperature distribution and where the temperature would exceed the limit. At 2.91W the initial design will certainly exceed the maximum temperature, but an insight into the temperature distribution in this design will help identify hot spots and so guide design improvements.

To build this model components that have a significant power dissipation were included in the model, along with electrolytic capacitors as these are both large, influencing the overall thermal behavior and also temperature sensitive. One key aspect of the electrical design is the need to isolate the high voltage mains supply from the low voltage output stage. This is achieved by leaving a large distance on the main PCB between the mains connected high voltage primary side, and the output connected low voltage secondary side. A small second PCB was added as not all secondary side components fit on the main PCB.

Each of these PCBs were expected to have four layers – two 35µm outer signal layers, and two 70µm internal power planes (mainly used for internal low voltage supply and ground). These boards were modeled with discrete layers, but with averaged material properties for each layer, assuming 30%

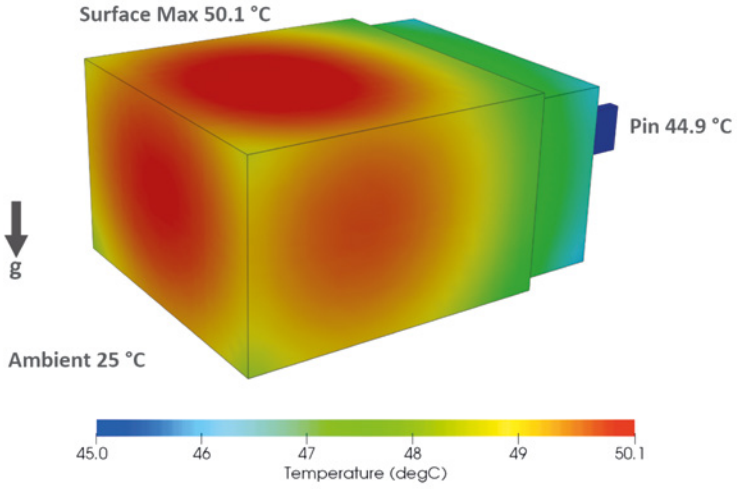

Figure 1. 18W Charger (inset) and simple block model of to determine maximum power dissipation for 25W charger

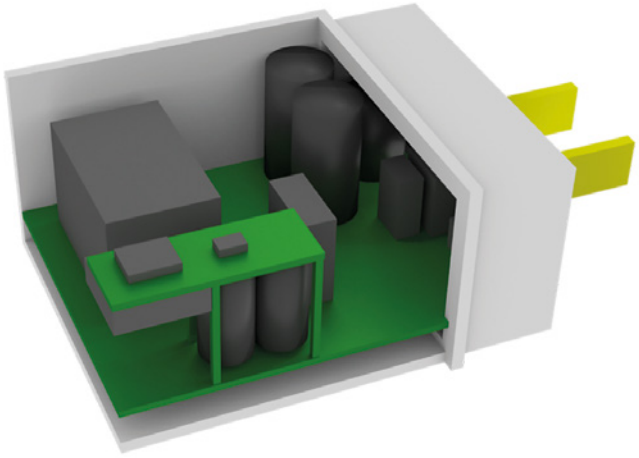

Figure 2. Initial Design

copper coverage for the signal layers and 70% for the power planes. A key decision in early design is how to model the components. Most components were modeled as discrete blocks having a uniform internal power dissipation and material properties, with values chosen based on the primary material for the part.

As expected, the case temperature exceeded 50°C in several places due to local hot spots inside the adapter, with a maximum surface temperature of 61°C.

Two key hot spots inside the adapter were the primary MOSFET on the main board, and the SyncRec MOSFET on the small second board, as shown in figure 4. Next to these two components, also the transformer has a high dissipation and gets warm, but has better cooling because of its larger size.

Rectifying these issues to achieve an acceptable thermal design within the deadline our customer required was very much a team effort, requiring close collaboration between the application engineers working on the sizing of the components, and providing power estimates; the layout engineer working on the PCB layout; and myself as the thermal engineer to also suggest thermal improvement suggestions and to investigate the thermal impact of these by performing FloTHERM simulations. The process involved sitting together and listing ideas, which I then tried out in FloTHERM, and based on the results we decided which ideas to accept and which to reject. We went through that cycle several times, described in more detail below, resulting in the final PCB layout that was simulated, and later assembled.

As the power dissipation was known to be too high, optimization of the circuitry and control was started, focusing mainly on reducing the

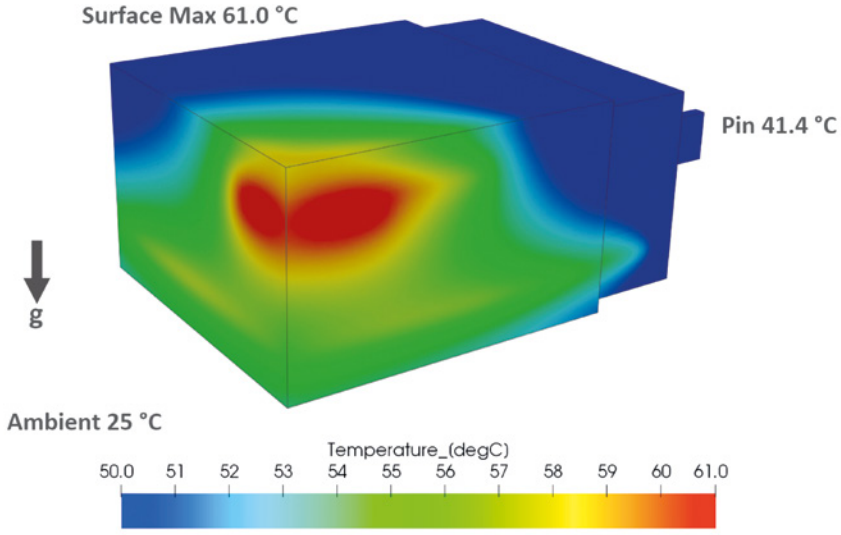

Figure 3. Initial design surface temperatures

power dissipation in the primary MOSFET, the transformer, and the Synchronous Rectifier MOSFET which was used in preference to a rectifier diode as it dissipates less power, but requires a drive signal provided by the Synchronous Rectifier IC. This work resulted in a total power dissipation of 2.2W, giving a 91.9% efficiency. As this is below the 2.7W theoretical maximum, cooling this should be possible provided hot spots on the casing can be reduced, so as well as forcing a redesign of the circuitry and control, thermal constraints also necessitated a significant redesign to alter the layout of the components to better spread the heat dissipation throughout the adapter.

Key changes were to move the primary MOSFET from the top of the main board to the bottom, and away from the transformer to separate these dissipating sources. The primary MOSFET was mounted flat onto the PCB to conduct away more of its heat. The four bridge diodes were then moved to the top of the main board. These changes meant

that one the large cylindrical capacitors on the main board also had to be moved. This was mounted on its side and raised off the main board.

The size of the small second board was increased to improve heat spreading, and the location of the connecting wires changed. Between the transformer and the secondary board a vertical plastic wall is added to help conduct heat from the secondary PCB down into the main board and into the transformer to help remove heat from the Synchronous Rectifier MOSFET. The final change was to move the USB connector to the main board, as the cable, which will be present when the adapter is charging, and so dissipating heat, will help remove heat from the adapter, as will the pins supplying mains power to the unit as these will conduct heat into the mains socket.

These changes dropped the temperature of the casing adjacent to the Synchronous Rectifier MOSFET by 10°C and with the maximum casing temperature averaged over a 2cm x 2cm area of 48.8°C, as measured by a FLIR infrared camera with a USB cable attached, thereby meeting the design requirements. This work illustrates the importance of thermal design for electronic products and the insights possible using FloTHERM. The simulation results of the final adapter design are in close agreement with the surface temperature maximum and distribution measured on the casing after the adapter had been fabricated, so no further design rework was necessary.

Had this design not met the requirements we still had the opportunity to increase the size of the charger slightly to improve the external cooling, and use more expensive electronic components internally, or increase the copper content of the PCB. All of these would have added cost to the final product and FloTHERM helped us to find what we believe is the lowest cost cooling solution for the product.

Acknowledgements:
The author would like to thank Frank van Rens for the optimization of the adapter's circuitry and control, and Thady Bruton for the thermal measurements performed on the adapter.

Reference:

[1] Ferdinand Sluijs (2017), "Thermal Modelling to Optimize Design in Mobile Charging Applications", Proc. of 23rd THERMINIC Workshop, Amsterdam NL, September 2017, pp. 1-5. ISBN 978-1-5386-1928-8

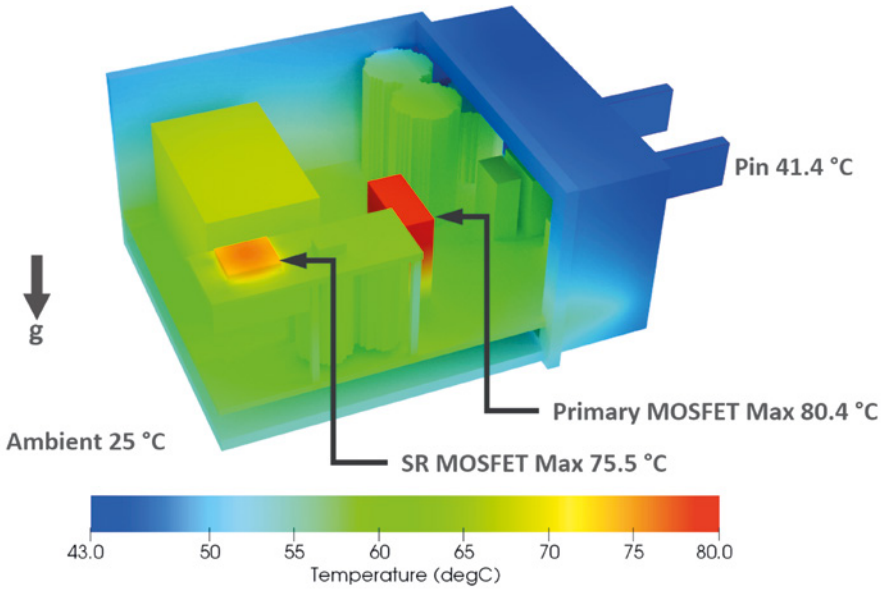

Figure 4. Initial design showing internal hot spots

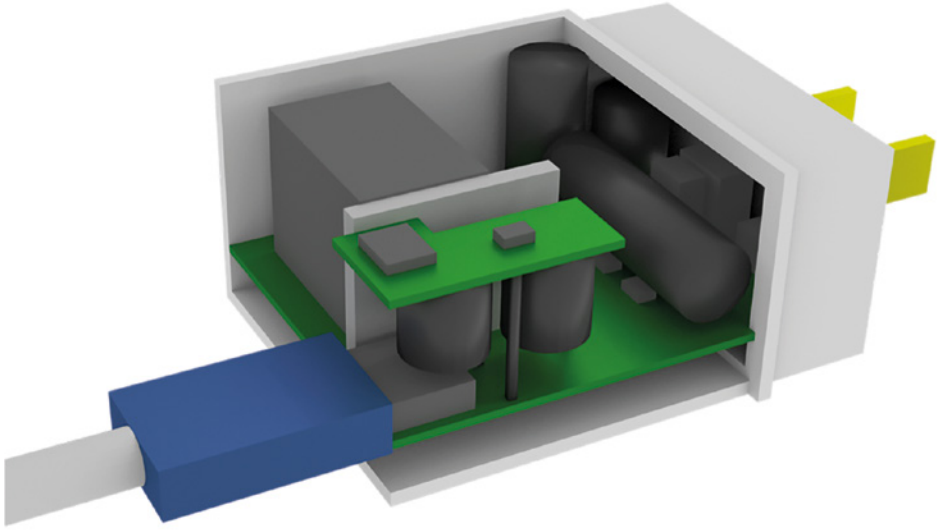

Figure 5. Final design

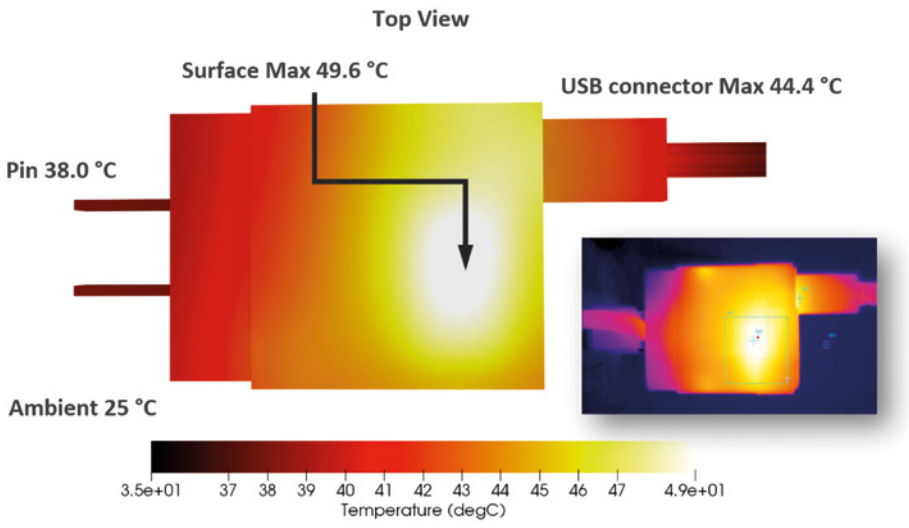

Figure 6. Final design showing surface temperature comparison with measurement

Understanding Die Attach **Thermal Performance**

By John Parry, Senior Industry Manager, Mentor, a Siemens Business
and Sujay Singh, Principal Reliability Engineer, ON Semiconductor

Power electronics components improve the energy efficiency of electric machines and motors across all industries and applications. Increasingly these power electronic components are being more densely packed together, positioned close to or on the motor itself, and so are affected by harshness and vibration in the application. Heat dissipation from these components has to be removed efficiently to prevent the premature failure of the device, or in the most serious cases thermal runaway as leakage currents increase with increasing temperature, further adding to the heat dissipation until the semiconductor die actually melts.

As part of ON Semiconductor's commitment to delivering the highest quality of product to its customers it is essential to understand how the new technologies such as wide band gap semiconductors can be introduced and produced with the lowest possible defects to achieve improved device performance and reliability in service, when the part experiences temperature swings.

