DIGITAL INDUSTRIES SOFTWARE

A complete guide to enclosure thermal design

Learn about 14 key considerations

Executive summary
Efficient thermal management is crucial for electronic products, especially at the system or enclosure level. As devices become more compact and powerful, the challenges of dissipating heat effectively intensify, impacting product reliability, performance and lifespan.

This paper provides a comprehensive framework for optimizing thermal performance, emphasizing a proactive, top-down codesign approach by integrating thermal considerations from early product development stages. It addresses crucial enclosure thermal design aspects, from optimizing airflow via fan placement and vent design, to leveraging advanced techniques like internal heat spreaders, phase change materials and thermal radiation effects.

By adopting best practices and utilizing advanced simulation tools, companies can significantly enhance thermal design workflows, creating more robust, cost-effective and higher-performing electronic products that meet stringent market demands.

Dr. John Parry
CEng, Siemens Digital Industries Software

Contents

Introduction

Before we look at 14 key considerations for enclosure thermal design, first contemplate why thermal issues need to be considered at the system or enclosure level.

A natural focus when designing an electronic product is inevitably on the electronics. The electronics need to work within some kind of enclosure, be it a standardized rack system,[1] a custom-built enclosure like the housing of a smartphone or tablet computer or part of another product like an automotive dashboard or aircraft flight deck.

In all cases, including standardized rack systems, engineers must design the electronics with the enclosure in mind, as the enclosure can act as either a barrier or a conduit for carrying the heat away to the ambient or possibly both. Cooling is a system issue, which is why we advocate a top-down approach, starting at the enclosure level [1].

Often the air flow inside the enclosure is complex, making computational fluid dynamics (CFD) the only viable approach to developing a successful design. The electronics and enclosure design affect the flow; hence, the electronics and enclosure interact to provide the application thermal environment. In most cases, the primary cooling medium is ambient air. Even in something as small as a notebook computer where there is limited space, designers look to circulate the air to improve cooling performance, although heat spreading and conduction to the casing will also be important.

Whenever engineers and designers need enclosure-level air flow prediction, Simcenter™ Flotherm™ software, shown in figure 1, has been the industry's primary tool of choice for almost 30 years.[2] Simcenter is part of the Siemens Xcelerator business platform of software, hardware and services.

Figure 1. Airflow distribution in a 1U server modeled in Simcenter Flotherm.

Simcenter Flotherm uses a direct modeling approach in computer-aided design (CAD), allowing users to graphically create, resize, reposition, etc., objects rather than define them parametrically by typing in dimensions (although this is possible). As such it is extremely fast for use early in the design phase. Users can build models in just a few minutes and solve them in seconds. In addition, Simcenter Flotherm represents objects in the exact same hierarchy order that is common in many product developments. There is a main assembly with subassemblies and parts, and all part attributes are directly linked to and accessible via the part object. There is no step to map different hierarchies onto each other, which translates into time savings and lower error rates. Miniaturization has increased importance of mechanical CAD (MCAD) in many electronics products, necessitating the need to work with native CAD geometry, which is defined parametrically.

1 Chassis designed to house third-party circuit card assemblies is a more specialized topic.

2 See references to the ThermaNator in 'Hot Air Rises and Heat Sinks' [2].

Optimizing the air flow

Enclosure thermal design is then all about optimizing the system air flow when installing the electronics. That does not mean it needs to be done as an afterthought after completing the electronics' design. On the contrary, to ensure the shortest design time and lowest design cost while arriving at a product that has a low-cost, high-reliability cooling solution, it is essential to take a codesign approach. The golden rule is to start early and start simple [1].

This codesign can start with the concept design phase. It can be a prerequisite for getting the architectural design for the electronics correct. Where there is flexibility over the enclosure design, using Simcenter Flotherm software provides an enclosure, Simcenter Smartparts software offering a fast, parametrically defined representation of the enclosure.

In other cases, the enclosure design may be largely predefined, but there may still be a need to optimize the enclosure for cooling the electronics. Using the Simcenter Flotherm MCAD Bridge module provides the ability to simplify and import native CAD and standard CAD neutral file formats.

The main design flows, MCAD and electronic design automation (EDA) or electronic CAD (ECAD) focus on the physical geometry, and in doing so, miss the most important part of most electronics designs from a cooling perspective – the air gaps. Hence, we have our first key consideration.

> "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."

