

Without Thermal Analysis, You Might Get Burned

Thermal analysis used to be an afterthought, but now many designers must consider it up front.

Remember when thermal analysis meant getting your prototype back and deciding if you might need to throw in a couple of heatsinks and a fan for good measure? Try that approach now and you may find yourself in deep and without a paddle. After all, heat can hamper electrical performance and ultimately reduce mean-time-between failures.

Back in my engineering heyday, I never put much thought into thermal analysis because it just wasn't necessary, and I know I'm not alone. But with semiconductors dissipating greater amounts of power (and therefore heat) per area than ever, coupled with continued system shrinkage over time, more system engineers who don't perform thermal analysis are winding up in hot water.

"A lot of functions that used to be spread across several components are now contained in a single component," says Dave Rosato, lead product manager for Ansys. So now, the heat density is much greater for those SoC-type (system-on-a-chip) components.

"The rules of thumb that engineers used to design a board five and 10 years ago just don't apply to today's designs," continues Rosato. "Years ago, the board was ignored as a heat transfer path. Now you must account for all heat transfer paths."

The "simple solution" is to perform thermal analysis sooner in the design cycle. How soon? At the least, you should perform a rudimentary analysis just after the block diagram stage. You'll need to download the datasheets for the components you plan to use and get a feel for future challenges from a thermal standpoint.

If that analysis points to potential trouble, you need to consider using some thermal-analysis simulation software and possibly even working with a materials company to determine if it can engineer something that will suit your design parameters.

"DANGER, WILL ROBINSON!"


I own a laptop that recently stopped working because the fan integrated with the heat-sink/heatpipe combination no longer gets powered correctly. Even with the case open and plenty of cool air all around, the unit won't power up and the "Fan error" message appears before it even performs the typical power-on self-test (POST).

It immediately shuts down when it senses the fan isn't powered on. The assumption is that the average laptop user won't pop the case open in a nice air-conditioned room, and thus the CPU will experience the often fatal "thermal runaway." The downside to this approach is that my entire system is shot because the fan (or the underlying power source to the fan) isn't working.

This is a good example of a laptop manufacturer deciding that under no circumstances is the CPU to ever run without forced air blowing on the attached heatsink. This design was engineered with these requirements because the laptop designers knew that improper thermal management meant imminent doom. In fact, Intel and AMD take this problem very seriously.

For example, "If the external thermal sensor detects a catastrophic processor temperature of 125°C (maximum), or if the THERMTRIP# signal is asserted, the V_{CC} supply to the processor must be turned off within 500 ms to prevent permanent silicon damage due to thermal runaway of the processor," says the January 2008 edition of the datasheet for Intel's Core 2 Duo Processor.

"Maintaining the proper thermal environment is key to reliable, long-term system operation. A complete thermal solution includes both component- and system-level thermal management features," according to the datasheet.



1. Flomerics' FloPCB software shown can simulate the junction temperatures of several components on a printed-circuit board. (courtesy of Anritsu)

“To allow for the optimal operation and long-term reliability of Intel processor-based systems, the system/processor thermal solution should be designed so the processor remains within the minimum and maximum junction temperature (T_J) specifications and the corresponding thermal design power (TDP) value,” it notes.

“Caution: operating the processor outside these operating limits may result in permanent damage to the processor and potentially other components in the system,” the datasheet concludes.


Why are companies taking such grand steps to curtail improper thermal management? “A lot of applications (systems) are getting smaller, e.g., Mac Air, and the thermal path is being both shortened and rearranged,” says Sara N. Paisner, senior microelectronics technology scientist at Lord Corp.

Generally, heatsinks are placed directly above the component. But the latest techniques move heat in alternative directions. “Now the heatsink may be behind the component, or heat may be dissipated through the board itself,” says Paisner.

Yet thermal management isn't so simple anymore. “Casing material is acting as both an EMF (electromagnetic field) shield and a heatsink, as the casing itself has become part of thermal path,” says Paisner. A typical printed-circuit board (PCB) includes a built-in heat path, causing systems engineers to rethink their design strategy. Everything is shrinking, and now several components share cooling responsibilities while heat transfers to a larger area.