These temperature swings can cause cracking of the solder die-attach or delamination between the die attach and the die, or the die-attach and the substrate. This in turn leads to an increase in the die temperature, as the heat flow path from the die to the ambient is disrupted, further elevating the magnitude of the temperature swing which further speeds up the rate of damage. Device lifetime is dependent on many factors, but based on the operating temperature of the die, ON Semiconductor know that a 10°C increase in temperature roughly corresponds to a two to threefold decrease in lifetime for the same duty cycle.

Voids that can occur in the solder die attach have the combined effect of making the die attach less thermally conductivity and act as sites where cracks can form, so solder die-attach voids are one of the major reliability concerns in power electronics packages. The impact of die-attach voids depends on the void type, pattern, and location of the voids within the solder, making it critical to really understand the influence these have on the thermal impedance of the die attach.

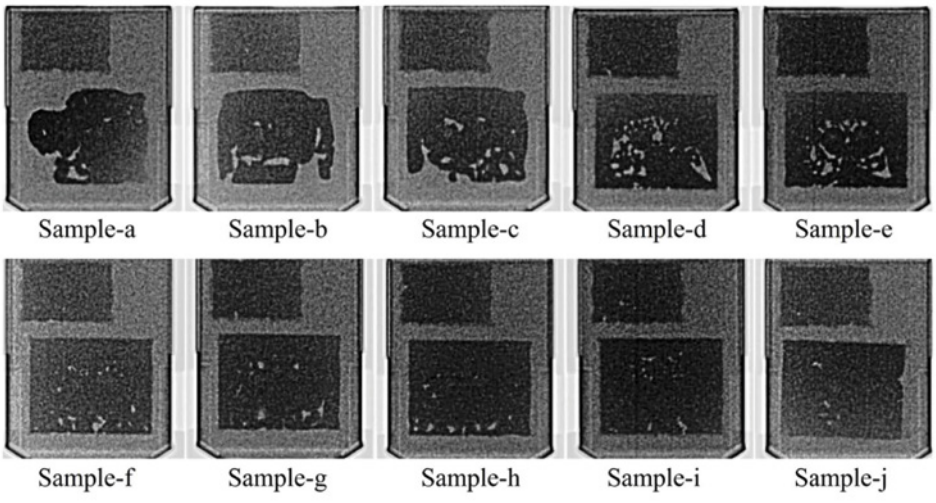

Figure 1. Samples in decreasing order of die-attach voiding under IBGT die

To investigate this, ON Semiconductor selected an insulated gate bipolar transistor (IGBT) co-packaged with a diode in a TO-247 package. Ten samples were selected having differing amounts of voiding in the die attach, detected using X-ray imaging, as shown in Figure 1.

These were categorized both in terms of the total amount of voiding, as a percentage of the area of the die, and in terms of the size of the largest void, typically caused by coalescence of voids around the periphery of the die.

Figure 2. MicReD Industrial 1500A Power Tester

Measuring the effect of the voids on the thermal impedance of the die-attach require a highly sensitive measurement system, and a means of identifying the contribution of the die-attach to the overall thermal resistance measured. That meant following the JEDEC JESD51-14 standard to measure the junction-to-case thermal resistance, Θ_{JC}, using the transient dual interface measurement (TDIM) method. The thermal impedance Zth measurements were performed using Mentor's MicReD Power Tester 1500A showing in Figure 2, which provides the necessary measurement fidelity, combined with inbuilt structure function analysis of the Zth curve to identify the partial thermal impedance due to the die attach.

The structure function is a 1D representation of the heatflow path and provides information regarding the spatial distribution of thermal properties in a system, presenting this information as a graph of the cumulative thermal resistance versus the cumulative thermal capacitance the heat experiences as it passes from the source on the die surface out to the cold plate on which the part is mounted in the Power Tester 1500A. The method, which fully conforms to the JEDEC JESD51-14 standard is non-invasive as it uses the electrical test method described in JEDEC JESD51-1 to both electrically heat the die and sense the die temperature during the measurement.

Figure 3 (a) and (b) show the electrical schematic for heating and sensing of the junction temperatures of the IGBT and diode respectively. The Zth measurements on the IGBT were done in saturation mode. Figure 3(c) shows the experimental fixture for the IGBT measurement. The fixture is clamped on the cold plate by a clamping torque of 6"/lb to ensure good thermal conduction. Figure 3(d) shows the package construction and the physical separation of the IGBT and diode which are both mounted on the leadframe. While the Power Tester 1500A is capable of delivering up to 1500A to a single package, the heating current used for the measurements was 20A. All measurements were done at a cold-plate temperature of 25°C with data automatically captured during the measurement by the Power Tester. Analysis of the data was done using the MicReD T3Ster Master software. The electrical transient that occurs during the first few microseconds of the measurement caused by the power being switched down from the 20A heating current to a measurement current of 0.1A was corrected by using the in-built "square-root" method that replaces the initial/parasitic transient with a curve extrapolated using a square-root fit within an appropriately selected time window.

The measurement of Θ_{JC} using JESD51-14 involves performing the measurement twice, once with the sample mounted to the cold plate without grease between the sample and the cold plate, and again with grease.

The microscopic roughness of the fixture and the package resists the heat flow for the dry case, whereas the TIM/grease minimizes the surface roughness and decreases the interfacial thermal resistance. Therefore, the separation between the two curves is due to the difference in the thermal path for the two measurements. The two curves start to separate as soon as the heat flux reaches the package-fixture interface, i.e., when flux leaves the case of the package. The splitting point of the two curves indicates Θ_{JC}. However, the two curves separate gradually rather than on a well-defined point, as shown in Figure 4(a). Therefore, it is necessary to define the point of separation in time more precisely. The task is relatively easier if the derivative of the curves is used, as shown in Figure 4(b). Noise on the graph is accounted for by fitting a resistor-capacitor ladder to the experimental response and noting where that intersects the trend line fitted to the difference in the derivatives at the start and end of the transient.

Using this method the separation point corresponding to the junction-to-case thermal resistance is 0.18K/W for Sample-a. T3Ster-Master software, used to post-process the temperature versus time response of the sample provides an alternative way to assess the junction-to-case thermal resistance using a structure function. The structure function is a graphical representation of the total, cumulative thermal capacitance as a function of the total, cumulative thermal resistance measured from the heat source or driving point (junction), that the heat encounters as is passes through the package to the ambient, in this case the cold plate. The structure function is computed by the mathematical transformation of the measured transient response that involves several steps not described here.

For a 1D heatflow path, the thermal structure function provides information about the thermal properties of the individual layers inside a packaged device. A change in the heat flow due to defects in the packaged device would change the

> "Thermal Structure Function analysis offers an effective and a non-invasive method to identify physical characteristics of the individual layers of a packaged device. This technique complements other non-destructive failure analysis techniques such as X-ray and C-SAM that visibly identify defects, however, these techniques do not provide any information regarding the impact of these defects on the thermal behavior of the part.
>
> **Sujay Singh, Principal Reliability Engineer, ON Semiconductor**

structure function. In other words, a change or shift in cumulative thermal resistance or cumulative thermal capacitance would correspond to a change in the thermal structure or material properties within the heat flow path. For example, if the thermal interface between the case and the fixture is changed from a dry contact to grease, then the structure functions for these two measurements would separate at the corresponding point with the thermal resistance at the separation point indicating the value of the junction-to-case thermal resistance, providing an alternative approach to its measurement.

One challenge, with either approach, is clearly determining at what point the curves have separated. In the case of the structure function approach, the difference in the cumulative structure functions also has some noise, so again there is no unique point at which the curves suddenly deviate from one another. While the method is standardized and documented in the JEDEC JESD51-14 standard, it does not contain clear guidelines for choosing the value of the difference in thermal capacitances at which the curves are judged to have separated. The difference is taken to be 5% of the capacitance value at which the curves appear to separate. Using a lower value leads to larger variations in the junction-to-case thermal resistance.

According to JESD51-14, the junction-to-case thermal resistance obtained through the structure function difference method for low thermal resistance packages is often impeded by numerical effects; therefore, the derivative delta method seems to be more reliable in this case as the TO-247 package has a low thermal resistance.

Having selected the most reliable way to measure the junction-to-case resistance, ON Semiconductor were able to turn their attention to examining the impact of voids in the die attach layer on the junction-to-case resistance. Previous studies have found that large coalesced voids that are hot in nature affect the thermal dissipation more severely when compared with distributed or random voids. To deliver high quality and cost-effective products it is essential to understand exactly how the size, position and distribution of voids affects the package's thermal performance.

The ten packages, Sample-a to Sample-j were measured and the junction-to-case resistance measured using the temperature versus time derivative method and the results were correlated with both the

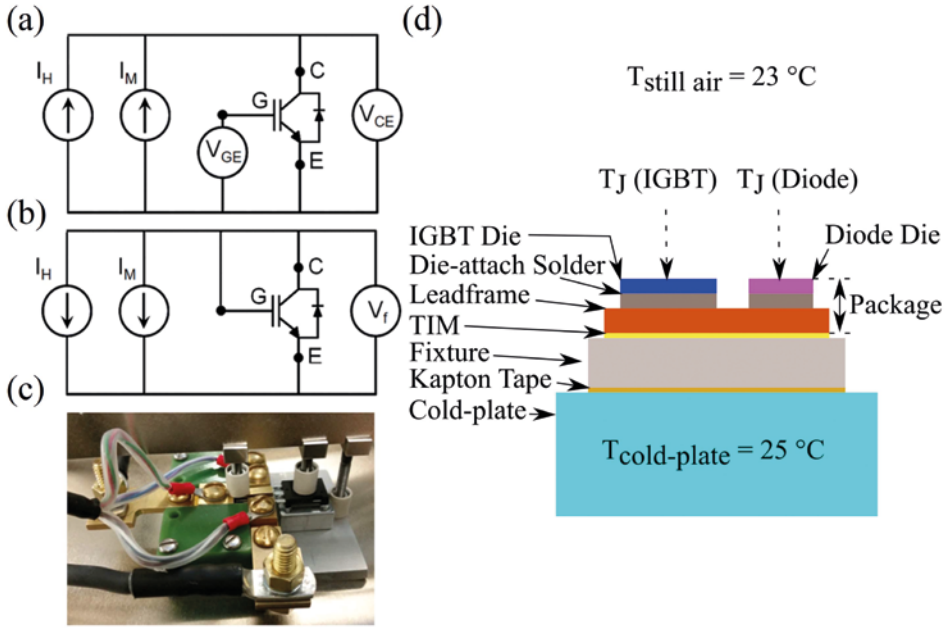

Figure 3. (a) and (b) Schematics for Θ_{JC} (IGBT) and Θ_{JC} (Diode) measurement, respectively. (c) Experimental setup for Θ_{JC} (IGBT) measurement. (d) Schematic showing the cross section of different layers and boundary conditions.

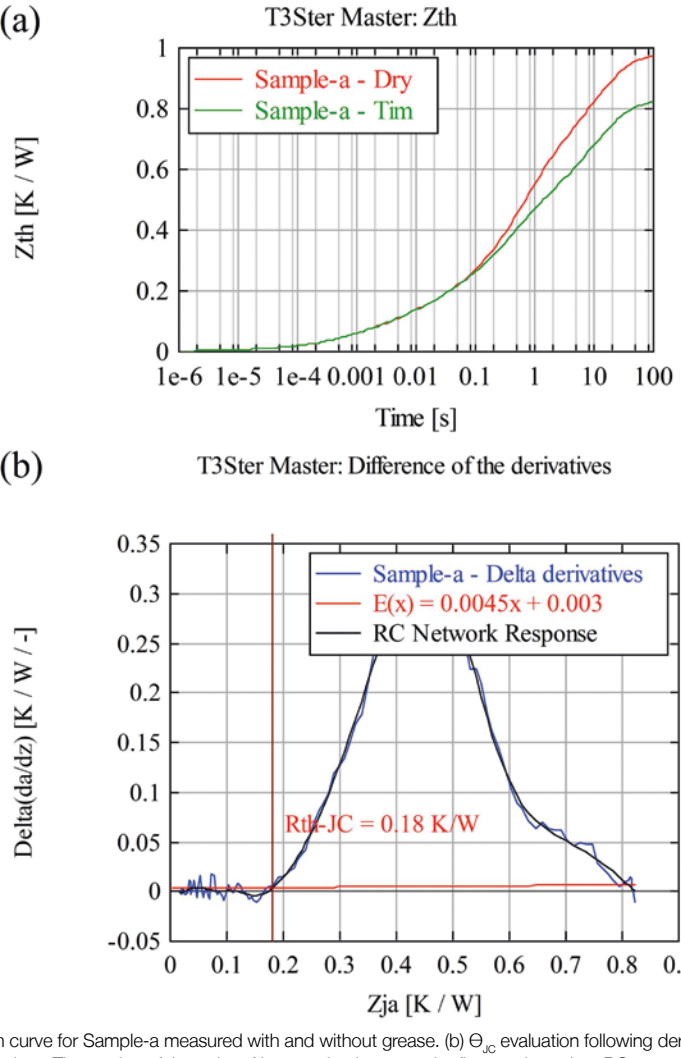

Figure 4. (a) Zth curve for Sample-a measured with and without grease. (b) Θ_{JC} evaluation following derivative delta method for Sample-a. The x value of the point of intersection between the fit curve based on RC-network response and the trend line defines Θ_{JC}

percentage area corresponding to the largest void and the total void percentage. The results of this are shown in Figure 6.

From the results it is clear that below about 10% total void area the junction-to-case resistance is not appreciably increased by the presence of voids. The conclusion from this is that the distributed voids included in the total void percentage do not significantly affect thermal dissipation severely. The largest void size, as a percentage of the total area has a more marked impact on the junction-to-case resistance.

The behavior of junction-to-case thermal resistance with increasing largest void percentage can be described by an exponential fit as shown in Figure 6(a). The reason for choosing an exponential function is to assess the nature of the acceleration behavior.

It is clear that the increase in void levels would decrease the effective area for heat conduction and that will lead to an increase in junction-to-case resistance. The change in resistance with an increase in void percentage could be dramatic for a higher void level, and the exact behavior is hard to predict as it depends on how these voids are distributed. Thermal dissipation of packaged device is severely affected by die-attach voids, and minimizing these voids is important from a reliability standpoint. The largest void percentage has greatest impact on the thermal impedance of the devices. Measurements on the diode junction-to-case resistance showed this is unaffected by voiding in the IGBT die attach, implying that the diode heat path is independent of the IGBT heat path.

