Thermal Engineering Overview

Imagine you have just finished making the last few changes to a presentation for one of your biggest clients, the most critical presentation of the year. Then, your worst fear comes true - your system suddenly fails. Data stored on your hard drive is lost. You pull the cover off the machine and you see that the miniature fan on top of the CPU has stopped running. You quickly realize that your microprocessor failed due to inadequate cooling. Promptly, you purchase an upgrade microprocessor and behold - it has a heat sink with no fan!

Confused? You're not alone. This is a puzzle encountered by people who want to ensure their computer investment is protected with the right products to keep it functioning for years to come. The present problem with cooling microprocessors is a relatively new issue -- about 4 years old. It stems from the collision of two conflicting trends -- end-user desire for more powerful microprocessors to run the next generation software, and the equally strong demand for smaller, more mobile computer form factors. With each, the introduction of faster next-generation semiconductors aimed at improving computing power, heat concentration problems increase.

Thermal energy generated within these devices can be compared to that of a stovetop burner. Today's generation of Pentium, Athlon, and Power PC chips can dissipate more than 100 watts of power. In straight forward terms, you could fry an egg on top of any of these chips.

Choosing the right cooling solution is a complex task requiring the coordination of several factors and requires good thermal engineering. So how is it that certain upgrade processors do not require a heat sink with a miniature fan solution?

To answer that question, we need to first understand the fundamental problem. Heat to be removed must travel from the microelectronic chip to the surrounding air stream. As heat flows, it encounters a series of thermal resistance that impedes overall heat removal. The laws of physics dictate that performance and reliability of semiconductor and integrated circuit devices are absolutely constrained by temperature. Mathematicians have worked out formulas that inform us for every 10°C rise of the junction temperature (this corresponds to any one of the 4 million transistors that may be on a microprocessor) the failure rate doubles. There are any number of ways device temperature related failures modes may be manifested on the electronic level that cannot be visually detected.

These failure modes include such items as gate dielectric strength, junction fatigue, electromigration diffusion, electrical parameter shifts, and thermal runaway any of which could ultimately result in a failed CPU. But what does all of this tech-jargon really mean to you and me? It means that computer manufacturers and suppliers of upgrade processors must work diligently to ensure that their products are engineered properly to provide sufficient cooling to the microprocessor for an extended length of time.

Thermal engineering is more than a way of controlling heat and temperature rise in a computer. By definition, thermal engineering requires designers to pay attention to every aspect of heat generation and removal in an electronics system. It is NOT putting the biggest heat sink or biggest fan on a semiconductor. Rather, it is the careful and purposeful selection of parts and components, with heat removal in mind. An ideal cooling design achieves the desired junction temperatures, is compact in size, low in cost, adaptable enough to fit into system designs.

Attempts to increase capacity by using larger heat sinks, higher air velocities and fan heat sink combinations will be wrought with diminishing returns. Contrary to market demands for smaller packaging and higher performance, these solutins add to the packaging volume with larger parts or decreasing the system reliability.

So how do we choose the best solutions? In early stages of system design as well as the package design for the processor, thermal issues need to be recognized and advanced engineering solutions brought to the design table. The design of upgrade processors allows thermally optimized packages to be utilized in non-optimized systems. This is due to upfront thermal engineering efforts where materials, manufacturing techniques, and thermal barriers all maximize a processor's performance and reliability at the lowest cost and smallest sizes. While these upgrade processors then offer advanced thermal solutions where they can replace less reliable fan/sink combinations, they represent a change to thermal engineering in the earliest stages of product design.

Why don't the system manufacturers implement these advanced solutions in the original product design cycle? Often this is due to the fast pace product development cycles, and often it is just because the solutions were not available at the time the system was ready for market. Will you likely see lower reliability solutions in future system designs? Not if thermal engineering takes place!