Wendy Luiten,
Thermal Specialist and DfSS Master Black Belt Consulting

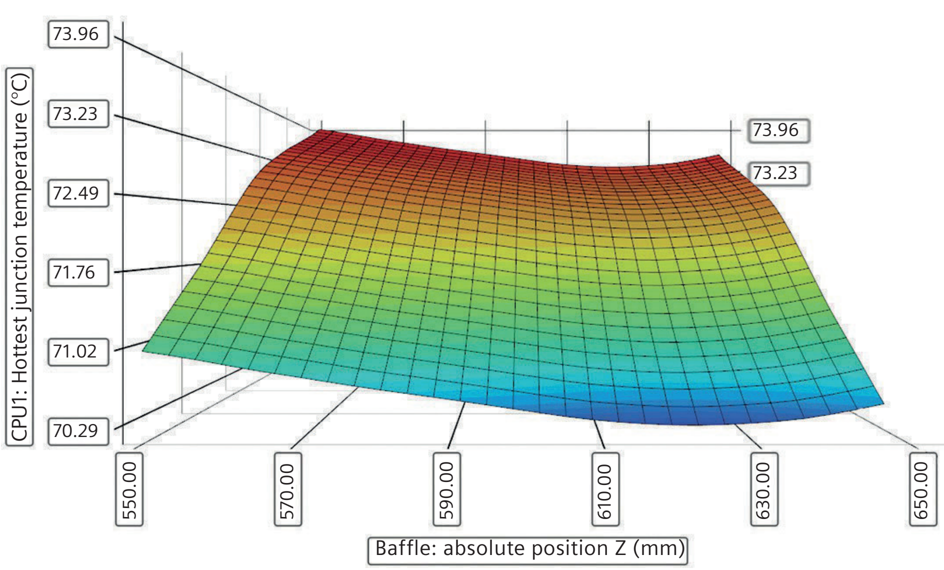

Figure 2. Response surface showing junction temperature variation with baffle position.

1. Protect the critical white space in your design

The white space is all the gaps between the geometry in the design. Not all of this will play an important role in cooling the electronics, but in many cases, providing an adequate space for the cooling air to pass through is an important part of the design. One problem thermal designers face is that the air is not part of the bill-of-materials (BOM) associated with the product. As such, designers can adjust and close gaps without raising any change request. Even if that does not happen, other people involved in the design activity are free to fill these spaces, for example with cables to connect the boards and other components together. As Wendy Luiten, Wendy Luiten Consulting, thermal specialist and design for sigma six (DfSS) master black belt puts it, "Beware of the cable guy."

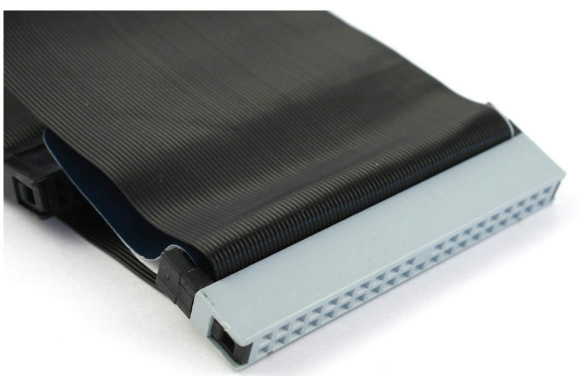

Figure 3. Strip cable for a computer hard disk drive.

The cables may not fill the space, but they reduce the area available for flow and increase the surface area, reducing the effective hydraulic diameter of the channel, and hence the flow rate through it and for systems with a fan, increasing the pressure drop. For fanless systems, a cable in an air duct can block flow in this duct completely. Strip cables, placed across the air flow, can have a major disrupting influence on the system-level air flow. Thus, when considering the design of a board from a thermal perspective, pay attention to the connectors not just the high-power dissipating components.

One approach to gaining control over this is to have the MCAD designers include dummy parts for the important air gaps, which then cannot be adjusted without a change request, with the thermal engineer involved in the approval process for changes to that part.

Open, forced convection systems
Due to their prevalence, we will first discuss the key considerations when designing open systems and then the closed or sealed systems.

2. Select the right fan type and size

Different types of fans are available – axial, centrifugal, radial and even mixed flow fans. These have various flow rate versus pressure characteristics.

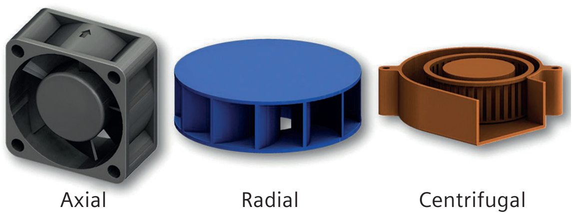

Figure 4. Common types of fan used for electronics cooling.