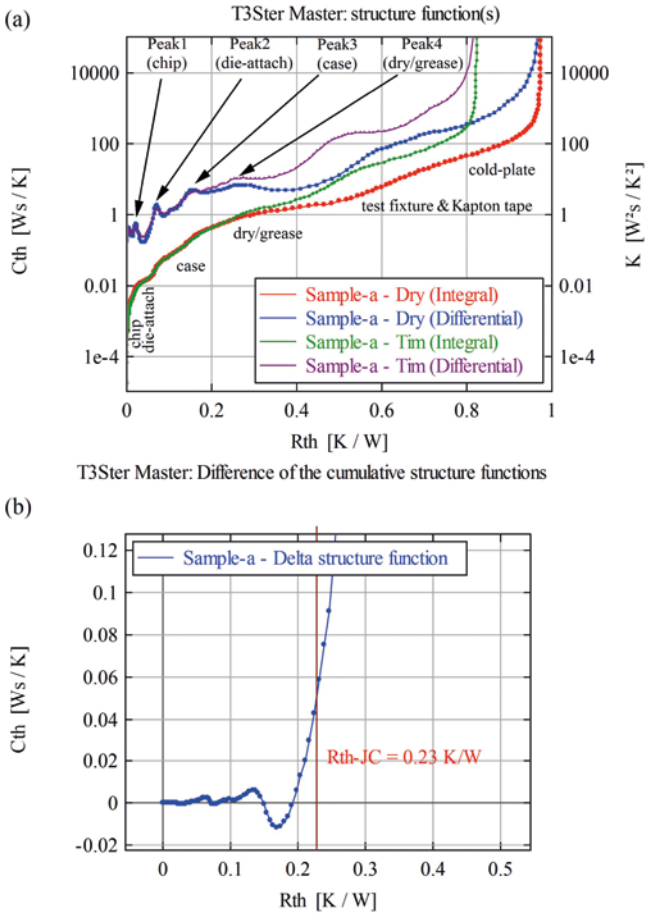

Figure 5. (a) Thermal capacitance versus thermal resistance for Sample-a showing difference in thermal paths due the change in the interface between the case of the package and

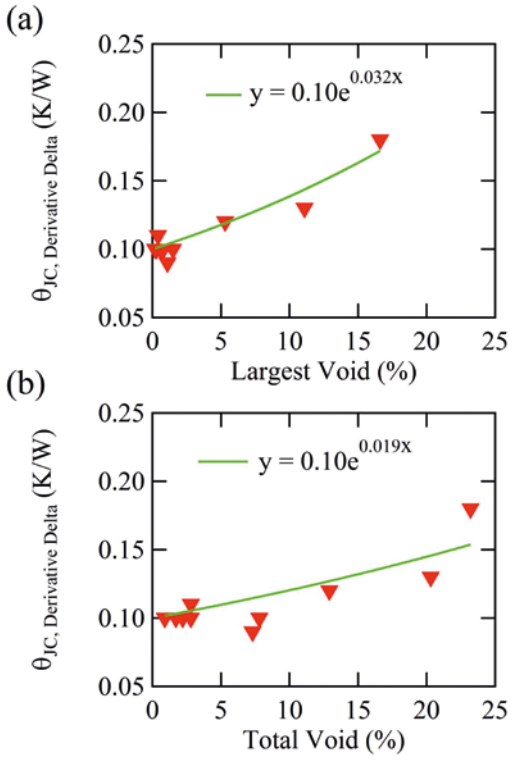

Figure 6. (a) and (b) Junction-to-case thermal resistance (for the IGBT dies) as a function of largest void as a percentage, and the total void percentage, respectively.

Reference:

[1] Sujay Singh, Jifa Hao, Daniel Hoffman, Thomas Dixon, Andrew Zedolik, Joe Fazio, and Thomas E. Kopley (2017) "Effects of Die-Attach Voids on the Thermal Impedance of Power Electronic Packages", IEEE Transaction on Components, Packaging and Manufacturing Technology, Vol. 7, No. 10, October 2017, pp. 1608-1616.

Rapid Beverage Cooling Analysis

By Jay Chien and Hank Chung, EFD Corporation

There is an increasing market for appliances that are designed to cool beverages many times faster than a conventional refrigeration approach. When the liquid is prepackaged in a can or bottle, there are cooling approaches that aren't possible when cooling items of arbitrary shapes and sizes. The same is true for beverage coolers that are designed with an integrated barrel or tank for the liquid as shown in Figure 1.

There are many cooling options to consider early in the design phase depending on the performance requirements. The design may involve thermoelectric coolers (TEC) that use the Peltier effect to extract heat from the fluid. Other options to absorb heat from the fluid could include liquid cooling loops or cold plates. Rejecting the heat to the environment might include a liquid to air heat exchanger or fan cooled heatsinks. The fluid within the barrel might also be mixed to improve the rate of liquid cooling. The thermal design tool used to consider all of the possible options must be able to model each of the approaches. FloTHERM XT was used to explore some of the cooling strategies used when designing a rapid beverage cooler.

Two designs with an integrated barrel were considered in FloTHERM XT to explore some of the alternatives for rapid beverage cooling. The two designs considered both use TECs to extract heat from the fluid.

The first design scenario used an air cooling strategy with three cooling assemblies and is shown in Figure 2. Each cooling assembly, as shown in Figure 3, consists of a fan, plate fin heatsink and two TECs. The cooling assemblies are connected together with four heat pipes. The TEC chosen for this design can pump up to 57W of heat and are colored pink for clarity in Figure 3.

The second design scenario used a liquid cold plate with an air cooled heat exchanger as shown in Figure 4. The cold plate is mounted to three TECs that pump up to 145W of heat each. The barrel in Design 2 also included internal fins and a rotating impeller, rotating at 60RPM, as shown in Figure 5.

Figure 1. Rapid beverage cooler with integrated barrel

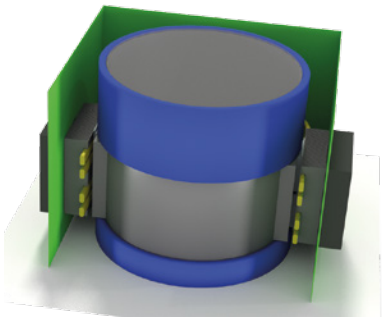

Figure 2. Design 1 cooling strategy

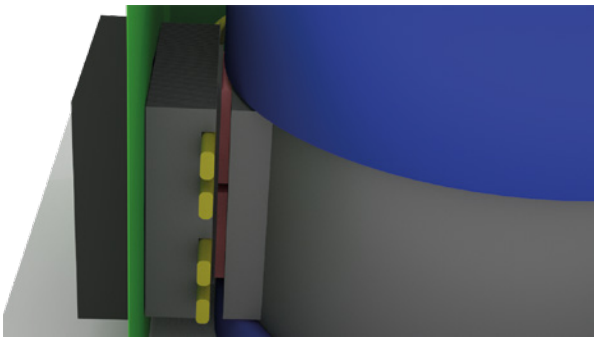

Figure 3. Design 1 cooling assembly with TEC

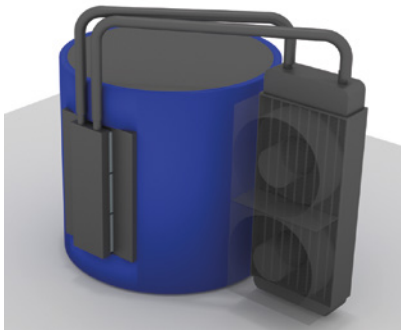

Figure 4. Design T cooling strategy

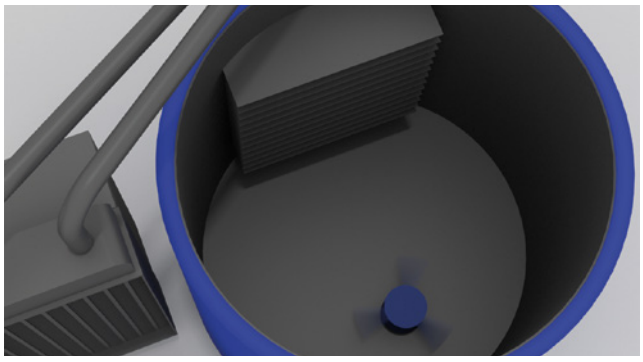

Figure 5. Design 2 internal fins with impeller

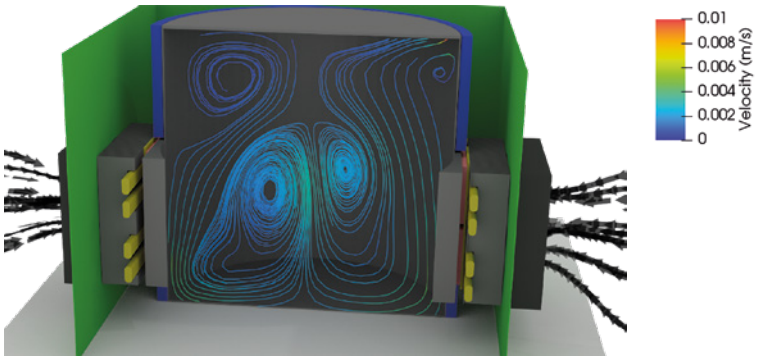

Figure 6. Design 1 flow streamlines at 30 minutes

Each design was simulated for a transient duration to assess the quality of each design in terms of rate of fluid temperature decrease. Figure 6 displays the flow streamlines for Design 1 at the end of the transient simulation. The fluid travels downwards near the cooling assemblies, along the bottom of the barrel and rises in the center. The average velocities in Design 1 are well below 0.01m/s. Design 2 included a rotating impeller and had peak velocities in the 0.4m/s range as shown in Figure 7. Also shown in Figure 8 is the cooling loop fluid temperatures with the return temperature from the heat exchanger at about 40°C.

Figure 8 shows the rate of temperature decrease at the center of the barrel as a function of time. Though Design 2 had a higher maximum possible TEC heat removal rate and impeller, the rate of temperature decrease was less than Design 1.

When designing a rapid beverage cooler there are many cooling strategies that need to be explored early in the design process to determine the most effective approaches to be considered for prototyping. FloTHERM XT is an easy to use design tool that allowed us to quickly consider design choices that could be used to develop an effective rapid beverage cooling design.

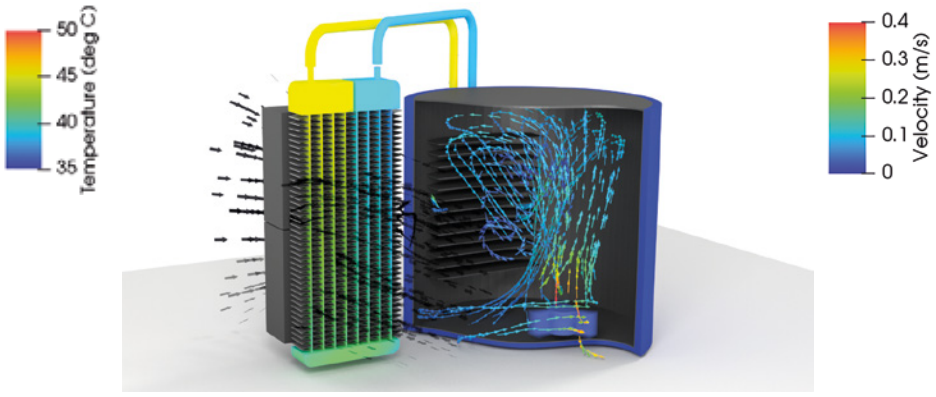

Figure 7. Design 2 Flow streamlines and cooling loop fluid temperature at 30 minutes

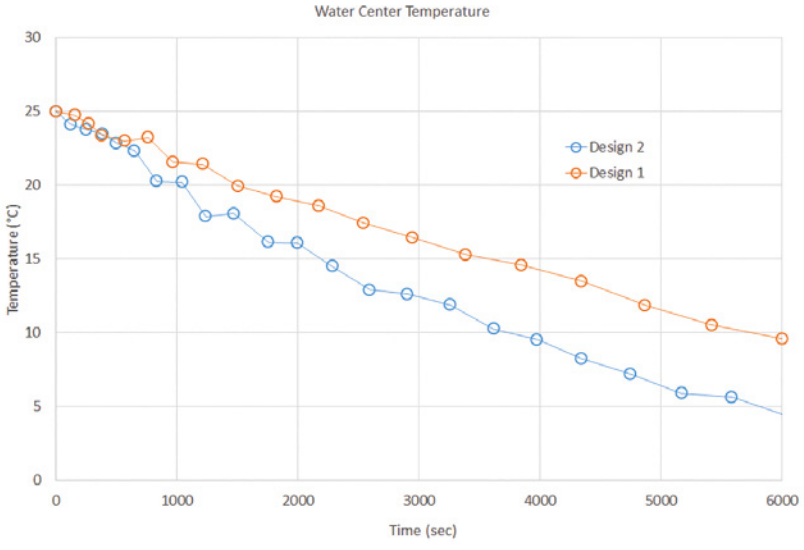

Figure 8. Water center temperature vs. time

Bitcoin Mining: A Thermal Perspective

By John Wilson, Technical Marketing Engineer, Mentor, a Siemens Business

Bitcoin is arguably the most well-known of the crypto-currencies in part because it has been around for nearly 10 years and has increased in value dramatically over the last few years. Another reason Bitcoin is in the news has to do with the high energy usage of the crypto-currency mining industry. For example, Iceland is a popular destination for Bitcoin mining because energy costs are very low, temperatures are cool, and internet speeds are high. It is estimated that the Bitcoin mining industry in Iceland will use about 840 gigawatt hours of electricity in 2018, which is more than all of Iceland's homes did in 2017 (Morris, 2018). Why does Bitcoin mining consume such large amounts of electricity and how is that related to thermal engineering?

What is Bitcoin Mining?
Simply put, Bitcoin miners run a hashing algorithm on computer hardware to determine the correct hash, a fixed length 64 character(256-bit) string, which represents a block of Bitcoin transactions. Once a miner has determined the correct hash, as verified by other miners, that block of data is added to the Bitcoin blockchain, which is an encrypted ledger of all Bitcoin transactions. The miner that determines the correct hash is rewarded Bitcoins (BTC), currently 12.5 BTC. The Bitcoin block mining reward halves every 210,000 blocks or about every 4 years.