The preventative measures taken by Intel and AMD with respect to proper thermal design are interesting from a chip perspective. To start with, Intel indicates that “The processor requires a thermal solution to maintain temperatures within operating limits.” It uses thermal diodes, digital thermal sensors (DTSs), and the Intel Thermal Monitor to monitor die temperature.

Used in conjunction with the thermal sensor, the thermal diode can be used to calculate silicon temperature. The DTS is an on-die sensor that continuously monitors and outputs data on the die temperature relative to the maximum thermal junction temperature. Temperatures that will cause catastrophic conditions can be detected when a special bit is set in the DTS.



2. The Nextreme OptoCooler module can be used to cool devices like this laser diode, which is housed in a TO-56 package.

The Intel Thermal Monitor helps control the processor temperature by activating a thermal control circuit when the silicon temperature reaches the maximum. This, in turn, modulates the core clock as needed to keep the silicon temperature in check.

Also, the monitor generates an external signal (PROCHOT#) if the processor is above the thermal trip point. It can generate an interrupt signal as well. If the monitor is deactivated, a special signal (THERMTRIP#) will be asserted, indicating imminent failure if the core voltage isn't switched off immediately.

AMD takes a similar approach. Its “Thermal Design Guidelines” whitepaper provides specifications such as the maximum length, width, and height of the heatsink, in addition to the heatsink and fan material requirements.

While CPUs are an easy target because they dissipate so much heat, other system components must not be overlooked. This is where some simple calculations come into play, as well as some basic thermal-management theory.

THE JUNCTION BONE'S CONNECTED TO THE SINK BONE

Thermal management moves heat from the semiconductor junction and into the surrounding ambient environment. Typically, heat is transferred from the semiconductor to the package, then to the heat spreader (sink), and finally to the ambient environment. Your design may not have a heatsink, or it may have more exotic technologies like fans and pipes.

Still, the general theory remains the same—spread heat from a small area to a large area. According to the basic theory

of thermal conductivity, the rate at which heat conducts through a material is proportional to the area perpendicular to the flow of heat and the temperature gradient.

Junction temperature (T_J) is the operating temperature (typically in °C) of the semiconductor junction, where most of the heat is generated. Thermal resistance is the effective temperature rise (typically in °C) per unit of power dissipation (typically watts) of a designated reference point (such as junction or case) above an external reference point, such as the lead, case, or ambient air.

Thermal resistance is expressed as $\Theta_{\text{Letter1Letter2}}$ (e.g., Θ_{CA} or Θ_{JA}). Letter1 is the designated reference point and the letter typically represents the initial for the reference (e.g., C = case; J = junction). Letter2 is the external reference point and has a similar representation structure (e.g., A = ambient).

BACK-OF-THE-ENVELOPE CALCULATION

When a formal thermal analysis is performed, the goal is to provide a complete understanding of how heat is both formed and moved throughout the system. However, a simple back-of-the-envelope calculation may be quite sufficient in the early stages of the development process.

The idea is to get a rough feeling of just how hot things are going to get after throwing the power switch. Another way to look at it is that you're preventing the inadvertent reduction of the mean time between failures by letting a device or the system overheat.

Once you perform the calculation, you should have a basic understanding of the level of sophistication needed for your thermal-management scheme. That is, are you looking at adding a simple heatsink to your bill of materials, a more exotic solution requiring a heatpipe, or some cutting-edge solution that uses a combination of heat spreading, forced air, and even new materials? Even if you can get away with something simple like adding thermal vias, it's much better to know up front and plan for it than getting burned later.


So how do you perform a back-of-the-envelope thermal analysis? According to Byron Blackmore, electronics cooling engineering supervisor for Flometrics Inc.,

one of the first numbers to crunch is the total power density on both surfaces of the board. "This can be determined by calculating the total power dissipation divided by the surface area," he says.