Radial fans generally produce the highest maximum pressure rise, but at the cost of delivering a lower maximum flow rate.

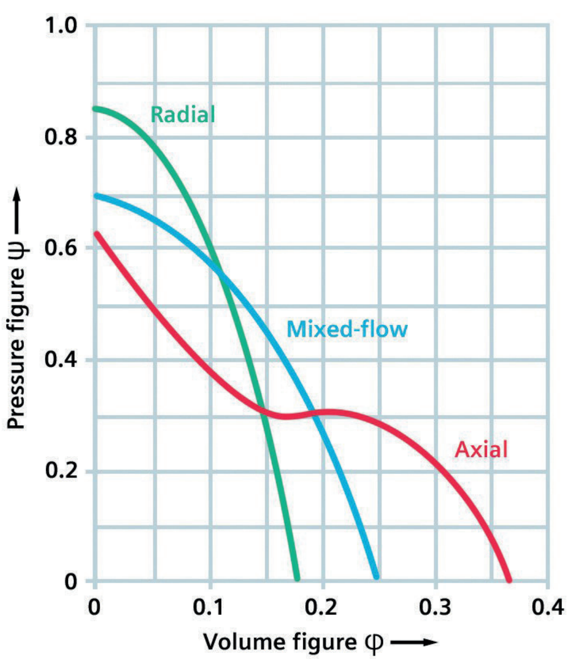

Figure 5. Normalized pressure versus volume flow rate for various fan types [3].

To achieve the same airflow, albeit with different flow rate and pressure rise characteristics and if space permits, designers may need to choose between using larger fans operated at a lower speed or smaller fans controlled at a higher speed.

There are benefits of using a larger fan operating at lower speeds since it will generate less aeroacoustic noise and will have a longer field life. Using a larger fan in early design work is conservative, as there is a risk that the initial power estimates may turn out to be on the low side. A smaller fan could then be used, whereas the reverse may not be the case. This is not only a question of finding the space to mount the larger fan itself, but also because a fan needs about one fan diameter of free space at the inlet and the outlet, and this typically is not available. If the powers you are designing for are lower than expected, you can reduce the fan size later on or replace the fan with another of the same size that delivers less flow for a lower cost and power consumption. It also allows scope for intelligence to be built into the cooling solution as we will discuss below. There is clearly a cost implication, as the larger fan is likely to be more expensive, and electromagnetic compatibility (EMC) emissions may be harder to contain, depending on how the equipment is shielded. Finger guards fitted to fans can be grounded to the casing to help reduce EMC issues.

3. Select the best fan flow arrangement

Another common system-level design consideration is whether to opt for a push or a pull type arrangement. For example, if fans are to be used to push air through the system or draw air out of the system. In general, pull type arrangements are the most common, found in most deskside, desktop and laptop computers. Here, the air is pulled into the enclosure through vents by a fan or fans that exhaust out of the enclosure.

In pull type arrangements, the electronics (for example, a rack of boards) act as a flow resistance, which tends to even out the flow through the card slots. For a bespoke system where the designs of the printed circuit board (PCB) the system will house will not change, engineers can optimize the board spacing to get the best cooling for each board. However, this is not recommended for modular systems designed to house a range of PCBs, often from different vendors, and with potential upgrades over the life of the system.

Some systems may require a dual push and pull arrangement to achieve higher volume flow rates.

If a pull type arrangement is used, the static pressure inside the system will be lower than the ambient. Consequently, any leaks in the system will allow air to bleed into the system and may bypass the electronics and therefore not contribute to cooling the system. This air will also bypass any inlet filters designed to prevent dust and dirt entering the system. Conversely, if a push type system is used, air can bleed out and again not contribute to the system cooling.

During the detailed design phase, or earlier if a preexisting enclosure is being repurposed, the mechanical design of the enclosure available as a CAD model should be used, paying particular attention to joins in the casing.

4. Use a plenum

In push type arrangements, the incoming air naturally forms a jet, with recirculation regions forming at the sides of the jet. However, we want the incoming air to flow uniformly over the downstream electronics, especially when forced air cooling a rack of PCBs where we desire equally apportioned airflow through the gaps between the cards. Engineers can introduce a pressure plate downstream of the fan to make flow more uniform. The downside is that this adds to the overall system pressure drop the fan must overcome. The higher the flow resistance, the more uniform the flow rate through it. Hence, the flow uniformity needs to be balanced with the cost, in terms of increased pressure drop, required to achieve it. Simcenter Flotherm has a general flow resistance and perforated plate Smartparts to assist with this task.