The total number of Bitcoins that will ever be produced is 21 million, of which about 80% have been mined. It is estimated that the last Bitcoin will be mined in 2140. The time estimate should be accurate since the Bitcoin mining difficulty of generating the correct hash is managed such that a new block is added to the blockchain about every ten minutes. As the rate of block solutions, or hashes, increases, so does the difficulty of generating the correct hash. The current hash rate of the Bitcoin mining network is about 25 EHash/s (1018 Hash/sec).

Evolution of Bitcoin Mining
From 2009, when the first Bitcoins were

mined, to 2012, the Bitcoin mining network hash rate increased from MHash/s (106) to THash/s (1012) to today's EHash/sec (1018) and as a result mining hardware has evolved. In the early days miners could use their CPUs to mine at a hash rate on the order of MHash/s. Desktop computers could also be used to mine or dedicated custom mining rigs like the one shown in Figure 1. Note the excessive use of cooling fans on each chip. Many mining' rigs are simply a collection of existing systems with little redesign with respect to cooling efficiency.

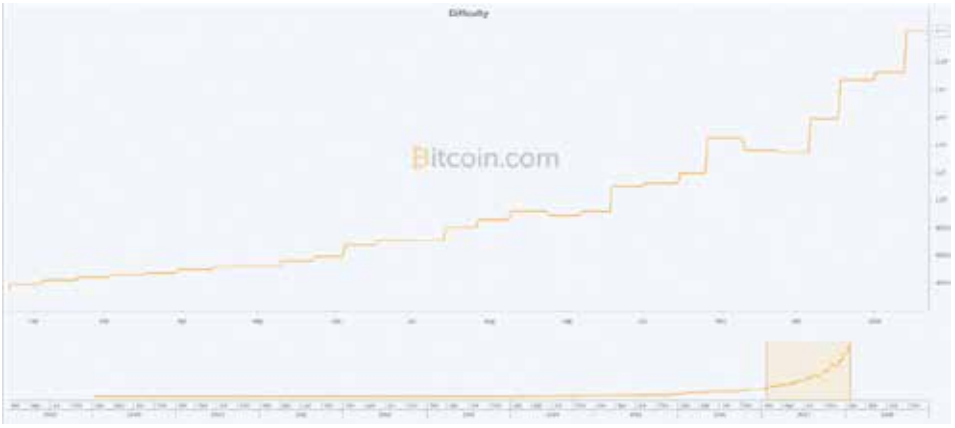

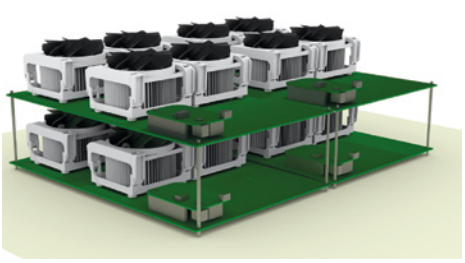

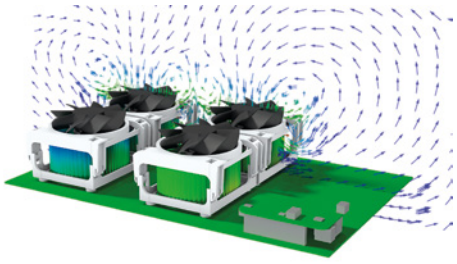

Figure 1. CPU mining rig

As the hash rate of the network has increased, two important metrics have emerged. The first metric is simply the hash rate of the mining hardware. Similar to a single ticket in a raffle, the odds of winning are greatly increased with more than one ticket, but it only takes one to win. The greater the hash rate the greater the chance of calculating the correct hash. It is therefore much more likely that many calculations will be required, so in addition to the hash rate it is important to consider the efficiency of each calculation, or the energy required to produce the hash. The second metric that is used when comparing mining hardware is the energy consumed per hash.

Because of low hash rates as compared to the Bitcoin network and efficiencies in the J/MH (Joule per Mega-Hash) range CPU, mining using CPUs became financially unviable (Mining - Bitcoin Wiki, 2018) and evolved to the use of GPUs. The popularity of GPUs was presumably aided by the availability of off-the-shelf graphics cards with integrated cooling solutions, as shown in Figure 2, and the ability to scale up to an entire mining rig fairly easily.

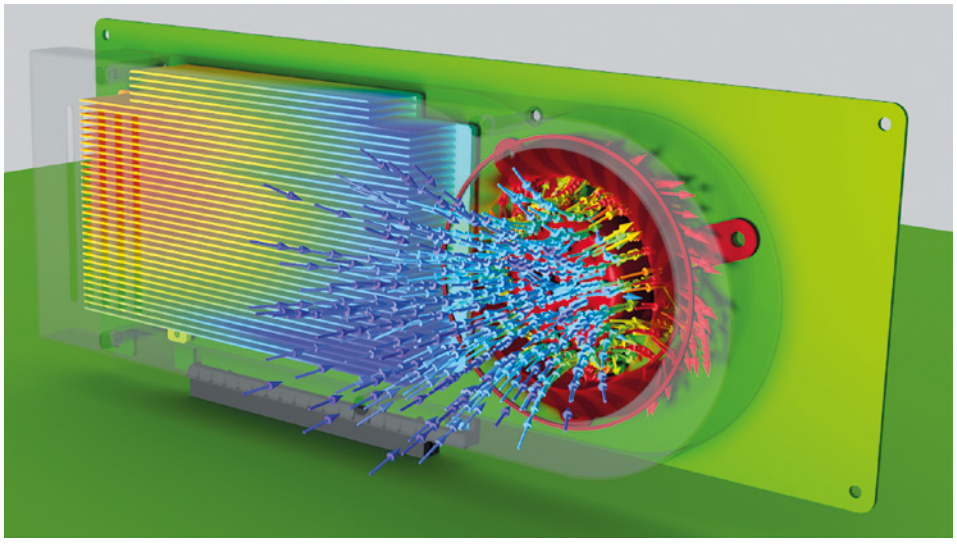

Figure 2. GPU graphics card

Bitcoin mining subsequently evolved to the use of FPGAs that offer hash rates similar to a GPU with increased efficiency due to lower power consumption as compared to a GPU. In 2013, the first ASIC specifically designed for Bitcoin mining was released (Mining - Bitcoin Wiki, 2018). The year the ASIC was introduced the Bitcoin mining network hash rate increased from THash/s (1012) to PHash/s (1015) range and BTC increased in the open market from about $100/BTC to around $1000/BTC by the end of the year (Bitcoin.com Charts, 2018). Note that it peaked at $20,000/BTC in December 2017 but has fallen to the $7,000/BTC mark recently.

For Bitcoin miners looking for an off-the-shelf solutions that use ASICs there are a number of options available that address varying investment comfort levels. At the lower end, a miner might invest in a USB miner, as shown in Figure 3. At the higher end there are mining rigs available like that shown in Figure 4 (from companies such as Bitmain with its AntMiner units).

The USB miner is clearly a low power solution that is limited by the USB port capability. Even at the 2.5W maximum power supplied by the USB port, the heatsinked ASICs on the device can get very warm. The ASIC temperatures are well within typical

operating limits but the heatsink would feel hot to the touch. You often see these devices operating with one or more fans directing air across them. USB miners offer hash rates in the 1GH/s range with efficiencies of about 1 J/GH, easily outperforming the CPU based miners from the early days of 2009.
The ASIC mining rig will have multiple PCBs filled with Bitcoin-hashing ASICs and will consume about 1 KW. Rather than a collection of existing systems this type of rig has more of a holistic cooling approach. The cooling strategy might therefore involve forced convection with two fans in a push-pull arrangement. The ASICs will have heatsinks, either a monolithic approach as shown in Figure 5, or simply a small heatsink mounted on each ASIC with thermal adhesive. Mining rigs offer fantastic performance with hash rates in the 1TH/s range and efficiencies of 0.1 J/GH although they do require a significant initial investment of ~$1,000 - $3,000 in equipment.

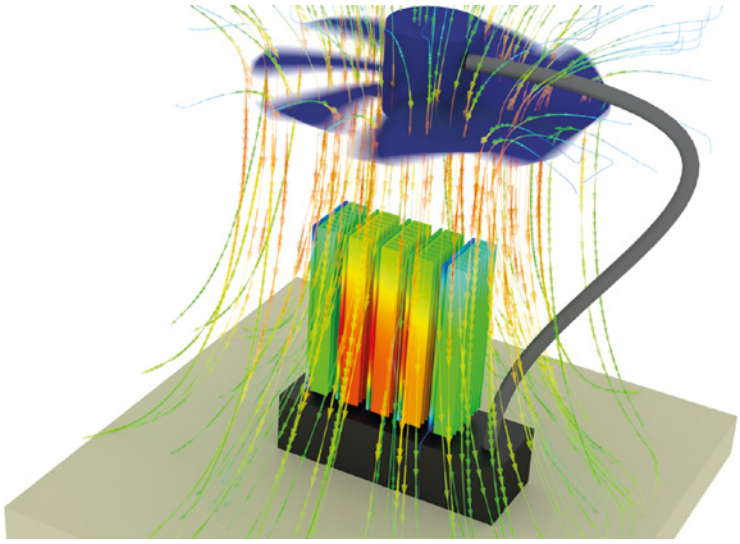

Figure 3. Typical Bitcoin ASIC USB miner

On an even bigger scale of Bitcoin mining there are mining farms. These large mining operations are located in regions that have low energy costs like Iceland (with its cheap geothermal energy and cold air cooling climate), or central Washington State in the United States (with its cheap hydropower). The buildings used in mining farms (Figure 6) are typically much more like a barn or warehouse or shipping container than a datacenter. These operations clearly use megawatts of electricity and cost $10,000,000s to build and millions of dollars a month just to maintain.

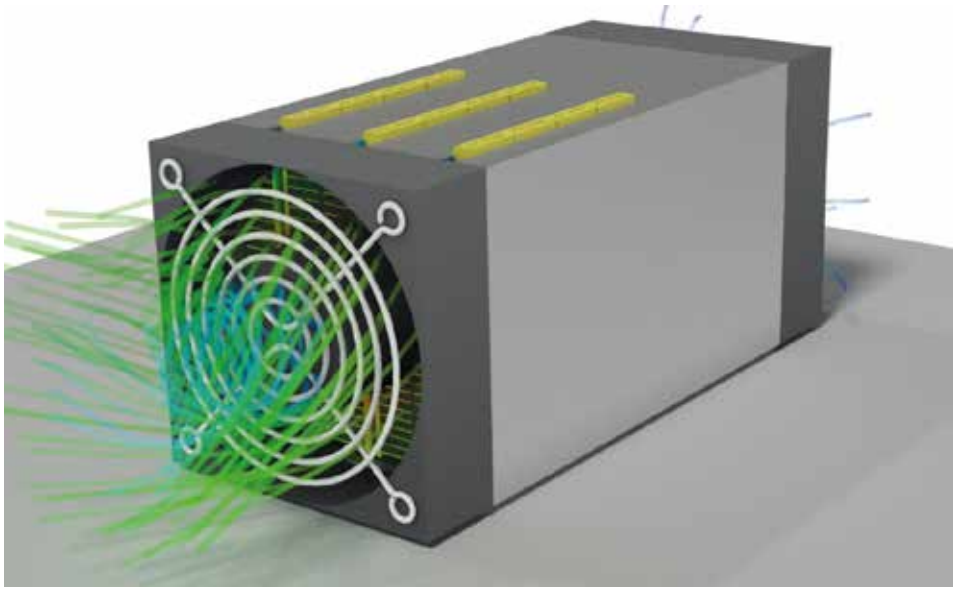

Figure 4. Typical Bitcoin ASIC mining rig

The Bitcoin farms in places like Iceland also benefit from the year-round cool air that can be drawn in via vents on the building exterior, and across the mining equipment, exhausting from the roof of the building (Figure 7).

Challenges and Opportunities

Nearly all of the technical innovation in the Bitcoin mining industry has happened at the IC package level with the development of the Bitcoin mining ASIC. Innovations continue at the IC package level that will no doubt provide higher hash rates but presumably at an increased power. Other areas of innovation will involve submerging the mining rigs in oil or engineered fluids suitable for electronics (i.e. liquid cooling). The advantage of this approach is the ASICs could be more densely packed and could operate in a much smaller facility than the current Bitcoin farm approach.

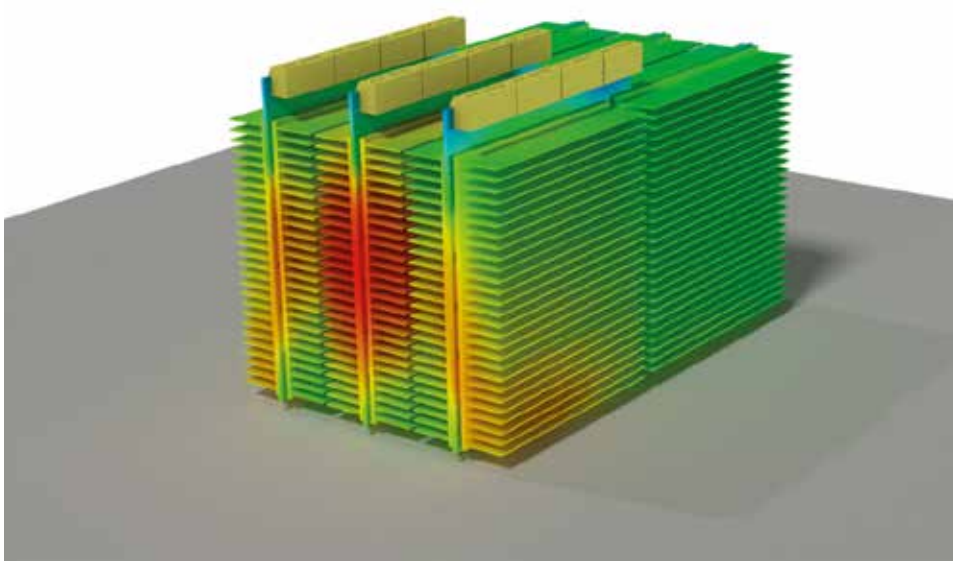

Figure 5. Typical Bitcoin ASIC mining rig internals

It was much easier to mine for Bitcoin when they had essentially no value. Now today BTC trades for thousands of dollars and the

mining network hash rate is at the EHash/s level which makes it very difficult for a small mining operation to compete, or even turn over a profit. As a result individual miners join mining pools where resources are pooled and rewards are split.