Blackmore also provided a rough rule of thumb by indicating that if your calculation reveals your design will dissipate more than 1.5 W/in.², you need to start thinking about additional measures to keep heat from creating downstream issues.

Paisner also chimed in with some guideline numbers. "One of the key determining factors for additional action is temperature," she says. Up to 85°C is acceptable, and 85°C to 100°C is probably okay, but proceed with caution. However, additional measures typically will be needed at 100°C and higher. Of course, in addition to the absolute temperature, you should worry about how the temperature changes as system conditions change.

How do you get there? "Take the maximum power dissipation of each component at the highest temp the board will run at and divide by the surface area, and then repeat for the other side of the board," says Blackmore. Then, you must research the thermal resistance (e.g., Θ_{JA}) and multiply by expected power dissipation



3. Ansys' IcePak tool provides thermal simulation of a complete system, including component, board, and cabinet-level thermal analysis.

to determine temperature rise above ambient. Now, compare that number to the maximum rated temperature for the component.

Note that the Θ_{JA} listed is for "stale air" and must be taken with a grain of salt, especially if you plan to have air moving through the system. Some datasheets may list the thermal resistance at a given airflow rate above the part (e.g., Θ_{JMA}). Obviously, if your design is pushing one of these limits, you probably need to consider additional thermal-management measures, and it may be time to think about simulation software.

These calculations may be sufficient for a given design, especially if you have a lot of leeway in regards to the system chassis. So when may additional thermal analysis be required?

"Optimally, you would like to do thermal analysis twice: once after the EE has a rough idea of the board size and components that will be used, and later when a preliminary route has been performed," says Rosato. Again, depending on your system, you may need to consider a much more accurate simulation using thermal-analysis software at this post-layout point (Fig. 1).

LAYOUT AND CHASSIS CONSIDERATIONS

Thermal analysis must be performed early and often. Some designers may even want to consider it before going after a patent, because if a product will fail due to a thermal problem, what's the point? But other factors impact the system design.


"[Systems] engineers must understand how different materials interact with various package sizes and types," says Paisner. "Companies like Lord Corporation work with customers to develop new materials to meet thermal requirements."

She used Apple's Mac Air notebook as an example of a product with significant design challenges, because designs like that likely don't have room for large heatsinks or other cooling technologies. As a result, the limitations of an extremely small form factor can be overbearing unless you're willing to spend some serious cash for exotic thermal solutions.

"The more complex the thermal path, the higher the cost," says Paisner. "Then you must figure out how you are going to get heat out of the system, and what material and layout tradeoffs are you willing to concede."

Additionally, component placement plays a major role during layout from a thermal perspective. "The preference for components dissipating a lot of heat is to place them near a vent, but that is not always possible, and other tradeoffs may be necessary," Rosato says.

In addition, components that dissipate a lot of power may generate "downstream" heat, which could easily affect other components. Another trick of the trade is to place heat-generating components side by side and normal to the air path. Also, "Diverters may be used to route airflow where necessary," notes Rosato.



From concept to completion,
We can find a solution!

- We optimize your heatsink designs.
- Complimentary designs, from CFD simulations to final product!
- Custom and standard solutions to solve your thermal needs.



radian
ISO 9001:2000 Certified

web: www.radianheatsinks.com · tel: 001-408-988-6200 · fax: 001-408-988-0683

Radian Heatsinks is a division of Intricast Company Inc.

READER SERVICE 116

PICO

New High Voltage Hi Power Regulated DC-DC Converters



High Voltage, Isolated Outputs
100-500 VDC

Output Voltages from 5VDC to 500VDC
High Power; to 50 Watts,
Efficiency to 90%

Miniaturized Size Package:
2.5" X 1.55" X 0.50"

Two Standard Wide Input Ranges

Safe: Short Circuit, Over/Under
Voltage, and Over Temperature
Protected

Options Available: Expanded
Operating Temperature, -55 to +85°C
Environmental Screening, Selected
from MIL Std. 883

Ruggedized for Operation in Harsh
Environments

External Bias Control: For Charge
Pump Applications

Custom Modules: Available to
Optimize Your Designs, Special
Input or Output Voltages Available