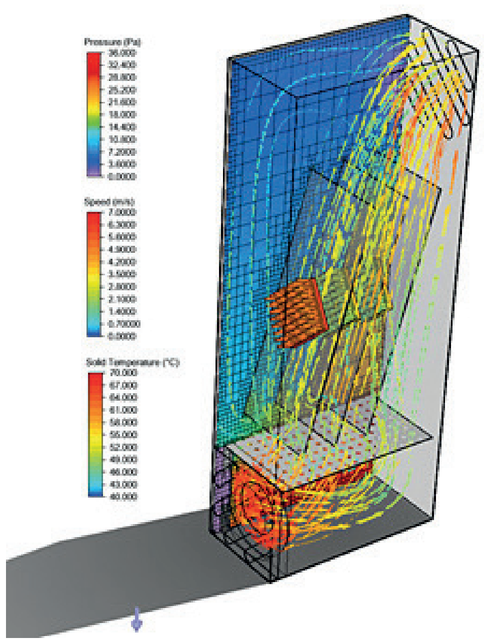

Figure 6. Flow resistance Smartpart in Simcenter model of a rack system.

Another factor to consider is that even axial fans generate nonaxial flow. At high-volume flow rate and low-pressure rise, axial fans predominantly exhaust air axially. At lower flow rates, the radial and tangential components of the air leaving the fan increase. Unless a plenum arrangement is used, this can change the near field flow pattern downstream of the fan and, hence, the local cooling.

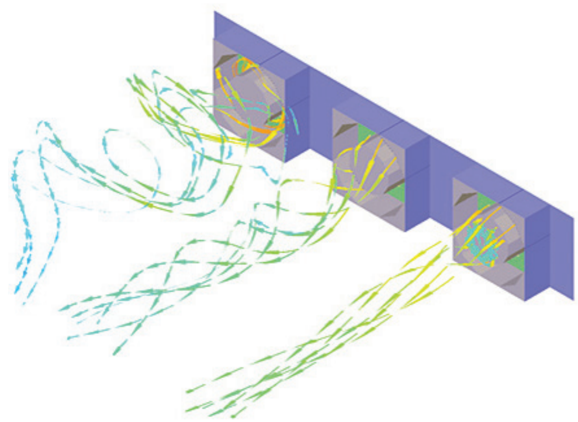

Figure 7. Simcenter Flotherm model of an identical fan operating at three different operating points and showing different spreads.

5. Size vents correctly

Regardless of whether a push or pull type fan arrangement is used, arrangements need to be made for air to enter or exit the system. The size of the vent or vents affects the volume of air flowing through the system. What is perhaps less apparent is that it will also affect the flow pattern within the system.

Similar considerations to fan sizing also apply to vent sizing. Larger vents will reduce the overall system pressure drop, which will increase the volume flow through the system. How much will depend on the flow resistance of all the other geometry. It will also affect the speed that air enters or exits the system via the vent. In pull type systems, the jet effect caused by the air entering at high speed can set up flow recirculation regions within the system, potentially leading to hot spots.

Vents in the system may also give rise to EMC issues, which in turn restrict the maximum hole size of the screen used for the vent. Engineers must consider the flow resistance of this screen. Honeycomb EMC screens offer excellent shielding and have a low blockage factor, since the aluminum sheet used is thin, giving an open area greater than 95 percent. However, the small diameter combined with the length of these honeycomb cells results in a surprisingly high pressure loss, comparable to that of a conventional drilled plate.

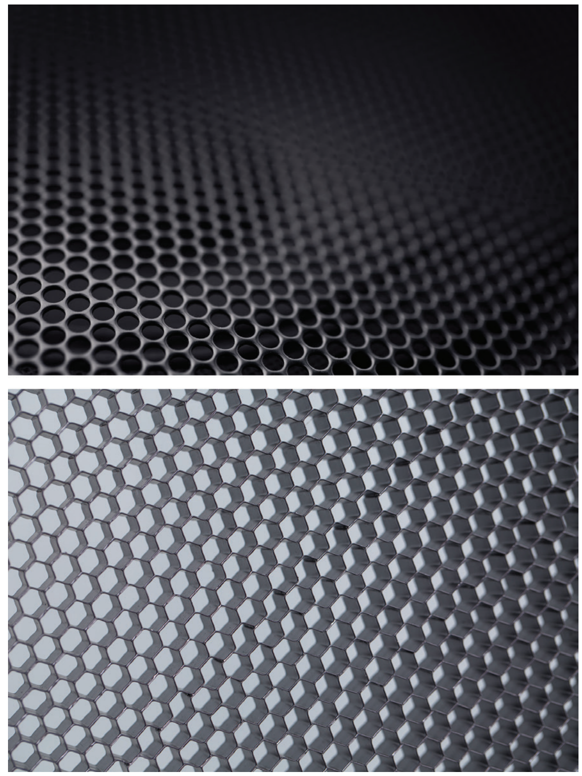

Figure 8. EMC screens conventional drilled plate (top) and honeycomb (below).