Our conclusion from a thermal perspective is that the cost of mining for cryptocurrencies in terms of electronics cooling capital costs and 24/7 electricity usage and maintenance costs are such that it might be a better use of your cash to buy a lottery ticket especially as the gold rush and diminishing number of coins to be mined accelerates. Moreover, there is a major ethical question over the use of our planet's power generation resources (which are stretched as it is) for something that essentially appears to be speculative mining to gain a fortune in cryptocurrencies in what may turn out to be a commodity bubble. Indeed, cryptocurrencies are known to be used extensively in the dark web, money laundering and drug deals. Only time will tell on the cryptocurrency mining boom, and of course the big traditional banks are actually better placed today with their massive datacenters to do hashing calculations for block chain creation than rough and ready mining operations. However, we do think that there will definitely be a place in the cashless world and banking sector of the future for block chain creation algorithms and computer datacenters because of its secure encryption capabilities.

References:

[1] Bitcoin.com Charts. (2018, March 28). Retrieved from https://charts.bitcoin.com/chart/hash-rate

[2] Mining - Bitcoin Wiki. (2018, February 2018). Retrieved from https://en.bitcoin.it/wiki/Mining

[3] Morris, C. (2018, February 13). Bitcoin Mining Uses More Energy Than Homes Do in Iceland. Retrieved from Fortune.com: http://fortune.com/2018/02/13/iceland-bitcoin-mining-electricity/

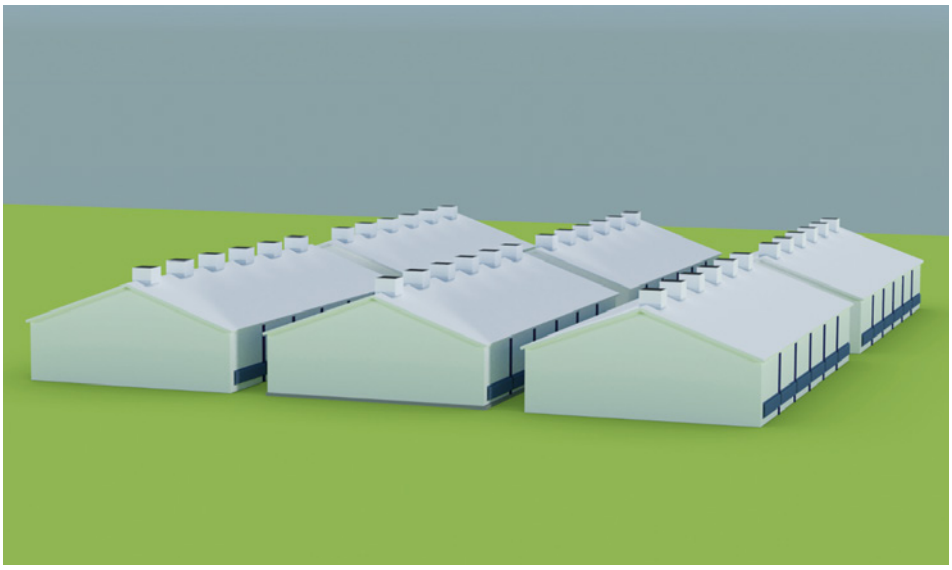

Figure 6. Typical Bitcoin mining farm (effectively a Datacenter)

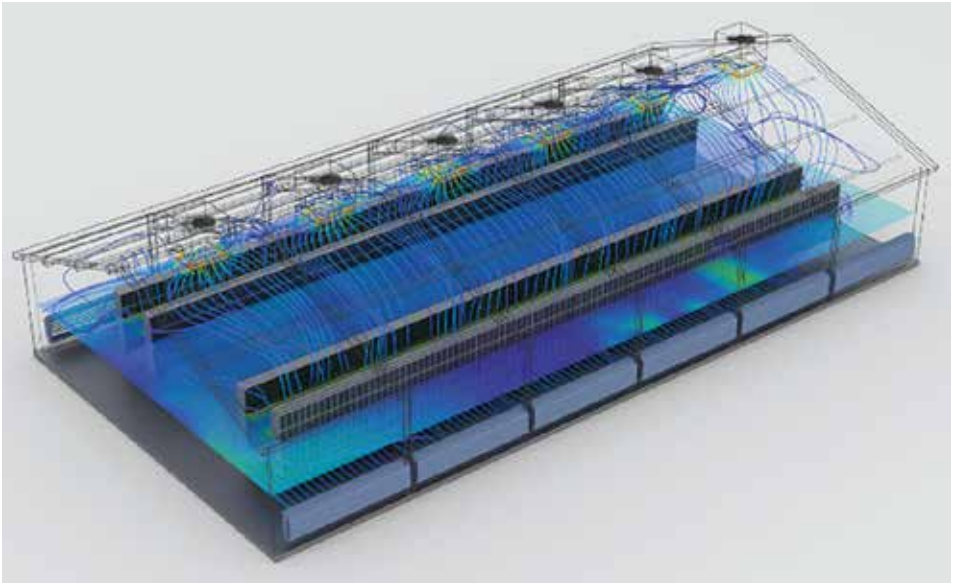

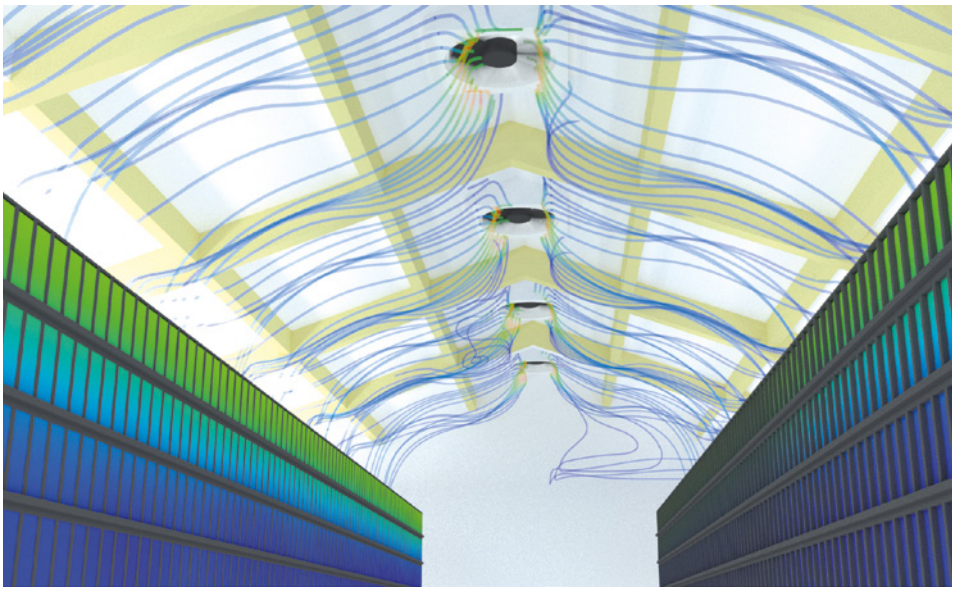

Figure 7. Typical Bitcoin Farm Ventilation Flow and Exhaust System

Delphi4LED - From Measurements to Standardized Multi-Domain **Compact Models of LEDs**

A Report on the Thermal Modeling Aspects

By Genevieve Martin, Phillips Lighting

The rise of LED technology is changing the ecosystem of the lighting industry. With the commoditization of LED, Delphi4LED, a European Union consortium, is providing a solution to this dynamic market. It is providing the EU LED Lighting industry with a set of methodologies and standards to enable the design and production of more reliable, cost effective and market-leading LED-based lighting solutions.

Although LEDs are very efficient, and therefore dissipate much less heat than conventional light sources such as incandescent, halogen or high-intensity discharge lamps, their efficiency is dependent on the LED temperature. This temperature depends on how effectively the LED dissipates its power to the surrounding environment, the amount of heat dissipated by the device and the effectiveness of the heat transfer paths inside the LED component. Thermal-optical-electrical parameters are inter-related but incomplete or even erroneous in datasheets of LED suppliers.

At luminaire developers, this results in considerable reverse engineering effort to get the proper product information about packaged LED. Ways to do this include material analysis, cross-section analysis, computer simulation and measurements. All these reverse engineering techniques combined lead to reasonable results but require long development time and cost a lot of money. For example the exact package dimensions and materials properties – specific features and parameters of the LED chip and phosphors – are unfortunately proprietary information of LED suppliers and cannot be expected to be publicly shared.

The Delphi4LED consortium includes 15 partners from seven countries and seeks to identify and exercise methods to extract and use multi-domain compact models of LEDs. A modular approach to such compact models [1], [2], [3], [4] would then be employed to enable lighting designers to meet both thermal reliability and optical design goals (Figure 1).

Such models take forward current as an input, and calculate emitted optical properties such as luminous flux as well as operating junction temperature for each LED in a design. The temperature dependence of the conversion of electrons to photons is considered, which impacts the resulting thermal power dissipation which in turn affects the temperature etc. It is this electro-optical-thermal coupling that necessitates all three domains to be considered concurrently in a single multi-domain modular compact model.

From the thermal perspective, although methods such as DELPHI exist to extract compact thermal models (CTMs) of monolithic single heat source IC packages, LEDs can have multiple heat sources. This is the case for White LEDs where additional heat is dissipated in phosphor layers due to Stokes shift light conversion losses. In addition LED packages may have multiple LED chips within them as is the case for Chip on Board (CoB). Temperature prediction at each chip junction, in each phosphor layer and also at solder points is required to ensure reliability in operation.

Dimming behavior and operation under AC conditions also necessitates that these CTMs are capable of dynamic prediction, i.e. dynamic compact thermal models (DCTMs).

To extract such DCTMs, Delphi4LED is determining a methodology that starts with a physical LED sample, performs a transient thermal measurement, uses that to calibrate a 3D detailed thermal model and from that a DCTM with an assumed nodal topology may be extracted (Figure 2).

T3Ster and TeraLED are used to perform the transient thermal response measurement in compliance with the latest JEDEC LED testing standards and recommendations of CIE, such as JESD51-51, JESD51-52 and CIE 225:2017 [5], [6], [7]. The total emitted optical power is measured by the TeraLED integrating sphere, the resulting thermal power dissipation (the difference between the electrical supplied power and the optically emitted power) is used to correct the resulting Zth profiles ((Tj-Ta)/ thermal dissipated power) vs. time. An important aspect of these combined thermal and optical measurements is that the LEDs' junction temperature is kept at known, constant value as also recommended by the CIE 225:2017 technical report. As a side product of such a test procedure, the temperature dependence of the light output properties of the LEDs can also be measured.

A FloTHERM model of an LED package is then constructed of the same configuration used for the T3Ster+TeraLED measurement (Figure 3).

FloTHERM's calibration feature is then used to fine tune material properties along the heat flow path until such time as simulated Zth responses (and corresponding structure function profiles) match the measurement. This entails nomination of those materials properties that are both most unknown and contribute most to the overall thermal resistance. In this case the chip submount, bottom die attach and FR4 thermal conductivities. Upper and lower bounds of those parameters are defined, a (computational) design of experiments set and solved then a gradient based optimization used to determine which parameter values result in the smallest deviation between measured and simulated transient thermal response curves (Figure 4).

Unlike the measurement that was limited to only being capable of recording the

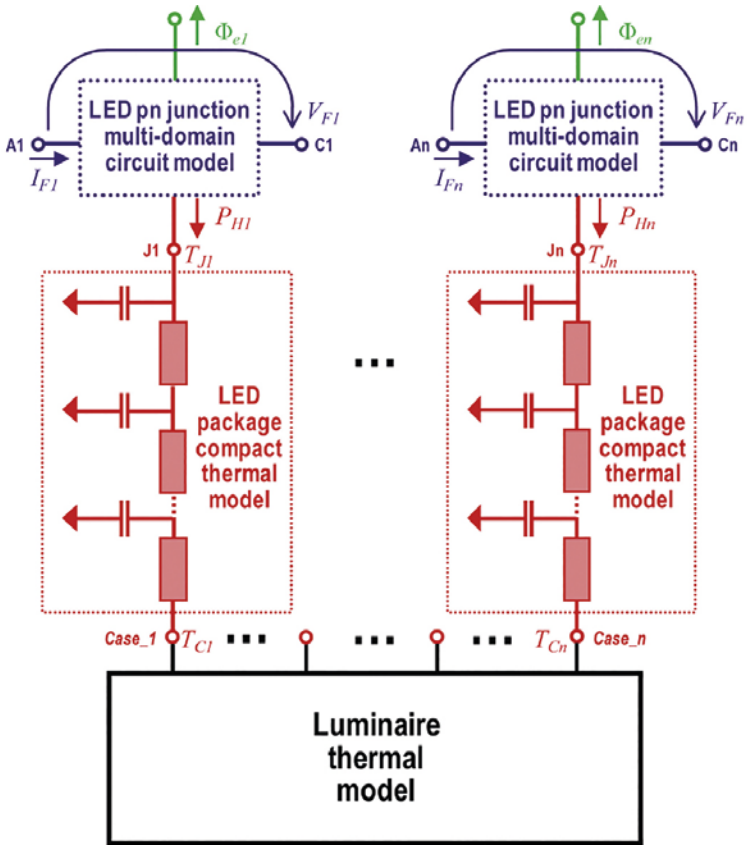

Figure 1. Multi-domain compact model of Multiple LEDs within their Operating Environment [3], [4]