PICO's QP Series compliments our 650
plus existing standard High Voltage
Modules. Isolated, Regulated,
Programmable, COTS and Custom
Modules available to 10,000 Vdc and
other High Voltage to 150 Watts!

www.picoelectronics.com
E-Mail: info@picoelectronics.com

or send direct for free PICO Catalog
Call Toll Free 800 431-1064

PICO Electronics, Inc.
143 Sparks Ave, Pelham, NY 10803-1837
FAX 914-738-8225

From a layout perspective, keep your eye out for stacked-die or stacked-chip configurations, as taller components tend to impede heat paths. Also, components that can be soldered directly to the PCB (and eliminate any air gap between the component and PCB) can rely on the PCB to act as a heat spreader. Furthermore, thermal vias may be designed in, but typically you'd like to know that you're implementing them before layout.

According to Blackmore, a good layout rule of thumb is to strive to put the "leading edge" of any cooling air on the largest power dissipater. It's also wise to spread components out to avoid pockets of hot air downstream. Lastly, "Tall components and connectors could cause a dead zone for air blockage downstream," he says. Therefore, any tall components or connectors should raise a thermal red flag that may require further analysis.

GARBAGE IN, GARBAGE OUT

Don't assume the maximum power dissipation for your component set. The maximum may be fine during the calculation stage to get a rough idea of where you stand. But you must insist on using more realistic numbers or your design will likely get over-engineered, adding unnecessary weight and cost.

If you have an FPGA, is all of the internal logic going to switch at the maximum speed all of the time? That's highly unlikely, so get the logic engineer to give you a reasonable estimate for the assumed set of operating parameters. Then, it's up to you to decide if you want to add a fudge factor.

Keep in mind, though, that the FPGA manufacturer probably already has three levels of fudge built in by the engineering team, the testing team, and the sales/marketing department. If you can get actual usage data and add some fudge to that, you will wind up in much better shape.

Companies may then go on to ask you all-important questions: What is the error percentage? How do the numbers provided correlate to "real-life data?" Are the numbers validated? Were they tested using real materials in the end environment?

Then, where actual thermal simulation tools are concerned, you can get a much better feeling for accuracy. "Thermal-analysis simulation tools should be able to read

in routing and board design information, including traces, planes, and via definitions from other EDA tools," says Rosato.

Simulations can also include system packaging, detailed component design parameters, and so on. "Simulation tools can predict operating temperatures to see if rated junction temps will possibly be exceeded and where your system may have 'stale air,'" adds Rosato. The simulation may also take on an iterative approach, where engineers can play around with various thermal-management scenarios, add heat-sinks, and rerun the simulation as needed.

Parameters like board outline and size and the relevant board stackup data, such as information on the metal layers, are also read in, says Blackmore. The remainder of the process involves the systems engineer describing the environment in which the system will operate, including information on the chassis, vents, power supplies, and other parts. All information is then combined to provide a thermal simulation.

WHERE TO GO FROM HERE

So you now understand the basic principles and importance of thermal analysis and good thermal-management techniques. But what happens when your design reaches or exceeds some of these limits, such as 1.5 W/in.², even after all other precautions have been considered?

You're likely aware of the basic tradeoffs between heatsinks, fans, heatsinks with integrated fans, and so on. But what about advanced solutions? Many companies offer thermal products and solutions.

"Conventional solutions are out of gas, and thus, there became a need to extend the performance range by adding other capabilities," says Seri Lee, CTO of Nextreme. For example, heatpipes have solid-state refrigeration and would be considered more advanced than heatsinks and fans alone, yet they're big, bulky, and expensive and often must be custom-made.

Nextreme has several chip-level innovations that actively remove heat using technology that's 10 to 20 times thinner and smaller than typical solutions, yet provides 10 to 15 times greater heat-pumping capability (Fig. 2). Bergquist manufactures several different thermal materials and thermal substrates. Asys offers tools for thermal simulation as well (Fig. 3). 