6. Optimize the fan and vent placements

Where a fan is placed on one wall of the enclosure, typical of deskside computers, rack mounted servers, etc., the fan's position will influence the flow pattern within the system, as will the location of the vent(s) that supply the fan, so engineers should optimize the fan and vent placements together.

In pull type systems, having the exhaust fan near to the top of the enclosure helps ensure it exhausts the hottest air from the system, as buoyancy effects are always present. Otherwise, a blanket of hot air will develop at the top of the enclosure. Conversely, in push type systems, engineers should position the inlet fan near the bottom of the enclosure to ensure cold air is brought into the bottom of the enclosure, making sure it passes over the electronics, assisted by buoyancy effects.

We advocate activating buoyancy in all simulations, as this adds to the realism of the results obtained at no cost in terms of the simulation time taken when using Simcenter software. It is essential for naturally ventilated systems, as it is only buoyancy that drives the airflow in such an arrangement. However, it is also essential in case the mechanical layout does not allow for optimal fan placement. Cases exist where a slowly rotating fan in the wrong location inadvertently balanced and counteracted natural convection, resulting in a pocket of still air and too high of temperatures.

Engineers must be careful with vent placement to prevent hot exhaust air being drawn back into the system. If in doubt, they can extend the model to include the air flow in the local environment. This is also recommended for high vertical structures, as this allows calculating the vertical temperature gradient in the outside air adjacent to the enclosure wall.

7. Optimize board placement

Something often overlooked is the potential to adjust the location of the PCBs within the enclosure. In a rack system, it is obvious that the card spacings can be altered, but in other systems this is perhaps less apparent. Altering the position of the PCB can adjust the amount of air that flows over and under the PCB. Minor changes can provide a cost-free improvement in the cooling performance, for example, by increasing the cooling off the back side of the board. This may arise from the interaction of the leading edge of a PCB with the near field flow around a fan or vent for example.

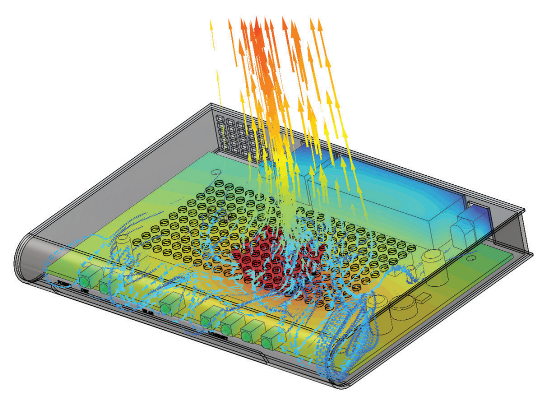

Figure 9. Buoyancy induced airflow in a model of a router.

Open, free convection systems
Free convection systems range from digital TV set top boxes to rack systems and historically telecoms racks. In naturally vented systems, all the convective cooling arises from buoyancy induced by air inside the system becoming hotter, and, therefore, lighter than the surrounding air.

8. Design vents carefully

Smooth uniform air flow is key to good design, and consideration needs to be given to the detailed design of the vent grilles to minimize pressure loss while maintaining EMC and electromagnetic compatibility immunity (EMI).

To design the inlet vent grill, it is often necessary to consider how the air approaches the vent, requiring the solution domain to be extended beyond the enclosure, to include the mounting surface for the unit.

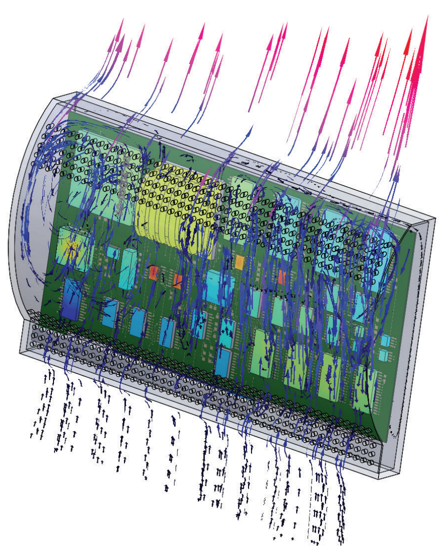

Figure 10. Flow distribution showing the optimized vent arrangement for a naturally ventilated, wall mounted electronics system.