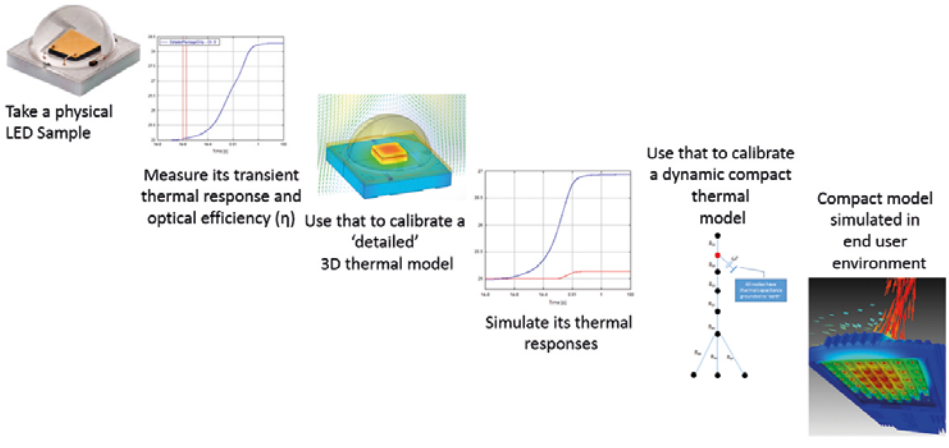

Figure 2. Dynamic compact thermal model extraction process

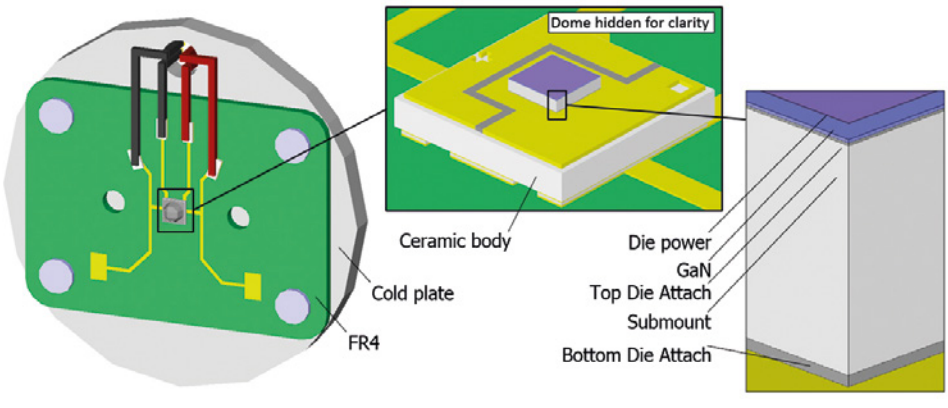

Figure 3. Cree XP-E2 Blue LED on FR4 on coldplate [8]

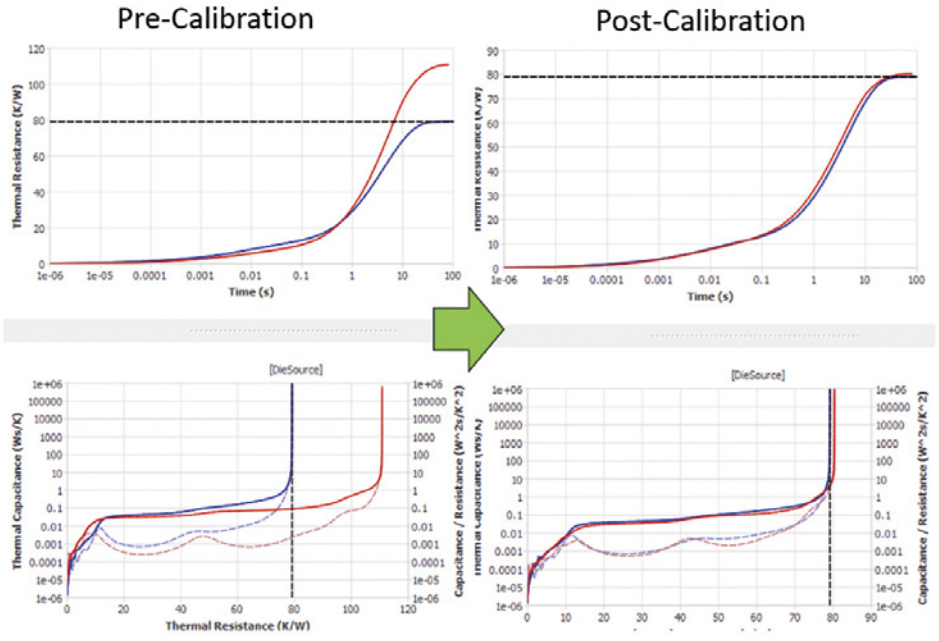

Figure 4. Pre- and post-calibration Zth (top) and Structure Function (bottom) comparisons. red = Simulated response. Blue = Measured response.

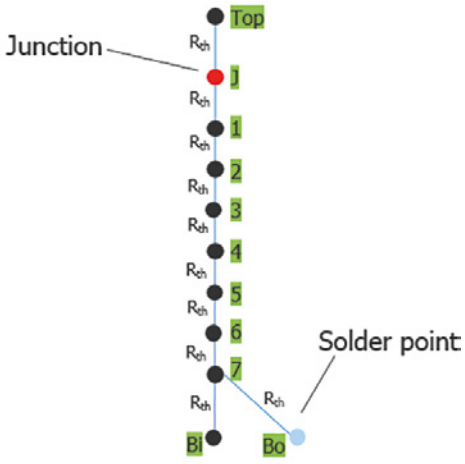

Figure 5. Nodal DCTM topology (thermal capacitance values at each node, not shown for clarity)

chip junction thermal response, the 3D FloTHERM model is capable of predicting temperature and 1000s of points. This additional simulated thermal response behavior is subsequently used to calibrate the thermal resistance and capacitance values at certain points in a nodal DCTM.

The topology of the DCTM has been determined having enough points so as to capture the dominant thermal time constants experienced along the heat flow path (Figure 5)

The heat flow path within the package, away from the package periphery, is boundary condition independent and can be represented with thermal resistance (R) and thermal capacitance (C) values taken directly from the measured structure function. The R and C values nearer the package periphery have to be determined so that the thermal response of the DCTM at both junction and solder point nodes match with the corresponding points of the 3D detailed model. The same calibration methodology is applied to achieve this, considering the DCTM in isolation with representative boundary heat transfer coefficients applied on top and bottom surfaces. In this case there are two pairs of curves that need to match, the driving point impedance at the junction as well as the transfer impedance of the solder point. The results of the DCTM calibration are shown in Figure 6.

A verification of the accuracy of the DCTM can be ascertained by comparing the resulting thermal response when the

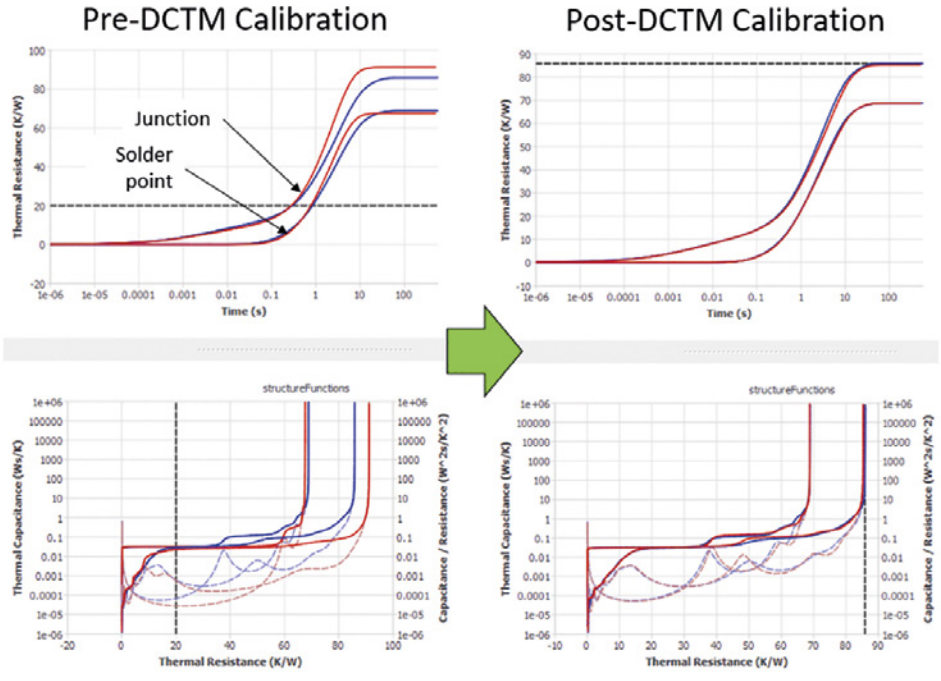

Figure 6. Pre- and post-DCTM calibration Zth (top) and Structure Function (bottom) comparisons. Red = DCTM response. Blue = 3D Detailed model response.

DCTM is placed back into the full system level model representing the original T3Ster+TeraLED measurement (Figure 6)

A maximum error of 4.5% in dT prediction is noted, but only during 1-10s of the transient thermal response. Steady state temperature prediction is highly accurate.

The same methodology is being extended in the Delphi4LED project to extract boundary condition independent DCTMs by calibrating

R and C values so that all thermal response at all points of interest are accurate over a set of differing top and bottom heat transfer coefficient values. The number of points of interest may be extended to incorporate other dissipating nodes (i.e. nodes representing phosphor conversion) and additional solder point temperatures.

Coupled with a chip level model of LEDs (extracted directly from measurement) describing the temperature dependent electrical and optical characteristics, the DCTM forms part of the multi domain compact model. An accurate representation of the LED capable of providing thermal and optical information quickly and easily so that lighting designers may meet both functional and reliability product specifications.

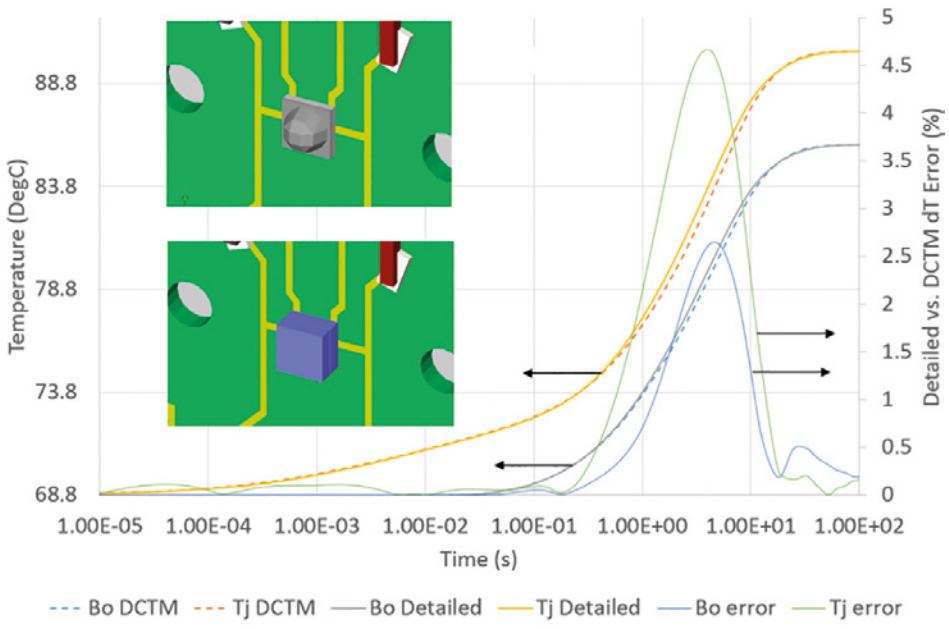

Figure 7. DCTM vs. 3D Detailed model comparisons

References:

[1] R. Bornoff, V. Hildenbrand, S. Lugten, G. Martin, C. Marty, A. Poppe, M. Rencz, W. Schilders, J. Yu, Delphi4LED - From Measurements to Standardized Multi-Domain Compact Models of LEDs: a New European R&D Project for Predictive and Efficient Multi-domain Modeling and Simulation of LEDs at all Integration Levels Along the SSL Supply Chain, Therminic 2016, Budapest, Hungary

[2] www.delphi4LED.org

[3] A. Poppe, J. Hegedűs, A. Szalai, R. Bornoff, J. Dyson, "Creating multi-port thermal network models of LED luminaires for application in system level multi-domain simulation using SPICE-like solvers", In: Proc. of the 32nd IEEE SEMI-THERM Symp., 14-17 March 2016, San Jose, USA, pp. 44-49, DOI: 10.1109/SEMI-THERM.2016.7458444

[4] A. Poppe, "Simulation of LED Based Luminaires by Using Multi-Domain Compact Models of LEDs and Compact Thermal Models of their Thermal Environment", MICROELECTRONICS RELIABILITY 72:(5) pp. 65-74. (2017), http://dx.doi.org/10.1016/j.microrel.2017.03.039

[5] JEDEC JESD51-51 Standard "Implementation of the Electrical Test Method for the Measurement of Real Thermal Resistance and Impedance of Light-Emitting Diodes with Exposed Cooling" (2012), https://www.jedec.org/system/files/docs/JESD51-51.pdf

[6] JEDEC JESD51-52 Standard "Guidelines for Combining CIE 127:2007 Total Flux Measurements with Thermal Measurements of LEDs with Exposed Cooling Surface" (2012), https://www.jedec.org/system/files/docs/JESD51-52.pdf

[7] Y. Zong, P-T. Chou, P. Dekker, R. Distl, K. Godo, P. Hanselaer, G. Heidel, J. Hulett, K. Oshima, A. Poppe, G. Sauter, M. Schneider, H. Shen, M.M. Sisto, A. Sperling, R. Young, W. Zhao, "Optical Measurement of High- Power LEDs", CIE 225:2017 (Technical Report) ISBN 978-3-902842-12-1, DOI: 10.25039/TR.225.2017

[8] R. Bornoff, G. Farkas, L. Gaal, M. Rencz, A. Poppe, "LED 3D Thermal Model Calibration against Measurement", In: Proceedings of the 19th International Conference on Thermal, Mechanical and Multiphysics Simulation and Experiments in Microelectronics and Microsystems (EuroSimE'18), 15-18 April 2018, Toulouse, France

A Cool **Emulator**

By Jean Paul Prevot, Mentor, a Siemens Business, Emulation Division

Mentor, a Siemens Business, Emulation Division designs, manufactures and sells a hardware emulation system under the Veloce® registered trademark. A wide range of semiconductor and systems manufacturers around the world verify their chip and system designs with the Veloce emulator.

They use the Veloce emulator because it is tens of thousands times faster than the alternative software simulation approaches thus boosting productivity in the verification of SoC and Intellectual Property (IP)-based designs. Veloce also accelerates block and full SoC RTL (Register Transfer Level) and gate-level simulations during all phases of the design process, and enables pre-silicon testing and debug at hardware speeds, using real-world data, while both hardware and software designs are still fluid.

The Veloce Strato platform, is the current generation of the Veloce platform (Figure 1).

Mentor's emulation division has been a leader in the development of emulators designed for the best energy use profile possible making the Veloce Strato emulator a very efficient and datacenter friendly tool. Despite this, tens of kWs in thermal energy, in total, are generated by the emulator's boards and AC/DC power modules during normal operation which need to be removed.

Since 2004, the Mentor Emulation Division has been using the FloTHERM® electronics cooling CFD product to optimize internal airflow and cooling for the Veloce emulation platform.

The development of the Veloce StratoM configuration presented a number of mechanical engineering challenges. One was to fit the Veloce StratoM into the same system layout as the previous generation Veloce2 (same external dimensions) even though it has longer logic boards and increased logic board pitch. Another challenge was to define a unique logic board pitch compatible with cooling and midplane connector positioning rules dictated by the use of mixed technologies.

Increasingly, numerous constraints needed simultaneous management to comply with all macropackaging requirements in terms of:

- boards interconnect
- cabling
- accessibility
- mechanical strength and stability
- AC and DC powering
- cooling, sub-assemblies positioning
- audio noise control
- pressure drop limitation
- air inlet and air outlet locations
- EMC shielding
- safety norms compliance
- customer site access capability
- aesthetic
- system configurations, etc.

All while complying with manufacturing requirements in terms of design for manufacturing, parts commonality, cost optimization, etc.

Regarding cooling aspects, several complementary explorations were used to deal with cooling of both the chip package, logic boards, board module sub-assemblies and global system.

The team started with rough estimations for the global airflow needed to insure global heat dissipation, and manual calculations to determine fan types, to pre-dimension grid air inlets and outlets, and to consolidate board pitch and package heat-sink draft design. Because of the wide range of inputs, we were quickly led to build several CFD models using FloTHERM software modules.

By using FloTHERM PACK at the package level, we were able to easily obtain a compact thermal model of our chip package enabling us to integrate it into our logic board CFD model and to obtain a good estimation of junction temperature. (Figure 2)

The compact model provided the advantage of being computationally less

Figure 1. The Veloce StratoM emulator

expensive while still insuring reliability and accuracy. Both “2 resistors” as well as “Delphi” compact models, currently being standardized by the industry standards body JEDEC, were generated and used.

By using FloTHERM PCB at the board level, we were able to build our logic board CFD model using data imported directly from Expedition®. PCB layers, components and DC/DCs connectors were automatically built and then parameters were detailed to “dress” the model. Identification of a sufficiently accurate power dissipation estimate for each key component is a crucial phase and results are dependent on this data. Once the board’s model was available, we studied the central board assembly located in a “three pitch boards” model. This allowed us to perform package heatsink optimization in terms of efficiency and pressure drop and DC/DCs or component repositioning on PCB layout.

One of the numerous advantages of FloTHERM is the ability to detail, in exacting granularity, some components as package or heatsink and decide to use them with a simplified representation to save computation time while keeping a reliable model. We could also add dissipative zones to represent the heating effect of any deliberately missing components. In addition, the use of localized grids sped up the calculation.

We were able to determine the optimal temperature sensor positioning on PCB’s (numbered 1 to 3 in Figure 3); thermal sensors used to monitor fans’ speed while emulation is running, enabling us to manage audible noise levels.

We were able to study temperature at both PCB level and component level using the cross section thermal planes shown in Figure 4.

By using FloTHERM at the chassis level, we were able to optimize sub-assembly positioning inside the chassis—primarily board modules, air inlets, lower and upper fan drawers and fan positioning inside the drawers.

It was important to balance the air intakes, taking into account the front and rear areas for both lower and upper levels. We needed to prevent dead air zones inside the system due to the excessive air velocity at inlet, and lack of air plenum between fans and the board’s edge.

The last step was to build the entire system model to validate global heat transfer

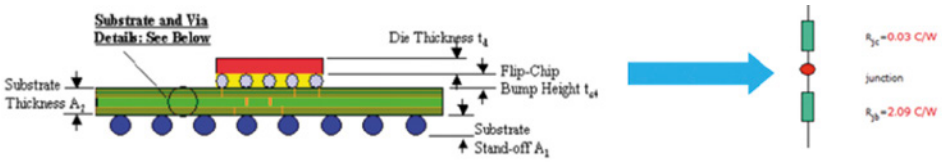

Figure 2. Example of compact model creation for one of Mentor’s packages

and verify the absence of dead areas or deleterious air recycling zones (inducing hot spots and/or audio noise disturbances) for different system configurations (board population, failed fans, etc.).

FloTHERM also helped us to dimension our copper DC busbar enabling us to take into account the Joule effect inside copper as it is carrying several hundreds of amperes, and to deal with current distribution all along the midplane.

Another highlight of our work was the use of FloVENT that let us study different system layouts at customer sites and analyse the benefits of using hoods to manage air exhaust through the ceiling.

To conclude, from the beginning of the Veloce Strato design, each step was taken to optimize air path, optimize air velocity and minimize pressure drop. Continuous effort was applied from the package to global system.

The main goal of using FloTHERM CFD software was to ensure optimum package junction temperature to reach the optimal electronic performance and reliably while using the minimum airflow. Thus allowing the system to operate within a reasonable temperature range.

FloTHERM CFD software is a great help in showing trends associated with parameter variations. Moreover, the “smart” pre- and post-processors offered by FloTHERM CFD modules help you understand your product so you can detect unexpected physical behavior and optimize the design. FloTHERM is like a bike; you never forget how to use it.

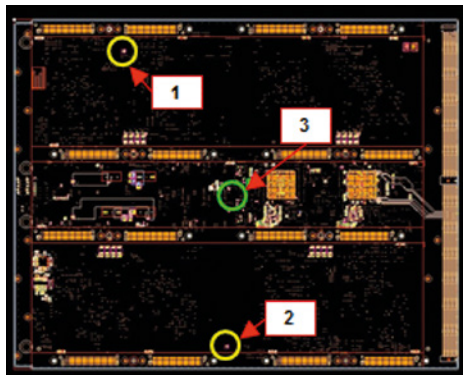

Figure 3. Mentor’s logic board thermal sensors positioning

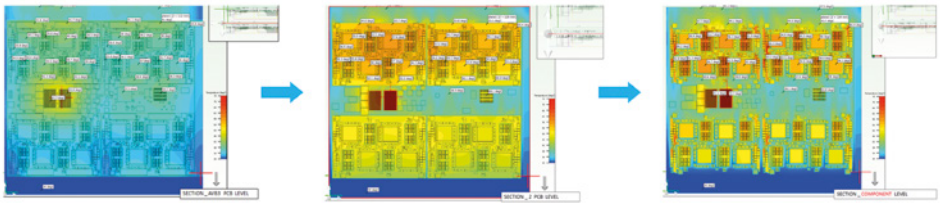

Figure 4. Mentor’s logic board thermal study - PCB / Components sections

Thermal Interface for Pluggable Optics Modules

By Bonnie Mack, Senior Thermal Engineer and Terence Graham, Senior Thermal Engineer, Ciena Corporation

ciena
Experience. Outcomes.

Power on pluggable optics modules (POMs) such as SFP+, QSFP+, QSFP28, CFP2, CFP8 has increased along with the demand for higher bandwidth. POMs give access at the faceplate to an optical signal. Existing Multi-Source Agreements (MSAs) specify physical form factor and electrical interfaces, which allow multiple manufacturers to make physically compatible products to promote competition, interoperability and multiple sources for systems vendors and end users. These MSAs also define power classes for POMs that are based on the supplied power and correspond to different internal processing levels and optical signal reach. POMs are designed to support various communication standards and their data rates range from ~1 Gb/s to 400 Gb/s with many data rates available in each form factor.

POMs are difficult to cool; all but the CFP, are housed in a cage which guides the modules to the connectors and contains EMC solutions for the faceplate ports. The modules extend through the faceplate, and can be hot swapped. There are air gaps, ~0.2mm to 0.3mm nominal, between the module case and the cage, and between the cage and PCB. This provides an inconsistent thermal resistance due to tolerances. The cages generally have small openings on the sides to allow ingress and egress of air for cooling purposes. With high power POMs, these openings do not provide sufficient cooling so an opening on the cage top is added which gives access for a spring-loaded riding heatsink. To date only a dry thermal interface between the two surfaces has been available because the pluggable feature has precluded the use of thermal interface materials between the heatsink and the POM case. The thermal interface between POM and heatsink is not consistently defined or controlled in the MSAs. The challenge is to permit the required sliding while providing a low interface thermal resistance. This is especially important in a telecom environment where equipment in NEBS [1] shelf level products must operate in an ambient of 55ºC. POMs temperature case limit is usually 70ºC. This results in a 15ºC temperature delta to cool the POMs, usually less when pre-heated by other POMs or components. Figure 1 shows typical thermally-important features of POMs and cages.

Heat transfer routes to and from the POMs are described in references [2,3]. Thermal evaluation requires detailed 3D conjugate heat transfer analysis software. Essential input for the analyses includes detailed geometric, power, and thermal properties of the POM, EMC gasket and cage, and contact interface as described in the OIF thermal interface specification IA# OIF-Thermal-01.0 [4].

IA# OIF-Thermal-01.0 specifies general resistance parameters for the thermal interface as a function of power density. For high power modules, the major path for heat removal is via the heatsink across the contact area with the POM. It stipulates the MSA to define the location and size of contact area for heat removal on the top surface of POM. It also describes factors affecting thermal interface resistance: flatness, surface finish, and heat spreading. Additionally, it defines a calibration method for the internal sensors and includes the requirement to identify the location of thermal monitor point(s).

The initial work done in support of IA# OIF-Thermal-01.0 IA included a study of the thermal interface resistance between a CFP2 lid and heatsink base including heat spreading effects. The study examined three contact

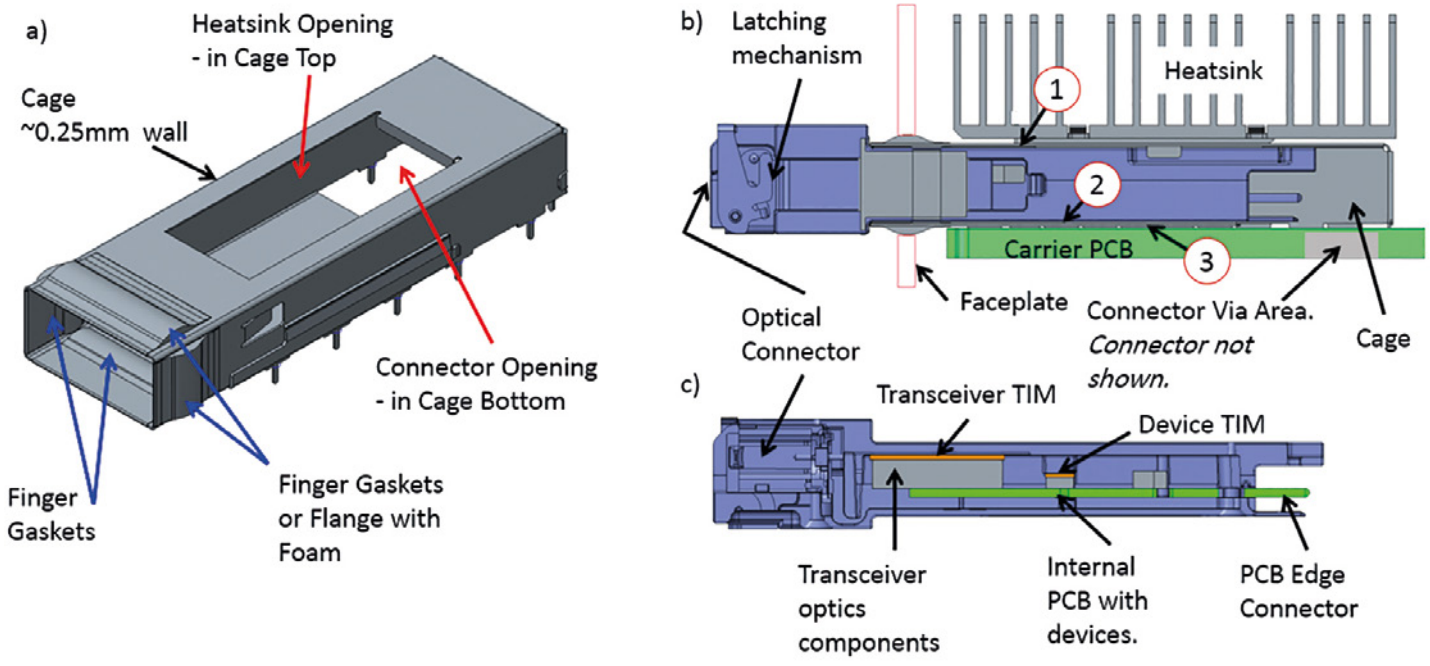

Figure 1. a) A typical POM cage, b) QSFP in cage section at inside edge of cage, c) QSFP section showing typical internal layout. Narrow air gap locations: 1) Module to top of cage, 2) Module to bottom of cage, 3) Bottom of cage to PCB, and not shown 4) sides of module to sides of cage.

scenarios: 1) a transverse bump in the center of the lid, 2) a transverse hollow in the center of the lid, and 3) a transverse contact in the center and ends of the lid. All three scenarios have the same contact area. These simplified contact geometries are depicted in Figure 2 and the contact gap ranges from what would be an extremely fine production surface flatness to 0.3mm, the maximum allowed by the MSA for the CFP2 contact surface for Power Class 1 and 2. It was assumed for the study that both surfaces had the same type of out of flatness so that modeling of the net gap between surfaces was easily implemented in FloTHERM® software used for the analysis.