In addition, location and open areas are especially important for free convection cooled systems. Air inlets should be placed low; air outlets should be placed high and open areas should be roughly the same size in the vents and along the flow path through the system.

Enclosed/sealed systems
In small sealed systems, the size of the enclosure is a major parameter affecting cooling performance. In larger enclosed systems, fans can be used to recirculate the air, for example, through an internal heat exchanger. When considering using a fan within an enclosed system, it is important to bear in mind that the fan itself, and the additional power draw on the power supply, will also add heat into the system. Thus, the improvement in cooling performance needs to outweigh the additional cooling required.

9. Do not ignore conduction to the enclosure

Conduction to the enclosure can be the primary cooling mechanism for some fully enclosed systems, for example, avionics boxes and smartphones. Consider how to remove heat from the casing. Often natural convection to the ambient or conduction to a human hand are the only ways heat can ultimately leave the system. In both cases, it is highly desirable to spread the heat as much as possible to make the best use of the casing's surface area. As such, the casing should be as thermally conductive as possible, so it extends the hot area in the same way a heatsink does. If not, heat loss from the casing will be reduced, and simultaneously, a hot spot will arise, which may lead to problems due to the safe limit for touch temperature [4].

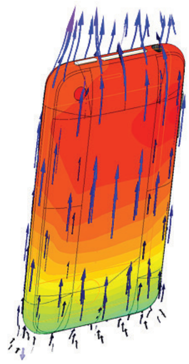

Figure 11. Case temperatures and buoyancy-driven airflow shown on a model of a smartphone.

In addition to conductive casings, miniature heat pipes are available for this purpose as well as graphite foil heat spreaders. The graphite foils have the added benefit of providing EMC screening.

As a system-level cooling solution, we should also mention that engineers can use conducting rails to carry heat to wedgelocks that clamp a PCB to an externally cooled cold wall, for example, as part of a ruggedized enclosure, it also provides mechanical stability for the board mounting. This approach only applies to a certain type of system, for example, a sealed avionics box, but it is sometimes possible to apply the same principle in other systems where the casing itself can be used as part of the cooling solution. Additionally, engineers can also insert gap pads between the top of board-mounted components and a metal enclosure to provide extra heatsinking.

10. Design in internal heat spreaders as necessary

A final point about conduction concerns using heat internal spreaders. Depending on the type of system, internal heat spreaders can be used to augment cooling. These can be either tightly mechanically integrated as part of the PCB's construction, as in the case of a metal core PCB (MCPCB), or a mechanically separate system-level cooling solution that provides heat spreading on the top or back of the PCB or memory module. Engineers usually need to design in these solutions from the outset.

Many material options are available, including graphite foils, highly oriented pyrolytic graphite (HOPG), various ceramics, silicon carbide (SiC) and flat heat pipe constructions. Material and processing costs vary considerably, and using an internal heat spreader also adds to the assembly cost as heat must be efficiently conducted into the heat spreader for it to be effective. Although internal heat spreaders can offer a simple means of improving thermal performance, engineers should weigh the benefit against other options.

11. Understand whether a phase change material can help

Phase change materials are a relatively new addition to a thermal designer's toolset. Phase change materials absorb heat by turning from a solid to a liquid over a relatively narrow temperature range, limiting the temperature rise until all the solid has become liquid. Phase change materials have been used as thermal interface materials for heatsink attachment for several years but have more recently found use in applications where engineers must design for high bursts of power, for example, watching a video on a mobile phone. Using a phase change material can delay the need to throttle back the performance to manage heat dissipation. When reducing the power, the temperature drops and after reaching the range of the phase transition, the phase change material starts to resolidify, rejecting the stored heat. The phase change material is typically contained in a bag that can conform to the shape of the cavity it is inserted into, providing good contact with the electronics, and in some cases, the enclosure, to assist with heat rejection. See reference 5 for a study into the effects of incorporating a phase change material into an existing design of a laptop computer.