The resistance in the contact areas between the heatsink and the module was modeled using the method described by Yovanovich et al. [5]. Where:

Joint Resistance $h_j = h_{contact} + h_{gap}$

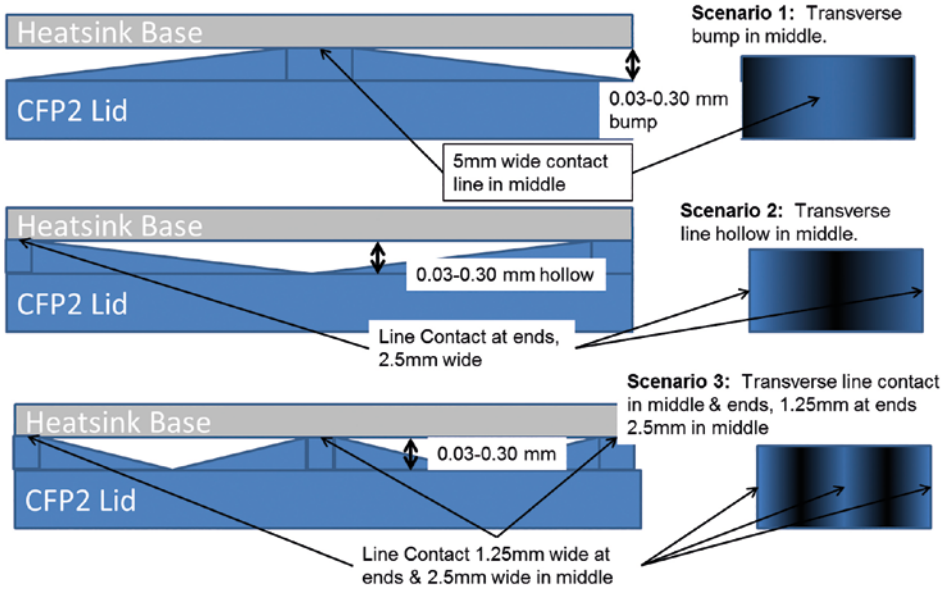

Figure 2. CFP2 Contact Interface Flatness Scenarios

A simplified CFP2 FloTHERM model was created with a T6063 aluminum case and up to six sources that can be set to dissipate power and to contact the lid as shown in Figure 3. Results were obtained for varying gaps due to out-of-flatness and different source locations. Intake air is 55°C at 1m/s across the enclosure cross-section upstream of the module. Total power for the CFP2 is 12W in all cases. Other model details are given in [2,3]. CFP2 lid temperature is monitored directly above the center of each source. Maximum lid temperatures are plotted in Figure 4 versus net flatness over the range of 0.03 mm to 0.3 mm flatness, for the three scenarios depicted in Figure 2.
Defining:

$T_{Lid\ max}$ = maximum lid temperature
$T_{HS\ ave}$ = average heatsink pad temperature
dT_{Lid} = temperature difference among lid locations A – F (indicative of thermal spreading resistance)
dT_{Lid} max to HS ave = TLid max, - THS ave

Spreading resistance in the lid was highest for scenario one, a center bump, which is

also effectively a single contact. Spreading resistance varies with the thermal conductivity and thickness of the lid. The interface resistance between the lid and heatsink is shown to be sensitive to the interface flatness. Temperature differences between lid and heatsink base can be reduced by 5°C to 8.5°C with flatness improvement on a CFP2 depending on heat source location and the nature of the out-of-flatness. This is a third to half of the CFP2 temperature budget in a NEBS environment! This is a huge improvement without increasing the size of the heatsink.

QSFP Model – Expanding on the Effect of Spreading Resistance

The QSFP form factor is a common POM that presently has the highest power density of all the form factors. Most of its heat is dissipated close to the faceplate and not directly underneath the heatsink contact area. A numerical wind tunnel study was conducted to explore the internal resistances and develop methods of reducing the QSFP temperatures. The numerical model, detailed in [2,6], was similar to that of [7] having a 5 W QSFP and a power density of 1.34 W/cm^2, class pd14. Cooling was via a typical aluminum off-the-shelf heatsink. The model was used to predict the effect of changes in the heat source locations relative to the heatsink contact area and surface finish in the interface area. Figure 5 gives the scenario descriptions. For scenario d) the 5mm extension of the heatsink contact area towards the transceivers, the power density decreases 15% to 1.14 W/cm^2 or to pd12. Two contact resistance values between the case and the heatsink were explored as well for two difference surface roughness and load conditions. Results were calculated for QSFP case material thermal conductivities of 116, 169, 385 and 1000 W/m-K corresponding to a zinc alloy, high grade aluminum casting, copper, and an ultra-high conductivity material respectively. FloTHERM Command Centre was used to solve these scenarios.

The results shown in Figure 6 illustrate the importance of the surface finish of the case and heatsink, and of locating the heat sources and the thermal interface area as closely together as possible. QSFP MSA cage dimensions [8-9] allow an increase in heatsink contact length by up to 5mm. In our model this larger contact resulted in a temperature decrease of more than 1.5°C with the lowest case conductivity and Rc1. If the case material conductivity is increased to 169 or 385 W/m-K, further decreases of 1°C to 2°C respectively could be achieved. While very high case conductivities representing an exotic material was examined, changing QSFP case material to Cu from zinc alloy can improve performance by 2 to 3°C. Decreasing the contact resistance to Rc2 would bring total improvement to ~5°C. These are significant when the overall ambient to case temperature budget could be 15°C or less.

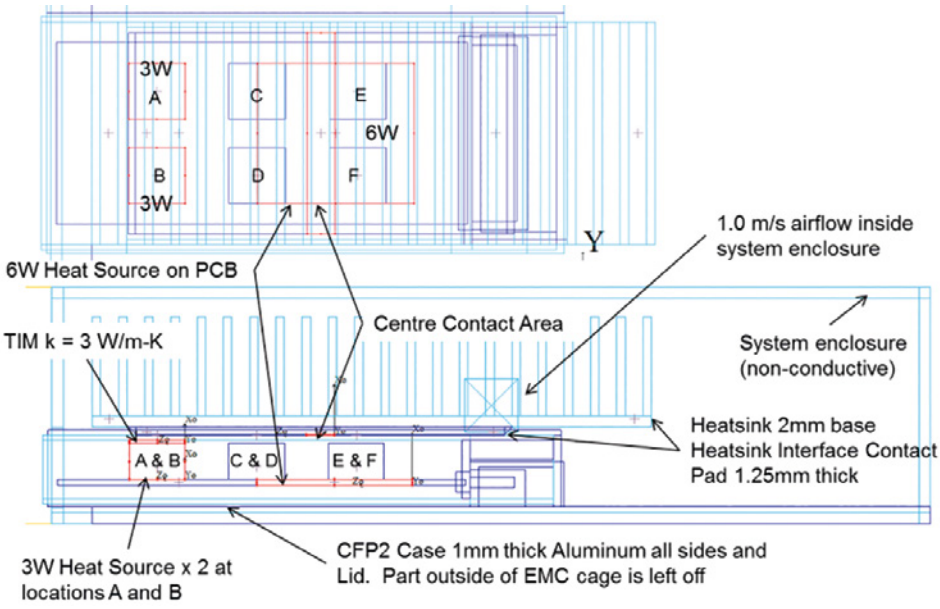

Figure 3. CFP2 flatness model layout showing front device, rear CFP2 PCB heat and central contact location

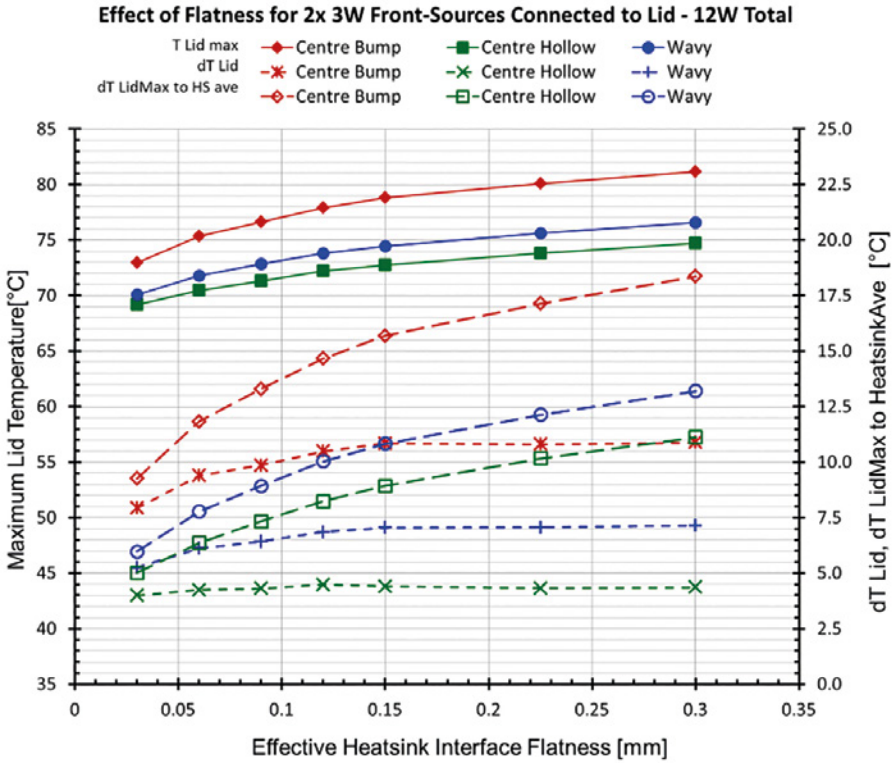

Figure 4. Range of flatness results with front device, and rear PCB heat dissipation and only the devices connected directly to the lid.

The expansion of the heatsink opening in the cage has been incorporated into the MSA for the much higher power QSFP-DD modules [10].

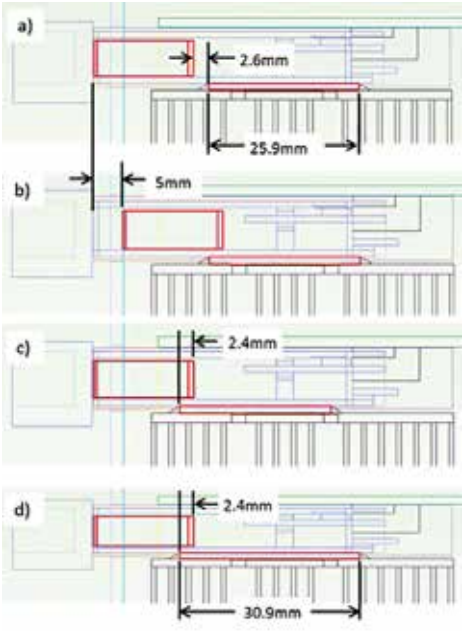

Figure 5. Modeling scenarios for the QSFP a) Original, b) Transceivers 5 mm closer to heatsink contact, c) Heatsink contact 5 mm closer to transceivers, d) Heatsink contact 5 mm longer towards transceivers. Transceiver location is the red rectangle on the left.

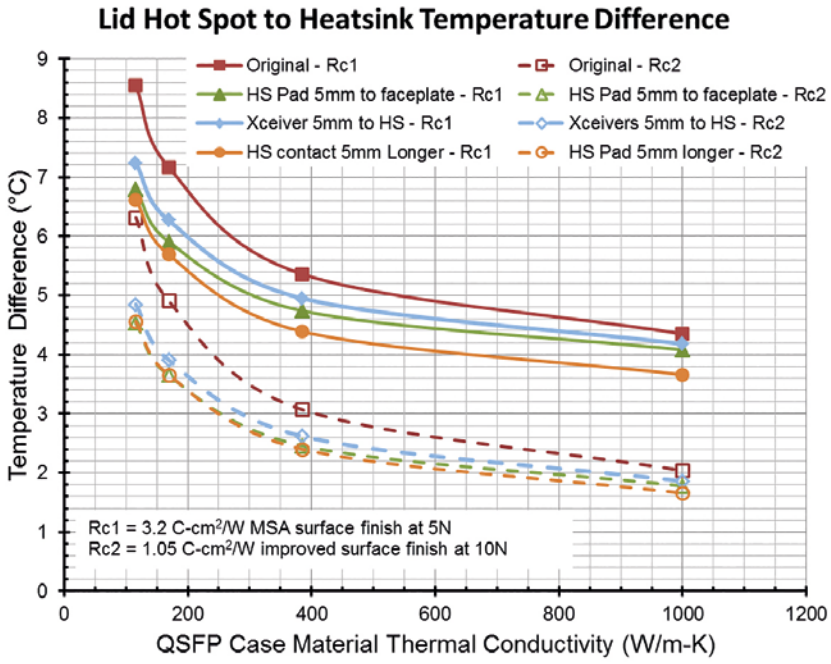

Figure 6. Temperature difference between QSFP hot spot monitor point and heatsink pad. Surface finishes of 0.6 µm Ra and 1.6 µm Ra on both heatsink and case

Final module thermal assessment is at present only available with a computationally expensive detailed model created from the information specified in [4] and either built by the system designer with information from the module supplier, or is a model supplied by the module vendor, likely under some type of non-disclosure agreement. An alternative to this that has not yet been explored is the development of a Delphi-type resistance network. A model of this type could be used to model the connections to the PCB, surrounding air and heatsink with distributed internal heat sources in a manner similar to that used for multi-junction integrated circuit devices. This could be provided by module vendors without giving internal details of the module.

References:

[1] TelcordiaTechnologies Generic Requirements, Generic Requirements for Telecommunications Data Center Equipment and Spaces, GR-3160-CORE, Issue 2, July 2013

[2] Mack, B., Graham, T., Thermal Specifications for Pluggable Optics Modules, 32 IEEE SEMI-THERM Symposium, March 2016.

[3] Mack, B., Graham, T., Pluggable Optics Module – Thermal Specifications, Part 1, Electronics Cooling, June 2016.

[4] IA# OIF-Thermal-01.0, Implementation Agreement for Thermal Interface Specification for Pluggable Optics Modules, Optical Internetworking Forum, May 18, 2015

[5] Yovanovitch, M.M., Culham, J.R., and Teerstra, P., Calculating Interface Resistance, Electronics Cooling, Article 3, May 1997, http://www.electronics-cooling.com/Resources/EC_Articles/May97/article3.htm

[6] Graham, T., Mack, B., Pluggable Optics Module – Thermal Specifications, Part 2, Electronics Cooling, September 2016.

[7] OIF-PLUG-Thermal-01.0, Thermal Management at the Faceplate White Paper, Optical Internetworking Forum, March 2012

[8] SFF Committee, SFF-8663 Specification for QSFP+ 28 Gb/s Cage (Style A), Rev 1.5 April 2014, ftp://ftp.seagate.com/sff

[9] SFF Committee, SFF-8661 Specification for QSFP+ 28 Gb/s 4x Pluggable Module, Revision 2.0, February 2014, ftp://ftp.seagate.com/sff

[10] Quad Small Form Factor Pluggable Double Density (QSFP-DD) Multi Source Agreement (MSA) Group, QSFP-DD 3.0 Hardware Specification, September 2017 www.qsfp-dd.com