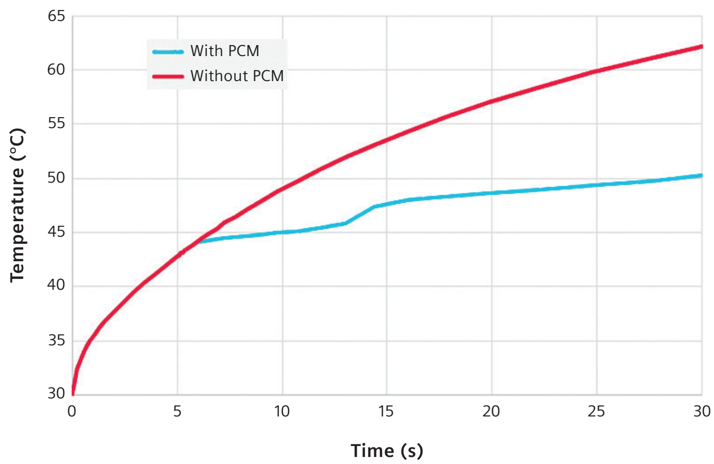

Figure 13. Component temperature rise for a system with and without phase change materials.

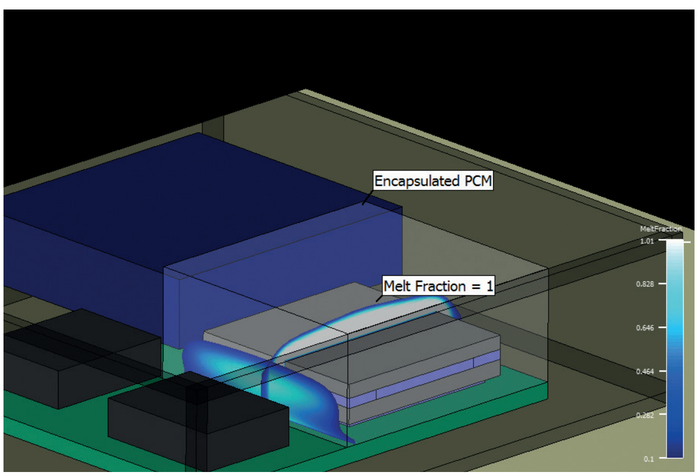

Figure 12. Melt fraction shown in an encapsulated phase change material heated by a package.

It is relatively easy to model a phase change material in a transient simulation with electronics cooling CFD, helping thermal designers decide if this should be part of the cooling solution during the architectural design phase, and optimizing the design later. The phase change material will require a heat spreader due to the relatively low thermal conductivity of the material itself. Introducing a phase change limits the temperature rise within the enclosure and can also alter heat flow paths, so that the temperature rise after the phase change material may be different to when there is no phase change material included.

12. Include thermal radiation effects

Designers often ignore thermal radiation effects in early design, as this is generally a conservative assumption, and often the model fidelity does not provide sufficiently accurate surface temperatures to warrant the addition computational overhead. If radiative heat loss is considered important, its effects can be represented as an augmentation to convective heat transfer coefficients, increasing them by about 5 W/m2K on materials having a high emissivity (for example, greater than 0.8). This is in the same order of magnitude as the heat transfer coefficient for natural convection. Thus, we advise incorporating thermal radiation in early design of natural convection cooled products.

Engineers should consider the radiation effects during detailed design to avoid general overdesign and problems. For example, a temperature-sensitive component on one board may be heated by radiation from a high-power dissipating component on an adjacent PCB. Radiation is particularly important when natural convection makes a significant contribution to cooling. In rack systems cooled by natural convection, for example, board to board radiation and radiation to the side panels of the sub-rack, can significantly reduce the temperature variation across the rack, preventing overdesign of the cooling for centrally located boards. Surface emissivity influences heat loss, and for metal enclosures it is worth testing the sensitivity of the simulation results to see if coating or painting the inside of the enclosure is worthwhile. In one case, painting the inside of the aluminum housing of an automotive electronic control unit (ECU) reduced component temperature up to 7°C, easily warranting the additional cost.

Solar radiation deserves a special mention at this point. For outdoor equipment that is exposed to sunlight, solar heating of the system can present a major challenge. Shielding the electronics from the heating effects of sunlight may require a dual skin design with controlled ventilation and with additional heating for frost protection.

Specialized cabinets are available for this purpose. Note that the color of the enclosure is important as this affects solar absorptivity.

13. Codesign the enclosure and electronics

In this white paper, we have discussed the design tasks associated with enclosure design. We have not yet talked about these tasks in the context of the overall product design process.

In conceptual design, a simple representation of the enclosure may be all that is possible. Alternatively, there is little constraint from other aspects of the design, so there is great scope for designing a highly thermally efficient enclosure that has low design, manufacturing and operating costs.

> "I find the Command Center 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."
>
> **Wendy Luiten,**
> **Thermal Specialist and DfSS Master Black Belt Consulting**

Leveraging the Simcenter Flotherm Command Center module provides an easy way to explore the design space in early design. Developing a complete understanding of the design space allows the thermal engineer to make informed design decisions. Users can deploy design optimization tools automatically in Command Center to systematically improve aspects of the product.

The enclosure design itself should be updated in the thermal model as soon as design details are committed to the CAD system and updated thereafter in step with updates to the CAD geometry to ensure these changes do not negatively impact the overall thermal design. The system-level model should be reran whenever there is a change to the PCB layout and heatsinking arrangements for any components and if information about the component powers changes. The aim should be to keep the thermal model of the system in sync with the design information held in main MCAD and EDA design flows.

14. Design in cooling intelligence

Increasingly it is impractical and not cost effective to design cooling solutions for systems based on the maximum thermal load. For example, multicore computers and servers often spend periods of time running with relatively low workloads. Today's equipment is designed to draw little power when idle, so the cooling can be throttled back for much of the time. This is a thermal opportunity, but also a challenge – as it directly links cooling performance to system performance.

Incorporating cooling intelligence makes good electronic cooling solutions more important, not less, as an inadequate cooling solution will cause earlier throttling and directly impact functional performance. This performance loss will negatively impact a portion of the user's experience (UX) of the product. It is difficult to relate that working point to use cases to determine how often this will occur, and when it does, how it will affect using the product.

In early design phases, engineers can use a representative thermal design power [6], allowing them to leverage steady-state calculations to investigate various aspects of the cooling solution discussed above. In the detailed design phase, however, it may be necessary to consider different use cases, each having various power levels for the critical components and, in some cases, unique power distributions on the surface of some of the packaged die, for example, in high-power density mobile applications.

The task is to optimize the cooling solution for a range of use cases and design the cooling solution to adapt to the various power levels it must cool, thereby minimizing the power the cooling solution requires. In the simplest case, this can mean having the fan normally run at slow speed and increasing the fan speed when a component temperature exceeds a predefined value.

Using Simcenter Flotherm can run design variants, and support simulating systems with dynamic cooling solutions such as temperature controlled variable speed fans. The Command Center is the go-to tool to run the various scenarios necessary to determine the transfer function between temperatures, cooling intelligence and system performance, enabling the design to work for operating points that fulfil end customers' expectations.

Conclusion

This white paper provides a complete guide to enclosure thermal design, covering all the main points that need to be considered. If you are responsible for enclosure-level thermal design, leveraging Simcenter provides a range of thermal design software, allowing you to select a product that is an ideal fit with your company's thermal design workflow and supply chain needs.

Acknowledgements

Dr. Ian Clark, Simcenter Flotherm Product Line Director
John Wilson, Simcenter Flotherm Product Specialist
Ir. Wendy Luiten, Principal at Wendy Luiten Consultancy
Byron Blackmore, Simcenter Flotherm Product Manager

References

1. Parry, J., Bornoff, R. and Blackmore, B. "Move your thermal strategy for air-cooled electronics up in the design flow," Engineering Essentials Series, ED Online 21417, 2009. www.electronicdesign.com.

2. Kordyban, T. "Hot air rises and heat sinks," ASME Press, 1998.

3. Angelis, W. "General aspects on fan selection and layout," Electronics Cooling Magazine, May 2001. http://www.electronics-cooling.com/2001/05/general-aspects-on-fan-selection-and-layout

4. Divakar, MP. "Surface temperatures of electronics products: appliances vs. wearables," Electronics Cooling Magazine, September 2016. https://www.electronics-cooling.com/2016/09/surface-temperatures-of-electronics-products-appliances-vs-wearables

5. Maltz, W., Moore, D. and Raghupathy, A. "Application of phase change materials in handheld computing devices," Electronics Cooling Magazine, 07 July 2016. https://www.electronics-cooling.com/2016/07/application-of-phase-change-materials-in-handheld-computing-devices

6. Thermal design power https://en.wikipedia.org/wiki/Thermal_design_power

Siemens Digital Industries Software

Americas (USA): 1-800-498-5351

EMEA (United Kingdom): 0800-279-0464

APAC (India): 1-800-202-6796

For additional numbers, click here.

Siemens Digital Industries Software helps organizations of all sizes digitally transform using software, hardware and services from the Siemens Xcelerator business platform. Siemens' software and the comprehensive Digital Twin enable companies to optimize their design, engineering and manufacturing processes to turn today's ideas into the sustainable products of the future. From chips to entire systems, from product to process, across all industries, Siemens Digital Industries Software – Accelerating transformation.

siemens.com/software

74414-D13 3/26 H